

Quest® Collaboration Services 3.9
User Guide



© 2019 Quest Software Inc.

ALL RIGHTS RESERVED.

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

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

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

Quest Software Inc.
Attn: LEGAL Dept.
4 Polaris Way
Aliso Viejo, CA 92656

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


Patents


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


Trademarks

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

Legend

-  **WARNING:** A WARNING icon indicates a potential for property damage, personal injury, or death.

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

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

Contents

Getting started with Collaboration Services	6
Introduction	6
Security considerations	6
Authentication	6
Authorization through administrative groups and Internet Information Server (IIS)	6
Web management console	7
Configuring Collaboration Services	8
Configure secure data communication	8
The service account mailbox	9
Recipient Update Service (RUS)	10
Configure Active Directory access	10
Set up logging	10
Set up debug logging	11
Configure alerting	11
Configure synchronization partners	12
Check license information	13
Organizing objects for synchronization	14
Introduction	14
Creating a collection for publication	14
Publish collection wizard	15
Enabling or disabling a collection	19
Editing a collection	20
Subscribing to a collection	20
Modifying subscription properties	21
Unsubscribing from a collection	22
Removing a collection	22
Viewing collection settings and synchronization statistics	22
Advance security of stub objects created by the collection	23
Resolving synchronization conflicts	24
What can cause synchronization conflicts?	24
How can conflicts be resolved?	24
Automatic conflict resolution	25
Deletion	25
Matching	27
Pre-Matching	28
Manual conflict resolution	29
Applying approved objects	30
Matching rules viewer utility	31
Managing the synchronization process	32

Configuring the synchronization process	32
General settings	32
Synchronization settings for Active Directory objects	33
Free/Busy information synchronization settings	36
Calendar information synchronization settings	38
Using the Packet Queue as interim data storage	39
Suspended object storage	39
Manual operations	40
Synchronization	40
Re-publish collections	40
Configuration database clean-up	41
Transport management	42
Overview	42
Send	42
Receive	43
Scheduling	43
Limiting traffic rates	43
Outgoing traffic	44
Incoming traffic	44
Fine-tuning and maintenance	45
Active Directory synchronization performance	45
Resource usage	45
Memory usage	45
Disk usage	46
Resource usage statistics report	47
Statistics update frequency	47
Backup, restore, and troubleshooting	48
Collaboration Services automatic backup	48
Regular file backup	49
How to re-deploy Collaboration Services to another computer from a backup	50
Troubleshooting	50
Repair or uninstall	51
Appendix A: Customizing the format of synchronized data	52
displayName mapper	52
Suffix mapper	52
inetOrgPerson to Contact mapper	53
SMTP filter mapper	53
User to Contact mapper	53
Group to Contact mapper	54
Mapping groups from Exchange 2007	54
Attribute filter mapper	54

Appendix B: Collaboration Services events	55
Event format	55
Using event filters	56
Predefined filters	56
Custom filters	56
Index	58
About us	60
Technical support resources	60

Getting started with Collaboration Services

- [Introduction](#)
- [Security considerations](#)
- [Web management console](#)
- [Configuring Collaboration Services](#)

Introduction

Quest® Collaboration Services™ is a centralized collaboration application that enables you to organize collaboration between multiple forests. Collaboration Services securely synchronizes Active Directory and Exchange data (such as user objects, distribution lists, free/busy and calendar information) between isolated forests in a multi-forest/multi-org deployment of Active Directory.

Collaboration Services provides a consolidated view of all collaboration processes, gives easy access to object management functionality, and assists corporate IT administrators in synchronizing data between separated forests. It can also be used on a continuous basis to reduce the costs and complexities associated with managing a decentralized, multi-forest network.

It can be deployed in a single company (for example, between divisions and subsidiaries) as well as between separate companies (partners, consultants, and vendors). Collaboration Services allows for granular and selective data synchronization, including global address list (GAL), free/busy and calendar information in Exchange 2007, 2010, 2013, 2016 and 2019.

Security considerations

Authentication

The first security level for Quest Collaboration Services is authentication. The authentication type should be specified by an administrator according to the company's environment and policy.

Authorization through administrative groups and Internet Information Server (IIS)

The second security level is authorization through Collaboration Services administrative groups and IIS.

To restrict access to administration tasks, Collaboration Services automatically creates and uses two local groups on every Collaboration Services server:

- ACS Administrator

Members of this group have unrestricted access to the service configuration and can perform all service administration tasks.

- ACS Users

Members of this group can only view current service settings and statistics; they do not have permission to change any settings or to invoke re-publication.

When the service is installed on a member server or a workstation, local computer groups are created. The administrator can redefine group membership using the Computer Management tool.

If the service is installed on a domain controller, domain local groups are created. Users can be added to or removed from these groups using the Active Directory Users and Computers administrative tool.

i | **NOTE:** Members of the local Administrators group on the Collaboration Services server are also implicitly allowed to perform all synchronization service administration tasks, even if they are not members of the Collaboration Services Administrators group.

Web management console

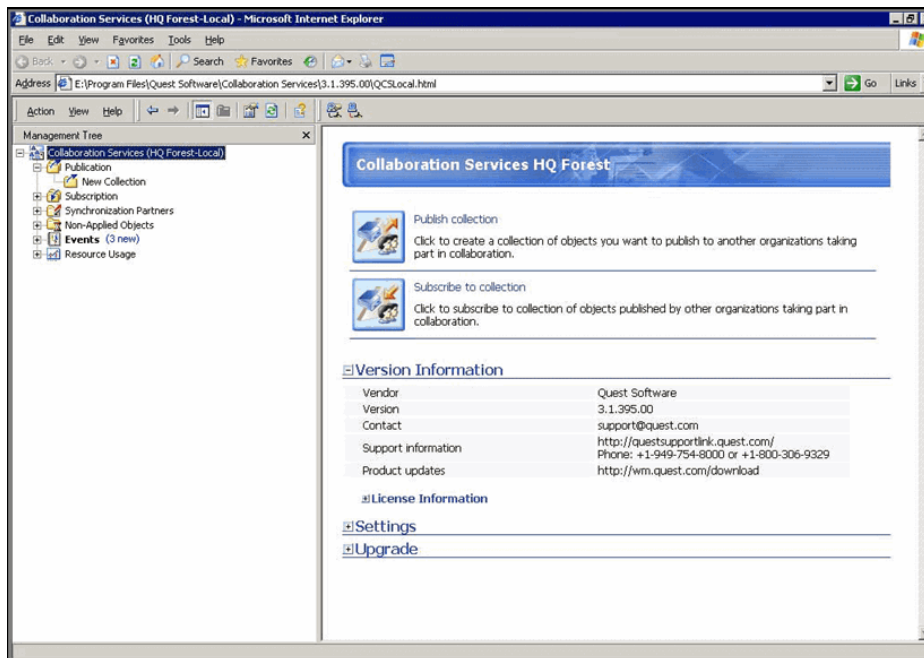
To perform administrative tasks, you need to connect to the Collaboration Services service using the web management console.

To open the Collaboration Services web management console

- Select **Start | Programs | Quest Software | Collaboration Services** on the computer where the Collaboration Services service is installed.

i | **NOTE:** For Windows Server 2012 and later versions, if you are not logged in as "Administrator", launch Internet Explorer with the option "Run as administrator".

Figure 1. The web management console main window



Configuring Collaboration Services

The Configure Branches Wizard, which allows for registering branches in the collaboration structure, automatically displays when you first open the user interface at the HQ forest. For details, see [Configuring branch forests](#) on page 12.

All the branches and the HQ forest are called synchronization partners.

Configure secure data communication

Collaboration Services provides secure communication between forests by using data encryption and signing. Collaboration Services uses public key files. Each public key file contains the following:

- Data encryption key
- SMTP address of the service mailbox for one forest

During Collaboration Services setup, you specify the following:


- File name and location where the HQ public key file should be saved
- Password to protect the file

All messages from the HQ forest are encrypted using the encryption key in the HQ public key file.

The public key file of each branch forest is required for the administrator of the HQ forest to register the branch forests in the collaboration structure. The Export Key Wizard, which allows for exporting the branch public key, automatically appears when you first open the user interface in a branch.

To change the settings for data encryption

- 1 Right-click the **Collaboration Services** node in the tree, select **Properties**, and click **General**.



If for some security reason your encryption key cannot be used any longer, you should generate a new key by clicking the Generate New button.

- NOTE:** After generating a new encryption key, you must generate a new public key file and distribute it to the other forests so they can update your forest's key using the Configure Branches or Configure HQ wizard. Until then, communication with other synchronization partners will not be possible.
- NOTE:** The HQ public key file is not created immediately; rather, it is automatically generated and saved upon setup completion, after the synchronization service is installed and started.

To export the existing public data encryption key to a public key file, click Export.

If running on Windows 2003 platform (or later), you can change the data encryption algorithm. By default, RC4 128-bit symmetric key encryption is used; you can choose to use an AES 256-bit encryption algorithm instead.

- NOTE:** All synchronization partners should use the same data encryption algorithm. If it is changed, then after all partners make the same change to the selection on this dialog box, all keys must be re-exported and re-registered, and then all collections must be re-published.

The service account mailbox

Instances of Collaboration Services installed in the HQ and branch forests communicate with each other using the Simple Mail Transfer Protocol (SMTP), which is a native protocol for Exchange. All communications between partners are performed through administrative mailboxes, one mailbox per forest. The data packets are compressed, grouped in regular email messages, and sent to other service account mailboxes.

- TIP:** The service account mailbox should be created before installing Collaboration Services. It is defined during the install procedure of Collaboration Services and should not be modified after unless necessary.

To view the service account mailbox information

- 1 Right-click the **Collaboration Services** node in the tree, select **Properties**, select **General**.
The name of Exchange server where the service mailbox resides is in Exchange server: field.

The mailbox name used by the service account is in the Mailbox: field.

- 2 Click **Test** to find out whether the specified mailbox exists on the Exchange server and whether it can be accessed by Collaboration Services.

i | **NOTE:** If you change the service mailbox, you must generate a new public key file and distribute it to the other forests so they can update your forest's key using the Configure Branches or Configure HQ wizard. Until then, communication with other synchronization partners will not be possible.

Recipient Update Service (RUS)

With Collaboration Services version 3.6 and later, the RUS is enabled by default if Exchange 2007 or higher is detected in the Exchange environment. If not selected, published objects will not appear in the Global Address List.

Configure Active Directory access

To configure Active Directory access settings

- 1 Right-click the **Collaboration Services** node in the tree, select **Properties**, select **General**, and **AD Options**.

The General account is used to read and update Active Directory data. By default, the service account is used as the general account, but you can use a different account if you want.

i | **NOTE:** The permissions required for that account are described in the Rights and Permissions section of the Quest Collaboration Services Deployment Guide. You can also find more information under Service Account Minimum Rights and Permissions in the Quest Collaboration Services Deployment Guide.

- 2 Use the **Test** button to test whether the newly specified account has sufficient permissions.

By default, Collaboration Services uses any available domain controller (DC) to read or update Active Directory data. To optimize the traffic and reduce the load on mission-critical DCs, you can configure the service to connect only to a particular domain controller (for example, you can choose to set up a dedicated DC for Collaboration Services to use).

For connection with a specific domain controller

- 1 Select the **Use preferred domain controllers** check box.
- 2 Select the domain, then click **Load**.
- 3 In the list of DCs loaded, select the check box next to the DC the service should use.

i | **NOTE:** If the selected DC is down, Collaboration Services will not be able to access the Active Directory data. A corresponding event will be logged to the event log in this case.

Set up logging

Collaboration Services service logs its activity to its own log, ACSAlerts.log, which is located in the ACSLOGS subfolder of the service installation directory, and (optionally) to the Windows System log.

To write events to the System log

- 1 In the management console, right-click **Collaboration Services** and select **Properties**.
- 2 Click **Logging | General**, and select whether to write events to the System log.

- 3 Adjust the level of events to be logged.

i **NOTE:** By default, all events are logged. It is recommended not to lower this setting soon after installation of Collaboration Services because log files are extremely helpful for troubleshooting. You can reduce the logging level after you are confident that the synchronization between partners is established correctly.

To clear the ACSAlerts.log

- Right-click **Collaboration Services** node in the tree, and use the corresponding shortcut menu command.

For Collaboration Services events details, see [Appendix B: Collaboration Services events](#) on page 55.

Set up debug logging

Collaboration Services service has debug logging when additional information is needed or when a support call is placed. Once turned on, a large number of files are generated which contain data from all areas within Collaboration Services. The logs are located in the following folder: \Quest Software\Collaboration Services\ACSLOGS\Debug.

The debug log setting will consume a lot of disk space. The logs are archived in the Archive folder found in the folder for the logs mentioned above. The logging will have an impact on the speed Collaboration Services processes both incoming and outgoing data.

To enable debug logging

- 1 In the management console, right-click **Collaboration Services** and select **Properties**.
- 2 Click **Logging | General**, and select the **Enable debug log** check box.

Configure alerting

To receive notifications about service activities, configure the required notification method on the Logging | Email tab, the Logging | Net send tab, or both. For example, you can configure email notifications to be sent when a specified percentage of free disk space is exceeded.

Table 1. Alerting

Notification method	Used to	Configuration required
email	Send events of the selected types and notifications to the specified email address.	<ol style="list-style-type: none"> 1. Select the mailbox to which the email messages will be sent. 2. Select whether you wish to receive Error and Warning events, Information events, or both. 3. Specify the subject for the email message.
Net Send	Send the events of the selected types and notifications using net send to the specified computers.	<ol style="list-style-type: none"> 1. Specify the NetBIOS names, DNS names, or IP addresses of the computers to which net send should deliver messages. Separate multiple entries by semicolons. 2. Make sure the Messenger service is running on the Collaboration Services server and on the recipient computer.

Configure synchronization partners

Using the Synchronization partners node of the management tree, you can:

- Configure new partners in the collaboration
- Track current synchronization partners and their statistics
- Manage the current collaboration structure and its forests

Configuring branch forests

The HQ forest administrator can manage the current collaboration structure in the following ways:

- Add or remove branch forests
- Include branch forests in the data synchronization process, or exclude them from it
- Update the data encryption key (if it was changed by a branch) and provide an update for the appropriate public key file

To manage the collaboration structure

- 1 Select the **Synchronization Partners** node in the management tree.
- 2 Right-click and select **Configure**.

The Configure Branches Wizard is started.

On the Specify Branches step, you can add partners to or remove partners from the list of registered branch forests:

- 3 To add a partner, click **Add** and browse for the key file (.AKF) of the branch you need. Supply the authentication password and click **OK**.
- 4 To specify whether each branch is currently enabled (involved in synchronization), select the **Enabled** check box.

Clearing the box temporarily excludes the branch from collaboration. All statistics and configuration information related to the branch will be preserved until the branch is enabled again.

i | **NOTE:** When you decide to re-enable the branch, be sure to request re-publication for all collections it is subscribed to; otherwise, all data included in updates sent during the time while the branch was disabled will be missing. The check box is not selected by default and should be left that way until the upgrade procedure is reviewed.

- 5 To specify whether upgrade packets should be sent to each branch, use the **Send Upgrades** check box.

i | **NOTE:** The check box is not selected by default and should be left that way until the upgrade procedure is reviewed.

For details about upgrading the product and about this option, refer to the Quest Collaboration Services Deployment Guide.

i | **NOTE:** Upgrade packages will not be sent to a disabled branch.

- 6 If any branch forest has changed its encryption key, you can update it by selecting the branch, clicking the **Update Key** button, and providing the appropriate public key file.
- 7 To remove a branch from the collaboration structure, select it in the list and click **Remove**.

i | **NOTE:** If you remove a branch, all statistics and settings related to it will be lost. If you want to temporarily exclude a branch forest from the collaboration, clear the Enable check box instead of removing the branch.

Alternatively, to add, edit, or remove synchronization partners, you can use the Branches page of the Collaboration Services Properties dialog box.

Branch forest administrator activity

Branch forest administrators can run the Configure HQ Wizard to edit, remove, or update HQ forest information.

Alternatively, the same settings can be changed using the HQ page of the Collaboration Services Properties dialog box.

Synchronization partner statistics

Collaboration partners involved in collaboration are displayed under the Synchronization Partners node in the management tree; the statistics report is shown in the right pane.

- To view statistics information for all collaboration partners, select the Synchronization Partners node of the management tree.
- To view statistics for a particular branch forest, expand this node and select the forest you are interested in.

The following information is displayed in the statistics pane of the Collaboration Services web interface:

Table 2. Statistics

Statistics section	Can be used for	Displays
Transport	Estimating transport service performance	Information about traffic bandwidth, current branch forest quotas, history load, last operation time, last operation result, last operation status, and transport status.
Sent messages Received messages	Estimating traffic size	Number and total size of service messages sent and received per specified period of time.
Quota	Verifying the bandwidth load and setting optimal transport options for the HQ forest	Information about currently assigned quotas.
Sent without acknowledgment	Tracking a potentially problematic branch	Number of messages for which no acknowledgment has been received.
Message queue load	Check the transport layer of the send queue. If the load is high, the partner is slow or not responding	The total size of files in "PacketHistory" directory.

Check license information

The License page allows you to view the data on the license currently being used by this forest. (You can also view this information about the license in the statistics pane when the Collaboration Services node of the management tree is selected.)

To view the license data

- 1 In the management console, right-click **Collaboration Services** and select **Properties**.
- 2 Click **General**, and select **License**.

To change the license

- Click **Browse** and select a new license file.

Organizing objects for synchronization

- [Introduction](#)
- [Creating a collection for publication](#)
- [Enabling or disabling a collection](#)
- [Editing a collection](#)
- [Subscribing to a collection](#)
- [Unsubscribing from a collection](#)
- [Removing a collection](#)
- [Viewing collection settings and synchronization statistics](#)
- [Advance security of stub objects created by the collection](#)

Introduction

The administrator of each forest can select the objects and attributes to publish to other synchronization partners. The objects to be synchronized are grouped into collections—sets of objects and their attributes. Administrators can either define the objects to be synchronized on a per-container (organizational unit) basis, or explicitly add groups, users, and contacts to collections. This information then becomes available to other forests. Administrators of synchronization partners can see new collections of objects and can choose whether to subscribe to them.

Publication settings can vary for different collections of objects. For example, you might choose to publish the mobile phone numbers of administrators so they can be contacted in case of emergency, while keeping the mobile phone numbers of other users (such as top executives) unlisted outside the organization. Predefined "Minimum" and "All" sets of attributes to be published are provided, and administrators can create their own predefined sets (presets).

Creating a collection for publication

Collaboration Services offers numerous methods for organizing objects into collections. For example, you can do any combination of the following:

- Publish a whole domain or some top-level OUs and exclude objects not to be published explicitly, through containers or using LDAP filters.
- Create a group with users to be published, and publish just this group with the Synchronize members option enabled. Later, if you want to add additional users to the synchronization, you can join them to the group and they will be published automatically.
- Select users and groups to be published explicitly.
- Import a list of objects to be published from a plain text file.

Publish collection wizard

The Publish Collection Wizard helps you create a new collection of objects and publish it. As soon as a collection is published, the administrators of other branch forests (who have appropriate permissions) can subscribe to the collection.

To run the wizard

- Right-click the **Publication** node in the management tree and select **Publish Collection**.

Step 1: Specify collection name, description, and synchronization type

In the first step of the wizard:

- 1 Specify the collection's name and description.

The name and description should be as informative as possible; this will be the only collection information available to the administrators of other forests until they subscribe to the collection.

- 2 Select what information will be published and synchronized:

- Active Directory objects (for GAL synchronization)
- Free/busy information
- Calendar information

i | **NOTE:** If you choose not to synchronize Active Directory objects with Collaboration Services, use some other directory object synchronization tool to synchronize Active Directory objects between forests, such as Microsoft Identity Integration Server (MIIS).

Step 2: Select objects to publish

Select the objects you want to include in the collection. Consider that only mail-enabled or mailbox-enabled objects can be published.

You can populate a collection in the following ways:

- Add a container; all objects in that container will be included in the collection.
- Add objects (user, groups, and contacts).
- Import a list of objects from a file. This file should list the distinguished names (DNs) of the objects, one DN per line.

To add a container

- 1 Click **Add Container**.
- 2 Click **Load** to retrieve the Active Directory containers tree.
- 3 Select the containers you need, and click **Add**.

To add objects

- 1 Click **Add Objects**.
- 2 To display the objects you need, use filtering:
 - Domain—displays objects of the selected domain only.
 - Additional filter—specifies an additional LDAP filter to reduce the scope of displayed objects. For example, to display only users from the list of Active Directory objects, specify the following LDAP filter: (objectClass=user).

- Maximum number of displayed objects—specifies the maximum number of objects to display.

NOTE: If you add a group to a collection, group members are not synchronized automatically by default. If you want them to be synchronized, select the added group, click **Settings**, and then choose the **Synchronize members** option in the **Settings** dialog box.

Nested groups are not included in a collection. For example, suppose Group 2 is a member of Group 1. If Group 1 is included in the collection, Group 2 will not be published. To publish Group 2, you must include it in the collection explicitly.


Configuring synchronization settings

From the **Select Objects** dialog box, you can configure synchronization options by clicking **Settings**. For example, for Active Directory containers, you can select object classes to be synchronized.

It is recommended that you use this dialog box to configure synchronization of group members (they are not synchronized automatically by default).

To synchronize group members

- Select the group you have added, click **Settings**, and then choose the **Synchronize members** option.



Step 3: Specify additional LDAP filters

If you have added containers to be published, you may need to specify additional LDAP filters for each class of objects included in these containers.

NOTE: Creating LDAP filters is a fairly complex task that should be performed by advanced users only. LDAP filters must be RFC2254-compliant.

For example, to exclude from publication all users who either belong to the HR department of your company or are members of the HRGroup group, apply the following filter to the User class:

```
(!(|(department=HR)(memberOf=CN=HRGroup,OU=Groups,DC=acme,DC=com)))
```

Step 4: Exclude well-known objects

Specify if any built-in accounts, containers, or single objects should be excluded from the collection.

- To exclude built-in groups or accounts (such as Domain Admins and Guest), select Exclude well-known users and groups.
- To exclude a container (and all objects inside), select the Exclude the following objects check box and click Add Containers. All objects in the containers you select will be excluded from the collection.
- To exclude particular objects, click the Add Objects button, and select the objects from the list.

Step 5: Select object attributes to be synchronized

Specify which attributes should be synchronized for each object class. You can do the following:

- Select attributes manually from the list
- Use predefined sets of attributes: All Attributes set or Minimum set
- Save the set of selected attributes for the future use as a preset
- The Minimum set includes the only attributes required to create a stub object in another forest.

To see what attributes comprise the minimum set for a class

- Expand the node for that class and click **Minimum**.

i | **NOTE:** All object attributes in the Minimum set are mandatory, so you are not allowed to clear the check boxes next to them.

If any object attribute included in the Minimum set is not valid, the object cannot be published.

To save a set of selected attributes

- Click **Save Preset**. If you need to publish an object with the set of attributes saved as a preset, click the **Load Preset** button.

To use an extended Active Directory schema

- Click **Add Attributes** and select the necessary attributes from the list.

i | **NOTE:** To synchronize extended schema attributes to other forests, their schemas must be extended with the same attributes as well.

Step 6: Select calendar attributes

On the Calendar Details page, specify which appointment attributes should be published and synchronized.

i | **NOTE:** This page is skipped if Calendar information was not selected on the Step 1 of the wizard.

By default, if the **Synchronize appointment details** check box is cleared, Collaboration Services publishes to other synchronization partners only the time to start and end the appointment.

Select the Synchronize appointment details check box to allow publication of the following additional parameters for each appointment:

- Time zone
- Location
- Recurrence attributes
- Appointment subject

i | **NOTE:** The Location and Subject appointment details will become available in Free/Busy information.

If you select the Synchronize appointment details option, then the Reviewer permission will be set for the Calendar folder (for Default on the subscriber side), so any user from the subscribed organization will be able to view synchronized calendar details.

Step 7: Apply mappers

Sometimes the publishing or subscribing administrators might want to change the format or appearance of the data they publish or subscribe to. To do that, you can use mappers. A mapper is a plug-in component that changes the appearance of data that you publish or subscribe to. For example, a mapper can be used to set the Active Directory objects' attributes values to those specified in the .ini file.

In this step of the Publish Collection Wizard, you can apply mappers to the published objects.

For more information about what mappers are available, see [Appendix A: Customizing the format of synchronized data](#) on page 52.

To enable the option

- 1 Select the **Apply mappers to synchronized objects** check box.

You can perform the following actions:

- a To add mappers, click **Add**. Select the mapper to be applied from the list of Available mappers and click **Add**. The mapper appears in the Selected mappers list. Repeat these actions to add more mappers if needed.
- b To organize the list of mappers use the Up and Down buttons.

i | **NOTE:** Make sure that frequently used mappers are at the top of the list because the topmost mappers are processed by Collaboration Services first.

- c You can edit properties of the added mappers or remove the mappers from the list.

Step 8: Allow for subscription and synchronization

In this step you can:

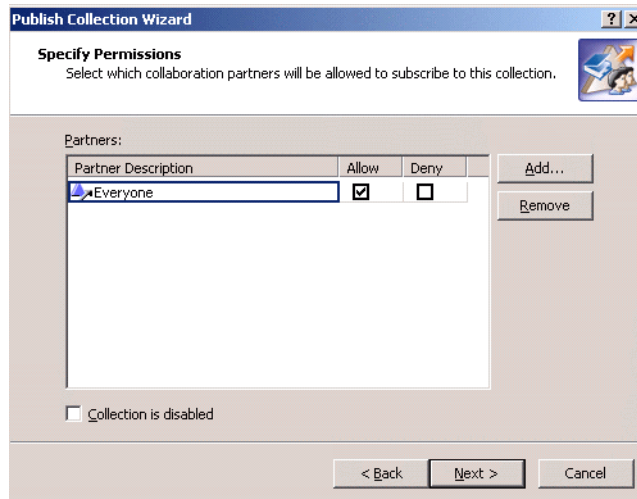
- Allow other synchronization partners to subscribe to the collection you are creating or prohibit partners from subscribing.
- Specify whether synchronization of collection members should start immediately after the collection is published (after you finish the wizard).

To allow or prohibit subscription

- 1 Click the **Add** button.
- 2 Check or clear the **Allow** or **Deny** check box for a partner in the list.

i | **NOTE:** Deny always overrides Allow. Thus, Deny selected for Everyone overrides all explicitly set Allow permissions.

Because the HQ forest operates as the service data packet distributor, it is allowed to subscribe to all collections in the collaboration structure by default.



If you do not want collection members to be involved in synchronization immediately after the collection is published, select the Collection is disabled check box. For details on enabling or disabling collections, see [Enabling or disabling a collection](#) on page 19.

Step 9: Apply settings

Wait until the collection is created and its settings are saved to the configuration database.

Collaboration Services will notify allowed synchronization partners that the collection is available for subscription. As soon as a collection is published, the administrators of other branch forests who have appropriate permissions will be able to subscribe to the collection.

- NOTE:** If a collection was created as enabled, then as soon as the synchronization partners subscribe to it, the collection's data will start being synchronized to the other forests.

Enabling or disabling a collection

Enabling a collection means all objects included in it should be immediately synchronized by Collaboration Services. By default, a collection is created as enabled and synchronization starts right after the Publication Wizard is completed.

However, you may need to create a collection without enabling it (that is, synchronization does not start upon publication), and other branches subscribe to it as required in advance. Then on the day when the synchronization needs to be started, all you need to do is enable the collection.

To enable a collection during creation

- Ensure that the **Collection is disabled** check box on the Set Permissions step of the Publication Wizard is cleared.

To enable an existing collection

- Right-click it in the management tree and clear the **Disabled** check box.

To disable a collection during creation

- Select the **Collection is disabled** check box on the Set Permissions step of the Publication Wizard.

To disable an existing collection

- Select the **Disabled** check box on its shortcut menu.

Editing a collection

After a collection is published, you may need to add objects to it, remove objects from it, or change some of its properties.

To update a collection

- Run the Edit Collection Wizard by selecting the **Edit Collection** icon from collection's shortcut menu.



Its steps are similar to those of Publish Collection Wizard. For details, see [Publish collection wizard](#) on page 15.

To access a collection's properties

- Right-click the collection and select **Properties**.

i | **NOTE:** On the Advanced tab, you can select whether to Apply approved objects only. For details, refer to [What can cause synchronization conflicts?](#) on page 24.

After you edit a collection, the updates are automatically sent to the HQ forest and distributed to all branch subscribers; then the modified collection is automatically re-republished.

Subscribing to a collection

To receive collection data, synchronization partners need to subscribe to collections. As soon as a forest subscribes to a collection, the Collaboration Services instance creates the following items for each member of the collection in this forest:

- Stub objects in the Collaboration Service OU created during the install. The stubs are configured to allow them to be published in the GAL. Any mail sent to these stubs will be redirected to the original object's email address.
- Free/busy information.
- Calendar information.

Use the Subscribe Wizard to subscribe to the collections published by the HQ and branch forests.

To subscribe to a collection

- 1 Select the **Subscription** node in the management tree, and select **Action | Subscribe**.
Alternatively, you can use the toolbar button or the Subscribe shortcut menu command to start the wizard.
- 2 Select the collection you want to subscribe to from the list.
- 3 You can change the format or appearance of the objects in a collection you have subscribed to by applying mappers.

i | **NOTE:** Before applying a mapper to a collection you are subscribing to, check with the collection originator to make sure this mapper was not already applied to the collection. Applying the same mapper twice is not necessary, and sometimes it can lead to undesired results, such as a suffix being added twice.


- a Select the collection you need and click the **Properties** button.
 - b Select the **Apply mappers to synchronized objects** check box.
 - c Click **Add** and select the mappers you want from the list of available mappers. For more information about mappers, see [Appendix A: Customizing the format of synchronized data](#) on page 52.
- 4 Wait for the wizard to apply your changes.
 - 5 Review the summary log of the subscription process, and click **Finish**.

Modifying subscription properties

To configure the collection to which you have subscribed

- 1 Right-click the subscription and open its properties.
- 2 On the Mappers tab, configure the mappers you want to apply. For details, see [Appendix A: Customizing the format of synchronized data](#) on page 52.
- 3 On the Advanced tab, select whether to Apply approved objects only. For details, see [Manual conflict resolution](#) on page 29.
- 4 If you subscribe to a collection for calendar information synchronization, you should also select the **Always create mailbox** check box and specify the Exchange server and mailbox store that mailbox will be created on.

In this case, a mail-enabled object on the subscriber side will be created, and the Calendar information will be presented as the corresponding Free/Busy.



Unsubscribing from a collection

- NOTE:** In Collaboration Services 3.6.1 and later the Always create mailbox check box is set by default for Calendar collection subscriptions. Note that the first Exchange server and the first mailbox store are selected automatically from the Exchange environment. If you do not want to use either of these then change them by selecting the drop down icon.
The Exchange server and Mailbox store can be changed later if needed. Once the server and store are changed, it may take up to 24 hours for the move to be completed by Collaboration Services.

To unsubscribe from a collection

- Select the **Collection** node in the management tree and select **Unsubscribe**.

- NOTE:** If you unsubscribe from a collection, the stub objects created in your forest for that collection's members are hidden from the GAL in your forest. The stub objects are not deleted from the Active Directory container they were created in. The stub objects will be automatically deleted after the time interval specified by the Retention time for removed objects parameter. The default period is 30 days.

To change the retention period for these objects

- In the management console, right-click **Collaboration Services** and select **Properties**.
- Select the **Synchronization | Active Directory | General** page.
- Modify the Retention time for removed objects parameter.

- TIP:** By configuring the Retention time for removed objects parameter, you can delay deleting the stubs for a specified number of days; this will help improve overall collaboration performance and decrease the load on the production environment because if the object is added to another collection during this period, the stub objects will be re-used.

Removing a collection

To remove a collection you published

- Right-click the collection name in the management tree and select **Remove**.

- NOTE:** A collection can be removed only by its publisher.

When a collection is removed by publisher, the following is performed:

- All subscribing partners get a notification.
- The stub objects created by Collaboration Services are immediately removed from the Global Address List in all subscribing forests.
- The stub objects are not deleted from the Active Directory container they were created in. The stub objects will be automatically deleted after the time interval specified by the Retention time for removed objects parameter. The default period is 30 days.

Viewing collection settings and synchronization statistics

The collection settings and synchronization statistics are displayed in the right pane when you select a collection in the management tree. To see the overall statistics for all collections currently published or subscribed to by your forest, select the Publication or Subscription nodes.

The report shows, in particular:

- General information about the collection: collection name, description, publisher, creation date.
- Collection membership information: number of objects included in the collection and membership scan statistics.
- Statistics on Active Directory objects synchronization (if it is configured for the collection): the number of objects added, modified or deleted per week, month, or year; synchronization status and state; and the date and time of the last synchronization.
- Statistics on Calendar and Free/busy information synchronization (if it is configured for the collection): the number of objects for which this information was synchronized, the number of updates sent, the synchronization status and state, and the date and time of the last synchronization. For more details on synchronization, see [Managing the synchronization process](#) on page 32.

Advance security of stub objects created by the collection

By default, Collaboration Services creates stub objects with the Password Not Required (PASSWD_NOTREQD) option enabled. This may violate your organization's security policy and may cause audit software to report a password issue. To resolve this, you will need to update the attribute. In addition, the password should be set to a random value so that the stubs cannot be used to exploit the environment through unauthorized access.

- Default setting: userAccountControl attribute is set to 546 (0x222): NORMAL_ACCOUNT = 512; ACCOUNTDISABLE = 2; PASSWD_NOTREQD = 32
- Updated setting: userAccountControl attribute set to 514 (0x202): NORMAL_ACCOUNT = 512; ACCOUNTDISABLE = 2

In Collaboration Services, you can use a registry key to clear the PASSWD_NOTREQD option and set a random password for each stub object as it is created. If the stub already exists, it will clear the option and set the password on the next republish of the collection.

To enable the random generation of a password

- Create a string value, REG_SZ, called StubAccountPassword in the registry under HKLM\SOFTWARE\Wow6432Node\Aelita\AelitaCollaborationServices

– OR –

HKLM\SOFTWARE\Aelita\AelitaCollaborationServices for 32 bit systems

The StubAccountPassword should comply with the password complexity policy. Collaboration Services will add a random string and digits to it and generate a password. For example, Z1f9j7Yab\$!.

As a best practice, the password template:

- should not contain common names or number combinations.
- should not contain meaningful word that will fail the password policy.
- can contain rules such as; Cannot include firstname, lastname, birthday, and/or userID.

Resolving synchronization conflicts

- [What can cause synchronization conflicts?](#)
- [How can conflicts be resolved?](#)
- [Automatic conflict resolution](#)
- [Manual conflict resolution](#)
- [Matching rules viewer utility](#)

What can cause synchronization conflicts?

A conflict occurs when two or more objects from different forests have the same email address. Most conflicts are due to manually-created contacts pointing to users in other forests.

While creating a stub object in the target forest, Collaboration Services populates the proxyAddresses attribute and the mail attribute of the stub with the appropriate attribute values of the corresponding source object. However, if existing contacts or other objects already have their properties populated with the same values, a valid stub object cannot be created.

Therefore, before creating a stub, Collaboration Services checks whether any of the proxy addresses or mail attributes of the existing forest objects are already populated with the same values as the attribute planned for a stub object. If a conflict is detected, Collaboration Services cannot create a stub object until the conflict is resolved.

- i** | **NOTE:** Collaboration Services is not capable of automatically resolving conflicts when there is more than one conflicting object; resolving conflicts involving more objects requires administrator intervention. Information about every conflict found during synchronization is recorded in the Collaboration Services log.

How can conflicts be resolved?

There are several different ways to resolve a conflict, object deletion, object matching, and pre-matching.

- Object Deletion

The recommended method of resolving conflicts is to delete the conflicting object and create a stub object that inherits deleted object's properties. This is the simplest way to resolve conflicts, and it is the only option available for resolving conflicts between objects of different classes (for example, user and contact).

- Object Matching

If both the conflicting object and the published object are of the same class (for example, both are users), then you can choose to match the published object into the existing one. In this case, the existing object is used as a stub for the time of synchronization. One significant difference between such a "temp stub"

object and a regular stub object is that the user object's account state remains unchanged. For example, if the account was enabled, it remains enabled after Collaboration Services makes a stub of it.

- **NOTE:** Matching needs to be used for deployment scenarios where some users have active user accounts in multiple forests. With matching, you can use active user accounts as stubs, so the users will still be able to log on in multiple forests with Collaboration Services deployed.

You can define attribute synchronization rules for matching published objects with existing objects. The rules allow you to specify which attributes of the matched objects should be overwritten by the published object's attributes and which should be left intact.

- **NOTE:** Attribute synchronization is useful if the user account is updated by the provisioning system in one of the forests and those changes need to be propagated to the same user's accounts in other forests.

- Pre-Matching

This method instructs Collaboration Services to match published objects into specified objects automatically in the selected OUs.

Forest administrators can configure Collaboration Services to handle conflicts using the following options: deletion, matching, and pre-matching. For that, the automation rules should be configured, and then the synchronization service will attempt to automatically resolve every conflict it encounters according to these rules.

Automatic conflict resolution


To set up automatic conflict resolution

- Open the Collaboration Services Properties dialog box, select **Synchronization | Automation**, and then select the conflict resolution method you need.

Deletion

Use this option to make Collaboration Services automatically delete conflicting objects in the selected OUs before creating the stub objects.

Figure 2. Deleting conflicting objects



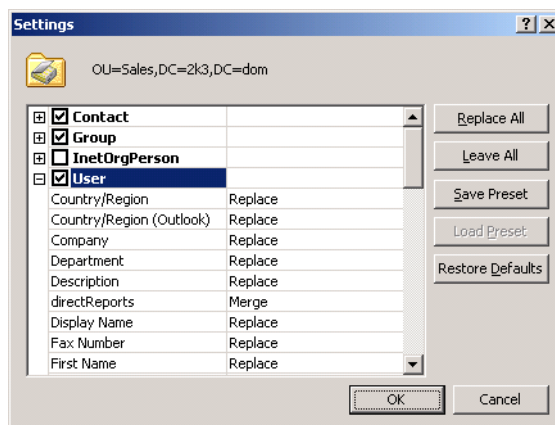
- NOTE:** This configuration should be implemented with care, since the objects will be deleted from the selected OUs automatically without further confirmations.
- NOTE:** Ensure that the Collaboration Services service account has full access to any containers which are selected. If not, the conflict resolution will not be completed and errors will be generated.

To configure the deletion

- 1 Select the **Allow deletion of conflicting objects from the following containers:** check box.
- 2 Click **Add**, and then **Load the list of containers from Active Directory**; select the containers from which conflicting objects will be deleted.

You can also import a list of containers from a text file. Click the **Import** button and choose the file containing the list of containers. The file should list the distinguished names (DNs) of the containers, one DN per line.

- 3 For each container, specify how the attributes of deleted objects should be handled. To configure attribute handling:
 - a Select the container and click **Settings**.
 - b Select the attribute synchronization rule for each attribute.



You can choose from among the following rules:

- Leave—Preserve the attribute of the deleted object.

- Replace—Overwrite the attribute with the published object's attribute (you can click Replace All if this rule should be applied to all selected attributes).
- Merge—Merge the attributes of the deleted and published objects (this option is available for multi-valued attributes only).

i **NOTE:** A stub object created after deletion of the conflicting object will inherit the following attributes from the conflicting object: LegacyExchangeDN, X500 addresses. Preserving these attributes prevents users from getting non-delivery receipts (NDRs) when replying to older email from a deleted object.

- c You can also do the following:
 - Save the configuration for future use by clicking Save Preset.
 - Load a preset.
 - Restore default settings.
- d Click **OK** to apply your changes.

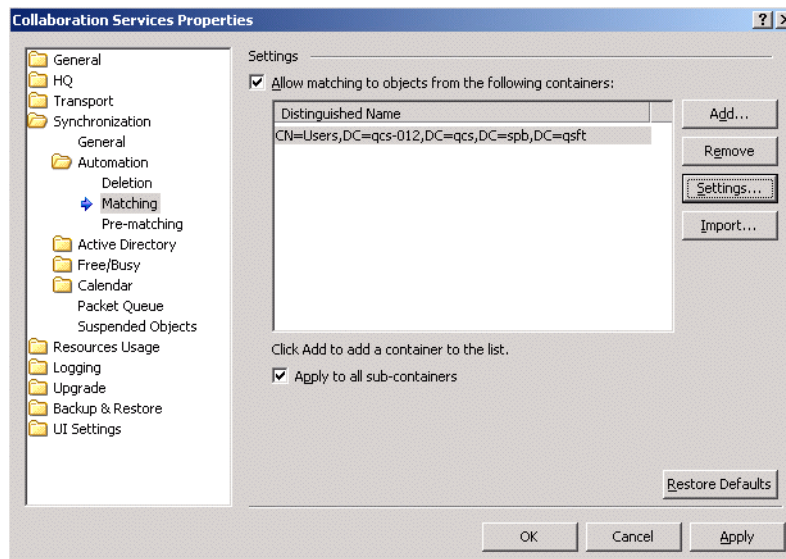
i **NOTE:** With the default settings, a stub object created by matching into the conflicting object will replace the targetAddress attribute value of the target object with the redirector address. If you choose to leave the targetAddress attribute of the target object intact, mail redirection might not function properly.

- 4 To delete the conflicting objects not only in the selected OUs but in the sub-containers as well, select the **Apply to all sub-containers** check box. In that case, parent OU settings will be propagated to child OUs settings.

Matching

Use this method to make Collaboration Services automatically match published objects into conflicting objects in the selected OUs, thus resolving the conflict.

Figure 3. Matching objects



To configure matching

- 1 Select the **Allow matching to objects from the following containers:** check box. Follow the procedure for configuring object attributes found within the section [Deletion](#) on page 25.

You can select containers and specify how the attributes of objects you match into are handled as described in Deletion section, and the same rules are available for the attributes of target objects.

NOTE: Passwords cannot be synchronized across forests due to the secure nature of Collaboration Services. Therefore, passwords need to be managed separately in each forest.

NOTE: With the default settings, a stub object created by matching into the conflicting object will replace the targetAddress attribute value of the target object with the redirector address. If you choose to leave targetAddress attribute of the target object intact, mail redirection might not function properly.

- 2 Select the **Apply to all sub-containers** check box to specify that the conflicting objects should be deleted not only in the selected OUs, but in the sub-containers also. In that case, parent OU settings will be propagated to child OUs settings.

Pre-Matching

On the Pre-matching page you can configure Collaboration Services to match published objects into specified objects automatically in the selected OUs. To match published objects into existing ones, pre-matching rules are used. A rule includes:

- Email address of the published object
- Distinguished names (DNs) of the object to match into


Whenever an object with a specified email is synchronized, it will be matched into the object whose DN is stated in the rule.

NOTE: Ensure that the Collaboration Services service account has full access to any containers which are selected. If not, the conflict resolution will not be completed and errors will be generated.

NOTE: The incoming object and the object you are matching into must be of the same object class; in most cases, they are user objects.

To configure pre-matching


- Select the **Use custom matching for the following objects:** check box to enable the option.



Creating a New Rule and the Pre-Matching Wizard

To create a new rule

- 1 Click the **Add** button and follow the steps of the Pre-Matching Wizard.
- 2 On the Welcome page, click **Next**.
- 3 On the Email Address page, click **Browse** to specify the email address of published object using the Select recipient dialog box and click **OK**.



- 4 On the Select Object page, specify the object you want to match into and click **Next**.
- 5 On the Attribute Synchronization page, choose how to handle the attributes of objects you match into. Select the **Replace**, **Leave**, or **Merge** (for multi-valued attributes only) attribute synchronization rule for each attribute of the object to match into and click **Next**.
- 6 In the message box that appears, click **Yes** if you want to proceed with creating the new pre-matching rule.
- 7 View the Summary and click **Finish**.

For each pre-matching rule you add, you can perform the following actions:

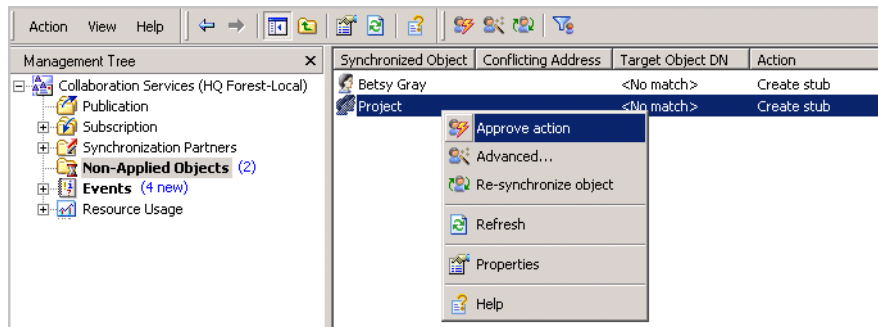
- Click **Remove** to delete the rule from the list.
- Click **Settings** to change how to handle matching object attributes.
- Click **Import** to import a list of pre-matching rules from a text file. The file should contain a list of pairs, one pair per line. Each pair must include the email address of the published object and the distinguished name (DN) of the object to match into, separated by a space.
- Click **Detect** to scan Active Directory and update the list of rules if any changes were made to Active Directory (for example, objects were moved).

Manual conflict resolution

If no automation rules are specified, or no rules can be used to resolve a particular conflict, the synchronization service marks the published object as non-applied.

In other words, non-applied objects are the ones whose changes will not be automatically applied to your forest. These objects are displayed in the Non-Applied Objects node in the management tree.

Figure 4. Non-Applied Objects



For the listed objects, you have to resolve conflicts manually, either by using the method Collaboration Services suggests or by using your own method (custom action).

To use the conflict resolution method suggested by Collaboration Services (for example, Create stub), you have to approve it manually

- Simply select the objects in the list, right-click the selection, and click **Approve Action**.

To resolve a conflict using a custom action

- 1 Right-click the object and choose **Advanced**.

This starts the Conflict Resolution wizard, which allows you to choose a custom action for resolving the conflict. In this wizard, you specify:

- The email address of the published object
- The DN of object you want to match into
- How the attributes of objects you match into are to be handled

Custom actions created with this wizard are stored in the configuration as pre-matching rules; if the collection is republished, the custom action is re-applied automatically.

- 2 After you resolve the conflict, you should republish the objects by clicking **Re-synchronize object** in the shortcut menu.

The objects will be re-published and applied.

Applying approved objects

You may want to manually resolve all conflicts for the objects your forest is subscribed to (that is, to approve objects applied to your forest manually), either one by one or all at once.

i | **NOTE:** Ensure that the Collaboration Services service account has full access to any containers which are selected. If not, the conflict resolution will not be completed and errors will be generated.

To manually approve objects

- 1 Right-click the collection of objects you need and select **Properties**.
- 2 Go to the **Advanced** tab and select the **Apply approved objects only** check box.

All collection objects will appear in the Non-Applied Objects list, where you can choose which objects you allow to be applied to your forest.

Matching rules viewer utility

If you have made a mistake when manually matching conflicts or you think the automatic conflict management may have been configured incorrectly, then there is a utility available that will let you look at the matches which have been made and correct any issues.

This utility, MatchingRulesViewer.exe, can be found in the Collaboration Services folder:

C:\Program Files (x86)\Quest Software\Collaboration Services\x.x.x.xxx

where x.x.x.xxx is the version number of Collaboration Services you have installed.

i | **NOTE:** Care should be exercised when using this utility as it could result in additional issues with conflict resolution.

Managing the synchronization process

- [Configuring the synchronization process](#)
- [Using the Packet Queue as interim data storage](#)
- [Suspended object storage](#)
- [Manual operations](#)

Configuring the synchronization process

The main task for Collaboration Services is keeping Active Directory data, calendar, and free/busy information synchronized between all synchronization partners so that users have access to up-to-date information about all objects participating in the collaboration. Use the synchronization settings to:

- Fine-tune the interaction between Active Directory, Exchange, and your Collaboration Services instance.
- Optimize the connectivity between the synchronization partners.

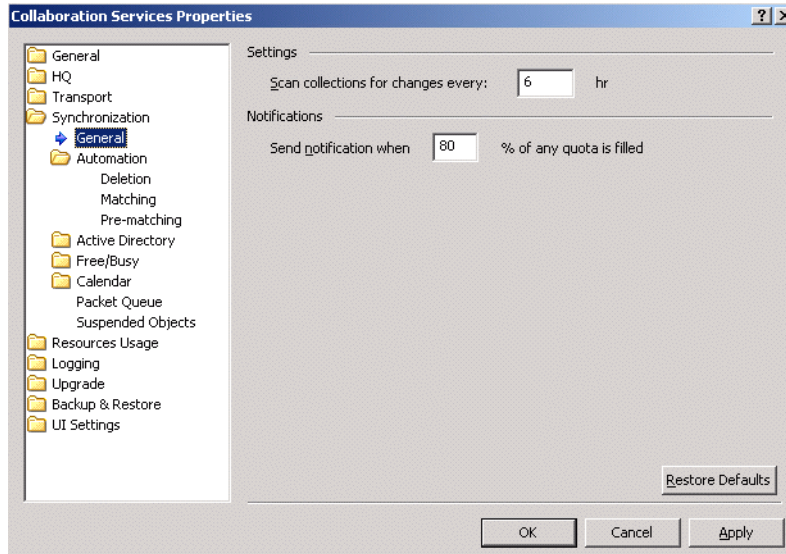
To access the synchronization settings for Collaboration Services service in your branch

- 1 Right-click the **Collaboration Services** node in the management tree, and select **Properties**.
- 2 Under the Synchronization node, select the required settings.

General settings

The General page of the Synchronization node contains common controls for all types of synchronization.

Figure 5. Synchronization general options



Collections re-scan frequency

The regular scan (which uses Active Directory's DirSync control) registers the majority of changes made to Active Directory objects. However, some changes cannot be registered by regular means, such as when an OU with collection members is moved to another domain or published objects are deleted.

Thus, a more thorough collections re-scan should be initiated periodically to register these types of changes; use the Scan collection for changes every option to specify how often the detailed re-scan should take place. The default period is 6 hours.

NOTE: This type of scan uses significantly more resources than a normal scan, so it is recommended not to set this parameter to a low value.

Alerting and notification

In the Notifications section, you can specify what percentage of any quota being filled should trigger notification of a forest administrator. The notification will state that the specified percentage of a quota is filled. The default value is 80%. For more information about quotas, see [Synchronization settings for Active Directory objects](#) on page 33.


Synchronization settings for Active Directory objects

The Active Directory object synchronization process is as follows:

- 1 Collaboration Services periodically checks (scans) whether any changes were made to Active Directory objects in your forest.
- 2 When Collaboration Services finds that the properties of one or more published objects have changed, it creates update packages and sends them to the HQ for distribution to all branch forests that subscribe to the collections those objects belong to.

Use the Active Directory tab of the Synchronization node to adjust the synchronization settings for Active Directory objects.

Figure 6. Active Directory synchronization options



General Settings

On the General page, you can set the time intervals for scanning Active Directory for changes and applying the changes received from collaboration partners.

Publication settings

At the time interval specified in the Query AD for changes every field, Collaboration Services checks whether any changes were made to Active Directory objects in your forest.

The sleep interval between the collection update sessions is specified in the Processing collections interval.

Subscription settings

At the interval specified in the Update AD every field, Collaboration Services applies incoming updates to Active Directory in the forest.

When an original object is removed from a collection you subscribe to, the deletion is automatically synchronized to the subscribing forests and the corresponding objects are hidden from the Global Address List. However, the stub objects themselves are not deleted immediately. Using the Retention time for removed objects option, you can delay deleting the stubs for a specified number of days; if the object is added to another collection during this period, the stub object will be re-used.

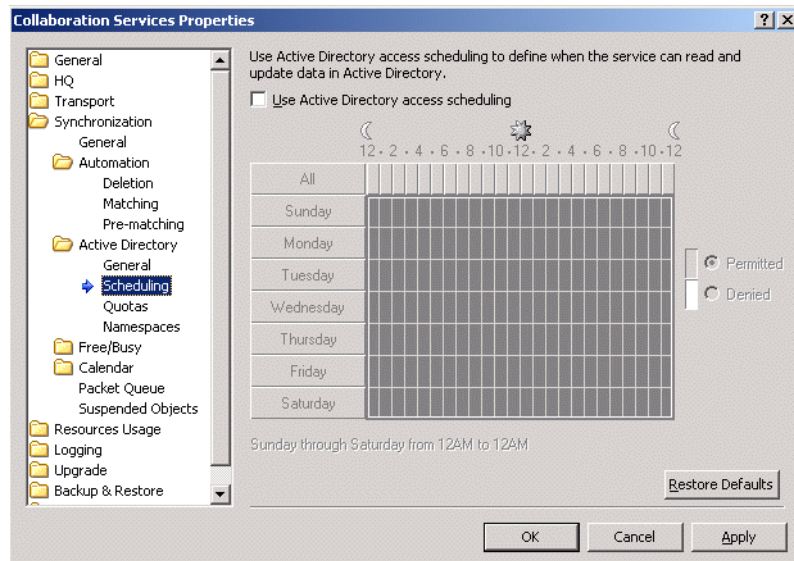
Scheduling access to Active Directory

To avoid excessive load on the DCs during production hours, schedule Collaboration Services access to Active Directory for the least busy time periods.

i | **NOTE:** Active Directory access schedules must be thoroughly coordinated, especially when some of the collaboration partners are located in different time zones; otherwise, synchronization performance may be poor.

To schedule Active Directory access

- 1 On the Scheduling page, select the **Use Active Directory access scheduling** check box.



- 2 Specify the time intervals when the service is allowed to read and update data in Active Directory.

Limiting the number of Active Directory objects to synchronize

When configuring Active Directory objects synchronization, you can specify the following limits on the Quotas page:

- Overall maximum number of synchronized Active Directory objects
- Maximum number of AD objects allowed to be synchronized daily

The maximum number of synchronized directory objects should be coordinated with the following:

- Hardware resources of the Collaboration Services server, and Active Directory and Exchange servers load in the HQ forest and in the branch forest.
- Number of the collaboration branches.
- Number of objects to be published in the collaboration structure.

i **NOTE:** When the specified maximum number of Active Directory objects is reached, no new objects will be added to Active Directory synchronization. In this case, it is recommended either to increase the specified value or to unsubscribe from some collections.

Namespace usage

For each original object, Collaboration Services creates a stub object in the forests that subscribe to the collection. The mail received by this stub should be redirected to the original object's email address.

Since mailbox-enabled object can have more than one proxy address, Collaboration Services should select which of them is to be used for mail redirection from stub object.

To choose the redirection address, Collaboration Services does the following:

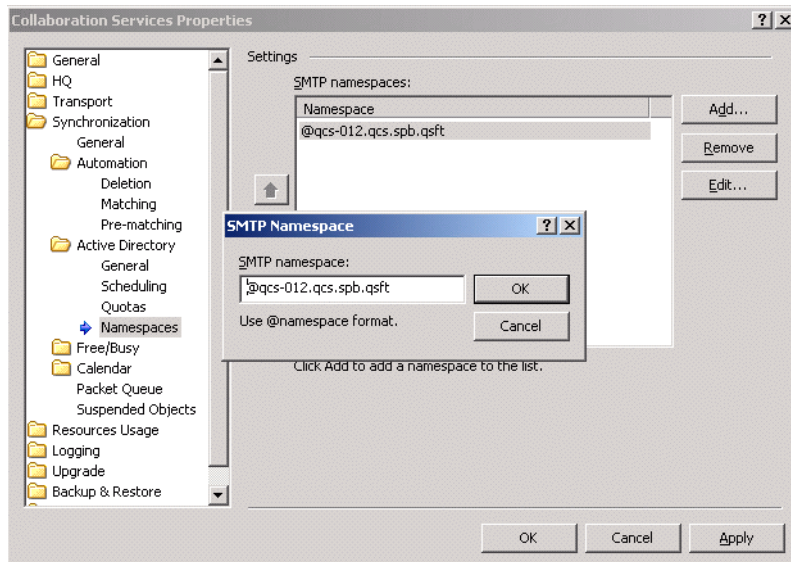
- 1 If the object's `targetAddress` attribute is populated, that will be used.
- 2 If the `targetAddress` attribute is not available, the primary SMTP address is used.

- 3 If neither address is available, redirection address will be set to the first matching proxy address found in the Namespaces page of the Collaboration Services Properties dialog box. The namespaces listed there are in priority order.

NOTE: If none of the retrieved addresses is encountered in the specified namespace, no stub object will be created.

The Namespaces page of the Collaboration Services Properties dialog box allows you to add new namespaces or edit existing namespaces if your policy has changed since Collaboration Services setup.

Figure 7. Namespaces



Free/Busy information synchronization settings

Synchronized free/busy information is kept on the Exchange server that stores the free/busy public folder for the administrative group containing the Exchange server where the Collaboration Services service mailbox resides.

So, if your Exchange deployment contains multiple administrative groups, the groups must access free/busy information using either public folder referrals or free/busy folder replication. Please see the appropriate Microsoft Exchange server documentation for more information.


The free/busy information synchronization process is performed as follows:

- 1 Collaboration Services periodically checks whether any changes were made to the free/busy information of the published objects.
- 2 When the service detects such changes, it creates update data packets and sends them to the HQ forest for distribution to subscribers.
- 3 Collaboration Services applies the free/busy information changes to the stub objects of the collections the forest is subscribed to based on the frequency settings you select.

General settings

Use the General page to adjust the frequency of scanning for updates and applying changes:

Figure 8. Free/busy general options



Publication settings

- Query Free/Busy for changes every:
Specify how often Collaboration Services should check for changes to the free/busy information of the published objects.

Subscription settings

- Update Free/Busy every:
Specify how often Collaboration Services should apply the free/busy information changes to the stub objects of the collections the forest is subscribed to.

Advanced subscription settings

- Wait for missing packets for
Specify how long the service should wait for missing data packets with free/busy information before automatically re-publishing this information.
i | **NOTE:** Increasing this value lessens the load on Collaboration Services and the Exchange server, but makes the free/busy information less accurate.
- Scan Active Directory for required objects every
In mixed collections (that is, intended for synchronization of several types of information), an object's free/busy information may be received before the object itself. In such cases, the free/busy information cannot be published right away.
Use this value to specify how often the service should scan Active Directory for the appearance of objects to which the free/busy information should be applied.

Scheduling

To avoid excessive load during production hours on the Exchange server where free-busy information is kept, schedule Collaboration Services to access the free/busy folder for the least busy time periods.

- i** | **NOTE:** Free/busy access schedules must be thoroughly coordinated, especially when some of the collaboration partners are located in different time zones; otherwise, synchronization performance may be poor.

To schedule free/busy folder access

- 1 On the Scheduling page, select the **Use Free/Busy access scheduling** check box.
- 2 Specify the time intervals when the service is allowed to read and update free/busy information on the Exchange server.

i | **NOTE:** If you allow access to the free/busy information in off hours only, the updated free/busy information will not be propagated to the other forests until the next day.

Limiting the number of Free/Busy messages to synchronize

The Quotas page allows you to specify the following:

- Overall maximum number of synchronized free/busy messages.
- Maximum number of messages allowed to be synchronized daily.

Calendar information synchronization settings

To support Exchange 2007, 2010, 2013, 2016 and 2019 deployments, Calendar information synchronization has been implemented in Collaboration Services.

Calendar Details data is stored in the users' mailboxes and not in the central location as with free/busy information. Free/busy information is stored in public folders.

Scan frequency

Use the General page to adjust the frequency of scanning for updates and applying changes:

- Query Calendar for changes every:
Specify how often Collaboration Services should check for changes to the Calendar information of the published objects.
- Update Calendar every:
Specify how often Collaboration Services should apply the calendar information changes to the stub objects of the collections the forest is subscribed to.

Scheduling

To avoid excessive load on the Exchange server during production hours, schedule Collaboration Services access to the Calendar information for the least busy time periods.

i | **NOTE:** Calendar information access schedules must be thoroughly coordinated, especially when some of the collaboration partners are located in different time zones; otherwise, synchronization performance may be poor.

To schedule Calendar access

- 1 On the Scheduling page, select the **Use Calendar access scheduling** check box.
- 2 Specify the time intervals when the service is allowed to read and update Calendar information on the Exchange server.

i | **NOTE:** If you allow access to the Calendar information in off hours only, the updated Calendar information will not be propagated to the other forests until the next day.

Using the Packet Queue as interim data storage

Collaboration Services use SMTP as a transport protocol. Though this protocol has many advantages and is a native protocol for Exchange, it introduces one problem; it does not provide guaranteed delivery. For Collaboration Services, this means that synchronization data packets do not necessarily arrive in the order in which they were sent. In fact, sometimes a packet may not arrive at its destination at all, and needs to be sent again.


To deal with this problem, Collaboration Services use a Packet Queue. This is a special location where all received data packets are kept until all the other data required to create a particular object arrive. After all data is received, the service extracts the necessary data from the queue and creates an object.

The Packet Queue is also used to store the value of linked attributes, such as the Manager attribute of a user object. The link cannot be applied until the object it points to is created, so the link is stored in the Packet Queue until then.

In addition to Active Directory data, the Packet Queue stores an object's free/busy information until the object is created and the free/busy information can be applied.

Packet Queue processes are fully automated and do not have to be managed. However, it is possible to fine-tune the packet queue behavior using the Packet Queue page.

Figure 9. Packet queue options



To adjust the time when the packet should be marked as obsolete

- On the Packet Queue page, select the time interval for the **Mark synchronization packet as obsolete if it could not be processed within:** setting. The default is one week and the setting has a range from one hour to one year.

Packets marked as obsolete will be deleted automatically if you select the Delete obsolete packets automatically check box. You can also clean up the packet queue manually by clicking the Clean Up Now button.


Suspended object storage

Objects excluded from synchronization due to various problems are placed into the Suspended Object Storage. If you notice that some objects are missing from the synchronization, you should first examine the log for errors and fix the problems that led to the objects being excluded from processing.

Once the problems have been fixed, open the Suspended Objects page in the Collaboration Services properties and click the Purge button to empty the Suspended Object Storage. This will add the excluded objects back into the synchronization queue to be processed again.

Re-examine the log for any errors, fix the problems and repeat the purge until the objects have been processed correctly.

Figure 10: Suspended object options



Manual operations

Synchronization

Manual synchronization is recommended when your collaboration partners must get the modifications made to objects in your forest before automatic synchronization is due.

Usually, Collaboration Services automatically synchronizes and updates collection data according to the time intervals specified on the Synchronization page and all sub-pages of the Collaboration Services properties. When you perform manual synchronization, all timeouts are minimized for a short period of time on both the publishing and subscribing forests, so all changes in published objects are synchronized much faster. This mode affects all available collections.

To start manual synchronization

- 1 Select the **Publications** node in the management tree.
- 2 Go to the **Action** menu and select **All Tasks | Synchronize Now**.

You can select whether to synchronize:

- Everything (Active Directory objects, Calendar information, and free/busy data)
- AD objects only
- Free/Busy only
- Calendar only

Re-publish collections

Re-publication is sending the entire data publication (all of the published objects data) again from the publisher. Re-publication performs exactly the same function as initial synchronization.

i | **NOTE:** Initial synchronization is resource consuming and may take significant time depending on the collection size. Therefore, re-publication should be performed only when you need to fully resynchronize all the data that was previously synchronized.

For example, you might need to re-publish data when Collaboration Services detects that an unrecoverable failure occurred. In this case, you will be prompted that re-publication is needed. You should not re-publish until you have fixed the cause of the error.

Re-publication should be limited in its use as an initial synchronization is performed. If numerous re-publications are performed, considerable time will pass until the publications are updated and synchronization is complete.

To re-publish a single collection

- 1 Select the collection in the management tree.
- 2 Open the **Action** menu, point to **All Tasks**.
- 3 Point to **Re-publish**, and choose what data you want to republish:
 - Everything (Active Directory objects, Calendar information, and free/busy data)
 - AD objects only
 - Free/Busy only
 - Calendar only

Alternatively, to run the Re-publish command, you can use the toolbar button or collections' shortcut menu commands.

To re-publish all collections

- 1 Select the **Publication** node in the management tree.
- 2 Open the **Action** menu, point to **All Tasks**.
- 3 Point to **Re-publish**, and choose what data you want to republish:
 - Everything (Active Directory objects, Calendar information, and free/busy data)
 - AD objects only
 - Free/Busy only
 - Calendar only

Initiating re-publication from the subscriber side

To initiate re-publication when you are a subscriber, select either the Subscription node or a particular collection and then use the Request Re-publication command from the shortcut menu.

i | **NOTE:** Re-publication of a collection may take significant time depending on its size and should be performed only when required.

Configuration database clean-up

In case of intensive Active Directory updates, your collections may turn out to contain some objects that have become obsolete. For example, your published collections may include a number of user accounts that have been already removed from Active Directory.

To prevent from synchronization conflicts that might occur due to obsolete objects still kept in the collections, you can periodically clean up Collaboration Services configuration data (new data will be stored to the configuration database at the next scan).

To perform a clean-up

- Right-click the **Publication** node in the management tree and select **Clean Up Database**.

Transport management

- [Overview](#)
- [Send](#)
- [Receive](#)
- [Scheduling](#)
- [Limiting traffic rates](#)

Overview

Instances of Collaboration Services installed in the HQ and branch forests communicate with each other using the Simple Mail Transfer Protocol (SMTP), which is a native protocol for Exchange organization.

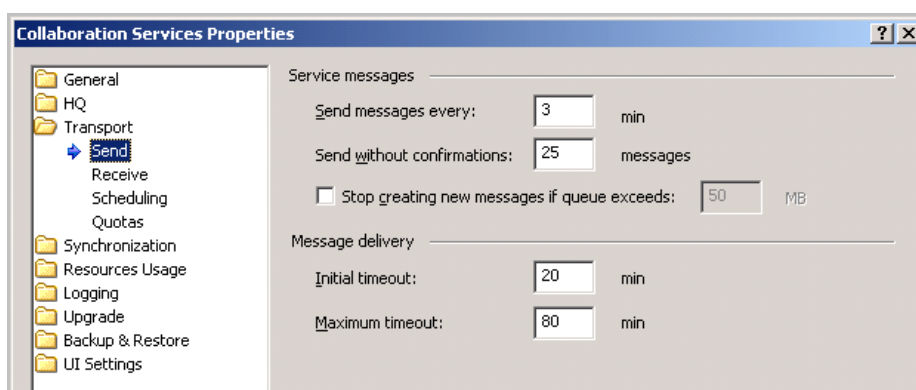
All communications between partners is performed through administrative mailboxes, one mailbox per forest. The data packets are compressed, grouped in regular email messages, and sent to another service mailbox.

To ensure optimal data flow between synchronization partners, you can adjust the transport settings to suit your needs. To view or edit transport settings, use the Transport page of Collaboration Services Properties dialog box.

Send

The Send page allows you to specify settings for sending messages to other synchronization partners. After a data packet to be sent is generated and compressed, the service performs the following procedure:

- 1 First, if the packet size exceeds the maximum allowed size, it is divided into multiple messages of the default size.
- 2 Then the messages are put in the message queue from which they are sent to the designated synchronization partner.
- 3 A message is kept in the queue until the service receives an acknowledgement of successful message delivery from the partner.



Collaboration Services guarantees data delivery by checking the delivery status of every message and re-sending any message for which delivery is not confirmed.

- The time interval at which Collaboration Services sends new messages from the queue can be specified with the `Send messages every`: option.
- You can specify the number of messages that may be sent before confirmation of delivery of a given message in the `Send without confirmations`: option.
- Use the `Stop creating new messages if queue exceeds`: option to specify the maximum queue size. Once this value is reached, no more new messages will be created. This option prevents you from running out of disk space due to long queues if a partner is not responding. As soon as a partner comes back online and starts processing the incoming messages, the queue will clean up and then new messages can be created.
- The time interval after which a message will be re-sent if is specified in the `Initial timeout`: option. Each time the message is re-sent, the timeout is automatically increased.
- Use the `Maximum timeout`: option to specify the maximum timeout allowed.

Receive

The Receive page allows you to specify the settings for receiving incoming messages.

- The time interval for checking the service mailbox for new messages is specified in the `Check for new messages every`: option.
- To specify a maximum number of messages that can be retrieved from the mailbox during one session, use the `Retrieve per session`: option.
- If the specified `Session timeout`: value is more than 0, then after retrieving the specified number of messages, Collaboration Services will stop the current session and wait for the specified length of time. Having an interval between sessions may help decrease the load on the Exchange server.
- If the `Session timeout`: value is 0, then Collaboration Services will retrieve messages continuously until all available messages are retrieved from the mailbox.

Scheduling

Message exchange between synchronization partners can be scheduled in order to decrease the load on the Exchange servers during production hours.

To schedule exchanges


- Select the **Use SMTP transport scheduling** check box and select the time intervals when the service is allowed to send and receive messages containing synchronization data.

i **NOTE:** Transport is the key to successful implementation of inter-forest synchronization. Transport, Active Directory, Calendar, and Free/Busy schedules should be carefully coordinated (especially if collaboration forests are located in different time zones) and provide a sufficient time interval for Collaboration Services to perform its activities (sending messages, accessing Active Directory, and accessing free/busy information). Otherwise, synchronization performance may be poor.

Limiting traffic rates

The Quotas page allows you to control the incoming and outgoing traffic while messages are being exchanged between forests.

Figure 11. Quota options



Outgoing traffic

To control outgoing traffic, the HQ and branch administrators can specify a maximum outgoing traffic size in megabytes per hour by selecting the Control outgoing traffic check box. As soon as the traffic exceeds the specified limit, sending of new packets is delayed until the traffic drops below the limit, and an appropriate message is logged to the Collaboration Services Event log.

Incoming traffic

To control the incoming traffic, the branch forests and an HQ forest can limit the number of incoming messages and the overall size of the incoming traffic.

Use the Control incoming traffic and Control incoming traffic by number of messages options to protect the service from being overloaded with large amounts of synchronization data:

- If the Control incoming traffic check box is selected, the service will process only a limited amount of synchronization data per hour.
- If the Control incoming traffic by number of messages check box is selected, the service will process only a limited number of incoming messages every hour.

i **NOTE:** You can use these options to control incoming mail flow to ensure safe service functionality. However, the use of these options may impact the initial synchronization performance significantly, so it is recommended to turn them on only after the initial synchronization is over for all collections.

The HQ forest web interface includes two additional options that can be used to set the bandwidth and quotas for branch forests. These options let the HQ forest administrator specify which message and traffic quotas should be assigned each newly-registered branch forest.

To ensure optimal synchronization performance between all synchronization partners, HQ constantly redistributes the message quota and the bandwidth between the branch forests. The redistribution is performed every hour for each branch.

Collaboration Services uses an adaptive algorithm when redistributing available bandwidth and message quotas between branches. This algorithm is quite complex and takes into account statistical information, such as when the branch was added and its average bandwidth usage.

Fine-tuning and maintenance

- [Active Directory synchronization performance](#)
- [Resource usage](#)

Active Directory synchronization performance

To improve Collaboration Services performance, first register schmmgmt.dll (located in the \\COMPUTERNAME\admin\$\SYSTEM32 subfolder). Then, using the MMC console, add the Active Directory Schema Snap-In to the program group of the Administrative Tool on a Domain Controller and ensure indexing is turned on for the following object attributes:

- proxyAddresses
- mail
- targetAddress
- mailNickName
- Extension attribute used by Collaboration Services for storing data (selected during the Collaboration Services setup procedure—by default, extensionAttribute10)

Resource usage

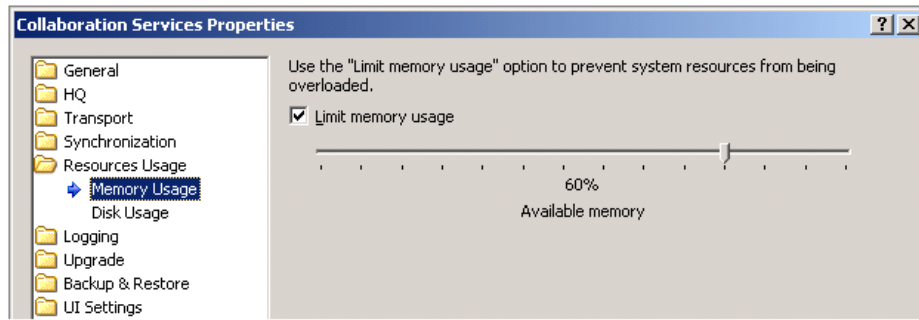
To ensure safe service functionality, Collaboration Services always monitors the usage of system memory and disk space. Using the Resource Usage page of Collaboration Services properties, you can adjust the related settings according to your needs.

Memory usage

This page allows you to specify whether the service memory usage should be limited on the Collaboration Services server. This helps you protect the system resources from being overloaded. Use the Available memory slider to adjust the maximum amount of system memory that can be used by Collaboration Services.

- i** **NOTE:** Though the Available memory slider can be set to the minimum value, there is a minimum amount of memory that the service is allowed to use. This amount is based on the system requirements and cannot be changed. Also, the memory usage for the first 60 seconds after the service starts is not limited to prevent the service from being cyclically restarted.

Figure 12. Memory usage information

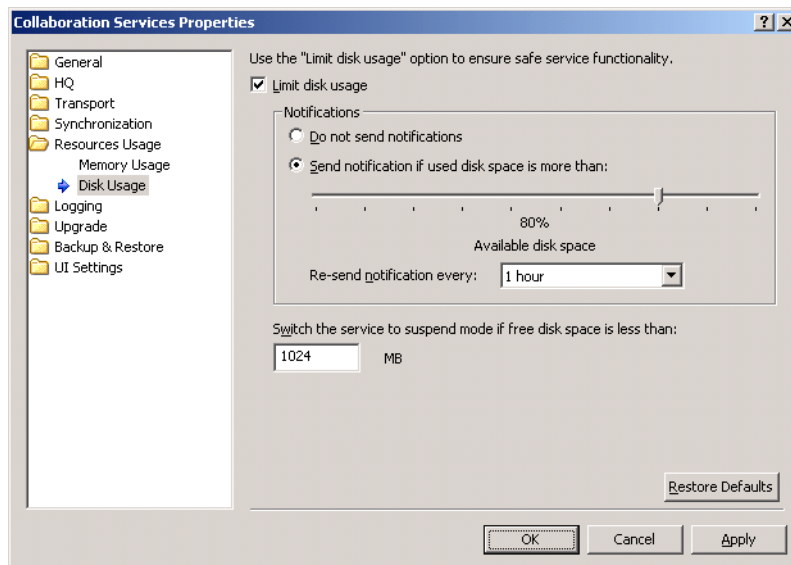


Decreasing the memory usage limit worsens the service performance, so this is not recommended unless other software is running on the same computer with Collaboration Services.

Disk usage

This page allows you to specify the service behavior if the computer the service is installed on runs out of free disk space.

Figure 13. Disk usage options



- Select the Limit disk usage to make the service monitor the disk space usage on the volume where the service is installed.

TIP: It is highly recommended to have this option turned on.

- To receive notifications when a specified percentage of free disk space is exceeded, select Send notifications if used disk space is more than option.
- The notification format and the notification service to be used can be configured on the Logging page of the Collaboration Services Properties dialog box.
- To specify the notification time interval, use the Re-send notification every drop-down box.
- The Switch the service to suspend mode if free disk is less than option allows you to choose the minimum amount of free disk space for the service to be allowed to run. If the free disk space drops below this threshold, the service is automatically switched to suspend mode. When the service runs in suspend mode,

no synchronization is performed. The service will automatically switch to normal mode as soon as the specified minimum amount of disk space is free.

Resource usage statistics report

To view statistics information on resource usage, select the Resource Usage node in the management tree. The following information is displayed in the statistics pane of the Collaboration Services web interface:

Table 3. Resource usage statistics

Statistics Section	Can Be Used For	Displays
General information	Checking the status of the synchronization service	Service state, which can have the following values: <ul style="list-style-type: none">• Running• Suspended
Memory resources	Checking the amount of memory currently used by the synchronization service	<ul style="list-style-type: none">• The time of last memory usage information update• Service memory usage• Service handles usage
Disk resources	Checking the amount of the available free disk space	<ul style="list-style-type: none">• The time of last disk usage information update• Working directory• Disk usage statistics
Storage integrity	Checking the condition of the configuration database	Configuration database integrity, which can have the following values: <ul style="list-style-type: none">• OK• Recoverable error detected• Fatal error detected• The time of last disk usage information update

Statistics update frequency

Use the UI Settings page to configure whether and how often Collaboration Services should automatically refresh the synchronization statistics. It is not recommended to set the refresh period to a low value, because the refresh may take significant time on slow connections or if the Collaboration Service server is too busy processing synchronization data.

Backup, restore, and troubleshooting

- [Collaboration Services automatic backup](#)
- [Troubleshooting](#)
- [Repair or uninstall](#)

Collaboration Services automatic backup

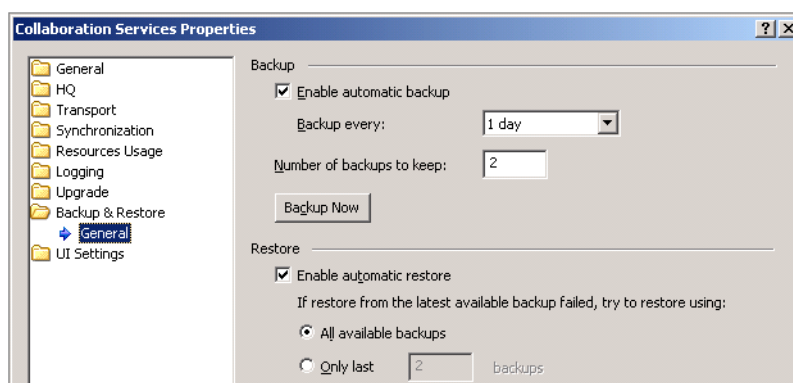
By default, to protect the configuration database from power and hardware failures, Quest Collaboration Services is configured to perform an automatic configuration backup and restore by default. All critical information is saved in the database file containing all existing collections, statistics, and the service configuration information.

If a problem is encountered by Collaboration Services, then the service will automatically detect whether any configuration data is corrupted and perform an automatic restore.

Backups can be performed as quickly as every 3 hours and as long as every 180 days. While there are no specific recommendations on how often to perform a backup, there are a number of factors to consider such as how often are Active Directory and Exchange changes made and what is the impact on operations in the event of power or hardware failure. Also, remember that backups performed at a quicker rate will impact the overall performance of Collaboration Services.

Use the Backup & Restore page of the Collaboration Service Properties dialog box to adjust the backup and restore options.

Figure 14. Backup and restore general options



In particular, you can specify:

- Whether and how often the automatic backup should be performed.
- How many backup files should be kept.

If you have applied significant changes to the Collaboration Services configuration, you may want to perform an immediate configuration backup. To do this, click the Backup Now button.

The restore options allow you to:

- Enable automatic restore (recommended).
- Select what backups should be used if restore from the latest backup fails.

The default settings include:

- The configuration data backup is located in the %Program Files%\Quest Software\Collaboration Services\DataBackup folder.

Folders are identified by the date and time the backup was performed.

The files that are backed up were originally located in the %Program Files%\Quest Software\Collaboration Services\Storage folder.

- The number of automatic backups to keep is set to two.

If disk space on the computer where Collaboration Services is installed is not an issue, then it is recommended to increase this 5 or 10 backups. This is because a considerable amount of time may pass before a problem is noticed.

It is also recommended that this period be longer than the recommended file system backup described under [Regular file backup](#) on page 49. This will ensure you are covered by the automatic Collaboration Services backup and restore and a regular file system backup to tape or disk.

Regular file backup

In addition to automatic backup of configuration data on each server running Collaboration Services, it is also recommended to back up the contents of the installation folder %Program Files%\Quest Software\Collaboration Services folder, Keys from HQ or Branches and the entire registry key:

- HKEY_LOCAL_MACHINE\SOFTWARE\Aelita\AelitaCollaboration Services and all subkeys

– OR –

HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Aelita\AelitaCollaborationServices and all subkeys

To backup the registry keys using regedit

- 1 Locate and select the key that you want to back up.
- 2 Click **File | Export**.
- 3 In the Save in box, select where you want to save the backup copy to, enter a name for the backup, and click **Save**.

Before backing up the Collaboration Services installation folder:

- Close the Collaboration Services console.
- Stop and disable all Collaboration Service services (Logging Service, Synchronization Service, and Watchdog Service) running on the computer using the Services snap-in. Stop the Watchdog service first as it will restart the other services as soon as you stopped them.

Performing a weekly (or longer) backup of the Collaboration Services folder and registry keys to tape or disk will allow you to cover any issues that the automatic Collaboration Services backup may not include based on the selections you have made.

How to re-deploy Collaboration Services to another computer from a backup

You may need to re-deploy Collaboration Services on a new computer if the original computer is unavailable due to disaster or hardware failure. To do this, you will need the backup copy of the Collaboration Services installation folder created during a file backup. See [Regular file backup](#) on page 49.

To re-deploy Collaboration Services on a new computer

- i** | **NOTE:** You must install Collaboration Services to the same path (drive and folder) as the previous installation and specify the same service account with the required permissions as described in the Deployment Guide. The rest of the settings you specify during setup, including creation of new public encryption key will not be in effect. Instead, the settings from the backup will be used.
- 1 Install Collaboration Services on a new computer.
 - 2 During installation, if a Branch use the latest HQ key from the backup. If an HQ, create a new encryption key file. We are not going to use this key file.
 - 3 Once Collaboration Services is installed, stop and disable all its services (Logging Service, Synchronization Service, and Watchdog Service) using the Services snap-in.
 - 4 Import the entire HKEY_LOCAL_MACHINE\SOFTWARE\Aelita\AelitaCollaborationServices registry key and all subkeys from the backup using regedit.
 - 5 Restore the contents of the Collaboration Services installation folder from the backup and overwrite all the files in the new installation.
 - 6 Enable and start all the Collaboration Services services from the Services snap-in.
 - 7 Start the Collaboration Services management console and create a new backup of configuration data. To do this, right-click the **Collaboration Services** root node and select **Properties**. In the Properties dialog box, click the **Backup and Restore** tab and click the **Backup Now** button.
 - 8 Re-publish all the collections from all the sides involved (HQ and all branches). To decrease load on HQ server, re-publish collections separately one by one.
- i** | **NOTE:** To republish all collections simultaneously, right-click the Publication node and select Re-publish | Everything.

Troubleshooting

This section describes the some issues you may face during Collaboration Services operations and explains how to resolve them:

Table 4. Troubleshooting

Symptoms	Possible reasons	solution
You see the following message in a particular synchronization partner's statistics: "Synchronization partner is not responding"	The SMTP connection between forests is configured incorrectly, - OR - The encryption key was changed for one of the forests.	Check whether the SMTP channel is working by sending a test email between the service mailboxes of both forests in both directions. Export the forests' public key files and re-register both forests using the newly-generated public key files.
Some objects are not available from a collection your forest subscribes to.	The initial synchronization is still being performed.	Wait until the initial synchronization is completed. It could take up to a few days for large collections, depending on the servers' hardware and network conditions.
One or more objects are not available from a collection your forest subscribes to, but the initial synchronization is already completed.	Some of the objects were accidentally deleted from the Collaboration Services container.	Initiate re-publication of the collection using the Request re-publishing command.
You cannot add newly-added custom attributes, or objects from a newly-added domain to the synchronization.	Collaboration Services reads a forest's schema only once, during the service start-up. Thus, on-going changes to the schema cannot be detected.	Restart the synchronization service using the Services MMC snap-in.

Repair or uninstall

To repair or remove Collaboration Services

- 1 From the Start menu, select **Control Panel | Add/Remove Programs | Quest Collaboration Services**, and select **Change**.
- 2 If you need to Repair, click the **Options** button, select the check boxes next to the required repairs, and then click **OK**.

The repair wizard will guide you through the repair process.

– OR –

To completely remove the application, select **Remove**.

Collaboration Services and the web management console (if installed locally) will be removed.

Appendix A: Customizing the format of synchronized data

Sometimes publishing or subscribing administrators might want to change the format or appearance of the data they publish or subscribe to. To do that, mappers can be used. A mapper is a plug-in component that changes the appearance of data that you publish or subscribe to.

If you use mappers when creating a collection, the mappers modify the data before it is published, so all subscribers get modified data. If you use mappers while subscribing to a collection, the mappers modify the published data before it is applied to your particular forest.

i | **NOTE:** Mappers applied during publication do not modify the original Active Directory objects; only the information to be published is modified before it is distributed to other forests.

Collaboration Services is shipped with the following mappers:

- displayName mapper
- Suffix mapper
- inetOrgPerson to contact mapper
- SMTP address filter mapper
- User to Contact mapper
- Group to Contact mapper

There is one additional mapper available only in Subscriptions:

- Attribute filter mapper

displayName mapper

The displayName mapper can be selected on the "Mappers" page for both publications and subscription properties. It has five tabs corresponding to the types of objects that Collaboration Services can create as a stub. You can modify Contacts, Groups, InetOrgPersons, Query-based Distribution Lists, or Users.

For the required object, select the Customize Display Name appearance for objects option, click the Add Attribute button and build the template for the object's Display Name. For example, to have FirstName.LastName enter: [GivenName].[SN].

i | **NOTE:** Ensure the attributes you select are being synchronized within the publication properties. If not, an error will be generated stating that attributes are not available.

Suffix mapper

This mapper adds a specified suffix to the display name of the published objects so that the users can distinguish published objects from the objects of their own forest. For example, you may want to add the [PartnerCo] suffix to the display names of objects from a partner company in the subscribed collection to emphasize that they belong to the other company.

inetOrgPerson to Contact mapper

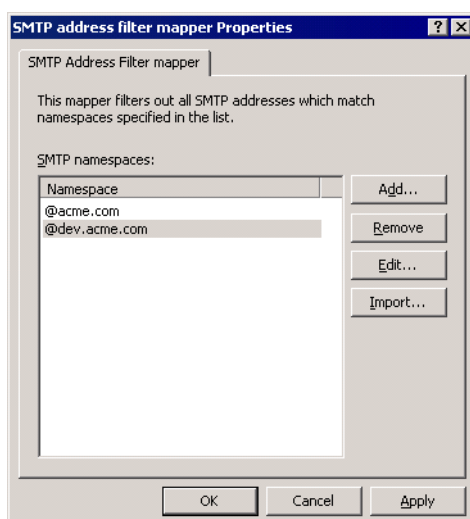
This mapper converts inetOrgPersons you may have in your Active Directory into contact stub objects. There are no parameters to configure with this mapper.

SMTP filter mapper

This mapper allows a subscribing administrator to defend against “address stealing.” By adding an SMTP address to the published object, the publishing administrator effectively assigns his forest’s users email addresses from another forest. In a merger or acquisition scenario, this is usually a requirement and is done on purpose. However, the same action can also be performed by a malicious administrator in a partner collaboration scenario, causing unsanctioned email address usage with all associated risks, including possible information leaks. By using this mapper and filtering out your own forest’s addresses from incoming objects, you protect your forest from this type of attack.

To configure the mapper, add all of your company’s SMTP email address namespaces (including the @ sign) to the list of namespaces to be filtered out. You can also import a list of namespaces from a text file by clicking the Import button and choosing the file. The file should list the SMTP namespaces, one namespace per line.

Figure 15. SMTP filter mapper



User to Contact mapper

This mapper converts users and inetOrgPerson objects into contacts. Converting these objects to contacts might be required to meet internal regulations. This mapper can also be useful if you have software that is licensed by the number of user objects in Active Directory and you do not want to buy additional licenses.

- NOTE:** In the Global Address List, contacts are marked with a globe next to their names, while users are not. To achieve full transparency, you should avoid using this mapper if possible, so that users from all forests look the same in the GAL. Otherwise, end users may be puzzled by the differences.

Group to Contact mapper

This mapper converts groups into contacts, allowing you to hide group membership from other forests. This is useful if you want users from other forest to be able to send mail to your distribution groups but do not want the group members to be revealed.

NOTE: When publishing a group as a contact, it is not necessary to publish all group members for distribution lists to function, as long as email messages sent to the contact are always redirected to the original group object in the group's "home" forest.

Mapping groups from Exchange 2007

If you want to use the Group to Contact Mapper for a source group from an Exchange 2007 organization, you must first turn off user authentication. Otherwise, message delivery to the target contact will fail.


To Disable User Authentication in Exchange 2007

- 1 Open that distribution group's properties in the Exchange 2007 Management Console.
- 2 On the Mail Flow Settings tab, double-click **Mail Delivery Restrictions**, and clear the **"Require that all senders are authenticated"** check box.

Attribute filter mapper

The Attribute Filter mapper is only available for subscriptions. This mapper allows a subscribing forest's administrator to filter out unnecessary or undesired attributes, thus controlling exactly which attributes are applied to the forest.

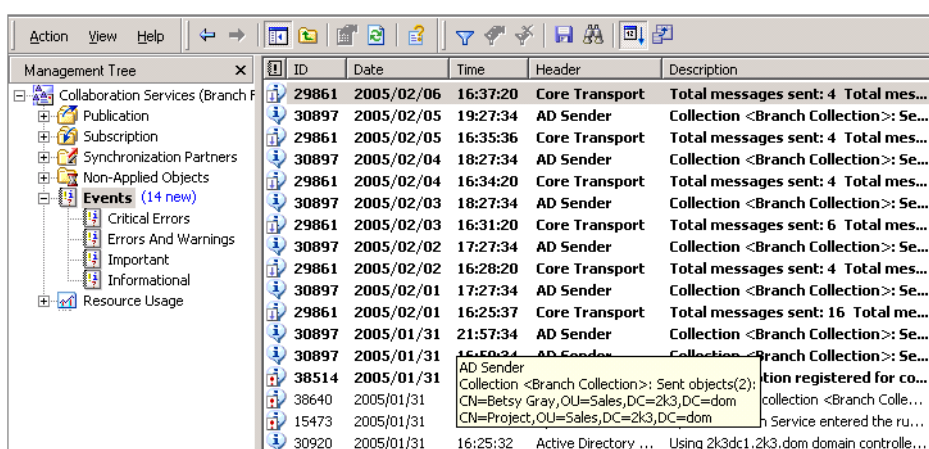
Figure 16. Attribute filter mapper



Appendix B: Collaboration Services events

To view the service events in the Collaboration Services web interface, select the Event node in the management tree.

Figure 17. Collaboration Services web interface



ID	Date	Time	Header	Description
29861	2005/02/06	16:37:20	Core Transport	Total messages sent: 4 Total mes...
30897	2005/02/05	19:27:34	AD Sender	Collection <Branch Collection>: Se...
29861	2005/02/05	16:35:36	Core Transport	Total messages sent: 4 Total mes...
30897	2005/02/04	18:27:34	AD Sender	Collection <Branch Collection>: Se...
29861	2005/02/04	16:34:20	Core Transport	Total messages sent: 4 Total mes...
30897	2005/02/03	18:27:34	AD Sender	Collection <Branch Collection>: Se...
29861	2005/02/03	16:31:20	Core Transport	Total messages sent: 6 Total mes...
30897	2005/02/02	17:27:34	AD Sender	Collection <Branch Collection>: Se...
29861	2005/02/02	16:28:20	Core Transport	Total messages sent: 4 Total mes...
30897	2005/02/01	17:27:34	AD Sender	Collection <Branch Collection>: Se...
29861	2005/02/01	16:25:37	Core Transport	Total messages sent: 16 Total me...
30897	2005/01/31	21:57:34	AD Sender	Collection <Branch Collection>: Se...
30897	2005/01/31	16:50:34	AD Sender	Collection <Branch Collection>: Se...
38514	2005/01/31		AD Sender	Collection <Branch Collection>: Sent objects(2): CN=Betsy Gray,OU=Sales,DC=2k3,DC=dom CN=Project,OU=Sales,DC=2k3,DC=dom h Service entered the ru...
38640	2005/01/31			collection <Branch Colle...
15473	2005/01/31			
30920	2005/01/31	16:25:32	Active Directory ...	Using 2k3dc1.2k3.dom domain controlle...

You can filter, sort, and save events and search for particular items using the commands from the Events node's shortcut menu.

To clear the current event view (all events displayed in the management console), use Clear All Alerts command (these records stay intact in the event log).

You can right-click any item in the right pane and use the following shortcut menu commands:

- Find—Enables you to search for specific things such as EventIDs, types.
- Sort and Select Columns—Changes the way the events are displayed.
- Mark As Read—Highlights the last viewed item.
- Refresh—Updates the event list being displayed.
- Properties—Enables you to view event details.

Event format

Collaboration Services events have the following format:

- Type—One of the following:
 - Informational
 - Warning
 - Error

- Level—Event severity level:
 - High
 - Normal
 - Low
- ID—Unique event ID
- Date, Time—When the event was generated
- Header—Event source
- Description—Detailed event description

To explore an event's details

- 1 Right-click the event in the right pane and select **Properties**.
- 2 To view the details of the next or the previous item in the list, click the up or down arrow buttons in this dialog box. To copy event data to the clipboard, click the **Copy** button.

Using event filters

Predefined filters

To help you view the events more efficiently, Collaboration Services allows you to use event filters. There are several predefined event filters:

- Critical errors
- Errors and Warning
- Important
- Informational

To use one of these filters, select the corresponding sub-node of the Events node in the management console.

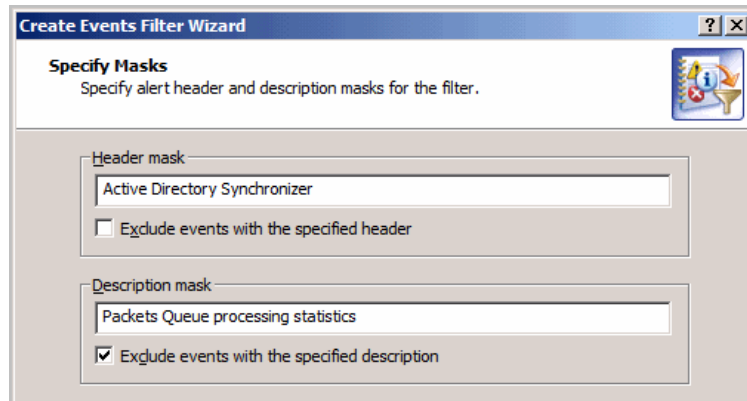
Custom filters

You can customize any filter by selecting it in the management tree under the Events node and selecting the Edit Filter command. The Create Events Filter wizard is started to help you with the procedure.

To create a new event filter

- 1 Run the Create Events Filter wizard.
- 2 Specify a name for the new event filter and select the template you want to base it on. The following templates are available:
 - Custom (allows you to make a filter of your own)
 - Informational
 - Important
 - Errors and Warnings
 - Critical Errors

- 3 Select which types, importance, levels, and IDs of events should be shown by the filter.
- 4 Set up the filter to include or exclude events with a specified header or description by using a corresponding mask.



If the Exclude check box is cleared, events containing the specified header or description are included in the view; otherwise, they are excluded from it.

- 5 You can configure the filter to show only those events that occurred during a specified period of time. To specify the beginning of this period, either select Events on and specify a particular date and time, or choose to show all events starting from the first one by choosing First event. Similarly, to define the end of the period, you can either specify the end date and time or choose to show all events up to the last one.

- A
 - Active Directory
 - limiting the objects to synchronize 35
 - synchronization performance 45
 - attribute filter mapper 54
 - automatic conflict resolution 25
- C
 - configure
 - Active Directory access settings 10
 - alerts 11
 - collection 21
 - synchronization 32
 - Synchronization Partners 12
 - conflicts 24
 - delete 26
 - manual resolution 29
 - matching 27
 - pre-matching 28
 - resolving 24
 - create
 - collection 14
- D
 - data storage 39
 - database clean up 41
 - delete conflicts 25
 - disable a collection 19
 - disk usage 46
- E
 - edit
 - collection 20
 - subscription properties 21
 - enable a collection 19
- Exchange 2007
 - using group to contact mapper with 54
- G
 - group to contact mapper 54
- L
 - logging 10
- M
 - mapper
 - attribute filter 54
 - group to contact 54
 - SMTP filter 53
 - suffix 52
 - user to contact 53
 - matching 27
 - memory usage 45
- P
 - pre-matching 28
 - Publish Collection Wizard 15
- R
 - receive messages 43
 - remove a collection 22
 - repair 51
 - resolving conflicts 24
- S
 - send messages 42
 - SMTP mapper 53
 - subscribe to a collection 20
 - suffix mapper 52
 - synchronization
 - Active Directory settings 33
 - calendar information 38
 - free/busy 36
 - general settings 32
 - manual operations 40
- T
 - transport
 - quotas 43
 - receive 43
 - scheduling 43
 - send 42
 - transport management 42
 - troubleshooting 50
- U
 - unsubscribe from a collection 22
 - usage statistics 47
 - user to contact mapper 53
- V
 - view collection settings 22

W
web management console 7

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

Technical support resources

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

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

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