FORCEPOINT

Directory Synchronization Client Administrator's Guide

Forcepoint Cloud Protection Solutions

Last modified: August 8, 2022

Contents

1 | Introducing the Directory Synchronization Client

> Related topics:
>
> - *Before you begin*, page 3
> - *Basic directory synchronization steps*, page 6

The Directory Synchronization Client collects user directory information from one or more directory servers for use by Forcepoint cloud-based services.

- For Forcepoint Email Security Cloud, the client synchronizes registered email addresses and groups (Mail synchronization).
- For Forcepoint Web Security Cloud, the client synchronizes user and group information (Group+User synchronization).
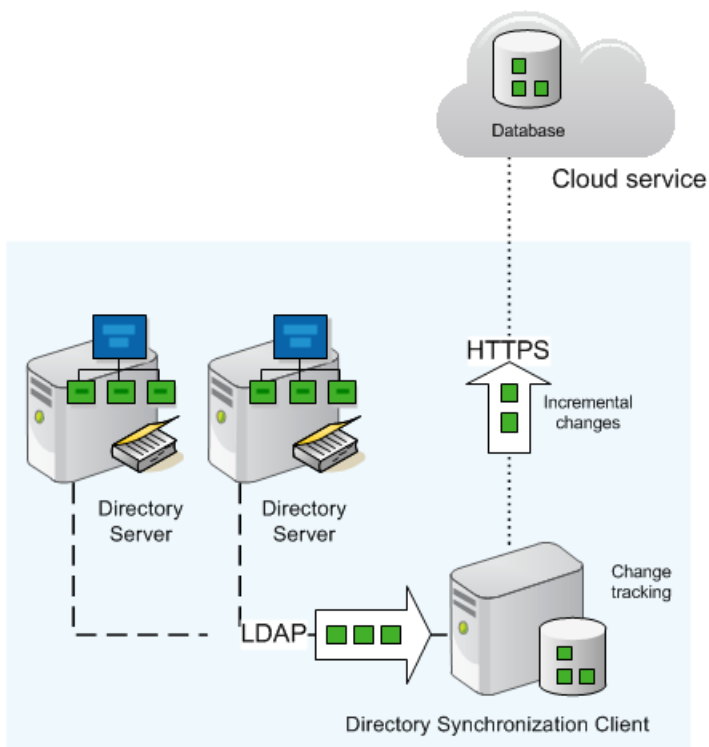
The Directory Synchronization Client supports on-premises LDAP-based directories such as Microsoft Active Directory and IBM Domino, as well as cloud-based directory services such as Microsoft Azure and Google Apps.

> **Note**
> Support for Directory Synchronization Client is limited to the most recent version and the version that immediately preceded it.

The Directory Synchronization Client runs either as a graphical or command-line application. Start by using the graphical application to create a configuration profile. You can then:

- Run the synchronization process from the graphical console or the command line.
  - The graphical console allows you to choose either a full upload of all data or an incremental upload.
  - By default, the command-line synchronization process passes only incremental changes since the last run.

    You can enable an option in the cloud portal to force a full update using the command-line process.
- Schedule the process to run automatically.
- Receive email notifications reporting the results of each synchronization run.

Email address registration

Forcepoint Email Security Cloud can protect against dictionary-type spam attacks by registering your valid email addresses and rejecting any email destined for invalid addresses. The Directory Synchronization Client helps you maintain your valid addresses by synchronizing the update of registered addresses with the cloud service. The task can be automated and, for example, integrated with Human Resources procedures for employees leaving or joining the company.

Registered addresses are synchronized using:

- A secure HTTP-based interface to the Forcepoint Email Security Cloud synchronization service
- The Directory Synchronization Client to extract address data from your directory sources and export it via the synchronization service

Group and user synchronization overview

In the Forcepoint cloud service, your directory information is used in applying web and email security policy rules to users and groups.

If you are synchronizing groups, you must also synchronize users.

● When you synchronize a group, only information about the group itself (such as the group name and any parent group) is transfered—not the contents of the group.

● User synchronization includes details of each group that users belong to.

When you apply a web policy or an email policy to a synchronized group, that policy is applied to all synchronized users who are members of that group.

Before you begin

Before installing and running the synchronization client, be clear about what directory information you do and do not need to send to the cloud service.

For example, users may be members of many groups—some global groups (like "All" and "All Sales"), a geographical group (like "London" or "New York"), a department (like "NY Telesales"), and so on.

It is important to synchronize only groups that are going to be useful in the cloud (for setting policies or exceptions, for example). Using the diagram in the next section:

● If members of the New York Telesales department use a policy that gives them special permissions, synchronize the "NY Telesales" group.

● If you do not use geographical policies, do not synchronize the "London" group, even though the London users might be using the cloud service.

The cloud service is designed to accept users with references to groups that are not synchronized to the service.

For specific information about preparing to synchronize directory information, see:

● *LDAP filter for users, groups, and email*, page 3

● *Multiple domains*, page 5

● *Synchronizing Dynamic Distribution Lists*, page 5

● *Renaming groups in your directory*, page 6

LDAP filter for users, groups, and email

In the Directory Synchronization Client, there are 3 synchronization types (groups, users, and email), each with its own LDAP search set up. The searches are independent of one another to give you flexibility in selecting the appropriate data.

For example, you can use the LDAP group attribute to select the users you want, even if you choose not to synchronize the group itself.

This is an example Active Directory schema:
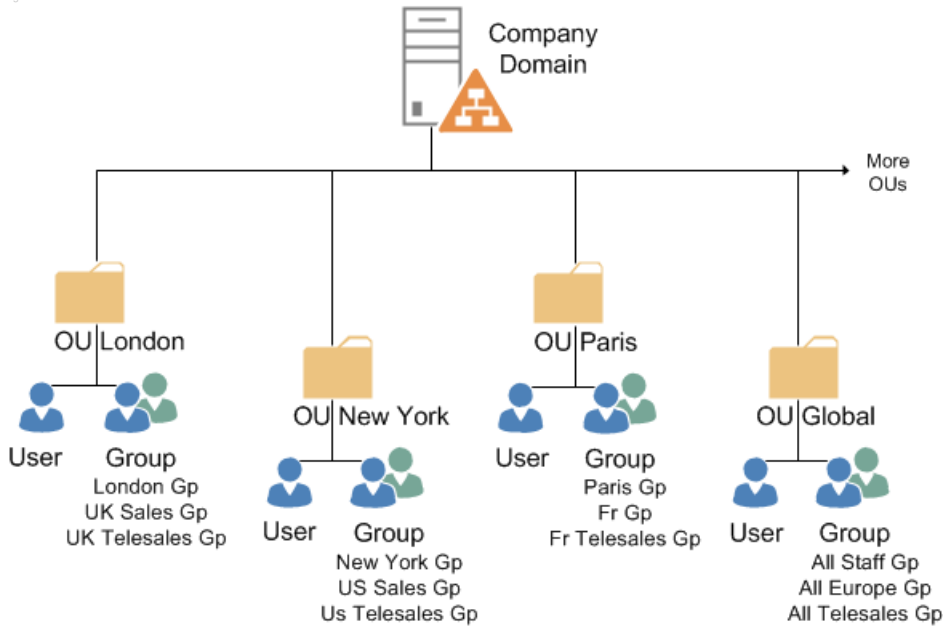
Below are some examples of the synchronization choices you might make based on this schema.

Group selection

If you require specific policies or exceptions in your cloud product for French and English telesales staff, select the "UK Telesales" and "Fr Telesales" groups for synchronization.

More information about selecting groups is provided in the configuration profile setup instructions (*Step 6: Selecting groups for synchronization*, page 32).

User selection

If your cloud product is currently available only for your European staff, synchronize only those users. You can achieve this by:

● Set the users LDAP search filter to search on European users by group.

  When setting up a users configuration, set the LDAP search base to the domain level. Then the search filter is set to something like the following:

  ```
  (&(objectCategory=person)(objectclass=user)(memberOf=CN=All
  Europe,OU=Global,dc=company,dc=com)
  (!(userAccountControl:1.2.840.113556.1.4.803:=2)) )
  ```

  This selects users that are members of the global Europe group, and that are enabled (strictly, that have accounts that are not disabled).

  More information about LDAP search filters is provided in the configuration profile setup instructions (*Step 4: Setting up the LDAP search configuration*, page 21).

- Select users from the relevant OUs by setting up multiple data sources for the LDAP search.

  When setting up your users configuration, on the **Configure data source** window check the **Advanced** box. Select another source, and then set the LDAP search base to be one of the European OUs (for example London or Paris). Leave the search filter as the default to load all users from that OU.

  Once you have configured that data source, repeat the process for each OU that you want to include. The Directory Synchronization Client merges all of the users from the various OU sources and synchronizes them with the portal.

  More information about multiple data sources is provided in the configuration profile setup instructions (*Step 2: Selecting your data source*, page 17).

Email selection

The valid email address list can be created from a completely different LDAP search. This may be especially useful in dealing with users who leave the organization.

For example, when a user leaves the organization, you are likely to disable their account immediately upon departure. If you use the default filter in the users synchronization, this removes the departed user from the cloud service (disabled accounts are not synchronized). You might, however, want to allow email messages to be received for a while after the employee's departure, so the email synchronization might still include the employee's address.

To include European email addresses from the above example and also include departed users in the valid address list, set the LDAP search base to the company domain, and set the search filter to:

```
(&(objectCategory=person)(objectclass=user)(memberOf=CN=All
Europe,OU=Global,dc=company,dc=com)
```

Multiple domains, dynamic distribution lists, and renaming groups

Multiple domains

In directory structures such as Active Directory Forests, or where multiple directory servers contain user data, you can use the multiple data source option. The Directory Synchronization Client searches the multiple sources and merges the data before sending it all to the cloud service.

More information about multiple data sources is provided in the configuration profile setup instructions (*Step 2: Selecting your data source*, page 17).

Synchronizing Dynamic Distribution Lists

If you have set up a Dynamic Distribution List in Microsoft Exchange, the default mail synchronization filter does not synchronize the email address of this list. If you

want to include the address of a Dynamic Distribution List in your synchronization, change your mail synchronization filter to:

```
(|(&(mailnickname=*)(objectCategory=person)(objectClass=user
))(objectCategory=group)(objectClass=msExchDynamicDistributi
onList))
```

Note that "(objectClass=msExchDynamicDistributionList)" has been added into the "or" part of the filter.

Renaming groups in your directory

Once you have synchronized your groups with the cloud service, we recommend that you do not rename them in the directory. If you do rename a synchronized group, the new group name is automatically sent to the cloud service on the next synchronization. If you have set up any group-based black and white list settings for Forcepoint Email Security Cloud, however, you must manually reapply those settings to the new group name in the cloud portal.

Basic directory synchronization steps

In the cloud portal

1. Configure directory synchronization for your account.
2. Set up authentication for the synchronization client. Create a dedicated user name and password for the client to gain access to the cloud service.

Refer to the portal Help for instructions on setting up the cloud portal for directory synchronization.

In your network

1. Download and install the Directory Synchronization Client. See *Installing the Directory Synchronization Client*, page 7.
2. Set up the client. See *Setting up the Directory Synchronization Client*, page 11.
3. Create a configuration profile. See *Creating and Modifying Configuration Profiles*, page 15.
4. Test the client to make sure it is returning the correct data from your LDAP server. If you are an existing customer switching to LDAP for the first time, you should compare the data with that which already exists on the portal. See *Testing an update*, page 43.
5. Initiate a synchronization. See *Performing a synchronization update*, page 44.
6. Schedule automatic synchronization. We suggest that you schedule the synchronization process to run twice a day. See *Scheduling the synchronization process*, page 47.

2 | Installing the Directory Synchronization Client

Install the Directory Synchronization Client on a Windows machine with:

- Internal network access to your directory system via LDAP (Lightweight Directory Access Protocol), or network access to your cloud-based directory service
- External network access to the cloud service via HTTPS

    This access may be via a proxy server, as described in your *Getting Started Guide*.

The Directory Synchronization Client builds a local database to track changes to your source data. Use a single instance of the client to synchronize any given set of source data. Using multiple synchronization configurations, or even using multiple installations of the client, can cause data on the cloud service to be overwritten.

Prerequisites

Before starting, ensure that you have the following information:

- Your cloud portal login details
- The address of your directory server and any authentication details you might need to be able to perform searches on it.

When you are ready, continue with *Install Directory Synchronization Client software*, .

Supported operating systems

The Directory Synchronization Client has been tested and is supported on the following Windows operating systems.

- Windows 7
- Windows 10
- Windows Server 2003
- Windows Server 2008
- Windows Server 2012
- Windows Server 2016

- Windows Server 2019

Install Directory Synchronization Client software

Downloading the client software

Download the Directory Synchronization Client from the cloud portal.

1. Log onto your cloud portal account.
2. Navigate to the **Account > Identity Management** page.
3. Select the version of the synchronization tool that applies to your selected installation machine.
4. Save the installation executable to your hard drive.

Installing the software

Before starting the Directory Synchronization Client installation, ensure all other applications on the machine are closed.

1. Double-click the installation executable.
2. Click **Next** in the welcome screen.
3. In the license agreement window, select **I accept the agreement** and click **Next** to continue.
4. To change the installation location, browse to the folder where you want to install the client. The installation path changes to the folder you select, with "DirSyncClient" appended.
   - The default installation location is either "C:\Program Files\DirSyncClient" or "C:\Program Files(X86)\DirSyncClient".
   - If you don't want to install in the "DirSyncClient" sub-folder, change the directory path before clicking **Next**.
5. Select where you want the Directory Synchronization Client shortcuts to appear on your Windows Start menu. If you want the tool to be available to all users, ensure that the box is checked. Click **Next** to continue.
6. If you want an icon to appear on the desktop in addition to the one accessible from the Start menu, check the **Create a desktop icon** box.
7. Click **Next** to start the installation.
8. Click **Next** to continue.

   A message announcing the successful completion of the installation is displayed.
9. Click **Finish** to exit the installer.

After the installation

Ensure your firewalls allow the necessary ports for the Directory Synchronization Client to contact:

- Your domain controllers on either port 389 or port 3268.

- The cloud service on port 443.

- Your mail server on port 25 (if you enable mail notifications).

Upgrading from an earlier version

When you upgrade the Directory Synchronization Client from an earlier version, the configuration and settings files are also upgraded to the new version. Backup copies of the original XML settings files are made in the same location as the original XML files. The backup files are named settings.xml.old and can safely be deleted if you are not planning to revert to the previous version of the application.

The application settings are upgraded when the application is first started. The configuration settings are upgraded when a configuration is saved.

3 | Setting up the Directory Synchronization Client

To set up the Directory Synchronization Client:

1. Open the tool and select **Edit > Settings**.
2. Select **Synchronizations** in the left navigation pane.
3. Select which synchronization types you want to use:
   - **Mail** to synchronize email addresses
   - **Groups+Users** to synchronize information for users and groups
4. Click **OK**.

After selecting one or more synchronization types, continue with *Configuring log settings*, page 11.

Configuring log settings

To define the log file details and the maximum number of days that logs are retained before they are automatically deleted:

1. In the left panel of the **Sync Client settings** window, select **Log Settings**.
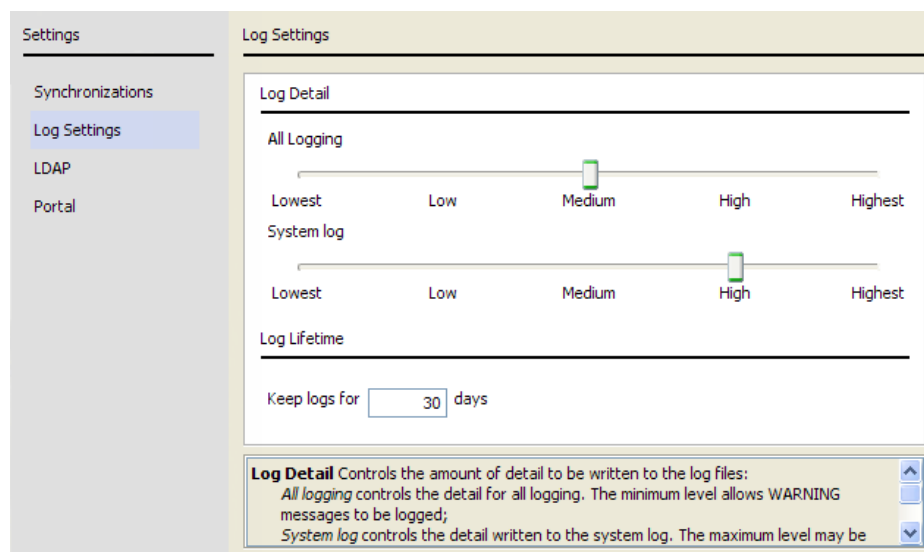
2. Under **Log Detail**, use the sliders to determine the level of detail held in the log files.

3. Under **Log Lifetime**, enter the number of days to keep log files. To hold the log files indefinitely, set this number to "0". If you do this, please check periodically that you have adequate disk space to allow new logs to be created.

4. Click **OK**.

For more information about logging, see *Directory Synchronization Client Log Files*, page 51.

To continue with your setup, see *Setting the LDAP search string*, page 12.

Setting the LDAP search string

LDAP search filters are used in 2 places in the Directory Synchronization Client:

●   Selecting which objects are returned when browsing for the search base.

●   Identifying which objects in your directory are examined, for example email address attributes or user attributes.

The LDAP page of the **Sync Client settings** window shows the first of these search filters.
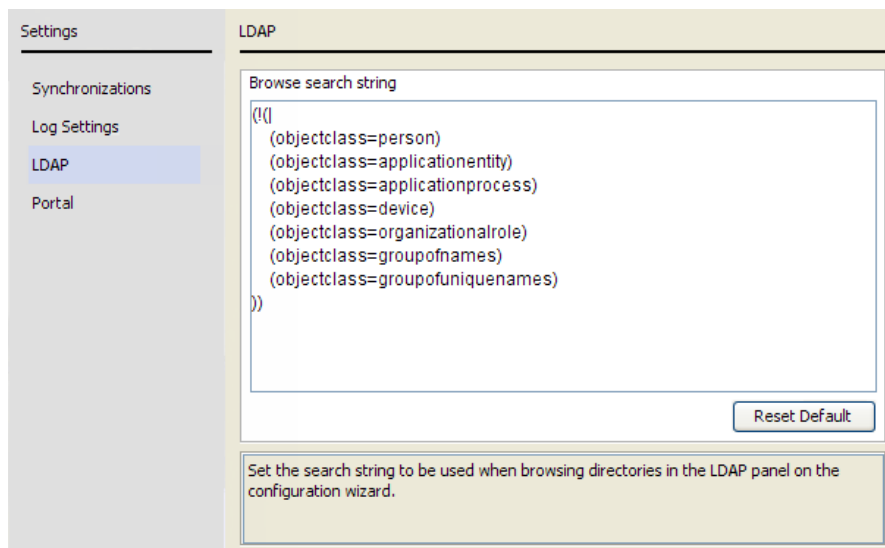
In this filter, the "!" character means **not** and the "|" character means **or**. This means the filter returns any objects that do not match any of the object classes shown in the list.

You should not need to modify this filter.

Continue with *Setting up your cloud service connection*, page 13.

Setting up your cloud service connection

> **!** **Important**
>
> When you first run the Directory Synchronization Client, you synchronize directory information to a file. This enables you to confirm that the synchronization is configured to your requirements before you upload your data to the cloud portal.

The Directory Synchronization Client connects to your cloud account using HTTPS.

Create a dedicated administrator account in the cloud portal with directory synchronization permissions to use solely for the synchronization process. Consider extending the password expiration date for this account, to avoid having to regularly update it. For more information, see "Working with LDAP Directories" in the portal Help.

To set up your connection in the **Sync Client settings** window:

1. In the left panel, select **Portal**.

2. Do not change the default value in the **Access URL** field unless instructed to do so. If necessary, click **Reset Defaults** to restore the original value of the field.

3. Enter the cloud portal administrator account user name and password, then click **Apply**.

When you click **Apply**, the Directory Synchronization Client attempts to connect to the cloud service using your settings. If you receive an error message, you may need to define proxy settings.
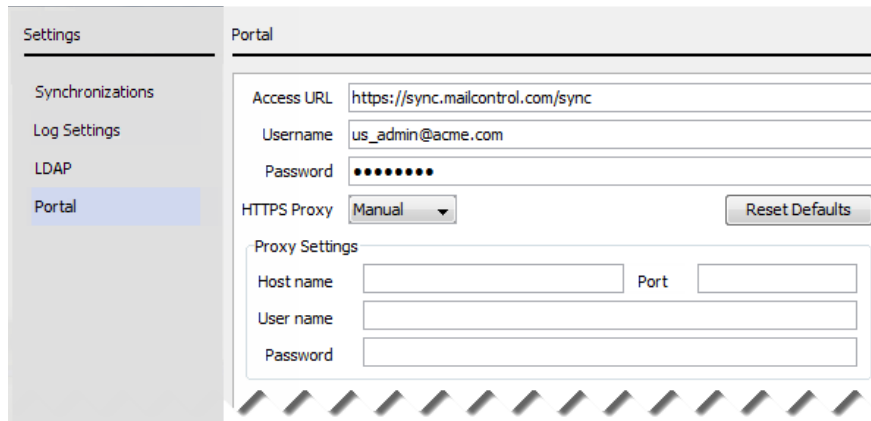
1. From the **HTTPS Proxy** drop-down list, select **Manual**.

2. Enter the connection details for your proxy server, then click **OK**.

> ✅ **Note**
> If the proxy server supplies its own certificate to decrypt and monitor traffic between the client and the server, see *Java Certificate Store*, page 59, for more information.

Continue with *Creating and Modifying Configuration Profiles*, page 15.

4   Creating and Modifying Configuration Profiles

Before you can use the Directory Synchronization Client, you need to create a configuration profile specifying details of the data source and destination systems, as described in the following sections:

> **Important**
>
> Both **Groups+Users** configurations and configurations that combine **Mail** and **Groups+Users** require setting up multiple synchronization types.
>
> Do not forget to set up **all** of your synchronization types before starting a directory synchronization!

> **Note**
>
> The example steps that follow are based on using a Microsoft Active Directory configuration. Using a generic LDAP-compliant system is covered in *Using Generic LDAP*, page 61. For information on setting up the Directory Synchronization Client with other directory services such as Microsoft Azure and Google Apps, please refer to the help file within the client. This is accessed via **Help > Contents**.

Step 1: Starting your configuration

The Directory Synchronization Client includes a configuration wizard that leads you through the process of creating or editing a configuration.

To create a configuration profile, click **New Configuration** on the client's landing page. This launches the configuration wizard to the **Name** screen.

1. Enter a unique **Configuration Name** for your profile, using alphanumeric characters.

    > **Note**
    >
    > To later modify an existing profile, select **View > Configuration**, then:
    >
    > 1. Use the **Configuration** drop-down list to make a selection.
    > 2. If your configuration contains multiple synchronization types, click the tab for the type you want to edit.
    > 3. Click **Modify**.
    >
    > To copy a profile that you are editing, enter a new configuration name and click **Save**.

1. If the **Synchronization Type** list appears, select a type.
    - This list is not displayed for those only configuring **Mail**.
    - If you are configuring **Groups+Users**, select Users first, then configure Groups afterward.

Existing configurations appear in the list with the appended text **(configured)**. To delete the current configuration for a specific synchronization type, select the entry from the drop-down list and click **Remove**.

2. Click **Next** to continue.

See .

Step 2: Selecting your data source

Note that your configuration profile name now appears at the top of the left panel, and is followed by your synchronization type in brackets.
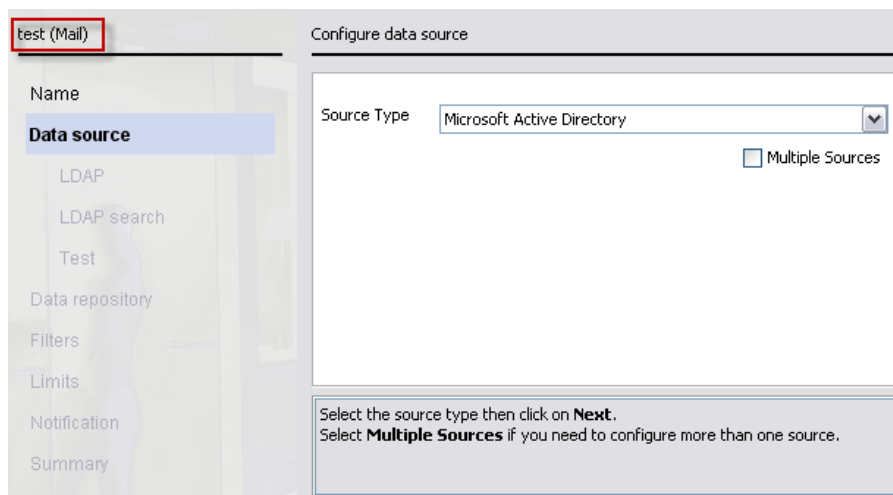
Each synchronization type can have one or more data sources.

If you use multiple directories, or want to search multiple, separate areas of a single directory, you might need to configure more than one data source. In this case, the Directory Synchronization Client consolidates your source data before sending it to the cloud service, avoiding the need for multiple synchronizations.

To configure multiple data sources:

1. Mark the **Multiple Sources** check box.
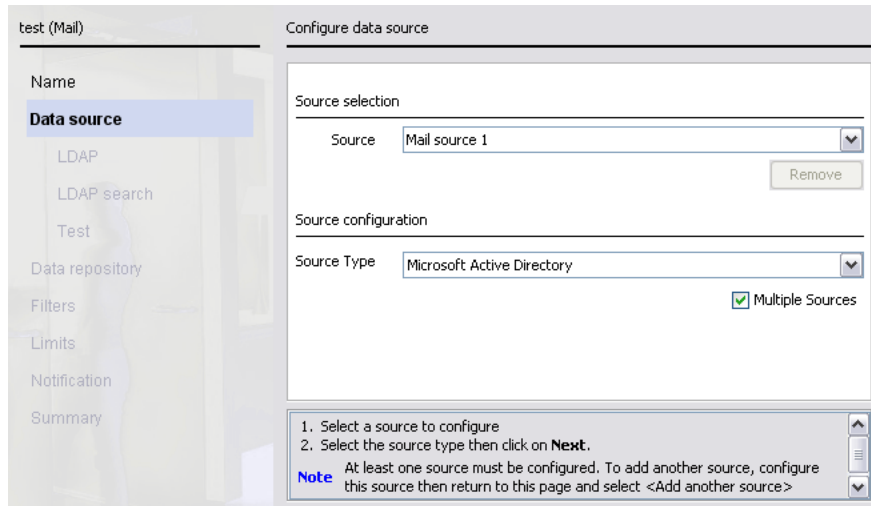
2. In the **Source** drop-down list, select **<Add another source>**.

   The source name defaults to the source type and a number, for example "Mail source 1" or "Group source 3". You can change these names.

3. Use the **Source Type** drop-down list to select a data source.

4. Use the wizard to associate each source with its own:

   ■ Server name and port

   ■ Top of the search point in the directory tree

   ■ Attributes to retrieve from the server.

   You must complete these fields for each data source before you can add further sources.

Once you have defined multiple sources, the **Multiple Sources** box is disabled. To switch back to the single data source definition window, you must delete all but one data source definition. To delete a data source definition, select it from the **Source** drop-down list and click **Remove**.

You cannot remove all data sources.

When you are finished, continue with *Step 3: Configuring your LDAP server*, page 19.

Step 3: Configuring your LDAP server

To set up your LDAP server:

1. Enter the **Host name** or IP address of your LDAP server.

2. Enter the **Port number** used for LDAP communication (**389** by default).

> **Note**
>
> If your server is an Active Directory Global Catalog server, you can specify port 3268 for a plain connection, or port 3269 for SSL. Some required attributes may be unavailable when searching the Global Catalog.

3. Select an **Authentication** option from the drop-down list:

    ■   **Anonymous**: No authentication details are required to access the LDAP server. Some servers restrict the results returned to anonymous users.

    ■   **Simple**: Enter the user name for your server, and optionally the password.

    ■   **Strong**: Enter the user name and password for your server.

    ■   **Strong/Certificate**: Enter or browse to the location of your certificate, and enter the certificate's password.

    > **Note**
    >
    > To retrieve data, we recommend that the user identity has read-only privileges equivalent to those of a domain administrator.

4. Select one of the following encryption types:

    ■   **Plain** sends unencrypted text over the connection. You cannot select this option for Strong/Certificate authentication.

- **SSL** provides cryptographically secure communication. It can either use a certificate signed by a trust point already held in the cacerts file, or a self-signed certificate that has been imported into the Java cacerts directory.

  To locate your Java cacerts directory, go to the Directory Synchronization Client installation directory, then navigate to /jre/lib/security.

- **TLS** (Transport Layer Security) offers another secured method of sending data, and requires a certificate.

5. Do one of the following:

    - Click **Advanced** to define further LDAP settings (see *Defining advanced LDAP server settings* below).

    - Click **Next** to continue to the LDAP search configuration (see *Step 4: Setting up the LDAP search configuration*, page 21).

If there are problems with your LDAP connection details, an error is shown in red at the bottom of the window. If an error occurs, click **Back** to amend your LDAP server settings.

Defining advanced LDAP server settings

Use the **Advanced** settings to specify:

- **Paging** may be used to address limits on the maximum number of results that can be returned at one time. For a server that has 220 results, a page size of 100 would retrieve entries 1 to 100 the first time, 101 to 200 the second time, and 201 to 220 the third time. The Directory Synchronization Client collects all results as if they were returned at once.

  To enable paging:

  1. In the **Paging type** drop-down list, select **page**.

  2. In the **Page Size** text box, enter the number of results to retrieve.

     The default page size is 100 results. This means that a maximum of 100 results are retrieved from the LDAP server at one time.

     > **Note**
     >
     > Not all LDAP servers support paging. If paging is unavailable, you get an error during the synchronization process.

- **Referral settings** determine how the Directory Synchronization Client handles LDAP referrals to other points on the server, or to points on other LDAP servers.
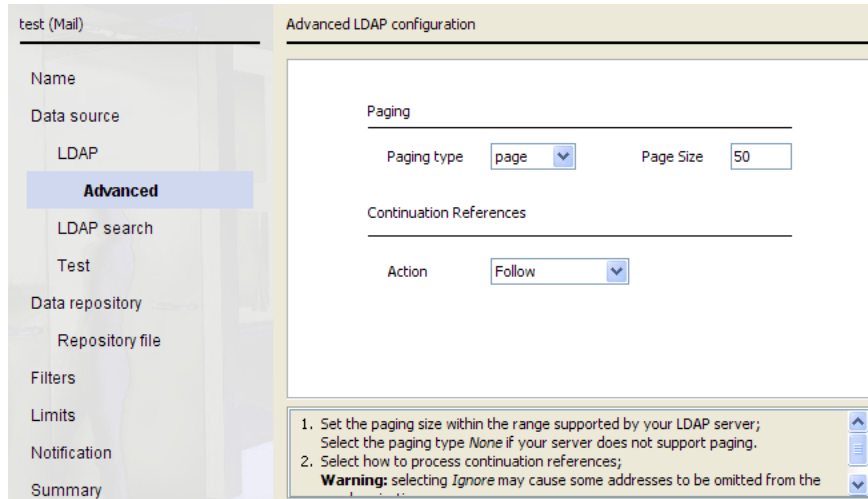
To define how to handle referrals, in the **Action** drop-down list, select one of the following:

■ **Follow** instructs the Directory Synchronization Client to follow any referrals to continue retrieving results.

■ **Ignore** instructs the Directory Synchronization Client to ignore the referral and continue the synchronization process.

■ **Abort update** instructs the Directory Synchronization Client to end the synchronization process and log this fact.

If a referral server is only intermittently available, you can set threshold limits to ensure that the update does not continue if there is a noticeable difference in the number of results returned. This means that if a referred server is not available, the threshold limit stops the update and prevents the potential loss of data from an entire server. For more information, see *Limits*, page 37.
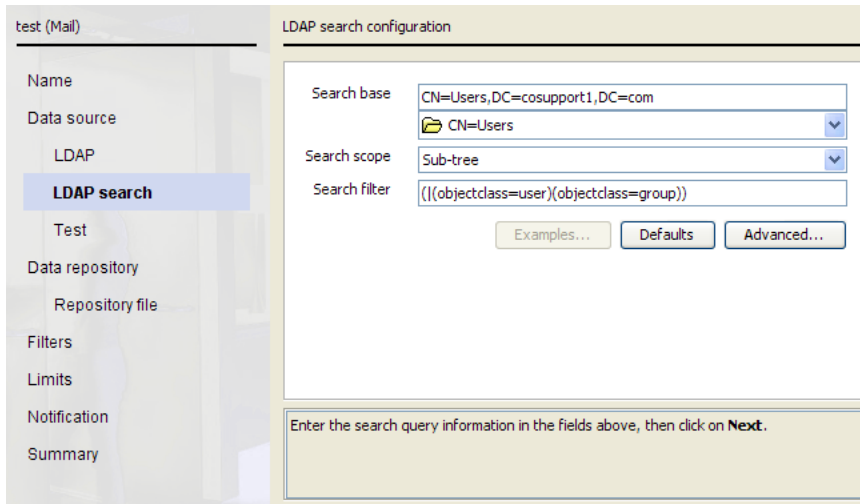
Note that the DNS name or IP address followed by the Directory Synchronization Client is the one seen by the machine running the client software. If you experience problems with the Directory Synchronization Client after referrals, make sure you can contact the referred servers via ping.

After configuring advanced LDAP settings, click **Next** to continue to the LDAP search configuration (see *Step 4: Setting up the LDAP search configuration*, page 21).

Step 4: Setting up the LDAP search configuration

The example below shows the LDAP search configuration window when configuring mail synchronization. You can select the level to search in the hierarchical tree structure on the LDAP server.

The text in the **Search filter** field may differ from the example shown below depending on your configured data source.

The following fields are common to Users, Groups, and Mail configurations:

- Search base
- Search scope
- Search filter

The **Search base** field and the drop-down list below it let you navigate through the LDAP directory. Some LDAP servers do not allow you to search for entries at their root, and you may need to enter a search base manually before you can browse further.

> **Note**
> The criteria for the returned object classes are defined in the LDAP search base filter. For more information, see *Setting the LDAP search string*, page 12.

To set up your search configuration for mail synchronization:

1. Do one of the following:

   ■ In the **Search base** field, enter your search starting point in the LDAP server tree structure.

   ■ Select an entry in the drop-down list. That entry appears in the **Search base** field.

The drop-down list now contains all the entries at the directory level specified in the **Search base** field. To move up the tree, click on the first item in the drop-down list.
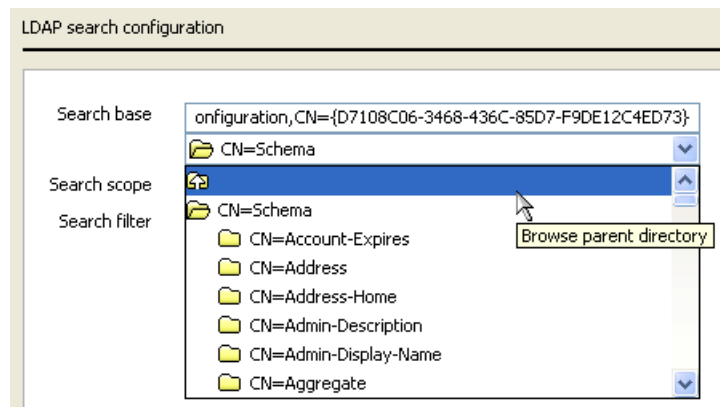
2. From the **Search scope** drop-down list, select one of the following:

    ■   **Object** searches for a single object specified by the search base.

    ■   **One level** searches for all objects at the level specified in the **Search base** field.

    ■   **Sub-tree** searches the LDAP server from the level specified in the **Search base** field downwards until the server restricts the results or the search reaches the bottom of the tree. Use this option to return the most results.

3. The **Search filter** field defines the type of object to return data on. You can leave the default filter or create your own.

    See *Search query filter*, page 23, for a description of the filter syntax and how to specify a different search filter from the ones available.

To revert to the original search settings or to return to the top of your LDAP server's tree, click **Defaults**.

To define specific attribute settings for your current synchronization type, click **Advanced**, then see:

●   *Defining mail attributes*, page 24

●   *Defining group attributes*, page 25

●   *Defining user attributes*, page 28

When you are finished, continue with *Step 5: Checking your search results*, page 30.

Search query filter

LDAP search filters are defined using a notation that is fully described in RFC 2254 "The String Representation of LDAP Search Filters". You can see this document at http://rfc.net/rfc2254.html.

To establish your own filters, you also need an understanding of your directory's schema. The schema defines the objects and their attributes that constitute your directory content.

Examples

The Directory Synchronization Client lets you define a search query filter that targets the objects in your directory that are examined for email address attributes.

If you want to include all objects in your search query, enter the following in the **Search filter** field:

```
(objectclass=*)
```

The following filter includes all Microsoft Exchange users that are currently enabled:

```
(&(objectclass=user)(msexchuserAccountcontrol=0))
```

The following filter includes all objects that define users and groups. This may include both security groups and mailing lists.
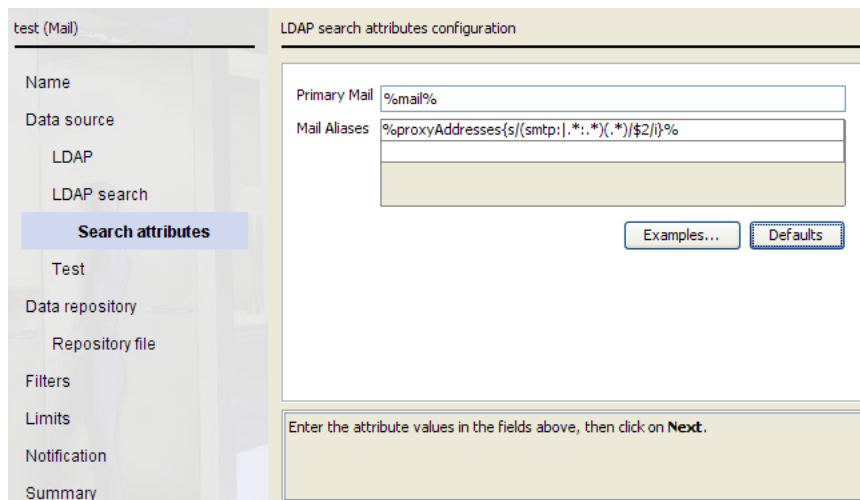
```
(|(objectclass=user)(objectclass=group))
```

If you want to exclude the system mailbox objects in Microsoft Exchange from the search described above, you could modify the filter as follows:

```
(&(|(objectclass=user)(objectclass=group))
(!(cn=SystemMailbox*)))
```

Defining mail attributes

For mail configurations, click **Advanced** to display and edit the mail search attributes. The default settings for these attributes are taken from your data source.

The **Primary Mail** field contains the mail address attributes within the object returned by the search filter.

If your LDAP data does not include users' email addresses, you can change the default attribute for the primary mail value in the Directory Synchronization Client as follows:
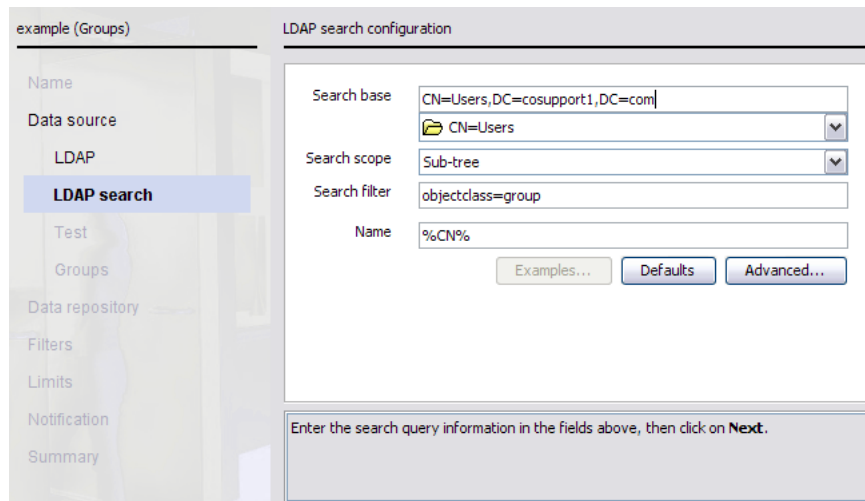
1. When creating or modifying the Users part of your configuration profile, go to the **Data source > LDAP search** page in the wizard. Click **Advanced** to display the Search attributes page.

2. In the Primary Mail field, replace %mail% with another attribute.

   For example, you could use %userPrincipalName% if configured, or create a fake email address using the sAMAccountName such as %sAMAccountName%@mydomain.com.

Optionally, you can enter alternatives to the mail attribute in the **Mail Aliases** field. If an email address returned from this attribute is prepended with "smtp:", this is automatically removed.

To view and select from a list of mail attribute examples, click in the **Mail Aliases** field and then click **Examples**.

Defining group attributes

For group configurations, there is an additional field on the LDAP search configuration window. The **Name** field defines a rule for constructing a textual name that is used to represent individual users and groups. The name can be constructed from other LDAP attributes using simple template replacement strings.

Note

If you want to synchronize groups with the same name from different domains (for example, **domain1/Admins** and **domain2/admins**), you must change the string in the Name field from the default **%CN%** to **%DC%/%CN%**.

Attribute names are delimited by percent (%) symbols. The special attributes **DN[n]** and **DC[n]** allow part of the object class distinguished name to be used. Anything not enclosed between % symbols is treated as literal text.

The number (n) following the DN or DC attribute is an index, starting from 1, from the least significant component. When used with DN, the index refers to all components of the distinguished name. When used with DC, the index refers to only the DC components of the distinguished name. If the number exceeds the actual number of components, an empty string is substituted. If *n* is a negative value, it refers to the components starting with the most significant component first.

For example, the table below shows how different replacement string templates would resolve for the following object class:

```
dn: cn=Marketing, ou=Security, dc=Forcepoint, dc=com
objectClass: group
Name: SecureMarketing
SamAccountName: SecurityMarketingServices
```

Template	Resolves To
%Name%\%DN[-2]%.%DN[-1]%	SecureMarketing\Security.Marketing
%DN[1]%\%SamAccountName%	com\SecurityMarketingServices
%Name%\%DC[1]%SamAccountName%	Com\SecurityMarketingServices
%DC[-1]%\%SamAccountName%	Forcepoint\SecurityMarketingServices

To see a list of examples that you can use for the name template, click **Examples**.

Click **Advanced** to edit the group attributes.
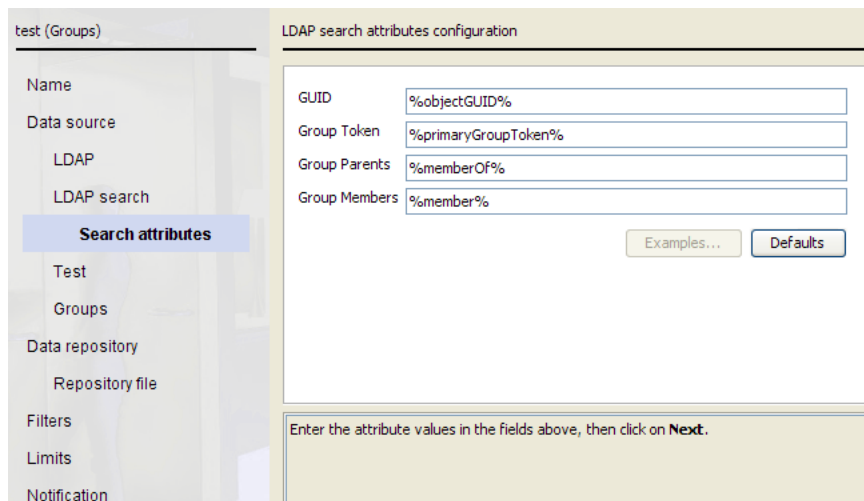
You can edit the following attributes:

- **GUID** is a unique identifier maintained by the LDAP server. Use this attribute if it is available on your server. Microsoft Active Directory supports GUID, but it is not supported by all servers. If you omit this attribute, the Directory Synchronization Client derives an identifier from the distinguished name (DN) of the object class.

  The disadvantage of using a DN is that if the group is renamed, the group entry is removed and re-added instead of modified. This means that any group associations in the cloud service are broken and must be re-established.

- **Group Token** is an optional attribute that holds the number this group is in. The value may be referred to by the "Primary Group attribute" in the user object class settings. If a user's primary group is set to a particular group token, then the user is part of that group. The group token is specific to Active Directory so may be unavailable in other directories. If unavailable, it should be left blank.

- **Group Parents** is used to relate a group to its parent group, if it exists. The optional attribute retrieved from the directory may consist of a single DN that contains the parent group.

- **Group Members** is a multiple-value attribute that holds the users (in DN form) who are part of this group.

  Active Directory maintains membership lists on both group and user objects so the Group Members attribute of the group object class lists all the users for the group and the Other Groups and Primary Group attributes list all the groups to which the user belongs. In theory, these should be equivalent. In practice, when the directory is modified, some tools may update one list but not the other. Specifying both attributes causes the lists to be merged.

> **Note**
>
> Group membership can be represented in the directory by use of either:
>
> **Group Members**: a list of users/groups belonging to a group.
>
> **Group Membership**: a list of groups to which a group/ user belongs.
>
> The client allows either convention. For users to be correctly associated with groups one of these must be specified. The Group Membership attribute is labeled Group Parents for group objects and both Primary Group and Other Groups for user objects.

Defining user attributes

> **Note**
>
> Be sure to set up a users search filter that includes the users for the groups you are synchronizing. Before you synchronize with the cloud service for the first time, test your synchronization by sending the results to a local file and carefully check that the contents match your requirements.

The **NTLM Identity** field defines a template for constructing the NTLM identity of the user. The default is the format "domain\username".

The **Name** field defines a template for constructing a name that is used by the cloud service to identify users. This is not required if the Relative Distinguished Name (RDN) of the user is a Common Name (CN), as that will be automatically be included. The Distinguished Name (DN) can be seen on the test page by selecting **Show Detail** or hovering over one of the results: the RDN is the first naming component of the DN.

Including only enabled user accounts

In Active Directory, it is possible to mark user accounts as disabled. You might do this if an employee is away for a short period of time. If you want to prevent disabled user accounts from being uploaded to the cloud service, you can filter them out of the search by searching for only enabled user accounts. At the next synchronization, any disabled user accounts are removed from the cloud service.

To search for enabled accounts only, add the following to the filter in the **Search filter** field:

```
(!(userAccountControl:1.2.840.113556.1.4.803:=2))
```
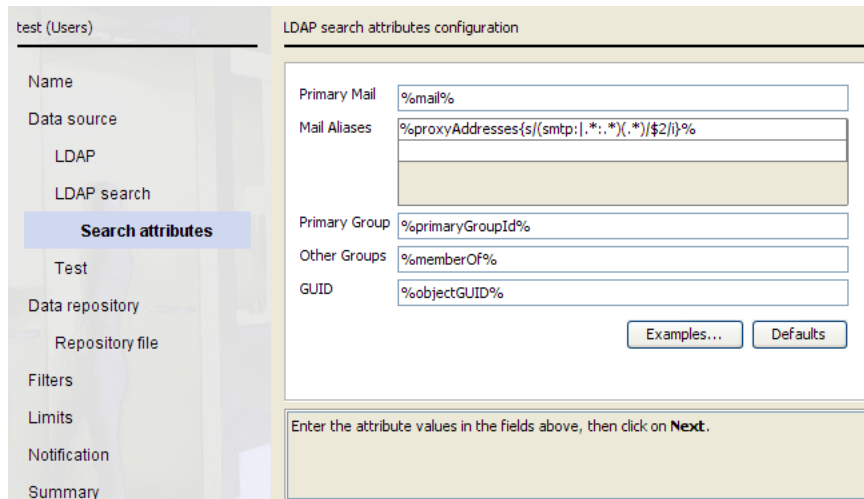
The complete filter might then look like this:

```
(&(objectCategory=person)(objectClass=user)(!(userAccountCon
trol:1.2.840.113556.1.4.803:=2)))
```

Advanced user attributes

Click **Advanced** to edit the user attributes.

You can edit the following attributes:

● **Primary Mail** is used to retrieve a user's email address.

● **Mail Aliases** are any aliases set up for a user. To view and select from a list of mail attribute examples, click in the **Mail Aliases** field and then click **Examples**.

● **Primary Group** is the token number attributed to a user. If this matches a group's **Group Token** value, it places this user in that group. This attribute should be considered as an extension to the **Other Groups** attribute for placing a user in a particular group. Not all LDAP directories offer support for it.

● **Other Groups** is the attribute name that describes the group or groups this user belongs to. For Active Directories this is symmetrical to the **Group Parent** attribute for the group object class that points from each group to its users. If you omit this attribute, or your directory does not support this feature, the Directory Synchronization Client searches for each user in the entire list of groups.

For users to be correctly associated with groups, one of the following must be true:

a. The **Group Parent** attribute exists, and the user **GUID** and group **GUID** attributes do not.

b. The **Group Members** attribute exists.

● **GUID** is a unique identifier assigned to each user in a similar manner to the GUID attribute for groups. If you omit this attribute, you should also omit the group GUID attribute.

Step 5: Checking your search results

Click **Next** to test your search settings.
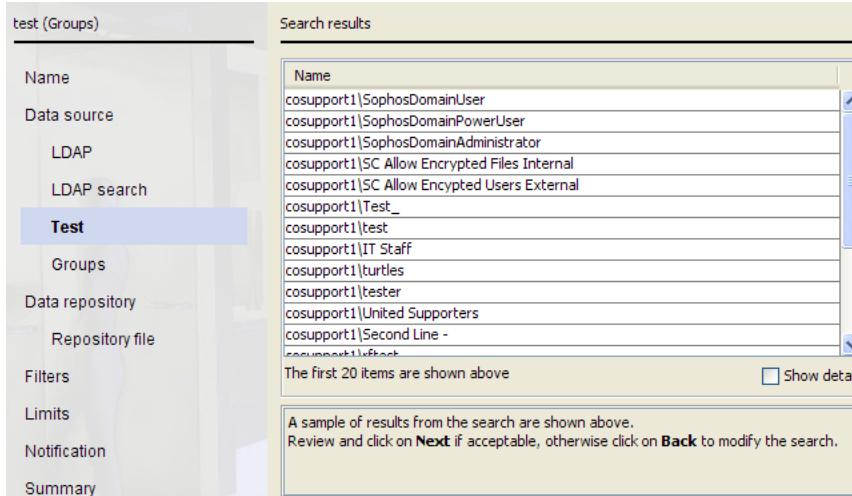
If you have changed any of the default attributes used for the LDAP search, you can use the test window to confirm that you have correctly retrieved the attribute you were expecting. To view full details of groups and users on this window, check the **Show detail** box. (This option is not available for Mail configurations.)

You can reorder columns in the table by clicking and dragging the top of the column.
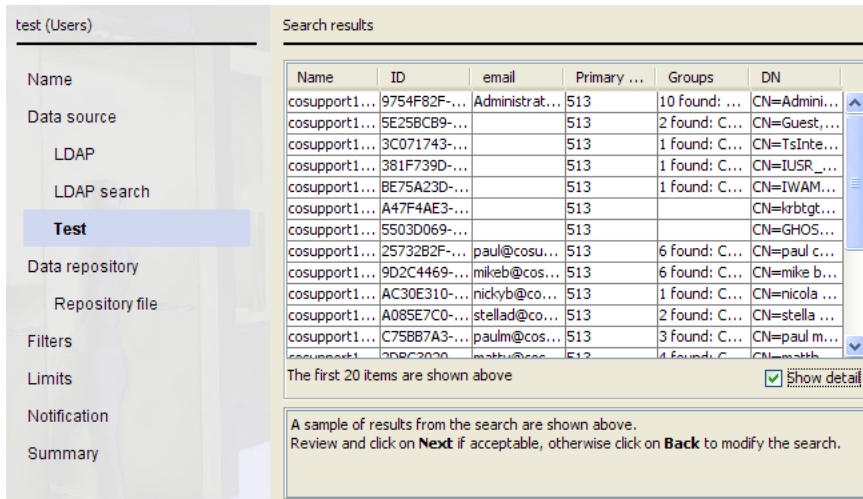
For groups, each line includes:

- The result from the name template after it has been changed using any template rules.
- The GUID. If there is no group **GUID** attribute, this is derived from the DN.
- The Group Token, retrieved using the **Group Token** attribute.

- The DN automatically retrieved by the Directory Synchronization Client.
- The number of parents that this group belongs to (normally 0 or 1) and the DN of the first of these groups. This is retrieved using the **Group Parents** attribute.
- The number of users in this group and the DN of the first of these users. This is retrieved using the **Group Members** attribute.

For users, each line includes:

- The result from the name template after it has been changed using any template rules.
- The GUID. If there is no user GUID attribute, this is derived from the DN.
- The email address retrieved using the **Primary Mail** attribute.
- The Primary Group retrieved using the **Primary Group** attribute.
- Groups that the user belongs to.
- The DN automatically retrieved by the Directory Synchronization Client.

For both groups and users, if the name, the GUID, or the DN is blank, you should correct the attribute names before starting a synchronization.

If you see no results in this window, check that:

- The source type on the **Configure data source** window is correct. For information on configuring multiple data sources and advanced details, see *Step 2: Selecting your data source*, page 17.
- The **Search scope** field on the **LDAP search configuration** window is set to **Sub-tree**. This returns the most results.
- The **Search base** field on the **LDAP search configuration** window is set to a suitable level in your LDAP server's hierarchy to find the mail addresses or groups and users that you want.
- The location specified in the **Search base** field exists in the LDAP server directory. If in doubt, return to the top of the LDAP server tree and then navigate to the location you want.
- You haven't changed the **Search filter** field. Click **Defaults** to reset this field setting.
- Your authentication settings are sufficient to return details from the LDAP server. If you selected **anonymous** from the **Authentication** drop-down list and no results are returned, try selecting **simple** and entering a username and password.

For users and groups, the search results display the names constructed from the **Name template** field. Check that these are representative names. For example, for users in a Microsoft environment, the names should represent the "domain\username" identity of individual users.

Click **Next** to continue. See *Step 6: Selecting groups for synchronization*, page 32.

Step 6: Selecting groups for synchronization

For the group synchronization type, you must specify which groups in your data source are to be synchronized.
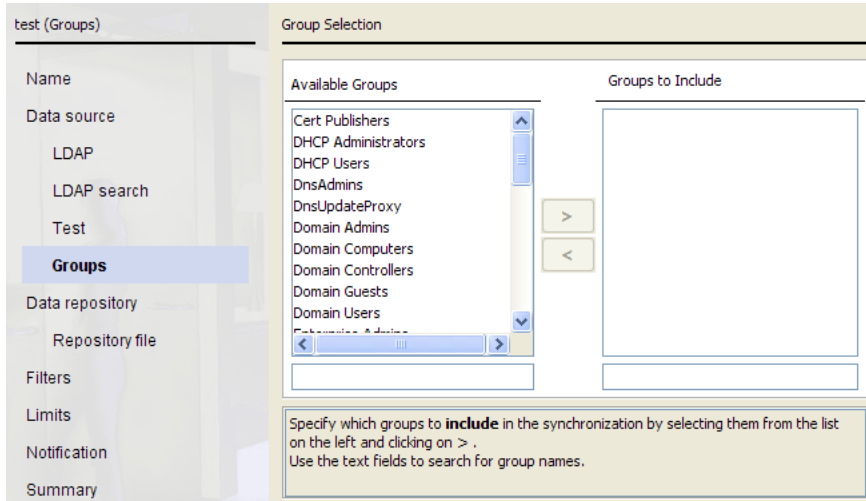
By default, no groups are synchronized. To include a group in the synchronization, select the group in the **Available Groups** list, and then click > to move it to the **Groups to Include** list.

> **Note**
>
> You do not have to synchronize all of your Active Directory groups. You only need to select the groups that you plan to use in policy assignment and/or exceptions in the cloud portal.

To filter a group list, enter text in the field below the list. For example, in the screenshot above, if you enter "dhcp" in the field below the **Available Groups** list, only the groups DHCP Administrators and DCHP Users are displayed.

Step 7: Setting up a data repository
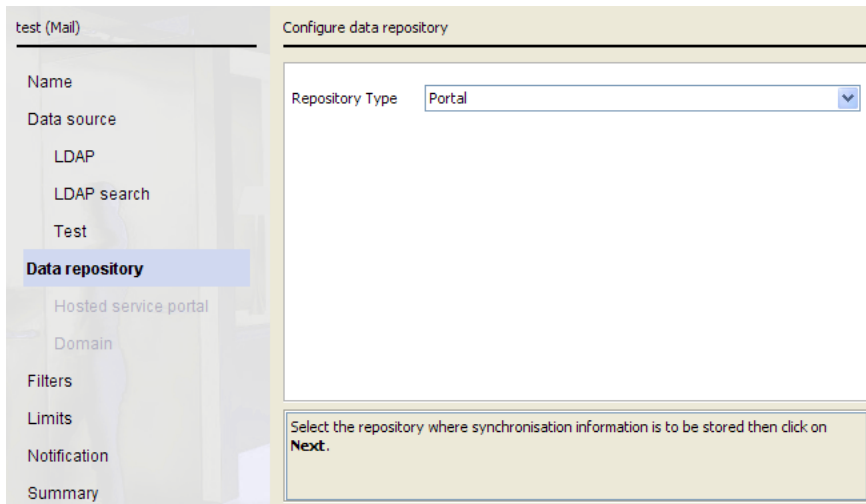
Select an option from the **Repository Type** drop-down list, then click **Next**:

- **Portal**: Your data is synchronized to and held on the cloud portal.
- **File**: Your data is held in a text file on your local system. See *Selecting a file repository location*, page 34, for details of the next step.

> **Note**
>
> Select **File** when you are setting up the Directory Synchronization Client for the first time. This lets you test the connection to your LDAP server and ensure the results are correct before setting up the connection to the portal.

When you are finished with your repository settings, see *Step 8: Optional settings*, page 35.

Customizing cloud service connection details

By default, the client uses the cloud connection details that you entered in the **Sync Client settings** window. You can optionally override your default cloud service username and password details, however, by marking **Custom account details for this synchronization** and entering the new details.

Configuring mail domains

If you are configuring a mail synchronization, you can upload email addresses from all available domains, or only the ones you specify. The list of domains is obtained

from the cloud portal. There may be a slight delay as the cloud service is contacted to retrieve your domain details.
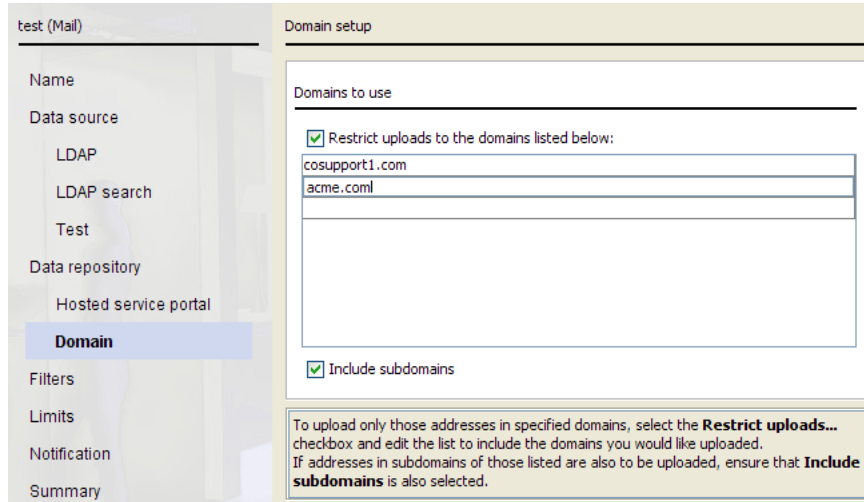
To restrict the domains you use:

1. Mark **Restrict uploads to the domains listed below**.
2. Edit the list of domains so that it includes only the ones you want to use.

    To include the subdomains of selected domains in the upload, mark **Include subdomains**.

Selecting a file repository location

If you select **File**, also specify the location of the data repository files. There are 3 files: a list of current email addresses, additions since the last synchronization, and deletions since the last synchronization.

The file holding the current list represents the state of the repository after the synchronization. The additions and deletions files show what would be sent to the cloud service in order to adjust the repository based on the source data.
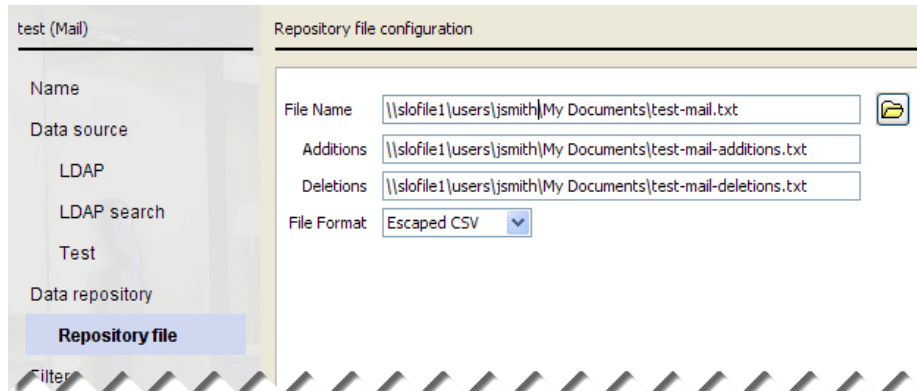
1. Click the **Browse** icon next to the **File Name** field, then browse to the location where you want to store your mail synchronization files.

2. Enter a file name and click Select.

   The Additions and Deletions fields are automatically filled in with file names based on the name you entered. For example, if your main file is named **test-mail.txt**, the additions file is named **test-mail-additions.txt**, and the deletions file is named **test-mail-deletions.txt**.

3. From the **File Format** drop-down list, select one of the following:

   ■ **Escaped CSV** creates a file in Comma Separated Variable (CSV) format. For mail, the file has one email address per line. For groups, each line includes the group name, the users in the group, the GUID, and any parent groups. For users, each line includes the username, any mail aliases, the GUID, the email address, and the member groups.

   ■ **LDIF** creates a file in the LDAP Data Interchange Format (LDIF). Each directory entry is represented as a record, including the dn and objectClass attributes.

4. Click **Next** to continue.

Step 8: Optional settings

At this point you have completed the basic source and destination configuration required for a single synchronization type. You can now set up any of the following optional settings:

● *Filters* exclude or change details found on your LDAP server before they are written to the cloud service.

● *Limits* allow you to protect an existing synchronization against accidental modification, for example if the Directory Synchronization Client is configured incorrectly or if your LDAP server returns incorrect results.

● *Notification settings* request the Directory Synchronization Client to send an email notification on completion of synchronization.

After completing any optional settings, click **Next**, then continue with *Step 9: Verifying your settings*, page 40.

Filters

Use the Filters window to:

● Exclude retrieved email addresses or names.

● Modify email addresses for mail synchronizations before the addresses are written to the destination data repository.

● Modify names for groups and users synchronizations before the names are written to the destination data repository.
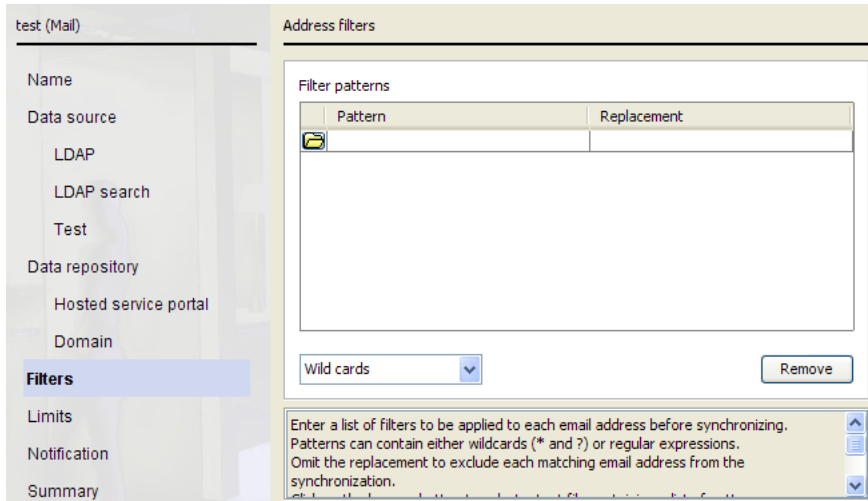
Each line can contain a different pattern to match against.

To exclude an entry, enter the pattern in the **Pattern** column and leave the **Replacement** column blank. When entries are found in the data source file or LDAP server, they are checked against patterns in this column and removed if a match occurs.

To modify entries, enter replacement text in the **Replacement** column. The replacement rule is applied against the matching pattern in the **Pattern** column.

To specify a text file of patterns, click the **Browse** icon at the start of the next blank line, and browse to the file. The file should contain a list of patterns, with each entry separated by a new line.

To remove a line from the filters list, select the line and click **Remove**.

The rules for matching the entries against the filters are determined by the drop-down list below the filters list. You can select one of the following:

● **Wild cards** match characters as follows:

■ Use an asterisk (*) to match 1 or more characters. For example, "*y@acme.com" matches the email address "andy@acme.com", but not the address "john@acme.com."

■ Use a question mark (?) to match a single character. For example, "?andy@acme.com" matches the email address "mandy@acme.com" and "sandy@acme.com" but not "andy@acme.com."

If you are modifying an entry, only the result of the first matching asterisk can be used in the replacement. In the above example, "*y@acme.com" with a replacement field of "*i@acme.com" would match "andy@acme.com" and replace it with "andi@acme.com."

● **Literal text** matches on the precise text in the filters list.

- **Regular expressions** are used for complex pattern matching and replacements. For a detailed description of regular expressions, see *Standard Regular Expression Strings*, page 55.

> **Note**
>
> Changing this setting affects how all patterns are implemented, as the client requires all filters to be of the same type. If you have previously used wildcard filters and you need to convert them to regular expressions, replace "*" with "(.*)" and ? with . in the pattern, and "*" with "$1" in the replacement string.

Limits

We strongly recommend that you use threshold limits to provide a safeguard against accidental deletion of entries in your data repository. Threshold limits warn you when the number of added or deleted entries exceeds a specified amount. This protects you from mistakes in your configuration, especially if you are using filters.

You can enter threshold limits as absolute numbers or as a percentage. Percentages must be followed by a percent (%) symbol.

If your LDAP servers contain referrals, we recommend you set threshold limits to ensure that if a referred server is not available, the update does not continue. for more information about referrals, see *Defining advanced LDAP server settings*, page 20.

- **Maximum added entries** is the largest number of entries from your data source that can be added. If this number is exceeded, the synchronization process is aborted.

- **Maximum deleted entries** is the largest number of entries from your data source that can be deleted. If this number is exceeded, the synchronization process is aborted. A percentage greater than 100% is treated as 100%.

- **Maximum changed entries** is the largest number of entries from your data source that can be modified. If this number is exceeded, the synchronization process is aborted. Note that this is not available for mail synchronizations.

- **Minimum replacement entries** is for actions that cause the entire contents of the repository to be overwritten. This is the minimum number of entries to be accepted. If fewer entries are found in the data source, the synchronization process is aborted.

If any of your threshold limits are exceeded, the Directory Synchronization Client displays a message asking if you wish to force the update.

When the Directory Synchronization Client is running in command line mode and a threshold limit is exceeded, the synchronization is not performed. If you have set up email notifications, you are notified of this.

Notification settings

After every synchronization process, you can automatically send an email message containing a summary of the process and whether or not it was successful.

The client generates an email message for each synchronization type, and lists all items that have been added, deleted, and updated, including totals.

We recommend you set up a notification summary for the following reasons:

- If you intend to schedule the Directory Synchronization Client to run automatically, the email notification gives you a historical record of your synchronizations.

- You can use the emails to confirm that synchronization with the cloud service completed successfully. The total numbers of additions and deletions in the email messages should match the totals in the cloud service logs.

To set up notifications:

1. From the **Email notifications** drop-down list, select the type of notification you want to send.

   ■ **Summary** sends a summary of the synchronization process.

   ■ **Summary + Log (WARNING)** sends a summary and all log messages with the level WARNING or higher.

   ■ **Summary + Log (INFO)** sends a summary and all log messages with the level INFO or higher.

   ■ **Summary + Log (FINE)** sends a summary and all log messages with the level FINE or higher.

   For a definition of log message levels, see *Directory Synchronization Client Log Files*, page 51.

2. Enter your mail server name (for example smtp.acme.com).

3. The Directory Synchronization Client supplies default message subject headers for success and failure notifications. Edit these if required.

4. In the **From** field, enter the address to use for originating the email.

5. In the **To** and **CC** fields, enter the email addresses to send the summary email to. Separate multiple recipients by commas.

An example email summary notification might look like this:

```
Sync Client Mail Synchronization Report
Replace operation to file completed
Time:        Fri Jun 27 10:38:12 GMT 2008
Host:        bloggs.acme.com
User:        fredbloggs
Configuration:example

Updated domains
    acme.com

Up-to-date domains
    test.acme.com

Unknown domains
    None

Updates
    4 additions
        lyndonb@acme.com
        stellad@acme.com
        miken@acme.com
        nickyb@acme.com
    2 deletions
        philm@acme.com
        waynek@acme.com

Invalid addresses
    None

Failed updates
    None
```

```
                     Addresses in domains not configured on the repository
                          None
```

Step 9: Verifying your settings

In the Summary window:
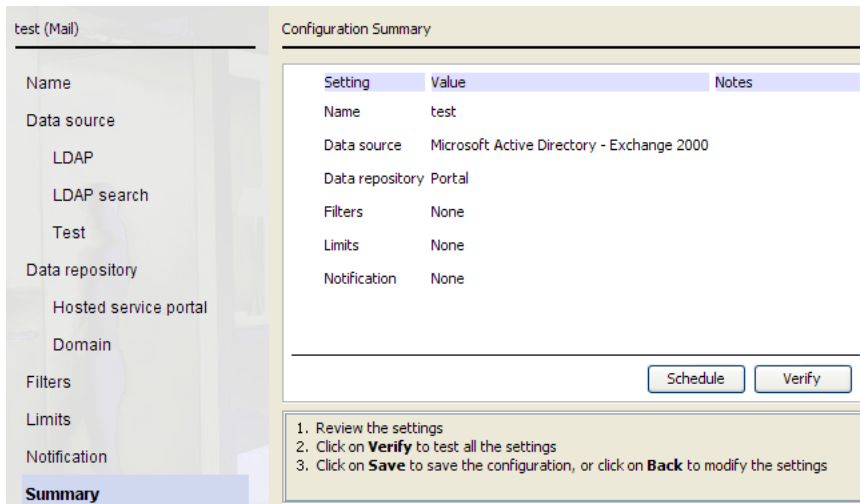
1. To confirm the settings you have chosen for this synchronization type, click **Verify**.
2. Click **Yes** to confirm.
    -   As each test is performed, the icon to the left of each setting changes from an hourglass to a green check mark. If any settings is marked with a red X, select the relevant option in the left panel to correct it.
    -   If you have multiple sources, each source is checked in turn.
    -   If your data repository is set to **Portal**, the Directory Synchronization Client contacts the cloud-based server during this process and, if configured, sends a test email notification.
3. To save your configuration profile, click **Save**.
4. Click **Next** to continue.

After completing configuration of one synchronization type, you can optionally configure another synchronization type.

-   To set up another synchronization type, see *Step 10: Setting up another synchronization type*, page 41.

    If you have set up a groups configuration, you must now set up a corresponding users configuration.

-   If you are finished with your configuration, click **Finish** to exit the wizard.

Step 10: Setting up another synchronization type

You can now set up multiple synchronization types as part of the same configuration profile.

If you have set up a groups configuration, you must also set up a users configuration.
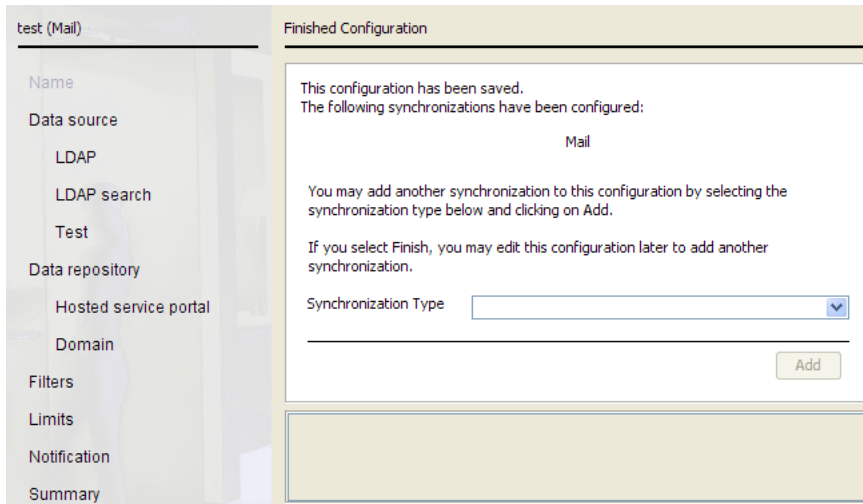
To add another synchronization type under this configuration profile, select the type from the **Synchronization Type** drop-down list, then click **Add**.

To finish this configuration, click **Finish** to exit the wizard.

5 | Synchronizing with the Cloud Service

Once you have created your configuration, you are ready to test and then synchronize your data.

> **!** **Important**
> Before you run the synchronization process for the first time, make sure your LDAP data is in the format that you want. If you are an existing Forcepoint cloud service customer switching to LDAP synchronization, review your current portal data structure. For more information, see "Working with LDAP Directories" in the portal Help.

Testing an update

Before synchronizing with the cloud portal, make sure synchronization returns the correct data from your LDAP server.

1. Open the Directory Synchronization Client.
2. Select **View > Configuration**.
3. From the **Configuration** drop-down list, select your configuration.
4. Click **Test Update.**

The Directory Synchronization Client looks at the email addresses, users, and groups on your LDAP server or source file, and lists the additions, removals, or exclusions without changing any details in your repository.

● The first time you run a test update, you should only see additions.
● Click a column entry to see additional information about the entry at the bottom of the window.

For users, this includes the user's email address and group membership.
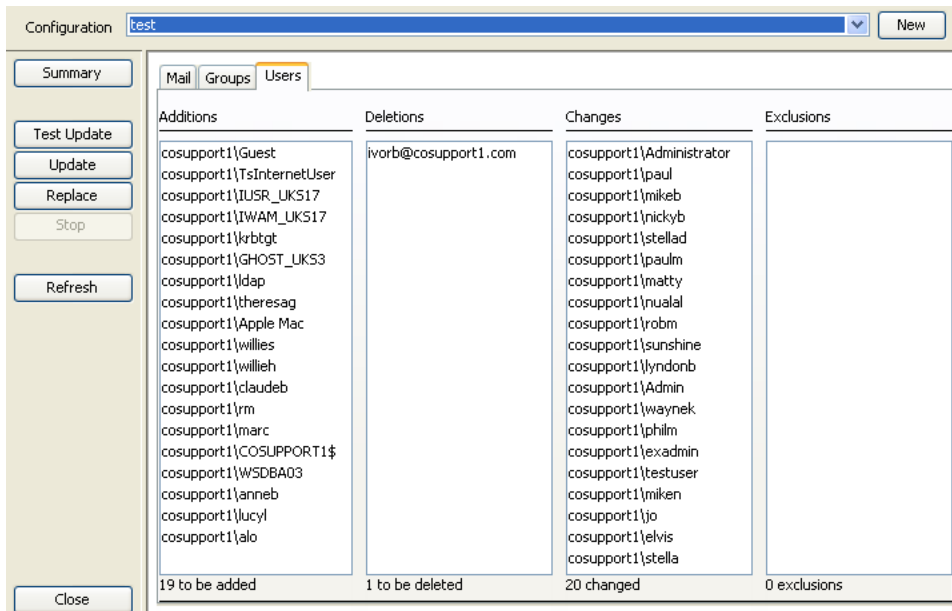
To view the test results at any time, select **View > Test Results**.

If the results of the test are as expected, continue with *Performing a synchronization update*, page 44.

Performing a synchronization update

To launch the synchronization process:

1. Open the Directory Synchronization Client.
2. Select **View > Configuration**.

3. Select your **Configuration**, then click **Update**.
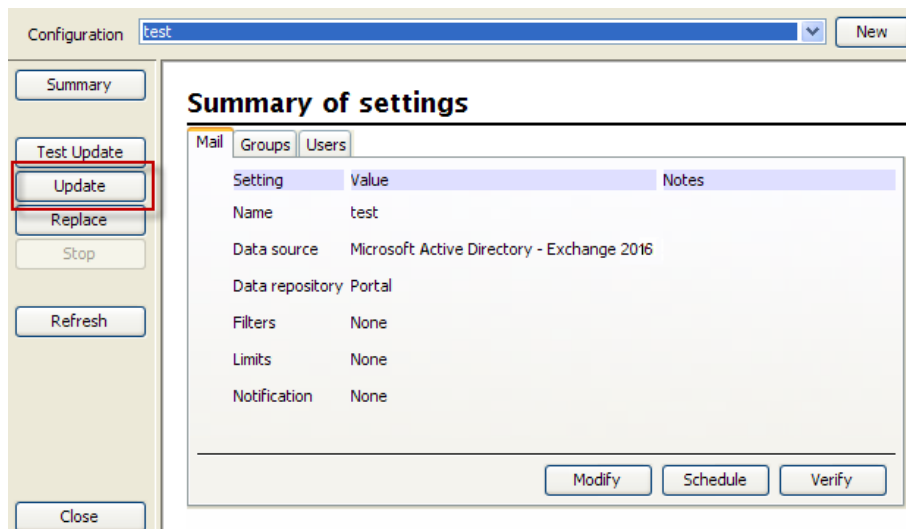
4. The Directory Synchronization Client compares the details held in your local change tracking database and only sends additions and removals. The incremental update is very efficient, and preserves any unchanged data held by the service.

   To abort the synchronization, click **Stop**.

While the client is querying your LDAP servers, the portal shows a status of "In progress." Once the extraction is complete and processed, the data is synchronized to the cloud service and the portal status is updated to show the number of amended records.

Once the update is complete, you can click on the **Mail**, **Groups**, and **Users** tabs to view the results. Click on a title bar to sort the list by that field.

To return to the configuration summary, click **Summary**.

Synchronization errors

If you receive a "Sync not enabled for account" error when you attempt to perform a synchronization:

● Use the **Account > Identity Management** page in the cloud portal to verify that **Enable directory synchronization** is selected.

● Use the **Account > Contacts** page in the cloud portal to verify that the administrator account used to connect to the cloud service has **Directory Synchronization** permissions. To do this, click the account user name, then click **Edit** and check the Account Permissions.

Replacing and refreshing data

Replacing your data in the cloud service

Periodically, you may want to do a full synchronization of your LDAP directory, rather than relying exclusively on incremental updates.

> ⚠ **Warning**
> This action overwrites all data in the cloud service with the contents of your LDAP directory.

To re-create your local change tracking database and resend all data to the cloud service, click **Replace**.

Refreshing your local data

To retrieve all data from the cloud service and re-create your local change tracking database, click **Refresh**. This allows subsequent updates to be based upon calculating changes from a current copy of the data held by the cloud service.
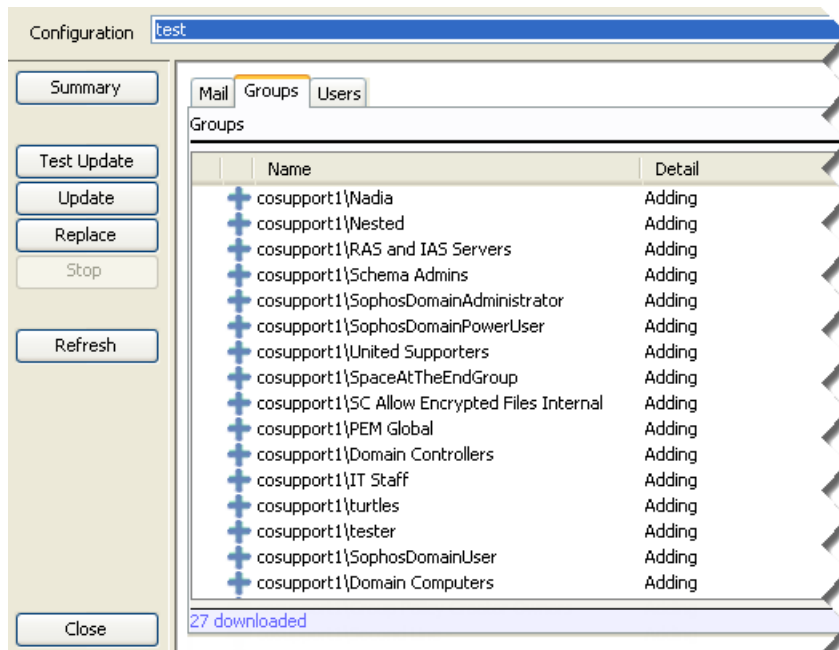
Scheduling the synchronization process

Once you have run the initial synchronization, you should set up a scheduled service to run automatic updates in the background.

Make sure Windows Task Scheduler is running

To schedule synchronization updates from the Directory Synchronization Client, first ensure that the Windows Task Scheduler service is started. To check this:

1. Select **Start > Administrative Tools > Services**.

2. In the Services window, scroll down to Task Scheduler.
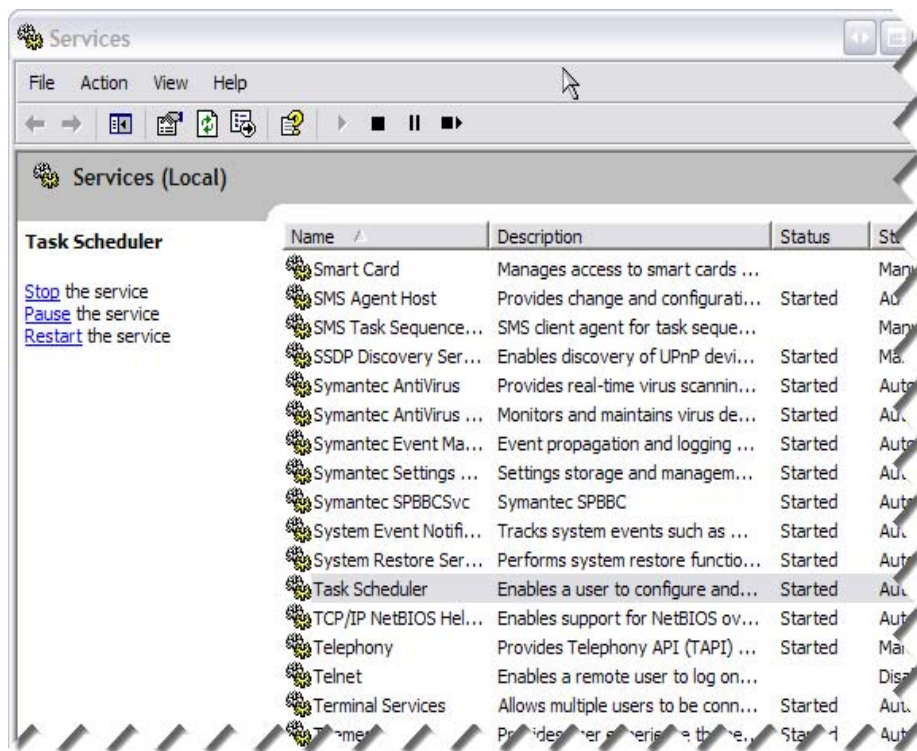
If the status is Started, you need do nothing. Otherwise, do one of the following:

- If the status column is empty, right-click Task Scheduler and select **Start.**

- If the status is Paused, right-click Task Scheduler and select **Resume**.

Create the synchronization schedule

To create a synchronization schedule in the Directory Synchronization Client:

1. In the Configuration window, click **Summary**, then click **Schedule**.

2. In the **Run at** field, enter the time you want the process to run.

3. Specify how often you want to run the process (in months, weeks, days, or hours).

4. Select whether to run the process on any day within your chosen schedule, or on specific days.

5. If you are running the process on specific days, do one of the following:

   ■ Mark **Day of Month**, then click the dates on which to run the process.

   ■ Mark **Day of Week**, then click the days on which to run the process.

   ■ Mark both boxes and then click the dates and days you want. For example, you can run the process each Friday, and on the 30th of each month.
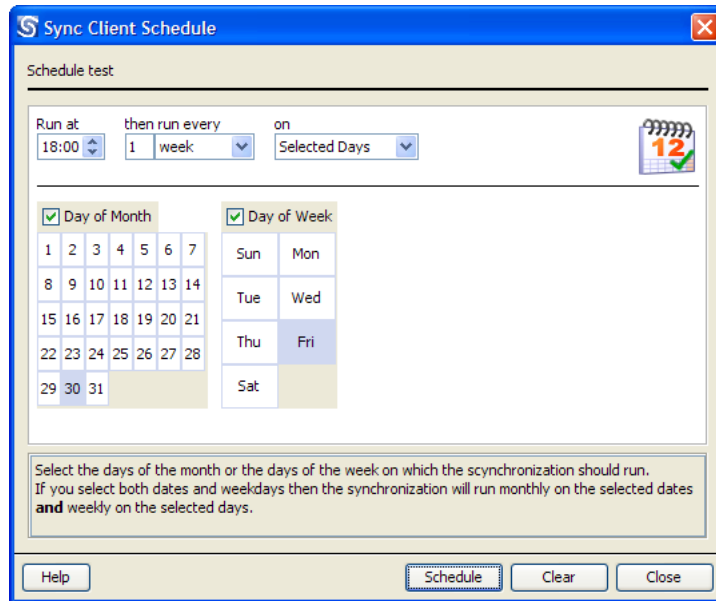
6. Click **Schedule**. The calendar icon changes from gray to color to show that the synchronization is scheduled.

> **Note**
> If the Directory Synchronization Client is running at the same time as the scheduled task, the synchronization will fail.

Removing the synchronization schedule

You can edit your synchronization schedule at any time in the Directory Synchronization Client. To cancel the schedule, use the Windows Control Panel.

1. Select **Start > Control Panel > Scheduled Tasks**.

2. Right-click the task "Sync Client-*example*", where *example* is the name of your configuration profile, then select **Delete**.

3. Click **Yes** to confirm.

Running the command-line synchronization client

You can optionally run the Directory Synchronization Client from the command line. This option performs a synchronization update only.

When you run the client from the command-line, use the **-config** parameter to specify the name of the configuration profile you want to use.

The command-line version of the Directory Synchronization Client is called **dirsyncclientc.exe** and is located in the installation folder.

To run the tool:

1. Ensure that the graphical version of the Directory Synchronization Client is not running.

2. Open a command prompt.

3. Change to the installation directory. For example:

    ```
    C:\>cd "\Program Files\DirSyncClient"
    ```

4. Enter the **dirsyncclientc -config** command, specifying the name of the configuration profile you want to use. For example:

    ```
    C:\Program Files\DirSyncClient>dirsyncclientc –config mycfg
    ```

Troubleshooting the synchronization process

To confirm the synchronization process has worked:

1. Check that the totals in your notification email messages look correct.

2. Log on to the cloud portal and go to the **Account > Identity Management** page.

3. Under **Recent Synchronizations**, check the numbers in the Additions and Deletions columns for the relevant synchronization match the totals in your notification email messages.

4. If the totals don't match:

    ■ If you are an existing cloud customer, you can restore the previous version of your data while you diagnose the problem.

    ■ Use the information on the **Groups** page and the **Search End Users** page in the cloud portal to compare the portal content against the expected results.

    ■ In the cloud portal, under **Recent Synchronizations** on the **Manage Directory Synchronization** page, click the date and time of the relevant synchronization. Check the logging information.

    ■ Check the contents of the logs in the Directory Synchronization Client.

Checking directory referrals

Referrals, or continuation references, are links from a location on an LDAP server to either a different directory server or another location on the same directory server. You can define how referrals are handled in your configuration profile: see *Defining advanced LDAP server settings*, page 20.

A missing or broken referral can result in fewer results being returned than expected, and can cause errors in your search.

A company spread across multiple sites is likely to have an LDAP server that contains referrals to other LDAP servers on different sites. If a referral is temporarily broken, all of the entries on the other server disappear when a synchronization takes place. To protect against such a situation, we recommend that you either set the action to abort the update, or set a maximum number of deleted entries on the threshold limits page. For more information on setting threshold limits, see *Limits*, page 37.

You can use the Directory Referral Checker diagnostic tool to check if there are any problems with following references.

1. In the Directory Synchronization Client, select a configuration profile from the **Configuration** drop-down list.
2. Select **View > Directory Referral Checker**.
3. If your configuration contains multiple synchronization types, click the tab for the type you want to check.
4. Click **Start**.
5. Examine the output to see if there are any problems.
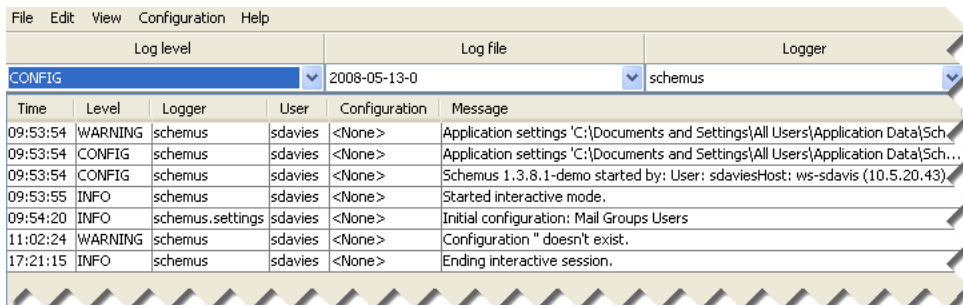
6 | Directory Synchronization Client Log Files

Both the graphical and command-line Directory Synchronization Clients produce log messages. Each message includes the:

- Time and date that the event occurred
- Logging level
- Configuration profile (if any) in use
- User logged in when the client was run
- The client component that is the source of the log message

To access the logging window in the Directory Synchronization Client, select **View > Logs**.

You can define the messages that you see in this window by selecting options from the drop-down lists:

- **Log level** sets the importance of displayed messages. Select one of the following:

Log Level	Description
SEVERE	Error message from the Directory Synchronization Client.
WARNING	Warning message, for example if a file does not exist.
INFO	Information message, for example logging when you start or exit the Directory Synchronization Client.
CONFIG	Extended information message, for example logging the Directory Synchronization Client version details and user details on startup.
FINE	Detailed message, for example logging when you try to connect to an LDAP server.
FINER	More detailed than FINE, for example logging the selection you make from a drop-down list.
FINEST	More detailed than FINER, for example logging system properties.

> **Note**
>
> The level of logging available depends on the log settings you have configured. For example, if you set the log detail to Lowest, only SEVERE and WARNING messages are available. For more information, see *Configuring log settings*, page 11.

The Directory Synchronization Client displays messages for the option you select, and also all messages higher in the list than that option. For example, if you select INFO, the window also displays SEVERE and WARNING logging levels.

- **Log file** is the name of the directory that stores the log messages. This directory is in the root logging directory:

```
Documents and Settings\All Users\Application
Data\DirSyncClient\application\log
```

The log file name comprises the year, month, day, and an extension, which is the number of the client instance that generated the message.

- **Logger** is the Directory Synchronization Client component that generated the message. The drop-down list allows you to restrict the displayed messages to a particular component and its sub-components. For example, selecting **dsc.sync** shows messages from the components **dsc.sync.source**, **dsc.sync.repository**,

dsc.sync.repository.add, and **dsc.sync.repository.remove**. Selecting
dsc.sync.source shows only messages from the **dsc.sync.source** component.

Component	Description
dsc	All components
dsc.settings	Creation of new configuration entries or changes to existing configuration
dsc.sync	All synchronization operations
dsc.sync.source	Operations to the source repository (normally an LDAP server)
dsc.sync.repository	All modifications to the destination repository
dsc.sync.repository.add	Entries added to the destination repository
dsc.sync.repository.remove	Entries removed from the destination repository

Click on a line in the message list to display any additional information at the bottom of the window.

Note

Messages with a level of INFO and higher importance are also logged to the application section of the event log.

A | Standard Regular Expression Strings

Regular expressions (regex) are a powerful way of matching a sequence of simple characters. You can use regular expressions in the Directory Synchronization Client to create filters (see *Filters*, page 35).

Regular expressions are case-sensitive: a lowercase "a" is distinct from an uppercase "A." You can enclose a range of characters in square brackets to match against all of those characters. For example:

Expression	Description
[tT]here	matches against "There" and "there"
[ ]	may also be used on a range of characters separated by a – character.
[0-9]	matches any digit.
[A-Z]	matches any uppercase alpha character
[A-Za-z0-9]	matches any alphanumeric character
^	is the "not" character, so [^0-9] matches against any character that is not a digit.

Although you can use ranges to specify a group of characters, you can also use the following shortcuts:

Expression	Description
.	matches against any character
\d	matches against a digit [0-9]
\D	matches against a non-digit [^0-9]
\s	matches against a whitespace character (such as a tab, space, or line feed character)
\S	matches against a non-whitespace character
\w	matches against an alphanumeric character [a-zA-Z_0-9]
\W	matches against a non-alphanumeric character
\x*hh*	matches against a control character (for the hexadecimal character *hh*)
\u*hhhh*	matches against a Unicode character (for the hexadecimal character *hhhh*)

Note

As the backslash character is used to denote a specific search expression, if you want to match against this character, you must enter a double backslash (\\).

To match against occurrences of a character or expression, you can use the following.

Expression	Description
*	matches against zero or more occurrences of the previous character or expression
+	matches against one or more occurrences of the previous character or expression
?	matches zero or one occurrences of the previous character or expression
(n)	matches n occurrences of the previous character or expression
(n,m)	matches from n to m occurrences of the previous character or expression
(n,)	matches at least n occurrences of the previous character or expression

You can provide text to replace all or part of your search string. To do this, you need to group together matches by enclosing them in parentheses so they can be referenced in the replacement. To reference a matched parameter, use **$n** where **n** is the parameter starting from 1.

For regular expression examples, see *Regular expression examples*, page 57.

Regular expression examples

Example 1: Filtering all addresses without a domain

Match string:	^@.*
Replacement string:	none
Input data:	fred.bloggs, fred.bloggs@acme.com
Output data:	fred.bloggs@acme.com

The string **@.*** matches against anything following an @ symbol, which indicates that a domain is present. The ^ symbol ensures that matching strings are excluded from replacement, and all data without a domain is removed.

Example 2: Appending "acme.com" to every email address

Match string:	\s*(\S*)
Replacement string:	$1@acme.com
Input data:	fred, jim
Output data:	fred@acme.com, jim@acme.com

The string **\s*** removes any whitespace at the start of the matching string, and **(\S*)** matches against the remaining non-whitespace characters. The parentheses allow you to reference this matching string as parameter 1 **($1)**.

In the replacement string, **$1** contains the text matched by **\S***, and then **acme.com** is appended to that text.

Example 3: Removing a selection of characters from an email address

Match string:	(.*)[#!_\.\s]*(.*)
Replacement string:	$1$2
Input data:	fred.bloggs@acme.com, #jim_bloggs@acme.com
Output data:	fredbloggs@acme.com, jim_bloggs@acme.com

The string **[#!_\.\s]*** matches against the pound (#), exclamation mark, underscore, period, and whitespace characters, with the final asterisk allowing multiple matches. The string **(.*)** on either side places all other characters in parameters 1 and 2. These parameters then form the replacement string, stripping out all instances of the matched characters.

Example 4: Converting .org domains to .com

Match string:	(.*@acme\.)org
Replacement string:	$1com
Input data:	fred.bloggs@acme.org
Output data:	fred.bloggs@acme.com

The match string **.*@acme\.** detects any address that contains the string **@acme.** and is preceded by 1 or more characters. The final full stop needs to be escaped (preceded by a backslash), to avoid being interpreted as any character.

The parentheses around this part of the match string ensure the string is placed in parameter 1. The final **org** in the match is outside the parentheses, hence it does not get placed in parameter 1. The replacement string contains the text **com** which is appended to the string matched by:

> **.*@acme\.**

Example 5: Filtering out all addresses from 2 domains

Match string:	^.*(?<![@]somewhere|[@]here)\.com$
Replacement string:	none
Input data:	fred.bloggs@somewhere.com, fred.bloggs@acme.com
Output data:	fred.bloggs@acme.com

These filters match and exclude any email address that ends in **@somewhere.com** or **@here.com**. The remaining email addresses are uploaded.

Changing wildcard filters to regular expressions

All filters in the Directory Synchronization Client must be of the same type, either regular expression or wildcard. To change from wildcard filters to regular expressions, convert your filters as follows:

- In the filter pattern, replace all instances of ***** with **(.*)**.
- In the replacement string, replace all instances of ***** with **$1**.

C | Working with Java

When you install the Directory Synchronization Client, you can choose whether or not to install the Java Runtime Environment (JRE).

You can install the JRE independently of the Directory Synchronization Client so that it is available to multiple applications. Alternatively, you can install a separate copy for each application on your system that requires a JRE.

The advantage of installing a JRE with each application is that if you remove or update the global JRE, your application does not stop working. The main disadvantage is that the JRE is several megabytes in size, and installing a copy for each application could consume disk space. We recommend installing the Directory Synchronization Client with its own JRE.

To check the current version of JRE in Windows, select **Start > Control Panel > Java**. Click **About** to display the version number.

To check the current version of JRE in Linux, go to the following Web site, which displays your Java version:

http://www.java.com/en/download/help/testvm.xml

Java Certificate Store

Java uses a certificate store, located in the **jre/lib/security** directory of your Java installation. If you are using the Directory Synchronization Client with its own Java Runtime, the jre directory is located in the directory where the client is installed.

For secure communications, the server provides a certificate which has been signed by a Certification Authority. The client checks the certificate store for the Certification Authority's certificate before allowing communication with the server. Because the certificate provided by the cloud service has been signed by a Certification Authority whose certificate is present in the standard Java certificate store, in most cases, no action need be taken to enable secure communications with the cloud service.

Proxy servers typically pass HTTPS traffic unaltered so no action is required when accessing the cloud service via a proxy. Some proxy servers, however, decrypt then re-encrypt the data before passing them to the destination. In this case, the proxy server,

rather than the cloud service, supplies the certificate used by the Directory Synchronization Client. If the proxy's certificate is self-signed or signed by a Certification Authority whose certificate is not in the standard cacerts file, the signing certificate will need to be imported.

If you need secure communications with an internal LDAP server, it is common for the certificate provided by the LDAP server to be either self-signed or signed by a Certification Authority whose certificate is not present in the standard cacerts file. In order to allow secure communications with such a server, you must import the signing certificate into the cacerts file as a trustpoint.

Importing a certificate

To add a certificate to the Java cacerts file, you can use the keytool application provided with the Java installation, located in the **jre/bin** directory. The following command imports a certificate from the file ldap-certificate.cer into the cacerts file:

```
keytool -import -trustcacerts -alias ldap-certificate -file
ldap-certificate.cer -keystore cacerts
```

> **Note**
> If you have not changed the keystore password, enter the default password, *changeit*, when prompted.

If you are using a system JRE and do not want to modify the system cacerts file, you can create the directory **application/lib/security** in the directory where the Directory Synchronization Client is installed, copy the system cacerts file to **application/lib/security/schemus-cacerts**, then modify the copy.

D | Using Generic LDAP

This appendix is intended for administrators who want to synchronize data from a generic LDAP format or from a file.

Generic LDAP enables you to query any LDAP-compliant system to return user, groups, and email data that can then be synchronized with the cloud service. Alternatively, you can extract your LDAP data to a comma-separated variable (CSV) or plain text file, and use that file as the input to the synchronization process.

It is important that you review the data you are about to synchronize before you synchronize it. The Requirements section describes a number of issues that you must take into account before your first synchronization. See:

- *Requirements*, page 61
- *Setting up the Directory Synchronization Client*, page 63

In addition, *Formats*, page 66, lists all of the LDAP elements that can be present in your input data, with examples.

Requirements

Note the following requirements for successful synchronization with your Forcepoint cloud-based product

Mail synchronization

- Email addresses can be synchronized independently of any other data.
- Each address to be synchronized must be of a valid email type (for example, jbloggs@acme.com), or it will be rejected by the cloud service.
- Each address must be globally unique, otherwise it will be rejected.
- Email addresses can be obtained from a number of objects, for example users, groups, distribution lists, or contacts.

Users and groups synchronization

Observe the following when synchronizing users and groups for the first time.

- Users and groups must be synchronized together.

- You must synchronize at least one group.

- You do not have to synchronize all members of a group, nor is it necessary for all groups containing members to be synchronized. The cloud service can handle references between users and groups that are not actually present in the portal.

To synchronize valid group and user data in the portal, ensure your LDAP data meets the following requirements.

- If you want to use group and user membership hierarchies, you must ensure the Group Parents, Group Members, and Other Groups fields are populated and return consistent data. For example, for the group/user relationship between GroupA and UserA to work correctly in your cloud security product, GroupA must have a 'member' attribute with a value of "UserA", and UserA must have a 'memberOf' attribute with a value of "GroupA".

- Group Parents and Group Members attributes must be accurately maintained in the directory for each group, and both attributes synchronized with the DN of the referenced groups.

- Users must have a MemberOf attribute with the DN of the group or groups of which they are members.

- Globally unique identifiers (GUIDs) are normally an attribute in each directory object, and are used as the primary key in the portal. If GUIDs are not synchronized, they are automatically generated by the Directory Synchronization Client based on the object's DN, and loaded into the portal. This is an acceptable solution unless the object's DN changes (for example, if someone gets married and changes their name, or their object is reorganized in the directory). In this case the auto-generated GUID changes, and the user is treated as a new user. This could mean that their old details on the portal are not picked up, and they have to re-enter passwords and possibly other details. It is recommended that you synchronize GUIDs if possible.

- There is considerable flexibility in the name attribute used when synchronizing to the portal. Use %CN% in the Name field to return the common name of the object; this is then synchronized. Most directories require the CN to be unique, which ensures the name is also unique on the portal. However, note that this is not enforced in all directories.

  %sAMAccountname, if provided, is also commonly used for the unique name of an object. We recommend that the Name is unique in the portal, although duplicates are tolerated.

- Each user must have a valid email address, and this must be unique. Any users without an email address are rejected by your cloud security product during the synchronization.

- For users, the Name attribute can be constructed dynamically to become the NTLM ID for the user object. A typical NTLM ID is domain\username.

  In directory terminology this could be constructed in a variety of ways, for example:

  - **ACME\\%CN%** would produce an NTLM ID with the domain=ACME, and username=common name of the object—for example ACME\JSmith

- ■ **%DC[-1]%\\%CN%** would produce an NTLM ID based on a DC and the CN of the object – for example, in the domain acme.com, this would produce acme\JSmith
- ■ **ACME\\%sAMAccountName%** would produce ACME\\JohnSmith. This is used in Active Directory schemas as it is used for the NTLM ID in Windows and is the recommended solution in those environments.

When constructing the NTLM IDs, it is important to ensure a match with the NTLM IDs used by the end users.

On the portal there is also a Name attribute in the Users record. This is always the CN of the object.

Setting up the Directory Synchronization Client

The instructions in this section describe how to set up a single synchronization type in the Directory Synchronization Client for either generic LDAP or an input file.

For full details of setting up a configuration profile, refer to *Creating and Modifying Configuration Profiles*, page 15.

Installing the client

Follow the instructions in *Installing the Directory Synchronization Client*, page 7, to download the Directory Synchronization Client from the cloud portal.

By default, the Directory Synchronization Client is not set up to use generic LDAP. You need the following additional files, available from Forcepoint:

- ● **datasources.xml** provides the required options in the Directory Synchronization Client
- ● **attributes.xml** overrides the default Name attribute, allowing this field to be blank in the Configuration Wizard if this is required

Place these files in the following folder on the same machine where you installed the Directory Synchronization Client:
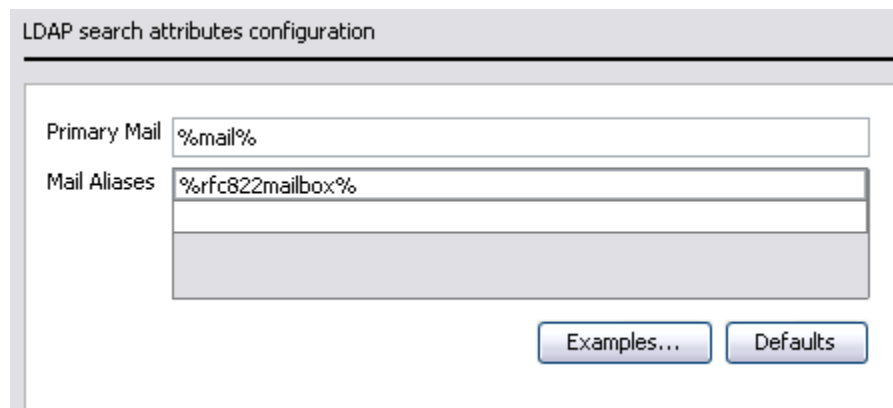
```
C:\Documents and Settings\All Users\Application Data\DirSync
Client\Application\Settings
```

Configuring generic LDAP

1. Run the Directory Synchronization Client.
2. To start the Configuration Wizard, click **New** (to the right of the **Configuration** drop-down list).
3. Enter a name for the configuration profile and select a synchronization type from the drop-down list, then click **Next**.
4. From the **Source Type** drop-down list, select **Generic LDAP**, then click **Next**.
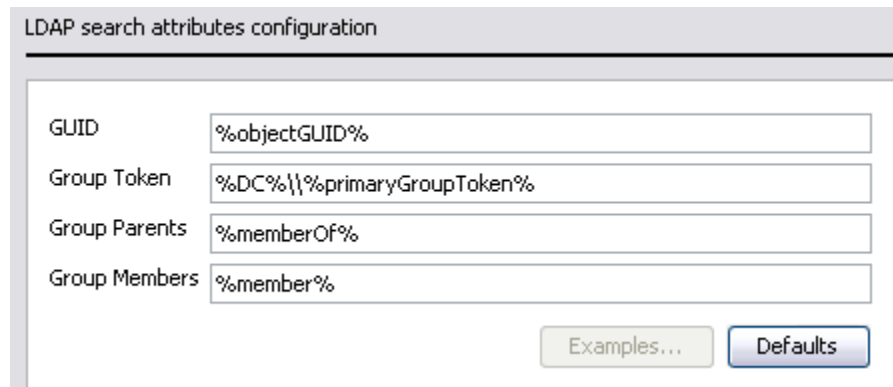
5. Set up your LDAP server as follows:

   a. In the **Host name** field, enter the host name of your LDAP server.

   b. Unless you know otherwise, leave the **Port number** field as the default value 389, which is the number used for communicating with an LDAP server in plain text mode.

   c. Select an authentication type, and if required enter a user name and password.

   d. Click **Next**.

6. On the LDAP search configuration page, click **Advanced**.

7. Enter search attributes that match the format described in *Generic LDAP format*, page 66, then click **Next**.

   The mail synchronization type should appear as follows:

   The groups synchronization type should look like this:

The users synchronization type should be as follows:

8. Work through the rest of the wizard, setting up your data repository and any filters, limits, and notifications.

Configuring a file input

1. Create your input file, using the format described in *File format*, page 67.

2. Save the file as a **.csv** or **.txt** file.

3. Run the Directory Synchronization Client, then click **New** to start the Configuration Wizard.

4. Enter a name for the configuration profile and select the required synchronization type from the drop-down list, then click **Next**.

5. From the **Source Type** drop-down list, select **File**, then click **Next**.

6. Click the Browse button, navigate to your input file, and click **Select**.

7. Click **Next**.

8. Work through the rest of the wizard, setting up your data repository and any filters, limits, and notifications.

Formats

Generic LDAP format

The table below describes how mail addresses, groups, and user information must be formatted in generic LDAP input.

Fields		Syntax	Format	Other
Mail fields				
Primary Mail	%mail%	Directory string Example: jsmith@acme.com	Text	Mandatory Globally unique
Mail aliases/ proxy addresses	%rfc822 mailbox%	Directory string Example: joe@acme.com smith@acme.co.uk	Text	Optional Globally unique
Groups fields				
Name	%CN%	Directory string Example: Name, CN, sAMAccountName, Display Name	Text	Mandatory Unique in account
GUID	%object GUID%	Hex string Example: 746B8515-C8FF-C940-9D905F053CB22D25	Hex 16 bytes	Mandatory Unique in account
Group Parents	%member Of%	DN Example: CN=AllStaff,OU=London, DC=acme,DC=com	Text	Optional Unique in account
Group Members	%member %	DN Example: CN=Sales,OU=London,DC =acme,DC=com	Text	Optional Unique in account
Users fields				
Name	%CN%	Directory string Can be constructed dynamically to become the NTLM ID for the user object. A typical NTLM ID is domain\username, for example acme\JSmith.	Text	Optional Unique in account

Fields		Syntax	Format	Other
Primary Mail	%mail%	Directory string Must be a valid SMTP email address.	Text	Mandatory Globally unique
Mail aliases/ proxy addresses	%rfc822mailbox%	Directory string Must be a valid SMTP email address.	Text	Optional Globally unique
Primary Group	%primaryGroupId%	Integer Not used	Text	Not used
Other Groups	%memberOf%	DN Example: CN=AllStaff,OU=London, DC=acme,DC=com	Text	Optional Unique in account
GUID	%objectGUID%	Hex string Example: 746B8515-C8FF-C940-9D905F053CB22D25	Hex 16 bytes	Mandatory Unique in account

File format

You can synchronize data to the cloud portal from a file source. The table below shows the required format of data in input files for mail, groups, and users. Each field must be separated by a comma, and each mail address, group, or user must start on a new line.

Fields		Syntax	Type	Other
Mail fields				
Email address		rfc822 mailbox format Example: joe.smith@acme.com	Text	Optional Globally unique
Group fields				
Group DN		DN Example: DN=CN=Telesales,OU=London,DC=acme,DC=com	Text	Mandatory Unique in account
,,	Blank	Two commas, no space between ,, (This field is not used in Groups, but commas must be included.)	Text	Mandatory

Fields		Syntax	Type	Other
GUID		Hex string Example: 746B8515-C8FF-C940-9D905F053CB22D25	Hex 16 bytes	Mandatory Unique in account
Name		String GroupName (for example, Sales)	Text	Mandatory Unique in account
Group Parents (MemberOf)		DN Example: CN=Sales,OU=London,DC=acme,DC=com	Text	Optional - can be many Unique in account
Users fields				
Username DN		DN Example: DN=CN=JSmith,OU=London,DC=acme,DC=com	Text	Mandatory Unique in account
Extra attribute (use ,, if not required)	Mailalias or ,,	Mailalias=<email> Can be used for any mail attributes by defining the attribute in the string itself: Proxyaddress=JSmith@acme.com,joe.smith@acme.com,joe@acme.net	Text	Optional Globally unique
GUID		Hex string Example: 746B8515-C8FF-C940-9D905F053CB22D25	Hex 16 bytes	Mandatory Unique in account
NTLM ID		Domain\username Example: Sales\JSmith	Text	Optional Globally unique
Primary email		Valid email address Example: joe.smith@acme.co.uk	Text	Mandatory Globally unique
Group Parents (MemberOf)		DN Example: CN=Sales,OU=London,DC=acme,DC=com	Text	Optional - can be many Unique in account

A single line in an input file for users would look like this:

```
<usernameDN>,,<GUID>,<NTLMID>,<Primary email>,<Memberof
Group1>,<Memberof Group2>,<Memberof Groupn>
```

Within each field, any backslashes or further commas must be escaped, using hexadecimal code **\0x002c** for a comma, and **\0x005c** for a backslash. For example:

```
dn=CN=Joe.Smith\0x002cOU=Salesoffice\0x002cDC=acme\0x002cDC=
com, ,746B8515-C8FF-C940-
9D905F053CB22D25,acmenet\0x005cjsmith,smith@acme.com,
CN=Sales\0x002cOU=salesoffice\0x002cDC=acme\0x002cDC=com,
CN=USemployees \0x002cDC=acme\0x002cDC=com
```

If you use the mailalias option, any commas in the alias list must be "double escaped," with the escape character escaped itself. For example:

```
mailalias=JSmith@acme.com,J.Smith@acme.co.uk
```

must be written:

```
mailalias=JSmith@acme.com\0x005c0x002cJ.Smith@acme.co.uk
```

The full line in the input file would look like this:

```
dn=CN=Joe.Smith\0x002cOU=Salesoffice\0x002cDC=acme\0x002cDC=
com,
mailalias=JSmith@acme.com\0x005c0x002cJ.Smith@acme.co.uk,746
B8515-C8FFC940-
9D905F053CB22D25,acmenet\0x005cjsmith,smith@acme.com,
CN=Sales\0x002cOU=salesoffice\0x002cDC=acme\0x002cDC=com,
CN=USemployees \0x002cDC=acme\0x002cDC=com
```