

## Coveo Enterprise Search 6.2

### Microsoft Exchange Connector

Coveo's Microsoft Exchange connector allows users to search Microsoft Exchange content. By using a unique special administrative account, all the mailboxes, mailbox archives, and public folders can be indexed.

#### Features

The Microsoft Exchange connector can be used to allow Coveo Enterprise Search (CES) to crawl and index the content of Microsoft Exchange Server 2003, 2007, or 2010 installations. Furthermore, it offers a more efficient search than Microsoft Exchange built-in search components.

The following details the features available in the Microsoft Exchange connector:

- Mailbox indexing;
- Mailbox archive indexing (beta release);
- Public folder indexing;
- Security indexing by combining Microsoft Exchange permissions set in Microsoft Exchange by users and set in Active Directory by administrators;
- Live indexing to periodically query Microsoft Exchange for the latest edits, keeping the index content up to date;
- Metadata collecting for emails, contacts, tasks and calendars;
- Supports two types of forms-based authentication for Outlook Web Access (OWA):
  - Standard Exchange forms-based authentication;
  - ISA Server forms-based authentication.

#### Requirements

The Microsoft Exchange connector supports the standard Microsoft Exchange front-end and back-end server configurations, as described in the following articles:

- [http://technet.microsoft.com/en-us/library/aa997482\(EXCHG.65\).aspx](http://technet.microsoft.com/en-us/library/aa997482(EXCHG.65).aspx).
- [http://technet.microsoft.com/en-us/library/bb124804\(EXCHG.65\).aspx](http://technet.microsoft.com/en-us/library/bb124804(EXCHG.65).aspx)

#### Installation

The Microsoft Exchange connector components are installed when you install CES. There is no other installation required.

## Configuration

### How to Set Up an Administrative Account

You need to set up a special account in the Microsoft Exchange Server to be able to crawl multiple mailboxes using a single source. This account must have read access on all the mailboxes and all the public folders that you want to index.

1. For a Microsoft Exchange Server 2003 setup, refer to the following link for information on how to set up such an account: <http://support.microsoft.com/default.aspx?scid=kb;EN-US;821897>.
2. For Microsoft Exchange Server 2007 or 2010 setups:
  - a. Open the Exchange Management Shell.
  - b. Type the following command to give sufficient rights to a user to crawl using WebDAV or WebServices:

```
get-mailbox -server ServerName -ResultSize Unlimited | Add-MailboxPermission -User 'domain\user' -AccessRights FullAccess
```

### How to Configure a User Identity Based on an Administrative Account

After setting up the Microsoft Exchange Server administrative account, you need to configure the CES user identity.

1. In the CES Administration Tool, access the **Administrators** page (**Configuration > Security**).
2. In the left navigation pane, click **User Identities**. The **User Identities** page is displayed.
3. Click **Add**. The **Modify User Identity** page is displayed.
4. In the **Name**, **User** and **Password** fields, enter the credentials of the special administrative account that you set up in the Microsoft Exchange Server.

The screenshot shows the 'Modify User Identity' page in the Coveo Administration Tool. The page title is 'Modify User Identity' and it includes a 'Help' link. Below the title, there is a description: 'Defines the parameters of this user identity.' On the left, there is a navigation pane with options: Administrators, Impersonators, Security Providers, **User Identities**, and Super User Access. The main content area contains the following fields and options:

- Name:** companyname\jsmith
- User:** companyname\jsmith
- Password:** [Masked with dots]
- Option:**  Support basic authentication
- Client certificate:**  Use client certificate
- Store:** Personal (dropdown menu)
- Certificate:** [Dropdown menu]
- Used By:** [Empty field]

At the bottom right, there are 'Save' and 'Cancel' buttons.

5. Click **Save**.

## Why Choose WebServices versus WebDAV

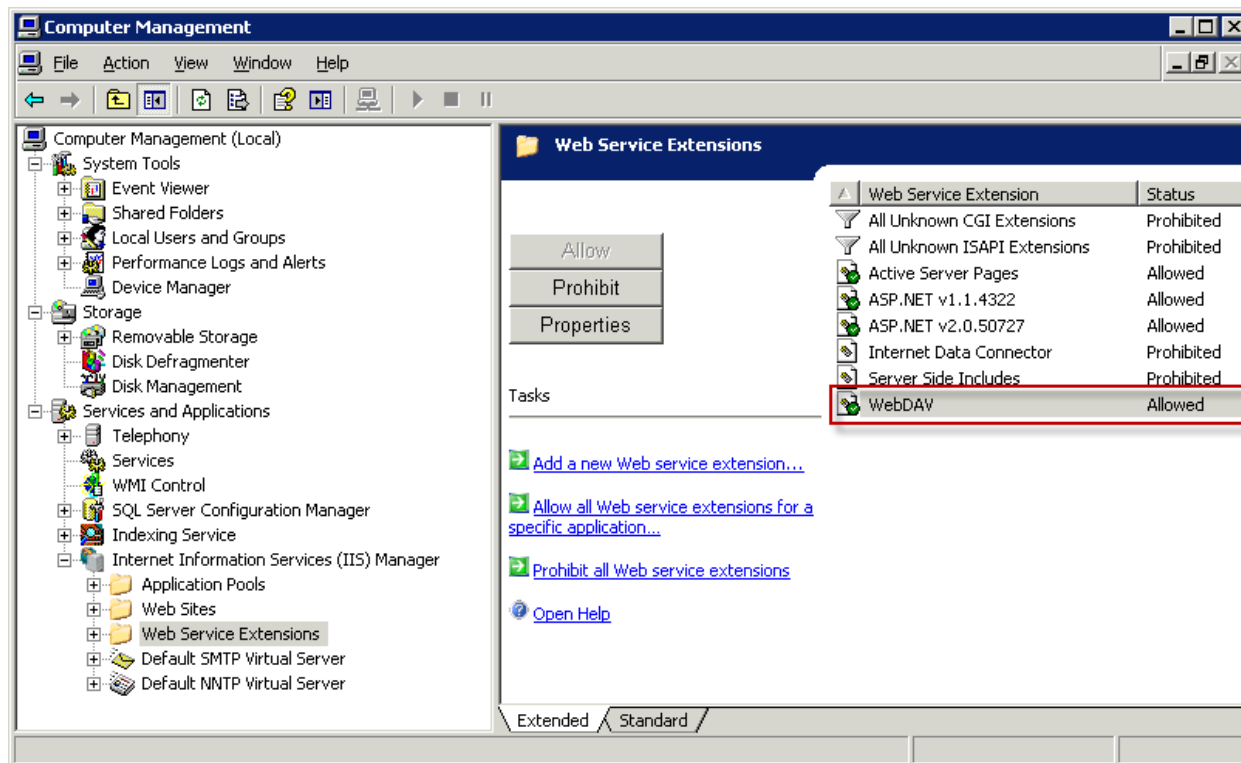
The Microsoft Exchange connector offers two options: **Exchange (WebDAV)** and **Exchange (WebServices)**. The following table indicates which options are supported by which version of Exchange.

|               | Exchange (WebDav)                                                     | Exchange (WebServices) |
|---------------|-----------------------------------------------------------------------|------------------------|
| Exchange 2003 | Supported                                                             | Not supported          |
| Exchange 2007 | Supported<br>(when the Exchange 2003 Compatibility mode is installed) | Supported              |
| Exchange 2010 | Not supported                                                         | Supported              |

You should choose the **Exchange (WebDav)** option only when indexing an Exchange 2003 server. For Exchange 2007 setups, it is recommended to use the WebServices option as it offers better performances and is simpler to configure.

### Important:

- Once a source is created, you cannot change the option (WebDAV or WebServices). The work around is to create another source with the other option.
- When you choose WebDAV, be aware that the WebDAV protocol is prohibited by default in IIS. You can allow WebDAV in IIS from the **Computer Management** tool (see the following illustration).



## How to Index a Microsoft Exchange Source

1. In the CES Administration Tool, access the **Sources and Collections** page (**Index > Sources and Collections**).
2. In the **Sources** section, click **Add**.
3. In the **Add Source** page, refer to the following table to enter the appropriate parameter values.

| Field                   | Description                                                                                                                                                                                                                                            | Example                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                    | Any descriptive name.                                                                                                                                                                                                                                  | All Microsoft Exchange mailboxes                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Source Type             | The option used by this source: <b>Exchange (WebDav)</b> or <b>Exchange (WebServices)</b> (see <a href="#">Why Choose WebServices versus WebDAV</a> ).                                                                                                 | <b>Exchange (WebServices)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Addresses               | List of starting points for the connector, one address per line.<br><b>Note:</b> Entering a specific server name will not restrict the crawled mailboxes to those located on that server (see <a href="#">What Are LDAP Searches</a> in the appendix). | Addresses are in the form of:<br>https://servername/exchange<br>https://servername/exchange/mailboxname<br>https://servername/public<br>To use a single source to index all Microsoft Exchange mailboxes, enter a single URL that stops at the mailboxes root (ex: https://owa.companyname.com/exchange).<br>To index all public folders, enter a URL that stops at the public folders root (ex: https://owa.companyname.com/public).<br>The connector supports both http and https |
| Max Number of Retries*  | Number of retries to perform when a request fails due to a Microsoft Exchange Server downtime or temporary unavailability.                                                                                                                             | The default value is 3.                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Retry Delay*            | Delay between each retry attempt.                                                                                                                                                                                                                      | The default value is 30 seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Exchange Folder         | Name of the folder hosting all the mailboxes.                                                                                                                                                                                                          | The default value is exchange. It corresponds to the segment of the URI that is after the server name and points to the mailboxes: https://servername/exchange/.                                                                                                                                                                                                                                                                                                                    |
| Use Form Authentication | Whether to use forms-based authentication or not.                                                                                                                                                                                                      | The default value is False. Necessary when <b>Forms-based Authentication</b> is enabled on the Microsoft Exchange server.<br>Supports two forms-based authentication types:<br>Standard Microsoft Exchange forms-based authentication.<br>ISA server forms-based authentication. When an ISA filter is put in front of the Microsoft Exchange OWA server.                                                                                                                           |

| Field                   | Description                                                                                                                                                                                                 | Example                                                                                                                                                                                                                                                                                            |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Index Exchange Security | Whether to index the securities specified in Microsoft Outlook or not.                                                                                                                                      | The default value is True. Useful to prevent Coveo from indexing user defined securities. Only Active Directory Securities are used.                                                                                                                                                               |
| Callback Service Url*   | Indicates the location of the Exchange Callback Service. Enhances the live indexing performance on Microsoft Exchange 2007 (see the How the Exchange Callback Service Improves Live Indexing Performances). | The default value (http://localhost:8181/) is appropriate when the Exchange Callback Service runs on the same machine as CES. When it runs on another computer, "localhost" must be changed accordingly, as well as the port if it was modified. Leave the parameter empty to disable the service. |
| Public Folder           | Name of the public folder.                                                                                                                                                                                  | The default value is public. It corresponds to the segment of the URI that is after the server name and that points to the public folders: https://servername/public/.                                                                                                                             |
| Index Junk E-mails      | Whether to index junk emails or not.                                                                                                                                                                        | Junk emails are not indexed by default (checkbox is cleared).                                                                                                                                                                                                                                      |
| Skip Archived Items     | Whether to skip items archived by Symantec Enterprise Vault or not.                                                                                                                                         | Useful to make the Microsoft Exchange connector ignore email stubs left behind by Symantec Enterprise Vault. The default value is true.                                                                                                                                                            |
| Index Archive           | Whether to index Exchange mailbox archives or not.                                                                                                                                                          | Exchange mailbox archives are indexed by default (checkbox is selected).                                                                                                                                                                                                                           |

\* Optional parameter

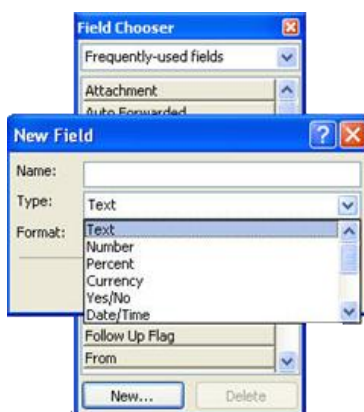
4. Click **Save and Start**.

## How to Modify Advanced Hidden Source Parameters


The **Add Source** page of the CES Administration Tool presents the parameters with which you can configure the connector for most Exchange setups (see How to Index a Microsoft Exchange Source). More advanced and more rarely used parameters are hidden. You can choose to make one or more of these parameters appear in the Administration Tool user interface so that you can change their default value.

1. Review the parameters described in the following list to see if you want to change the value of one or more of them.
  - **LDAPSearchRoot (string)**: Root of the search in Active Directory for the current domain (see What Are LDAP Searches in the appendix).
  - **LDAPFilters (string)**: Filter applied to the search of Active Directory in order to find mailboxes (see What Are LDAP Searches in the appendix).
  - **MaxNumberPrefetchChanges (integer)**: Maximum number of items that can be stored in the prefetch cache. This cache is useful to reduce the delay between each query to Microsoft Exchange. The default value is 200. Increasing this value increases performance, but also the amount of memory used by the connector. If the connector is running on a server with limited memory resources, lowering this number to 100 or 50 can prevent out of memory issues.

- **BatchSize (integer):** Number of items to fetch per request made to the Microsoft Exchange Server. The default value is 50. The minimum value is 1 and the maximum 512. A small value forces the connector to make small but frequent queries to Microsoft Exchange. A larger value leads to larger and less frequent connections.
- **IndexJunkAttachments (boolean):** When set to true, the attachments of junk emails are indexed, otherwise they are ignored. The default value is false.
- **SkipEmbeddedImageAttachment (boolean):** When set to true, the images embedded directly within the text of an email (like signature images) are skipped, otherwise they are indexed. The default value is true.
- **UserDefinedFields (string):** Can be added to retrieve user defined fields from Microsoft Exchange items. This string is in the following format: *FieldType:FieldName;FieldType:FieldName*, etc. The FieldType value is taken directly from the Microsoft Exchange Form Designer. The FieldName value is the name of the field as entered in the Form Designer (see the following illustration).



- **ConnectionLimit (integer):** Can be added to increase performance. If you plan on increasing the number of threads in the connector, you will also need to increase this value. The suggested value is the sum of all the threads you want to be able to connect to Exchange, in all your sources. Increasing this value will increase the load on the Exchange Server. The default value is 2.
- **IndexOutbox (boolean):** Can be added to crawl the outbox folder. This folder usually only temporarily contains emails before they are sent by the Exchange server. The default value is false.
- **IgnoreUnresolvedDeniedSecurities (boolean):** Can be added to allow the connector to keep crawling even after finding denied securities that could not be resolved. This behavior could create a security hole, so this parameter should be used with caution. The default value is false.
- **ResolveDisplayValuesWithContacts (boolean):** Can be added to allow the connector to perform a resolution of the email address values found on emails with the mailbox contacts. For example, if the connector finds the “To” value “jsmith <jsmith@companyname.com>”, it looks into the contacts folder of this mailbox and tries to form a new value using the first and last name of the contact: “John Smith <jsmith@companyname.com>”. The default value is false.

2. In the CES Administration Tool, access the **Modify the parameters of the additional connector** page (**Configuration > Connectors > Additional Connector > Exchange (WebDav) or (WebServices)**).
3. For each parameter that you wish to modify, click **Add Parameter**, and then enter the following parameters:
  - a. In **Type**, select the parameter type (ex.: String)
  - b. In **Name**, type the parameter name exactly as it appears in the previous list. Parameter names are case sensitive.
  - c. In **Default Value**, enter the default value specified in the previous list, not the value that you wish to use. You will be able to change the value for each source later.
  - d. In **Label**, enter the label that you wish to see for this parameter.
  - e. Optionally in **Quick Help**, enter the help text that you wish to see for this parameter when clicking the  icon beside the parameter value.
  - f. Select **Option** when you wish to identify this parameter as an optional parameter.
  - g. Click **Apply Changes**.  
The hidden parameter now appears in the **Add Source** page for Exchange sources, where you can change the parameter value.

## How to Index Exchange 2010 Mailbox Archives

Microsoft Exchange 2010 has a new user mailbox archiving feature that can be enabled. By default, the Exchange Connector attempts to crawl these mailbox archives. Results for both the mailbox and the mailbox archives are merged, making it easy for users to search for emails found in either. The mailbox archives are indexed by default.

1. To disable the indexing of mailbox archives, clear the **Index Archive** checkbox (see How to Index a Microsoft Exchange Source).
2. To only index a mailbox archive:
  - a. In the CES Administration Tool, access the **Sources and Collections** page (**Index > Sources and Collections**).
  - b. In the **Sources** section, click **Add**.
  - c. In the **Add Source** page, for the **Addresses** parameter, use the following starting address pattern:  
http://server/exchange/onlinearchivename  
where onlinearchivename is the value found in the Active Directory field msExchArchiveName. This points the connector directly into the mailbox archive, skipping the normal mailboxes.

## Exchange Callback Service

### ▶ How the Exchange Callback Service Improves Live Indexing Performances

For Microsoft Exchange Server 2007 and 2010, the Coveo Exchange WebServices connector uses an Exchange API to request notifications from Exchange when something is modified within a monitored mailbox. With this information, the Exchange WebServices connector can immediately update the good folders and items on the next live indexing run, instead of polling for every folder in every mailbox. This drastically improves live indexing performances and reduces bandwidth usage. This service is installed with the CES install kit, and is started if an Exchange WebServices connector

requests it. By default, a new Exchange WebServices source is ready to use the Exchange Callback Service.

The **Callback Service Uri** parameter (see How to Index a Microsoft Exchange Source) indicates the location of the Exchange Callback Service.

**Important:**

Ensure to open the port 8181 in the firewall. The Exchange Server has to connect to the Exchange Callback Service to send notifications.

A User Identity must be defined on the source or the Exchange Callback Service will not work.

When an Exchange WebServices source starts, a message appears in the log to mention whether the Exchange Callback Service is used or not:

- When the Exchange Callback Service is used: “SourceName – Using Push Notifications.”
- When the Exchange Callback Service is not configured correctly, or cannot be started, an error appears along with the following message: “SourceName – Using Standard Mode.”

► **How to Modify the Exchange Callback Service Configuration File**

The Exchange Callback Service is configured to run correctly in most cases using default parameter values. You can however manually configure advanced parameters in the Exchange Callback Service configuration file to override out-of-the-box behaviors.

1. Consult the following table of available configuration file parameters to see if you need to change the default value of one or more of them.

| Parameter       | Description                                                                                                  | Example                                                                                                    |
|-----------------|--------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|
| LogFileLifeTime | Duration before a log file will be automatically deleted by the service.                                     | Default is 30 days.                                                                                        |
| LogMaxSize      | Maximum size before splitting a log file.                                                                    | Default is 10 MB (10 485 760 bytes)                                                                        |
| MaxConnections  | Maximum number of concurrent connections that the service will create when calling the Exchange WebServices. | Default is 4 connections.                                                                                  |
| ShowAllCounters | Displays the usage statistics about the service in the CES console and logs before each live indexing run.   | Default is false.                                                                                          |
| DataFilePath    | Location of the data files used by the service.                                                              | By default, data files are located in the %installPath%\CES6\Index\Crawlers\EchangeCallbackService folder. |
| LogFilePath     | Location of the log files created by the service.                                                            | By default, log files are saved in the same folder as the other CES log files.                             |

| Parameter           | Description                                                                                                                                                                            | Example                                                 |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------|
| SubscriptionTimeout | Duration before an Exchange subscription will be considered invalid. This value should not be changed unless there are issues using the default value.                                 | Default is 30 minutes. Maximum allowed is 1440 minutes. |
| HostName            | Host name that will be used internally to build the web service URL. Changing this value affects directly the <b>Callback Service Uri</b> parameter that is set on an Exchange source. | Default is the computer name.                           |
| PortNumber          | Port number used internally to build the web service URL. Changing this value affects directly the <b>Callback Service Uri</b> parameter that is set on an Exchange source.            | Default is 8181.                                        |

- On the CES computer, browse to the CES bin folder (%installPath%\Coveo Enterprise Search 6\bin).
- Using a text editor, open the CESEExchangeCallbackService.exe.config file.
- In the content of the file, locate the <appSettings> section under <configuration> (see the following illustration).

```

<?xml version="1.0" encoding="utf-8"?>
<configuration>
  <appSettings>
    <add key="MaxConnections" value="6"/>
  </appSettings>
  <system.serviceModel>
    <bindings>
      <wsHttpBinding>
        <binding name="WSHttpBinding_INotif.
    
```

- Edit the file to add one custom parameter per line using the following format:

```
<add key="parameter_name" value="value">
```

## Appendix

### What Are LDAP Searches

Lightweight Directory Access Protocol (LDAP) searches are an important part of the Microsoft Exchange connector. For example, when the starting addresses `https://servername` or `https://servername/exchange` are used, a query is performed against the Active Directory of the current domain to find the available mailboxes.

The results should reflect all the users in the Active Directory who own a mailbox. Depending on the configuration of the Active Directory, it is possible that the results do not match the anticipated results. Two parameters of the connector allow to fine tune the search (see [How to Modify Advanced Hidden Source Parameters](#)):

- **LDAPSearchRoot:** Specifies to the connector where to start looking in Active Directory. When this parameter is not specified, the connector looks at the root of Active Directory, which can be extremely large. By specifying a value, it is possible to refine the search. For example, to search only within the organizational unit (OU) "companynameOU" within the domain "corp.companyname.com", enter: "LDAP://OU=companynameOU, DC=corp, DC=companyname, DC=com".
- **LDAPFilters:** Allows adding properties to the results the connector is looking for. The default filter is: "(&(mail=\*)(objectclass=user)(objectclass=person))". This AND operation is performed on a few properties to find the maximum number of mailboxes without receiving any unwanted ones. If this filter is not specific enough, you can complete the filter by adding a value to the AND operation on the filter. For example, if the value "cn=JohnSmith" is entered in the parameter, the resulting filter is: "(&(mail=\*)(objectclass=user)(objectclass=person)(cn=JohnSmith))".