



Coveo Enterprise Search for Intranet

Deployment Guidelines

Contents

Requirements	3
Database Server Considerations.....	3
Platform Requirements.....	3
Single Server/Small SharePoint Farm	3
Single SharePoint Server	3
Small SharePoint Farm	4
Medium/Large SharePoint Farm	4
Handling Large Number of Users	5
Indexing Large Amounts of Documents	5
Determining the Number of Slices.....	5
▶ Single or Multiple Front-End Web Servers, Single CES Server on a Dedicated Machine	6
▶ Multiple Front-End Web Servers, Multiple CES Mirrors on Dedicated Machines	7
▶ Multiple Front-End Web Servers, Multiple CES Mirrors with Multiple Slices	10
Naming Remote Slices	15
Conclusion	16

This document provides guidelines to deploy *Coveo Enterprise Search for Intranet* using *SharePoint* for the various topologies defined by *Microsoft*. It presents step-by-step procedures describing the tasks to perform in order to complete the setup for each supported topology.

Before installing *Coveo Enterprise Search for Intranet*, the installation and configuration of the *SharePoint* server farm must be completed and fully functional.

Requirements

Database Server Considerations

Coveo Enterprise Search for Intranet does not have any special requirements regarding the database server powering the *SharePoint* installation (either *MSDE* or *SQL Server*). However, *Coveo* recommends following the guidelines provided by *Microsoft* for the topology being deployed. CES does not directly use a database; in fact, it frees up *SQL Server* resources previously used for *SharePoint* search.

Note: The procedures discussed in this document can be applied to *Microsoft Office SharePoint Server* (MOSS).

Platform Requirements

Coveo Enterprise Search for Intranet performs adequately with the minimum hardware requirements if it is the only application running on the server. However, if other processes are running in parallel or the query activity is high and mission-critical, a server with the minimum required capacity may not be sufficient. In these cases, *Coveo* recommends the following to achieve the desired performance. The *Coveo* Technical Support team can help determine the appropriate server configurations for the intended application and desired performance levels. Additionally, refer to the [How to Speed up CES](#) help entry of the Administration Tool for a complete list of actions to help keep CES running at maximum speed.

Recommended hardware and software:

- CPU Dual 2.0 Ghz or better processors
- RAM: 4 GB or more
- 150 MB disk space for *Coveo*
Disk space for the search index varies based on the size of the original documents. Rule-of-thumb: allow for 50% of the size of the original email (e.g., 1 million documents can take ~40GB disk).
- *Microsoft Windows Server 2003* or *2008*
- *.NET Framework 3.5* or later
- *Microsoft Internet Explorer 6.0* or higher, *Mozilla Firefox 2* or higher
- Mobile Search: *Windows Mobile 5* or higher with *IE mobile*, *BlackBerry OS 4* or higher with the browser (optional)

Important: When the installation kit is launched, *Coveo Enterprise Search 6.0* automatically installs *Internet Information Services (IIS)* and *.NET Framework* with all the appropriate options on all supported versions of *Windows* (if they are not already installed). This facilitates the deployment of CES on *Windows 2008 Server* as well as on other *Windows* versions.

Single Server/Small SharePoint Farm

Single SharePoint Server

In this scenario, where *SharePoint* runs on a single server, it is recommended to install *Coveo Enterprise Search for Intranet* on the same *SharePoint* server. For a successful installation, run the *Coveo Enterprise Search for Intranet* installation kit on the same server as *SharePoint*, using the default options.

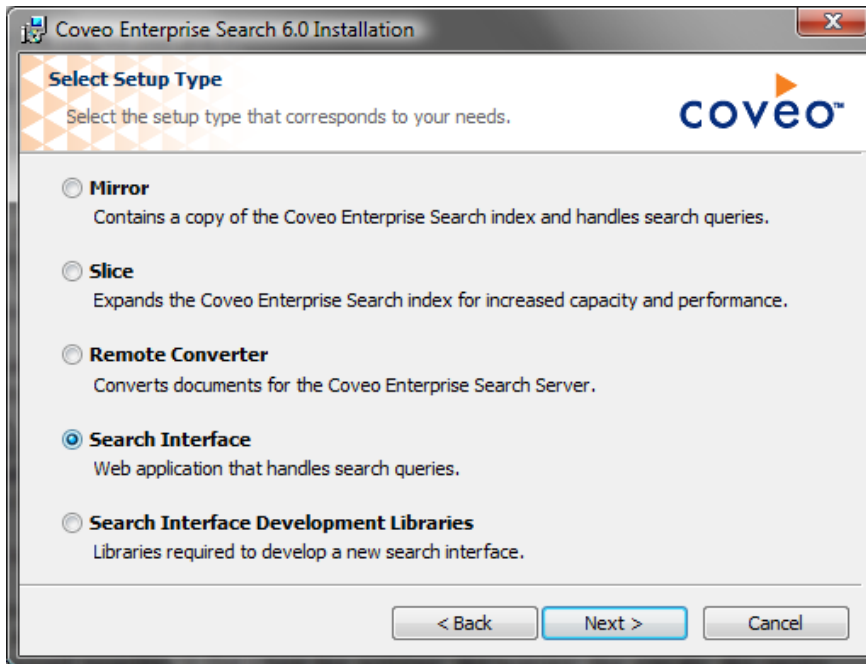
Small SharePoint Farm

In this configuration, the *SharePoint* installation is a single *SharePoint Web Front-End* server and a separate server running *SQL Server*. It is recommended to install CES on a separate, dedicated server using the following procedure:

1. Run the *Coveo Enterprise Search for Intranet* installation kit on the dedicated server, using the default installation options, and perform the initial setup using the Administration Tool.
2. Run the *Coveo Enterprise Search for Intranet* installation kit on the front-end Web server:



- a. Select **(Advanced)**, then Search Interface.



- b. Click **Next >** and enter the name of the server where the CES server was installed in Step 1.
- c. Proceed with the default options.

Medium/Large SharePoint Farm

In this scenario, the *SharePoint* setup can contain multiple *SharePoint Web Front End* servers, search servers and index servers along with a separate *SQL Server* which can be clustered. These more complex *SharePoint* configurations can require the deployment of one or more advanced CES performance configurations in order to scale and meet the high availability and performance requirements.

[CES mirrors](#)¹ and [slices](#)² are designed to distribute the search engine's work load for queries, as well as optimize content indexing in order to maintain optimal performance. Two factors must be considered when choosing which deployment scenario is best: the number of documents to index and the number of users.

¹ Mirrors are installed on separate servers to maintain fast search result response time while supporting more simultaneous user queries.

² Slices are installed on the same server or a separate server, to speed up content indexing and accommodate large (over 5-10 million document) indexes.

Handling Large Number of Users

For *SharePoint* deployments with large number of users, it is essential to spread the query load among multiple servers using [CES mirrors](#). A CES mirror holds a complete copy of the search index, and thus is able to perform queries independently from the other CES servers. A default setup only has a single copy of the index, meaning that there is only one mirror. Adding mirrors to your setup increases the number of users that can be handled effectively.

Microsoft recommends different topologies depending on the number of users (<http://office.microsoft.com/en-us/sharepointportaladmin/HA011647621033.aspx>). *Coveo's* recommendations are based on the numbers and formulas provided by *Microsoft*.

Users	SharePoint Topology	CES Topology
< 10,000	Single server with <i>SQL Server 2000</i>	1 CES server on a dedicated machine.
< 25,000	Small farm: Front-End Web server (1), SQL (1+, optionally clustered)	1 CES server on a dedicated machine.
< 100,000	Medium farm: Front-End Web server, index or job (1), SQL (1+, optionally clustered)	2+ CES mirrors, each running on a dedicated machine.
> 100,000	Large farm: Front-End Web server (2+), SQL (1+, optionally clustered)	2-3+ CES mirrors, each running on a dedicated machine.

Indexing Large Amounts of Documents

If an important number of documents is indexed, it is important to split each mirror's index into multiple slices, which can optionally be hosted on separate servers.

A slice holds a subset of the documents contained in a mirror. In a default installation, after performing the initial setup, the CES system only has one slice. In order to increase indexing performance as well as the maximum number of documents to index, it is possible to add other slices to the system. The new slices can either be on the same server or a remote machine.

All the mirrors in the same CES installation are required to have the same number of slices; therefore, when adding or deleting a new one through the Administration Tool, the user is only required to do it once for all mirrors in the system.

Determining the Number of Slices

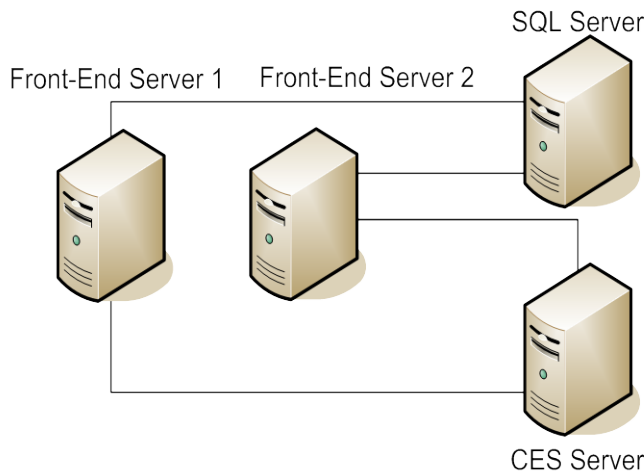
Slices are useful when indexing an important number of documents. Although a single slice can hold up to 16 million documents, for best performance it is recommended not to exceed 5 million.

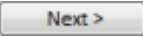
A single machine can host several slices—the exact amount depending on the hardware characteristics of the server; however, *Coveo* recommends a maximum of 4 slices per machine. For larger deployments, it is possible to host the slices on dedicated servers.

Note: In order to optimize performance when several slices are hosted on the same server, each slice should use a separate physical disk and/or array.

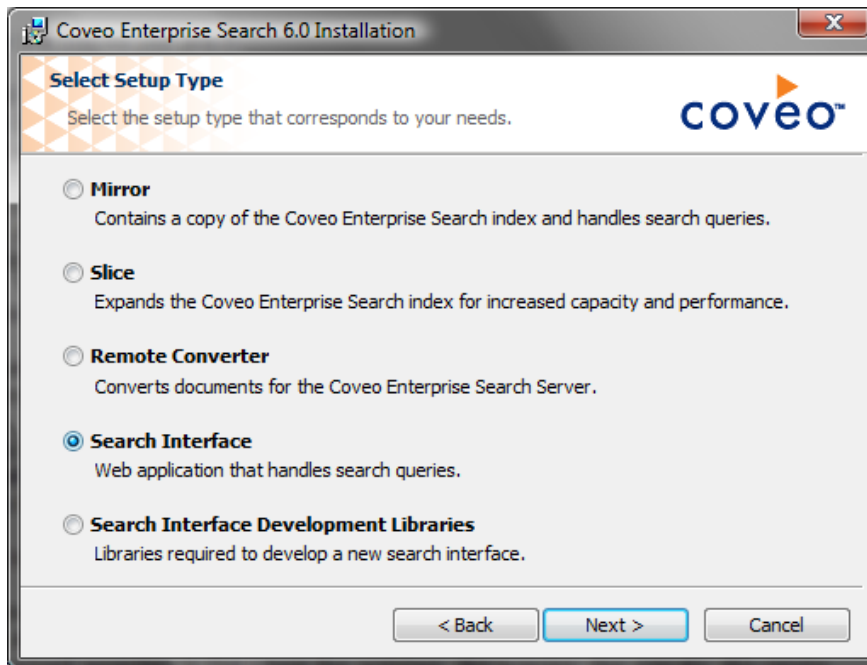
Number of Documents	Recommended Setup
< 5 M documents	One slice per server (the default).
5-10 M documents	Two slices per server.
10-15 M documents	Three slices per server, running on the same machine.
15-20 M documents	Four slices per server. For higher performance, consider using 2 servers.
20 M+	Contact <i>Coveo</i> and consider 64-bit for more scalability.

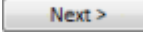
▶ **Single or Multiple Front-End Web Servers, Single CES Server on a Dedicated Machine**

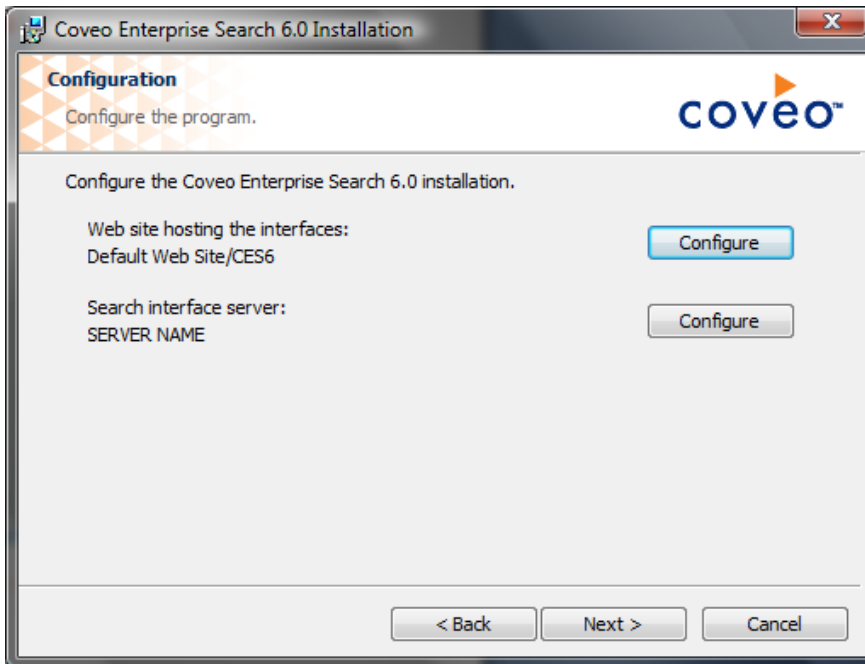


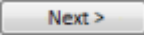
1. Install CES on the dedicated server.
 - a. Run the CES installation kit on the dedicated server, using the default options, and perform the initial setup using the Administration Tool.
2. Install the Search Interface on the front-end server.
 - a. Run the CES installation kit on each Front-End Web server.
 - b. Click . The **Choose Setup Type** page is displayed.

- c. Select the  (**Advanced**) setup.



- d. Select the **Search Interface** option and click . The **Configuration** page is displayed.



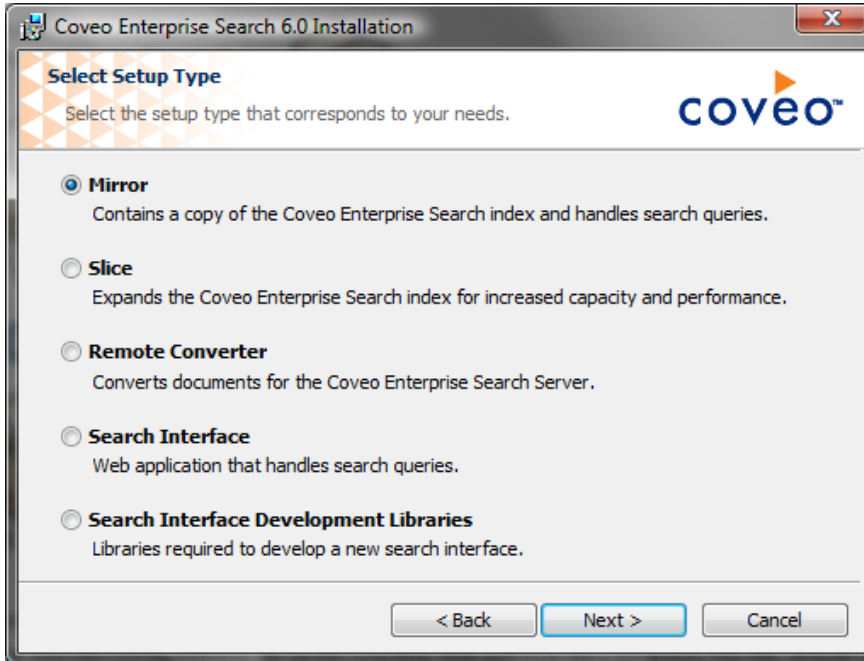
- e. Click the **Configure** button next to **Web site hosting the interfaces**.
 - f. Select the appropriate option.
 - i. **Create a new Web site:** Enter the appropriate information in the **Web site name** and **Web Site port** fields.
 - ii. **Use an existing Web site and create a new virtual directory:** Click the appropriate existing Web site and virtual directory name.
- Note:** For an easier administration, it is suggested to create a new Web site and use the same Web Interface Configuration on all front-ends.
- g. Click .
 - h. Click the **Configure** button next to **Search interface server** and enter name of the back-end server on which the CES index resides.

▶ **Multiple Front-End Web Servers, Multiple CES Mirrors on Dedicated Machines**

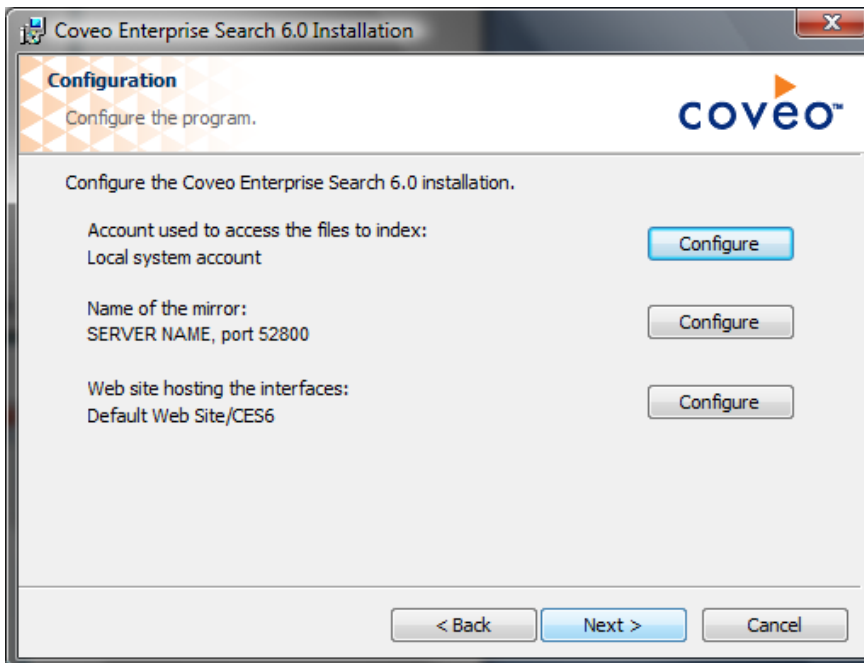
1. Install the master CES server on a dedicated machine.
 - a. Run the CES installation kit on the machine that will host the CES server. Use the default options and perform the initial setup using the Administration Tool.
2. Install the mirror CES servers on their dedicated machines.
 - a. Run the CES installation kit on each machine that will host an additional CES mirror. Select



(Advanced), then the **Mirror** option.



- b. Click . The **Configuration** screen is displayed.



- c. Click the **Configure** button next to **Account used to access the files to index**.
 d. Select **Use this account to index local and network files**.
 e. Enter the user account, the password and the domain that will run this service.

Important: It is essential to use the same account as the one used by the master index.

- f. Click the **Configure** button next to **Name of the mirror**. The **Mirror Configuration** page is displayed. The name of the server is the mirror name. It is suggested to leave it as-is.
 g. Click the **Configure** button next to **Web site hosting the interfaces**.

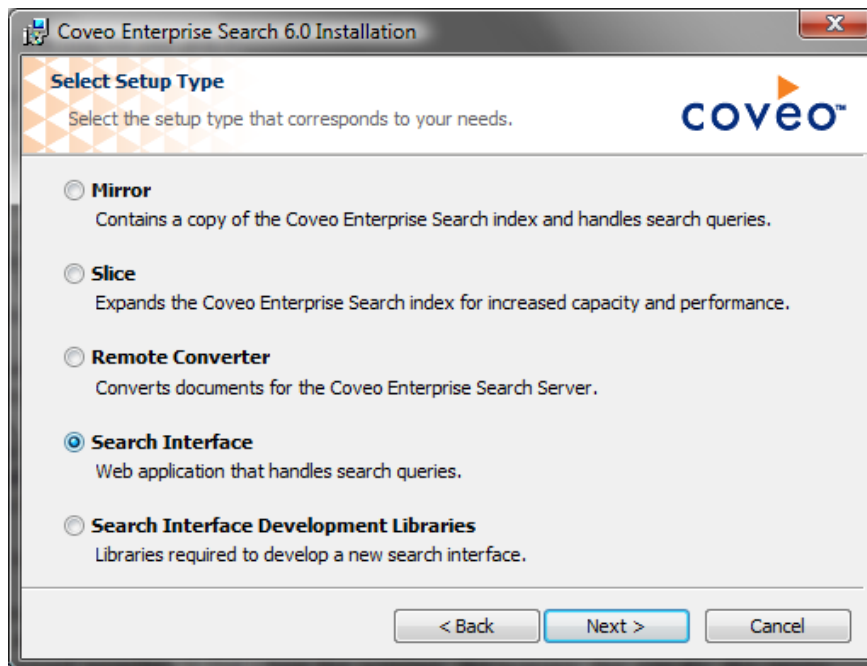
- h. Select the appropriate option, either Create a new Web site or Use an existing Web site and create a new virtual directory.
 - i. **Create a new Web site:** Enter the appropriate information in the **Web site name** and **Web Site port** fields.
 - ii. **Use an existing Web site and create a new virtual directory:** Click the appropriate existing Web site and virtual directory name.

Note: For an easier administration, it is suggested to create a new Web site and use the same Web Interface Configuration on all front-ends.

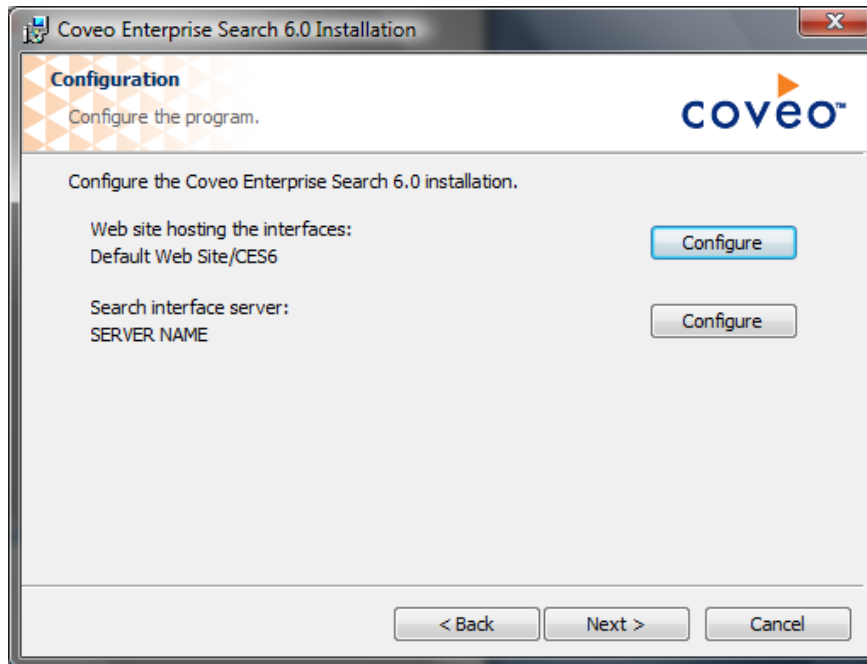
- b. Click .
- 3. Configure network load balancing between the front-end Web servers and CES servers.
 - a. Using a software or hardware load balancer, create an NLB cluster that contains the master CES server and all the mirrors. Setup a DNS name that points to this cluster.
 - b. A typical setup uses *Microsoft Network Load Balancing* (NLB). For more information, refer to <http://go.microsoft.com/fwlink/?LinkId=18371>.
- 4. Install the Search Interface on the front-end servers.



- a. Run the CES installation kit on each front-end Web server. Select **Search Interface** option. (Advanced), then the **Search Interface** option.

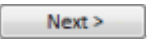


- b. Click . The **Configuration** page is displayed.



- c. Click the **Configure** button next to **Web site hosting the interfaces**.
- d. Select the appropriate option.
 - i. **Create a new Web site:** Enter the appropriate information in the **Web site name** and **Web Site port** fields.
 - ii. **Use an existing Web site and create a new virtual directory:** Click the appropriate existing Web site and virtual directory name.

Note: For an easier administration, it is suggested to create a new Web site and use the same Web Interface Configuration on all front-ends.

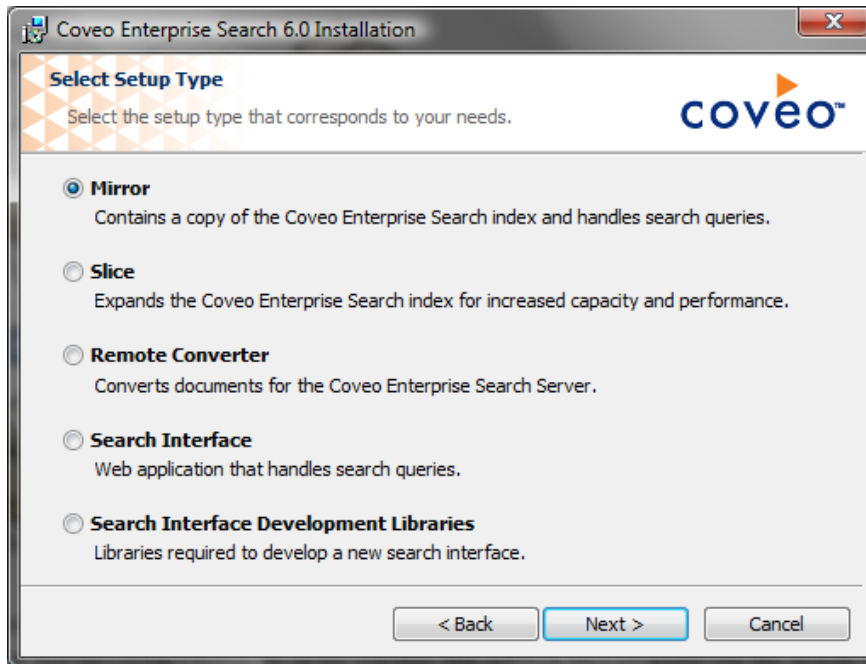
- e. Click .
- f. Click the **Configure** button next to the **Search Interface server** and enter name of the back-end server on which the CES index resides.

▶ **Multiple Front-End Web Servers, Multiple CES Mirrors with Multiple Slices**

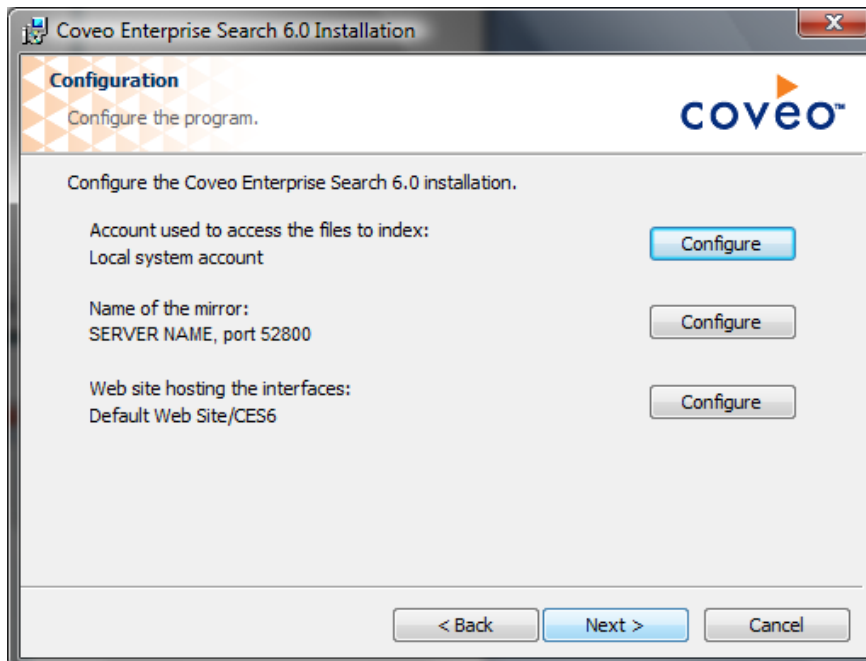
1. Install the master CES server on a dedicated machine.
 - a. Run the CES installation kit on the machine that will host the CES server. Use the default options and perform the initial setup using the Administration Tool.
2. Install the mirror CES servers on their dedicated machines.
 - a. Run the CES installation program on each machine that will host an additional CES mirror. Select



(Advanced), then the **Mirror** option.



- b. Click . The **Configuration** page is displayed.



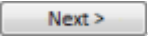
- c. Click the **Configure** button next to **Account used to access the files to index**.
 d. Select **Use this account to index local and network files**.
 e. Enter the user account, the password and domain that will run this service.

Important: It is essential to use the same account as the one used by the master index.

- f. Click the **Configure** button next to name of the mirror. The **Mirror Configuration** page is displayed. The name of the server is the mirror name. It is suggested to leave as-is.
 g. Click the **Configure** button next to Web site hosting the interfaces.

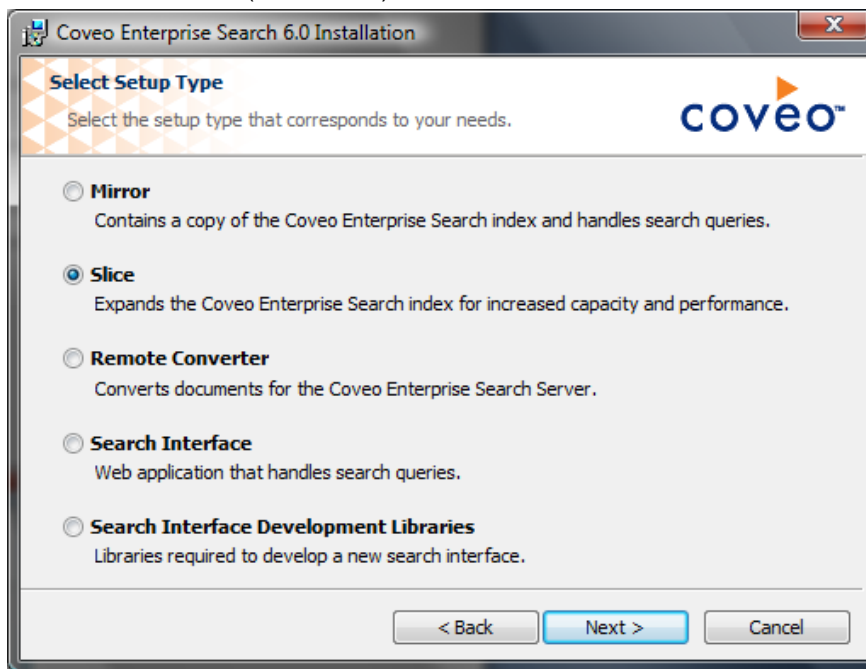
- h. Select the appropriate option.
 - i. **Create a new Web site:** Enter the appropriate information in the **Web site name** and **Web Site port** fields.
 - ii. **Use an existing Web site and create a new virtual directory:** Click the appropriate existing Web site and virtual directory name.

Note: For an easier administration, it is suggested to create a new Web site and use the same Web Interface Configuration on all front-ends.

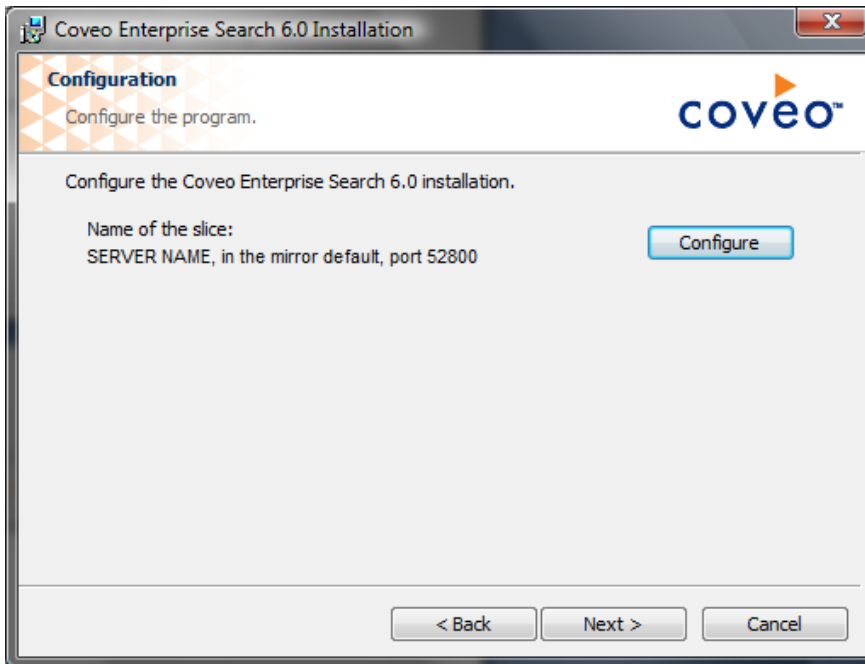
- i. Click .
3. If you are deploying remote slices, install the slice CES servers on their dedicated machines.
 - a. Run the CES installation kit on each machine that will host one or several remote CES slices.



- b. Select **(Advanced)**, then **Slice**.



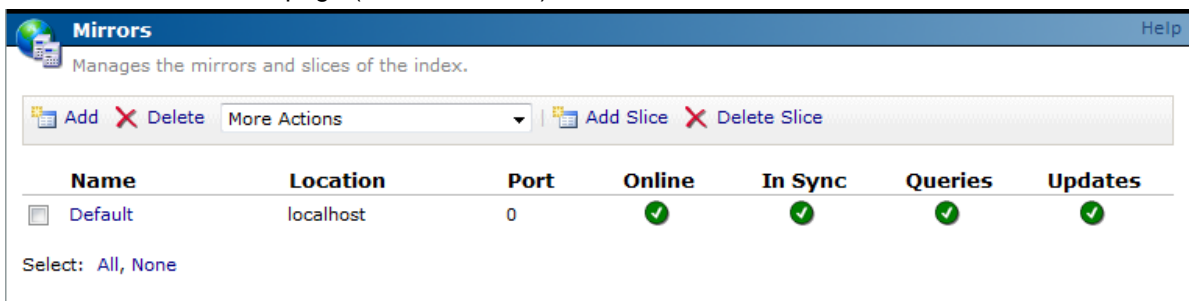
- c. In the **Configuration** page, click the **Configure** button next to **Name of the slice**.



- d. Enter the name of the first slice on this machine. Proceed with the default settings.
4. If some servers host more than one remote slice, create them using the command prompt (only one can be created using the installation kit).
 - a. Open a command shell and enter `cd C:\Program Files\Coveo Enterprise Search x\Bin.`
 - b. Create the additional slice using this command: `CESSliceService.exe -register -mirror <mirror name> -slice <slice name> -User <domain\username> <password> -automatic -verbose.`

Note: `<domain\user>` `<password>` should be the same as in the installation of the master/mirror server.

 - c. Access the **Windows Services MMC** tab and confirm that a new service has been registered for the new slice.
 - d. Start the service if it was not done automatically.
5. Register the additional slices in the Administration Tool.
 - a. Open the CES Administration Tool on the master CES server.
 - b. Access the **Mirrors** page (Index > Mirrors).



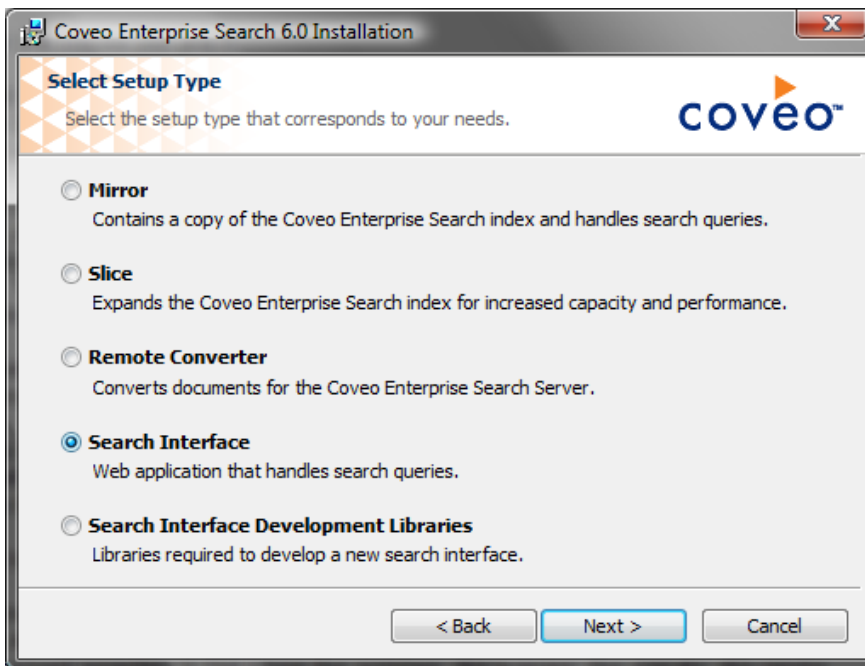
Because each mirror must contain the exact same number of slices, it is not necessary to register each physical slice individually. Instead, for each slice added, you will be asked for parameters for each mirror.

- c. Click **Add Slice**.
- d. Enter the name of the slice.

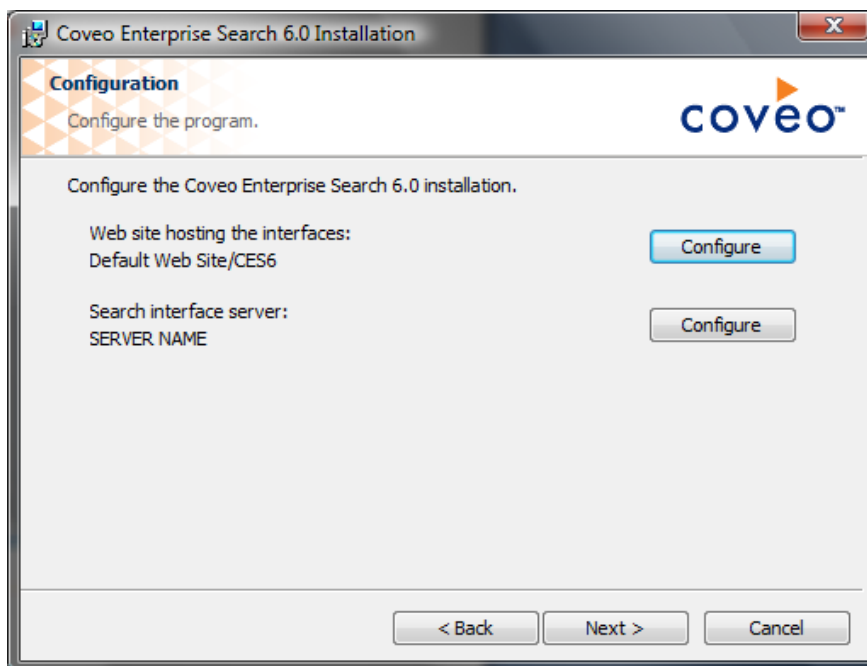
- e. For each mirror, provide the parameters for the slice. If the slice is on the same machine as the mirror, use the default settings. If the slice is remote, it is important to enter the name of the server that hosts it, as well as the path in which its files should be saved. By default, a remote converter will also be created on the remote machine to increase the indexing speed.
6. Configure network load-balancing between the front-end Web servers and the CES servers.
 - a. Using a software or hardware load balancer, create an NLB cluster that contains the master CES server and all the mirrors (**not** the dedicated slice servers).
 - b. Setup a DNS name that points to this cluster.
 - c. A typical setup uses *Microsoft* NLB. For more information, refer to <http://go.microsoft.com/fwlink/?LinkId=18371>.
7. Install the Search Interface on the front-end servers.



- a. Run the CES installation kit on each front-end Web server. Select **(Advanced)**, then the **Search Interface** option.

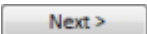


- b. Click . The **Configuration** page is displayed.



- c. Click the **Configure** button next to **Web site hosting the interfaces**.
- d. Select the appropriate option.
 - i. **Create a new Web site:** Enter the appropriate information in the **Web site name** and **Web Site port** fields.
 - ii. **Use an existing Web site and create a new virtual directory:** Click the appropriate existing Web site and virtual directory name.


Note: For an easier administration, it is suggested to create a new Web site and use the same Web Interface Configuration on all front-ends.

- e. Click .
- f. Click the **Configure** button next to the **Search Interface server** and enter name of the back-end server on which the CES index resides.


Naming Remote Slices

The names of the slices in a CES system can be different from the name of the server on which they are hosted. Each slice in a multi-mirror deployment must have the same name as the corresponding one in any other mirror. The names of the remote slices are specified during the installation procedure on the server that hosts the slice. The following displays a typical setup:

CESS Master + Default Slice



CESS Mirror + Default Slice



In this setup, each mirror has 3 slices. One of them is local; therefore, it lives directly in the mirror's process. The two other are remote. Since each corresponding slice in a mirror must share the same name as the ones from other mirrors, and that the name for the slices are 'A' and 'B', the administrator must install a CES slice named 'A' on two machines, and a CES slice named 'B' on two others.

Conclusion

Coveo's personnel is available to discuss any part of this document, as architecture decisions are made. *Coveo Enterprise Search for Intranet* has been designed to function seamlessly with *SharePoint*, in any server configuration. Again, contact *Coveo* for any questions you might have.