



Feith Install Guide
for
Elasticsearch

Updated 6/20/2016

**Feith Install Guide
for
Elasticsearch**

© Copyright 2016 Feith Systems and Software, Inc. All Rights Reserved. No part of this publication may be reproduced, transmitted, stored in a retrieval system, or translated into any language in any form by any means, without the written permission of Feith Systems and Software, Inc.

All information in this work is subject to change and reflects software current at the time of publication.

FDD and the FDD logo are trademarks of Feith Systems and Software, Inc. All other trademarks, product names and company names may be trademarks or registered trademarks of their respective holders.

Feith Systems and Software, Inc.
425 Maryland Drive
Fort Washington, PA 19034
Tel (215) 646-8000
Fax (215) 540-5495
www.feith.com

Table of Contents

Introduction	4
Requirements.....	4
Install Elasticsearch	5
Start Elasticsearch	21
Customize Feith Authentication Plugin Credentials	22
Review and Update Stop Words	23

Introduction

Elasticsearch is a third-party engine used to power full text search in Feith's client applications (e.g. FDD, WebFDD). Full text search provides an alternative to structured metadata searches, allowing users to search for documents by the text in their pages.

This installer includes:

- The Elasticsearch product
- The Feith Authentication Plugin for Elasticsearch.

Once Elasticsearch is installed, you may want to:

- [Set the Elasticsearch service to start automatically](#)
- [Customize the credentials for the Feith Authentication Plugin](#)
- [Review and maybe update the stop words ES uses](#)

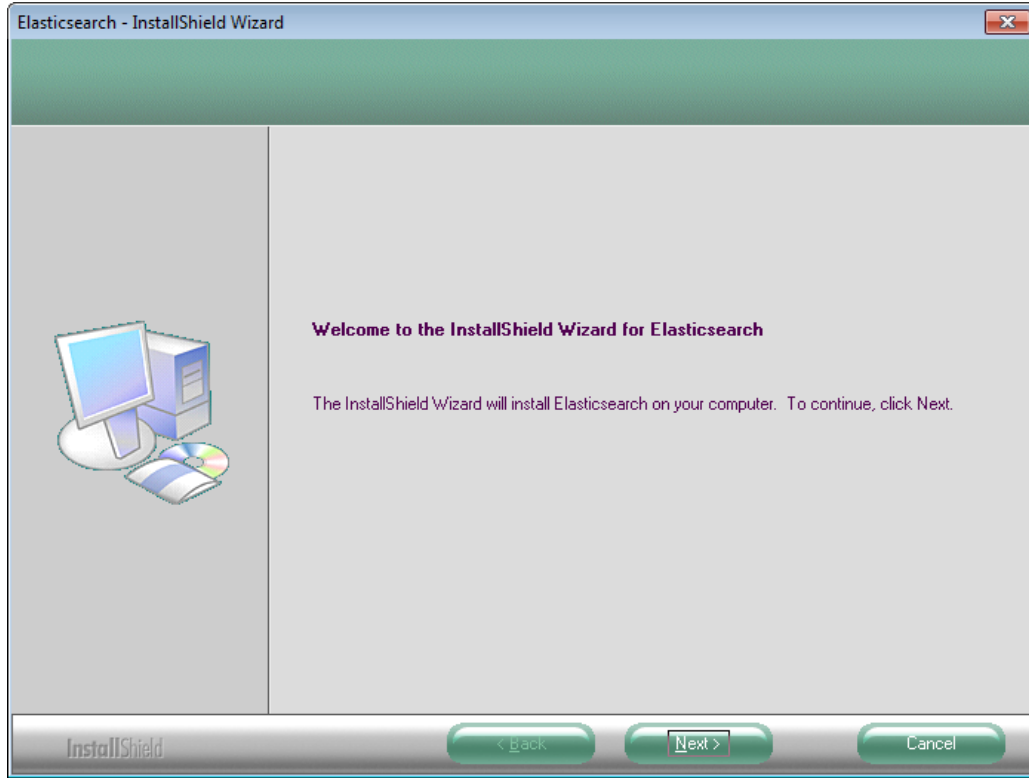
Requirements

Java is required and the JAVA_HOME must be set.

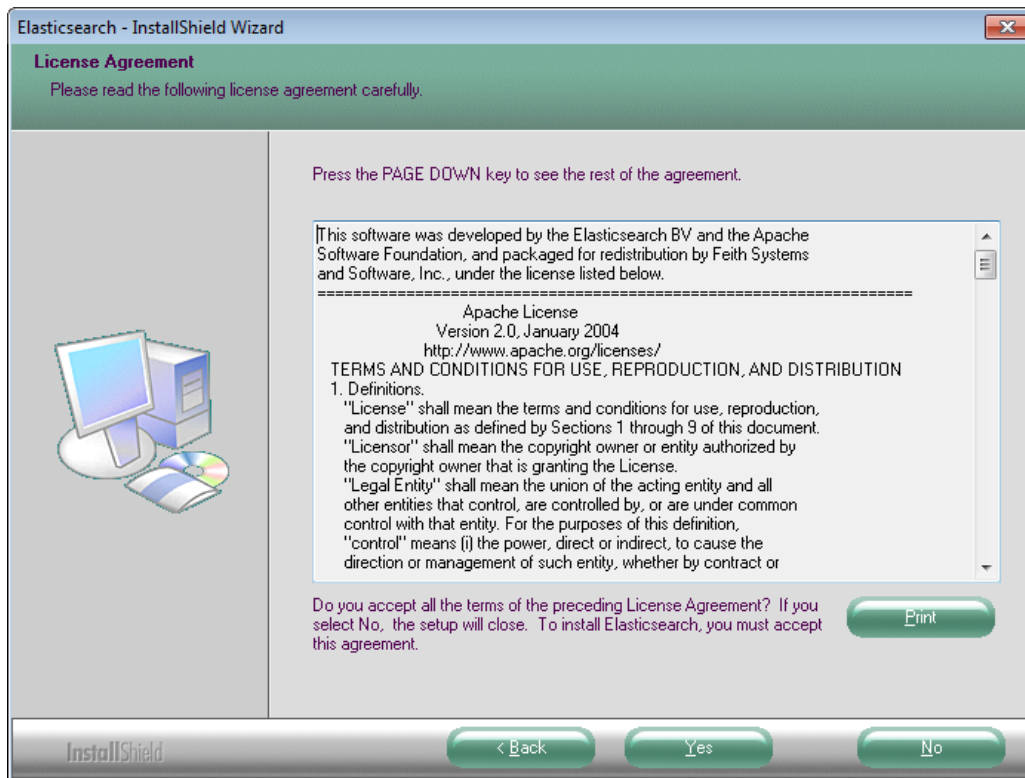
Install Elasticsearch

To install Elasticsearch:

1. Run the Elasticsearch (ES) installation package. The **Welcome to the InstallShield Wizard for Elasticsearch** opens. Click **Next** to continue.

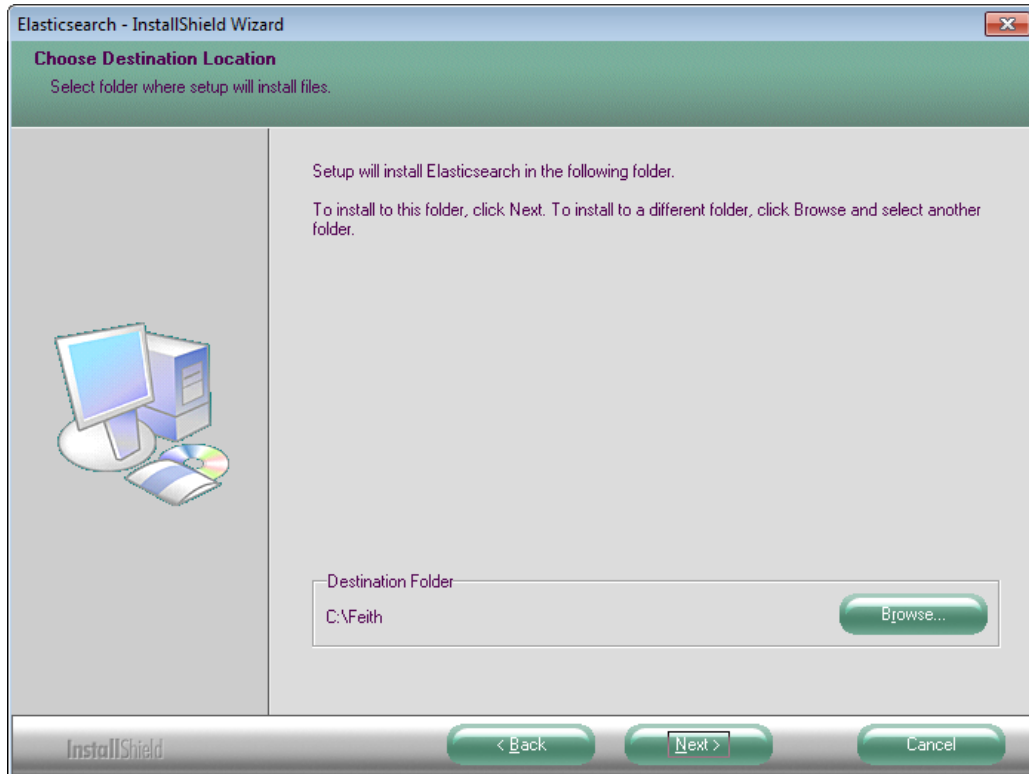


2. On the **License Agreement** screen, click **Yes** if you agree to the terms.



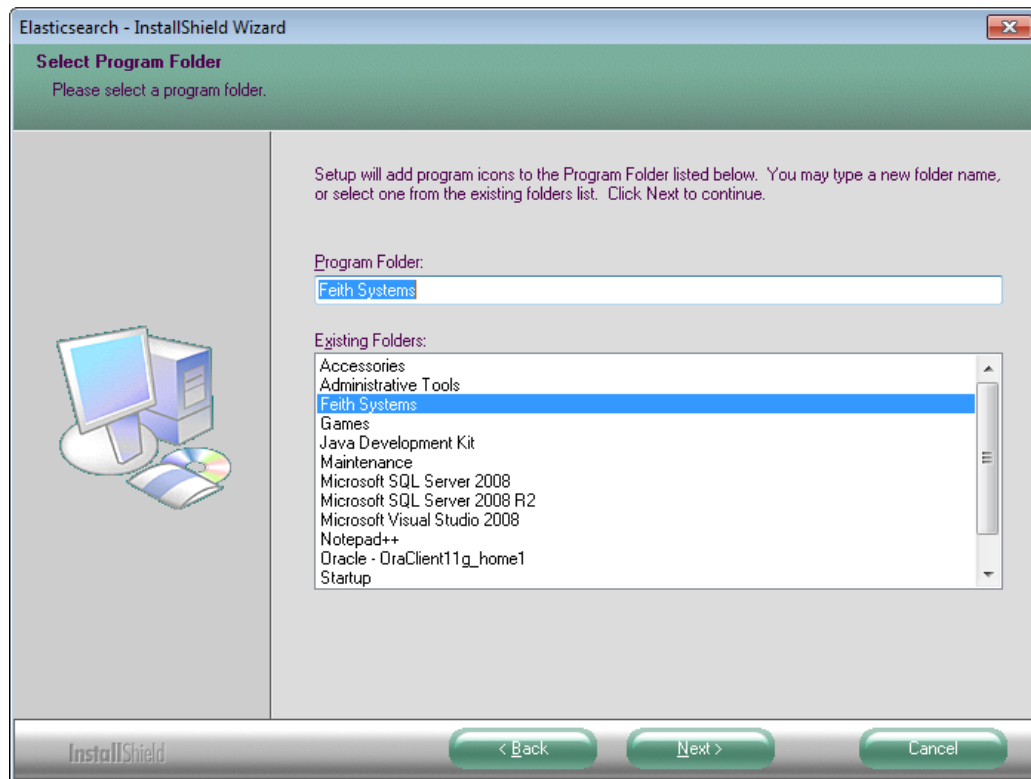
3. On the **Choose Destination Location** screen, choose the folder where Elasticsearch will be installed. Click **Next** to continue.

The default location is **C:\Feith**. To choose another location, click **Browse** and select another folder.



4. On the **Select Program Folder** screen, choose where you want Elasticsearch in the Programs list. Click **Next** to continue.


The default program folder is **Feith Systems**.



5. A message informs you that the next few settings are required to configure Elasticsearch.

In a typical installation you can accept the defaults for the ES settings that follow. If you are unsure what to enter, just accept the defaults. Terms like “cluster” and “node” and more will be explained along the way.

Click **OK** to continue.

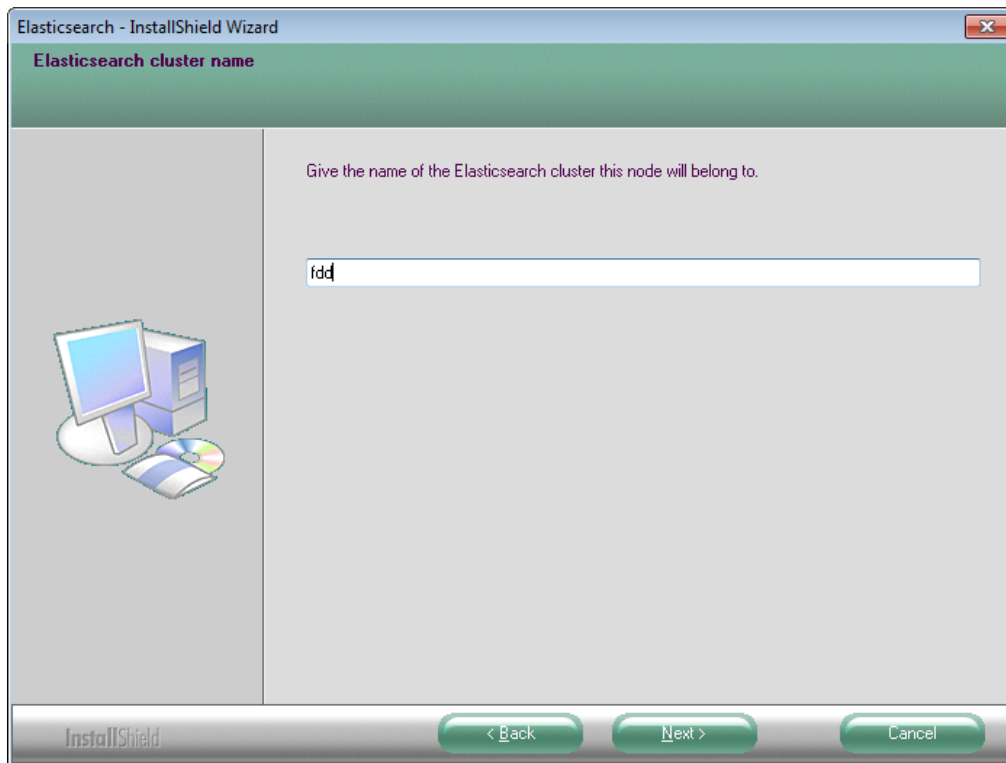


6. On the **Cluster Name** screen, enter the name of the ES cluster to which this node belongs. Click **Next** to continue.

A “node” is a machine where Elasticsearch is installed. The node may function on its own or, if needed, you can scale up ES and have multiple nodes work together in a “cluster” sharing the storage and work.

Usually you will have only one node in one cluster. Most systems do not need to scale up to multiple nodes.

In a typical installation you can accept the default. If you are unsure what to enter, just accept the default.

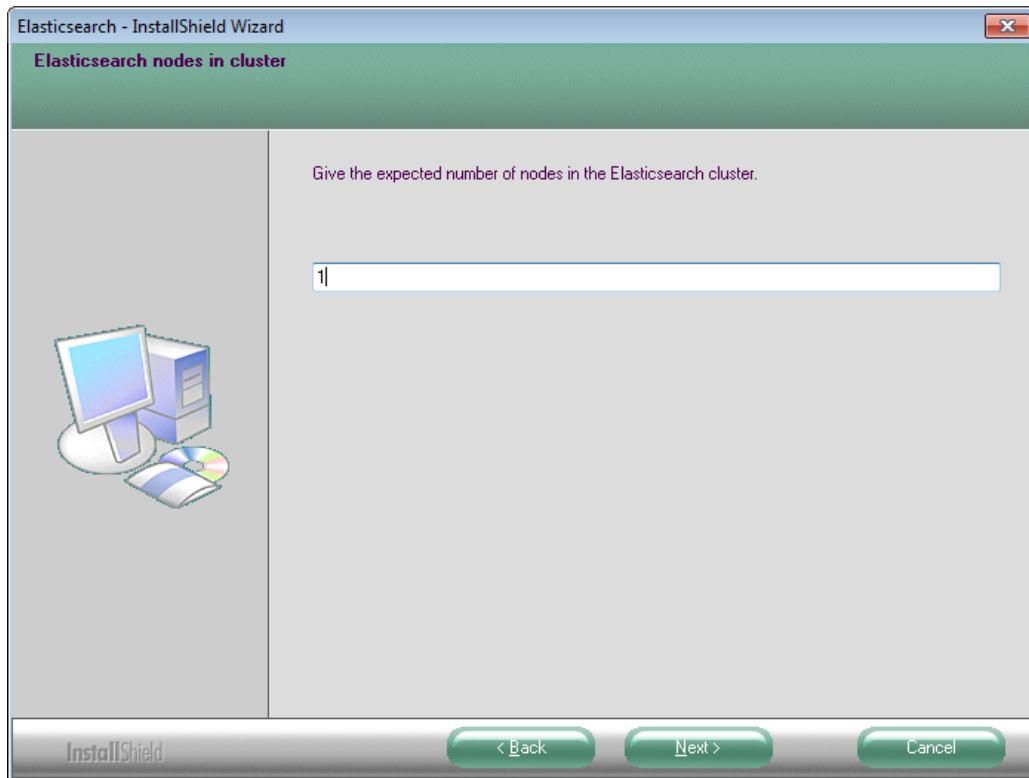


7. On the **Nodes in Cluster** screen, enter the number of nodes that are expected to be in the Elasticsearch cluster.

Enter the number of “nodes” i.e. the number of machines running Elasticsearch, that you expect to be in the ES cluster. You can increase this number later if needed.

In a typical installation you can accept the default. If you are unsure what to enter, just accept the default.

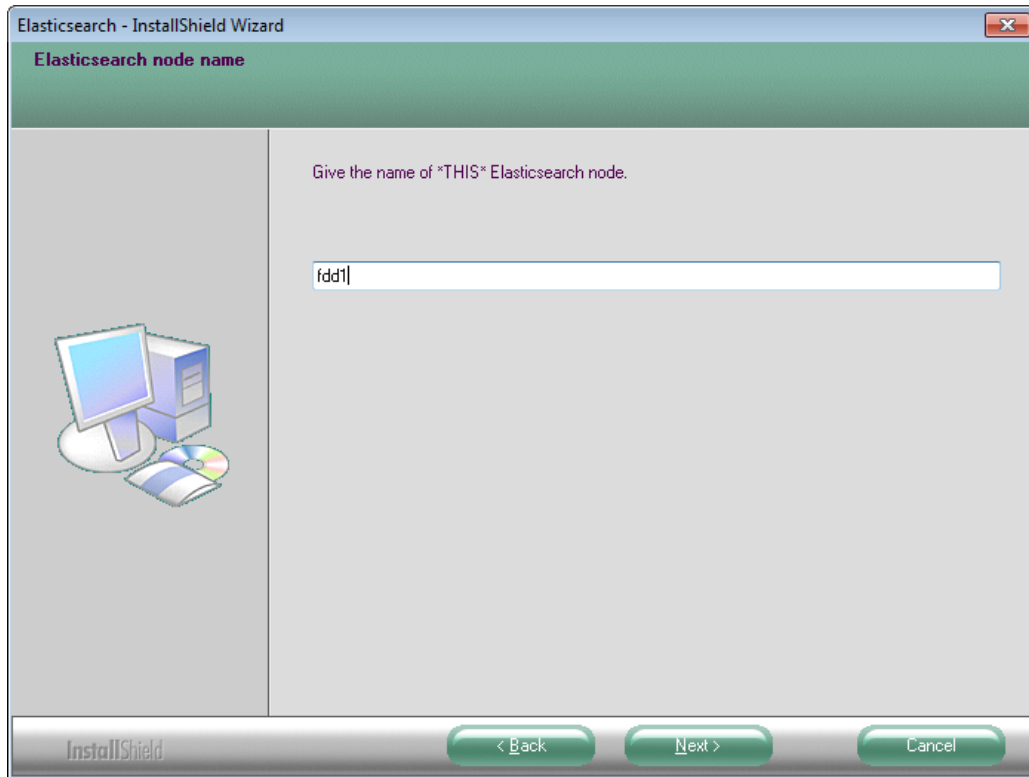
Click **Next** to continue.



8. On the **Node Name** screen, enter the name of this specific node. If there are other nodes in the cluster, this name should be unique from the other nodes' names.

In a typical installation you can accept the default. If you are unsure what to enter, just accept the default.

Click **Next** to continue.




9. On the **Other ES Hosts** screen, enter the names of the other machines – if any – hosting the Elasticsearch nodes in this cluster. Each name is enclosed in double quotes and separated by commas.
For example: **“fdd2”, “fdd3”**

If this is the only Elasticsearch host (“node”) in this cluster, just leave this blank.

In a typical installation you can accept the default. If you are unsure what to enter, just accept the default.

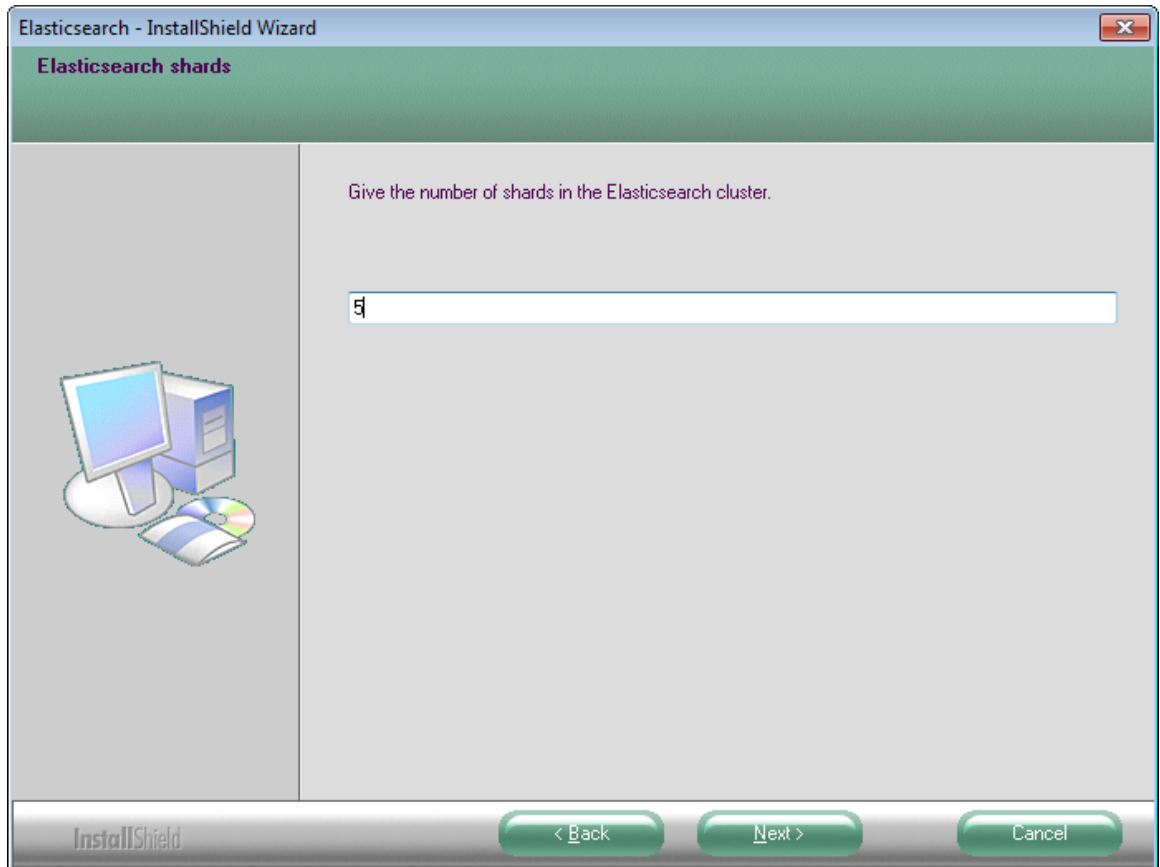
Click **Next** to continue.



10. On the **Shards** screen, enter the number of shards that the Elasticsearch index will be broken up into and leave room for growth. Once this number is set here, it cannot be changed later.

An “index” is the ES database that will contain the data from your FDD database to power full text searching. An index can be broken up into “shards,” or chunks, for a system that needs to scale up and spread the work across multiple machines. For a typical system, five shards should be fine.

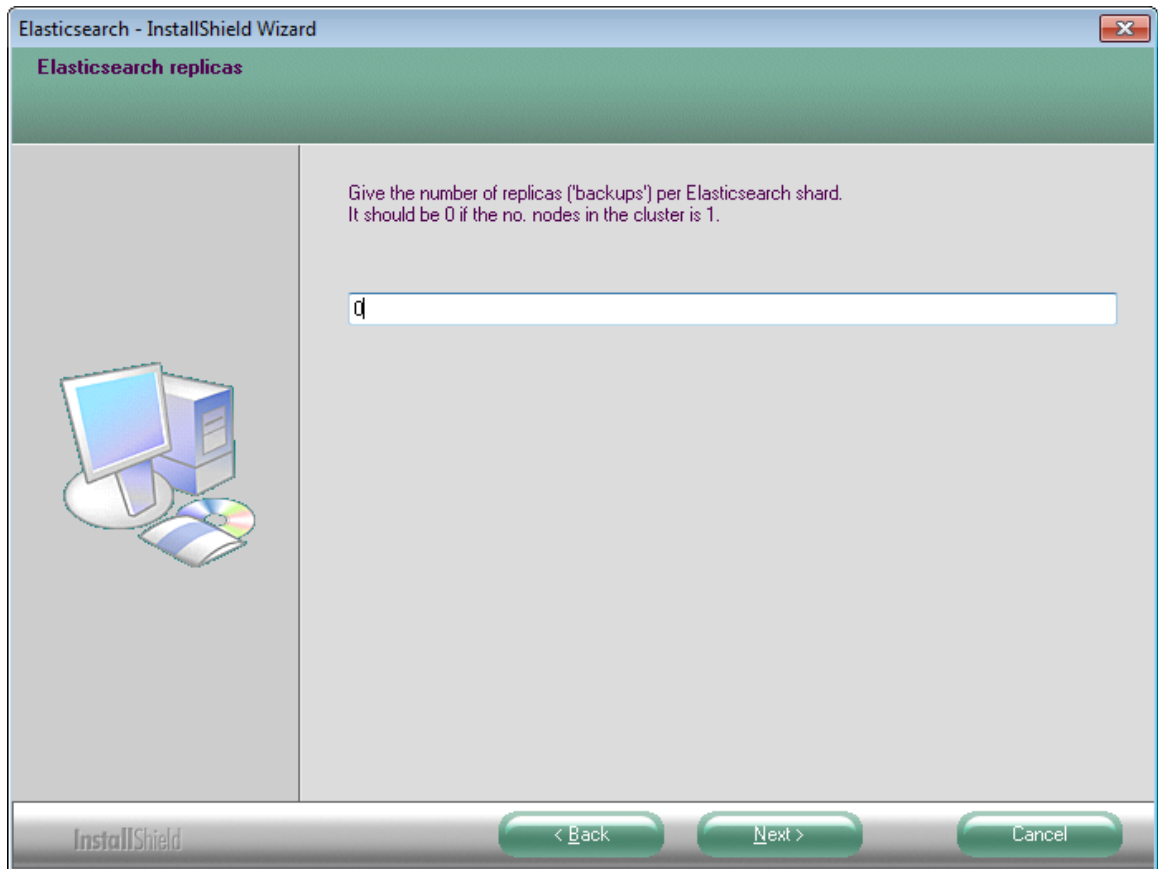
In a typical installation you can accept the default. If you are unsure what to enter, just accept the default.



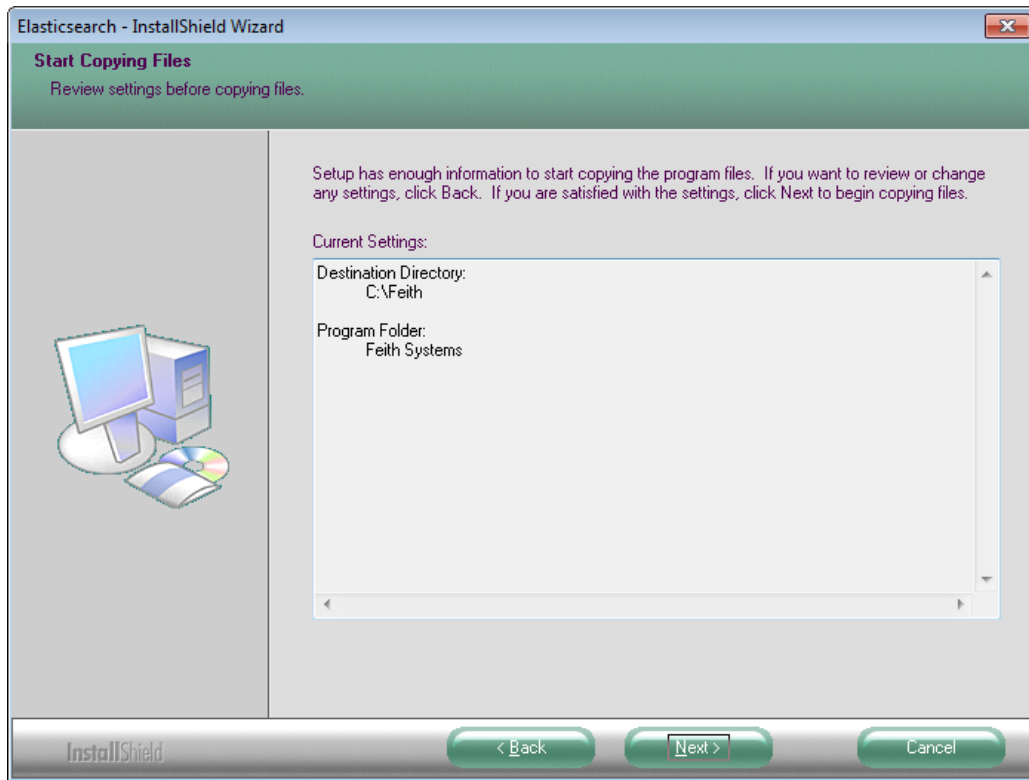
11. On the **Replicas** screen, enter the number of replicas to store for each shard.

A “replica” is a copy of the data in the shard, like a backup. If you have just one node, you cannot have any replicas. If you have multiple nodes, you can set up replicas. If one node goes down the system may be able to continue to run normally using the replica until the node is restored.

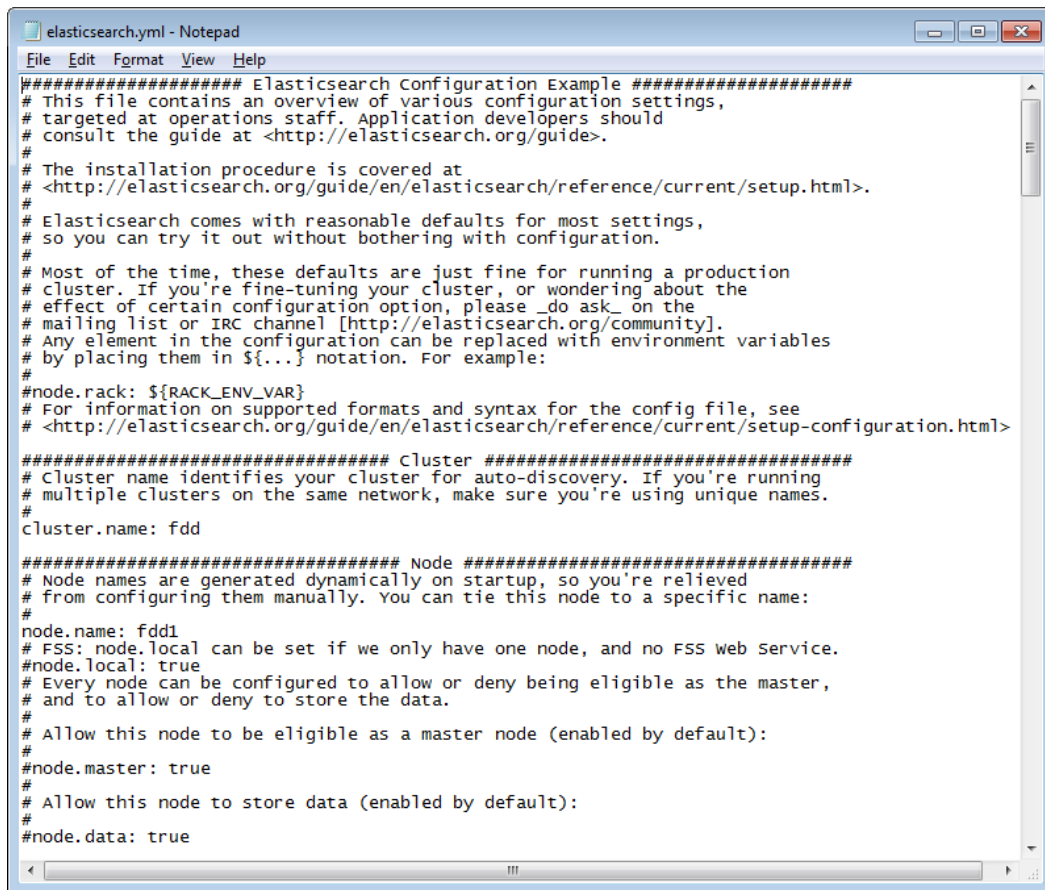
In a typical installation you can accept the default. If you are unsure what to enter, just accept the default.



12. On the **Start Copying Files** screen, click **Next** to start installing Elasticsearch.



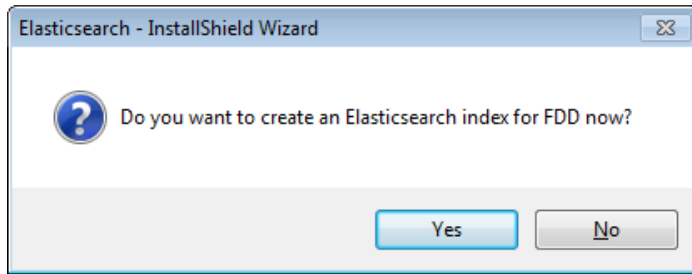
14. The Elasticsearch configuration file (elasticsearch.yml) displays. Make changes if desired. Save and close the file.



```
elasticsearch.yml - Notepad
File Edit Format View Help
##### Elasticsearch Configuration Example #####
# This file contains an overview of various configuration settings,
# targeted at operations staff. Application developers should
# consult the guide at <http://elasticsearch.org/guide>.
#
# The installation procedure is covered at
# <http://elasticsearch.org/guide/en/elasticsearch/reference/current/setup.html>.
#
# Elasticsearch comes with reasonable defaults for most settings,
# so you can try it out without bothering with configuration.
#
# Most of the time, these defaults are just fine for running a production
# cluster. If you're fine-tuning your cluster, or wondering about the
# effect of certain configuration option, please do ask on the
# mailing list or IRC channel [http://elasticsearch.org/community].
# Any element in the configuration can be replaced with environment variables
# by placing them in ${...} notation. For example:
#
#node.rack: ${RACK_ENV_VAR}
# For information on supported formats and syntax for the config file, see
# <http://elasticsearch.org/guide/en/elasticsearch/reference/current/setup-configuration.html>
##### Cluster #####
# Cluster name identifies your cluster for auto-discovery. If you're running
# multiple clusters on the same network, make sure you're using unique names.
#
cluster.name: fdd


##### Node #####
# Node names are generated dynamically on startup, so you're relieved
# from configuring them manually. You can tie this node to a specific name:
#
node.name: fdd1
# FSS: node.local can be set if we only have one node, and no FSS web service.
#node.local: true
# Every node can be configured to allow or deny being eligible as the master,
# and to allow or deny to store the data.
#
# Allow this node to be eligible as a master node (enabled by default):
#
#node.master: true
#
# Allow this node to store data (enabled by default):
#
#node.data: true
```

15. A message asks if you want to create an ES index for your FDD database now. Click **Yes** to make the index now. Click **No** to skip this step and do it later.



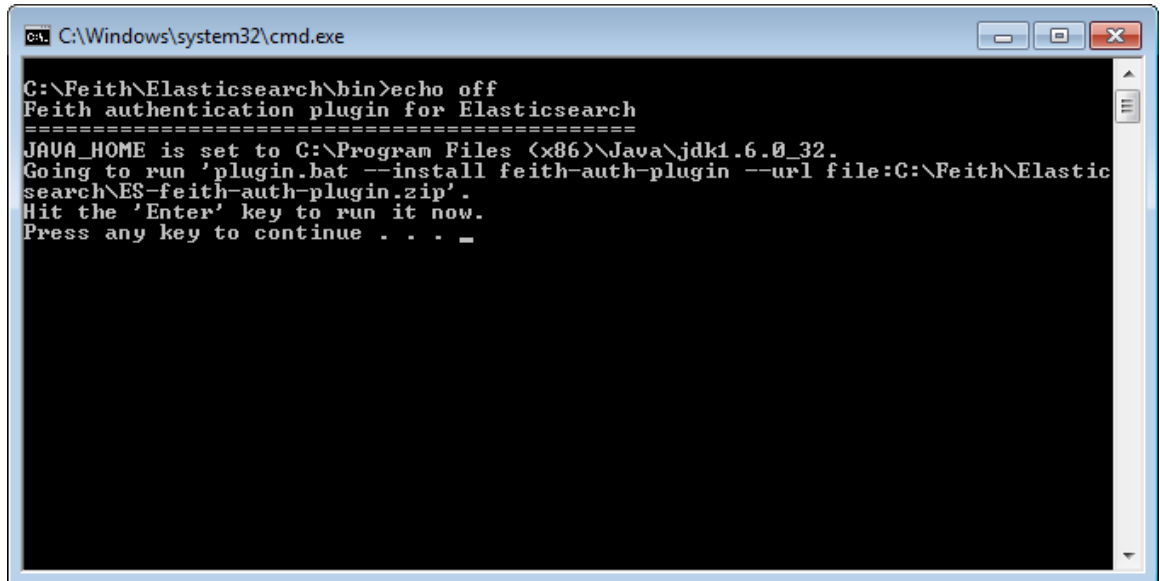
If you chose **Yes**:

- a. A command prompt opens and asks what you want to name the new index.



- b. Enter the desired name. Using a name identical to your FDD database's connect string is usually a good choice.
- c. Hit the **ENTER** key. The index is created and a **Status code: 200** is returned to confirm success.
- d. Hit **ENTER** to exit the script.

16. A command prompt opens to confirm installation of the Feith Authentication Plugin for ES. Hit the **ENTER** key to install the plugin. The plugin is installed.



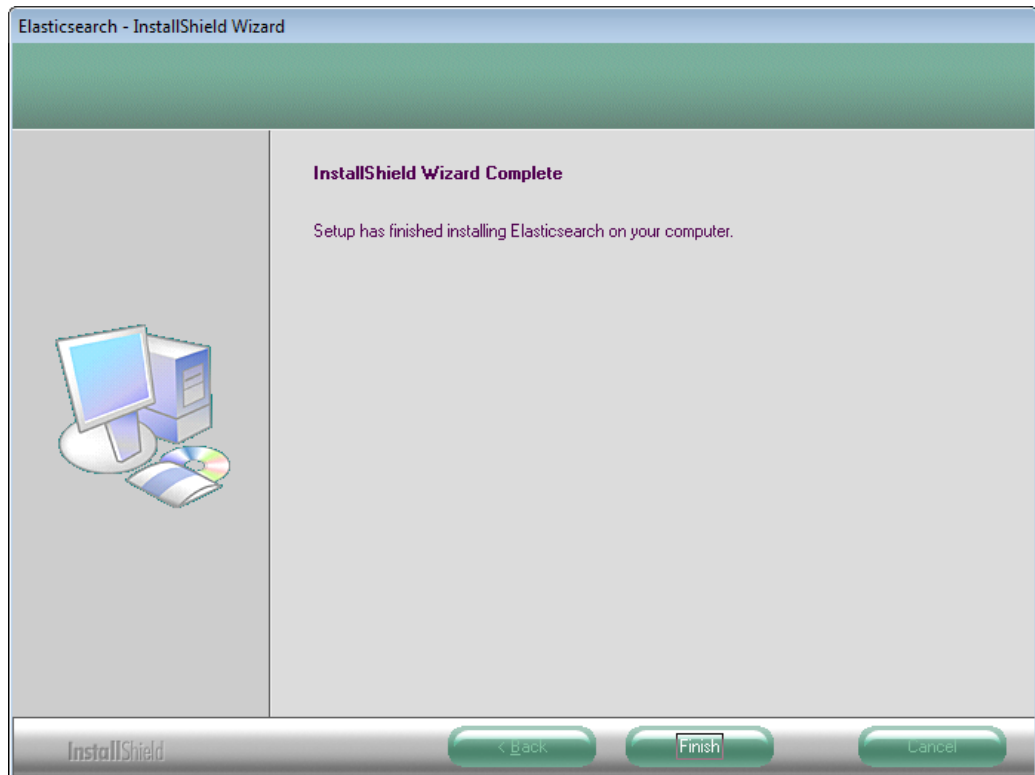
```
C:\Windows\system32\cmd.exe

C:\Feith\Elasticsearch\bin>echo off
Feith authentication plugin for Elasticsearch
=====
JAVA_HOME is set to C:\Program Files (x86)\Java\jdk1.6.0_32.
Going to run 'plugin.bat --install feith-auth-plugin --url file:C:\Feith\Elastic
search\ES-feith-auth-plugin.zip'.
Hit the 'Enter' key to run it now.
Press any key to continue . . . _
```

17. On the **InstallShield Wizard Complete** screen, click **Finish** to complete the installation.

You may want to:

- [Set the Elasticsearch service to start automatically](#)
- [Customize the credentials for the Feith Authentication Plugin](#)
- [Review and maybe update the stop words ES uses](#)



Start Elasticsearch

By default, the Elasticsearch service is set to a Manual start, instead of Automatic.

To start ES on Windows, open the **Services** module, select the service called **Elasticsearch**, and start the service.

To set ES to start automatically when the server machine starts:

1. In the Windows services module, right-click the **Elasticsearch** service and select **Properties**. The **Properties** dialog opens.
2. Change the **Startup type** to **Automatic**.
3. Click **OK**. The **Elasticsearch** service will start automatically when the server starts.

Customize Feith Authentication Plugin Credentials

When you use the Feith installer for Elasticsearch, the Feith Authentication Plugin is installed and credentials must be given in order to index and search for documents in ES.

The credentials are at the bottom of ES's configuration file, **elasticsearch.yml**, stored in ES's **config** directory (e.g. C:\Feith\ElasticSearch\config). For example:

```
com.feith.es.plugin.token: a1b2c3d4e5  
com.feith.es.adminUsers:  
admin : z26y25x24
```

The password token is used by programs that do full text searches (e.g. FDD Client). The administrator credentials are used by UTR.

We recommend you change the default user name and password to the credentials you want to use.

To customize the password token for clients that do full text searches:

1. Change the password token given in **com.feith.es.plugin.token** in ES's configuration file **elasticsearch.yml**. For example

```
com.feith.es.plugin.token: 14m5n6
```

Note that the password is encrypted, so you need to encrypt the password using the **fencrypt** program installed in ES's **bin** directory (e.g. C:\Feith\Elasticsearch\bin).


2. Change the ES security info keys in the system information table to match the new value in ES. Update **fdd.system_info** for the row with obj_name **ES_SecurityInfoKeys** and set that row's obj_desc to the new password you just set in ES. For example:

OBJ_TYPE	OBJ_NAME	OBJ_DESC
M	ES_SecurityInfoKeys	14m5n6

To customize the administrator credentials for UTR:

1. Change the user name and password given in **com.feith.es.adminUsers** in ES's configuration file **elasticsearch.yml**. For example:

```
com.feith.es.adminUsers:  
administrator : o7p8q9
```



Note that the password is encrypted, so you need to encrypt the password using the **fencrypt** program installed in ES's **bin** directory (e.g. C:\Feith\Elasticsearch\bin).

2. Change the **ESUser** and **ESEncryptedPassword** in UTR's initialization file **utr.ini** to match the new user name and password you just set in ES. For example:

```
ESUser administrator  
ESEncryptedPassword o7p8q9
```

See **UTR Manual** for more information on configuring ESUser and ESEncryptedPassword.

Review and Update Stop Words

The “stop words,” which Elasticsearch ignores, are stored in a file called **fss_stopwords.txt**. The stop words file is stored in ES’s **config** directory (e.g. C:\Feith\ElasticSearch\config). ES’s configuration file, **elasticsearch.yml**, has a setting at the bottom that points to the stop words file. For example:

```
stopwords_path: C:/Feith/Elasticsearch/config/fss_stopwords.txt
```

You may review the stop words list in a text editor (e.g. Notepad) and modify it if desired. Each word should be on a separate line. The list of stop words we recommend are those already in fss_stopwords.txt. In addition to these words, you may add words that are excessively common in your organization (e.g. your organization’s name may be in every one of your archived emails).

Revision History

Date	Version	Author	Description
6/20/2016	1.7.4	GJG	Added clarification for customizing security and reviewing stop words.
3/11/2016	1.7.4	GJG	New documentation.