

# Content Manager

Software Version 9.3

## Elasticsearch Integration Installation and Configuration Guide



Document Release Date: August 2018

Software Release Date: August 2018

## Legal notices

### Copyright notice

© Copyright 2017-2018 Micro Focus or one of its affiliates.

The only warranties for products and services of Micro Focus and its affiliates and licensors (“Micro Focus”) are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Micro Focus shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

### Trademark notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

## Documentation updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To verify you are using the most recent edition of a document, go to <https://softwaresupport.softwaregrp.com/manuals>.

You will also receive new or updated editions of documentation if you subscribe to the appropriate product support service. Contact your Micro Focus sales representative for details.

To check for new versions of software, go to <https://www.hpe.com/software/entitlements>. To check for recent software patches, go to <https://softwaresupport.softwaregrp.com/patches>.

The sites listed in this section require you to sign in with a Software Passport. You can register for a Passport through a link on the site.

## Support

Visit the Micro Focus Software Support Online website at <https://softwaresupport.softwaregrp.com>.

This website provides contact information and details about the products, services, and support that Micro Focus offers.

Micro Focus online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support website to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Access the Software Licenses and Downloads portal
- Download software patches
- Access product documentation
- Manage support contracts
- Look up Micro Focus support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require you to register as a Passport user and sign in. Many also require a support contract.

You can register for a Software Passport through a link on the Software Support Online site.

To find more information about access levels, go to

<https://softwaresupport.softwaregrp.com/web/softwaresupport/access-levels>.

# Contents

Getting started with Elasticsearch .....	5
Environment overview .....	5
Install Content Manager 9.2 or later .....	7
Upgrading from earlier versions .....	7
New Installations .....	7
Install Elasticsearch .....	8
Download and install the Java Runtime Environment (JRE) .....	8
Ensure the JAVA_HOME environment variable is configured .....	8
Install Elasticsearch on your nominated servers .....	9
Edit the Elasticsearch configuration file .....	9
Memory Management and configuring the Java Virtual Machine (JVM) .....	12
Configure Elasticsearch to run as a service .....	13
Run Elasticsearch .....	14
Adjusting the JVM heap size after installing Elasticsearch as a service .....	15
Confirm the Elasticsearch service and cluster are operational .....	16
Configure an Elasticsearch Content Index .....	17
Reindex your document store with Elasticsearch .....	17
Configure Content Manager Enterprise Studio to use Elasticsearch .....	22
Remove existing IDOL content indexes .....	22
Create a new Elasticsearch index .....	23
Confirm your Elasticsearch content index is operational .....	26
Troubleshooting common issues .....	27
Logging .....	27
Diagnosing “operation timed out” errors .....	27
Proxy servers and firewalls .....	28
Useful Information .....	29
Kibana .....	29
cURL .....	29
Elasticsearch Reference .....	29
Hardware requirements .....	29

## Getting started with Elasticsearch

With the release of Content Manager 9.2 support for Elasticsearch has been introduced.

Elasticsearch is an open source search and content analytics engine built on top of Apache Lucene™.

If you're an existing customer, we'll help you make the transition from IDOL to Elasticsearch by supporting integration with both products. If you're installing Content Manager for the first time you will only be able to use the Elasticsearch integration.

Unlike IDOL, Elasticsearch will not be shipped with Content Manager (i.e. non-OEM). Instead, you will be required to download Elasticsearch from [elastic.co](https://elastic.co) if you wish to be able to search the content of documents in your Content Manager repository.


This guide will help you understand what's required to establish a basic Content Manager 9.2 and Elasticsearch instance after upgrading from a previously supported release or installing Content Manager 9.2 for the first time.

Before you begin, it's important to note Elasticsearch can be configured and scaled in many different ways. This means specialist Elasticsearch configuration advice and support cannot be provided by the Content Manager team. If, after reading this guide, you believe your organization requires specialist Elasticsearch configuration advice and support, you should consult the Elastic team directly regarding your requirements.

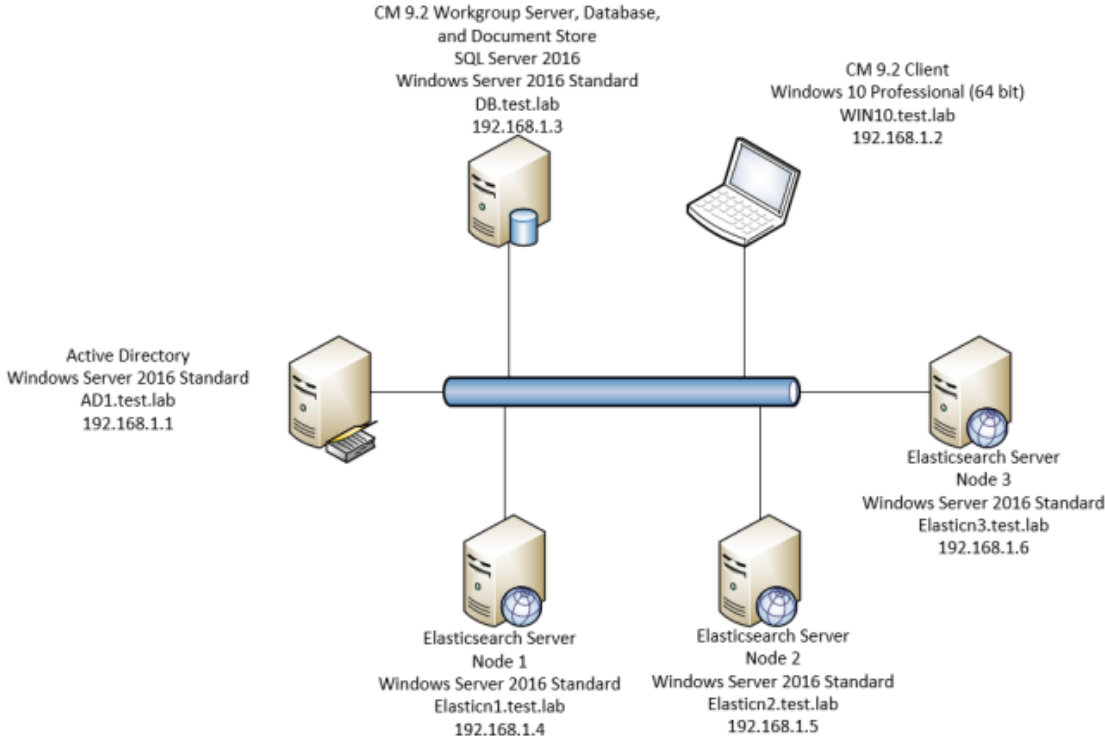
## Environment overview

This guide will begin by running you through a scenario where a customer is upgrading from Content Manager 9.1 and IDOL to Content Manager 9.2 and Elasticsearch.

The starting environment will be a simple Content Manager 9.1 and IDOL implementation. We assume you already know how to install Content Manager 9.1 as outlined in Installation Guide (**CM9.1\_Install.pdf**), for example:



Whether you're upgrading or doing a new installation the resulting Content Manager 9.2 and Elasticsearch environment, in this example scenario, will include three additional "node" servers in an Elasticsearch cluster. This is inline with the Elastic team's best practice guidelines regarding the minimum number of nodes you should have in a production cluster:



## Install Content Manager 9.2 or later

### Upgrading from earlier versions

If you are running an earlier version of Content Manager we recommend upgrading to a Content Manager 9.2 (or later) and IDOL environment before attempting to install and configure Elasticsearch. This will allow you to confirm your environment and content searches are working as expected prior to reconfiguring your environment to use Elasticsearch integration instead of IDOL.

Follow the instructions outlined in the Installation and Setup Guide (**CM9.3\_Install.pdf**) to upgrade your environment to Content Manager 9.2 or later.

### New Installations

If you're installing Content Manager 9.2 or later for the first time you should do so as outlined in the Installation and Setup Guide (**CM9.3\_Install.pdf**) before you implement Elasticsearch.

## Install Elasticsearch

When you have upgraded to or installed Content Manager 9.2 you can proceed to installing Elasticsearch on each of the servers that will make up the cluster.

We do not recommend installing Elasticsearch on a Workgroup Server as it can use ports that are used by Content Manager and adds unnecessary complexity when you're trying to troubleshoot a problem in the future. The Elastic team are very clear on keeping it simple. We suggest you follow this approach when designing and setting up your environment.

We do recommend you dedicate a separate server for each Elasticsearch node as we've done in the above network diagram.

The Elastic team recommend a minimum of three nodes for production environments so that if any single node fails no data will be lost.

## Download and install the Java Runtime Environment (JRE)

The JRE is the only prerequisite needed for Elasticsearch. If you don't already have it installed you can download the latest version from Oracle's Java Platform page.

You should install the JRE on every server you intend to install Elasticsearch on.

When you install the JRE ensure you run the installation package as an administrator e.g. right mouse click -> Run as administrator

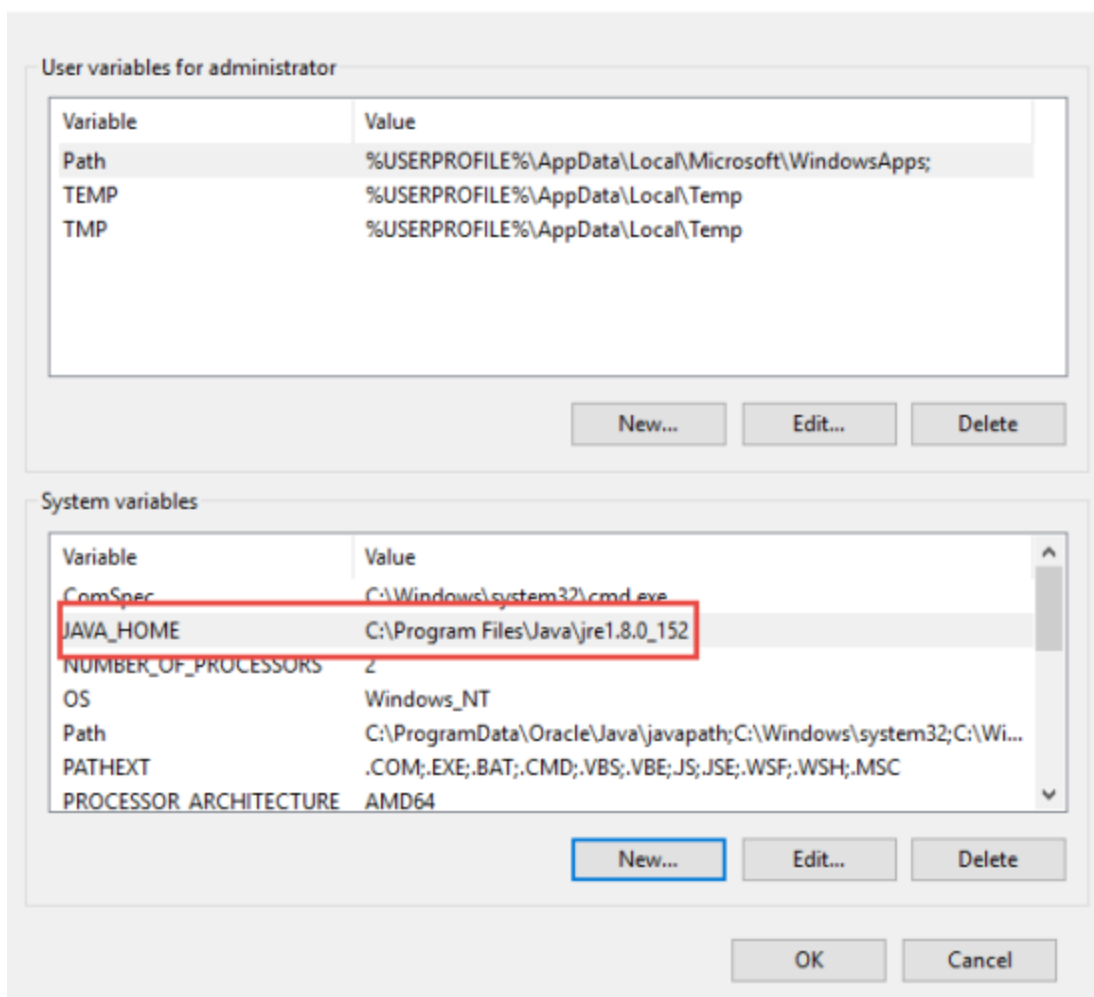
## Ensure the JAVA\_HOME environment variable is configured

When you've installed the JRE, make sure the JAVA\_HOME environment variable is configured. To edit environment variables, right mouse click over the Start Button and then select:

**System -> Advanced system settings -> Advanced tab -> Environment variables**

Add an environment variable JAVA\_HOME and the path to the JRE installation directory (e.g. C:\Program Files\Java\jre-1.8.0\_152)





## Install Elasticsearch on your nominated servers

Download Elasticsearch 6.2.x from [elastic.co](http://elastic.co) and unzip it your preferred directory.

In this example, we downloaded the zip package and unzipped to C:\Elasticsearch\ on Elastic1, Elastic2, and Elastic3.

## Edit the Elasticsearch configuration file

Elasticsearch is configured via a configuration file located in the directory you've just unzipped it to. In this scenario it's located at C:\Elasticsearch\elasticsearch-6.2.x\config\elasticsearch.yml

You can open, edit, and save this file using text editor like TextPad or Notepad.

The most common settings are already in this file and you'll find they're commented out via the # symbol at the start of each line. The following table lists the ones that are most commonly used:

**node.name** - Name of the server node. This name is just for reference. If you leave this blank Elasticsearch will choose a random default name for you.

**cluster.name** - For multi node systems, each node needs to have the same cluster name. All nodes with the same cluster will use a network discovery protocol to find the other nodes in the cluster, and automatically configure themselves. The default name is *elasticsearch*, but you should change it to an appropriate name which describes the purpose of the cluster. A cluster does not become active until this is set.

**path.data** - Path where Elasticsearch stores the index data. The default is the 'data' directory in the Elasticsearch folder.

**path.logs** - Path where Elasticsearch stores the log files. The default is the 'log' directory in the Elasticsearch folder.

**network.host** - The node will bind to this hostname or IP address and publish (advertise) this host to other nodes in the cluster. Accepts an IP address, hostname, a special value, or an array of any combination of these.

**discovery.zen.ping.unicast.hosts** - In order to join a cluster, a node needs to know the hostname or IP address of at least some of the other nodes in the cluster. This setting provides the initial list of other nodes that this node will try to contact. Accepts IP addresses or hostnames.

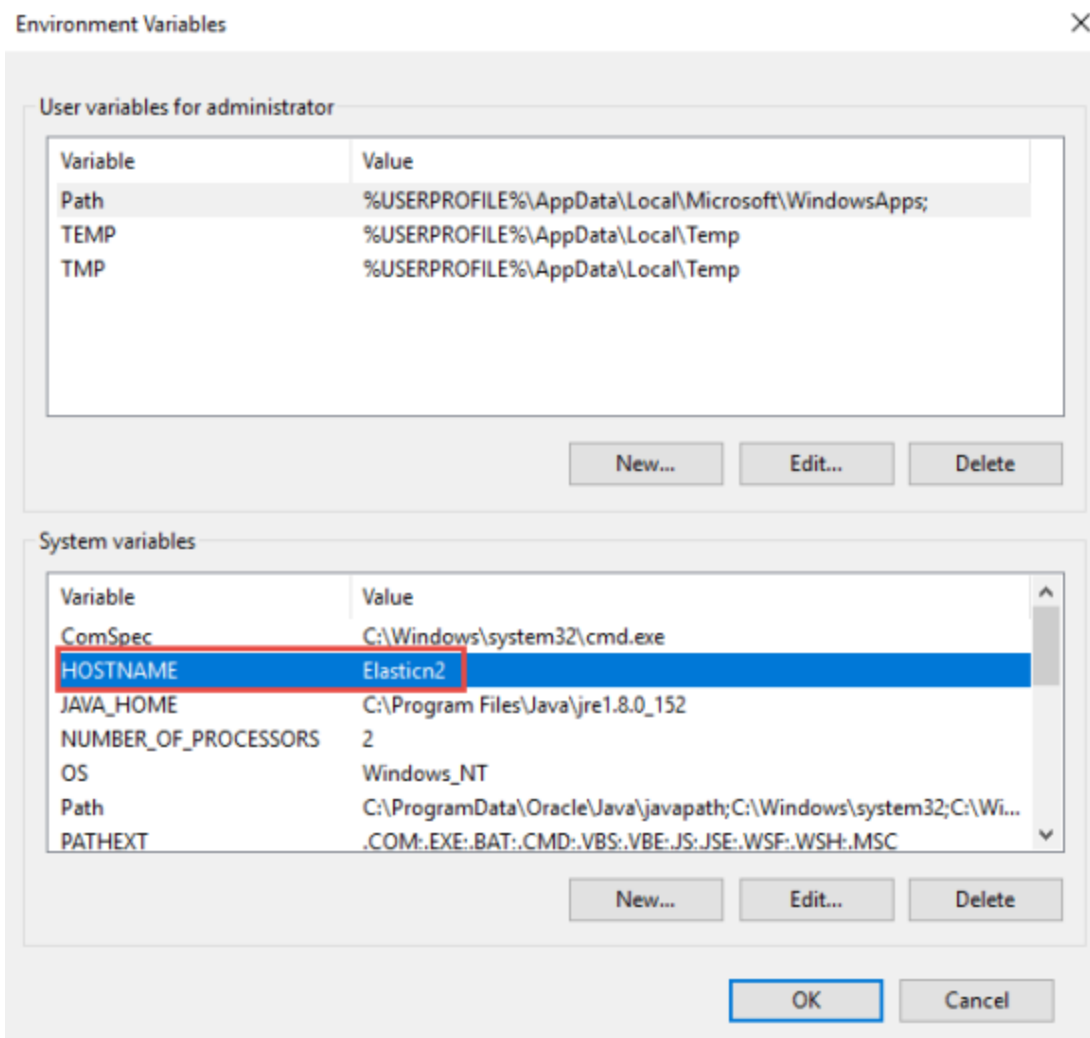
To setup the three-node cluster in our environment we'll be changing these settings in the configuration file on each server:

**node.name** - We'll set this to be the server's hostname by removing the # and using the entry `${HOSTNAME}`

```
# ----- Node --
#
# Use a descriptive name for the node:
#
node.name: ${HOSTNAME}
#
# Add custom attributes to the node:
#
#node.attr.rack: r1
#
```

To use this entry, you'll also have to enter the HOSTNAME environment variable like you did for JAVA\_HOME.

This will be Elasticn1, Elasticn2, and Elasticn3 on each of the node servers.



**cluster.name** - We'll set this by removing the # and using the entry contentmanager

```
# ----- Cluster -----  
#  
# Use a descriptive name for your cluster:  
#  
cluster.name: contentmanager
```

**network.host** - We'll set this by removing the # and using the entry `_site_` which will set the value to any site-local addresses on the system, for example 192.168.1.4

```
# ----- Network -----  
#  
# Set the bind address to a specific IP (IPv4 or IPv6):  
#  
network.host: _site_  
#  
# Set a custom port for HTTP:  
#  
#http.port: 9200  
#  
# For more information, consult the network module documentation.  
#
```

**discovery.zen.ping.unicast.hosts** - We'll set this by removing the # and replacing the host entries with the IP addresses of the three node servers that will make up the cluster (192.168.1.4, 192.168.1.5, and 192.168.1.6). We could have also used hostnames instead of IP addresses (Elasticn1, Elasticn2, and Elasticn3).

```
# ----- Discovery -----  
#  
# Pass an initial list of hosts to perform discovery when new node is started:  
# The default list of hosts is ["127.0.0.1", ">:::1"]  
#  
discovery.zen.ping.unicast.hosts: ["192.168.1.4", "192.168.1.5", "192.168.1.6"]  
#
```

**TIP:** If you change these settings after your environment is operational you'll have to stop and start the Elasticsearch service to implement the changes.

## Memory Management and configuring the Java Virtual Machine (JVM)

Getting your memory allocation configured correctly is critical to the successful operation of Elasticsearch. You must size your hardware and configure your environment appropriately to handle the indexing load your organization will be generating.

There are three memory variables you have to consider when setting up Elasticsearch:

- **Installed RAM**

The Elastic team advise a server that has 64 GB of RAM is the ideal amount in a production environment. The Elastic team do not recommend having less than 8 GB of RAM as it can lead to a need for many small servers.

- **Heap memory**

This is the memory allocated to the JVM in which Elasticsearch runs. The Elastic team recommend giving 50% of the available memory to Elasticsearch heap, while leaving the other 50% free.

When setting the heap memory, it is also recommended you do not allocate more than 26 GB. This is due to a Java limitation that becomes apparent at 32 GB and can have a negative impact on performance. This limitation is described in detail in the **Heap: Sizing and Swapping** section of **Elasticsearch: The Definitive Guide** on the [elastic.co](https://www.elastic.co) web page.

• **Off heap memory**

This is the portion of memory used by the operating system and Lucene (the underlying engine used by Elasticsearch) outside of the JVM heap.

When configuring the JVM heap you should ensure there's enough memory left to accommodate off heap processes. In line with the above recommendation, this should be around 50% of the available memory.

Elasticsearch provides a way to manage JVM settings via the `jvm.options` file which has the default heap size set to 2 GB. You can open, edit, and save this file using text editor like TextPad or Notepad. You'll find this file in:

`<Drive>:\<Elasticsearch installation directory>\config\`

In our test environment, we'll be setting the initial and maximum heap size to 1g (1 GB).

```
#####
## IMPORTANT: JVM heap size
#####
##
## You should always set the min and max JVM heap
## size to the same value. For example, to set
## the heap to 4 GB, set:
##
## -Xms4g
## -Xmx4g
##
## See https://www.elastic.co/guide/en/elasticsearch/reference/current/heap-size.html
## for more information
##
#####

# Xms represents the initial size of total heap space
# Xmx represents the maximum size of total heap space

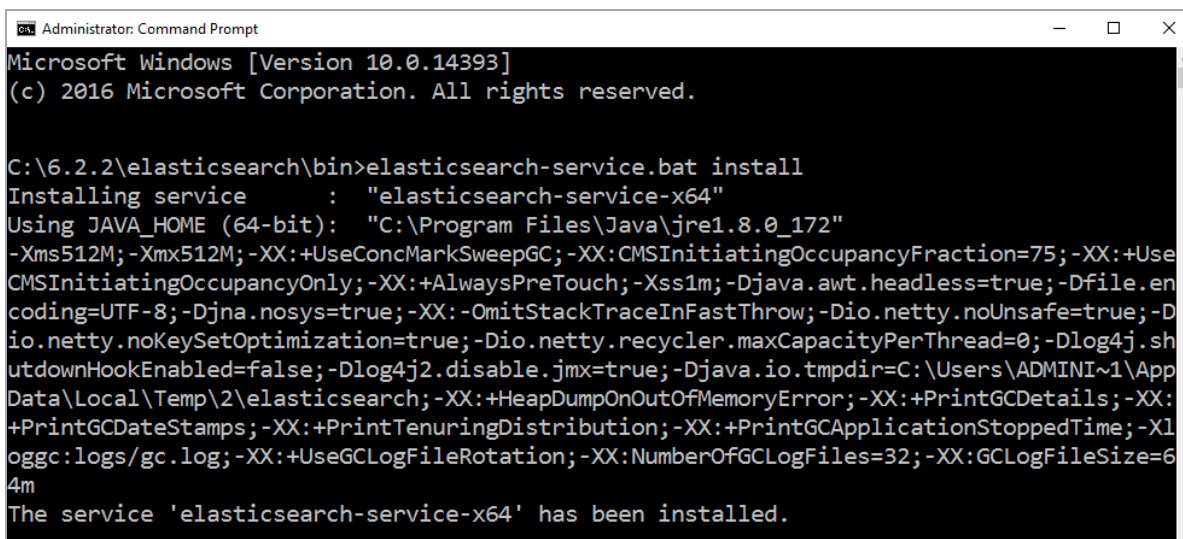
-Xms1g
-Xmx1g
```

## Configure Elasticsearch to run as a service

You can use the supplied batch script `<Drive>:\<Elasticsearch installation directory>\bin\elasticsearch-service.bat` to install Elasticsearch as a Windows service.

To do this, open an administrative command prompt, navigate to the bin directory, and run the following command on each Elasticsearch server:

```
elasticsearch-service.bat install
```



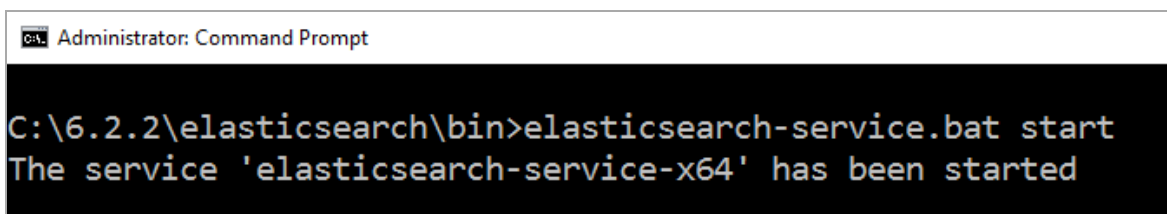
```
Administrator: Command Prompt
Microsoft Windows [Version 10.0.14393]
(c) 2016 Microsoft Corporation. All rights reserved.

C:\6.2.2\elasticsearch\bin>elasticsearch-service.bat install
Installing service      : "elasticsearch-service-x64"
Using JAVA_HOME (64-bit): "C:\Program Files\Java\jre1.8.0_172"
-Xms512M;-Xmx512M;-XX:+UseConcMarkSweepGC;-XX:CMSInitiatingOccupancyFraction=75;-XX:+Use
CMSInitiatingOccupancyOnly;-XX:+AlwaysPreTouch;-Xss1m;-Djava.awt.headless=true;-Dfile.en
coding=UTF-8;-Djna.nosys=true;-XX:-OmitStackTraceInFastThrow;-Dio.netty.noUnsafe=true;-D
io.netty.noKeySetOptimization=true;-Dio.netty.recycler.maxCapacityPerThread=0;-Dlog4j.sh
utdownHookEnabled=false;-Dlog4j2.disable.jmx=true;-Djava.io.tmpdir=C:\Users\ADMINI~1\AppData\Local\Temp\2\elasticsearch;-XX:+HeapDumpOnOutOfMemoryError;-XX:+PrintGCDetails;-XX:
+PrintGCDateStamps;-XX:+PrintTenuringDistribution;-XX:+PrintGCApplicationStoppedTime;-Xl
oggc:logs/gc.log;-XX:+UseGCLogFileRotation;-XX:NumberOfGCLogFiles=32;-XX:GCLogFileSize=6
4m
The service 'elasticsearch-service-x64' has been installed.
```

## Run Elasticsearch

Start the Elasticsearch service on each server via an administrative command prompt using the following command:

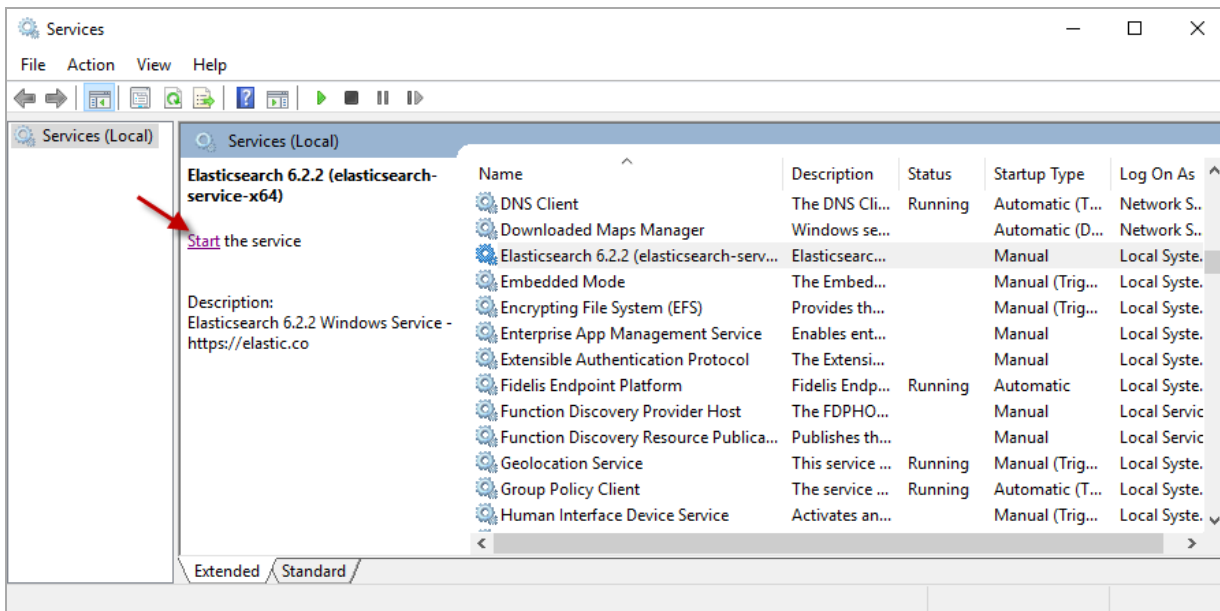
```
elasticsearch-service.bat start
```



```
Administrator: Command Prompt

C:\6.2.2\elasticsearch\bin>elasticsearch-service.bat start
The service 'elasticsearch-service-x64' has been started
```

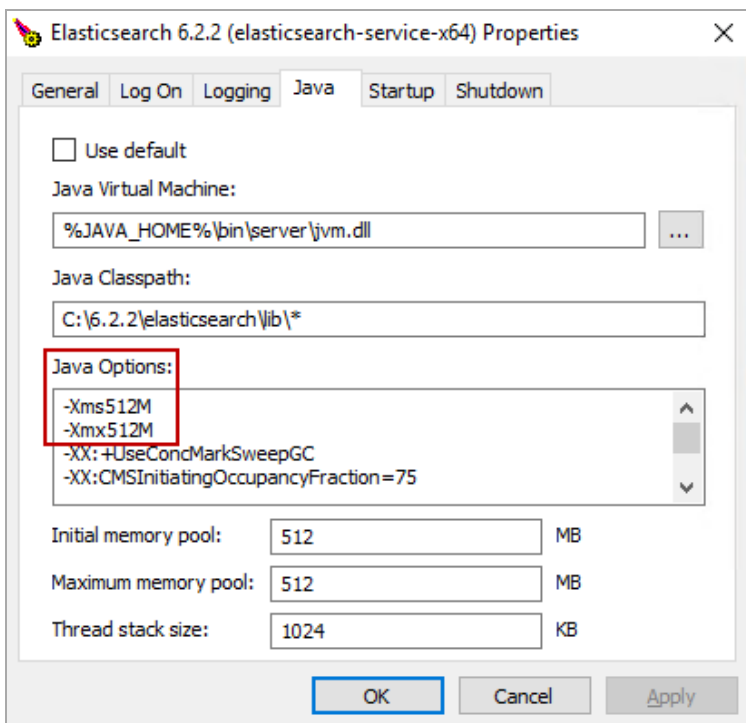
You can also start the Elasticsearch service via the Windows Services Management Console (Start Button->Run->services.msc)



## Adjusting the JVM heap size after installing Elasticsearch as a service

To adjust the JVM heap size after installing Elasticsearch as a service you'll need to use the service manager GUI located in the ...\\bin directory. Editing `jvm.options` will not pass on heap changes to the service once it's installed.

You can open the service manager window by invoking `elasticsearch-service.bat` manager from the command line:



Any changes you make via the service manager will require you to re-start the service which can be done from the **General** tab.

## Confirm the Elasticsearch service and cluster are operational

Check the health of your cluster by opening `http://<ip address or hostname>:9200/_cat/health?v` in a browser. The IP address or hostname used can be any Elasticsearch server in your environment.

Elasticsearch uses a traffic-light scheme for server health: green, yellow if there are warnings (e.g. disk space more than 85% full), and red for serious errors.

The Elasticsearch service can take some time to start if you've had a reason to restart it e.g. after changing configuration values. During this period, you may see the health status as **red**. If this occurs, you should wait until the service is finished its start-up sequence and try the health check again.





## Configure an Elasticsearch Content Index

Content Manager 9.2 or later will only allow you to have one type of content indexing engine configured. If you upgrade from earlier versions your existing IDOL content index will continue to be available in 9.2 or later.

However, if you decide to use Elasticsearch as your content index your IDOL content index will no longer be available.

**NOTE:** If you decide to go back to an IDOL content index, you will need to delete your Elasticsearch content index and create a new IDOL index. See **Setting up document content indexing and searching** in the Content Manager Enterprise Studio help file.


Now we have a working Elasticsearch environment in place we can move on to enabling it within Content Manager 9.2 or later.

## Reindex your document store with Elasticsearch

If you're setting up a new instance of Content Manager there's no need to perform this step yet. You can go straight to [Create a new Elasticsearch index, on page 23](#) and follow the steps to setup a new Content Index in Enterprise Studio.

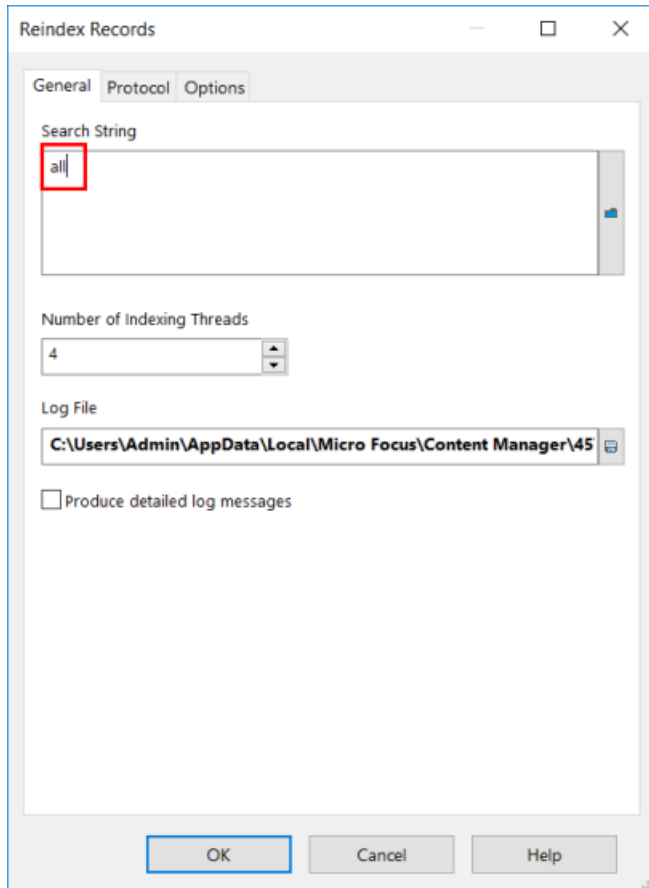
Open the Content Manager client and select the **Administration** tab.

From the **Administration** tab, in the **Elasticsearch Index** group, click **Records**.



The **Reindex Records** dialog will appear.

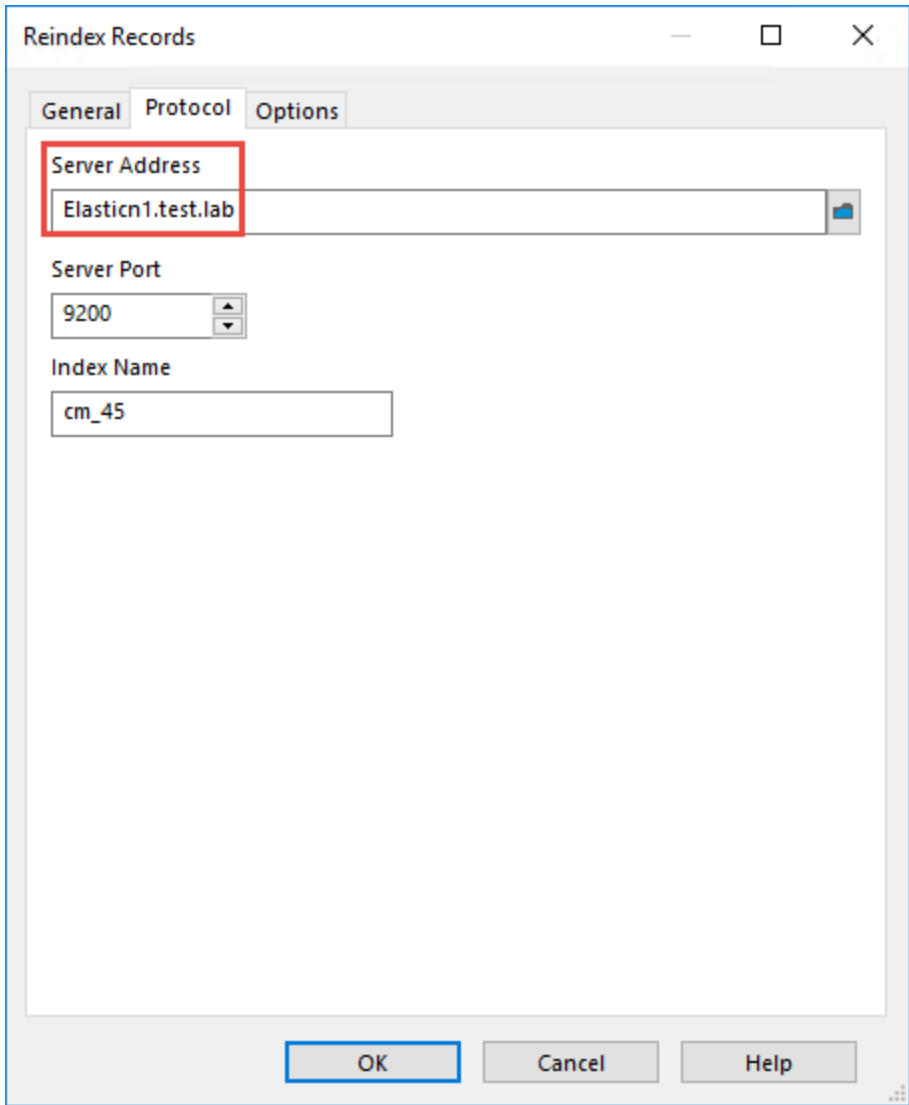
In the **Search String** field on the **General** tab, type all.



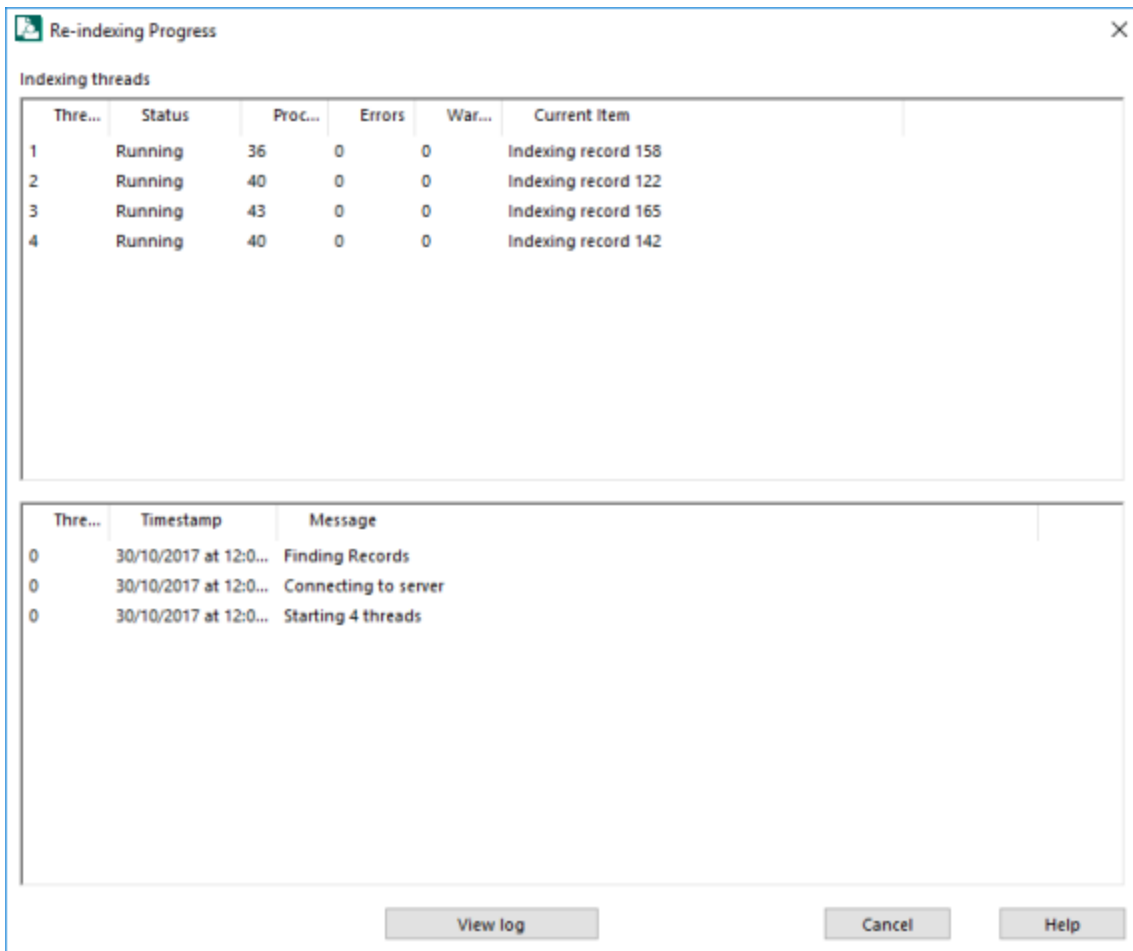
**IMPORTANT:** To successfully complete a full re-index, make sure you are using a Content Manager location profile that has sufficient access to bypass all current security settings you may have in place.

You also have the opportunity to potentially improve performance by increasing the **Number of Indexing Threads**. This is largely dependent on your hardware. You'll have to trial a number of different thread settings to find out exactly how many threads it's capable of handling. For this scenario we'll use the default value of 4.

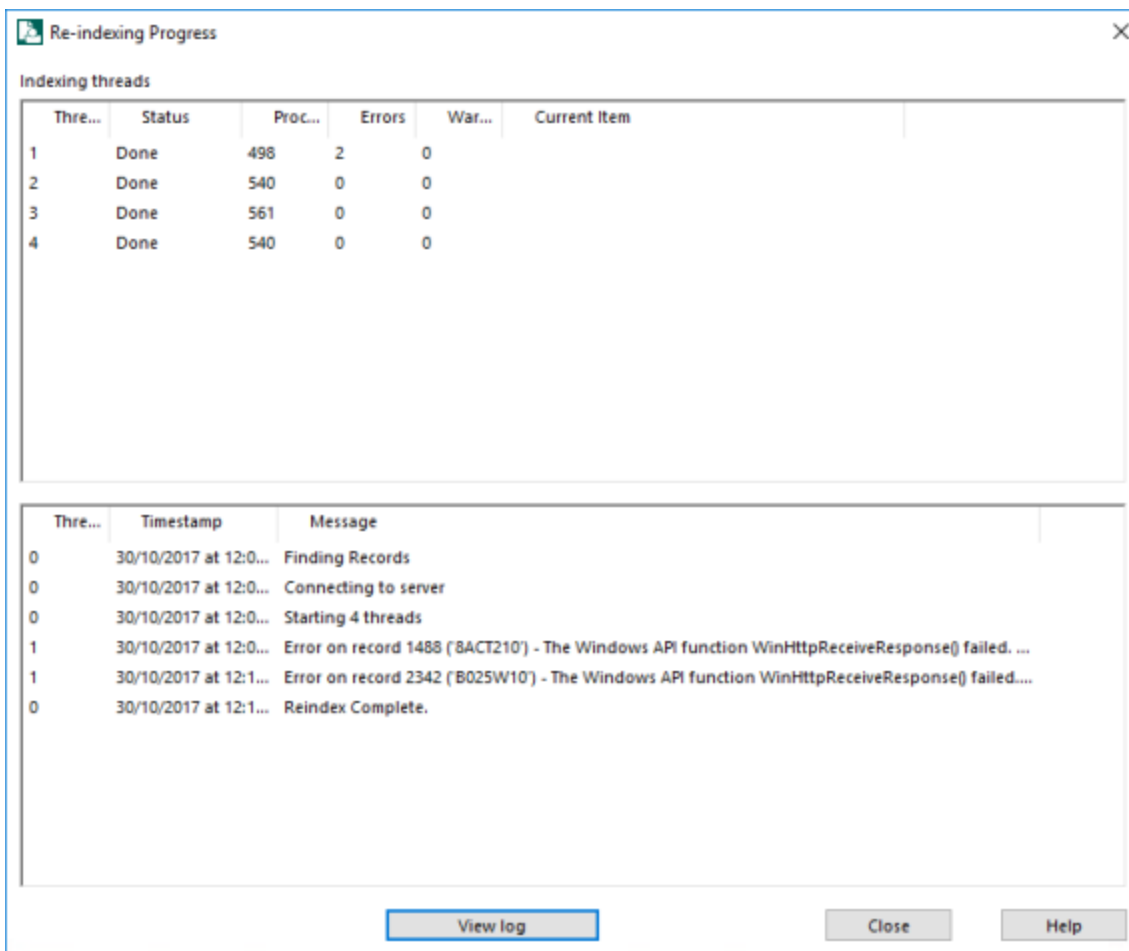
In the **Server Address** field on the **Protocol** tab browse to, or type, the Elasticsearch server name or its IP address. This can be any Elasticsearch server in the cluster.



Click **OK** and the re-indexing process will start.



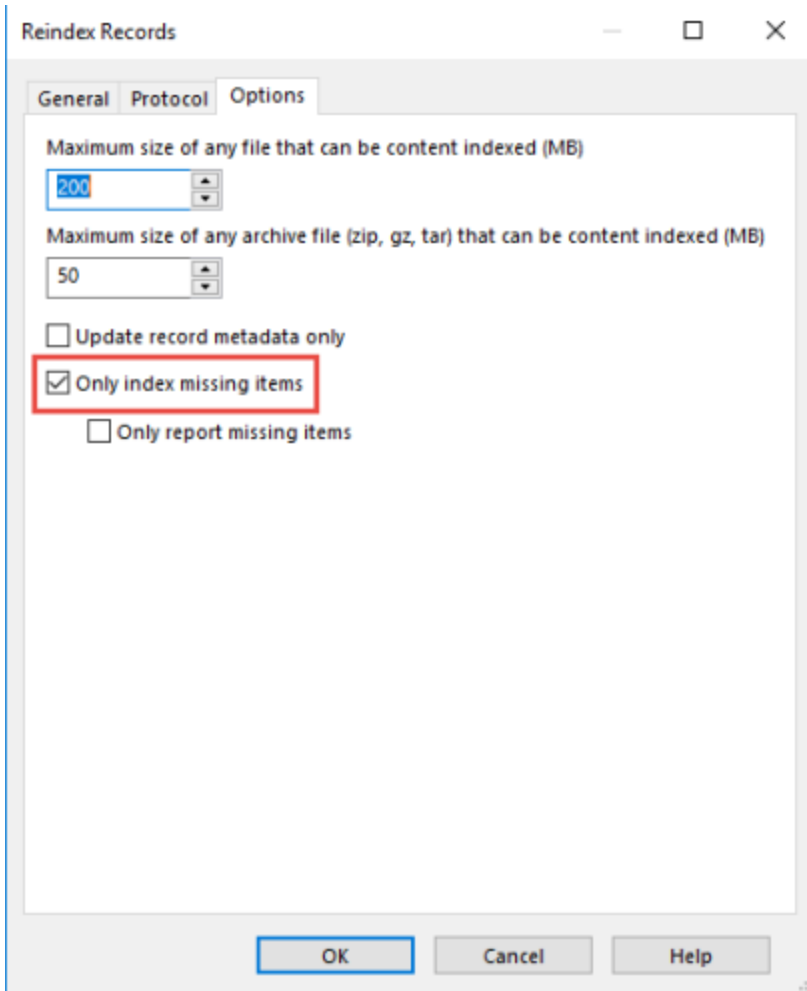
When the re-indexing process is complete a **Reindex Complete** entry is displayed along with any errors it encountered.



In this example, the process reported two errors. Click **View Log** to get additional information.

The log entries show records 1488 and 2342 have failed because “**The operation timed out (0x00002ee2)**”. Timeout errors usually occur when Elasticsearch is under load and cannot process the documents sent to it. Given this is a small test environment with limited resources this is the most likely cause.

To rectify this, we can re-run the re-index but this time only index those records that are missing by selecting the **Only index missing items** on the **Options** tab.

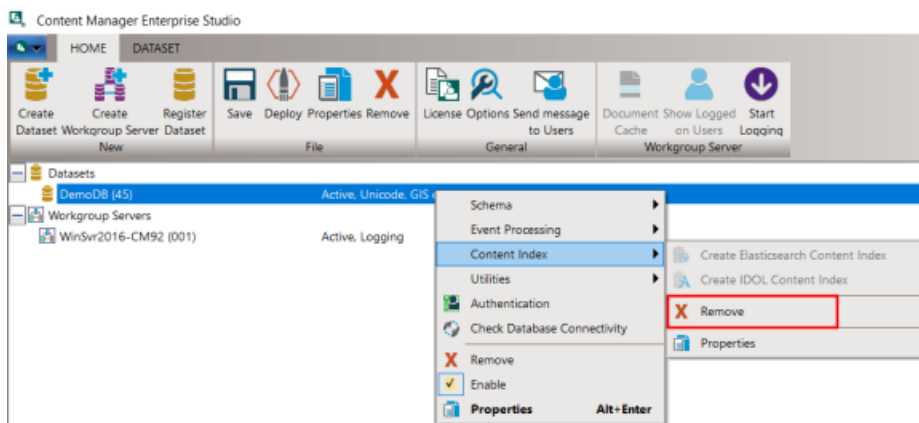


## Configure Content Manager Enterprise Studio to use Elasticsearch

### Remove existing IDOL content indexes

Now we've successfully created an Elasticsearch content index the existing IDOL index can be removed.

On your Workgroup Server open Enterprise Studio -> right mouse click on your Dataset -> **Content Index** -> **Remove**.



Select **Delete Content Index from Server** then select **OK**.

## Create a new Elasticsearch index

On your Workgroup Server open Enterprise Studio -> right mouse click on your Dataset -> **Content Index** -> **Create Elasticsearch Content Index**.

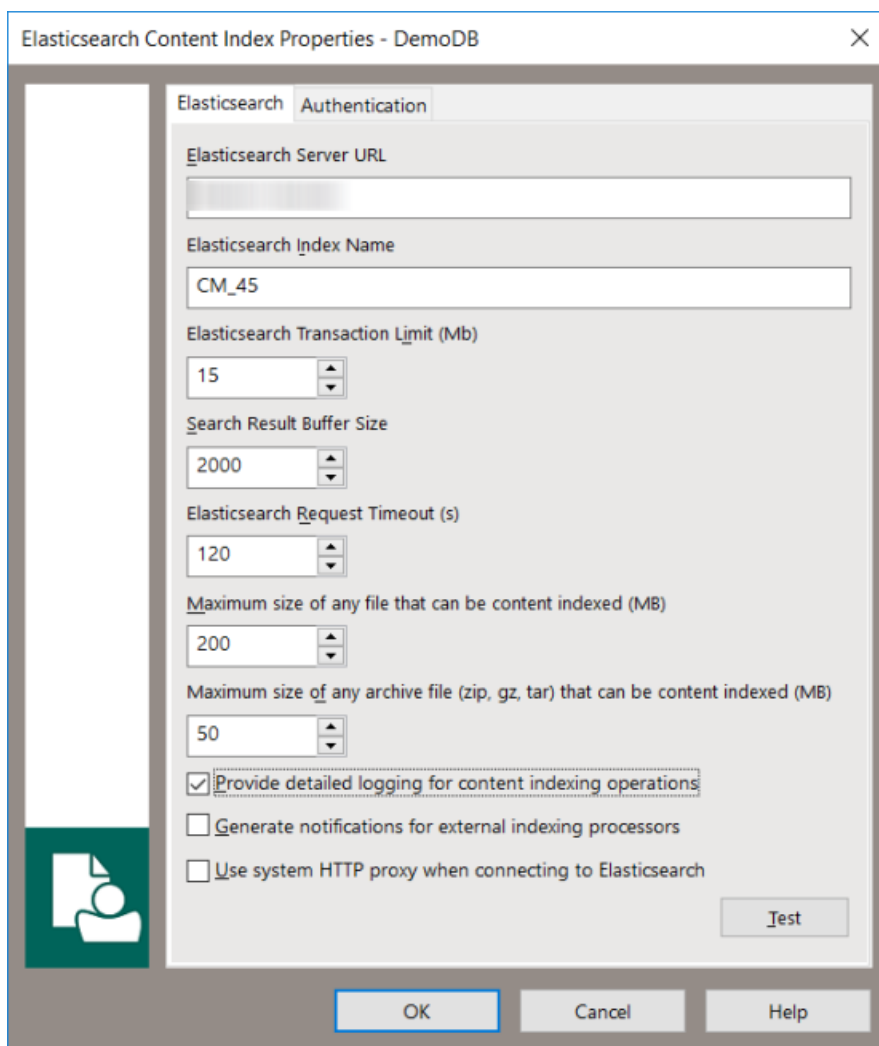
Fill in the fields on the Content index properties dialog:

On the **Elasticsearch** tab:

- **Elasticsearch Server URL** - enter the Elasticsearch Server URL.
- **Elasticsearch Index Name** - Name of the Elasticsearch index to use for this CM dataset. The characters in the name must be alphanumeric characters or underscores.
- **Elasticsearch Transaction Limit (MB)** - Limits how much indexing data can be sent to Elasticsearch in a single request.
- **Elasticsearch Request Timeout (s)** - Limits how long CM will wait for Elasticsearch to process a single request.
- **Maximum size of any file that can be content indexed (MB)** - default: 200. Content Manager does not index the content of documents whose file size is greater than this number in megabytes.
- **Maximum size of any archive file (zip, gz, tar) that can be content indexed (MB)** - default: 50. Content Manager does not index the content of archive files whose file size is greater than this number in megabytes.
- **Provide detailed logging for content indexing operations** - select this option to generate detailed log files.
- **Generate notifications for external indexing processors** - if you are configured to use Elasticsearch but you wish for a CFS connector to be notified of indexing events, then enable this

option. This is for installations that use ControlPoint.

- **Use system HTTP proxy when connecting to Elasticsearch -**



On the **Authentication** tab, fill in the fields:

**Enable Elasticsearch X-Pack authentication** - select this option to enable X-Pack authentication. You can either specify a user name and password or a certificate.

- **User Name** - enter the user name to connect to the Elasticsearch server.
- **Password** - enter the password for the defined user.
- **Client Certificate** - if using a certificate for the X-Pack authentication, the certificate must be installed to the Personal store of the Local Computer account. Once installed, the certificate will appear in the **Client Certificate** drop-down list. Select the required certificate from the drop-down



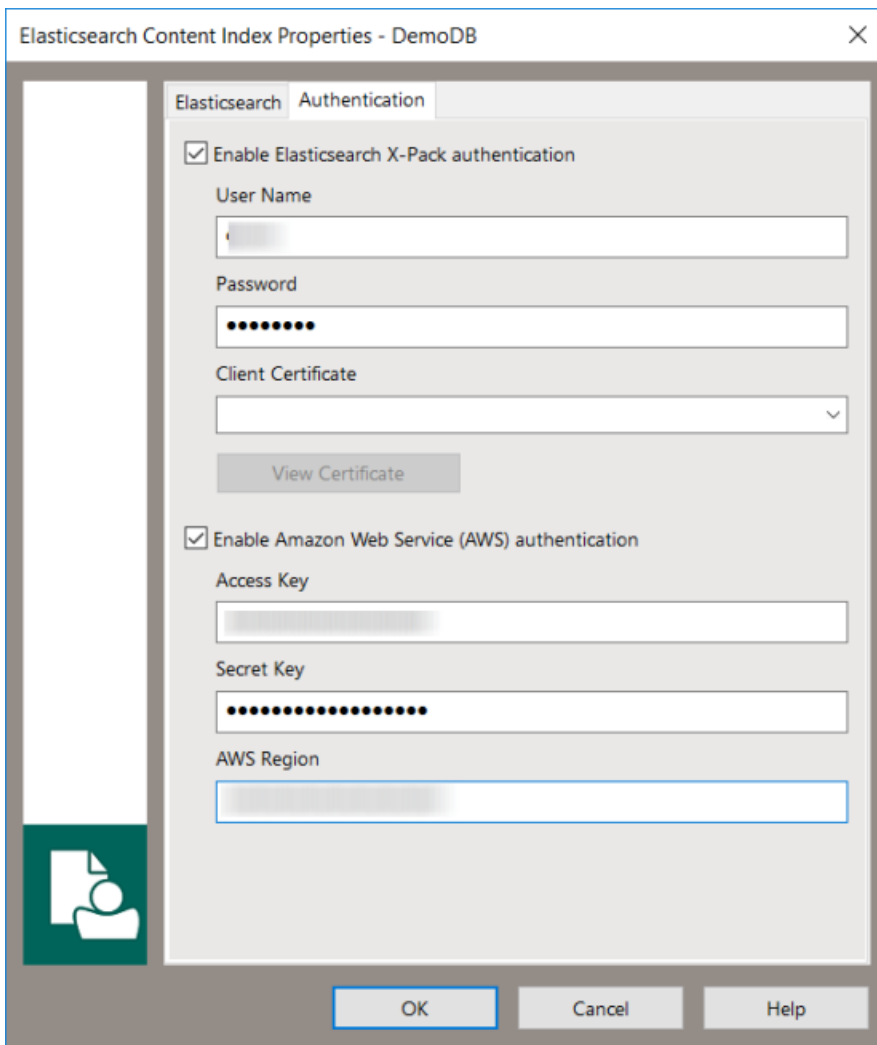
list.

- **View Certificate** - click to view the selected Client Certificate.

**TIP:** If you get certification validation errors, ensure that the local computer trusts the certificate authority of the certificate that the Elasticsearch server is presenting.

**Enable Amazon Web Service (AWS) authentication** - select this option to use the Amazon Web Service (AWS) version of the Elasticsearch service.

- **Access Key** - enter the AWS Access Key. This is the username for the AWS IAM user that has access to the Elasticsearch services inside AWS.
- **Secret Key** - enter the AWS Secret Key. This is the password for the Access Key user.
- **AWS Region** - enter the AWS Region name.




If you've setup your environment correctly the **Test** button should return a **Succeeded** response.

Click **OK** to close the **Content index properties** dialog.

If you encounter the following prompt: **The index 'CM\_45' already exists on the Elasticsearch server. Do you want to delete it and create and new one?** Select **No**.

**Save** and then **Deploy** your new settings in Enterprise Studio.



## Confirm your Elasticsearch content index is operational

Performing a simple document content search from a Content Manager client is the best way to confirm your content index is operational.

In this case, we selected a document with unique content so only one result was returned. The use of unique content confirms the searching mechanism can interrogate the entire index and return the correct result.

Search for		Search by		Matching criteria	
Records	Document Content				Polyfluoroalkyl
Record Type	Record Number	Title		Date Created	
Document	G17/19	An Overview of Perfluoroalkyl and Polyfluoroalkyl Substances and Interim ...		9/10/2017 at 9:15 AM	

## Troubleshooting common issues

### Logging

When it comes to troubleshooting general content indexing issues there are three key logs you should be looking at:

Log file name	Default Log file location	Useful for...
Elasticsearch cluster log	ES_HOME\logs\[cluster_name].log	Identifying issues to do with the health of the Elasticsearch environment. Located on each Elasticsearch server in your cluster.
Workgroup server document content indexing event log	<Drive>\Micro Focus Content Manager\ServerLocalData\TRIM\Log\ DCI_<DB ID>_<YYYY>_<MM>_<DD>.log	Whenever an electronic record is created, a new version is committed or an electronic document is destroyed, deleted or removed, an event is created to indicate that the document content index should be updated. This log file may help you narrow down where a particular problem is occurring when compared to other logs listed in this table.
Client Reindex log	<Drive>\Users\<username>\AppData\Local\Micro Focus\Content Manager\<DB ID>\Log\ElasticReindex_<YYYY>_<MM>_<DD>.log	Identifying problematic documents or timeout issues when reindexing from the Content Manager client.

### Diagnosing “operation timed out” errors

The most common issue you’ll probably come across when indexing content is **operation timed out** errors in the client **Reindex** log.

This error doesn’t always mean there’s an issue. Sometimes, it can just mean the client has waited for a response from the Elasticsearch server beyond the **Request Timeout** value you’ve set in Enterprise

Studio. Often, you'll find Elasticsearch has continued to index the content after the client has reported a timeout event. You can check if the content has been indexed by doing a search based on content. If it doesn't return a result you should then attempt to run a re-index on only missing items.

If you continue to observe timeout errors and content isn't being indexed you can then proceed to examining the logs identified in the previous section. Key areas you should be looking at in this scenario are:

- **Enterprise Studio – Transaction Limit**

Content Manager sends documents to Elasticsearch using a bulk-index API, this value sets the maximum size of the memory buffer before we send the bulk import request to Elasticsearch. When performing re-indexing, using a larger value will speed up indexing. However, you don't want it to be too large as it can result in Elasticsearch taking longer to return a result and/or the JVM running out of heap memory if resources are limited. We recommend you keep the transaction limit set to 5-15 MB (the default transaction limit value is 9 MB).

- **Enterprise Studio – Request Timeout**

The request timeout value sets the time Content Manager will wait for Elasticsearch to respond every time a HTTP request is sent. Every environment will perform differently depending on the resources it has available. If there's adequate resources simply extending the timeout value may resolve many of the timeout errors observed (the default timeout value is 60 seconds).

- **Elasticsearch – available heap and available nodes**

Timeout errors may also indicate there's isn't enough memory allocated to the JVM, your servers aren't large enough to handle the load, or you don't have enough nodes in your cluster. If the cluster log indicates this you should revisit the memory configuration and sizing recommendations outlined in [Memory Management and configuring the Java Virtual Machine \(JVM\)](#).

## Proxy servers and firewalls

Content Manager Clients and Workgroup Servers need direct access to your Elasticsearch servers. Make sure your firewall has the appropriate ports open (default port range is 9300-9400) and your proxy server is configured to not cache traffic bound for Elasticsearch hosts.

## Useful Information

### Kibana

The Elastic team recommend that as your Elasticsearch environment grows to include more nodes and clusters you should consider using Kibana to manage it. Kibana not only allows you to manage your clusters more efficiently it also gives you the ability to analyze and make sense of your data. Visit [elastic.co](https://elastic.co) to get more information on Kibana.

### cURL

The Elastic team recommend cURL, which is an open source command line utility that can be used to directly query and manage Elasticsearch via its API. If you don't decide to use Kibana you can use cURL commands to perform a large range of tasks e.g. delete an index, get node and cluster health data. The Elasticsearch documentation contains many examples of cURL scripts.

cURL can be downloaded, free of charge, from the official cURL site.

### Elasticsearch Reference

The Elasticsearch Reference is a detailed knowledge base of all aspects of Elasticsearch. You should consult this when considering the use of Elasticsearch in your organization.

You can find the Elasticsearch Reference in the Docs section on [elastic.co](https://elastic.co)

### Hardware requirements

The Content Manager Support team regularly gets asked for assistance to size a customer's environment for content indexing. The Elastic team do a good job of explaining how to approach this problem in the **Hardware** section of **Elasticsearch: The Definitive Guide** on [elastic.co](https://elastic.co)

Microsoft also present an excellent overview of Elasticsearch architecture and how to scale it on Azure which can be applied to any environment. You can find this information in the **Elasticsearch on Azure** section of the Azure online documentation.