Secure Elasticsearch Deployment For Sage X3

Raheel Khan - 4th December 2024

Contents

Secure Elasticsearch Deployment for Sage X3

What is Elasticsearch

Elasticsearch is an opensource search and analytics engine that helps you quickly find and analyse large amounts of data. Originally built to handle complex search tasks, it's now widely used for everything from powering website searches to monitoring real-time data. Its ability to index and search data fast makes it invaluable for applications that need instant insights, while its open-source nature makes it versatile and accessible for a wide range of uses.

- Elasticsearch is a powerful search and analytics engine designed for quick data retrieval.
- Organizes data using "indexing," making searches almost instant, regardless of data size.
- Commonly used in applications that require fast searching, such as:
  - Websites and apps
  - Log and event data analysis
  - Real-time business data insights
- Built on open-source technology, making it accessible and adaptable to various needs
- Java full-text search library Lucene

Elasticsearch Components

In Elasticsearch, data storage and search efficiency rely on four core components in the Elasticsearch deployment

1. **Index:** Organizes related documents or data
2. **Shard**: Breaks down the index into smaller parts
3. **Node:** Hosts shards and processes search requests
4. **Cluster**: Which can combine multiple nodes to work as a single system.

In an Elasticsearch system, data is organized within **indexes**, which are like databases for specific sets of information. Each **index** is divided into smaller pieces called **shards** to allow efficient storage and faster search operations. These shards are distributed across multiple **nodes**, which are individual servers in the system. Together, the nodes form a **cluster**, which is a group of servers working as a single, unified search and analysis engine. This architecture enables Elasticsearch to scale horizontally, handle large data volumes, and provide high availability.

Set up a cluster for high availability | Elasticsearch Guide

Sage X3 Integration

By connecting Sage X3 to an Elasticsearch instance, businesses can index large volumes of transactional, financial, and operational data from Sage X3. This integration allows users to perform fast, full-text searches across Sage X3 data, enabling quick retrieval of information.

- Elasticsearch enables the search bar at the top of the dashboard once you have logged into Sage X3
- This gives the user the capability for rapid searches and results of data specific to their business function
- The associated X3 screens can then be accessed directly from the search output results
- The Sage X3 development teams will define what data is indexed out of the box
- Customers can modify the fields that can be indexed and therefore searched, in order to satisfy their own business requirements (Set the class to be ‘searchable’, set the field to be searchable in the class the associated representation also has to be set to be ‘used in search results’)

Elasticsearch Architecture

Elasticsearch can be installed alongside any of the Sage X3 technical components. The Sage X3 Syracuse server defines the connection and index settings to the Elasticsearch application as well as retrieving the search results. Due to both these applications being resource-intensive the recommendation in the X3 Architecture diagram recommends Elasticsearch installed on a separate server.

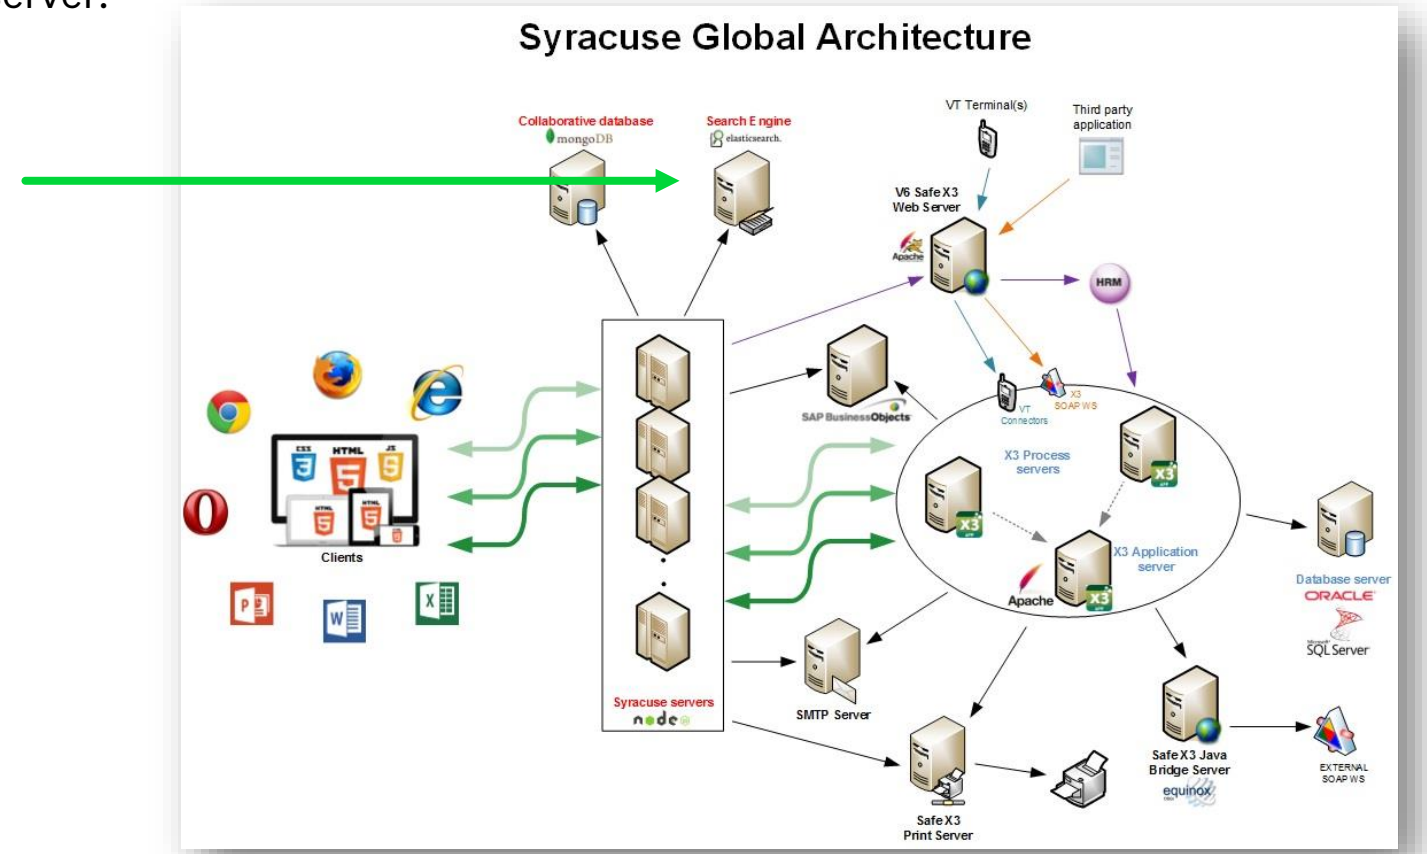

Architecture schema

Elasticsearch Sage X3 Support

From Sage X3 2023 R2, Sage X3 now supports Elasticsearch 8 and the Elastic cloud offering with a new dedicated page in Sage X3 to set up Elasticsearch security and connection security. This is accessible in Sage X3 via **Administration > Usage > Search > Search server settings**.

- Prior to 2023 R2 The Elasticsearch connection was configured in the nodelocal.js with the limited options it was not possible to use Elasticsearch 8 or the Elastic Cloud offering. Elasticsearch 8 introduces
- **Connection security:** Data stream encryption using certificates
- **Elastic search credentials**: User and password protection for the Elasticsearch database
- **Certificate authentication:** adds an extra layer of security

Component	Product version	Component version
Elasticsearch	From Syracuse 12.3 before release 2020 R1/V12.0.21	6.4
	From Syracuse 12.6 in release 2020 R1/V12.0.21 to release 2020 R3/V12.0.24	6.8
	From Syracuse 12.9 in release 2020 R4/V12.0.24 to release 2023 R1/12.0.33	7.16
	From Syracuse 12.19 in release 2023 R2/12.0.34	Latest ElasticSearch 8 Version

Sage X3 Prerequisites Overview

Installing Elasticsearch 8+

The Install Steps Overview

The install steps that we will cover

1. Prerequisites are fulfilled
2. Downloading Elasticsearch
3. Decide on a Default installation or not
4. Installing Elasticsearch 8
5. Security Setup
6. Testing connection

Elasticsearch Prerequisites

To install Elasticsearch on a Windows Server, you'll need to ensure the following prerequisites are met:

- **Java Runtime Environment (JRE):**
  - Elasticsearch requires Java 8 or higher. It's recommended to use the bundled OpenJDK that comes with ElasticsearchThis bundled JDK is optimised and tested specifically for Elasticsearch, ensuring compatibility and performance

- **Sufficient Memory and Disk Space:**
  - Ensure your server has enough RAM and disk space to handle Elasticsearch's requirements. The exact amount depends on your use case, but having at least 4GB of RAM and 10GB of disk space is a good starting point

- **Network Configuration:**
  - Ensure that the necessary ports (default is 9200 for HTTP and 9300 for transport) are open and not blocked by firewalls between **Syracuse & Elasticsearch**

- **User Permissions:**
  - The user running Elasticsearch should have the necessary permissions to read and write to the installation directories and data paths.

Once these prerequisites are in place, you can proceed with downloading and installing Elasticsearch using the Windows .zip archive

Install Elasticsearch with .zip on Windows

The Install

1. Download the Installation **.zip** file from Elastic search note that the **.msi is no longer available**. Download Elasticsearch

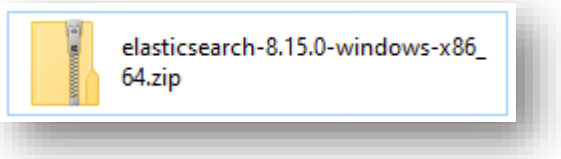

2. Extract the contents to the root of the C drive

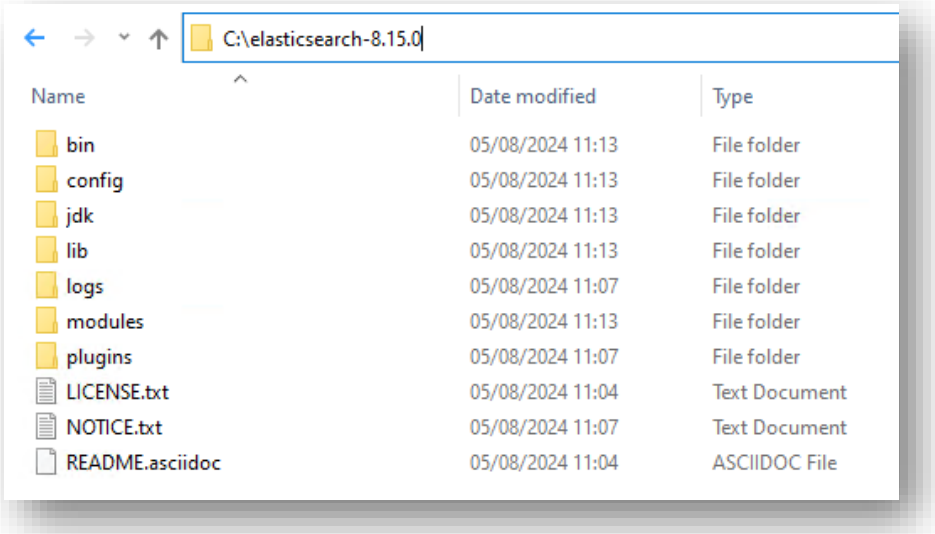

3. Execute the batch file which launches the elasticsearch-service.bat with the required parameters (as Administrator)

REM Set the name of the Elasticsearch service
set "mzNewServiceName=elasticsearch-8.15.0"

REM Set the home directory for the Elasticsearch installation
set "ES_HOME=C:\elasticsearch-8.15.0"

REM Set the path to the JDK if you want to specify your own JDK
REM set “ES_JAVA_HOME=C:\elasticsearch-8.15.0\jdk"

REM The following lines should not need to be changed
REM Change the current directory to the 'bin' folder of Elasticsearch
cd /d "%ES_HOME%\bin"

REM Set the path to the configuration directory for Elasticsearch
set "ES_PATH_CONF=%ES_HOME%\config"

REM Set the Elasticsearch service to start automatically with Windows
set "ES_START_TYPE=auto"

REM Define the display name for the service
set "SERVICE_DISPLAY_NAME=%mzNewServiceName%"

REM Provide a description for the service
set "SERVICE_DESCRIPTION=%mzNewServiceName%"

REM Install the Elasticsearch service with the specified configurations
elasticsearch-service.bat install

The Install

4. Once the installation has been completed you will see the message in the command window

```
Select C:\Windows\System32\cmd.exe
C:\elasticsearch-8.15.0\bin>REM Install the Elasticsearch service with the specified configurations

C:\elasticsearch-8.15.0\bin>elasticsearch-service.bat install
warning: ignoring JAVA_HOME=C:\Program Files\Zulu\zulu-8-jre; using bundled JDK
Installing service : elasticsearch-service-x64
Using ES_JAVA_HOME : C:\elasticsearch-8.15.0\jdk
The service 'elasticsearch-service-x64' has been installed

C:\elasticsearch-8.15.0\bin>
```

Note: JAVA_HOME=C:\Program Files\Zulu\zulu-8-jre is ignored because we specified no JDK path so the bundled JDK is used.

5. The Elasticsearch service will have been created

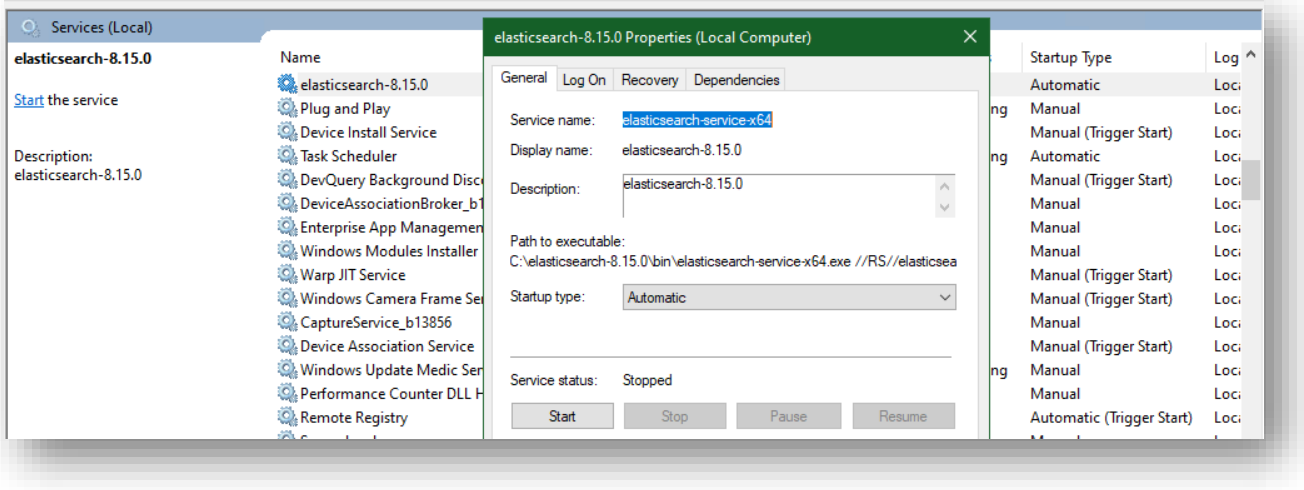

[Install Elasticsearch with .zip on Windows | Elasticsearch Guide [8.16] | Elastic]

Default Elasticsearch Install

- By launching Elasticsearch 8 as an application using bin**elasticsearch.bat** the Elasticsearch instance is automatically configured with default parameters which include security-enabled.
- The certificates created are not in PEM format
- Certificate client authentication is not enabled only user authentication

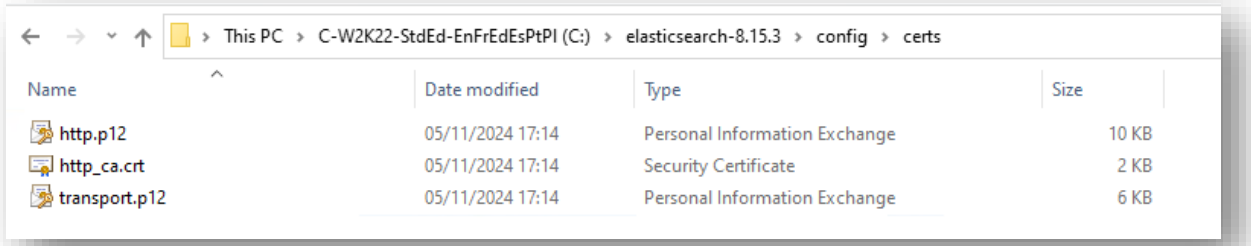

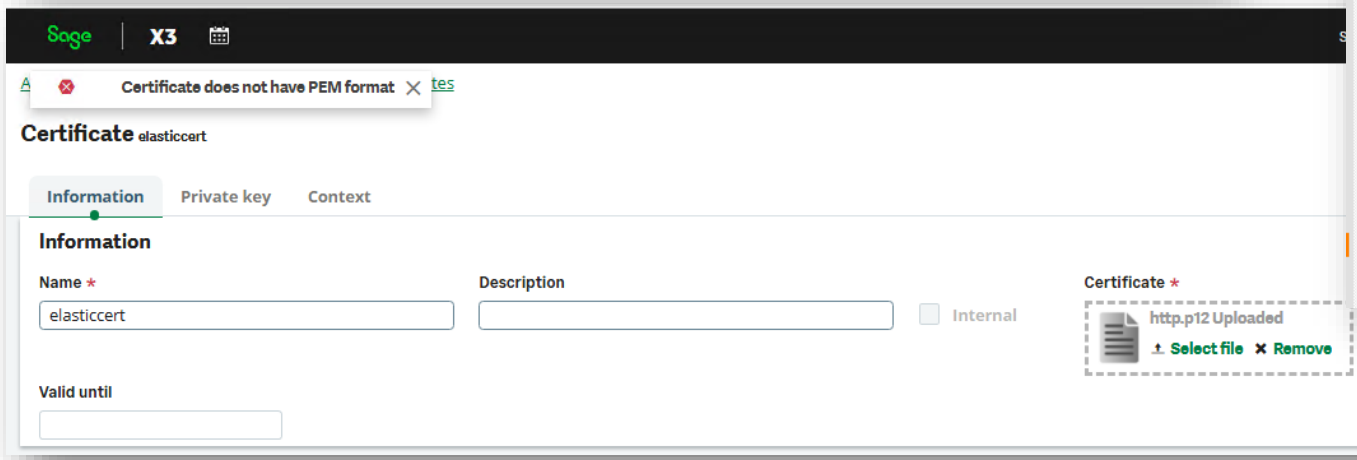

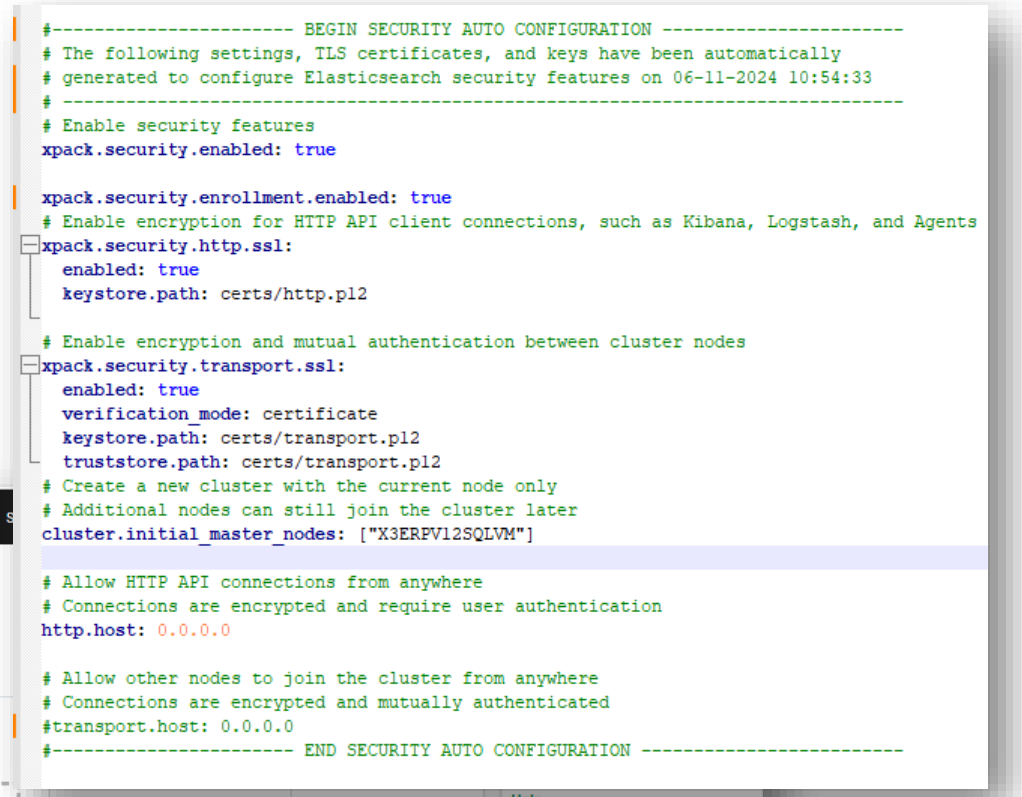

```
#---------------------- BEGIN SECURITY AUTO CONFIGURATION ----------------------
# The following settings, TLS certificates, and keys have been automatically
# generated to configure Elasticsearch security features on 06-11-2024 10:54:33
# -------------------------------------------------------------------------------
# Enable security features
xpack.security.enabled: true

xpack.security.enrollment.enabled: true
# Enable encryption for HTTP API client connections, such as Kibana, Logstash, and Agents
xpack.security.http.ssl:
  enabled: true
  keystore.path: certs/http.p12

# Enable encryption and mutual authentication between cluster nodes
xpack.security.transport.ssl:
  enabled: true
  verification_mode: certificate
  keystore.path: certs/transport.p12
  truststore.path: certs/transport.p12
# Create a new cluster with the current node only
# Additional nodes can still join the cluster later
cluster.initial_master_nodes: ["X3ERPV12SQLVM"]

# Allow HTTP API connections from anywhere
# Connections are encrypted and require user authentication
http.host: 0.0.0.0

# Allow other nodes to join the cluster from anywhere
# Connections are encrypted and mutually authenticated
#transport.host: 0.0.0.0
#---------------------- END SECURITY AUTO CONFIGURATION ------------------------
```

Start the Elastic Stack with security enabled automatically

Elasticsearch 8+ Security

Elasticsearch Security Overview

Elasticsearch introduced several new security features in version 8.0+ to address the need for robust data protection and compliance in modern IT environments. The features were introduced to secure data, and network communication in self-managed clusters.

Features include

- **TLS for Encryption**: Securing communication between nodes and clients using HTTPS
- **User Authentication**: Ensuring that only authorised users can access the Elasticsearch cluster
- **TLS Authentication**: Ensuring that only authorised users can access the Elasticsearch cluster
- **Security is enabled by default** when you start Elasticsearch for the first time, **Note that we recommend Elasticsearch to run as a service for Sage X3, if you start as a service security will not be configured**.

Set up minimal security for Elasticsearch

Security settings in Elasticsearch

Elasticsearch Communication HTTPS

We can use TLS, enable a HTTPS configuration to ensure that communication between clients and the cluster is always encrypted using HTTPS. We can use the **elasticsearch-certutil utility** which is a command-line utility in Elasticsearch that simplifies the creation and management of certificates

The Steps

1. Generate a Certificate Authority (CA) using elasticsearch-certutil utility

2. Generate Certificates and Keys elasticsearch-certutil utility

3. Move the generated files to the Elasticsearch configuration directory (e.g. /etc/elasticsearch/certs/).

4. Configure Elasticsearch by editing the elasticsearch.yml file to enable HTTPS and specify the paths to your certificates this will be covered in the elasticsearch.yml configuration section

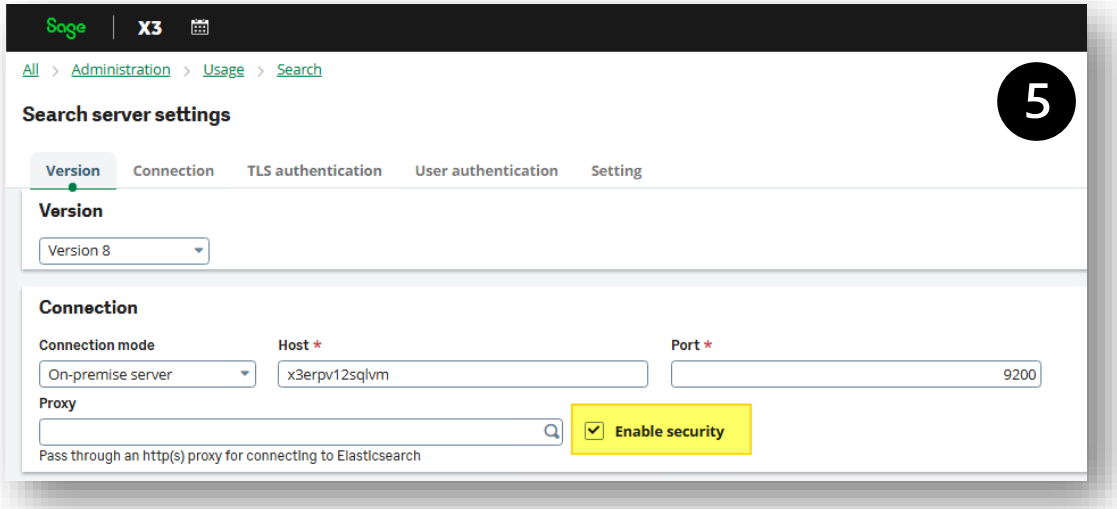

Elasticsearch TLS Authentication

TLS certificate authentication for users in Elasticsearch 8 involves using Public Key Infrastructure (PKI) to authenticate clients based on digital certificates. This method ensures that both the client and server verify each other's identity through certificates issued by a trusted Certificate Authority (CA).

Certificates have to be imported into Syracuse using the **certificates** function for the certificate and **certificates of certification authorities for** the CA certificate.

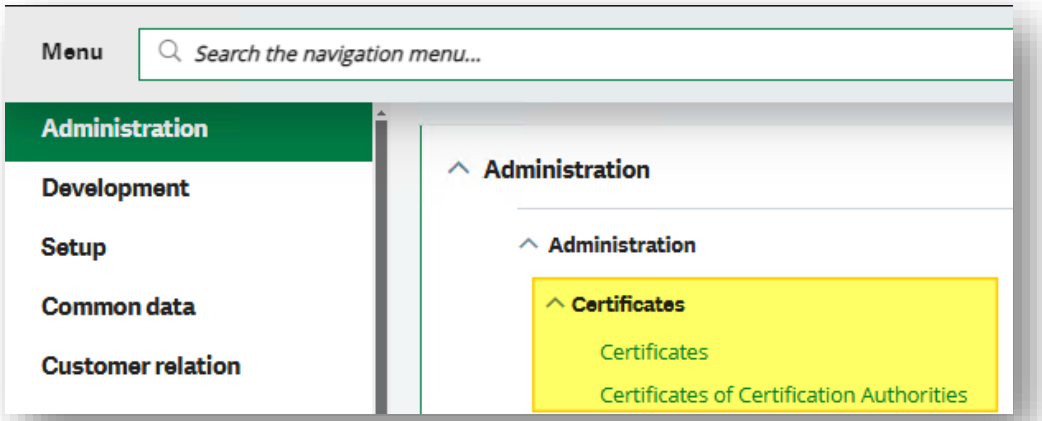

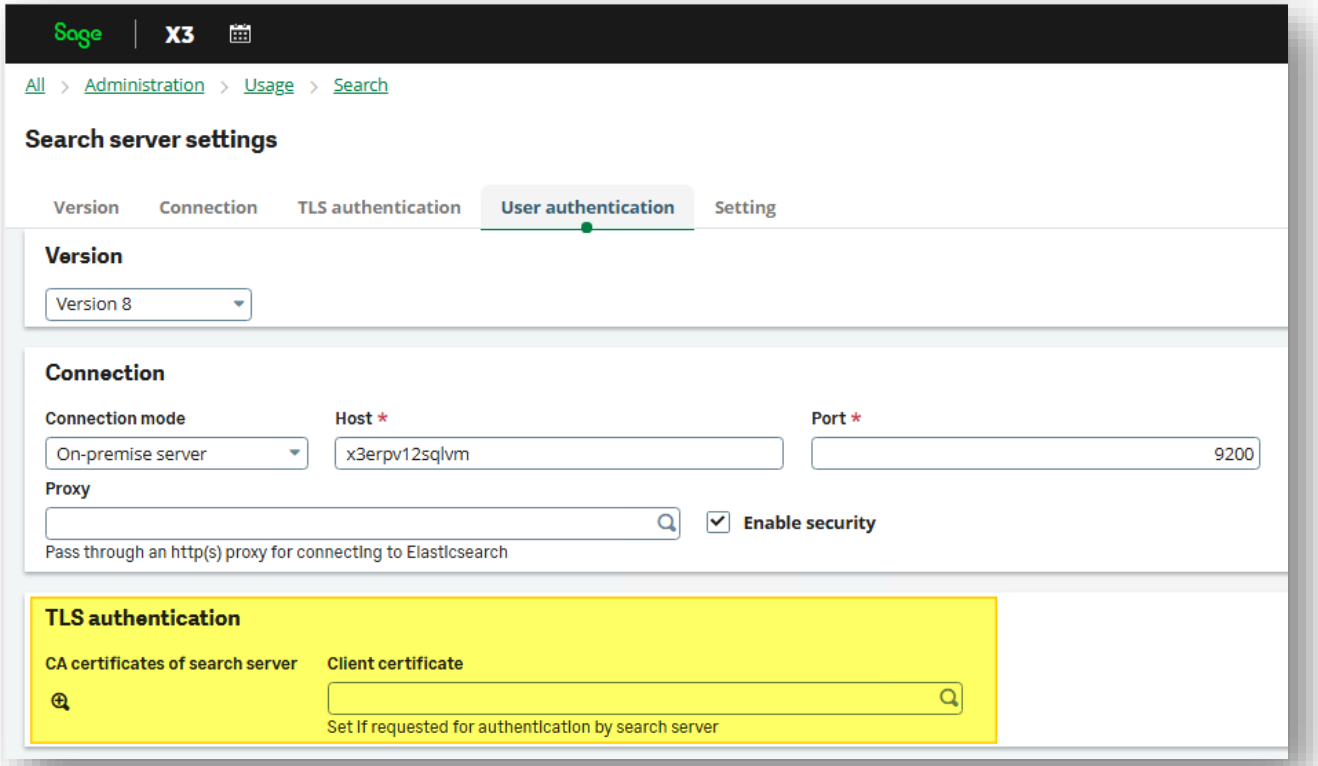

Elasticsearch User Security

Authentication in Elasticsearch 8 ensures that only authorised users can access the cluster. Elasticsearch supports various authentication methods, including **basic** authentication (username and password), **API key authentication**, and **JSON Web Token (JWT) authentication**. When security features are enabled, users must provide their credentials to access the cluster

Elasticsearch comes with several built-in users, each serving specific purposes these built-in users have predefined roles and privileges. The elastic user is a superuser with full access to the cluster and can manage security settings.

For X3 we have the option to use basic or API Key authentication mode for the user authentication

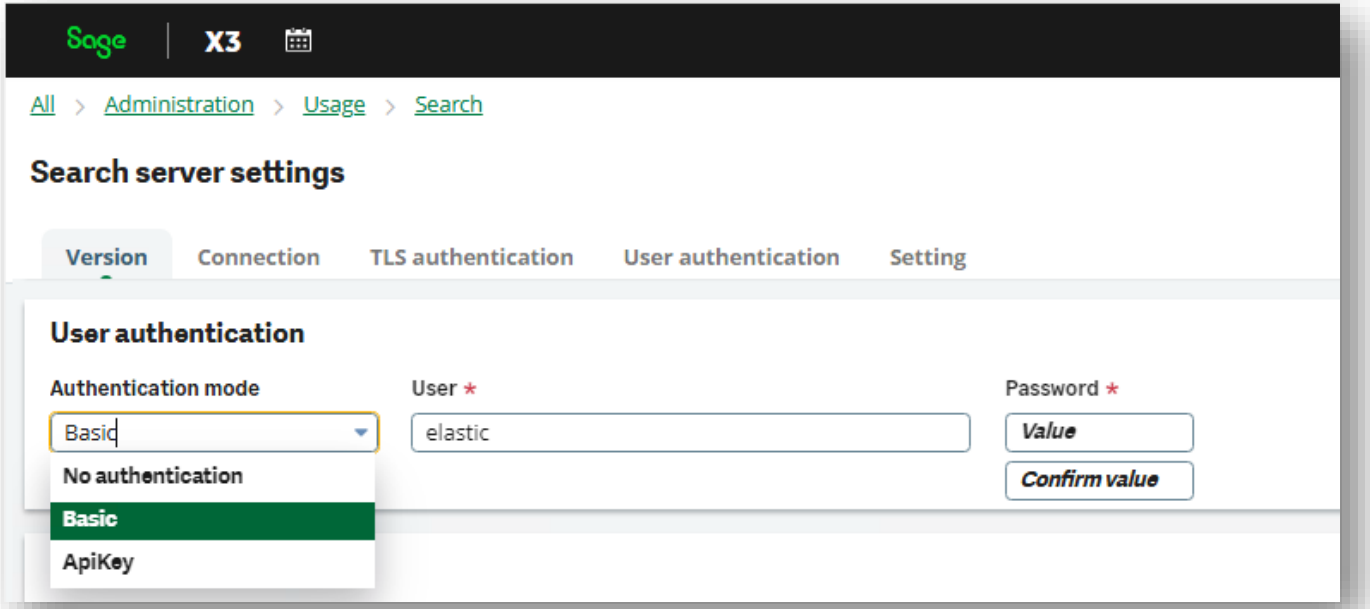

Note: The elastic user is a built-in superuser in Elasticsearch, best practice would be to create your own user with the appropriate roles

Built-in users
Elasticsearch-users

Elasticsearch Security Configuration

Security Setup/Configuration

Once the installation is completed, we can move on to configure our instance the steps we need to complete are

- **CA certificate Generation**
- **Certificate generation and deployment for SSL/HTTPS and client authentication**
- **Update Elasticsearch security configuration in YML**
- **Elastic User Password Setup**
- **Test a remote authentication to the Elasticsearch server\service**

Certificates For Communication And Authentication

To generate the certificates for HTTPS communication we will use the elasticsearch-certutil utility. All the Elasticsearch utilities are found in the \bin directory of the Elasticsearch application. You can use your Certificate Authority (CA) or generate a self-signed CA; the CA validates client certificates in HTTP SSL connections

To generate the Certificate Authority (CA)

We can use the command

bin\elasticsearch-certutil ca --pem --ca-dn CN=esserver

elasticsearch-certutil : This is a tool included with Elasticsearch to assist with certificate generation

CA : This flag specifies that we are generating a Certificate Authority (CA) type certificate .

--pem : This flag specifies that the output should be in PEM format, a base64-encoded format for encoding certificates.

ca-dn CN=esserver : This option sets the Distinguished Name (DN) of the CA to "CN=esserver". "CN" stands for Common Name, part of the CA's identity.

This command will create a new Certificate Authority with the common name “esserver" and produce the necessary files in PEM format.

Certificates For Communication And Authentication

To generate a certificate signed by the CA to secure communication

We can use the command

bin\elasticsearch-certutil cert --pem --ca-cert config/certs/ca.crt --ca-key config/certs/ca.key --dns esserver --ip 127.0.0.1,10.1.19.150 --name elasticclient

cert --pem: This specifies that you want to generate a PEM formatted certificate.

--ca-cert config/certs/ca.crt: This points to the CA certificate file.

--ca-key config/certs/ca.key: This points to the CA key file.

--dns esserver: This sets the DNS name for the certificate.

--ip 127.0.0.1,10.1.19.150: These are the IP addresses to be included in the certificate.

--name elasticclient: This sets the name for the certificate

Generated Certificates

In the last steps in the certificate generation, all certificates are created as .zip files. Unzip all the certificates in the cert's directory

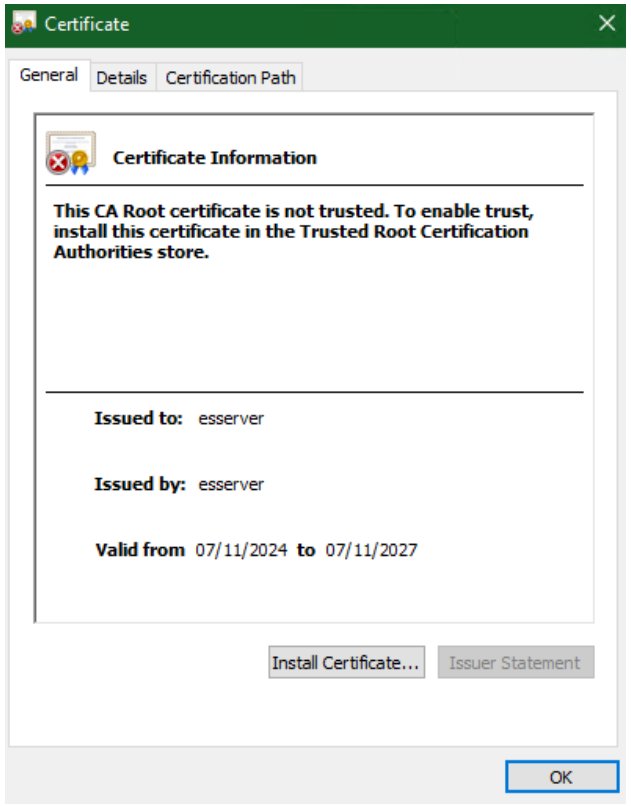

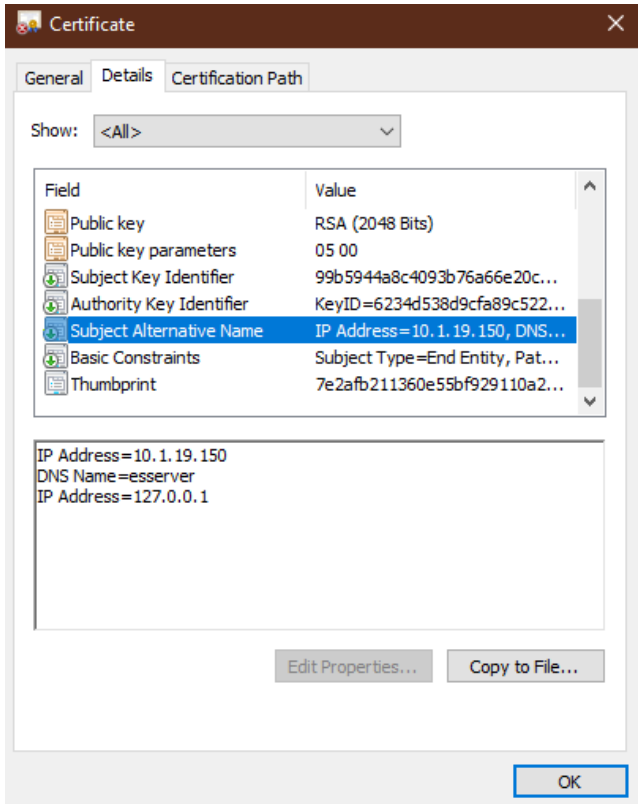

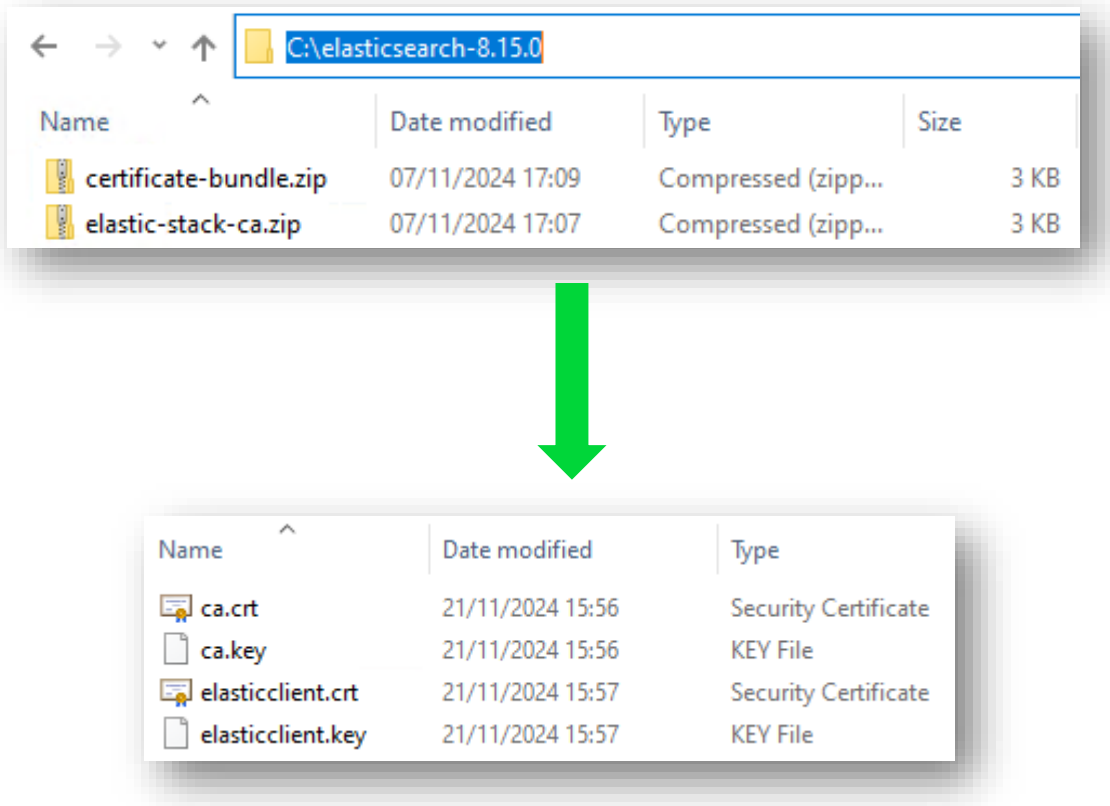

Elasticsearch YML

The elasticsearch.yml file is the main configuration file for Elasticsearch. It is used to set various parameters and options that control the behaviour of your Elasticsearch instance

```
xpack.security.enabled: true

xpack.security.http.ssl.enabled: true
xpack.security.http.ssl.certificate_authorities: [ "certs/ca.crt" ]
xpack.security.http.ssl.certificate: certs/elasticsearch.crt
xpack.security.http.ssl.key: certs/elasticsearch.key
xpack.security.http.ssl.client_authentication: required

xpack.security.transport.ssl.enabled: true
xpack.security.transport.ssl.certificate: certs/elasticsearch.crt
xpack.security.transport.ssl.key: certs/elasticsearch.key

cluster.name: Esserver
node.name: node-1
network.host: 0.0.0.0
discovery.seed_hosts: []
cluster.initial_master_nodes: [node-1]
```

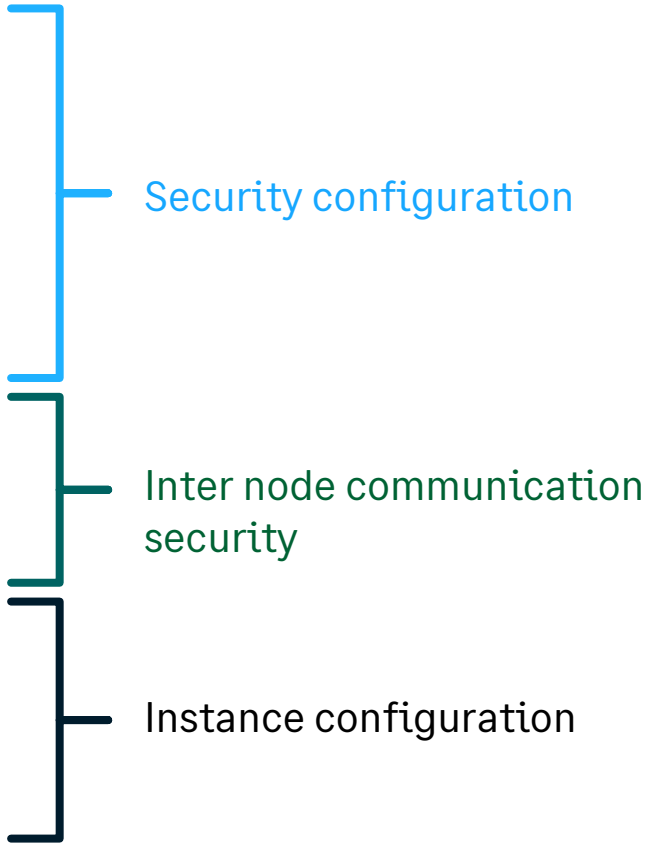

[Configuring Elasticsearch | Elasticsearch Guide [8.15] | Elastic]

Elasticsearch YML

The security parameters enable security for the instance, HTTPS using the generated certificates as well as certificate authentication for users connecting to the instance of Elasticsearch

Security Parameter	Description
xpack.security.enabled: true	Enables X-Pack security features
xpack.security.http.ssl.enabled: true	Enables SSL/TLS encryption for HTTP connections, ensuring data sent to and from the Elasticsearch server is encrypted
xpack.security.http.ssl.certificate_authorities: ["certs/ca.crt"]	specifies the path to the certificate authority (CA) file that validates client certificates in HTTP SSL connections
xpack.security.http.ssl.certificate: certs/elasticsearch.crt	Defines the path to the server certificate file for HTTP SSL, which identifies the Elasticsearch server during connections.
xpack.security.http.ssl.key: certs/elasticsearch.key	Specifies the path to the private key for the HTTP SSL certificate, used to establish secure HTTPS connections.
xpack.security.http.ssl.client_authentication: required	Requires client authentication for HTTP connections, ensuring only clients with valid certificates can access the server.

Elasticsearch YML

The TLS parameters for inter-node communication are settings for inter node communication that I needed to add as Elasticsearch is installed on a separate server and X3 needs to connect remotely

TLS Parameters	Description
xpack.security.transport.ssl.enabled: true	Enables SSL/TLS encryption for internal transport communication between nodes in the Elasticsearch cluster.
xpack.security.transport.ssl.certificate	Specifies the path to the server certificate file for securing transport communication between Elasticsearch nodes.
xpack.security.transport.ssl.key: certs/elasticsearch.key	Specifies the path to the private key for the transport SSL certificate, used for secure node-to-node communication.

Elasticsearch YML

Cluster Parameter	Description
cluster.name: Esserver	Sets the name of the Elasticsearch cluster, useful for identifying and managing multiple clusters.
node.name: node-1	Specifies the name of this particular node within the cluster, which helps in identifying it among other nodes.
network.host: 0.0.0.0	Binds Elasticsearch to all available network interfaces, making it accessible to remote connections.
discovery.seed_hosts: []	Defines the list of other nodes in the cluster used during node discovery. Here, it's empty, meaning this might be a single-node setup.
cluster.initial_master_nodes: [node-1]	Lists the initial set of master-eligible nodes during the first cluster setup.

Elasticsearch Security (users)

User security in Elasticsearch is crucial for protecting your data and ensuring that only authorized users can access and interact with your Elasticsearch cluster

For data access we will use the ‘elastic’ user, to reset the password

1. Navigate to the Elasticsearch bin Directory **cd C:\elasticsearch-8.15.0\bin**
2. For an interactive password reset **elasticsearch-reset-password -u elastic -i**
3. For an auto-generated password **elasticsearch-reset-password -u elastic –a**

```
C:\elasticsearch-8.15.0\bin>elasticsearch-reset-password -u elastic -i
warning: ignoring JAVA_HOME=C:\Program Files\Zulu\zulu-8-jre; using bundled JDK
This tool will reset the password of the [elastic] user.
You will be prompted to enter the password.
Please confirm that you would like to continue [y/N]y


Enter password for [elastic]:
Re-enter password for [elastic]:
Password for the [elastic] user successfully reset.
```

https://www.elastic.co/guide/en/elasticsearch/reference/current/reset-password.html

Elasticsearch Memory Settings

To set the memory settings for Elasticsearch 8 on a Windows Server, follow these steps:

1.Locate the JVM Options File:

- Navigate to the config directory within your Elasticsearch installation folder.
- Open the jvm.options file in the \config directory.

2.Set the JVM Heap Size:

- Add or modify the following lines to set the minimum and maximum heap size (replace 2g with your desired heap size): for example –Xms1g –Xmx1g
- Ensure that both Xms and Xmx values are the same

3.Using Elasticsearch Service Manager:

- Open a Command Prompt as an administrator.
- Navigate to the bin directory of your Elasticsearch installation.
- Open the Elasticsearch Service Manager using command: **elasticsearch-service.bat manager**
- In the Service Manager, go to the Java tab.
- Set the Initial Memory Pool and Maximum Memory Pool to your desired values (e.g., 2048 MB for 2 GB).

Elasticsearch Memory Settings

elasticsearch-service.bat manager & JVM Options file

```
################################################################
## IMPORTANT: JVM heap size
################################################################
##
## The heap size is automatically configured by Elasticsearch
## based on the available memory in your system and the roles
## each node is configured to fulfill. If specifying heap is
## required, it should be done through a file in jvm.options.d,
## which should be named with .options suffix, and the min and
## max should be set to the same value. For example, to set the
## heap to 4 GB, create a new file in the jvm.options.d
## directory containing these lines:
##
## -Xms4g
## -Xmx4g
##
## See https://www.elastic.co/guide/en/elasticsearch/reference/8.15/heap-size.html
## for more information
##
################################################################
```

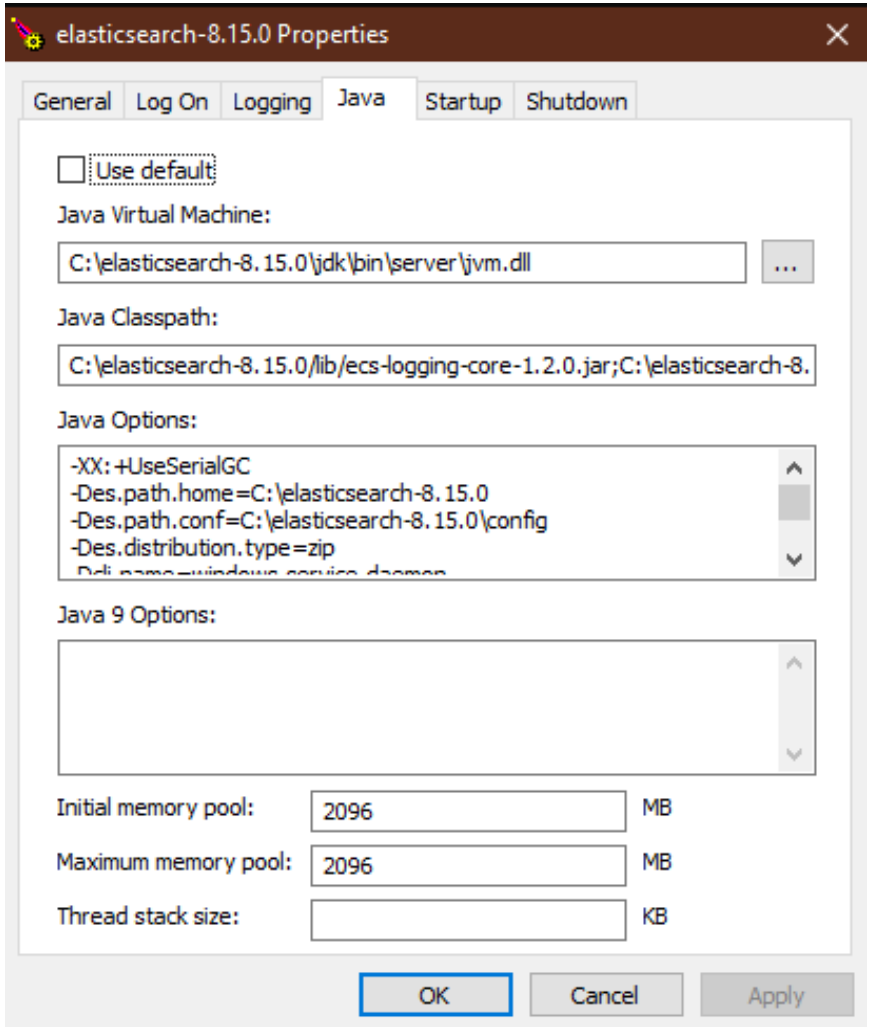

[Performance tuning your Sage X3 system: Elastic Search – Sage X3 UK Support & Insights - Sage X3 UK - Community Hub]

Starting Elasticsearch

Start the service using services.msc

Check logs in elasticsearch-8.15.0\logs to make sure the cluster health is GREEN

Note: The first time the log takes longer to populate as its setting up the node

vice-x64.2024-09-24.log

```
a.f.FileUserRolesStore] [node-1] users roles file [C:\elasticsearc
s.SecurityIndexManager] [node-1] security index does not exist, cr
MetadataCreateIndexService] [node-1] [.security-7] creating index,
a.AllocationService] [node-1] updating number_of_replicas to [0] f
s.SecurityMigrationExecutor] [node-1] Security migration not neede
a.AllocationService] [node-1] current.health="GREEN" message="Clus
ealth="YELLOW" reason="shards started [[.security-7][0]]"
a.f.FileUserPasswdStore] [node-1] users file [C:\elasticsearch-8.1
a.f.FileUserRolesStore] [node-1] users roles file [C:\elasticsearc
a.f.FileUserPasswdStore] [node-1] users file [C:\elasticsearch-8.1
a.f.FileUserRolesStore] [node-1] users roles file [C:\elasticsearc
a.f.FileUserPasswdStore] [node-1] users file [C:\elasticsearch-8.1
a.f.FileUserRolesStore] [node-1] users roles file [C:\elasticsearc
```

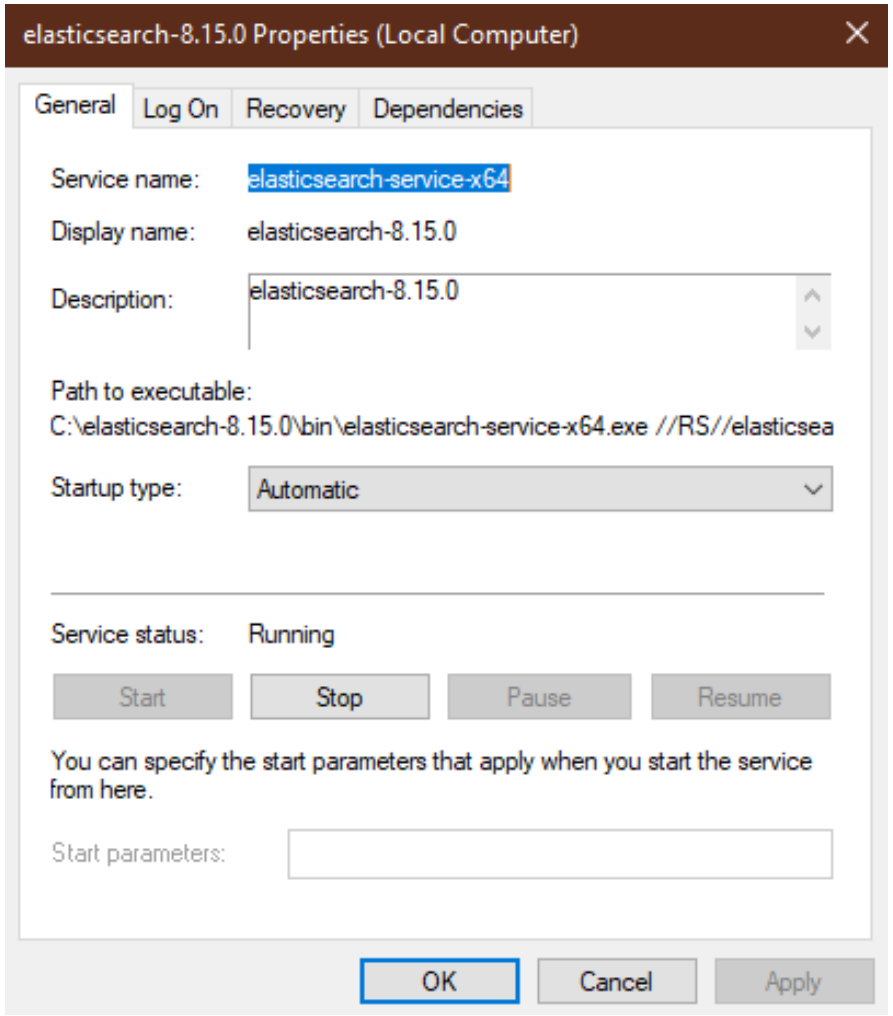

Test Connection To Elasticsearch

Because we have set **xpack.security.http.ssl.client_authentication required** so only clients with valid certificates can access the server. The browser is unable to provide the certificate for authentication. We can use the Postman application which allows us to provide certificates for authentication.

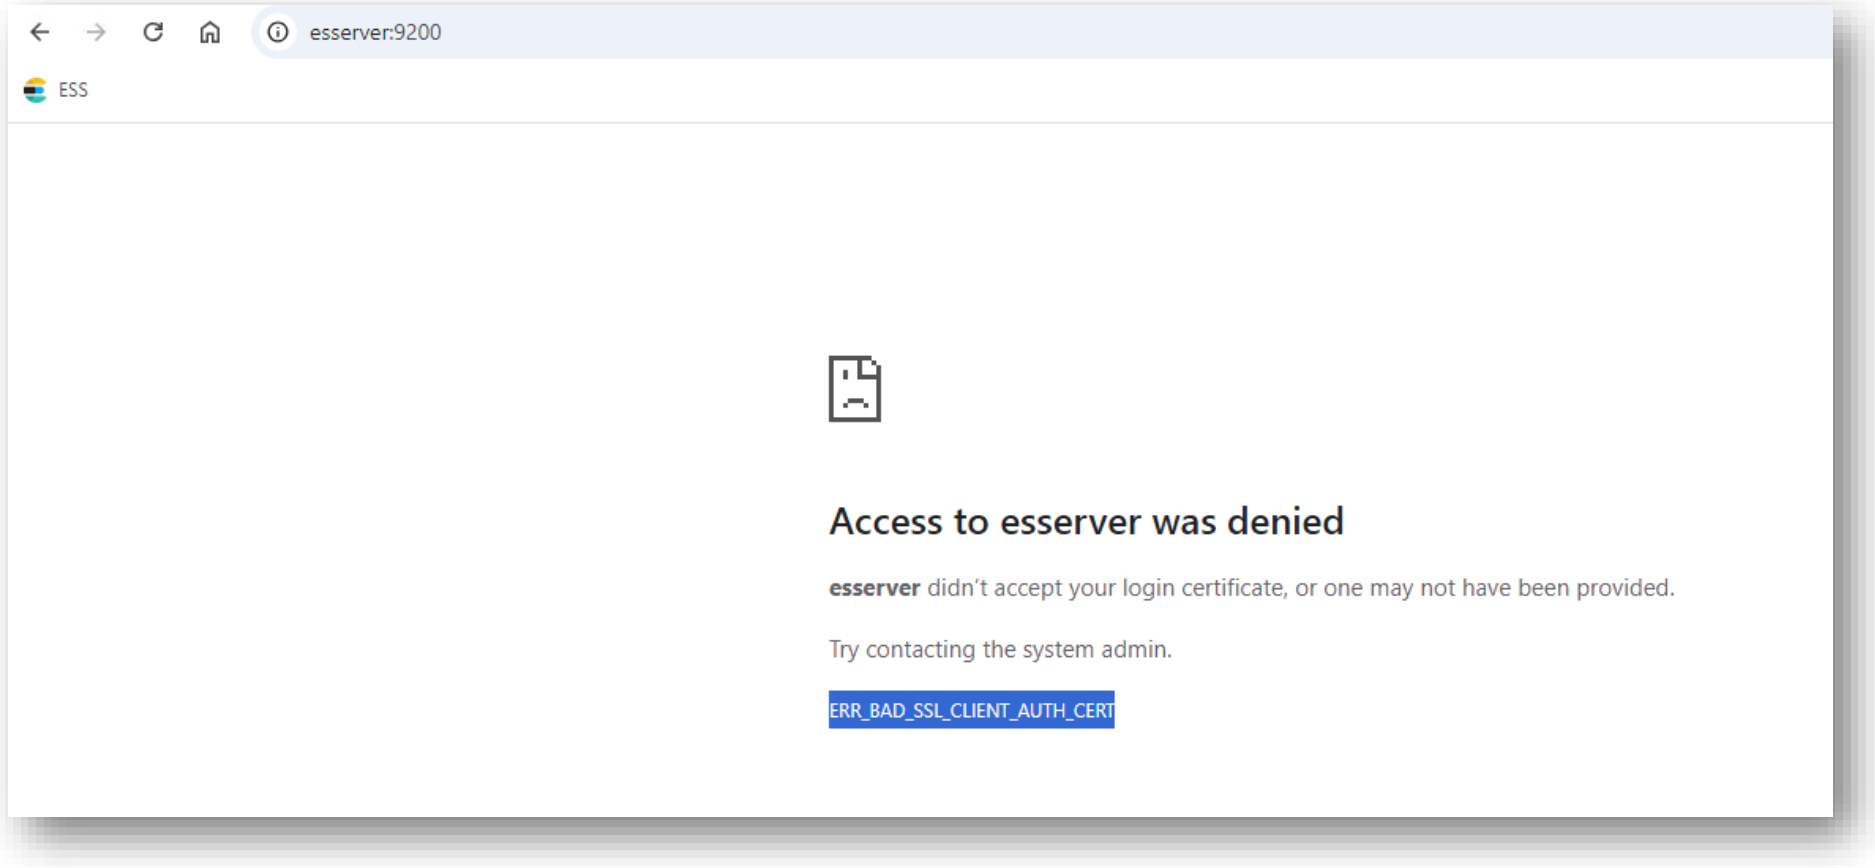

Test Connection To Elasticsearch

Test the connection using the Postman application which allows us to use certificate authentication

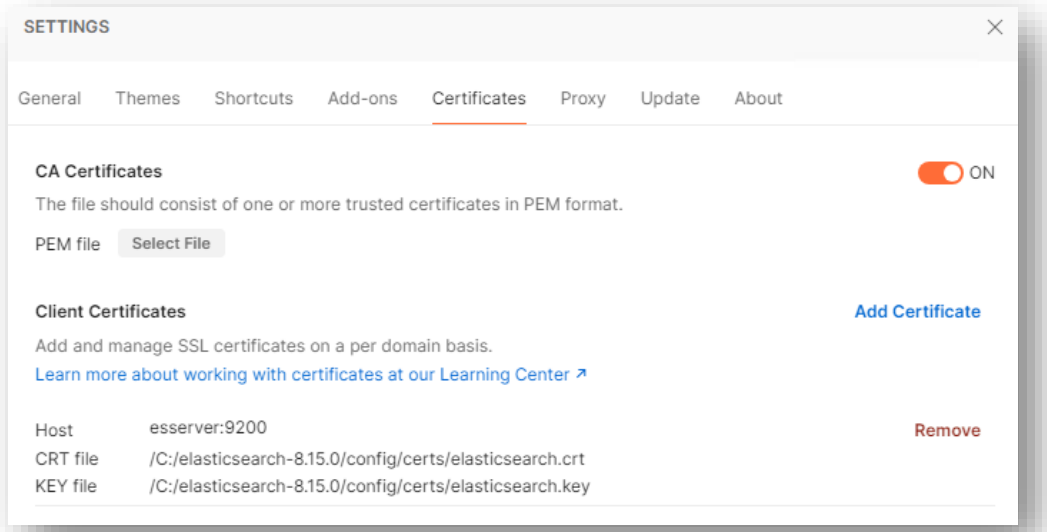

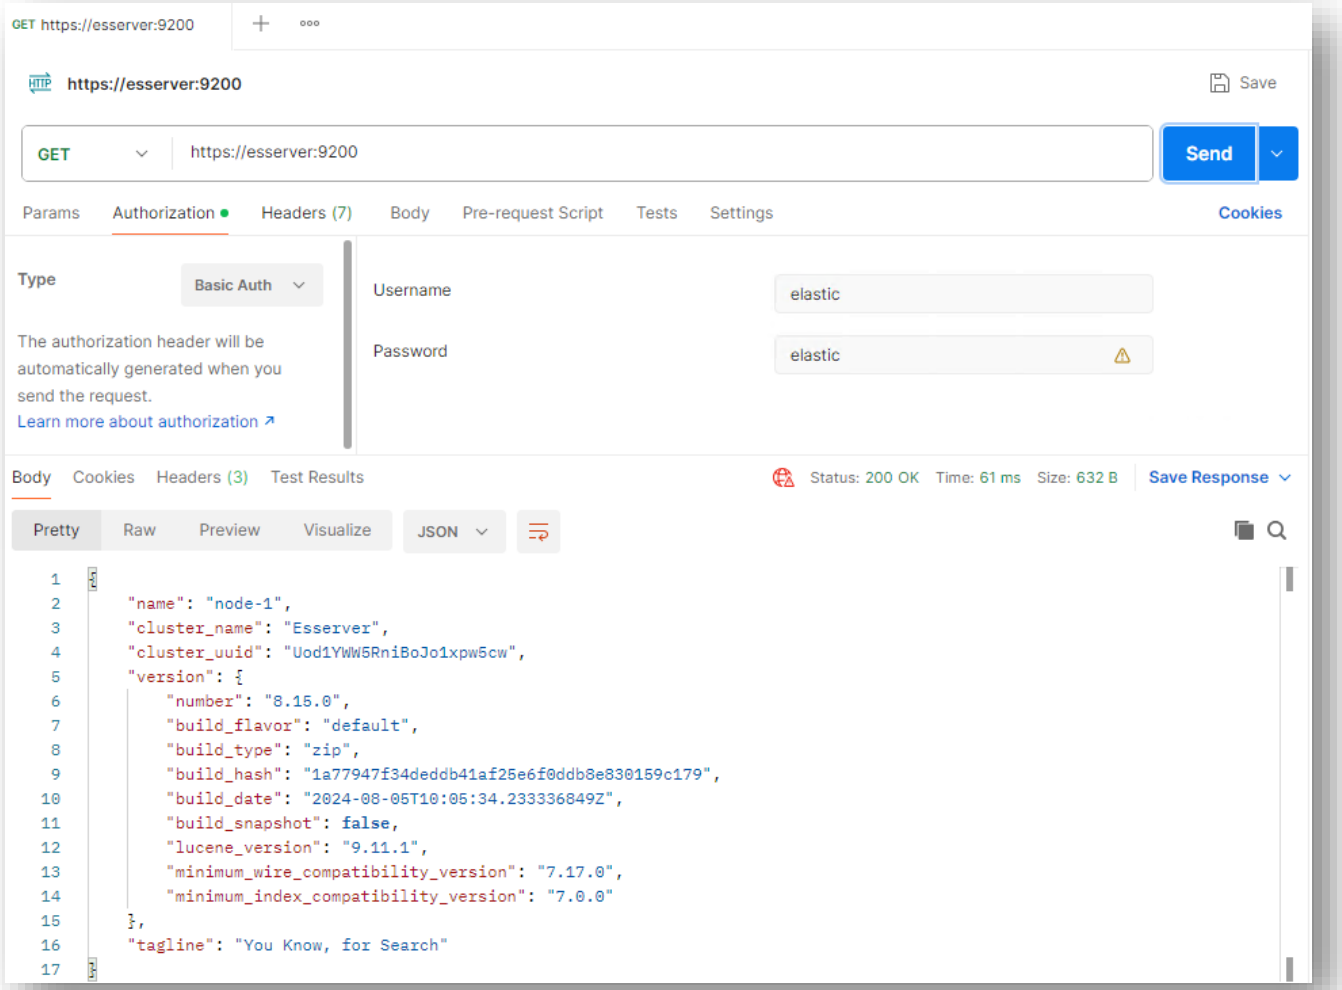

1. Settings > Certificates – Upload the Client certificates for the Elasticsearch host esserver:9200

2. Specify the user & password in your request on the authorization tab.

3. Send your GET Request , See the response

Connecting Sage X3 To Elasticsearch

Importing CA Certificate

Upload CA Certificate from your Elasticsearch server to Syracuse server

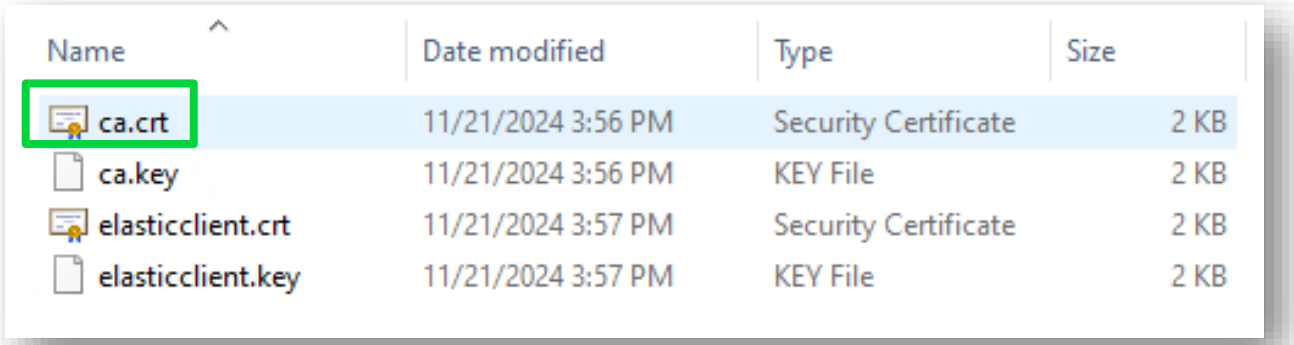

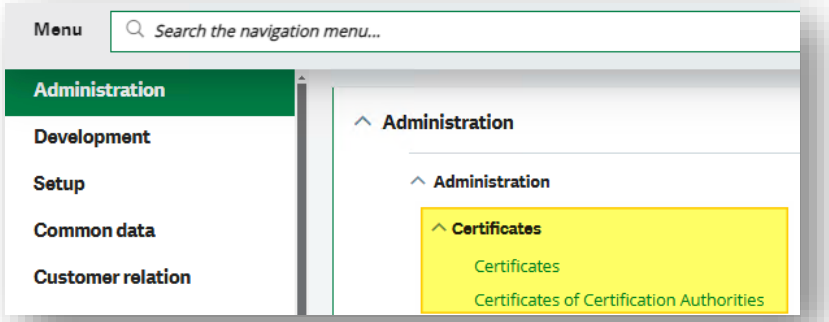

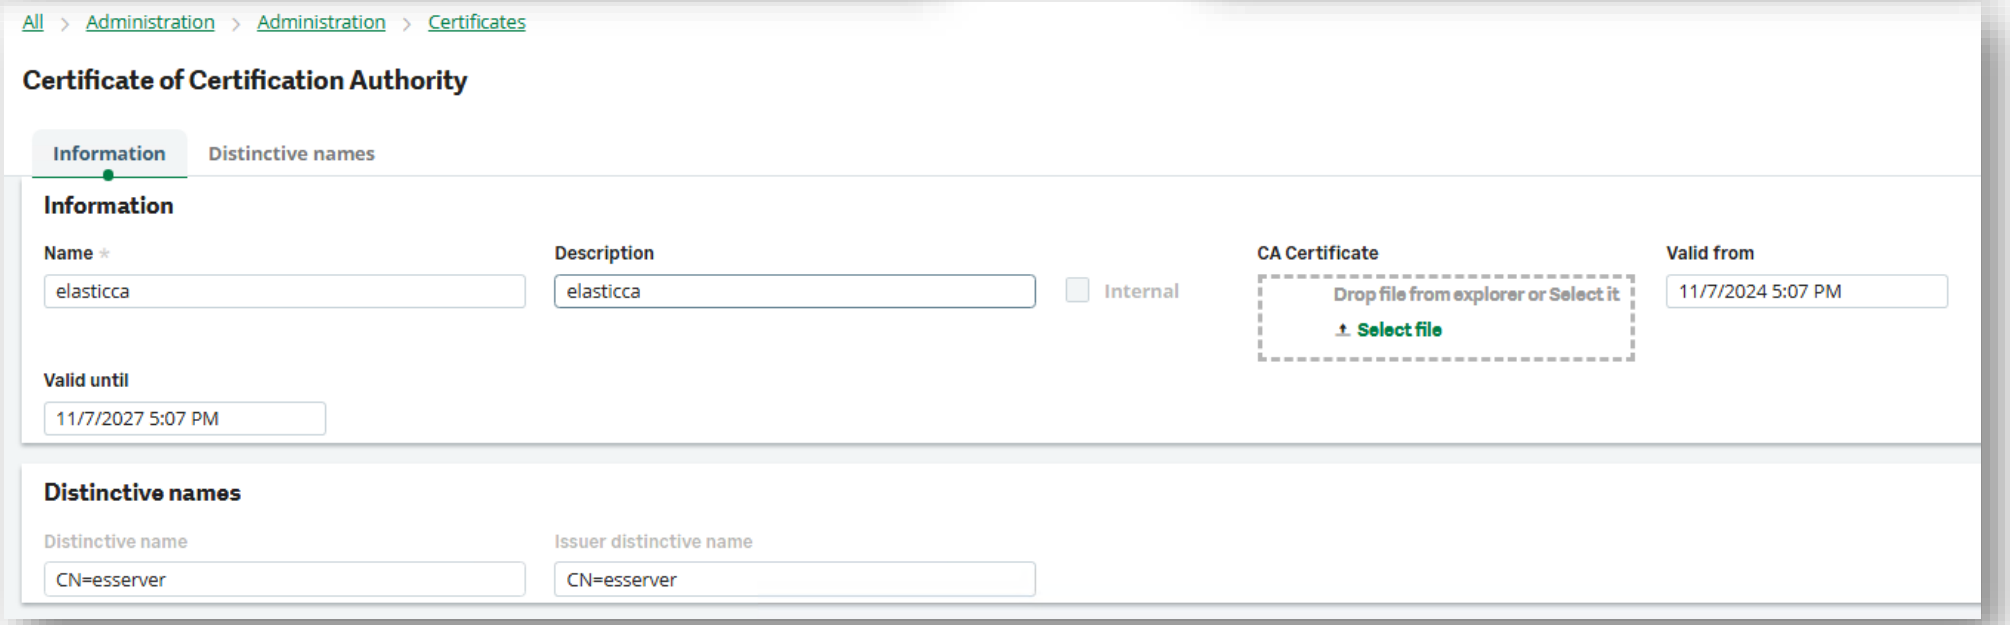

Importing Client Certificates

Upload Client Certificate & key from your elastic server to Sage X3 Syracuse Server

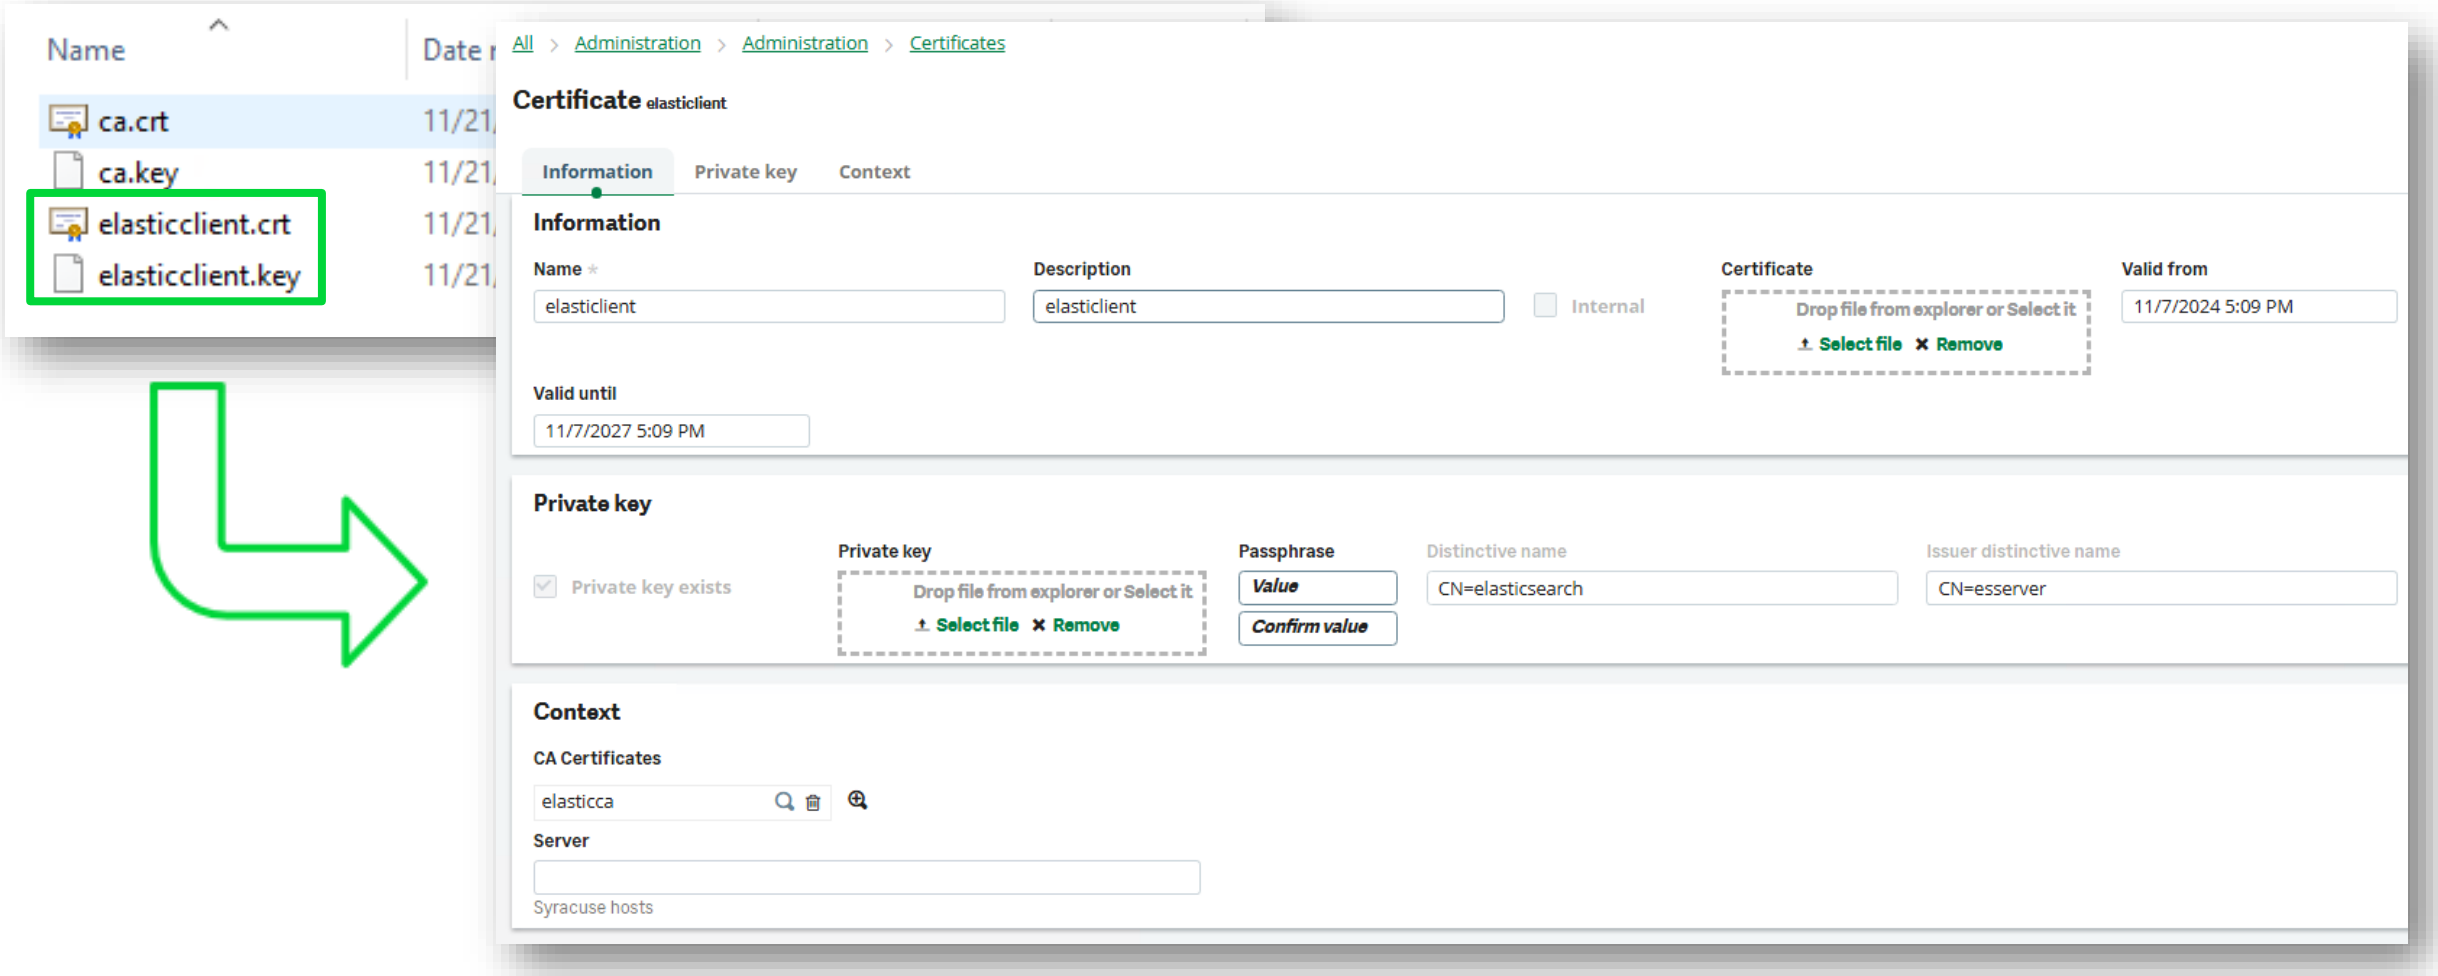

Elasticsearch Connection

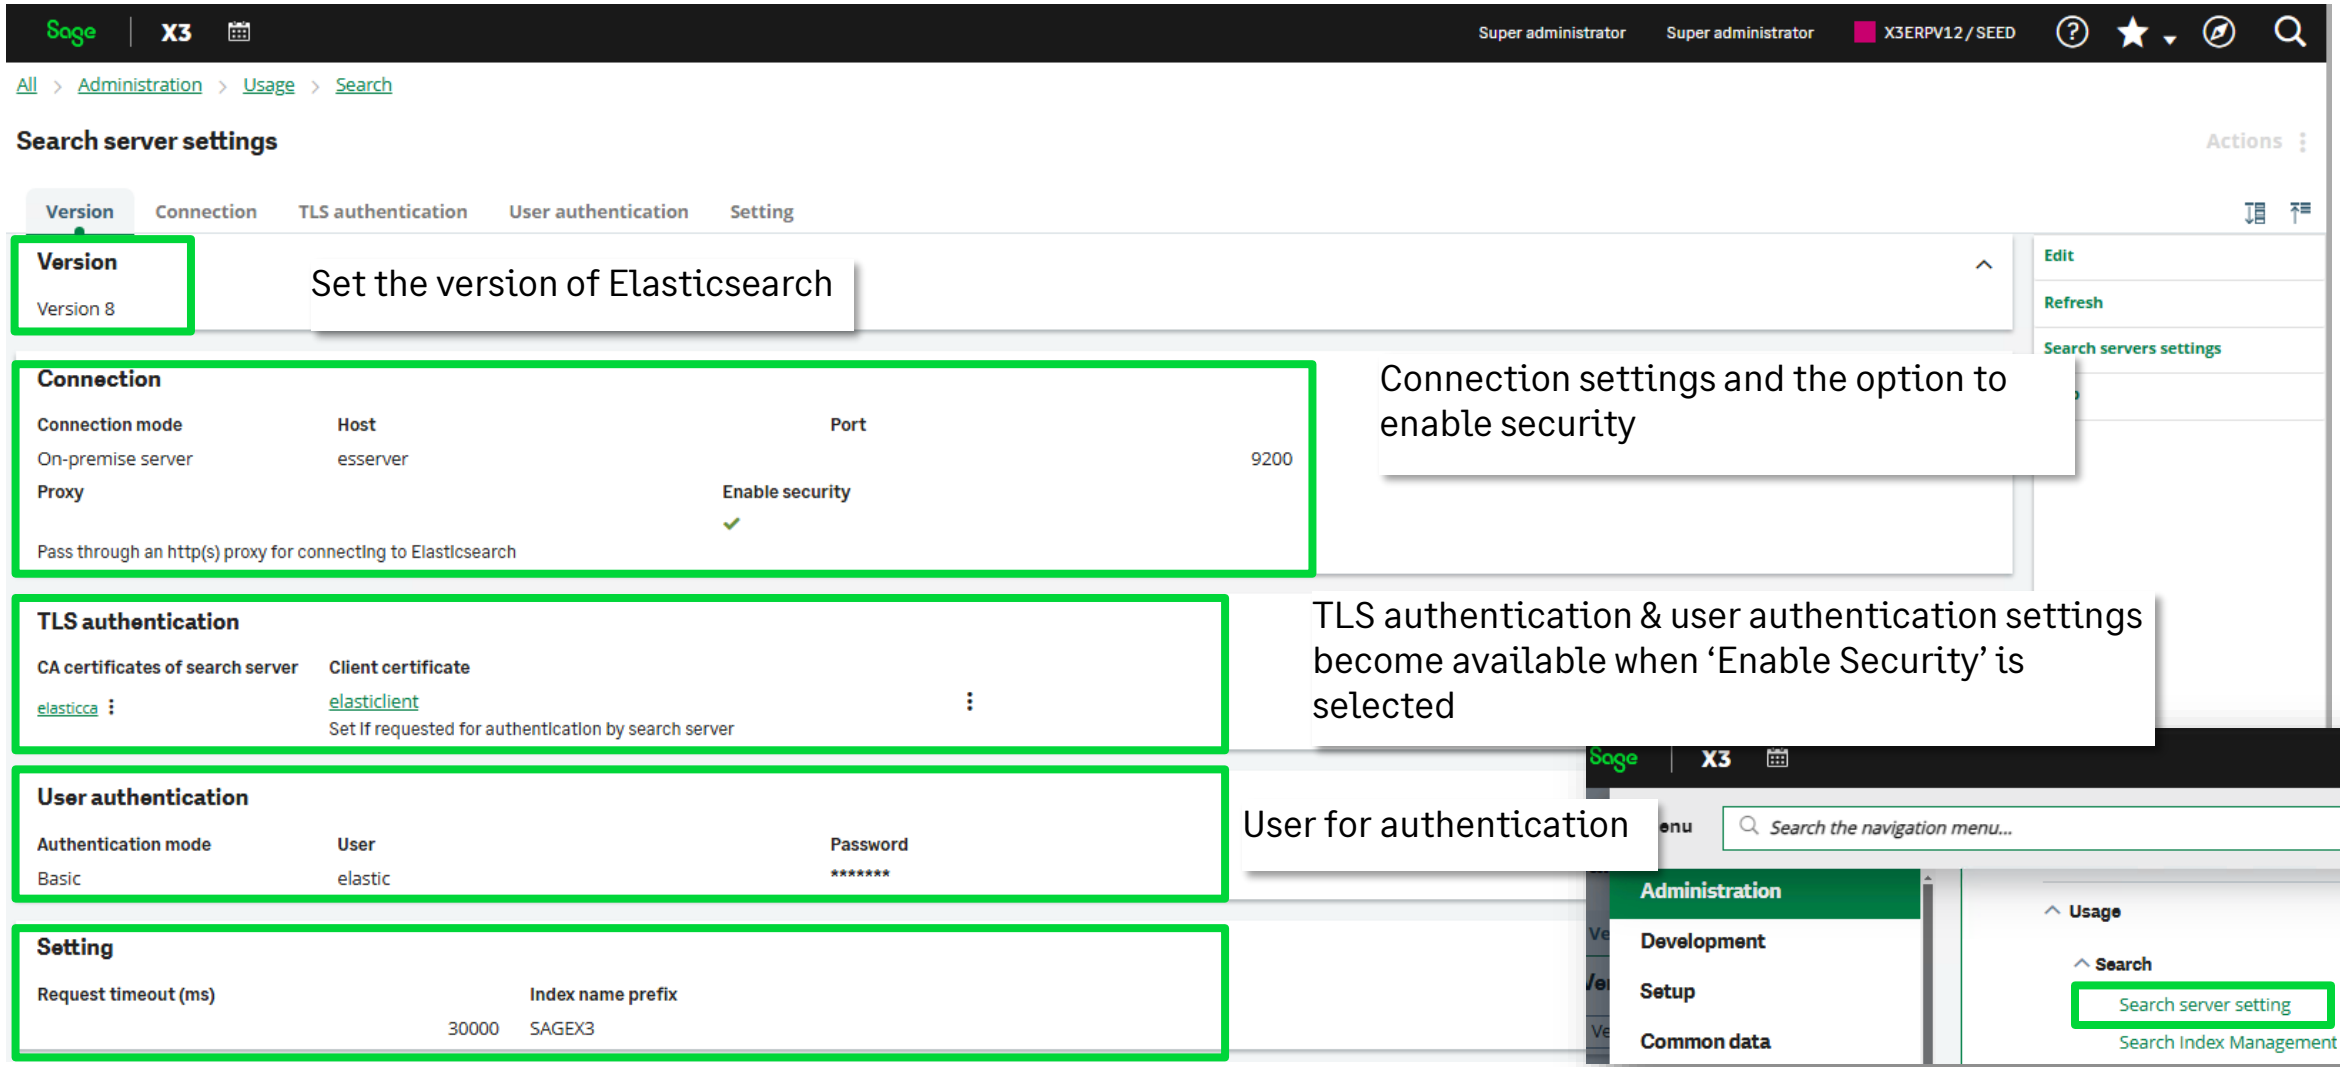

Search Server Connection

Configured Search server settings

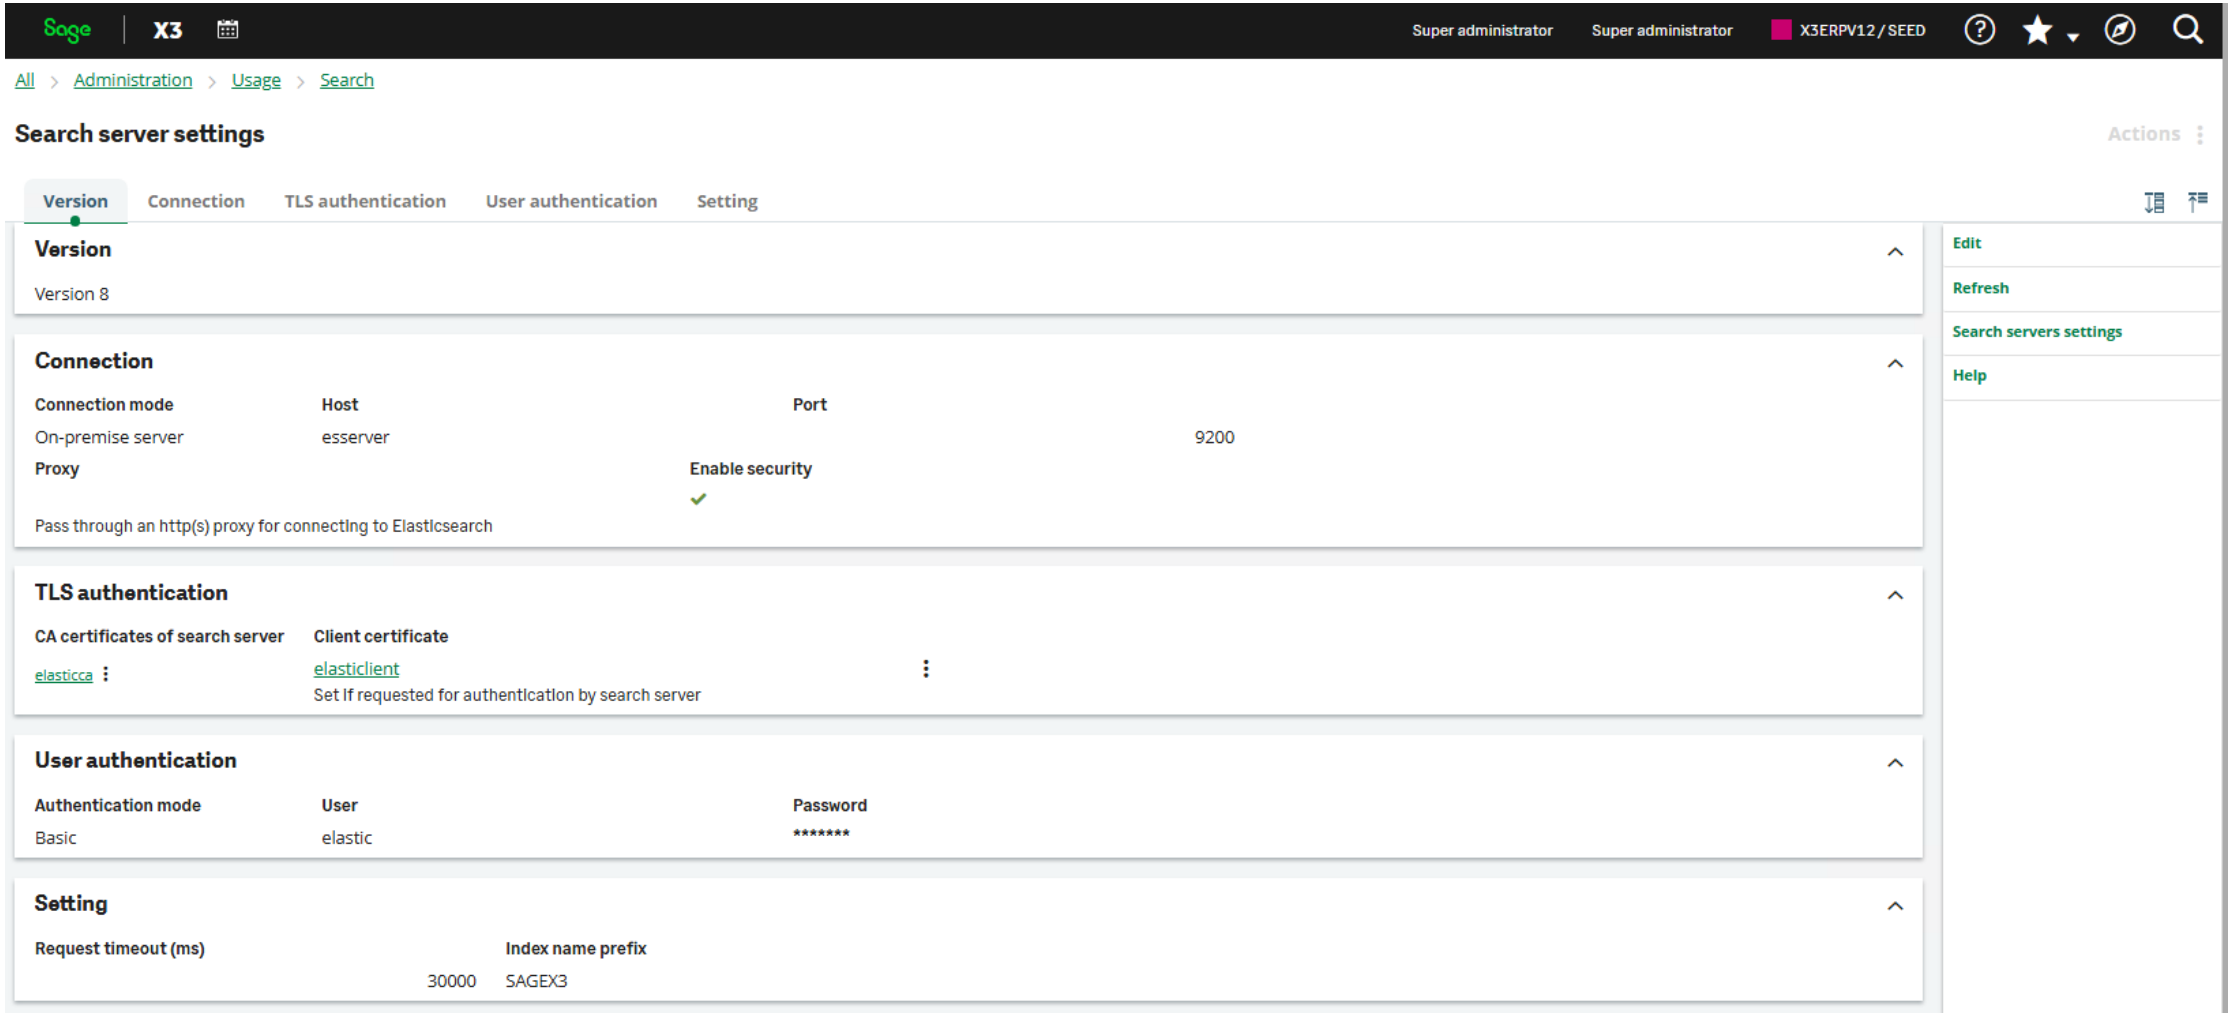

Indexing Management

Elasticsearch indexes are not dynamically updated Indexing and Index management can be performed from the search index management function.

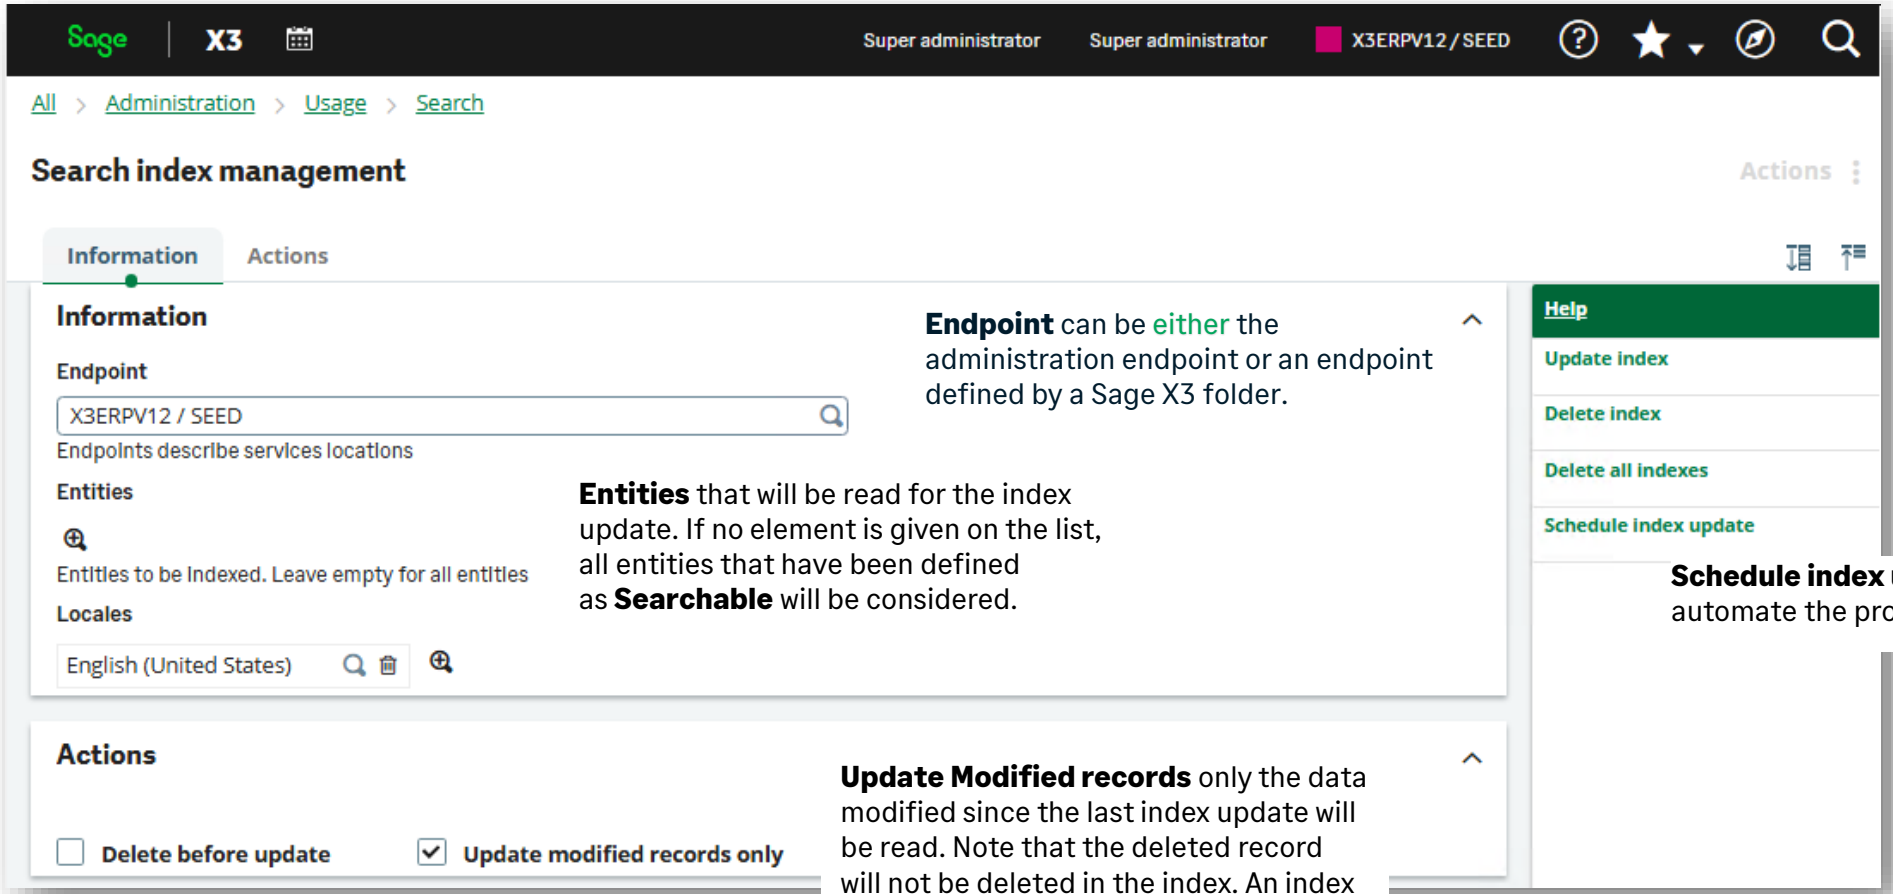

Endpoint can be either the administration endpoint or an endpoint defined by a Sage X3 folder.

Entities that will be read for the index update. If no element is given on the list, all entities that have been defined as **Searchable** will be considered.

Delete before update index operation will delete the endpoint index. This is important when a complete reconstruction of the index must be planned.

Update Modified records only the data modified since the last index update will be read. Note that the deleted record will not be deleted in the index. An index delete followed by an index rebuild must be done.

Schedule index update to automate the process of indexing

Indexing Automation

Schedule index update in the indexing function opens a popup window that lists the existing automation schedules. The task defined by the index update request can be attached to one or more automations to be executed in batch with the right schedule.

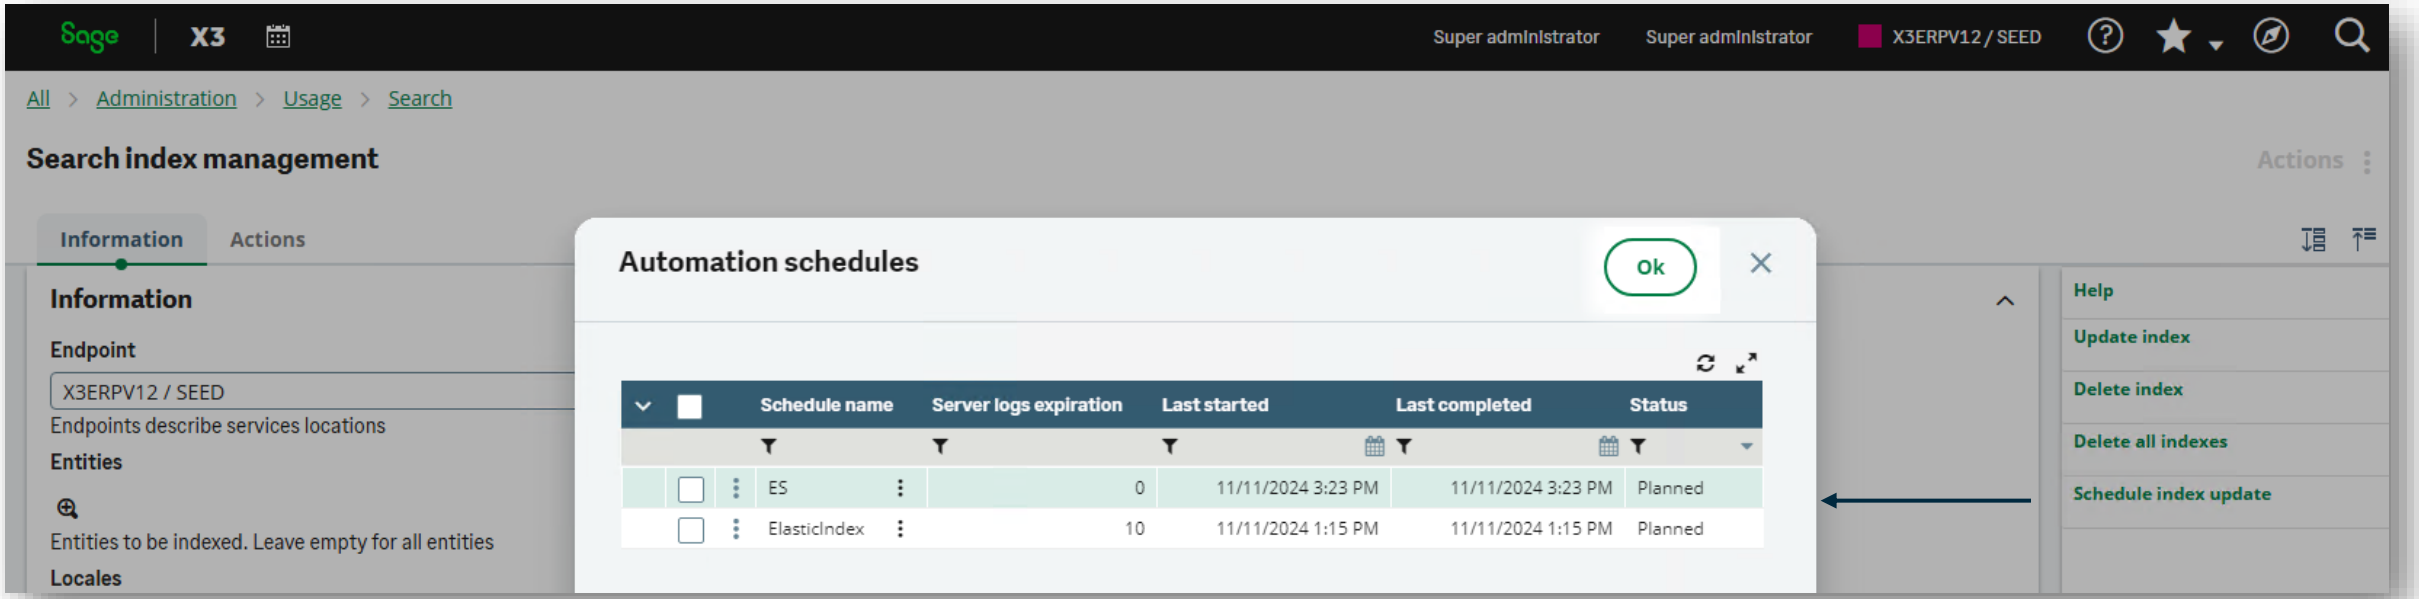

Automation Schedules

Automation is a planned process of automatically triggered events executed by the node.js server in the background. In the automation schedule you can set the required parameters.

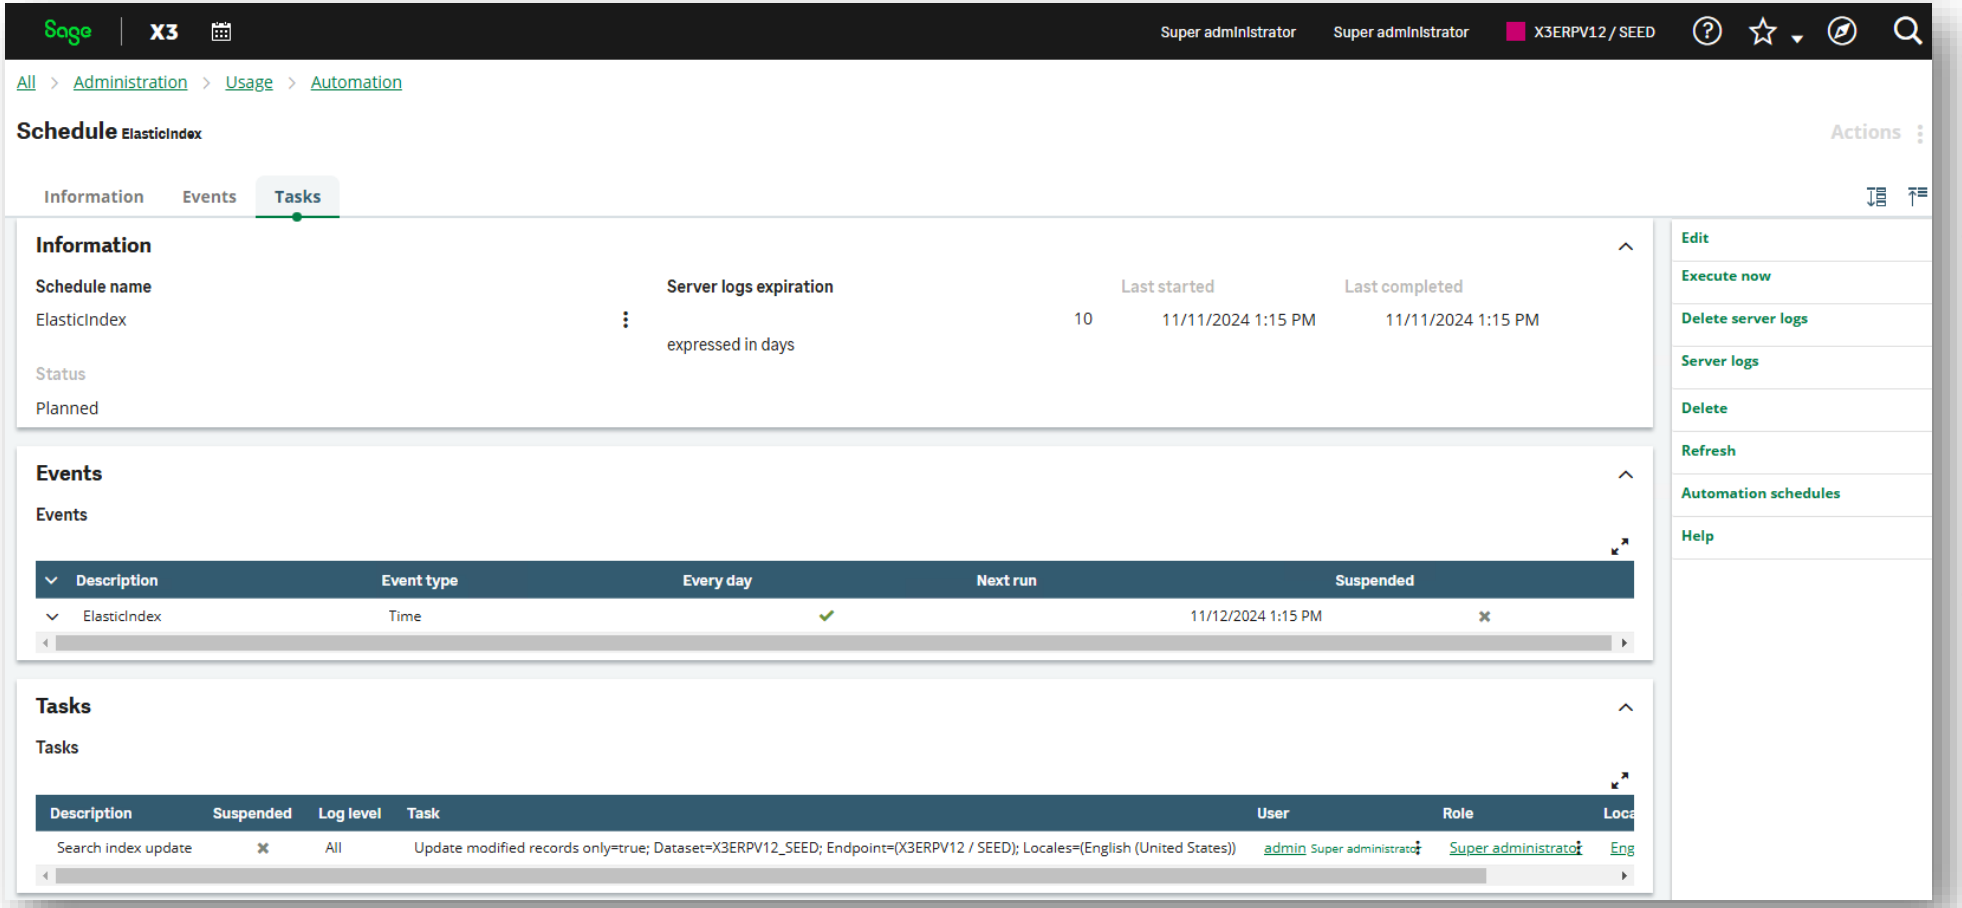

- Set the days flags by selecting every day or a specific day
- Set a list of times when the event will execute. A time picker is available, that gives all the hours, but any including minutes can be entered.
- Task lines cannot be entered in the task list, this is done from the function to be scheduled (index management)

Indexing

We can now test that the connection to the Elasticsearch server is successful by attempting to update index (create our index) The dialogue shows the connection status and the phase of the started job.

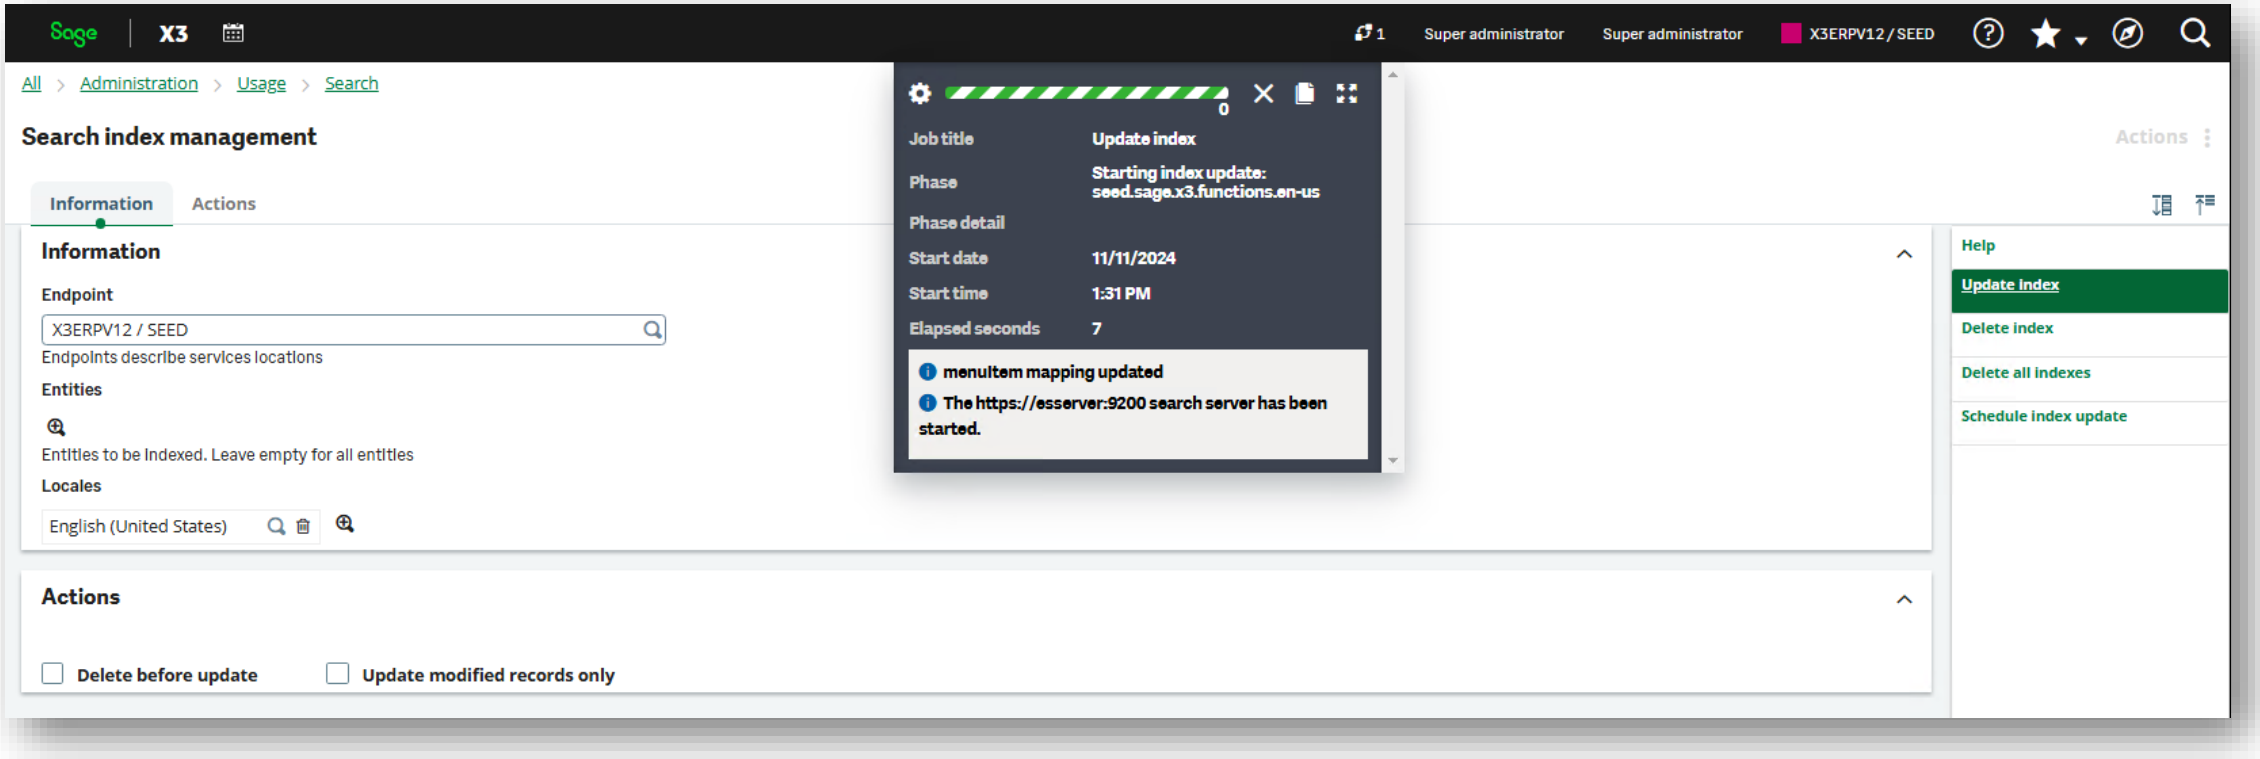

Searching Data

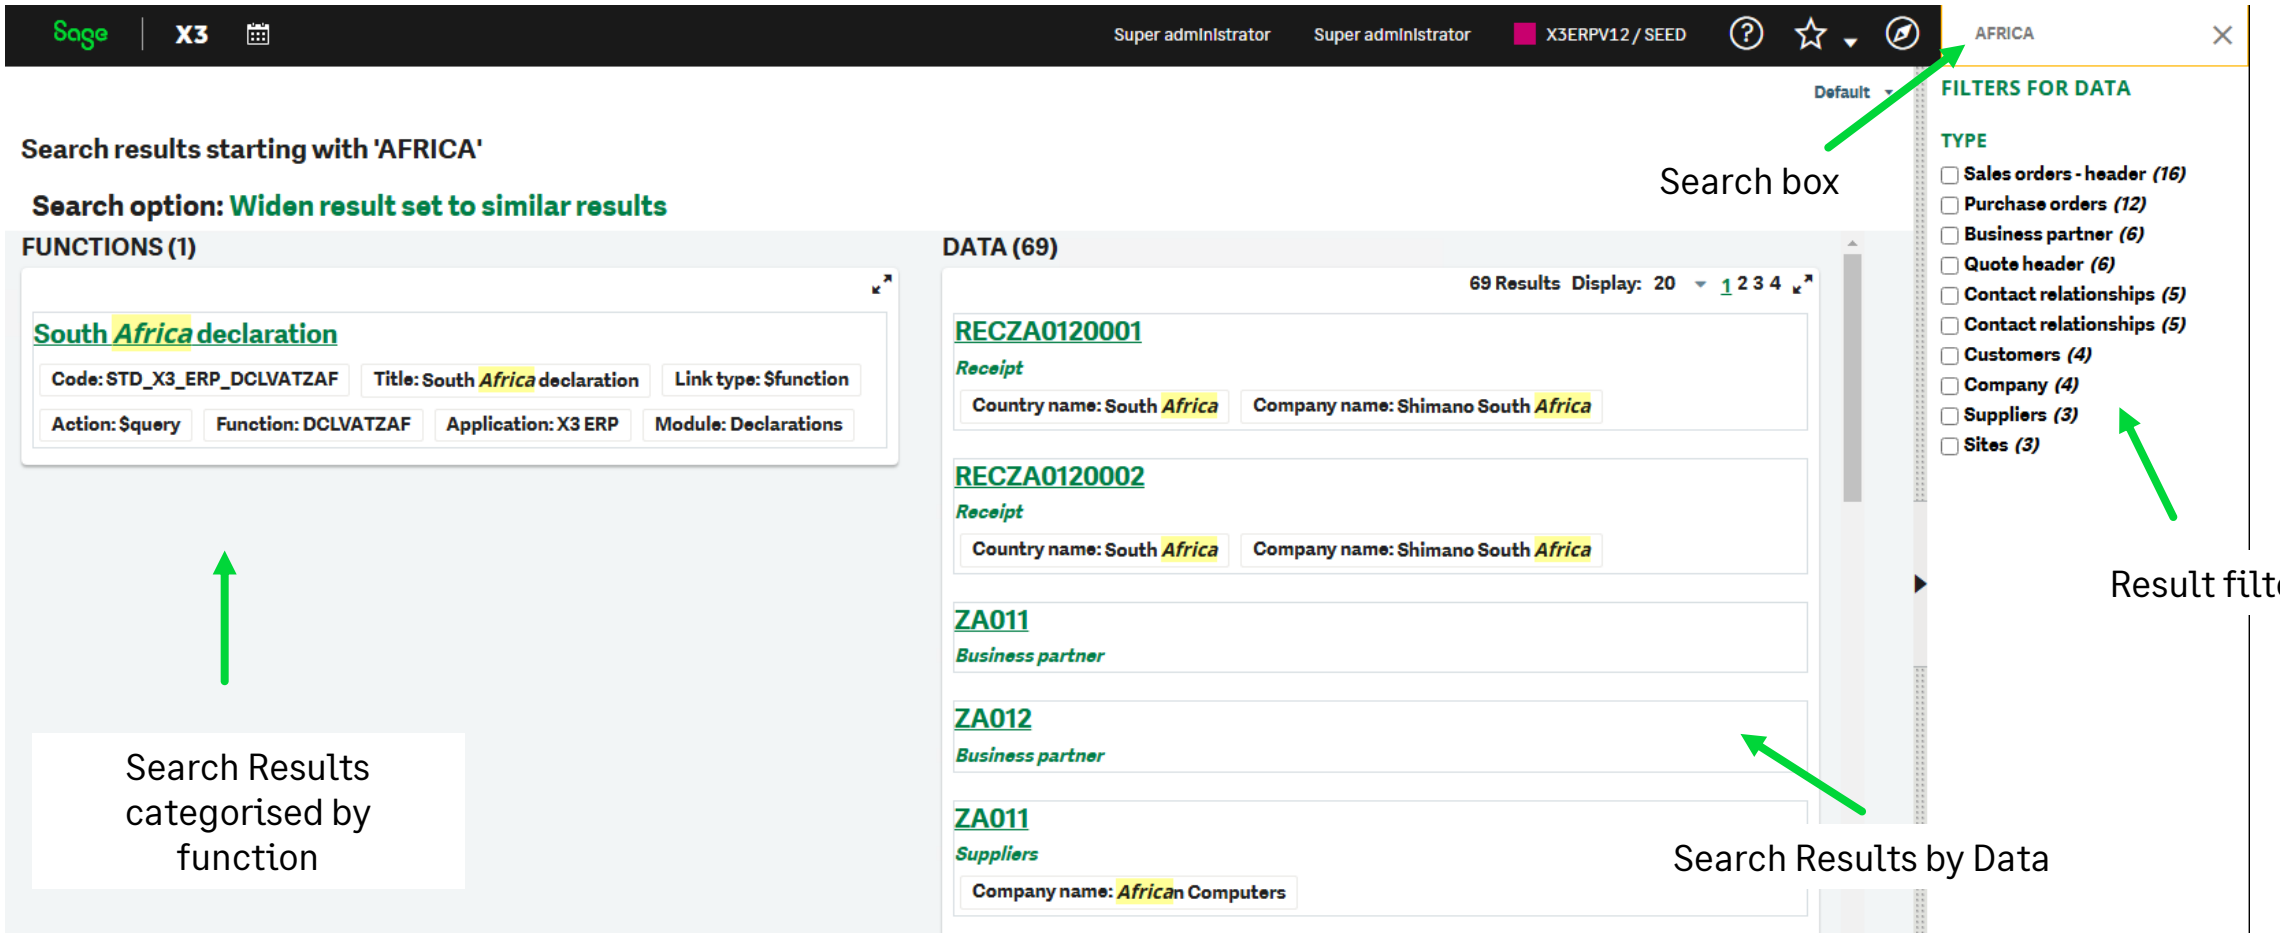

What does elastic search do anyway

Logging

Elastic search Logs default location is \elasticsearch-8.15.0\logs directory. The log file location can be modified in the in the elasticsearch.yml. Generally, the logging level is sufficient to troubleshoot but if you do need to increase the logging levels this can be done by modifying the \elasticsearch-8.15.0\config\log4j2.properties example

```
# Set a logger level for Elasticsearch indexing
logger.org.elasticsearch.index.level = debug
```

X3 Logs can help if you experience any connection issue or interruptions to the Elasticsearch server these can be enabled from Administration > Settings > Global settings.

The logs have five levels Error, which traces errors only. Debug, which is more verbose and returns more information. The logs Will be written to the standard Syracuse logs.

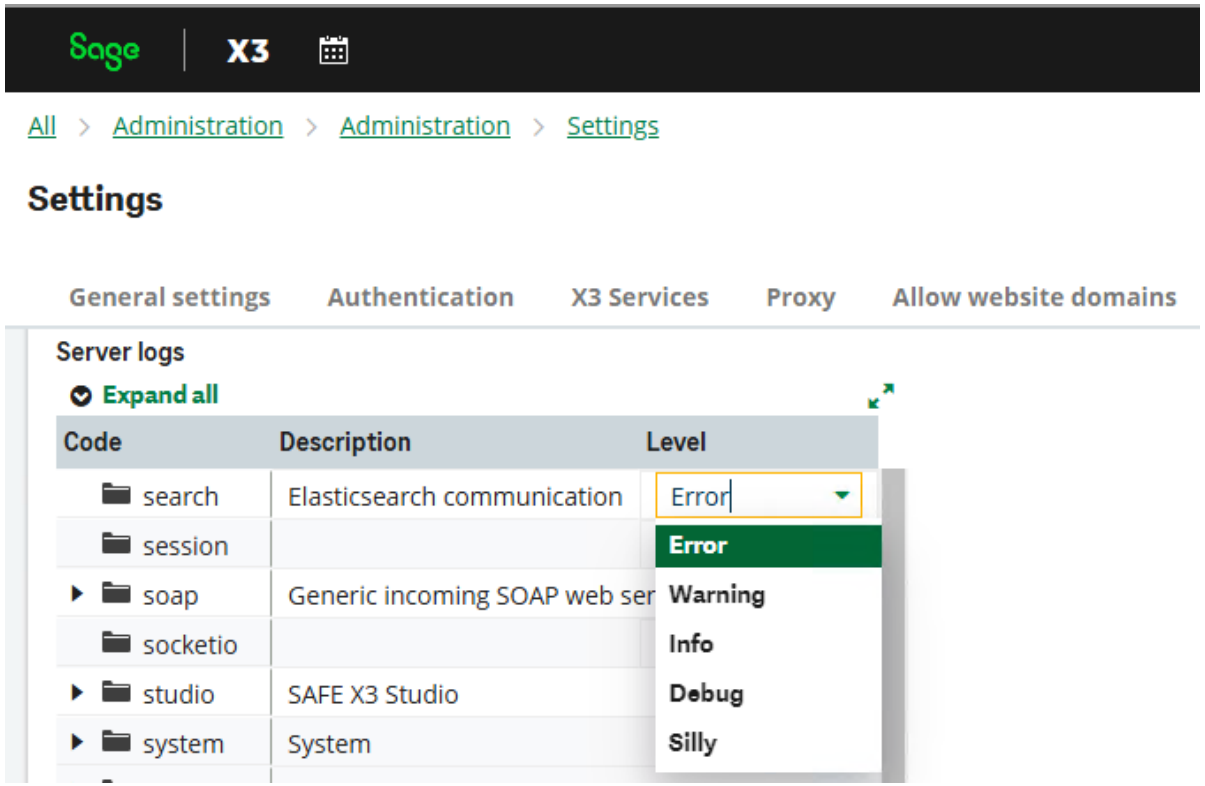

Live Environment Run-through

Appendices

Useful Links

Sage Resource	Link
Blog post	What does Elastic Search do anyway?
Blog post	Utilizing Elastic Cloud with Sage X3
Blog post	Performance tuning your Sage X3 system: Elastic Search
Blog post	Securing ElasticSearch 8.14 with HTTPS Certificates
Tips & Ticks 2021	Elasticsearch 7 and Sage X3

External Resources	Link
Elastic Help	Secure the Elastic Stack | Elasticsearch Guide [8.16] | Elastic
Elastic Help	Installing Elasticsearch | Elasticsearch Guide [8.16] | Elastic

Thank you!