Analysis Cockpit Elasticsearch Cluster Manual

Markus Meyer

Oct 08, 2023

CONTENTS

In this Manual we will describe how you can set up your Analysis Cockpit with an Elasticsearch Cluster. Please follow the sections thoroughly to get the desired result.

CHAPTER

ONE

INTRODUCTION

1.1 Analysis Cockpit Architecture

The ASGARD Analysis Cockpit uses an Elasticsearch database to store all event data. Each day worth of incoming events uses a single Elasticsearch index.

Normally, Elasticsearch is running locally on the Analysis Cockpit Server. However, when required Elasticsearch can easily be extended to become a cluster of almost arbitrary size.

When running in Cluster mode, the Analysis Cockpit runs the underlying metadata database and acts as the cluster master, while all data is stored on the additional nodes.

1.2 When to consider Clustering

You should consider extending the Elasticsearch installation to become a cluster if:

- there is significant performance degradation
  - for searches that cover multiple days and/or
  - for adding events to cases.
- performance cannot be sufficiently improved by adding more CPU cores or faster disks (RAM is supported up to 32GB)
- disk size of the analysis cockpit cannot be increased but retention period requires additional storage

1.3 Performance

Benchmarks suggest there is a communication overhead of 10% - 20% for a cluster compared to a single node in cases where a single node would be sufficient for the given load.

As logs of one day are stored in one index and indices are distributed over cluster members the performance gain will also depend on the number of days stored in the cluster.

In a cluster configuration the former Analysis Cockpit will act a master and will hold no data. Therefore, the minimum reasonable cluster size is three. In such a minimum configuration we expect a performance gain of 60% given we have at least 60 days of logs.

CHAPTER

TWO

ANALYSIS COCKPIT SETUP

2.1 Prerequisites

Elasticsearch Cluster setup requires:

- A fully functional installation of Analysis Cockpit version 3.4
- At least two additional nodes with a similar high-end spec
- High-performance low-latency networking between all nodes

2.2 Analysis Cockpit preparation

After installation, the Analysis Cockpit runs with a single local Elasticsearch instance as usual. To prepare it for use with a cluster, run `es-cluster-install.sh`:

```
nextron@cockpit3:~$ sudo /etc/nextron/analysiscockpit3/es-cluster-install.sh
```

The script will configure Elasticsearch in the following way:

- The Analysis Cockpit node continues to be the master node but data is automatically moved away from it once possible.
- SSL certificates are used for authentication of nodes.
- Any number of data nodes can be added with exactly the same configuration and certificate (as long as they are reachable).

2.3 Resulting Elasticsearch configuration

The Elasticsearch configuration can be found in `/etc/elasticsearch/elasticsearch.yml`. It will look like the following:

```
cluster.name: elasticsearch
cluster.routing.allocation.exclude._name: elastic-test-01.nextron
path.data: /var/lib/elasticsearch
path.logs: /var/log/elasticsearch
node.roles: [ master, data, ingest ]
http.host: _local:ipv4_
transport.host: _site:ipv4_
```

(continues on next page)

(continued from previous page)

```
discovery.seed_hosts: [ elastic-test-01.nextron ]
discovery.zen.minimum_master_nodes: 1
cluster.initial_master_nodes: [ elastic-test-01.nextron ]
xpack.security.transport.ssl.enabled: true
xpack.security.transport.ssl.verification_mode: certificate
xpack.security.transport.ssl.client_authentication: required
xpack.security.transport.ssl.keystore.path: elastic-certificates.p12
xpack.security.transport.ssl.truststore.path: elastic-certificates.p12
```

The configuration:

- designates the Analysis Cockpit node as the (only) cluster master.
- automatically moves existing data away from the Analysis Cockpit node, and distributes it across the other nodes.
- TLS security is enabled so that nodes authenticate by certificate.

2.4 Cluster Node configuration script

In addition to reconfiguring the Analysis Cockpit, `es-cluster-install.sh` will create a script es-node-install.sh that contains the required configuration for additional nodes to join the cluster.

2.5 Restarting Elasticsearch

Finally, restart elasticsearch so that it picks up the new configuration:

```
nextron@cockpit3:~$ sudo systemctl restart elasticsearch
```

CHAPTER

THREE

CLUSTER NODE SETUP

3.1 Prerequisites

The following prerequisites have to be given:

- Server must be suitable for the Nextron base image.
- All nodes must be able to reach each other by resolving the fully qualified host name.
- TCP port 9300 must be open between all nodes (Note: API port 9200 is only used locally).

3.2 Elasticsearch node installation

Install the server from the Nextron ISO base image as you normally would when installing the Analysis Cockpit itself, but **DO NOT** run the Nextron Installer.

Instead, copy `es-node-install.sh` to the new node and run it:

```
nextron@es-node1:~$ chmod +x es-node-install.sh
nextron@es-node1:~$ sudo ./es-node-install.sh
```

The script will automatically install Elasticsearch and configure the node to join the cluster with the Analysis Cockpit host as its master.

3.3 Resulting Elasticsearch configuration

The Elasticsearch configuration can be found in `/etc/elasticsearch/elasticsearch.yml`. It will look like the following:

```
cluster.name: elasticsearch
cluster.routing.allocation.exclude._name: elastic-test-01.nextron
path.data: /var/lib/elasticsearch
path.logs: /var/log/elasticsearch
node.roles: [ data, ingest ]
http.host: _local:ipv4_
transport.host: _site:ipv4_
discovery.seed_hosts: [ elastic-test-01.nextron ]
discovery.zen.minimum_master_nodes: 1
xpack.security.transport.ssl.enabled: true
```

(continues on next page)

(continued from previous page)

```
xpack.security.transport.ssl.verification_mode: certificate
xpack.security.transport.ssl.client_authentication: required
xpack.security.transport.ssl.keystore.path: elastic-certificates.p12
xpack.security.transport.ssl.truststore.path: elastic-certificates.p12
```

3.4 Enabling the node

After the installation, start elasticsearch and watch it becoming healthy:

```
nextron@es-node1:~$ sudo systemctl restart elasticsearch.service
nextron@es-node1:~$ watch curl http://127.0.0.1:9200/_cluster/health
```

The node should automatically join the cluster.

CHAPTER

FOUR

ELASTICSEARCH NODE MAINTENANCE

4.1 Performing Updates

When updates are applied to the Analysis Cockpit, you also need to update all additional cluster nodes by running:

```
nextron@es-node1:~$ sudo apt update
nextron@es-node1:~$ sudo apt upgrade
```

It is recommended that you update one node at a time, in particular when a reboot is required. It is not necessary to remove the node from the cluster for the update.

4.2 Checking Elasticsearch status

You can check elasticsearch status and index distribution on any of the nodes:

```
nextron@es-node1:~$ curl http://127.0.0.1:9200/_cat/health
nextron@es-node1:~$ curl http://127.0.0.1:9200/_cat/nodes
nextron@es-node1:~$ curl http://127.0.0.1:9200/_cat/shards
```

4.3 Removing Elasticsearch nodes

Before temporarily or permanently removing a node, you should reconfigure the cluster to move away any shards from that node.

You can tell Elasticsearch to remove all indexes from a node (change the placeholder value of “node_to_remove” to the actual node name):

```
nextron@es-node1:~$ curl -XPUT “http://127.0.0.1:9200/_cluster/settings” -d '{"transient
↪": {"cluster.routing.allocation.exclude._name": "node_to_remove"} }'
```

Then wait until the node has no shards left:

```
nextron@es-node1:~$ curl http://127.0.0.1:9200/_cat/shards
```

Once no shards are assigned to the node, it is safe to shut it down. When you have replicas of each index (number_of_replicas >= 1), the cluster should automatically cope with the removal of any single node. Refer to Elasticsearch documentation!

For obvious reasons, you must not remove the Analysis Cockpit node itself from the cluster but it is ok to shut it down or restart it for maintenance.

CHAPTER

FIVE

INDEX

- genindex

INDEX