



# Configuring an Elasticsearch Cluster



<https://pandorafms.com/manual!/current/>

Permanent link:

[https://pandorafms.com/manual!/current/en/documentation/pandorafms/technical\\_annexes/37\\_pfms\\_elasticsearch\\_cluster](https://pandorafms.com/manual!/current/en/documentation/pandorafms/technical_annexes/37_pfms_elasticsearch_cluster)

2024/06/10 14:36



# Configuring an Elasticsearch Cluster

We are working on the translation of the Pandora FMS documentation. Sorry for any inconvenience.

## Requirements

- First you must follow the steps of [installation and configuration in each node](#).
- The minimum size of an Elasticsearch server cluster is 3 nodes and it should always grow by odd numbers to make use of the quorum system and ensure data integrity.
- Make sure you have connectivity between all 3 nodes and that ports 9200 and 9300 are accessible between each and every node.

Remember to configure the firewall (firewall) of each node to allow connection through these port numbers.

## Settings

Stop the Elasticsearch service on each and every node:

```
systemctl stop elasticsearch.service
```

Modify in the configuration file `/etc/elasticsearch/elasticsearch.yml` the following lines:

```
#discovery.seed_hosts: ["host1", "host2"]  
#cluster.initial_master_nodes: ["host1", "host2"]
```

**Uncomment** the lines and add the IP addresses or URLs of each of the nodes:

```
discovery.seed_hosts: ["host1", "host2", "host3"]  
cluster.initial_master_nodes: ["host1", "host2", "host3"]
```

Example with IP addresses:

```
discovery.seed_hosts: ["172.42.42.101", "172.42.42.102", "172.42.42.103"]  
cluster.initial_master_nodes: ["172.42.42.101", "172.42.42.102",  
"172.42.42.103"]
```

Make sure the line `cluster.initial_master_nodes` is defined only once in the configuration file, in some cases the same line appears in two different blocks of the same

file.

Before starting the service, because the nodes started for the first time by themselves (standalone), the contents of the data folder (by default `/var/lib/elasticsearch/`) must be deleted. to be able to start the cluster for the first time. Do it with the command:

```
rm -rf /var/lib/elasticsearch/*
```

Now it's time to start the services on each and every node. Start and verify that they are running with the commands:

```
systemctl start elasticsearch.service && systemctl status elasticsearch.service
```

You should get an output similar to:

```
[root@rocky8-node1 ~]# systemctl status elasticsearch.service
● elasticsearch.service - Elasticsearch
   Loaded: loaded (/usr/lib/systemd/system/elasticsearch.service; disabled; vendor preset: disabled)
   Active: active (running) since Thu 2022-05-12 08:23:05 UTC; 49s ago
     Docs: https://www.elastic.co
   Main PID: 3334 (java)
    Tasks: 67 (limit: 11401)
   Memory: 1.3G
   CGroup: /system.slice/elasticsearch.service
           └─3334 /usr/share/elasticsearch/jdk/bin/java -Xshare:auto -Des.networkaddress.cache.ttl=60
             └─3619 /usr/share/elasticsearch/modules/x-pack-ml/platform/linux-x86_64/bin/controller

May 12 08:22:53 rocky8-node1 systemd[1]: Starting Elasticsearch...
May 12 08:23:05 rocky8-node1 systemd[1]: Started Elasticsearch.
lines 1-13/13 (END)
```

Once the services have started, you must confirm that the 3 nodes are joined to the cluster correctly, so when you run the following command on any of the nodes, it should give the same response:

```
curl -XGET http://127.0.0.1:9200/_cat/nodes
```

```
[root@rocky8-node1 ~]# curl -XGET http://127.0.0.1:9200/_cat/nodes
172.42.42.102 46 89 3 0.16 0.23 0.17 cdfhilmrstw - rocky8-node2
172.42.42.103 39 90 3 0.48 0.17 0.12 cdfhilmrstw * rocky8-node3
172.42.42.101 15 93 0 0.00 0.00 0.00 cdfhilmrstw - rocky8-node1
[root@rocky8-node1 ~]#
```

Recheck the firewall configuration, always taking into account that the nodes must communicate through ports 9200 and 9300 and that the port must be accessible from the PFMS server and the PFMS Web Console 9200 too. With these steps you will already have the Elasticsearch cluster

installed and configured to be used as the Pandora FMS logs storage engine.

## Data Models and Templates

Before putting an environment into production, whether it is a single node or a data cluster, it is recommended to apply the corresponding configurations to this node or cluster depending on its use. In the case of the indices generated by Pandora FMS, the most effective way to make them is by defining a template (template) to define the configuration of the fields and the stored data.

Templates or templates are settings that are only applied at index creation time. Changing a template will have no impact on already existing indices.

- To create a basic template follow the instructions in “[Data models and templates for a node](#)”:
- To define a template multinode you must take into account the following information:
  - When configuring the template (JSON format), you need to configure as many search as nodes you have, however to correctly configure the replicas you must subtract 1 from the number of nodes in the environment .

For example, in a Pandora FMS environment with Elasticsearch with 3 nodes configured, when you modify the fields `number_of_search` and `number_of_replicas` it should be as follows:

```
{
  "index_patterns": ["pandorafms*"],
  "settings": {
    "number_of_shards": 3,
    "auto_expand_replicas" : "0-1",
    "number_of_replicas" : "2"
  },
}
```

This is a very basic definition, in order to define coCorrectly sizing the Elasticsearch environment it is recommended to take into account the factors described in this article:

- <https://www.elastic.co/blog/how-many-shards-should-i-have-in-my-elasticsearch-cluster>

From the command line you can list the environment templates by running:

```
curl -X GET "localhost:9200/_cat/templates/*?v=true&s=name&pretty"
```

You can also see the details of a template, for example the one we have created for pandorafms by running:

```
curl -X GET "localhost:9200/_template/pandorafms*?pretty"
```

which will return in JSON format the configuration that it has defined.

You can perform these operations through the Elasticsearch interface in Pandora FMS using the native Elasticsearch commands.

- `PUT _template/<name_of_template> {json_data}`: allows you to enter the data of the template to be created.
- `GET _template/><name_of_template>`: allows you to view the template created.

Elasticsearch Interface **WARNING**

This is a view to interface with Elasticsearch directly from WEB console. Please note that you can damage your Elasticsearch if you don't know exactly what are you are doing. This view is intended to be used only by users with a knowledge of Elasticsearch .




## Query

1 GET \_template/pandorafms|

## Results

```
{
  "pandorafms": {
    "order": 0,
    "index_patterns": [
      "pandorafms*"
    ],
    "settings": {
      "index": {
        "number_of_shards": "1",
        "auto_expand_replicas": "0-1",
        "number_of_replicas": "0"
      }
    },
    "mappings": {
      "properties": {
        "agent_id": {
          "type": "long",
          "fields": {
            "keyword": {
              "ignore_above": 256,
              "type": "keyword"
            }
          }
        },
        "group_id": {
          "type": "long",
          "fields": {
            "keyword": {
              "ignore_above": 256,
              "type": "keyword"
            }
          }
        },
        "group_name": {
          "type": "text",
          "fields": {
            "keyword": {
              "ignore_above": 256,
              "type": "keyword"
            }
          }
        },
        "utimestamp": {
          "type": "long"
        },
        "source_id": {
          "type": "text",
          "fields": {
            "keyword": {
              "ignore_above": 256,
              "type": "keyword"
            }
          }
        },
        "suid": {
          "type": "text",
          "fields": {
            "keyword": {
              "ignore_above": 256,
              "type": "keyword"
            }
          }
        }
      }
    }
  }
}
```

Execute query 



[Back to Pandora FMS documentation index](#)