OpenSearch Installation and Configuration

OpenSearch Installation and Configuration

Online installation and requirements (recommended method) are specified in the topic “Pandora FMS Installation” in its corresponding section.

To configure Pandora FMS with OpenSearch, see “Log collection and monitoring” and “SIEM monitoring”.

Manual Installation and Advanced OpenSearch Configuration

Before running OpenSearch, memory paging and swap must be disabled on the host to improve performance and increase the number of memory maps available for OpenSearch. See “Important settings” for more information:

https://opensearch.org/docs/latest/install-and-configure/install-opensearch/index/#important-settings

```
# Disable memory paging and swapping.
sudo swapoff -a

# Edit the sysctl config file that defines the host's max map count.
sudo vi /etc/sysctl.conf

# Set max map count to the recommended value of 262144.
vm.max_map_count=262144

# Reload the kernel parameters.
sudo sysctl -p
```

For Rocky Linux 8, installation via RPM package is recommended. Once OpenSearch is installed, access to OpenSearch must be checked from Pandora FMS. Before performing this test, each node or cluster must be configured. To check the installation, run:

```
curl -X GET https://< ip_opensearch >:< opensearch_port > -u 'admin:admin' --insecure
```

You should get a response similar to:

```
{
   "name" : "hostname",
```

```
    "cluster_name" : "opensearch",
    "cluster_uuid" : "6XNc9m2gTUSIoKDqJit0PA",
    "version" : {
        "distribution" : "opensearch",
        "number" : <version>,
        "build_type" : <build-type>,
        "build_hash" : <build-hash>,
        "build_date" : <build-date>,
        "build_snapshot" : false,
        "lucene_version" : <lucene-version>,
        "minimum_wire_compatibility_version" : "7.10.0",
        "minimum_index_compatibility_version" : "7.0.0"
    },
    "tagline" : "The OpenSearch Project: https://opensearch.org/"
 }
```

By default, OpenSearch installation enables SSL, username, and password, which is a best practice; it is also recommended to change the default username and password.

Node Configuration

First, the configuration file `/etc/opensearch/opensearch.yml` must be edited, and then the OpenSearch service will be restarted.

This file contains the configuration for all OpenSearch service parameters; see the official documentation for more information:

https://opensearch.org/docs/latest/install-and-configure/configuration/

Minimum configurations required to start the service and use it with Pandora FMS.

- Port number.

```
# ---------------------------------- Network
# Set the bind address to a specific IP (IPv4 or IPv6):
network.host: 0.0.0.0
# Set a custom port for HTTP:
http.port: 9200
# For more information, consult the network module documentation.
```

- Location of stored data and logs:

```
# ---------------------------------- Paths
# Path to directory where to store the data (separate multiple locations by
comma):
path.data: /var/lib/opensearch
```

```
# Path to log files:
path.logs: /var/log/opensearch
```

It will also be necessary to *uncomment* and define the following lines:

```
cluster.name: pandorafms
node.name: ${HOSTNAME}
network.host: 0.0.0.0
```

- `cluster.name`: This will be the name assigned to the group or cluster.
- `node.name`: To name the node using the `${HOSTNAME}` system variable; it will automatically take the host name.
- For `network.host`, the value `0.0.0.0` allows OpenSearch to “listen” on all network interfaces (NICs); *to use a specific NIC, enter the corresponding specific value*.

If working with a single node, the following line must be added to the configuration file to allow *single node* startup:

```
discovery.type: single-node
```

In case of working with a cluster, you need to complete the `discovery.seed_hosts` parameter:

```
discover.seed_hosts : ["ip:port", "ip", "ip"]
```

In the most recent versions of OpenSearch, Java® virtual machine memory management is done automatically, and it is recommended to let it be managed this way in production environments, making it unnecessary to modify JVM values.

To start OpenSearch, run:

```
systemctl start opensearch.service
```

To restart use `restart`, to stop use `stop`, and use `status` to check the status.

If the service does not start, check the logs located in `/var/log/opensearch/` (in this case, the `pandorafms.log` file or the name given to the node).

Remember that to check the installation and operation of OpenSearch, you can run:

```
curl -X GET https://<node-ip> -u 'admin:admin' --insecure
```

OpenSearch Cluster Configuration

For OpenSearch cluster configuration, it is recommended to follow the official documentation:

https://opensearch.org/blog/optimize-opensearch-index-shard-size/

A configuration guide for Pandora FMS is also provided.

OpenSearch User Management

To change the default `admin` password, follow a series of steps. First, export the variable to use the Java® JDK installed by OpenSearch to use any of the tools:

```
export OPENSEARCH_JAVA_HOME=/usr/share/opensearch/jdk
```

Then, to generate the hashed password to be placed in the OpenSearch configuration file, use the following script (replace `< password >` with the password to be used):

```
/usr/share/opensearch/plugins/opensearch-security/tools/hash.sh -p <password>
```

For example:

```
[root@test ~]# /usr/share/opensearch/plugins/opensearch-security/tools/hash.sh -p pandora
************************************************************************
** This tool will be deprecated in the next major release of OpenSearch **
** https://github.com/opensearch-project/security/issues/1755          **
************************************************************************
$2y$12$aoOrXV/hLZ88gGrwobXuM.61K1HWmpLqXHiPQKwRmgEJDe5ncecn6
```

Then open the `/etc/opensearch/opensearch-security/internal_users.yml` file with the vim or nano text editor to modify the password for the required user or users.

It is recommended to leave only the `admin` user for use with Pandora FMS; it is unnecessary to maintain any other user.

Example file:

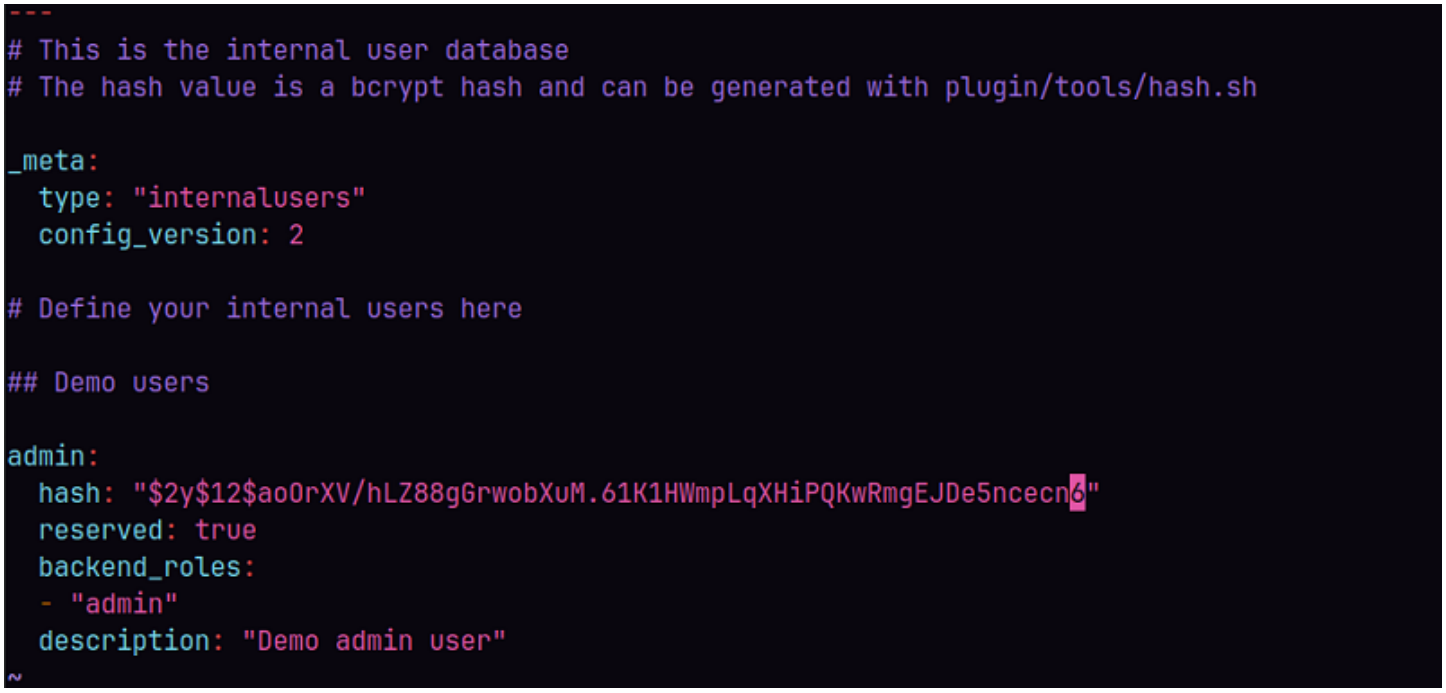

```
---
# This is the internal user database
# The hash value is a bcrypt hash and can be generated with plugin/tools/hash.sh

_meta:
  type: "internalusers"
  config_version: 2

# Define your internal users here

## Demo users

admin:
  hash: "$2y$12$aoOrXV/hLZ88gGrwobXuM.61K1HWmpLqXHiPQKwRmgEJDe5ncecn6"
  reserved: true
  backend_roles:
  - "admin"
  description: "Demo admin user"
~
```

To make the changes effective, run:

```
cd /usr/share/opensearch/plugins/opensearch-security/tools
```

```
OPENSEARCH_JAVA_HOME=/usr/share/opensearch/jdk ./securityadmin.sh -cd
/etc/opensearch/opensearch-security/ -cacert /etc/opensearch/root-ca.pem -cert
/etc/opensearch/kirk.pem -key /etc/opensearch/kirk-key.pem -icl -nhnv-t
internalusers   -icl   -nhnv   -cacert /etc/opensearch/root-ca.pem   -cert
/etc/opensearch/kirk.pem   -key /etc/opensearch/kirk-key.pem
```

A final `Done with success` message should be displayed; to check the new password (following the previous example with `pandora` used):

```
❯ curl https://10.235.50.104:9200 -ku 'admin:pandora'
{
  "name" : "node-1",
  "cluster_name" : "my-application",
  "cluster_uuid" : "3MDB9QFtS5OBPhK9AWn6Yg",
  "version" : {
    "distribution" : "opensearch",
    "number" : "2.11.0",
    "build_type" : "rpm",
    "build_hash" : "4dcad6dd1fd45b6bd91f041a041829c8687278fa",
    "build_date" : "2023-10-13T02:56:26.505314582Z",
    "build_snapshot" : false,
    "lucene_version" : "9.7.0",
    "minimum_wire_compatibility_version" : "7.10.0",
    "minimum_index_compatibility_version" : "7.0.0"
  },
  "tagline" : "The OpenSearch Project: https://opensearch.org/"
}
```

For more information on user management in OpenSearch:

- https://opensearch.org/docs/latest/security/configuration/yaml/
- https://opensearch.org/docs/latest/security/access-control/users-roles/#create-users

Pandora FMS Configuration with OpenSearch

To configure Pandora FMS with OpenSearch, see the topic “Log collection and monitoring”.

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 based on its usage. In the case of indexes generated by Pandora FMS, the most effective way to create them is by defining a template to define field settings and stored data.

Templates are configurations that only apply at the time of index creation. Changing a template will not have any impact on existing indexes.

To create a basic template, you only need to define the following fields:

```
curl -X PUT -ku 'admin:admin' https://<node_ip>:9200/_index_template/pandorafms
-H 'Content-Type: application/json' -d'
{
 "index_patterns": [
   "pandorafms*"
 ],
 "template": {
   "aliases": {
     "pandorafms_logs": {}
   },
   "settings": {
     "number_of_shards": 1,
     "auto_expand_replicas" : "0-1",
     "number_of_replicas": "0"
   },
"mappings" : {
    "properties" : {
      "agent_id" : {
        "type" : "long"
      },
      "group_id" : {
        "type" : "long"
      },
      "group_name" : {
        "type" : "text"
      },
      "logcontent" : {
        "type" : "text"
      },
      "source_id" : {
        "type" : "text"
      },
      "suid" : {
        "type" : "text"
      },
      "type" : {
        "type" : "text"
      },
      "utimestamp" : {
        "type" : "long"
      },
      "@timestamp": {
          "type": "date"
        }
    }
  }
}
}
'
```

Through the Pandora FMS (menu) interface, you will be able to upload said template:

- PUT _template/<template_name>: in this example `PUT _template/pandorafms` .

You can also check the templates through the same Pandora FMS interface:

- GET _template/<template_name>: in this example `GET _template/pandorafms` .

Multi-node Templates

To define a multi-node template, you must take the following information into account:

- When performing the template configuration (JSON format), you need to configure as many shards as you have nodes; however, to correctly configure replicas, you must subtract 1 from the number of nodes in the environment.

For example, in a Pandora FMS environment with 3 configured nodes, when you modify the `number_of_shards` and `number_of_replicas` fields, it should look like this:

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

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, created for `pandorafms`, by running:

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

which will return the defined configuration in JSON format.

You can perform these operations through the Pandora FMS interface:

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

To configure Pandora FMS with OpenSearch, see “Log collection and monitoring”.

SIEM Installation Requirements

To be able to use the Pandora FMS SIEM functionality, the following additional steps are required:

1. It is necessary to have at least one OpenSearch database on an extra dedicated instance (server).
2. It is necessary to have log collection enabled.
3. At a minimum, the OpenSearch server must have 4 cores, 8 GB of RAM, and 100 GB of high-speed SSD disk.
4. Recommended for higher performance: Deploy a dedicated instance (server) for the SIEM events server.
5. A license with the SIEM functionality enabled.

Additionally, SIEM Events and the SIEM server must be activated from the Console.

←Back to Pandora FMS documentation index