



Elastic Search Backup



Permanent link:
https://pandorafms.com/manual/!current/en/documentation/pandorafms/technical_annexes/16_elastic_search_backup
2024/06/10 14:36





Elastic Search Backup

Backup and restore of ElasticSearch (ELK)

Data migration from an ElasticSearch server using Snapshots is relatively quick. First, a backup of the server's data is made and then saved to a repository for later restoration.

Snapshot

The machine where the backup will be made will be called the “source machine” and the machine where the restoration will be made will be called the “target machine”.

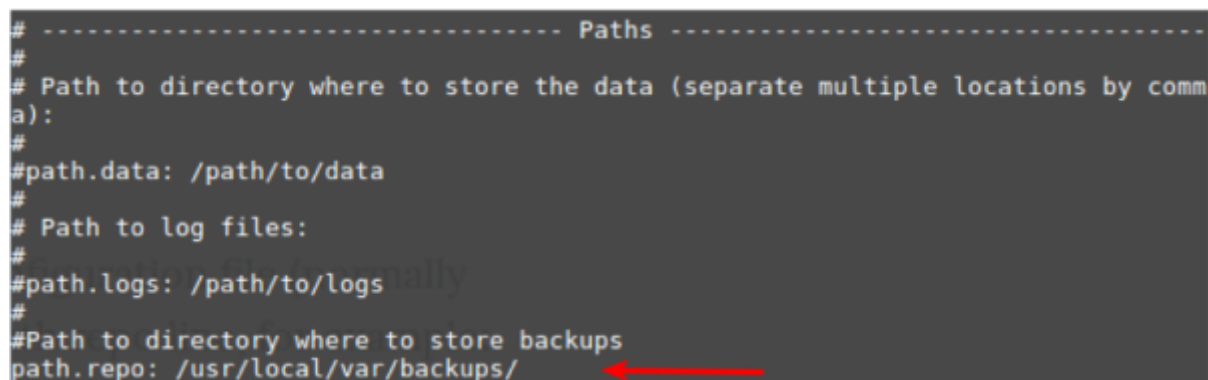
- In the origin machine

1) We modify the configuration file of “elasticsearch.yml”:

```
vi /etc/elasticsearch/elasticsearch.yml
```

And we add the following line:

```
path.repo: /usr/local/var/backups/
```



```
# ----- Paths -----  
# Path to directory where to store the data (separate multiple locations by comma):  
#  
#path.data: /path/to/data  
#  
# Path to log files:  
#  
#path.logs: /path/to/logs  
#  
#Path to directory where to store backups  
path.repo: /usr/local/var/backups/
```

2) We create the directory previously added to the configuration file:

```
mkdir -p /usr/local/var/backups/
```

3) We give read and write permissions to the directory and user:

```
chmod 700 /usr/local/var/backups  
chown elasticsearch:elasticsearch /usr/local/var/backups
```

4) We restart the service:

```
/etc/init.d/elasticsearch restart
```

5) Create the backup with the following command:

```
curl -XPUT http://localhost:9200/_snapshot/my_backup -d '{"type": "fs",  
"settings": {"compress": "true", "location": "/usr/local/var/backups/"}}'
```

6) We compress the previously generated backup:

```
cd /usr/local/var/  
tar -zcvf elastic_backup.tar.gz backups/
```

7) From the destination machine where we are going to make the restoration, we copy the compressed backup of the source machine.

- *On the target machine*

```
scp -P 41122 root@<ipOrigen>/root/elastic_backup.tar.gz /home/user/backup
```

To use the 'scp' command you must have an ssh server installed on the source machine and at least one ssh client on the target machine.

It is important that the version of Elasticsearch on the importing machine supports data export, i.e. in this case your local machine must have the same or higher version. If not, you must first update Elasticsearch.

Restore Backup

- *On the target machine*

1) We modified the configuration file of "elasticsearch.yml" in the same way we did when we created the backup in the first machine:

```
vi /etc/elasticsearch/elasticsearch.yml
```

And we add the following line:

```
path.repo: /usr/local/var/backups/
```

```
# ----- Paths -----  
#  
# Path to directory where to store the data (separate multiple locations by comma):  
#  
#path.data: /path/to/data  
#  
# Path to log files:  
#  
#path.logs: /path/to/logs  
#  
#Path to directory where to store backups  
path.repo: /usr/local/var/backups/
```

2) We create the directory added previously to the configuration file:

```
mkdir -p /usr/local/var/backups/
```

3) We give read and write permissions to the directory:

```
chmod 700 /usr/local/var/backups  
chown elasticsearch:elasticsearch /usr/local/var/backups
```

4) We restart the service:

```
/etc/init.d/elasticsearch restart
```

5) We decompress the backup that we import from the source machine:

```
tar -xzvf /home/user/backup/elastic_backup.tar.gz -C /usr/local/var/backups
```

6) We create the repositories where the snapshots are located:

```
curl -X PUT "localhost:9200/_snapshot/my_backup" -H 'Content-Type: application/json' -d'
```

```
{  
  "type": "fs",  
  "settings": {  
    "location": "/usr/local/var/backups"  
  }  
}
```

7) We close the indexes:

```
curl -XPOST http://localhost:9200/<indexname>-*/_close
```

The asterisk shows all the indexes that start with that name.

8) We import the backup:

First we copy the backup to the repository:

```
cp <name of the snapshot.dat> my_backup_location/
```

We renamed the file without capital letters:

```
mv my_backup_location/<name of snapshot.dat> my_backup_location/snap1
```

Finally it matters:

```
curl -X POST  
"localhost:9200/_snapshot/my_backup/snap1/_restore?wait_for_completion=true"
```

9) Finally, we reopen the indexes:

```
curl -XPOST http://localhost:9200/<indexname>-*/_open
```

[Go back to Pandora FMS documentation index](#)