Single Instance Deployment in Nagios Log Server 2024

Overview

Nagios Log Server is a clustered application, specifically it is the Elasticsearch database used by Nagios Log Server that provides the cluster functionality. The Elasticsearch database spreads the log data across the instances in the cluster to provide redundancy. If one of the instances in the cluster died then the log data on that instance is not lost.

The redundancy is provided by shards and replicas that are created for each index. This is explained in more detail in the Managing Nagios Log Server Indices documentation.

When you have a single instance cluster, the replica shards can never be put on another instance as such instance doesn't exist. The cluster will always have a Yellow health status.

The purpose of this documentation is to provide you with the steps to disable the replica shards in your cluster, this will result in your cluster having a Green health status. This documentation also provides the steps to re-enable the replica shards in the event that you wish to add another instance in the future.

Nagios Enterprises strongly recommends that you have at least two instances in your cluster. A single instance cluster provides no redundancy and this can lead to a total loss of your log data.

Cluster Yellow Status

This first screenshot demonstrates a single instance cluster that has a Yellow health status. To view the cluster status navigate to Admin > System > Cluster Status.

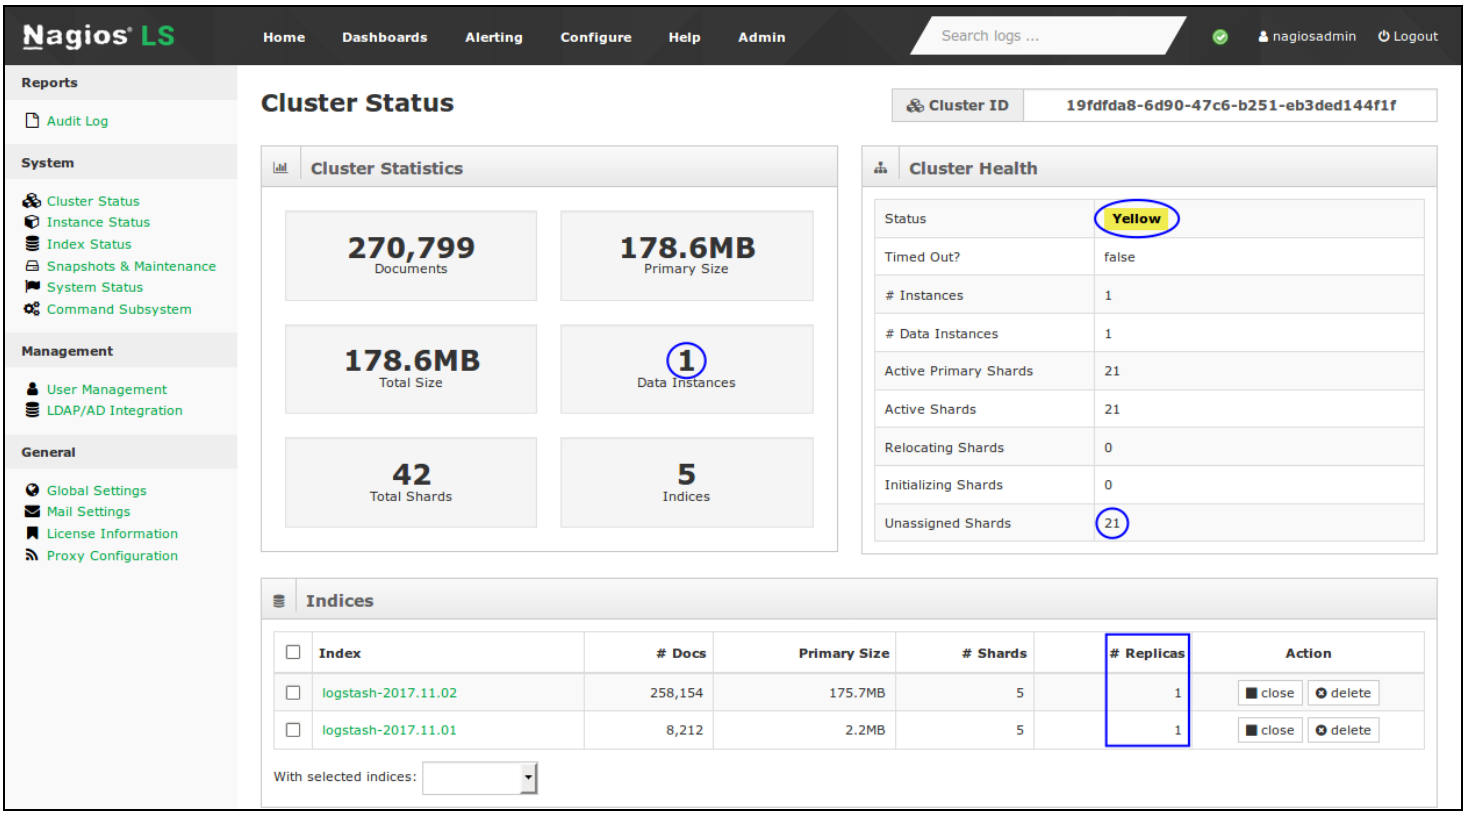

In the Indices table you can see that each index has 1 replicas. In the Cluster Health table you can see there are 21 Unassigned Shards and the cluster has a Yellow health status. It is these unassigned shards waiting for replication that is causing the cluster to remain with a Yellow health status.

Disable Replicas

The following steps will disable replicas being created for every new index and will also disable the replicas on the existing indices.

If you were to add another instance to your cluster after following these steps you will need to follow the steps in the Enable Replicas section of this document.

Nagios Enterprises strongly recommends that you have at least two instances in your cluster. A single instance cluster provides no redundancy and this can lead to a total loss of your log data.

You will need to establish a terminal session to your Nagios Log Server as the root user to disable replicas.

The first step is to disable replicas on newly created indices. You will need to edit the elasticsearch.yml file using the vi editor with this command:

```
vi /usr/local/nagioslogserver/elasticsearch/config/elasticsearch.yml
```

When using the vi editor, to make changes press i on the keyboard first to enter insert mode. Press Esc to exit insert mode.

Find this line:

index.number_of_replicas: 1

You will need to remove the # from the beginning of the line and set the value to 0:

index.number_of_replicas: 0

When you have finished, save the changes in vi by typing:

```
:wq
```

and press Enter.

The next step is to restart the Elasticsearch database using one of these commands:

RHEL | CentOS | Oracle Linux | Debian | Ubuntu

```
systemctl restart elasticsearch.service
```

The last step is to disable replicas on all existing indices using the following command:

```
curl -XPUT 'http://localhost:9200/_settings' -d '{ "number_of_replicas" :
"0" }'
```

This will produce output of:

{"acknowledged":true}

Now that these changes have been made, return to the Nagios Log Server web interface and navigate to Admin > System > Cluster Status.

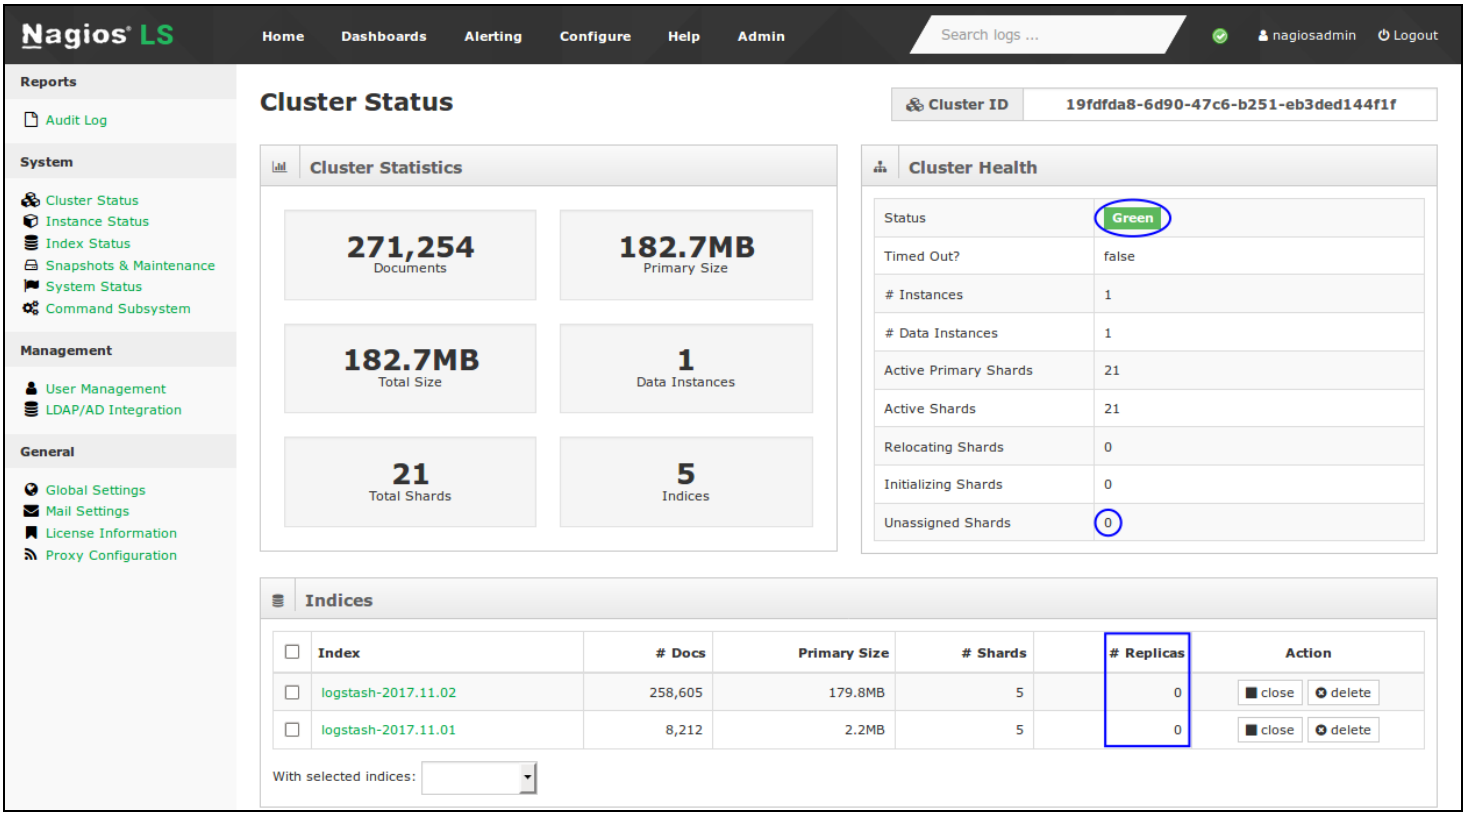

In the Indices table you can see that each index has 0 replicas. In the Cluster Health table you can see there are 0 Unassigned Shards and the cluster has a Green health status.

Your cluster is now configured so that it does not create replicas and hence will have a Green health status.

If you were to add another instance to your cluster after following these steps you will need to follow the steps in the Enable Replicas section of this document.

Nagios Enterprises strongly recommends that you have at least two instances in your cluster. A single instance cluster provides no redundancy and this can lead to a total loss of your log data.

Enable Replicas

If you previously disabled replicas on your Nagios Log Server cluster then you will need to re-enable them if you decided to add another instance to your cluster.

You will need to establish a terminal session to your Nagios Log Server as the root user to enable replicas.

The first step is to enable replicas on newly created indices. You will need to edit the elasticsearch.yml file using the vi editor with this command:

```
vi /usr/local/nagioslogserver/elasticsearch/config/elasticsearch.yml
```

When using the vi editor, to make changes press i on the keyboard first to enter insert mode. Press Esc to exit insert mode.

Find this line:

index.number_of_replicas: 0

You will need to set the value to 1:

index.number_of_replicas: 1

When you have finished, save the changes in vi by typing:

```
:wq
```

and press Enter.

The next step is to restart the Elasticsearch database using one of these commands:

RHEL | CentOS | Oracle Linux | Debian | Ubuntu

```
systemctl restart elasticsearch.service
```

The last step is to enable replicas on all existing indices using the following command:

```
curl -XPUT 'http://localhost:9200/_settings' -d '{ "number_of_replicas" :
"1" }'
```

This will produce output of:

{"acknowledged":true}

Now that these changes have been made, return to the Nagios Log Server web interface and navigate to Admin > System > Cluster Status. In the Indices table you can see that each index has 1 replicas.

You will need to wait while the unassigned shards are replicated to your other instance(s) in the cluster. Depending on the size of the indices it may take a while to finish replicating. Once completed the cluster should have a Green health status.