

Installing a SonarQube cluster

Running SonarQube as a Cluster is only possible with a [Data Center Edition](#).

The Data Center Edition allows SonarQube to run in a clustered configuration to make it resilient to failures. For more details about the Data Center (DC) Edition, please see this [page](#).

Requirements

App Servers

You need five servers dedicated to SonarQube. Servers can be VMs, it's not necessary to have physical machines.

Servers must be co-located (geographical redundancy is not supported).

You can find the operating system requirements for servers in the [Requirements](#) page.

DB Server

The Data Center Edition supports PostgreSQL, Oracle, and Microsoft SQL Server. If your data is currently stored in MySQL you can use the [Sonar DB Copy Tool](#) to move it.

Load Balancer

In addition to the five SonarQube servers, you must configure a reverse proxy / load balancer to load balance traffic between the two application nodes. The precise configuration of the reverse proxy / load balancer will vary by vendor, but the SonarQube DC Edition requirements are very simple:

- Share requests (load) between the two application nodes configured in the SonarQube cluster
- If you are using HTTPS, ensure you are meeting the requirements set out in [Securing SonarQube Behind a Proxy](#)
- **Note:** there is **no** requirement for the load balancer to preserve sessions; this is handled by the in-built JWT mechanism

License

You need a dedicated license to activate the DC Edition. If you don't have it yet, please contact the SonarSource Sales Team.

Support

Don't start this journey alone; as a Data Center Edition customer SonarSource will assist with the setup and configuration of your cluster. Get in touch with [SonarSource Support](#) for help.

Cluster Topology

There are two types of nodes:

- an **application** node responsible for handling web requests from users (**WebServer** process) and handling analysis reports (**ComputeEngine** process)
- a **search** node that is an Elasticsearch process that will store indices of data

In order to achieve high availability, the **only supported configuration** for the DC Edition comprises 5 servers:

- 2 **application** nodes containing both WebServer and ComputeEngine
- 3 **search** nodes that host Elasticsearch. For performance reasons, SSD are significantly better than HDD for these nodes

With this configuration, a node can be lost without impacting the service. More precisely, one application node and one search node can be lost without impacting users.

Here is a schema for the supported topology:

Machines

Here are the type of machines we used to perform our validation with a 200M Issues database. This could be used as a minimum recommendation to build your cluster.

- Search Node made of [Amazon EC2 m4.2xlarge](#): 8 vCPUs, 32GB RAM - 16GB allocated to Elasticsearch
- App Node made of [Amazon EC2 m4.xlarge](#): 4 vCPUs, 16GB RAM

Installation Details

Prepare your personalized SonarQube package:

1. Install SonarQube as you would for a normal installation following [Installing the Server](#)
2. Once you can confirm the connectivity with the DB is working well
3. Install all the other plugins you may want. If you do this step manually be sure to check compatibility with your SonarQube version using the [Plugin Version Matrix](#)
4. Now, it is time to install the DC edition:
 - a. Manually: stop SonarQube, and expand the DC Edition zip into `SONARQUBE_HOME/extensions/plugins`
 - b. Automatically: use the Marketplace to automatically install the DC Edition into your server, restart to properly deploy the plugins in the server, then stop SonarQube again
5. Double-check there are no duplicated or outdated plugin versions in `extensions/plugins` - keep only the latest versions
6. Zip the directory `SONARQUBE_HOME`
7. You now have your own personalized SonarQube Data Center Edition package

Deploy your SonarQube package on the 4 other nodes:

It is expected that all **application** nodes are identical in term of hardware and software (same JVM build). Similarly, all **search** nodes should be identical to each other. But application and search nodes can differ. Generally, **search** nodes are configured with more CPU and RAM than **application** nodes.

1. Unzip your personalized SonarQube package from the previous step on the 4 others nodes
2. Congratulations, you now have 2 applications nodes and 3 search nodes completely identical in term of SonarQube softwares

Configuration

Now you have 5 machines with the same SonarQube software it's time to start configuring them to specialize them. Some will become **application** nodes, some **search** nodes.

The idea is as follows: you need to edit the `sonar.properties` file on each node to configure the node's specialization. In the following example we will assume:

- The VMs (server1, server2) having IP addresses ip1 and ip2 are going to be **application** nodes
- The VMs having IP addresses ip3, ip4 and ip5 (server3, server4 and server5) are going to be **search** nodes

The default configuration to be applied for each node is the following:

server1

```
sonar.cluster.enabled=true
sonar.cluster.hosts=ip1,ip2,ip3,ip4,ip5
sonar.cluster.search.hosts=ip3,ip4,ip5
sonar.cluster.node.type=application
sonar.auth.jwtBase64Hs256Secret=Vf4TRrfS6tvsFIHsQlgfhDUkiw3r8=
```

server2

```
sonar.cluster.enabled=true
sonar.cluster.hosts=ip1,ip2,ip3,ip4,ip5
sonar.cluster.search.hosts=ip3,ip4,ip5
sonar.cluster.node.type=application
sonar.auth.jwtBase64Hs256Secret=Vf4TRrfS6tvsFIHsQlgfhDUkiw3r8=
```

server3

```
sonar.cluster.enabled=true
sonar.cluster.hosts=ip1,ip2,ip3,ip4,ip5
sonar.cluster.search.hosts=ip3,ip4,ip5
sonar.cluster.node.type=search
sonar.search.host=ip3
```

server4

```
sonar.cluster.enabled=true
sonar.cluster.hosts=ip1,ip2,ip3,ip4,ip5
sonar.cluster.search.hosts=ip3,ip4,ip5
sonar.cluster.node.type=search
sonar.search.host=ip4
```

server5

```
sonar.cluster.enabled=true
sonar.cluster.hosts=ip1,ip2,ip3,ip4,ip5
sonar.cluster.search.hosts=ip3,ip4,ip5
sonar.cluster.node.type=search
sonar.search.host=ip5
```

The full set of cluster parameters is listed [here](#).

Once this configuration is done, take a break and a coffee, then you can [Operate your Cluster](#).