



This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

SIP Feature Server Deployment Guide

Appendix: Add new Datacenter when running Feature Server with Embedded Cassandra

Contents

- 1 Appendix: Add new Datacenter when running Feature Server with Embedded Cassandra
 - 1.1 Prerequisites
 - 1.2 Using PropertyFileSnitch
 - 1.3 Using GossipingPropertyFileSnitch

Appendix: Add new Datacenter when running Feature Server with Embedded Cassandra

If your Cassandra Cluster of Feature Servers *already meets the prerequisites*, then follow the steps to add a new data center to the cluster.

Prerequisites

Important

This is not a to-do list; these actions must have been completed before (or when) enabling your Cassandra Cluster. You cannot do them "now".

- All Feature Servers must be installed in Standalone mode with embedded Cassandra cluster.
- The Cassandra cluster must be using Network Topology Strategy.
The `replicationStrategyClassName` option in the `TServer/Cassandra` section of the master, Feature Server must have been set to `NetworkTopologyStrategy` before the schema was created (before the first/initial start of the 'master' Feature Server).
- The Cassandra cluster must be using the `PropertyFileSnitch` type or the `GossipingPropertyFileSnitch` type of the endpoint snitch.
You define this option in the `cassandra.yaml` configuration file(s) of all the Feature Servers in the cluster. See [the Cassandra documentation](#) for details.

Do this only when the new data center is created:

- Set one of the new data center's Feature Server nodes to `confSync`.

Using PropertyFileSnitch

cassandra-topology.properties file configuration

```
# Cassandra Node IP=Data Center:Rack
#Data Center One
999.999.99.99=us_west:RAC1
999.999.99.98=us_west:RAC2
#Data Center Two
999.999.99.97=us_east:RAC1
999.999.99.96=us_east:RAC2
#Data Center Three
999.999.99.95=eu_east:RAC1
999.999.99.94=eu_east:RAC2
#Data Center Four
999.999.99.93=eu_west:RAC1
999.999.99.92=eu_west:RAC2
# default for unknown nodes
default=us_west:RAC1
```

1. Deploy new Feature Server nodes for the new data center. [Use this process](#) but ignore the instructions for single data centers.
2. Update Cassandra topology in the `cassandra-topology.properties` file on every Feature Server within the Cassandra Cluster that includes the new nodes.
[+] Read instructions:

In a Cassandra cluster, each Feature Server is a node, and each has an identical `cassandra-topology.properties` file that describes the network topology. When you created the Cassandra Cluster, you created this file and placed a duplicate in the `/resources/` directory of each Feature Server deployment.

You must now update each of these files to enable the new data centers.

The example below began as a single-node configuration. It defined only Data Center One. To configure a new data center, you simply add those same defining lines—modified to contain the correct data. The example below **(modifications in red)** now defines four data centers.

```
# Cassandra Node IP=Data Center:Rack
# Data Center One
999.999.99.99=us_west:RAC1
999.999.99.98=us_west:RAC2
<span style="color: red;"># Data Center Two
999.999.99.97=us_east:RAC1
999.999.99.96=us_east:RAC2
# Data Center Three
999.999.99.95=eu_east:RAC1
999.999.99.94=eu_east:RAC2
# Data Center Four
999.999.99.93=eu_west:RAC1
```

```
999.999.99.92=eu_west:RAC2 </span>
# default for unknown nodes
default=us_west:RAC1
```

-To finish your configuration, restart the Feature Servers in existing data center(s).

Read the Cassandra documentation [for additional details](#).

3. Restart the existing Feature Server nodes.
4. Change the value of the endpoint_snitch type to PropertyFileSnitch in newly deployed Feature Servers, then start the new Feature Server nodes [using these instructions](#).
5. Change the Cassandra cluster's keyspace replication options to accommodate the added new data center: specify the replication factor. Be certain that you understand [the Cassandra Keyspace Properties documentation](#) before undertaking this action.
6. The replication options change in the previous step automatically triggers the rebuild process for each new Feature Server/node.
This process:
 - Runs automatically on all new nodes that you configured in step 2.
 - Streams data for the new nodes from existing nodes.

Using GossipingPropertyFileSnitch

cassandra-rackdc-properties file configuration for node1 in new datacenter (DC2)

```
# dc=dc name
# rack=rack name
dc=DC2
rack=RAC1
```

cassandra-rackdc-properties file configuration for node2 in new datacenter (DC2)

```
# dc=dc name
# rack=rack name
dc=DC2
rack=RAC2
```

1. Deploy new Feature Server nodes for the new data center.
2. Add `cassandra-rackdc.properties` files to the new Feature Server nodes with data center and rack information specific to that node as mentioned in the example below.
Example: The figure depicts `cassandra-rackdc.properties` files for two Feature Server nodes in a new data center (DC2) to be placed in the Feature Server nodes.
3. Configure `endpoint_snitch` type to `GossipingPropertyFileSnitch` in the **`cassandra.yaml`** file on the newly deployed Feature Server nodes.
4. Modify the **`replicationOptions`** option to accommodate the newly added data center in the **[Cassandra]** section of the Feature Server application. Be certain that you understand [the Cassandra Keyspace Properties documentation](#) before undertaking this action.

Important

When adding a new node to an existing Cassandra cluster:

- Add only one node at a time.
- New nodes must be in the same snitch as the existing nodes snitch. Genesys does not recommend using mixed snitch mode clusters.