



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 Administration Guide

Appendix: Backing up and restoring embedded Cassandra data

Appendix: Backing up and restoring embedded Cassandra data

You can back up your Cassandra storage folders and use the saved storage folders to restore your Cassandra data if needed.

Important

Feature Server is capable of restoring its data by replicating it from other nodes in the Cassandra cluster. The restoration procedure described below is an exceptional point-in-time recovery measure. As with any backup procedure, **do not** make restoration a part of your regular maintenance process.

Backing up Cassandra data

Backing up your Cassandra storage folder requires you to:

- flush the Feature Server Cassandra keyspace on all Feature Server nodes, then
- save the Cassandra **storage** folder of each node

Tip

On the Linux platform, you can flush all the nodes in a datacenter by running the flush command using a parallel ssh utility such as pssh.

For each Feature Server instance:

1. Change to the Feature Server deployment location.
2. Verify the vms-port parameter in **launcher.xml**. The default value is 8080.
3. Change to the **lib** directory:
 - For Linux:
`cd work/jetty-0.0.0.0-8080-fs.war-_fs-any-/webapp/WEB-INF/lib`
 - For Windows:
`cd work\jetty-0.0.0.0-8080-fs.war-_fs-any-\webapp\WEB-INF\lib`
4. Verify the JMX port parameter in launcher.xml. The default value is 9192.
5. While Feature Server is running, run the nodetool flush command:

- For Linux:

```
java -cp libthrift-0.7.0.jar:cassandra-thrift-1.1.12.jar:commons-cli-1.1.jar:cassandra-all-1.1.12.jar org.apache.cassandra.tools.NodeCmd -h localhost -p 9192 flush sipfs
```



```
java -cp libthrift-0.7.0.jar:cassandra-thrift-1.1.12.jar:commons-cli-1.1.jar:cassandra-all-1.1.12.jar org.apache.cassandra.tools.NodeCmd -h localhost -p 9192 flush system
```
- For Windows:

```
java -cp libthrift-0.7.0.jar;cassandra-all-1.1.12.jar;cassandra-thrift-1.1.12.jar;commons-cli-1.1.jar org.apache.cassandra.tools.NodeCmd -h localhost -p 9192 flush sipfs
```



```
java -cp libthrift-0.7.0.jar;cassandra-all-1.1.12.jar;cassandra-thrift-1.1.12.jar;commons-cli-1.1.jar org.apache.cassandra.tools.NodeCmd -h localhost -p 9192 flush system
```

6. Copy the Cassandra **storage** folder to a safe location.

Restoring Cassandra data

Important

Use this recovery procedure only when absolutely necessary. Do not make it a part of your regular maintenance process.

To restore your Cassandra data from your backed-up storage folders:

1. **Stop all Feature Server instances.**
2. Copy the **storage** folder backups to their original location on each Feature Server.
3. **Restart each Feature Server instance**, beginning with the master instance.