



# Identifying key files and folders

DataStax Academy:

Installing and configuring Cassandra

Leo Schuman

# Continue your Free Self-Paced Apache Cassandra Training at DataStax Academy

- Instructional Videos
- Learning Exercises
- Tutorials
- Demos

**<http://academy.datastax.com>**

# Can Cassandra be installed by package?

- Package installations define various folders for
  - data directories
  - log files
  - configuration files
  - binaries
- Package types include
  - RPM on \*nix using *yum*
  - DEB on \*nix using *apt-get*
  - MSI on Windows

<http://datastax.com/download>

<http://planetcassandra.org/cassar>

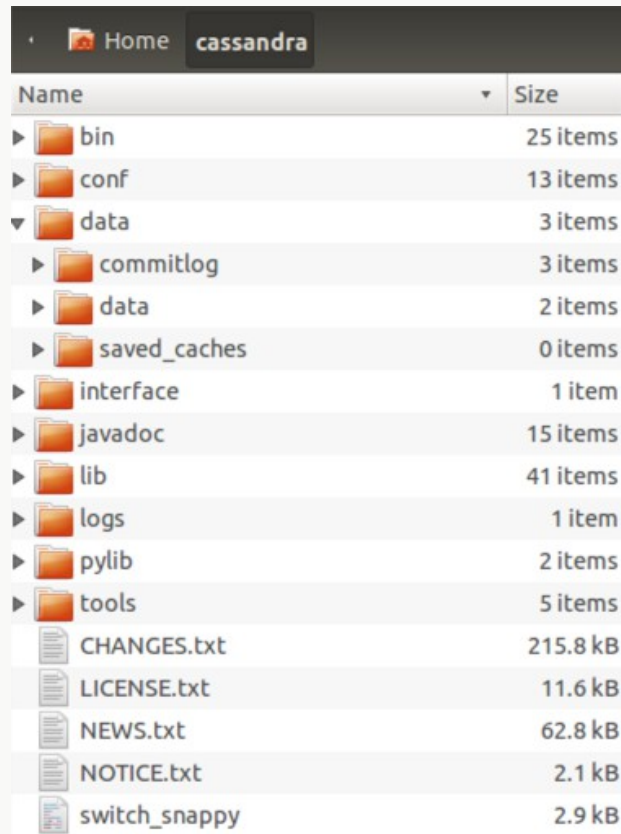
| Configuration Files                        | Locations                         |
|--|-----------------------------------|
| <code>cassandra.yaml</code>                | <code>/etc/cassandra</code>       |
| <code>cassandra-topology.properties</code> | <code>/etc/cassandra</code>       |
| <code>cassandra-rackdc.properties</code>   | <code>/etc/cassandra</code>       |
| <code>cassandra-env.sh</code>              | <code>/etc/cassandra</code>       |
| <code>cassandra.in.sh</code>               | <code>/usr/share/cassandra</code> |

| Directories                           | Description            |
|---------------------------------------|------------------------|
| <code>/var/lib/cassandra</code>       | Data directories       |
| <code>/var/log/cassandra</code>       | Log directory          |
| <code>/var/run/cassandra</code>       | Runtime files          |
| <code>/usr/share/cassandra</code>     | Environment settings   |
| <code>/usr/share/cassandra/lib</code> | JAR files              |
| <code>/usr/bin</code>                 | Optional utilities     |
| <code>/usr/bin</code>                 | Binary files           |
| <code>/usr/sbin</code>                |                        |
| <code>/etc/cassandra</code>           | Configuration files    |
| <code>/etc/init.d</code>              | Service startup script |
| <code>/etc/security/limits.d</code>   | Cassandra user limits  |
| <code>/etc/default</code>             |                        |

# What is in the tarball folder?

- Tarball installations create these folders in a single location

- **bin**: executables (*cassandra, cqlsh, nodetool*, etc)
- **conf**: Configuration files (*cassandra.yaml*)
- **data**: commit log, caches, and data files
  - ▶ **commitlog**: 3 items
  - ▶ **data**: 2 items
  - ▶ **saved\_caches**: 0 items
- **interface**: thrift protocol support
- **javadoc**: C\* source documentation
- **lib**: Library dependencies (jars)
- **pylib**: Python libraries (e.g. for *cqlsh*, written in Python)
- **tools**: Additional tools (e.g. *cassandra-stress*)

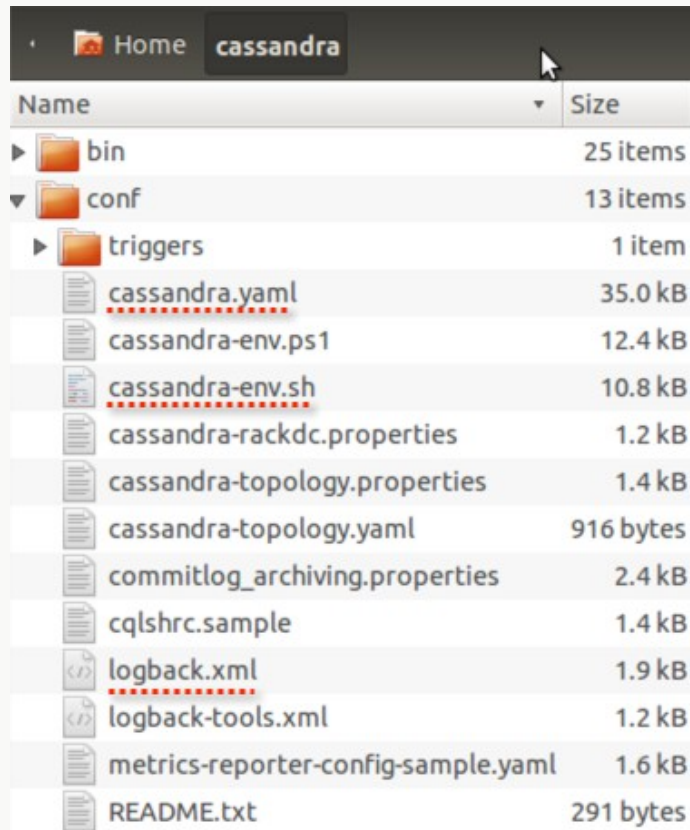


| Name           | Size     |
|----------------|----------|
| ▶ bin          | 25 items |
| ▶ conf         | 13 items |
| ▼ data         | 3 items  |
| ▶ commitlog    | 3 items  |
| ▶ data         | 2 items  |
| ▶ saved_caches | 0 items  |
| ▶ interface    | 1 item   |
| ▶ javadoc      | 15 items |
| ▶ lib          | 41 items |
| ▶ logs         | 1 item   |
| ▶ pylib        | 2 items  |
| ▶ tools        | 5 items  |
| CHANGES.txt    | 215.8 kB |
| LICENSE.txt    | 11.6 kB  |
| NEWS.txt       | 62.8 kB  |
| NOTICE.txt     | 2.1 kB   |
| switch_snappy  | 2.9 kB   |

# What configuration files may be used?

- Configuration files include

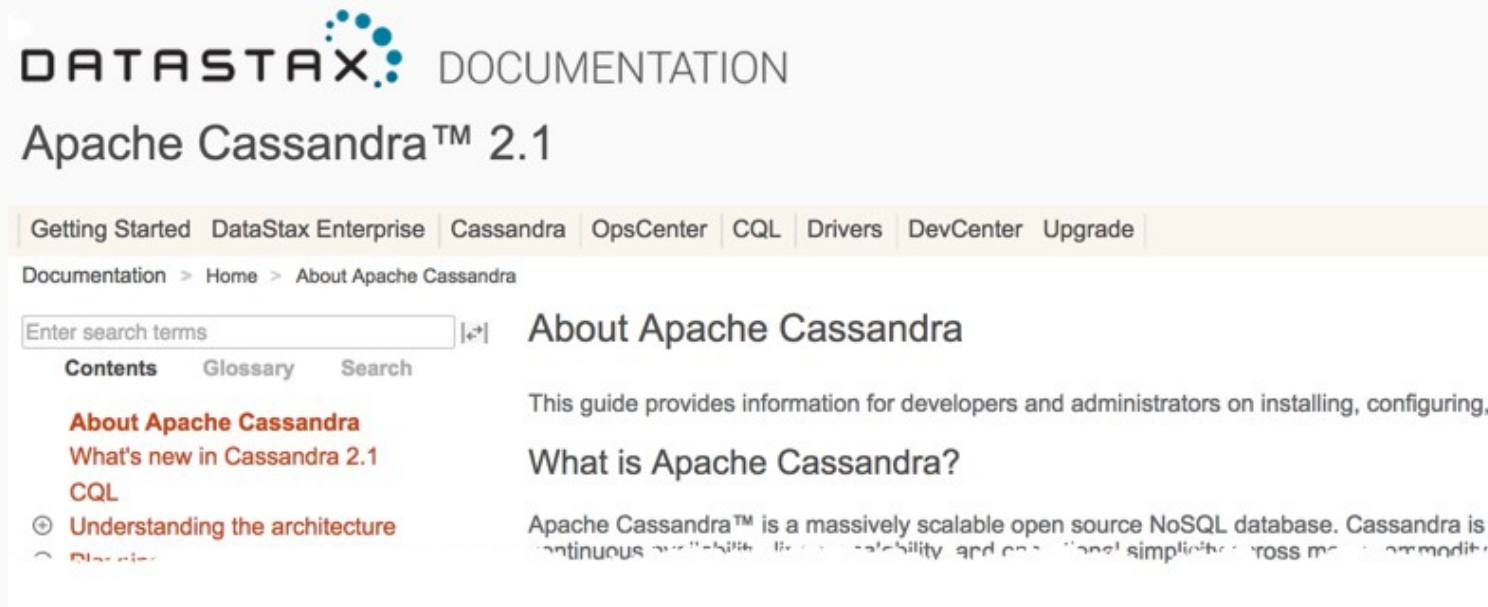
- `cassandra.yaml`: primary config file for each instance
- `cassandra-env.sh`: Java environment config (e.g. `MAX_HEAP_SIZE`, etc.)
- `logback.xml`: log config
- `cassandra-rackdc.properties`: config to set the Rack and Data Center to which this Node belongs
- `cassandra-topology.properties`: config IP addressing for Racks and Data Centers in this cluster
- `bin/cassandra-in.sh`: `JAVA_HOME`, `CASSANDRA_CONF`, `CLASSPATH`



| Name                                | Size      |
|-------------------------------------|-----------|
| bin                                 | 25 items  |
| conf                                | 13 items  |
| triggers                            | 1 item    |
| <u>cassandra.yaml</u>               | 35.0 kB   |
| cassandra-env.ps1                   | 12.4 kB   |
| <u>cassandra-env.sh</u>             | 10.8 kB   |
| cassandra-rackdc.properties         | 1.2 kB    |
| cassandra-topology.properties       | 1.4 kB    |
| cassandra-topology.yaml             | 916 bytes |
| commitlog_archiving.properties      | 2.4 kB    |
| cqlshrc.sample                      | 1.4 kB    |
| <u>logback.xml</u>                  | 1.9 kB    |
| logback-tools.xml                   | 1.2 kB    |
| metrics-reporter-config-sample.yaml | 1.6 kB    |
| README.txt                          | 291 bytes |

# Where do I find more information?

- Detailed documentation is just a click away ...  
[datastax.com/documentation/cassandra/2.1/](http://datastax.com/documentation/cassandra/2.1/)



The screenshot shows the DataStax documentation page for Apache Cassandra 2.1. The page features the DataStax logo and the title "DOCUMENTATION Apache Cassandra™ 2.1". A navigation bar includes links for "Getting Started", "DataStax Enterprise", "Cassandra", "OpsCenter", "CQL", "Drivers", "DevCenter", and "Upgrade". Below the navigation bar, there is a breadcrumb trail: "Documentation > Home > About Apache Cassandra". A search bar is present with the placeholder text "Enter search terms". The main content area is titled "About Apache Cassandra" and contains a paragraph: "This guide provides information for developers and administrators on installing, configuring, and using Apache Cassandra." Below this, there is a section titled "What is Apache Cassandra?" with a paragraph: "Apache Cassandra™ is a massively scalable open source NoSQL database. Cassandra is designed for continuous availability, linear scalability, and operational simplicity across multiple commodity hardware nodes." On the left side, there is a table of contents with links for "About Apache Cassandra", "What's new in Cassandra 2.1", "CQL", "Understanding the architecture", and "Planning".

# Demo: Setting up a learning environment

# Continue your Free Self-Paced Apache Cassandra Training at DataStax Academy

- Instructional Videos
- Learning Exercises
- Tutorials
- Demos

**<http://academy.datastax.com>**