

Changes and New Features in Cassandra 2.1

Monday, October 13, 2014 20:58

URL To find details:

<http://www.datastax.com/documentation/cassandra/2.1/cassandra/features2.html>

Changes

Some of the noteworthy changes include:

- Cassandra 2.1 does not support pre-Cassandra 2.0 SSTables.
To upgrade to Cassandra 2.1 from a previous release that stored data in Cassandra 1.2.x SSTables, start the node on Cassandra 2.0 and use the [sstableupgrade](#) tool after upgrading. Upgrade SSTables even if you do not perform a rolling upgrade. Resolve schema disagreements if any exist, and restart each node. For more upgrade information, see "[Upgrading Cassandra](#)" and [NEWS.txt](#).
- The shuffle utility for migrating to virtual nodes (vnodes) and the nodetool taketoken command have been removed. To migrate to vnodes, [bootstrap a new data center](#).
- Cassandra 2.1 bundles and enables JNA. If JNA fails to initialize, you can disable JNA by using the `-Dcassandra.boot_without_jna=true` option to start Cassandra.
- Cassandra rejects USING TIMESTAMP or USING TTL in the command to update a counter column, and now generates an error message when you attempt such an operation.
- Configurable properties have been added to manage counter writes.
- A configurable counter cache reduces lock contention and helps with concurrency.
- In Cassandra 2.1, the CQL table property `index_interval` is replaced by `min_index_interval` and `max_index_interval`. The `max_index_interval` is 2048 by default. The default would be reached only when SSTables are infrequently-read and the index summary memory pool is full. When upgrading from earlier releases, Cassandra uses the old `index_interval` value for the `min_index_interval`.
- CASSANDRA-6504 has been backported to Cassandra 2.0.5 so you can perform a rolling upgrade of a database having counters to Cassandra 2.1.
- Default data and log locations have changed for tarball installations and source checkouts. By default, the data file directory, commitlog directory, and saved caches directory are in `$CASSANDRA_HOME/data/data`, `$CASSANDRA_HOME/data/commitlog`, and `$CASSANDRA_HOME/data/saved_caches`, respectively. The log directory now defaults to `$CASSANDRA_HOME/logs`. If not set, `$CASSANDRA_HOME`, defaults to the top-level directory of the installation. Deb and RPM packages continue to use `/var/lib/cassandra` and `/var/log/cassandra` by default.
- Cassandra 2.1 maintains data consistency during bootstrapping. As you [bootstrap](#) a new node, Cassandra streams the data for the new node from an existing node that is free from [range movement](#). If data inconsistency issues are present in the cluster, the improvement to bootstrapping handles these issues. Data inconsistency commonly occurs after frequent data deletions and a node going down.
- To inhibit the new Cassandra 2.1 bootstrapping behavior, and make Cassandra 2.0 behavior effective, start the node using the [Dconsistent.rangemovement=false](#) property:
- Package installations: Add the following option to `/usr/share/cassandra/cassandra-env.sh` file:
`JVM_OPTS="$JVM_OPTS -Dconsistent.rangemovement=false`
- Tarball installations: Start Cassandra with this option:
`$ bin/cassandra -Dconsistent.rangemovement=false`
To replace a dead node, you also need to [specify the address](#) of the node from which Cassandra streams the data.

For a complete list of fixes and new features, see the Apache Cassandra 2.1.0 [CHANGES.txt](#).

New Features

- User-defined types:

Cassandra 2.1 supports user-defined types. A user-defined type facilitates handling multiple fields of related information in a table. Applications that required multiple tables can be simplified to use fewer tables by using a user-defined type to represent the related fields of information instead of storing the information in a separate table. The address type example demonstrates how to use a user-defined type. You can create, alter, and drop a user-defined type using these commands:

- CREATE TYPE: http://www.datastax.com/documentation/cql/3.1/cql/cql_reference/cqlRefCreateType.html
- ALTER TYPE: http://www.datastax.com/documentation/cql/3.1/cql/cql_reference/cqlRefAlterType.html
- DROP TYPE: <http://www.datastax.com/documentation/cql/3.1/cql/ddl/ddlIndexColl.html>

The `cqlsh` utility includes these commands for describing a user-defined type or listing all user-defined types:

- DESCRIBE TYPE: http://www.datastax.com/documentation/cql/3.1/cql/cql_reference/describe_r.html#reference_ds_vyl_gns_xj_describe_type
- DESCRIBE TYPES: http://www.datastax.com/documentation/cql/3.1/cql/cql_reference/describe_r.html#reference_ds_vyl_gns_xj_describe_types

The scope of a user-defined type is the keyspace in which you define it. Use dot notation to access a type from a keyspace outside its scope: keyspace name followed by a period followed the name of the type. Cassandra accesses the type in the specified keyspace, but does not change the current keyspace; otherwise, if you do not specify a keyspace, Cassandra accesses the type within the current keyspace.

- Collection indexes: **In Cassandra 2.1 and later, you can index collections and query the database to find a collection containing a particular value.**
- Better implementation of counters that makes them safer, simpler, and typically faster
- New listsnapshots and reloadtriggers nodetool commands
- Improved metrics reporting through the use of the metrics-core library

Performance improvements:

- Faster reads and writes than previous releases: <http://www.datastax.com/dev/blog/cassandra-2-1-now-over-50-faster>
- Improved row cache
- Reduced heap used by memtables
- New counters implementation

Compaction and repair improvements:

- Post-compaction read performance
- A configurable percentage of cold SSTables can be ignored
- Incremental node repair

Other notable changes:

- Improved Hadoop support
- Unique table IDs
- Improved logging using logback
- New configuration options for allocating and managing memtable memory
- Improvements to bootstrapping a node that ensure data consistency
- Bundled JNA
- A number of other CQL and `cqlsh` changes