

Apache Cassandra 4.0 Prerelease Testing

Goal: Discuss testing required in support of qualifying Apache Cassandra 4.0.

This document will be migrated to Confluence after live collaboration.

Reference:

<https://cwiki.apache.org/confluence/display/CASSANDRA/4.0+Quality%3A+Components+and+Test+Plans>

Testing Methodologies

1. Replay Testing
 - a. Execution by several contributors who operate large clusters.
 - b. Diff testing
2. In-JVM Dtest
3. Cluster Fuzz
4. Targeting Fuzzing (Integration)
 - a. Have all the targeted fuzz test tooling executed as a pre-release task prior to all new builds being published.
 - b. Integration with Circle CI.
 - c. Ability to launch for execution on any platform (launch script) to enable at-scale execution on elastic compute infra.
5. Circle CI runs on commit and prior to release
6. Extra unit tests
7. Performance testing
8. Netflix Fuzz testing
 - a. Writes data with checksum, later reads out to validate.

Testing by Component

1. Internode
 - a. Significant progress: Completed as part of validation of CASSANDRA-15066, with exhaustive fuzz testing runs and load testing
 - b. Pending: large messages, combinations of consistency level, backpressure.
<https://issues.apache.org/jira/browse/CASSANDRA-14746> (pre-beta dependency).
 - i. Volunteer: Joey Lynch
2. Local Read/Write Path
 - a. Replay Testing

- i. Volunteer: Yifan Cai
- b. Upgrade / Diff Testing (cassandra-diff)
 - i. Volunteer: Yifan Cai
- c. Audit IndexInfo changes in 3.7 (source review)
 - i. Volunteer: Jordan West
- d. Cluster Fuzz
 - i. Volunteer: Benedict Elliott Smith, Alex Petrov
- e. Targeted Fuzz
 - i. Memtable vs sstables (similar to Blake Eggleston's work during winter 2018)
 - 1. Volunteer: Alex Petrov, Blake Eggleston
 - ii. Compaction
 - 1. Volunteer: Marcus Eriksson, Jordan West
 - 2. JBOD / CASSANDRA-6696
- f. Audit File Transaction log
 - i. Volunteer: Blake Eggleston
- g. Unit tests
 - i. Dropping tombstones
 - 1. Volunteer: Aleksey Yeshchenko, if does not exist already.
- h. Chunk cache
 - i. Volunteer: TBD
- 3. Commit Log
 - a. Audit / add unit tests
 - i. Volunteer:
 - b. Drain verification (in-JVM dtest)
 - c. Coverage for CDC (methodology TBD)
 - i. Volunteer: Sumanth Pasupuleti
- 4. Cluster Membership (goal: no regression from 3.0.x)
 - a. Replay Testing
 - i. Expansion and decommission
 - 1. Volunteer: Yifan Cai
 - b. Targeted Fuzzing
 - i. TokenMetaData
 - 1. Volunteer: Jon Meredith
 - c. Gossip
 - i. Performance / stability
 - 1. Volunteer: Joey/Vinay
- 5. Distributed Read/Write Path: Coordination, Replication, and Read Repair, Short read protection, speculative reads
 - a. Cluster Fuzz:
 - i. Volunteer: Benedict Elliott Smith, Alex Petrov
 - b. In-JVM Dtests (convert SRP dtests to in-JVM dtests)

- i. Volunteer: Alex Petrov, Jon Meredith
 - c. Replay testing (covers implicitly)
 - i. Volunteer: Yifan Cai
 - d. LWT Fuzzer (similar to Ariel Weisberg's work during spring 2019)
 - i. Volunteer: Alex Petrov
- 6. Repair
 - a. Cluster Fuzz
 - i. Volunteer: Alex Petrov, Benedict Elliott Smith
 - b. Tooling / integration: Reaper, Priam
 - i. Priam Volunteer: Joey Lynch (Priam)
 - ii. Repair service (ref:14346): Vinay Chella
 - iii. Reaper Volunteer: TBD (Jon Haddad?)
 - c. Preview Repair (as a method of asserting correctness of repair)
 - d. Off-heap merkle trees: preview repair, cassandra-diff test @ CL:ONE
 - i. Volunteer: Aleksey Yeshchenko
 - e. Incremental Repair
 - i. Volunteer: Blake Eggleston
- 7. Client Protocol V5
 - a. Targeted Fuzz - follow methodology employed in internode testing for CASSANDRA-15066
 - i. Serialization, CAS batch overflow
 - 1. Volunteer: Sam Tunnicliffe, Aleksey Yeshchenko
 - b. Backpressure (availability / performance)
 - i. Volunteer: Sumanth + Joey (client and server)
- 8. Secondary Index
 - a. Targeted Fuzz (volunteer: Jordan)
 - i. Index selection / query planning
 - ii. Consistency with indexed table
 - b. Performance / stability
 - i. Volunteer TBD
 - c. Audit and add unit tests, consider conversion to in-JVM dtests
 - i. Volunteer TBD
- 9. Streaming
 - a. Audit (Volunteer: Aleksey Yeshchenko)
 - b. Bootstrap
 - i. Volunteer: Dinesh Joshi
 - c. Cluster fuzz and/or replay test + "ship of Theseus" whole-cluster replacement
 - i. Yifan Cai
- 10. Java 11 Support
 - a. Runtime stability / performance
 - i. Volunteer: Joey Lynch
 - b. Mixed-version diff test (Java 8 + Java 11)
 - i. Volunteer: Yifan Cai

- c. In-tree tooling (all commands)
 - i. Compare output between j8 and j11 (they should be the same)
 - ii. Volunteer: Dinesh Joshi
- 11. Mixed-Version Operation
 - a. Multiple replay test with ongoing upgrade operation
 - i. Volunteer: Yifan Cai
 - b. Upgrade / diff test
 - i. Volunteer: Yifan Cai, Vinay Chella
- 12. Operational
 - a. Metrics / naming (updates to NEWS.txt)
 - i. Volunteer: Dinesh Joshi
 - b. Audit log
 - i. Volunteer: Vinay Chella
 - c. Not testing, but pre-release task: ensuring backwards-compatibility of interface change made to SeedProviders.
 - i. Volunteer: Joey
 - d. Starting / Stopping clusters
 - i. Volunteering: Joey/Vinay(CASSANDRA-15181)
- 13. System tables
 - a. Tracing bug (table exists on 4.0 side, but not 3.0 side). @Aleksey to file.