

Apache Cassandra Technology Test Plan

Document Identifier: ICF_CTP
Document Version: 1.0
Issue Date: 24 APR 2019

Table of contents

1	Instaclustr Certification Framework	3
1.1	Description	3
1.2	Document Identification Information.....	3
1.3	Copyright Statement	3
2	Introduction	4
2.1	Overview	4
2.2	Purpose.....	4
2.3	Scope	4
3	Functional Testing.....	5
3.1	Overview	5
3.2	Unit Tests	5
3.3	Distributed Tests	5
4	Performance Testing	6
4.1	Overview	6
4.2	Stress Tests	6
4.2.1	Test Configuration	6
4.2.2	Test Payload.....	7
4.2.3	Test Procedure	7
4.3	Soak Tests	8
5	Integrations Testing	10
5.1	Overview	10
5.2	Driver Tests.....	10
5.2.1	Test Configuration	10
5.2.2	Test Scope	10
5.2.3	Test Procedures	10
Annex A	- Performance Testing Definitions	11

1 Instaclusr Certification Framework

1.1 Description

The **Instaclusr Certification Framework for Open Source Software (ICF-OSS)** provides an analytical basis for selecting and recommending various open source technologies for production usage. This is achieved by:

- a) assessing various open source projects to gain a level of assurance that the project has the foundation and capability to build and sustain production-grade software, and
- b) testing specific versions of open source technologies for function, performance, and interoperability.

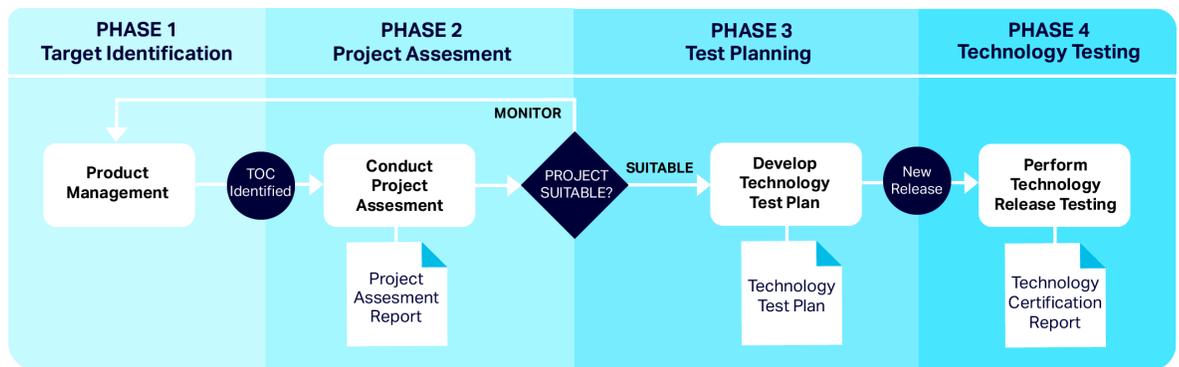


Figure 1: Instaclusr Certification Framework Process

1.2 Document Identification Information

Identifier	ICF_CTP - Apache Cassandra
Description	This document is the Technology Test Plan (TTP) for Apache Cassandra as part of the Instaclusr Certification Framework for Open Source Software (ICF-OSS).
Target of Certification	Apache Cassandra http://cassandra.apache.org/

1.3 Copyright Statement

© Instaclusr Pty Limited, 2019

Except as permitted by the copyright law applicable to you, you may not reproduce, distribute, publish, display, communicate or transmit any of the content of this document, in any form, but any means, without the prior written permission of Instaclusr Pty Limited.

Instaclusr Use Policy:

<https://www.instaclusr.com/company/policies/acceptable-use-policy/>

2 Introduction

2.1 Overview

This document is a Technology Test Plan that has been developed as Phase 3 of the certification process for Apache Cassandra.

Apache Cassandra is a distributed, wide column store, NoSQL database management system. The database is open source and is managed and governed by the Apache Software Foundation under its license agreements.

The database is designed to handle large amounts of data across many commodity servers, providing high availability with no single point of failure. The software provides support for clusters that can span multiple data centers with asynchronous masterless replication.

2.2 Purpose

The purpose of these testing efforts is to provide a level of assurance that a specific release of Apache Cassandra is suitable for inclusion in the Instacluster Managed Platform and is suitable for production integration and deployment.

The testing efforts will result in a certification report that provides the results of the testing effort and identify any performance regression, or potential caveats, as well as some further guidance on how to use the release.

2.3 Scope

The scope of any testing effort is for a specific release or version of Apache Cassandra software. Any limitations or recommendations will be specified in the certification report for specific releases of Apache Cassandra.

3 Functional Testing

3.1 Overview

The aim of functional testing for Apache Cassandra is to ensure that the version under test has had all the features and components tested in accordance with the developers.

Functional testing is achieved using the following test methods:

- The unit tests that are included for each component of the software.
- The distributed tests that test the distributed functionality of the software.

For any failures that are identified during the functional testing process a detailed analysis of the risks associated with these functions will be provided.

3.2 Unit Tests

Unit tests are maintained as part of the Apache Cassandra project. These tests are focused on the functionality of the various components of the Apache Cassandra codebase within a specific instance.

The unit tests will be performed using the following testing procedures:

1. All unit tests provided with the release of Apache Cassandra that is under test will be run by the certification team.
2. Each unit test will have a result of either PASS or FAIL.
3. If a unit test has a result of FAIL then further explanation and relevant details will be provided.
4. Both the “All” and the “long” tests are to be performed.

3.3 Distributed Tests

The distributed tests (also known as DTests) are maintained as part of the Apache Cassandra project. These tests are specifically focused on the functionality of the software in a multi-node cluster.

For our certification process we run the distributed tests branch that corresponds to the relevant release of Apache Cassandra.

The distributed tests will be performed using the following testing procedures:

1. All distributed tests that corresponds to a release of Apache Cassandra will be run by the certification team.
2. Each distributed test will have a result of either PASS or FAIL.
3. If a distributed test has a result of FAIL then further explanation and relevant details will be provided.

4 Performance Testing

4.1 Overview

The aim of performance testing for Apache Cassandra is to determine that the version being tested delivers suitable performance under stress and with production-grade workloads.

Performance testing is delivered through the following specific methods:

- Stress testing aims to test the performance of the software on various infrastructure under increasing workloads.
- Soak testing aims to test the performance of the software under sustained production-grade workloads.

4.2 Stress Tests

The aim of the stress testing is to benchmark the read and write latency of various infrastructure types running Apache Cassandra under a range of relevant workloads.

4.2.1 Test Configuration

The baseline infrastructure for the testing efforts is performed on AWS and are as specified in the following table.

Identifier	Node Type	Cluster Size	Node Specifications
CS01	AWS R4-I	3 Nodes	<ul style="list-style-type: none">• 2 CPU cores• 15.25GB RAM• 3200 GB SSD
CS02	AWS R4-XL	3 Nodes	<ul style="list-style-type: none">• 4 CPU cores• 30.5GB RAM• 3200 GB SSD
CS03	AWS R4-2XL	3 Nodes	<ul style="list-style-type: none">• 8 CPU cores• 61GB RAM• 3200 GB SSD
CS04	AWS R4-XL	9 Nodes	<ul style="list-style-type: none">• 4 CPU cores• 30.5GB RAM• 3200 GB SSD

4.2.2 Test Payload

The test payload is created by providing a stress **Spec yaml**, and the list of newly created nodes, to **cassandra-stress**. From here, we run **cassandra-stress** with the required parameters to perform the current test run.

We defined the column specifications for the test run inside a table named **typestest**. For the full YAML file, see 5.2.3 Annex A- Performance Testing Definitions.

The following table described the operations that are used. Each operation type is performed for each of the infrastructure tests and the following WRITE operation sizes are used:

- **SMALL**: Approx. 0.6kb per row
- **MEDIUM**: Approx. 12kb per row

WRITES	Writes	<ul style="list-style-type: none">• partitions: fixed(1)• batchtype: UNLOGGED• select: uniform(1..10)/10
READS	A combination of the two queries at a rate of 10:1	<ul style="list-style-type: none">• Simple: select * from typestest where name = ? and choice = ? LIMIT 1• Range: select name, choice, uid from typestest where name = ? and choice = ? and date >= ? LIMIT 10
MIXED	A combination of writes, simple reads, and range reads at a ratio of 10:10:1	<ol style="list-style-type: none">1) Writes: 102) Simple: 103) Range: 1

4.2.3 Test Procedure

The stress tests will be performed using the following procedures:

1. Fill the disk to approximately 20% full using a combination of “Small” and “Medium” size writes. This is done to ensure that the entire data model is not sitting in memory, and to give a better representation of a production cluster.
2. Wait for the cluster to finish all compactions. This is done to ensure that all clusters are given the same starting point for each test run, to make test runs comparable and verifiable.
3. For each both **SMALL** and **MEDIUM** writes, and waiting for compactions to finish between each step:
 - a. Perform a stress test with fixed operations per second to record time taken to perform the operations. Then perform and record the following:
 - b) Perform writes for 1 hour at a given OPS/Second to determine write latency.

- i. Perform reads for 1 hour at a given OPS/Second to determine if read latency has changed.
- ii. Perform a combination of reads and writes at a given OPS/Second to determine mixed latency.

Instance	# of Nodes	Write OPS/Sec		Reads OPS/Sec		Mixed OPS/Sec	
		Small	Medium	Small	Medium	Small	Medium
R4-l	3	5700	400	1800	400	3000	350
R4-xl	3	12000	450	12500	2500	11000	450
R4-2xl	3	17000	450	25000	15000	9000	400
R4-xl	9	37000	450	42000	18000	31000	1200

- a. Perform multiple stress tests on the cluster, increasing the number of operations per second each test, until the cluster is overloaded. The cluster will be considered overloaded if either of two conditions are met:
 - i. The latency of the operation is above the threshold of 5 ms Median latency for writes, or 20 ms Median latency for reads, indicating that the cluster is unable to maintain the the current load.
 - ii. If a node has more than 20 pending compactions a minute after the test completes, which indicates that the cluster is not maintaining currency with compactions and the load is not sustainable.
- b. The following tests are then run to measure maximum throughput in an overloaded state:
 - i. Perform a write operation in 30 minute test runs, increasing threads until a median write latency of 5ms is reached.
 - ii. Perform a read operation in 30 minute test runs, increasing threads until a median read latency of 20ms is reached.
 - iii. Perform a combination of reads & write operation in 30 minute test runs, increasing threads until a read median latency of 20ms is reached.

4.3 Soak Tests

The aim of soak testing is to test how a cluster will perform under sustained production like conditions.

This test is performed by running fixed operations per second with a mixed workload for 24 hours, while simulating node outages and maintenance operations (repairs) to confirm that Apache Cassandra operates as expected under typical production conditions.

The CS04 configuration is used for infrastructure and the following testing procedure is used:

1. Each disk of each node is filled to 20% capacity using medium sized writes.

2. Cluster compactions are performed.
3. Once compactions are completed, run a mixed load performance test for 24 hours with the following activities:
4. Remove node 1 in the first rack every hour for 5 minutes.
 - a. Remove an additional node in the first rack every 2 hours for 5 minutes.
 - b. Remove a third node in the first rack every 3 hours for 5 minutes
 - c. Run a full repair after 12 hours.
5. Collect and record latency, operations per second, failures, and Garbage Collection metrics.

5 Integrations Testing

5.1 Overview

The purpose of our integration testing is to confirm that the version of Cassandra under test has not broken any compatibility with the major Cassandra drivers. In this way, we are able to indicate that your application should work with minimal changes.

5.2 Driver Tests

Driver testing aims to confirm interoperability between major Cassandra language drivers and the new release of Apache Cassandra under testing.

5.2.1 Test Configuration

We create a 3 node cluster which runs the version of Cassandra under test. We then open the Cassandra firewall port to our test box and attempt to perform operations on the cluster from various Programming Languages.

5.2.2 Test Scope

The following drivers are to be tested:

1. CQLSH
2. The Datastax Java Driver
3. The Datastax Python Driver
4. The Datastax Ruby Driver
5. The Gocql Go Driver
6. The Datastax Scala Driver
7. The ALIA Clojure Driver

5.2.3 Test Procedures

For each of the drivers to be tested the following test procedures are used:

1. Connect to the new 3 node M4I-250 cluster which is running the Cassandra Version under test, and the language & driver we are testing
2. Create a basic keyspace with RF 3
3. Create a basic table with two text columns, and a decimal column. The first text column is the primary key
4. Insert a small row into the table which contains two text columns, and a decimal column
5. Query the written data by its primary key
6. Confirm that the queried data matches the written data.

If the driver version is able to successfully complete all operations then the test result will be PASS. If the connection does not work as expected it will be marked as FAIL and details will be provided in the certification report.

Annex A - Performance Testing Definitions

keyspace: stresscql2<Small|Medium>

```
keyspace_definition: |
    CREATE KEYSPACE stresscql2<Small|Medium> WITH replication =
        {'class': 'NetworkTopologyStrategy', <datacentre>': 3};
```

table: typestest

```
table_definition: |
    CREATE TABLE typestest (
        name text,
        choice boolean,
        date timestamp,
        address inet,
        dbl double,
        lval bigint,
        ival int,
        uid timeuuid,
        value blob,
        PRIMARY KEY((name,choice), date, address, dbl, lval, ival, uid)
    ) WITH compaction = { 'class': 'LeveledCompactionStrategy' }
    AND comment='A table of many types to test wide rows'
```

columnspec:

```
- name: name
size: fixed(48)
population: uniform(1..<Number of Writes>)
- name: date
cluster: uniform(20..1000)
- name: lval
population: gaussian(1..1000)
cluster: uniform(1..4)
- name: value
size: fixed(<Operation Size in Bytes>)
```

insert:

```
partitions: fixed(1)
batchtype: UNLOGGED
select: uniform(1..10)/10
```

queries:

```
simple1:
    cql: select * from typestest where name = ? and choice = ? LIMIT 1
    fields: samerow
range1:
    cql: select name, choice, uid from typestest where name = ? and
choice = ? and date >= ? LIMIT 10
    fields: multirow
```