

Annex 1 - Functional Testing

For a test to be successful, normal operations should be restored after every test. The default master server should regain master role with all replicators and data sources ONLINE.

Manual and Automatic Failover

Failover Test 1 - Administrator uses Tungsten to promote new master

Scenario: [LOGICAL] /usa> switch

Expectation:

- The master role will move a different server. Remaining servers will be reconfigured accordingly.
- Examine how the application performs during the process
- Data should be consistent

Failover Test 2 – Manually kill the master database process

Scenario: [LOGICAL] /usa> service master/mysql stop

Expectation:

- The database server stops and the master role is moved to another server.
- Examine how the application performs during the process
- Data should be consistent

Failover Test 3 – Remove power from the master database server

Scenario: Pull the power plug on the master server or run a restart command if that is not an option.

Expectation:

- The database server stops and the master role is moved to another server.
- Examine how the application performs during the process
- Data should be consistent

Failover Test 4 – Manually kill a slave database process

Scenario: [LOGICAL] /usa> service slave1/mysql stop

Expectation:

- The slave data source is marked as FAILED.
- Examine how the application performs during the process

Backup and Restore

Backup Test 1 – Take a backup of a slave and restore it to the same server

Scenario: [LOGICAL] /usa> datasource slave1 backup

[LOGICAL] /usa> datasource slave1 restore

Expectation:

- The commands should complete successfully

Backup Test 2 – Restore the backup to another server

Scenario:

slave1\$ rsync -avze ssh /opt/continuent/backups/ slave2:/opt/continuent/backups/

[LOGICAL] /usa> datasource slave2 restore

Expectation:

- The commands should complete successfully

Backup Test 3 – Take a backup of the master and restore to a slave

Scenario: [LOGICAL] /usa> datasource master backup

master\$ rsync -avze ssh /opt/continuent/backups/ slave2:/opt/continuent/backups/

[LOGICAL] /usa> datasource slave2 restore

Expectation:

- The restore command should complete successfully, or
- Run `treectl online -from-event #####:#####` to bring the slave ONLINE

Connectivity

Connectivity Test 1 – Connect to the connector and verify master host

Scenario: `$> mysql -h`hostname` -P9999 -uapp -p -e"select @@hostname for update"`

Expectation:

- The master hostname is returned

Connectivity Test 2 – Verify access to slaves

This test is only required if read/write splitting has been enabled. It should be run from a connector running on a server other than the master.

Scenario: `$> mysql -h`hostname` -P9999 -uapp -p -e"select @@hostname"`

Expectation:

- A hostname is returned
- Repeat the process on other hosts until a slave hostname is returned

Connectivity Test 3 – Verify access to the master before and after a switch

Scenario: `[LOGICAL] /usa> switch`

Expectation:

- Run Test 1 again and verify the new master hostname is returned

Performance

Performance Test 1 – Run a load test against the cluster

Scenario: Run load tests of some variety against the cluster to ensure the Connector and Replicator properly handle the load.

- HammerDB
- sysbench

Expectation:

- Solution can handle at 1TB of data with a minimum of 24k reads per minute and 1k writes per minute, over three tests:
 - Insert only from multiple locations
 - Read/write test
 - Heavy read test with few writes

Replicator Testing

Scenario: Evaluate system responsiveness in conjunction with performance tests [above]. This is not a test of network speed

Expectation:

- Quick replication latency between regions (<100ms greater than the current ping latency)

Network Partition Testing

Scenario: Simulate a partitioned network (for example, by modifying security group rules in AWS), and continue to do reads and writes on multiple clusters for 30 minutes.

Expectation:

- After resolving the partition, clusters should resync.

Scenario: Create a similar network partition, and write the same record on both sides of the partition

Expectation:

- After resolving the partition, the replicators on both sides should report errors. Demonstrate possible resolutions:
 - Remove record on side, skip record on the other side, and bring replication back online. The record not removed will be replicated
 - Remove both records and skip both transactions
 - Modify records and bring replicators online.