



Lustre Test Plan

Large Scale Test Plan

Adaptive Timeouts

Author	Date	Description of Document Change	Client Approval By	Client Approval Date
Elena Gryaznova	11/10/08	draft		
Nathan Rutman	11/13/08	review		
Elena Gryaznova	11/16/08	add correction in according to Nathan review		



I. Test Plan Overview

Executive Summary

- Statement of the problem trying to solve:
Test at scale Adaptive Timeouts of HEAD, b1_8_gate branches
- Required inputs:
 - The packages: build for HEAD branch
 - The packages: build for b1_8_gate branch
- Hardware to be used:
 - User clusters
- Expected output:
 - The current AT status of the listed branches
 - The following tickets will be used for summary and a status of testing

Problem Statement

This test plan describes the tests of large scale systems level testing of Adaptive Timeouts feature.

Goal

Verify AT functionality on a large system.

Success Factors

All tests need to run successfully.
No performance regressions.

Testing Plan

Define the setup steps that need to happen for the hardware to be ready? Who is responsible for these tests?

Get system time at User on a large system. Pre-feature testing has been completed and signed off by SUN QE for this feature.
--

Specify the date these tests will start, and length of time that these test will take to complete.

Date started: TBD

The time estimation for new test creation by QE: 4 days

The time estimation of 1 run: Common recovery tests: 4 days
--



* In the case of defects found the tests should be repeated. The estimated time of completed testing depends on:

- the number of defects found during testing;
- the time needed by developer to fix the defects;

Specify (at a high level) what tests will be completed?

Functional tests: new acceptance-small tests: recovery-scale (without Lustre reformatting)

Specify how you will restore the hardware (if required) and notify the client your testing is done.

We will need feedback from the User, recommend we use bugzilla for outputs.

The bugzilla ticket is filled for each failure.

Summary and status report are printed in the bug that we create for this testing.

Test Cases

Test Cases

Post-gate landing

COMMON scale tests

All these tests are (will be) integrated into acceptance-small as recovery-scale.sh (RECOVERY_SCALE).

Requirements:

1. Installed Lustre build packages on all cluster nodes.
2. installed Lustre user tools (lctl).
3. Shared directory with lustre-tests build on all clients
4. Formatted Lustre file system, mounted by clients
5. The configuration file in according to a formatted Lustre system
6. Installed dbench, iofzone
7. Installed mdsrate, IOR programs

To run these recovery scale tests over acceptance-small:

```
SETUP=: CLEANUP=: FORMAT=: ACC_SM_ONLY=RECOVERY_SCALE
NAME=<config_file> sh acceptance-small.sh
To run all recovery scale tests:
  SETUP=: CLEANUP=: FORMAT=: NAME=<config_file> sh recovery-scale.sh
To run the recovery scale tests separately:
  SETUP=: CLEANUP=: FORMAT=: NAME=<config_file> DURATION=<duration>
recovery-mds-scale.sh
  SETUP=: CLEANUP=: FORMAT=: NAME=<config_file> recovery-double-scale.sh
```



I.

Scale recovery tests (NC:1M:MO) (the test based on test11/17 from CMD3 project)

Scale recovery tests (NC:1M:MO) (the test based on test11/17 from CMD3 project)

(Collect recovery time stats for all these tests, verify that there are no unreasonably long recovery time, bug 16860).

1.a (was test11 in cmd3 project)

For defined duration (1-24 hours) repeatedly fail an MDS at defined (5-10 minutes) intervals and verify that no application errors occur. Load of clients: mdsrate, IOR, dbench, iozone.

Verify that clients loads are large enough to generate timeouts on some of the clients (from syslogs). Increase the loads on clients if needed.

recovery-mds-scale.sh

Example:

```
SETUP=: CLEANUP=: FORMAT=: NAME=<config_file> DURATION="3600" sh
recovery-mds-scale.sh
```

1.b (was test17 on cmd3 project)

Fail a random pair of nodes at defined (5-10 minutes) intervals and verify that no application errors occur. Load of clients: mdsrate, IOR, dbench iozone.

Verify that clients loads are large enough to generate timeouts on some of the clients (from syslogs). Increase the loads on clients if needed.

- 1: failover MDS, then OST
- 2: failover MDS, then 2 clients
- 4: failover OST, then another OST
- 5: failover OST, then 2 clients
- 6: failover OST, then MDS
- 7: failover 2 clients, then MDS
- 8: failover 2 clients, then OST
- 9: failover 2 clients, then 2 different clients

recovery-double-scale.sh

Example:

```
SETUP=: CLEANUP=: FORMAT=: NAME=<config_file> sh recovery-mds-scale.sh
```

II.

Additional AT large-scale tests

No additional tests.

Benchmarking

No additional benchmarks will be done.



II. Test Plan Approval

- Review date for the Test Plan review with the client:
11/10/08 – reviewed by J.D. Neumann
- Review date for the Test Plan by developers
(sent to Nathan Rutman).
- Date the Test Plan was approved by the client (and by whom)
- Date(s) agreed to by the client to conduct testing

III. Test Plan – Final Report

Test Results

Benchmarking

Conclusions

Summary of the test:

•
•

Next Steps

•
•