



OST Pools Support Large Scale Test Plan

Author	Date	Description of Document Change	Approval By	Approval Date
Minh Diep	10/20/08	First draft		
Minh Diep	10/26/08	Second draft with additional test cases		



I. Test Plan Overview

This test plan describes various testing activities and responsibilities that are planned to be performed by Lustre QE. Per this test plan, Lustre QE will provide large-scale, system level testing for OST Pools project.

Executive Summary

- Create a test plan to provide testing for OST Pools project for a large-scale cluster
- Required input from developers
- Require customer large cluster lab
- The output will be all tests are passed

Problem Statement

We need to test the OST Pools feature on a large-scale cluster to make sure the feature is scalable.

Goals

- Verify that file creation performance is the same with and without the OST Pools feature.

Success Factors

- All tests must be passed.
- Minimum performance regression. If there is a performance regression, the results must be reviewed and approved by RMG.

Testing Plan for OST Pools feature testing at large scale.

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

- 1) Get system time at customer lab.
- 2) Install Lustre RPMs.
- 3) Configure the Lustre file system and start running the tests.

The QE team in the Lustre group is responsible for setting up the test environment, running the tests, vetting and reporting the test results.

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

Date start: 2008-10-30

Estimated time for install and setup file system: 5 hours

Estimated time for 1 run:

It is difficult to estimate the time for a run because we have not tested on a large scale.

Specify (at a high level) what tests will be completed? New, Exist tests, manual or automate

Mdsrate

Existing test, automate

Test Cases

Large Scale Testing

Large Scale: all large scale tests for OST Pools will be integrated into acceptance-small as largescale.sh

To run this large scale test:

1. Install lustre.rpm and lustre-tests.rpm on all cluster nodes.
2. Specify the cluster configuration file, see cfg/local.sh and cfg/ncli.sh for details.

3. Run the test as:

```
ACC_SM_ONLY=LARGE_SCALE NAME=<config_file> sh acceptance-small.sh
```

or

```
NAME=<config_file> sh large-scale.sh
```

no.	Test Case	Description
1.	Run mdsrate without pools	<p>Create a file system with all of the OSTs and set stripe across all. Run mdsrate from all clients and record the performance numbers.</p> <p>Manual steps to run mdsrate: mkdir /mnt/lustre/single cd /usr/lib64/lustre/tests ./mdsrate.x86_64 --create --time 600 --dir /mnt/lustre/single --filefmt 'f%%d' ./mdsrate.x86_64 --unlink --time 600 --nfiles 840000 --dir /mnt/lustre/single --filefmt 'f%%d'</p>
2.	Run mdsrate with pool on all OSTs	<p>Create a file system with a pool of all of the OSTs and set stripe across all. Run mdsrate from all the clients and record the performance numbers.</p> <p>Instruction for creating a pool manually: 1. From mds: lctl pool_new <fs name>.<pool name> 2. lctl pool_add <fs name>.<pool name> lustre-OST[0-N] (where N is the total number of OST – 1) 3. Mount the file system on one client and run lfs setstripe -c -1 -s 1048576 -p <pool name> /mnt/lustre 4. Mount the file system on all clients 5. Start run mdsrate on all clients</p> <p>Manual steps to run mdsrate: mkdir /mnt/lustre/single cd /usr/lib64/lustre/tests ./mdsrate.x86_64 --create --time 600 --dir /mnt/lustre/single --filefmt 'f%%d' ./mdsrate.x86_64 --unlink --time 600 --nfiles 840000 --dir /mnt/lustre/single --filefmt 'f%%d'</p>
3.	Run mdsrate with pool on 50% OSTs	<p>Create a file system with a pool of half of all OSTs and set stripe across all. Run mdsrate from all the clients and record the performance numbers.</p> <p>Instruction for creating a pool manually:</p>



		<ol style="list-style-type: none"> 1. From the MDS: <code>lctl pool_new <fs name>.<pool name></code> 2. <code>lctl pool_add <fs name>.<pool name> lustre-OST[0-N]</code> (where N is the total number of OST/2) 3. Mount the file system on one client and run <code>lfs setstripe -c -1 -s 1048576 -p <pool name> /mnt/lustre</code> 4. Mount the file system on all clients. 5. Start run <code>mdsrates</code> on all clients. <p>Manual steps to run <code>mdsrates</code>:</p> <pre>mkdir /mnt/lustre/single cd /usr/lib64/lustre/tests ./mdsrates.x86_64 --create --time 600 --dir /mnt/lustre/single --filefmt 'f%%d' ./mdsrates.x86_64 --unlink --time 600 --nfiles 840000 --dir /mnt/lustre/single --filefmt 'f%%d'</pre>
4.	Evaluate creation/deletion of large pools	<p>Manually create a pool with the maximum number of OSTs available. Record the time it takes to add such large number of OSTs.</p> <ol style="list-style-type: none"> 1. From mds: <code>lctl pool_new <fs name>.<pool name></code> 2. Time <code>lctl pool_add <fs name>.<pool name> lustre-OST[0-N]</code> (where N is the total number of OST – 1) <p>Manually delete the pool in above configuration. Record the time it takes to delete such large number of OSTs.</p> <ol style="list-style-type: none"> 3. Time <code>lctl pool_remove <fs name>.<pool name> lustre-OST[0-N]</code> (where N is the total number of OST – 1)
5.	verified overlapping pools	<ol style="list-style-type: none"> 1. Create a pool (named abc) with 75% of the number of OSTs. <code>lctl pool_new lustre.abc</code> <code>lctl pool_add lustre.abc OST[0-2]</code> 2. Create a second pool (named 123) with 50% overlapping with pool abc and 25% non-overlapping OSTs. <code>lctl pool_new lustre.123</code> <code>lctl pool_add lustre.123 OST[1-3]</code> 3. On the clients, mount lustre FS under <code>/mnt/lustre</code>, then create two directories. <code>mkdir /mnt/lustre/abc /mnt/lustre/123</code> 4. Set stripe the directories to different a pool. <code>lfs setstripe -c -1 -s 1048576 -p abc /mnt/lustre/abc</code> <code>lfs setstripe -c -1 -s 1048576 -p 123 /mnt/lustre/123</code> 5. Create several files under each directory and verify the pool information is correct on each file. <p>Touch <code>/mnt/lustre/abc/a</code>, <code>getstripe</code> sample below is correct. <code>lfs getstripe -v a</code> OBDS:</p>



```
0: lustre-OST0000_UUID ACTIVE
1: lustre-OST0001_UUID ACTIVE
2: lustre-OST0002_UUID ACTIVE
3: lustre-OST0003_UUID ACTIVE
```

a

```
Imm_magic:      0x0BD30BD0
Imm_object_gr:  0x5
Imm_object_id:  0x800001cfc005
Imm_stripe_count: 3
Imm_stripe_size: 1048576
Imm_stripe_pattern: 1
Imm_pool_name:  abc
```

obdidx	objid	objid	group
2	3173	0xc65	5
0	2869	0xb35	5
1	3199	0xc7f	5

Touch /mnt/lustre/123/1, getstripe sample below
fs getstripe -v ./1

OBDS:

```
0: lustre-OST0000_UUID ACTIVE
1: lustre-OST0001_UUID ACTIVE
2: lustre-OST0002_UUID ACTIVE
3: lustre-OST0003_UUID ACTIVE
```

./1

```
Imm_magic:      0x0BD30BD0
Imm_object_gr:  0x5
Imm_object_id:  0x800001cfc008
Imm_stripe_count: 3
Imm_stripe_size: 1048576
Imm_stripe_pattern: 1
Imm_pool_name:  123
```

obdidx	objid	objid	group
2	3176	0xc68	5
3	3179	0xc6b	5
1	3202	0xc82	5

Benchmarking

Not applicable. This is not a performance project

II. Test Plan Approval

- Review date for the Test Plan review with the client:
- 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

Not applicable

Conclusions

Next Steps