


# 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		
Minh Diep	02/02/09	Updated with Cray's comments		

## 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 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 OST Pools feature on large scale cluster to make sure the feature is scalable.

## **Goals**

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

## **Success Factors**

- All tests must be passed
- Minimum performance regression. If there is performance regression, 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 Lustre file system and start running the tests

QE team in 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: Undefined

Estimated time for install and setup filesystem: 5 hours

Estimated time for 1 run: 5 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? 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 configuraion 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 filesystem with all the OSTs and set stripe across all. Run mdsrate from all the 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 and IOR with pool on all OSTs	<p>Create a filesystem with a pool of all 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 &lt;fs name&gt;.&lt;pool name&gt;            2. lctl pool_add &lt;fs name&gt;.&lt;pool name&gt; lustre-OST[0-N] (where N is the total number of OST – 1)            3. mount the filesystem on one client and run lfs setstripe -c -1 -s 1048576 -p &lt;pool name&gt; /mnt/lustre</p> <p><b>Note: additional scenarios can be done by setting the stripecount to 1 or 4.</b></p> <p>4. mount the filesystem on all clients            5. Start run mdsrate on all clients            6. Start IOR run 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>

		<p>Example ommand to run IOR</p> <pre>mpirun -np &lt;client * core-per-client&gt; -machinefile &lt;client file&gt; IOR -i 1 -a POSIX -b 4G -o /mnt/lustre/ior/iorData -t 2m -v -v -w -r -F</pre>
3.	Run mdsrate with pool on 50% OSTs	<p>Create a filesystem with a pool of half of all 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:</p> <ol style="list-style-type: none"> <li>1. From mds: <code>lctl pool_new &lt;fs name&gt;.&lt;pool name&gt;</code></li> <li>2. <code>lctl pool_add &lt;fs name&gt;.&lt;pool name&gt; lustre-OST[0-N]</code> (where N is the total number of OST/2)</li> <li>3. mount the filesystem on one client and run <code>lfs setstripe -c -1 -s 1048576 -p &lt;pool name&gt; /mnt/lustre</code></li> <li>4. mount the filesystem on all clients</li> <li>5. Start run mdsrate on all clients</li> </ol> <p>Manual steps to run mdsrate:</p> <pre>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'</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"> <li>1. From mds: <code>lctl pool_new &lt;fs name&gt;.&lt;pool name&gt;</code></li> <li>2. <code>time lctl pool_add &lt;fs name&gt;.&lt;pool name&gt; lustre-OST[0-N]</code> (where N is the total number of OST – 1)</li> </ol> <p>Manually delete the pool in above configuration. Record the time it takes to delete such large number of OSTs</p> <pre>time lctl pool_remove &lt;fs name&gt;.&lt;pool name&gt; lustre-OST[0-N] (where N is the total number of OST – 1)</pre>
5.	verified overlapping pools	<ol style="list-style-type: none"> <li>1. Create a pool (named abc) with 75% number of OSTs  <pre>lctl pool_new lustre.abc lctl pool_add lustre.abc OST[0-2]</pre> </li> <li>2. Create a second pool (named 123) with 50% overlapping with pool abc and 25% non-overlapping OSTs  <pre>lctl pool_new lustre.123 lctl pool_add lustre.123 OST[1-3]</pre> </li> <li>3. On the clients, mount lustre FS under /mnt/lustre, then create two directories  <pre>mkdir /mnt/lustre/abc /mnt/lustre/123</pre> </li> <li>4. Set stripe the directories to different pool</li> </ol>

```

lfs setstripe -c -1 -s 1048576 -p abc /mnt/lustre/abc
lfs setstripe -c -1 -s 1048576 -p 123 /mnt/lustre/123

```

5. Create several files under each directories and verify the pool information and ost used are correct on each file.

Touch /mnt/lustre/abc/a, getstripe sample below is correct.

```

lfs getstripe -v a
OBDS:
0: lustre-OST0000_UUID ACTIVE
1: lustre-OST0001_UUID ACTIVE
2: lustre-OST0002_UUID ACTIVE
3: lustre-OST0003_UUID ACTIVE
a
lmm_magic:      0x0BD30BD0
lmm_object_gr:  0x5
lmm_object_id:  0x800001cfc005
lmm_stripe_count: 3
lmm_stripe_size: 1048576
lmm_stripe_pattern: 1
lmm_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

```

lfs 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
lmm_magic:      0x0BD30BD0
lmm_object_gr:  0x5
lmm_object_id:  0x800001cfc008
lmm_stripe_count: 3
lmm_stripe_size: 1048576
lmm_stripe_pattern: 1
lmm_pool_name:  123
      obdidx      objid      objid      group
        2         3176      0xc68       5
        3         3179      0xc6b       5
        1         3202      0xc82       5

```

6. Test with each directory per OST/pool

```

1. Create a pool for every OST
lctl pool_new lustre.ost0
lctl pool_add lustre.ost0 OST0
...
lctl pool_new lustre.ostN

```

		<pre>lctl pool_add lustre.ostN OSTN</pre> <p>2. Set stripe one directory per pool</p> <pre>lfs setstripe -c -1 -s 1048576 -p ost0 /mnt/lustre/ost0</pre> <p>...</p> <pre>lfs setstripe -c -1 -s 1048576 -p ostN /mnt/lustre/ostN</pre> <p>3. Copy or dd some data to a file in each directory</p> <pre>dd if=/dev/zero of=/mnt/lustre/ost0 bs=1024 count=5</pre> <p>...</p> <pre>dd if=/dev/zero of=/mnt/lustre/ostN bs=1024 count=5</pre> <p>4. Verify that each file have the correct information</p> <pre>lfs getstripe -v /mnt/lustre/ost0</pre> <p>...</p> <pre>lfs getstripe -v /mnt/lustre/ostN</pre>
7.	<p>Test that the use of OST Pools do not affect roll-back to pre-1.8 releases.</p> <p>Note: This is not a scale testing.</p>	<p>1. On 1.8 release, create a pool and create some files with pool specific and some without pools</p> <p>2. Client on 1.6.6 can access files created with pools. Older client will get an EIO.</p> <p>3. Roll back to 1.6.6 server and verify that all files (pools and without pools) are accessible from 1.6.6 and 1.8 clients.</p>

### Benchmarking

Not applicable. This is not a performance project

## I. 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

## I. Test Plan – Final Report

### Test Results

### Benchmarking

Not applicable

### Conclusions

### Next Steps