

Lustre OSS Read Cache Feature SCALE Test Plan

Author	Date	Description of Document Change	Client Approval By	Client Approval Date
Jack Chen	10/08//08	Add scale testing to Feature Test Plan		
Jack Chen	10//30/08	Create Scale Test Plan, draft		
Jack Chen	11/04/08	Add performance with OSS read cache enabled/disabled		
Jack Chen	11/06/08	Add test goal by Alex's suggestion		
Jack Chen	02/11/09	Q&A, add new 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 the OSS read cache feature.

Executive Summary

- Create a test plan to provide testing for the OSS read cache 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 OSS read cache feature on a large-scale cluster to make sure the feature is scalable.

Goal

Verify that OSS read cache functions with a large system.
Compare performance data between OSS read cache enabled/disabled.

- 1) See improvement in those specific cases.
- 2) Prove there is no noticeable regression in the other cases.

Success Factors

All tests need to run successfully.

Performance data of OSS read cache enabled should be better then OSS read cache disabled.

Testing Plan

Pre-gate landing

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

Get system time 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: 2008-11-05

The time estimation for new test creation: 1 week



The time estimation of 1 run:
 OSS read cache feature tests (large-scale) : 2 ~ 3 days
 Performance tests: 2 ~ 3 days

* In the case of defects found, the tests should be repeated. The estimated time to complete testing depends on:
 -- the number of defects found during testing;
 -- the time needed by a developer to fix the defects;

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

IOR, PIOS Existing test, automate

Specify what requirements for a large scale testing

MDS number	8
OSS number	32~128, Depend on how many memory size on each OSS node.
Client numberss	More than 512
Memory size on MDSs	Normally, the more the better
Memory size on OSSs	OSSs_memory_size more than ½ Clients_memory_size, the more, the better
Memory size on Clients	Not specify

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 that we use bugzilla for outputs.

The bugzilla ticket is filed for each failure.

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

Test Cases

Test Cases

Post-gate landing

All these tests are (will be) integrated into acceptance-small as large-scale.sh (LARGE_SCALE).
 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 without lustre reformatting as:
 SETUP=: CLEANUP=: FORMAT=: ACC_SM_ONLY=LARGE_SCALE NAME=<config_file>
 sh acceptance-small.sh



or

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

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 dd, tar, dbench, iozone, IOR, PIOS

Report:

output file.

Test Cases

Improvement testing

Using kernel build to measure if OSS read cache improve read performance.

Requirements:

1. Installed Lustre build packages on all cluster nodes.
2. Installed Lustre user tools (lctl).
3. Formatted Lustre file system, mounted by clients
4. Copy Lustre source to Lustre mount piont
5. Disable OSS read cache, see Performance test cases
6. Build Lustre
 - 1) download linux kernel source
 - 2) set start time
 - 3) compile linux kernel
 - 4) set end time
- 7, Enable OSS read cache, repeat #6
8. Compare the time between enabled/disable time

Excepted result:

Enabled OSS read cache will spend less time than disable OSS read cache

Performance compare test data between cache enabled and disabled. Test performance regressions

No.	Test Case	Description
1. IOR	1)Run IOR with OSS cache disabled	Create a file system with all of the OSTs and set stripe across all. Run IOR using \${MPIRUN} from one client and record the performance numbers. Install and configure Lustre first. Manual steps to run IOR: 1) disable OSS cache



		<pre>lctl set_param -n obdfilter.*.cache 0 or lctl set_param -n obdfilter.*.read_cache_enable 0 lctl set_param -n obdfilter.*.writethrough_cache_enable 0</pre> <p>2) disabled debug /sbin/sysctl -w lnet.subsystem_debug=0 /sbin/sysctl -w lnet.debug=0</p> <p>3) set sripe count to -1 /usr/bin/lfs setstripe \${MOUNTPOINT} 1048576 -1 -1 run IOR as mpiuser with/without -C for "file_per_process" \${MPIRUN} -machinefile \${MACHINE_FILE} -np \${NPROC} \${IOR} - t \$bufsize -b \$blocksize -f \$input_file Note: Options: -C -v -w -r -k (-T \${15*60}) MACHINE_FILE list all Lustre clients we used NPROC is the number of MPI processes, the better number is the client number BUFSIZE is 1M BLOCKSIZE depends on how many memory blocks are on all Lustre client nodes, \$Total_client_memory X 2 = \$BLOCKSIZE x \$NPROC</p> <p>4) set stripe count to 1 run IOR as mpiuser with/without -C for "file per process" \${MPIRUN} -machinefile \${MACHINE_FILE} -np \${NPROC} \${IOR} - t \$bufsize -b \$blocksize -f \$input_file Note: Options: -C -v -w -r -k (-T \${15*60}) MACHINE_FILE list all Lustre clients we used NPROC is the number of MPI processes, the better number is the client number BUFSIZE is 1M BLOCKSIZE depends on how many memory blocks are on all Lustre client nodes, \$Total_client_memory X 2 = \$BLOCKSIZE x \$NPROC</p> <p>5) enable debug, rerun 4) and 5)</p> <p>6) collect "vmstat 1" and "oprofiles" for all the runs - this is very important data.</p>
	2) Run IOR with OSS cache enabled	<p>The different step is 1) enable OSS read cache first.</p> <p>1) enable OSS cache lctl set_param -n obdfilter.*.cache 1 or lctl set_param -n obdfilter.*.read_cache_enable 1 lctl set_param -n obdfilter.*.writethrough_cache_enable 1</p>
2. PIOS	Run PIOS with OSS read cache disabled	<p>Create a file system with all of the OSTs and set stripe across all. Run PIOS using pdsh from all of the clients and record the performance numbers.</p>



		<p>Install and configure Lustre first.</p> <p>Manual steps to run PIOS:</p> <p>1) disable OSS cache <code>lctl set_param -n obdfilter.*.cache 0</code> or <code>lctl set_param -n obdfilter.*.read_cache_enable 0</code> <code>lctl set_param -n obdfilter.*.writethrough_cache_enable 0</code></p> <p>2) set stripe <code>/usr/bin/lfs setstripe \${MOUNTPOINT} 1048576 -1 -1</code></p> <p>3) run PIOS as mpiuser <code>\$PIOS -t 1,8,16,40 -n 1024 -c 1M -s 8M -o 16M -p \$LUSTRE</code> <code>\$PIOS -t 1,8,16,40 -n 1024 -c 1M -s 8M -o 16M -p \$LUSTRE verify</code> <code>\$PIOS -t 1,8,16,40 -n 1024 -c 1M -s 8M -o 16M -L fpp -p \$LUSTRE</code> <code>\$PIOS -t 1,8,16,40 -n 1024 -c 1M -s 8M -o 16M -L fpp -p \$LUSTRE verify</code></p> <p>4) collect “vmstat 1” and “oprofiles” for all the runs - this is very important data</p>
	Run PIOS with OSS read cache enabled	<p>The different step is 1) enable OSS read cache first.</p> <p>1) enable OSS cache <code>lctl set_param -n obdfilter.*.cache 1</code> or <code>lctl set_param -n obdfilter.*.read_cache_enable 1</code> <code>lctl set_param -n obdfilter.*.writethrough_cache_enable 1</code></p>

Benchmarking

No benchmarks will be done.

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 report requirements of result

1. *Lustre configuration*
 - 1) *Node information: Processor model, Processor count, Memory size, Disks, Network type*
 - 2) *Disk read speed*



- 3) How many OSSs and clients used
2. For Improvement testing and Post-gate testing, report test output files
3. For Performance testing, report as spreadsheet file recorded test data, output file and "vmstat 1"/"oprofiles" file for each run.

Example:

IOR (xfer size: 1M)	Number of processes		
	2	4	8
File-per-process-Max write (MB/sec)	57.21	54.29	51.37
File-per-process-Max read (MB/sec)	72.66	67.51	63.23
Single-share-file-Max write (MB/sec)	54.64	55.23	49.27
Single-share-file-Max read (MB/sec)	66.47	83.22	68.48

PIOS Measurement Scenario (MB/s)	T:1,4,8,16,32,40 N:1024 C:1M S:16M O:32M					
	1 Cycle : Threads					
	1	4	8	16	32	40
First Write Single File	36	38	32	36	37	37
Read & Verify Single File	24	43	36	41	42	42
First Write Multiple Files	11	74	72	72	69	63
Read & Verify Multiple Files	21	22	38	54	61	60

Benchmarking

Conclusions

Summary of the test:

:

Next Steps

:

