

**Demonstration Milestone Completion for the
Inode Iterator and Object Index Scrub Subproject
on the
Lustre File System Checker Project of the
SFS-DEV-001 contract.**

Revision History

Date	Revision	Author
5/17/2012	Original	R. Henwood
5/21/2012	Rate-limit added	R. Henwood

Contents

Contents	2
Introduction	3
Performance Baseline	3
Test Methodology for Performance	3
Test filesystem configuration	3
Load Generation.....	3
Hardware platform	4
Test Results	4
Acceptance Tests.....	4
OI Scrub performance in the urgent case of missing all (64/64) OI files	4
OI Scrub performance in the urgent case of MDT restored from file-level backup	5
Performance impacts for create by urgent OI Scrub with cases of missing (N/64) OI file(s).....	5
Performance impacts for create by background OI Scrub with kinds of speed limit.....	7
Conclusions	8
Appendix A: Acceptance Test Results.....	9

Introduction

This document describes demonstration of Subproject 3.1 – Inode Iterator and Object Index Scrub subproject within the OpenSFS Lustre Development contract SFS-DEV-001 signed 7/30/2011.

The Inode Iterator and OI Scrub code is functionally complete. The purpose of this Milestone is to verify that the code performs acceptably in a production-like environment. In addition to completing all the Test Scenarios (recorded as Acceptance Criteria in the OI Scrub Solution Architecture), OI Scrub Performance will be measured as described below.

Performance Baseline

Inode Iterator and OI Scrub is novel development work. No baseline measurement exists. The following test will be performed to create a baseline measurement.

- OI Scrub performance in the urgent case of missing all (64/64) OI files.
- OI Scrub performance in the urgent case of MDT restored from file-level backup (if the results are identical to rebuild in case of missing OI file this test will be omitted from future testing).
- Impact on create performance by urgent OI Scrub with (N/64) missing OI file(s).
- Impact on create performance by background OI Scrub with speed limit.

An *urgent* scrub is one where there is detected severe corruption or loss of one or more OI files. In the case of a corrupted/missing OI file, there is loss of MDT functionality due to the inability to look up any MDT inode by its File Identifier (FID). In this case, the completion of the OI scrub is critical for proper MDT operation, so the OI scrub is run at the highest rate possible in order to return the MDT to a fully-functional state quickly.

A *non-urgent*, or *background*, OI scrub is one that is run by the administrator to periodically verify that the OI table contains valid data, but should minimize performance impact on the running filesystem. In the future the inode iteration used by a background OI Scrub will also be used for distributed Lustre coherency checking.

Test Methodology for Performance

The intention of this phase of testing is to demonstrate the Inode Iterator and OI Scrub behaves in accordance with typical real-world usage scenarios on a realistic environment. Completing these tests will provide a baseline for future performance enhancements. To perform the tests, the following tools are required:

Test filesystem configuration

Single MDS driven by echo_client directly

Load Generation

Load will be generated on the MDS directly by echo_client.

Hardware platform

Demonstration will take place on the Whamcloud cluster Toro. The hardware specification include:

- MDS node is of type 'Fat Intel Node' (fat-intel-2)
 - 2 x Intel Xeon(R) X5650 2.67GHz Six-core Processor (2-HT each core)
 - 16GB DDR3 1333MHz Memory
 - 2PT 40Gb/s 4X QSFP InfiniBand adapter card (Mellanox MT26428)
 - 1 QDR IB port on motherboard
 - SSD as external journal device (INTEL SSDSA2CW120G3), SATA II Enterprise Hard Drive as MDT (single disk, WDC WD2502ABYS-02B7A0)

Test Results

Acceptance Tests

Acceptance tests are defined in the [OI Scrub and Inode Iterator Solution Architecture](#).

Results from a successful run of these tests are included in Appendix A: Acceptance test results.

These tests are available as:

```
./lustre/tests/sanity-scrub.sh
```

at <http://review.whamcloud.com/#change,1910>

The referenced command line for tests:

```
sh sanity-scrub.sh
```

OI Scrub performance in the urgent case of missing all (64/64) OI files

File Count (in millions)	10	20	40	80	125*
Scrub Speed (files/sec)	108938	103961	100641	91430	86571

* Device space exhausted

A script to perform this test is available as:

```
./lustre/tests/scrub-performance.sh
```

at <http://review.whamcloud.com/#change,1910>

The referenced command line for tests (on Toro fat-intel-2):

```
MDSOPT="--mkfsoptions='-J device=/dev/sdc1' --mountfsoptions='user_xattr'"  
MDSDEV=/dev/sdb1 MDSDEV1=/dev/sdb1 MDSSIZE= 250000000 NTHREADS=64  
UNIT=2560000 EJOURNAL="/dev/sdc1" sh scrub-performance.sh 0 10240000  
163840000 2
```

OI Scrub performance in the urgent case of MDT restored from file-level backup

There is a small, but measurable difference in performance compared to performance of urgent OI rebuild in the case of missing OI file (above.)

File Count (in millions)	1	2	4	8	16
Scrub Speed (files/sec)	93096	85338	81923	81923	80316

Comparing remove/recreate and backup/restore: In the remove/recreate case, the OI Scrub is starting with a clean slate, the first creates can be inserted into the OI table directly in cache. When the OI is large (e.g. after restore, since it contains all of the old FID->inode entries) it is slower to read-modify-write.

NOTE: this test relies on GNU tar. Tar exhibits slow performance if millions of files require processing. This means only a reduced number of files can be created (1-16M) in backup/restore test compared with remove/recreate test.

A script to perform this test is available as:

```
./lustre/tests/scrub-performance.sh
```

at <http://review.whamcloud.com/#change,1910>

The referenced command line for tests (on Toro fat-intel-2):

```
MDSOPT="--mkfsoptions='-J device=/dev/sdc1' --mountfsoptions='user_xattr'"  
MDSDEV=/dev/sdb1 MDSDEV1=/dev/sdb1 MDSSIZE= 250000000 NTHREADS=64  
UNIT=2560000 EJOURNAL="/dev/sdc1" sh scrub-performance.sh 1 1024000 16384000  
2
```

Performance impacts for create by urgent OI Scrub with cases of missing (N/64) OI file(s).

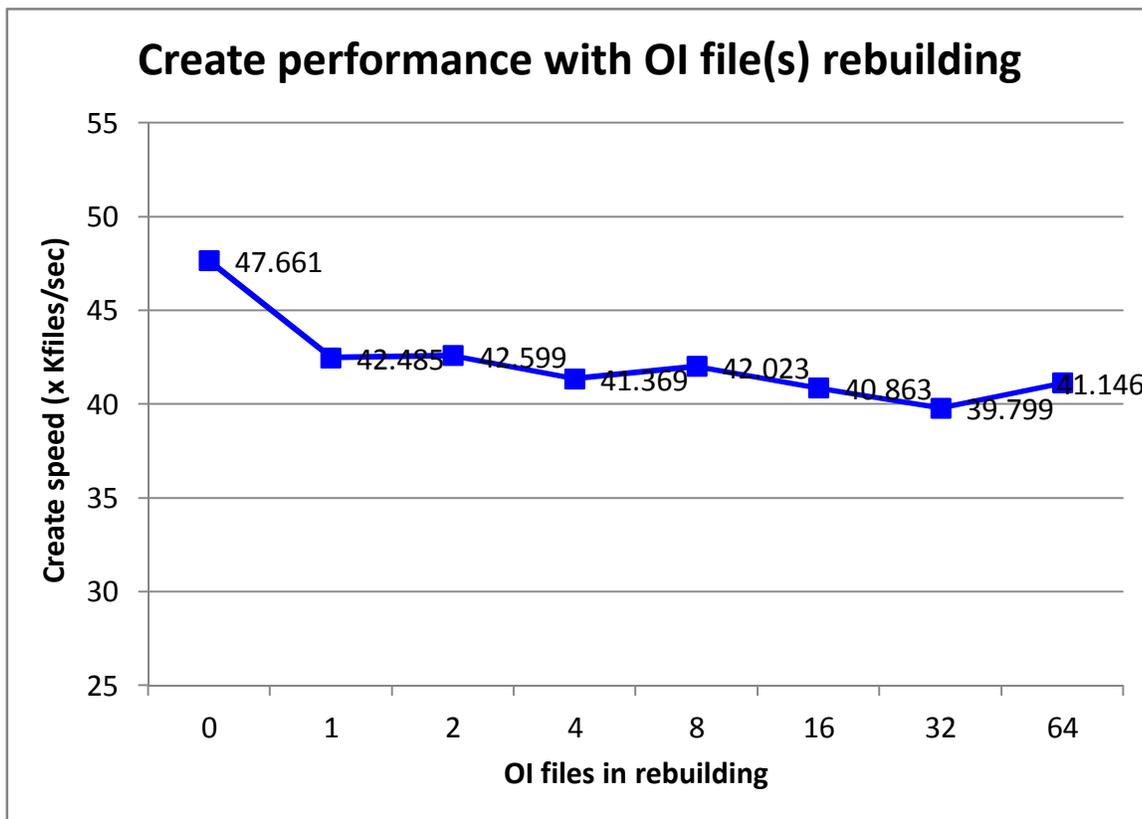
The method to perform these tests includes:

- 1) Create 2,560,000 files (zero-striped) driven by echo_client with 64 threads under per-thread based directories.
- 2) Repeat 1) for 2 times to create 5,120,000 files.
- 3) Remount the MDT device as type "ldisksf", and remove N/64 OI files.
- 4) Remount the MDT device as type "lustre", the removed OI files will be recreated automatically, run OI Scrub to rebuild those N OI files.
- 5) Create another 2,560,000 files (zero-striped) driven by echo_client with 64 threads under per-thread based directories.
- 6) Repeat 1) – 5) with different "N" (N=1,2,4,8,16,32,64). Between any two cases, reformat the MDT device to guarantee the disk layout is similar for each test.

These tests were repeated six times, with the mean performance reported:

OI files in rebuilding	Create performance	Performance %
0 (baseline: w/o OI Scrub)	47,661	100
1	42,485	89.14
2	42,599	89.38
4	41,369	86.80
8	42,023	88.17
16	40,863	85.74
32	39,799	83.50
64	41,146	86.33

There is a measurable performance impact of running urgent OI Scrub on a loaded system. The slow-down is typically measured to be approximately 10% ~ 17%. Such slow-down with urgent OI Scrub running is judged to be acceptable at this stage in the project. These measurements are illustrated below:



This figure shows that as the number of OI files to be rebuilt increases, the performance impacts for create do not substantially increase.

Performance impacts for create by background OI Scrub with kinds of speed limit

The method to perform these tests includes:

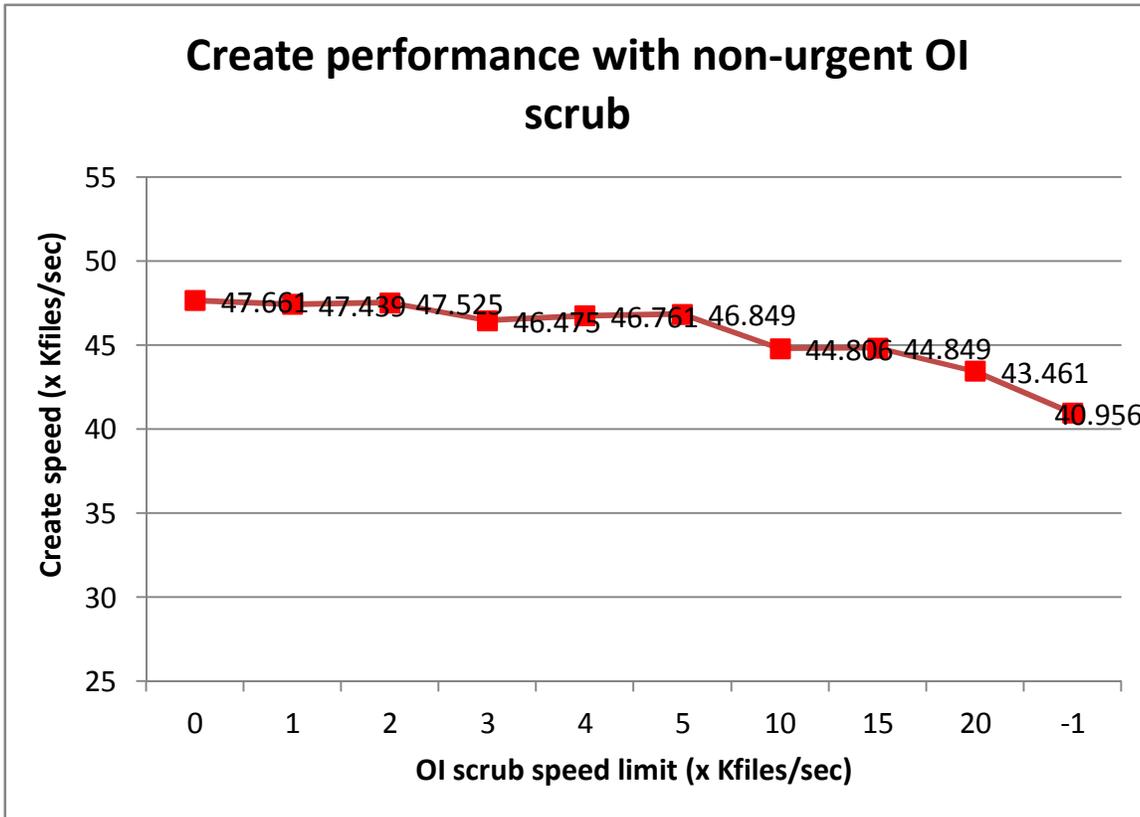
- 1) Create 2,560,000 files (zero-striped) driven by echo_client with 64 threads under per-thread based directories.
- 2) Repeat 1) for 2 times to create 5,120,000 files.
- 3) Run OI files consistency routine check with kinds of speed limit "R". In fact, all OI files are consistent.
- 4) Create another 2,560,000 files (zero-striped) driven by echo_client with 64 threads under per-thread based directories.
- 5) Repeat 1) - 4) with different "R" (R=1K/sec, 2K/sec, 3K/sec, 4K/sec, 5K/sec, 10K/sec, 15K/sec, 20K/sec). Between any two cases, reformat the MDT device to guarantee the disk layout is similar for each test.

These tests were repeated six times, the mean value shown in the table below:

OI Scrub rate limit (1000 inodes/s)	Create Performance	Performance %
0 (baseline: w/o OI Scrub)	47,661	100
1	47,439	99.53
2	47,525	99.71
3	46,475	97.51
4	46,761	98.11
5	46,849	98.30
10	44,806	94.01
15	44,849	94.10
20	43,461	91.19
-1 (OI Scrub w/o speed limit)*	40,956	85.93

* The real performance of background OI Scrub w/o speed limit during create (with full speed on the test node) is approximately 20K/sec.

There is a measurable performance running background OI Scrub on a loaded system as the scrub rate limit is increased. The slow-down grows from approximately 0% to 14% as the scrub rate limit is increased from 1k inodes/s to 20k inodes/s. Such slow-down with background OI Scrub running is judged to be acceptable at this stage in the project. These measurements are illustrated below:



As you can see, when OI Scrub speed is limited within 25% (5K / 20K) of the full speed of OI Scrub, the performance impacts for create is less than 3%. Means the performance impacts for create by low-speed OI Scrub is negligible.

As for dynamic rate control of OI Scrub, which runs at full speed on an idle MDT and is rate limited when the MDT is under load, will be implemented as part of LFSCK Phase IV.

Conclusions

OI Scrub and Inode iterator has successfully completed both functional Acceptance and Performance tests. Performance tests illustrate that the OI Scrub performs as expected online and under load. The overall performance of OI Scrub is particularly encouraging given this is the first stage of a project where functionality, not performance is the focus. LFSCK performance will be addressed in Phase IV of this project.

Appendix A: Acceptance Test Results

This test log is available as an attachment to: [OSF-87: OI Scrub completes all acceptance tests.](#)

```
Logging to shared log directory: /tmp/test_logs/1336703641
Checking servers environments
Checking clients fat-intel-2.lab.whamcloud.com environments
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
quota/lquota options: 'hash_lqs_cur_bits=3'
Setup mgs, mdt, osts
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting ost1: -o loop /tmp/lustre-ost1 /mnt/ost1
Started lustre-OSTffff
Starting ost2: -o loop /tmp/lustre-ost2 /mnt/ost2
Started lustre-OSTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Starting client fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Started clients fat-intel-2.lab.whamcloud.com:
fat-intel-2.lab.whamcloud.com@tcp:/lustre on /mnt/lustre type lustre (rw,user_xattr,flock)
Using TIMEOUT=20
disable quota as required
setting all flavor to null
already have total 5 null connections
excepting tests:
```

```
== sanity-scrub test 0: Do not auto trigger OI Scrub for non-backup/restore case == 19:34:26 (1336703666)
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre opts:
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre2 (opts:)
Stopping /mnt/mds1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost2 (opts:-f) on fat-intel-2.lab.whamcloud.com
waited 0 for 5 ST ost OSS OSS_uid 0
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
Formatting mgs, mds, osts
Format mds1: /tmp/lustre-mdt1
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
Format ost1: /tmp/lustre-ost1
Format ost2: /tmp/lustre-ost2
Checking servers environments
Checking clients fat-intel-2.lab.whamcloud.com environments
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
Setup mgs, mdt, osts
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting ost1: -o loop /tmp/lustre-ost1 /mnt/ost1
Started lustre-OSTffff
Starting ost2: -o loop /tmp/lustre-ost2 /mnt/ost2
Started lustre-OSTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Starting client fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Started clients fat-intel-2.lab.whamcloud.com:
fat-intel-2.lab.whamcloud.com@tcp:/lustre on /mnt/lustre type lustre (rw,user_xattr,flock)
Using TIMEOUT=20
disable quota as required
10.10.4.129@tcp:/lustre /mnt/lustre lustre rw,flock,user_xattr 0 0
```

```
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping /mnt/mds1 (opts:) on fat-intel-2.lab.whamcloud.com
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDT0000
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Resetting fail_loc on all nodes...done.
PASS 0 (50s)
```

== sanity-scrub test 1: Trigger OI Scrub when MDT mounts for OI files remove/recreate case == 19:35:16 (1336703716)

```
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre opts:
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre2 (opts:)
Stopping /mnt/mds1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost2 (opts:-f) on fat-intel-2.lab.whamcloud.com
waited 0 for 5 ST ost OSS OSS_uid 0
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
Formatting mgs, mds, osts
Format mds1: /tmp/lustre-mdt1
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
Format ost1: /tmp/lustre-ost1
Format ost2: /tmp/lustre-ost2
Checking servers environments
Checking clients fat-intel-2.lab.whamcloud.com environments
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
Setup mgs, mdt, osts
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting ost1: -o loop /tmp/lustre-ost1 /mnt/ost1
Started lustre-OSTffff
Starting ost2: -o loop /tmp/lustre-ost2 /mnt/ost2
Started lustre-OSTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Starting client fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Started clients fat-intel-2.lab.whamcloud.com:
fat-intel-2.lab.whamcloud.com@tcp:/lustre on /mnt/lustre type lustre (rw,user_xattr,flock)
Using TIMEOUT=20
disable quota as required
10.10.4.129@tcp:/lustre /mnt/lustre lustre rw,flock,user_xattr 0 0
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping /mnt/mds1 (opts:) on fat-intel-2.lab.whamcloud.com
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDT0000
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Resetting fail_loc on all nodes...done.
PASS 1 (54s)
```

== sanity-scrub test 2: Trigger OI Scrub when MDT mounts for backup/restore case == 19:36:10 (1336703770)

```
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre opts:
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre2 (opts:)
Stopping /mnt/mds1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost2 (opts:-f) on fat-intel-2.lab.whamcloud.com
waited 0 for 5 ST ost OSS OSS_uid 0
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
Formatting mgs, mds, osts
Format mds1: /tmp/lustre-mdt1
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
Format ost1: /tmp/lustre-ost1
```

```
Format ost2: /tmp/lustre-ost2
Checking servers environments
Checking clients fat-intel-2.lab.whamcloud.com environments
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
Setup mgs, mdt, osts
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting ost1: -o loop /tmp/lustre-ost1 /mnt/ost1
Started lustre-OSTffff
Starting ost2: -o loop /tmp/lustre-ost2 /mnt/ost2
Started lustre-OSTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Starting client fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Started clients fat-intel-2.lab.whamcloud.com:
fat-intel-2.lab.whamcloud.com@tcp:/lustre on /mnt/lustre type lustre (rw,user_xattr,flock)
Using TIMEOUT=20
disable quota as required
10.10.4.129@tcp:/lustre /mnt/lustre lustre rw,flock,user_xattr 0 0
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping /mnt/mds1 (opts:) on fat-intel-2.lab.whamcloud.com
/root/SCRUB/lustre-release/lustre/tests
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
/root/SCRUB/lustre-release/lustre/tests
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Resetting fail_loc on all nodes...done.
PASS 2 (56s)
```

```
== sanity-scrub test 3: Do not trigger OI Scrub when MDT mounts if 'noscrub' specified == 19:37:06 (1336703826)
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre opts:
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre2 (opts:)
Stopping /mnt/mds1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost2 (opts:-f) on fat-intel-2.lab.whamcloud.com
waited 0 for 5 ST ost OSS OSS_uuid 0
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
Formatting mgs, mds, osts
Format mds1: /tmp/lustre-mdt1
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
Format ost1: /tmp/lustre-ost1
Format ost2: /tmp/lustre-ost2
Checking servers environments
Checking clients fat-intel-2.lab.whamcloud.com environments
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
Setup mgs, mdt, osts
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting ost1: -o loop /tmp/lustre-ost1 /mnt/ost1
Started lustre-OSTffff
Starting ost2: -o loop /tmp/lustre-ost2 /mnt/ost2
Started lustre-OSTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Starting client fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Started clients fat-intel-2.lab.whamcloud.com:
fat-intel-2.lab.whamcloud.com@tcp:/lustre on /mnt/lustre type lustre (rw,user_xattr,flock)
Using TIMEOUT=20
disable quota as required
```

```
10.10.4.129@tcp:/lustre /mnt/lustre lustre rw,flock,user_xattr 0 0
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping /mnt/mds1 (opts:) on fat-intel-2.lab.whamcloud.com
/root/SCRUB/lustre-release/lustre/tests
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
/root/SCRUB/lustre-release/lustre/tests
Starting mds1: -o loop,user_xattr,noscrub /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre2 (opts:)
Stopping /mnt/mds1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost2 (opts:-f) on fat-intel-2.lab.whamcloud.com
waited 0 for 5 ST ost OSS OSS_uid 0
Resetting fail_loc on all nodes...done.
PASS 3 (82s)
```

== sanity-scrub test 4: Trigger OI Scrub automatically if inconsistent OI mapping was found == 19:38:28 (1336703908)

```
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre2 (opts:)
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -pingr
gss/krb5 is not supported
Formatting mgs, mds, osts
Format mds1: /tmp/lustre-mdt1
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
Format ost1: /tmp/lustre-ost1
Format ost2: /tmp/lustre-ost2
Checking servers environments
Checking clients fat-intel-2.lab.whamcloud.com environments
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -pingr
gss/krb5 is not supported
Setup mgs, mdt, osts
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting ost1: -o loop /tmp/lustre-ost1 /mnt/ost1
Started lustre-OSTffff
Starting ost2: -o loop /tmp/lustre-ost2 /mnt/ost2
Started lustre-OSTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Starting client fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Started clients fat-intel-2.lab.whamcloud.com:
fat-intel-2.lab.whamcloud.com@tcp:/lustre on /mnt/lustre type lustre (rw,user_xattr,flock)
Using TIMEOUT=20
disable quota as required
10.10.4.129@tcp:/lustre /mnt/lustre lustre rw,flock,user_xattr 0 0
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping /mnt/mds1 (opts:) on fat-intel-2.lab.whamcloud.com
/root/SCRUB/lustre-release/lustre/tests
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
/root/SCRUB/lustre-release/lustre/tests
Starting mds1: -o loop,user_xattr,noscrub /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Resetting fail_loc on all nodes...done.
PASS 4 (37s)
```

== sanity-scrub test 5: OI Scrub state machine == 19:39:05 (1336703945)

```
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre opts:
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre2 (opts:)
Stopping /mnt/mds1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost2 (opts:-f) on fat-intel-2.lab.whamcloud.com
waited 0 for 5 ST ost OSS OSS_uid 0
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
```

```

subsystem_debug=all -lnet -lnd -pingr
gss/krb5 is not supported
Formatting mgs, mds, osts
Format mds1: /tmp/lustre-mdt1
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
Format ost1: /tmp/lustre-ost1
Format ost2: /tmp/lustre-ost2
Checking servers environments
Checking clients fat-intel-2.lab.whamcloud.com environments
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -pingr
gss/krb5 is not supported
Setup mgs, mdt, osts
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting ost1: -o loop /tmp/lustre-ost1 /mnt/ost1
Started lustre-OSTffff
Starting ost2: -o loop /tmp/lustre-ost2 /mnt/ost2
Started lustre-OSTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Starting client fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Started clients fat-intel-2.lab.whamcloud.com:
fat-intel-2.lab.whamcloud.com@tcp:/lustre on /mnt/lustre type lustre (rw,user_xattr,flock)
Using TIMEOUT=20
disable quota as required
total: 1500 creates in 0.95 seconds: 1573.95 creates/second
10.10.4.129@tcp:/lustre /mnt/lustre lustre rw,flock,user_xattr 0 0
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping /mnt/mds1 (opts:) on fat-intel-2.lab.whamcloud.com
/root/SCRUB/lustre-release/lustre/tests
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
/root/SCRUB/lustre-release/lustre/tests
Starting mds1: -o loop,user_xattr,noscrub /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
fail_loc=0x190
10.10.4.129@tcp:/lustre /mnt/lustre lustre rw,flock,user_xattr 0 0
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
fail_loc=0x80000191
Stopping /mnt/mds1 (opts:) on fat-intel-2.lab.whamcloud.com
Starting mds1: -o loop,user_xattr,noscrub /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Stopping /mnt/mds1 (opts:) on fat-intel-2.lab.whamcloud.com
fail_loc=0x190
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
fail_loc=0x80000192
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
fail_loc=0x190
File: `/mnt/lustre/d0.sanity-scrub/d5/f51000'
Size: 0          Blocks: 0          IO Block: 2097152 regular empty file
Device: 2c54f966h/743766374d      Inode: 144115205255726131  Links: 1
Access: (0644/-rw-r--r--)  Uid: (  0/   root)   Gid: (  0/   root)
Access: 2012-05-10 19:39:51.000000000 -0700
Modify: 2012-05-10 19:39:45.000000000 -0700
Change: 2012-05-10 19:39:51.000000000 -0700
fail_loc=0
Resetting fail_loc on all nodes...done.
PASS 5 (80s)

== sanity-scrub test 6: OI Scrub resumes from last checkpoint == 19:40:25 (1336704025)
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre opts:
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre2 (opts:)
Stopping /mnt/mds1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost2 (opts:-f) on fat-intel-2.lab.whamcloud.com
waited 0 for 5 ST ost OSS OSS_uuid 0

```

```
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
Formatting mgs, mds, osts
Format mds1: /tmp/lustre-mdt1
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
Format ost1: /tmp/lustre-ost1
Format ost2: /tmp/lustre-ost2
Checking servers environments
Checking clients fat-intel-2.lab.whamcloud.com environments
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
Setup mgs, mdt, osts
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting ost1: -o loop /tmp/lustre-ost1 /mnt/ost1
Started lustre-OSTffff
Starting ost2: -o loop /tmp/lustre-ost2 /mnt/ost2
Started lustre-OSTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Starting client fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Started clients fat-intel-2.lab.whamcloud.com:
fat-intel-2.lab.whamcloud.com@tcp:/lustre on /mnt/lustre type lustre (rw,user_xattr,flock)
Using TIMEOUT=20
disable quota as required
total: 1000 creates in 0.65 seconds: 1543.82 creates/second
10.10.4.129@tcp:/lustre /mnt/lustre lustre rw,flock,user_xattr 0 0
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping /mnt/mds1 (opts:) on fat-intel-2.lab.whamcloud.com
/root/SCRUB/lustre-release/lustre/tests
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
/root/SCRUB/lustre-release/lustre/tests
Starting mds1: -o loop,user_xattr,noscrub /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
fail_loc=0x190
fail_loc=0x80000192
fail_loc=0x190
File: `/mnt/lustre/d0.sanity-scrub/d6/f6800'
Size: 0 Blocks: 0 IO Block: 2097152 regular empty file
Device: 2c54f966h/743766374d Inode: 144115205255725931 Links: 1
Access: (0644/-rw-r--r--) Uid: ( 0/ root) Gid: ( 0/ root)
Access: 2012-05-10 19:41:14.000000000 -0700
Modify: 2012-05-10 19:41:08.000000000 -0700
Change: 2012-05-10 19:41:14.000000000 -0700
10.10.4.129@tcp:/lustre /mnt/lustre lustre rw,flock,user_xattr 0 0
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
fail_loc=0x80000191
Stopping /mnt/mds1 (opts:) on fat-intel-2.lab.whamcloud.com
fail_loc=0x190
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
fail_loc=0
Resetting fail_loc on all nodes...done.
PASS 6 (75s)
```

```
== sanity-scrub test 7: System is available during OI Scrub scanning == 19:41:40 (1336704100)
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre2 (opts:)
Stopping /mnt/mds1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost2 (opts:-f) on fat-intel-2.lab.whamcloud.com
waited 0 for 5 ST ost OSS OSS_uuid 0
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
```

```
Formatting mgs, mds, osts
Format mds1: /tmp/lustre-mdt1
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
Format ost1: /tmp/lustre-ost1
Format ost2: /tmp/lustre-ost2
Checking servers environments
Checking clients fat-intel-2.lab.whamcloud.com environments
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -pinger
gss/krb5 is not supported
Setup mgs, mdt, osts
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting ost1: -o loop /tmp/lustre-ost1 /mnt/ost1
Started lustre-OSTffff
Starting ost2: -o loop /tmp/lustre-ost2 /mnt/ost2
Started lustre-OSTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Starting client fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Started clients fat-intel-2.lab.whamcloud.com:
fat-intel-2.lab.whamcloud.com@tcp:/lustre on /mnt/lustre type lustre (rw,user_xattr,flock)
Using TIMEOUT=20
disable quota as required
total: 500 creates in 0.33 seconds: 1511.19 creates/second
10.10.4.129@tcp:/lustre /mnt/lustre lustre rw,flock,user_xattr 0 0
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping /mnt/mds1 (opts:) on fat-intel-2.lab.whamcloud.com
/root/SCRUB/lustre-release/lustre/tests
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
/root/SCRUB/lustre-release/lustre/tests
Starting mds1: -o loop,user_xattr,noscrub /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
fail_loc=0x190
  File: `/mnt/lustre/d0.sanity-scrub/d7/f7300'
  Size: 0          Blocks: 0          IO Block: 2097152 regular empty file
Device: 2c54f966h/743766374d      Inode: 144115205255725431  Links: 1
Access: (0644/-rw-r--r--)  Uid: (  0/   root)   Gid: (  0/   root)
Access: 2012-05-10 19:42:27.000000000 -0700
Modify: 2012-05-10 19:42:21.000000000 -0700
Change: 2012-05-10 19:42:27.000000000 -0700
fail_loc=0
Resetting fail_loc on all nodes...done.
PASS 7 (58s)
```

```
== sanity-scrub test 8: Control OI Scrub manually == 19:42:38 (1336704158)
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre opts:
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre2 (opts:)
Stopping /mnt/mds1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost2 (opts:-f) on fat-intel-2.lab.whamcloud.com
waited 0 for 5 ST ost OSS OSS_uuid 0
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -pinger
gss/krb5 is not supported
Formatting mgs, mds, osts
Format mds1: /tmp/lustre-mdt1
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
Format ost1: /tmp/lustre-ost1
Format ost2: /tmp/lustre-ost2
Checking servers environments
Checking clients fat-intel-2.lab.whamcloud.com environments
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -pinger
gss/krb5 is not supported
Setup mgs, mdt, osts
```

```
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting ost1: -o loop /tmp/lustre-ost1 /mnt/ost1
Started lustre-OSTffff
Starting ost2: -o loop /tmp/lustre-ost2 /mnt/ost2
Started lustre-OSTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Starting client fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Started clients fat-intel-2.lab.whamcloud.com:
fat-intel-2.lab.whamcloud.com@tcp:/lustre on /mnt/lustre type lustre (rw,user_xattr,flock)
Using TIMEOUT=20
disable quota as required
10.10.4.129@tcp:/lustre /mnt/lustre lustre rw,flock,user_xattr 0 0
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping /mnt/mds1 (opts:) on fat-intel-2.lab.whamcloud.com
/root/SCRUB/lustre-release/lustre/tests
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
/root/SCRUB/lustre-release/lustre/tests
Starting mds1: -o loop,user_xattr,noscrub /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
fail_loc=0x190
Started LFSCK on the MDT device lustre-MDT0000: noop
Stopped LFSCK on the MDT device lustre-MDT0000.
Started LFSCK on the MDT device lustre-MDT0000: noop
fail_loc=0
Resetting fail_loc on all nodes...done.
PASS 8 (64s)
```

```
== sanity-scrub test 9: OI Scrub speed control == 19:43:42 (1336704222)
```

```
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre2 (opts:)
Stopping /mnt/mds1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost2 (opts:-f) on fat-intel-2.lab.whamcloud.com
waited 0 for 5 ST ost OSS UUID 0
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
Formatting mgs, mds, osts
Format mds1: /tmp/lustre-mdt1
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
Format ost1: /tmp/lustre-ost1
Format ost2: /tmp/lustre-ost2
Checking servers environments
Checking clients fat-intel-2.lab.whamcloud.com environments
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -ping
gss/krb5 is not supported
Setup mgs, mdt, osts
Starting mds1: -o loop,user_xattr /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Starting ost1: -o loop /tmp/lustre-ost1 /mnt/ost1
Started lustre-OSTffff
Starting ost2: -o loop /tmp/lustre-ost2 /mnt/ost2
Started lustre-OSTffff
Starting client: fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Starting client fat-intel-2.lab.whamcloud.com: -o user_xattr,flock fat-intel-2.lab.whamcloud.com@tcp:/lustre
/mnt/lustre
Started clients fat-intel-2.lab.whamcloud.com:
fat-intel-2.lab.whamcloud.com@tcp:/lustre on /mnt/lustre type lustre (rw,user_xattr,flock)
Using TIMEOUT=20
disable quota as required
total: 8000 creates in 5.17 seconds: 1548.12 creates/second
10.10.4.129@tcp:/lustre /mnt/lustre lustre rw,flock,user_xattr 0 0
Stopping client fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping /mnt/mds1 (opts:) on fat-intel-2.lab.whamcloud.com
/root/SCRUB/lustre-release/lustre/tests
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
```

```
/root/SCRUB/lustre-release/lustre/tests
Starting mds1: -o loop,user_xattr,noscrub /tmp/lustre-mdt1 /mnt/mds1
Started lustre-MDTffff
Started LFSCK on the MDT device lustre-MDT0000: noop
Started LFSCK on the MDT device lustre-MDT0000: noop
Resetting fail_loc on all nodes...done.
PASS 9 (97s)
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre (opts:)
Stopping clients: fat-intel-2.lab.whamcloud.com /mnt/lustre2 (opts:)
Stopping /mnt/mds1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost1 (opts:-f) on fat-intel-2.lab.whamcloud.com
Stopping /mnt/ost2 (opts:-f) on fat-intel-2.lab.whamcloud.com
waited 0 for 5 ST ost OSS OSS_uuid 0
Loading modules from /root/SCRUB/lustre-release/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -pinger
gss/krb5 is not supported
Formatting mgs, mds, osts
Format mds1: /tmp/lustre-mdt1
Format ost1: /tmp/lustre-ost1
Format ost2: /tmp/lustre-ost2
== sanity-scrub sanity-scrub.sh test complete, duration 699 sec == 19:45:40 (1336704340)
```