



Acceptance Small (acc-sm) Testing on Lustre

Author	Date	Description of Document Change	Approval By	Approval Date
Sheila Barthel	10/21/08	First draft		
Sheila Barthel	10/28/08	Second draft		
Sheila Barthel	11/11/08	Third (final) draft		



Overview

The Lustre QE group and developers use acceptance-small (acc-sm) tests to catch bugs early in the development cycle. Within the Lustre group, acc-sm tests are run on YALA, an automated test system. This document is being published to document acc-sm testing and encourage wider acc-sm testing in the Lustre community.

What is acc-sm testing and why do we use it for Lustre?

Acceptance small (acc-sm) testing is a suite of test cases used to verify different aspects of Lustre functionality.

What tests comprise the acc-sm test suite?

Each CVS branch contains a lustre/tests sub-directory; all acc-sm tests are stored here. The acceptance-small.sh file contains a list of all tests in the acc-sm suite. To get the list, run:

```
$ grep TESTSUITE_LIST acceptance-small.sh
```

The acc-sm tests are listed below, by CVS branch.

b1_6 branch:

```
$ grep TESTSUITE_LIST acceptance-small.sh
export TESTSUITE_LIST="RUNTESTS SANITY DBENCH BONNIE IOZONE FSX SANITYN LFSCK
LIBLUSTRE REPLAY_SINGLE CONF_SANITY RECOVERY_SMALL REPLAY_OST_SINGLE
REPLAY_DUAL INSANITY SANITY_QUOTA PERFORMANCE_SANITY"
-- 17 test suites
```

b1_8_gate branch:

```
$ grep TESTSUITE_LIST acceptance-small.sh
export TESTSUITE_LIST="RUNTESTS SANITY DBENCH BONNIE IOZONE FSX SANITYN LFSCK
LIBLUSTRE REPLAY_SINGLE CONF_SANITY RECOVERY_SMALL REPLAY_OST_SINGLE
REPLAY_DUAL REPLAY_VBR INSANITY SANITY_QUOTA PERFORMANCE_SANITY"
-- 18 test suites
```

HEAD branch:

```
$ grep TESTSUITE_LIST acceptance-small.sh
export TESTSUITE_LIST="RUNTESTS SANITY DBENCH BONNIE IOZONE FSX SANITYN LFSCK
LIBLUSTRE REPLAY_SINGLE CONF_SANITY RECOVERY_SMALL REPLAY_OST_SINGLE
REPLAY_DUAL INSANITY SANITY_QUOTA SANITY_SEC SANITY_GSS
PERFORMANCE_SANITY"
-- 19 test suites
```



To see the test cases in a particular acc-sm test, run:

```
$ grep run_ <testunit_script>
```

For example, to see the test cases that comprise the SANITY test:

```
$ grep run_ sanity.sh | tail -3
```

```
run_test 130c "FIEMAP (2-stripe file with hole)"
run_test 130d "FIEMAP (N-stripe file)"
run_test 130e "FIEMAP (test continuation FIEMAP calls)"
```

For each acc-sm test, what does it measure or show?

Here is a brief description of each acc-sm test.

RUNTESTS
A basic regression test with unmount / remount.

SANITY
Verifies Lustre operation under normal operating conditions.

DBENCH
Dbench benchmark for simulating N clients to produce the filesystem load.

BONNIE
Bonnie++ benchmark for creation, reading and deleting many small files.

IOZONE
IOzone benchmark for generating and measuring a variety of file operations.

FSX
File system exerciser.

SANITYN
Verifies operations from two clients under normal operating conditions.



LFCK

Tests e2fsck and fsck to detect and fix filesystem corruption.

LIBLUSTRE

Runs a test linked to a LibLustre client library.

REPLAY_SINGLE

Verifies recovery after an MDS failure.

CONF_SANITY

Verifies various Lustre configurations (including wrong ones), where the system must behave correctly.

RECOVERY_SMALL

Verifies RPC replay after communications failure.

REPLAY_OST_SINGLE

Verifies recovery after an OST failure.

REPLAY_DUAL

Verifies recovery from two clients after a server failure.

INSANITY

Tests multiple concurrent failure conditions.

SANITY_QUOTA

Verifies filesystem quotas.



How do you get the acc-sm tests?

The acc-sm test suite is stored in the lustre/tests sub-directory on each CVS branch (b1_6, b1_8, and HEAD).

Do you have to run every acc-sm test?

No. You can choose to run only specified acc-sm tests. Tests can be run either with or without the acceptance-sm.sh (acc-sm.sh) wrapper script. Here are several examples:

- To only run the RUNTESTS and SANITY.sh tests:
`ACC_SM_ONLY="RUNTESTS" sh acceptance-small.sh`

- OR -

```
sh runttests
```

- To only run test_1 and test_2 of the SANITY.sh tests:
`ACC_SM_ONLY="SANITYN" ONLY="1 2" sh acceptance-small.sh`
- To only run the replay-single.sh test and except the test_3* and test_4* tests:
`ACC_SM_ONLY="REPLAY_SINGLE" REPLAY_SINGLE_EXCEPT="3 4" sh acceptance-small.sh`
- To only run the conf-sanity.sh script (without the acceptance-small.sh wrapper script):
`CONF_SANITY_EXCEPT="$ (seq 15) " sh conf-sanity.sh`

Do the acc-sm tests have to be run in a specific order?

The test order is defined in the acceptance-small.sh script and in each test script. Users do not have to (and should not) do anything to change the order of tests.

Who runs the acc-sm tests?

Currently, the QE group and Lustre developers run acc-sm as the main test suite for Lustre testing. Acc-sm tests are run on YALA, the automated test system, with test reports submitted to Buffalo (a web interface that allows for browsing various Lustre test results). We welcome external contributions to the Lustre acc-sm test efforts – either of the Lustre code base or new testing platforms.



What type of Lustre environment is needed to run the acc-sm tests? Is anything special needed?

The default Lustre configuration for acc-sm testing is 1 combined MGS/MDT, 1 OST and 1 client. Several clusters with this default configuration are available in YALA, the automated Lustre test system.

To run the acc-sm test suite on a non-default Lustre configuration, you have to modify the default settings in the acc-sm configuration file, `lustre/tests/cfg/local.sh`. The configuration variables include `mds_HOST`, `ost_HOST`, `OSTCOUNT`, `MDS_MOUNT_OPTS` and `OST_MOUNT_OPTS`, among others.

To create your own configuration file, copy `cfg/local.sh` to `cfg/my_config.sh`:

```
cp cfg/local.sh cfg/my_config.sh
```

Edit the necessary variables in the configuration file (`my_config.sh`) and run `acc-sm`.

What are the steps to run acc-sm?

There are two methods to run the acc-sm tests.

1. Check out a Lustre branch (b1_6, b1_8 or HEAD).
2. Change directory to `lustre/tests`:
`cd lustre/tests`
3. Run `acc-sm` on a local, default Lustre configuration (1 MGS/MDT, 1 OST and 1 client):
`sh acceptance-small.sh 2>&1 | tee /tmp/output`

- OR -

1. Install the `lustre-tests` RPM (available at `lts-head:/var/cache/cfs/PACKAGE/rpm/lustre`).
2. Change directory to `lustre/tests`:
`cd /usr/lib/lustre/tests`
3. Create your own configuration file and edit it for your configuration.
`cp cfg/local.sh cfg/my_config.sh`
4. Run `acc-sm` on a local Lustre configuration.

Here is an example of running `acc-sm` on a non-default Lustre configuration (MDS is `sfire7`, OST is `sfire8`, `OSTCOUNT=1`, etc). In this example, only the SANITY test cases are being run.

```
ACC_SM_ONLY=SANITY mds_HOST=sfire7 ost8_HOST=sfire8 MDSDEV1=/dev/sda1
OSTCOUNT=1 OSTDEV1=/dev/sda1 MDSSIZE=5000000 OSTSIZE=5000000
MDS_MOUNT_OPTS="-o user_xattr" OST_MOUNT_OPTS=" -o user_xattr"
REFORMAT="--reformat" PDSH="pdsh -S -w" sh acceptance-small.sh
```



How do you run acc-sm on a mounted Lustre system?

To run acc-sm on a Lustre system that is already mounted, you need to use the correct configuration file (according to the mounted Lustre system) and run acc-sm as:

```
SETUP=: CLEANUP=: FORMAT=: NAME=<config> sh acceptance-small.sh
```

How do you run acc-sm with and without reformat?

By default, the acc-sm test suite does not reformat Lustre. If you want to reformat Lustre, run acc-sm with REFORMAT="--reformat":

```
REFORMAT="--reformat" sh acceptance-small.sh
```

If needed, you can specify WRITECONF="writeconf", and then run acc-sm with WRITECONF="writeconf":

```
WRITECONF="writeconf" sh acceptance-small.sh
```

How do you run acc-sm in a Lustre configuration with several clients?

The default configuration file for acc-sm is cfg/local.sh, which uses only one client (local). To use additional remote clients, specify the RCLIENTS list and use the cfg/ncli.sh configuration file (or your own copy of ncli configuration).

```
NAME=ncli RCLIENT=<space-separated list of remote clients> sh acceptance-small.sh
```

For example:

```
NAME=ncli RCLIENT="client2 client3 client11" sh acceptance-small.sh
```

What is the SLOW variable and how is it used with acc-sm?

The SLOW variable is used to run a subset of acc-sm tests. By default, the variable is set to SLOW=no, which causes some of the longer acc-sm tests to be skipped and acc-sm test run to complete in less than 2 hours. To run all of the acc-sm tests, set the variable to SLOW=yes:

```
SLOW=yes sh acceptance-small.sh
```

What is the FAIL_ON_ERROR variable and how is it used with acc-sm?

The FAIL_ON_ERROR variable is used to "stop" or "continue" running acc-sm tests after a test failure occurs. If the variable is set to "true" (FAIL_ON_ERROR=true), then acc-sm stops after test_N fails and test_N+1 does not run. If the variable is set to "false" (FAIL_ON_ERROR=false), then acc-sm continues after test_N fails and test_N+1 does run.

```
FAIL_ON_ERROR=false sh acceptance-small.sh
```



What is the PDSH variable and how it is used with acc-sm?

The PDSH variable is used to provide remote shell access. If acc-sm is run on a Lustre configuration with remote servers, specify PDSH like this:

```
PDSH="pdsh -S w" sh acceptance-small.sh
```

If the client has no access to the servers, you can run acc-sm without PDSH, but the tests which need PDSH access are skipped. A summary report is generated which lists the skipped tests.

What is the CMD configuration for HEAD?

For the HEAD branch, specify the MDSCOUNT variable (number of MDTs). By default, the variable is set to 1. If you have a Lustre configuration with several MDT nodes, they need to be specified in the configuration file as mds1_HOST, mds2_HOST, ...

By default, all of these variables are set to the mds_HOST value.

What do we do with the acc-sm test results?

Acc-sm test results are sent to Buffalo, a web interface for Lustre test results. The default Buffalo display shows a summary of tests run on different hardware configurations for various CVS branches for the past 24 hours, with links to the various reports. For more information on reporting test results to Buffalo, see [Buffalizing Tests](#).

If an acc-sm test fails, then the failure is investigated. If the investigation reveals there is a Lustre defect, then a bug is opened in Bugzilla to fix the problem.

