Milestone Completion for Implementation Milestone 1of the Distributed Namespace Project of contract SFS-DEV-001.

Revision History

Date	Revision	Author
Date (fixed)	Original	R. Henwood

Contents

Introduction....4
Subproject Description....4
Milestone Completion Criteria....4
Sanity.sh passed....4
mdsrate-create passed....4
Regression tests implemented and passed....5
Upgrade demonstrated....6
Downgrade demonstrated....8
Conclusion....8
Appendix A: sanity.sh screenshot....9
Appendix B: Downgrade demonstration....10

Introduction

The following milestone completion document applies to Subproject 2.1 – Remote Directories subproject within the OpenSFS Lustre Development contract SFS-DEV-001 signed 7/30/2011.

Subproject Description

Per the contract, DNE 1: Remote Directories is described as follows:

> *This subproject distributes the Lustre namespace over multiple metadata targets (MDTs) under administrative control using a Lustre-specific mkdir command. Whereas normal users are only able to create child directories and files on the same MDT as the parent directory, administrators can use this command to create a directory on a different MDT. The contents of any directory remain limited to a single MDT. Rename and hardlink operations between files and directories on different MDTs return EXDEV, forcing applications and utilities to treat them as if they are on different file systems. This limits the complexity of the implementation of this subproject while delivering capacity and performance scaling benefits for the entire namespace in aggregate.*
>
> *Metadata update operations that span multiple MDTs are sequenced and synchronized to create and/or increment the link count on a MDT object before it is referenced by the remote directory entry and to update the remote directory entry before decrementing the link count and/or destroying the MDT object it referenced. Although this may result in an orphan MDT object under some failure conditions, it ensures that the Lustre namespace remains intact under any and all failure scenarios. All the other metadata operations avoid synchronous I/O and execute with full performance.*
>
> *This project includes the implementation of OST FIDs (File Identifiers). These are required to overcome a limitation in the current 2.x Lustre protocol that would otherwise prevent a single file system from having more than 8 MDTs. Addressing this technical debt in the first subproject of DNE avoids protocol compatibility issues that would arise if this feature were implemented after Remote Directories were used in production.*

Milestone Completion Criteria

Per the contract, three Implementation milestones have been defined by the Whamcloud. This document is concerned with completing the first Implementation milestone which is agreed as: “Demonstrate working DNE code. The sanity.sh and mdsrate-create tests will pass in a DNE environment. Suitable new regression tests for the remote directory functionality will be added and passed, including

functional Use Cases for upgrade and downgrade.” These requirements are enumerated as:
Demonstrate working DNE code

- sanity.sh passed.
- mdsrate-create passed.
- Regression tests implemented and passed.
- Upgrade demonstrated.
- Downgrade demonstrated.

These requirements are demonstrated below.

Sanity.sh passed

The results from the sanity.sh run are recorded in maloo at:
https://maloo.whamcloud.com/test_sets/4c945302-9bf0-11e1-8837-52540035b04c
A screenshot of the page is available as Appendix A: sanity.sh screenshot.

mdsrate-create passed

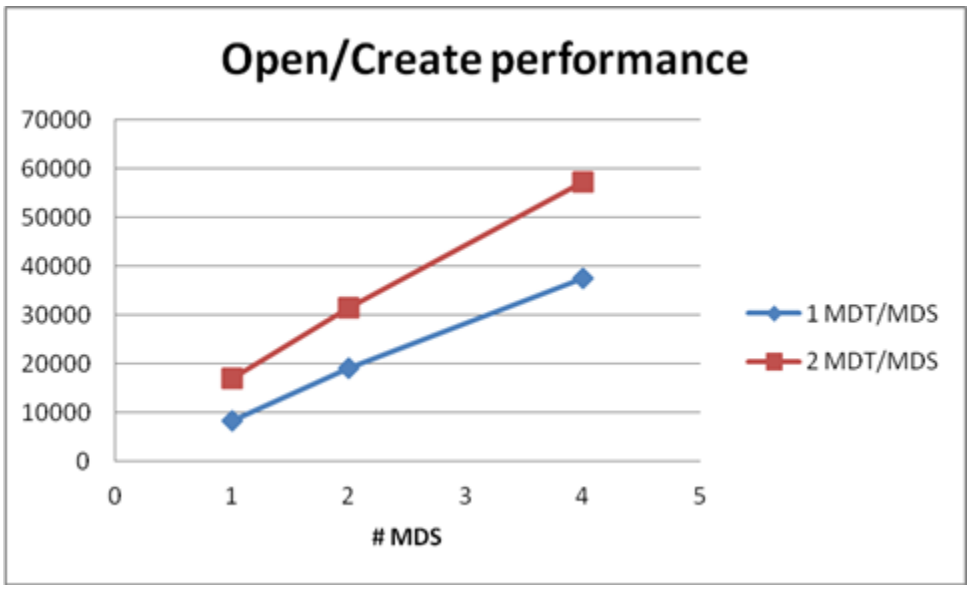

Data was collected from a Hyperion test run between April 12th and April 17th 2012. The test configuration included 100 clients, 4 threads on each client, each thread with an individual mount point. Each thread performs 10000 file open/create within unique directories. The units of the y-axis are completed file operations per second.

Regression tests implemented and passed

test_230 has been created to test DNE Remote Directories functionality:

```
test_230() {
       [ "$MDTCOUNT" -lt "2" ] && skip_env "skipping remote directory test" && return
       local MDTIDX=1

   mkdir -p $DIR/$tdir/test_230_local
       local mdt_idx=$($GETSTRIPE -M $DIR/$tdir/test_230_local)
       [ $mdt_idx -ne 0 ] &&
                  error "create local directory on wrong MDT $mdt_idx"

       $LFS setdirstripe -i $MDTIDX $DIR/$tdir/test_230 ||
              error "create remote directory failed"
       local mdt_idx=$($GETSTRIPE -M $DIR/$tdir/test_230)
       [ $mdt_idx -ne $MDTIDX ] &&
                  error "create remote directory on wrong MDT $mdt_idx"

       createmany -o $DIR/$tdir/test_230/t- 10 || error "create files on remote directory
failed"
       mdt_idx=$($GETSTRIPE -M $DIR/$tdir/test_230/t-0)
       [ $mdt_idx -ne $MDTIDX ] && error "create files on wrong MDT $mdt_idx"
       rm -r $DIR/$tdir || error "unlink remote directory failed"
}
run_test 230 "Create remote directory and files under the remote directory"
```

Upgrade demonstrated

Upgrade functionality is demonstrated as part of the test test_32c in conf-sanity: https://maloo.whamcloud.com/test_sets/5cfc5278-9d2e-11e1-8587-52540035b04c

```
== conf-sanity test 32c: Upgrade with writeconf ================== 14:09:40 (1336932580)
Loading modules from /work/orion_release/orion/lustre-dev/lustre
../libcfs/libcfs/libcfs options: 'libcfs_panic_on_lbug=0'
debug=-1
subsystem_debug=all -lnet -lnd -pinger
../lnet/lnet/lnet options: 'networks=tcp accept=all'
gss/krb5 is not supported
/work/orion_release/orion/lustre-dev/lustre/utils/tunefs.lustre
arch
commit
kernel
list
mdt
ost
sha1sums
Upgrading from disk2_2-ldiskfs.tar.bz2, created with:
  Commit: 2.2
  Kernel: 2.6.32-220.el6_lustre.g4554b65.x86_64
    Arch: x86_64
debug=-1
mount old MDT ....
mkfs new MDT....
mkfs.lustre: Warning: default mount option `errors=remount-ro' is missing
mount new MDT....
Mount client with 2 MDTs
Create the local directory and files on the old MDT
total: 10 creates in 0.04 seconds: 230.77 creates/second
Verify the MDT index of these files...Pass
Create the remote directory and files on new MDT
total: 10 creates in 0.05 seconds: 188.40 creates/second
Verify the MDT index of these files...Pass
Skip b1_8 images before we have 1.8 compatibility
Skip b1_8 images before we have 1.8 compatibility
Resetting fail_loc on all nodes...done.
```

Downgrade demonstrated

Downgrade demonstration is recorded on: http://jira.whamcloud.com/browse/OSF-102

A detailed record of the demonstration is recorded in Appendix B.

Conclusion

Implementation phase 1 has been completed according to the agreed criteria.

Appendix A: sanity.sh screenshot

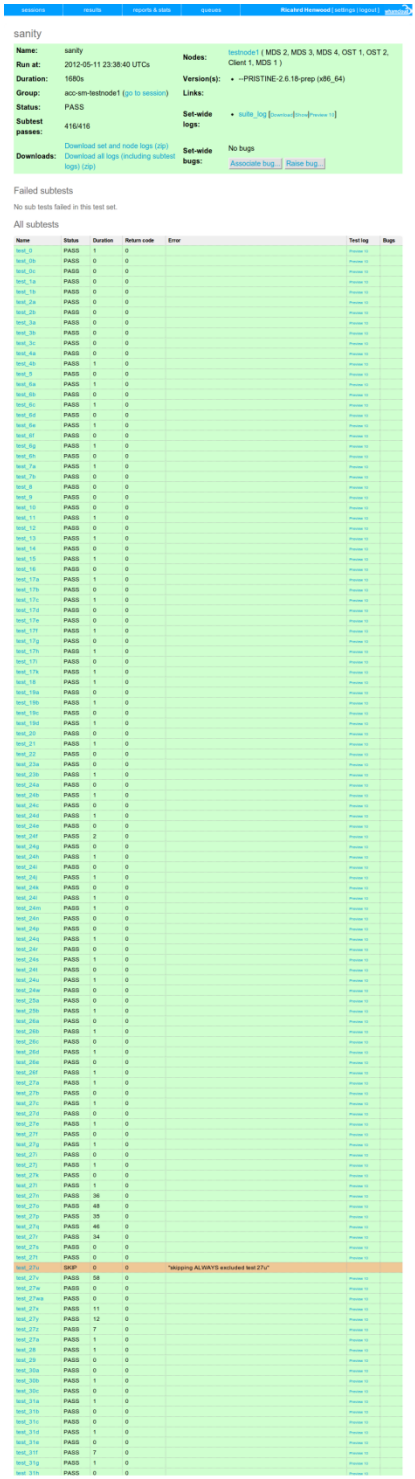

sessions | results | reports & stats | queues | Ricardo Henwood [settings | logout] | whamcloud

sanity

Name: sanity
Run at: 2012-05-11 23:38:40 UTCs
Duration: 1680s
Group: acc-sm-testnode1 (go to session)
Status: PASS
Subtest passes: 416/416
Downloads: Download set and node logs (zip) Download all logs (including subtest logs) (zip)

Nodes: testnode1 (MDS 2, MDS 3, MDS 4, OST 1, OST 2, Client 1, MDS 1)
Version(s): • --PRISTINE-2.6.18-prep (x86_64)
Links:
Set-wide logs: • suite_log [Download] [Show] [Preview 10]
Set-wide bugs: No bugs Associate bug... Raise bug...

Failed subtests

No sub tests failed in this test set.

All subtests

Name	Status	Duration	Return code	Error	Test log	Bugs
test_0	PASS	1	0		Preview 10	
test_0b	PASS	0	0		Preview 10	
test_0c	PASS	0	0		Preview 10	
test_1a	PASS	0	0		Preview 10	
test_1b	PASS	0	0		Preview 10	
test_2a	PASS	0	0		Preview 10	
test_2b	PASS	0	0		Preview 10	
test_3a	PASS	0	0		Preview 10	
test_3b	PASS	0	0		Preview 10	
test_3c	PASS	0	0		Preview 10	
test_4a	PASS	0	0		Preview 10	
test_4b	PASS	1	0		Preview 10	
test_5	PASS	0	0		Preview 10	
test_6a	PASS	1	0		Preview 10	
test_6b	PASS	0	0		Preview 10	
test_6c	PASS	1	0		Preview 10	
test_6d	PASS	0	0		Preview 10	
test_6e	PASS	1	0		Preview 10	
test_6f	PASS	0	0		Preview 10	
test_6g	PASS	1	0		Preview 10	
test_6h	PASS	0	0		Preview 10	
test_7a	PASS	1	0		Preview 10	
test_7b	PASS	0	0		Preview 10	
test_8	PASS	0	0		Preview 10	
test_9	PASS	0	0		Preview 10	
test_10	PASS	0	0		Preview 10	
test_11	PASS	1	0		Preview 10	
test_12	PASS	0	0		Preview 10	
test_13	PASS	1	0		Preview 10	
test_14	PASS	0	0		Preview 10	
test_15	PASS	1	0		Preview 10	
test_16	PASS	0	0		Preview 10	
test_17a	PASS	1	0		Preview 10	
test_17b	PASS	0	0		Preview 10	
test_17c	PASS	1	0		Preview 10	
test_17d	PASS	0	0		Preview 10	
test_17e	PASS	0	0		Preview 10	
test_17f	PASS	1	0		Preview 10	
test_17g	PASS	0	0		Preview 10	
test_17h	PASS	1	0		Preview 10	
test_17i	PASS	0	0		Preview 10	
test_17k	PASS	1	0		Preview 10	
test_18	PASS	1	0		Preview 10	
test_19a	PASS	0	0		Preview 10	
test_19b	PASS	1	0		Preview 10	
test_19c	PASS	0	0		Preview 10	
test_19d	PASS	1	0		Preview 10	
test_20	PASS	0	0		Preview 10	
test_21	PASS	0	0		Preview 10	
test_22	PASS	0	0		Preview 10	
test_23a	PASS	0	0		Preview 10	
test_23b	PASS	1	0		Preview 10	
test_24a	PASS	0	0		Preview 10	
test_24b	PASS	1	0		Preview 10	
test_24c	PASS	0	0		Preview 10	
test_24d	PASS	1	0		Preview 10	
test_24e	PASS	0	0		Preview 10	
test_24f	PASS	2	0		Preview 10	
test_24g	PASS	0	0		Preview 10	
test_24h	PASS	1	0		Preview 10	
test_24i	PASS	0	0		Preview 10	
test_24j	PASS	1	0		Preview 10	
test_24k	PASS	0	0		Preview 10	
test_24l	PASS	1	0		Preview 10	
test_24m	PASS	1	0		Preview 10	
test_24n	PASS	0	0		Preview 10	
test_24p	PASS	0	0		Preview 10	
test_24q	PASS	1	0		Preview 10	
test_24r	PASS	0	0		Preview 10	
test_24s	PASS	1	0		Preview 10	
test_24t	PASS	0	0		Preview 10	
test_24u	PASS	1	0		Preview 10	
test_24w	PASS	0	0		Preview 10	
test_25a	PASS	0	0		Preview 10	
test_25b	PASS	1	0		Preview 10	
test_26a	PASS	0	0		Preview 10	
test_26b	PASS	1	0		Preview 10	
test_26c	PASS	0	0		Preview 10	
test_26d	PASS	1	0		Preview 10	
test_26e	PASS	0	0		Preview 10	
test_26f	PASS	1	0		Preview 10	
test_27a	PASS	1	0		Preview 10	
test_27b	PASS	0	0		Preview 10	
test_27c	PASS	1	0		Preview 10	
test_27d	PASS	0	0		Preview 10	
test_27e	PASS	1	0		Preview 10	
test_27f	PASS	0	0		Preview 10	
test_27g	PASS	1	0		Preview 10	
test_27i	PASS	0	0		Preview 10	
test_27j	PASS	1	0		Preview 10	
test_27k	PASS	0	0		Preview 10	
test_27l	PASS	1	0		Preview 10	
test_27m	PASS	36	0		Preview 10	
test_27n	PASS	48	0		Preview 10	
test_27p	PASS	35	0		Preview 10	
test_27q	PASS	46	0		Preview 10	
test_27r	PASS	34	0		Preview 10	
test_27s	PASS	0	0		Preview 10	
test_27t	PASS	0	0		Preview 10	
test_27u	SKIP	0	0	"skipping ALWAYS excluded test 27u"		
test_27v	PASS	58	0		Preview 10	
test_27w	PASS	0	0		Preview 10	
test_27wa	PASS	0	0		Preview 10	
test_27x	PASS	11	0		Preview 10	
test_27y	PASS	12	0		Preview 10	
test_27z	PASS	7	0		Preview 10	
test_27za	PASS	1	0		Preview 10	
test_28	PASS	1	0		Preview 10	
test_29	PASS	0	0		Preview 10	
test_30a	PASS	0	0		Preview 10	
test_30b	PASS	1	0		Preview 10	
test_30c	PASS	0	0		Preview 10	
test_31a	PASS	1	0		Preview 10	
test_31b	PASS	0	0		Preview 10	
test_31c	PASS	0	0		Preview 10	
test_31d	PASS	1	0		Preview 10	
test_31e	PASS	0	0		Preview 10	
test_31f	PASS	7	0		Preview 10	
test_31g	PASS	1	0		Preview 10	
test_31h	PASS	0	0		Preview 10	
test_31i	PASS	1	0		Preview 10	
test_31j	PASS	0	0		Preview 10	
test_31k	PASS	1	0		Preview 10	
test_31m	PASS	0	0		Preview 10	
test_32a	PASS	0	0		Preview 10	
test_32b	PASS	1	0		Preview 10	
test_32c	PASS	0	0		Preview 10	
test_32d	PASS	1	0		Preview 10	
test_32e	PASS	1	0		Preview 10	
test_32f	PASS	0	0		Preview 10	
test_32g	PASS	1	0		Preview 10	
test_32h	PASS	0	0		Preview 10	
test_32i	PASS	0	0		Preview 10	
test_32j	PASS	1	0		Preview 10	
test_32k	PASS	0	0		Preview 10	
test_32l	PASS	0	0		Preview 10	
test_32m	PASS	1	0		Preview 10	
test_32n	PASS	0	0		Preview 10	
test_32o	PASS	1	0		Preview 10	
test_32p	PASS	1	0		Preview 10	
test_32q	PASS	0	0		Preview 10	
test_32r	PASS	0	0		Preview 10	
test_33	PASS	1	0		Preview 10	
test_33a	PASS	1	0		Preview 10	
test_33b	PASS	0	0		Preview 10	
test_33c	PASS	2	0		Preview 10	
test_34a	PASS	0	0		Preview 10	
test_34b	PASS	1	0		Preview 10	
test_34c	PASS	0	0		Preview 10	
test_34d	PASS	1	0		Preview 10	
test_34e	PASS	0	0		Preview 10	
test_34f	PASS	0	0		Preview 10	
test_34g	PASS	1	0		Preview 10	
test_34h	PASS	4	0		Preview 10	
test_35a	PASS	0	0		Preview 10	
test_36a	PASS	0	0		Preview 10	
test_36b	PASS	1	0		Preview 10	
test_36c	PASS	0	0		Preview 10	
test_36d	PASS	1	0		Preview 10	
test_36e	PASS	0	0		Preview 10	
test_37	PASS	1	0		Preview 10	
test_38	PASS	1	0		Preview 10	
test_39	PASS	2	0		Preview 10	
test_39b	PASS	1	0		Preview 10	
test_39c	PASS	2	0		Preview 10	
test_39d	PASS	1	0		Preview 10	
test_39e	PASS	1	0		Preview 10	
test_39f	PASS	3	0		Preview 10	
test_39g	PASS	2	0		Preview 10	
test_39h	PASS	2	0		Preview 10	
test_39i	PASS	2	0		Preview 10	
test_39k	SKIP	0	0	"skipping ALWAYS excluded test 39k"		
test_39l	PASS	2	0		Preview 10	
test_39l	PASS	7	0		Preview 10	
test_39m	PASS	3	0		Preview 10	
test_40	PASS	0	0		Preview 10	
test_41	PASS	2	0		Preview 10	
test_42a	SKIP	0	0	"skipping ALWAYS excluded test 42a"		
test_42b	SKIP	0	0	"skipping ALWAYS excluded test 42b"		
test_42c	SKIP	0	0	"skipping ALWAYS excluded test 42c"		
test_42d	SKIP	0	0	"skipping ALWAYS excluded test 42d"		
test_42e	PASS	47	0		Preview 10	
test_43	PASS	3	0		Preview 10	
test_43a	PASS	1	0		Preview 10	
test_43b	PASS	1	0		Preview 10	
test_43c	PASS	1	0		Preview 10	
test_44	PASS	1	0		Preview 10	
test_44a	PASS	2	0		Preview 10	
test_45	SKIP	0	0	"skipping ALWAYS excluded test 45"		
test_46	PASS	3	0		Preview 10	
test_48a	PASS	0	0		Preview 10	
test_48b	PASS	1	0		Preview 10	
test_48c	PASS	1	0		Preview 10	
test_48d	PASS	0	0		Preview 10	
test_48e	PASS	0	0		Preview 10	
test_50	PASS	0	0		Preview 10	
test_51a	PASS	1	0		Preview 10	
test_51be	PASS	0	0		Preview 10	
test_51bb	PASS	6	0			
test_51d	SKIP	0	0	"skipping ALWAYS excluded test 51d"		
test_52a	PASS	1	0		Preview 10	
test_52b	PASS	1	0		Preview 10	
test_53	PASS	2	0		Preview 10	
test_54a	PASS	0	0		Preview 10	
test_54b	PASS	0	0		Preview 10	
test_54c	PASS	1	0		Preview 10	
test_54d	PASS	0	0		Preview 10	
test_54e	PASS	0	0		Preview 10	
test_56a	PASS	1	0		Preview 10	
test_56g	PASS	1	0		Preview 10	
test_56h	PASS	0	0		Preview 10	
test_56i	PASS	1	0		Preview 10	
test_56j	PASS	0	0		Preview 10	
test_56k	PASS	0	0		Preview 10	
test_56l	PASS	1	0		Preview 10	
test_56m	PASS	0	0		Preview 10	
test_56o	PASS	1	0		Preview 10	
test_56p	PASS	2	0		Preview 10	
test_56q	PASS	1	0		Preview 10	
test_56s	PASS	1	0		Preview 10	
test_56v	PASS	2	0		Preview 10	
test_56x	PASS	1	0		Preview 10	
test_56t	PASS	1	0		Preview 10	
test_56u	PASS	2	0		Preview 10	
test_56w	PASS	10	0		Preview 10	
test_57a	PASS	1	0		Preview 10	
test_57b	PASS	5	0		Preview 10	
test_58	PASS	0	0		Preview 10	
test_59	PASS	5	0		Preview 10	
test_60a	PASS	55	0		Preview 10	
test_60b	PASS	0	0		Preview 10	
test_60c	PASS	1	0		Preview 10	
test_61	PASS	5	0		Preview 10	
test_64a	PASS	0	0		Preview 10	
test_65a	PASS	0	0		Preview 10	
test_65b	PASS	1	0		Preview 10	
test_65c	PASS	1	0		Preview 10	
test_65d	PASS	0	0		Preview 10	
test_65e	PASS	0	0		Preview 10	
test_65f	PASS	0	0		Preview 10	
test_65g	PASS	1	0		Preview 10	
test_65h	PASS	0	0		Preview 10	
test_65i	PASS	0	0		Preview 10	
test_65la	PASS	1	0		Preview 10	
test_65lb	PASS	1	0		Preview 10	
test_65lc	PASS	16	0		Preview 10	
test_65j	PASS	2	0		Preview 10	
test_65k	PASS	6	0		Preview 10	
test_65l	PASS	1	0		Preview 10	
test_66	PASS	1	0		Preview 10	
test_68b	SKIP	0	0	"skipping ALWAYS excluded test 68b"		
test_69	PASS	2	0		Preview 10	
test_72a	PASS	7	0		Preview 10	
test_72b	PASS	1	0		Preview 10	
test_74a	PASS	1	0		Preview 10	
test_74b	PASS	1	0		Preview 10	
test_74c	PASS	1	0		Preview 10	
test_76	SKIP	0	0	"skipping ALWAYS excluded test 76"		
test_77a	PASS	5	0		Preview 10	
test_77b	PASS	1	0		Preview 10	
test_77c	PASS	5	0		Preview 10	
test_77d	PASS	1	0		Preview 10	
test_77e	PASS	1	0		Preview 10	
test_77g	PASS	2	0		Preview 10	
test_77h	PASS	2	0		Preview 10	
test_77i	PASS	2	0		Preview 10	
test_77j	PASS	4	0		Preview 10	
test_79	PASS	1	0		Preview 10	
test_80	PASS	3	0		Preview 10	
test_81a	PASS	0	0		Preview 10	
test_81b	PASS	1	0		Preview 10	
test_99a	PASS	2	0		Preview 10	
test_99b	PASS	3	0		Preview 10	
test_99c	PASS	8	0		Preview 10	
test_99d	PASS	0	0		Preview 10	
test_99e	PASS	2	0		Preview 10	
test_99f	PASS	1	0		Preview 10	
test_100	PASS	1	0		Preview 10	
test_101b	PASS	90	0		Preview 10	
test_101c	PASS	37	0		Preview 10	
test_101d	PASS	114	0		Preview 10	
test_101e	PASS	90	0		Preview 10	
test_101f	PASS	5	0		Preview 10	
test_102a	PASS	0	0		Preview 10	
test_102b	PASS	0	0		Preview 10	
test_102c	PASS	1	0		Preview 10	
test_102d	PASS	1	0		Preview 10	
test_102f	PASS	0	0		Preview 10	
test_102h	PASS	1	0		Preview 10	
test_102ha	PASS	1	0		Preview 10	
test_102i	PASS	0	0		Preview 10	
test_102j	PASS	0	0		Preview 10	
test_102k	PASS	1	0		Preview 10	
test_104a	PASS	1	0		Preview 10	
test_104b	PASS	1	0		Preview 10	
test_105a	PASS	0	0		Preview 10	
test_105b	PASS	1	0		Preview 10	
test_105c	PASS	0	0		Preview 10	
test_105d	PASS	11	0		Preview 10	
test_105e	PASS	1	0		Preview 10	
test_106	PASS	0	0		Preview 10	
test_107	PASS	2	0		Preview 10	
test_110	PASS	1	0		Preview 10	
test_116	PASS	66	0		Preview 10	
test_117	PASS	0	0		Preview 10	
test_118a	PASS	3	0		Preview 10	
test_118b	PASS	2	0		Preview 10	
test_118c	PASS	5	0		Preview 10	
test_118d	PASS	6	0		Preview 10	
test_118f	PASS	2	0		Preview 10	
test_118g	PASS	2	0		Preview 10	
test_118h	PASS	12	0		Preview 10	
test_118i	PASS	8	0		Preview 10	
test_118j	PASS	1	0		Preview 10	
test_118k	PASS	13	0		Preview 10	
test_118l	PASS	1	0		Preview 10	
test_118o	PASS	0	0		Preview 10	
test_119b	PASS	10	0		Preview 10	
test_119c	PASS	1	0		Preview 10	
test_119d	PASS	4	0		Preview 10	
test_120a	PASS	0	0		Preview 10	
test_120b	PASS	0	0		Preview 10	
test_120c	PASS	1	0		Preview 10	
test_120d	PASS	1	0		Preview 10	
test_120e	PASS	0	0		Preview 10	
test_120f	PASS	0	0		Preview 10	
test_121	PASS	0	0		Preview 10	
test_123a	PASS	36	0		Preview 10	
test_123b	PASS	26	0		Preview 10	
test_124a	SKIP	0	0	"skipping ALWAYS excluded test 124a"		
test_124b	SKIP	0	0	"skipping SLOW test 124b"		
test_125	PASS	0	0		Preview 10	
test_126	PASS	1	0		Preview 10	
test_127a	PASS	1	0		Preview 10	
test_127b	PASS	1	0		Preview 10	
test_128	PASS	0	0		Preview 10	
test_129	PASS	6	0		Preview 10	
test_130a	PASS	0	0		Preview 10	
test_130b	PASS	0	0		Preview 10	
test_130c	PASS	1	0		Preview 10	
test_130d	PASS	1	0		Preview 10	
test_130e	PASS	18	0		Preview 10	
test_131a	PASS	1	0		Preview 10	
test_131b	PASS	1	0		Preview 10	
test_131c	PASS	1	0		Preview 10	
test_131d	PASS	0	0		Preview 10	
test_131e	PASS	1	0		Preview 10	
test_132	SKIP	0	0	"skipping ALWAYS excluded test 132"		
test_133a	PASS	5	0		Preview 10	
test_133b	PASS	0	0		Preview 10	
test_133c	PASS	15	0		Preview 10	
test_133d	SKIP	0	0	"skipping ALWAYS excluded test 133d"		
test_140	PASS	1	0		Preview 10	
test_150	PASS	2	0		Preview 10	
test_151	PASS	1	0		Preview 10	
test_152	PASS	2	0		Preview 10	
test_153	PASS	1	0		Preview 10	
test_154	PASS	1	0		Preview 10	
test_155a	PASS	1	0		Preview 10	
test_155b	PASS	1	0		Preview 10	
test_155c	PASS	1	0		Preview 10	
test_155d	PASS	1	0		Preview 10	
test_155e	PASS	16	0		Preview 10	
test_155f	PASS	17	0		Preview 10	
test_155g	PASS	19	0		Preview 10	
test_155h	PASS	16	0		Preview 10	
test_156	PASS	7	0		Preview 10	
test_160	PASS	1	0		Preview 10	
test_161	PASS	12	0		Preview 10	
test_162	PASS	1	0		Preview 10	
test_163	PASS	2	0		Preview 10	
test_169	PASS	1	0		Preview 10	
test_170	PASS	2	0		Preview 10	
test_171	PASS	4	0		Preview 10	
test_180a	PASS	1	0		Preview 10	
test_180b	PASS	1	0		Preview 10	
test_181	PASS	46	0		Preview 10	
test_200a	PASS	11	0		Preview 10	
test_200b	PASS	10	0		Preview 10	
test_200c	PASS	1	0		Preview 10	
test_200d	PASS	0	0		Preview 10	
test_200e	PASS	0	0		Preview 10	
test_200f	PASS	1	0		Preview 10	
test_200g	PASS	1	0		Preview 10	
test_200h	PASS	0	0		Preview 10	
test_201a	PASS	2	0		Preview 10	
test_201b	PASS	0	0		Preview 10	
test_201c	PASS	7	0		Preview 10	
test_204a	PASS	1	0		Preview 10	
test_204b	PASS	0	0		Preview 10	
test_204c	PASS	1	0		Preview 10	
test_204d	PASS	0	0		Preview 10	
test_204e	PASS	0	0		Preview 10	
test_204f	PASS	1	0		Preview 10	
test_204g	PASS	0	0		Preview 10	
test_204h	PASS	1	0		Preview 10	
test_212	PASS	2	0		Preview 10	
test_213	PASS	1	0		Preview 10	
test_214	PASS	5	0		Preview 10	
test_215	PASS	1	0		Preview 10	
test_216	PASS	1	0		Preview 10	
test_217	PASS	0	0		Preview 10	
test_218	PASS	1	0		Preview 10	
test_219	PASS	0	0		Preview 10	
test_220	PASS	31	0		Preview 10	
test_221	PASS	1	0		Preview 10	
test_222a	PASS	1	0		Preview 10	
test_223b	PASS	1	0		Preview 10	
test_223	PASS	1	0		Preview 10	
test_224a	PASS	2	0		Preview 10	
test_224b	PASS	8	0		Preview 10	
test_225a	SKIP	0	0	"skipping ALWAYS excluded test 225a (base 225)"		
test_225b	SKIP	0	0	"skipping ALWAYS excluded test 225b (base 225)"		
test_230	PASS	0	0		Preview 10	
test_900	PASS	73	0		Preview 10	

Appendix B: Downgrade demonstration

```
[root@testnode1 tests]# sh -vx upgrade_downgrade.sh
#Upgrade tests
#go to lustre 2.3 make old ldiskfs
cd /work/lustre-2.3/lustre/tests
+ cd /work/lustre-2.3/lustre/tests
testnode=${testnode:-"`hostname`"}
hostname
++ hostname
+ testnode=testnode1
LOAD=y sh llmount.sh
+ LOAD=y
+ sh llmount.sh
Loading modules from /work/lustre-2.3/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -pinger
../lnet/lnet/lnet options: 'networks=tcp accept=all'
gss/krb5 is not supported
quota/lquota options: 'hash_lqs_cur_bits=3'
set -vx
+ set -vx
../utils/mkfs.lustre --reformat --mgs --mdt --device-size=1048576 /tmp/lustre-mdt-2.3
+ ../utils/mkfs.lustre --reformat --mgs --mdt --device-size=1048576 /tmp/lustre-mdt-2.3
Permanent disk data:
Target: lustre-MDTffff
Index: unassigned
Lustre FS: lustre
Mount type: ldiskfs
Flags: 0x75
(MDT MGS needs_index first_time update )
Persistent mount opts: user_xattr,errors=remount-ro
Parameters:
formatting backing filesystem ldiskfs on /dev/loop0
target name lustre-MDTffff
4k blocks 262144
options -I 512 -i 2048 -q -O dirdata,uninit_bg,dir_nlink,huge_file,flex_bg -E
lazy_journal_init -F
mkfs_cmd = mke2fs -j -b 4096 -L lustre-MDTffff -I 512 -i 2048 -q -O
dirdata,uninit_bg,dir_nlink,huge_file,flex_bg -E lazy_journal_init -F /dev/loop0 262144
Writing CONFIGS/mountdata
../utils/mkfs.lustre --reformat --mgsnode=$testnode --ost --device-size=1048576
/tmp/lustre-ost-2.3
+ ../utils/mkfs.lustre --reformat --mgsnode=testnode1 --ost --device-size=1048576
/tmp/lustre-ost-2.3
Permanent disk data:
Target: lustre-OSTffff
Index: unassigned
Lustre FS: lustre
Mount type: ldiskfs
Flags: 0x72
(OST needs_index first_time update )
Persistent mount opts: errors=remount-ro,extents,mballoc
Parameters: mgsnode=192.168.122.162@tcp
formatting backing filesystem ldiskfs on /dev/loop0
target name lustre-OSTffff
4k blocks 262144
options -I 256 -q -O extents,uninit_bg,dir_nlink,huge_file,flex_bg -G 256 -E
resize=4290772992,lazy_journal_init -F
mkfs_cmd = mke2fs -j -b 4096 -L lustre-OSTffff -I 256 -q -O
extents,uninit_bg,dir_nlink,huge_file,flex_bg -G 256 -E
resize=4290772992,lazy_journal_init -F /dev/loop0 262144
Writing CONFIGS/mountdata
mount -t lustre -o loop,user_xattr,acl /tmp/lustre-mdt-2.3 /mnt/mds
+ mount -t lustre -o loop,user_xattr,acl /tmp/lustre-mdt-2.3 /mnt/mds
mount -t lustre -o loop /tmp/lustre-ost-2.3 /mnt/ost1
+ mount -t lustre -o loop /tmp/lustre-ost-2.3 /mnt/ost1
```

```
mount -t lustre $testnode:/lustre /mnt/lustre
+ mount -t lustre testnode1:/lustre /mnt/lustre
cp /etc/fstab /mnt/lustre
+ cp /etc/fstab /mnt/lustre
cp /etc/hosts /mnt/lustre
+ cp /etc/hosts /mnt/lustre
umount /mnt/lustre
+ umount /mnt/lustre
umount /mnt/ost1
+ umount /mnt/ost1
umount /mnt/mds
+ umount /mnt/mds
losetup -d /dev/loop0
+ losetup -d /dev/loop0
losetup -d /dev/loop1
+ losetup -d /dev/loop1
losetup -d /dev/loop2
+ losetup -d /dev/loop2
ioctl: LOOP_CLR_FD: No such device or address
losetup -d /dev/loop3
+ losetup -d /dev/loop3
ioctl: LOOP_CLR_FD: No such device or address
losetup -d /dev/loop4
+ losetup -d /dev/loop4
ioctl: LOOP_CLR_FD: No such device or address
losetup -d /dev/loop5
+ losetup -d /dev/loop5
ioctl: LOOP_CLR_FD: No such device or address
losetup -d /dev/loop6
+ losetup -d /dev/loop6
ioctl: LOOP_CLR_FD: No such device or address
losetup -d /dev/loop7
+ losetup -d /dev/loop7
ioctl: LOOP_CLR_FD: No such device or address
LOAD=y sh llmountcleanup.sh
+ LOAD=y
+ sh llmountcleanup.sh
Stopping clients: testnode1 /mnt/lustre (opts:-f)
Stopping clients: testnode1 /mnt/lustre2 (opts:-f)
osd_ldiskfs 296768 0
fsfilt_ldiskfs 119600 0
mdd 426496 3 osd_ldiskfs,cmm,mdt
ldiskfs 354264 2 osd_ldiskfs,fsfilt_ldiskfs
jbd2 101384 3 osd_ldiskfs,fsfilt_ldiskfs,ldiskfs
crc16 35328 1 ldiskfs
obdclass 1109104 29
llite_lloop,lustre,obdfilter,ost,osd_ldiskfs,cmm,fsfilt_ldiskfs,mdt,mdd,mds,mgs,mgc,lov,
osc,mdc,lmv,fid,fld,lquota,ptlrpc
lvfs 72256 22
llite_lloop,lustre,obdfilter,ost,osd_ldiskfs,cmm,fsfilt_ldiskfs,mdt,mdd,mds,mgs,mgc,lov,
osc,mdc,lmv,fid,fld,lquota,ptlrpc,obdclass
libcfs 344320 24
llite_lloop,lustre,obdfilter,ost,osd_ldiskfs,cmm,fsfilt_ldiskfs,mdt,mdd,mds,mgs,mgc,lov,
osc,mdc,lmv,fid,fld,lquota,ptlrpc,obdclass,lvfs,ksocklnd,lnet
exportfs 39296 2 fsfilt_ldiskfs,nfsd
modules unloaded.
echo "go to DNE branch do upgrade"
+ echo 'go to DNE branch do upgrade'
go to DNE branch do upgrade
cd /work/lustre-dne/lustre/tests
+ cd /work/lustre-dne/lustre/tests
LOAD=y sh llmount.sh
+ LOAD=y
+ sh llmount.sh
Loading modules from /work/lustre-dne/lustre/tests/..
../libcfs/libcfs/libcfs options: 'libcfs_panic_on_lbug=0'
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -pinger
../lnet/lnet/lnet options: 'networks=tcp accept=all'
gss/krb5 is not supported
mount -t lustre -o loop,user_xattr,acl,abort_recov,write_conf /tmp/lustre-mdt-2.3
/mnt/mds1
```

```
+ mount -t lustre -o loop,user_xattr,acl,abort_recov,write_conf /tmp/lustre-mdt-2.3
/mnt/mds1
../utils/mkfs.lustre --reformat --mgsnode=$testnode --mdt --device-size=1048576 --index
1 /tmp/lustre-mdt-new
+ ../utils/mkfs.lustre --reformat --mgsnode=testnode1 --mdt --device-size=1048576
--index 1 /tmp/lustre-mdt-new
Permanent disk data:
Target: lustre:MDT0001
Index: 1
Lustre FS: lustre
Mount type: ldiskfs
Flags: 0x61
(MDT first_time update )
Persistent mount opts: user_xattr,errors=remount-ro
Parameters: mgsnode=192.168.122.162@tcp
formatting backing filesystem ldiskfs on /dev/loop1
target name lustre:MDT0001
4k blocks 262144
options -I 512 -i 2048 -q -O dirdata,uninit_bg,dir_nlink,huge_file,flex_bg -E
lazy_journal_init -F
mkfs_cmd = mke2fs -j -b 4096 -L lustre:MDT0001 -I 512 -i 2048 -q -O
dirdata,uninit_bg,dir_nlink,huge_file,flex_bg -E lazy_journal_init -F /dev/loop1 262144
Writing CONFIGS/mountdata
mount -t lustre -o loop,user_xattr,acl,abort_recov,write_conf /tmp/lustre-mdt-new
/mnt/mds2
+ mount -t lustre -o loop,user_xattr,acl,abort_recov,write_conf /tmp/lustre-mdt-new
/mnt/mds2
mount -t lustre -o loop,abort_recov,write_conf /tmp/lustre-ost-2.3 /mnt/ost1
+ mount -t lustre -o loop,abort_recov,write_conf /tmp/lustre-ost-2.3 /mnt/ost1
mount -t lustre $testnode:/lustre /mnt/lustre
+ mount -t lustre testnode1:/lustre /mnt/lustre
diff /mnt/lustre/fstab /etc/fstab || { echo "the file is diff1" && exit 1; }
+ diff /mnt/lustre/fstab /etc/fstab
diff /mnt/lustre/hosts /etc/hosts || { echo "the file is diff1" && exit 1; }
+ diff /mnt/lustre/hosts /etc/hosts
../utils/lfs setdirstripe -i 1 /mnt/lustre/test_mdt1 || { echo "create remote directory
failed" && exit 1; }
+ ../utils/lfs setdirstripe -i 1 /mnt/lustre/test_mdt1
mdt_idx=$(../utils/lfs getstripe -M /mnt/lustre/test_mdt1)
../utils/lfs getstripe -M /mnt/lustre/test_mdt1
++ ../utils/lfs getstripe -M /mnt/lustre/test_mdt1
+ mdt_idx=1
[ $mdt_idx -ne 1 ] && { echo "create remote directory on wrong MDT" && exit 1; }
+ '[' 1 -ne 1 ']'

mkdir /mnt/lustre/test_mdt1/dir
+ mkdir /mnt/lustre/test_mdt1/dir
mdt_idx=$(../utils/lfs getstripe -M /mnt/lustre/test_mdt1/dir)
../utils/lfs getstripe -M /mnt/lustre/test_mdt1/dir
++ ../utils/lfs getstripe -M /mnt/lustre/test_mdt1/dir
+ mdt_idx=1
[ $mdt_idx -ne 1 ] && { echo "create remote directory on wrong MDT" && exit 1; }
+ '[' 1 -ne 1 ']'
cp /mnt/lustre/hosts /mnt/lustre/test_mdt1/dir/hosts
+ cp /mnt/lustre/hosts /mnt/lustre/test_mdt1/dir/hosts
cp /mnt/lustre/fstab /mnt/lustre/test_mdt1/dir/fstab
+ cp /mnt/lustre/fstab /mnt/lustre/test_mdt1/dir/fstab
echo "downgrade DNE to single MDT"
+ echo 'downgrade DNE to single MDT'
downgrade DNE to single MDT
mkdir /mnt/lustre/test_mdt1_backup
+ mkdir /mnt/lustre/test_mdt1_backup
cp -R /mnt/lustre/test_mdt1/ /mnt/lustre/test_mdt1_backup/
+ cp -R /mnt/lustre/test_mdt1/ /mnt/lustre/test_mdt1_backup/
../utils/lctl dk > /tmp/debug.out
+ ../utils/lctl dk
umount /mnt/lustre/
+ umount /mnt/lustre/
umount /mnt/mds2
+ umount /mnt/mds2
umount /mnt/mds1
+ umount /mnt/mds1
```

```
umount /mnt/ost1
+ umount /mnt/ost1
LOAD=y sh llmountcleanup.sh
+ LOAD=y
+ sh llmountcleanup.sh
Stopping clients: testnode1 /mnt/lustre (opts:-f)
Stopping clients: testnode1 /mnt/lustre2 (opts:-f)
osd_ldiskfs 343056 0
fsfilt_ldiskfs 43776 0
ldiskfs 354392 2 osd_ldiskfs,fsfilt_ldiskfs
mdd 338320 3 osd_ldiskfs,cmm,mdt
fld 113776 7 osd_ldiskfs,lod,obdfilter,cmm,mdt,lmv,fid
obdclass 1221936 28
llite_lloop,lustre,osd_ldiskfs,osp,lod,obdfilter,ost,cmm,mdt,mdd,mgs,mgc,lov,osc,mdc,lmv
,fid,fld,ptlrpc
lvfs 59024 22
llite_lloop,lustre,osd_ldiskfs,fsfilt_ldiskfs,osp,lod,obdfilter,ost,cmm,mdt,mdd,mgs,mgc,
lov,osc,mdc,lmv,fid,fld,ptlrpc,obdclass
libcfs 344192 24
llite_lloop,lustre,osd_ldiskfs,fsfilt_ldiskfs,osp,lod,obdfilter,ost,cmm,mdt,mdd,mgs,mgc,
lov,osc,mdc,lmv,fid,fld,ptlrpc,obdclass,lvfs,ksocklnd,lnet
jbd2 101384 2 osd_ldiskfs,ldiskfs
crc16 35328 1 ldiskfs
modules unloaded.
losetup -d /dev/loop0
+ losetup -d /dev/loop0
losetup -d /dev/loop1
+ losetup -d /dev/loop1
losetup -d /dev/loop2
+ losetup -d /dev/loop2
losetup -d /dev/loop3
+ losetup -d /dev/loop3
ioctl: LOOP_CLR_FD: No such device or address
losetup -d /dev/loop4
+ losetup -d /dev/loop4
ioctl: LOOP_CLR_FD: No such device or address
losetup -d /dev/loop5
+ losetup -d /dev/loop5
ioctl: LOOP_CLR_FD: No such device or address
losetup -d /dev/loop6
+ losetup -d /dev/loop6
ioctl: LOOP_CLR_FD: No such device or address
losetup -d /dev/loop7
+ losetup -d /dev/loop7
ioctl: LOOP_CLR_FD: No such device or address
cd /work/lustre-2.3/lustre/tests
+ cd /work/lustre-2.3/lustre/tests
LOAD=y sh llmount.sh
+ LOAD=y
+ sh llmount.sh
Loading modules from /work/lustre-2.3/lustre/tests/..
debug=vfstrace rpctrace dlmtrace neterror ha config ioctl super
subsystem_debug=all -lnet -lnd -pinger
../lnet/lnet/lnet options: 'networks=tcp accept=all'
gss/krb5 is not supported
quota/lquota options: 'hash_lqs_cur_bits=3'
/work/lustre-2.3/lustre/utils/tunefs.lustre --writeconf /tmp/lustre-mdt-2.3
+ /work/lustre-2.3/lustre/utils/tunefs.lustre --writeconf /tmp/lustre-mdt-2.3
checking for existing Lustre data: found CONFIGS/mountdata
Reading CONFIGS/mountdata
Read previous values:
Target: lustre-MDT0000
Index: 0
Lustre FS: lustre
Mount type: ldiskfs
Flags: 0x5
(MDT MGS )
Persistent mount opts: user_xattr,errors=remount-ro
Parameters:
Permanent disk data:
Target: lustre-MDT0000
Index: 0
```

```
Lustre FS: lustre
Mount type: ldiskfs
Flags: 0x105
(MDT MGS writeconf )
Persistent mount opts: user_xattr,errors=remount-ro
Parameters:
Writing CONFIGS/mountdata
/work/lustre-2.3/lustre/utils/tunefs.lustre --writeconf /tmp/lustre-ost-2.3
+ /work/lustre-2.3/lustre/utils/tunefs.lustre --writeconf /tmp/lustre-ost-2.3
checking for existing Lustre data: found CONFIGS/mountdata
Reading CONFIGS/mountdata
Read previous values:
Target: lustre-OST0000
Index: 0
Lustre FS: lustre
Mount type: ldiskfs
Flags: 0x2
(OST )
Persistent mount opts: errors=remount-ro,extents,mballoc
Parameters: mgsnode=192.168.122.162@tcp
Permanent disk data:
Target: lustre-OST0000
Index: 0
Lustre FS: lustre
Mount type: ldiskfs
Flags: 0x102
(OST writeconf )
Persistent mount opts: errors=remount-ro,extents,mballoc
Parameters: mgsnode=192.168.122.162@tcp
Writing CONFIGS/mountdata
mount -t lustre -o loop,user_xattr,acl,abort_recov /tmp/lustre-mdt-2.3 /mnt/mds
+ mount -t lustre -o loop,user_xattr,acl,abort_recov /tmp/lustre-mdt-2.3 /mnt/mds
mount -t lustre -o loop,abort_recov /tmp/lustre-ost-2.3 /mnt/ost1
+ mount -t lustre -o loop,abort_recov /tmp/lustre-ost-2.3 /mnt/ost1
mount -t lustre $testnode:/lustre /mnt/lustre
+ mount -t lustre testnode1:/lustre /mnt/lustre
diff /mnt/lustre/fstab /etc/fstab || { echo "the file is different" && exit 1; }
+ diff /mnt/lustre/fstab /etc/fstab
diff /mnt/lustre/hosts /etc/hosts || { echo "the file is different" && exit 1; }
+ diff /mnt/lustre/hosts /etc/hosts
diff /mnt/lustre/test_mdt1_backup/test_mdt1/dir/fstab /etc/fstab || { echo "the file is
different" && exit 1; }
+ diff /mnt/lustre/test_mdt1_backup/test_mdt1/dir/fstab /etc/fstab
diff /mnt/lustre/test_mdt1_backup/test_mdt1/dir/hosts /etc/hosts || { echo "the file is
different" && exit 1; }
+ diff /mnt/lustre/test_mdt1_backup/test_mdt1/dir/hosts /etc/hosts
umount /mnt/lustre
+ umount /mnt/lustre
umount /mnt/mds
+ umount /mnt/mds
umount /mnt/ost1
+ umount /mnt/ost1
```