

Create and Mount a Lustre Filesystem

Purpose

Describe the steps you need to get a Lustre system (MGS, MDT, MDS, OSS, OST, client) up and running on a cluster of three machines.

Prerequisite

- Three machines, networked together: client-10, client-11, client-12
- Able to use fdisk to create primary partitions from unallocated disk space.

Overview

The chosen platform is:

- Centos 5
- Lustre 2.x, most recent build from build.whamcloud.com.

The Lustre topology will be:

- lustre file system name is `lustrewt` (lustre walk through)
- `client-12` (IB: 192.168.4.12) for the OSS and OST.
- `client-11` (IB: 192.168.4.11) for the MGT, MDS and MDT.
- `client-10` (IB: 192.168.4.10) for the client.

This document is based on: section 4.1 [Configuring the Lustre File System](#) from the Lustre File System, Operations Manual - Version 2.x

Procedure

Provision Machines

This is covered in the pages:

- [Walk-thru- Build Lustre 1.8 on CentOS 5.5 or 5.6 from Whamcloud git](#)
- [Walk-thru- Build Lustre MASTER on RHEL 6.4/CentOS 6.4 from Git](#)

Set up the MGS, MDT, and MDS

The node `client-11` is going to be the MGS, MDT, and MDS.

1. connect to client-11

```
rlogin -l root client-11
```

2. check to see which networks are suitable

```
cat /etc/modprobe.conf
```

This should show

```
options lnet networks="o2ib0(ib0)"
```

configured for lnet.

3. Use

```
fdisk /dev/sda
```

to create a primary partition from the unallocated disk space. This partition is

```
/dev/sda4
```

4. Create the MGS/MDT file system:

```
mkfs.lustre --fsname=lustrewt --mgs --mdt /dev/sda4
```

You should see:

```
Permanent disk data:
Target:      lustrewt-MDTffff
Index:       unassigned
Lustre FS:   lustrewt
Mount type:  ldiskfs
Flags:       0x75
(MDT MGS needs_index first_time update )
Persistent mount opts: iopen_nopriv,user_xattr,errors=remount-ro
Parameters:  mdt.group_upcall=/usr/sbin/l_getgroups

checking for existing Lustre data: not found
device size = 205220MB
2 6 18
formatting backing filesystem ldiskfs on /dev/sda4
target name  lustrewt-MDTffff
4k blocks    52536566
options      -J size=400 -i 4096 -I 512 -q -O dir_index,uninit_groups -F
mkfs_cmd = mke2fs -j -b 4096 -L lustrewt-MDTffff -J size=400 -i 4096 -I 512 -q -O dir_index,
uninit_groups -F /dev/sda4 52536566
Writing CONFIGS/mountdata
```

5. Create a mount point:

```
mkdir /mnt/mdt
```

6. Mount the MGS/MDT file system

```
mount -t lustre /dev/sda4 /mnt/mdt
```

7. Install lnet module: `modprobe lnet`
8. Start-up lnet: `lctl network up`
9. Check the NID: `lctl list_nids` this should return `192.168.4.11@o2ib`
10. Make lnet install during boot.

Set up the OST and OSS

The node `client-12` is going to be the OST and OSS.

1. connect to client-12

```
rlogin -l root client-12
```

2. check to see which networks are suitable

```
cat /etc/modprobe.conf
```

This should show `options lnet networks="o2ib0(ib0)"` configured for lnet.

3. Use `fdisk /dev/sda` to create a primary partition from the unallocated disk space. This partition is `/dev/sda4`.
4. Create the OST:

```
mkfs.lustre --ost --fsname=lustrewt --mgsnode=192.168.4.11@o2ib0 /dev/sda4
```

you should see:

```

Permanent disk data:
Target:      lustrewt-OSTffff
Index:       unassigned
Lustre FS:   lustrewt
Mount type:  ldiskfs
Flags:       0x72
(OST needs_index first_time update )
Persistent mount opts: errors=remount-ro,extents,malloc
Parameters: mgsnode=192.168.4.11@o2ib

checking for existing Lustre data: not found
device size = 205220MB
2 6 18
formatting backing filesystem ldiskfs on /dev/sda4
target name  lustrewt-OSTffff
4k blocks    52536566
options      -J size=400 -i 16384 -I 256 -q -O dir_index,extents,uninit_groups -F
mkfs_cmd = mke2fs -j -b 4096 -L lustrewt-OSTffff -J size=400 -i 16384 -I 256 -q -O dir_index,extents,
uninit_groups -F /dev/sda4 52536566
Writing CONFIGS/mountdata

```

5. Create a mount point: `mkdir /ostoss_mount`
6. Mount the OST: `mount -t lustre /dev/sda4 /ostoss_mount`

Setup the client.

1. connect to client-10

```
rlogin -l root client-10
```

2. Load the Lustre client module

```
modprobe lustre
```

3. NOTE: The client is available from build.whamcloud.com
4. Mount the lustre filesystem

```
mount -t lustre 192.168.4.11@o2ib:/lustrewt /mnt
```

Making Lustre servers persist through a reboot.

1. Add mount command for MGS, MDS, and MDT to `/etc/fstab` on client-11

```
/dev/sda4          /mgsmtdt_mount    lustre defaults,_netdev    0 0
```

2. Add mount command for OST and OSS to `/etc/fstab` on client-12

```
/dev/sda4          /ostoss_mount     lustre defaults,_netdev    0 0
```

ENDS-