

Walk-thru- Creating a Lustre client

Purpose

This walk-thru is intended to provide simple instructions to create a client of your Lustre filesystem.

Prerequisite

1. You have a Lustre filesystem available on your local network. This walk-thru assumes the server and client are both running Lustre MASTER code. In this walk-thru the MGS is assumed to be at: 192.168.1.1 on a tcp network.



Provisioning a Lustre filesystem

Lustre server components can either be [built from scratch](#) or [installed from Whamcloud RPMs](#).

2. You have a Linux distribution installed that uses one of the [Whamcloud supported client kernels](#).
3. In this walk-thru the Linux distribution is assumed to be RHEL 6, kernel=2.6.32_131.6.1.el6, on a host with IP address: 192.168.1.100.
4. iptables should be disabled on the MGS node.

Procedure

Fetching to correct RPMs

1. Visit the [Whamcloud build server](#) and choose the correct build from the matrix. In this case it is 'client', 'x86_64, el6, inkernel'. A direct link to this page is: http://build.whamcloud.com/job/lustre-master/arch=x86_64,build_type=client,distro=el6,ib_stack=inkernel/.



Network stacks and Lustre

[Lustre supports a wide variety of networks](#). Whamcloud builds two version of Lustre, one with Infiniband drivers that are shipped with the kernel (inkernel) and the other with Infiniband drivers that are provided by OpenFabrics Alliance (ofa)

2. Save `lustre-client-modules` and `lustre-client` from build.whamcloud.com to your host.

Installing the client RPMs

The RPMs for the client provide the kernel modules to allow the client to mount the Lustre filesystem.

1. Use `yum` to install the RPMs.

```
yum --nogpgcheck localinstall lustre-client-2.1.52-2.6.32_131.6.1.el6.x86_64_ga296e94.x86_64.rpm lustre-client-modules-2.1.52-2.6.32_131.6.1.el6.x86_64_ga296e94.x86_64.rpm
```

2. Install the module into the kernel.

```
modprobe lustre
```

Mounting the Lustre filesystem.

1. Use `mount` to mount the filesystem:

```
mount -t lustre 192.168.1.1@tcp:/lustre /mnt
```

2. A basic test that the filesystem is working:

```
# mount
/dev/mapper/VolGroup-lv_root on / type ext4 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
tmpfs on /dev/shm type tmpfs (rw)
/dev/vda1 on /boot type ext4 (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
192.168.122.51@tcp:/lustre on /mnt type lustre (rw)
# dd if=/dev/zero of=/mnt/test.dat bs=1K count=100K
102400+0 records in
102400+0 records out
104857600 bytes (105 MB) copied, 2.6479 s, 39.6 MB/s
```

Troubleshooting

Q Mount returns:

```
mount.lustre: mount 192.168.122.51@tcp:/lustre at /mnt failed: Input/output error
Is the MGS running?
```

A

- Check the MGS is available on the network.
- Check that iptables are disabled on the MGS machine.