

**S-2508-B: C3 Lustre Client Software Installation
Guide**

Contents

Introduction.....	3
C3 Lustre® Client Software Packages.....	3
Pre-Installation Tasks.....	4
Gather Required System Information.....	4
Download C3 Lustre® Client Software.....	4
Install C3 Lustre® Client Software on Linux Systems.....	6
Install C3 Lustre® Client Software on a CCS System.....	7
Troubleshoot Installation Problems.....	9

Introduction

This guide describes procedures to install the Cray® Cluster Connect™ (C3) Lustre® Client software on supported Linux® systems. It assumes that Linux and Lustre server software and Lustre networks are configured and operating. The intended audience is experienced Linux system administrators.

C3 Lustre® Client Software Packages

C3 Lustre® Client software is distributed as compressed tape archive files (tarballs). Each tarball package includes the complete product of a Lustre client build. Two binary RPMs, the client RPM and the modules RPM, required to install and run the C3 Lustre Client software are included. Several other binary and source RPMs as well as other build artifacts, which some users might find useful, are also included.

Tar Files

C3 Lustre client software tarballs (.tgz) adhere to the following naming convention:

```
lustre-client-lustre_version-CLE_version-IB_version-distro_kernel.x86_64.tar
```

lustre_version Lustre version number (e.g., 1.8.6 or 2.4.0)

CLE_version Cray Linux Environment (CLE) version (e.g., 4.2UP02_4.2.34 or 5.1UP00_5.1.18)

IB_version InfiniBand version (e.g., mlnx_1.5.2_2.1.0 or mlnx_2.0_3.0.0)

distro OS distribution (e.g., cent_5.6 or rhel_6.4 or sles_11.2)

kernel Kernel version (e.g., 2.6.32_358)

Examples of C3 Lustre client tarball file name conventions:

Lustre Version: 2.4.0 From CLE 5.1UP00

Operating System: RHEL 6.4, Kernel 2.6.32_358.el6.x86_64

InfiniBand version: Mellanox MLNX_OFED_LINUX-2.0-3.0.0

Filename: lustre-

client-2.4.0-5.1UP00_5.1.18_mlnx_2.0_3.0.0_rhel_6.4_2.6.32_358.el6.x86_64.x86_64.tgz

Lustre Version: 2.4.0 From CLE 5.1UP00

Operating System: CentOS 6.4, Kernel 2.6.32_358.el6.x86_64

InfiniBand version: Mellanox MLNX_OFED_LINUX-2.0-3.0.0

Filename: lustre-

client-2.4.0-5.1UP00_5.1.18_mlnx_2.0_3.0.0_cent_6.4_2.6.32_358.el6.x86_64.x86_64.tgz

RPM Files

The file names for both the client and modules RPMs also include the Lustre version, CLE version, IB version, and OS kernel version strings as follows:

Client RMP:

```
lustre-client-lustre_version-CLE_version-IB_version-distro_kernel.x86_64.rpm
```

Modules RMP:

```
lustre-client-modules-lustre_version-  
CLE_version-IB_version-distro_kernel.x86_64.rpm
```

Pre-Installation Tasks

C3 Lustre® client packages are built to support specific combinations of the operating system (OS), kernel, and Open Fabrics Enterprise Distribution (OFED) versions. Gather the version information for each of these components and then download the correct Lustre client software prior to beginning the installation procedure.

Gather Required System Information

The information gathered in these steps is used to select the correct software package to download. Examples displayed in these steps are for reference only. Site-specific output will vary.

1. Display the OS version.

- a) For CentOS and RHEL:

```
# cat /etc/redhat-release
```

```
CentOS release 5.8 (Final)
```

- b) For SLES:

```
# cat /etc/SuSE-release
```

```
SUSE Linux Enterprise
Server 11 (x86_64)
VERSION = 11
PATCHLEVEL = 2
```

2. Display the kernel version.

```
# uname -a
```

```
3.0.80-0.5.-default
```

3. Display the InfiniBand version string.

```
# ofed_info | head --lines=1
```

```
MLNX_OFED_LINUX-1.5.3-3.1.0 (OFED-1.5.3-3.1.0) :
```

Download C3 Lustre® Client Software

C3 Lustre® Client software packages are available for download from CrayPort™, the Cray customer portal. A login account is required to access CrayPort. Customers with active support contracts can register directly on the CrayPort website to obtain a login account.

1. Navigate to <http://crayport.cray.com>.
2. Log in using a valid customer account or register to obtain a new customer account.
3. Navigate to the software download page: **Software > Order Cray Software**.
4. Locate the customer's site and, if applicable, the Cray system serial number under **Site Serials**.
5. Use the system versioning information (system OS, kernel, and OFED version combination) to download the appropriate C3 Lustre Client tarball from the **Cray Software Distribution** database to an installation directory on the local system.
6. Extract the C3 Lustre Client software from the tarball.

```
# tar -xvzf client_tarball.tgz
```

The RPMs used in this procedure are placed in the `./rpmbuild/RPMS/x86-64` directory.

Install C3 Lustre® Client Software on Linux Systems

Prerequisites:

- Complete the steps described in [Pre-Installation Tasks](#) on page 4.
- This procedure requires `root` authority.

Use this procedure to install CS Lustre® Client software on Linux systems, excluding the CCS environment. If installing C3 Lustre® Client software in a CCS environment, see [Install C3 Lustre Client Software on a CCS System](#) on page 7.

1. Change directory to `./rpmbuild/RPMS/x64_64`.

```
cd ./rpmbuild/RPMS/x64_64
```


2. Install the C3 Lustre Client RPMs.

```
# rpm -Uvh ./rpmbuild/RPMS/x86_64/client_rpm \
./rpmbuild/RPMS/x86_64/modules_rpm
```

Where `client_rpm` and `modules_rpm` are the client and modules RPM files, respectively, extracted from the downloaded tarball file in [Download C3 Lustre Client Software](#) on page 4.

For example:

```
# rpm -Uvh ./rpmbuild/RPMS/x64_64/lustre-client-2.4.0-5.1UP00_5.1.18_\
mlnx_2.0_3.0.0_rhel_6.4_2.6.32_358.el6.x86_64.x86_64.rpm \
./rpmbuild/RPMS/x86_64/lustre-client-modules-2.4.0-5.1UP00_5.1.18_\
mlnx_2.0_3.0.0_rhel_6.4_2.6.32_358.el6.x86_64.x86_64.rpm
```

 **Note:** For readability, long RPM file names are broken across multiple lines in this example.

3. (SUSE Linux only) Verify that `allow_unsupported_modules` is set to 1 in the `/etc/modprobe.d/unsupported-modules` file prior to loading the software. This is necessary because SLES restricts the installation of some third-party modules.

- a) Edit `/etc/modprobe.d/unsupported-modules`.

```
# vim /etc/modprobe.d/unsupported-modules
```

- b) Verify that the following line exists, or add it if necessary.

```
allow_unsupported_modules 1
```

4. Load the C3 Lustre Client software onto the system.

```
# modprobe -v lustre
```

Install C3 Lustre® Client Software on a CCS System

Prerequisites: Complete the steps described in *Pre-Installation Tasks* on page 4.

Use this procedure to install C3 Lustre Client software on Cray® Cluster Solutions (CCS) systems. If installing C3 Lustre Client software in a non-CCS environment, see *Install C3 Lustre Client Software on Linux Systems* on page 6.

1. Log on to a management server.
2. List the servers available in the cluster to select an available compute server on which to build a prototype software image.

```
[root@mgmt1~]$ ace servers
```

3. Stop the selected server to shut it down and remove it from the cluster. In this example, `prod-0010` is the hostname of the selected server, and the command line option `1` specifies 1 client.

```
[root@mgmt1~]$ ace stop prod 1 prod-0010
```

4. Clone the current (active) software image to create a new software image revision. In this example, `server-0014` is the cluster name of the server with the hostname `prod-0010`, and the command line option `4` specifies the revision number.

```
[root@mgmt1~]$ ace checkout prod 4 server-014
```

5. Copy the C3 Lustre client RPMs from their downloaded location to the server where the changes will be made.

```
[root@mgmt1~]$ scp lustre*.rpm root@prod-r04:/root
```

6. Log on to the server.

```
[root@mgmt1~]$ ssh root@prod-r04
```

7. Install the C3 Lustre Client RPMs.

```
[root@prod-r04~]$ rpm -ivh lustre*.rpm
```

8. Create the `/etc/modprobe.d/lnet.conf` file. This is a site-specific configuration file for the Lustre client driver. If necessary, check with appropriate installation personnel to get the site-specific `lnet.conf` file. The following is an example of a `lnet.conf` file and is not meant to be appropriate for all sites.

```
options lnet networks=o2ib1(ib0)
options lnet avoid_asym_router_failure=1
options lnet dead_router_check_interval=60
options lnet live_router_check_interval=60
options lnet check_routers_before_use=1
options ko2iblnd timeout=100 credits=2048 ntx=2048
options ko2iblnd peer_credits=63 peer_timeout=0 keepalive=30
```

9. Create the Lustre client mount point. The mount point name differs for each site-specific configuration.

```
[root@prod-r04~]$ mkdir -p local_mnt_path
```

10. Load the C3 Lustre client driver onto the system.

```
[root@prod-r04~]$ modprobe -v lustre
```

11. Mount the file system.

```
[root@prod-r04~]$ mount -t lustre \
mds_ip_address@o2ib_network_name:remote_mnt_path local_mnt_path
```

12. Verify that the file system is mounted and readable/writeable, then unmount it and remove the `lustre` module.

```
[root@prod-r04~]$ umount local_mnt_pt
[root@prod-r04~]$ modprobe -r lustre
```

13. Exit to the mgmt server.

```
[root@prod-r04~]$ exit
```

14. Check in the new version if satisfied with the changes to the prototype image. This example yields version 5.

```
[root@mgmt1~]$ ace checkin prod 4 "Installed lustre-client-new-rev"
```

15. Activate the new version.

```
[root@mgmt1~]$ ace activate prod 5
```

16. Edit the `/etc/fstab` file to mount the Lustre file system at boot. The path name to `/etc/fstab` is site-specific; verify the correct location.

```
[root@mgmt1~]$ vi /etc/fstab
```

17. Add the mount point (as one line) to the `/etc/fstab` file. Ensure that `noauto` is specified as a parameter.

```
mds_ip_address@o2ib_network_name:remote_mnt_path local_mnt_path lustre \
rw,flock,noauto 0 0
```

18. Add the following command to the `0200-modprobe` file; this ensures that the Lustre driver is loaded before the Lustre file system is mounted during boot. The path to `0200-modprobe` is site-specific.

```
modprobe lustre
```

19. Add the following command to the `1900-mount-lustre` file; this mounts the Lustre filesystem during boot. The path to `1900-mount-lustre` is site-specific.

```
test -d local_mnt_path && /bin/grep -q local_mnt_path && \
mount local_mnt_path
```

20. Reboot all nodes with active image revision.

Troubleshoot Installation Problems

Common troubleshooting scenarios and solutions are described below. If a failure occurs that is not resolved by these suggestions, contact a Cray service representative.

Client Mismatch Failures

If the C3 Lustre® Client software is not compatible with the software environment, failures will occur. See [Pre-Installation Tasks](#) on page 4 for further details.

SLES 11 Restrictions

SUSE Linux Enterprise Server (SLES) restricts the installation of some third-party modules. If the client software load fails on a system running SLES, it may be due to this restriction. See the SUSE Linux only step in [Install C3 Lustre Client Software on Linux Systems](#) on page 6 for details.

Check the Network

The Lustre network must be configured and operational before loading the client software with the `modprobe` command. If not, messages similar to the following are displayed.

```
[root@demo~]# modprobe -v lustre
insmod /lib/modules/2.6.18308.el5/updates/kernel/fs/lustre/ptlrpc.ko
WARNING: Error inserting ptlrpc
(/lib/modules/2.6.18-308.el5/updates/kernel/fs/lustre/ptlrpc.ko) :
Input/output error WARNING: Error inserting osc
(/lib/modules/2.6.18-308.el5/updates/kernel/fs/lustre/osc.ko) :
Unknown symbol in module, or unknown parameter (see dmesg)
WARNING: Error inserting mdc
(/lib/modules/2.6.18-308.el5/updates/kernel/fs/lustre/mdc.ko) :
Unknown symbol in module, or unknown parameter (see dmesg)
WARNING: Error inserting lov
(/lib/modules/2.6.18-308.el5/updates/kernel/fs/lustre/lov.ko) :
Unknown symbol in module, or unknown parameter (see dmesg)
FATAL: Error inserting lustre
(/lib/modules/2.6.18-308.el5/updates/kernel/fs/lustre/lustre.ko) :
Unknown symbol in module, or unknown parameter (see dmesg)
```