

Resources Available to You

Production Systems on the Open Computing Facility (OCF)

Multiprogrammatic & Institutional Computing Resources (M&I)

Ansel: Linux cluster of Intel Xeon X5660 CPUs for parallel processing. Composed of 324 nodes, each with 12 CPUs (2.8 GHz) and 24 GB memory.

Aztec: Linux cluster of Intel Xeon X5660 CPUs for serial or on-node processing. Composed of 96 nodes, each with 12 CPUs (2.8 GHz) and 48 GB memory.

Herd: Linux cluster of AMD Opteron 6128 CPUs for parallel processing. Composed of 7 nodes, each with 16 or 32 CPUs (2.0 GHz) and 32 or 512 GB memory.

Sierra: Linux cluster of Intel Xeon X5660 CPUs for parallel processing. Composed of 1,944 nodes, each with 12 CPUs (2.8 GHz) and 24 GB memory.

RZCereal: Linux cluster of Intel Xeon E5530 CPUs for parallel processing. Composed of 34 nodes, each with 8 CPUs (2.4 GHz) and 24 GB memory.

RZZeus: Linux cluster of Intel Xeon E5530 CPUs for parallel processing. Composed of 288 nodes, each with 8 CPUs (2.4 GHz) and 24 GB memory.

ASC/M&I

Cab: Linux cluster of Intel Xeon E5-2670 CPUs for parallel processing. Composed of 1,296 nodes, each with 16 CPUs and 32 GB memory.

Edge: Linux cluster of Intel Xeon X5660 CPUs and NVIDIA Tesla M2050 GPUs for visualization and data analysis. Composed of 216 nodes, each with 12 CPUs (2.8 GHz) and 96 GB memory and 2 GPUs with 3 GB memory.

Hera: Linux cluster of AMD Opteron 8356 CPUs for parallel processing. Composed of 864 nodes, each with 16 CPUs (2.3 GHz) and 32 GB memory.

OSLIC: Lustre interface cluster for moving files to and from HPSS archival storage.

RZMerl: Linux cluster of Intel Xeon E5-2670 CPUs for parallel processing. Composed of 162 nodes, each with 16 CPUs (2.6 GHz) and 32 GB memory.

RZSLIC: Lustre interface cluster for moving files to and from HPSS archival storage.

Production Systems on the Secure Computing Facility (SCF)

Advanced Simulation and Computing (ASC) Resources

CSLIC: Lustre interface cluster for moving files to and from HPSS archival storage.

Graph: Linux cluster of AMD Opteron 8423 CPUs for visualization and data analysis. Composed of 576 nodes, each with 24 CPUs (2.0 GHz) and 64 or 128 GB memory.

Inca: Linux cluster of Intel Xeon X5660 CPUs for serial or on-node processing. Composed of 96 nodes, each with 12 or 16 CPUs (2.8 GHz) and 48 or 128 GB memory.

Juno: Linux cluster of AMD Opteron8354 CPUs for parallel processing. Composed of 1,152 nodes, each with 16 CPUs (2.2 GHz) and 32 GB memory.

Muir: Linux cluster of Intel Xeon X5660 CPUs for parallel processing. Composed of 1,296 nodes, each with 12 CPUs (2.8 GHz) and 24 GB memory.

Zin: Linux cluster of Intel Xeon E5-2670 CPUs for parallel processing. Composed of 2,916 nodes, each with 16 CPUs (2.6 GHz) and 32 GB memory.

Accessing LC Resources

An RSA OTP (one-time password) is used for authentication to all OCF LC resources. A CRYPTOCARD is used for authentication to the rzgw (gateway to the RZ), then the RSA OTP is then used to authenticate to the RZ resource. Either an RSA OTP or a Kerberos password is used for authentication to all SCF LC services.

Remote Access Accounts

From outside llnl.gov, you will also need a remote access account to access OCF systems through the llnl.gov firewall.

- **VPN (Virtual Private Network)/VPN-C (VPN for collaborators):** Install the VPN or VPN-C client on the machine from which you are coming. The service gives you a virtual internal llnl.gov IP address. VPN is needed to access RZ resources.
- **SNL/LANL Access:** Sandia and Los Alamos users may access OCF systems without VPN using port 922 from an SNL or LANL IP address.

More information about remote access accounts is available at <https://access.llnl.gov/> or contact the OpenLabNet Help Desk at 422-4090.

Information on LC access issues is available at <https://computing.llnl.gov/access/>

Secure Shell (SSH)

Secure Shell is the only login method for LC systems. An SSH client is necessary on any host used to access LC systems.

From an llnl.gov host to access LC host lc_host.llnl.gov:

```
ssh <lc_host>
```

For more information about using SSH to access LC systems, go to https://computing.llnl.gov/access/content/ssh_outside.php

Files and File Systems

Home Directory

Your home directory is shared globally across all LC systems—two on OCF (CZ/RZ) and one on SCF. The home directory is backed up and is not subject to purge, but it is subject to quota limits. There is an online backup of the past 48 hours in your home directory. From your home directory, you can `cd .snapshot` to see those files.

The types of user login files (“dot” files) in your home directory are:

- Master dot files that are executed upon login to all LC systems.
 - tsh, csh users: `.cshrc`, `.login`
 - ksh users: `.kshrc` or `.kshenv` and `.profile`
 - bash users: None. Users must create their own.
- System-specific dot files that are executed only upon login to specific hosts or types of hosts. Example: `.login.linux` is executed only upon login to the Linux machines.

For details about these and other LC systems, including hardware and operating system version, see <https://computing.llnl.gov/resources/>

For information about graphics support, visualization, and simulation, see <https://computing.llnl.gov/vis/>

Temporary File Systems

There are two temporary file system choices on LC systems. These file systems are targeted for serial I/O.

/nfs/tmp Large file system, globally available to all LC OCF and SCF systems.

`/nfs/tmp2`: 2 TB crash barrier per user

/var/tmp (/usr/tmp) Each node has a small `/var/tmp` (`/usr/tmp`) accessible only from that node.

All temporary file systems share these characteristics:

- They are not backed up.
- Files are subject to purge as needed. The purge may remove files that meet a criterion, often more than 10 days since last access (or 5 days when necessary) to make room on the file system.
- For long-term storage of files, use the archival storage facility.

Temporary file system quota and purge information is described in the news item accessible by typing `news purge.policy`.

Parallel File Systems

For resources with parallel file systems, these temporary file systems are typically found in a directory named `/p/l*`, where `l*` indicates Linux Lustre file systems. These file systems are targeted for parallel I/O.

There are no quotas on the parallel file systems, but they are subject to purging as described in the news item accessible by typing `news purge.policy`.

Type `quota -v` to see current file system usage. For more information on LC file systems, see the EZFILES manual at <https://computing.llnl.gov/LCdocs/ezfiles>.

Storage (High Performance Storage System)

HPSS is an archival storage facility on both the OCF and SCF. This facility provides long-term storage for each user. All LC users are given a storage account. Type `ftp storage.llnl.gov` to access your storage account. Use FTP commands to put files into storage and to retrieve them.

There is a yearly growth quota on HPSS storage. Type `quota` to reveal your current yearly growth and the pools from which you have quota allocations.

For information about other higher-performance interfaces to storage see the EZSTORAGE manual at <https://computing.llnl.gov/LCdocs/ezstorage>.

Development Environment

LC provides support for compilers, debuggers, performance analysis tools, and parallel tools and libraries on all LC platforms. Languages include C, C++, and Fortran; parallel APIs include MPI, OpenMP, and Pthreads. For additional details, see:

- **Supported Software and Computing Tools**
https://computing.llnl.gov/code/content/software_tools.php
- **Available Compilers**
<https://computing.llnl.gov/code/content/compilers.php>
- **MPI/OpenMP Usage**
<https://computing.llnl.gov/mpi/libraries.html>

LC Batch System

User work on LC systems is done within LC's batch system, Moab. You can use `msub` to submit jobs and `showq` or `checkjob` to monitor jobs. For more information:

- See man pages for `msub`, `showq`, and `checkjob`
- Computing resource management documentation is available at <https://computing.llnl.gov/jobs/content/crm.php>.
- Type `news job.lim.<cluster-name>` on any LC system for information about batch and interactive job limits (e.g., `news job.lim.hera`).

File Interchange Service

The File Interchange Service (FIS) can be used to move files from the OCF CZ/RZ to the SCF. You must request a FIS account from the LC Hotline to use this service. For information about FIS, refer to the File Interchange Service manual at <https://computing.llnl.gov/LCdocs/fis/>.

Additional Information

- **Message of the day (MOTD)**
The announcements print out when you log in to an LC machine.
- **News postings**
The latest unread news items are listed following the MOTD and can be read by typing `news <item_name>`. To list all news items, type `news -n`.
- **System status**
CZ/RZ machine and file system status are available from the System Status navigation link at <https://computing.llnl.gov/>.
- **Technical Bulletins**
<https://lc.llnl.gov/computing/techbulletins/>
- **Forms**
<https://computing.llnl.gov/forms/>
- **Training** (tutorials)
<https://computing.llnl.gov/training/>
- **Documentation**
<https://computing.llnl.gov/documentation/>
- **E-mail lists**
You will be alerted via e-mail to changes, scheduled events, and problems on those system(s) on which you have an account.

LC Quick Guide

User Information

<https://computing.llnl.gov/>

LC Hotline Technical Consultants

OCF e-mail: lc-hotline@llnl.gov
SCF e-mail: lc-hotline@pop.llnl.gov

Phone: (925) 422-4531, option 1

LC Hotline Account Specialists

OCF e-mail: lc-support@llnl.gov
SCF e-mail: lc-support@pop.llnl.gov

Phone: (925) 422-4531, option 2

Walk-in: B453, Room 1103
Monday–Friday, 8:00 a.m.–noon, 1:00–4:45 p.m.

LC Operations

(Available 24 hours a day/7 days a week)

OCF e-mail: lc-operator@llnl.gov
SCF e-mail: ops@pop.llnl.gov

Phone: (925) 422-4531, option 3



at
Lawrence Livermore National Laboratory

UCRL-TB-148428 This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344.

Revised November 6, 2012