

The LWA User Computing Facility
LWA Memo #193
Ver. 6

J. Dowell*

January 20, 2021

Contents

1	Introduction	2
2	Available Hardware and Software	2
3	Networking	3
4	Storage	3
5	User Management and Access	3
6	Usage and Storage Policy	4
7	Document History	5

*University of New Mexico. E-mail: jdowell@unm.edu

1 Introduction

The LWA User computing facility (UCF) is a cluster of three computing nodes and one data storage node available to LWA users for reducing LWA data. The cluster is located in two racks in the old correlator room in the VLA control building. This location was selected in order to decrease the turn around time for data processing by moving the processing equipment closer to the data. This document describes the hardware and software setup of the cluster, and provides information about user access.

2 Available Hardware and Software

Each of the computing nodes consists of:

- two eight core Intel(R) Xeon(R) Silver 4215R CPUs with clock rates of 3.20GHz;
- a SuperMicro X11DPG-QT system board;
- 192 GB (nodes 1 and 2) or 384 GB (node 3) of RAM;
- two Nvidia GeForce GTX 980 cards with
 - 4 GB of RAM,
 - 2048 CUDA cores, and
 - CUDA compute capability 5.2;
- a two port 10 GbE fiber card; and
- 18 TB of RAID10 storage.

In addition to the three computing nodes there is also a data storage node for raw LWA data. This machine consists of:

- an Intel(R) Xeon(R) E5620 eight core CPU with a clock rate of 2.40GHz;
- a Supermicro X8DT6 system board;
- 16 GB of RAM;
- a single port 10 GbE fiber card; and
- 138 TB of RAID6 storage.

Each computing node runs the desktop version of Ubuntu 18.04 LTS¹ and has a variety of software installed by default, including bbcp, CASA, FFTW, PRESTO, CUDA, LSL², OpenMPI, TEMPO, TEMPO2, PSRCHIVE, and DSPSR. Users requiring additional software should install the needed software locally in their home directories or contact the cluster administrator, Jayce Dowell, via e-mail. The storage node runs the server version of Ubuntu 18.04 LTS and is not intended for access by users of the cluster.

¹<http://releases.ubuntu.com/bionic/>

²The LSL Commissioning, Pulsar, and SessionSchedules extensions are also available on each node under `/usr/local/extensions/`.

3 Networking

The UCF nodes are physically connected by combination of SFP+ direct attach cables and 10Gbase-SR 10 GbE fiber modules with 50/125 multimode fiber patch cables via a switch. The cluster is connected to the LWA1 Monitor and Control (MCS) 10.1.1.* network via a 10 GbE connection to the shelter. See [1] for details of the network configuration. The nodes can be accessed by their short hostnames in the form of “lwaucf#” where “#” represents the number of the node, one through three, the user wishes to connect to. In addition to the main cluster and MCS networks there is also an internal 1 GbE management network that is used to provide an alternate means of communication between nodes in the case of a problem with the fiber connections.

4 Storage

Data storage for the UCF is divided into three tiers: user-specific storage, temporary local storage, and temporary massive data storage.

1. The first tier is implemented through an NFS mounted home directory available on all of the nodes. The total storage under `/home/` is 3.6 TB which is shared among all users. This mount point is intended to store the user’s personal files as well as any final data products waiting to be transferred out of the cluster. Backups of the home directory are irregular and users are encouraged to actively backup files they are interested in preserving.
2. The second tier is implemented as a temporary “scratch space” in `/data/local` where users of the node can store intermediate data products that require further manipulation. This directory is available on each node and is world writable for all users of that particular node.
3. The final tier of storage is a Gluster file system³ with a current total capacity of 138 TB. This file system, mounted on `/data/network/`, is located on the storage node of the cluster. Since this file system is intended to only store raw data it is world readable but write permissions are restricted to the LWA operators.

Although the file system is read only for users of the UCF there are a few tools provided for basic file management: `mkdirrd`, `mvrdd`, and `rmrdd`. `mkdirrd` is similar to the standard `mkdir` utility and is used to request directories to be created in the `/data/network/` area. `mvrdd` is similar to `mv` and allows files to be renamed or moved inside of `/data/network/` while `rmrdd` is used to request deletion of files in this directory. These three commands work using a queue system. The `mkdirrd` and `mvrdd` command queues require between one to two hours to process, while the `rmrdd` command queue requires between 24 and 48 hours to process.

Data and MCS0030 metadata [2] files are copied over to the Gluster file system through mount points on each of the data recorders at the LWA stations. The copying process is initiated and monitored by the LWA operators on duty or by using the “auto-copy” feature. This feature allows observation files and metadata to be automatically copied to the cluster within a few hours of the observation ending. This is requested through session definition files (SDFs) generated with the `sessionGUI.py` script that is available as part of the `SessionSchedules` LSL extension.

5 User Management and Access

Users of the UCF are given individual unprivileged accounts which work on all nodes via a light-weight directory access protocol (LDAP) server . All remote access to the cluster will be han-

³<http://www.gluster.org/>

dled through secure shell (SSH) and X2Go⁴ sessions. Since the cluster is located within the MCS network, users of the cluster will first need to connect to one of the UNM machines, e.g., `lwalab.phys.unm.edu`, `hercules.phys.unm.edu`, or `virgo.phys.unm.edu` and use the OpenVPN⁵ tunnel to access the cluster nodes via SSH. The cluster LDAP information will also be available on `lwalab`, `hercules`, and `virgo` so that users will only need one set of credentials to access the cluster.

Time on one or more of the computing nodes is allocated based on requests in the LWA observing proposals. See [3] for details on how time is requested. Users needing to copy data off the cluster should contact the cluster administrator to schedule a window for the copy. This is required to ensure that data transfers do not interfere with the operation of LWA1 or other user requests. Files can be copied off using a number of programs such as `scp`, `rsync`, and `bbcp`.

User accounts will remain active as long as they are being actively used. Accounts that have been inactive for more than three years will be locked. Any users wishing to access their locked account should contact the cluster administrator via e-mail.

6 Usage and Storage Policy

Since the cluster is a shared resource, it is possible for multiple users to access the same node concurrently. Users are responsible for monitoring their processes to ensure that the nodes they are using do not become overburdened or hang. The state of the machines can be checked using standard UNIX tools, such as `ps`, `top`, and `free`, or via the LWAUCF Operations Screen page at <http://lwalab.phys.unm.edu/CompScreen/cs.php>.

Users are also responsible for monitoring their disk space usage on the `/home`, `/data/local`, and `/data/network` drives. Currently there is a combination of quotas and file age limits that are enforced for each storage tier:

- `/home` — quotas but no file age limits,
- `/data/local` — quotas and file age limits, and
- `/data/network` — quotas and file age limits.

Files stored on `/home` are subject to per-user quotas. By default each user receives a quota of 200 GB.

Files stored on `/data/local` are subject to per-user quotas and file age limits. By default each user receives a quota of 3 TB in `/data/local`. These directories are monitored on a monthly basis to identify files between 547 and 730 days old. Users with files meeting this criterion will be notified by e-mail in order to determine whether or not the indicated files can be deleted. Files that are older than 730 days (two years) are subject to automatic deletion within two weeks after the user is notified by e-mail. Any files that the user wishes to save should be transferred either to the `/home` directory or off the cluster.

Files stored on `/data/network` are subject to per-user quotas and file age limits. By default each user receives a quota of 10 TB in `/data/network/` although additional storage may be requested by e-mailing the LWA director. Since there is additional pressure on the space in this directory, the automatic file age monitoring system notifies users via e-mail about files between 180 and 270 days old. Files older than 270 days (nine months) are subject to automatic deletion two weeks after the user is notified by e-mail. Any user with files in this category can request that the files be saved by notifying the observatory staff. For further details about the management of data files see [3].

⁴<https://wiki.x2go.org/doku.php>

⁵<http://openvpn.net/index.php/open-source/overview.html>

7 Document History

- Version 6 (Jan 20, 2021): Updated for the new cluster configuration and quota information.
- Version 5 (Sep 13, 2019): Fixed a typo that gave node 4 GTX 980 GPUs instead of GTX 690 GPUs.
- Version 4 (Aug 20, 2019): Updated for recent changes to the GPU available on the cluster. Updated the list of pulsar software titles available on the cluster.
- Version 3 (Dec 21, 2018): Updated for changes to the hardware, software, and network setup. Added information about the data storage node. Updated the data retention policy for all three classes of storage. Added information about when accounts will be locked due to inactivity. Removed the listing of software packages.
- Version 2 (May 22, 2014): Updated for the new home directory size and monitoring of the data directory. Added information about the “toasters” on each node. Added a list of `.deb` packages that are installed. Added a usage and storage policy.
- Version 1 (Nov 25, 2012): First version.

References

- [1] J. Craig et al., “LWA1 Local Area Network Configuration”, Ver. 7.0, MCS Memo No. 33, Oct 25, 2012.
- [2] S. Ellingson, “LWA Station-Level Observing Procedure and Associated Metadata”, Ver. 5.0, MCS Memo No. 30, Apr 13, 2011.
- [3] G. Taylor et al., “LWA Data Management”, Ver. 4.0, LWA Memo No. 177, Dec 21, 2018.