The LWA User Computing Facility
LWA Memo #193
Ver. 3

J. Dowell*

December 21, 2018

Contents

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

1 Introduction

The LWA User computing facility (UCF) is a cluster of six 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:

- an Intel(R) Core(TM) i7-3930K six core CPU with a clock rate of 3.20GHz;
- an ASUS P9X79 Deluxe (nodes 1-3, 5, and 6) or P9X79 system board;
- 32 GB (nodes 1 through 4) or 64 GB (nodes 5 and 6) of RAM;
- an Nvidia GeForce GTX 680 card (nodes 1-3) or two GTX 690 cards (nodes 4-6) with
    - 2 GB of RAM,
    - 1536 CUDA cores, and
    - CUDA compute capability 3.0;
- an Nvidia Tesla C2050 card (nodes 2 and 3) or C2075 card (nodes 1, 4-6) with
    - 2 GB (C2050) or 6 GB (C2075) of RAM,
    - 448 CUDA cores, and
    - CUDA compute capability 2.0;
- a four port 10 GbE fiber card; and
- two 3 TB hard drives (nodes 1 through 4) or a 3 TB and a 6 TB hard drive (nodes 5 and 6).

In addition to the six 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 14.04 LTS[1] and has a variety of software installed by default, including ATLAS, bbcp, CASA, FFTW, PRESTO, CUDA, LSL[2], OpenMPI, TEMPO, and TEMPO2. Users requiring additional software should contact the cluster administrator, Jayce Dowell, via e-mail, or install the needed software locally in their home directories. The storage node runs the server version of Ubuntu 14.04 LTS and is not intended for access by users of the cluster.

---

[1]`http://releases.ubuntu.com/trusty/`
[2]The LSL `Commissioning`, `Pulsar`, and `SessionSchedules` extensions are also avalible on each node under `/usr/local/extensions/`.

3 Networking

The UCF nodes are physically connected by 10Gbase-SR 10 GbE fiber modules and 50/125 multi-mode 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 six, 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 and bare 3 TB drives known as "toasters". 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" where users of the node can store intermediate data products that require further manipulation. There are two areas available for this purpose, `/data/local` and `/mnt/toaster`. Both of these areas are 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 system[3] 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`, `mvrd`, and `rmrd`. `mkdirrd` is similar to the standard `mkdir` utility and is used to request directories to be created in the `/data/network/` area. `mvrd` is similar to `mv` and allows files to be renamed or moved inside of `/data/network/` while `rmrd` is used to request deletion of files in this directory. These three commands work using a queue system. The `mkdirrd` and `mvrd` command queues require between one to two hours to process, while the `rmrd` 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.

---

[3]`http://www.gluster.org/`

5   User Management and Access

Users of the UCF are given individual unprivileged accounts which work on all nodes via a lightweight directory access protocol (LDAP) server running on node 1. All remote access to the cluster will be handled through secure shell (SSH) 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 OpenVPN[4] 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`, `/mnt/network`, and `/data/network` drives. Currently there is a combination of quotas and file age limits that are enforced for each storage tier:

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

Files stored on `/data/local` and `/mnt/toaster` 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 of the cluster.

Files stored on `/data/network` are subject to per-user quotas in addition to 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

---

[4]`http://openvpn.net/index.php/open-source/overview.html`

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].

7   Document History

- 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.