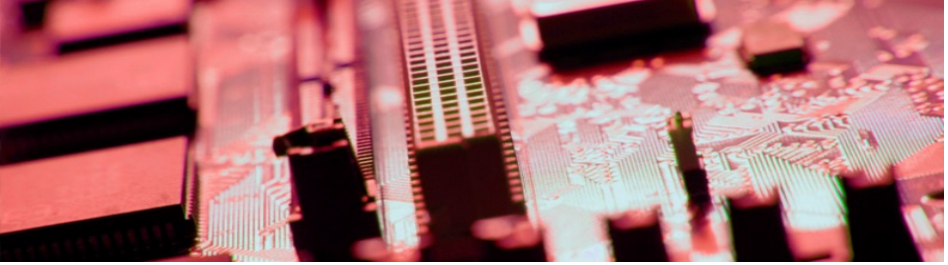

LCI workshop
on high performance clustered computing

Linux Clusters Institute: Lustre Hands On Exercise

Georgia Tech, August 15th – 18th 2017

J.D. Maloney | Storage Engineer
National Center for Supercomputing Applications (NCSA)
malone12@illinois.edu

NCSA

Goal of Hands on Exercise

- Create Lustre File System
- Bring in Sample Data
- Run Robinhood Scan
- Enabling Changelog consumption
- Run Purge on sample data driven by Robinhood

Lay of the Land

- You should have 3 storage servers; 1 for metadata, 2 for data; and 1 Robinhood Server
- Metadata Server has small disks; Object Servers have large disk
- Lustre packages should be installed and ready
- All severs have root ssh keys set between them

Creating Lustre File System

- Format & Mount the MGS

- Format & Mount the MDT

Creating Lustre File System

- Format and Mount the OSTs

- Setup the Clients

Bring In Test Data

- Bring in folder 1 of test data

Standing Up Robinhood

- Robinhood RPMs are already installed

- Edit Robinhood configuration file with correct parameters

Standing Up Robinhood

- Configure Robinhood Database

- Register with the MDS

Robinhood Scan & Change Logs

- Begin a Full Scan of the File System & Changelog Consumption

- Enable Change Logs

Robinhood Scan & Change Logs

- Bring in folder 2 of test data once the change logs are consuming (should be done while full scan is still running)

Running a Purge

- Take a look at the example purge policy file, make sure you understand what it's doing; ask questions if you have any

- Put purge policy file in place and run the purge

Running a Purge

- Watch the purge run

- Verify results, the old files should now have been removed from the system

Wrap Up

- Further Exploration
  - Other Lustre commands or Robinhood play you find interesting
- When done
  - Can leave things as they are, no need to tear down