

Livermore Computing User Meeting

August 21, 2018

Tim Fahey

LC Customer Services Group Leader
Lawrence Livermore National Laboratory



LLNL-PRES-XXXXXX

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Agenda

- /p/lustre* File System Rollout
 - Tim Fahey, Customer Services Group Leader
- LC Systems Status
 - David Smith, LC Systems Administration Group Leader
- Sierra Software Status
 - John Gyllenhaal, LC Development Environment Group
- RealVNC
 - Cameron Harr, LC Systems Administration Group

/p/lustre*

- LC will be deploying two new 15 PB file systems within the next 30 days, the first on the SCF and the second on the CZ.
- These file systems will both be called /p/lustre1
- The primary differences between these file systems and our existing /p/lscratch* file systems are as follows:

/p/lustre* cont.

- Purging: Unlike /p/lscratch file systems,
these new file systems will never be purged.
- Quotas: LC will implement quotas for each user on these new file systems.
- An “AutoDelete” directory will be available to manage usage and aid in remaining under quota
 - More on this later.

Quotas

- All users will get the following initial quota:
 - 15TB soft quota/20TB hard quota (storage)
 - 900k soft quota/1M hard quota (# of files) [UPDATED: 11/27/18]
- The command `quota -v` will show a user's current usage
- When a user hits their soft quota a 10-day timer starts during which time they can continue to write to the file system
- If after 10 days they have not reduced usage to under quota they will be temporarily unable to write to the file system until they get under quota
- If a user hits the hard quota they will be temporarily unable to write to the file system until they get under quota

Quotas cont.

- Users needing an increase of their initial quota can request one by sending a justification to lc-hotline@llnl.gov.
- Further increases to a users quota will require additional justification, including the duration the quota increase is required. These increases require approval by LC (OCF) or WCI (SCF) and will also be processed by sending email to the LC Hotline
- LC Hotline and LC Operations (24x7) will have the ability to temporarily increase quotas (DAT run, programmatic deadline....)

AutoDelete directory

- Every user will have an AutoDelete directory on each new Lustre filesystem
 - e.g. /p/lustre1/fahey2/0_LC_AutoDelete
 - yes that is a leading number zero in the directory name.
- LC wants to help users manage their data so they can work comfortably within their quota. To help with this LC will be available to help users delete data

o_LC_AutoDelete directory

- Initially any directories a user moves (mv UNIX command) into the AutoDelete directory can be deleted by contacting the hotline. Hotline staff can initiate a parallel remove to delete data located there for you.
- Upon deployment of new SLIC clusters later this fall, directories placed in the AutoDelete directory will automatically be deleted (scan to be run hourly).

Deployment Schedule

- SCF /p/lustre1 online Thursday 8/23
- OCF/CZ /p/lustre1 online Thursday 9/13
- 9/2018-10/2018 will be a delete campaign for /p/lscratchh and /p/lscratch2
- 11/2018 lsh and ls2 will both be renamed/remounted as /p/lustre2 with quotas implemented. Users will have the same quotas on /p/lustre2 that they currently have on /p/lustre1. No data will be lost.
- Symlinks will be created for lscratch2 and lscratchh pointing to lustre2

Additional Information

- LC has developed a number of file utilities that users can use to change permissions, copy files, remove files and list files in parallel. For example, drm (remove) run on 16 nodes across 256 tasks executes in under 3 minutes which serially would have taken 5 hours.
- Details and examples can be found on the LC Confluence site here:
<https://lc.llnl.gov/confluence/display/deg/mpiFileUtils>

Additional Information cont.

- /p/lscratchrza will be renamed to /p/lustre1 at a future date yet to be determined. When this happens, on rzslic the two /p/lustre1 file systems (CZ and RZ) will be mounted as /p/czlustre1 and /p/rzlustre1

Questions





Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.