



EOS at the Fermilab LHC Physics Center (LPC)

Dan Szkola - Fermi National Accelerator Laboratory

EOS Workshop 2019

04 Feb 2019

The LHC Physics Center At Fermilab

The LHC Physics Center (LPC) at Fermilab is a regional center of the **Compact Muon Solenoid Collaboration**. The LPC serves as a resource and physics analysis hub primarily for the seven hundred US physicists in the CMS collaboration. The LPC offers a vibrant community of CMS scientists from the US and overseas who play leading roles in analysis of data, in the definition and refinement of physics objects, in detector commissioning, and in the design and development of the detector upgrade. There is close and frequent collaboration with the Fermilab theory community. The LPC provides outstanding computing resources and software support personnel. The proximity of the Tier-1 and the Remote Operations Center allow critical real time connections to the experiment. The LPC offers educational workshops in data analysis, and organizes conferences and seminar series.

[LHC Physics Center At Fermilab - https://lpc.fnal.gov/index.shtml](https://lpc.fnal.gov/index.shtml)

[CMS Experiment - https://cms.cern](https://cms.cern)

EOS Back Then At Fermilab LPC

- Needed POSIX compliant online area for LPC analysis data
- EOS testbed built at Fermilab around June 2012
- Initially 1 MGM and 3 FST nodes
- Access was by FUSE mount and XROOTD

By May 2013, more than 600 TB was in use with EOS still not being officially in production

EOS Today At Fermilab LPC

- LPC cluster is a 4500 core user analysis cluster
- LPC cluster supports several hundred users annually from CMS users across the U.S.
- EOS is used for LPC user data which tends to be small files with very random access
- EOS storage is approximately 7.12 PB (about 40% increase over last 18 months)

EOS Space Allocation, Usage, and Growth




1 year ago



Today

EOS Growth



EOS Hardware and Layout

- 2 MGM nodes (a third is standing by for the QuarkDB namespace)
- 44 FST nodes
- 96 filesystems
- 4 groups

type	name	status	N(fs)	dev(filled)	avg(filled)	sig(filled)	balancing	bal-shd	drain-shd
groupview	default.0	on	24	1.98	69.97	0.51	idle	0	0
groupview	default.1	on	25	1.86	68.25	0.51	idle	0	0
groupview	default.2	on	24	1.97	70.81	0.63	idle	0	0
groupview	default.3	on	23	1.99	72.78	0.64	idle	0	0

MGM Hardware

MGM servers each have an individual IP address and hostname. An 'instance' IP address and hostname is defined and a virtual NIC is brought up on the MGM currently defined as the master using this instance name and IP address.

- Dual Intel Xeon CPUs @ 2.10 GHz
- 256 GB RAM
- 1 TB system disk
- 2 TB SSD (for /var/eos)
- 10 Gb Ethernet

```
eth0    cmseosmgm01.fnal.gov
eth0:0  cmseos.fnal.gov
```



```
eth0    cmseosmgm02.fnal.gov
eth0:0  cmseos.fnal.gov
```


FST Hardware

FST hardware varies as FST nodes have been added and removed over time. Typically they will have:

- Dual or Quad CPU (usually AMD Opteron)
- 64 GB RAM
- 1 - 2 TB system disk
- 10 Gb Ethernet
- 2 or 3 Nexsan volumes, most are ~71TiB a few are ~50TiB, all formatted as xfs

Older FST hardware usually had 3 volumes assigned for EOS use. Newer hardware is assigned 2 volumes.



How Is EOS Space Allocated?

- Most users get a 2 TB logical (4 TB physical) area enforced by quota
- For groups (usually associated with experiments or projects), a user account is created and a quota is set based on their need for space.
- Some of the EOS space is used to hold rotated EOS logs
- `/lustre/unmerged` - This area is used to hold job output files that are later merged into bigger (4 - 5 GB) files.
- A temp area is defined to hold initial output of user analysis jobs.

LRU Usage

A directory hierarchy exists in EOS for the initial output of CRAB jobs. This user analysis data is later picked up by a separate process and moved to a user-defined area.

LRU rules are defined to clean up this job data after a week or so.

```
attributes.sys.lru.expire.match=*:86400s
```

```
attributes.sys.lru.expire.match=*:1w
```

Access To Files In EOS

- XROOTD (xrdcp, etc.)
- GridFTP
- FUSE mount

The gridFTP service runs on all FST nodes. An F5 load balancer front-ends the gridFTP service. BeSTMan2 was eliminated with the citrine upgrade, as it has not been updated in years and no longer works properly with SL7. There are FUSE mounts on all LPC interactive nodes. On CMS worker (job) nodes, users use XROOTD to access EOS files.

Upgrading from 0.3.268 on SL6 to 4.4.10 on SL7

- First things first
 - All FST nodes were upgraded to SL7 and 0.3.268 is reinstalled on them.
 - We drained each node before upgrading it, which lengthened this step to about 3 months. Late in the process we added 6 new FSTs which gave us a bit of extra headroom and sped things up slightly.
 - No surprises here

Upgrading from 0.3.268 on SL6 to 4.4.10 on SL7 (cont.)

CHANGE PLAN

- Update our local EOS RPM repo with the 4.4.10 RPMs and dependencies
- Shut down FSTs and MGMs
- Compact the namespace with the *eos-log-compact* command
- Copy /var/eos hierarchy to another node, just in case
- Copy scripts, cron-jobs, etc to another node
- Turn over MGM nodes to be updated to SL7
- Make any puppet changes necessary for SL7 and EOS Citrine

Upgrading from 0.3.268 on SL6 to 4.4.10 on SL7 (cont.)

CHANGE PLAN (cont.)

- MGM nodes should now be at SL7 and have EOS 4.4.10 RPMs installed
- Update FST nodes to EOS 4.4.10

So How Did It Actually Go?

- I had allotted 12 hours (8am - 8pm) for the whole upgrade
- MGM upgrade went smoothly
 - I missed a config item or two in puppet, but this was easily remedied
- Most of the upgrade work was finished by about 12:30pm
- The FST upgrades went a bit less smoothly
 - The nodes came up fine, but did not seem to recognize that a full metadata sync was required. Instead an empty LevelDB was created.
 - This was easily fixed by booting the filesystems by hand with the `--syncmgm` option.

What Do LPC Users Say About EOS?

- Many users very upset with changes to the EOS ‘find’ command
- Lack of wild card access over EOS/xrootd commands causes significant difficulty
- No tab completion for interactive commands
- FUSE mount can be unstable, especially with heavy use
 - LPC interactive nodes have been running the citrine client for some time now. I have not heard whether things are better, but LPC staff still try to keep FUSE activity low
 - Also watching work on FUSEX

EOS In The Near Future At Fermilab LPC

- Already have a test stand running QuarkDB, will move production to QuarkDB when it is feature complete
- Still using FUSE, we will look at FUSEx
- We have an eye towards IPv6
- Eliminate or limit the amount of single replica space
- Some users have expressed interest in a CERNBox type implementation
- EOS growth is expected to be approximately 1 PB per year if funding permits

EOS - What Could Be Improved

- Still many ‘mysterious’ MGM crashes
- Documentation could be more complete and is sometimes out of date
- Output of some of the commands is cryptic and is not explained anywhere
- No complete list of config statements for xrd.cf.* files or /etc/sysconfig files and no explanation for some options without digging into the source
- Might like to see easier methods for identifying/dealing with users who are ‘misbehaving’

EOS - What Could Be Improved

- Community support
 - The EOS Community site is a big improvement
 - Good level of participation
 - CERN devs and admins often answer questions
 - Could maybe pin some useful information at the top such as...
 - Officially supported or recommended procedures for upgrades (to citrine, to QuarkDB, etc.)
 - Changelogs, release notes for latest supported release
 - Any useful tidbits, scripts, configuration suggestions admins might find useful
 - As previously mentioned, description of all available config directives

Contributors

Thanks to the following people at Fermilab who provided information for this Presentation:

- Bo Jayatilaka
- David Mason
- Marguerite Tonjes