Experiences at Oak Ridge

EOS Workshop, CERN
5 Feb. 2018

Pete Eby - ebypi@ornl.gov

Slides: http://cern.ch/go/S8dM

Overview of EOS at ORNL

- Deployed a T2 site in spring of 2015
  - Went into production with EOS
  - 970TB initially presented

- During first year and a half expanded storage twice
  - 1PB - 04/2016
  - 1.5PB - 05/2017
  - 2.5PB - Fall 2018 - Installed (not yet available, using for conversions)

- Started seeing issues in Feb 2017

- We'll cover our experiences investigating these issues and share what we have learned.

Hiccups along the way

Over the last year we've encountered a few issues running EOS and in the course of tackling them learn some things that might be useful to the community. When things don't work well it's a great time to understand.

We'll cover each of the following:

- Design decisions and pre-production planning
- Performance and stability Issues as capacity & concurrency increased
- File system conversions project
- Sharing recommended best practices
  - If we'd only known...
- EOS performance tests

But first:

Underneath it all.. Storage configs matter

Many wonderful things happening with EOS, advances with Citrine, scheduler, draining, distributed metadata - great stuff!

- Underneath the stack storage configuration is key, affecting everything above..
- *Issues may not present until much later into production* (as we will see)
- More emphasis and clarity, with best practices for deploying, expanding, and validating storage configs would benefit the EOS community, and are critical for non CERN projects to effectively deploy EOS.

Design Decisions

We designed a multi PB, general purpose, storage environment to perform well under mixed use cases, though lacked specifics on what EOS IO access characteristics and levels of concurrency. What we built wasn't what EOS wanted.

- Design topics encompassed
  - Filesystem, blocksize, per fs capcity, raid alignment
  - Difficulty testing and validating eos environments, validation was to test in production - by then it was too late.
  - Questions about capacity expansion and lifecycle management
  - Our thanks to the EOS dev team!
- Lots of back and for with the experts early on
  - But we were still flying in the dark

Performance and Stability

- Issues arose as capacity & concurrency increased
  - First year 1/2 all was great! But then the magic ended.
  - Aside: Theory of large modern HDs, thus larger pools, and such mapping poorly to EOS concurrency and IO characteristics

- Failing add / get tests
  - Difficulty in reproducing add tests, identifying which layer of the stack it is
  - Put availability report misleading when storage is full

Performance and Stability - Pre

Three year GET availability test graph (the redshift is strong)
Cf three year capacity used graph.

First 1 ½ year (7/15 - 2/18) availability was >97%!

Then entropy crashed the party.

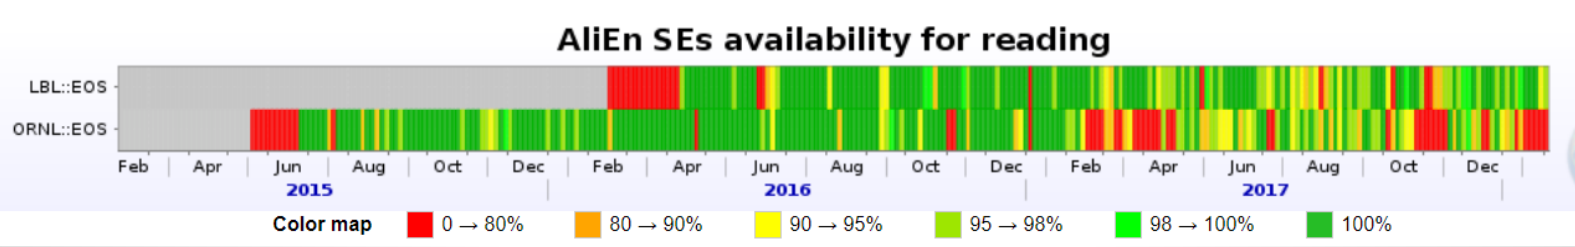

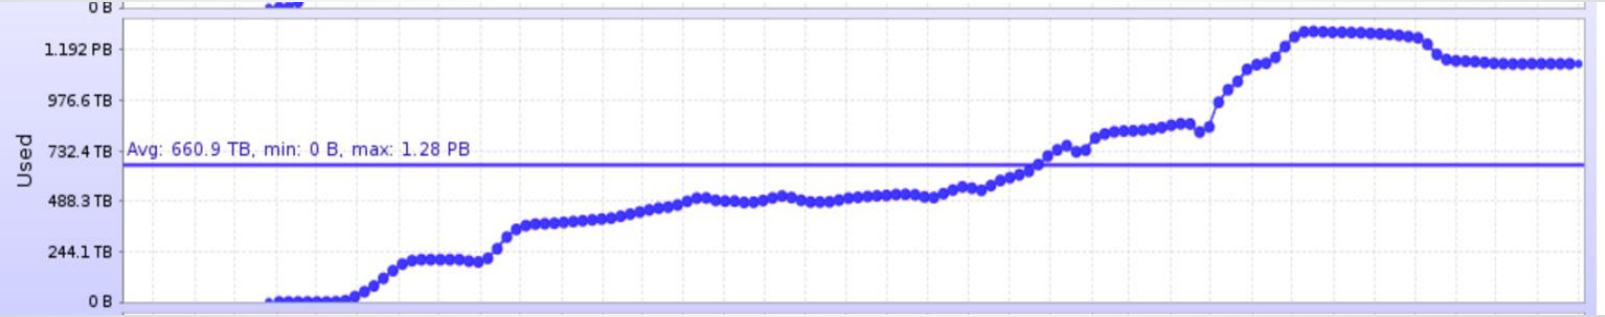

Performance and Stability - Pre

Three year graph of all FSTs
http://alimonitor.cern.ch?2380

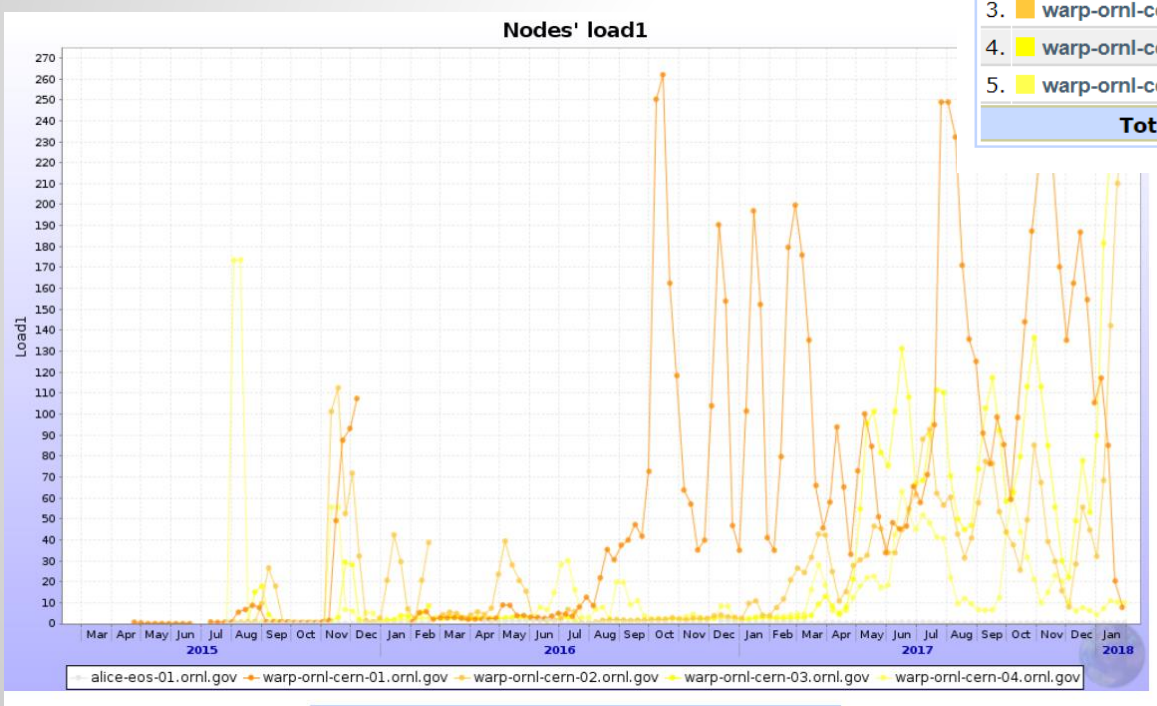

Nodes' load1

	Series	Last value	Min	Avg	Max	Total
1.	alice-eos-01.ornl.gov	0.325	0	0.418	32.01	4040249
2.	warp-ornl-cern-01.ornl.gov	7.799	0	64.09	1558	682541233
3.	warp-ornl-cern-02.ornl.gov	242	0	27.14	2091	262507124
4.	warp-ornl-cern-03.ornl.gov	247.4	0	31.32	2049	302958922
5.	warp-ornl-cern-04.ornl.gov	9.955	0	14.01	2082	135538564
Total		**507.4**		**137**		**1387586094**

Performance and Stability - Pre

Three year graph Job Efficiency
http://alimonitor.cern.ch?2382

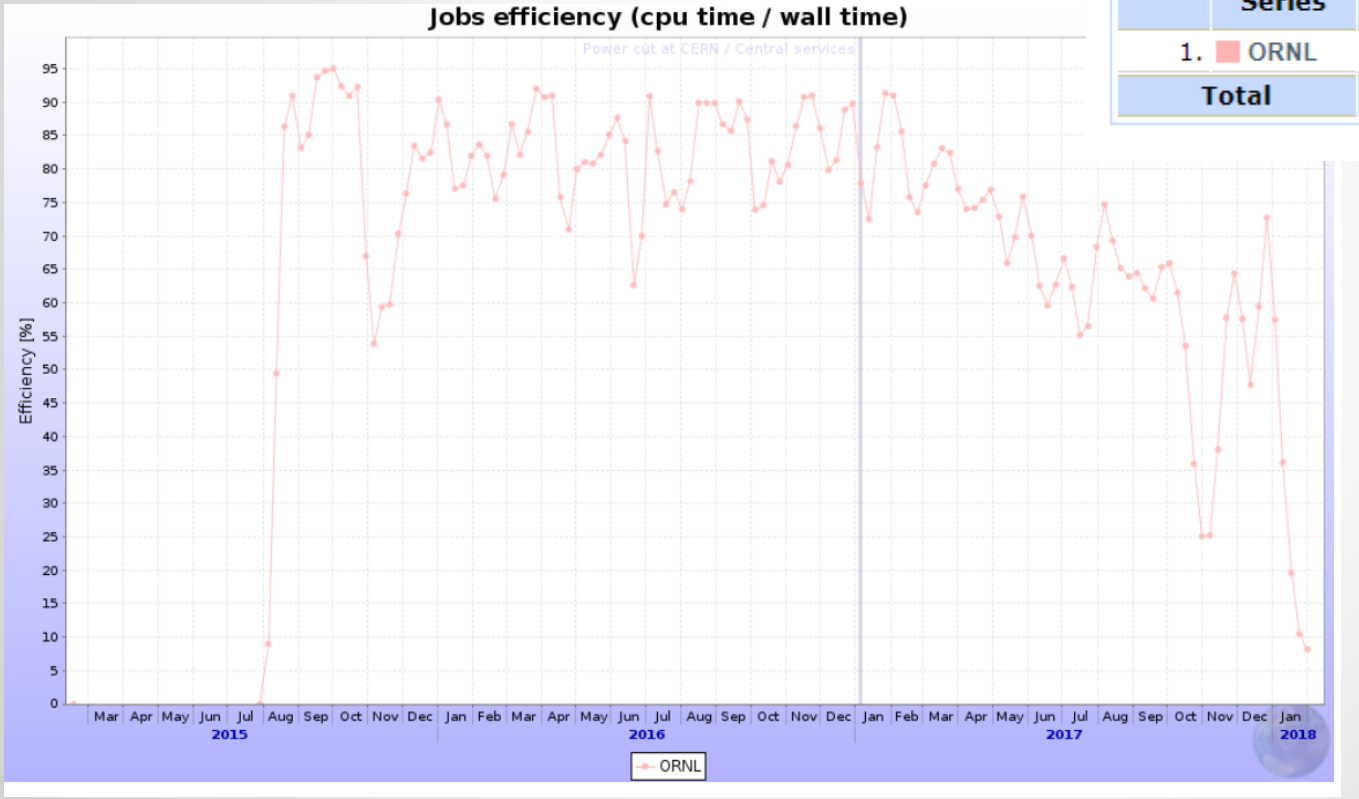

Jobs efficiency (cpu time / wall time)

	Series	Last value	Min	Avg	Max
1.	ORNL	8.152	0	71.92	100
Total		**8.152**		**71.92**	

File System Conversion Project

A long story of the road from 1M to 16k blocksize, and new volume layout.

- Iteratively tested many filesystems with different parameters
  - Fs type (zfs vs. xfs), blocksize, volume layouts, parity levels, readahead, descendent filsystems per pool, etc.

- 4K is “best” for EOS - but you lose capacity
  - Converting one FST = ~20 space bloat
  - Blocksize is critical to fs performance. The EOS desired (4k?) should be shared.

- Moving from 1M to 16k blocksize - when you have no space
  - Top Tip: Always keep some space. (eos quota, or hold a few fsid or a group back from production.)

Performance and Stability - Post

Three month comparison serving the same ten eos fsids

Yellow = 1MB blocksize, black = 16k (after migrated)

http://alimonitor.cern.ch?2375

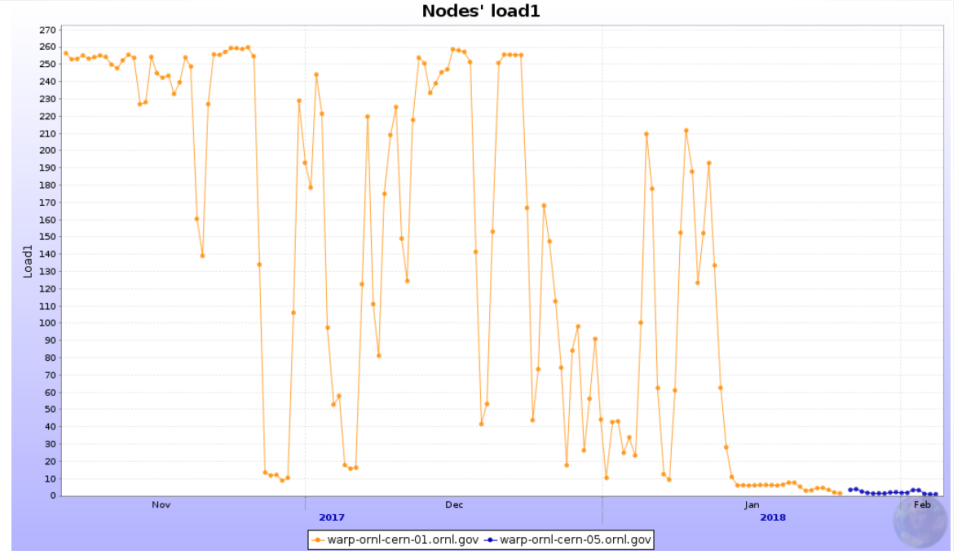

Nodes' load1

	Series	Last value	Min	Avg	Max	Total
1.	warp-ornl-cern-01.ornl.gov	1.452	0.03	139.2	278.3	118342433
2.	warp-ornl-cern-05.ornl.gov	0.945	0.18	2.038	15.45	191043
	Total	**2.397**		**141.3**		**118533476**

Performance and Stability - Post

Last month ORNL::EOS Network Traffic

http://alimonitor.cern.ch?2381

Orange was old FST, Dark blue new one serving converted 16k fsids

IE: The FST serving the most data now has lowest system load

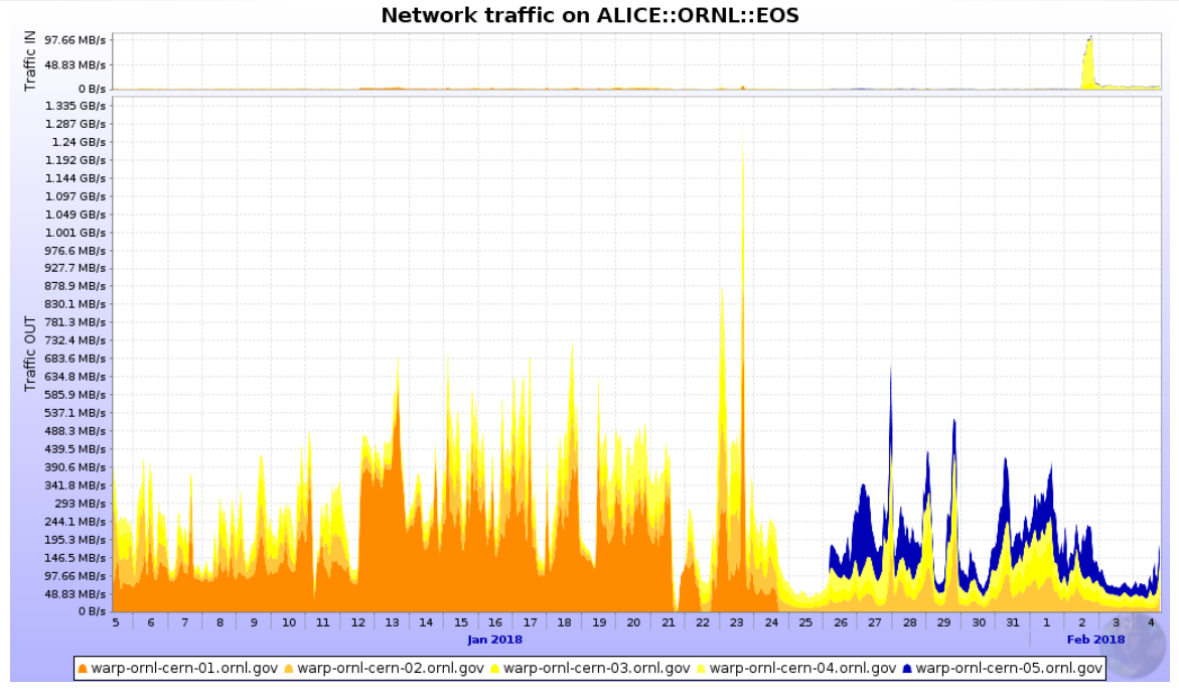

Sharing Recommendations

Knowledge is locked in, and distributed among, EOS community members.

How to best share this invaluable information for new deployments, expansions, troubleshooting, etc?

- EOS can benefit from a “Suggested Practices” guide for storage configuration
  - Share the IO characteristics EOS expects - this is huge.
  - Suggested storage layouts

- Forecast for RAIN
  - CERN-IT is moving to RAIN and one disk per eos fsid
  - There are reasons for this, but not well communicated

EOS Standardized Performance Tests

The O2 Disk Buffer WP15 Project is developing a test framework. Information is available at:

Slides: http://cern.ch/go/CNF6
Gitlab: https://gitlab.cern.ch/alice-o2-wp15

Test Framework Goals

Validate O^2 disk buffer storage environment through the development of an EOS testing framework which uses synthetic (fio, etc.) and simulated O^2 workloads under expected levels of concurrency for standardized, reproducible results and SE performance analysis.

Our Next Steps...

- Complete our fsid conversions
  - Considering preliminary success of fsids converted thus far
  - Retain some empty fsids for future drains, conversions, etc.
- Upgrade to Citrine
- Concurrently work on the EOS Testing Framework
- Would like to contribute to an EOS storage design blueprint or suggested practices guide (in gitlab, the wiki, etc.) if there is interest.