

Data Handling Feedback from NOvA

Alex Himmel

FIFE Workshop


June 21st, 2017

Introduction

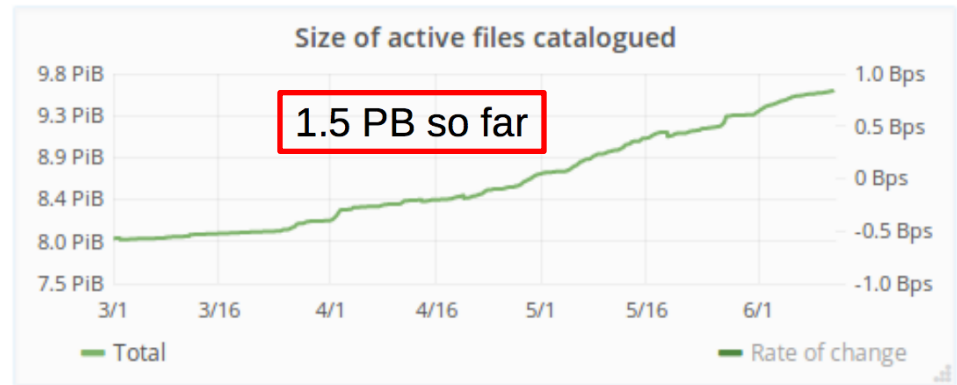
From 30,000 feet

A slide from Jeremy Wolcott, our production convener at the last NOvA collaboration meeting.


From 30,000 feet, I'd call that generally successful.



Running Current: 9.72 K
Idle Current: 760 Held Current: 499
Quota Current: 2.40 K



▼ SAM



We've been busy!

Production

- Generally, things are running smoothly.
 - Handling large scale running smoothly, sometimes 15k jobs at a time.
 - In addition to typical resources, we have a dedicated pool for staging raw FD data from tape.
- #2 complaint – the slowness of SAM.
 - If SAM queries reliably executed in seconds, working in production would be *much* easier.
- #3 Pre-staging
 - The samweb run-project command finishing does not seem to mean that all files are available on dCache.
 - The only reliable way to check this is to check each file, which can take 8 hours or more on the types of datasets that need pre-staging.
 - Can we implement some asynchronous notification once a SAM dataset has been fully pre-staged?
 - We believe this has been at the root of some low-efficiency problems we have had recently.

Analysis Users

- Analysis users on NOvA have been (mostly) moved to dCache for file access on the grid.
 - We keep mirrors of small analysis ntuples on persistent dCache as well as tape so most users do not need to worry about staging.
- Under most circumstances, this works smoothly.
 - xrootd and SAM tools have been tightly integrated into our analysis framework, both interactively and on the grid.
- But, that brings me to our #1 storage issue...

- #1 Periodically, things fall over due to load.
 - Every 1-2 months we have users complain about numerous jobs failing with xrootd timeouts.
 - We have a ticket open now on this issue.
 - Discussion there has turned to reducing the number of movers so that the timeouts are on the client rather than server end.
- The success rates listed below seem very low for what looks like a pretty reasonable job loads.
 - Do more resources need to be directed to the persistent dCache pools?
 - Asking everyone to move to dCache means dCache needs to be able to handle the load.

