MINERVA FIFE REPORT

6/16/14

Minerva experience with FIFE products

Overview

- Minerva has been taking data since 2009
- 300 TB of bluearc
- Now moving to dcache for production
- SL5 for now – SL6 development in progress
- Use 500-2000 batch slots
- Rest of ths talk will go through the products we know about and are using

Standard Products that we use

- jobsub
- samweb_client
- lfdh
- dcache
- fts
- Pnfs
- Sfa
- lfbeam
- fifemon

Jobsub

Standard job submission interface for IF

- We use this for most of our job submissions
- Generally very successful
- Remaining issues are bugs
  - mainly due to need to accommodate minerva features that predate jobsub – like e938 group instead of minerva group
  - And short term problems due to dcache migration
- Currently using v1_2p – need to test more recent version

samweb_client

- Interface to the sam db and file delivery systems
- We use this very successfully as a file catalog
- Much easier to use than the old sam API
- Major issues:
  - A bit slow – doing bulk changes to metadata can take days when dealing with 10^6 files
  - We still have a lot of legacy sam code left to convert.
  - Still testing a metadata.py to json converter
- Future:
  - We are not using sam delivery methods in production yet

ifdh

- Intensity Frontier Data Handling
- Uniform interface for data transfers
- Has been very helpful in moving to dcache reasonably transparently
- Issues:
  - We still have legacy cpn's hardcoded in some scripts
  - Campaign to find and change them
  - File ownership issues in some cases.
  - We will be very glad to see the end of cpn LOCKS which lead to very inefficient processing

dcache

- The new disk cache
- Much faster transfers and no more locks!!!
- Issues:
  - NOT a normal disk system – this implies major rewrites of many of our disk dependent products
  - Access via ifdh works great
  - Have figured out how to use root standalone
  - Many products need to be changed to allow dcache to be used transparently – Event Display!

cpn activity for minerva

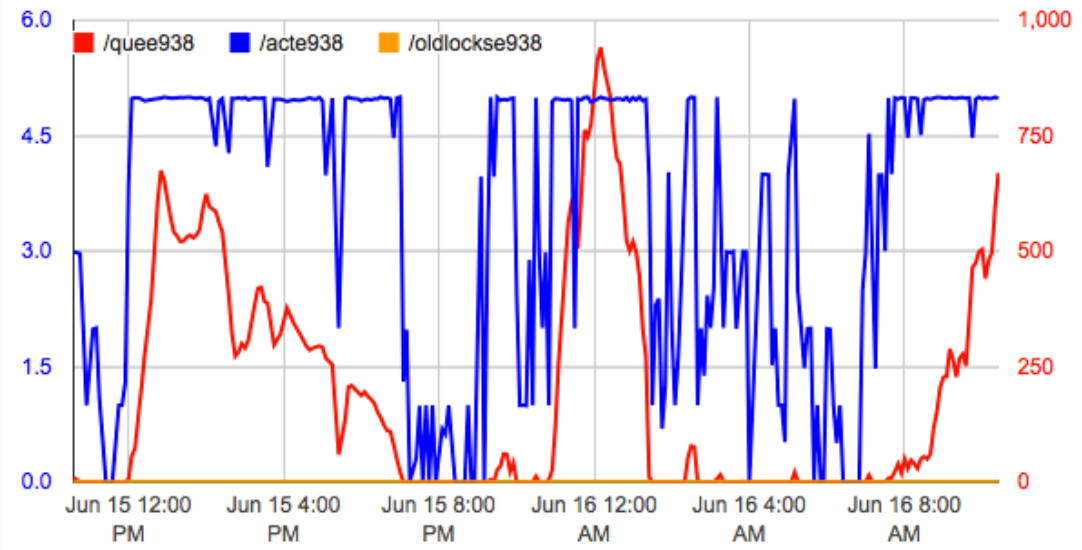

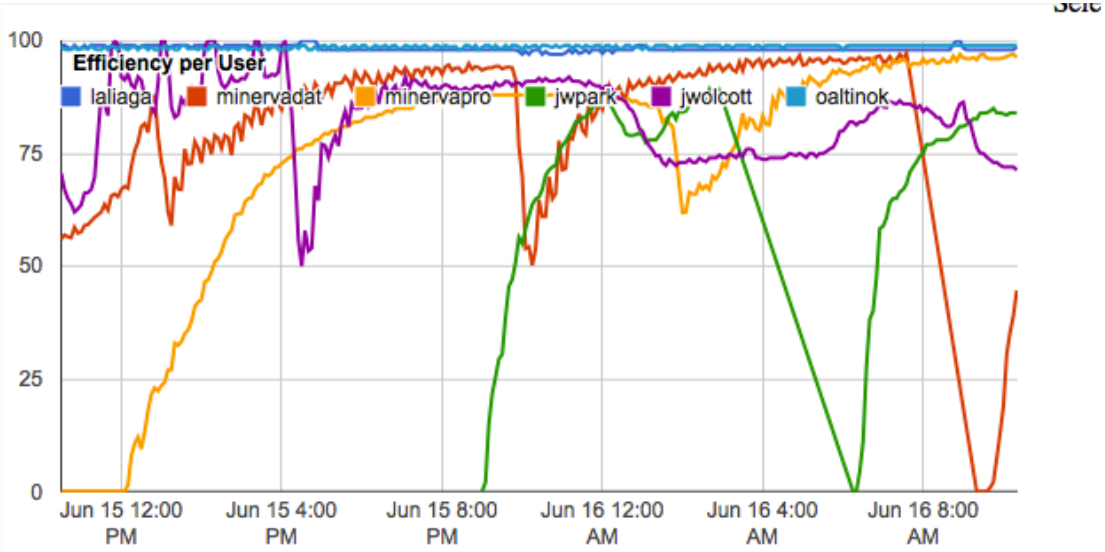

http://fermigrid.fnal.gov/scoreboard/

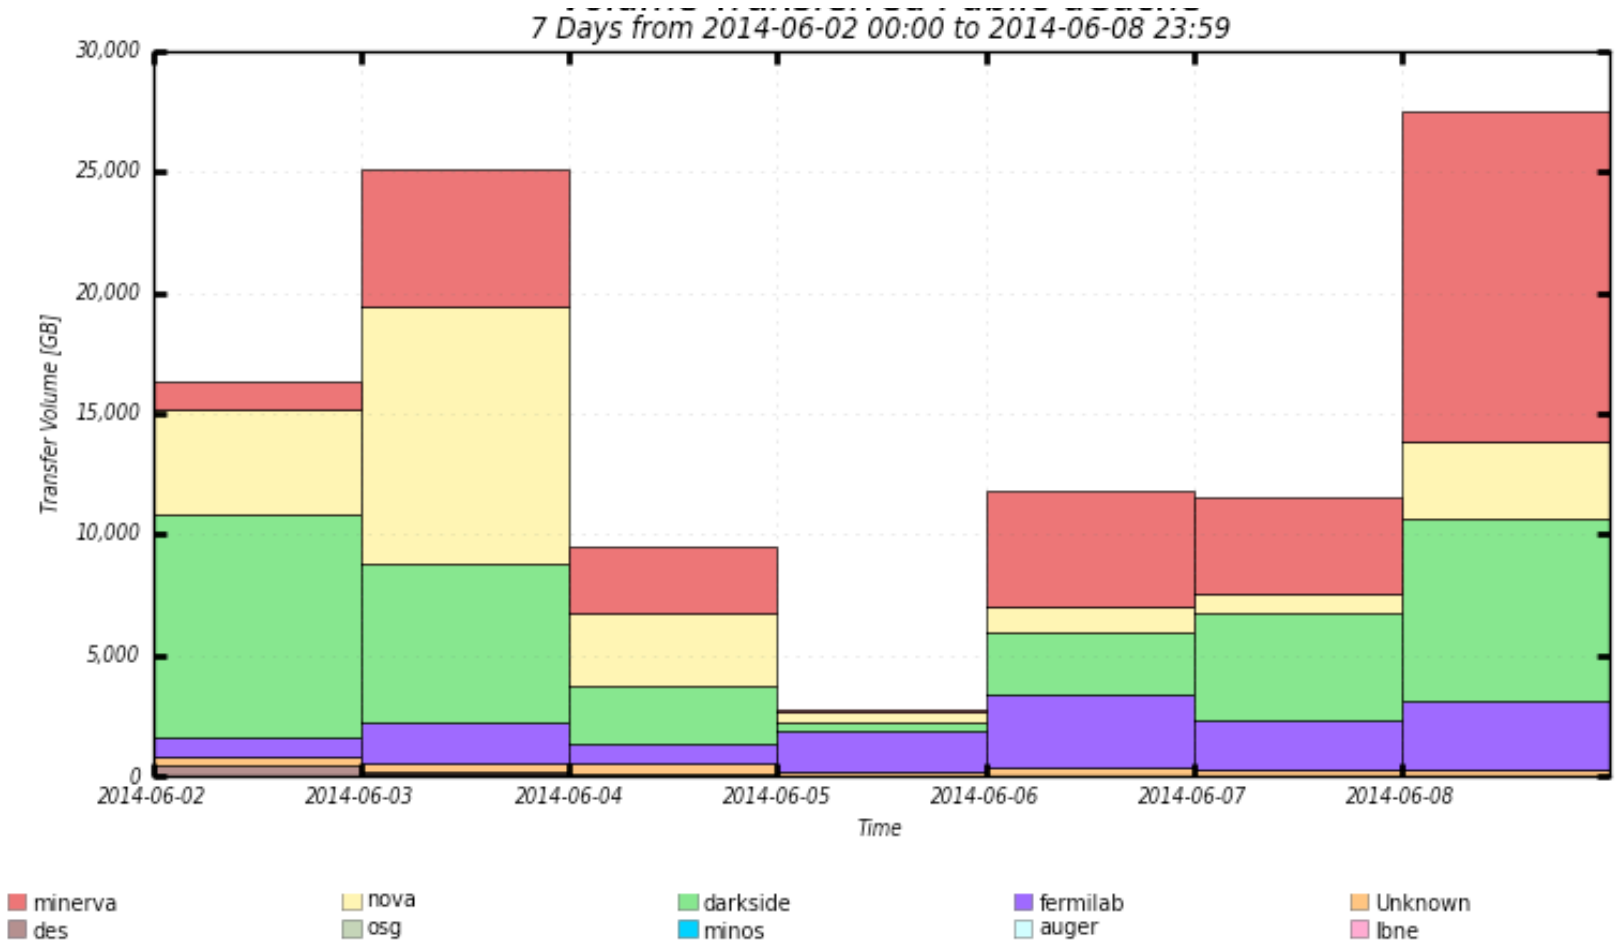

Fts

- File transfer system
- We run two instances
  - Bluearc for raw data/early processing
  - Dcache for reconstruction/MC
- No major issues except the confusion that results from running two of them
- Not autogenerating metadata
- Put simlinks in the dropbox – file deletion is separate.

SFA

- Small File Archive
- Our typical files are 100-500 MB
- Had some problems when files are only on tape
  - Largely resolved by storage group bulk staging the files.
  - Files are now stored via dcache so are cached already.

pnfs

- One issue here for users – no idea what they are talking to

- /pnfs/minerva/data
- /pnfs/minerva/scratch

- Are VERY different beasts – scratch is disk only, data migh be tape only
- Perhaps alias the different pools more transparently?

Fifemon

- The fife monitoring pages
- Very convenient and useful
- May want to add more monitors to them – in particular FTS transfer summary from http://samweb.fnal.gov:8480/sam/minerva/samdftsgpvm01.fnal.gov/fts/status

lfbeam

- The beams database
- It works….

ECL

- Electronic Control Room Logbook
- We depend heavily on this!

Archiving?

- We're old enough to have lots of legacy stuff from users
- Have /pnfs/minerva/archive/users/ area which is a tape file family in SFA
- Plan to dccp (ifdh?) dead user areas to that – should be write once, read never.

Major issue is documentation

- Most fife products are still in development
- Documentation is HARD when things keep changing!
- We strongly support the effort to provide professional documentation for FIFE and are happy to help.