FIFE Workshop 2015
CVMFS news, best practices

DAVE DYKSTRA

News since last year's workshop

- Most VOs have switched to locally-hosted CVMFS opensciencegrid.org repositories
  - Installation performance is much improved
  - Local kerberos authentication
  - Server directly mounts bluearc source, can auto-sync
  - Currently there are 9 VOs plus "fermilab": darkside, des, gm2, lariat, lsst, minos, mu2e, nova, and seaquest
  - Phasing out use of oasis.opensciencegrid.org and *cfs.fnal.gov repos
- Small projects share fermilab.opensciencegrid.org with common ups/upd products
  - Currently larsoft and genie
  - Auto-syncs each night from /grid/fermiapp
  - Projects can also do update on demand

News since last year's workshop

- FNAL Stratum 1 upgraded to faster, larger disk space machines that are now managed by Distributed Computing Services (DCSO) group
- Support has been added to GlideinWMS pilots to probe for CVMFS errors to avoid job "black holes"
  - Plans are being made to add per-experiment customization
- cvmfs & cvmfs-server 2.1.20 released
  - includes garbage collection option for nightly build repos (used by three cern.ch VOs but would still require some operations discussion for OSG VOs)

OSG CVMFS/OASIS news

- A couple of egi.eu repositories now directly imported to OSG, more can be added on request of an OSG VO
- fermilab VO now officially known to EGI, but we are still awaiting friendly sites to accept the VO before EGI imports requested OSG repositories
- OSG GOC oasis servers' software now provided by development team (which includes me)

CVMFS end user best practices

- If a site cannot mount your CVMFS repository, submit a FIFE support service desk ticket
  - All OSG sites are supposed to be able to mount all repositories but some sites aren't configured correctly
- Work around cases where updates can appear to be delayed for hours on some but not all batch nodes
  - In those cases a cvmfs repository was mounted but idle, and the cvmfs client is taking one minute to flush buffers while letting the application continue
  - Workaround: run /cvmfs/grid.cern.ch/util/cvmfs-uptodate <acvmfspath> at beginning of job (if pilot doesn't do it)

CVMFS repository maintainer best practices

- Set up .cvmfsdirtab with wildcards matching every software release directory for application and external packages
  - Avoids catalogs getting too large (keep under 200K files each)
  - Avoids loading info about files that will not be used
  - Avoids generating as much garbage when anything in a catalog changes
- To sync from another filesystem use cvmfs_rsync
  - Avoids subtle problem when removing old releases with catalogs
  - /grid/fermiapp/cvmfsfermilab/sbin/cvmfs_rsync
- Make sure all files are world readable

CVMFS repository maintainer best practices

- Avoid data files that not similar sizes and access patterns as executable software
  - All jobs in a batch of jobs should generally access the same files
  - Typically the total read per job should be about a Gigabyte compressed or less
  - Larger amounts of data or randomly accessed different data should go into high bandwidth storage (dCache)
  - Tar files, etc, are better if they are unpacked
- We have had much success with large data files in CVMFS cached on high bandwidth storage
  - Still a research project, however
  - Not yet extendable to all OSG sites

CVMFS repository maintainer best practices

- Also generally best to avoid source files in CVMFS
  - Does not affect client or squid performance, but it multiplies the number of small files on Stratum 1s which affect their performance
    - In the future, different implementations on Stratum 1s may mitigate this affect
  - Not a requirement, a best practice
  - Source required for compiling is fine, but avoid very rarely accessed source such as for debugging if possible

FIFE CVMFS Documentation

- https://cdcvs.fnal.gov/redmine/projects/fife/wiki/Introduction_to_FIFE_and_Component_Services#OASISCVMFS