

SW, Conditions Data and CreamCE

- Rod Walker, LMU Munich

- ATLAS SW installation
- Conditions Data access
- Cream CEs

ATLAS SW

- Disk space
 - base release ~ 20GB. Currently have 20, ~ 400GB
 - some sites have problem to provide this space
 - especially if high performance system needed to scale for more WNs. E.g. mirror NFS area.
- Setup & build performance
 - CMT touches $O(10k)$ files during build
 - even setup to run athena is very heavy
 - Panda and gliteWMS default is build once, run many
 - but users can choose to build in every job
 - this recently crippled PIC
 - too rare to rush some protection.
 - all we can do is educate users

ATLAS SW(2)

- CVMFS – CERNVM file system
 - Fuse intercepts file operations and directs to local disk cache, or first pulls file via http
 - only cache files which are used
 - identical files(by hash) are only cached once
 - 90%* of files in ATLAS releases are common
 - all releases will fit in 50GB cache*
 - probably much smaller cache fine for personal use
 - local Squid proxy provides site cache for multiple cvmfs instances
 - Feasible to have cvmfs on WNs?

* - this number is made-up


ATLAS SW(3)

- Possible to NFS export a cvmfs area
 - run cvmfs with cache on SSD
 - export to WNs via NFS, i.e. no change to WN
 - very nice solution, but not tested
 - needs recent kernel to export via NFS, i.e not SL5
- All releases instantly available
 - until recently the updates were manual and weekly
 - now done by SW release/cache shifter

Conditions Data access

- Oracle access via Frontier
- Flat file access via direct access to SRM
 - large blobs unsuitable for Db
 - referred to by guid from the Db

Frontier with T2 Squid cache



- Frontier server with Squid(accelerator)
 - CERN, KIT, PIC, LYON, BNL, TRIUMF
 - access local Oracle RAC
- Squid caching web proxy at T2/3

Failover strategy

- Jobs will try 2 Frontier servers, via 2 Squids, and then try direct to Frontiers
 - not really known how often jobs failover
 - squid and frontier logs not monitored
 - total failure is very rare
- CMS have single Frontier instance
 - If T2 Squid fails, only failover is direct to Frontier
 - log is closely monitored, and Squid problem reported
- I like the chaotic, silent failover
 - but should improve monitoring

Flat CD files

- Also need flatfile conditions data
 - 500GB/year. Accessed via remote io to SRM.
 - PoolFileCatalog maintained on site with TURLS
- AdeS sw install builds PFC for files in HOTDISK
- Athena searches ATLAS_POOLCOND_PATH
 - PoolFileCatalog.xml in
\$VO_ATLAS_SW_DIR/local/conditions/poolcond
 - has dcap or rfio turls
 - also install/configure various fixes for remote io

Flat CD files(2)

- Are a real pain
 - is Db but needs completely separate machinery
 - DDM subscriptions to HOTDISK to distribute – fine
 - frequently build PoolFileCatalog used by athena to get turl for a guid.
 - many problems of stale or corrupt PFC
 - need to add SAM test
 - remedial PFC with http turls
 - used when rfio not working(RAL,ASGC)
 - works very well, with some performance hit

ATLAS SAM tests

- Alessandro(di G) added an ATLAS SAM test for Frontier access from the WN
 - part of CE test
 - extracts serverurl and proxyurl, for primary and backup, from FRONTIER_SERVER env
 - run fngget 4 times for all permutations (2Frontier, 2Squid)
 - if all fail, it`s an error
 - if any fail then it`s a warning
- Running for weeks, but does anybody look?
- Need PFC test

CREAM CEs

- CondorG submits panda pilots for prod and analysis
 - Need CondorG submission to Cream CE
 - is supported but ...
 - required gridftp server running as root on submit host
 - WN pulls and pushes file directly from/to submit host
 - no control can lead to high loads on submit host
 - saw jobs block WNs for hours on stage-out
 - few Cream CEs in test – maybe misinterpreted result
 - e.g, Glasgow problem maybe related to NAT tcp issue
 - subsequent tests at low scale due to lack of prod jobs
 - must test in controlled way with fake jobs

Cream CE(2)

- Condor 7.5.3 removes need for gridftp server
 - supposedly gridmanager initiates transfers
 - but this implies to/from CE, not WN
 - don't understand this so need to test
 - just installed 7.5.3 and trying cream submission
- what does this mean for the sites?
 - Cream CE looks better for submit rate and status updates
 - install cream but don't stop LCG CE yet

Final thought

- Squid needed for many things we do
 - frontier, CVMFS, flat CD access, various wget/curl from WN(pilotcode download, analy code)
 - other VOs do some of these, or maybe other use-cases
 - or will do ... several data demos want proxies
 - **should be a required WLCG service**
 - updates, monitoring, cache peer config
- partial file caching would help, e.g. vector read
 - open task, without developer, in Squid opensource project http://devel.squid-cache.org/stale_projects.html -anyone?