

AAA Federation

Brian Bockelman

“Any Data, Any Time, Anywhere”

Fabrizio's List of Questions

- As a guide, here are the things I'm trying to answer:
 - Their ultimate goal in using the technology
 - Current usage of the tech: status, interests, caveats (e.g. name translations)...
 - Sites that are participating, which ones are joining
 - A message to the sites that should install their federated endpoints in the next year
 - What should they do to join
 - What level of effort is expected (e.g. manpower, debugging, babysitting, ...) for the site
 - Which are the parts that sites could share among experiments

Goals of AAA


- Increase the data accessibility for physics.
- Deliver tools to decrease the barriers between physicists and data.
- Increase the portability of the CMS environment.
- Remove the data locality requirement and increase the number of sites where a job can run.

Current Usage

- We have a multi-layered hierarchy.
- US redirector containing all US T1/T2 sites.
- EU redirector containing a smaller subset of European sites, including EOS. Sites from Italy, UK, Germany. Finland is working on joining.
- Doesn't cover all CMS files, but probably has 90% of those relevant to analysis.

Example from Accounting

Volume of Gigabytes Transferred By Facility
14 Days from 2012-07-08 to 2012-07-22



Legend: GLOW_Internal (red), Nebraska (yellow), MIT (green), USCMS-FNAL-WC1 (purple), UFL (orange), UCSD (brown), Purdue (light green), T2_IT_Bari (cyan), GLOW (light blue), Vanderbilt (pink)

Maximum: 51,728 GB, Minimum: 13,686 GB, Average: 29,930 GB, Current: 16,216 GB

Current Usage

- We invite sites to (optionally) participate in our monitoring and accounting.
- The monitoring we receive allows us to do deep, meaningful studies of I/O performance of our studies.
- The only CMS-specific aspect is namespace translation. In CMS, this is done via an XML-based mapping file; we have a plugin to parse this file and map the namespace accordingly.

Overflow and Fallback

- If a job encounters an error when opening a file, we will fall back to the regional redirector.
- When a job in the US has been in queue for more than 6 hours, we will run it at any US site, regardless of data locality.
- If it runs at a site without the input data, fallback will be invoked immediately.
- About 5-10% of total analysis, depending on the day.

Joining the Federation

- All CMS sites are encouraged to join the federation, or at least setup fallback.
- We recommend a T2 have at least 3 servers participating (2 for load-balancing, 1 for failover).
- Export the CMS namespace.
- Implement Authz and Authn according to CMS policy.
- Join their respective region, or just the testbed if they want to gain confidence in the system
- For more details, open a support ticket with the xrootd Savannah squad.
- <https://twiki.cern.ch/twiki/bin/view/Main/CmsXrootdArchitecture>

Effort Levels

- The effort levels are fairly minimal; for most site configurations, similar to running a web server.
- Mostly related to setup, as sites need to make configuration decisions. The rest “just works”, to the extent that any storage service works.
- For US sites, software, documentation and basic support comes from the OSG.
- Periodically, we’ll send out requests to upgrade the software, restart services (if we are monitoring them), or change configurations.
- Keep the server on, make sure the processes are running, keep watch over Nagios.

Common Among Experiments

- The only CMS-specific piece is the namespace translation.
- The security plugins are likely only used by CMS, but there's no particular limitation. Just an LCMAPS callout.
- A single Xrootd daemon *could* export to multiple federations.
- But the more pragmatic install would be to configure one daemon per VO. The xrootd config file format can support multiple instances per config; the init scripts are setup to do this cleanly also.