

Data Management cluster summary

David Smith

JRA1 All Hands meeting, Catania,

7 March 2007

- **Current DPM status**
 - Around 90 sites using DPM
 - Version 1.6.3 has just been released and certified
 - First certified version to include the SRM 2.2 server
- **LFC version 1.6.3**
 - Support for bulk query for replicas from list of GUIDs
- **Next DPM/LFC release(s)**
 - Support for secondary VOMS groups (DPM & LFC)
 - Development is done and is undergoing internal testing
 - Support for ACLs on disc pools (DPM only)
 - To be done after final clarification of proposed solution

- **Other features of the last release**
 - DPM 1.6.3 also has more *admin friendly* logs files (some complaints from sites about the previous messages)
 - Move to 1.6.3 of the DPM has required a schema change
 - further schema changes also required for the forthcoming support for VOMS secondary groups and pool ACLs:
 - *not decided whether to make one release with both new features, or two releases (if, for example the ACLs on pools takes longer than expected to finalise)*

- **FTS production status**
 - CERN and all WLCG T1 sites on FTS v1.5
 - (from gLite release 3.0)
 - 6 PB exported from CERN since SC4

- **FTS 2.0 in testing on CERN FTS pilot service**
 - Target for production deployment at CERN is (start of) April 2007

- **FTS 2.0**
 - Delegation of proxy from the client to the FTS service
 - Improved monitoring capabilities
 - More data retained in the database, some new methods to access them in the admin API

- **FTS 2.0 continued**

- *Alpha* SRM 2.2 support (people need to get experience with SRM 2.2 in general, not that the work is alpha quality)
- Some changes in schema for future functionality
- Currently deployed on the pilot service at CERN
 - Under internal testing since December
 - Next step: open pilot to experiments to verify full backwards compatibility with experiment code [mid-March]
- Aim to deploy at CERN early April
 - Subject to successful verification
 - Implies release to certification during the next couple of weeks
- Rollout to WLCG T1 sites a month after that
 - Expect at least one update will be needed to the April version

- **FTS - Further work**

- Evolve the SRM 2.2 code as we understand the SRM 2.2 implementations (and aspects of the specification)
- Improve service monitoring
 - FTS will have the capacity to give very detailed measurements about the current service level and problems currently being observed with sites
 - Integration with experiment and operations dashboards
- Site grouping in channel definition
 - To make it easier to implement the computing models of CMS and ALICE
- Better service administration tools
- SRM/gridFTP split
- Notification of job state changes
- Pre-staging support using SRM 2.2

- **Not planned**

- Support for a non-Oracle version
- Support for more complex data movement orchestration

- **GFAL and lcg_utils, no update for some time but now preparing a new version**
 - Undergoing internal testing now, will then submit a patch
- **Several bug fixes**
 - GFAL fix bug 24399 (aka GGUS 9332): 'lcg_utils API functions break LCG_GFAL_VO environment variable' of particular concern for DILIGENT
 - Other fixes:
 - GFAL: 23844, 21916, 15887
 - lcg_utils: 24245, 23844, 21030, 18311
- **New functionality**
 - Support for SRM 2.2, including pre-staging

- **BIOMED**

- BIOMED VO would like certain data to be 'visible on the grid', with the ability to replicate it to ordinary sites with ordinary SEs
- Have their own storage system (comprising a server and database) called the DICOM server
- It is required that data be encrypted after leaving the DICOM server and before going elsewhere on the site or the grid.

- **Currently BIOMED are**

- Using a modified version of dCache with gLite I/O and Fireman to provide access the DICOM service (all together called the *DICOM SE*)
- Hydra is used as a keystore for encryption keys, with the hydra and EDS clients (which in turn use the gLite I/O client) to access the data, retrieve the key and decrypt or encrypt the data.

- **Requirement summary**

- Files/pictures are in the DICOM server
- LFN can be study/series/ID, which can be used with DICOM (no need for hierarchical directories and custom file names)
- SRM interface
- Efficient access, e.g. gridFTP or https - either are OK
- Mandatory encryption
- ACL on each file or picture (synchronized between SEs)

- **Possible solutions - two seen**
- **#1 - 'DPM' solution**
 - Need to develop the DPM staging backend
 - Need to write the DICOM staging backend plug-in, with a callout to Hydra, to copy between DICOM server and the DPM
 - ACL synchronization needs to be solved
- **#2 - from scratch solution**
 - Write an implementation of a DICOM SRM service
 - Not trivial but it is independent of DPM
 - DICOM module with a call-out to Hydra for permission check and encryption key, plus the actual encryption or decryption
 - No separate staging, and one set of ACLs removed (but there is still need of an ACL sync procedure)

- **ACL synchronization**

- ACL synchronization between ACLs on access to Hydra keys and access to the data stream out of the DICOM SE
 - In 'DPM' solution this is between Hydra and DPM
 - In 'from scratch' this is just Hydra
- But also need consistency between ACLs on replicas at arbitrary sites
- All the above is currently provided by the combination of Hydra, Fireman and glite I/O at any involved storage elements
- We believe we will need to provide explicit synchronization between LFC, Hydra, (DICOM SE) and other SEs.

- **Hydra client**
 - Have prepared a version that uses GFAL for I/O rather than gLite I/O
 - In combination with EDS this will give clients access to the data and encryption/decryption capability once the new DICOM SE is provided.
- **Evaluating possible methods of ACL synchronization**
 - For example looking at ActiveMQ as a possible way to message storage elements about ACL changes
- **Ongoing work to provide https data access for DPM**
 - Provides a useful data access protocol in the 'DPM' solution
 - Also useable for the 'from scratch' solution

- **Great deal of work on build with ETICS, for SL4**
 - SL4 32bit builds of clients (and some of the servers) now done
 - Working for the rest
 - Then 64 bit client (& server) builds
- **gSOAP - We wanted to synchronise versions**
 - But different components have various issues, so currently
 - 2.6.2 for DPM SRMs, LFC DLI, (FTS v1.5)
 - 2.7.6b for FTS (v2.0)
 - Possible target is 2.7.9(x) for all components
 - but only after the current problem in gSOAP 2.7.9(a-d) is solved
 - see sourceforge item: 'Content-length should be optional in input'
 - http://sourceforge.net/tracker/index.php?func=detail&aid=1672548&group_id=52781&atid=468021

- **Plan**
 - Complete tasks related to the build
 - Deal with any issues with SRM2.2 (may affect DPM, FTS, GFAL and lcg_utils)
 - Testing of new functionality (e.g. DPM, FTS and GFAL)
 - BIOMED design choices and work

- **Plus**
 - Service and user support, bug fixes