

Upcoming Features and Roadmap

Ricardo Rocha


(on behalf of the LCGDM team)




- We are now able to release more often
 - 1 month between 1.8.4 and 1.8.5
 - Probably 1.8.6 still in December
- Smaller releases, less intrusive
- Available early in epel-testing (more later)

- But we can even do more
 - Independent component releases
 - DMLite plugins, individual frontends
 - Should we keep a global number anyway?

- GLUE2 info publishing
- New endpoints publishing on the infosys
 - HTTP/DAV, XROOT
- Some small bug fixes



- Fix for a report on a security vulnerability
- Database connection pooling on the old daemons
 - Default enabled only on the LFC
- Other small bug fixes
- Beta: Plugin taking into account current I/O load
 - More later



- Update for GridFTP/XROOTD/NFS frontends
 - With versions using DMLite
 - Faster... and able to handle multiple pool types
- Dropping pool accounts requirement
 - We never actually make use of them
- Update of DMLite
 - Improved replica handling methods, lots of bug fixes
- Beta state
 - Tool for hot file replication / rebalancing
 - Plugins for HDFS and S3 support (already there since a while, but first official release)

Some features worth discussing




- Some sites have seen issues when under load
- Too many files on a single disk server
 - DPM redirects a high number of requests there
 - Filesystem then blocks under I/O wait
- Solutions
 - Poor man: DMLite plugin taking into account recent disk server access history (for replica selection)
 - Better: DMLite tool taking care of rebalancing the nodes (obsoletes filesystem weights?)
 - Also better: DMLite tool taking care of replicating hot files

- HTTP, XROOT, NFS... all support client redirection
 - Which fits our model... separate metadata from data
- GridFTP doesn't... or our plugin doesn't
 - Not a problem if using the SRM... but bad if not!
 - But the specification does (or at least it looks like it!)
- Possibility 1: SPAS (striped access)
 - Should be enough for redirection
 - Should give us parallel data access too
 - Client support?
- Possibility 2: GridFTP v2 extensions
 - dCache seems to support extensions for this
 - Globus on the way to implement them
 - The calls should look the same as for SPAS


- RFIO is still the default almost everywhere
 - But we've seen the performance numbers
 - We should slowly phase it out in favor of HTTP or XROOT (or later NFS)
- Not an issue when clients directly use protocols
 - Not the case when requesting via the SRM
- Thanks to
 - ASGC / ShuTing for the numbers
 - And UK / Wahid / Sam for the XROOT federation testing
- We now need larger scale deployment

- To allow HTTP as an alternative
 - Together with ASGC/ShuTing, we added support for it in Panda (ATLAS analysis)
 - ShuTing also following with Athena for some additional fixes required
 - Added HTTPS support in ROOT
- But we still need to improve ROOT HTTP further
 - Add support for metalink, for replica failover and parallel reads
 - Add multi-stream access (read and write)



- Number one plugin request for DMLite
- Idea
 - Expose a locally mounted clustered filesystem via all the grid protocols, with grid auth/authz, etc
- Implementation
 - DMLite plugin implementing Catalog and I/O
 - Redirection to *gateways* for the I/O
- Status
 - Catalog implementation is there
 - It could be picked up by anyone (even with no knowledge of DMLite) and finalised
- Volunteers?

- [Check here](#)
- Lots and growing... you can help! Really!!
 - Testing of new pool types (HDFS, S3, ...)
 - Improved log analyzer
 - VFS plugin
 - Pool managing in the dmlite shell
 - Better drain process with dmlite tools
 - Dark data detection tool (integration)
 - Evaluating mysqltuner.pl
 - ...



- Expect more frequent releases
 - But you don't have to update each time
- Refactoring effort is over
 - A lot is in production, the rest just around the corner
- Next couple of months focus will be in the long standing fixes for I/O access issues
- Lots of parallel work on tools and plugins
 - You don't need to know many details to contribute
 - If you have an idea/request, have a look at the APIs