



DPM Collaboration & Plans

Oliver Keeble
IT/SDC
On behalf of the DPM Collaboration

09/07/2013



EMI->DPM Collaboration

- The DPM Collaboration has been formed to support the ongoing development and maintenance of DPM.
- Started 2nd May.
- Collaboration Meeting on 23rd April, with representatives from :
 - CERN
 - Czech Rep.
 - France
 - Italy
 - Japan
 - Taiwan
 - UK

Decisions

- Approval of a Collaboration Agreement
- Pledges of effort (over 4FTE)
 - In return for influence on direction
- Approval of distribution of responsibility
 - Development, testing, support
 - inc all derived and related services
 - LFC and dynamic federations
 - HTTP/DAV, NFS, xrootd, gridftp frontends
 - S3, HDFS, ClusterFS backends
 - Memcached for the nameserver
 - Some new developments (eg cluster FS support)
- Best effort support will be provided to non-member communities

Why is this important?

- Protects everyone's existing investments
- Recognition of the importance of the technology
 - Existing installations
 - Future potential
- Middleware is growing up
 - Exploit the open source collaborative model
 - Real delegation of responsibility
 - ASGC now own the NFS frontend
 - CESNET proposed a project and joined the collaboration

Partner responsibilities

- To give a flavour of tasks (incomplete list)
 - ASGC: NFSv4.1, ClusterFS support, testing
 - Czech: ClusterFS support, GSSAPI/X509
 - France: Support, testing, admin interface
 - Italy: Testing
 - Japan: Testing
 - UK: Admin tools, support, performance evaluations

DPM Release strategy

- DPM will release to EPEL
 - Metapackages, oracle stuff -> EMI
 - We hope sites will not have to change repo settings
 - Working with EMI to ensure this
 - First release to EPEL-testing, collaboration partners will then check
 - Talking to EGI about making this the “staged rollout”
- We can do component based releases from now on, ie more frequent, smaller updates
 - Raises questions about versioning

1.8.7

- Big improvements to the puppet configuration
- DB Connection pooling in the DPNS/LFC daemon
 - Improves performance 3X-5X and reduces hiccups
- The EGI/OGF Storage Accounting Record producer (StAR)
- Improved replica handling methods, lots of bug fixes
- Polishing the new Xrootd 3.3 interface, brand new using native DMLite (good feedback on the prev one, from ALICE and ATLAS FAX)
- Dynamic HTTP Federations (implemented as a dmlite plugin called “UGR”)

12 Month Roadmap

- Roadmap entirely determined by user requirements
 - Let us know...
- HTTP ecosystem
 - Ensure that HTTP/DAV is a high-performance, fully featured interface that HEP can use
 - Part of a coherent drive which touches many other components
 - LFC, FTS3, gfal2, xrd-http, federations, davix, root
- Management of legacy components
 - Review RFIO and legacy daemons
 - Client library access via new interfaces

12 Month Roadmap

- New interfaces
 - Support for the GridFTP redirections
 - Finalisation of NFSv4.1
- Standard stuff
 - dpm-drain, rebalancing, i/o limits
- Backends
 - Cluster file systems (“VFS plugin”)
 - Evaluate new deployment scenarios (S3, HDFS...)
- Configuration
 - Support puppet
 - Phase out yaim
- SRM
 - Support all necessary SRM-like actions through alternative interfaces

Summary

- DPM maintenance and development will continue post-EMI through a collaboration of stakeholders