
CERN

HTTP Deployment TF
GDB Report

Oliver Keeble

Why?

- All the pieces are now in place for HEP exploitation of the HTTP protocol.
- Atlas and LHCb (at least) are developing their HTTP infrastructures.
- The majority of storage systems are able to offer an HTTP interface.
- ROOT support is available via davix, and studies have shown that this can allow HTTP to reach a performance comparable with that of the xroot protocol.
- Projects are exploring how cloud data resources could be exploited, and in many scenarios this involves HTTP (S3), so the time is right to shake out the last problems.
- At the moment there are many site misconfigurations and instabilities which could be tracked and fixed at the WLCG level for all experiments

Mandate

- Define the minimum set of useful functions which must work for a site to "support HTTP for WLCG". Document what further functionality is desirable over HTTP, under what circumstances, and with what priority.
- Put in place the necessary monitoring and validation tools.
- Summarise advice for sites on deployment. Create baseline version information for the storage systems. Identify necessary changes to supporting services.
- Track and support the deployment on the infrastructure until the point where maintenance is transferred to standard experiment operations.

Status

- Twiki & mailinglist set up
  - https://twiki.cern.ch/twiki/bin/view/LCG/HTTPDeployment
- Membership
  - Experiments - Atlas, CMS & LHCb
  - Sites/Infrastructures - ASGC, BNL, CERN, GridPP, KIT, PIC & TRIUMF
  - Storage systems - dCache, DPM, EOS, StorRM, xrootd
  - WLCG monitoring is represented
- 3 meetings have occurred
  - 1st – setup, scope
  - 2nd – functional description of necessary
  - 3rd – monitoring & validation

Functionality

- Full text
  - https://twiki.cern.ch/twiki/bin/view/LCG/HTTPTFStorageRecommendations
- TOC
  - Basic Methods
  - Authentication and Authorisation
  - Space Reporting
  - Checksums
  - 3rd party copy
  - ACL Management
  - Traffic Monitoring
  - BDII Integration

Access monitoring

- Transfers and access using gridftp and xroot are currently monitored on the infrastructure
  - http://dashb-wlcg-transfers.cern.ch/ui/#
- HTTP will have to be represented here too.
- The TF agreed on two valid approaches:
  - A properly formatted UDP stream, as implemented in the xrootd f-stream.
  - Access summaries conforming to a documented json schema delivered over the messaging system
    - http://wdtmon.web.cern.ch/wdtmon/dash/http.html (under review)
- DPM (and xrootd of course) support the xrootd f-stream solution. dCache and StoRM will consider their options.

Monitoring/Validation

- A shared probe will be provided for integration with the experiments' SAM instances
  - A first version is nearly complete
- Tests will reflect the functionality identified as "required" by the TF (more detail in twiki)
  - Plus a check that access monitoring is working
  - Plus auxiliary checks to aid operational debugging (tbc)

Next Steps

- Infrastructure monitoring platform is evolving
  - Updated SAM/Nagios scheduler
  - SAM meeting next week with experiments to discuss integration with topology dbs
- Preprod prototype to be set up (within a few weeks)
  - HTTP TF will join this
  - Complete the shared probe
- Proceed to visualisation via SAM3
- Establish an operational plan