

Data Management I and II: recommendations

Simone Campana
CERN IT-ES

- LST and Demonstrators
 - LST demonstrated being appropriate framework for evaluation of storage technologies and data management solutions
 - DirectIO test should be included (besides functionality)
 - Analysis over the WAN should also be included
- EOS
 - Deployment of EOS in production at CERN should proceed. EOS should be provided as a production service
 - WLCG should evaluate if EOS can be exported outside CERN
- BNL demonstrator based on BlueArc
 - Will provide file:// access in the LAN (pNFS)
 - Will be deployed as DDM endpoint (SRM+gridFTP) + xroot(?)
 - endpoint will be provided by mid May and will become production if feedback is positive
- General
 - ATLAS should initiate a discussion (other experiments, main computing centers) to plan for the evolution of WLCG data management (but this was supposed to be Amsterdam, no?)
 - Focus on protocols (LAN and WAN; gridFTP+SRM, xroot, https)
 - This will have infrastructure consequences like the evolution of FTS, the end-of-life of SRM, etc..

- Breaking cloud boundaries
 - We should proceed with Fernando's proposal and start commissioning links through the sonar
 - Define few “bigT2s” and commission links with T1s
 - Leave to cloud support links within the cloud (to start with) and to other T1s. Follow the “channel priority” in the monitoring
 - Push SSB monitoring consolidation ASAP
 - Reuse CMS experience in links commissioning
 - Ask OPN to clarify procedure to debug channels
 - Ask OPN clarification about routing through T1s.
- LFC consolidation
 - Plan is understood both at the IT and ATLAS level

- PD2P
 - We accept Graeme's proposal and start replication based on thresholds (technical details up to Tadashi).
 - In parallel we also try the approach of replicating based on popularity and number of replicas (technicalities up to Cedric)
- Space organization
 - MCDISK and DATADISK consolidation largely agreed.
 - GROUPODISK should be also consolidated (need development in DDM, approx 6 weeks)
 - We offer group space to all groups at all T1s
 - » Group defines how many copies at T1s
 - » Decision of which T1 is up to DDM
 - PD2P should include group data
 - We remove group space at T2s (controversial)
 - » Alternatively, we offer T2 group space at “bigT2s” (same as T2s)