



ALICE comments on TEG outcomes

WLCG workshop

May 19, 2012

v0.3

Operations and Tools – WG1

- ALICE monitoring versatile & homogeneous
 - Workload & data management, services
 - Common framework – MonALISA
 - History data since mid 2006, predating Dashboard
 - Applications instrumented with common library
 - VOboxes collect site specific monitoring info
 - Info exposed to clients via web and subscriptions
- SAM complementary
 - Common WLCG view
 - Should be able to import MonALISA results!

Operations and Tools – WG3 - 1

- Build using SVN, CMake and native tools
 - Nightly and on demand
 - CDash and CTest coming
- Dependencies managed by AliRoot
- Pre-build distributions for SLC5, Ubuntu, MacOSX
 - Tested and working for other Linux flavors, including SL(C)6

Operations and Tools – WG3 - 2

- SW distribution based on BitTorrent
 - Shared areas being decommissioned
- Pilot jobs fetch what is needed and serve other pilots at the same site
- Architecture has no single point of failure
- Sites need to deploy nothing on the WN
 - Just allow for well-constrained Torrent traffic
- Works for non-WLCG, opportunistic / cloud resources
 - Where WLCG-specific means may not be available

Data Management & Databases

- See “shopping list”
 - https://twiki.cern.ch/twiki/bin/view/LCG/WLCG_TEGDataManagement_ALICE
 - Single AliEn catalogue model – all files, all replicas, all metadata – since 2003 !
 - As of today – 900M entries (LFNs/PFNs)
 - Data placement – SE autodiscovery
- Databases
 - Offline Conditions DB framework
 - AliEn file catalogue + ROOT files

Storage Management

- XRootD for everything – since 2007 !
 - Data placement, storage autodiscovery
 - Federation
- Tape SEs for raw data and old productions
 - Privileged access
 - Current year's Pb data should be fully online
- Disk SEs used by everyone
 - Automatic, user specified replication

Workload Management

- VObbox submits pilots to its own site
- CREAM functionality quite satisfactory
 - Pioneered by ALICE!
- Whole nodes, multi-core jobs
 - Good for job efficiency, development foreseen
- Virtualization – “don’t ask, don’t tell”
- Clouds
 - Not urgent, waiting for business model in HEP
 - Grid usage already cloud-like
 - T1-T2 distinction blurred for storage & brokering

Security

- AliEn security model uses envelopes
 - Describe which user can access (read/write) which files on which SEs
 - No need nor desire for general-purpose user proxies on the WN
- gLExec could be called with a specially crafted proxy only usable for that purpose
 - Will be investigated this year
 - Another idea to be pioneered
- AliEn LDAP server handles DN changes