

Experiment Summary

LHCb-I

- CPU time Normalization
 - LHCb needs that sites publish correctly the parameters for computing normalized CPU time and scale reference, both for match making to submit jobs and for accounting (CPUScalingReferenceSI00, GlueHostBenchmarkSI00).
 - In order to move to input files of 3GB might need longer queues at T1 sites, by a factor 2 (Currently 18000 HS06 minutes is the requested CPU time for the longest queue and input files size is 2GB)
- Shared area
 - Access to shared area: scalability and reliability issues observed at GRIDKA, IN2P3 and more Tier2 sites

LHCb -II

■ Data Management

- Data upload from UK sites to CERN
- Data access at IN2P3: newly transferred data to IN2P3 T1D0 storage are reported as UNAVAILABLE by the SRM
- Previous problems of data access at CERN with rfiio solved by using xrootd instead. Now is fine.
- Scalability issue for access to data at some sites: disk servers should be deployed accordingly to the data volume to access. Otherwise, considered to throttle job submission per site in Dirac.
- StoRM at CNAF: some files reported as non existing
- FTS replication to RAL from SARA (GGUS 59397)

LHCb-my perspective

- Worried regarding CPU accounting
- Increase in Queues needed ???
- Coping with higher pile up.

CMS session (I)

- Resource utilization and performance monitoring
 - Important as resource balancing is manual
 - Compare Tier-1 sites usage w.r.t. pledges
 - Good job success rate observed (90%)
 - New CMSSW versions will improve CPU efficiency
 - Weekly reviews of the plots
- Data production rates in Tier-1 workflows
 - Most demanding: redigi/rereco = 2 MB/(s·job)
 - E.g. at FNAL 4 TB/hour
 - Long migration queues may become an issue

CMS session (II)

- Modeling transfers and network requirements
 - to adapt computing model to changes in conditions
 - to track resource deployment at sites
 - Larger AOD sizes imply larger T1-T1 traffic
 - T1-T2 traffic more difficult to predict (depends on Physics groups)
 - Work in progress, but most sites have ample resources
- Custodial data storage
 - ~1/1 ratio between custodial/non custodial data
 - FNAL: ~ 5 PB on date, other T1's: ~ 1 PB
 - Data distribution and quantity more or less follows the pledges

CMS session (III)

- Tier-2 utilization
 - CMS uses full mesh of T1-T2 links, which is challenging
 - T2-T2 transfers becoming significant
 - In average 300 MB/s
 - >2 PB of data at T2's managed by physics groups
 - Main challenges
 - Need to often refresh data at T2 due to frequent reprocessing
 - End user activity rising independently of available data
 - 500 active users are the norm now
- Communications between CMS and WLCG
 - Several internal meetings with strong site participation
 - Several tools to evaluate site performance (Site Readiness, etc.)
 - 24/7 computing shifts, various primary and secondary "CMS centers"
 - Savannah most used but GGUS used increasing via bridge and for TEAM and ALARM tickets
 - CMS site contacts have an essential (and well defined) role
 - Understand CMS needs, follow up problems with local experts

CMS-my perspective

- Data taking good
- Review of TDR with initial data has allowed better planning
- 2.5Gbps for large T2s?
- Tape access at The T1s still worrisome.
- 1000 slot T1 needs 5Gb/s network for write from WN

ATLAS-1

- SW installation:
 - ~20 releases installed, ~400GB: trying to reduce
 - CMT touches $O(10k)$ files, very heavy: optimization work in progress
 - Releases installed on CVMFS-CERNVMFS
- Cream CE: please install Cream-CEs but keep the LCG CE
- Squid (for Conditions Data access): Proposition to have it as a WLCG service

ATLAS-II

- Storage resources:
 - Effort to optimize storage resources made available by sites
 - Tools for monitoring data storage utilization
 - Tools for keeping most popular data on sites
 - To ATLAS Grid sites: please deploy the pledged storage space!
 - ?Who are the contact persons if pledged disk is not provided in a federation?
 - ?Who should enforce that the disk space is deployed?

ATLAS-III

- Distributed Data Analysis Functional and Stress Testing:
 - HammerCloud responsible for Data Analysis functional and stress tests
 - New template model for HC tests:
 - more user friendly
 - not much room for error!
 - HammerCloud results available in different places and formats (HC web, Email reports, SAM, PandaMon, SiteStatusBoard coming soon)

ATLAS-my perspective

- “Where is all our space??”
- “we can clean up space now”
- “We have a new model for data placement”
- HC can help sites improve their performance
- Squid/CVMFS to help with SW/Cond’ Data
- We are trying to sort our SW area

Alice-I

- Data Taking Achievements 2010:
 - All sites migrated to SLC5 and CREAM (all the sites have at least one CREAM-CE for Alice). Except at CERN, where there is a dual submission system due to the fact that there are 20 LCG machines against 3 CREAM ones.
 - Put file a jobs quotas in the AliEnv2.18 to avoid user's abuse
 - Taking the closest SEs once the file has been registered
 - SAM to Nagios migration for the VOBox service already done
- Analysis:
 - Trains Analysis: Reduce load on storage servers grouping many analysis tasks in a common data set
 - Chaotic Analysis: Problems with memory consumption (in production only at CERN)
 - Facilities Analysis: pre-stage file available in the working nodes (in production only at CERN)

Alice-II

- No remarkable Issues about CREAM 1.6 from the point of view of the experiment, although small instabilities found at the level of site operations
- Extensive discussion about the use of xrootd as T0-T1 transfer protocol. Experiment decision discussion out of the scope of this forum (no experiment management could be present at this meeting)

Alice-III

- T2/T3 feedback on setting-up and operating
 - The importance of internal network setup for analysis activities and maybe also other activities
 - Guidelines to balance worker nodes and storage and choose an appropriate network setup
 - The use of local monitoring tools in addition to MonaLisa and also in order to confront the figures
 - How to look at the site performance in a efficient/correct way (what to look)
 - Priorization of site messages
 - Sites using link aggregation (1Gb not enough and 10Gb not necessary)
 - Moving from DPM-xrootd to xrootd only motivation: the evolution of DPM with xrootd was stopped. But now, much effort is done in the set up of xrootd in DPM

Alice-IV

- Alice at Prague T2
 - In general, Alice approach is to lighten the site requirements in terms of services and aims at their high performance
 - Good experience with the creamCE 1.6/gLite 3.2 @ SL5/64bit and the VOBox on gLite3.2 and SL5.4/64bit.
 - Usage of virtual machines for the services.
 - Using of more than 1 oboe and CREAM-CE at the site. We would also be willing to do it in Prague
 - Very good experience with and an extensive usage of the MonALISA.
 - Very good experience with the support from the CERN team

Alice-V

- ALICE Tier1 test bed at KISTI-NSDC
 - KISTI presented the evolution of the site in terms of services and support to provide Alice with a T1 service infrastructure

ALICE-my perspective

- XROOTD good (why SRM?)
- New T1 in Korea possibly
- No problem to install new middleware stack at the start of data taking
- T2 site setup
 - Similar issue to other Sites

Commonalities

- T1s Good.
- Worried about space/CPU/resource
- T2s
 - Transfers
 - Space
- Monitoring Tool Aggregation
- Communication
- Job completion
 - Async stage out (Job Recovery?/pCache?)

Summary of VOs

- ...
- .