Summary

GDB

Chaired by John Gordon, STFC-RAL

GDB meeting @CERN December 14th 2011

http://indico.cern.ch/conferenceDisplay.py?confId=106651

Forthcoming Events

GDB

- **WLCG TEG co-located and joint FtF meetings, 23, 24-25 january 2012, NIKHEF, Amsterdam**
  - **TEG-OPS : https://indico.cern.ch/conferenceDisplay.py?confId=161833**
  - **TEG-STORAGE&DATA : https://indico.cern.ch/conferenceDisplay.py?confId=165687**
- **LHCOPN/LHCONE, LBL, 30-31 January 2012**
- **ISGC, Taipei, 26 February - 2 March 2012**
- **EGI User Forum, Munich, 26-30 March**
- **Also EMI Technical Forum**
- **HEPiX Spring Meeting, Prague, 23-27 April 2012**
  - **https://indico.cern.ch/conferenceDisplay.py?confId=160737**
- **WLCG Workshop, NYC, 19-20th May 2012**
  - **http://indico.cern.ch/conferenceDisplay.py?confId=146547**
- **CHEP2012, NYC, 21-25th May 2012**
- **EGI Technical Forum, Prague, September 2012**
- **HEPiX Fall Meeting, IHEP, October 2012**

LCG-CE / CREAM

GDB

- Sites no longer required to run an LCG-CE.
- Proposed removing from Availability calculation in October
  - Provided ACE worked OK (it did)
  - But reliable CREAM for SGE not yet rolled out
- MB discussed again in November and decided to wait longer for the anticipated release of CREAM for SGE.
  - Released in November EMI1 update 11
  - End of year to change calculation
  - **Is it working?**

AOB

- **Disk pledges :** Flooding in Thailand - Any feedback from sites on disk availability?
- **Whole node scheduling : needs for precision related to use cases …to be addressed by Workload Management TEG**
  - Some sites have provided configuration but use cases need to be precised
  - I.Fisk : CMS use : schedule multiple copies of the same application
  - J.Gordon sees an other use case : a more efficient use of hardware by an appropriate scheduling combining I/O tasks with memory consumming task etc…connected to pb of memory consumption
  - Whole node scheduling and WN virtualisation ?
  - Virtualisation of WN more a site issue - Whole node scheduling is a VO isssue
- **Vidyo service** (in replacement of EVO) : Ready to start in January ?
- **GDB Chair** : to be renewed (2 years mandate)

Mware - EMI 1 status
Cristina Aiftimiei (INFN)

GDB

- Update 10 – 24.11.2011 :
  - Major release CREAM/GE module
- Update 11 – 15.12.2011 :

http://bit.ly/EMI1UpdatesReady

- Early availability :

https://twiki.cern.ch/twiki/bin/view/EMI/EMIBetaAcceptanceTesting

- EMI products not yet in UMD
  - WMS : Some tesing outside stage-rollout process
  - M.Litmaath : Close interactions between dev. , rollout and UMD...to get WMS ready - some bugs were founded
- gLite security updates (APEL and TORQUE) released as soon as they are ready

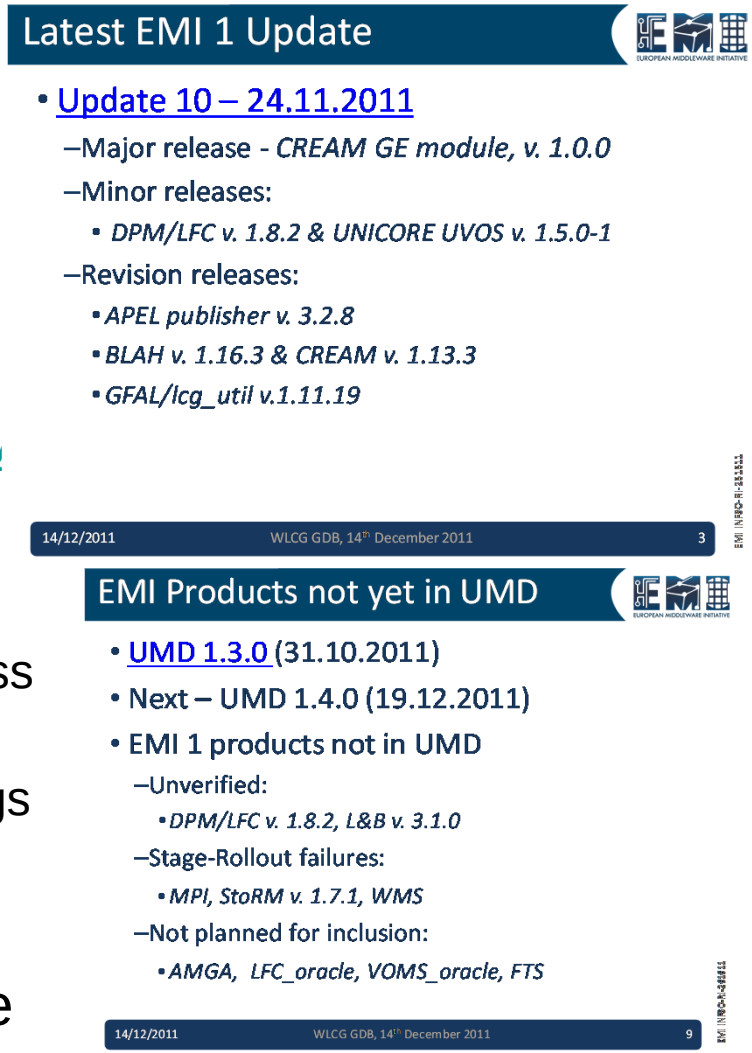

GDB

Experiment operations for Q4 2011

Short summary

were asked to address

- SLC6 migration
- Memory requirements

ALICE
Latchezar Betev (CERN)

GDB

- ALICE is happy – Thanks to sites
  - 2011 HI Run: high luminosity and stable beams => 1.3 PB raw data
  - 70 % of RAW replicated during data taking (internally limited to 800 MB/s)
  - Quasi online data processing, QA, calibration
  - Processing of HI data is ongoing (Pass2 about to start), MC and user activities at stable level
- ALICE is running out of disk space => cleanup of replicas (of older productions) at all centres
- Memory overruns because of higher luminosity - More pileup than expected and simulated
- Efficiency :
  - Condition Data system overload + high number of end-user analysis jobs with varying efficiency
  - Remedies : strengthen cond. Data infrastructure , include user analysis in groups (analysis train)
  - Effect CPU/Wall efficiency has risen an acceptable level of 70 % - 80 %

ALICE
Latchezar Betev (CERN)

GDB

- Efficiency :
  - Condition Data system overload + high number of end-user analysis jobs with varying efficiency
  - Remedies : strengthen cond. Data infrastructure , include user analysis in groups (analysis train)
  - Effect : CPU/Wall efficiency has risen an acceptable level of 70 % - 80 %
- SLC6 : migration as soon as possible - already seen centres using SL6 (disk servers)
  - Also OK with a more aggressive move to latest Ubuntu (immediate gain x 2 on memory - due to latest kernel version or gcc …???)
- Memory requirements
  - ALICE realizes that asking for more memory cost more momey
  - Would prefer more storage
  - Alice do not ask for more memory

LHCb
Elsia Lanciotti(CERN)

GDB

- Activity during last quarter : Reprocessing campaign
- Next :
  - MC campaign has just started
  - Re-stripping before 2012 data taking in march
- Operations split by T0/T1 and T2s

T1 – T2 association

“Not all UK T2s have been associated to RAL”

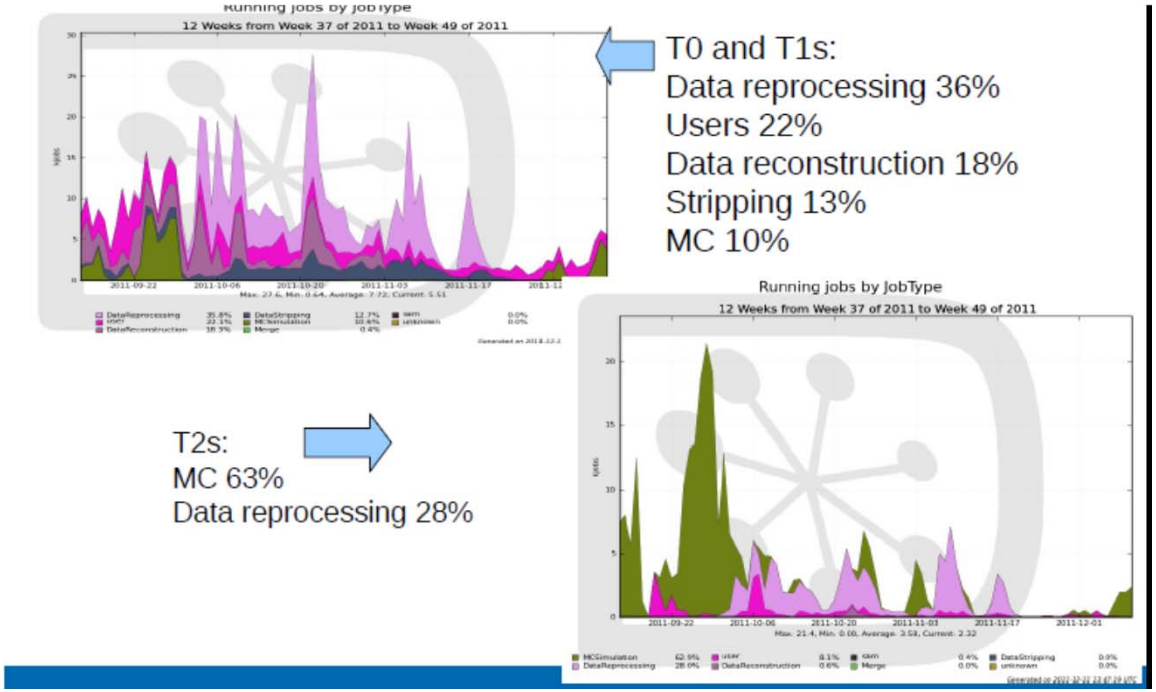

LHCb
Elsia Lanciotti(CERN)

GDB

- For Next re-stripping campaign for stripping jobs , the possibility of downloading input files will be tested
- Required local space from 10 GB => 20 GB is this acceptable for Tier-1s ???
- During pre-stagging, some problems due to disk cache size
  - Better monitoring would be appreciated …
- Huge MC campaing has started at the beginning of december
  - Manual Data cleaning before the MC campaign
  - dCache sites have to be asked via GGUS tickets to re-allocate the free space to the new space tokens
  - Data migration from the old space tokens schema to the 2011 schema ???

LHCb
Elsia Lanciotti(CERN)

GDB

- LHCb not ready for the migration to SLC6

  … python 2.7 gcc 4.6.2 ????? (the current forecast from the architect forum to be confirmed)
- DISCUSSION Memory requirements : 2 GB /Core
- Site Procurement : is it required to put more physical memory ? NO
- CERN physical mem. 3 GB/ core for procurement - no limitation in vmem usage
- LHCb : 4 GB of virtual memory ….usage of virtual memory may affect efficiency
  - Whole node could help to optimize the memory usage and comsumption
  - Too many different approaches from sites , could be helpful to clarify the configuration of site batch scheduling

ATLAS
Cédric Serfon (Univ. Munich)

GDB

- Last 3 months Activity
  - MC production : 70 %
  - Central group production : ~8 %
  - CPU comsumption in T2s : 60 %
  - Top 3 clouds : US, UK, DE (FR 4th position) (last 3 months)
- FOCUS on network

2011 major change

« Breaking the clouds Boundaries »

Network congestion

- Since 6 months, we have started to break the clouds boundaries by setting up T2Ds (i.e. T2s qualified to get their data directy from T1s). Right now more than 50% of all ATLAS T2s are T2Ds.
- Start to see some problems :
  - 1Gb/s bandwidth of TRIUMF research network fully saturated. Will be increased.
  - Problem with transfers from US T2s to KIT (firewall on the general 10Gb/s IP uplink overloaded).
  - Degraded connectivity between IN2P3 and BEIJING/TOKYO and other foreign T2s under investigation.
- Other problems observed : degradation of transfers CERN→AGLT2 due to change of the path from CERN to AGLT, or PIC→T2 due to switch to LHCOne.

Discussion

- Problems observed : degradation of PIC => T2

  "*due to switch to LHCONE"*
- Gonzalo Merino : PIC - 1 Gbps link to LHCONE …. In testing mode

- Ian Bird : LHCONE is not in production. It is a prototype infrastructure.
- Scaling issues with the current infrastructure. So The migration should be handled with care.
- ATLAS want to be informed . Is Roll- back possible in case of performances degradation ?
- Not a VO specific activity to monitor network performances….
- Different understanding of LHCONE….?????

=> Info à Xavier Jeannin

En France, Renater offre un service LHCONE de production avec une connectivité au moins égale à celle sur l'IP générique

ATLAS
Cédric Serfon (Univ. Munich)

GDB

- LFC Consolidation
  - 4 LFCs migrated to CERN : SARA, Taiwan, GridKa, CNAF
  - If no scaling problem, the merges will continue (2/month) after Oracle 11g migration. Still 14 LFCs to merge…..

CC-IN2P3 not in the list…
Evaluation ????

Limited link commissionning effort on a subset of pair - ATLAS do not want to start to exercise full-mesh

perfSONAR-PS deployment – not a VO tool
==> ATLAS wiki page …..
Cross VO Page

LHCOne

- The ATLAS approach is to evaluate LHCONE in a managed way and gain expertise in liaison with Network providers.
- Agreement on a list of sites that we want to use for initial evaluation of LHCOne :
  - T1s : ASGC, BNL, CERN, PIC, SARA, TRIUMF.
  - T2s : AGLT2, (DESY-HH), LAL, (LRZ-LMU), Napoli, Prague, Tokyo, Toronto
- Those sites must have perfSONAR-PS deployed + Need to setup dedicated FTS channels.
- Measurements will be done before and after establishing connectivity to LHCOne (some sites are already connected to LHCOne).

ATLAS
Simone Campana (CERN)

GDB

- ATLAS mem. Requirements
  - 32 bits under 2GB
  - For native 64 bits, memory consumption is not under control
- ATLAS not prepared for SL6

CMS OPS report
Ian Fisk (FNAL)

- Increasing pile-up makes the computing problem more challenging
- Very low efficiency observed at Tier-0 - using 5 cores out of 8 and consumming the overall mem. Memory leaks...
- Whole node scheduling at Tier-0 / Tier-1 may help
- HI Run - memory requirements for HI are currently lower than high pile-up pp due to differences in tracking
- Zero suppression has reduced Tier-0 I/O
- CMS closed to 100% of Tier-1 pledge over all 2011
  - Mostly reprocessing and MC simulation

CMS OPS report
Ian Fisk (FNAL)

GDB

- CMS closed to 100% of Tier-1 pledge over all 2011
  - Mostly reprocessing and MC simulation

??? T1-FR-IN2P3 ???
CMS Activity not fully reported to CMS dashboard pb related to GE migration ????(I.Fisk)
WLCG/EGI accounting reports should be OK.

Tier-1s

- Reasonable balance across Tier-1s
- Recent reprocessing passes everyone has contributed
- Some issues getting all the pledge at ASGC during this time