

Table of Contents

Summary of GDB meeting (May 9, 2012)	1
Agenda.....	1
Welcome.....	1
TEGs: status and what next.....	1
Experiment Resources in the Coming Years.....	1
LHCOPN/ONE Status and Future Directions.....	2
Federated Identity Management.....	3
KISTI, a new T1 for ALICE.....	3
HEPiX Prague Summary.....	3
WLCG Workshop.....	4
HEPiX WG Report on trusted virtual images.....	4
: CNAF experience with virtualized WNs.....	5
Cloud Resources in EGI.....	5
ATLAS viewpoint.....	6
LHCb viewpoint.....	6
Discussion.....	7

Summary of GDB meeting (May 9, 2012)

Agenda

- <https://indico.cern.ch/conferenceDisplay.py?confId=155068>

Welcome

- Meeting organization
 - ◆ every meeting should have one or more minute takers
 - ◆ Meeting summaries: a TWiki area should be set up (WLCGGDBDocs)
 - ◆ someone in the room should monitor the chat area
- Frederique: autumn remote meeting might be held in Annecy, will check
- Main topics for next GDBs
 - ◆ June GDB: TEGs, post-EMI plans, RFC/SHA-2, gLExec, ...
 - ◆ July GDB: experiment reports, ... May change depending on final June agenda
 - ◆ Sep GDB: IPv6, ...
- Actions to follow from previous meeting
 - ◆ perfSONAR tracking

TEGs: status and what next

- big TEG meetings finished, just focused meetings on specific topics
- WLCG workshop preceding CHEP:
 - ◆ priorities
 - ◆ open questions
 - ◆ unfinished areas
 - ◆ Initial proposal of working groups that WLCG should set up
 - ◆ Proposal for those general topics that could be dealt with in HEPiX
- June pre-GDB:
 - ◆ DPM future
 - ◆ Conflict with LHCC week: some people not available
- Claudio: glxexec deployment desired at all CMS sites, with WLCG support
 - ◆ Ian B: could be a task for the requested WLCG operations support team
- Jeremy: status of relocatable glxexec?
 - ◆ Maarten: recipe provided, proven to work, but requires compilation per site; may make it easier in the future
- Davide: WM TEG has many recommendations to be followed up

Experiment Resources in the Coming Years

General remarks from C-RSG

- priorities should be set according to resource availability
- pileup is an issue for 2012 data
- ALICE: low CPU efficiency for chaotic analysis
- keep requests reasonable for funding agencies

C-RSG main conclusions

- Increase use of HLT farm in 2013 for reprocessing and simulation
- Look at the possibility of using external sources for some particular activities, like MC
- C-RSG would like to keep a balanced usage of different Tiers. Ensuring such a balance will maintain a healthy collaboration

Other actions

- Collection of installed capacity data too complex: will do it manually using REBUS (some development needed first)
- Need to progress with storage accounting

Experiments disappointed by the lack of interaction with C-RSG this year: probably the year with the least interaction

- Used to be considered useful in the past
- C-RSG made some recommendations not exactly appropriate or difficult to interpret
 - ◆ try separating organized and chaotic jobs for accounting? would require non-trivial developments in various places

Discussion

- Hans (Atlas):
 - ◆ grateful to sites
 - ◆ pileup much higher, different energy, not like 2011
 - ◆ 2013 extra requirements not a luxury, well justified
- Helge: some criticism on the C-RSG can be defended, but various input data arrived rather late

LHCOPN/ONE Status and Future Directions

LHCOPN

- LHCOPN operations: it just works!
- new T1s may come
- dashboards, perfSONAR deployment improving, also for LHCONE

LHCONE

- LHCONE operational and progressing
- L3 VPN symmetric routing requirement
- various project updates, e.g. GLORIAD, GEANT, NORDUnet, DANTE
- L2 point-to-point service investigations

Discussion

- Michel: how to enforce perfSONAR everywhere on LHCONE? is it scalable?
 - ◆ John: most instances dormant for troubleshooting, some actively used for monitoring
- Michel: how useful if mostly dormant? perfSONAR is used to establish baseline for further troubleshooting...
 - ◆ John: guidelines on TWiki, core testing sites decided per VRF (routing domain)
- Michel: who is responsible for taking action on issues?
 - ◆ John: sites should set up alarms for themselves
- Michel: site will contact NREN --> other NREN etc; same discussion as for OPN...

Federated Identity Management

- remove identity management from services, allowing SSO
 - trust needed between service providers and identity providers, like with IGTF
 - communities also have attribute authorities, e.g. VOMS
 - examples: IGTF, educational (national, international), social networks (e.g. Google ID)
 - collaborative effort involves:
 - ◆ photon & neutron facilities
 - ◆ social science & humanities
 - ◆ high energy physics
 - ◆ climate science
 - ◆ life sciences
 - ◆ fusion energy
 - current summary document: <https://cdsweb.cern.ch/record/1442597>
 - common requirements, some non-trivial
 - common vision statement
 - recommendations to research communities:
 - ◆ risk analysis
 - ◆ pilot project
 - recommendations to technology providers:
 - ◆ separate AuthN from AuthZ
 - ◆ revocation
 - ◆ attribute delegation to the communities
 - ◆ levels of assurance
 - recommendations to funding bodies:
 - ◆ funding model
 - ◆ governance structure
 - not only grid, also other collaborative tools (TWiki, Indico, mailing lists, ...)
 - pilot study foreseen at CERN, e.g. TWiki
 - how to involve IGTF?
 - MB endorsement needed (Ian B: next meeting)
-
- Matteo: relation to cloud computing? input welcome for EGI work in that area
 - Dave: cloud implications have to be considered as well

KISTI, a new T1 for ALICE

- T1 ascension procedure now officially documented
 - ramp-up milestones: candidate --> associate --> full T1
 - KISTI plan being prepared
 - Russian project: progress expected later this year
-
- Michel: might extra resources in one place allow for a reduction elsewhere?
 - ◆ Ian B: normally not; ALICE in particular still have much less than their nominal requirements

HEPiX Prague Summary

- very full program
- new business continuity track
- others: IT infrastructure, storage, grid/cloud virtualization, network & security, ...
- energy efficiency: future meeting
- fabric management changes
 - ◆ Puppet

- ◆ Quattor
- ◆ Nagios --> Icinga
- batch:
 - ◆ PBS/Torque scalability issues
 - ◆ SLURM, Condor rising
 - ◆ xGE forum
- clouds on the horizon of realism; OpenNebula, OpenStack
- storage:
 - ◆ federation
 - ◆ what comes after RAID
- Federated Identity Management
- IPv6
- Working Groups: virtualization, IPv6, storage, benchmarking
- HEPiX very healthy!

WLCG Workshop

- CHEP/LHC schedule mismatch for future workshops?
 - ◆ workshops can be standalone (again)
- New York: TEG recommendations + exciting new developments
 - ◆ reserve a slot for glxexec deployment timeline and hints
- loose agenda to allow a lot of discussions
- Gantt chart for Run-2 preparation?

- Ian B: will send draft comments + questions to TEG chairs
 - ◆ Jamie: focus on explicit, time-bound recommendations
 - ◆ Ian B: do not repeat TEG discussions

HEPiX WG Report on trusted virtual images

Mandate almost fulfilled

- image endorsement: approved JSPG policy
- framework for publishing and distribution of images that can be integrated with any image repository implementation
 - ◆ integrated with StratusLab Marketplace
 - ◆ being integrated with OpenStack Glance
- Technical arrangements have been defined for site contextualization and for exchange of information between site infrastructure and a running VM
 - ◆ E.g. remaining lifetime...
- CERNVM images compliant and reviewed
- experiment-specific images could directly connect to pilot framework
 - ◆ Probably the most desirable option

- Dan: experiment should also be ascertained the image is what it expects, i.e. not replaced/updated by the site
 - ◆ This is by design!
- Matteo: Glance vs. StratusLab?
 - ◆ Ulrich: site choice

: CNAF experience with virtualized WNs

- WNoDeS in production since Nov 2009 at several Italian sites, incl. T1
- included in EMI-2
- mixed mode: use physical nodes as traditional batch workers and for VMs in parallel
 - ◆ some pros and cons
- upcoming features: interactive, OCCl, web interface, dynamic private VLANs, federated access, storage
- end of EMI timeline may have some impact

- Matteo: mixed mode - jobs on hypervisor might spy on network traffic
 - ◆ Davide: an exploit would still be needed; already an issue without VMs today
- Michel: usage by WLCG experiments?
 - ◆ Davide: high-memory VMs were deployed for ALICE; no use case yet for the others; some other VOs using special VMs, created by the site

Cloud Resources in EGI

- resource providers and communities interested in clouds for various reasons
- create community platform alongside grid infrastructure
 - ◆ interface to commercial providers
- WGs to address technical issues and engagement; testbed
- goals: blueprint, dissemination
- standards and validation
- resource typologies, heterogeneity, provider agnosticism
- task force consisting of 23 institutions from 13 countries
 - ◆ stakeholders
 - ◆ technologies
- federated testbed, living blueprint document
- demo was given at EGI CF 2012
- many consolidation activities next 6 months

- Philippe: relation with HEPiX?
 - ◆ Michel: different areas
 - ◇ EGI: federated clouds
 - ◇ HEPiX: trust infrastructure for virtual images
 - ◆ Matteo: complementary; e.g. interested in Marketplace/Glance integration
- Ulrich: EC2 support?
 - ◆ Matteo: yes, but most standard user interface is OCCl, not EC2
- Michel:
 - ◆ HEPiX WG was started because of experiment wish for controlled (virtual) environment
 - ◆ sites: how can cloud-like resources be used transparently?
- Jeff: Dutch communities have been asking for cloud resources, but not federated: why are federated resources needed?
 - ◆ Matteo:
 - ◇ we asked various communities and got different requirements
 - ◇ federated offer: user can handpick where to run jobs/VMs without knowing details of implementation; EGI can tailor
 - ◆ Michel: small communities may not ask for federated clouds, just resources with some implementation
- Matteo: relation between private and public clouds; some communities already using Amazon because it is easier, less expensive

ATLAS viewpoint

- various cloud integration and testing activities
- Jeff: why?
- Fernando:
 - ◆ some sites interested in cloud infrastructure
 - ◆ MC production in Amazon etc.
 - ◆ want to be ready for cloud resources
- contextualization strategies
 - ◆ golden image expensive to maintain
 - ◆ HEPiX CDROM approach?
 - ◆ Puppet? how much can the image be changed?
- image management issues: presently entirely manual, no signing of images

- Tony: why/how does the image need to be contextualized by the VO?
- Fernando:
 - ◆ install some packages + configuration files
 - ◇ not everything in CVMFS (e.g. Condor, Ganglia)
 - ◆ certificate handling
- Dan: you need to give the image a secret to pull in jobs from Condor, typical workflow is:
 - ◆ we boot image running for a long time as a Condor worker; shut down when no work
 - ◆ credential needs to be passed and renewed --> handled by Condor
 - ◆ use of Condor also convenient for joining PanDA infrastructure
 - ◆ Condor also solves whole-node/multi-core problem
 - ◆ Ulrich: pass it as user data; got it to work with Puppet
- Tony: site should have last word in contextualization for logging etc; put Condor in CVMFS!
- Michel:
 - ◆ Contextualization initially designed to be done by site exclusively on an image built/tailored by the VO
 - ◆ Now VOs want to rely more on CVMFS to avoid too many images: issue with sustainability of image management without an image catalog
 - ◆ VO contextualization is a new use case, needs more thought but can probably be integrated in the model
- Philippe: LHCb needs very little in CVMFS, e.g. mount point and script to set up environment
- Ulrich: image needs to be bootstrapped with VO parameters passed as user data, image itself should not need to be touched
- Philippe: indeed
- Ulrich: potential issue with long-lived images, e.g. for SW updates
- Tony: let the batch system shut them down as needed, see proposed mechanisms

LHCb viewpoint

- not yet at the level of ATLAS; interested in CVMFS, lxcloud, commercial clouds (DIRAC extension)
- consider cloud as yet another batch system? create overlays with pilots as usual
- fair share mechanism vs. adding/removing VMs on demand
- single-core VMs not OK
- multi-cores: run N jobs in parallel, mix of CPU and I/O bound, or parallel Gaudi job
- LHCb would like to couple virtualization and whole node scheduling
- account VOs on wall-clock time, not CPU time

Discussion

- Accounting
 - ◆ Jeff: mismatch with cloud world; we need limits on wall-clock-HEPSPEC-hours!
 - ◆ Tony: wall-clock time accounting makes sense now
- Jeff: ATLAS could take more when others are quiet and then run out of quota faster?
- Michel:
 - ◆ it is an important question but this is not the right time for theoretical discussions on how scheduling will work in the cloud
 - ◆ need to get some small-scale workflows going to find and fix the issues in the whole chain
- Dan: some effort available for projects with limited scope, e.g. using the HEPiX tools
- Helge:
 - ◆ looking into different ways to set up cloud services, which should not concern experiments
 - ◆ also trying to get endorsed images to work
 - ◆ single- vs. multi-core is orthogonal
- Dan: root access to VM image?
 - ◆ Ulrich: many/most sites cannot handle that, so all that is needed should be done beforehand or with trusted contextualization plugins
- Dan: VM shutdown announcement needed for job cleanup
 - ◆ Ulrich: being looked into, some mechanisms already available
 - ◆ Jeff: imitate fuel gauge!
- Michel: many ideas and questions to be further discussed in WGs etc.

-- MaartenLitmaath - 09-May-2012

This topic: LCG > GDBMeetingNotes20120509

Topic revision: r8 - 2012-05-14 - JamieShiers



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback