

Table of Contents

Summary of GDB meeting, June 14, 2014 (CERN)	1
Agenda	2
Introduction - M. Jouvin	3
Bitcoin Mining.....	3
DPHEP Update - J. Shiers	5
MW Readiness WG - M. Dimou	7
Report - H. Meinhard	9
Actions in Progress	10
Ops Coordination Report - A. Forti.....	10
HS14 Preparation - M. Alef.....	11
IPv6 pre-GDB Summary - D. Kelsey	12

Summary of GDB meeting, June 14, 2014 (CERN)

Agenda

<https://indico.cern.ch/event/272622/>

Introduction - M. Jouvin

Still looking for more volunteers to take notes: some of the contributors cannot anymore, need to renew the team

- 2-3 volunteers would be welcome
- Thanks to Stefan for taking the notes today

Next GDBs

- July will be replaced by WLCG workshop
- August no meeting
- No proposal yet for an outside CERN GDB meeting so far, contact Michel if interested to organise, **decision to be made before the summer**

WLCG workshop

- deadline extended to June 17

upcoming pre-GDBs topic proposals

- Clouds in September
- None in October if Michel participation is required...
- Volunteer computing in November?
- Topics may emerge from the WLCG workshop

Actions in Progress:

- Migration to GFAL2/FTS3 clients: plan to decommission old ones in October
- tree like middleware clients in CVMFS (grid.cern.ch repo) now available
 - ◆ LHCb will test it soon
- ARGUS support: being actively discussed, INFN solicited to take over from SWITCH, no decision yet

Bitcoin Mining

Attention raised by R. Wartel.

- Recent event with a large publicity at the NSF
 - ◆ From http://www.cio.com/article/753803/US_Researcher_Banned_for_Mining_Bitcoin_Using_University_S

"The U.S. National Science Foundation (NSF) has banned a researcher for using supercomputer resources to generate bitcoin. In the semiannual report to Congress by the NSF Office of Inspector General, the organization said it received reports of a researcher who was using NSF-funded supercomputers at two universities to mine bitcoin.

The computationally intensive mining took up about \$150,000 worth of NSF-supported computer use at the two universities to generate bitcoins worth about \$8,000 to \$10,000, according to the report. The universities told the NSF that the work was unauthorized, reporting that the researcher accessed the computers remotely, even using a mirror site in Europe, possibly to conceal his identity. "

WLCG, working with other, should raise user awareness on the consequences

- Coordinate with EGI : not a HEP specific issue

Should also help sites to detect and handle cases

- Write down some documentation helping sites to identify illegitimate/criminal use
- Stress importance of configuring sites for central banning

DPHEP Update - J. Shiers

Vision: by 2020 all HEP archived data shall be easily findable and fully usable by designated communities

- Need clear targets and metrics between funding agencies, sites offering services and experiments
- DPHEP rebranded "Data sharing in time and space"

DPHEP portal prototype based on Invenio

- Not easily seen from outside CERN...
- Really a prototype : information largely static
- Would like to receive comments, suggestions
- Main goal: find/ass a common look and feel for different data
- This summer a portal will be populated using sources from past CERN experiments and the CERN grey book (current and past CERN experiments).

A fellow will start January 2015 to help populating the archive

- Another task will be to describe analysis workflows for the data to ensure that the analysis code can be reused

Bit preservation

- For LHC data this is under control: see G. Cancio's talk at HEPiX
- For 2nd/3d copies of data (+ necessary metadata to know where useful data are): still to be discussed
 - ◆ Take into account that metadata by 2020 may be the same size as the data today

C-RSG & RRB news: recommendation to distinguish ability to read/analyse old data from requirements from open data

- Long term analysis of data should be part of the experiment computing costs
- Open access is clearly an additional cost
 - ◆ How much do we really need to provide?
 - ◆ Do what we can afford: only a few TB could be enough to start as a sliding window

Collaboration agreement

- A version acceptable to 9 parties is now ready, still waiting for signatures

H2020: 3 calls of particular relevance to long term data preservation are currently open

- 2 closing Sept. 2, 3d one on Jan. 14
- More calls in the next years
- They are complementary: complementary proposals expected
 - ◆ WLCG should use its long experience in generic and specific solutions to target these calls
 - ◆ CVMFS is an example of a technology developed by WLCG that solves issues face by all communities
 - ◆ EINFRA-1: a project being prepared including bit preservation and hooks for VRE-level services, under the umbrella of EU-T0
- **Looking for volunteers to help review the proposals**

Targets and metrics: Data Seal of Approval is a starting point, focused on infrastructure. Also need metrics for VRE-oriented services

- Open Access Data for Educational Outreach
- Analysis reproducibility
- Full scientific potential of data?

CHEP2015: only option is during the week-end before, not a very good option (competition with the WLCG workshop planned the same week-end), close in time if not in space with RDA workshop

- May not have a DPHEP workshop during next CHEP

Discussion

- Why only X parties have agreed on the collaboration?
 - ◆ Jamie: No clear obstacles for the other parties to sign as it was discussed last year, just details that took longer than expected (needed?)
- What are the other communities HEP is discussing these issues with?
 - ◆ Jamie: Many different ones, in particular astrophysics and astronomy, but also cultural heritage that has a problematic very close to our
- Also medical communities?
 - ◆ Jamie: Less contact to those.

MW Readiness WG - M. Dimou

Goal: assess readiness of certified/validated new MW versions to be deployed in production and integrated into experiment workflows

Preparation phase done

- List of products that will be taken into account established
 - ◆ Mainly storage products + VOMS
 - ◆ Additionally if needed CVMFS, Condor...
- Experts identified in each experiment
- List of volunteering sites: they need to configure specific resources for this effort that will appear separately in monitoring results
 - ◆ How they will appear in the monitoring still being discussed

The WLCG Package Reporter: identifying the MW versions running at sites

- Proposal [\[?\]](#) and implementation by L. Cons
- Sites report the RPMs (related to MW) they use
- Checked against a baseline for production or testing
- Prototype expected by beginning of September

HC used to run tests: tests for release candidates/testing are distinguished from production version

- Results kept during 1 year to allow comparison over time

MW Officer (A. Manzi, new role created) will in charge to maintain the baseline version, to communicate with sites not compliant with baseline and to liaise with other bodies like EGI/UMD

- Also has the role to identify the versions that need to enter readiness verification
- Ensure that the new versions of clients are in CVMFS grid.cern.ch to evaluation by experiments
- Check with volunteering sites which version is tested by which VO
- Results announced through the Ops Coord

First use case to assess the workflow: DPM 1.8.9

- Target date: end of June/beg. of July
- Use this first case to refine the process

Second use case planned: Condor

- Will allow to involve OSG

To get more information and participate

- Twiki page
- Next meeting: 2nd of July, 4:00 pm CEST
- egroup

Discussion

- Shall sites wait for the readiness report before deploying new versions.
 - ◆ Maria: it is not mandatory, some sites decide to deploy certain versions of packages early b/c of e.g. close collaboration with the development team. Nevertheless if a package version is

- marked as baseline version it shall be regarded safe to be deployed.
- ◆ Ian C: volunteer sites usually have the possibility to test a wide range of functionalities of a certain product, so not everybody shall test everything all the time. This is the reason to coordinate this effort.
 - ◆ Michel: the role of this working group is the coordination of all the testing and validation efforts to avoid duplication and ensure proper coordination with experiments
 - ◆ Christina: should be underlined that the baseline version is the minimum recommended version for sites to have deployed.
 - ◆ Michel: baseline in the sense of this WG should be seen as the **recommended version**.
 - ◆ Andrea M.: baseline is currently the minimum recommended version. In the future, we'll add a recommended version that describes the version which has successfully been tested by VOs and should be deployed at each site.
- Michel: during the winter, it was discussed to use Pakiti for detecting the package version used at site. What are the reasons that lead to the development of WLCG package reporter ?
 - ◆ Maria: Pakiti was described as one of the candidates during the discussion but in a discussion was decided that the intended WLCG package reporter will better suite needs (also discussed in the upcoming meeting).
 - ◆ Michel: interested to hear technical arguments why this was chosen. Need to consider the consequences in term of manpower of splitting the effort rather than joining it...
 - ◆ Vincent: Pakiti gets all rpms, the WLCG tool only needs a subset of packages. But Pakiti could be improved to deal with this...
 - ◆ Tim B: propose to try Pakiti and only if it doesn't serve our purpose we develop something ourselves.
 - ◆ Michel propose to follow-up by email with the relevant people, after checking Lionel's presentation as suggested by Markus, and to have this may be rediscussed at the July 2 meeting of the WG.
 - Simone: Do you have a timeline for the setup of the test infrastructure? Worrying about the impact of adapting HC to this use case on experiment (Atlas) manpower... Has to be planned in advance.
 - ◆ Maria: plan to have the first version of the package reporter for September, testing infrastructure in the following months

Report - H. Meinhard

Global organization of service managers and support staff in HEP

- Main characteristics: "transparent" exchange of experience

Last meeting in Annecy in May: large participation

- Strong participation from France, many newcomers
- Many after-talk and offline discussions: the real distinctive value of HEPiX
- Trip report by Helge attached to Indico

Agenda: the usual tracks, pretty well balanced

- Networking & Security: IPv6 progressing nicely, community increasingly involved
 - ◆ Also perfSonar now fully deployed
- Storage: CEPH becoming increasingly popular, in particular for cloud storage and dropbox-like services (ownCloud)
- Batch & Computing: HTCondor hype, several sites considering using it
 - ◆ Also early benchmarking activities of Atom-based servers
 - ◆ HS06 successor preparation work
- Basic IT services: Puppet a clear winner for configuration management but Quattor actively maintained and improved
 - ◆ Log analysis: ELK (ElasticSearch, LogStash, Kibana) everywhere
- Scientific Linux: CentOS still not stable enough to take a firm decision, agreement that HEP should participate in CentOS discussions and to launch a Scientific SIG with contributions by at least FNAL and CERN
 - ◆ Each team will start to build SL7 based on RHEL7/CentOS7 until the situation has been clarified
 - ◆ Decision if possible at the Fall meeting: preference for a common decision across the community
- Other OS topics: WXP being phased out everywhere, W8.1 already the recommended platform at several sites
- IT Facilities: new facilities at Orsay and Wigner
- Grid & Clouds: increasing production use of private clouds
 - ◆ Also several experiences with commercial providers, including Helix Nebula

WG news

- Puppet: volunteers sought for writing Puppet modules replacing YAIM
- Energy efficiency: insufficient interest again, going into passive/observing mode
 - ◆ The right people are not necessary part of HEPiX

American co-chair replacement at next Fall meeting: a search community formed, call for nominations launched

Next meetings

- Fall meeting: Univ. of Nebraska, Lincoln, Oct. 13-17
- Spring 2015: Oxford Univ., March 23-27

Actions in Progress

Ops Coordination Report - A. Forti

Next meeting dates: August 7 canceled

Working on changing meeting organizations to avoid duplication between meetings and make them more interesting for everyone

- Get experiment and TF reports 1h in advance at least
- Experiment reports tailored for sites
- Introduce T1 and T2 feedback sections
- Improve reporting to allow a better tracking of progress

DPM 1.8.8 released: new gridftp component had a backward compatibility issue (not working with FTS2)

- patched dpm-gsiftp released in EPEL and metapackage updated in EMI, UMD coming soon

T0

- SLC5 decommissioning: stop of grid submissions to SLC5 June 19
- ALICE still observing low efficiencies on SLC6 nodes
- LFC decommissioning for ATLAS

Experiments

- ALICE: KIT overload tracked down to the use of an old ROOT version: campaign in progress to get users upgrading their JDL
- ATLAS: issue with RFC proxy support in Condor, LFC -> Rucio migration in the US completed, Rucio full chain testing started 3 weeks ago
- CMS: glxexec test critical since May 19, xrootd fallback not critical yet
 - ◆ Operations is moving to GGUS: Savannah - GGUS bridge to be deprecated
- LHCb: dCache pb with brazilian certs understood and fixed
 - ◆ **CVMFS: new baseline 2.1.19**, 20 sites so far have deployed it for LHCb

Tracking tools: next release (July 16) will not support creation of tickets by email anymore

FTS3: new version with minor bug fixes expected in July, will move to EPEL

- ATLAS: no issue
- CMS: still 10 sites using FTS2 in PHEDEX debug instance, tickets opened

SHA-2: UK CA switched to SHA-2 on May 28

- No issue so far
- VOMS servers: blocking issue fixed on May 27th, waiting for UMD release to deploy it
- RFC proxies: ATLAS and CMS tests failed for different reasons

glxexec: still 10 problematic sites

WMS decommissioning: final date moved from end of June to end of August

- Waiting for the new SAM Condor-G probes

Machine/job features: good progress, more sites involved

- A Nagois probe is being developed

Network & Transfer Metrics: preparation work to organize the WG in progress, kickoff meeting in June?

Multicore: ATLAS activity stopped, CMS progressing with the tests at T1s

HS14 Preparation - M. Alef

HS14 candidates

- New SPEC suite
- GEANT4
- Also looking at a "light" version to do a quick estimate of performance of an assigned job slot in virtualized environments
 - ◆ Suggestions welcome!

Currently collecting apps representative of HEP experiment workloads

Discussion with WLCG Architects Forum about the compiler flags to use

New volunteers welcome: contact Manfred or Michele

IPv6 pre-GDB Summary - D. Kelsey

WG main motivation: all WLCG outward facing services should be run on dual-stack **to support IPv6-only clients**

- Also CERN is running out of IPv4 addresses soon, others likely to follow soon(ish)

pre-GDB well attended: ~30 in the room + ~20 on Vidyo

- Demonstrated that IPv6 is fully/easily usable at CERN!

Main achievement of this year: strong experiment involvement (all experiments)

- Critical mass reached for something to be done
- ALICE encouraging sites to deploy dual-stack services asap (xrootd v4 required)
 - ◆ ALICE well advanced with some central services and few site VOBoxes running dual stack.

Reports on IPv6 at CERN, testbed activities, status of storage/batch services

- Also site status and experiences, including monitoring (perfSonar, Nagios)
 - ◆ Perfsonar is IPv6 compliant, some 8 sites are currently being monitored.
 - ◆ Nagios, some new sensors are needed
- WLCG Apps Readiness tracked at <http://hepix-ipv6.web.cern.ch/wlcg-applications>
 - ◆ Good news: all storage solutions now have an IPv6-ready version
 - ◆ Batch systems: situation a bit more confused but not critical

Site survey run during 10 days end of May: got a large number of responses despite of this

- Response received from T0, all T1 except 1, 85/155 T2s
- Sites can continue to respond: results will be updated
 - ◆ After sending a reminder, may consider opening GGUS tickets to get the site plans
- IP connectivity: full connectivity at T0 and 16 T2s, partial 7 T1s and 9 T2s
 - ◆ No T1 anticipating shortage of IPv4 addresses, 5 T2s already facing this shortage
 - ◆ 5 T1s and 8 T2s have plans in the next year... but 45 T2s with no plans currently

Next testing plans

- Continue the file transfer mesh testbed: get more sites, in particular all T1s, involved
- FTS IPv6 pilot
- More storage options
- ATLAS: BigPanDA dual-stack, Rucio with IPv6
- CMS: dual-stack glideinWMS, IPv6 in AAA
- LHCb: dual-stack DIRAC services

Next milestone: dual-stack production services

- Several sites already did it successfully
- Need to find more volunteering sites
 - ◆ Joining perfSonar IPv6 testbed can be used as a first step
- Should be done in coordination with experiments: role of the WG
- Test IPv6-only WNs as soon as we have enough sites IPv6 ready: must not wait for all, prepare to run a in mixed environment

Need to create documentation based on early adopters experience

This topic: LCG > GDBMeetingNotes20140611

Topic revision: r3 - 2014-06-11 - MichelJouvin



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback