

Table of Contents

Summary of GDB meeting, March 11, 2015 (CERN)	1
Agenda	2
Introduction - M. Jouvin	3
SAM3 Update - Rucio Rama	5
EGI Future Plans and WLCG - P. Solagna	6
European Procurement Update - I. Bird	8
WLCG Operational Costs - J. Flix	9
Actions in Progress	11
Report - J. Flix.....	11
RFC Proxies - M. Litmaath.....	12
Summaries	13
Discussion with Other Sciences - J. Templon.....	13
Cloud Issues - M. Jouvin.....	13

Summary of GDB meeting, March 11, 2015 (CERN)

Agenda

<https://indico.cern.ch/event/319745/other-view?view=standard>

Introduction - M. Jouvin

Planning for 2015 in Indico

- No meeting in April (workshop in Okinawa)
- October will probably be cancelled
 - ◆ Idea of co-locating with HEPiX on HEPiX side: echo not overwhelmingly positive (except if it was held on the Sunday before)
 - ◆ Budget constraints due to WLCG workshop in Okinawa on GDB side
 - ◆ Still the option to move the GDB week but Michel not in favor
 - ◇ Proved not to work well in the past
 - ◆ Decision in May

Pre-GDBs planned in the coming months: May and June at least

- batch systems
- volunteer computing, accounting
- to be clarified by mid April

WLCG workshop: agenda pretty final

ARGUS

- Collaboration meeting last week
- Indigo Datacloud project approval expected to help
 - ◆ ARGUS in the cloud to use federated identities rather than X.509
- New release with patches already in use at some sites being prepared
- No new problems reported
- Preparing for Java 8 support

Data preservation

- training course on digital repositories at CERN 15-19 June
 - ◆ <https://indico.cern.ch/event/376809>
- DPHEP collaboration workshop 8-9 June at CERN
 - ◆ <https://indico.cern.ch/event/377026>
 - ◆ 1st workshop since MoU has been signed
 - ◆ report to be given at GDB on 10 June

Actions in progress

- list of "class 2" services used by VOs: NIKHEF agreed to start a twiki page with the list they are aware of for the 3 VOs they support (ATLAS, ALICE, LHCb)
 - ◆ Will ask CMS to provide the missing information when the initial list has been created
- Multicore accounting: still 15% of used resources not reporting the core count
 - ◆ Difficult to find how many sites are concerned
- perfSonar
 - ◆ 15% of the instances in bad shape: look at http://grid-monitoring.cern.ch/perfsonar_report.txt
 - ◆ T1s requested to dual-stack their instances by April 1 (not a joke!)

Discussion

- Jeff: where are we with the possibility to run IPv6-only WN? NIKHEF interested (wants to use containers).
 - ◆ Michel: better to talk directly with the IPv6 WG for details, should not be very far from making it possible
 - ◆ Ulf/Mattias: NDGF already doing much v6 for ATLAS. Main issue is storage. Still a few potentially problematic configuration between FAX and dCache.

SAM3 Update - Rucio Rama

SAM3 in production since last November.

- More power to experiments
- Increased flexibility in algorithms used
- VOfeed used to aggregate services into sites and implement VO naming convention
- Profiles used to define resources and algorithms to use for each VO service

Draft A/R report created at the end of each month: 10 days for asking for corrections/recomputation

- Recomputation can be triggered by experiments
- Can set manually the site A/R in case of problems not related to site
 - ◆ Wrong data can be set to unknown and be ignored in A/R calculation

Common schema with SSB, combine several UIs like myWLCG and SUM

Recent fixes to ALICE profile to fix issues with sites not appearing (neither CREAM nor ARC)

- Also NDGF T1 not appearing as a unique site

New profile for ATLAS: AnalysisAvailability

- Simpler algorithm
- Evaluated every 2h

Future developments

- NoSQL storage
- New operator: NOT
- Numerical metrics
- Combine data from several SSB instances

Discussion

- NIKHEF and SARA would like to appear as one site
 - ◆ possible, ask experiments
- Integration into site Nagios: see PIC component presented at a past GDB (mid 2014)

EGI Future Plans and WLCG - P. Solagna

EGI Engage funded: engage EGI community towards Open Science Commons

- Not only EGI: to be done in collaboration with other infrastructure projects (EUDAT, PRACE...)
- Easy and integrated access to data, digital services, instruments, knowledge and expertise
- User-centric approach: 40% of the project user-driven
 - ◆ Federated HTC and cloud services
 - ◆ Support of 7 RIs in ESFRI roadmap
- 8 Meuros (1/3 of EGI Inspire), 30 months, 1169 person-months, 42 beneficiaries

Strong focus on federation

- Security: evolution of AII infrastructure to enable distributed collaboration between diverse authn/authz technologies
 - ◆ Collaboration with AARC project
- Accounting, monitoring, operation tools
- PID registration service
- Computing and data cloud federation
 - ◆ Including PaaS managed by EGI if any need/use case
 - ◆ Virtual appliance library (AppDB)
 - ◆ Federated GPGPU infrastructure
- Service discoverability in EGI marketplace
- Collaboration with EUDAT2020 and INDIGO DATA CLOUD

Exploration of new business models

- Pay for use
 - ◆ Currently EGI doing brokering/match making between site price advertised and potential customers
 - ◆ Not yet clear if EGI will play a role as a "proxy" to charge the customers: currently direct relationships between sites and customers
 - ◆ EGI will provide sites tools to do the billing
- SLAs in a federated environment
- Cross-border procurement of public services
- Big data exploitation in various selected (private) sectors
- Investigating the potential impact on EGI governance

Distributed Competence Center: support for ESFRI RIs

- Help their VRE integration within EGI solutions
- Co-development of solutions for specific needs
- Promote RI technical services: training, scientific apps...
- Foster reuse of solutions across RIs
- Build a coordinated network of DCCs: European Open Knowledge Hub (EGI, ESFRI RIs, e-Infra...)

Prototype of an open data platform: federated storage and data solution providing sharing capabilities integrated with a federated cloud IaaS

- Includes a dropbox-like service: plan to reuse an existing, proven solution
- Deploy a best-of-breed existing tool as a prototype infrastructure: not necessary EGI only, not enough resources
- Collaboration with OSG and Asia-Pacific partners

Discussion

- Jeff: is the pay-per-use really the role of EGI?
 - ◆ Currently no actual enforcement of pay-per-use: just an indicative billing
 - ◆ Not clear if EGI will play a role in the billing process or just offer a service to do the match making between offers and demands
 - ◆ pay-per-use is not intended for all communities: clearly not for WLCG (pledges are used to match offer and demand) but some communities, like ESA, say this would be their preferred mode
 - ◆ Need to have an added value to commercial cloud providers: not our role to compete directly with them
- Jeff: why EGI has to deal long-tail science users, should be the role of NGIs
 - ◆ Peter: wording may be ambiguous but EGI is supporting NGIs rather than long-tail science users directly. But sometimes initial contact is going through EGI (during conferences for example) instead of NGIs. Also some countries/regions with no NGI or a weak NGI.
- EGI clearly addressing new communities, not clear what space is there for a large existing community like WLCG
 - ◆ Operations and AAI R&d/evolution are important topics for collaboration
 - ◆ WLCG sites offering services to other communities important as well : ensure that procedures for WLCG and EGI resource provisioning don't diverge more than necessary else it will become a problem for sites

European Procurement Update - I. Bird

Several presentation about the European procurement idea during last Fall but not much positive feedback but funding agencies insisted about the need to make progress on this idea

- Paper attached to agenda summarises the situation and the potential

European Science Cloud pilot projet

- Bring together many stakeholders to buy workload capacity for WLCG at commercial cloud providers
 - ◆ Commercial resources to be available through GEANT, integrate with federated identities,
- Funded by H2020 ICT8 call as Pre-Commercial Procurement (PCP) proposal to EC in April 2015 (14)
 - ◆ A group of research organizations pledge procurement money to the European Science Cloud
 - ◆ The project defines the technical requirements
 - ◆ PCP is the approach taken for LHC magnets where the products not yet existed: allows an exploration phase for defining the design and a prototype phase. Also a wrapping phase to prepare the project follow-up. In this case, 6 months for preparation, 18 months for implementation, 6 months for wrapping up.
 - ◆ EU funding is proportional to the project member contributions: reimburse at the end of the project up to 70% of member contributions (members need to fund the total budget initially).

Early works in the experiment and in HELIX NEBULA demonstrated the feasibility

- Also some quotes at the end of Helix Nebula demonstrated the prices of commercial cloud services was closer to in-house resources for some use cases (in particuar simulation)

Buyers group: public orgnisations from WLCG collaboration

- Procured services will count towards the buyers pledges in WLCG
 - ◆ Initially, participation proposed to all T1s
- Other communities could benefit from procured services (~20%)

Timescale: project starting in Jan 2016, implementation by end of 2017

- Would be in place for the second part of Run2

Discussion

- Do we have an initial list of interested partners?
 - ◆ Ian: not yet, still in discussions

WLCG Operational Costs - J. Flix

~100 answers to survey

- 1 (anonymous) answer per site

5 areas surveyed

- FTE effort spent on operation of various services
- Service upgrades and changes
- Communication
- Monitoring
- Service administration

Supported VOs

- Most sites either dedicated to 1 LHC VO or supported most of them (3 or 4)
- T2 typically support ~10 VOs but large distribution

FTE effort quantification

- Aware of the potential inconsistency between sites but most obvious mis-interpretations fixed. Still need to be careful with conclusions.
- Ticket handling effort: no clear correlation between the FTE spent on VO support and the number of LHC VO supported
 - ◆ A bit surprising... but inline with the grid promess!
- T0/T1: FTE dominated by storage systems and "other WLCG tasks (experiment service, OS and configuration...)"
 - ◆ Average of 12.8 FTEs per T1
- T2: storage and other WLCG tasks also among the largest fraction but not in the same proportion as at T1. APEL is a major area for FTE effort at T2.
 - ◆ Average of 2.8 FTE/T2
 - ◆ Small effort for participation to WLCG TF and coordination
- FTE effort seems to be clearly correlated to site size (based on the HS06 or PB delivered by site)
 - ◆ Less clear for storage than for CPU
- Core grid/experiment services take more effort at T0/T1 than T2
 - ◆ APEL is the most often mentioned service at T2
- Networking effort similar in T1s and T2s

Communication

- Importance of experiment requests coming from WLCG Ops: no clear indication that something should be changed
 - ◆ Future analysis: may be interesting to correlate site responses with site size (dedicated or multi-VO sites)
- Possible improvements suggested
 - ◆ Better distinction between official requirements and suggestions
 - ◆ Blessing/endorsement of new service/protocol requirements by WLCG MB before making them a formal request
 - ◆ WLCG Ops bulletin. Maarten: we already have the WLCG Ops meeting minutes... Collect more feedback from sites before making new requests...
- Encourage more participation to both HEPiX and GDB
- Create site service specific e-groups
- Consolidate information into open WLCG wikis

- ◆ Currently often in experiment (protected) wikis
- WLCG OpsCoord meeting: low regular participation from T2 but the majority reading the minutes
 - ◆ Still a small fraction not reading the minutes: need to address it
 - ◆ Suggestion for a shorter, more focused meeting (1h)
 - ◆ Time slot not entirely convenient for US and doesn't allow asian participation
 - ◆ Put more information from sites in the minutes
- WLCG TF seen as useful
 - ◆ Most non participating sites said that it was because of the lack of manpower
- Sites happy with GGUS
 - ◆ Easy programmatic access to current and historical contents would be welcome
 - ◆ Support for every MW component should be through GGUS
- WLCG broadcast and GGUS tickets seen as the best channels to pass requests to sites
 - ◆ Reducing the number and the duplication of broadcasts make them more effective
 - ◆ Michel: a bit surprising compared to the experience where only tickets tend to get the actions done

Conclusion: some improvements needed, but generally things not too bad

Actions in Progress

Report - J. Flix

VOMRS finally decommissioned March 2!

- Experiments acknowledge efforts by CERN-IT and VOMS-Admin developers

Savannah was decommissioned on Feb. 19

- Inactive project archived
- Others migrated to JIRA

Baselines

- UMD 3.11.0: APEL, CREAM-CE, GFAL2 and DPM
- dCache: various bug fixes for different versions
- New argus-papd (1.6.4) fixing issues seen with recent Java version
- FTS 3.2.32: activity shares fixed

Freak vulnerability classified as low risk

LFC-LHCb decommissioned March 2

- LFC to DIRAC migration successful
- The only LFC instance left at CERN is the shared one: discussing the future with EGI

Experiments

- ALICE: high activity
- ATLAS: cosmic rays data taking, MC15
 - ◆ Tricky pb with FTS shares understood and now fixed by developers
- CMS: cosmic rays data taking, global Condor pool for Analysis and Production deployed
 - ◆ Also tape staging tests at Tier-1s ongoing
- LHCb: restripping finished

2nd ARGUS meeting: see minutes² and Michel introduction

Oliver K. proposed the creation of http deployment TF

- Mandate approved: identify features required by expts, providing recipes and recommendations to sites
- ATLAS, CMS and LHCb support the TF
 - ◆ ALICE currently not interested

glexec

- Finishing Panda validation campaigned (63 sites covered)

IPv6

- T1s requested to deploy dual-stack perfSonar by April 1
- FTS3 IPv6 testbed progressing

- CERN CMVFS Stratum 1 working well in dual-stack

Multi-core deployment

- Successfully shared resources between ATLAS and CMS

MW Readiness WG

- Participating sites asked to deploy Package Reporter: progressing well
- MW database view will indicate versions to use

Network Transfer and Metrics

- perfSonar: see Michel introduction
- Integration into experiments: LHCb pilot, extending the ATLAS FTS perf study to CMS and LHCb
- Network issue between SARA and AGLT2 being investigated

RFC Proxies - M. Litmaath

Difference between legacy and RFC proxies: latter better supported, while legacy proxies have already given rise to issues

- Should switch to RFC proxies this year

Status on service side

- CMS have moved months ago
- ALICE: Switching VOboxes to RFC proxies now
- ATLAS and LHCb checking
- Other players: SAM-Nagios proxy renewal needs an easy fix
- Anything else?

UI clients

- legacy proxies are still the default
- RFC proxies could become default later this year (to be coordinated with EGI and OSG)

Discussion

- P. Solagna: EGI shares the goal of moving to RFC proxies asap, plan proposed seems realistic, no major problem foreseen. Happy to coordinate with WLCG on this
 - ◆ Change of default this year is probably okay for EGI

Summaries

Discussion with Other Sciences - J. Templon

Coorganized with the Netherlands eScience Center NLeSC

- Introduce HEP to NLeSC and other sciences to HEP
- NLeSC: help scientific communities to address their computational challenges and use efficiently e-Infrastructure
 - ◆ Part of an ecosystem with e-Infrastructure and computer science: NLeSC doesn't operate any resource
 - ◆ Project based: provide expert manpower to a project for a certain duration
 - ◆ Interested into turning project developments into more generic solutions/services

Data challenge in Astronomy with next generation experiments (SKA): no possibility to keep on disk intermediate data products

- Streaming one algorithm to another one, almost realtime
- Close to challenges seen in LHC experiments

Data challenge

- Strong move in HEP in adopting industry standards
- HEP has experience in handling huge volume of data: 1 PB/week to tape...

Everybody interested by the contact

- NLeSC interested further in further contacts, visit their site
- NLeSC involved in SoftwareX which hosts a SW repository: why not to publish ROOT, GEANT4 or other HEP SW

Also see [<https://twiki.cern.ch/twiki/bin/view/LCG/GDBMeetingNotes20150210>][summary].

Cloud Issues - M. Jouvin

Attendance: some 25 local, many remote

- No experiment representatives in Amsterdam but a few remotely connected

Review of work in progress after the last meeting in September

- Dynamic sharing of resources: Vcycle looks promising, a lot of improvements in the last 6 months
 - ◆ Possibly complemented by fair-share scheduler for OpenStack
- Accounting: still a lot of work to do but most solutions agreed
 - ◆ Still a potential issue about double counting resources as grid and cloud
- Traceability: already some work done after the initial meeting one month ago
- Data bridge very interesting: opening a way for using federated identity to access storage

Discussion about EGI federated cloud

- Already a collaboration on accounting
- Potential interest for the EGI monitoring infrastructure but requirement of OCCI may be an obstacle: more thoughts required

- Should work in common on integration of federation identities

Also see summary

This topic: LCG > GDBMeetingNotes20150311

Topic revision: r3 - 2015-03-23 - MichelJouvin



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)