

Table of Contents

Notes of GDB meeting, July 11th 2012.....	1
Welcome (Michel Jouvin).....	2
CVMFS deployment status (Ian Collier/Stefan Roiser).....	3
Summary of Pre-GDB Meeting on CE Extensions (David Salomoni).....	4
Post EMI Discussion (Michel Jouvin).....	5
IS Evolution Follow Up (Maria Alandes Pradillo).....	6
GL Exec Deployment (Maarten Litmaath).....	7
Federated Storage (Fabrizio Furano).....	8
DPM Community (Oliver Keeble).....	9
SHA-2 and RFC Proxy Support (Maarten Litmaath).....	10
Introduction to ARGUS and central banning (Valery Tschopp).....	11
Operations Coordination Team (Maria Girone).....	12
Jobs with High Memory Requirements.....	13

Notes of GDB meeting, July 11th 2012

See:

* <http://indico.cern.ch/conferenceDisplay.py?confId=155070>

Welcome (Michel Jouvin)

Meeting schedule:

- Meetings will continue on the 2nd Wednesday each month to mid-2013 (no exceptions)
- Pre-GDBs one every month in next few months. Will be confirmed with meeting notice 3+ weeks in advance
- August meeting is cancelled
- October meeting (and pre-GDB) is planned in Annecy

Next Meeting agenda items (September):

- pre gdb=data/storage protocols)
- IPV6
- LS1 plans
- WG presentations
- HW trends (bernd)
- Other issues to follow up See actions in progress in twiki area
- Data storage working groups
- EMI-2 WN testing to be re-discussed in September

Extended run confirmed – still wait for details. Consequences on pledges unclear but probably to late for any significant adjustments.

See presentation for other relevant dates.

CVMFS deployment status (Ian Collier/Stefan Roiser)

Current deployment status is:

- ATLAS – 78/104 sites have deployed. Plan to have all sites migrated by 2012Q4 – WN disk space main obstacle to deployment.
- LHCb 36/86 sites have deployed. **Officially request all sites supporting VO to deploy CVMFS.** Tier-2 sites with CVMFS will be preferred for higher workload.
- CMS (5 T1s migrated) **encourages sites to move to CVMFS**

Many of the German sites are constrained by small WN disks. One site has problems when cache 50% full of pin files. Reported as a bug. Not easily fixed in current branch. New branch will handle differently and reduce catalogue space requirements. Typical experience is that a 10GB ATLAS cache contains 3GB pin files – need to monitor as when it hits 5GB ... problems.

Work on a shared cache client is ongoing (Jakob Blomer). A handful of sites testing are testing an NFS client as an alternative solution. Expect a beta release to be available by the end of August and plan a slow ramp up of sites testing the new features. **Request sites volunteer for testing new version** as progress limited without test sites. New branch has significant features that may resolve outstanding issues.

Summary of Pre-GDB Meeting on CE Extensions (David Salomoni)

See <http://indico.cern.ch/conferenceDisplay.py?confId=196743>

Meeting goals:

- Review possible extensions – particularly whole node/multi-core (n-m) + memory requirements.
- Agree development plan and timeline

Implies changes in CE and experiment frameworks. The meeting focused on Glue 2 solution (not 1.3). Agreed best to publish features at queue level rather than site level. The CE developers willing to develop this type of extension.

- **ARC CE** already supports whole node/ fixed or multi-core through RTE scripts.
- **CREAM** – supports whole node/fixed multi core. Need to know how to specify requests in JDL etc.
- **OSG** – in middle of decision making process for long term CE plans. Has been able to get whole machine and n-core jobs already on all OSG LRMS. Do we need variable number of cores

Environment variables will be needed to tell job what resources it may actually use. Details of the spec for environment variables available on wiki page. The team are **looking for a test LSF site** . NIKHEF working on Torque tests. There was discussion about the inconsistency between using scripts and files to set environment variables. See Tony Cass's presentation in last GDB and Ulrich's WM TEG wiki page. **Any comments on implementation details should be sent to WM TEG mailing list please.**

Main issue is how to let CE know what site supports. See slides for enhancement steps needed in order to get whole chain working.

- CMS are prepared to test whole chain.
- LHCb willing to test workload management.

Need limited number of volunteer sites. Start soon after summer with larger tests.

ATLAS have no requirement for (n-m) feature. Believe it is extra complexity not needed at the moment. Should this be implemented on level of CE rather than PANDA? WM TEG believes it should be.

Jeff – Vos need to realize that they must relate pilot factory to payload. Site will not be happy if they have to reserve cores that pilot factory then does nothing. Some VOs do this, but others leaving pilots hanging around doing nothing.

This will be discussed again at the September GDB

Post EMI Discussion (Michel Jouvin)

Initial meeting with EMI/EGI and OSG. Jamie was concerned that not all interested parties had been invited to the meeting. The objective was to identify issues relating to WLCG and the end of EMI and supporting projects (not attempt to solve them). Foster operational relationship between European and North American part of WLCG. Mainly focused on the middleware – EGI not much discussed.

Issues:

- Globus – need to address support and packaging. Support can only come from the community. Packaging from IGE on best efforts for medium term. Need to identify who are the support community.
- EMI middleware. Most WLCG components have a long term support plan. HUC products not part of EMI list – discussion in progress.
- OSG discussion in progress with external providers.
- Middleware validation and provisioning – reference repositories? Baseline definition etc etc. Markus will write initial document on middleware. Identify areas where OSG and Europe could collaborate).

Jamie concerned about loss of HUC support (25 FTE has been expended in this area). HUC has made a significant contribution to experiments – big impact when it goes. HUC Wasn't discussed at the post EMI discussion. Needs urgently a discussion with all four experiments to agree what is needed. Jamie will organize meeting by September – mainly an internal CERN discussion.

IS Evolution Follow Up (Maria Alandes Pradillo)

Working to identify the best top level BDIIIs for sites to use. Described how a discovery method was being used to understand how sites were configured to use the BDII service – via SAM system. Has queried 526 CEs/WNs. Most WNs use only 1 BDII. Only 11% contain at least 3 BDIIIs.

The next step will be to check most heavily used BDIIIs for quality of service and service configuration (will contact sites concerned). Intention is to provide recommended configuration information around baseline services page – eg optimal configurations and practice.

GL Exec Deployment (Maarten Litmaath)

There has been very slow progress. glExec mainly set up and working at the Tier-1s but overall, the number of CEs supporting glExec increased by a few (83 to 85). Some sites disappeared/appeared.

Maarten plans to do another broadcast after the holidays to stimulate more progress. CMS raising tickets against sites. Hope to see ramp up over coming months. Some sites respond to CMS – some not. Eventually will become a critical test for CMS.

Federated Storage (Fabrizio Furano)

Federated storage one of the main recommendations of the TEG. We need to make concept better known. Understand issues, reduce differences amongst VOs. A Federated Storage working group is being formed. The working group membership will include:

- Tech providers
- Experiments (1 or 2)
- Sites (2 or 3 site managers)

Fabrizio has contacted various groups looking for volunteers. Need experience in topic. Hoping for a first meeting in time for Federated data stores workshop in Lyon.

There was discussion how ARC systems fits into this. There is a project ongoing at the moment for WAN data access to ARC systems. Would like to integrate with this federation project if possible.

DPM Community (Oliver Keeble)

The DPM/LFC project is heavily reliant on European funding – soon ending. The project has never been in better shape.

- **DPM** - Many new features now or coming soon. CERN does not run a DPM instance and believes collaboration of stake holders will be necessary to fully support future support and development. Worldwide there is over 36PB (10 over 1PB) of DPM capacity. There are over 200 sites in 50 regions, over 300 VOs have access. France and UK biggest users by capacity.
- **LFC** - ATLAS and LHCb plan to move away from using the LFC but no timescale set yet. However numerous non HEP VOs continue to use it. Ongoing support still needed.

The DPM project is appealing to these big user regions to come forward to discuss support model. Sites are invited to express their interest in a collaboration (with MoU). Timescale to identify interest by end of the year.

CERN will not abandon DPM at end of EMI. If the community does not step up to help support DPM then Plan B would be to manage critical bug fixes and security fixes but no significant development (unlike recent past).

SHA-2 and RFC Proxy Support (Maarten Litmaath)

IGTF intends that CAs will be free (or even recommended) to commence issuing certificates with SHA2 from 1st January 2013. There are various pieces of middleware and experiment-ware that still need to be made ready to support SHA-2 or RFC proxies. Even after release, it would be many more weeks before products are available in UMD. Looks impossible for WLCG to prepare for SHA-2 in 5 months (particulatrly with extended run). Planned phases and milestones indicate cannot deploy SHA-2 CAs until Jan 2014! Can we persuade IGTF to postpone for 1 year?

In the event that IGTF will not delay and a supported WLCG CA moves to SHA-2 we need a plan B. It is proposed that we create a new CA at CERN issuing short lived SHA-1 certificates (under / namespace) to any WLCG member who cannot get a SHA_1 certificate from their home CA. A significant effort and will need to start soon if we are to be ready by January 2013. Discussion about the viability of this plan took place.

We need IGTF to reconsider timeline by September if we are to deploy a new CA by January. We will have major problems if a major CA moves to SHA-2 in January. **Dave Kelsey will pursue this matter with IGTF.**

Introduction to ARGUS and central banning (Valery Tschopp)

- Valery presented an overview of ARGUS Authorisation engine, PAP language and PAP admin tool.
- CERN are running central ARGUS service. **Sites should check how they would link their local system to central banning service.** Ongoing discussion with OSG as to how they can interface to a central banning list without running ARGUS servers.

Operations Coordination Team (Maria Girone)

- TEG recommends service coordination/commissioning team. Computing as a Service. Goal by the end of LS1. Team needs representation/knowledge of sites/regions. Persistent WG. **Volunteers needed from sites (both Tier-1s and Tier-2s).**
- Will relate to existing structures such as daily ops. Short term task forces will be created to address specific deployment / de-commissioning issues. Members of team will be Service Coordinator On Duty – SCOD. Looking for team of 15-20 core members– in team, not a full time commitment but able to take on actions (SCOD role maybe 20% when on duty). Replace T1SCM with Fortnightly Operations Coordination.

Jobs with High Memory Requirements

Experiments provided a brief summary of their status and issues:

- CMS (Steve Gowdy) CMS Jobs kill themselves if they exceed 2.3GB real memory. Vast majority of jobs not >> 2GB.
- ALICE (Maarten Litmaath) discusses problems with allocate/deallocate memory in order to keep memory footprint down. Problems particularly compiled under SL5 – SL6 much better.
- LHCb – highlights that different sites have different policies which makes situation difficult to manage. Explanation from IN2P3 explaining why constraints on VMEM per job required by sites.
- ATLAS – physical memory requirements still 2GB / core but still need some jobs > 3.5GB VMEM

Experiments were concerned that the sites follow different strategies. Some sites were concerned that if job resource limits were not implemented then there was a risk that problem jobs would cause disruption to service. VOs would like few or no limits. There was discussion on a range of issues but little general agreement about the way forward or under what forum. One conclusion look at PSS to see if it is a useful metric that can be used by batch system. Suggest batch system experts look at this. Maarten proposes we set up a TWIKI page to survey site memory management policies.

This topic: LCG > GDBMeetingNotes20120711

Topic revision: r6 - 2012-07-21 - MichelJouvin



Copyright &© 2008-2022 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