

GDB Notes

John Gordon

April 2009

The April GDB was held on 8th April 2009. The agenda is here <http://indico.cern.ch/conferenceDisplay.py?confId=45474>

Security Service Challenge – Run 2

Sven Gabriel of Nikhef who ran the second of these challenges against Tier1s, presented the preliminary results. This was a development of the previous challenge run last year. In this the site was asked to trace a job back from WN, through CE, WMS, to the submitting UI. To ban a particular user, and to trace certain storage operations. Nine Tier1s (NIKHEF **and** SARA, excluding OSG sites and ASGC)_ were tested plus one Tier2 Prague who volunteered.. The tests were announced in the OSCT phone conferences and the evaluation criteria were described in the alert email. Six out 9 sites equalled or exceeded the maximum score (bonus points were possible). One of the others and the Prague T2 scored >90%. The improvement for the sites who were previously poor to middling has been considerable. The exception was the INFN T1 at CNAF. It took three attempts before there was any response at all and even that was poorer than last year. The EGEE Security Officer will discuss in detail with CNAF but **the MB should be concerned at the apparent inability of this T1 to react to standard procedures.**

This test is currently being run against Tier2s as well. UK and South-East Europe have completed, Asia Pacific and Benelux are in progress and NDGF and OSG are preparing. The aim is to have completed all regions in time to report to the EGEE09 conference in September.

STEP09

The LHCC reviewers shared the concerns of the T1s that the experiments had not stressed the T1s with the full range of anticipated workflow at the same time. At the WLCG workshop in Prague a window in May-June was identified when hopefully the real work of the experiments could coincide with each other to exercise the T1s with data transfer, reconstruction including staging from tape, simulation, and analysis where appropriate. Jamie stressed that this should be nothing special, all sites should currently be running in full production but reaction to the inevitable problems needs to improve.

The four experiments presented their plans. Alice are ready to run from late May until the end of June. CMS plan to run through the work in May and then repeat 8-19 June to align with ATLAS. ATLAS wish all T1s and T2s to participate. No exceptions for T1s, T2s to make excuses by the 15th of May. Preparation last week in May full run first two weeks in June. Could slip by one week to align with others. LHCb will run their FEST09 during this period.

So we seem to have an overlap period 8-12th June, and possibly the following week too.

EGEE Authorisation Service

Christoph Witzig gave an update on progress. The service has been tested at 4 sites: CNAF, HIP, NIKHEF and SWITCH. A 3 day test was carried out with remote command line clients. Results were satisfactory, much less than the expected total glxexec time. The planned deployment plan starts with

1. glxexec on WN
2. OSCT Banning list
3. CREAM
4. WMS

Since all of these can be chosen by the site there should be nothing disruptive in sites choosing to deploy it one at a time and, like CREAM, widespread deployment can wait until it has proven itself performant and reliable at a smaller number of sites.

Identity Management

David Kelsey reported on recent developments in international authentication. Our current grid model relies on the persistence and uniqueness of UserDNs. If a CA issues a DN then we rely on the fact that it has a long lifetime and that the CA will never reissue it to some other individual later. We also get a 'warm and fuzzy feeling' from the presence in the DN of an identifiable human name. Today some countries are moving from small CAs run primarily for grid use to large-scale academic federations. New IGTF authentication profiles are being proposed including some based on short lived credentials. While the certificates thus issued are technically compatible with those we use today we cannot make the same long-term use of them.

David asked for GDB's backing in his negotiations, putting the WLCG position. He proposed:

- Face to face identity vetting
- Persistence of subject DN
- Appropriate presentation of real name in DN

The GDB supported his proposals.

Distributed Monitoring in EGEE

Steve Traylen gave an update on the work of the EGEE Operations Automation Team (OAT) which is working towards the distribution of many EGEE operational services to the EGEE regions in preparation for further distribution to countries as NGIs. In particular he reported on the checkpoint from a workshop the previous week. Among the topics covered were:

- Redefining the definitive sources of various types of information between GOCDDB, BDII, and other sources.
- AN Aggregated Topology Provider (ATP) which combines various sources of information into specific views – e.g. Clouds of ATLAS.
- The replacement of the central SAM test service with Nagios running in each region testing its own sites and sending the results by the Message Bus to a variety of places including a central Metric Repository to replace the central SAM database.

This should mainly be in place by the end of this year.

Distributed Monitoring in OSG

Brian Bockelman described OSG monitoring. The OSG strategy is to centralise as little as possible so for example Gratia keeps all data locally at a site and reports to a central place where the information is aggregated. Gratia includes its own transport mechanism which is similar in design to the EGEE Message Bus but open source solutions like ActiveMQ were not considered mature enough when it was designed. Monitoring of data transfers has started and more sites are joining. RSV monitors services using a condor-based cron-like mechanism. Plans to move this into Nagios, similar to EGEE plans. Both GlideInWMS and Panda are developing sophisticated job-level monitoring. Looking at unifying logging formats to allow machine analysis of logging data. It would be good to work on this with EGEE as some middleware is shared.

Middleware

Andreas Unterkircher gave the latest news on his timeline for SL5 releases of middleware components. <https://twiki.cern.ch/twiki/bin/view/EGEE/SL5Planning>

The WN was released on 23/3/2009 but there has been little take up so far. **T1s should confirm their plans for deployment.** NIKHEF, at least, is waiting for UI before moving WN.

Pilot Jobs

Little recent progress on frameworks. Would like to hear from ATLAS and LHCb on their glexec testing experiences. SCAS testing has started. More detailed news in May.

SRMv2.2 MoU Extension

Andrea Sciaba reviewed progress on the implementation of the features defined in the extension to the MoU. The following will not be available for 2009/2010 data taking

- Client-driven tape optimization
- VOMS support in CASTOR
- Space selection in read operations in dCache

One Castor feature is only available in release 1.2.8 and so will only be available at CERN once this release is in general production and at Tier1s if they upgrade before data taking. A few dCache features are only available in 1.9.3 which is not yet available. Unclear whether sites will have time to test and upgrade before data taking.

Experiments should comment on acceptability.

EGI

John Gordon gave his view of current progress. There are still separate strands of work towards an EGI.

EGEE is writing a second year workplan to anticipate EGI. Anticipating no-cost extension if required. Keep services running for EGI as priority

EGI Policy Board took the decision to host EGI.eu in Amsterdam. Setting up formal EGI organisation with Council and MoU. Appointing a Director, forming Teams to develop transition plan who will also work on proposals

EU and many NGIs stress that EGI must deliver for LHC. John understood that to mean that no-one is interested in an EGI that doesn't deliver for LHC computing. Ian Bird understood it to be a call for WLCG to support EGI. WLCG needs to define what it means for EGI fully to support LHC. Also to propose an SSC for PP as part of EGI. IB showed a slide defining all the services required. JG showed one of the EGI.org operational services with their EGEE equivalents. Long discussion. **Tier1s to be asked for details on how existing EGEE support will be delivered in their country/region post-EGEE.** Some in the room gave brief updates.

Pre-GDB Meetings

A number of pre-GDB topics have been identified.

- Tier2 Storage using DPM, StoRM, and possibly Hadoop, following on from the dCache workshop held in February. Still discussing suitability of May
- Virtualisation – In June. following on from a half-day session at HEPiX in Umea at the end of May.
- Site Management Tools – July

In addition a meeting of SRM developers is envisaged to consider the MoU amendment.

Next GDB

Wednesday 12th May at CERN