September 2010 – GDB minutes

Introduction (John Gordon)

Recently: DM & storage jamboree; EGEE III final review; LHC OPN meeting and WLCG workshop.

For information, next meeting is 13th October. Will continue with the second Wednesday of each month in 2011. Lyon offered to host in March.

Upcoming meetings: EGI Technical Forum; LHCOPN; CHEP 10 and HEPiX.

On pilot jobs. Asked all T1s to install glexec/SCAS etc so experiments could test and regular tests on status can be run. Would like to hear on progress from ATLAS and CMS – LHCb and ALICE have given some feedback on their changes already.

??: CREAM – running CE alongside LCG-CE. Can not really call it a production service, certainly not at same stability with CREAM.

JG: Yes, the test results just shown were for the nagios glexec test of CREAM.

A new version of condor was provided that fixes the condor-CREAM submission issues reported by ATLAS.

Graeme Stewart: Confirmed. Got a new version before summer break and from initial tests things look better.

DM and virtualisation – at least two experiments have mentioned desire for reserving whole WNs. There is also a proposal for a better managing of the Information System and data.

Steve Traylen: Are they taking suggestions? It is not just the top-BDIIs it is also the number of sites covered.

JG: Laurence open to suggestions. Also would restrict number of sites queried and have someone manage the WLCG information tree.

MS: Some mechanisms to address this have actually increased the amount of information – e.g. Glue localisation attribute. So there is information that we need to review as to whether it is needed. Having a smaller set of top-BDIIs helps with follow up and management of problems.

Installed capacity: Checking numbers suggests site confirmation needed. Each T1 need to go away and check gstat figures with what they have. NIKHEF for example has 3 times pledge installed? No. … also few sites publishing the nearline figures.

KIT figures are correct on the gstat-prod instance.

IB: Can someone report this to Laurence?

On the US sites? JG: Not sure how publishing works.

IB: US has their own BDII that they manage.

Action: All T1 sites to check their figures – correct/verify this information by the end of September.

Action: Will ask for report next time on the status of the VO SAM tests

Shared Software Area

An alternative model to distribute VO specific software to WLCG sites: a prototype at PIC based on CERNVM file system (Elisa Lanciotti PIC)

Looked at this because of a number of limitations of the current model for software distribution – e.g. NFS scalability, limited quota per VO…

RW: Is this over plain http to download the files? Do you plan to download plain text files in this way in production?

Testing CernVM-FS scalability at RAL Tier1 (Ian Collier)

IC: Suspect use of squids will not generate a lot of traffic. Need to monitor how when people start using it the load increases on the proxy servers at CERN.

??: For CERN VM use the proxy servers CVN – point back to CERN proxies. Additional layer helps to cope with higher loads.

IC: Also to note, benefits to network/server load – these would be true of any caching file system (eg. NFS4). The part that is not unique is the file based duplication. With AFS, flexible read-write but being used to be read-only for one group and write-only for another – so the setup and config is complicated. Here however the central management is the only place the config needs to happen.

JT: Looks promising. Some questions. This will require some cache on each WN. Experiments also want some cache. So how big does the cache need to be?

AF: How do you publish the releases.

GS: There would be one central store so no need to publish. The validation is one area that is though still in need of discussion.

APEL Status and Plans (Cristina del Cano Novales)

JT: Did not use APEL parser. Have only database and injected directly to RGMA using a direct schema. Do not know how to do this via activeMQ.

CN: You are region C. We will provide the publisher.

JG: Could either move to be region A like others or …mysql

Very few issues. Have improved documentation. Some cert issues. Now clarified through GGUS.

ST: CERN gave up with the RGMA publishing. Move for us no problem only an issue with joining the data.

Helger: Difficult bit was gathering the data. ActiveMQ works better. Have option at CERN of grid and local submission. Creates problems at accounting stage because the support only support the grid submitted work and that means gymnastics for us.

IB: How is publishing done now at sites? Grid job vs async local job….

JG: No grid ID with local jobs. Cristina's suggestion to publish with "local" tag. Had problems with integrity. Could look at CERN solution.Looked like non-grid jobs were phasing out.

Holger: CMS for example does not appear large in usage because much of the work is local submission at CERN.

Cristina – jobs appear local unless you can join it. Often jobs appear local but end up grid based and duplication can occur. Will work with CERN to find a separate solution.

ST: Time out is the issue – how long do you wait before trying to join?

WLCG Accounting Requirements (John Gordon)

OK: There is an EMI working group now looking at storage accounting.

JG: Would be useful to have input from that group at the meeting next week.

GS: About pilot jobs – ATLAS have interface that allows us to see what user submitted what jobs. From our point of view I would not say we need to see the DN of the submitter in APEL but the site may want to see it.

JG: Accounting is needed for two things. For management do not need the indentiy. And sites are not so concerned about individual use, more about the VO allocation vs usage.

IB: The simpler we can make it the better.

JG: So you are saying you may not need user accouting at the DN level?

IB: Yes in the cases where pilot jobs are used.

JG: General message is that you have most of what you want and local jobs are the only additional thing you want in.

Middleware Update (Andrew Elwell)

There are products that are not adequately tested in staged-rollout, so the components get released to production as there is no negative feedback.

Concerning SL4 functional updates. DPM should be on list. When does this list timeout?

JG: SO nobody is suggesting a critical cut-off for a given product? I mean stopping/banning support.

MS: Highlights problems. There is no guidance for the infrastructure and nobody to make the decision. EGI should have some say in it.

JG: There is a WLCG baseline release. If others want to run it that is fine.

IB: From an effort point of view the point when we can stop critical updates.

JG: I might have said security patches… if critical means ATLAS has a DPM problem is that critical.

OK: Documented in EGEE by Nick Thackray. Included the concept of what a critical update means. For the dates, gLite can take this and come back with dates.

JG: Nick's approach looked at when each product was fully functional etc. There is a

JT: We are one of the sites that want DPM supported on SL4. This is due to hardware. For this could go for gLite 3.2 on SL4. Happy to upgrade if can install the OS.

JG: Raised this at MB and were shouted down. Ok, before putting things in baseline we should consider against the criteria in Nick's wiki.

AE: Is there a timeline for phasing out 3.1?

JG: The feedback and therefore dates will vary.

LHCOPN update (John shade)

JG: The requirement for this came out of one of our previous meetings so it is good to see this moving forward. What is TOMS?

JS: It is James Casey

Alessandro: Checking now, from experiment support point of view we had a problem between INFN and BNL yet the dashboard is green. I see from your slides you are oly taking into account a one-way delay.

JS: Start with something simple and work from there. Now collect information about total bandwidth so extra functionality foreseen.

KB: Same issue, there is a problem with the interface with the experiments. CNAF-BNL issue mentioned at every meeting .. it should be reported once and escalated.

JG Problem is multi-layering

KB: There is no feedback or commication back about what is happening.

SC: Remark on dashboard. Current view good. Historical views would be useful.

JS: This is foreseen.

SC: When do you expect to have something to present to experiments? Had many point to point issues and never know at which point to start.

JS: Not sure about the TOM version of the dashboard. Would want something for the next OPN meeting.

LUNCH

Experiment Operations

ALICE (Latchezar Betev)

No questions

LHCb (Roberto Santinelli)

Q: Why did you call it xroot? It should be xrootd.

Is there an operational view as to why "A run == a single site"?

RS: Do not remember.

LFC for LHCb is a multi-replicated system so it is not critical?

RS: The agent performing the transfer – wrapper around the client.

GM:

The shared software area is somewhere where everyone needs to work hard. Another option is to put the software on each WN. …

JG: You were not here this morning but we did cover options this morning.

At GridKa…Have recently implemented dedicated NFS servers… investigated what LHCb is doing. LHCb have many small files and LHCb software frequently touches all files/directories which is effectively a software area DNS. If you would use a tarball and include files this would greatly help performance.

GS: A lot of meta-data operations …

GKa: So effectively a file system not NFS problem

JG: GODCB is the place where services should be defined.

GS: On compiling – know the PIC issue was a single user issue.

On RAL disk servers…

JG: Lot of investigation. Different sort of jobs doing different things has made it difficult to understand the load. There was a rumour that downloads were changed to WNs at the end of August. Is that correct? Noticed that jobs at KIT were also creating high-file load on disk servers. Seems similar… DIRAC kills the jobs because it thinks jobs have stalled.

JT: What is the problem you want to solve with the VO ID card?

RS: There is an issue about throttling jobs. Perhaps we are running too many small jobs on the site and that causes problems. We have 20 disk servers each with 200 slots. Considering our transfers/channels etc. this would mean not more than 400 jobs running concurrently

JT: Seems to me that you will end up with a mess of configuration information that will need to be maintained. What you really want is a way to get the site to handle what you want to give them.

JG: LHCb were good at defining their data flows. Our guys thought the main flows are fine but the chaotic user analysis flows are not understood.

CMS (Ian Fisk)

You menion condor-G submission to CREAM. Are you expecting something from the developers?

Claudio: The people submitting by condor-G are submitting to the LCG-CEs. As far as I know the people in CMS are waiting for something new to come out of the testing that ATLAS is doing.

GM: DO you have any plans for deploying on this VM software system?

IF: Plan wouldl be to implement VM file system and use the CERN based fs to deploy the Tier-3s… and experiment with scale at that level.

FD: The way the information publishing should have been used was documented in the installed capacity document. The only implementors for the IPs was

CASTOR – information providers were provided by Jens Jensen. Job taken over by Jan Ivan.

MJ: Every information provider is software and there could be a bug. All things are well defined though so if CASTOR is not compliant then it needs to be fixed. Installed capacity should report was is deployed whether available or not (i.e. if disk servers offline then should be counted).

Stephen Burke: Depends not on site but implementation of the SE. For DPM there could be bugs. For dCache Paul Miller imeplemented it. Not sure if STORM existed when this was formulated but it is probably true that they are implanting. CASTOR did this but only for RAL – CERN and Taiwan doing different things. So overall should be about right.

JT: DPM makes specific assumptions about how your system is configured and if you implement differently then the number may not be correct.

SC: SO the officer will take care of this?

JG: Well still to define the officer but yes!

ATLAS report on summer issues (Stephance Jezequei)

Unavailability/failure greater than one day difficult to compare because there are different limitations. Need balance between high availability and level of data

GS: Would be interesting to see if the other sites have similar rates of failure – Andrew's figures were quite interesting. Can we get log information on disk failures from other sites.

Old systems are particularly problematic. Can not throw the kit out and buy new stuff so we have to stick with it. Sometimes it is bad hardware or firmware that you can not fix yourself.

JG: Other option is perhaps to do what CERN do and backup to tape.

TC: Not sure that happens now. Disk servers are mirrored not RAID5.

IB: Need to make individual failures invisible to the experiments. Not going to get much more reliable hardware in the near future.

RT: Dedicated hardware/storage. What happened was an unscheduled and unexpected reboot of the server. Do not know if this is a hardware issue. Primary focus was to get the backup.

JG: Sounds like the Oracle server problems we had about 6 months ago. Do not remember the underlying problem.

Concerned about ASGC

??: Recent instabilities are SRM and DB. The SRM was a misconfiguration of the CMS production role. ATLAS was affected. The second was the DB issue. Some CASTOR data should be put to tape but found some inconsistency in DB. On September 1st Eric ? involved.

Conclusiona/Summary (John Gordon)

1) Discuss with LHCb about their ID card or whether it isbetter to meet needs with direct site communication.
2) Sharing between Tier-1s/experiments about how to respond to disk server loss. What instructions do we follow. Produres to follow.
3) The issue of publishing of data and how we respond to that – including the CIO role.

On EVO:

Jeremy Coles

Alessandra Forti

Mario David

Jose Hernandez

Fabio Hernandez

Richard Gokieli

Ron Trompert

Massimo Sgaravatto

Denise Heagerty

Ian Collier

Josep Flix

Derek Ross

Mike Kenyon

EVO chat:

[08:43:42] **CERN 31-3-004** Hello Mario. I thought you were on holiday, Jon
[08:55:45] Alessandra Forti joined [08:57:46] **Mario David** that was until

yesterday [08:58:27] **Mario David** couldn't miss the stop of support of

glite31 [08:58:28] Jose Hernandez joined [09:02:44] Fabio Hernandez joined [09:02:59] Fabio Hernandez left [09:03:38] **Alessandra Forti** yes [09:03:39] **Mario David** loud and clear [09:04:53] Richard Gokieli joined [09:05:29] Ron Trompert joined [09:05:43] Massimo Sgaravatto joined [09:06:35] Denise Heagerty joined [09:07:26] **Mario David** yep [09:07:39] **Mario David** yes [09:14:05] **Massimo Sgaravatto** Condor provided a new version of Atlas fixing the problems reported by atlas for Condor --> CREAM submission [09:14:28] **Massimo Sgaravatto** Condor team provided a new version of condor fixing the problems reported by atlas for Condor --> CREAM submission [09:16:48] Ian Collier joined [09:18:05] **Massimo Sgaravatto** Would also like to understand which are the many ggus tickets assigned to cream [09:21:48] Josep Flix joined [09:27:04] Derek Ross joined [09:28:25] Mike Kenyon joined [09:31:37] **Alessandra Forti** please use the mic [09:39:11] **Ian Collier** John, did you get my sides? [09:58:05] Oxana Smirnova joined [09:58:38] **Derek Ross** yes [10:10:50] **Alessandra Forti** how will the releases be published in this model? [10:16:57] **Alessandra Forti** fuse should be part of SL5 [10:16:59] **Alessandra Forti** now [10:17:13] **Alessandra Forti** (5.5) [10:20:37] Ian Collier left [10:24:12] Stephen Burke joined [10:27:27] Ian Collier joined [10:28:04] luciano gaido joined [10:34:44] **Alessandra Forti** Manchester hasn't had any problem only a couple of bugs that were already in check. [10:35:17] Ruth POrdes joined [10:40:37] Pete Gronbech joined [10:40:53] Wahid Bhimji joined [10:42:29] Mario David left [10:44:12] Jose Hernandez left [10:46:19] Mario David joined [10:48:57] **Alessandra Forti** it was a vo requirement [10:53:53] **Massimo Sgaravatto** info providers, yaim [10:55:55] Fabio Hernandez left [10:58:42] **Mario David** what about the myproxy hydra and amga?? [11:08:18] **Mario David** if someone sees this? [11:08:30] **Mario David** what is the plan for myproxy?? [11:08:51] **Mario David** supose that hydra and amga doesn't matter much to wlcg [11:09:07] **Stephen Burke** amga started as an LCG (ARDA) project! [11:09:31] **Mario David** it's a PT in EMI, I think [11:09:38] **Stephen Burke** myproxy comes from VDT - I seem to remember that John White was going to package it for 3.2, but I haven't seen any activity [11:11:29] **Stephen Burke** The thing I don't understand is the batch system support - how can EMI provide middleware for CEs with no batch system interface?! [11:12:38] **Mario David** Gonçalo was asking about the glue2 for the SGE (in particular) [11:12:43] **Mario David** no answer up to now [11:14:00] **Stephen Burke** It isn't obvious to me who is responsible for the CE info providers for either glue 1 or glue 2 [11:15:28] **Mario David** Gonçalo and Co offered for the SGE, but he has to know oficial vars/values (I think) [11:16:48] **Stephen Burke** The batch system plugins presumably belong to whoever supports that batch system, but most of the code is generic [11:25:09] **Massimo Sgaravatto** stephen: the changes for glue-1 --> glue 2 I guess involves the info provider and yaim [11:25:17] **Massimo Sgaravatto** which are both part of lrms-utils [11:25:29] **Massimo Sgaravatto** torque-utils, lsf-utils, etc. [11:25:41] Wahid Bhimji left [11:26:01] Alessandra Forti left [11:28:09] Ian Collier left [11:28:32] Derek Ross left [11:29:04] luciano gaido left [11:35:10] **CERN 31-3-004** starting again t 1400 [11:50:49] Andres Aeschlimann joined [12:33:36]

Mike Kenyon left [12:59:17] Gang Qin joined [13:02:32] Ruth POrdes left [13:02:38] Ruth POrdes joined [13:04:47] Alessandra Forti joined [13:11:16] Claudio Grandi joined [13:18:24] Derek Ross joined [13:36:02] **Ron Trompert** Hi Roberto, the LHCb LFC at SARA is back. Cheers, Ron [13:42:30] **Massimo Sgaravatto** #61558 is a problem of the info provider I guess [13:46:20] Wahid Bhimji joined [14:37:52] Josep Flix left [14:47:49] **Wahid Bhimji** https://twiki.cern.ch/twiki/pub/LCG/WLCGCommonComputingReadinessChallenges/WLCG_GlueSchemaUsage-1.8.pdf [15:07:25] **Mario David** as the T" deploy storm 1.5, we will enable the checksum [15:31:21] Wahid Bhimji left [15:31:24] Massimo Sgaravatto left [15:31:25] Ron Trompert left [15:31:27] Claudio Grandi left [15:31:33] Mario David left [15:31:46] Stephen Burke left [15:32:32] Andres Aeschlimann left [15:32:58] **CERN 31-3-004** Meeting closed. See ou in October. [15:33:11] Derek Ross left [15:33:15] CERN 31-3-004 left [15:35:18] Richard Gokieli left