

GridPP Operations Meeting – Tuesday 5th November

Agenda: <https://indico.cern.ch/conferenceDisplay.py?confId=281944>

Attendees: See chat log.

LHCb

Raja not present. Report in absentia:

1. RAL is in scheduled downtime
2. Sheffield seems to have problems uploading data to LHCb SEs

CMS

Daniela - no problems this week. Problems with the Xrootd redirector last week, and some CVMFS issues, but nothing on the site side.

ATLAS

Elena - monthly availability report for October out (see <https://indico.cern.ch/materialDisplay.py?contribId=0&materialId=slides&confId=281944>) .

Availability was affected by the moves to SL6. For Cambridge, there was also a problem with setting the pilot factory. For Glasgow and QMUL the availability for SL5&SL6 queues was >90% - still a question for Manchester and Oxford, who need to explain their lower efficiency.

Jeremy noted that they should cross-check the availability between the Ops Availability and the ATLAS Availability and the reasons given for low availability. Elena clarified that the Availability she speaks of is *not* the ATLAS SAM test availability, but is instead derived from the fraction of successful ATLAS analysis jobs. Jeremy noted that there is still a correlation between the two types of availability.

ATLAS Software & Compute meeting: mostly concerning storage (Rucio move - sites asked to move to DPM 1.8.7 for xrootd and WebDAV support. There's also the new dq2 clients for sites and Tier3s. Changes and details on how spacetokens are full and sites are blacklisted.) and the move to ATLAS setup process. Will try to get Alisatair along to the next ops meeting to review these areas/requests. See the talk attached to the agenda: [CSmeetingOct2013.pdf](#)

Brian noted that he is looking at UK Spacetoken allocation and its use as part of a general review.

OTHER VOs

Chris - LFC at RAL is still outstanding. Backup VOMS servers still need to be configured by several sites (UCL, Durham, RALPP and Sussex). Chris has ticketed sites concerning this (and has also ticketed sites external to the UK who support t2k, and are therefore affected by this change). There was a suggestion to use the VO Nagios to monitor progress. Jeremy asked if there were any limitations to this and Kashif said he was waiting to upgrade the deployed version first (see chat)

David C noted that the backup VOMS issue is resolved at Glasgow.

--

Tier-1

RAL Downtime

Gareth S: top BDII, MyProxy, small VO's LFC is up, FTS services are up. Hoping to get CASTOR back later today, depending on timing of electrical work. Batch is down until tomorrow.

Jeremy noted that the UK Certificate Authority is also marked at risk.

-

Middleware Verification

Maarten Litmaath joined the meeting to talk through the WLCG ops coordination task force background and discuss any concerns/suggestions.

Opening remarks: Progress is currently at the location seen on the page. The name of the taskforce is still under discussion. Trying to attain a "catchy summary" of the purpose of the taskforce.

Maarten noted that Markus' talk, linked, is subtly different to the present thoughts on the matter.

(Talking through the presentation given at the July MB – first linked with name [WLCG-MB-MiddlewareValidationProcess.pdf](#))

Process started early last year, Markus tasked by the MB to think how we can continue after the EMI project ends in keeping infrastructure in a reasonable state (and possibly also with reduced EGI support for software).

We can still count on EMI a bit until Spring next year. "Typical" site in Europe uses EMI or UMD (usually EMI). We track baseline minimum releases for software. "EGI Staged Rollout does add value", but "does not have all services sufficiently covered. The experiments don't sign off on any of this."

Most Product Teams currently release into EPEL, when they're happy (via EPEL-testing, 2 weeks later -> EPEL-stable). There is also stuff that cannot be released to EPEL. EPEL is also a two edged sword as you can't stop things released into it.

The gLite -> EMI/EGI WN move was a big exercise that we learned from. We'd like to apply that experience to future processes. We think that we just need enough sites to cover the configuration space of sites (at least one of each SE, at least one of each experiment etc). The naive idea was that 5 to 10% of the infrastructure might be permanently on pre-prod (such as happens for WNs at CERN), for testing, grid-wide. "Various admins were quite worried about this increasing instabilities of sites, worsening availability figures."

New concept: to avoid this "big step", (that production services are only updated when absolutely necessary, which leads to chicken-and-egg situation of who updates first, and ad-hoc update process for the early adopters).

So, the idea is now is to round up all the activities that are already there (DPM, for example, has the DPM Collaboration which typically deploy updates in advance, in production, before the official release happens for everybody. This is part of their internal procedures. dCache works similarly.). Once we discover all these activities, we try to fill in the gaps or weak spots with our effort. There will be a defined baseline in the twiki but also ideally information on the latest version verified by the experiments.

We have a WLCG-stable repo that can be configured safely on any node type (won't break nodes by pulling unneeded packages).

FTS middleware testing in "pretty good shape". Sites and experiments involved in a pilot infrastructure to validate it. Note that this is not easily extensible to other services. Data management validation is trickier than Compute tool validation, as Data Management services are not ephemeral in their effects (a broken CE is working as soon as you replace the software, broken Data Management can have long term effects on historical data placed when it was broken).

Assertion: validation effort would be a small fraction of effort available on collaboration (and some of this effort is already being done). Overall, we are not too worried about getting sites to sign up (because many already are signed up to specific testing consortia). The bigger problem is getting the Experiments explicitly involved.

We need to do gap analysis to work out the places where we need effort & the experiments to sign off on things.

Timeline:

The idea is to try to get the new process working while we still have

something to fall back on (before EMI goes away, etc next year).

Questions from people:

Steve Jones: 2 issues: of the models we use, it's important that at every stage we settle on one model, and only one model. This implies that the people designing and planning our deployment models make their changes and policies clear at every stage. So there is no chance of any misinterpretation, these should also be transparently documented.

Maarten noted that the current ideas behind the middleware provisioning architecture rely, implicitly, on EGEE and EMI/EGI infrastructures taking care of most of the middleware. We would need to do something because we will no longer be able to rely on these other kinds of activities (other than the adhoc, poorly visible, ones). So we should definitely document things a bit. We should have one central repository of information [the twiki] that the sites and Experiments consult on a regular basis, as a matter of course. [We can also broadcast particularly important changes, etc]

Ewan: Why does the proposed flow use UMD as the preferred final step, rather than from EPEL-stable or WLCG-stable? De-facto, people tend not to use UMD at the moment.

Maarten: the thinking has evolved on this since the slide was written. As long as EGI is able to provide UMD, we wanted to be able to take advantage of it while we can. We won't force sites to pull from UMD.

Alessandra: We have used widespread test queues for many transitions recently (SL6 migration, EMI migrations etc), for ATLAS [versus a small number of test sites]. Since CMS is also moving to PaNDA, we could also set things up for them.

Update AF: *The concept of test queues is valid in some cases, but not all cases historically.* The reason this didn't work is because some experiments (Maarten mentioned also LHCb) are against permanent tests queues because they think they don't need them. He added the test queues are a valuable tool that we should use where available (atlas for example) and try to get other experiments to use them if conditions have changed.

Maarten noted that he didn't expect CMS to use PaNDA any time soon. He was not against this kind of process, where it is felt that the investment is not disproportionate to the benefits. The concept of test queues is valid in some cases, but not all cases historically. We may discover that this is easier to do now, but we also need to have alternatives. We should document these resources, so we can identify that we cover, for example, all the combinations of middleware that we want to test in the available test queues. Maarten noted that the working group can be, should be, shaped by all

the community involved.

Brian: "When it comes to the VOs interaction, from my perspective of the FTS and things like WebDAV for the t2s and ATLAS; the VOs are very good at getting involved when they can see there's an obvious benefit to them. How do we get the VOs more involved in the day-to-day validations where there are no obvious problems to fix for them. Also, for the FTS, it is deployed on a relatively small number of sites, so it is easier to test, compared to software that arrives on the Tier2s, for example [as there is wider variation in Tier2s]."

Maarten noted that it is an unfortunate reality that every site can look different to all the others, and the final test is always in production. We can approximate the ideal scenario, there are obvious things that ought to be tested [ie that WNs are compatible with all SEs]. The migration that we had last time, now we are safely in EMI-land, would probably be less of a site problem, as the Experiments tend to bring their own specific middleware. The ball therefore is also in the court of the Experiments to validate that their software releases (for example, of gfal) work on all the SEs they need to.

The Experiments need to foresee that a stable environment needs an investment in testing and validation of their own software and services. It needs regular attention. (We already have a decade of experience. We've never had a large scale outage of a majority of the infrastructure - and the services that could break with those consequences tend to be at T0 or T1, which also have the most experience and resources to fix issues.) Need to identify obvious gaps in the machinery that tests things.

Alessandra: Worry about the support - the DPM community always gets quoted as a good example, but the community is structured and backed by people who are funded centrally for support. This isn't the case for other Products.

Maarten: each product has its own particular worries and peculiarities. These worries, for products without central funding, are offset by the host institutions often having strong connections with WLCG. We were lucky, for the past decade, that many of the areas of testing and provision covered by European projects for us. Part of our working group is to evaluate how to provide support.

We could conclude that some products should be phased out on the basis of evidence collected. We need to avoid significant disruption of the existing production infrastructure, but in the long term we might move to different places.

Chris noted that the StoRM model also works like the DPM model (StoRM<->QMUL). Agile response to bugs is very helpful, in comparison to some other products.

Maarten re-emphasised that the aim is to avoid things sliding to a worse situation as we lose various established kinds of support.

Ewan emphasised that the key thing about the testing process is that we need to explicitly have it noted that when an early deployment site tests a release, then the developers (and experimetns) should engage in a detailed way with that site to provide bug fixes. (Sometimes what happens instead is less integrated.)

Next steps: For sites interested to join the taskforce. See more information on the twiki page:

<https://twiki.cern.ch/twiki/bin/view/LCG/MiddlewareReadinessVerification>.

Maarten: Things have been slowed by the huge number of meetings recently (CHEP etc). Hopefully we will get moving more solidly now we have the time to hash out details and send out emails – the list is setup. We expect to make some announcements before Christmas this year, but we do not expect a huge amount of work to be given to people.

The CHEP poster with the latest view on matters is

<https://indico.cern.ch/getFile.py/access?contribId=473&sessionId=9&resId=0&materialId=slides&confId=214784>.

-

Last minute tickets discussion:

Matt noted the ticket at Glasgow concerning CMS. Daniela will try to get stuff done from the CMS end.

RAL CVMFS ticket. Waiting for RAL to move Stratum 1 to 2.1.x releases.

Matt noted that CMS tickets work differently, and so setting them solved if you think you'd solved them, without waiting for responses from CMS ticketers themselves. (There was some discussion as to if this is actually what we're supposed to do anyway. The ticketer is supposed to Verify the Solved state but it is automatic after 2 weeks.)

--

Chat Log:

10:58:17] Elena Korolkova joined

[10:58:52] Rob Fay joined

[10:59:49] John Hill joined

[11:00:02] John Hill left

[11:00:12] Daniel Traynor joined
[11:00:23] Chris Brew joined
[11:01:45] John Hill joined
[11:02:09] Sam Skipsey joined
[11:02:28] Linda Cornwall joined
[11:02:53] Alessandra Forti joined
[11:03:39] Steve Jones joined
[11:04:02] RECORDING Sam joined
[11:04:10] Matt Doidge Those with the lowest scores start having to buy drinks at face to face meetings?
[11:04:15] Pete Gronbech joined
[11:04:31] Brian Davies joined
[11:04:40] Queen Mary, U London London, U.K. joined
[11:04:41] Andrew Washbrook joined
[11:04:56] H323 Queen Mary, U London London, U.K. joined
[11:05:16] Andrew Washbrook System Message to user Andrew Washbrook: SeeVogh has detected significant packet loss in your incoming video stream.
[11:05:16] Andrew Washbrook System Message to user Andrew Washbrook: SeeVogh has reduced the quality of your incoming video to optimize your performance.
[11:05:37] Jeremy Coles Sam and I will take minutes today.
[11:05:47] Robert Frank joined
[11:06:41] Ewan Mac Mahon joined
[11:06:56] SIP Call [Inbound call] joined
[11:07:06] Andrew Washbrook System Message to user Andrew Washbrook: SeeVogh has detected significant packet loss in your incoming video stream.
[11:07:06] Andrew Washbrook System Message to user Andrew Washbrook: SeeVogh has turned off your incoming videos due to poor network performance.
[11:07:31] Andrew Washbrook System Message to user Andrew Washbrook: SeeVogh has detected significant packet loss in your incoming video stream.
[11:07:33] Andrew Washbrook System Message to user Andrew Washbrook: SeeVogh application will connect to another server to see if it improves
[11:09:30] Duncan Rand joined
[11:10:08] Gareth Smith joined
[11:10:21] Govind Songara joined
[11:12:46] Alessandra Forti not all sites go to the atlas weekly meeting. I think it's good to repeat the main points here.
[11:13:12] Mark Slater joined
[11:13:25] Alessandra Forti btw arcproxy is now the default in the pilot.
[11:14:00] Chris Brew I'm working on it now! As we speak!
[11:14:12] SIP Call [Inbound call] joined
[11:14:33] Elena Korolkova I forgot to mention that RAL is in DT today and UK cloud is offline

[11:15:51] Ewan Mac Mahon Speaking of which, what's all this about mutli-cloud-iness?

[11:16:32] Mohammad kashif I will add new VO after updating VO nagios box

[11:17:14] Mohammad kashif New release came last week so I will update in next 2-3 weeks

[11:18:20] Christopher Walker joined

[11:18:35] Elena Korolkova this was discussed today in email exchange. AFAIUI after move to sl6 MAN and QM are not MCP sites ATM

[11:19:34] Jeremy Coles

[11:19:34] Jeremy Coles <http://indico.cern.ch/getFile.py/access?contribId=3&resId=4&materialId=slides&confId=217952>

[11:20:00] Alessandra Forti all sites that had multicloud setting lost them in the new queues

[11:20:47] Alessandra Forti it is possible to update the status with a curl command but not via AGIS

[11:21:25] John Bland joined

[11:21:36] Steve Jones where is the intro?

[11:21:40] Jeremy Coles

[11:21:40] Jeremy Coles <https://twiki.cern.ch/twiki/bin/view/LCG/MiddlewareReadinessVerification>

[11:21:45] Steve Jones cheers

[11:21:50] Elena Korolkova are there the records which sites were MCP sites and for which clouds?

[11:23:38] Steve Jones which presentation? What is tht title

[11:23:44] Jeremy Coles

[11:23:44] Jeremy Coles <http://indico.cern.ch/materialDisplay.py?contribId=3&materialId=slides&confId=217952>

[11:24:00] Steve Jones what is the title of the presentation, please?

[11:24:18] Jeremy Coles Hi Steve - I have put a direct link into the chat window twice above. It is

[11:24:18] Jeremy Coles <http://indico.cern.ch/materialDisplay.py?contribId=3&materialId=slides&confId=217952>

[11:24:38] Jeremy Coles The title is WLCG Mddleware Verification

[11:24:50] Ewan Mac Mahon That's a link to a list of three files, hence, I imagine, Steve's question.

[11:24:56] Steve Jones thnaks. There were two v. simialr ones.

[11:25:04] Jeremy Coles The first one. Sorry.

[11:25:16] Ewan Mac Mahon WLCG-MB-MiddlewareValidationProcess.pdf ?

[11:25:24] Jeremy Coles Yes WLCG-MB-MiddlewareValidationProcess.pdf

[11:25:27] Steve Jones I think so, Ewan.

[11:25:41] Jeremy Coles Each time I pasted it went to the talk....

[11:25:59] Linda Cornwall left

[11:27:28] Linda Cornwall joined

[11:30:09] Ewan Mac Mahon @Jeremy, ah, that is odd. I got a list of the right presentation in PDF format, and one with a different name in PDF and PPTX. Looks like they're all actually the same content though. Anyway, I think we're probably all on the same page now.

[11:30:24] Sam Skipsey They're actually subtly different.

[11:30:35] Sam Skipsey Some of the slides are slightly revised.

[11:35:39] Jeremy Coles Thanks Sam. Hopefully Maarten's verbal update will smooth any subtle differences.

[11:40:41] Christopher Walker System Message to user Christopher Walker: SeeVogh has detected significant packet loss in your incoming video stream.

[11:41:53] Christopher Walker System Message to user Christopher Walker: SeeVogh has reduced the quality of your incoming video to optimize your performance.

[11:42:53] Christopher Walker System Message to user Christopher Walker: SeeVogh has detected significant packet loss in your incoming video stream.

[11:42:53] Christopher Walker System Message to user Christopher Walker: SeeVogh has turned off your incoming videos due to poor network performance.

[11:43:18] Brian Davies big difference is also the number of expected deployment sites of a middleware. FTS is expected to be deployed at a very few number of sites. therefore site dependent issues are very small . c.f WN deployment which will be at 100s of sites

[11:44:48] Christopher Walker System Message to user Christopher Walker: SeeVogh has detected significant packet loss in your incoming video stream.

[11:44:51] Christopher Walker System Message to user Christopher Walker: SeeVogh application will connect to another server to see if it improves

[11:51:40] Ewan Mac Mahon I have a Q about slide 10 - why are sites expected to pull from UMD, rather than from EPEL-stable and WLCG-stable? I'm not seeing the value in the extra step.

[11:52:29] Christopher Walker Executive summary: rollout as it is at present, but shine a bit more light on it - and scrap the idea of sites in permanent preprod state.

[11:53:02] Christopher Walker Oh - and more participation of the experiments.

[11:53:24] Alessandra Forti we have permanent testbeds

[11:53:43] Alessandra Forti which IMO are a better solution than only 6 sites

[11:55:11] Steve Jones Yes - not too verbose, but complete and rigorous.

[11:56:55] Christopher Walker System Message to user Christopher Walker: SeeVogh has detected significant packet loss in your incoming video stream.

[11:59:05] Sam Skipsey Slide 10 has ghostly dashed lines from EPEL-Stable and WLCG-stable in one version of the presentation, but not the other.

[11:59:21] Sam Skipsey So, I think this is partly a different revisions subtlety, Ewan.

[11:59:36] B have comments about VO interaction

[12:02:58] Steve Jones

[12:05:33] Jeremy Coles This looks like the CHEP poster:
<https://indico.cern.ch/getFile.py/access?contribId=473&sessionId=9&resId=0&materialId=slides&confId=214784>

[12:14:36] Jeremy Coles Does anyone have any pressing questions from what they have heard or can see in the talk or the CHEP poster?

[12:15:30] Jeremy Coles Are there other concerns to be mentioned before moving on?

[12:15:41] Christopher Walker A comment: The dpm example of developers working closely with sites and sending pre release version has also worked well for me with StoRM. It is much more rewarding to see bugs fixed quickly.

[12:15:57] Jeremy Coles Does anyone have any suggestions for improvements to the proposed approach?

[12:16:21] Christopher Walker Also, the UMD3/SL6 not being tested by ATLAS prior to the window for moving was an oversight.

[12:16:37] Ewan Mac Mahon Ditch the whole thing and just run cloud interfaces? No middleware, no hassle.

[12:16:55] Christopher Walker That's something to make sure we avoid in future

[12:17:54] Andrew Washbrook left

[12:20:30] Ewan Mac Mahon We've certainly found some of the direct engagement with DPM devs to be very useful, too.

[12:20:45] Jeremy Coles The task force twiki page:
<https://twiki.cern.ch/twiki/bin/view/LCG/MiddlewareReadinessVerification>

[12:26:55] SIP Call [Inbound call] left

[12:27:31] Matt Doidge https://ggus.eu/ws/ticket_info.php?ticket=98253

[12:28:40] Linda Cornwall left

[12:30:51] Ian Loader left

[12:30:54] Mark Slater left

[12:30:54] Robert Frank left

[12:30:54] John Hill left

[12:30:56] Alessandra Forti left

[12:30:57] John Bland left

[12:30:58] Elena Korolkova left

[12:30:58] Gareth Smith left

[12:30:59] Brian Davies left

[12:31:01] Chris Brew left

[12:31:01] Ewan Mac Mahon left

[12:31:02] Daniel Traynor left

[12:31:02] Matt Doidge left

[12:31:07] Rob Fay left

[12:31:09] Steve Jones left

[12:31:28] Daniela Bauer left

[12:31:36] Duncan Rand left

[12:31:37] Govind Songara left

[12:31:47] Mohammad kashif left

[12:32:08] Pete Gronbech left

[12:32:10] David Crooks left

[12:32:21] Q