

GridPP Ops 25/10/22

Attending: Matt, Sam S, Gerard H, Vip D, Winnie, Dan T, Steven S, Gordon, JamesW, Alexander R, Jose, Mike L, Linda, Luke, Dan W, Tom D, Tom B, Wenlong, Jyoti, Daniela, Darren M, Raul L, Rob C, Dwayne, Alessandra

Apologies: PeteC, Duncan, DavidC, Emanuele, Mark

Urgent Business

NTR

Actions from previous meetings.

None

VO Updates

CMS (Daniela):

[monit link](#)

Tokens on ARC:

CMS (prompted by Raul) has just noticed Argus doesn't really support tokens:

https://ggus.eu/index.php?mode=ticket_info&ticket_id=158175 (Brunel ticket)

https://ggus.eu/index.php?mode=ticket_info&ticket_id=159270 (Argus)

CMS now notes that Argus is "in maintenance mode only". (So Imperial is just ahead of the curve having it abandoned 5 years ago ;-)

RALPP's token ticket looks abandoned:

https://ggus.eu/index.php?mode=ticket_info&ticket_id=158175

The Toms weighed in, Tom B testing with tokens at the Tier 1. Dan T notes in chat about ARC7 for tokens. Getting token auth working for ARC6 bypasses argus anyway.

Daniela asks for working token examples. Gordon has some of these.

Tom B notes ARC7 won't natively support argus integration.

Tom B bravely shows his configs live. In short need to put token stuff first in arc.

Request to put info on a wiki page. Need to keep an eye on this.

ATLAS (James):

- RAL:

- Largely recovered from Network problems of last week. Echo (kept going), Antares, back ok; Compute: now running above pledge
- T2:
 - A few continuing HC tests flapping for ECDF_CLOUD,
 - BHAM failing most jobs, also in/out of test

LHCb: (Mark/Alexander):

Introducing Alexander, the new VO rep.

- RAL:

- Recovered from major network outage. Some problems with transfers may still be present, probably due to long lasting issues.
- The following tickets are opened for ral:
 - Slow stats https://ggus.eu/index.php?mode=ticket_info&ticket_id=159137 . The problem is being debugged
 - Request to enable https as a write protocol for echo https://ggus.eu/index.php?mode=ticket_info&ticket_id=159146, seems to be solved, waiting for confirmation
 - Vector read problem https://ggus.eu/index.php?mode=ticket_info&ticket_id=142350 . Problem seems to be identified, waiting for the code patches.

-T2:

- pilots are failing at Royal Holloway College, ticket opened https://ggus.eu/index.php?mode=ticket_info&ticket_id=159303

“Other” VOs:

DUNE (Wenlong):

Links: [glideInWMS configuration for the different sites](#)

[SAM tests](#)

[Monit link for history](#)

CRIC: <https://dune-cric.cern.ch/core/experimentsite/list/>

[Storage monitoring site](#)

[Job efficiency monitoring site](#)

[DUNE job status](#)

- Workflow tests: <https://wfs.dune.hep.ac.uk/dashboard/?method=wtf-results>

Winnie notes that the long retired lcgce01.phy.bris.ac.uk is still in the [dune etf tests](#), and still in the [DUNE CRIC feed](#) - Winnie asks if this can be fixed?

-Wenlong will pass along to Andrew and Raja

SKA:

NTR

LSST/Rubin:

NTR

EUCLID:

["eucliduk.net"](#) ["voms.gridpp.ac.uk"](#) "15518"

["/C=UK/O=eScience/OU=Manchester/L=HEP/CN=voms.gridpp.ac.uk"](#) ["eucliduk.net"](#)
and the usual 3 GridPP voms servers.

The VO configuration is also available in [/cvmfs/grid.cern.ch](#)

Some discussion on storage.

Other VOs:

General Updates/Discussion

NTR

Meeting Updates

WLCG Ops Coordination last Thursday:

<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsMinutes221020>

The aforementioned concerns for Argus came up here.

HEPiX in Sweden, 31st October <https://indico.cern.ch/event/1200682/>

Registration will likely close soon (21st of October by the looks of it). Registration is required for remote attendance.

Future meetings:

SWIFT-HEP Manchester, 1-2 November <https://indico.cern.ch/event/1184802/>

WLCG meeting in Lancaster in November (Please do take advantage of this being in the UK if you want to attend!):

<https://indico.cern.ch/event/1162261/>

With the rucio workshop in the same week and location:

<https://indico.cern.ch/event/1185600/>

Reminder: CHEP 2023: <https://www.jlab.org/conference/CHEP2023>

Abstract submission closes on the 17th of November.

Computing Insight UK 2022 in Manchester in December, will be discussion on Energy Benchmarking on Heterogeneous Systems. <https://www.scd.stfc.ac.uk/Pages/CIUK2022.aspx>
Also Lustre UK user group meeting <https://www.lustrefilesystem.com/> on the Wednesday.

Tier 1 Status

On Monday 17th October 2022 at just after 15:00 local time (14:00 UTC), there was a network incident at [RAL](#)-LCG2. For context, over the last 2 years [RAL](#) has been replacing its network. The failure occurred on the legacy network and while this meant that many host were unaffected, there is enough hardware still on the legacy network to cause significant disruptions to services.

At the time of writing, we are still in the process of rebuilding parts of the network. While we have some understanding of what happened, we haven't done a full analysis yet to understand the cause. This will be done as part of the serious incident review. What we do know is that there were multiple failures at approximately the same time (Note that we have considered the possibility of a malicious actor and have found no evidence that this is the case):

- One of the two core switches for the legacy Tier-1 network failed (it lost its operating system).
- The switch that the legacy database servers were connected to failed (hardware failure). Our RT ticket system as well as the FTS and GOCDDB databases were attached to this switch.
- One of the switches with 25 Echo storage nodes crashed.
- We observed a large numbers of spanning tree updates on the management network (which is used by both the legacy and new Tier-1 networks).

Due to problems with the management network, which provides both monitoring and remote access to all servers, it was difficult to assess the initial magnitude of the problem. The initial recovery focused on getting the management network working, however by the end of the working day the management network had not been recovered and the decision was made to keep it offline as we weren't sure if it may have been the cause of the outage. On Monday evening, the switch that the Echo storage nodes were attached to was recovered and Echo started working again. The loss of one of the core routers halved our internal network connectivity and meant that while Echo functioned many transfers were slow / failing as the remaining network was being heavily used to rebalance Echo.

The GOCDB at [RAL](https://gocdb.hartree.stfc.ac.uk) went down due to the outage. We have a read only mirror off-site (<https://gocdb.hartree.stfc.ac.uk>) which was unaffected, however the DNS service which would allow us to point goc.egi.eu to that server had also been taken down. Due to the confusion the outage was causing we announced that the GOCDB was unavailable even though direct access to the mirror was working. As part of the serious incident review, we will be looking at ensuring that the handling of the GOCDB outage is improved.

On Tuesday 18th October we continued to work on the network and replaced the failed switch connecting the legacy databases. To help with the debugging, we didn't want to put any additional load on the network so while services were brought back up we didn't start production work on them. By the end of the day it was decided to rebuild the management network to segment it (which had been a long term planned) as we were unable to recover it. We also decided to start bringing more services back online:

- Antares, our new tape endpoint had been entirely built on the new network and should have been unaffected. Unfortunately, a single VM, the CTA frontend server which communicates between the EOS frontend and CTA backend was running on the legacy networking. Once this machine was brought back Antares started working again. We will be moving the CTA frontend host to the new network as soon as we can.

- The GOCDB was available in read/write mode towards the end of the day and at that point we declared an unscheduled outage (https://goc.egi.eu/portal/index.php?Page_Type=Downtime&id=33056) to the batch farm and FTS to end on Wednesday 19th at midday local time (11:00 UTC).

On Wednesday 19th October we started to bring the batch farm online. As we brought the farm online, we started to see network saturation again and transfer failures started to increase. We still haven't recovered the legacy core network switch (and at time of writing still aren't sure if this is possible). We therefore decided to double any saturated links to the working core switch. This started to alleviate problems although we have found that we can

only run the batch farm at about 60% occupancy. This just about meets our pledge to the LHC experiments.

For the remainder of the week, we continued to rebuild the management network and re-establish stability. For this week we will finish rebuilding the management network and will be reviewing the status of the legacy core network switches. It might be possible to bring the failed one back however we believe that the best course of action in the medium to long term will be to migrate hardware to the new network. We may schedule a downtime in the next month or so to reconnect hardware as that would be much faster than trying to do it transparently while services are in production.

Tier-1 Network upgrade (latest update on top)

Pete asked about LHCONE?

-There is a plan. Alastair sent an email.

13-09-22 - Tom B notes a problem with host firewalls on LHCONE nodes.

09-09-22 - IPv6 Traffic now enabled on the super spine.

02-08-22 - intention to make 2022 WNs (rocky 8 gen) hope to put these on lhcone.

17-05-22 Technically joined the LHCONE but at the moment it has only one host (the new perfSonar box). Currently we have resource checking and conforming routing between endpoints and IPV6 allocations.

05-04-2022 Resolution of new storage node access and central security work delayed LHCONE activity. Being re-timetabled.

Security Brief

New home for advisories: <https://advisories.egi.eu>

- Operational update [standing item]
- Pakiti status - should be working
- Site Security Survey
 - (good place to note things like Carbon Black@Bristol)

AAI

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

In particular, if you would like to join the discussion the egroup is "wlcg-resource-trust-evolution"

Auth Working Group Twiki: <https://twiki.cern.ch/twiki/bin/view/LCG/WLCAuthorizationWG>

Indico Category: <https://indico.cern.ch/category/68/>

Mailing list egroup: project-lcg-authz

Note - AuthZ meetings moved to fortnightly Friday's 14-15:00 UK.

Meeting this Friday. The recordings and minutes for the preGDB workshop are being uploaded.

Networking News

[UK Mesh](#)

[Check_MK monitoring](#)

https://www.gridpp.ac.uk/wiki/Perfsonar_refresh

[Dune Mesh](#)

RALPP looking better.

Storage and Data Management News

Minutes: [GridPP Storage](#)

Reminder: Storage Meeting is now @ 1030 [not 10am] to make it slightly easier for everyone to attend.

Xrootd 5.5.1 is now out.

DPM "retirement" plans becoming increasingly important. (Matt and I have been chatting to some of the sites with bigger DPM instances.)

Technical Update

(to include any Token News, Data Challenge work, "New" OS experiences)

Standing subjects:

-ESCAPE credentials and use of tokens

-"Carbon Black" @ Bristol and similar technologies elsewhere

-HEPSCORE experiences.

– call last week for a "hackathonised" HEPSCORE tech meeting.

Duty Report

Dashboard fixed. Nothing too exciting.

Tickets

[UK GGUS tickets](#)

49 Open tickets - back to 1 page!

https://ggus.eu/index.php?mode=ticket_info&ticket_id=159182 - this RAL test ticket has probably served its purpose?

Site News

NTR

AOB/Feedback for/from the PMB

Quarterly report reminder.

-please take the time to double check site figures.

Meetings the next two weeks? I think we can go ahead next week, but will want to organise backup chairs.

Actions/Decisions from This Meeting

No hard actions. argus concerns will go into a standing item.

Chat Window:

11:06:01 From Daniel Traynor To Everyone:

I guess I use arc7 token only? other VOs (non LHC) will have to update

11:08:10 From Daniela Bauer To Everyone:

Well, we are working on basic token support for the GridPP DIRAC instance, but it's slow. But maybe everyone could start with CMS and see how it goes ?

11:08:10 From Luke Kreczko To Everyone:

So do we phase out ARGUS (since it is in maintenance mode only)?

11:08:35 From Raul Lopes To Everyone:

I have configured one CE without Argus, but token authentication with CMS is failing

11:08:36 From Daniela Bauer To Everyone:

We stopped using Argus in 2017 and never looked back.

11:09:01 From Luke Kreczko To Everyone:
That's sensible. The fewer services needed for a site, the better

11:09:33 From Daniel Traynor To Everyone:
<https://wlcg-authz-wg.github.io/wlcg-authz-docs/token-based-authorization/>

11:10:47 From Daniela Bauer To Everyone:
@Dan: Dowsn't help with the ARC/ARGUS issue though ? AT least I don't see it.

11:10:48 From Daniel Traynor To Everyone:
<https://wlcg-authz-wg.github.io/wlcg-authz-docs/token-based-authorization/doma-testbed/>

11:11:36 From Daniel Traynor To Everyone:
can have both (argus and token(single)) enabled at the same time

11:12:23 From Daniela Bauer To Everyone:
@Thomas: Could you dump this on the gridpp Wiki by any chance ?

11:13:11 From Daniela Bauer To Everyone:
I object to the Wiki nothaving any love, the DIRAC users page is lovingly maintained !!! ;-)

11:15:58 From Jose Caballero Bejar - STFC UKRI To Everyone:
Could you remind me the link to that wiki page?

11:16:22 From Daniel Traynor To Everyone:
map_with_plugin=all 30 /usr/libexec/arc/arc-lcmaps %D %P liblcmaps.so /usr/lib64
/etc/lcmaps/lcmaps.db arc

11:16:37 From Daniela Bauer To Everyone:
@Jose: The DIRAC one?

11:16:56 From Jose Caballero Bejar - STFC UKRI To Everyone:
The one where Tom can put the snippet configuration from RAL

11:17:43 From Daniela Bauer To Everyone:
HE should have access to the gridpp Wiki, just make a page and link it from the main page,
so it; gets picked up by google ? If he can't link it, I think the wiki is still running out of
Manchester

11:17:56 From Jose Caballero Bejar - STFC UKRI To Everyone:
ok

11:18:20 From Daniela Bauer To Everyone:
<https://www.gridpp.ac.uk/wiki/mynewpagehere>

11:18:42 From Daniela Bauer To Everyone:
if you wanted to try it live (while we are doing 'live demos ;-)

11:19:12 From Thomas Birkett - STFC UKRI To Everyone:
The documentation found here:
<http://www.nordugrid.org/arc/arc6/admins/reference.html#map-with-plugin> defines that value
being sequenced, 'This is sequenced option.'

11:19:26 From Thomas Birkett - STFC UKRI To Everyone:
Try:

11:19:35 From Thomas Birkett - STFC UKRI To Everyone:
map_with_plugin=all arc 30 /usr/libexec/arc/arc-lcmaps %D %P liblcmaps.so /usr/lib64
/etc/lcmaps/lcmaps.db

11:19:45 From Thomas Birkett - STFC UKRI To Everyone:

Sorry:

11:19:53 From Thomas Birkett - STFC UKRI To Everyone:

map_with_plugin=arc all 30 /usr/libexec/arc/arc-lcmaps %D %P liblcmaps.so /usr/lib64
/etc/lcmaps/lcmaps.db

11:20:59 From Thomas Birkett - STFC UKRI To Everyone:

Some variation will give auth priority depending on the sequence, when Ive tested I can
report back with my findings

11:34:48 From Thomas Dack - STFC UKRI To Everyone:

<https://indico.cern.ch/event/1185598/>

11:39:58 From Alessandra Forti To Everyone:

on it