

GridPP Ops 22/11/22

Attending: Sam S, Jose, Gerard H, Raul L, Steven S, Winnie L, Alex R, Dan W, Brian D, James W, Vip D, Tim Chown, Rob F, Jyoti B, Thomas B, Gordon, Daniela B, Duncan R, Emanuele, Linda, Wenlong, Dan T, Luke K, Ian L, Tom D, Pete G,

Apologies: Matt, DavidC, DarrenM

Urgent Business

None

Actions from previous meetings.

None

VO Updates

CMS (Daniela):

Reasonably quiet.

Chris has been working on the AREX ticket:

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

Should we involve the ARC developers ?

Tom B has been talking to the ARC devs - who are generally focussed on ARC7 nowadays. Note that ARC6 does have hard-coded checks for x509 proxies in the submission process [which you have to manually remove if you're using tokens].

Tom will follow up with Chris Brew to see what's happening.

Matching RAL-LCG2 ticket also unresolved (though this seems to have reached the token step):

https://ggus.eu/index.php?mode=ticket_info&ticket_id=158161

[monit link](#)

ATLAS (James):

- Brief mass HC blacklisting Monday PM, due to rucio / config issue;
(following a previous HC blacklisting the previous Monday)
- RAL
 - Antares; Problem with degraded transfer efficiency from Cern to RAL (direct root transfer); under investigation
 - "Zombie" containers on batch farm last week; fixes on Friday working well.
 - ATLAS above pledge for the week (due to reduced LHCb prod.)

- T2
 - ECDF; LDAP issue resolved?
 - QMUL: reduce numbers of running jobs (and spikes in job failures); networking issues (planning on work to enable 100Gbit/s has been suspended due to the college link issue causing this)
 - Limited Sim jobs, affecting Glasgow and Lancaster last week (and noted on other clouds).

LHCb: (Mark/Alexander):

All sites:

- Throughout the week the number of running jobs was small, due to insufficient MC production requests
 - A lot of productions were accepted on Monday evening, right now everything looks OK

RAL:

- Farm restart on Friday, running jobs were lost.
- Long lasting issues in progress:
 - [Slow stats](#)
 - Stats may have nothing to do with the problem - may be slow checksums in fact?
 - [Deletion problems](#)
 - [Vector read problems](#)

T2:

- A lot of failed pilots at Brunel University
 - Problems with WN wn-b6-11-00.brunel.ac.uk ?
- Pilot submission fails at University of Edinburgh
 - Recent downtime, not recovered yet? [see LDAP issues mentioned in ATLAS update]
- [Migration to ceph](#) in progress at University of Glasgow

Dan T asked how much data T2Ds need to provide for LHCb now? 600TB? 300TB?

“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>
- Investigating `Rucio upload` fails to some RSEs, e.g. MANCHESTER, QMUL.
Suspecting some RSEs really needs // instead of / in the Rucio config for xrootd

SKA:

LSST/Rubin:

Other VOs:

General Updates/Discussion

Meeting Updates

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

Abstract submission extended to the 28th of November.

Note from Sam: can everyone planning on (or already) submitting an abstract to CHEP please let me know. (We'd like numbers so we can look at the travel budget etc)

There's a lot of abstracts submitted already before the first deadline. I assume that more people will be submitting abstracts for the second one - please still let me know.

Computing Insight UK 2022 in Manchester next week, there will be discussion on Energy Benchmarking on Heterogeneous Systems. <https://www.scd.stfc.ac.uk/Pages/CIUK2022.aspx> (Dan T is attending.)

Technical Meeting on Token Transition this Friday.

Tier 1 Status

Antares

On 16th November, CERN upgraded their FTS instances and introduced an intermittent bug. This caused many transfers going to tape to be stuck in the archiving state. The FTS file ID is converted from a uint64_t to a int32 (the bug) and if the resulting number is negative then FTS does not start to poll for the file's archive state. We received a ticket from ATLAS (https://ggus.eu/index.php?mode=ticket_info&ticket_id=159550) on Friday evening and separately Katy spotted this for CMS and talked to the FTS developers who were able to fix the bug.

Echo

We have a 4 hour downtime scheduled for 6th December for security patching and a reboot.

There was a failure of one of the gateways on the 14th November which caused a few hours of failures (to that node) before it was removed from production, this caused test failures.

There was a bug in the latest version of XRootD which could cause a multi-stream transfer to crash an Echo gateway. This was patched and pushed out to all external nodes.

LHCb deletion errors (https://ggus.eu/index.php?mode=ticket_info&ticket_id=159395). Occasional spikes of errors observed, overall failure rate is <2%. Alex is waiting for debugging to be turned up on a node to try and catch failures as they are occurring.

LHCb slow stats (https://ggus.eu/index.php?mode=ticket_info&ticket_id=159137). When LHCb are about to run certain types of job campaigns, they run a bulk stat of files at sites first. Occasionally this appears to run slowly. Alex is setting up further tests to see if he can replicate the problem.

LHCb vector reads (https://ggus.eu/index.php?mode=ticket_info&ticket_id=142350). XRootD version continues to be kept up to date. Debugging of the code is ongoing.

Currently the external gateways all have the same configuration. There is a plan to make the XrootD and WebDav aliases both point to the same complete set of nodes. This should spread load. Darren to follow up that basic sanity checks have been done.

Batch Farm

There was a critical security alert last week. Once this was patched, the VO noticed that the number of jobs they were running on the farm was decreasing even though the farm appeared full. We also noticed that the batch farm was reporting very low job efficiency even though the VOs were not observing a reduction in efficiency. The problem was tracked down to an issue between docker and the patched Kernel which wasn't able to effectively kill completed jobs. Therefore, jobs from HTCondor's point of view were taking up the farm and not doing anything while in reality, the job had completed. A further patching and reboot campaign was done at the end of last week.

LHCb ran out of jobs across the Grid and have mostly drained out of the RAL farm. ATLAS were using 33k job slots (now ~18k).

FTS test instance upgraded to 3.12, production will be done shortly (may postpone due to issue

Tier-1 Network upgrade (latest update on top)

22/11/22 - hopefully some updates this week. Central core hardware routers *have* been replaced.

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.eji.eu>

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

<https://www.jisc.ac.uk/get-involved/cyber-security-community-group>

Please respond to the current CVEs as prompted. EGI CSIRT may chase sites.

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/WLCGAuthorizationWG>

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.

Auth working group postmortem for ATLAS using IAM on the 18th - issues due to database timeouts <https://indico.cern.ch/event/1219515/>

Networking News

[UK Mesh](#)

[Check_MK monitoring](#)

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

[Dune Mesh](#)

perfSONAR certificate status: https://www.gridpp.ac.uk/wiki/Perfsonar_Cert_Status

Assessment by Chris Walker:

“I thought I'd have a look at whether it's just QMUL that has certificate issues. It's not - there are various certificate problems across the UK. Here's the results of a quick look.

Jisc-London OK

Jisc-Slough OK

RAL-LCG2 - t1-ps01-bw.gridpp.rl.ac.uk is down (or not responding) - due to ARP issues

UKI-LT2-IC-HEP 404 Not Found (ic.ac.uk) [this is because the server does not have a web interface enabled - but does have a cert!]

UKI-LT2-QMUL perfSONAR Toolkit | perfsonar-latency.esc.qmul.ac.uk Not secure

UKI-LT2-RHUL OK

UKI-LT2-UCL-HEP NET::ERR_CERT_AUTHORITY_INVALID (cert expired Sep 11, 2020)
UKI-Northgrid-Lancs-HEP NET::ERR_CERT_AUTHORITY_INVALID (certificate expired 329 days ago)
UKI-Northgrid-Liv-HEP OK
UKI-Northgrid-shef-HEP OK
UKI-scotgrid-Dur-hep OK
UKI-scotgrid-ecdf expired 62 days ago
Glasgow - OK
BHAM - Not trusted
Brisol - Not trusted
Oxford - OK
RALPP - doesn't respond (blocked by firewall?)
Sussex - This server couldn't prove that it's grid-perfsonar.hpc.susx.ac.uk; its security certificate is from *.hpc.sussex.ac.uk.

There's thus a mix of problems across lots of sites: Some certificates have expired. Some aren't trusted (self signed, or signed by the grid CA).

Sussex, the domain is incorrect - presumably the university has changed from susx to sussex - I guess someone ought to update the perfsonar grid (at which point we'll note the certificate has expired)."

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.

A note that ceph configuration is very network config sensitive, and in particular the management nodes(MON, MGR, MDS) need to know about subnet divisions.

Bristol issue... might actually be the Bristol firewall doing DDoS protection.

http://ggus.eu/index.php?mode=ticket_info&ticket_id=156902

<https://github.com/xrootd/xrootd/issues/1736>

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

-ARGUS concerns.

We put in a request for a Tech Meeting on this. David C approved.

Duty Report

The Ops Portal's timing out today.

Tickets

[UK GGUS tickets](#)

50 Open Tickets this week.

https://ggus.eu/index.php?mode=ticket_info&ticket_id=159506 Is this Tier 1 test ticket done with?

https://ggus.eu/index.php?mode=ticket_info&ticket_id=159494 This QMUL cms ticket looks like it might have been missed.

https://ggus.eu/index.php?mode=ticket_info&ticket_id=159380 Similar with this ECDF lhcb ticket.

Site News

AOB/Feedback for/from the PMB

Actions/Decisions from This Meeting

Soft actions (from last week)- tell Sam if any power management/infrastructure issues also tell Sam if you plan to go to CHEP.

Chat Window: