

GridPP Ops 23/4/24

Attending: Sam, Daniela, Rob F, Dan T, Darren M, Duncan R, Alex R, Brij K, Mark S, Vip, Thomas B, Brian D, Mike L, Emanuele, Dan W, Gerard H, Steven S, Ian L, Dwayne S, Wenlong Y, Gordon, Bruno

Apologies: Matt, David C, Linda,

Urgent Business

None

Actions from previous meetings.

270224-01 - Matt, others. Documentation “spring clean”/purge/move to confluence
Matt still to survey and flag outdated docs on the wiki.

Alastair has given Darren responsibility for steering this (so please let Darren know - and he will follow up with responsibilities).

270224-02 - Sam, Matt - liaise with Alastair over future tech meeting subjects.

-need to try to give more than a few days' notice for future meetings.

Next [™] is Harwell Open Week focused

General Updates/Discussion

EL7 EOL

Less than 10 weeks to go.

- Tom B notes that this is an opportunity to consider containerising important services to make migration easier in future.

Sussex Help

Thanks to those that responded off list.

QMUL certificates

Any luck getting in touch with the helpdesk?

Dan has certificates with SANs, which need CLI renewals. Dan renewed with the web interface previously - which stripped the SANs - preventing him from renewing it with the CLI subsequently. (Solution, revoke the web-renewed one, and then request from scratch).

But no responses from CA helpdesk to any of this, or in TB-SUPPORT.

HEPOSlib

Vip discovered telnet was still in the heposlib dependencies.

Reply from: Andrea Valassi

Hi Vip, thanks for the message.

This was added after a request from someone in ATLAS, <https://its.cern.ch/jira/browse/SPI-2171> and https://gitlab.cern.ch/linuxsupport/rpms/HEP_OSlibs/-/issues/6

That said, my impression is that this was for some 'personal' usage rather than experiment productions.

The motivation was "The usage is for worker nodes to connect back to login nodes to facilitate Dask-like job orchestration".

Do you have a major concern with this (e.g. for security reasons?). If this is the case I can cross check back with him and a few other people about removing it, maybe.

I am always somewhat reluctant to the idea of removing packages, but if there is a good reason we can discuss that.

Thanks

Andrea

VO Updates

ATLAS (Brij/Jyoti):

UK-ATLAS-Cloud real time status links

- [Sites status board](#)
- [weekly status](#)
- [Open ggus tickets](#)

- [Last week's snapshot](#)

RAL T1

- RAL ATLAS Tier1(RAL-LCG2) Availability/Reliability for the past week - **100%**
- CPU slot occupancy average vs pledge during last week - **~65%**
 - Batch farm draining due to firmware update
- Walltime job efficiency **~97%**
- Data Transfer efficiency **~97%**
- RAL-LCG2_DATATAPE(TB) Total Capacity - 27077.78, Occupancy - 22309.39
 - Occupancy at **~82%** [Link](#)
- RAL-LCG2_MCTAPE(TB) Total Capacity - 44309.22, Occupancy - 36167.73
 - Occupancy at **~82%** [Link](#)
- RAL-LCG2-ECHO_DATADISK(TB) Total Capacity-25494.00, Occupancy - 25175.26
 - Occupancy at **~98%** [Link](#)
- No open tickets

T2s

- Open ggus tickets - **in progress**
 - Error in [test trigGPU mc v1Dev grid](#)
 - [UKI-LT2-QMUL failing all transfers and deletions due to certificate problem](#)
 - [UKI-SCOTGRID-GLASGOW transfers errors](#)
 - [UKI-SCOTGRID-GLASGOW_ARM bad cvmfs](#)
 - [UKI-LT2-QMUL: Enable token support for storage](#)
- Open JIRA ticket **in progress**
 - [Upcoming QMUL Data Centre Refurbishment](#)
 - [Lost data @DURHAM – rebuilding storage elements](#)

General news

- [Notes from ADC operations daily](#)
 - Significant tape writes to RAL from rpg backlog [campaign](#),
 - ~24PB being distributed among T1s as per the pledge and occupancy.

Mark update on Xcache stuff - Ilija responded, and we've investigated - the CA certs for the XCache servers were out of date, which didn't break anything from the UK [where we

normally get files] but did for other servers [where Ilija's tests run and he uses to blacklist XCaches in VP]

CMS (Daniela):

Nothing to report.

Imperial has odd intermittent problems with its federation tests:

https://ggus.eu/index.php?mode=ticket_info&ticket_id=166165

We've asked dcache which have so far ignored the issue and it doesn't happen often enough to enable debugging at speed, but we are keeping an eye on it.

LHCb (Alexander):

RAL T1:

- A few WN gens are draining again this week, for IPv6 setup
- Vector read optimisation (follow-up on GGUS [ticket](#))
 - No updates

T2:

- [Failed](#) pilots at Glasgow
 - Seems to be a problematic WN
- Jobs are [failing](#) at Sheffield
 - Further limit decrease seems to be needed, since ~30% of jobs are still failing
 - Debug in progress
- ETF tests are [failing](#) for Lancaster
 - Still waiting for migration
- Preprod ETF tests are [failing](#) for Glasgow
 - Looks like a permission problem
- Token-based job submission is [being set up](#) for RALPP
 - Auth issues prevent successful job submission
- [Failed](#) uploads from Glasgow WNs to other sites
 - Network issue
 - Reappears periodically (last time – on Saturday)
- Issues with FTS transfers to QMUL
 - Certificate issue?
- ETF tests sometimes fail for Glasgow and Durham
 - Seems like test jobs spend too much time in the queue (tests have strict limits on this)
 - Do sites prioritise ETF test jobs?

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 (JustIN) tests url:
<https://justin.dune.hep.ac.uk/dashboard/?method=awt-results>

Other VOs:

(From Daniela): Moedal now has an IAM server which will double as the voms server once CERN gets around to streamline this (similar to Atlas/CMS/LHCb). An email went out via the portal, could the few sites supporting moedal please *add* the additional server (The old server is not decommissioned yet.) ?

```
$ cat /etc/grid-security/vomsdir/vo.moedal.org/voms-moedal-auth.cern.ch.lsc
/DC=ch/DC=cern/OU=computers/CN=moedal-auth.cern.ch
/DC=ch/DC=cern/CN=CERN Grid Certification Authority
=====
$ cat /etc/vomses/vo.moedal.org-voms-moedal-auth.cern.ch
"vo.moedal.org" "voms-moedal-auth.cern.ch" "443"
"/DC=ch/DC=cern/OU=computers/CN=moedal-auth.cern.ch" "vo.moedal.org"
=====
```

Tokens: We are now rolling out pilot job tokens for all VOs supported on the GridPP DIRAC server.

Please find the configuration info here:

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

I made it a Wiki, so if someone has some words of wisdom on particular quirks of the CE configurations, please add them here.

Done: Oxford, RHUL, Birmingham, Imperial, Glasgow, CERN (NA62), QMUL, Lancaster, RALPP

Meeting Updates

WLCG Workshop in May 13-17 in Hamburg (DESY) [WLCG/HSF Workshop 2024 \(13-17 May 2024\): Overview · Indico \(cern.ch\)](#)

Early bird ended. There is a survey out asking sites what they would like discussed.

The next IRIS Collaboration meeting has been scheduled for **July 2 & 3** at Imperial College. (query when registration opens)

HEPiX in Paris last week: <https://indico.cern.ch/event/1377701>

- Much Centos7 EOL discussion

- Much discussion on NetZero/sustainable computing. Dwayne's two talks, but also many other presentations about optimising site hardware replacement schedules to trade off embodied carbon and reduced power use, DC tuning etc

- Many talks from Seagate - HAMR, next gen SSDs, "embedded carbon, sustainability of production trade-offs - v pro the newer HDDs". Recycling.

- Tokens - security, "upgrade-processes".

- Site reports - toolbox session

CHEP call for abstracts open (deadline 10th of May).

GridPP52 - usual week in August, 28th-30th, at Ambleside.

FTS XRootD dev workshop hosted at RAL (Cosners) this year, 9th-13th of September. (registration is not open yet).

Research Network Engineering Community meeting at 2pm this Friday Katy and Alessandra will talk on WLCG Data Challenge 24, see:

<https://www.jisc.ac.uk/get-involved/research-network-engineering-rne-community-group>

for details.

Tier 1 Status

RAL is currently reporting no major operational issues (which has now majorly jinxed everything!).

Our XMA 2021 tranche of worker nodes had been out of production between the 8th and 18th April. This tranche was drained to have their BIOS upgraded and to enable IPv6. When

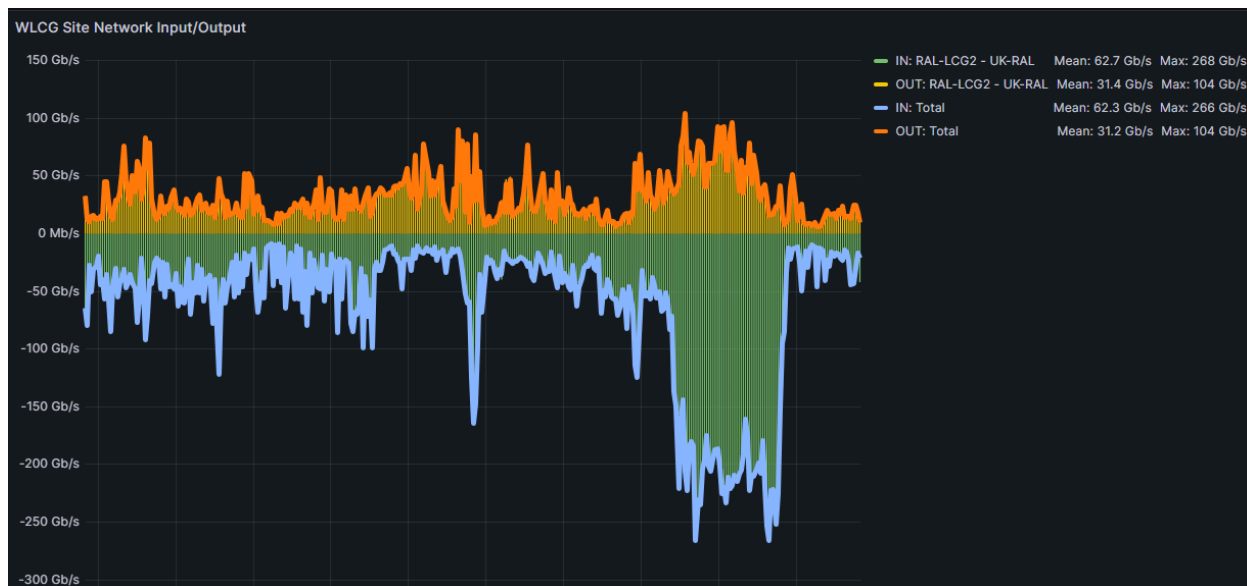
deploying IPv6 across an entire generation, a problem was found with the configuration management system (not always selecting the correct network interface) and this has taken time to resolve. This tranche is now back in full production.

As of this meeting, the whole batch farm is enabled for IPv6 routing.

There have been no operational problems with Antares in the last week, however the services has been ingesting significantly more data in the last few weeks than ever before. e.g. ATLAS wrote 176TB into the atlmc23AOD tape pool on Saturday, with a writing rate of ~10GB/s.

Over 200Gbit/s rates sustained for a week (so far) over the network (peaking at 266!)

<https://monit-grafana-open.cern.ch/d/Mwuxgoglk/wlwg-site-network?orgId=16&var-site=RAL-LCG2&from=1711281601759&to=1713870001759&var-binning=5m&var-bin=10m>



Reviewing gsiftp decommissioning from ECHO. Aware SNO+ in FTS still using it.

Security Brief

Advisories: <https://advisories.egi.eu>

- Operational update [standing item]
- DRI Cybersecurity Workshop (this week)
- Migration from EL7

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

The next INDIGO IAM Hackathon will take place on the **29th and 30th May, 2024** at **CNAF, Bologna**, event indico at: <https://indico.cern.ch/event/1401472/>

Networking News

[UK Mesh](#)

[Check MK monitoring](#)

[Dune Mesh](#)

[perfSONAR certificate status](#)

[perfSONAR issues](#) (the newer doc)

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.

<https://uofglasgow.zoom.us/j/94094366747>

IPv6 Everywhere

A call for v6 on all services is [here](#):

For our records: [Terry's v4/v6 translation work](#)

Twiki for WN IPv6 status - <https://twiki.cern.ch/twiki/bin/view/LCG/Wlclpv6#IPv6Comp>

Lancaster looking at a NAT66 solution.

Technical Update

Standing subjects:

-RHEL8/9 clone experiences:

Rob C is working on a wiki page listing rhel9 tweaks - very useful:

https://www.gridpp.ac.uk/wiki/RHEL9_systems

Also see Rob's work on ARC:

https://www.gridpp.ac.uk/wiki/ARC6_EL9

-HEPSCORE

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

Alessandra has asked for the DN of one person from each site, in the format of the output of
`openssl x509 -noout -in user.crt.pem -subject -nameopt RFC2253`

Tickets/ROD

[UK GGUS tickets](#)

Lots of tickets, note latest campaign to update the .lsc files for the Ops VO.

News

Dwayne's last Ops Meeting :(

AOB/Feedback for/from the PMB

Actions/Decisions from This Meeting

Chat Window:

From Thomas Birkett - STFC UKRI to Everyone 11:07 AM

And to test your backups

From Daniel Traynor to Everyone 11:10 AM

still need to maintain the container

From Thomas Birkett - STFC UKRI to Everyone 11:12 AM

Oh indeed but you can test the container in isolation to a greater degree of confidence

CI/CD pipelines and automated build processes help solve the maintenance issue

From Daniel Traynor to Everyone 11:15 AM

I have 6 certs in total

From Thomas Birkett - STFC UKRI to Everyone 11:20 AM

Security through obscurity is NOT security :)

From Daniel Traynor to Everyone 11:36 AM

can't id ETF tests vs normal jobs so can't really set priority, also we want them to go through the same chain as normal jobs so they go through the same queue.

From Alexander Rogovskiy - STFC UKRI to Everyone 11:37 AM

They use a dedicated cert, so can be potentially identified by DN.

From Daniel Traynor to Everyone 11:42 AM

to note: with tokens by default we map all jobs from the same vo to the same user. And users are where we can prioritise in the batch system.

From Duncan Rand to Everyone 11:55 AM

<https://monit-grafana-open.cern.ch/d/Mwuxgoglk/wlwg-site-network?orgId=16&var-site=RAL-LC-G2&from=now-30d&to=now&var-binning=5m&var-bin=10m>

<https://monit-grafana-open.cern.ch/d/dsY3tf5nk/uk-ral?orgId=16&from=now-7d&to=now&var-source=raw&var-interfaces=All&var-bin=5m>

From Daniela Bauer to Everyone 12:01 PM

Sorry, I've got to go

From Daniel Whitehouse (IC) to Everyone 12:04 PM

Am I right in thinking that the cern monit pages are a good indication that perfsonar is broadly working ok?

Thanks!

From Thomas Birkett - STFC UKRI to Everyone 12:09 PM

Best of luck Dwayne!