

## GridPP Ops 27/2/24

**Attending:** Matt, Rob F, Haruna A, Gordon, Tony, Emanuele, Vip, Mike L, Raul, Steven S, Rob-C, Alex R., Dwayne, Gerard, Linda C, Mark S, Wenlong

**Apologies:** David C, Daniela (wrangling IRIS resource requests), PeteC (at the same meeting)

Urgent Business

None

### **Actions from previous meetings.**

Action Free

### **DC24**

Gets its own section.

Mattermost channel for DC24 prep:

<https://indico.cern.ch/event/1350973/>

Dates for DC24 are 12th-23rd February 2024

Site traffic:

<https://monit-grafana-open.cern.ch/d/Mwuxgoglk/wlcg-site-network?orgId=16&from=now-12h&to=now&var-site=All&var-bin=5m>

Dashboard:

DC24 is over. Any preliminary thoughts?

### **General Updates/Discussion**

**GRIDPP51/SWIFT-HEP JOINT COLLAB MEETING - 26th to 27th MARCH, SHEFFIELD.  
(25th for PMB)**

**REGISTRATION IS CLOSED (AGENDA currently in progress...).**

Let Sam and Jill know if you need to change stuff.

<https://indico.cern.ch/event/1366954/>

## **APEL issues.**

Gordon notes in discussion over the Liverpool accounting that token based jobs are not correctly accounted for. To quote “ARC7, please hurry up...”

Daniela notes that the WLCG seems to have done very little testing before telling everyone “token or else” ?!

More accounting issues at RALPPD, filed by CMS:

[https://ggus.eu/?mode=ticket\\_info&ticket\\_id=165373](https://ggus.eu/?mode=ticket_info&ticket_id=165373)

(Chris has found the issue and is working on solving it.)

## **HEPIX TechWatch WG**

At the last GDB the rebirth of the TechWatch HEPiX WG was announced: (

<https://indico.cern.ch/event/1356131/contributions/5783188/attachments/2800731/4886071/HEPiX-GDB-14-02-2024.pdf> ). The presentation is interesting - but it would also be good to have one or two people from the UK in this WG - and they're currently recruiting.

<https://indico.cern.ch/category/10621/> ← TechWatch Indico. Next meeting is tomorrow (21st) @

1500 UK time.

Did anyone get to this meeting?

-a few people from RAL went.

Discussion on TB-SUPPORT about ARC user mappings post-lcmaps. Rob F made the sensible suggestion of recording this somewhere like the proposed confluence. Also yet another possible tech meeting subject.

-I'll take an action to do that (or at least start the page).

## **VO Updates**

### **ATLAS (Jyoti):**

- General news:
  - Rucio:
    - Disabled the use of tokens across all RSEs.
    - With FTS having disabled their refresh service, all already-submitted transfers queued for longer than six hours are expected to fail.
  - DC24:
    - ATLAS DC24 ended on Friday evening (the plan was to prolong it to Monday if the global rates were not achieved).
    - Last week, ATLAS FTS was overwhelmed by too many concurrent submitted transfers: there were 3M file transfers queued. Many of these transfers were obsolete and could be deleted from the DB.
      - They were deleted on Thursday

- Rucio also was hot patched to cope with the rate of injection and cleaning to avoid race conditions
    - ATLAS had a full mesh with >1200 links with injections (~2000 including the standard production) and FTS never had to support such a load
      - fts3-pilot was redeployed with higher memory to support the T2s.
      - Links and storage limits of concurrent transfers were increased to 5000 for T1s and 7000 for T0
      - Tokens refreshes were disabled after the first week and then completely
    - Friday 2.5 Tb/s (all experiments) were achieved for quite a few hours.
      - ATLAS passed 1.5 Tb/s for 4 consecutive hours.
    - Actions to revert FTS configurations to pre-DC24.
      - Default values reverted this morning
    - Plan another test from CERN to T1s.
    - [ATLAS DC24 retrospective](#).
  - Harvester disk utilisation:
    - Since one out of four nodes submitting as aipanda159 stopped submitting due to high load, the rest got their disks filled with logs. Back to nominal configuration and manually cleaned up logs.
  - Switcher3 ALMA9 upgradation tests: [CSOPS-2204](#).
  - First draft for VOMS decommissioning is there (getting circulated among the ADC experts).
- Tier-1:
    - Early RAL DATADISK full pledge increase for 2024/2025 – done, [ATLDDMOPS-5681](#).
    - DC24:
      - Data carousel was disabled for RAL in the second week to avoid the gateways being overwhelmed by staging activity
      - Tuned storage GWs multiple network/OS settings.
      - LHCOPN connection repaired amidst the Challenge.
      - What do other sites have for ECN ( Explicit Congestion Notification)?
    - [GGUS-165358](#):
      - Deletion attempts at RAL-LCG2 take a long time; **solved and closed**.

- As expected, with the end of DC24, the load decreased and average deletion duration dropped below 1s – [link](#).
    - Internally there are actions to improve the deletion process; may take months, however in general there is not anymore a load similar to DC24 so it shouldn't cause a problem.
  - The deletions have caught up with the additional writes from DC24, and ATLAS is now under the new pledge – see the weekly plot [here](#).
  - Yesterday, rolling out the permission mask to the Antares directory structures was successfully performed.
- Tier-2s:
  - Recently closed/solved GGUS tickets:
    - [GGUS-165393](#): DE FZK-LCG2: Transfer failures as destination.
    - [GGUS-165394](#): UKI-NORTHGRID-LANCS-HEP: Huge transfer failures as destination.
  - [GGUS-163552](#):
    - Timeouts in transfers as a destination and in deletions at UKI-SCOTGRID-GLASGOW; **in progress**.
    - The Tier-2s got hot hard by ATLAS with a large amount of transfers on 22 February.
    - Actions expected this/next week (as the DC24 is over).
  - Brunel:
    - For quite some time after its downtime was over, there were no new ATLAS jobs.
    - An error was found in the Arc Map for ATLAS and got corrected for dc2-grid-[21,26,28].
    - New ATLAS jobs started to register last Saturday; looking good now.
  - Lost data @DURHAM – rebuilding storage elements, [ATLDDMOPS-5680](#).
    - It needs some updates (later) this week!
  - Pending: LIV SAM3 test failures!
  - Pending at the ADC level: Glasgow - to return as nucleus at some point.
    - Not mentioned when that “some point” is going to be!
  - ECDF\_CLOUD HammerCloud exclusions go on, sometimes resulting in blacklisting.
  - RHUL\_VP is also experiencing frequent HammerCloud exclusions (especially the 1013 template).

- Job Accounting (UK Cloud): all UK sites saw a drop at 07:00 AM today – the nature of the drop is identical; [link](#).

### CMS (Daniela):

CMS trying to have more debugging information included in their ARC-CE SAM tests:

[https://ggus.eu/?mode=ticket\\_info&ticket\\_id=165402](https://ggus.eu/?mode=ticket_info&ticket_id=165402)

(And there was me thinking the WLCG experiments were directly involved in developing these tests, but alas no.)

Bristol was flagged for failing SAM tests:

[https://cmssst.web.cern.ch/sitereadiness/report.html#T2\\_UK\\_SGrid\\_Bristol](https://cmssst.web.cern.ch/sitereadiness/report.html#T2_UK_SGrid_Bristol)

[https://monit-grafana.cern.ch/d/A\\_1kxGmMz/wlwg-sitemon-latest-tests?orgId=20&var-vo=cms&var-dst\\_tier=2&var-dst\\_country=UK&var-dst\\_federation=All&var-dst\\_experiment\\_site=T2\\_UK\\_SGrid\\_Bristol&var-service\\_flavour=All&var-dst\\_hostname=All&var-metric=All&var-status=All](https://monit-grafana.cern.ch/d/A_1kxGmMz/wlwg-sitemon-latest-tests?orgId=20&var-vo=cms&var-dst_tier=2&var-dst_country=UK&var-dst_federation=All&var-dst_experiment_site=T2_UK_SGrid_Bristol&var-service_flavour=All&var-dst_hostname=All&var-metric=All&var-status=All)

Luke is on holiday:

[https://ggus.eu/index.php?mode=ticket\\_info&ticket\\_id=164092](https://ggus.eu/index.php?mode=ticket_info&ticket_id=164092)

Bristol currently has 8 open tickets. I don't want to be Luke when he gets back.

-back now

–"the storage might never come back", better to move on, EL9, new storage, new HTCondor.

Wrt last week's attempt to locate the CMS token tests for the Tier3s.

I think I have located them on the \*preprod\* nagios server:

<https://etf-cms-preprod.cern.ch/>

A quick scan of the ARC sites shows:

[org.cms.WN-01basic-/cms-ce-token](http://org.cms.WN-01basic-/cms-ce-token)

QMUL seems to fail with: "cn018.htc.esc.qmul: CRITICAL: insufficient free /tmp space" (seems to also be true for other WNs)

Glasgow has an ipv6 Warning <shrug>

RALPP passes.

However the next text up ([org.sam.CONDOR-JobSubmit-/cms-ce-token](http://org.sam.CONDOR-JobSubmit-/cms-ce-token)) fails for RALPP as well, with the job ending up in HELD state -> I assume the CMS against WLCG ticket mentioned at the start of the CMS section relates to this.

If you decide to have a go at debugging this, please let me know by email.

Katy DC update-

CMS using tokens by preference at 25 sites. Very wide scope, so token numbers not very large.

Tokens thus worked well. ~50% throughput done via token - IC and RALPP used them, Tier-1 didn't, but since DC ended token tests all green so likely was a load or tuning issue.

Some token mini-challenges needed in the future, especially when move to using tokens in a full way.

DC was a success, major targets were being hit. Got a few hours of 1Tb/s. Target at Tier 1 reached (more or less), but needed the tuning to get there.

Further analysis to come. IC link wasn't fully exercised. Fewer FTS issues than atlas due to using fewer links.

ECN being switched off (at Tier 1) was again a big part of the tuning. Alex gives a description of ECN. (date of this was ~21st).

Faster data tunings also applied: <https://fasterdata.es.net/host-tuning/linux/>

Duncan will follow up.

The OPN being down was problematic, and any issues could be blamed on it.

-Section to discuss this at GridPP51

## LHCb (Alexander):

General:

- DC24 is over

RAL T1:

- Token [setup](#)
  - Tokens were not used for RAL during DC24
- DC24
  - Writing part
    - Target rate for RAL was 3.96GiB/s (this is for average)
    - 2.68GiB/s was [achieved](#) on EOS->RAL Disk [link](#)
      - Instantaneous speed crossed the target only for 4 hours
    - 2.64GiB/s was [achieved](#) on RAL Disk -> RAL Tape [link](#)
      - This link was much faster than the previous one, low average is due to lack of data to transfer (because of relatively slow first link)
    - Very good efficiency on both links (almost no errors)
    - OPN link was down during the writing part
  - Staging part
    - Target rate for RAL was 2.40GiB/s
    - 3.28GiB/s [achieved](#)
    - Efficiency is not that great, because of the ECHO gateways overload
  - 4 files are missing at RAL after the writing part, most probably deleted by FTS
- Very [slow transfers](#) between SARA and RAL
  - Network tuning was done at RAL, that fixed the problem (and improved other things)
- Vector read optimisation (follow-up on GGUS [ticket](#))
  - New configuration was rolled-out to preprod farm
- Long-lasting issues

- [Slow](#) “stat” calls

T2:

- Jobs are [failing](#) at Sheffield
  - Pilot output is unavailable at the site, that complicates debugging
- ETF tests are [failing](#) for Lancaster
  - Still waiting for a few sites to fix their SEs to allow storage tests from the new tests suite to pass
- Pilot submission [problem](#) at Durham
  - Ticket is still in the “assigned” state (for almost two weeks!), needs reaction
- Pilot submission [problem](#) at Manchester
  - Looks OK
- [Failed](#) pilots at Bristol
  - Looks better since the 23rd of Feb
  - Ticket is still in the “assigned” state, needs reaction
- Some files are [inaccessible](#) at Manchester
  - DFC dump is ready, site contacted to provide storage dump
- Inefficient user jobs at Oxford
  - One user is submitting single-core jobs, while requesting 8 cores
- A lot of failed pilots at Lancaster since yesterday evening
  - `LRMS error: (309) Job failed`
  - No pilot output available

Matt will take a look, and if there’s no quick fix a ticket will get opened.

**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>
- Lancaster xrootd write issue
  - (the problem is with xroot TPC, Matt suspects it never got configured properly for their system)
- Dump file list for consistency check  
[https://ggus.eu/?mode=ticket\\_info&ticket\\_id=165419](https://ggus.eu/?mode=ticket_info&ticket_id=165419)  
[https://ggus.eu/?mode=ticket\\_info&ticket\\_id=165420](https://ggus.eu/?mode=ticket_info&ticket_id=165420)

Other VOs:

**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](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

CERN dragging their heels on providing an IAM server for Moedal:

[https://cern.service-now.com/service-portal?id=ticket&table=u\\_request\\_fulfillment&n=RQF2470783](https://cern.service-now.com/service-portal?id=ticket&table=u_request_fulfillment&n=RQF2470783)

Tier 1 now also ticketed, but no ETA yet.

26/02/24: No progress.

Note that na62 use tokens in production (see earlier conversation on apel accounting and tokens).

Gordon explained last time his “active shim” for adjusting the ARC accounting databases.

## Meeting Updates

Tech Meeting last Friday on Token Transition

<https://indico.cern.ch/event/1375893/>

It caught a lot of us by surprise, so we should keep a link to the TM indico in here:

<https://indico.cern.ch/category/4454/>

(none planned for this week)

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

HPC SIG meeting at RAL next week:

<https://hpc-sig.org.uk/index.php/event/hpc-sig-meeting-rutherford-appleton-laboratory/>

Online registrations still available for remote attending.

The next IRIS Collaboration meeting has been scheduled for June 3 & 4 at Imperial College.



## Tier 1 Status

DC24 completed successfully 23/2/24.

Initial thoughts:

- Overall, it has been considered reasonably successful from the RAL perspective with all VO's achieving target data transfer rates. We did encounter some ECHO issues during the challenge but these were resolved and did not become "show stoppers".
  - Improved GW network tuning and HA proxy setting.
  - BAU activities were stopped for some periods but most of the time reweighting and rebalancing continued as normal).
  - Deletion rates under high load /high occupancy to be investigated further.
- Suggest review of FTS channel settings for VOS (and how they interact with each other )
- Antares ran with no interventions ( aware ATLAS stopped RAL being in Data Carousel for part of the period.)
- Aware there may have been increased WN Job failures due to high I/O rates in Echo gateways (to be investigated further)
- Recovered LHCOPN connection halfway through DC24

As is our standard practice, we will be running a retrospective to gain the lessons learnt as part of our process of continual improvement.

## Security Brief

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

- Operational update [standing item]
- SOC Hackathon registration open: <https://indico.cern.ch/event/1370544/>
- DRI Cybersecurity Workshop (23rd and 24th of April)
- 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

Note - AuthZ meetings moved to fortnightly Thursdays

## Networking News

[UK Mesh](#)

[Check\\_MK monitoring](#)

[Perfsonar refresh - GridPP Wiki](#) (old)

[Dune Mesh](#)

[perfSONAR certificate status](#)

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

We followed up on sites with issues after the release of 5.0.8. Many improved until only RAL-LCG, Imperial, Glasgow and Durham were not showing results on the Maddash mesh. Now I see that Liverpool and Bristol seem to have regressed. Upgrading the pS host OS to EL9 is recommended as was done at Glasgow recently.

Rob F is looking at their box, it seemed to have crashed.

Bristol looks okay, doesn't appear to be storing results. Trying a restart.

Some discussion of the toolkit-less version at IC, Duncan will talk to the devs about this.

## 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>

Meeting tomorrow.

## IPv6 Everywhere

A call for v6 on all services is here:

<https://indico.cern.ch/event/1341866/contributions/5649115/attachments/2744834/4775731/IPv6%20compute%20deployment%20ops%20coord.pdf>

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

See network forward look.

HEPIX IPv6 WG call last Wednesday - <https://indico.cern.ch/event/1352163/>.

Next call is on 18 January.

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

A Tech Meeting was called for discussion on v6 solutions, now we're in the ADC24 era we'll arrange.

## 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](https://www.gridpp.ac.uk/wiki/RHEL9_systems)

Also see Rob's work on ARC:

[https://www.gridpp.ac.uk/wiki/ARC6\\_EL9](https://www.gridpp.ac.uk/wiki/ARC6_EL9)

-HEPSCORE

[https://www.gridpp.ac.uk/wiki/HEPSCORE\\_Site\\_Status](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`

-XRootD Shoveller Adoption/Other XRoot Monitoring

-the endpoint we're supposed to be shovelling to seems a bit unstable.

## Tickets/ROD

[UK GGUS tickets](#)

64 tickets, all "interesting" ones appear to have already been mentioned again.

ROD dashboard is uninteresting,

## Analysis Facilities:

[white paper is finalised](#) and HSF AFF is collecting endorsers.

-jyoti gives the SWIFT-HEP mailing list in the chat.

## News

None

## AOB/Feedback for/from the PMB

None

## Actions/Decisions from This Meeting

- Action on Matt to ARC doc page in confluence
- Action on Matt/Sam to talk to Alastair about future Tech Meetings

## Chat Window:

11:17:13 From Alessandra Forti To Everyone:

Sorry there was indeed an improvement on Friday.

11:17:29 From Alessandra Forti To Everyone:

We reached the target also from CERN

11:17:48 From Katy Ellis - STFC UKRI To Everyone:



11:29:25 From Samuel Skipsey To Everyone:

Effectively, the logical effect of ECN marked packets is the same as "dropping a packet" on TCP's rate backoff - but of course you don't actually drop a packet

11:29:42 From Samuel Skipsey To Everyone:

(So ECN lets you "warn that you might drop packets" before you actually do)

11:30:13 From Katy Ellis - STFC UKRI To Everyone:

<https://fasterdata.es.net/host-tuning/linux/>

11:30:35 From Thomas Birkett - STFC UKRI To Everyone:

Curious to know what the negative effect of this change may be, could we crash a router or DoS it for instance? I'm always nervous of turning off features designed to keep networks stable

11:33:18 From Raul Lopes To Everyone:

If your site has a perFONAR box on EL9, the TCP stack configuration is the fasterdata recommended one

11:35:09 From Alessandra Forti To Everyone:

I love the balance...

11:41:20 From Alessandra Forti To Everyone:

@Alex I'll provide the dump today

11:59:05 From Dan Whitehouse (IC) To Everyone:

Am I right in saying we (IC) are all grey on the first graph,

11:59:17 From Dan Whitehouse (IC) To Everyone:

However that's not the case for some of the others?

12:01:43 From phxIk To Everyone:

Bristol node is failing to archive:

Feb 27 12:01:10 dice-io-37-00 python3[12976]: 2024-02-27 12:01:10,081 - werkzeug - INFO - 127.0.0.1 - - [27/Feb/2024 12:01:10] "#033[31m#033[1mPOST / HTTP/1.1#033[0m" 405 -

Feb 27 12:01:10 dice-io-37-00 archiver WARNING 53173175: Failed to archive  
<https://dice-io-37-00.acrc.bris.ac.uk/pscheduler/tasks/22e37b16-ec7c-4b29-8be4-55d8d2c716fb/runs/e14b3e1f-10fe-4e11-ac60-ed9e1030f6e6> to esmond: Archiver permanently abandoned  
registering test after 1 attempt(s): 405: Invalid JSON returned

Feb 27 12:01:10 dice-io-37-00 archiver WARNING 53173175: Gave up archiving  
<https://dice-io-37-00.acrc.bris.ac.uk/pscheduler/tasks/22e37b16-ec7c-4b29-8be4-55d8d2c716fb/runs/e14b3e1f-10fe-4e11-ac60-ed9e1030f6e6> to esmond

12:01:49 From Dan Whitehouse (IC) To Everyone:

I think the thing is called "testpoint"

12:02:04 From Dan Whitehouse (IC) To Everyone:

I think that either we or someone else needs to host an additional service

12:02:35 From Chris Brew - STFC UKRI To Everyone:

Nuke it from orbit, it's the only way to be sure.

12:06:15 From Brian Davies - STFC UKRI To Everyone:

FYI dice certificate at Bristol appears to have expired so maybe being denied

12:08:50 From phxIk To Everyone:

thanks, will check and update

12:10:38 From VipDavda To Everyone:

I've to go. cheers Vip

12:11:55 From Jyoti Biswal - STFC UKRI To Everyone:

UK-SwiftHep@cern.ch