

## GridPP Ops 21/1/2020

Attending: Matt, Elena, Gordon, Raul, Daniela, Emanuele, Sam, James P, RobertF, Mark Slater, Duncan, Andrew, WinnieL, Brian, Teng, Pete C, Wenlong  
Apologies:

### Action from previous meetings.

\*190618-02 Duncan - form a plan for the future of perfsonar for GridPP sites.

-Technical meeting a few weeks back

<https://indico.cern.ch/event/876065/>

Follow up around March time.

\*191008-01 David C, Plan security team F2F followed by security day

Pre-security day team F2F being arranged.

\*191119-01 - Sites mentioned in dzero ticket [143926](#) to remove dzero from publishing.

#### For the CEs I see:

ce3.dur.scotgrid.ac.uk:2811/nordugrid-SLURM-ce3

ce4.dur.scotgrid.ac.uk:2811/nordugrid-SLURM-ce4

#### and for the SEs I see:

hepgrid11.ph.liv.ac.uk

heplnx204.pp.rl.ac.uk

se01.dur.scotgrid.ac.uk

**Matt has poked Durham about this, who have it as a round tuit. Ste is on it at Liverpool.**

**Matt will poke RALPP.**

\*191210-01 - Action on Matt to submit a ticket to find out about gfal2 python3 support.

[https://ggus.eu/index.php?mode=ticket\\_info&ticket\\_id=144444](https://ggus.eu/index.php?mode=ticket_info&ticket_id=144444)

Finally a response from the devs - python3 support will be indeed ungoing. This action can be closed. Matt will put gfal2-python3 into the tarball - but a follow up will be to ask this to be included in the UMD WN and UI metapackages.

\*2000107-01 - Action on all - follow up Sheffield and Sussex status.

Thanks to Elena and Patrick for getting back to me and updating their tickets. Things are looking reasonably under control, but there's a lesson here in using downtimes to avoid attracting the attention of the powers that be.

[https://ggus.eu/?mode=ticket\\_info&ticket\\_id=144719](https://ggus.eu/?mode=ticket_info&ticket_id=144719) (Sheffield)

[https://ggus.eu/index.php?mode=ticket\\_info&ticket\\_id=144720](https://ggus.eu/index.php?mode=ticket_info&ticket_id=144720) (Sussex)

We can likely close this action too.

## VO Updates

### Atlas (Elena):

In the event of any problems I suggest people pop along to the Thursday cloud support meetings:

<https://indico.cern.ch/category/4620/>

From WLCG meeting yesterday:

ATLAS [reports](#) ( [raw view](#) ) -

- Work in progress with Grand Unification more optimal decisions for Analysis vs. Production
  - [https://indico.cern.ch/event/877661/contributions/3703256/attachments/1969036/3274934/2020-01-14\\_Grand\\_Unification.pdf](https://indico.cern.ch/event/877661/contributions/3703256/attachments/1969036/3274934/2020-01-14_Grand_Unification.pdf)
  - Unification of Production and Analysis queues to a single PanDA Queue
  - Requirements for sites:
    - Harvester will submit pilots with pilot and production roles to the same queue on the same CE
    - No hard partition between analysis vs production jobs
- AGIS config for Grand Unified PanDAQueue:
  - same as production UCORE queue, plus:
  - set new queue type='unified'
  - associate SCRATCHDISK endpoint for write\_lan\_analysis and read\_lan activities
  - set direct\_access\_lan as needed for analysis jobs (it won't affect prod jobs)
  - optionally, set resource\_type\_limits.XXX params, where XXX is a job type to limit e.g. SCORE
  - Pending items
    - HammerCloud doesn't know yet about 'unified' queue type
      - Currently: manually setting unified queues to ONLINE
      - Short-term workaround: use only PFT results for HC auto-blacklisting/whitelisting
      - Next step: blacklist if (PFT or AFT) fail, whitelist if (PFT and AFT) succeed
    - Separation of analysis vs production jobs on unified queues in monitoring - implementation in progress
  - Migration progress

- CERN tested since December, working well
      - now fully grand unified: old production and analysis queues set to BROKEROFF and draining
    - Gaining experience with one more pull and one push queue
      - IN2P3-CC
      - LRZ-LMU\_UCORE
    - Then ATLAS will start to contact sites/clouds to ramp up migration
  - 3 CERN FTS with failing transfers [INC2274140](#)
    - No FTS server has updated the transfer status the last 900 seconds. Probably stalled
  - Few issues with slow queries in Rucio (related to new features from October release)

[GGUS tickets:](#)

### **UKI-SCOTGRID-GLASGOW**

ggus [144759](#) **State:**in progress High traffic from GLASGOW on RAL CVMFS Stratum1

ggus [144688](#) : timeout transfer errors

Atlas should transfer less files

### **UKI-SOUTHGRID-SUSX**

ggus [142329](#) **State:**on hold :CentOS7 migration

Glasgow is moving to Ceph storage

<https://its.cern.ch/jira/browse/ADCINFR-152>

Shrinking UKI-SCOTGRID-GLASGOW\_DATADISK

<https://its.cern.ch/jira/browse/ADCINFR-162>

Sheffield is going storageless:

<https://its.cern.ch/jira/browse/ADCINFR-147>

- Decommission UKI-SCOTGRID-ECDF-RDF storage

<https://its.cern.ch/jira/browse/ADCINFR-161>



### **CMS:**

All quiet- Bristol made a comeback - RALPP data transfer issues.

### **LHCB:**

Andrew couldn't spot anything in the ops minutes.

## “Other” VOs:

DUNE :

glideInWMS configuration for the different sites :

<http://gfactory-2.opensciencegrid.org/factory/monitor/factoryEntryStatusNow.html>

Data moving from RDF. RDF closes in August.

LSST:

Still lots of waiting jobs at Lancaster (although we've hit throughput issues at the site).

QM had jobs being spread over several nodes, there may be some job failures.

James - seeing jobs running at QM now, so picture improving.

Teng has SAM tests working - page for this coming.

SKA:

Anna asked for new meerkat VO (due to privacy reasons). Working on it in the voms, sites will need to enable it.

Very high memory jobs are coming, need to handle “precious” resources (like GPUs) - behind its own queue. Daniela can make a tag for it - there is an exclusive tag. Needs own queue.

“New” VO status:

No news.

## General Updates

**Survey on CREAM CEs and HTCCondorCE replacements:**

<https://www.surveymonkey.co.uk/r/62YVJ3C>

**This will be used to plan HTCCE training events, so worth taking - it is a quick one.**

## Meeting Updates

GDB last week:

<https://indico.cern.ch/event/813757/>

*Operational intelligence:* proposal to automatize some of the shifter work by using the enormous amount of log files we produce and ML techniques. VOs should start writing data in similar

format to make it easier to use the same system for everyone. The plan is to have a microservices infrastructure, so it will be based on k8s.

*Compute provisioning ideas:* proposal to have a common API among VOs to prioritize workflows. Idea born from the LSST request to increase their fairshare for few weeks because they have a peak load.

*Updates on OSG/WLCG perfSONAR Network Monitoring and Analytics:* plans to use analytics on perfsonar data. Close to ops intelligence. Lots of extra work going on, but more than 10 years of perfsonar we are still not using it for intelligent brokering.

*What CPUs to support, at what cost:* new cpus are coming up with different instruction sets and this talk was to trigger discussion about which CPU would be better to buy. There wasn't any proposal and experiments should express their opinion and different compilation flags should be recommended to make the most of the CPU power. Perhaps worth reading the chep talk that triggered this meta-proposal

*CHEP 2019: C. Marcon (Lund U): Impact of different compilers and build types on Geant4 simulation execution time*<https://indico.cern.ch/event/773049/contributions/3473317/>

With a LHCOPN/LHCONE pre-GDB:

<https://indico.cern.ch/event/828520/>

Summary slides (from GDB):

[https://indico.cern.ch/event/813757/contributions/3698521/attachments/1969608/3276108/LHCOPNE-20200114-CERN-meeting-report-v1.0.pptx\\_1.pdf](https://indico.cern.ch/event/813757/contributions/3698521/attachments/1969608/3276108/LHCOPNE-20200114-CERN-meeting-report-v1.0.pptx_1.pdf)

*Summary from AF:*

The core of the meeting was on how to handle multi-VO traffic beyond LHC. There were several discussions on how to isolate each VO traffic, paint packets, and orchestrate traffic. This is not only to avoid each VO traffic clogs the network, but also (mostly?) related to funding. All the solutions proposed will make T2 life very difficult and might require special hardware. The packet painting to identify the traffic will instead require all the experiments either implementing something at the core of their software or use special libraries that do not exist yet.

The experiments presentations contain also interesting future numbers but they are likely to change in the next few years so I wouldn't use them as definitive just as ball park.

Technical Meeting on Friday was on Google Summer of Code student ideas:

<https://indico.cern.ch/event/876083/>

Were any conclusions reached? Minutes in the meeting.

Technical meeting this Friday on HTCondorCEs:

<https://indico.cern.ch/event/876269/>

## **Tier 1 Status**

Castor downtime 22/01/20 for DB HW upgrade.

CMS job efficiency has dramatically increased (reason as yet unknown - no changes on either side).

## **Security Brief**

NTR

## **Storage and Data Management News**

A lot of retiring going on.

## **On Duty Report**

NTR.

**Technical Update** (was Tier 2 evolution, Accounting, Monitoring, Documentation, Services)

NTR

## **Tickets**

45 Open Tickets this week.

No movement on the Sussex ticket last week, but lots of work behind the scenes and all tests were green so I closed it:

[https://ggus.eu/?mode=ticket\\_info&ticket\\_id=144626](https://ggus.eu/?mode=ticket_info&ticket_id=144626)

It's come round to a few IPv6 tickets needing an update - Liverpool and Oxford. It's annoying when there's no news, but hopefully a once a quarter ticket update isn't too onerous.

Although can everyone please check their tickets- we're far enough into the New Year now that no ticket not "On Hold":

[GGUS tickets ordered by date.](#)

## Site News

We'll devote some time to helping Patrick at Sussex.

## AOB

NTR

## Actions/Decisions from This Meeting

## Chat Window:

<https://indico.cern.ch/event/872989/>

I can submit jobs as an atlas user to susx and the job runs. Don't see any smoking guns.

gm-jobs.log

this is arc6, I think it has different log file names.

prob equivalent arex-jobs.log

Sorry, but I've got to switch to another meeting.

do the all accounts exit in gridengine? su to prd/pilot atlas user and try and submit a job on the arc ce.

apologies other meeting at DC

Got to go to another meeting

in arc-jobs there should be an entry with the job "... lrmid: XXXXXX" with the gridengine job number if the arc ce is able to pass the job on to the batch system.

the hi memory nodes are not behind their own queue atm

you get to them by request enough memory

@Robert: Then it doesn't work.

ok, then we'll set up a queue then

Can you please set them behind it's own queue ? There is no way to advertise memory correctly and/or match it.

they used to be behind their queue

thanks

sorry sorry, gotta go

wait

Workflow change??