

GridPP Ops 25/1/22

Attending: Matt, Gerard H, Jose, Brian D, Vip, Mike L, Raul Lopes, Emanuele, SamS, JamesW, Wenlong, RobC, Linda C, Steven S, Gordon, RobertF

Apologies: Mark S, Darren

Action from previous meetings.

20210907-01 - David C- put forward a person to be the site security threat intelligence contact. Last act will be to get the Site Contacts round a virtual table, this has been pushed to this year.
-Will be a Technical Meeting. [ACTION DC]
-Alastair being messaged to arrange, penciled in for the 11th of Feb.

VO Updates

CMS (Daniela):

What happened at RAL yesterday ?!

https://cmssst.web.cern.ch/sitereadiness/report.html#T1_UK_RAL

-Brian suggests that Katy did some rucio cleanup yesterday which might have not on things Tier 2s all green.

How many CMS computing people does it take to copy three missing test files to a site ?!

Brunel celebrates the resolve of the eternal test failure ticket:

https://ggus.eu/index.php?mode=ticket_info&ticket_id=155294

Imperial now has GPUs for CMS available, however there is no great demand.

If you do want a CMS GPU queue, please open a GGUS ticket along these lines:

https://ggus.eu/?mode=ticket_info&ticket_id=155551

-Some discussion on getting GPUs up and running. Dan W notes that setting gpus=1 works okay directly, but this bit was getting stripped off going through the CEs. Needed to set an elaborate start variable to get things to work.

Not fully understood why the HTCCEs were stripping this.

IC are using version < > Bristol are waiting for 9.X for the enhanced GPU support.

A(nother) technical meeting on this? Sam wisely suggests we see the outcome of the Analysis Facility Technical Meeting this Friday before deciding. <https://indico.cern.ch/event/1117141/>

https://monit-grafana.cern.ch/d/m7XtZsEZk4/wlcg-sitemon-historical-tests?orgId=20&var-vo=cms&var-dst_tier=2&var-dst_country=UK&var-dst_federation=All&var-dst_experiment_site=All&var-service_flavour=All&var-dst_hostname=All&var-metric=All&var-status=All

ATLAS (James):

- Rucio issues last week(end) causing jobs failures and sites to go into HC test
 - Believed load / DB related; various updates to k8 / VM /loadbalancers, etc. to mitigate the issue
- Next Week: ATLAS S&C: <https://indico.cern.ch/event/1100871/>
- RAL:
 - Four hosts now included in webdav alias;
 - Continues to mainly do Reads (achieved peak 7GB/s over weekend)
 - Some FTS tuning to optimise the throughput.
 - Some transfer issues seen between RAL and SWT2; although not exclusively to RAL
 - HC sometimes unable to get itself out of test, due to jobs stuck in "Starting" state.
 - Trying to identify exact cause
- ECDF - Awaiting outcome of disk rebuild
- Brunel - also attempting to restore ATLAS test files

Jose notes that in the days of AGIS you could force a site to be online, assumes this possible in CRIC?

James confirms this is what happened Friday evening, which had some interesting knock on effects.

LHCb: (Mark)

RAL:

- FTS transfer failures via gsiftp. Seems it was caused by heavy traffic at RAL and the an RRCKI issue. Issue can now be closed (https://ggus.eu/index.php?mode=ticket_info&ticket_id=155605)
- Problems found in the XrdCeph plugin that sets all files as 'Offline = 1'. This then breaks command line root access as it believes the files can't be accessed and doesn't set the size of the file. Discussions ongoing as to how to proceed (https://ggus.eu/index.php?mode=ticket_info&ticket_id=155570)
Will discuss this tomorrow, this could hit other sites.

Tier 2:

- Glasgow:
 - FTS transfer failures (https://ggus.eu/index.php?mode=ticket_info&ticket_id=155153). Any luck getting

the disk server (svr018.gla.scotgrid.ac.uk) going again? It's been going again for several days!

“Other” VOs:

DUNE (Wenlong):

glideInWMS configuration for the different sites :

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

SAM tests :

https://etf-dune-preprod.cern.ch/etf/check_mk/

--- Now also testing for the stash cache

--- following up with the sites

--- [Monit link for history](#) : Shows historical timeline of tests now

--- CRIC information ready. Using it for ETF tests now.

DUNE site naming agreed: UK-Sitename format.

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

Wenlong is the DUNE UK data manager

Storage monitoring site:

<https://dune.monitoring.edi.scotgrid.ac.uk/app/dashboards#/view/318dbad0-727a-11ea-9164-a38a629544b4>

DUNE collaboration meeting this week.

Job efficiency is getting better since last week, after final upgrade of dCache, ipv6... There is a stashcache redirector at NIKHEF, need confirm UK jobs use NIKHEF.

Job efficiency monitoring site:

<https://fifemon.fnal.gov/monitor/d/000000069/user-efficiency-details?orgId=1&from=now-30d&to=now-5m&var-cluster=fifebatch&var-user=dunepro>

Efficiency boosted to at least 50%. If using NIKHEF might not need to have RAL spin one up. Wenlong will find out more at next week's Dune Ops (thanks!).

SKA:

Looking into token authorization for rucio/FTS/storage using ESCAPE credentials.

-starting at dcache at DESY and STORM at CNAF. Going forward we want to look into getting involved.

Another topic for the storage meeting?

LIGO:

NTR

LSST:

LSST “Data Facility Workshop” last week.

Other “New” VO status:

NTR

General Updates/Discussion

NTR

Meeting Updates

Analysis Facility Technical Meeting this week <https://indico.cern.ch/event/1117141/>

OMB meeting last Thursday: <https://indico.eji.eu/event/5764/>

Nothing very exciting, although the information system evolution plans seemed very unambitious.

ATLAS S&C week 31/01/22 - 4/2/2022

- General <https://indico.cern.ch/event/1100871/>
- ADC only <https://indico.cern.ch/event/1115413/>

Tier 1 Status

Further development of Antares ongoing.

Webdav improvements for ATLAS (and others)

“Cloud Team” is happy to explore further usage of jupyter notebooks . Contact (FAO Jacob Ward) via cloud-support@gridpp.rl.ac.uk

Tier-1 Network upgrade (latest update on top)

(08-11-21)

Moving a batch of WNs to new network.

-Jose notes that in preparation these nodes will be drained, there will be a small reduction in capacity over the weekend.

Adding new Storage Nodes for Echo into the new network

Plan to Join LHCONE 1st/2nd Week December

Internal changes to current network needed for this ongoing 10/11/21 17/11/21 22/11/21

(25-10-21) New WN transition to the new network due soon- this week or next (we think).

(12-10-21) New network online, but no hosts on it yet. Working in progress. New worker nodes and storage will be put on it. Part of work is making the migrations easier.

(24-08-21) Connections and configuration of new RAL Tier1 network ongoing. Connections and configuration of Antares network ongoing.

RAL Campus network upgrades complete which should allow for extra capacity , better diagnostics , and increase in IPv6 bandwidth to be planned.

Progress (but not complete_ configuration of new perfSONAR hosts).

Joining the LHCONE is part of the larger Tier-1 Network upgrade that we are doing. The upgrade plan was presented at the technical meeting on January 29th[1] and it remains on schedule. The next update is scheduled at the LHCOPN-LHCONE meeting on the 23rd March[2]. Access to the machine room is expected to be allowed this week. The cabling for the Tier-1 network is expected to be completed between the 16th and 19th March. The CTA network is currently being configured.

[1]<https://indico.cern.ch/event/996031/contributions/4185620/attachments/2180652/3683446/Tier1Network20210129.pdf>

[2] <https://indico.cern.ch/event/983436/>

Brian- internal testing of CTA progressing

Security Brief

- Operational update [standing item]
- Post-break update

- Setting agenda for 2022

For our bookmarks: <https://csirt.egi.eu/news/>

Threat intelligence meeting penciled in for the 11th of Feb. Open to all (not just contacts).
Next IRIS Security Workshop due mid-Feb, ideas/thoughts/volunteers welcome.

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

Tom Dack is the current chair of the WLCG Authorisation Group.

Networking News

UK Mesh:

<https://psmad.opensciencegrid.org/maddash-webui/index.cgi?dashboard=UK%20Mesh%20Config>

Check_MK monitoring:

https://psetf.opensciencegrid.org/etf/check_mk/index.py?start_url=%2Fetf%2Fcheck_mk%2Fview.py%3Fhostgroup%3DUK%26opthost_group%3DUK%26view_name%3Dhostgroup

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

Dune Mesh:

<https://psmad.opensciencegrid.org/maddash-webui/index.cgi?dashboard=DUNE%20Mesh%20Config>

Duncan shared this monitoring last week:

<https://monit-grafana-open.cern.ch/d/000000505/perfsonar-e2e-performance?orgId=16&var-src=UKI-SCOTGRID-ECDF&var-dst=UKI-NORTHGRID-LANCS-HEP&var-dst2=UKI-LT2-IC-HEP&var-dst3=UKI-SOUTHGRID-RALPP&from=now-3h&to=now>

Lots of grey one way for sites. Sites admins still not sure how to fix their perfsonars.

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.

Technical Update

(to include any Token News, Data Challenge work, “New” OS experiences)
Future need to investigate ESCAPE.

Duty Report

NTR

Some discussion of the Glasgow APEL ticket.

Tickets

[UK GGUS tickets](#)

33 Open UK tickets this week. Generally GGUS is looking well pruned.

BRISTOL

https://ggus.eu/index.php?mode=ticket_info&ticket_id=154227

This Bristol CMS ticket is stalled, but I think on the CMS side. Could do with a prod.

GLASGOW

https://ggus.eu/index.php?mode=ticket_info&ticket_id=152780

I think the core problem of this apel alarm ticket is being zeroed in on, with one CE drastically over reporting. It would be interesting to hear what the result is. Just discussed.

IPv6 Tickets

https://ggus.eu/index.php?mode=ticket_info&ticket_id=131612 (BIRMINGHAM)

https://ggus.eu/index.php?mode=ticket_info&ticket_id=131611 (GLASGOW)

https://ggus.eu/index.php?mode=ticket_info&ticket_id=131606 (LIVERPOOL)

https://ggus.eu/index.php?mode=ticket_info&ticket_id=131603 (RHUL)

Please can all these tickets get an update, even another null one would be useful.

Site News

NTR

AOB/Feedback for/from the PMB

Luke asks about Carbon-Black security software on nodes. Has anyone had any experience with this? <https://www.vmware.com/uk/products/whats-new/carbon-black.html>

No one does, Bristol testing on older nodes.

David C would like to stay in the loop about this. Also a reminder that if anyone has any similar conversations going on at their institutions he'd love to get in on it.

-Action on Luke to keep us apprised on his progress.

Actions/Decisions from This Meeting

Carbon-Black action on Luke.

Chat Window:

11:03:00 From Luke Kreczko To Everyone:

Yay, discord :)

11:03:31 From DarthMarianne To Everyone:

I changed my name to my discord name.

11:03:43 From DarthMarianne To Everyone:

best of both worlds

11:05:20 From Luke Kreczko To Everyone:

debugging is like a complex puzzle game

11:10:54 From Daniel Traynor To Everyone:

gpu story sounds familiar

11:13:13 From Luke Kreczko To Everyone:

so jobrouter stripped it?

11:16:01 From Whitehouse, Dan To Everyone:

Looks like we are on 9.0.9

11:16:28 From Luke Kreczko To Everyone:

thanks!

11:17:44 From Whitehouse, Dan To Everyone:

```
START = ($(NodeFSOK) && $(NodeUsersOK) && !$(SafeDrainNode) &&  
ifThenElse((MY.GPUs > 0) == TRUE, (TARGET.RequestGPUs != undefined) &&  
(TARGET.RequestGPUs > 0), TRUE))
```

11:17:52 From Whitehouse, Dan To Everyone:

That's our current START if its useful to anyone.

11:18:13 From Chris Brew - STFC UKRI To Everyone:

@Luke. My 8.8 condo allows for more than one GPU per host, but they've all been the same.
Do you have heterogeneous ones?

11:18:41 From Whitehouse, Dan To Everyone:

And we have something like this too:

11:18:42 From Whitehouse, Dan To Everyone:

```
# GPU Config
MACHINE_RESOURCE_INVENTORY_GPUS=$(LIBEXEC)/condor_gpu_discovery -extra
-properties $(GPU_DISCOVERY_EXTRA)
ENVIRONMENT_FOR_AssignedGPUs=GPU_DEVICE_ORDINAL=/(CUDA|OCL)//
CUDA_VISIBLE_DEVICES=/CUDA//
ENVIRONMENT_VALUE_FOR_UnAssignedGPUs=10000
USER_JOB_WRAPPER = /usr/bin/job-wrapper.sh
```

11:18:58 From Whitehouse, Dan To Everyone:

Our job wrapper simply unloads and reloads the nvidia stuff.

11:23:55 From Luke Kreczko To Everyone:

@Chris: I will check my notes

11:25:28 From Luke Kreczko To Everyone:

I have marked HTCondor 8.9.13 as the first useful release for GPU due to:

> HTCondor now detects instances of multi-instance GPUs.

<https://htcondor.readthedocs.io/en/latest/version-history/development-release-series-89.html>

11:26:23 From Luke Kreczko To Everyone:

ok, so multi-instance GPUs not multiple GPUs

11:27:22 From Chris Brew - STFC UKRI To Everyone:

Yes, I _think_ that's for cards like the A100 that can be partitioned for different concurrent users.

11:27:30 From Luke Kreczko To Everyone:

yup

11:27:45 From Chris Brew - STFC UKRI To Everyone:

But that doesn't really work until some point in the 9.X series

11:28:04 From Luke Kreczko To Everyone:

ok, that's important info

11:29:08 From Luke Kreczko To Everyone:

I think our way format will be a dockerized & tagged deployment for HTCondor 9.X

11:29:58 From Luke Kreczko To Everyone:

and we will try a few

11:36:41 From Daniel Traynor To Everyone:

If it works at CNAF should work at QM with similar config

11:59:18 From Chris Brew - STFC UKRI To Everyone:

Have to drop out now, sorry.

12:00:48 From Winnie Lacesso To Everyone:

Have to leave now - SORRY!

12:05:14 From Terry Froy To Everyone:

I've just kicked ours (QM) to update to 4.4.2.

12:06:48 From Alessandra Forti To Everyone:

Alastair confirmed he is going to talk to the cloud team

12:14:01 From Luke Kreczko To Everyone:

I have some AoB regarding this security software: Carbon-Black

(<https://www.vmware.com/uk/products/whats-new/carbon-black.html>)

Context: UoB IT Services want to run this security software on EVERY node.

It can have some bad behaviour for highly forked processes - adding lots of CPU use.

Has anyone experience with running this on grid worker nodes?

12:19:03 From Daniel Traynor To Everyone:

There's a CHEP paper in it

12:19:17 From Terry Froy To Everyone:

https://www.reddit.com/r/sysadmin/comments/l7w5io/has_anyone_used_vmwares_carbon_black/

12:20:45 From Alessandra Forti To Everyone:

It looks like this is the only on premises product

12:20:46 From Alessandra Forti To Everyone:

<https://www.vmware.com/products/endpoint-detection-and-response.html>

12:21:01 From Alessandra Forti To Everyone:

It'd be interesting to know what it really does

12:21:43 From Alessandra Forti To Everyone:

At any rate it might really be detrimental because it takes nodes offline if it sees something suspicious

12:22:12 From Alessandra Forti To Everyone:

I can imagine our chaotic workflows might bring everything down 😊

12:22:24 From Luke Kreczko To Everyone:

yeah, that is my concern

12:24:53 From Terry Froy To Everyone:

It could be interesting to run it on a small subset of worker nodes, purely as an experiment, to see how it impacts on production Grid workloads *and* whether the centralized logging infrastructure operated by the central IT folk is able to cope.

12:25:22 From Luke Kreczko To Everyone:

we've overwhelmed the latter before (network monitoring)