

GridPP Operations team and sites meeting 21st November 2017

Attended

Alessandra Forti (AF), Andrew Washbrook (AW), Chris Brew (CB), Daniel Traynor (DT), Elena Korolkova (EK), Gareth Roy (GR), Gordon Stewart (GS), Govind (GS), Jeremy Coles (JC), John Hill (JH), John Kelly (JK), Kashif Mohammed (KM), Linda Cornwall (LC), Matt Doidge (MD), Paige Lacesso (PL), Raul Lopes (RL), Peter Gronbech (PG), Raja Nandakumar (RN), Rob Currie (RC), Robert Frank (RF), Sam Skipsey (SS), Steve Jones (SJ), Teng Li (TL), Vip Davda (VD)

Experiment problems/issues

LHCb

LHCb computing workshop this week

RN: - Running normally across to the Grid - Glasgow (and other sites) has problems with LHCb user who ran multi-threaded jobs by mistake. Strongly worded email sent and user put on low priority list - Exercise to reprocess last couple of years data (?)

CMS

Site Readiness reports:

- T1:
https://cms-site-readiness.web.cern.ch/cms-site-readiness/SiteReadiness/HTML/SiteReadinessReport.html#T1_UK_RAL
- T2:
https://cms-site-readiness.web.cern.ch/cms-site-readiness/SiteReadiness/HTML/SiteReadinessReport.html#T2_UK_London_Brunel

Current CMS GGUS tickets:

- https://ggus.eu/?mode=ticket_info&ticket_id=131565 (RALPP)
- https://ggus.eu/?mode=ticket_info&ticket_id=131909 (RALPP)
- https://ggus.eu/?mode=ticket_info&ticket_id=131987 (Bristol)
 - Raised this morning - segmentation violation

IPv6 status for CMS: <https://twiki.cern.ch/twiki/bin/view/CMS/IPv6Status4Sites> - Dont seem to care yet about T3

Singularity discussion

- DB: CMS requires Singularity to be deployed by Feb 2018

- GR: Queried version of CVMFS and version of singularity expected at sites? 2.2 version on EPEL is out of date, security patches. Concerns for deploying on sites
- DB: No comments from CMS on version just that need to process SL7 and SL6 need singularity. Ask which version have in mind and respond next week
- AF:
 - The WLCG baseline is 2.3.2 and already in WLCG repo but not in EPEL due to hold up resolving patched. Will be discussed tomorrow in WLCG working group
 - CMS want to put bind parts in the image and --contain which would not work on CentOS6 will be discussed
 - WLCG container working group (created in May) but never met until now. Alessandra to present ATLAS view, Dave Dykstra to present CMS view
- GR: Does not have overlayFS bind mounts have to be all present, "--contain" without overlayFS causes problems for CentOS6/SL6
- CB: Require CentOS7 workflows on CentOS6 sites but there is some confusion on message. Probably EL7 *or* singularity
- AF:
 - Claim only a handful of bind mounts are needed only a handful and sites are similar.
 - Site have CVMFS Only need 3 directories (CA certs, CVMFS, batch system job dir)
- GR: Possible path problems with certificate chains. /tmp at Glasgow is a logical volume
- AF: binding mount causes problems with many sites if they offer their own solution. Prefer overlay or "pwd" approach.

ATLAS

EK: - QM ticket - Transfer degradation, trans-atlantic transfers look problematic - DT raised internal ticket with QM network - RALPP 139182 - Transfers affected by power cut, still recovering from this - Sunday there as a DB problem - Italian T1 incident - files lost - ATLAS cleaning AGIS next year planning to switch to new information system - AD emailed TB-SUPPORT on 15/11 ATLAS plans for multiple Panda queues to single Panda queue (UCORE) - One example is ECDF with 2 DDM endpoints

PG: - Requests change to ATLAS monitoring measuring a non-existent IPv6 endpoint. How to remove from ATLAS monitoring setup? Who to speak to or to ticket?

Other

No updates. Oversight committee report in two weeks so will need update before then.

GridPP DIRAC status

Nothing needs to be followed up this week.

SKA VO - delay in GridPP DIRAC SAM test jobs being processed at Manchester (last successful 9th November). Worth checking.

Meetings and Updates

General updates

- CHEP 2018 abstract submission is open.
 - travel policy as previous years (50% GridPP /50% local institute/alternative funding)
- Germany is leaving EGI from January 2018. GGUS will continue to be provided for WLCG.
- TB-SUPPORT/general discussions:
 - ATLAS PanDA Queue consolidation - UCORE. The VAC question....
 - Postpone VAC discussion until next week with AM is present
 - AF:
 - VAC cannot work with UCORE, have to be separated
 - Question on underlying concept, HTcondor works differently. Site set up for MCORE requires 2 queues (1 SCORE, 1 MCORE) - would not work for UCORE for ATLAS
 - GRIDPP2: Deployment and support of SRM and local storage management
 - SS: Glasgow not publishing XRootD endpoint but publishing in GLUE2 should be fixed.
 - DB: strange bug drop srm and publish to XRootD no matter what
 - Waiting on XRootD update at ECDF
- CERN Mattermost chat - concept and groups
 - Raise awareness of service
 - GR:
 - Scotgrid used slack but doesn't have message memory for free subscription so switch to mattermost
 - Seems to work as good as Slack as part from long hashname for users with CERN lightweight accounts
- LHCb multi-thread user jobs
- Mounting EOS outside of CERN. (See [this article](#)).
- The next LHCOPN/LHCONE meeting will be at RAL 6-7th March 2018
 - Connect RAL, for Tier-2s watch this space
- There is a CERNVM users workshop 30th Jan to 1st Feb.
- Minutes from Monday 20th Nov WLCG ops meeting are available.
 - Issues with Oracle DB
 - All mention CNAF incident

WLCG ops coordination

- The next WLCG ops coordination meeting will be on 7th December. Do we have any topics we would like covered?
 - Topics include Singularity
 - No other topics raised

Tier-1 status

A reminder that there is a weekly Tier-1 experiment liaison meeting. Notes from the last meeting [here](#)

- There was a brief power outage at the RAL site on Monday 20th November. Most services have now recovered.
 - Approx 2 minute outage
 - Diesel generators did not start automatically, investigation underway
 - Lost some batch worker nodes, some head nodes and castor worker nodes
- The FTS service (both instances) is experiencing problems which are not fully understood.
- The database behind the FTS-production instance is being updated today.
- Re-distribution of data in Echo onto the 2015 capacity hardware is now complete. There are now 8PBytes of usable space in Echo. However, there is some data re-balancing to be done before upping the quotas so the VOs can make use of the extra space.
- Castor was be unavailable for a few hours on Tuesday 14th Nov for patching of the nodes.
- The LHCb Castor SRMs were successfully updated on Wednesday 15th Nov. This new version (2.1.16-18) respects the user-specified bringOnline timeout rather than ignoring it and defaulting to 4 hours. It is planned to upgrade the other SRMs on Tuesday 5th December.
- Both Production and Test FTS services have IPv6 enabled (i.e. are now dual-stack).

Storage and data management

- SS: Mostly XRootD proxy cache discussion with full rundown of activities at each of the sites including Chris' work at RAL and progress at Edinburgh. Things moving forward.
- CB: Noticed by accident that server with purely pass through proxy had 4x throughput than caching. New plans to systematic test proxies.

Tier-2 Evolution

Accounting

Documentation

- VOMS servers located at CNAF have been knocked out due to some flooding. It is not known when they might be recovered. Hence, at present, there are no VOMS servers at all for both planck or ipv6.hepik.org. And a single new VOMS server has been provided for enmr. See the VOMS records in the Approved VOs document for details.
- Sjl: Limited impact due to low activity from these VOs

Interoperation

There was an EGI Operations Meeting today ([agenda](#))

- UMD3 deprecation: WMS to be dismissed by end of year, EGI Operations planning sending a note about end of support for UMD3 in line with WMS deprecation
 - DB: MICE using DIRAC. T2K.ORG still using grid?
 - DT: T2K working on that at QM, in Japan at the moment but will check. Using direct submission rather than WMS or DIRAC
- Preview repository new versions released on 15/11
- New weights for the NGIs average A/R values, based on Computation Power

- webdav probes in production
 - Ticket for UKI-NORTHGRID-MAN-HEP (solved)
- Storage accounting deployment
 - Ticket for UKI-NORTHGRID-SHEF-HEP (solved)

Monitoring

On-duty

Rollout

- IPv6: https://www.gridpp.ac.uk/wiki/IPv6_site_status (up-to-date as of November)
- Batch systems and WN moves to SL7/CentOS7:
 - https://www.gridpp.ac.uk/wiki/Batch_system_status (added column for date of last update)
 - Reminder for sites to complete last column in batch system

Security

- A few sites have showed up in pakiti with high risk vulnerabilities (CVE-2017-1000364). If you haven't been prodded about this already then expect an email from someone in the security team soon
 - MD:
 - A few sites show up in Pakiti with high risk vulnerability
 - For shared sites this can be problematic, ask when expected to be resolved
- LC:
 - Security discussion list on JISCMAIL not used for a year. Be good to get this going. Everyone should have got a prompt email - anyone did not get this?
 - At CSIRT F2F (Helsinki) discussed how to go forward with EOSC (EGI/EU.dat) collaborations. For example, for vulnerability handling have EOSC list

Services

- perfSONAR: No news.
 - KM and PG going to perfsonar course
 - MD: suggest knowledge transfer at next hepsysman
- GridPP DIRAC: DPM & xrootd version. Upgrade due 16th November.
- RIPE ATLAS probes: Status & tests TBC
- The next LHCOPN/LHCONE meeting will be at RAL 6-7th March 2018. The UK position on LHCONE is being reviewed. RAL T1 will (likely) connect via the LHCOPN in the coming months. T2s watch this space!

Tickets

- [link to tickets](#)
- MD:
 - Number of Tier 1 tickets going stale and red, just needs some culling
 - BHAM http SAM tests - Mark soon going to upgrade DPM doubt it will happen before Christmas

- Sussex - webdav front end ticket - having trouble getting xroot working
- DT: site firewall issue

Tools

- The WLCG dashboard is available here: <http://dashboard.cern.ch/>.
- Please review the results on [Steve's test pages](#) .

VOs

- Tom Whyntie has requested (and been granted) access to the GridPP VO to get some pipelines working for large-scale processing and analysis of - MRI scans associated with the UK Biobank project.
- All VOs in the incubation page being prompted for updates by the end of November (required input for OC documents).
- QMUL (Steve L) is following up on the biomed MoU. GridPP want to be cited in research papers for the support our resources/sites provide.

Site updates

Manchester

- AF: Mostly working on singularity
- RF: Planning transition to Foreman, puppet3 to puppet4. Investigating ZFS performance issues.

Edinburgh

- AW: Feasibility testing of Cloud scheduler on ECDF openstack
- RC: Storage activity update xrootd, ZFS performance related tasks

RAL PP

- CB:
 - Working on doubling firewall bypass links
 - Gradually replacing SL6 service nodes with CentOS7 (ARC CE, ARGUS servers, BDII) thinking about C7 worker nodes.
 - Ipv6 perfsonar work

QMUL

- DT:
 - Commissioning new storage
 - Failover Lustre setup using ZFS, performance ok, commissioning it
 - New deployment based on Ainsible
 - Looking into proxmox for managed VM hosts setup

Sheffield

- EK: Move Worker Nodes onto SL7

Glasgow

- GR:
 - Containers development work
 - Volume driver from CERN
 - Evaluating systems for migration to our new datacentre. Using foreman
 - Chasing IPv6 with network team
 - ZFS performance
 - Accounting details from Condor

RHUL

- GS:
 - Network upgrades and Worker Node upgrades
 - Question on how to test setup of Force 10 S4810
 - Minimise downtime to get spare switch

Cambridge

- JH:
 - Moving WN to VAC
 - Pool nodes DPM 1.9
 - Tackle SE in week or two
 - XRootD to latest version

Oxford

- KM:
 - Working on BRO, attending SOC meeting
 - Planning DPM head node move to C7
- VD:
 - Working on ARC CE and worker nodes to SL7

Lancaster

- MD:
 - Xcache proxy work
 - move worker node to C7 then service nodes
 - Will attend SOC (remotely)
 - Trouble having SPEC test on C7 (using v1.1)

Liverpool

- SJ:
 - Upgrading storage to centos7 (?)
 - IPv6 progress depending on central computing group