

IRIS Operations Status

In preparation for the IRIS TWG Operations Meeting of Tuesday 6th August 2019.

Actions

MM.YY.##	Who	What
07.19.01	AS	Chase up at RAL end why LIGO file transfer-in is an issue
07.19.02	SdW	<i>Closed</i>
07.19.03	AS	<i>Closed</i>
07.19.04	AC	<i>Closed</i>
07.19.05	AS/AC	Andrew and Adrian to discuss update frequency of APEL stats
07.19.06	AC/JK/AS	Get access to DiRAC accounting data
07.19.07	AS	Investigate getting seedcorn access for GPU work (c.f. KNL issues at Cam)
07.19.08	AS/AMcN	identify any resource constraints on DUNE maximising usage on GridPP.
08.19.01	JY	We need to find a TWG rep for CTA

Activity Status

Activity	Status (dd/mm/yy [reported by] text)
EUCLID	<p>25/06/19 [Mark] SLURM as a service Up and running at Cambridge and RAL.</p> <p>9/07/19 [Mark] Now expanding at RAL. Agreed most details of how to create vacuum-style VMs to run at Manchester. Next step is to connect to VPNs and then start submitting slurm jobs.</p> <p>22/07/19 [Mark] – The Saas cluster is up and running at Cambridge at capacity. We hit issues trying to deploy at scale at RAL, and hence are not operational there while we await a fix (there is an odd issue with disk image generation that causes failures when trying to spin up more than a couple VMs at a time). Troubleshooting that problem and then me being on holiday meant we didn't get any progress on a Manchester partition.</p> <p>6/08/19 [Mark/AM] - Expect to be scaling at RAL from now on and then development at Manchester restarting next week</p>
Gaia	<p>25/06/19 Cambridge talking to Gaia – a development activity to investigate Spark workload</p>

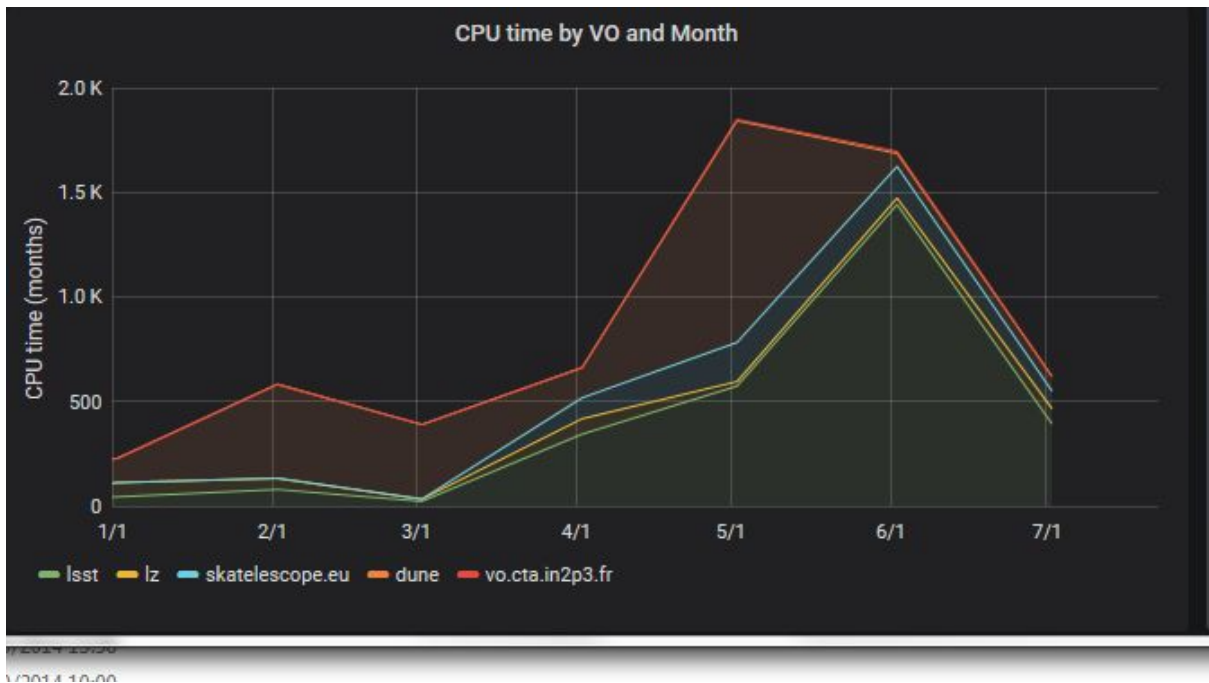
	<p>9/07/19 [AS] JY organised a meeting last Fri with Leicester, Edinburgh and Cambridge; there were discussions about how this development will be handled. This is expected to be reviewed in a couple of weeks time.</p> <p>22/07/19 Now have access to OpenStack at Cambridge</p> <p>6/08/19 Getting up to speed in deploying at Cambridge. Useful meeting more generally last week</p>
LSST	<p>25/06/19 Three distinct activities underway. Production work running on GridPP resources is going well (2000 processor months in May and June). LSST also working with Cambridge to deploy SLURM as a service. Jupyter notebook activity at RAL also underway.</p> <p>22/07/19 Grid use currently paused - a discrepancy (3%) noticed compared to NERSC Cori (running Knights Landing), which is being investigated. Cambridge stuff seems to be going OK, but status under SAFE not visible. Further, Richard McM has requested a small amount of GPU time, which is still to be confirmed. On RAL OpenStack, encountered issue, due to Mellanox drivers. This issue rectified with a driver, so Jupyter notebook work back on track.</p> <p>6/08/19 [GB] Three distinct activities are on-going:</p> <ol style="list-style-type: none"> 1. Grid-based simulations restarted briefly, at the beginning of August, but now paused again--as potentially serious issue detected in output from simulations. Expect no significant runs for the remainder of August. 2. We are still waiting for feedback on our request for access to GPU nodes on the IRIS system at Cambridge though, otherwise, no exceptions for that service. 3. We continue to see stability problems on RAL OpenStack system, which are significant enough to impact user experience. Working with Alex D to resolve these.
e-MERLIN	6/08/19 No update
SKA	<p>25/06/19 No update but modest usage in May and June on GridPP resources</p> <p>22/07/19 No update, but AS re-establishing link with SKA as Rohini has moved on within SKA. AMcN: there is now a PhD student who is doing lots of stuff now.</p> <p>6/08/19 More staff movement, but jobs still being run.</p>
ALMA	<p>9/07/19 (<i>Anthony Holloway</i>) The ATOMIUM project processing is underway and is now getting the larger data sets delivered from the ALMA telescope in Chile.</p> <p>22/07/19 No update</p> <p>6/08/19 Himem nodes at Manchester in heavy use</p>
LZ	<p>25/06/19 No update – GridPP workflow</p> <p>9/07/19 [DB] Still stuck getting their s/w for their latest MC data challenge. They have IRIS allowance for when DIRAC is working with the cloud again and their s/w is finished.</p> <p>22/07/19 No update</p> <p>6/08/19 Running on grid (~2700 jobs)</p>
CTA	<p>25/06/19 No update</p> <p>9/07/19 [JY] Julian Osborne has the token to get further help when he is ready to run stuff.</p> <p>6/08/19 No update</p>
aLIGO	<p>25/06/19 Working with Catalin Condurache to recommission Grid for LIGO</p> <p>9/07/19 Trying getting it running at RAL, although works elsewhere - jobs go to RAL but then file copy in fails.</p> <p>6/08/19 No update</p>
DUNE	<p>9/07/19 [<i>Peter Clarke</i>] Running well. About to begin a second reprocessing which will hopefully saturate available resources.</p> <p>23/07/19 [<i>Andrew McNab</i>] Starting reprocessing campaign and using (all?) available CPU resources. Would sites please let us know if you believe you should be running more, and we can look for bottlenecks.</p>

	<p>6/08/19 [Andrew McNab] Vcycle/Vac VMs demonstrated running DUNE production jobs (ie so we can backfill OpenStack with DUNE work). AAI + site/cluster naming policy has to be agreed with GlideinWMS team at FNAL. Reprocessing campaigns on going (up and down in volume.) Some limits at Manchester identified and removed: waiting for volume to go up generally to verify this. Need to repeat for other UK sites.</p>
DLS	<p>25/06/19 Contractor actively working for Frank von Delft to build SLURM clusters in OpenStack at RAL (and cambridge?). Also trying to make use of object store.</p> <p>9/07/19 SLURM work ongoing; having issues with data xfer from a data centre in Iceland Echo storage hasn't progressed. Has tooling for synching object-stores, to be determined whether this can be used with Echo storage.</p> <p>23/07/19 CR+RS: Things have moved along a bit, workaround found for Ansible, but with STFC to sort out why stuff isn't being spun up. Also discussions around use of S3 at RAL - not going to use Dynafed - instead building on top of native S3</p> <p>6/08/19 Galaxy workflow tool being used, directing work into STFC cloud. Working on VPN set up. Planning to do ECHO benchmarking, need to resolve some firewall issues first.</p>
PANOSC	<p>23/07/19 AS: PANOSC now has a storage allocation, pending an STFC recruitment</p> <p>6/08/19 no update</p>
ISIS	<p>25/06/19 iDaaS is in production for ISIS and busy now during data taking run. Using RAL OpenStack</p> <p>9/07/19 Appears to be working fine</p> <p>23/07/19 AS: RAL on top of this, beam run finished last week, onboarding more activities for the next run which is planned for early September.</p> <p>6/08/19 no update</p>
CLF	<p>25/06/19 Talks at fairly early stage</p> <p>9/07/19 Nothing to add</p> <p>23/07/19 Fraser B has been speaking to CLF regularly this week, looking at IDAAS for them</p> <p>6/08/19 Alex Robinson will be the CLF rep on the IRIS DB in the future - should he be in TWG too?</p>
CCFE	<p>9/07/19 [JY] On Cambridge things are running fine and they are ready for access to their GridPP and SCD allocations. RAL is probably the next one to go for. SdW will talk to SCD and then afterwards to JY and DB about using the other resources.</p> <p>23/07/19 AL: CCFE have been happily using DIRAC resources but nothing showing in SAFE. Don't need SCD resources yet.</p> <p>06/08/19 [AL]: Obtained access to OpenStack at IC today.</p>
JINTRAC	<p>9/07/19 JY to talk to Florian</p> <p>23/07/19 [AL] For the past few weeks have been working on updating JAMS (Jet Application Management System) to make use of PROMINENCE and hence use any IRIS resources (both cloud and DiRAC). Initial version almost complete and are running real test jobs on the SCD cloud and will start using the IC cloud and Cambridge DiRAC soon. Echo is being used for transient storage of input and output data and container images.</p> <p>6/08/19 no update</p>

Usage

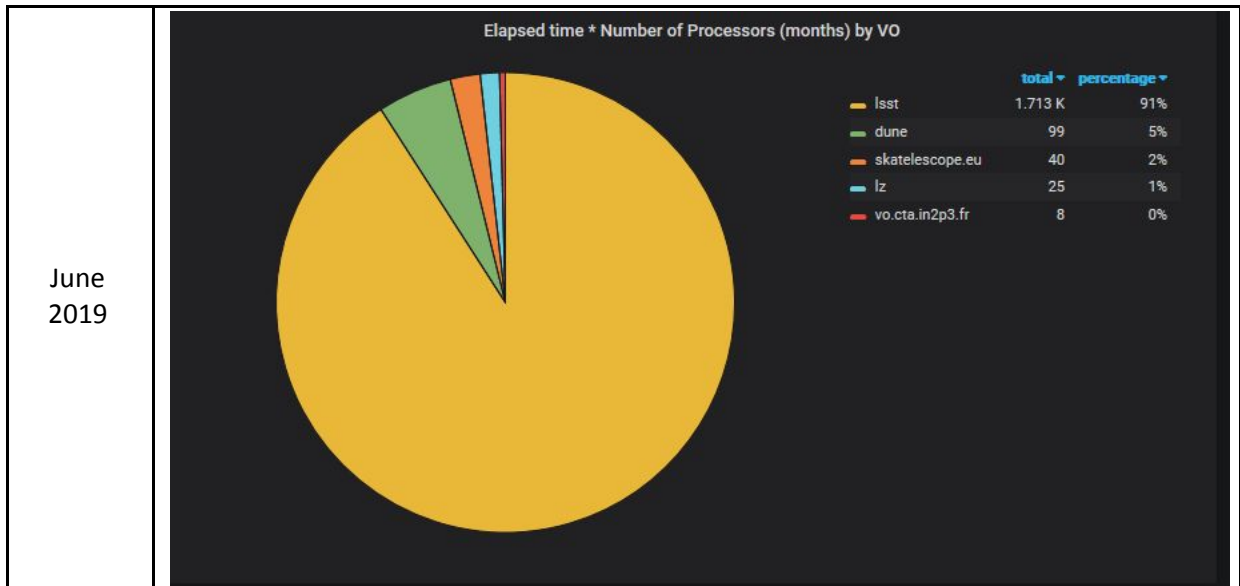
Accounting data still being manually updated monthly at present. Last run for this month was on 19th July, so July data does not yet represent the full monthly total. Intention is to automate and also provide SCD cloud accounting data by the end of July.

Year to date (stacked graph)




Monthly data breakdown.

Grid only accounting at present



JULY
2019



Infrastructure Status

We should also pick up here any operational issues with major IRIS components

Tickets/Support	23/7/19 [Andrew McNab] described current setup of ticket triage system.
RUCIO	
DIRAC	
Vcycle	23/7/19 [Andrew McNab] : Deployed version with improved scalability (multiple co-operating Vcycle instances.) Deployed multiprocessor DIRAC VMs for LSST (x4) and SKA (x8). Testing CernVM4 images (~CentOS7) for SKA etc. Next: restart EUCLID VM work, and get mixed VM tenancies at Cambridge and RAL working (currently using donated OpenStack capacity at CERN for testing.) 6/08/19 [Andrew McNab] : ready to deploy CernVM4/CentOS7 images for all GridPP DIRAC users/sites including OpenStack backfilling (SKA, LSST, LZ, ...). VMCondor simple batch system (almost) ready for testing by interested users.
AAI Proxy	6/08/19 [JJ] being registered with UK AMF to “bless” SRTFI/etc
APEL	6/08/19 [Adrian Coveney (via JK)] : We have some RAL cloud data loaded for August - we have further data waiting but are doing sanity checks on it first before we include it in the portal. Most of the kinks in the messaging have been sorted, so I will be moving on to gathering data from the other IRIS cloud sites.

Open Tickets

Latest on tickets

[GGUS Link for iris.ac.uk VO](#)