

IRIS Operations Status

In preparation for the IRIS TWG Operations Meeting of Tuesday 3rd September 2019.

Actions

MM.YY.##	Who	What
07.19.01	AS	Chase up at RAL end why LIGO file transfer-in is an issue [DC] unresolved - PH to contact support@gridpp.rl.ac.uk to provide information.
07.19.05	AS/AC	Andrew and Adrian to discuss update frequency of APEL stats
07.19.06	AC/JG/PJ/AS	Get access to DiRAC accounting data 3/9/2019 CLOSED - no progressing
07.19.07	AS	Investigate getting seedcorn access for GPU work GB: needed by October
07.19.08	AMcN	Identify any resource constraints on DUNE maximising usage on GridPP. 3/9/2019 CLOSED - moved to DUNE below
08.19.01	JY	We need to find a TWG rep for CTA 3/9/2019 CLOSED - CTA's responsibility to find someone
09.19.01	PJ	Get a handle on the allocation process for DiRAC

Activity Status

Usage info format xxx 6130 core months last calendar month - (this month - extrapolate to full month) - [allocation]. See infrastructure section to see accounting limitations. This includes account information from SCD, GridPP and DiRAC(OpenStack Only) (excludes DiRAC SLURM)

Activity/usage	Status (dd/mm/yy [reported by] text
EUCLID 80 (994) [3100]	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. 6/08/19 [Mark/AM] - Expect to be scaling at RAL from now on and then development at Manchester restarting next week

	<p>13/08/19 [Mark] - Work at RAL on scaling up continues, issues with the creation/attachment of volumes is still preventing scripted deployment at scale. Deployment at Manchester this week is somewhat contingent on if I'm selected during jury service today...</p> <p>20/08/19 - There have been problems attaching volumes at RAL. A workaround for the issue is in progress and we should be fully online there within a week. Work on onboarding at Manchester will begin in earnest later this week.</p> <p>3/9/12 - RAL partition is fully operational. Work on getting a partition running at Manchester has begun in earnest.</p>
Gaia 0 (11) [0]	<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> <p>20/08/19 [NH] Dev team familiarising with Spark SW stack and working on scripted deployment using ansible</p> <p>3/9/19 [NH] Nothing new to report at this time.</p>
LSST 845 (318) [1312]	<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. <p>20/08/19 [GB] A modest amount of DESC-related simulation work has restarted on GridPP, though throughput hampered by issue with LSST VO certificate. Recent RAL power outage was very unfortunate timing, as coincided with a presentation of the LSST Jupyter Notebook service to interested users. Also, has raised the need for us to think carefully about service availability, in light of some of the real-time telescope feeds we operate. LSST@Cambridge team have begun using cluster time on Cambridge site, plus have confirmed they will be looking to use GPU seedcorn funding from October onwards, if possible.</p> <p>3/09/19 [GB] No exceptions. Grid use is low just now: we finished a few test simulations and are awaiting feedback. No problems with cloud infrastructure (at RAL) over the last couple of weeks. Use of DiRAC (Cambridge) seems to be ramping up.</p>
e-MERLIN 0 (59) [96]	<p>20/08/19 No update</p>
SKA 725 (2805) [1920]	<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> <p>20/08/19 [AMcN] Rohini still involved in regional centre work and hence involved with IRIS.</p> <p>03/09/19 [JMT] I said at the last Operators meeting that I would explore SKA-SDP use for IRIS. The SDP current requires access to specific GPU resources as well as DASK (Scalable Data Analytics Platform https://dask.org/). These are currently supported on the STFC-AlaSKA SDP Prototype Platform and as such are continuing to use the AlaSKA with "best efforts" support to meet a deadline for the SKAO in Sept. I have also had a chance to talk to Ben Mort at Oxford who is responsible for the Systems Integration Prototype part of the SDP. This has now been migrated over the Summer from a</p>

	Docker-Swarm to a k8s environment. I said I would explore use of IRIS resources for this, with a mitigation of continued use of AlaSKA.
ALMA 0 (0) [0]	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. 22/07/19 No update 6/08/19 Himem nodes at Manchester in heavy use 20/08/19 No further update - high profile stuff planned for the near future
LZ 1103 (1246) [300]	22/07/19 No update 6/08/19 Running on grid (~2700 jobs) 20/08/19 There is now a plan - anticipating over-running their quota, starting to run soon. 03/09/19 (DB) LZ has used 350k CPU hours (DIRAC accounting, not normalized to HEPSPEC) on the Grid in the last month. That's ~470 jobs running 24 h/day averaged over the month.
CTA 0.4 (0) [285]	25/06/19 No update 9/07/19 [JY] Julian Osborne has the token to get further help when he is ready to run stuff. 20/08/19 No update 03/09/19 (DB) Luisa A (in2p3) is back from holiday and has been in contact with Matt D (Lancs) to hunt down some configuration issues. Julian has reiterated the need for using the resources now.
aLIGO 0 (0) [270]	25/06/19 Working with Catalin Condurache to recommission Grid for LIGO 9/07/19 Trying getting it running at RAL, although works elsewhere - jobs go to RAL but then file copy in fails. 6/08/19 No update 20/08/19 No new progress - need to make contact with Catalin's replacement. 03/09/19 Starting to resume debugging with James and RAL helpdesk
DUNE 812 (269) [1000]	23/07/19 [Andrew McNab] 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. 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. 20/08/19 [Andrew McNab] Second reprocessing campaign starting ?later this week? Will again review usage of IRIS sites and contact people in case less than expected. Jobs will request 3500MB of memory although they will typically use less than this (~3000MB?) This needs to be in the GlideinWMS config at FNAL and accepted by your CE. 03/09/19 Making progress in pinning down the config for the RAL resource issues - in contact with Fermilab
DLS 1561 (1256) [2100]	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. 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. 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


	<p>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> <p>20/08/19 [RS] Currently still only use from within FvD group (DLS), running Openshift. Chris wants to move to a native Kubernetes deployment to control access and applications. Issues with spinning up VMs seems to have been solved, but we will monitor. Some issues potentially remain with VMs becoming unresponsive, but we will monitor the status of this over the next week whilst we run benchmarking against other cloud-instances of one of our workflow applications (Galaxy)</p> <p>03/09/19 No update</p>
PANOSC 0 (0) [0]	<p>23/07/19 AS: PANOSC now has a storage allocation, pending an STFC recruitment</p> <p>3/9/19 AS: Work underway at STFC to interface PANOSC to Echo storage system. Have spoken to PANOSC PI about who will update IRIS-TWG with PANOSC status. details TBC</p>
ISIS 902 (1183) [175]	<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> <p>20/08/19 [AS] Fraser gave a good talk to SCD yesterday about DaAAS and the progress with ISIS</p>
CLF 2 (141) [68]	<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> <p>20/08/19</p>
CCFE 19 (1) [1250]	<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> <p>20/08/19 [AL]: Carrying out performance testing on STFC & IC clouds, DiRAC and comparison to a local HPC cluster. Note that most jobs to be run in this allocation are multi-node MPI jobs.</p> <p>03/09/19 [AL]</p> <ol style="list-style-type: none"> 1. Continuing tests in working out how best to run applications requiring low-latency networking (LAMMPS, VASP) on IRIS resources without low-latency networking. Found CPU performance variations when running multiple jobs on different VMs at the same time at IC. Simon Fayer resolved this with some OpenStack adjustments. Also experienced CPU performance issues when running multiple MPI jobs within the same VM, also resolved by Simon Fayer (thanks!) 2. Two users are running real work on IC. 3. Found substantial CPU performance variations across VMs on the STFC cloud when running multiple processes or threads within a VM. Some VMs are over 120 times slower than others (!) Submitted a ticket about this. 4. IRIS DiRAC August allocation fully used up by 2nd September. More has now been allocated.

JINTRAC	9/07/19 JY to talk to Florian
0	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
(0)	6/08/19 [AL] no update
[200]	03/09/19 [DB] Just for the record: JINTRAC had access to the IC IRIS Cloud for about a month. 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.
	03/09/19 [AL] Still waiting for users to try running their jobs on IRIS resources.

Usage


Accounting data still being manually updated monthly at present. Full data for August is now available. Due to the very limited amount of data September has not been forecast. Intention is to automate and also provide Cambridge data in the near future.

Year to date (stacked graph)



Monthly data breakdown.

Grid plus cloud



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	<p>20/8/19 [Ian Johnson]: Operation of the Rucio instance at RAL has been slightly disrupted by recent problems at RAL (including a power outage affecting some OpenStack hosts this Sunday, 18/08). The functionality of the cluster nodes has been recovered quickly by the SCD Fabric team. and the service returned to users' satisfaction, in all cases.</p> <p>03/09/19 [Ian Johnson] Update of RAL cluster nodes to latest Rucio software version delayed - database schema upgrades do not work on the PostgreSQL server. Conferring with Rucio developers on this issue. (Operation of Rucio at RAL is unaffected.)</p>
DIRAC	20/8/19 [DB] Working on DIRAC -> Workflow clusters - proof of concept
Vcycle	<p>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.)</p> <p>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.</p> <p>20/08/19 [Andrew McNab]: VMCondor ready for testing and using Manchester, Cambridge and donated CERN capacity. See https://github.com/iris-ac-uk/iris-vmcondor/wiki for Users' and Sites' Guides.</p>
AAI Proxy	<p>6/08/19 [JensJ] being registered with UK AMF to "bless" SRTFI/etc</p> <p>20/08/19 [JensJ] successfully registered with federation but doesn't like authentication attribute from STFC's IdP. Jens & Tom investigating. Try it yourself on https://iris-iam.stfc.ac.uk/ and let us know if it does/does not work (delete as approp.)</p> <p>03/09/19 [JensJ] I-it lives! Some usability issues need resolving before going full prod - & needs docs for service integrators</p> <p>[TomD] Currently working with Alex Dibbo to test and verify usage of groups for authz to SCD cloud, with the idea being group per science activity - authn however proven and works. Opening process to look at Dynafed/Echo storage connections.</p> <p>Work to be presented at FIM4R meeting at FermiLab 12/09.</p> <p>Aware of some issues (AUP, IdP values overwriting) and looking at these.</p>

APEL	<p>Feedback can be sent to: apel-admins@stfc.ac.uk</p> <p>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.</p> <p>20/08/19 [AC]: All ral-cloud data included (though also shows DAFNI project's usage at the moment - will be filtered out eventually). Messaging tweaks needed before we roll out automatic updating of the portal and some more changes might be required to integrate Cambridge etc.</p> <p>03/09/19 [AC]: DAFNI data now filtered out. Total of SCD + GridPP provision (i.e. what's currently accounted for) shown as green line. Messaging tweaks should be rolled out this week with automatic updating following subsequently.</p>
------	--

Open Tickets

Latest on tickets

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