

IRIS Operations Status

Preparing for the IRIS TWG Operations Meeting of Tuesday 12th November 2019.

Actions

mm.yy.##	Who	What
07.19.07	AS	Investigate getting seedcorn access for GPU work GB: needed by October 15/10/19 AS has contacted mark Wilkinson 26/11/19 [GB] w/ Richard McM
09.19.01	PJ	Get a handle on the allocation process for DiRAC [Complete] TO BE CLOSED
10.19.01	PJ	Setup a resource review meeting [Will schedule with RSAP] In-hand TO BE CLOSED
10.19.03	IC	IC to find out what this is about and then talk to Brian about IPv6 access RESOLVED: TO BE CLOSED
10.19.05	AC	Chase up the problem Daniela is having with Caso [26/11/19] No news back from developers yet.

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)

Note: We have some concerns that the monthly usage data is inaccurate and until we can validate it it has been removed from the report. Data column in black.

Activity/usage	Status (dd/mm/yy [reported by] text
EUCLID [3100]	1/10/19 - Running fully at RAL and Cambridge, but we're hitting the limits of our shared storage there and need to sort out a solution to keep going full on. Still awaiting the "permanent" gateway/IAL/ceph node at Manchester and can't get started there until that's available. 15/10/19 - Expanded Saas on RAL SCD, now up to 1728 cores. All running well at Cambridge. Just starting on the long-live Manchester node to get the VacuumVM stuff going.

	<p>29/10/19 - All running well at RAL and Cambridge. A few initial bits of work done to get started at Manchester.</p> <p>12/11/19 - All running well. Manchester work delayed due to other commitments.</p> <p>26/11/19 - All running well.</p>
Gaia [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>1/10/19 [NH] Not much to report. Dev team diverted by other projects during September, but have ansible-scripted deployment working locally and are planning to start rolling out on OpenStack next week.</p> <p>15/10/19 Nothing to report</p> <p>29/10/19 Now working closely with JG &co on OpenStack deployment details</p> <p>11/11/19 [DM] Working on automated deployment of Spark; telecon tomorrow with JG & co to discuss back-end file system options and other implementation details.</p> <p>26/11/19 [DM] Learning how to use OpenStack Magnum to deploy Kubernetes clusters.</p> <p>Manual deployment of Spark/Yarn/HDFS up and running; successfully exercised with small dataset in conjunction with simple scala and python example scripts.</p>
LSST [1312]	<p>26/11/19 [GeorgeB]</p> <ul style="list-style-type: none"> - The next DESC campaign began on Thursday 21st. We have approx. 4M core hours of simulations we would like to complete before the end of the year. Unfortunately, while throughput is good, approximately 40% of our jobs fail after a few hours. Initially, this looked to be related to a DiRAC watchdog process incorrectly identifying some jobs (with low CPU utilisation) as stalled. However, the situation seems more complicated, with different versions of Singularity running at different sites causing problems. Work continues to maximise throughput, but also reduce failure rate. - The Lasair service, running on RAL OpenStack, was yet again impacted by a DNS failure. This looks to be a repeat of the problem we reported previously. Alex D advised that it looks to be a particular network element, which has now been scheduled for replacement. - No exceptions for simulations running on DiRAC @ Cambridge.
e-MERLIN [96]	<p>20/08/19 No update</p> <p>15/10/19 Nothing to report</p> <p>29/10/19 Nothing to report</p>
SKA [1920]	<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 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.</p> <p>1/10/19 No update (SKA meeting)</p> <p>15/10/19 Nothing to report</p>

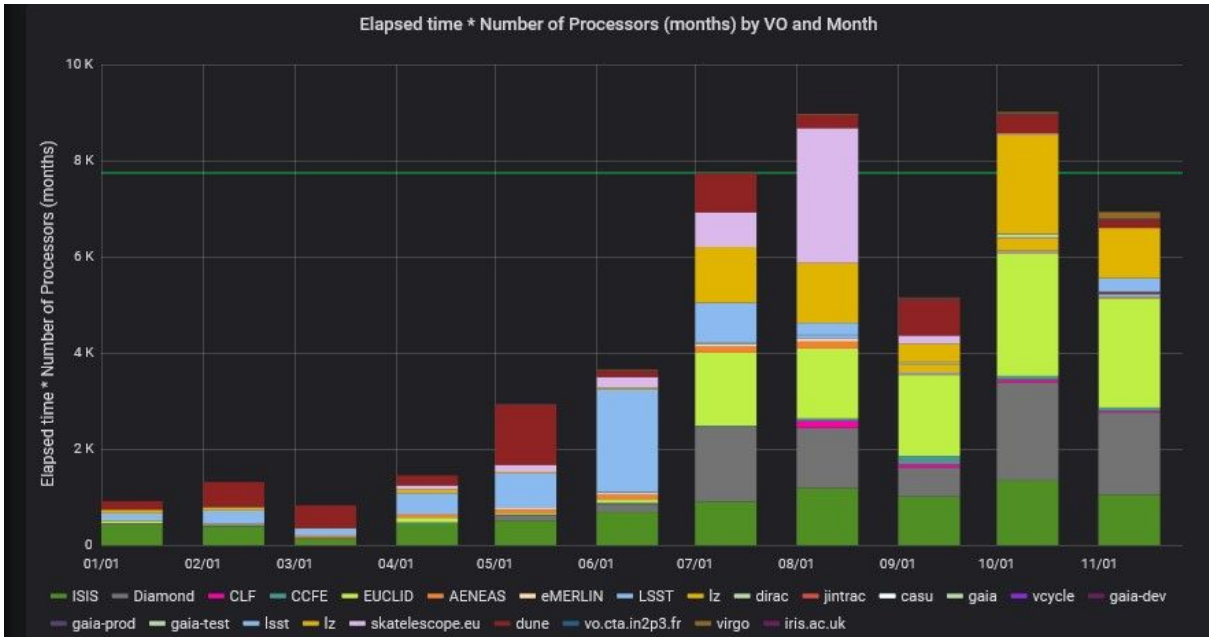
	<p>29/10/19 [RJ] looking into SKA SRC test payloads, created docker image that packages DIRAC client and a test payload that checks for site suitability for an SKA SRC payload. Debugging Miguel's RMSynthesis use case, trying to reproduce the error that is seen when DIRAC jobs submitted in bulk. Attempting to unpack an astronomy application docker image on CVMFS.</p>
ALMA	<p>6/08/19 Himem nodes at Manchester in heavy use 20/08/19 No further update - high profile stuff planned for the near future</p>
[0]	<p>15/10/19 Nothing to report 29/10/19 Nothing to report</p>
LZ	<p>15/10/19 (DB) LZ running production on Grid and Imperial IRIS cloud. 29/10/19 (DB) LZ running production on Grid and Imperial IRIS cloud. Some records now in the accounting portal after emailing them to Adrian. Accounting issues remain unsolved.</p>
[300]	<p>12/11/19 (DB) LZ running production on Grid and Imperial IRIS cloud. Some records now in the accounting portal after emailing them to Adrian. Accounting issues remain unsolved. (Yes, it's cut and paste.) 26/11/19 (DB) Same as before.</p>
CTA	<p>01/10/19 (DB via Matt D & Luisa) <u>Matt</u>: A quick check shows that we've run 488 CTA jobs over September, and only 3 failed (according to the batch system), with a little over 2000 hours of wallclock time clocked up. <u>Luisa</u>: Lancaster is well configured now and jobs run successfully, however there are no large productions currently running in this moment, so what you see are essentially some users jobs or small test productions. As soon as we will have productions running, you should see an increase of activity, but this will happen probably in a few months. At the end of the year, we will have a teleconf with all sites supporting the CTA VO, where we will show the plans for the next year productions.</p>
[285]	<p>12/11/19 [MD] Daniela is in contact with them</p>
aLIGO	<p>03/09/19 Starting to resume debugging with James and RAL helpdesk 17/09/19 Only just send email to James and RAL helpdesk. Enquired about sending dedicated UK LIGO cluster accounting info to IRIS to better track overall usage shift from dedicated to grid computing. I notice that LIGO usage at RAL still showing as "virgo".</p>
[270]	<p>01/10/19: From James Adams: The real underlying issue is that ligo jobs are being started as the root user rather than the ligo user (but without root privileges) which causes the wrong (aka no) environment to be sourced. [...] Investigations continue. 15/10/19 [PH] Waiting for RAL to make changes to pilot environment. 29/10/19 [PH] Still waiting for RAL to fix environment. Work ongoing with LIGO and OSG to provide a clean environment independent of site configuration. 12/11/09 [PH] No news from RAL. LIGO work to clean environment (I think) complete. Setup scripts to upload Cardiff cluster accounting using APEL, waiting for credentials. 26/11/09 [PH] Still no news from RAL. Still waiting for credentials to upload accounting.</p>
DUNE	<p>03/09/19 Making progress in pinning down the config for the RAL resource issues - in contact with Fermilab</p>
[1000]	<p>15/10/19 [AM] Successfully running productions as they appear but some sites are seeing jobs taking more processors than allocated. FNAL DUNE team investigating and limiting problematic jobs to FNAL for now.</p>

	<p>12/11/19 [AM] Processor-stealing production jobs should be fixed now. Any jobs doing it now are from naughty users: please let us know and they will be dealt with appropriately... DUNE jobs now running in updated DUNE VMs, about to roll out to OpenStack sites with Vcycle access. Requesting a renaming of UK sites in FIFE platform at FNAL used by DUNE, uboone, g-2, ... (eg "London" -> "UK-Imperial") to ensure clearer naming and get UK recognition on plots at workshops etc.</p>
DLS [2100]	<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> <p>17/09/19 CR Work done on K8s deployment on DLS prem using kubeadm. Not yet ready for a deployment on IRIS infra. Waiting on routing issue to be resolved inside STFC network to provide Echo access. [AS mentioned that routing issue had been addressed and further work was underway]</p> <p>01/10/19 CR No further progress in testing Echo access. Catchup with Xchem/Frank scheduled.</p> <p>15/10/19 Some progress being made. Xchem jobs running OK from Freiburg system. (AS - suggest maybe talk at TWG about this work)</p> <p>29/10/19 CR Echo access being looked at again by STFC. Diamond putting in application to ALC for position to develop software infrastructure for re-processing MX data on IRIS.</p>
PANOSC [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> <p>15/10/19 [IJ]: As part of WP6 pilot data archive task, attempting simple GridFTP transfer into ECHO from ILL. Currently, STFC GridFTP server believes a CRL related to ILL has expired, and closes the connection. IJ to investigate further.</p> <p>A suggestion arising from RAL PaNOSC meeting 25/09: create a VO for PaNOSC collaboration. Any thoughts about this?</p>
ISIS [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> <p>[17/09/19]AS reported that ISIS data taking was just starting</p> <p>12/11/2019 [JR] The service has been stable, no major/significant issues. Beam run starting today. We are still working on getting more groups on board, this will happen in February 2020.</p> <p>26/11/2019 [JR] There were two issues two weeks ago:</p> <ul style="list-style-type: none"> - IDAaaS was affected as a result of a recent change to the code of Facilities BusApp Visit system. First report occurred on Tuesday night and the issue was fixed by BusApp team on Thursday night (there is a Post Mortem analysis for this). - On Friday 15/10 all the cephfs mounts on all the VMs had locked up. Marcello from Data Services team provided support during that evening. As a follow-up on this issue on Monday 18/11 Marcello and Alison told

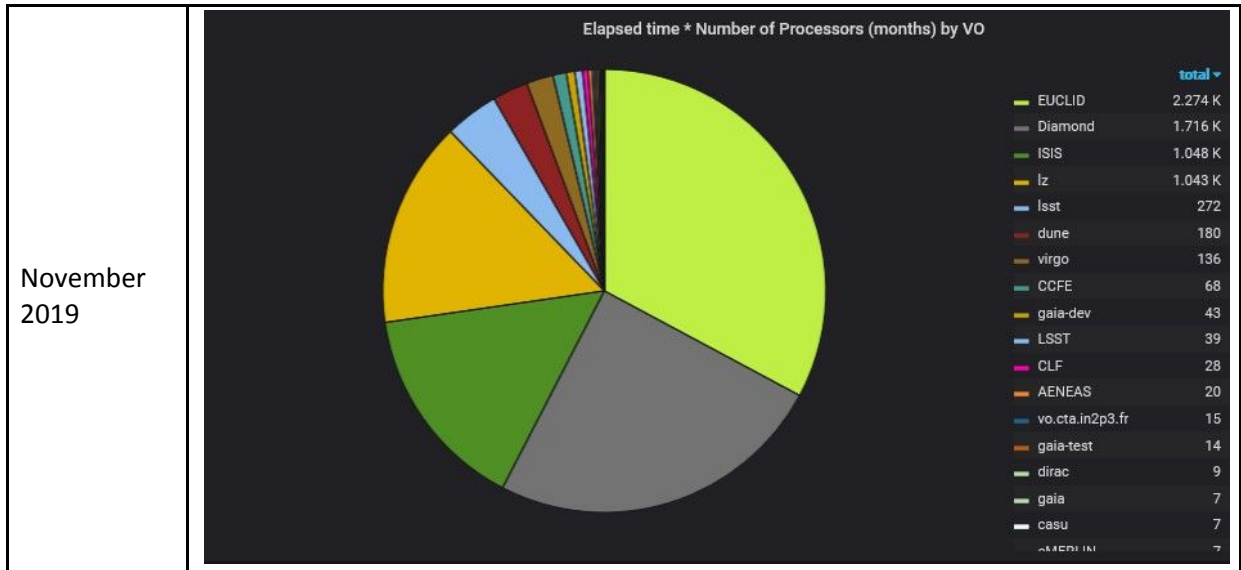
	<p>me that one of the metadata servers fell over and it may have been due to the number of small files. Also the cluster is close to full. This has been reported to Gordon and Lamar.</p> <p>We have been experiencing performance issues. This was fixed on Monday 25/11/2019 by the Cloud team.</p>
CLF [68]	<p>23/07/19 Fraser B has been speaking to CLF regularly this week, looking at IDAaaS 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 no update</p> <p>15/10/19 Nothing to report</p>
CCFE [1250]	<p>16/09/19 [AL] Continued use of DiRAC, some usage of IC. Difficult to track usage of DiRAC - currently the usage is larger than the allocation according to SAFE while "mybalance" gives very different numbers. There's no way to get usage as a function of time. [SJR] mybalance should give actual usage; for historical usage please use sacct or gstatement (or ask support@hpc.cam.ac.uk for guidance).</p> <p>01/10/19 [AL] DiRAC usage continues, occasional usage of IC.</p> <p>15/10/19 [AL] No change - as before</p> <p>29/10/19 [AL] No change.</p> <p>26/11/19 [AL] No change.</p>
JINTRAC [200]	<p>03/09/19 [DB] Just for the record: JINTRAC had access to the IC IRIS Cloud for about a month.</p> <p>03/09/19 [AL] Still waiting for users to try running their jobs on IRIS resources.</p> <p>16/09/19 [AL] There have been enough resources available at CCFE to run JINTRAC over the past few months so no external resources have been required.</p> <p>15/10/19 [AL] No change.</p> <p>29/10/19 [AL] Due to a cluster at CCFE becoming very busy several users have started running jobs on IRIS resources.</p> <p>26/11/19 [AL] There has been some (very) small-scale usage during November.</p>

Usage

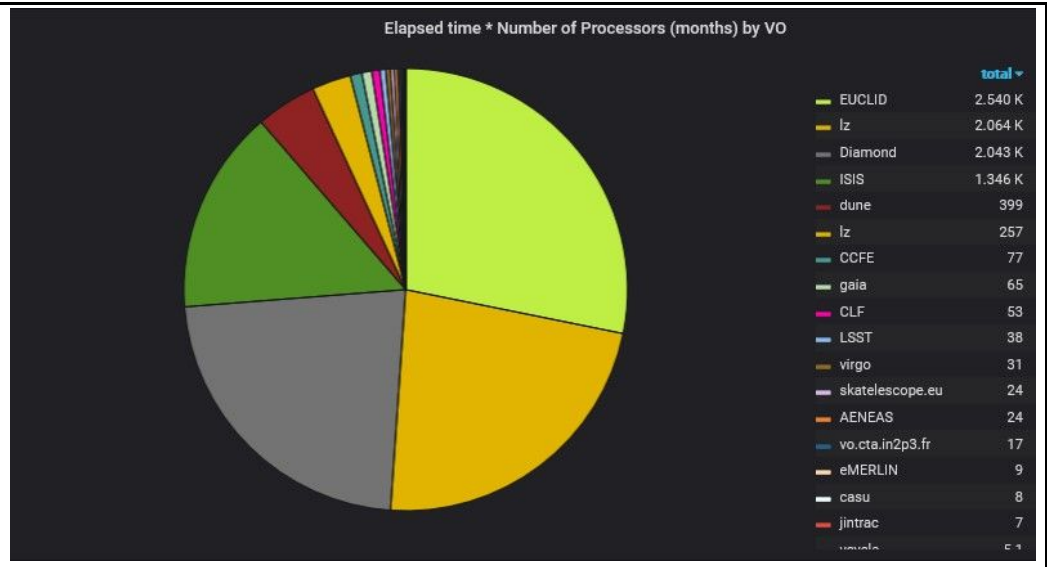
Year to date (stacked graph)



Monthly data breakdown. Grid plus cloud



October
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	15/10/19 [Ian Johnson] Graduate working on Kibana dashboards for Rucio transfer monitoring. Using one of the RAL Rucio cluster nodes as an ElasticSearch host, pending restricted public access to the Tier-1 ES cluster. 29/10/19 [Ian Johnson] Kibana dashboard prototype deployed for monitoring SKA transfer activity and shown to project partners.
DIRAC	20/8/19 [DB] Working on DIRAC -> Workflow clusters - proof of concept 15/10/19 [DB] LZ is now using DIRAC for cloud access at Imperial. LSST has used it at Cambridge. 12/11/19 [DB] All is good, nothing major to report.
Vcycle	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. 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.
AAI Proxy	15/10/19 [TomD] On-boarded Warren and Martin as Group Managers to test user roll out with the SCD cloud and users. Need to follow up Cert issues with John G and Cambridge and installation of OIDC endpoints for Dynafed. 12/11/19 [TomD] Basic AuthN working on the Dynafed test instance, need to modify AuthZ plugin to handle tokens. Setting up a test instance of Rucio to work on RUCIO/OIDC integration. Andrew Lahiff has signed up to IAM and will look at on boarding some CCFE users for IAM-Cloud testing. 26/11/19 [TomD] Working on DynaFed OIDC plugin, and will be testing setting up OIDC to a Rucio test instance provided by Ian Johnson. Working through STFC Services Review questions & planning IAM status update for F2F & EoY report.
APEL	Feedback can be sent to: apel-admins@stfc.ac.uk

	<p>12/11/2019 [AC]: Combined site and VO filters, fixed colour mapping, and IAM integration rolled out - IAM only accessible by STFC users for the moment while some security issues are looked into. Planning to move lower graphs to separate pages and add a more detailed capacity plot (the thin green line).</p> <p>26/11/2019 [AC]: Waiting for certificate for dashboard before opening up IAM sign-in option. Tim working on better allocation and capacity visualisations.</p>
--	--

Open Tickets

Latest on tickets

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

Ticket	Submitted by/submit Date/last update	Problem/Status
143436	db 1/10/19 4/11/19	<p><i>Accounting record size limit exceeded in AMS at IC</i></p> <p>11/10/19 APEL team referred problem to AMS team</p> <p>28/10/19 Updated to very urgent and assigned to caso</p> <p>4/11/19 Specific request from APEL team for default to be the same as for APEL (1000 records)</p>
143438	db 1/10/19 18/10/19	<p><i>Bug report for caso - accounting record uploader for openstack - local workaround applied</i></p> <p>17/10/19 Fix provided by developer (cause was local edge case)</p> <p>18/10/19 Fix tested by imperial. IC requests documentation update and rpm.</p>
143490	ph 3/10/19	<p><i>LIGO at Cardiff would like to upload accounting data</i></p> <p>3/10/19 APEL team report they are short staffed</p> <p>11/11/19 Work ongoing to upload data from Cardiff</p>