

GridPP Status & Outlook

DIRAC & RUCIO Workshop 2023

Daniela Bauer & Simon Fayer

Overview

1. Who we are: GridPP et al.
2. Who we support: Small to medium size communities.
3. What makes us special: This is not the WLCG.
4. What we work on: Mostly a consequence of points 1. - 3.
5. What Federico wanted to know: Answers to the questionnaire.
6. What matters: Conclusions (mine, yours might differ).

GridPP DIRAC Stakeholders

- Main Stakeholder: GridPP (Grid computing for particle physics - the UK arm of the WLCG): <https://www.gridpp.ac.uk/>
 - GridPP has always allocated 5-10% of resources to “non-LHC” VOs
- Imperial College (HEP group): <https://www.imperial.ac.uk/high-energy-physics>

GridPP and Imperial College fund personnel to run the service; GridPP also contributes the DIRAC hardware.

- IRIS (Digital research infrastructure for eScience): <https://www.iris.ac.uk/> (note: this is not IRIS-HEP as mentioned during the RUCIO talk)
- SWIFT-HEP (SoftWare and Infrastructure Technology for High Energy Physics): <http://swift.hep.ac.uk/>

IRIS and SWIFT-HEP grants are for project work only: Typically 6-18 months long, must have a finished product (read: code) at the end.

GridPP DIRAC Overview

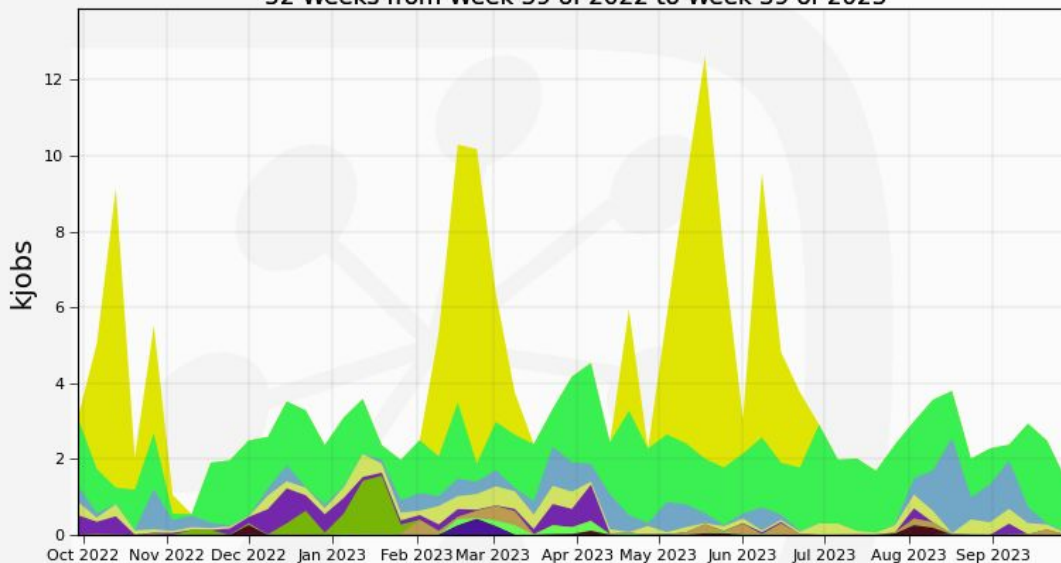
- The GridPP DIRAC instance supports small to medium size communities with UK participation:
 - Neutrinos: T2K, HyperK, SoLid, SNOPLUS
 - Dark Matter: LZ
 - Astronomy: Euclid (simulation workflows)
 - Phenogrid
 - Muon decays: COMET, Mu3e
 - Smaller CERN based experiments:
 - NA62 (rare Kaon decays)
 - Moedal (magnetic monopoles)



GridPP DIRAC - last year

Running jobs by UserGroup

52 Weeks from Week 39 of 2022 to Week 39 of 2023



eucliduk.net_user	38.0%	lz_user	1.5%	gridpp_user	0.0%
na62.vo.gridpp.ac.uk_production	37.8%	vo.moedal.org_user	0.7%	pheno_user	0.0%
lz_production	9.3%	hyperk.org_user	0.5%	snoplus.snolab.ca_user	0.0%
snoplus.snolab.ca_production	5.4%	solidexperiment.org_production	0.4%	lsst_user	0.0%
t2k.org_user	4.0%	vo.northgrid.ac.uk_user	0.0%	na62.vo.gridpp.ac.uk_user	0.0%
comet.j-parc.jp_user	2.4%	mu3e.org_user	0.0%		

Generated on 2023-10-03 14:06:05 UTC

- We typically have 6-8 experiments actively using the GridPP DIRAC instance at any one time (~10 over the course of a year).
- These numbers might look small from a WLCG perspective, but often represent a large part of these activities' science programmes.



GridPP DIRAC: From the questionnaire - part 1

- In the **last year**, what has been the DIRAC usage in terms of jobs ran, CPU (or wall time) used, and data transfers?
 - 9.1 million jobs executed.
 - ~424 800 000 HEPSPEC06 hours.
 - LZ generated and stored 4.25 PB of data using DIRAC data management tools.
 - t2k.org: replicated 370 TB of data (via DIRAC and FTS).

How do (some of) our users use the GridPP DIRAC instance ?

- T2K/Hyper-K: Please see Sophie King's talk on [Tuesday](#)
- LZ: GridPP DIRAC forms a large part of their processing work chain in the UK:
 - Used for grid and cloud computing
 - 12.5 mio files registered in the DIRAC File Catalog
 - LZ built a [experiment specific frontend](#) using the DIRAC python API for bulk processing, user analysis jobs use DIRAC API directly
 - A lot of DIRAC Cloud developments were driven by this project.
- Snoplus: Monte-Carlo production, using Ganga as their interface
- NA62: Monte-Carlo production, [experiment specific frontend](#)
- Euclid: Galaxy shear simulations (see [details](#) if you are interested in what we are up against)

Supporting ‘small’ communities is different

- These communities typically do not have dedicated computing teams, and where they exist, it’s often a team of one.
- The GridPP DIRAC team is a team of two (with day jobs), plus a developer, yet we still dwarf the computing teams of most (all?) the communities we support.
- For these communities the GridPP DIRAC instance is their interface to the grid:
 - We do a lot of triage when the underlying infrastructure breaks.
 - **WLCG decisions (tokens, operating systems, and even the BDII) are very much our problem:**
 - Operating systems: We cannot advise our communities to just “stick their code in a container”, unless we work with them to make this happen.
 - Similarly: “What’s a token and why do you keep changing stuff ?”

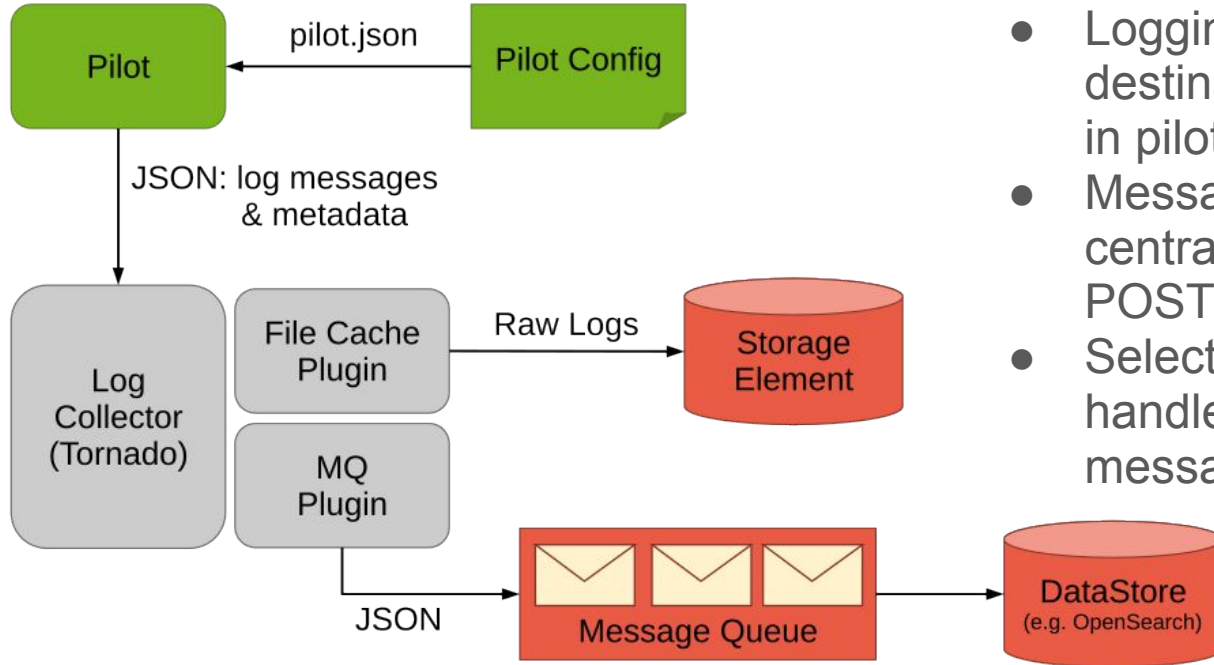
Supporting small experiments is different: how we run our service

- **We emphasize reliability over new features:**
 - One major version upgrade per year:
 - We offer a CVMFS DIRAC UI, but often our activities have their own.
 - We do not deploy anything without a use case (no proof-of-principle on the production server), as this tends to either increase the maintenance burden or fall into disrepair:
 - We really tried to sell RUCIO to the communities we thought it would be suitable for, but no takers yet.
 - We keep sprouting pre-prod servers: DIRAC(s), Rucio, IAM. We test a lot before deploying.
 - We use extensive monitoring: Every time something breaks, we add a new nagios and/or pre-prod test for it.

GridPP DIRAC pet projects - 1 & 2

1. Standardizing DIRAC's cloud interfaces using Apache libcloud: [Talk on Wednesday](#)
2. Pilot Logging:
 - When things go wrong, we need the logs from the pilot job for debugging
 - Currently pilot logs are distributed on a pull model using whatever interface the front end provides:
 - Logs are often only available for a short time or possibly not at all (e.g. on clouds or if there are problems with the CE)
 - Working on a new push model:
 - Pilot logs are written in regular intervals to DIRAC service
 - Can be enabled on demand to minimise load on the central log collector

GridPP pet projects - 2: Pilot Logging



- Logging enabled by destination URL + enable flag in pilot.json
- Messages sent periodically to central collector using HTTP POST
- Selectable plugin then handles further processing of messages.

GridPP DIRAC pet projects - 3

- Multi-VO tokens - it's not so simple (cf single VO tokens):
 - A CE might only accept tokens for a subset of VOs
 - VOs typically do not deploy tokens everywhere at the same time
 - **Most of our users have never heard of tokens:**
 - Start with the pilots, do users last
 - Users are used to three geographically distributed VOMS servers and 100% uptime:
 - IAM is not there yet
 - Currently we are using a pre-prod IAM server with a pre-prod DIRAC instance for testing:
 - Progress is limited by available hours in a day

The questionnaire - part 2

- What we like about DIRAC 😊
 - It mostly works. For all its faults, we get an enormous amount of work (and science!) done with it.
 - It's not a closed-shop system
 - Though we cannot always magic up the people to write all the features we would like
 - There is a culture of (trying to) systematically improving code quality, deployment etc
 - Apart from one incident, DIRACOS2 has worked really well.
 - The python API is just about the amount of computing interaction our users can cope with

The questionnaire - part 3



- What is your biggest frustration with DIRAC?
 - [#worksforlhcb](#):
 - We are still seeing new features (tokens!) being commissioned that are inherently single VO only.
 - We don't have an IT department that deploys the underlying infrastructure for us. Using 'standard' tools/software doesn't make it easier if you have to deploy Kubernetes, Grafana, OpenSearch etc etc yourself, just so your users can submit a job. (Note: We need to deploy DIRAC in a way that we can hot fix it, refer to pages 8 & 9 for the why.)
 - [#doesntevenworkforlhcb](#):
 - With the advent of diracx the development branch of DIRAC has been in a non-working state for months, this makes any testing beyond CI impossible.

The questionnaire - part 4

- Any notable operations incident in the last year?
 - No, though we had to back out of some surprise pilot updates a couple of times. I think this has been noted by the project management ;-)
- What can be improved for communication in the project?
 - It generally works, but the dissemination of known issues could be more proactive:
 - If I file an issue, I don't want to hear: "We've seen the same in LHCb, too" – just tell me immediately when you find it !
 - Like all science projects, it works on collaboration and limited people power, so we need to work within these limitations.
- What are your expectations with regard to DiracX?
 - Plus ça change.

Conclusion

- GridPP DIRAC is the workhorse for a number of small to medium size communities
- GridPP takes an active role in ensuring DIRAC meets to needs of these communities
 - Token deployment and WN operating system upgrades present challenges* to these communities which are not always obvious from a WLCG perspective: We are working on it, but like everyone we have limited effort assigned to the project.

* You might not consider this a DIRAC problem, however the communities we support do.

- Thanks for having us here at KEK :-)