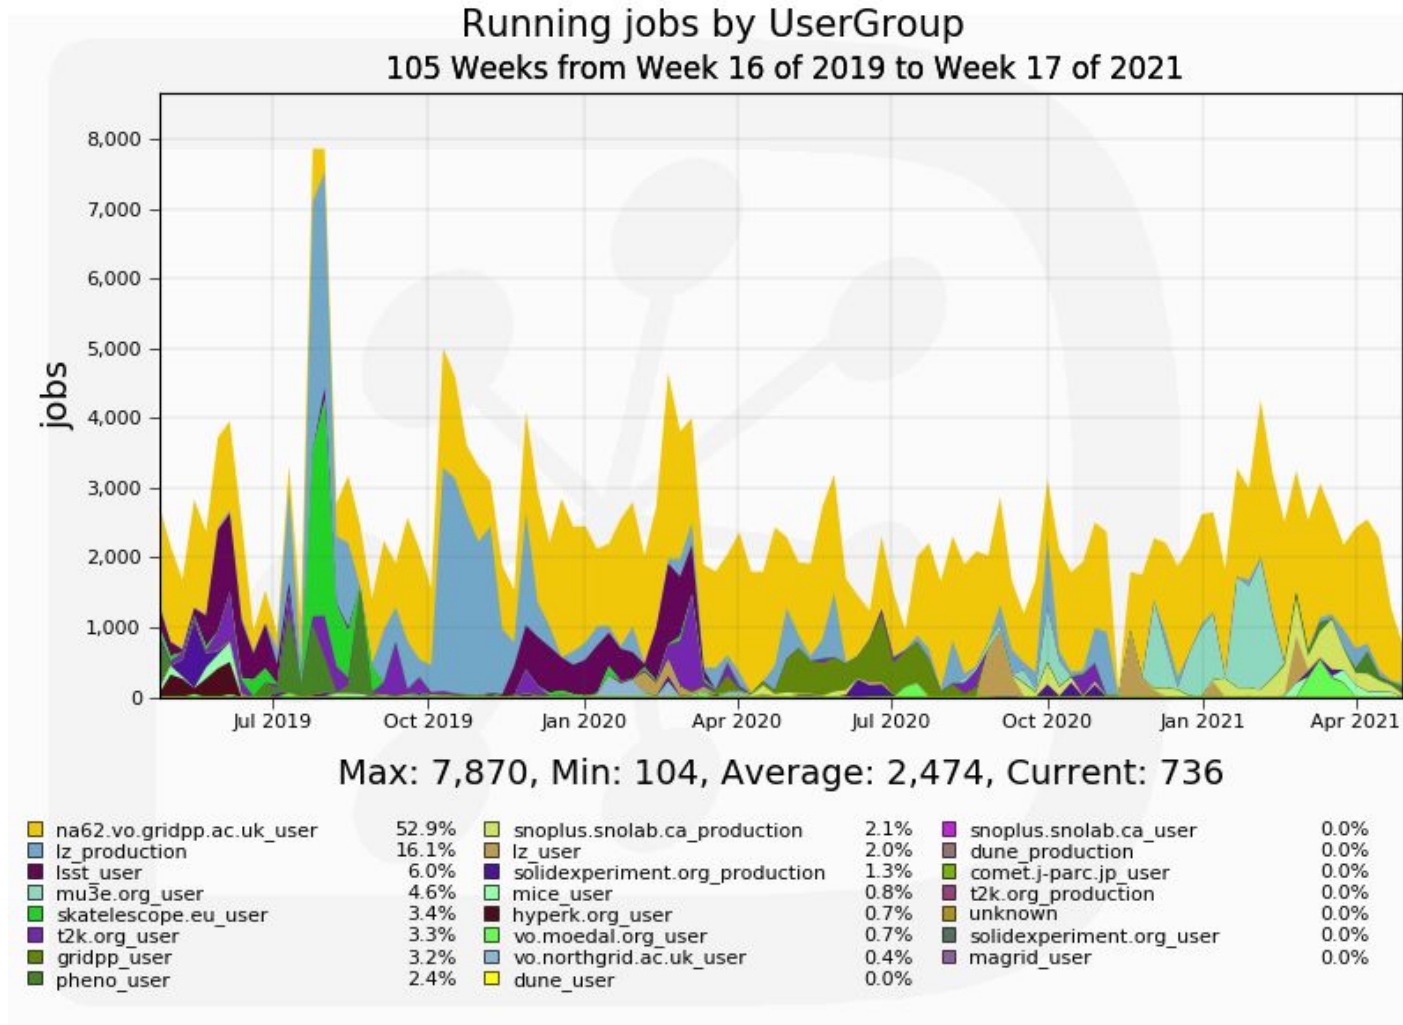


GridPP

Daniela Bauer, Simon Fayer & Janusz Martyniak


What have we done the last 2 years ?



Generated on 2021-05-06 14:40:52 UTC

CPU usage

CPU used by FinalMajorStatus
105 Weeks from Week 16 of 2019 to Week 17 of 2021

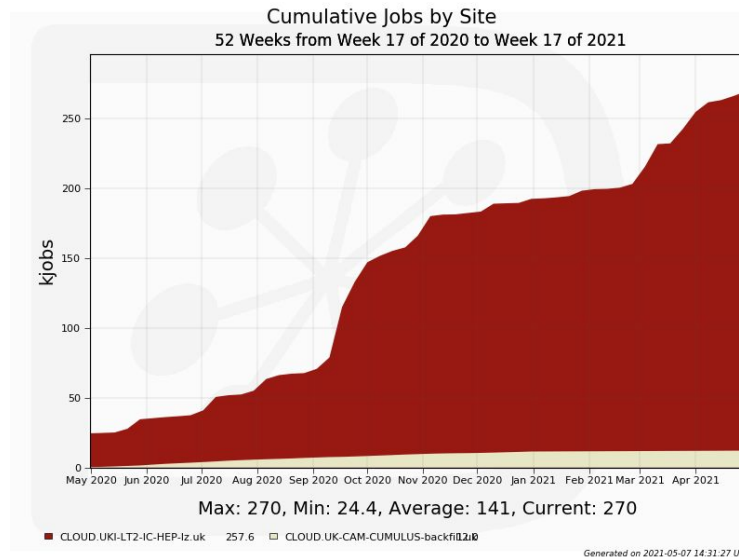


Max: 5,692, Min: 63.7, Average: 3,049, Current: 5,692

Done 4643.0 Failed 628.6 Completing 421.1 Rescheduled 0.2

VMDIRAC

- GridPP uses VMDIRAC in cloud-init mode to access two OpenStack Clouds in the UK.
- Our main customer is currently LZ, to create VMs with more memory than the average grid node.
- Setting it up is still non-trivial, though documentation is improving.
 - Simon is happy to answer questions.
- Once it's setup, it works well.



Admittedly this one is more a proof of principle.

Other things we worked on

Our motto: “If it’s single we make it multi !”

- Multi-VO metadata (completed)
- Multi-VO RSS (in v7r3)
- Multi-VO DIRAC to Multi-VO RUCIO (work in progress)

Survey Questions

What do you do you use DIRAC for, and which DIRAC functionalities you don't use, and why?

- How we use DIRAC:
 - “DIRAC as a service:“ GridPP provides a DIRAC instance as a service to the communities it supports, but there is no direct involvement of the experiments in the running/maintenance/development etc of the instance.
 - All communities use it for job submission.
 - A subset also uses it for basic data management.
- What we don't use:
 - RSS (yet)
 - Request Management System (possibly because our users aren't aware of it)
 - Transformation/Production system (yet, more on this later)

Survey Questions continued

Do you have a DIRAC extension? Why?

- Yes.
- Mainly to tweak the autoconfig on the server side to match user expectations (e.g. selecting “EL6” vs “EL7” during the EL6 to EL7 transition)

If yes, do you think some of it could become part of the vanilla projects?

- No, I think we pushed everything back we can (for now)

What is your biggest frustration with DIRAC?

- Error handling:
 - Pilot logs not available (more on that later)
 - Error messages not forwarded through the stack: Users (and admins!) see identical error messages for unrelated problems.
 - E.g: “Failed to submit cloud job.” (Yes, I can see that. And now ???) and the perennial favourite among our users:
 - ERROR: Error: failed to upload /myvo/myproductionfile to UKI-GENERIC-SITE-HEP-disk (catalogue problem ? storage element problem ? auth error ?)
 - Especially frustrating when debugging jobs that have already finished, and have run without the debug option turned on.

Survey - the end

Any notable operations incident in the last year?

Not really, but the “diracweb.cern.ch” server being down causing all pilots to fail and the ‘//’ in the LFN sending the file catalogue into an infinite loop (both now fixed) certainly caused some consternation at the time.

To support your "Grid", do you have to use other systems than DIRAC?

- Rucio: No actual use case yet, but fashionable.
- Custom frontends (“LZ Job Submission Interface”)
- FTS

How would you rate the communication?

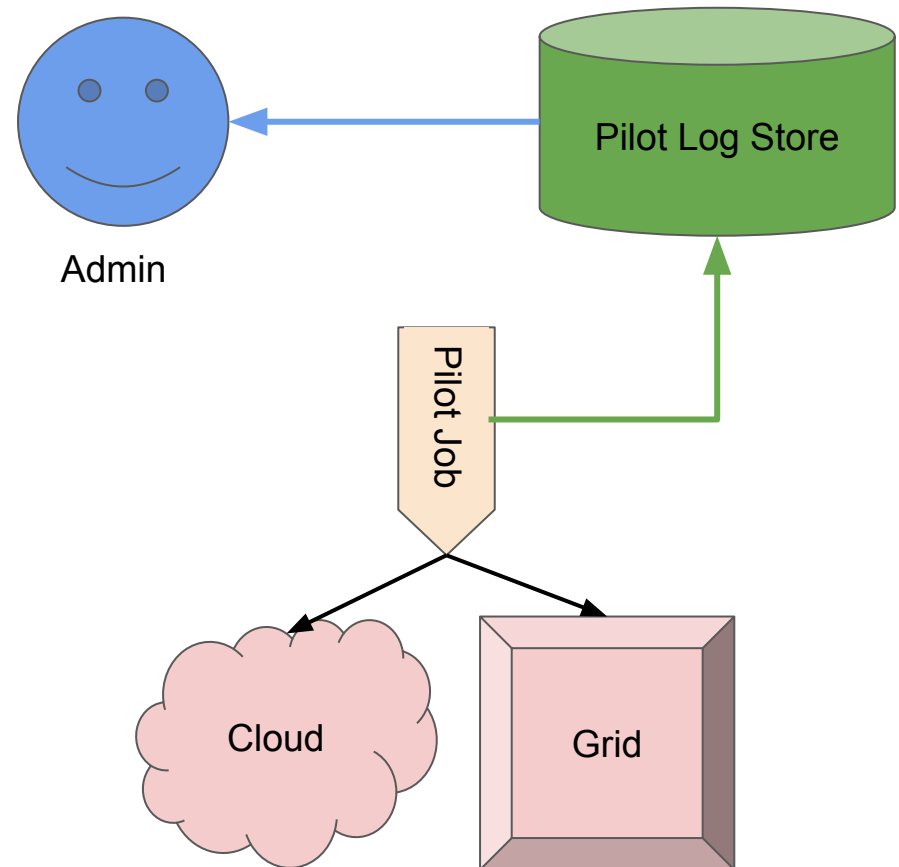
- BiLD meetings/hackathons really help. Should be advertised to any new DIRAC admin.

Projects and Plans

- As part of the SoftWare InFrastructure and Technology for High Energy Physics experiments (SWIFT-HEP) project we have been allocated 50 % of a Research Software Engineer post (i.e. Janusz) to work on DIRAC related projects for the next two years.
- The focus is improvements needed relating to GridPP projects and the GridPP DIRAC instance.
- We aim to commit as much as possible back to core DIRAC.
- The following slides are our initial ideas on how to approach these issues.

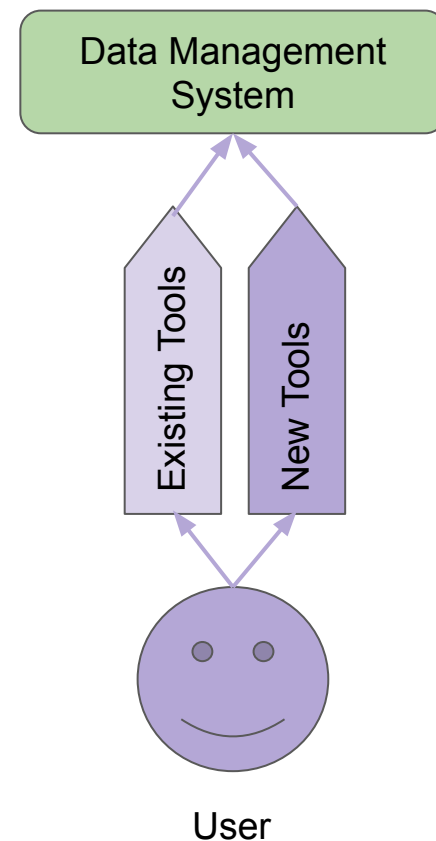
SWIFT-HEP: Pilot Log System

- Pilot job logs are stored in an technology dependent way at the execution resources.
- Retention policies vary by technology and site:
 - Some logs only kept while job running!
 - Others kept 3 days - 1 month depending on admin settings.
- Log can be lost in cases where job crashes (i.e. exceeding batch system limit).
- Task: Implement a central pilot log store and allow the pilot jobs to write logs there directly.
- At peak times this service needs to cope with a large amount of traffic in a fault-tolerant way.



SWIFT-HEP: DIRAC DMS High Level Commands

- Target are medium size communities for whom e.g. RUCIO would be too much overhead.
- These communities often already use the DIRAC File Catalogue and basic DIRAC data management tools, so the threshold for adoption is quite low.
- **Develop user-friendly tools for the most common use cases and make them available to all users as part of DIRAC.**
Translation: Integrate our “dirac tools” scripts that we have been handing out to our users as “starting points” for years properly into DIRAC.



SWIFT-HEP

- Admin override tools: E.g. reassign jobs to a different site, improved pilot killers, etc (we are still discussing those)
- Making the transformation system work for GridPP
 - Experience in single VO mode
 - Step 1: Set up Transformation System on Production server
 - Fix any issues that arise
 - Step 2: Get (at least) two communities to test it
 - Fix any issues that arise
 - Commit any bug fixes back to core DIRAC :-)

Conclusions

Business as usual, pandemic or not.

Questions ?

Contact: lcg-site-admin@imperial.ac.uk