GridPP DIRAC Status - Supporting non-LHC VOs in the UK

Daniela Bauer & Simon Fayer

Imperial College London

GridPP
UK Computing for Particle Physics

Business as usual - just more of it (Success!)

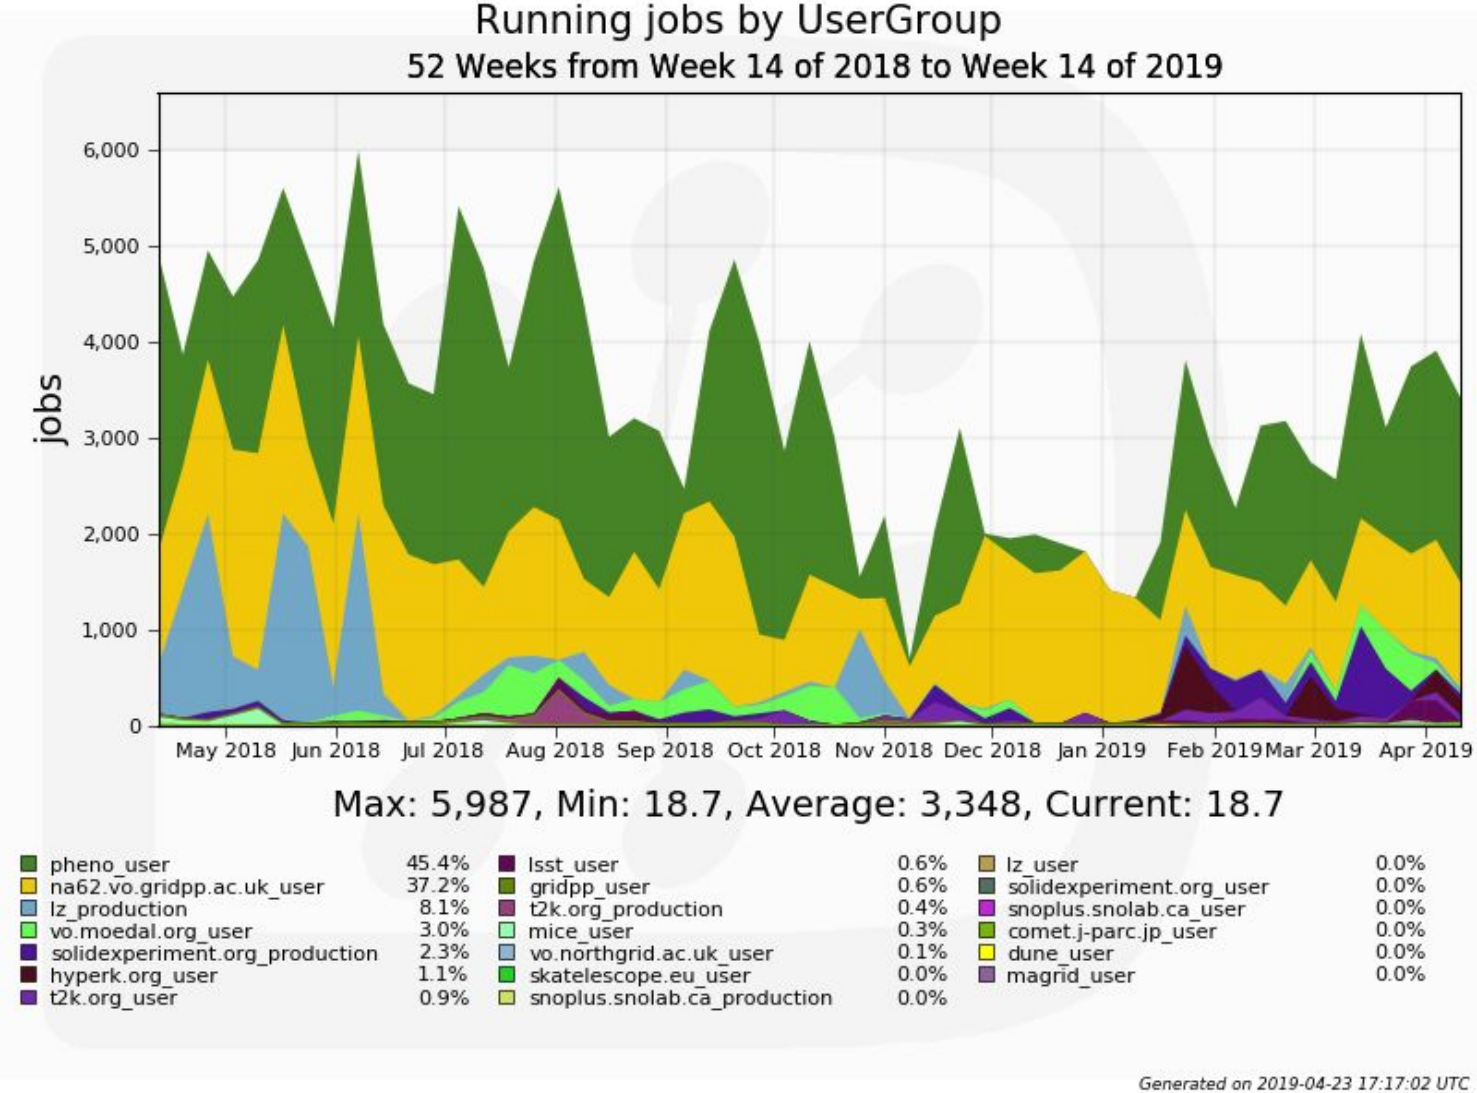

Where do the jobs go ?

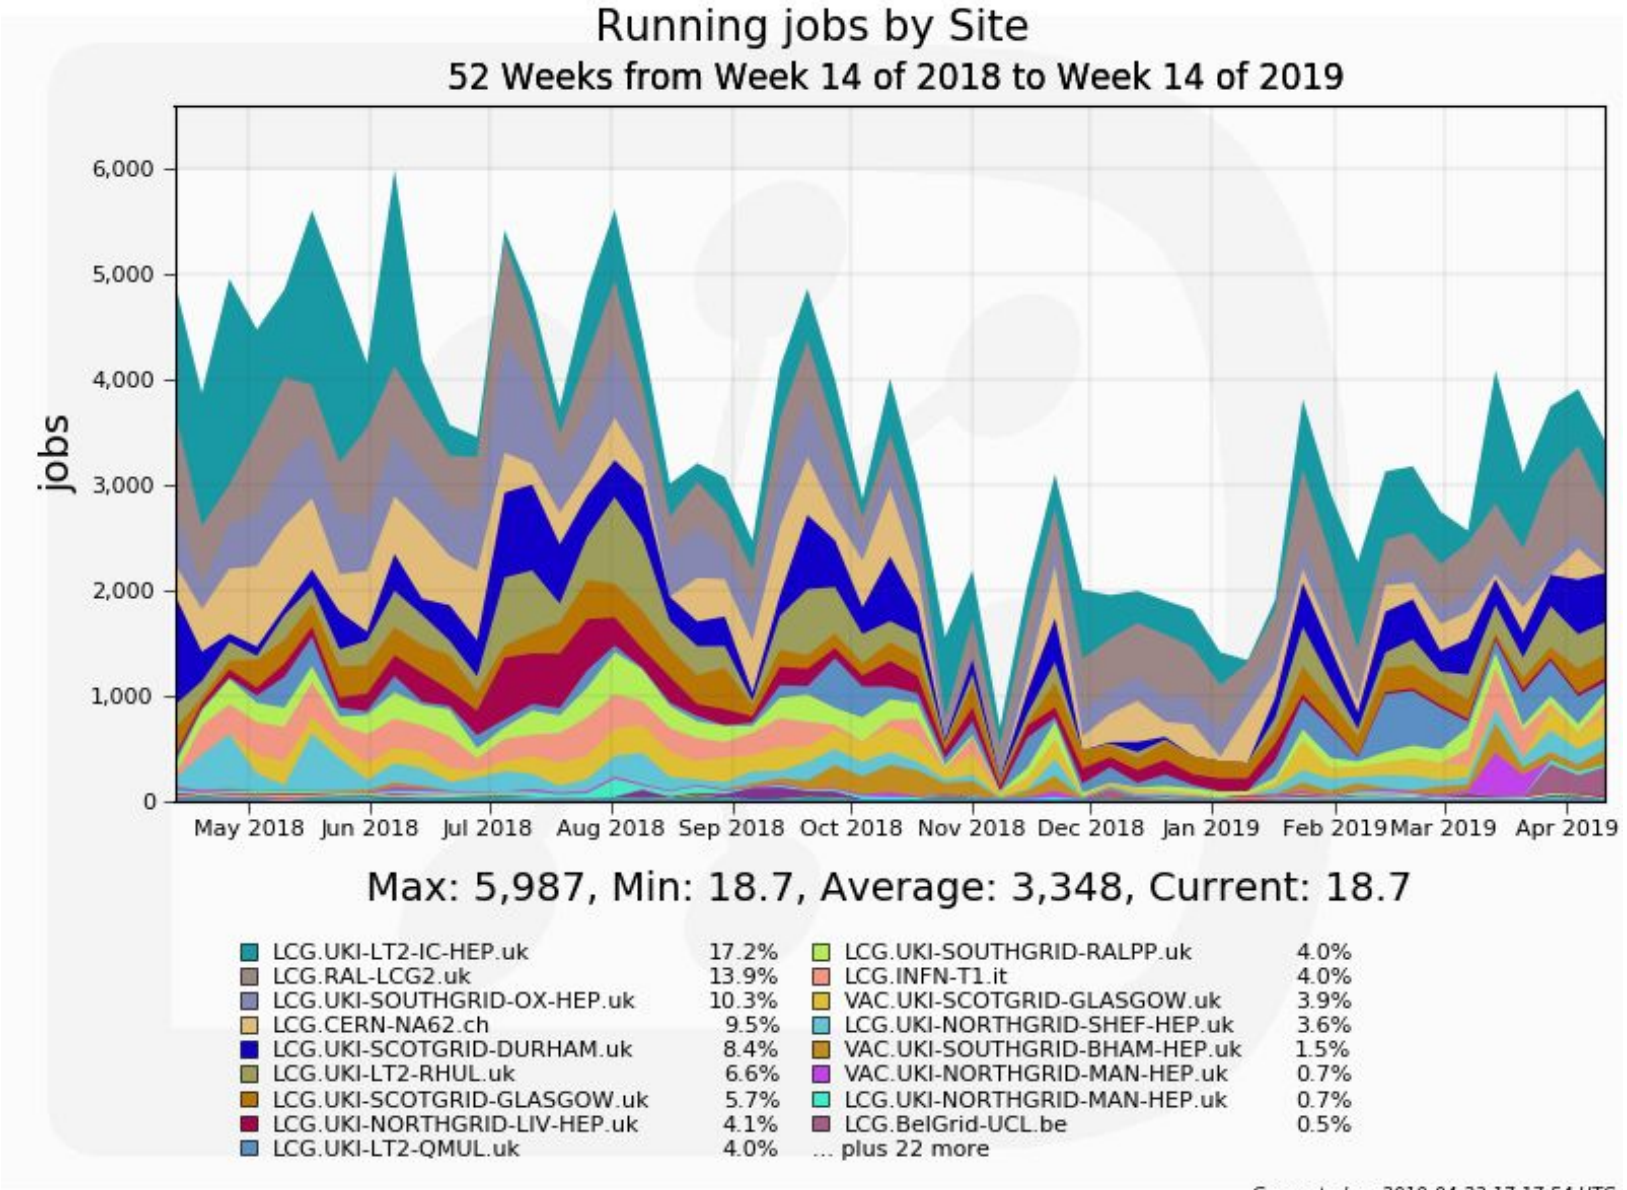

What did we do all year ?

- Administration is daily work.
  - Upgrades
  - Site changes
- Users: Why does my job not run ?
  - User issue
  - Target site issue
  - DIRAC issue
  - All of the above
- Debugging/making stuff work takes time:
  - Limited information from the Grid frontends
  - DIRAC only releases limited information to users
  - Sites configure new VOs blindly (no local users):
    - But: Most VOs don't know how to debug a site.
  - As most of this information is not accessible to everyone due to security reasons we rely on the cooperation of all parties concerned.
  - Adding even more people in the mix won't help (plus I actually like that bit of my job)

But I **want** to help ….

- Yes: **Write some code.**
  - VO specific: Get familiar with the DIRAC API and make a custom interface for your favourite VO.
  - For the adventurous: Become a DIRAC developer - it's open source. Fix your favourite issue right at the source :-)
- If you are doing "new" stuff at a site, involve the non-LHC VOs and GridPP DIRAC early.

Making DIRAC better

- Use IRIS digital asset grant to implement features needed by IRIS supported communities
- DIRAC and Rucio
  - Not just a UK issue, e.g. work done by Belle.
  - Still very much in the development phase.
  - In the UK mainly driven by SKA.
- DIRAC Workflow management system:
  - Originally requested by SKA.
  - Hopefully to finalize it during the DIRAC workshop in May.
- DIRAC Resource Status system:
  - Automatically ban sites on a VO by VO basis:
    - Take workload off DIRAC central and VOs
- DIRAC and clouds (IRIS!): VAC and non-VAC, we intend to support them all

New Communities

- This is working fairly well given the (wo)man power we have.
- Communities get help at sites, Simon and me provide some centralized backup.
- Involving the VOs directly and keeping them up-to-date on a regular basis seems to greatly increases VO happiness.

Challenges: LFC to DFC

- RAL is retiring their LFC (“logical file catalogue”) instance
- T2K: 8 million files registered (~4.7 on RAL TAPE) on 10 sites
- DIRAC offers an inbuilt file catalogue, not too dissimilar to LFC
- Migration looked fairly straight forward
  - most data needed moving due to different handling of LFNs by LFC and DFC
    - Mostly done by sites
- 5 months later we have seen:
  - dark data (lots, though apparently some of it is important)
  - data without checksums
  - file replicas with different checksums
  - lost data
- We are nearly there….

LFC to DFC

- Nothing in this project required any special privileges beyond being a t2k user and a dump of the LFC database.
- Now we just have MICE and SNO+ to contend with.
- MICE is work in progress.
  - Much smaller scale.
  - MICE will adapt their data mover.
- The eternal question: Should GridPP be doing this work ?
  - We've adopted these VOs, we shouldn't pull the rug underneath them.
  - We need to keep in mind that even with all the hand holding these VOs will still have to adapt their workflows.
  - It's definitely beyond the GridPP DIRAC scope, so any GridPP volunteer will do.

Non-LHC VOs and GridPP (& IRIS?)

- Alistair wanted a discussion…
- Due to my work in GridPP DIRAC and IRIS I'm almost automatically an advocate for the non-LHC VOs.
  - It's not personal, I just want to make it work.
- There is a slow shift from the LHC mindset ("but it works for Atlas!") to "we need to support all our users equally".
- Possibly helped by IRIS funding.
- To the sites:
  - Please be selective: If you cannot support a certain VO to their specifications, please pick the ones that fit your site. (Unless IRIS tells you otherwise.)
  - Please keep in mind we are running a production service. This is not the place to try out new features *unless requested*.

Engaging the (power) users

- DIRAC workshop 14-17 May in London.
- We've invited the non-LHC experiments with a UK presence to send their computing people.
- Last year SKA was present at the workshop, this proved very useful.
- Goal: Taking out the middle woman: I *try* and convey the experiments experiences & requirements, but why not talk to the DIRAC developers themselves ?
- If you think you know of someone who would profit from this, please let them know.
  - If you are quick we can still arrange a topic you are interested in.
  - Or you could give a talk about your work :-)

DIRAC Workshop

- DIRAC workshop: https://indico.cern.ch/event/756635/
- Learn how other communities use DIRAC.
- Learn about new DIRAC developments.
- Talk to the developers.
- Tutorials/Hacking sessions
- There's still time to put your favourite item on the agenda.

Imperial College London

Conclusions

- We are in a fairly good place with DIRAC
  - Not that we have many alternatives
- Non-LHC VO support does not come free
  - More communities means more support requirements