GridPP DIRAC

Background information:
The target user groups of the GridPP DIRAC instance are small to medium size experiments who typically have small (~1 person) to non-existent computing support teams.
We currently support a dozen VOs, about 6 of which typically actively use the instance at any one time. I include a usage plot for last year as an example. We onboard a new community approximately every 18 months.

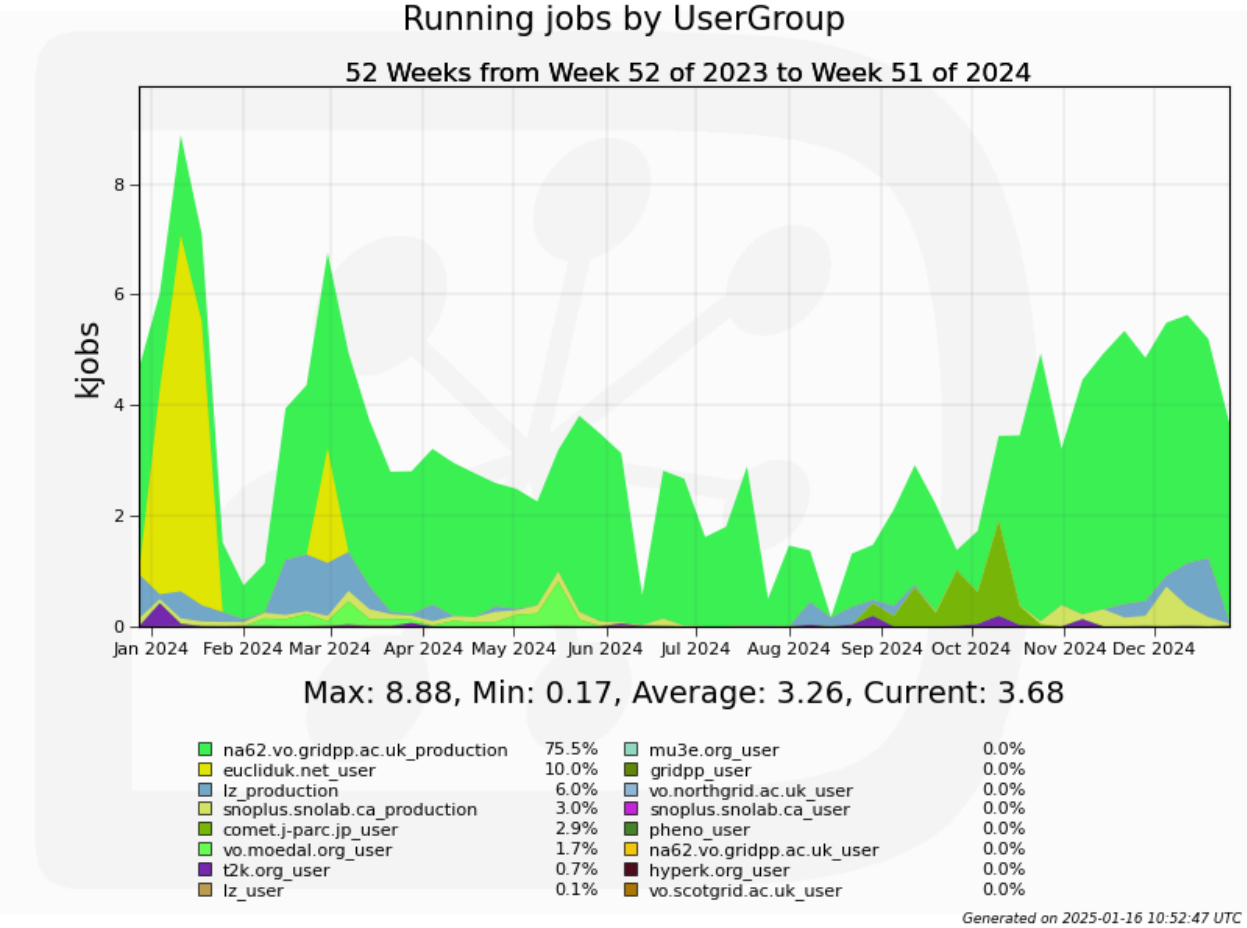

The main usage by these communities is the use of DIRAC as an extended batch system, i.e. the **Workload Management** component of DIRAC. A small number of them (NA62, LZ, T2K) have incorporated DIRAC into their higher level production systems. These systems are (very) experiment specific. The GridPP DIRAC team typically assists with queries (interfaces) resulting from the development of these production systems, but has no active involvement in them otherwise.
A number of VOs (LZ, T2K, moedal) use the DIRAC **Data Management** component. DIRAC has an interface to RUCIO and currently there is a shift to use RUCIO as a data management tool for communities that are large enough to support their own RUCIO server. At this time, none of the GridPP supported VOs is in this position, although it is believed that NA62 and possibly some of the neutrino experiments are investigating this possibility. We have no

intention of dropping the DIRAC data management component as it is sufficient for most of our use cases, but are happy to support the RUCIO interface should the need arise; we have expertise for this in-house.

We have deployed the **Workflow Management** system ("Transformation System") on the GridPP DIRAC instance, but we currently have no users that use this in production. T2K is looking to use it intermittently, initially for automated data transfers using the DIRAC File Catalogue Metadata, but due to limited developer time (on T2K's side), this remains in a test state. We are not aware of any communities that want to use the Workflow System for Monte Carlo Production (i.e. to chain different steps of a Monte Carlo production workflow), but our tests indicate that this is feasible on our system as it is now.

We do not use (and have not deployed) the **Production System** (Production System Documentation) which is the high level Workflow Management System that sits on top of the basic Workflow Management system. One of the main drivers of the development of this system is CTAO which uses it for large scale and fairly complex workflows. I believe the following CHEP 2024 talk might be highly relevant to this CMS review as their use case is closer to CMS than GridPPs: CTAO Production System Status and Development
I have spoken to one of the CTAO DIRAC project leads (Luisa Arrabito <arrabito@in2p3.fr>) about the CMS review and she is happy to discuss the CTAO setup if you would find that useful.

To address the specific questions by the committee:

1. Can you describe the architecture of your WM solution? What are its components and how do they interact?

I believe this is addressed in the background information.

2. What workflows are supported in the workflow management solutions? (Example Prompt reconstruction (formerly known as Tier-0), offline data processing, data analysis, MC production, Ntuple creation, end user analysis….)

None of our communities runs multi-step workflows. The focus of most of them are large (on a small/medium VO definition of large) Monte Carlo productions and raw data reprocessing campaigns. The only community that currently uses the GridPP DIRAC for (limited) user analysis is LZ. It is recognised that the DIRAC user interface is not suitable for inexperienced users and something similar to CRAB is needed to package user code in a relocatable way and enable easy recovery from failures.

3. How many developer FTEs are supporting your solution?

We typically have 0.5 FTE of **developer** time allocated to work on issues specific to our particular use case:
In the past we have focused on making individual DIRAC components (including the Transformation System) multi-VO compatible and we have taken the lead in the development of the DIRAC to Cloud interface. Our current focus is on a more persistent pilot logging as we rely heavily on this for debugging.

4. How many people are involved in operating the system and documenting it?

I would estimate the **operator** effort to be around 0.5 FTE distributed over two people. This includes system upgrades, onboarding and user support. We use custom nagios probes for monitoring.

5. Where does funding / person power for the solution come from?

Operational support is funded by GridPP. Development support is funded by various projects; in the last 5 years that included GridPP, IRIS and SwiftHEP. These grants are typically limited in time and scope, i.e. to develop a specific feature, rather than general development. Bugs and corresponding fixes found in deployment are covered by operator effort.

6. What are the evolutionary plans for this component?

I am assuming this question refers to DIRAC as a whole. Imperial College is a member of the DIRAC consortium and therefore committed to support DIRAC for the foreseeable future. We only deploy and develop new features in response to changes in user or operator requirements. Currently, we only make minor contributions to DIRACx (the overhaul of the core DIRAC code to become closer to industry standard frameworks and protocols) through things like testing and maintaining an IAM server for the development project. In the current political climate it is unlikely that the UK will commit any funds to developments that are not of immediate benefit to the UK Science community. However, the development of DIRAC for CMS would, of course, be of direct benefit to the UK Science community.

7. Is your solution used by other organizations and does it constrain evolution?

DIRAC itself is used by a number of organizations apart from the ones mentioned in the introduction and LHCb, these include CTAO, Belle, ILC and IHEP (to support a number of Chinese experiments). In over a decade GridPP has not encountered a situation where constraints coming from other projects have impacted on GridPP operations.

8. What is going well and what functionality is missing?

The Workload Management System is the workhorse of the GridPP DIRAC instance and the use of a multi-VO setup is a very efficient use of the operators' effort. We believe we would not be able to offer the same support and resulting access to resources to small to medium VOs on other currently available solutions with the FTEs available to us.
Basic data management works well and does not require too much specialized knowledge by users, but the Workflow system in its current incarnation is not very user and/or operator friendly.

9. What are operators/users most unhappy about?

Operators: Core DIRAC is undeniably developed at CERN and will often assume a CERN like infrastructure that may not be present outside CERN. Typical examples include underlying or external components (kubernetes, opensearch) being deployed by outside teams.
Users: DIRAC was developed as a system to handle large scale workflows handled by production operators and not individual users. This is mainly handled by expectation management.

Users and operators: Debugging can be a challenge, though in the operators' experience of the error reporting it is no better or worse as what we can see from CMS.

10. What progress have you made in testing the scalability of your solution to Run4 requirements?

This does not apply to the GridPP DIRAC instance. Having said this, experience has shown that at least on the scales we operate, individual DIRAC components scale fairly well. Users will submit 30k jobs in one go and then go on holiday, this does not seem to cause issues. However the backend (sites available to run workloads) is limited to mainly UK sites, two clouds and CERN, so we are unlikely to ever reach typical WLCG compute and data volumes.

GridPP vs LHCb

For comparison this is the LHCb plot of running jobs for the last 90 days.
Everything except "User" is managed by the Workflow (Transformation) System.
WGProduction which is afaik centralized analysis (I guess what CMS would call "ntuple production"?) submitted by the physics groups. The physics groups do not interact directly with DIRAC, instead there is a github based interface to which the groups submit their code, which after validation is forwarded to DIRAC. The Monte-Carlo Simulation/Production is handled by ProductionManagers who use DIRAC directly.