

# RAL Tier-1 strategy - Growing the UK community

*Alastair Dewhurst*

# Introduction

- I became the RAL Tier-1 manager in April this year.
  - Previously I worked for ATLAS.
- GridPP is the name of the project that funds the UK's contribution to LHC Computing.
  - Tier-1 + 12 Tier-2 sites around the UK.
  - We are working on the next GridPP bid that will provide funding from 2020 - 2024.
- IRIS is the name of the new project that aims to provide computing resources to the “long tail” of science.
  - All resources must be capital.
  - £4 million available each year for the next 4 years.



# Strategy 2020 - 2024

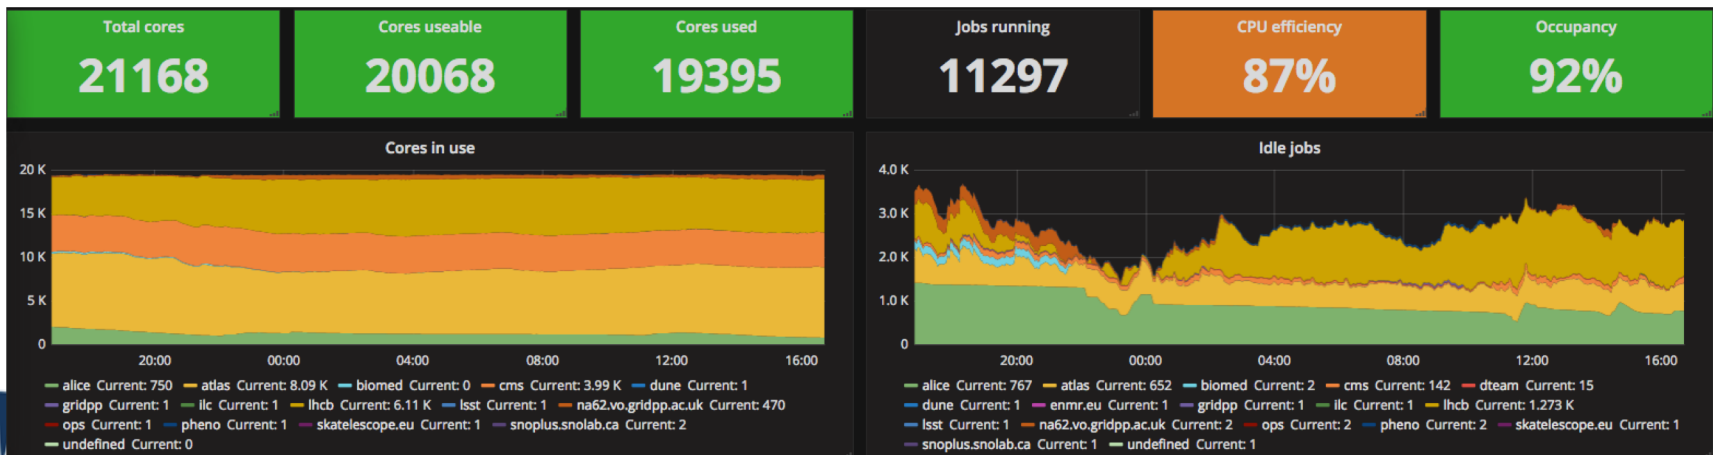
---

- Tier-I selling points are:
  - Tape system.
  - **Resources to provision services for high performance & availability.**
- I want the Tier-I to provide services that:
  - **Provide complete solutions to the common challenges facing VOs.**
  - **Harmonize diverse resources to simplify user experience.**
  - **Naturally encourage correct usage of resources.**
- Examples:
  - CVMFS - Software (and some data) distribution to all sites.
  - Rucio - Rule based data management.
  - **HTCondor - One submission point for a variety of Tier-1 resources.**
- Cloud infrastructure will allow VOs to run their own services.
  - Share pain of running crap software.



# Batch Farm

- As John mentioned:
  - We have a 20000+ core batch system and HTCondor is running brilliantly.
  - HTCondor can do everything we need: Containers, multi-core, different memory, GPUs, cloud bursting, preemptable jobs etc.
- What could possibly go wrong?!



# Euclid - an IRIS usecase

---

- Euclid is a space telescope that is currently under development that aims to accurately measure the accelerating expansion of the universe.
- They wanted to run a Monte Carlo campaign over the summer that needed 10 million CPU hours.
  - They have one person trying to run this.
  - He had built his workflow using his local batch system that happened to run SLURM.
- Somebody in IRIS setup a SLURM cluster in the cloud for him to submit too.
  - This (mostly) worked, so everybody in IRIS with cloud resources was asked if they could setup SLURM.
- **We need to have the ability to teach users to use HTCondor!**



**Things for the future?**

# Educating Users

- How do we educate users?
  - Better documentation, worked examples, code in github?!
- I believe the US has ‘dedicated’ people to help new users with HTCondor?
  - Is there a community mail list?
  - Tier-1 has Liaisons for the LHC VOs, one suggestion is to evolve their roles in future to focus on services.
  - Can we create a ‘digital asset’ that would help new users with HTCondor?



# Rucio integration

- Rucio appears to be becoming a very popular choice for distributed data management.
  - We have a test instance at RAL currently for SKA.
  - We have money from IRIS to make a multi-VO production quality instance.
- I have heard rumours that HTCondor will integrate with Rucio?
  - I am not sure what that might look like?
  - Is there anything else I need to apply for grant money for?!





# HTCondor CE

- RAL setup ARC CEs before HTCondor CE existed.
- ARC has worked very well.
  - Via ARC Control Tower, ATLAS run very complex workflows.
  - We have managed to get all other VOs to use them.
- Currently not enough benefits from HTCondor CE to justify effort to move.
  - Tier-2s often wait for Tier-1 to sort out problems before switching.
  - What killer features are we missing?



# Small sites

- HTCondor has a reputation for being complex!
- GridPP is consolidating the number of sites that run storage but is providing small amount of money to utilize CPU at lots of sites.
- ‘VAC’<sup>[1]</sup> is being recommended:
  - Very simple to setup.
  - Spins up VMs for each VO that know how to pull in jobs.
- Can we do better?
  - Central submission service at RAL.
  - VMs spin up HTCondor and join pool?

[1] <https://www.gridpp.ac.uk/vac/>



# Questions?



Alastair Dewhurst, 6<sup>th</sup> September 2018