

# The DiRAC Cloud Pilot: Application Matching to eInfrastructure

Where do I run my code/app/workflow

Jeremy Yates (DiRAC), James Hetherington (Turing)

# Understanding Our Needs

- HPC is a complex space. Data Movement between nodes is the interpretive key

Strongly parallel	Communications rate = processor rate	>100Gbs, < 1microsec	N
Weakly parallel	Comms - need quick comms to not waste cycles	>100Gbs, < 1 microsec	N
Loosely parallel	Small amount of data to be transferred e.g. fitting data	>10Gbs, >10 microsec	Y
Embarrassingly parallel	No data transferred	1Gbs, >10 mmicrosec	Y
Data Intensive	IO comms equiv to node-to-node comms	>100Gbs, < 1 microsec	?

# DiRAC wants to be a Cloud

- DiRAC-3 wants to be a cloud
- Portability, lowering barrier to access, hybrid models, users can configure resources much better for workflows, and being able to federate resources
- However the low latency issue, in which virtual VLANs perform very badly compared to bare metal, is a serious problem
- In Q1 2018 we will take 6 codes that cover the cases above
  - Alaska - openstack
  - Azure and Google Platform
- Measure
  - Faff factor - how easy to get running and port code
  - Optimisation for running on clouds
  - What breaks and why...and how do we fix it
  - Having a bare metal fix is a good fix, but lacks a certain elegance and why bother?

# Part of a greater whole

- The idea for this project came from the RCUK Cloud WG as well as DiRAC
- The main output is to build up skills and experience among RSEs and our sys admins
- We can learn how to characterise workflows, so we can offer advice and tools to make best use of resources
- Linked to STFC programme to build infrastructure for the Nel (PI Pete Clarke.
- Bottom-up joined up-ness.....

# In the context of UKT0

- The UKT0 Cloud project to run Openstack instantiations at different locations and then federate them is linked to this Project
  - If we wish to make DiRAC a cloud – starting with the Leicester-Cambridge Data Intensive Facility – then this particular project is part of the roadmap to delivering that wish.
- Measuring Application performance and solving the issues for weak and strong parallelism is very important
  - As increasingly applications (e.g. ML) will fall into this domain (Boyle et al 2018)
  - Data parallelism is heavily used
  - Parallelism is to set to be much more used.
  - UKT0 has the opportunity to help solve this problem and make the UK a leader in this type of cloud technology

# Progress so far

- Built the team from UCL, Leicester and Imperial.
- Supported by Turing, Edinburgh, Durham and Cambridge
- Measurements will be done via git-hub and ansible environment created by UCL
- Applications have been handed out to RSEs
- RSEs have run codes on DiRAC systems
- Had workshop with Azure engineer to get going
- Google now offering resources on their platform
- Then the Strike happened and this has inserted a 3 week delay into it.
  - I can say that the strike has been rock solid....