UK Computing in DUNE

Andrew McNab

University of Manchester

Overview

- UK computing organisation in DUNE
- Leverage as a theme
- Computing capacity delivered to DUNE
- Work done to use UK capacity effectively
- Development activities
- 2022 Data Challenge

DUNE UK Computing organisation

- Computing = hardware at sites and middleware which connects it
  - So not physics software. Not LArSoft etc.
- Development funded as part of WP1 of the UK construction project
  - At Edinburgh, Manchester, RAL-PPD
  - Non-funded institutes (Bristol, Imperial, Lancaster, QMUL, …) also provide computing capacity and support, attend ops meetings etc
- Monthly meetings to
  - Report status of development work
  - Review use of UK CPU and storage capacity by DUNE
- UK also contributes
  - DUNE Computing Contributions Board chair (Peter Clarke)
  - DUNE Computing International Technical Lead (Andrew McNab)

Leverage as a theme

- One of the themes of DUNE UK computing is the amount of leverage we are getting
  - Software (eg Rucio) and tools (eg GGUS tickets) from outside that we extend/adapt rather than write from scratch
  - CPU/storage capacity funded by IRIS that we apply for annually
  - Site-level operations that are funded by GridPP
- All this requires some DUNE UK effort to give them the necessary DUNE “flavour” and maintain that as things change
  - But we supply a lot more to DUNE globally than is funded by the computing part of the UK construction project
- LBNC noted the success of computing leverage in the March report, and this process was led by DUNE UK people

DUNE prod jobs: 2022 Q1/Q2

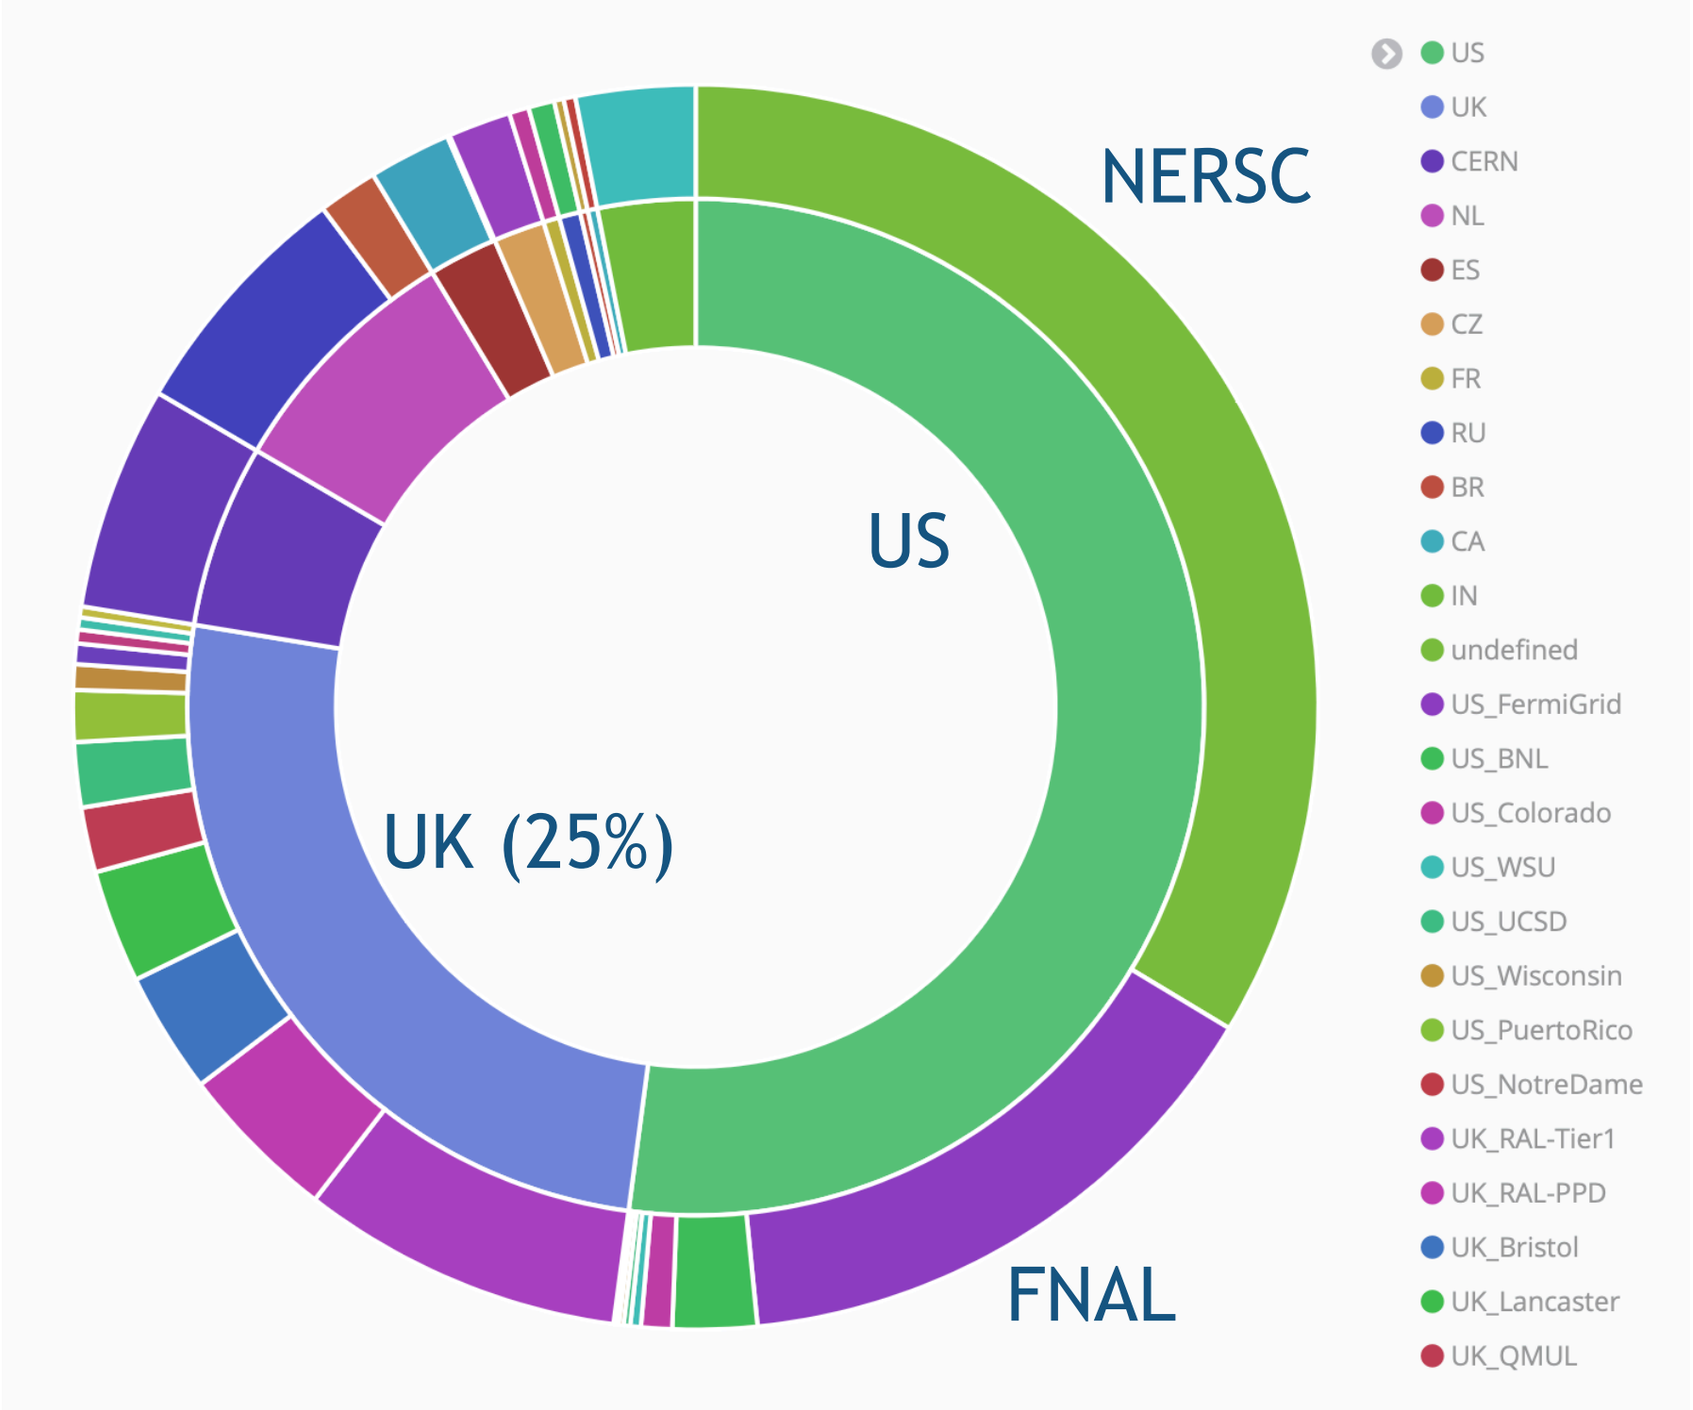

DUNE files in Rucio

- Storage used by DUNE files known to Rucio last week
  - So NOT just the amount of space provided by the sites!!!
- Some space has been cleared for the Data Challenge, especially at CERN

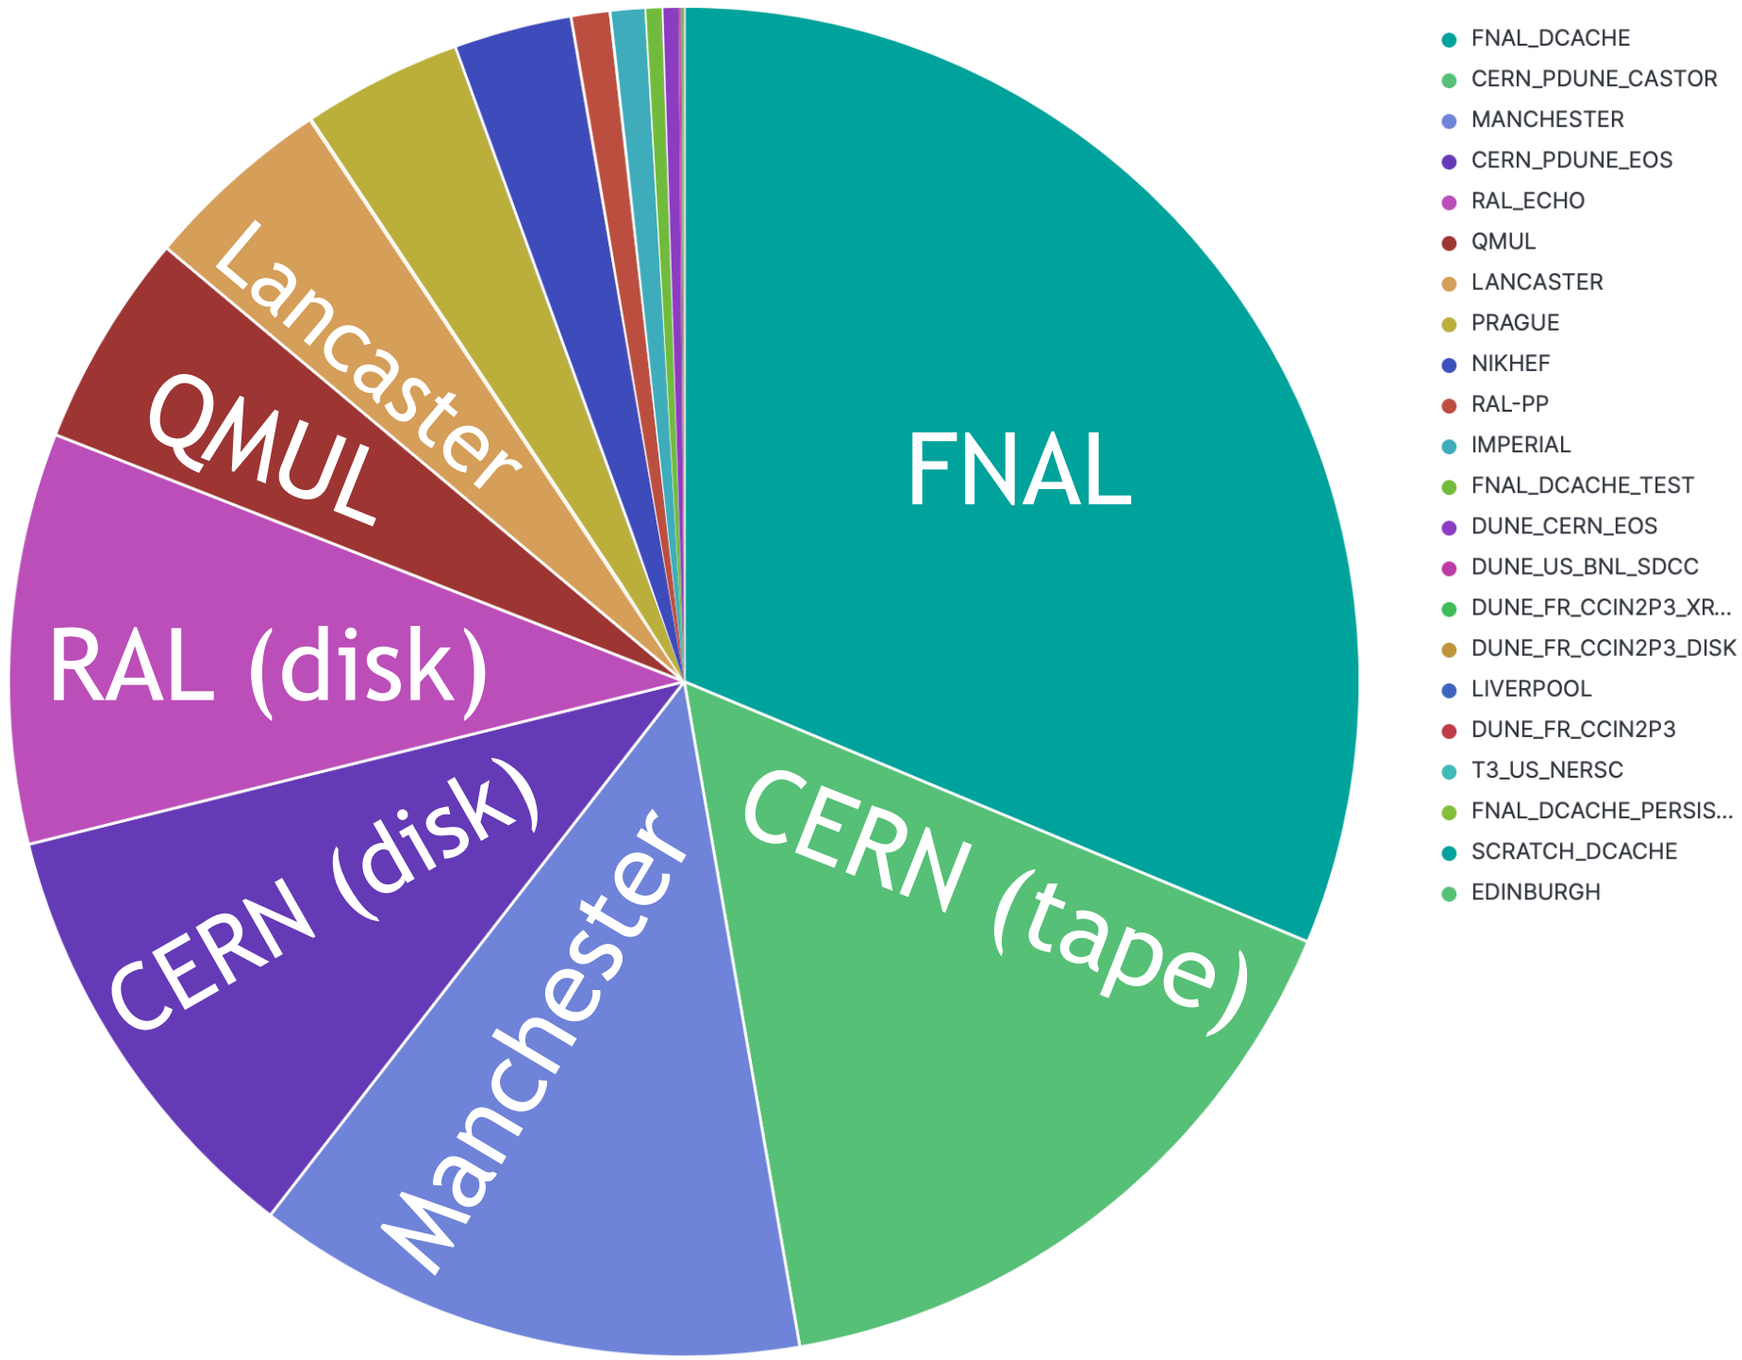

Effective use of offsite computing

- Several pieces of work fed into the UK slices in those charts
- We started weekly DUNE/UK sites meetings which non-DUNE staff at sites could join
  - Have become the weekly DUNE Computings Ops meetings
- We got DUNE set up on the GGUS ticketing system that WLCG/ GridPP sites expect experiments to use
  - Before that, no good way for DUNE staff at FNAL to contact UK sites about problems
- Set up CRIC info system and ETF site testing framework to make it easier for FNAL etc to understand state of non-US sites
- UK staff represent DUNE on various WLCG boards, including the Management Board
- All this has helped other countries too, but UK retains initial benefit

Leveraging IRIS funding

- All the work in setting up the infrastructure allows us to be credible users of computing capacity funded from outside the project
- In particular, DUNE UK staff were involved in the creation of IRIS and UKT0 which preceded it
  - IRIS funds computing for non-LHC STFC sciences (PP, Astro, Diamond users etc)
- We helped formalise annual DUNE computing requirements estimates, especially via the DUNE Computing Contributions Board
  - Needed for FNAL too, but we derive the UK estimates from it
  - These numbers **and justification** form the basis of the annual submissions to the IRIS Resource Allocation Panel
- These allocations are fulfilled by GridPP with IRIS funding
  - So this also brings in GridPP effort at sites to keep the hardware working, do software updates we need, help debug problems etc

Development activities

- We have two broad areas of development responsibility within DUNE Computing Consortium
- First, Data Management (Edinburgh)
  - This is centred on Rucio which DUNE uses to keep track of where copies of files are on the grid
  - Rucio is a well-established standalone project, but we add core functionality DUNE needs and DUNE plugins
  - See Rob's talk
- Second, Workflow System (Manchester and RAL-PPD)
  - Discovering lists of files to process, how to process them, and making sure the files are processed at the "right" site
- Plus other smaller projects: eg RAL-PPD extended ETF for DUNE

2022 Data Challenge

- A two phase data challenge is being done this year to check readiness of new components for protoDUNE data taking
- Phase I, Data Transfers: ~now
  - Registering and copying files from CERN to FNAL
  - Rely on Rucio file catalogue rather than SAM
    - With UK-contributed plugin to ensure DUNE Metadata catalogue and Rucio are consistent
- Phase II, Data Processing: September
  - Process data at multiple sites, using Rucio file catalogue to discover where input files are
  - The Workflow System will be used to run jobs “near” unprocessed files to maximise efficiency

Workflow System

- Hosted on RAL cloud by RAL-PPD
  - Have begun process of containerising the services and agents
    - Easier deployment and scaling (eg with Kubernetes)
- Manchester developed the prototype, which allows the complete process of submitting and processing requests
  - Input files defined by a metadata query
  - File locations discovered from Rucio
  - Monitoring progress via a web dashboard
  - Agents collect info, submit jobs to sites via jobsub/ GlideinWMS at FNAL, manage the lifecycle of requests
  - Workflow Allocator matches requests to sites with unprocessed files “nearby”

Workflow System components

See the CDR draft if you really want to know the details of all this!

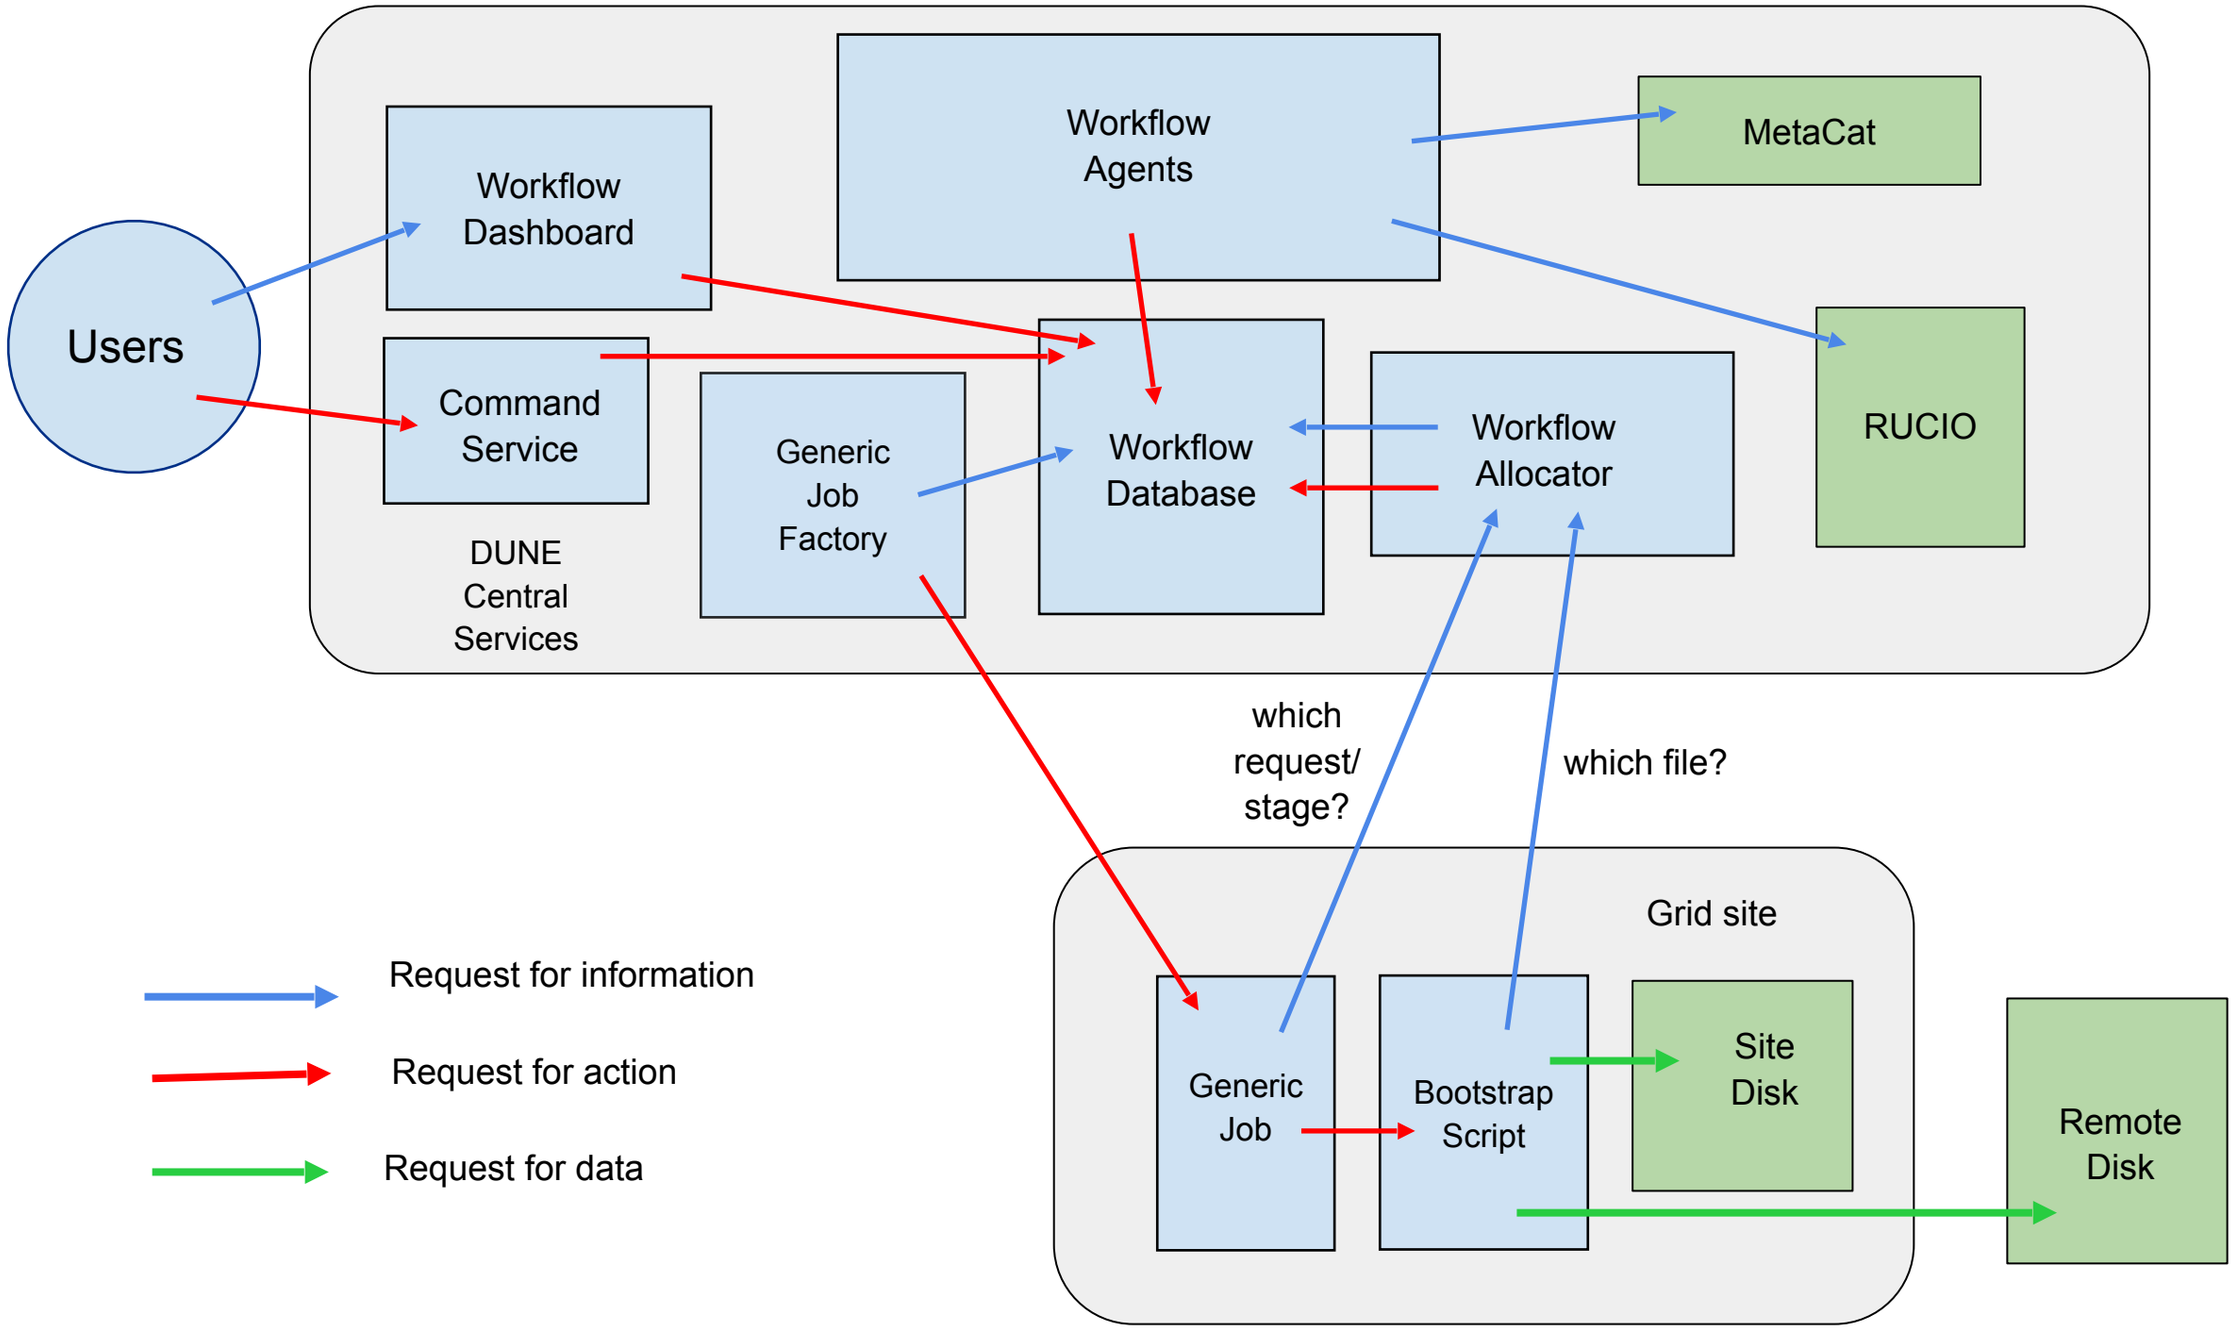

Summary

- Approach of identifying key areas where we can have lots of impact via leverage is working out
- The UK is making a significant and visible contribution to the CPU and storage supplied to DUNE
  - Underpinned by STFC funding for hardware and staff via IRIS and GridPP
- We are active members of the Computing Consortium, providing a Technical Lead (Andrew McNab) and the Chair of the Computing Contributions Board (Peter Clarke)
- As part of the construction project, we are providing development effort and leadership in Data Management and in Workflow
- The resulting systems are key components of the Data Challenge in preparation for protoDUNE data taking

Backup

RAL-PPD: ETF monitoring

- It is essential to have a way of monitoring the "health" of sites DUNE jobs run at
  - Things change, updates happen, hardware breaks
- RAL-PPD has adapted ETF monitoring used by LHC experiments for this, with support from ETF devs at CERN
- Sends out short test jobs to all DUNE sites and reports outcomes to dashboard
- Site by site, and service by service views

Local site etf

state	Host	Icons	OK	Wa	Un	Cr	Pd
UP	arc-ce01.gridpp.rl.ac.uk		3	0	0	0	0
UP	arc-ce04.gridpp.rl.ac.uk		3	0	0	0	0
UP	ce508.cern.ch		3	0	0	0	0
UP	ce511.cern.ch		3	0	0	0	0
UP	gpce03.fnal.gov		0	0	0	2	1

state	Host	Icons	OK	Wa	Un	Cr	Pd
UP	arc-ce02.gridpp.rl.ac.uk		3	0	0	0	0
UP	arc-ce05.gridpp.rl.ac.uk		3	0	0	0	0
UP	ce509.cern.ch		3	0	0	0	0
UP	etf-11.cern.ch		19	0	0	0	0
UP	gpce04.fnal.gov		0	0	0	2	1

state	Host	Icons	OK	Wa	Un	Cr	Pd
UP	arc-ce03.gridpp.rl.ac.uk		3	0	0	0	0
UP	ce503.cern.ch		3	0	0	0	0
UP	ce510.cern.ch		3	0	0	0	0
UP	etf-dune-preprod.cern.ch		8	0	0	0	0

refresh: 30 secs

Remote vs local input data

Jobs running at Liverpool

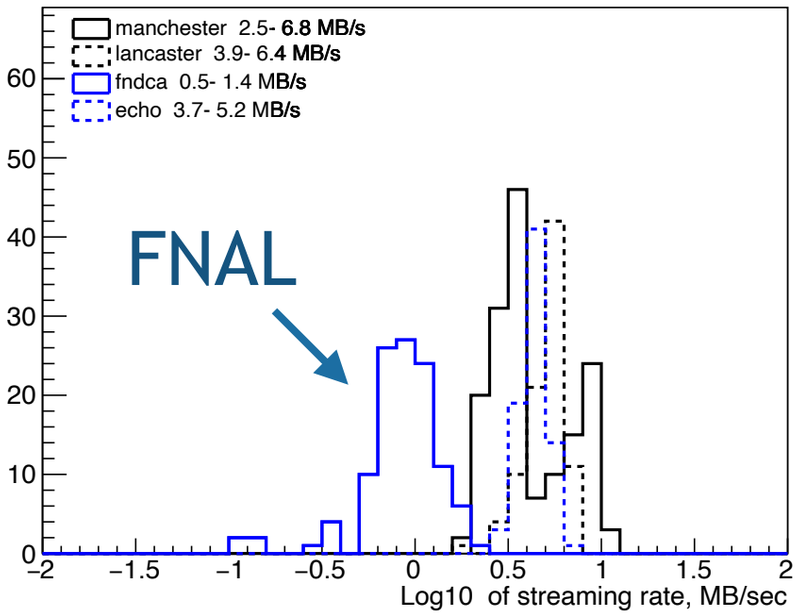

Jobs running at Manchester

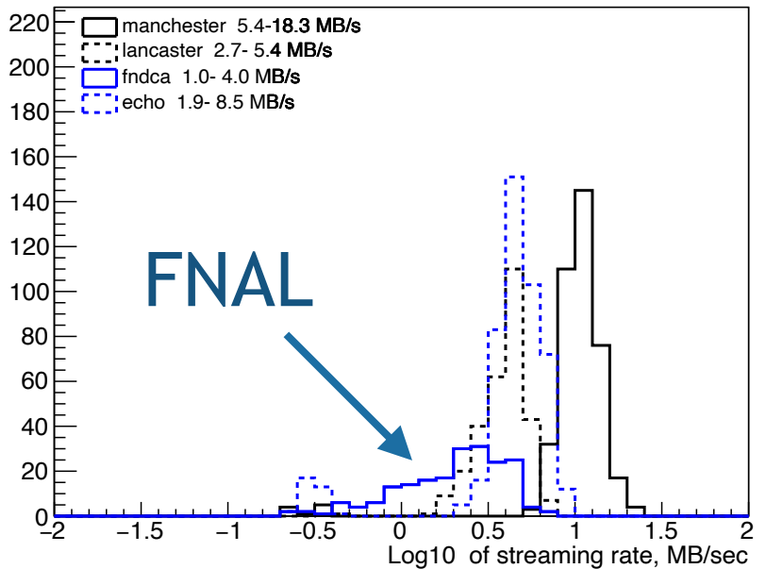

Three charts from a 2022 study on DUNE data access from 38 sites.

During these tests: it was best for Liverpool jobs to access any UK site but not FNAL; best for Manchester jobs to access Manchester; best for RAL-PPD jobs to access RAL Tier1. FNAL varies from much worse (Liverpool) to just like a UK site (RAL-PPD).

Jobs running at RAL PPD

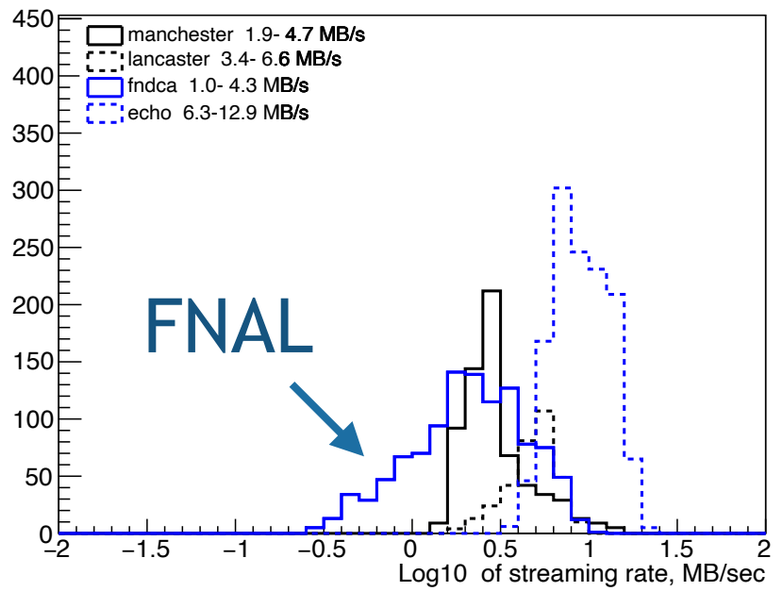

DUNE Rucio monitoring

- DUNE Rucio monitoring dashboards hosted by Edinburgh
  - ~1.5PB data moving to UK sites in last 3 months

[rucio] total replicas - UK

737,219 Total replicas

2.18PB Total bytes

[rucio] RSE allocation and usage - UK

Export

RSE	RSE Quota	IRIS Allocation	Used	Free	Free(%)
RAL_ECHO	1.0PB	1PB	1.1PB	-110.4TB	-12.139%
RAL-PP	99.0TB	(0.5PB)	97.4TB	1.4TB	1.587%
QMUL	200.0TB	1PB	81.2TB	108TB	59.387%
MANCHESTER	1.08PB	1PB	652.7TB	386.4TB	39.427%
LANCASTER	549.76TB	0.5PB	498.5TB	46.6TB	9.329%

[rucio] Replicas pie - UK

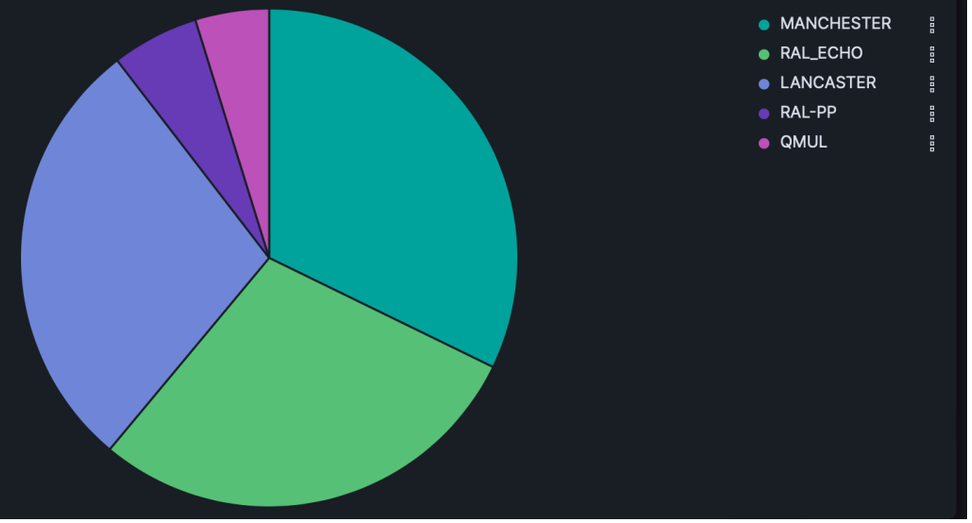

[Rucio] SRR Used History