

# DiRAC-4 design workshop

Welcome!

on behalf of the DiRAC Technical Directorate

20th March 2025

York

**DiRAC**

High Performance  
Computing Facility

# Welcome

- 11:30 - Discussion: Constructing the technical case
- 12:30 - Lunch
- 13:30 - Discussion: Scaling requirements
- 14:00 - Discussion: RAM and storage requirements
- 14:45 - Discussion: Post-processing, user-centric services, data management
- 15:15 - Discussion: Required information gathering
- 15:45 - Wrap up: next steps

Afternoon refreshment available around 14:45.

The zoom room should remain open throughout.

# DiRAC bespoke design

- We design systems to efficiently match the science, e.g.
  - TURSA: Extreme scaling for QCD
  - COSMA: Memory intensive for cosmology
- Enabling World-leading science at low cost
- DiRAC-4 should continue to deliver this bespoke system design



# The design process

- *Science Case* → **Technical Case** → *Technical Design*
- Data gathering: Science case, current projects, user survey, scoping workshop, design workshop, discussions, pilot studies, software strategy, additional inputs
- Requirements definition

# Net-zero

- High embodied CO2
- High electricity consumption
- We should only design what we really need
- We should design a Carbon-effective facility

# Capability systems

- Science that cannot (feasibly) be done elsewhere



# Optimised design

- Cost
  - Maximise science output per capital/operation/carbon cost
  - Network fabric can affect performance
    - Depending on job size and type
  - Storage needs to be suitable for the task
- Operational cost
- Net-zero

# Information gathering

- See summary document
  - Per science-case
  - Per project
  - Scoping meeting
  - Survey
  - This includes some assumptions
    - If codes have been flagged as being portable to GPU, we will assume that will happen
    - If information is unclear and clarifications not received, we have made assumptions
  - Missing information and clarifications can be added



# Notes

- Performance Analysis Methods: Workshop Series 2025

# Mentimeter

- Go to [menti.com](https://menti.com) and enter the code:
- 4353 2558

# Constructing the DiRAC-4 Technical Case 11:30-12:30

- Questions about the process
- Building the evidence base to define technical requirements implied by science drivers
- Does the science case reflect your needs? If not, what is missing?



# Scaling requirements 13:30-14:00

- How local are your communications?
- Which fraction of your workflows is Flop/s dominated?
- Which fraction of your workflows is bandwidth dominated?
- Is a roofline analysis a good approximation for your workflows? If not, why not?
- On a 200 PF system, what is the largest fraction you would use with one job?
- Is interconnect latency a constraint for your workflows? Are there ways to relax this constraint (e.g. payload bundling?)
- Do you think that shared memory between GPU and CPU would enable better scaling or easier porting to GPU architecture?
  - e.g. true shared memory or a global address space with NUMA domains

# Requirements for RAM and storage

## 14:00-14:45

- What is the largest amount of RAM you will need in one job with DiRAC-4?
- What is the largest amount of data you will need to I/O in one job with DiRAC-4?
- Do you think you can benefit from on-node fast temporary storage?
- Have you investigated domain-specific data compression, in-memory and on disk?
- Have you investigated hierarchical caching of a job data (storage-RAM-device)?
- Do you require checkpointing of large amounts of data (> 10 TB) in a single job?

# Post-processing, user-centric, data management 14:45-15:15

- If DiRAC-4 services are geographically distributed, would you rely on multi-site workflows? If yes, what bandwidth between sites would you require?
- Will you need the capability to move large data sets (>10TB) into or out of DiRAC sites? If yes, what scale of data and bandwidth will you need? Are there specific features of data transfer software tools that you need to have?
- Do you have strong requirements for data post-processing? Are these different from your main production requirements?
- Do you require domain-specific web services (Continuous Integration/Continuous Development, workflow monitoring, Databases, ...)
- Would you benefit from being able to deploy container-based applications?
- What requirements do you have regarding sharing data outside of DiRAC?
- Would you benefit from object-based storage?



# Information gathering to evidence requirements 15:15-15:45

- Are you able to provide evidence (e.g. benchmarking data, roofline analysis, etc) to demonstrate your requirements?
- Have you already assessed whether your code is using optimal implementations of algorithms? If yes, can you provide evidence to support this?
- What level of human resources would you need to make progress with the two points above in the next 1 to 2 years?
- Do you have standard, absolute figures of merit for your codes/workflows that could be used in procurements?

# Wrap up - next steps