DiRAC

High Performance Computing Facility

DiRAC 3 Phase 2

Biagio Lucini

On behalf of the DiRAC Technical Directorate

The system upgrade process

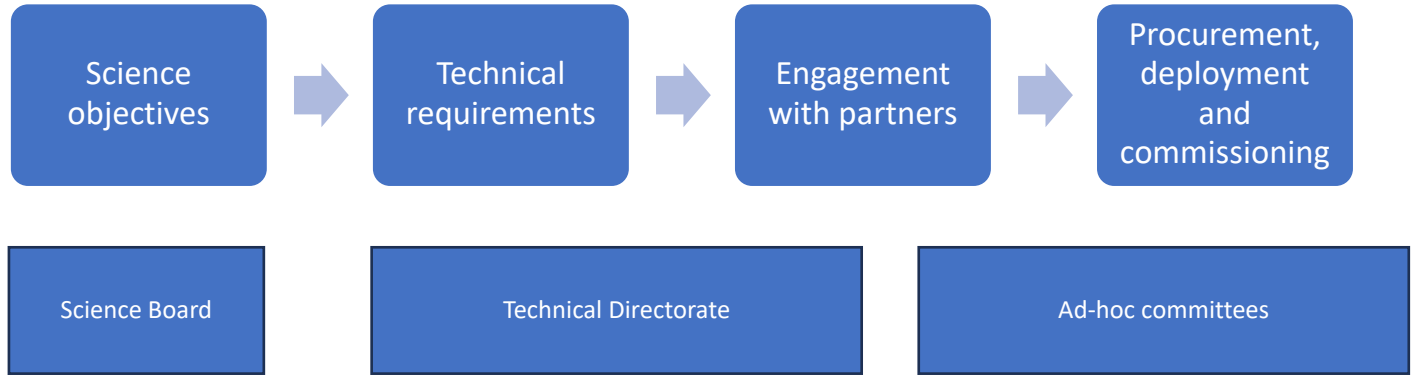

The Technical Directorate

- Established in 2020
- Membership: Co-chairs; Director of Innovation & Technology; 1-3 additional DiRAC community experts
- Current members: Alastair Basden, Antonin Portelli, Jeremy Yates, Kieran Leach, Biagio Lucini
- Relevant duties
  - Develop the technical case for DiRAC
  - Approve procurement plans
  - Oversee the smooth running of the facilities
  - Discuss future directions in software and hardware (including strategic partnerships)

DiRAC 3 Phase 1 (and before) systems

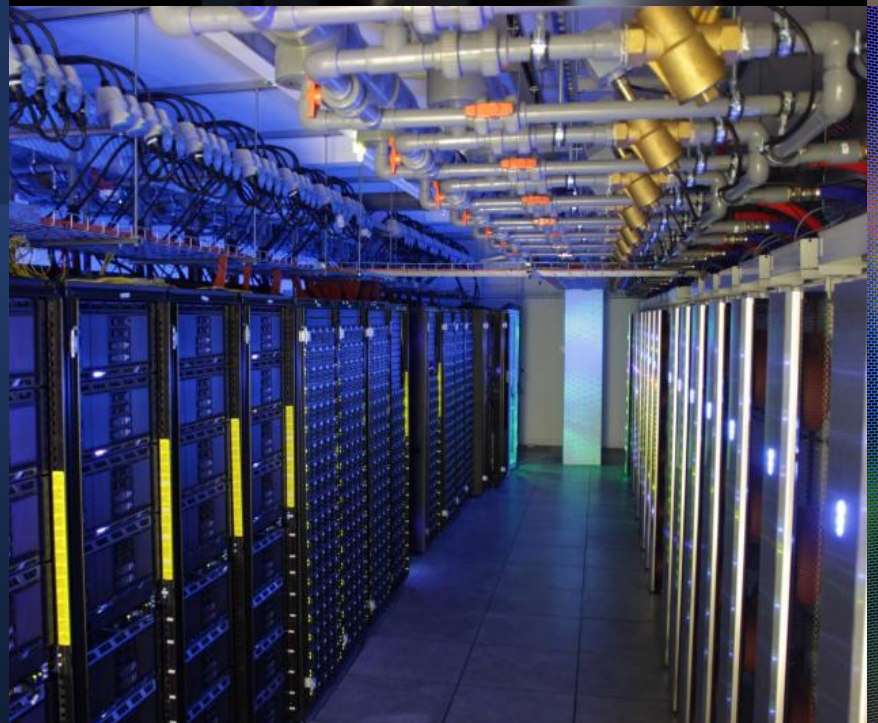

Timeline

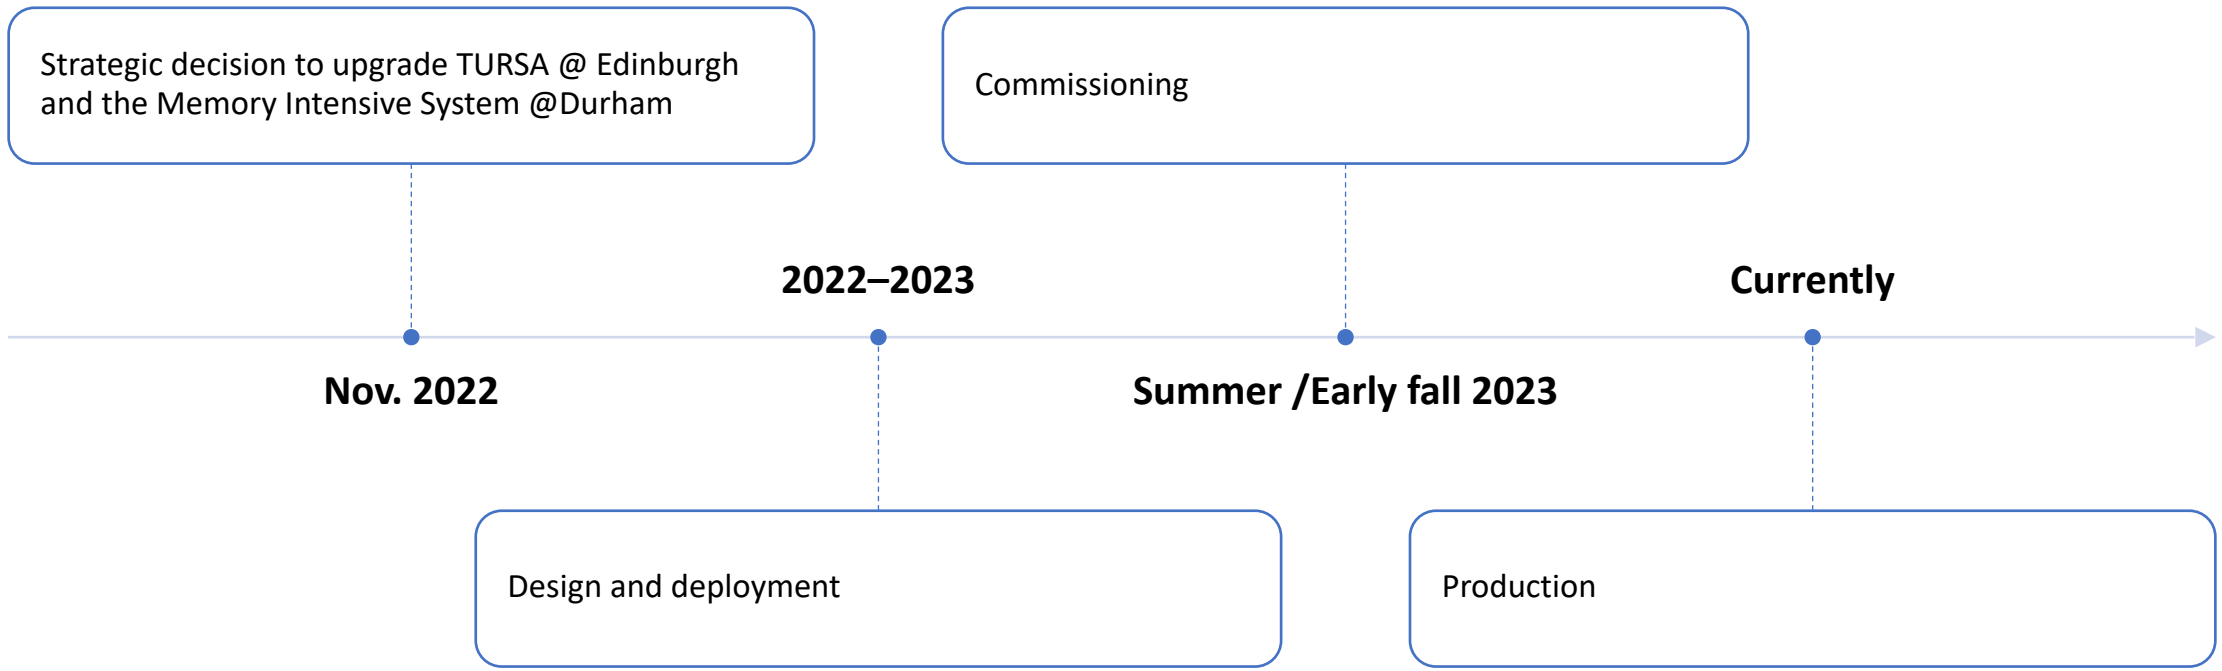
Strategic decision to upgrade TURSA @ Edinburgh and the Memory Intensive System @Durham
Commissioning
2022–2023
Currently
Nov. 2022
Summer /Early fall 2023
Design and deployment
Production

Extending access to DiRAC services (Credits: M. Wilkinson)

- Recently deployed additional hardware at two sites:
  - Edinburgh – Extreme Scaling Service – 256 additional A100-80GB GPUs
  - Durham – Memory Intensive Service – 168 nodes (21.5k cores) of AMD Milan CPUs
- Between Oct 2023 and Mar 2024:
  - 50% of new hardware allocated via Director's Discretionary Call for DiRAC projects
    - 400k GPU hours and 40M CPU core hours were allocated
    - Call was 2x over-subscribed
  - 50% being made available to non-STFC groups:
    - CompBioMed consortium
    - HECBioSim and CCPBioSim consortia
    - PAX ExCALIBUR project (materials modelling)
    - MRC-University of Glasgow centre for virus research
    - UCL Centre for AI
    - UKAEA
- **Goal of non-STFC allocations:** Offer sufficient resource allocation to allow production-scale calculations which would otherwise not be possible on a UK service
    - Explicit effort on the DiRAC side and within consortia, to inform a future cost model

COSMA 8

current status

COSMA8 COMPRISES

528 COMPUPTE NODES. (360 WITH 2x AMD ROME PROCESSORS AND 168 WITH AMD MILAN 7763 PROCESSORS), 1TB RAM AND NON-BLOOCKING HDR200 INFINIBAND NETWORK

2x 2TB LOGIN NODES WITH 64 CORES (DUAL AMD 7h12 ROME 7542 PROCESSORS)
TWO FAT NODES WITH 4TB RAM AND 128 CORES

GPU NODES WITH NVIDIA A100, V100 AND AMD MI200 AND MI100 GPUs – 7 NODES WITH 18 GPUs13PB BULK LUSTRE STORAGE

1.2 FAST SCRATCH STORAGE (~350GBytes/s)
HIGH-CACHE NODE (768MB CACHE MILAN-X) WITH 127 CORES AND 1TB RAM

26PB TAPE ARCHIVAL FACILITY

Additional capabilities include

- 168 nodes and 6 PB of Lustre storage
- Ability to run 512 node jobs on 64k cores

TURSA

current status

TURSA COMPRISES

4272 AMD CPU CORES RUNNING AT 2.6/3.3GHz

178TB OF SYSTEM MEMORY

712 * A100 NVIDIA GPUs

200GB/s HDR IB NON-BLOCKING INTERCONNECT

8PB TAPE BACKUP

Additional capabilities include

- 64 nodes with a total of 256 A100-80 GPUs
- 1.5 PB of Lustre storage

Key takeaway messages

The hardware follows the science goals

The strategy is always proactive, never reactive

Performance is always measured by user codes

The user community plays a central role in defining the strategy

The approach is anticipated to stay in place for future opportunities

The approach is open and cooperative