

Durham H&ES Testbeds

Alastair Basden et al

Alastair Basden
DiRAC / Durham University



DiRAC
High Performance
Computing Facility



ExCALIBUR H&ES @ Durham

- Has been an enabling factor for the HPC Hardware Lab @Durham
 - Providing new technology access for users across the UK
 - With a single login: low entry barrier
- Multiple ExCALIBUR H&ES systems
 - Supporting various ExCALIBUR software projects
 - And beyond: other use cases
- HPC system as a Living Lab
 - Not a static procurement

ExCALIBUR: Preparation for Exascale

- H&ES aimed to provide access to potential Exascale compute hardware
 - Technologies relevant for the UK HPC journey
 - Code porting, testing and profiling
 - Benchmarking for future HPC procurements

Benchmarking systems

- H&ES has enabled in Durham (since 2019):
 - GPU: All the data centre generations from the 3 major vendors
 - H100, A100, A30, V100, PVC, MI300, MI210, MI100, MI50
 - CPU: All the data centre generations from the major vendors
 - Bergamo, Genoa, Milan, Rome, Emerald Rapids, Sapphire Rapids, Ice Lake, Cascade Lake, Sky Lake, Grace
 - Accessible to users UK-wide

CXL-precursor systems

- Liquid composable infrastructure
- Gen-Z testbed (obsolete: merged with CXL)
- CerIO composable system
- Not user-visible
 - RTP expertise/experience
 - This is valuable
 - But used by users

Data Processing Units (Smart NICs)

- 2 generations of NVIDIA BlueField
 - Embedded ARM processors
 - Running CentOS/Ubuntu
 - Used in host-separated mode
 - MPI offloading and task stealing
 - 24-node cluster
 - Suitable for at-scale testing and small production workloads
 - Note: Good to be able to fund >1!
 - Scale is key in HPC!

Quantum computing

- Hosting of H&ES access to DWAVE
 - Quantum annealing system
- ExCALIBUR access to QuEra
 - Neutral atom system
 - QEVEC project

Networking / Storage

- Various H&ES-funded projects
 - Not necessarily visible to end-user
 - Somewhere to store or transport their data
 - RTP expertise and experience are key

Cooling tech

- Not visible to users:
 - Again RTP experience and expertise is key!

Views on H&ES

- Good source of funding to advise new procurements
- Light touch approach was appreciated
- Ability to create digital assets (FTEs) was essential
 - Lack of overheads was unfortunate
- Short spending periods make it difficult to plan
 - Annual calls meant that timescales weren't always appropriate
- Finished too early:
 - £1-2m urgent funding required for next FY (from April)
 - To keep people in employment

H&ES-2 thoughts

- Longer spend periods (not FY limited)
- More frequent calls
 - Annual does not match well to dynamic tech timelines
- Focus on high TRLs
 - Lower TRLs should be/are funded elsewhere
 - Emphasis on placing kit into production / making it available to users
 - No shelfware
- Related calls for FTEs to make use of hardware
 - Remove disconnect with the software projects

MI300X system

- Latest addition:
 - 8-way GPU system (Dell)
 - Global address space
 - Available for UK researchers to test/port/profile/benchmark
 - Also watch out for MI300A system coming soon

A photograph of a server rack with a purple geometric overlay on the right side. The server rack is composed of vertical panels with circular ventilation holes and small square indicator lights. The image is slightly blurred, and the purple overlay is a solid color with a sharp edge.

Alastair Basden
DiRAC / Durham University