DiRAC & IRIS

IRIS TWG Presentation

23rd January 2019

Mark Wilkinson

Director, DiRAC HPC Facility

DiRAC Services

Extreme Scaling "Tesseract" (Edinburgh)

2 Pflop/s to support largest lattice-QCD simulations

Data Intensive "DlaL" and "CSD3" (Leicester & Cambridge)

Heterogeneous architecture to support complex simulation and modelling workflows

Memory Intensive "COSMA" (Durham)

230 TB RAM to support largest cosmological simulations

DiRAC Services

Service	Characteristics (by February 2019)
Extreme Scaling "Tesseract" (Edinburgh) multi-core	32736 cores; 20 GPUs; > 1.8 Pflops
Memory Intensive "Cosma" (Durham) high-memory, multi-core	Max. memory footprint: 230 TB; 12600 cores; 0.44 Pflops
Data Intensive "CSD3_CPU & DIaL" (Cambridge & Leicester) multi-core	29564 cores; 2.3 Pflops, 192/384 GB RAM per node; Fat nodes: 3x1.5TB Shared Memory: 6TB SuperDome flexn; ~0.5PB flash storage
Data Intensive Xeon Phi "CSD3_KNL" (Cambridge) many-core	65 Tflops (jobs up to 500 Tflops can be run)
Data Intensive GPU "CSD3_KNL" (Cambridge) many-core	130Tflops (jobs up to 1 Pflops can be run)

IRIS Objective B

"Augmentation of the e-Infrastructure to allow the UK to design, construct and establish positions of global leadership for future exploitation of approved "in-build" investments."

Potential DiRAC/IRIS Opportunities:

1. Access to novel hardware for prototyping and benchmarking:
   - Arm (4000+ cores)
   - Large-volume (0.5PB) flash-accelerated storage
   - Shared Memory (6TB)
   - Fat nodes (1.5TB)
   - Large-scale GPU
   - Xeon-based system with OPA *and* Ethernet
   - Range of RAM 4-18 GB/core
2. Seedcorn/discretionary allocations available

DiRAC Innovation and Industrial Engagement

Service	Activities
Data Intensive: Cambridge	• Data Accelerator: Large-scale flash storage for accelerated I/O; • Support for multi-architecture workflows (Xeon, Xeon Phi, GPU). • Openstack development work for cloud middleware
Data Intensive: Leicester	• Arm system installed in January 2019 in collaboration with HPE/Arm/Suse Catalyst UK programme; • Large-scale flash storage to complement SMP workflows
Extreme Scaling (Edinburgh)	• Deep co-design work with Intel on DOE Aurora 2021 hardware development, including network work that influenced multithreaded extensions to their MPI • Collaboration with USQCD DOE Exascale Computing Project, including Nvidia and DOE Summit GPU portability project. • Joint patent application filed with Intel relating to new Deep Learning hardware ideas.
Memory Intensive (Durham)	• Large-scale flash storage for checkpointing of whole-system simulations; • HSM pilot; Two proof of concept systems being installed;
Project Office (UCL)	Data Transfer project; Public cloud pilot study; Durham/ Cambridge network upgrade project

IRIS Objective D

"Enhanced federation of computing across STFC for improved cost effectiveness and sharing of expertise, and as a prerequisite for developing a National eInfrastructure."

Potential DiRAC/IRIS Opportunities:

1. Shared hardware deployments - e.g. Cambridge OPA/Ethernet
2. Shared tape storage/archive/data curation service
3. Access to transient S3 storage
4. Joint training/RSE workshops to share experience
5. Automation of production/overflow access to IRIS/DiRAC resources
6. Capability for workflows that move from IRIS to DiRAC resources and back (or vice versa)
7. Digital assets work discussed this morning

Issues for discussion:

1. Required access modes:
   - ssh/batch - available at all DiRAC sites
   - containers - pot. available at Cambridge, Durham, Leicester
   - OpenStack - available at Cambridge
2. AAAI
3. Network improvements to facilitate cross-usage of services
4. Tools required to facilitate use of services