

Tier 2 Updates

Isambard system specification:

- Cray system
- **10,000+** ARMv8 cores
- Cray software tools
 - Compiler, math libs, tools...
- Technology comparison:
 - x86, Xeon Phi, Pascal GPUs
- Phase 1 installed March 2017
- The ARM part arrives early 2018
 - Early access nodes from Sep



I.K. Brunel 1804-1859

Codes we'll focus on for Isambard:

- The top 10 most heavily used codes on Archer:
 - VASP, CP2K, GROMACS, CASTEP, HIPSTAR, UM, ONETEP, LAMMPS, WRF, Oasis
 - Note 8 of these 10 codes is written in **FORTRAN**
- Additional codes relevant to Isambard project partners:
 - OpenFOAM, OpenIFS, ...
- Want to collaborate wherever possible!
 - **Let's avoid duplicating effort**

-
- Isambard RSE support:
 - 2 RSEs in total
 - Split 0.5 FTE across each of Bristol, Bath, Exeter and Cardiff
 - Will communicate via Slack etc
 - Will initially focus on porting/optimising our target codes, then later, supporting the open service
 - When service fully open later in 2018, can apply for time and also ask for RSE support

For more information:

- <http://gw4.ac.uk/isambard/>
- <https://www.epsrc.ac.uk/blog/isambardhpc/>
- Twitter: @simonmcs
- Please get in touch to collaborate on optimising applications:
simonm@cs.bris.ac.uk
- Bristol is hiring faculty staff in CS to focus on HPC



The
University
Of
Sheffield.


JADE and UoS

Twin Karmakharm

RSE, University of Sheffield
26th June 2017

The JADE System

- 22 NVIDIA DGX-1
 - 3.740 PetaFLOPs (FP16)
 - 2.816 Terabytes HBM GPU Memory
- 1PB filestore
- P100 GPUs - Optimised for Deep Learning
 - NVLink between devices
 - PCIe to Host (dense nodes)






The University
Of
Sheffield.

The JADE System

- Use cases
 - 50% ML (Deep Learning)
 - 30% MD
 - 20% Other

NVIDIA DGX-1 Delivers 56X More Performance



CPU is dual socket Intel Xeon ES-2697 v3. 170TF is half precision or FP16.



Procurement, Hosting and Access

- ATOS have been selected as the preferred bidder
 - Following procurement committees review from tender
 - Running costs to be recouped through selling time to industrial users
- To be hosted by STFC Daresbury
 - Will run SLURM scheduler for scheduling at the node level
- Resource allocation
 - Open to all without charge
 - Some priority to supporting institutions
 - Light touch review process (similar to DIRAC)



The
University
Of
Sheffield.

Co-Investigators, Governance and RSE Support

- All CIs have committed RSE support time for their local institutions
 - To support local users of JADE system
 - Role of TIER2 RSE network to be agreed
- Training
 - Some commitment to training offered by some CIs (Alan Gray at EPCC, Paul Richmond EPSRC RSE Fellow)



The
University
Of
Sheffield.

Co-Investigators, Governance and RSE Support

- Governance via steering committee
 - Prof. Anne Trefethen (chair), Oxford University CIO and PVC
 - an EPSRC representative
 - Alison Kennedy, Director of STFC Hartree Centre
 - Prof. Simon McIntosh-Smith, Bristol
 - a representative of the machine learning community
 - a representative of the molecular dynamics community



The
University
Of
Sheffield.

Current progress

- Pilot access call has gone out to each institution

At the UoS

- Investigating container technology (Singularity)
 - Workstation - ShARC (local cluster) - JADE
- GPU management on SGE
- Specialised support for Deep Learning, Simulation and GPU programming/optimisation
- Training & tutorials




The University Of Sheffield.



UNIVERSITY OF OXFORD



Queen Mary University of London



The University Of Sheffield.



University of BRISTOL



A WORLD TOP 100 UNIVERSITY




circrus

ARCHER Champions, 27 June 2017

Andy Turner


a.turner@epcc.ed.ac.uk



- Flexible HPC system
 - Support use modes that would be difficult on ARCHER: e.g. ISV applications, parallel data science
- Integration with HPC around the UK
 - Access to RDF and Tier-2 RDF to allow workflows to span multiple resources in a simple manner
- 280 node HPE(SGI) ICE XA
 - 10,080 cores (2x 18-core Broadwell) per node
 - 128 GiB memory per node
 - DDN Lustre file system
 - Single rail FDR Infiniband hypercube
- Tier-2 RDF
 - Based on DDN Web Object Scalar Appliances
 - Total 1.9 PiB usable
- Coordination across EPSRC Tier-2 sites
 - Technical working group
 - SAFE integration and development (if wanted by sites)



Callum Bennetts/Maverick Photography



RSE Staffing and Organisation

- 2.5 FTE RSE support available
 - Notionally split into HPC and Data Science but we plan to use as an integrated whole to support user community as required
- RSE support staffed by EPCC
 - Provided by multiple staff members depending on skills required at that moment
- Organisation
 - Close collaboration with EPCC ARCHER CSE Team
 - Cirrus RSE staff attend ARCHER CSE team meetings to ensure knowledge exchange
 - Slack instance including Cirrus RSE's and ARCHER CSE team




Cirrus RSE Work Structure

- Respond to service desk queries requiring technical support:
 - Second/third-level service desk support
 - Complete Technical Assessments for Cirrus access
 - Input to service operations group, change/problem/incident mgmt. and continual service improvement
- Technical project to support service users. Examples:
 - Evaluation of new technology/software
 - Benchmarking and profiling
 - Resource usage analysis
- Training
 - Running, collaborating on, and developing training



Accessing RSE support

- Query through service desk
 - Most common way to access support
 - User has a technical issue that needs to be resolved, needs a piece of software installed or has a question that needs answering
 - All users treated equally – no distinction based on where they work
- Suggestions for technical project areas
 - Will consult users on a yearly basis to assess which service development areas are important to them
 - Consultation will feed into design of technical projects for RSE staff
- Requests for local training
 - Will feed into UK HPC training in collaboration with other sites



Thomas: Tier 2 Hub in Materials and Molecular Modelling



UCL

Purpose

- Dedicated for UK Materials and Molecular Modelling community
- Primary target job size 48-120 cores (2-5 nodes)

Hardware

- 720 Lenovo x86-64 nodes, 17.2k cores total
- 24 cores per node
- 128GB RAM per node
- Intel OmniPath interconnect
 - 1:1 nonblocking in 36 node blocks (864 cores)
 - 3:1 between blocks and across system

Getting a ~2x speedup over current clusters at UCL & Imperial, likely OmniPath vs InfiniBand. (LAMMPS Rhodopsin benchmark, ONETEP)






RSE involvement

- On a per-institute basis
- UCL Research Software Development, led by Dr James Hetherington
 - will provide programming & tuning support for UCL researchers targeting Tier 2 platforms (not just Thomas)
 - One FTE provided for all UCL Tier 2 users, from across the RSD group
 - Following existing model for UCL:
 - Specific programming projects costed and funded from grants
 - General advice, support, training for free
- On the agenda to be discussed by the other institutes in the next few months




Cambridge Service for Data Driven Discovery (CSD3)


- **Tier-2 Theme:** Data Intensive Computation and Analytics
- **EPSRC PI:** Prof. Paul Alexander
- **Technical Director:** Dr. Paul Calleja
- **RAC lead:** Prof. Mike Payne
- **Early Science:** mid July 2017, policies under definition
- **General Service:** October '17 (first join Tier-1/Tier-2 call for proposal)




Computational
Chemistry



Material
Science



CFD and
Combustion



Smart Cities



Health
Informatics

CSD3 – HW components and Delivery

Ready for Early Science

- Peta4-KNL (Intel Phi 7210)
 - 21,888 cores, 508.9 TFlop/s, #404 WW Top500 (June 2017)
- Wilkes-2 (NVIDIA GPU P100)
 - 360 GPU, 1.19 Pflop/s, #100 WW Top500 & #5 WW Green500 WW (June 2017)

Under Deployment

- **Storage:** new 5 PB Lustre, up to 10 PB tapes
- **SkyLake system:** 24,576 cores, ~1.0 PFlop/s
- **Hadoop:** 50 nodes (600 TB) + 12 nodes (288 TB) (+ 2 head nodes)

Under Deployment

- **Storage co-design:** fast SSD Lustre with SLURM enhancements for stage in/out

CSD3 – Support & RSE

- 3 FTE in Cambridge already present, fully allocated by first half 2018
 - Co-development & optimization via RAC, deep support, training and outreach
 - It is *effort*, not *people*. Capacity planning managed by Head RSE
 - Cambridge currently recruiting 1 new person to increase projects bus factor
- 5.5 FTE promised as support from other institution
 - First point of contact in institution X
 - happy to leave coordination to “home” institutions, just communicate availability to RAC
 - Access (if needed) user support tools, contributing to training material & documentation
- Tier-2 “Champion” / Cambridge RSE contact: Mr. Filippo Spiga

Aston University



HPC Midlands Plus

Prof. Steven Kenny

Department of Materials

Loughborough University



Centre Facilities

- 14,336 x86 cores – consisting of 512 nodes each with 2 x Intel Xeon E5-2680v4 cpus with 14 cores per cpu and 128 GB RAM per node
- 3:1 blocking EDR Infiniband network, giving 756 core non-blocking islands
- 1 PB GPFS filestore
- 5 x 20 core POWER8 systems each with 1 TB RAM connected to the Infiniband network
- Dedicated 10 TB filestore for prestaging files

RSE Support

- 4 FTE total
 - (Loughborough 1 FTE, all other sites 0.5 FTE)
- RSE Network coordinated by Louise Brown (EPSRC RSE Fellow). Based at individual sites, but will collaborate on projects and work together
- Mechanism for Mid+ sites to bid for RSE time, with coordinates across the centre to work strategically and avoid duplication.
- Will publish what the RSEs are working on across the Tier-2 centres
- No free external RSE support

National RSE Resources



Facilities Access Training Applications Contact



ABOUT US

We bring together different organisations involved in HPC in the UK to provide information to users of HPC facilities.



HPC-UK provides information for users and providers on UK HPC facilities. This includes information on facilities available, how to gain access, what training courses are available and links to other useful resources. We are made up of a variety of organisations in the UK including EPCC and UK RSE.

About HPC-UK

<http://www.hpc-uk.ac.uk>

The screenshot shows a Slack channel named #ukrse. The channel has 13 members, including Christopher Woods, Iain Smith, Andrew Williams, Chris Eadall, Dave Axeman, Filippo Sgalla, Favian Rodriguez, James Grant, James Price, Keith Price, Keith Baker, Louise Brown/eng@hpc-uk.ac.uk, and Phil Kemp. A pinned message from Christopher Woods, dated March 29th, 11:51 AM, discusses the Tier-2 RSE standard. The message states that the standard is the product of a discussion arranged by the EPSC, involving the RSEs and the Tier-2 RSEs. It mentions that the UK's standing on HPC is strong, with the UK Research Council (UKRC) and the Engineering and Physical Sciences Research Council (EPSRC) leading the effort. The message also notes that the standard is a result of a discussion arranged by the EPSC, involving the RSEs and the Tier-2 RSEs. The message is pinned and has 1 reply from Keith Baker, dated March 29th, 10:48 AM, which says "Hi Chris".

RSE Leaders Meetings (network of leaders of RSE groups, subset of UKRSE).

Held twice a year, next one is 18th July in Sheffield.

Open to people who lead RSE groups. Ask if you would like an invitation.

<https://ukrse.slack.com>

Tier 2 RSE Discussion

- Where should we share the names of codes that we are working on? <http://www.hpc-uk.ac.uk>?
- How should Tier-2 RSEs ask for help from each other, or otherwise talk to each other? UKRSE Slack?
- How can Tier-2 RSEs collaborate on similar codes?
- How can Tier-2 RSEs collaborate with Tier-1 or Tier-3?
- Do Tier-2 RSEs want to meet up regularly, i.e. via Tier-1/Tier-2 Champions, RSE 2017 Conference or other venues?
- Any hints and tips on organising RSE work?