Tier-2

Six New Tier-2 Centres

- EPSRC provided ~£20m to fund a new wave of Tier-2 HPC centres
  - Tier-2 sits between Tier-1 (Archer) and Tier-3 (Uni)
- EPSRC provided capital for the hardware
- Regional centres provide funding for OpEx, e.g.
  - Power and hosting
  - System administrators
  - Research Software Engineers

Lightning talks

- Cirrus – Andy Turner
- Jade – Paul Richmond
- Peta-5 - Filippo Spiga
- Thomas – Heather Kelly
- HPC Midlands Plus – Steven Kenny
- Isambard – Christopher Woods

ARCHER Champions, 10 Feb 2017

Andy Turner

a.turner@epcc.ed.ac.uk

epcc

THE UNIVERSITY OF EDINBURGH

Cirrus System

- 280 node 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
- 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

cirrus

epcc

THE UNIVERSITY OF EDINBURGH

Access

- 30% resources for EPCC
  - Industrial projects
  - Research projects
- 70% EPSRC
  - Apply via regular lightweight calls (similar to ARCHER RAP)
  - Independent panel awards resources
  - Access free for EPSRC researchers
  - Other RC researchers charged by resources requested
- Coordination between ARCHER/Tier-2 sites:
  - Multiple systems for research workflows
  - Exchange projects between Tier-2 and ARCHER to ensure best fit
  - Avoid applicants applying for same resources twice

cirrus

epcc | THE UNIVERSITY OF EDINBURGH

Support and Timescales

- In depth technical support from EPCC experts provided as part of service
  - Technical advice
  - Porting, performance analysis and optimisation
  - Software installation and support
- Online training planned (initially similar to ARCHER Driving Test)
- Cirrus Champion: Andy Turner (a.turner@epcc.ed.ac.uk)
- Full user access 1 April 2017
- Documentation and information already available:
  http://cirrus.readthedocs.io

cirrus

epcc | THE UNIVERSITY OF EDINBURGH

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

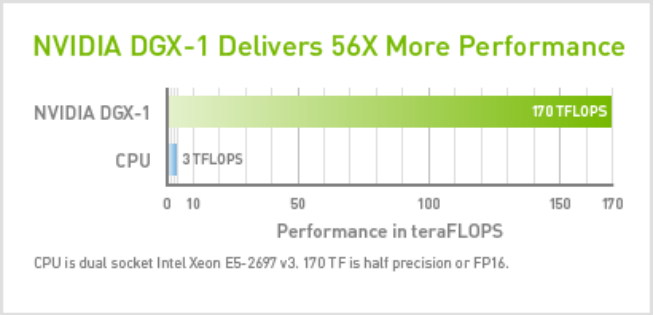

Procurement, Hosting and Access

- 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)

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 come Cis (Alan Gray at EPCC, Paul Richmond EPSRC RSE Fellow)
- 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


UNIVERSITY OF
OXFORD


UNIVERSITY OF
Southampton


Queen Mary
University of London


University of
BRISTOL


THE UNIVERSITY OF EDINBURGH


UCL


Disce Doce
RERVM COGNOSCERE CAVSAS
The
University
Of
Sheffield.


KING'S
College
LONDON

Cambridge Tier-2 “Peta-5” – Current Status

- **Tier-2 theme:** Data Intensive Computation and Analytics
- **EPSRC PI:** Prof. Paul Alexander
  - **Local Director:** Dr. Paul Calleja
  - **Co-PI:** G. Pullan, L. Drumright, MC. Payne, S. McIntosh-Smith, O. Kenway, M. Giles, C. Schönlieb, SJ. Cox, PD. Haynes, F. Fraternali, M. Wilkinson
  - **Institutions involved:** Cambridge, Oxford, Bristol, KCL, UCL, ICL, Southampton, Leicester
- **RAC:** under definition
- **Total expenditure:** ~£10m (£5m EPSRC funded)
- **Procurement:** concluded before Christmas, contract signed.
- **Deployment:** KNL & GPU - April ’17, x86 SkyLake & “co-design I/O” - July ‘17
- **General Service:** July / August ‘17

Cambridge Tier-2 “Peta-5” – HW components and Delivery

- **KNL:** 21,888 cores, ~0.5 PFlop/s (+ 4 login nodes)
  - 342 nodes, 7210 (64 cores), 96 GB RAM, 120 GB SSD, OPA 2:1
- **GPU:** 1,080 cores, 360 GPU, ~1.0 PFlop/s (+ 4 login nodes)
  - 90 nodes, 12c E5-2650 v4, 4x GPU P100 16 GB, 96 GB RAM, 120 GB SSD, EDR 2:1
- **SkyLake:** 24,576 cores, ~1.0 PFlop/s (+ 8 login nodes)
  - (low mem) 384 nodes, 2x 16c “6142”, 384 GB RAM, 120 GB SSD, OPA 2:1
  - (high mem) 384 nodes, 2x 16c “6142”, 192 GB RAM, 120 GB SSD, OPA 2:1
- **Lustre:** 5 PB usable
  - 16 OSS, 2 MDS, 4 MDT, 8 LNET + IEEL support (over OPA)
- **Tape:** housing up to 10 PB tapes
- **Hadoop:** 50 nodes (600 TB) + 12 nodes (288 TB) (+ 2 head nodes)

Filippo SPIGA <fs395@cam.ac.uk>

Cambridge Tier-2 “Peta-5” – Support & RSE

- Pool of 3 FTE RSE available to help porting new applications and exploit community codes on heterogeneous systems
  - Cambridge RSE time handled by RAC, coordinated by Cambridge team
- Main training activities
  - Training on GPU (OpenACC, CUDA FORTRAN, CUDA C)
  - Training on Intel Many-Core (in collaboration with DELL and Intel)
  - Training on Parallel I/O (*under development*)
- Happy to open to other Tier-2 RSE groups for testing purposes
  - Rules of engagement yet to define, limited allocations to be reviewed
- Tier-2 “Champion” / Reference RSE contact: Mr. Filippo Spiga

Tier 2 Hub in Materials and Molecular Modelling

- Who?
  - Thomas Young Centre (Imperial, Kings, UCL and Queen Mary), Queen's University of Belfast, Cambridge, Kent, Oxford and Southampton
- What?
  - "Thomas", 720 x86-64 nodes (24 cores, 128GB RAM), OmniPath
- When?
  - Hardware mostly in place, early access in late April, full production in May
- Software
  - Legion/Grace software stack, adding ~7 packages available on ARCHER
- Allocation
  - 20% top slice available for external award, EPSRC defined
- Champion
  - Notionally: Heather Kelly, UCL Research Computing Applications & Support

Tier 2 Hub in Materials and Molecular Modelling

- How?

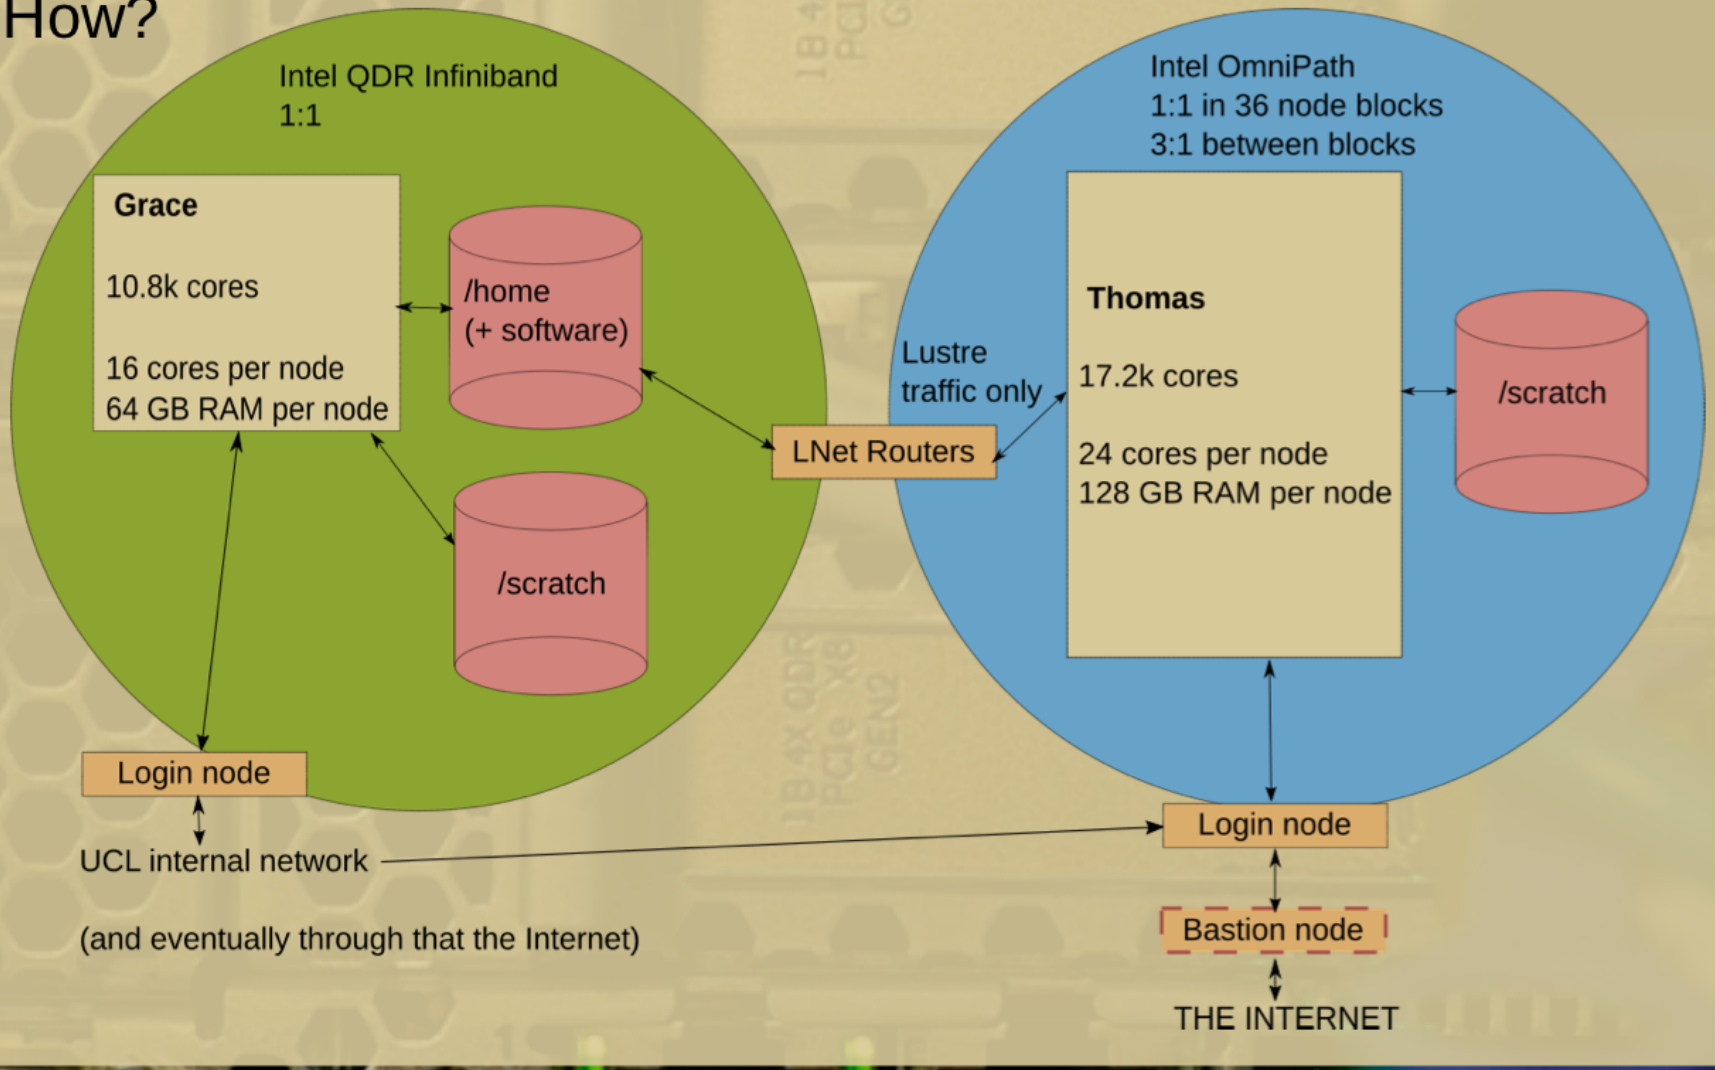

Aston University

UNIVERSITY OF BIRMINGHAM

UNIVERSITY OF LEICESTER

HPC Midlands Plus

Prof. Steven Kenny

Department of Materials

Loughborough University

Loughborough University

The University of Nottingham

UNITED KINGDOM · CHINA · MALAYSIA

Queen Mary University of London

WARWICK

THE UNIVERSITY OF WARWICK

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

Centre Setup

- Centre installation complete by end of March
- Pilot service in the early part of April
- Full service by the end of April
- 15% time on the system released as part of Tier 2 national service
- Research software engineers at all of the sites across the centre
- Training courses across the centre

GW4 Isambard

- Team
  - Universities of Bath, Bristol, Cardiff and Exeter
  - Met Office and Cray
- Hardware
  - Stage 1: Technology comparison: x86, KNL and Pascal
  - Stage 2: Cray CS-400, 10,000+ ARMv8 cores
- Procurement
  - Going well – installing Stage 1 in mid-March
  - Stage 2 will be installed between March-December

GW4 Isambard

- 20—25% will be made available to the community via resource allocation calls
- Plan to also make it available via a continuous integration platform
- 2+ RSEs hired to support code development and porting, plus Cray will provide engineer time
- Training materials will be made available
- Tier-2 Champion is Christopher Woods

Tier-2 Opportunities

- Large number of RSEs being hired to support lots of code porting and optimisation
- Avoid duplication of effort
  - Let's share knowledge and work together
- Lots of training will be provided
  - How will this be advertised?
- Lots of allocation calls and time available
  - How will this be advertised?
- How can a researcher or RSE at University X find out what is available, get training, get support and apply for access?

Tier-2 Champions

- Proposed Tier-2 Champions. They will help share information
- Provide a named person at each Tier-2 centre
- Will ensure information regarding training, support, access and documentation is shared with the wider community
- Will help RSEs at different Tier-2 centres to talk to each other, share knowledge, and avoid duplication of effort

Tier-2 Champions Website

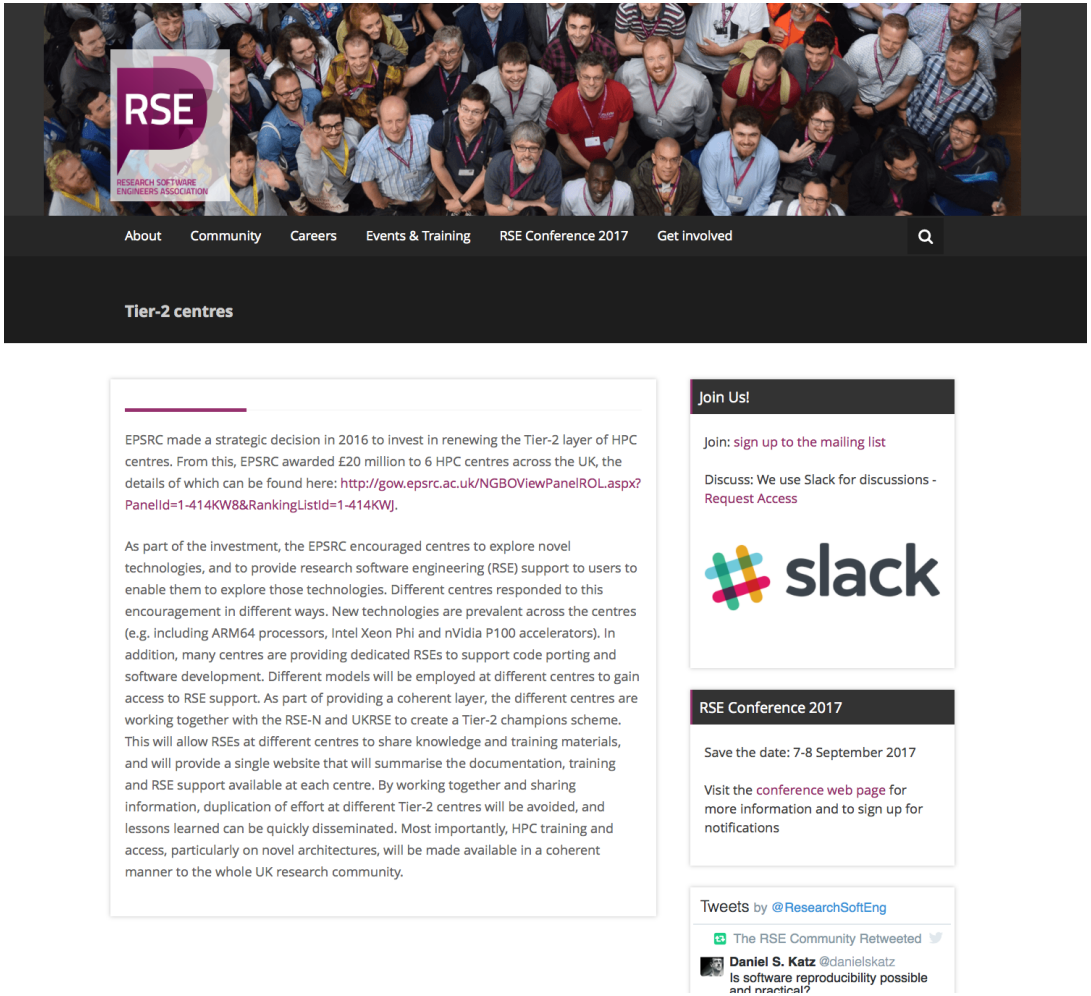

Main portal for sharing information

Hosted on UKRSE website (http://rse.ac.uk)

WordPress so easy to add and edit content

Tier-2 Champions Website

- Contains
  - Short description of what is provided by each centre, with links to each centre's website
  - Links to training materials provided by each centre
  - List of training courses and links to the sign-up pages
  - List of upcoming resource calls and links to the sign-up pages
  - Information on how to get access to time for the purpose of running external training workshops
    - e.g. gaining access to Isambard from University X for running a training course at University X

Tier-2 Champions : Questions

- How can Tier-2 integrate with Tier-1?
- How can Tier-2 champions integrate with Archer champions?
- What else is needed on the Tier-2 website to support external users of each centre?
- Is it ok to start the website hosting under UKRSE?
  - Who else would want to run the website?
- What forums are available to allow Tier-2 RSEs to communicate to share best practice and avoid duplication of effort?