HPC WALES • CYMRU

Powering Success
Pweru Llwyddiant

HPC Wales – An Integrated HPC service

Powering Success through Partnership

Christine Kitchen, David Elcock & Martyn Guest

Ewrop & Chymru:
Buddsoddi yn eich dyfodol
Cronfa Datblygu Rhanbarthol Ewrop

Europe & Wales:
Investing in your future
European Regional Development Fund

ERDF

Llywodraeth Cymru
Welsh Government

Introduction and Overview

- HPC Wales today
  - Current status & the critical role of Partnerships
  - Perspective of Welsh Companies using HPC
  - The Welsh Priority Sector Areas
- A pan-Wales Distributed Computing Network
- Phase 1 and the Fujitsu Westmere Clusters
  - Phase 2 and Sandy Bridge Installations
- Support Services
  - Role of Scientific Gateways – Deep or Wide?
  - Advanced User Support Services

An Integrated HPC service: Powering Success through Partnership

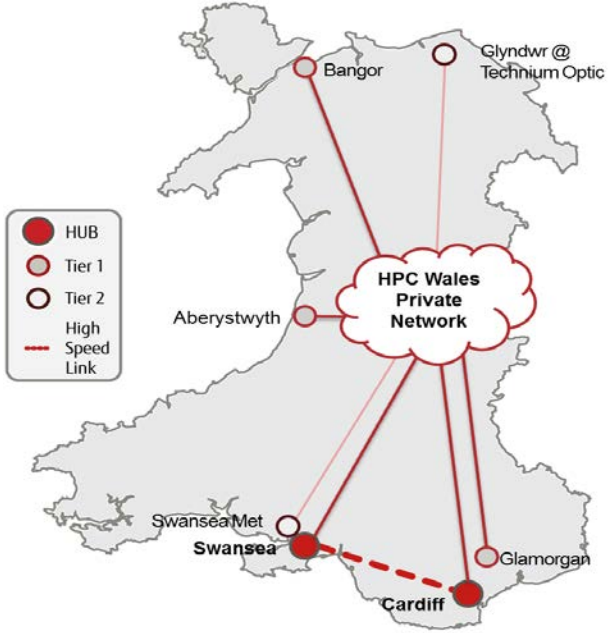

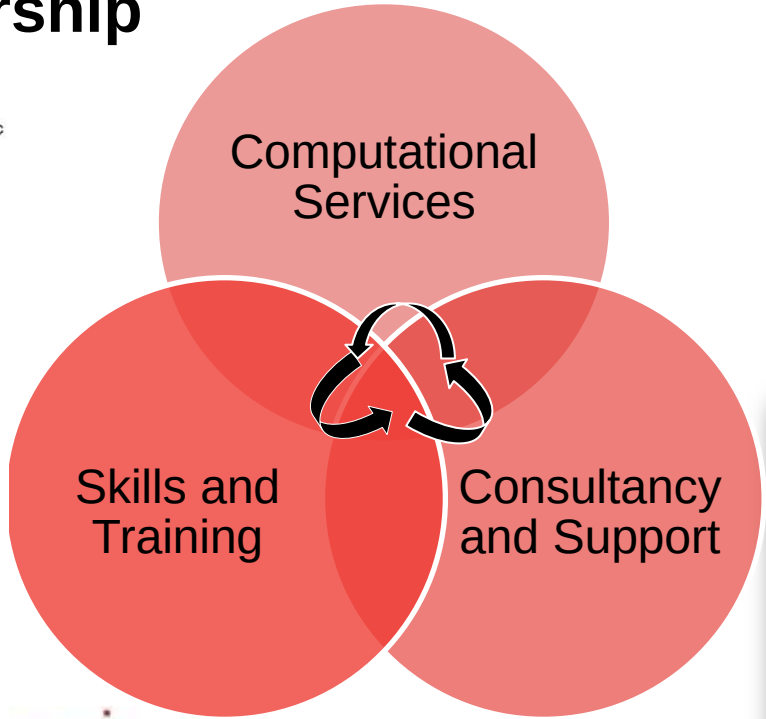

Llywodraeth Cymru
Welsh Government

BIS Department for Business Innovation & Skills

Cyngor Cyllido Addysg Uwch Cymru
Higher Education Funding Council for Wales
hefcw

Ewrop & Chymru: Buddsoddi yn eich dyfodol
Europe & Wales: Investing in your future
ERDF
European Regional Development Fund
Llywodraeth Cymru Welsh Government

Microsoft
nag
FUJITSU
intel
XSEDE
Extreme Science and Engineering Discovery Environment

Benefits and Demand

- Quicker
  - Significantly reduced time to market using our supercomputing facilities
- Quantity / Quality
  - Ability to handle larger data sets or run more realistic simulations and models in the same amount of time, meaning companies can take on more work
- Cost
  - If in the EU convergence zone, then it could be fully funded

(Omnibus Survey of 500 SMEs - 2011)

	Total	< £1.1m T/O
Large Dataset Analysis	21%	40%
Computer Design / Visualisation	18%	41%
Consumer Trends Analysis	16%	39%
Simulation / Modelling	6%	21%
Large Scale Quantitative Research	4%	13%

The Priority Sectors

Energy and Environment

- Designing and producing efficient and renewable energy sources
- Assessing convenient reservoir & well locations
- Researching & modelling climate change & weather patterns
- Exploring alternative waste disposal methods

Live Sciences

- Aiding in the design of new drugs
- Assisting in the development and use of remote surgery techniques
- Allowing advanced brain scanning to monitor/help combat Alzheimers, Autism and Stroke Damage
- Visualising large quantities of multi-omics data to understand the function of complex processes in living organisms

Advanced Materials &Manufacturing

- Improving the efficiency of product development (CAE and PLM)
- Increased speed of computationally demanding tasks such as rendering, analysis and simulations
- Improving testing and development through FEA and CFD simulations

Creative Industries

- Speeding up creative design when used as a render farm
- Production of high-end 2D & 3D visualisations and high-definition animations

Financial and Professional Services

- Speeding up analysis of customer spending patterns
- Running thousands of calculations in seconds to determine risk exposure
- Enabling monitoring and analysis of markets
- Providing predictions on future trends

Welsh companies using HPC

- Zeeko
- Morvus
- Allerna Therapeutics
- Control 2K
- Heritage Medical
- Geolang
- NHS Swansea
- EADS
- Procter & Gamble
- Phytovation Ltd
- Fugro Robertson
- Fulcrum Direct
- Xodus Group Ltd
- CLC Bio
- Corlan Hafren
- Tata Steel
- Tidal Energy
- Johnson Matthey
- Motion Blur
- Nextek Innovations
- Arthur M Ltd
- Gwynedd Archeological Trust
- Geosho
- CYF Acclimatise
- GP Solutions
- Vibetv Ltd
- Alcatel Lucent
- Fibrespeed
- Calon Cardio - Technology Ltd
- Knowtra Ltd

What our customers say

"HPC Wales provided access to the software (ANSYS) and the computing power we needed to refine the design of our implantable pumps, ***speeding up our research process tenfold and enabling us to carry out product trials within a competitive timeframe***."

Dr Graham Foster, Calon Cardio - Technology Ltd.

"HPC is invaluable for us in our production work as it enables us to render high end graphics and animation much more quickly than on regular systems. This means we ***have a quicker turn around, can undertake more projects and increase our turnover***. Supercomputing is definitely the way forward!"

Sue Powell-Reed, VIBETV

"Through collaboration with Bangor University and access to HPC Wales' services, supercomputing provides us with the *potential* ***to develop innovative organic healthcare products***. It is invaluable to us as a small company, helping us to maintain growth in an increasing competitive market."

Lynn Heritage, Heritage Medical

"As an SME delivering oceanographic consultancy to international projects, Knowtra relies on modern computing to undertake simulations and data processing for our clients. HPC Wales has allowed Knowtra *to* ***enhance the quality of simulations***.

Dr Steve Spall, Director, Knowtra Ltd

"In partnership with OSTC, we are using supercomputers to perform macroeconomic modelling to help us understand economic trends, and to inform their trading activities. HPC Wales is vital to the partnership due to the compute-intensive nature of the tests, ***dramatically reducing the time it takes to run thousands of simulations***."

David Meenagh

Distributed HPC infrastructure & Support Services

PHASE 1 COMPLETE	• Cardiff Hub (HTC) plus 3 X Tier-1 sites and the Swansea Analytics System
PHASE 2 – Q4 2012	• Fujitsu Intel Sandy Bridge Technology • Cardiff and Swansea sub-systems (Capability, Capacity & Accelerators)
PHASE 3 – 2014	• Technical specification to be driven by requirements of the User community
LINKED VIA PSBA NETWORK	• Access via internet and web browser based SynfiniWay
LOCAL ACCESS & SUPPORT	• Training, Outreach & Technical Teams

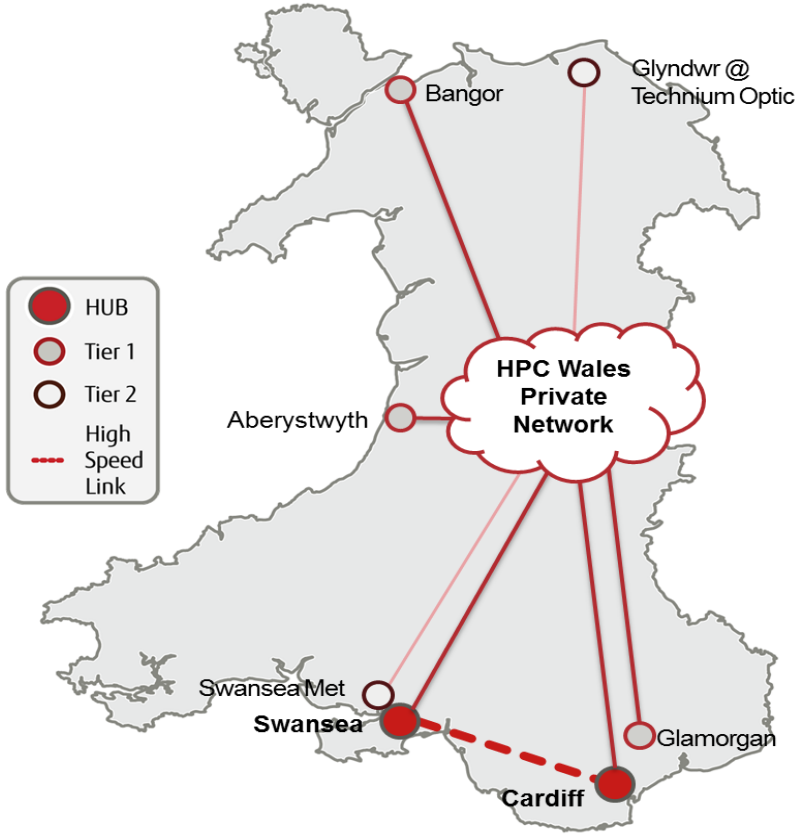

HPC Wales Phase 1 - Westmere Systems

Cardiff HTC System: First HPC Wales system

- 162 x BX922 dual- processor nodes
  - 12 cores & 36GB memory per node
  - 1944 × Intel Xeon Westmere X5650 2.67 GHz cores
- ***Dual-boot Windows HPC nodes***
- Infiniband non-blocking QDR network (1.2 usec latency, 40Gbps B/W)
- 200 TB Lustre high- performance temporary file storage; 100 TB reliable storage for permanent data
- Large memory systems - Intel Xeon Nehalem X7550 2.00 GHz
  - 4 × RX600 dual-processor nodes - 16 cores and 128GB memory each
  - 1 × RX900 octal-processor node - 64 cores and 512GB memory

Tier-1 Systems: Bangor, Glamorgan and Aberystwyth

- Development systems; each includes 54 x BX922 dual-processor nodes
  - 12 cores & 36GB memory per node
  - 648 Westmere X5650 2.67 GHz,
- 11TB storage; Infiniband non-blocking QDR network

Swansea Data Analytics Systems

- IBM Power7 server technology
- IBM DB2 data warehouse

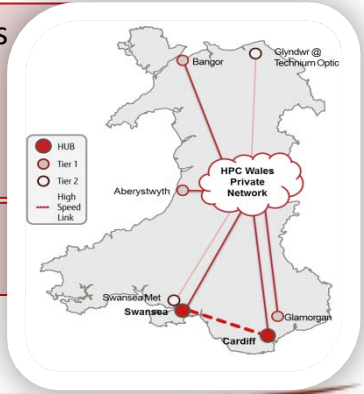

HPC Wales – Top 5 Projects (Sept. 2012)

Computer Simulation of supported nano-particle catalysts for the production of chemical feedstocks from plant waste

- HPCW070/1; Adam Lee, Karen Wilson and David Willock. Cardiff University [**Nextek Innovations Ltd**]

HPC for comparative genomics & genome assembly

- HPCW012, Denis Larkin Jitendra Narayan & Martin Swain, Bangor University [**Welsh Water**]

Development of a Next Generation Sequence HPC-based analysis suite

- HPCW001; Farzana Rahman, Francis Hunt, Negusse Kitaba, Tanya Tatarinova [**National Botanic Garden of Wales**]

Functional impact of mutations and polymorphic variation in whole genome sequences:

- HPCW087; Jonathan Mullins, Swansea University [**Allerna Ltd**.]

HPC for Computational Rheological Modelling

- HPCW035, Hamid Tamaddon & Mike Webster, Swansea University [**Tata Steel (Wales), Wellstream-GE, Procter & Gamble, Haemair (Swansea), NHS Swansea**]

Fujitsu Partnership

20 PhD Studentships

R&I Development Projects

Sandy Bridge at the Hubs

Phase 2 – Q4 2012

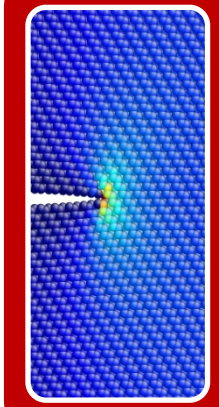

Swansea Capability system:

- 256 x CX250S1 E5-2690 dual professor 16-core nodes
- 240 nodes with 64GB memory; 16 nodes with 128GB memory
- 4096 × 2.9 GHz cores with 4/8 GB memory / core
- Infiniband non-blocking QDR network
- 89 Tflops peak performance
- 400 TB Lustre file storage and 100 TB NFS storage

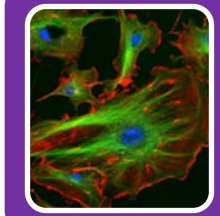

Novel Architecture Sub- systems:

- 16 x CX270 nodes with GP-GPU capability
- 256 x 2.6 GHz cores, 4 GB memory/ core
- each node with a M2090 GP-GPU card
- Infiniband non-blocking QDR network

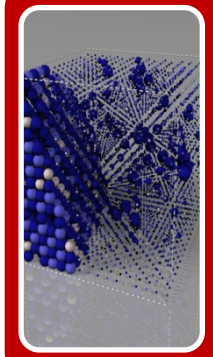

Swansea Capacity system:

- 128 x CX250S1 E5-2670 dual processor nodes; 16 cores & 64GB memory per node
- 2048 × 2.6 GHz cores with 4 GB memory / core
- Infiniband non-blocking QDR network
- 43 Tflops peak performance
- Filestore shared with Capability sub-system

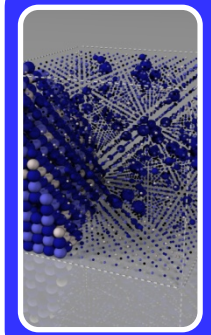

Cardiff Capacity system:

- 384x CX250S1 E5-2670 dual processor nodes; 16 cores & 64GB memory per node
- 6144 × 2.6 GHz cores with 4 GB memory / core
- Infiniband non-blocking QDR network
- 128 TFlops peak performance
- Filestore shared with HTC sub-system

Sandy Bridge Application Benchmarks

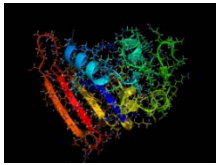

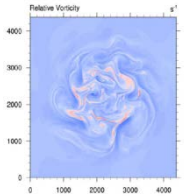

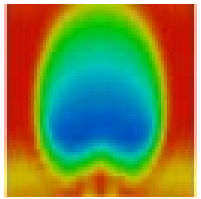

Quantum Espresso	GRIR443	Fujitsu CX250 e5-2690 2.9 GHz QDR	1.85
		Fujitsu CX250 e5-2670 2.6 GHz QDR	1.69
GROMACS	DPPC/pme	Fujitsu CX250 e5-2690 2.9 GHz QDR	1.29
		Fujitsu CX250 e5-2670 2.6 GHz QDR	1.21
WRF	4dbasic	Fujitsu CX250 e5-2690 2.9 GHz QDR	1.43
		Fujitsu CX250 e5-2670 2.6 GHz QDR	1.35
OpenFOAM	cavity3d-3M	Fujitsu CX250 e5-2690 2.9 GHz QDR	1.99
		Fujitsu CX250 e5-2670 2.6 GHz QDR	1.82

Capacity	*1.52*
Capability	*1.64*

0.00 0.50 1.00 1.50 2.00

Relative Performance

A Complete Service

Access Methods

- Transparent job execution & data movement enabled by the SynfiniWay-enabled Gateways
- Job monitoring & data access via web browser

New HPC and Windows Users:

- Web browser based SynfiniWay
- HPC applications provisioned as workflows
- Point and Click interface to specifying inputs and outputs to launch jobs

Experienced HPC and Linux Users:

- SSH based access
- Command-line control and job submission

Filesystems:

- /home – NFS; /scratch – Lustre; /tmp – local disk

Operating System:

- Redhat (Enterprise Linux Server release 5.5)
- Compilers: Intel v12.1, PGI 12.5 & PathScale

Libraries:

- Intel MKL
- ACML,HDF5, NetCDF, FFTW etc.
- MPI – Intel-MPI, OpenMPI & Platform-MPI

Technical Services and Support:

- Complete pipeline making HPC work for users
- Software installation & control
- Job execution, Debugging and Optimisation
- Training and support services

Collaboration with XSEDE

Scientific Gateways

- Training in deploying the HUBzero solution on the HPC Wales resources – provision of technical support and assistance in troubleshooting the deployment of the software.

Skills Academy

- Select the Top 5 courses which would provide the foundations of the knowledge transfer workshop with the key instructors.
- Train identified HPC Wales staff in the delivery of the course material / relevant instructor guides

Campus Champions Programme

- Advice and guidance on establishing an HPC Wales Campus Champion Programme.
- Participation in the XSEDE12 Campus Champions workshop with access to the campus champions training material

User productivity at first login

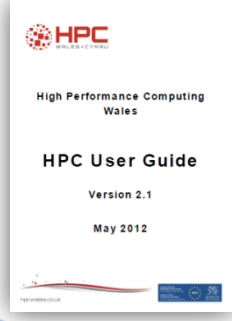

Single user identity

- Remove complexity of using multiple resources
- Uniform connect from any access point

Common reusable services

- Shared distributed repository of application services
- Accessible everywhere across the backbone

Optimised user workflows

- Best-practice processes
- Published to all partners

Development of dedicated portal

- Extend backbone usage to widest community
- Harmonise and simplify HPC access

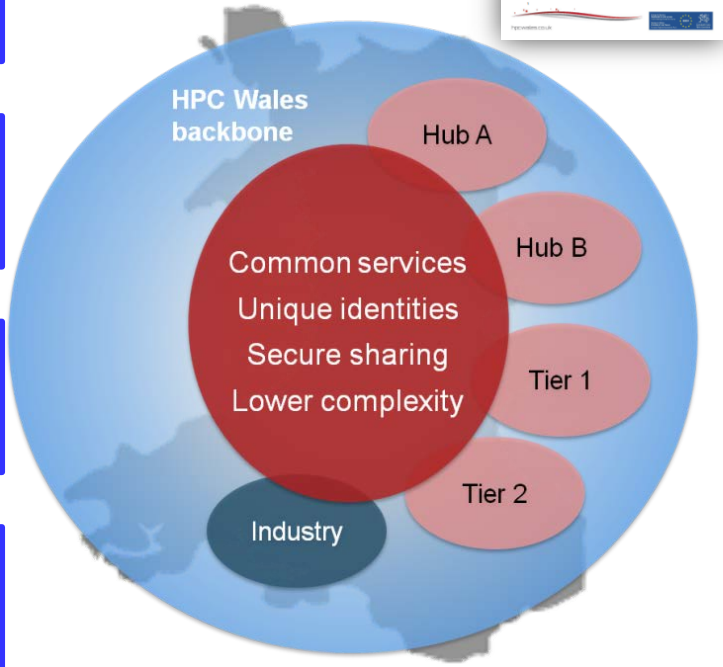

Scientific Gateways – “Wide”

Science gateways are community-designed interfaces (often web-based) that provide a user-friendly interface to many services. Depending on the community, these might include data collections, user work spaces, collaboration spaces, and computation and visualization capabilities.

A “wide” approach is unlikely to be possible based solely on SynfiniWay, but will require Open Source gateway capabilities e.g. HubZero, CIPRES

Welsh Government Priority Sectors

- advanced materials and manufacturing
- creative industries
- energy and environment
- financial and professional services
- information & communication technologies
- life sciences
- food and farming
- construction
- tourism

Advanced User Support Services

Comp. Phys. Commun. 183 (2012) 520–529

Develop a general approach to understanding scaling behaviour

- breakdown of scaling with increasing core counts in terms of code components
- Aim to develop a methodology applicable for any code on any system

Evaluation of open source & proprietary performance analysis tools

- Open Source: Open|speedshop, TAU with Paraprof & JumpShot, IPM (Integrated Performance Monitoring) & Ploticus, HPCToolKit
- Proprietary: Allinea OPT, ITAC (Intel Trace Analyzer and Collector), CrayPat

Findings: CrayPat & TAU identified as most appropriate tools

Scaling analysis procedure using both tools developed

Application Profiling & Scaling Analysis

Summary

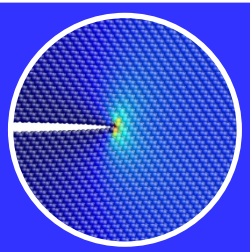

Speed of innovation is critical to business growth and profitability – many SMEs don't understand the impact high performance computing can have

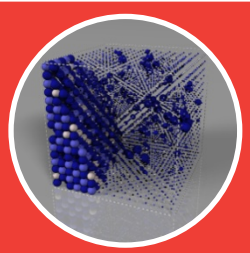

Research and innovation needs long term investment in infrastructure, software, and people to grow the economy

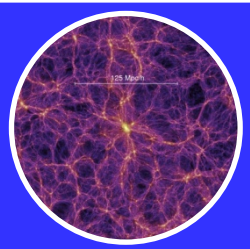

The best solutions are complete: researchers and developers need access to hardware, software, support and training