



The Scientific Computing Department Today

A high level overview of SCD. **David Corney**

SCARF a local HPC cluster supporting STFC scientists, collaborators and users from SCD, ISIS, CLF, RAL Space and Diamond. **Derek Ross**

Data analysis and data interpretation through modelling and simulation.
Barbara Montanari

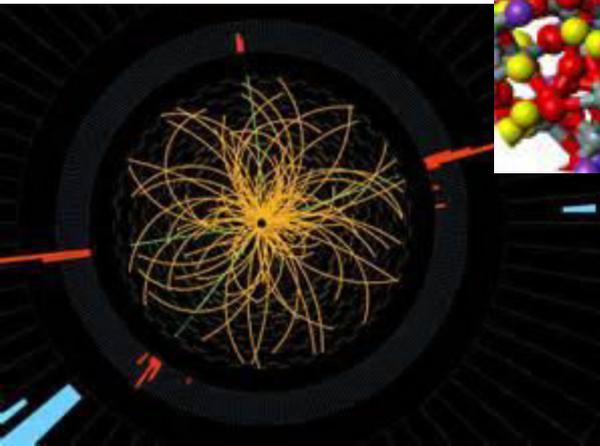
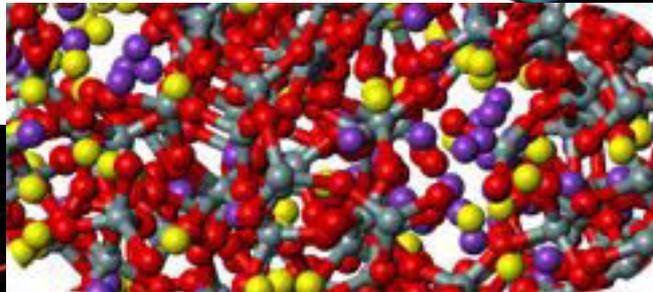
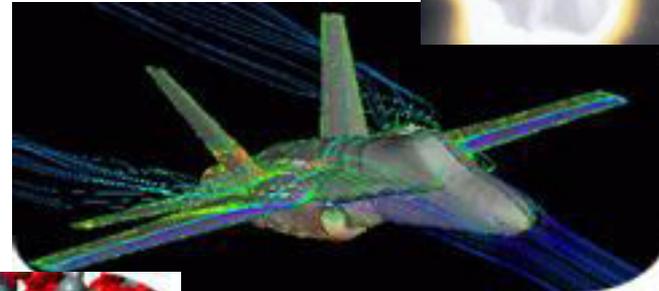
ATLAS@50

13 Nov 2014



Established 1st April 2014

- 180 staff supporting over 7500 users
- Applications development and support
- Compute and data facilities and services
- Research: over 100 publications pa
- Deliver over 3500 training days pa
- Systems administration, data services, high-performance computing, numerical analysis & software engineering.



**World leading
computational
expertise to support
Science and UK
Industry**

Scientific Computing Department

Acting Director



David Corney

APPLICATIONS

Division Head



Paul Sherwood

Group Leader

Martyn Winn
Biology and Life Sciences

Dave Emerson
Engineering and Environment

Barbara Montanari
Theoretical and Computational Physics

TBC
Computational Chemistry

CORE ACTIVITIES

Group Leader

Shirley Miller
IT Administration

RAL Administration

Finance

CECAM

DATA

Division Head



Juan Bicarregui

Group Leader

Jens Jensen
Data Services

Brian Matthews
Research Data

Chadwick and RAL Libraries
Group Leader

Debbie Franks
Daresbury Laboratory

Rutherford Appleton Laboratory

SYSTEMS

Acting Division Head



Andrew Sansum

Group Leader

Dave Cable
High Performance Systems

Nick Hill
Research Infrastructure

Andrew Sansum
Peta Scale Computing and Storage

TECHNOLOGY

Division Head



Peter Oliver

Group Leader

Jennifer Scott
Numerical Analysis

Chris Greenough
Software Engineering

Mike Ashworth
Application Performance Engineering

Martin Turner
Visualisation

HARTREE CENTRE

Division Head

Cliff Breerton
Hartree Centre

GRADUATES

HONORARY SCIENTISTS

VISITING SCIENTISTS

VISITORS

NSCCS

NSCCS (National Service Computational Chemistry Software)

Providing National and International Compute, Training and support

EPSRC Mid-Range Service

SGI Altix UV SMP system, 512 CPUs, 4TB shared memory

Supports over 100 active users

~70 peer reviewed papers per year

Over 40 applications installed

Portal to submit jobs

to allow access to less computationally aware chemists

EPSRC

Pioneering research
and skills

**Imperial College
London**

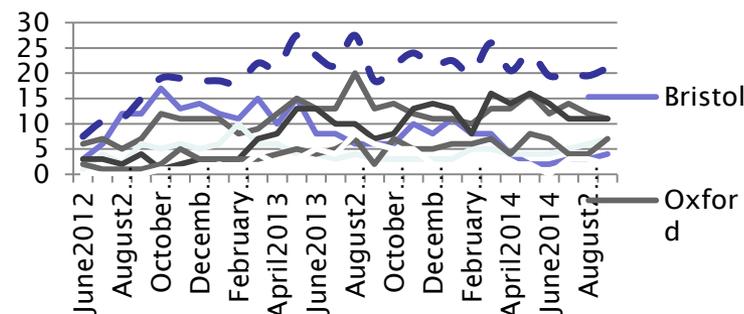


372 Nvidia M2090 GPU cluster purchased in 2012
Run for a consortium of universities – Bristol, Oxford, Southampton and UCL and a small number of external users from academia and industry

Wide range of scientific usage from drug discovery to climate modelling:

<http://www.cfi.ses.ac.uk/show-cases/>

Two functions – large scale GPU resource and also supporting porting codes to GPU

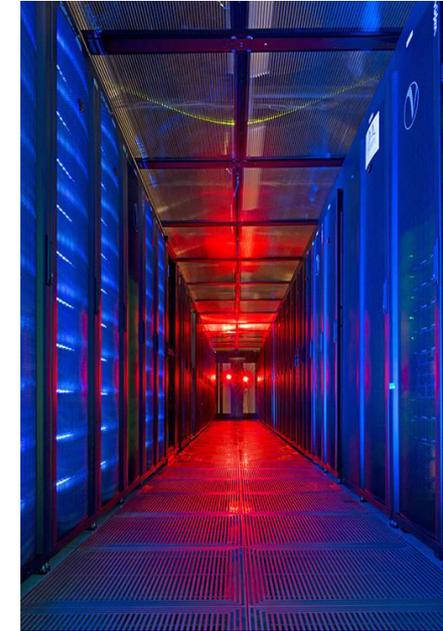


- Data Archive for BBSRC
- Data Support Service for MRC



The JASMIN “super-data-cluster”

- Bringing the compute to the data for all NERC sciences.
 - Data : Measured and Simulation
- Facilitating further comparison and eval. of models with data
- Initially UK and Worldwide climate and weather modelling community.
 - CEMS (Climate and Environmental Monitoring from Space) with UKSA and Industry
 - CMIP5 / IPCC (Climate Change) Data Analysis (~1PB)
 - MetOffice, ECMWF etc
- Now all of NERC environmental sciences since JASMIN2
 - Eg Genomics, Hydrology, Oceanography, Oil & Gas...
- Supports 10,000 UK and World users via ~2PB CEDA archive (Centre for Environmental Data Archive).
 - http,ftp,scp, etc + helpdesk services
- Supports ~500 UK and World direct login users
- Supporting ‘long tail’ users via Cloud SaaS and PaaS Cloud technologies. Eg “Environmental Workbench”.
- High speed dedicated network links to :
MetOffice, ARCHER, Leeds Uni., KNMI Holland, ESGF (Australia and USA)



**National Centre for
Atmospheric Science**

NATURAL ENVIRONMENT RESEARCH COUNCIL



**National Centre for
Earth Observation**

NATURAL ENVIRONMENT RESEARCH COUNCIL

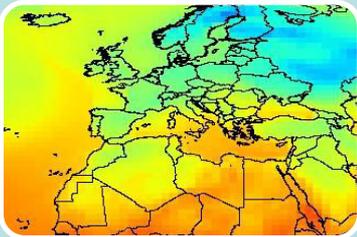


JASMIN's Purpose



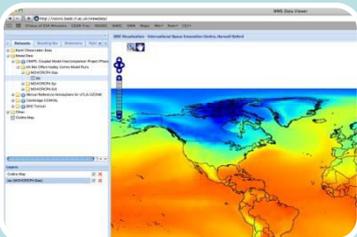
CEDA data storage & services

- Curated data archive
- Archive management services
- Archive access services (HTTP, FTP, Helpdesk, ...)



Data intensive scientific computing

- Global / regional datasets & models
- High spatial, temporal resolution
- Private cloud

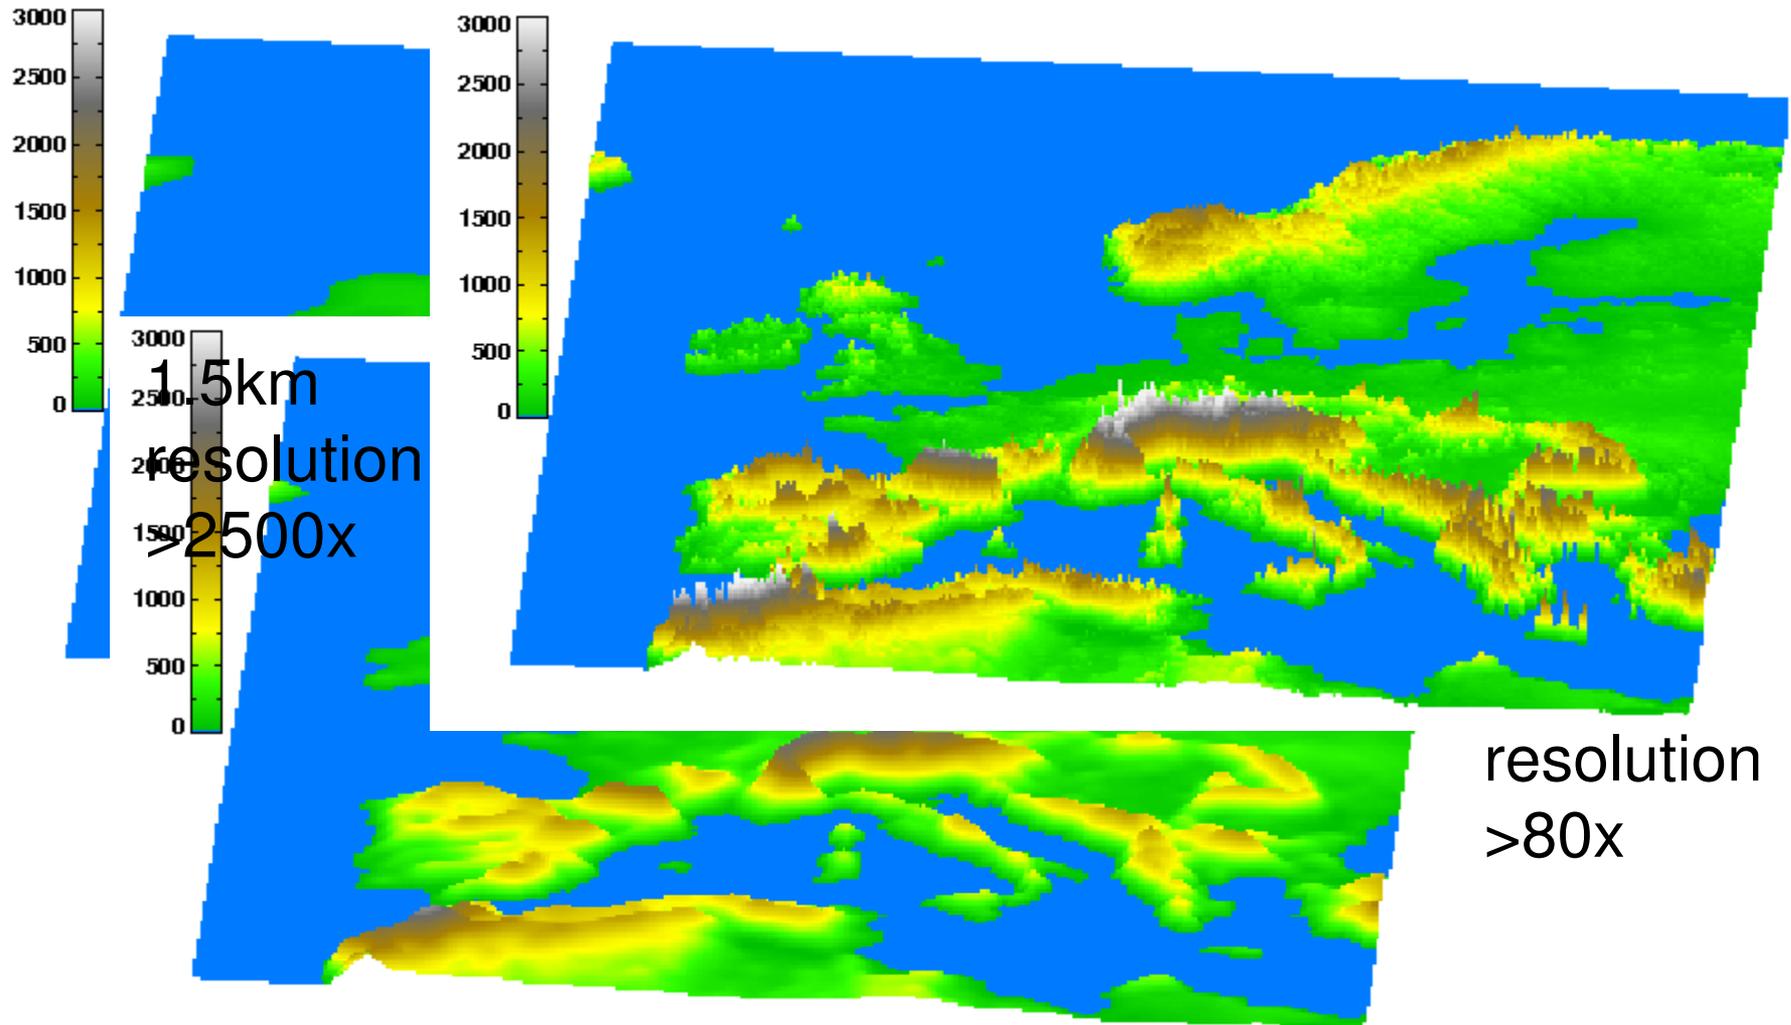


Flexible access to high-volume & complex data for climate & earth observation communities

- Online workspaces
- Services for sharing & collaboration

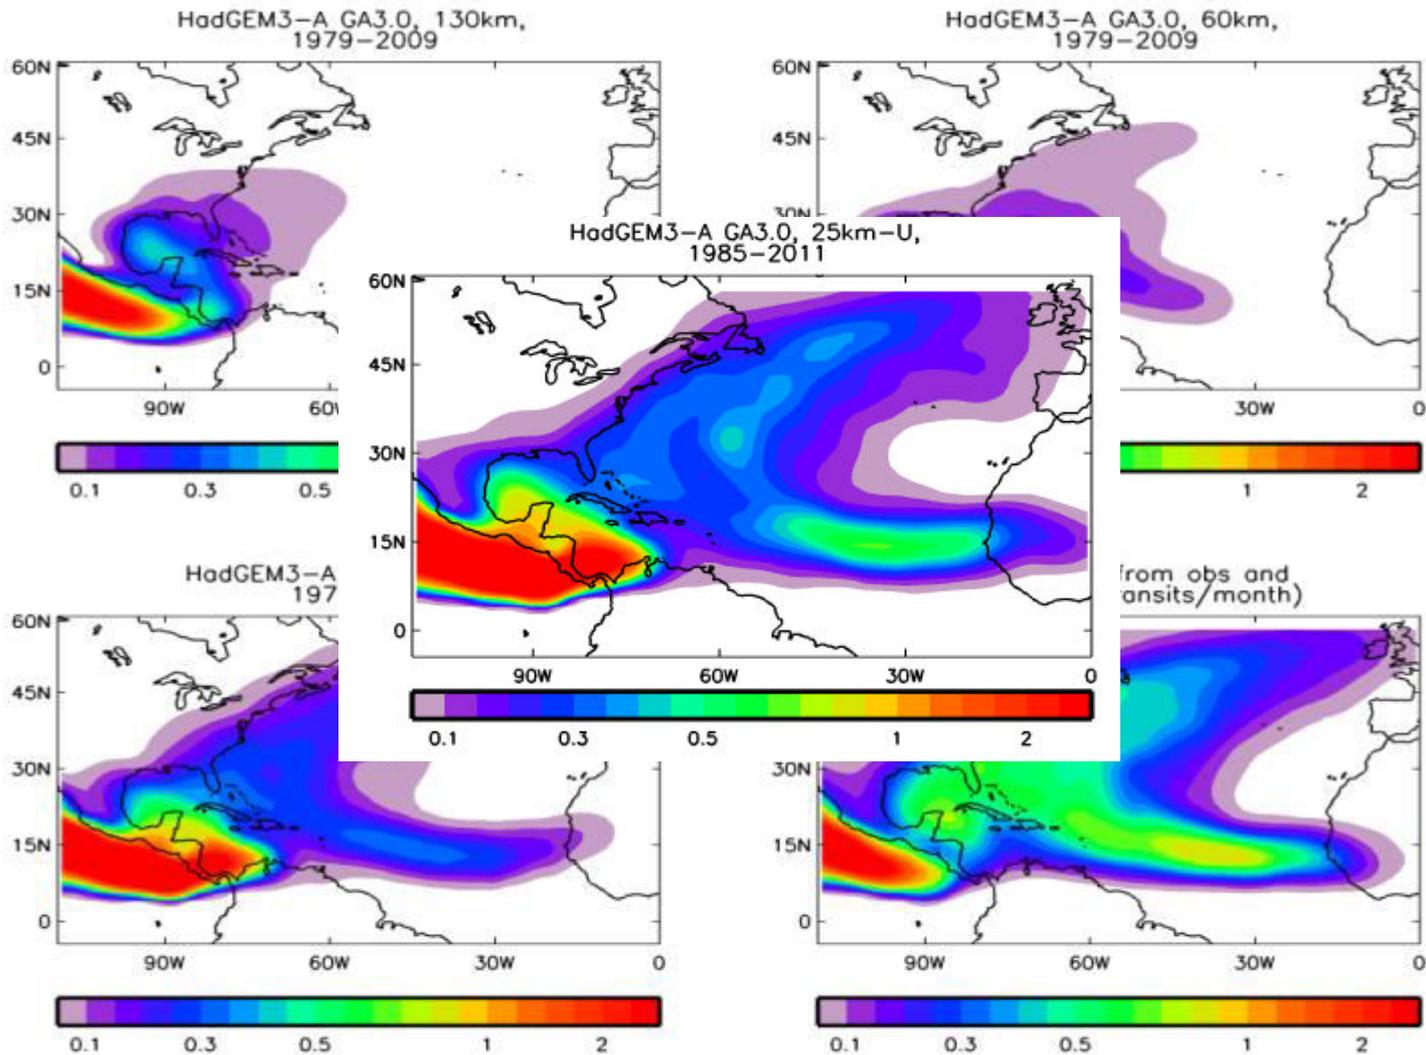
Value of model resolution – orographic height

*Slide by permission of
Pier Luigi Vidale, NCAS (UoR)
Malcolm Roberts and
Matthew Mizieliński, UK Met Office*



Tropical cyclone track density (transits per month)

Slide by permission of
Pier Luigi Vidale, NCAS (UoR)
Malcolm Roberts and
Matthew Mizielinski, UK Met Office





•Technology:

•12PB Parallel Storage Panasas at STFC (Largest in the world)

- Fast Parallel IO to Compute servers
- Largest capacity Panasas installation in the world
- Arguably one of top ten IO systems in the world (~250GByte/sec)
- “Game Changing” 100x throughput capabilities
- Plans to expand to 30PB by 2020



•7PB On Tape (StorageD, CASToR)

- Growing to >30–80PB by 2020 (Demand for 300PB by 2020)

•Virtualised and Physical Compute (~3,5000 cores)

- Physical Batch HPC compute: “LOTUS”
 - Mostly HTC. Some but not heavy parallel HPC/MPI.
- User + Admin provisioned Cloud of virtual machines
 - Currently VMware/vCloud. Likely OpenStack at next renewal



•Complex low latency , zero contention L3 ECMP Networking

- >1200 x 10GbE ports. Few at this performance/scale.

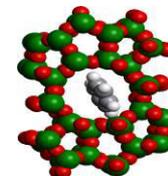
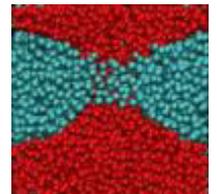
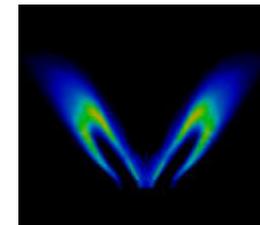
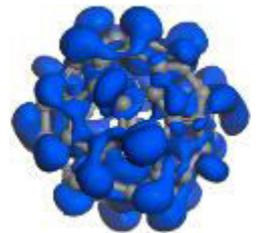
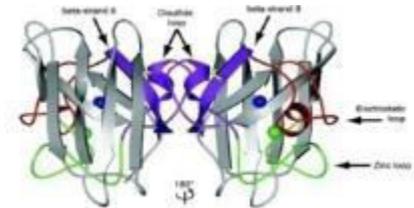
•Multiple Data xfer private network links to UK and World sites





Applications Division

- Four groups developing and applying computational science software packages to solve problems in the physical and biological sciences.
 - **Computational Biology (Martyn Winn)**
 - including structural biology, molecular simulation and bioinformatics
 - **Theoretical and Computational Physics (Barbara Montanari),**
 - electronic structure of the solid state and surfaces, atomic and molecular physics
 - **Computational Engineering (David Emerson)**
 - HPC solutions in fluid flow modelling, with particular strength in turbulence and microfluidics
 - **Computational Chemistry (Ilian Todorov)**
 - molecular dynamics, quantum chemistry and QM/MM techniques, and mesoscale methods

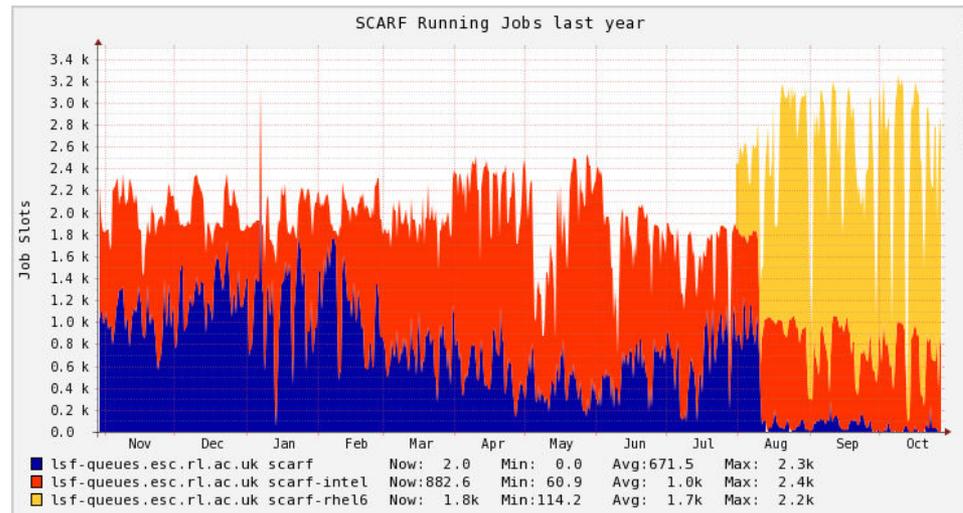


Collaborative Computational Projects

CCP4	Prof David Brown	Macromolecular Crystallography
CCP5	Prof Stephen Parker	The Computer Simulation of Condensed Phases
CCP9	Prof Mike Payne	Computational Electronic Structure of Condensed Matter
CCP12	Prof Stewart Cant	High Performance Computing in Engineering
CCP-ASEArch	Prof Mike Giles	Algorithms and Software for Emerging Architectures
CCP-BioSim	Prof Adrian Mulholland	Biomolecular Simulation at the Life Sciences Interface
CCP-EM	Dr Martyn Winn	Electron Cryo-Microscopy
CCPi	Prof Phillip Withers	Tomographic Imaging
CCPN	Prof Geerten Vuister	NMR
CCP-NC	Dr Jonathan Yates	NMR Crystallography
CCPP	Dr Tony Arber	Computational Plasma Physics
CCPQ *	Prof Tania Monteiro	Quantum Dynamics in Atomic, Molecular and Optical Physics
CCP-SAS	Prof Steve Perkins	Analysis of Structural Data in Chemical Biology and Soft Condensed Matter
CCPForge	Prof Chris Greenough	Collaborative Software Development Environment Tool

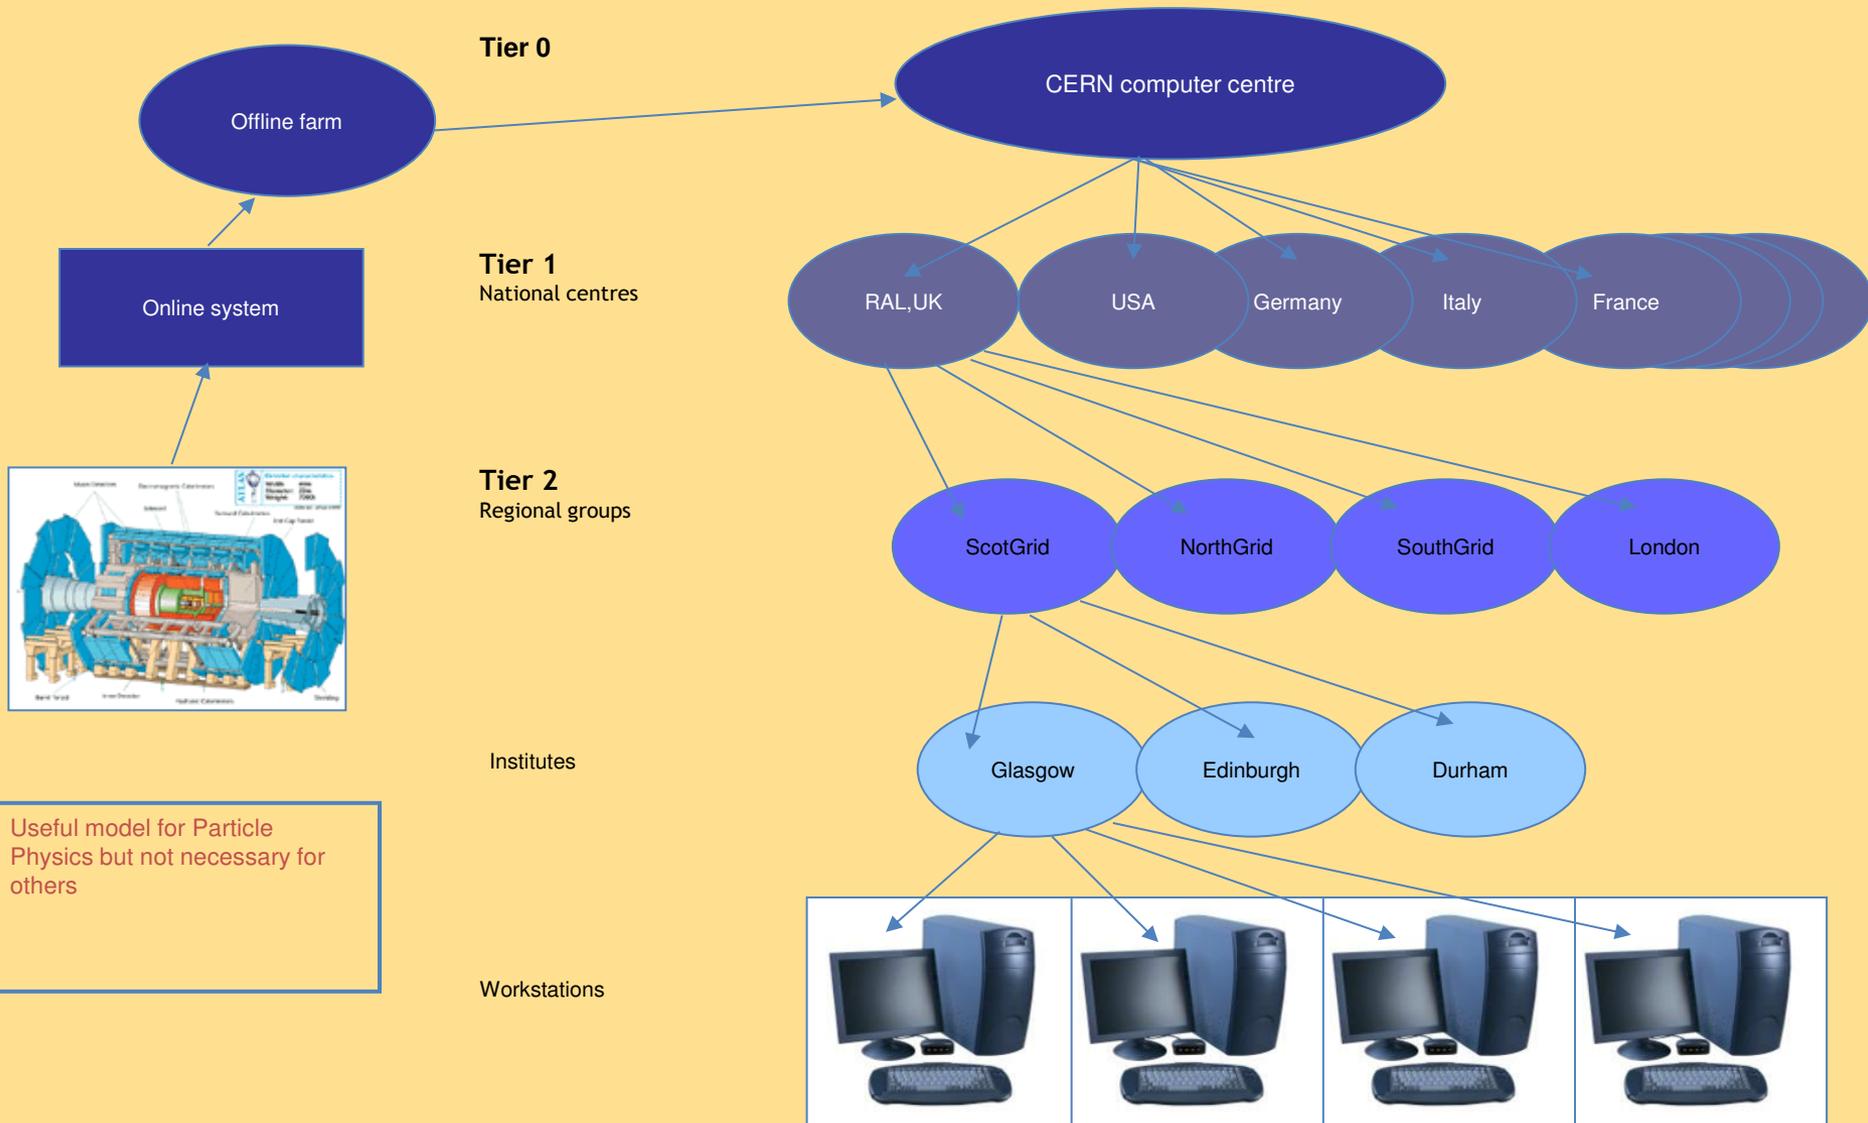


- STFC local HPC cluster supporting STFC scientists, collaborators and facility users
 - Users from SC, ISIS, CLF and others including RAL Space and Diamond
- 4400 cores, new equipment purchased yearly, decommissioned after 5-6 years
 - 3500 generally available, remainder are for particular user communities (CLF Plasma Physics, ISIS IBIS)
 - Recent years have seen additional capital investment from ISIS and CLF, each typically matching SC input
- 350 registered users

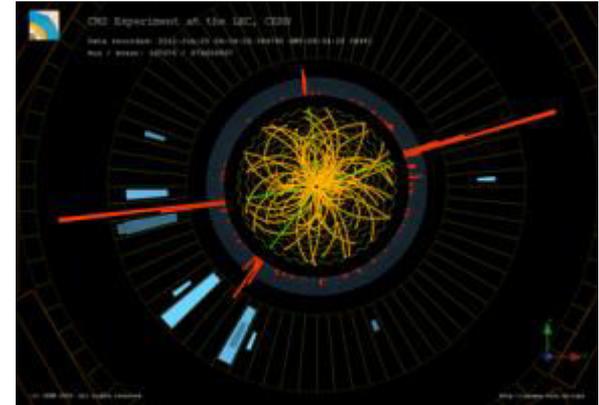


The LHC Tier 1 @ RAL

Context: WLCG architecture



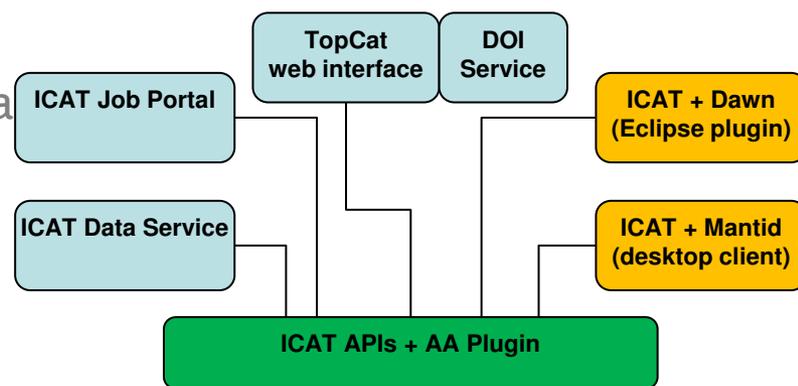
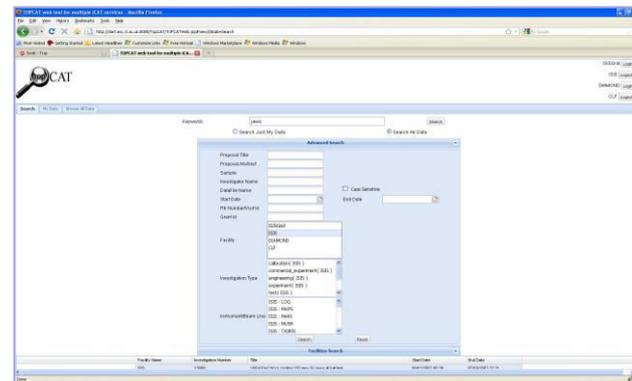
- High capacity service for hunting Higgs Particle
- Remote job submission (Grid)
- Low cost commodity hardware
- Open source middleware
- Delivered 99.2% availability in 2013 (70 hours lost all causes)
- Moves 50 petabytes/year globally (wide area transfers)
- Moves 300 petabytes/year internally (data processing)
- Preparing for Run 2 (data rates and volumes double)



- 13.6 PB disk
- 16 PB tape
- 10,000 CPU cores
- 2000 servers
- 40Gb network
- 10Gb/s direct optical link to CERN

- Tier-1 Fabric team funded by facilities program to support some facilities hardware using Tier-1 infrastructure management:
 - Tier-1 virtualisation framework (HyperV)
 - Tier-1 configuration management system (Quattor)
 - Tier-1 Production Team part funded to monitor and action facilities service exceptions
 - Also use Tier-1's change management process
 - Facilities share funding with Tier-1 for CASTOR
-
- Tier-1 developing cloud access to CPU farm and large Object store (CEPH) to generalise service, giving opportunistic access to non Grid user communities (such as facilities.)

- Integrated data management pipelines for data handling
 - From data acquisition to storage
- A Catalogue of Experimental Data
 - ICAT Tool Suite: *Metadata as Middleware*
 - Automated metadata capture
 - Integrated with the User Office system
- Providing access to the user
 - TopCat web front end
 - Integrated into Analysis frameworks
 - Mantid for Neutrons, DAWN for X-Ray
- In daily production use :
 - CLF, ISIS, DLS
- Also internationally:
 - ESRF, ILL, SNS, ...
 - PaNData Consortium

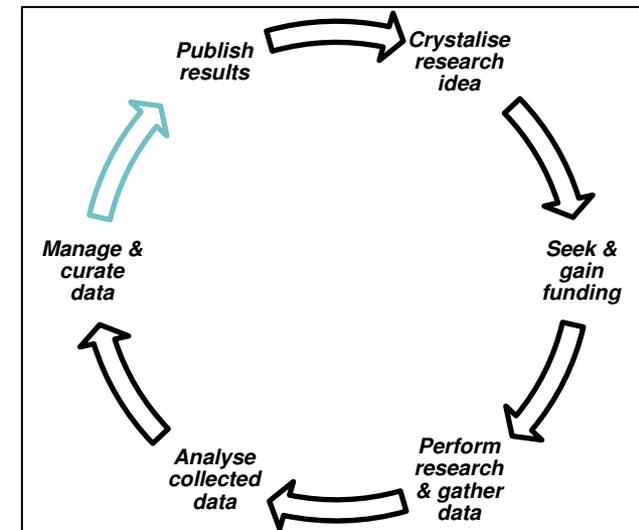


Growing opportunities to build on the data management infrastructure

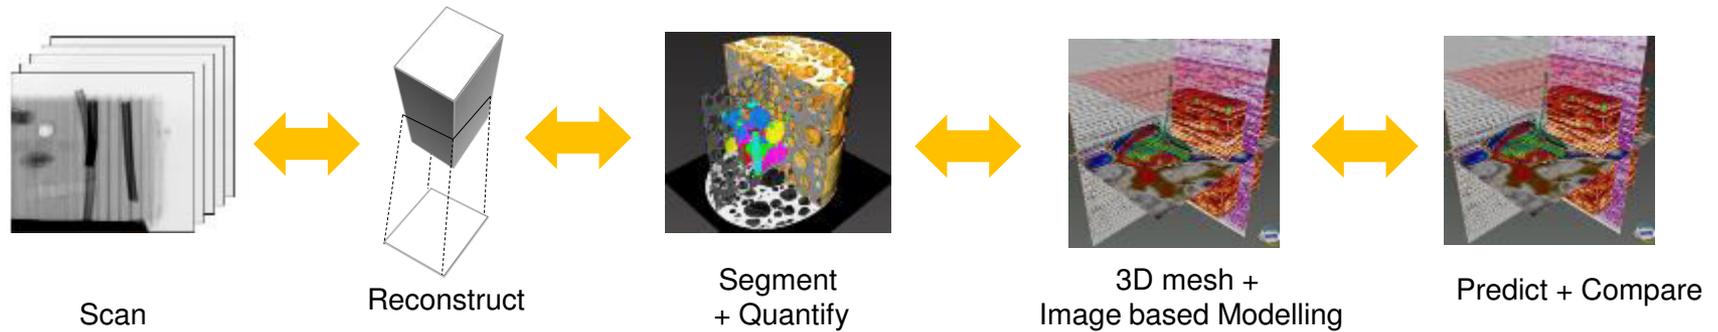
Extending the integrated data management system across the scientific lifecycle;

From Proposal to Publication

- Support post-experimental analysis
 - Image reconstruction with ISIS (IMAT) and DLS:
 - data from experiment, user access to SCARF
 - ICAT Job Portal for the LSF Octopus Facility
 - Integrated job submission
- Support Data Publication
 - DOIs issued for Data
 - Data Preservation
- Support Publications
 - STFC Epubs Repository
 - Linking data and publications to provide a record of science

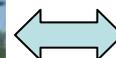
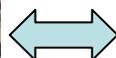


Prototyping Post-experimental support



Infrastructure + Software + Expertise!

ISIS:IMAT



DLS:I12/I13

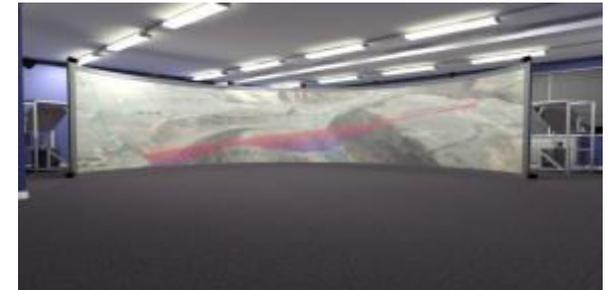


Tomography:

Dealing with high data volumes – 200Gb/scan, ~5 TB/day (one experiment at DLS)



Facilities and Resources of The Hartree Centre



Projects and codes developed on state of the art systems:
BlueGene/Q – (Was) Fastest UK machine and world's largest software
development platform
Over 5 PB disc and 15 PB tape stores
iDataplex cluster
Data Intensive systems
Visualisation System
Energy Efficient Computing program



Client Projects include

- Engineering & Manufacturing
 - Vehicle Design & Testing
 - Consumer Electronics Design
 - Consumer Packaged Goods Design
- Environment
 - Weather modelling
- Life Sciences
 - Genomics for better crop yields
- Energy
 - Advanced Battery Cell Design
 - Efficient Well Head Oil extraction
- Financial Services
 - Risk Management
 - Service Modelling

Accelerating the product discovery process at Unilever

Accelerating product design and development at Jaguar Land Rover

Smaller, affordable particle accelerators for healthcare and security

Supporting new product design at Bentley

Researchers at Lancaster University are using the supercomputing capability of the Hartree Centre to accurately simulate the flow of complex fluids to improve oil extraction techniques.

Work with us

The Science & Technology Facilities Council (STFC) is a leading partner in world class high performance computing (HPC) and supercomputing (SC) facilities. We offer a range of services to support our clients in their research and development. Our facilities are located at the Science and Technology Facilities Council (STFC) sites across the UK. We offer a range of services to support our clients in their research and development. Our facilities are located at the Science and Technology Facilities Council (STFC) sites across the UK.



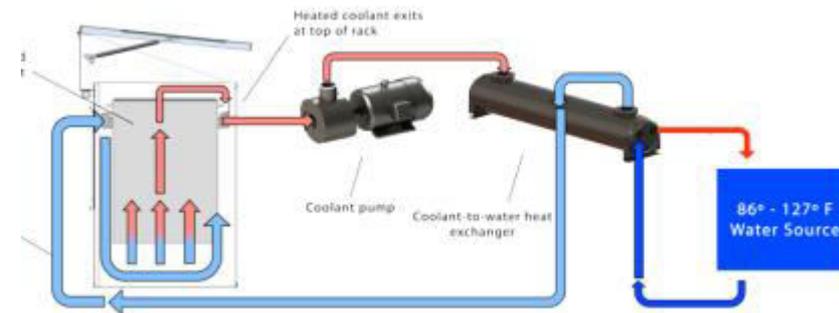
- SCARF a local HPC cluster supporting STFC scientists, collaborators and users from SCD, ISIS, CLF, RAL Space and Diamond. (**Derek Ross**)
- Data analysis and data interpretation through modelling and simulation. (**Barbara Montanari**)





Energy Efficient Computing

- IBM
 - NeXtScale, Idataplex, Intel Phi (£5M)
 - 10,000 cores, 42 intel Phi
 - 3.2PB file store
 - BG Active Storage (£5M)
 - Low Power Processors (£0.7M)
 - NeXtScale + ARM CPUs
 - Big Data & Data Analytics (£4M)
- Insight/Clustervision
 - Novel Cooling Technology (£1M)
- Viglen/Maxeler
 - Dataflow Technology (£1M)





Energy Efficient Hardware: Latest semiconductor technology, Energy saving processor & memory technologies, Special hardware or accelerators.

Energy Aware Management Software: Monitor the energy consumption of the compute system and the building infrastructure, Use energy aware system software to exploit the energy saving features of the platform.

Energy Efficient Infrastructure: Reduce power loss in the power supply chain, Improve cooling technology, Reuse waste heat from systems.

Energy Aware Management Software: Develop efficient algorithms and tools, Optimise libraries, Use most efficient programming paradigm.



HC Clients include



Unilever



acal
energy

Clean affordable power



LLOYD'S



SunChemical

a member of the DIC group



BENTLEY



BAE SYSTEMS



JPMorgan



dyson



SKA
SQUARE KILOMETRE ARRAY



Met Office



OCF



Infineum



IBM



syngenta



DragonHPC



RENUDA



JAGUAR



TECH-X
SIMULATIONS EMPOWERING



ROLLS
ROYCE



GENOYS



LAND-ROVER



ucb Pharma

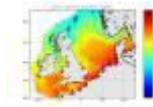
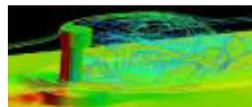
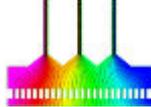
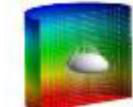
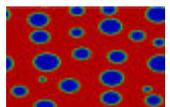
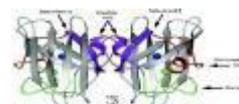
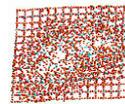
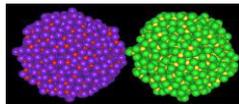
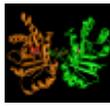
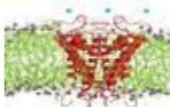
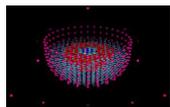


OPTIS

Scientific Computing Department

Major funded activities

- 180 staff supporting over 7500 users
- Applications development and support
- Compute and data facilities and services
- Research: over 100 publications pa
- Deliver over 3500 training days pa
- Systems administration, data services, high-performance computing, numerical analysis & software engineering.
- **Expertise across the length scales from processes occurring inside atoms to environmental modelling**



The LHC Tier 1 @ RAL



↔ To/from CERN:
Up to 40Gb/s

↔ To/from SuperJanet5:
Up to 40Gb/s

- 10,000 batch cores
- ~15PB disk storage
- ~16 PB tape storage
- ~450kW power

Higgs discovery

- 11 PB received by RAL
- 21 PB sent by RAL
- 60 PB processed by RAL



HIGGS EVENT
Nov. 24th 2010 2pm
Anatomy Theatre
Strand Campus
King's College London

My life as a boson - Peter Higgs
(U. of Edinburgh)

How to look for Higgs - John Ellis
(King's College & CERN)

IC Higgs - Tejinder Virdee
(Imperial College)

UC Higgs - Jon Butterworth
(UCL)

If not Higgs, what? - Christophe Grojean
(Saclay & CERN)

Imperial College London KING'S COLLEGE LONDON UCL

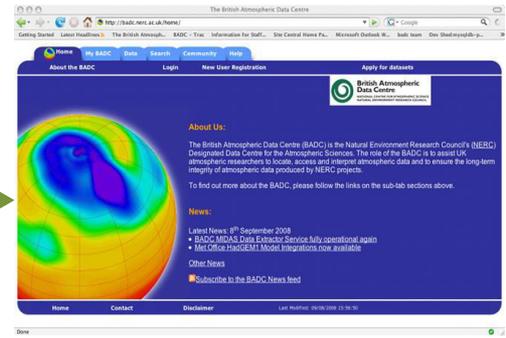
For more information, please see www.kcl.ac.uk/physics
All participants are asked to pre-register to paul.d.le_long@kcl.ac.uk

Common Components Architecture for Facilities Data Service



Synchrotron
experiment
data

CMIP5
data from
CEDA



Storage-D
File
(de-) aggregator

CASTOR
HSM

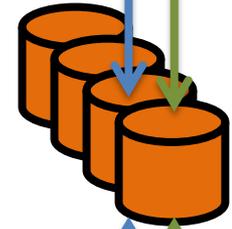
XML
metadata
ingest

ICAT
Metadata
Catalogue

TopCAT web
browser/FUSE

CEDA: 2.1 PB in 120
 $\times 10^6$ files at
September 2014

DLS: >2 PB in
 548×10^6 files at
September 2014



Select data
for download

Catalogue
query &
selection

Selected
experiment
data

Catalogue
query &
selection

Storage-D
File
(de-) aggregator

CASTOR
HSM

Synchrotron
experiment
data

CMIP5
data from
CEDA

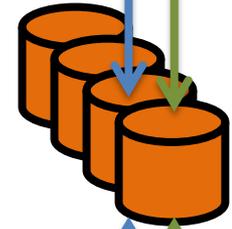
XML
metadata
ingest

ICAT
Metadata
Catalogue

TopCAT web
browser/FUSE

CEDA: 2.1 PB in 120
 $\times 10^6$ files at
September 2014

DLS: >2 PB in
 548×10^6 files at
September 2014



Select data
for download

Catalogue
query &
selection

Selected
experiment
data

Catalogue
query &
selection

Storage-D
File
(de-) aggregator

CASTOR
HSM

Scientific Computing Department

David Corney
Acting Director

Andrew Sansum
Acting Head of
Systems Division

Juan Bicarregui
Head of Data
Services Division

Cliff Brereton
Director of the
Hartree Centre

Peter Oliver
Head of
Technology
Division

Paul Sherwood
Head of
Applications
Division



UK(1 of 2)

Imperial College
London



UNIVERSITY OF
BIRMINGHAM





Science & Technology
Facilities Council

THE UNIVERSITY OF
WARWICK



University
of
St Andrews



THE UNIVERSITY of EDINBURGH

One of the world's top 20 universities



University of Dundee

University of
HUDDERSFIELD
Inspiring tomorrow's professionals

University of Chester

US

University of Sussex



The University of
Nottingham

UNITED KINGDOM · CHINA · MALAYSIA



University
of Glasgow

UNIVERSITY of York



Brunel
University
London



UNIVERSITY OF
LIVERPOOL



UNIVERSITY OF
WESTMINSTER



UNIVERSITY OF
OXFORD

University of
Strathclyde
Glasgow



The
University
Of
Sheffield.



PRIFYSGOL
BANGOR
UNIVERSITY



University of
Leicester

UEA

University of East Anglia



- We have completed 131 access projects for use of our machines
- Including use by the team of Nobel Laureate (Physics) - Prof. Konstantin Novoselov
- These projects have delivered free of charge 379 million Core / Hours
- The total value would have been £4.2M had they had to purchase these commercially
- Many attendees at Hartree Training Courses and Summer Schools
- £40K funding for PhD at Oxford/Culham

BLUE WONDER



IDATAPLEX
3,648 CORES
7.2 TB RAM
48 GPUS



IDATAPLEX
SCALEMP
8 X 4TB SYSTEMS
8X 512 CORES



NEXTSCALE
8640 CORES



IDATAPLEX
2016 CORES
16 X NVIDIA K20
42 X INTEL PHI



BLUE JOULE



IBM BG/Q
6 RACKS, 1.23PFLOP/s
98,304 CORES
9.8 PB RAM



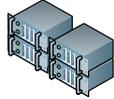
IBM BG/Q BGAS
2 RACK, 450TFLOP/s
32,768 CORES
3.2 PB RAM
256TB FLASH MEMORY



FILESTORE
GPFS
9 PB



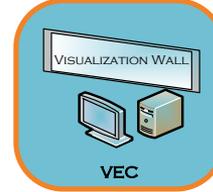
TAPE STORE
15PB



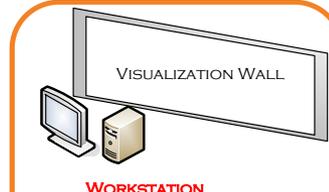
REMOTE GRAPHICS
SERVER



BIG INSIGHTS (912 CORES, 2.6TB RAM)
STREAMS (128 CORES, 1TB RAM)
INFOSPHERE DATA EXPLORER (96 CORES, 192GB RAM)
SPSS (16 CORES, 128GB RAM)
INFOSPHERE CONTENT ANALYTICS (8 CORES, 32GB RAM)
COGNOS (24 CORES, 192 GB RAM)



VEC

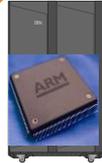


WORKSTATION

MERRISON LECTURE THEATRE



NOVEL COOLING
+ CLUSTERVISION/GRC
1920 CORES, 7.6 TB RAM



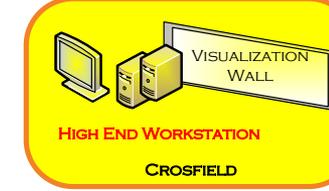
NEXTSCALE
ARM 64 BIT



DATAFLOW
MAXELER
12x CORES, 128 GB RAM
40 MAIA DATAFLOW ENGINES



X50
TRAINING STATIONS



HIGH END WORKSTATION

CROSFIELD

DL

RAL

SCARF
JASMIN

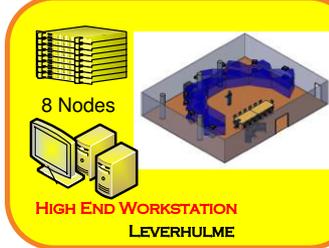
EMERALD
NCCS



USER



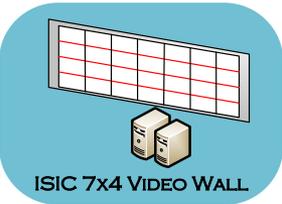
INTERACTIVE TABLE



8 Nodes

HIGH END WORKSTATION

LEVERHULME



ISIC 7x4 VIDEO WALL



HIGH END WORKSTATION
ISIC VISUALIZATION



HIGH END WORKSTATION
ATLAS VISUALIZATION WALL

