



enabling bigger questions

# Advanced Research Computing

*(previously know as Oxford Supercomputing Centre)*

*Things to do with data: How to do something with data in the ARC*

Dr Andrew Richards

An IT Services and Oxford e-Research Centre Partner facility

# Who are we?

- University central High Performance Computing
- Free training on how to get started
- Access for all AND provide access to priority services for paying customers (internal or external)







# Going Beyond Personal Computing...



# Computing to Compete

- IT in all forms increasingly underpins research activities
- Resources of the scale required can not always be provided locally at research institutes.
- A clear need for regional, national and international collaboration on e-infrastructure



# HPC Matters ?

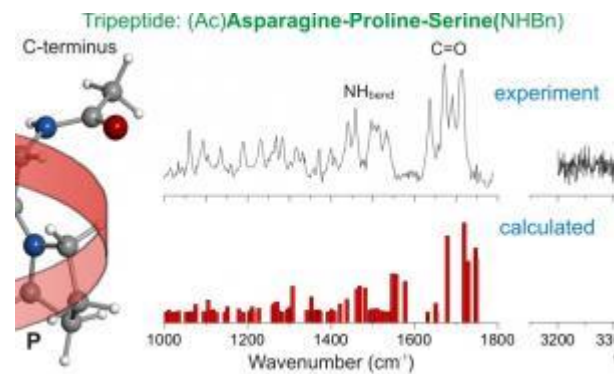


- Analysing data - see structure in the chaos
- Modelling endless combinations of molecules to find the cure for *<insert disease>*
- Turn back the clock 14 billion years
- Modelling the path of storms
- Every Day life
  - Investments, cars, shopping
- HPC Matters
  - <http://sc14.supercomputing.org/about-sc14>



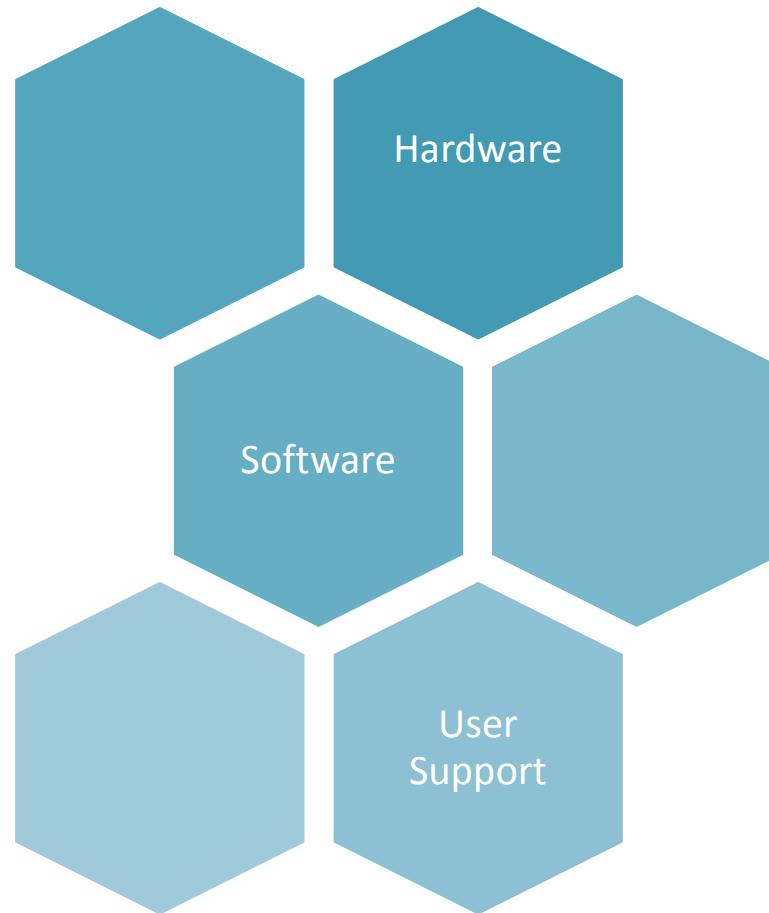
# Engine room of Research...



An IT Services and Oxford e-Research Centre Partner facility



# Provide access to






# ARC Services



- Local HPC
  - *ARCUS, JADE, SAL, HAL, CARIBOU, SKYNET, RUBY, PHILEAS...*
- Regional HPC
  - SES-5 e-Infrastructure
    - EMERALD (GPU) and IRIDIS (x86)
- Training, Support, Consultancy



# Research Computing Pyramid



**Tier 0:** Europe wide with users from multiple countries, e.g. PRACE.  
Tier-0 for the particle physics community is the HPC Data Centre at CERN.

**Tier 1:** National facility.  
e.g. ARCHER facility  
For particle physics users it is the LHC Tier-1 Centre at RAL.

**Tier 2:** Regional Centres  
e.g. SES Centre for Innovation

**Tier 3:** main institutional computing service such as ARC

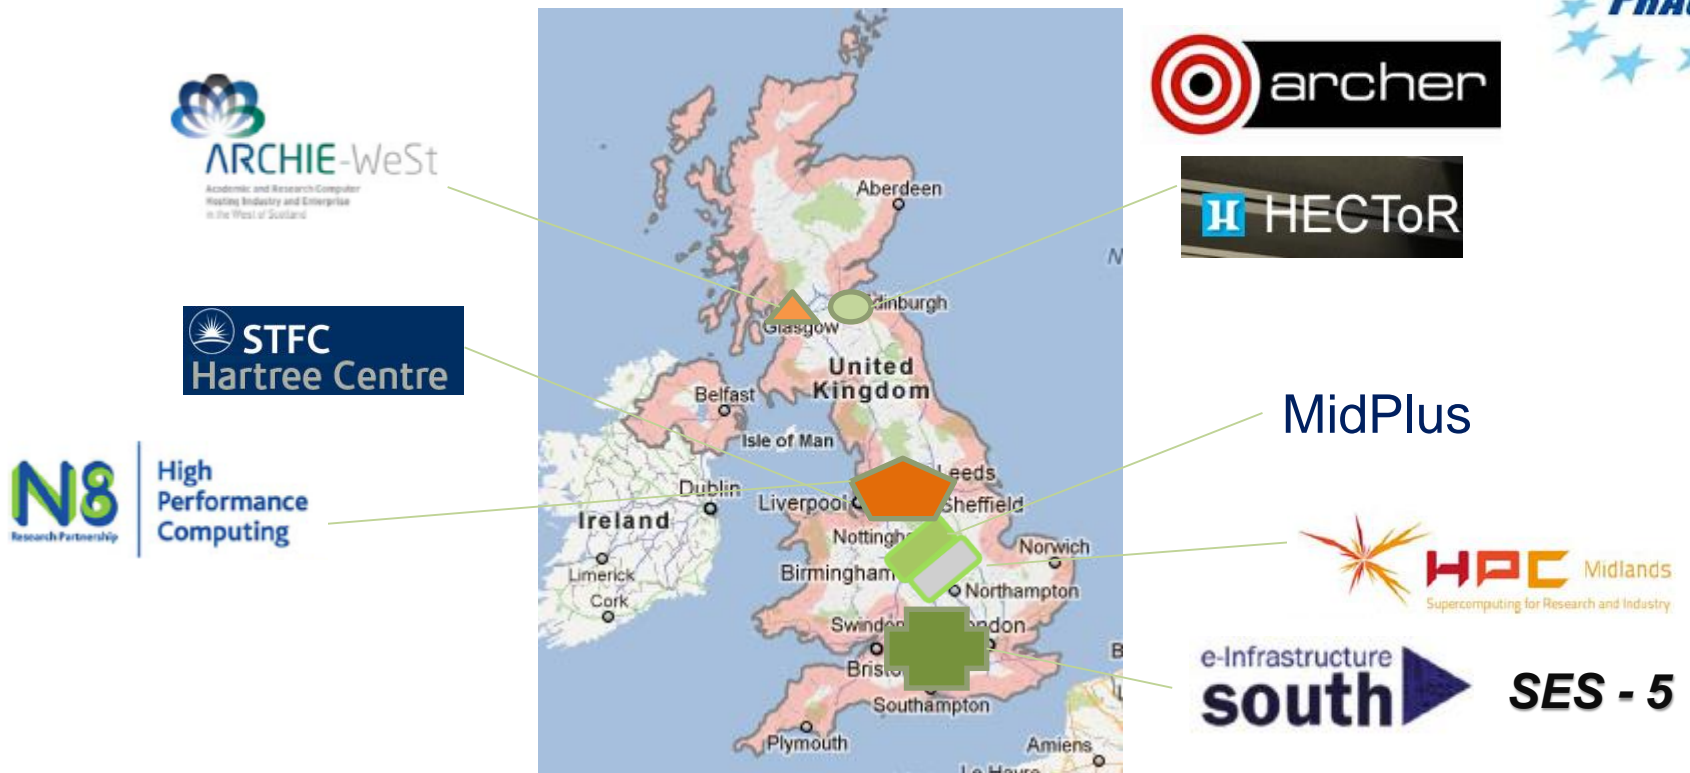
# HPC Landscape in the UK



# UK e-Infrastructure



UK Government decided there was a need for regional research infrastructure to link into national facilities



# HPC Terminology?



High Throughput  
Computing  
(Capacity)

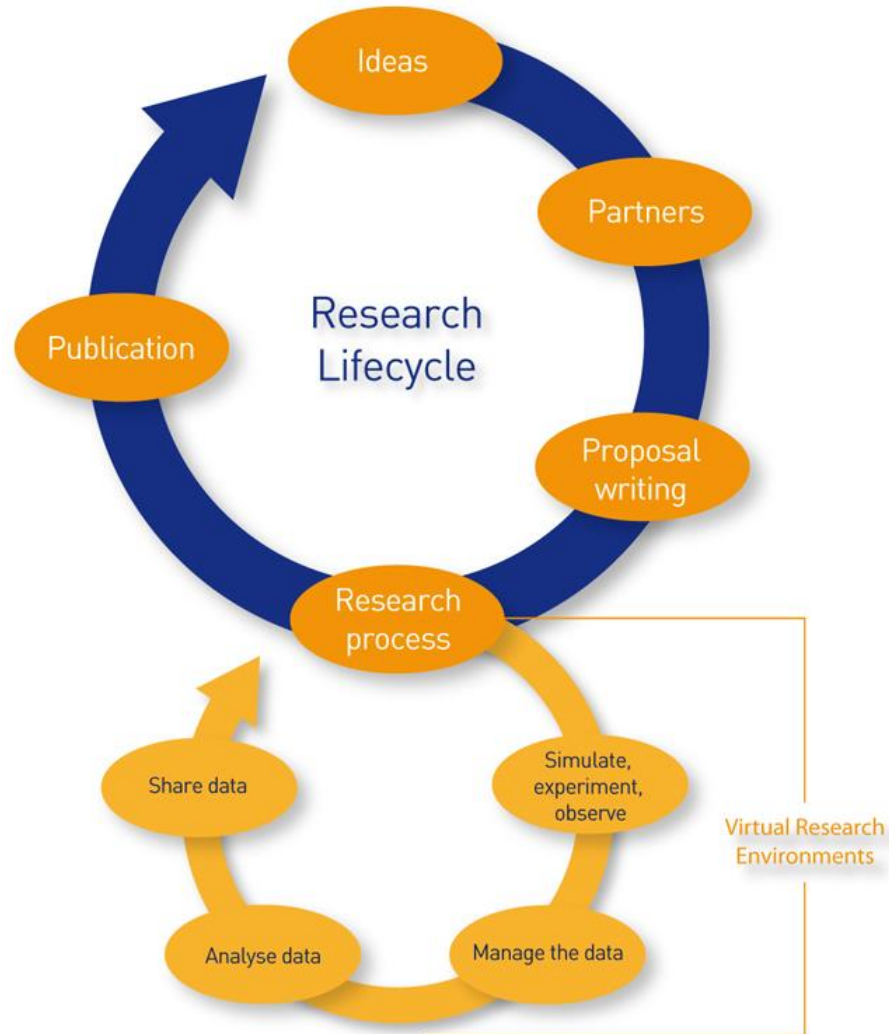
High Performance  
Computing  
(Capability)

# Capacity vs Capability Computing



Capacity	Capability
Using computing power to process many problems simultaneously.	Using maximum computing power to solve a large problem that no other computer can
Individual tasks, no communication between processes	One single individual task comprised of many child processes all communicating with each other
No specialist hardware interconnects needed between nodes in a cluster	Specialist, high performance, low latency, hardware interconnects needed between nodes in a cluster

# Research Data Lifecycle



# Data at the heart of computing



An IT Services and Oxford e-Research Centre Partner facility



# Big Data



# Big data and in-storage processing

- In storage processing



# Data Analytics




The screenshot shows a web browser window with the following content:

- Page Title:** IBM Calls on Watson Tech
- URL:** www.sci-tech-today.com/news/IBM-Calls-on-Watson-To-Manage-Storage/story.xhtml?story\_id=030003405HSO
- Article Title:** IBM Calls on Watson Technology To Manage Storage
- Author:** By Barry Levine
- Date:** May 12, 2014 3:32PM
- Summary:** The IBM Elastic Storage technology virtualizes storage, so a variety of systems and applications can share the same storage, which means applications do not need to be modified or storage management added. IBM's analytics are used to move seldom-used data to low-cost tape drives, and to move frequently accessed data to flash systems.
- Related Topics:** Watson, IBM, Software Defined Storage, Cloud
- Latest News:** IBM Calls on Watson To Manage Storage, Is Exercise the Fountain of Youth?, Dinosaurs Healed Bone Injuries, Astronauts Go Underwater for Test, Tech Not To Blame for Teen Woes
- Advertisement:** Eliminate costly downtime! Find out how in our expert white paper, and enter for a chance to WIN a Samsung Galaxy Note 3. Learn more!
- Main Text:**
  - IBM is out with a software-defined storage it is calling Elastic Storage -- and its roots trace back to the Jeopardy-winning supercomputer Watson. The company said the new approach can reduce storage costs by as much 90 percent by automatically moving it to the cheapest storage device.
  - The patented technology, a product of the company's esteemed Research Labs, is designed to "exploit -- not just manage" the growth in big data resulting from the proliferation of devices, sensors, processes and social networks. The technology was introduced Monday at an event in Boston. Initially, the service will be available only on-premises, but later this year it will become available through the company's SoftLayer cloud.
  - Tom Rosamilia, senior vice president of the IBM Systems and Technology Group, said in a statement that, with digital data growing so rapidly, "traditional storage systems used to house and manage it will eventually run out of runway." He added that the IBM tech is intended to use advances in speed, scalability and cost savings to accommodate the boom in data storage and the need for data access.
- Section: 200 Million Pages**
  - Using software-defined storage, data is automatically managed locally and globally, which the company said can provide dramatic speed improvements for access, easier administration and the ability to scale quickly. The technology, IBM said, can work with any company's storage systems.
  - Watson, which played on live TV against two human champions of Jeopardy, had access to 200 million pages of structured and unstructured data, as well as to all of Wikipedia. A technology similar to Elastic Storage was used to load 5 terabytes of knowledge -- which is what the 200 million pages represented -- in minutes into Watson's memory.
- Image:** A cartoon illustration of a dinosaur-like creature with a computer monitor for a head, standing next to a server rack. Text: MEET THE BEASTS, WATCH THE VIDEO > Kodak alaris
- Section: Also Visit:**
  - Billions of Bytes
  - Mobile Device Now
  - Apple Info Center
  - TopTechWire.com
- Text:** Salesforce.com is the market and technology leader in Software-as-a-Service. Its award-winning CRM solution helps 82,400 customers worldwide manage and share business information over the Internet. Experience CRM success. Click here for a FREE 30-day trial.
- Section: Innovation**
  - IBM, Universities Team on Watson
  - Cisco Invests \$150 million in IoT
  - Does Google Glass Cost \$80 To Make?
  - Gesture Control Ring 'Nod' Debuts

# How to do something with data (locally on ARC)

# How much computing knowledge do I need ?



**case study**

## Collisions and Complexes of Free Radicals

Sarandis Marinakis from the Department of Chemistry used the ARC facilities for his project: *Collisions and Complexes of Free Radicals*. Details about this research can be found on the ARC website. Below he responds to questions about his experience of using the ARC facilities.

**What were the main reasons you started using the facilities at ARC ?**

My programmes sometimes need more than 64 GBytes, so it was the access to a very large amount of RAM memory that was key for me. I was relieved to discover there were computing facilities on that scale within the University. And that there's such a lot of support to get your programmes running.

**What kind of support have you received from the Team at ARC ?**

I was given a lot of help compiling a very tricky programme – HLRidon. The ARC Team gave me invaluable advice in one-to-one meetings whenever I needed it.

**What other benefits have you experienced from using the ARC facilities?**

Something that was very useful was that I was able – with the help of the Team – to try out programmes that were not already installed. It was also fantastic that a few software companies allowed us to try out some expensive programmes at ARC for free.

**What would you say to researchers who are thinking of using the facilities at ARC ?**

I'd recommend they contact the Team and discuss their plans. Some people may be put off if they can't see the software they want to use on the list of those available, but ARC will get hold of new software if there's sufficient demand for it (and it's not too expensive). So the list is continually expanding in response to the needs of users.

**case study**

## Mapping the Muslim population of Uttar Pradesh

Raphael Susewind is an Associate of the Contemporary South Asian Studies Programme at Oxford while he completes his PhD research into the politics and poetics of Muslim belonging.

I first started using the facilities at ARC when I was in India working on my PhD research and wanted to create a map showing the distribution of the Muslim population in Lucknow. Since such data is not officially available, I had to make informed guesses based on the religious connotations of names listed in each neighbourhood's public electoral register. Later I wondered whether – rather than just doing this for Lucknow with two million registered voters – I could scale up my algorithm to the whole state of Uttar Pradesh which has a population of close to 200 million. I soon realised that my desktop would take about four months for this task – and most likely longer because something would be bound to go wrong at some point. So I contacted the Team at ARC – with their resources I could do this work in less than 10 days.

I got my algorithm running, but I was wondering why it wasn't processing as quickly as I had expected when I received an email from ARC pointing out (very nicely) that I wasn't using the resources to the full. I was only using 1 CPU per node rather than 14. I had simply asked the ARC high-performance computer to do the same as my desktop, rather than exploiting the power of parallel computing.

The Team helped me rethink my approach and things speeded up considerably after that. I'm so grateful that they checked up on how my work was progressing – it's this kind of attention to supporting users that I value most about ARC. The computers are not simply very big desktops; because they can handle masses of data, they offer the potential for researchers to think in a different – and much more imaginative – way about the work they want to undertake.

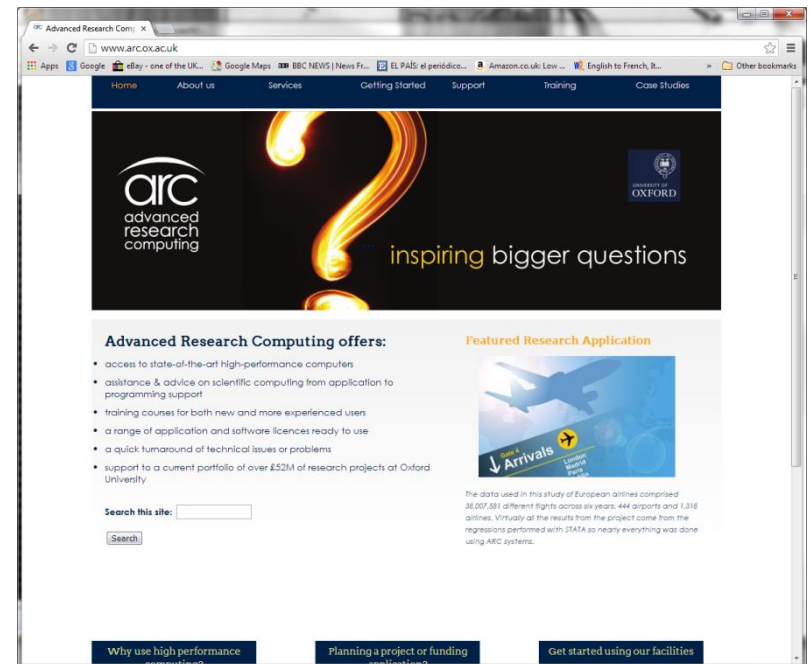
# How do I get on the Systems?

ARC is open to all researchers at the University of Oxford.

- To use the ARC you must first set up a new project and then apply for an individual user account.

Projects can be research groups, individuals or whatever best suits your local set up.

- For example, a research group may have two areas of research with different funding streams.
- It might make sense to create separate projects and have different user accounts for each project.
- Users can decide on the most sensible way to approach this.



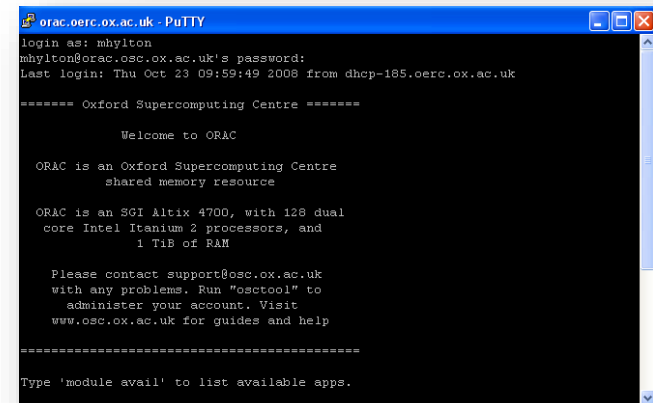
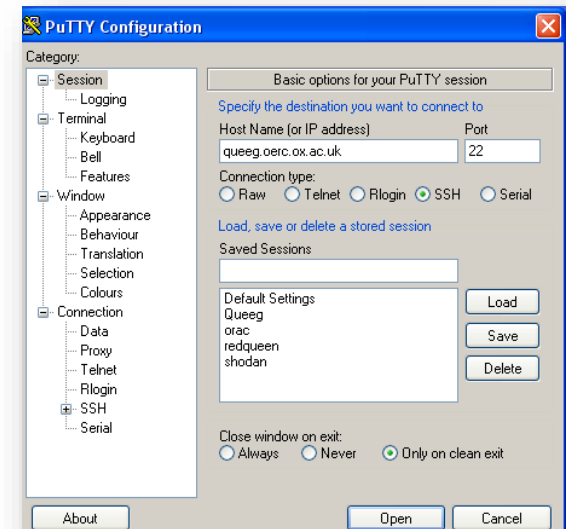
# Logging onto ARC Systems

## Remote access

- Linux, Mac (and other Unix/Unix-like) users should use **ssh** to connect to the ARC systems from a terminal
- `ssh -X username@systemname.arc.ox.ac.uk`
- Windows users should download and install an application called PuTTY and Xming for X11 support

## Batch Jobs

- Never run intensive jobs on ARC systems without using the job scheduler
- Jobs are submitted to the scheduler using the "qsub" command and a job submission script



# Data Storage Policy

The ARC makes every effort to ensure the integrity of data stored on our facilities

- However, we are under no obligation to guarantee the integrity or availability of data - ***this is the responsibility of the individual user.***

No Backups (limited snapshots of home and data)

- The ARC does not accept any liability, financial or otherwise for loss of data.
- We recommend that users employ standard industry practice for their important data and store it at sites other than the ARC , for example, in their department.

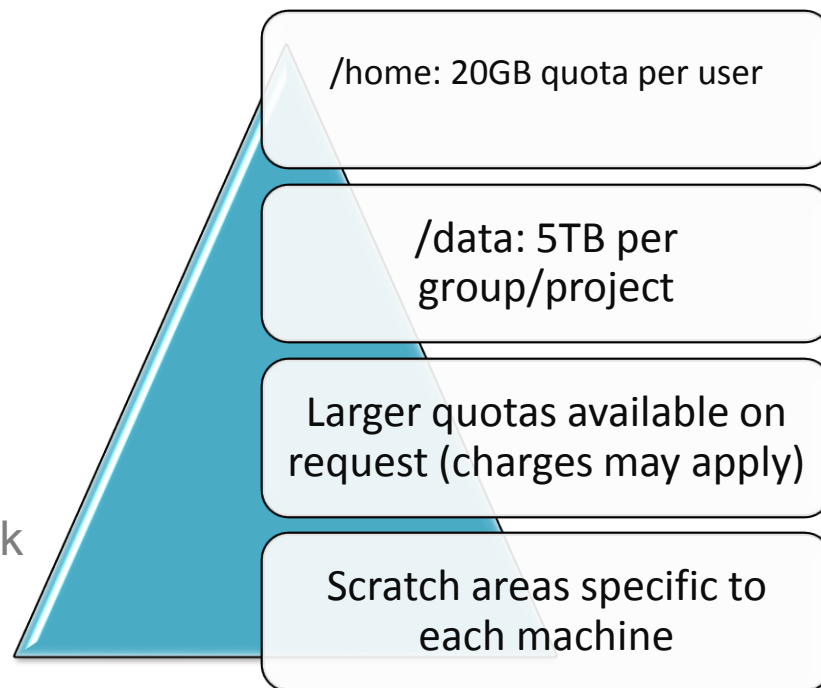


# Storage

ARC provides three types of storage

- *active data storage for ARC projects only*
- *limited backups of home and data*

- Home areas (/home)
  - mid-term storage
- General purpose (/data)
  - mid-term storage
- Scratch space (/scratch)
  - High performance storage
  - Jobs which perform significant disk read/write use the scratch disks and not /home
  - Short term (duration of job)



Storage Management: <http://www.arc.ox.ac.uk/content/storage-management>

# Transferring data

## Internal connectivity

- High speed, low latency Infiniband fabric
  - Mixture of QDR 40Gb/s (ARCUS) and FDR 56Gb/s (ARCUS-GPU). Fat tree
  - 40Gb/s connection (trunk) to storage
  - *Typically see 1GB/s transfer rates at present to storage infrastructure.*

## External connectivity

- 1Gb/s and 10Gb/s (rate limited at present) connection
- User tools: scp, sftp, *globus online*

# Using the right resource

First, move the data to ARC storage !

- Small data, big data?
- Input small output large ?
- Processing data from disk?
- Processing data in-memory?
  
- **Cluster nodes max memory 128GB**
- **SMP service max memory 1TB**

# Process data

## **Simulate / Experiment / Observe / Analyse**

- A large (> 200) applications already installed for processing data.
- Applications such as Matlab, Python, R, Abaqus are commonly used and available.
- We can help with installing custom applications or advise on 'home-grown' applications.
- Support high throughput (serial or 'capacity') and high performance ('capability') type jobs

# Visualization of data

- NoMachine NX server



# Contact

- [www.arc.ox.ac.uk](http://www.arc.ox.ac.uk)
- [support@arc.ox.ac.uk](mailto:support@arc.ox.ac.uk)
- [theteam@arc.ox.ac.uk](mailto:theteam@arc.ox.ac.uk)

