




Science and  
Technology  
Facilities Council

# STFC Cloud

Alastair Dewhurst



# Introduction

- STFC runs an OpenStack Cloud which aims to provide computing resources for STFC researchers and their collaborators.
  - Its not for running corporate services.
- <https://openstack.stfc.ac.uk/>
  - Can be accessed by IRIS IAM or STFC federal accounts.
- It has proved extremely popular and has grown exponentially since it was conceived in 2014.
- It has not yet been used extensively by the HEP community.
  - A project was setup for ExCALIBUR

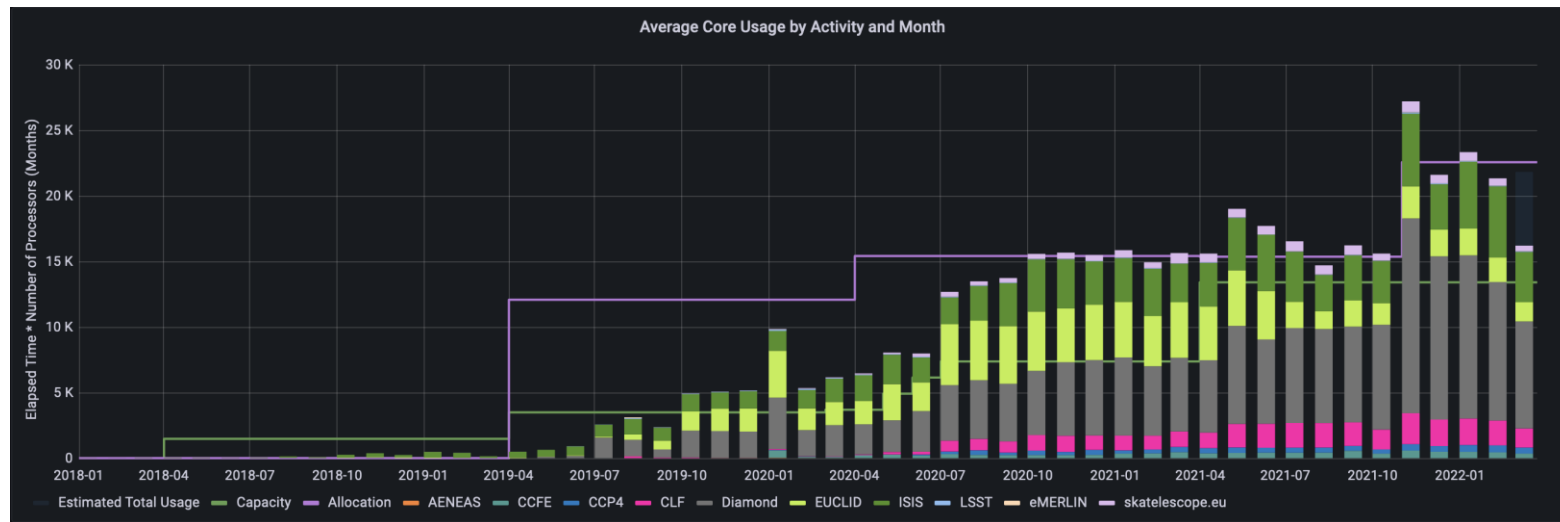
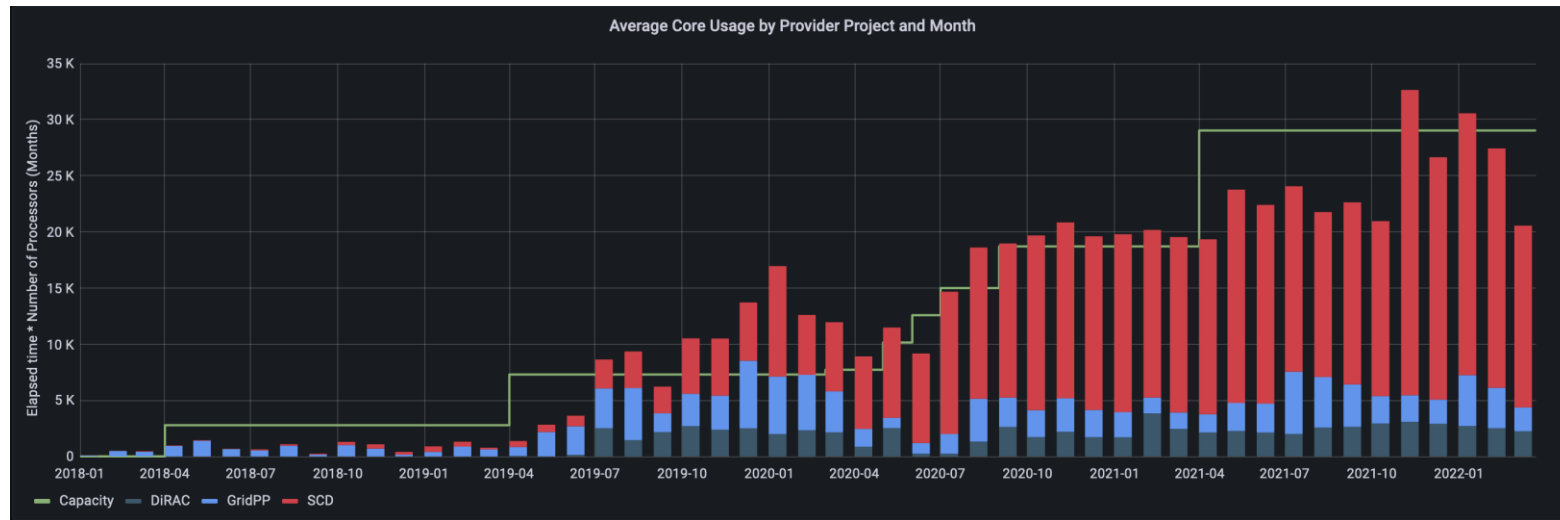
# Cloud compute

- The current capacity in CPU/GPU cores performance requirements and utilisation.
  - ~60k logical cores. ~90% utilization.
  - 240 Compute GPUs: 4 Tesla P100, 124 Tesla V100 and 112 Tesla A100s. ~97% utilisation.
  - 228 Visualisation GPUs: 108 Quadro P4000 and 120x Quadro RTX4000s. ~70% utilisation.
- Recently the 6 x FPGAs came online (that were ordered at the end of 2020...)
  - 3 x Intel PAC D5005, 32GB
  - 3 x Xilinx Alveo U200

# Cloud storage

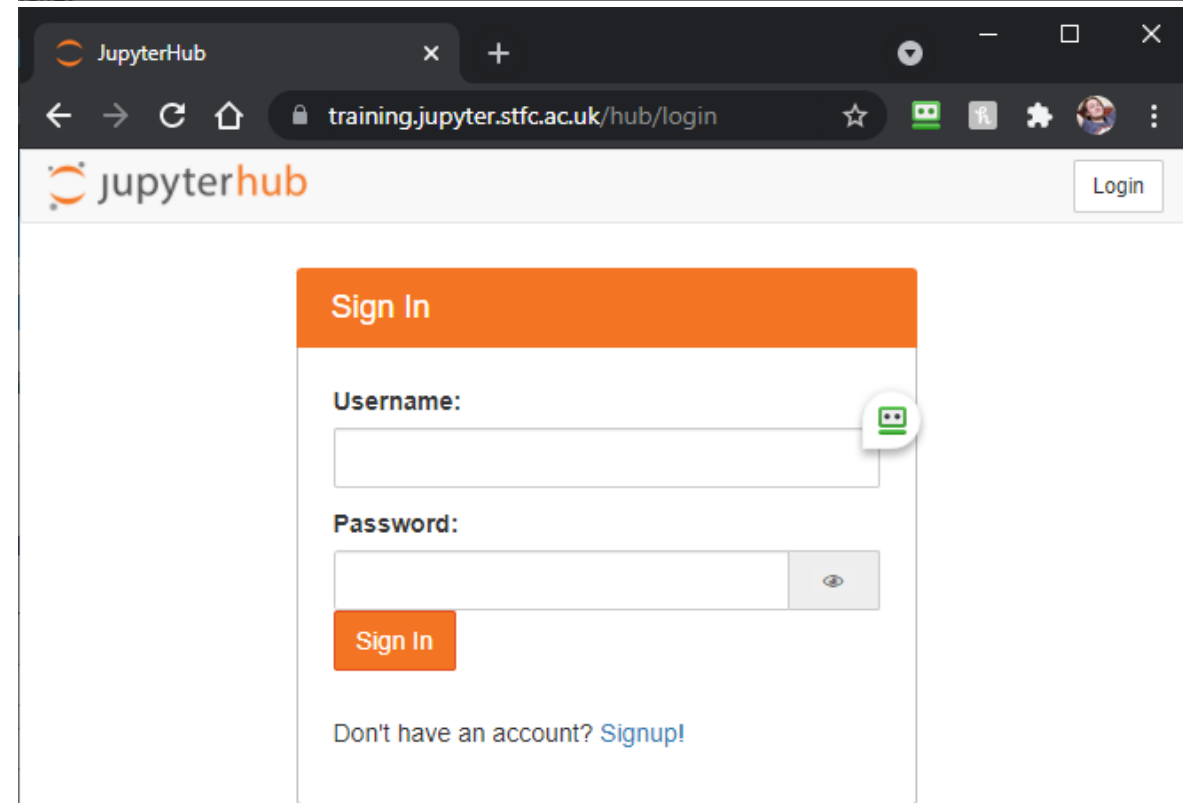
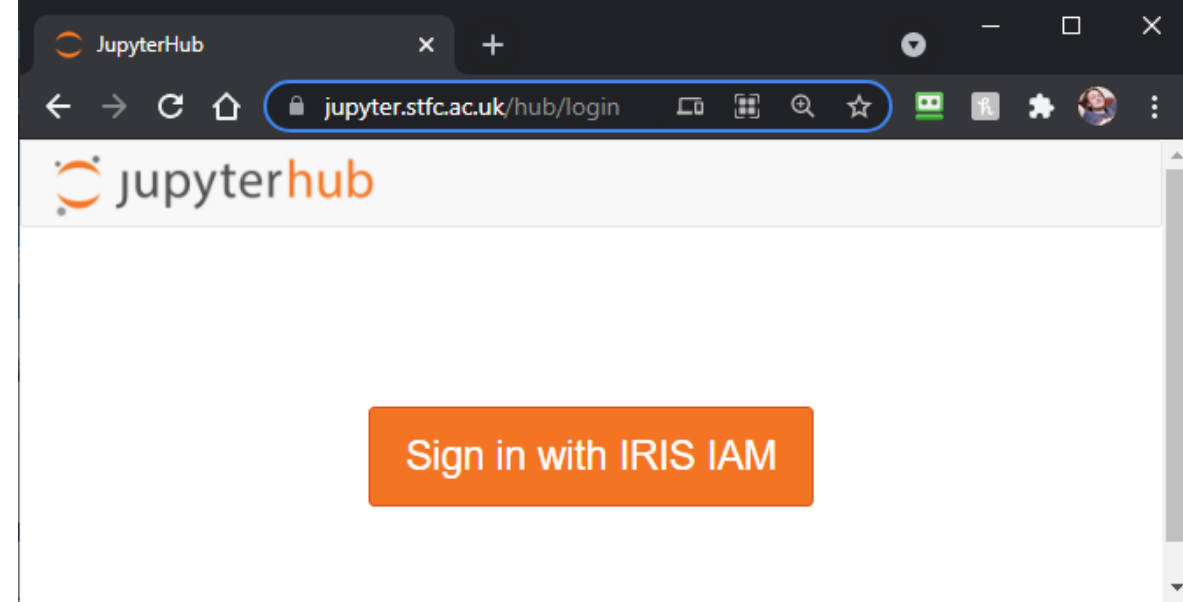
- There are a variety of storage endpoints that can be accessed via Cloud.
- Underpinning the Cloud is the Sirius Ceph cluster.
  - 768TB of NVMe block storage for high availability VMs.
- The Deneb Ceph cluster provides users with a shared file system.
  - 5PB of HDD in 8+3 Erasure Coding.
- New Arided Ceph cluster will provide a high performance shared file system across the cloud.
  - 1.5PB of SSD in (probably) 3 x replication.
- Users have access to Echo (via S3 or Swift) interfaces.
  - 400TB of usage.

# Cloud Usage



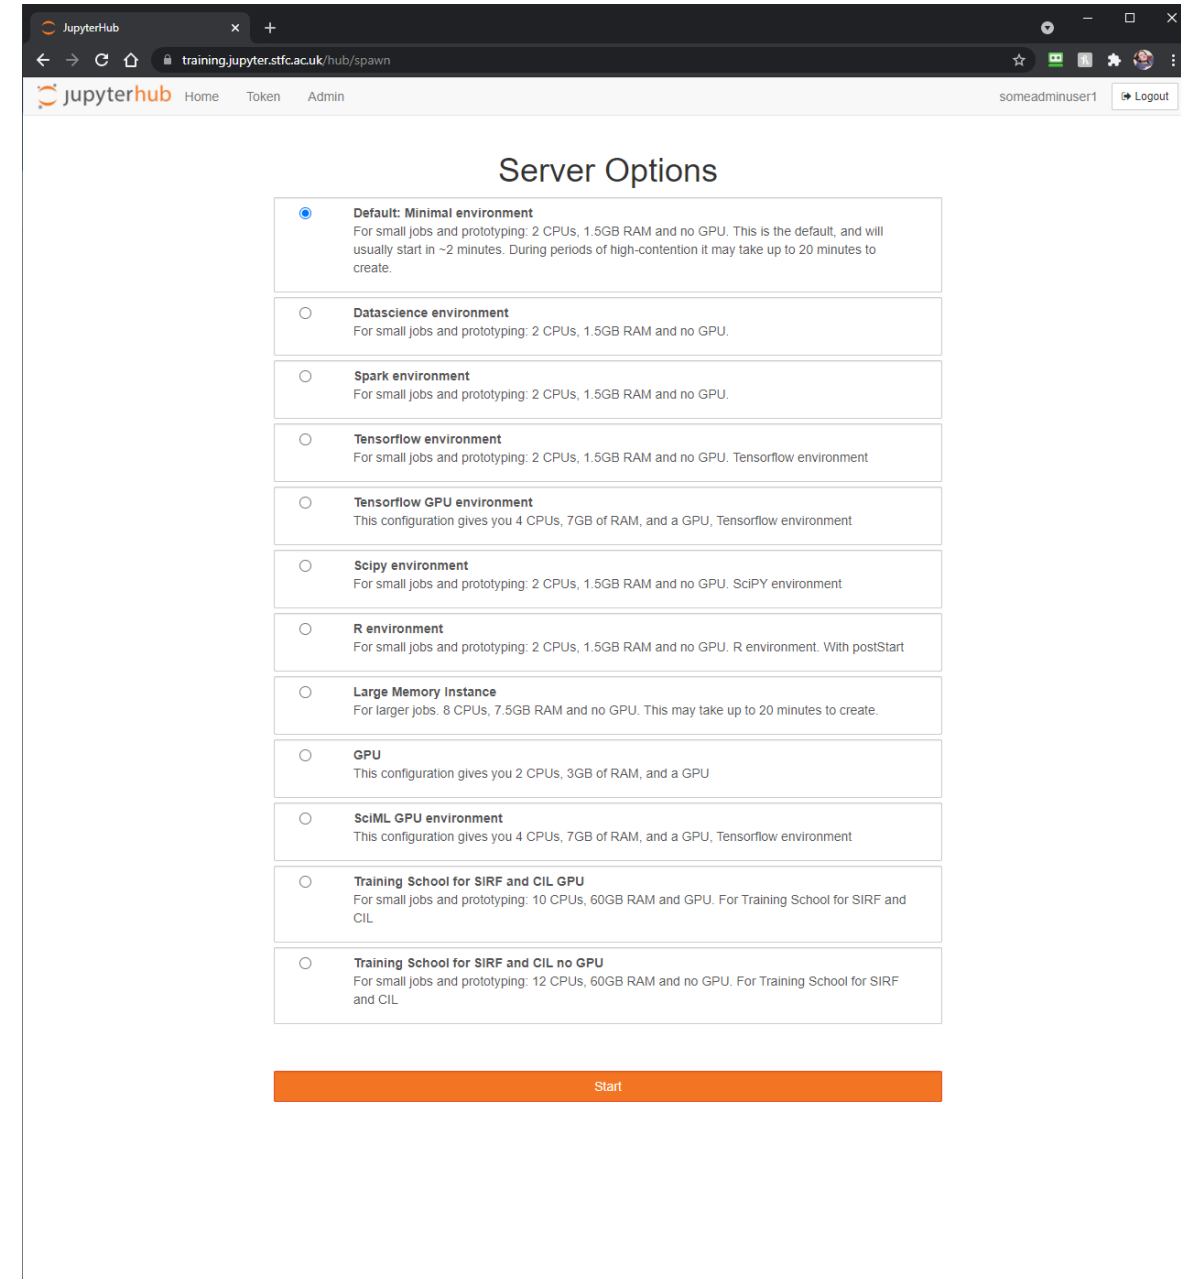
# JupyterHub

- For the last few years there has been a strong desire for a generic Jupyter infrastructure within STFC and IRIS
- This has been especially true for training courses.
- 2 infrastructures:
  - 1 with IRIS IAM login  
<https://jupyter.stfc.ac.uk>
  - 1 with local credential login (for training courses) <https://training.jupyter.stfc.ac.uk>



# Deployment

- Deployed on top of OpenStack Magnum
- CPU and GPU support
- All standard Jupyter environments supported
  - Also customised environments
- Users get 10GB of storage
- Shared storage available for training materials on training deployment.



The screenshot shows the JupyterHub interface for selecting a server environment. The browser address bar shows 'training.jupyter.stfc.ac.uk/hub/spawn'. The page title is 'Server Options'. The selected option is 'Default: Minimal environment'. Other options include 'Datascience environment', 'Spark environment', 'Tensorflow environment', 'Tensorflow GPU environment', 'Scipy environment', 'R environment', 'Large Memory Instance', 'GPU', 'SciML GPU environment', 'Training School for SIRC and CIL GPU', and 'Training School for SIRC and CIL no GPU'. A 'Start' button is located at the bottom of the list.

Environment	Description
<input checked="" type="radio"/> Default: Minimal environment	For small jobs and prototyping: 2 CPUs, 1.5GB RAM and no GPU. This is the default, and will usually start in ~2 minutes. During periods of high-contention it may take up to 20 minutes to create.
<input type="radio"/> Datascience environment	For small jobs and prototyping: 2 CPUs, 1.5GB RAM and no GPU.
<input type="radio"/> Spark environment	For small jobs and prototyping: 2 CPUs, 1.5GB RAM and no GPU.
<input type="radio"/> Tensorflow environment	For small jobs and prototyping: 2 CPUs, 1.5GB RAM and no GPU. Tensorflow environment
<input type="radio"/> Tensorflow GPU environment	This configuration gives you 4 CPUs, 7GB of RAM, and a GPU, Tensorflow environment
<input type="radio"/> Scipy environment	For small jobs and prototyping: 2 CPUs, 1.5GB RAM and no GPU. SciPY environment
<input type="radio"/> R environment	For small jobs and prototyping: 2 CPUs, 1.5GB RAM and no GPU. R environment. With postStart
<input type="radio"/> Large Memory Instance	For larger jobs. 8 CPUs, 7.5GB RAM and no GPU. This may take up to 20 minutes to create.
<input type="radio"/> GPU	This configuration gives you 2 CPUs, 3GB of RAM, and a GPU
<input type="radio"/> SciML GPU environment	This configuration gives you 4 CPUs, 7GB of RAM, and a GPU, Tensorflow environment
<input type="radio"/> Training School for SIRC and CIL GPU	For small jobs and prototyping: 10 CPUs, 60GB RAM and GPU. For Training School for SIRC and CIL
<input type="radio"/> Training School for SIRC and CIL no GPU	For small jobs and prototyping: 12 CPUs, 60GB RAM and no GPU. For Training School for SIRC and CIL

# Coyote


- The Tier-1 batch farm uses HTCondor.
  - There is a feature called Rooster which allows HTCondor to “wake up” machines when they are needed by the batch system.
  - Coyote’s eat Roosters...
- Coyote was set of scripts written by the Cloud team that allowed the Tier-1 to create VMs on the Cloud which would be included in the batch farm.
  - The Cloud would “re-claim” these VMs when demand for interactive machines grew.
- The Cloud is very full and without at least some effort to optimise things you couldn’t run data intensive jobs on these machines.
- With some effort Coyote could be used to spin up more exotic VMs for use by the Tier-1 batch farm.





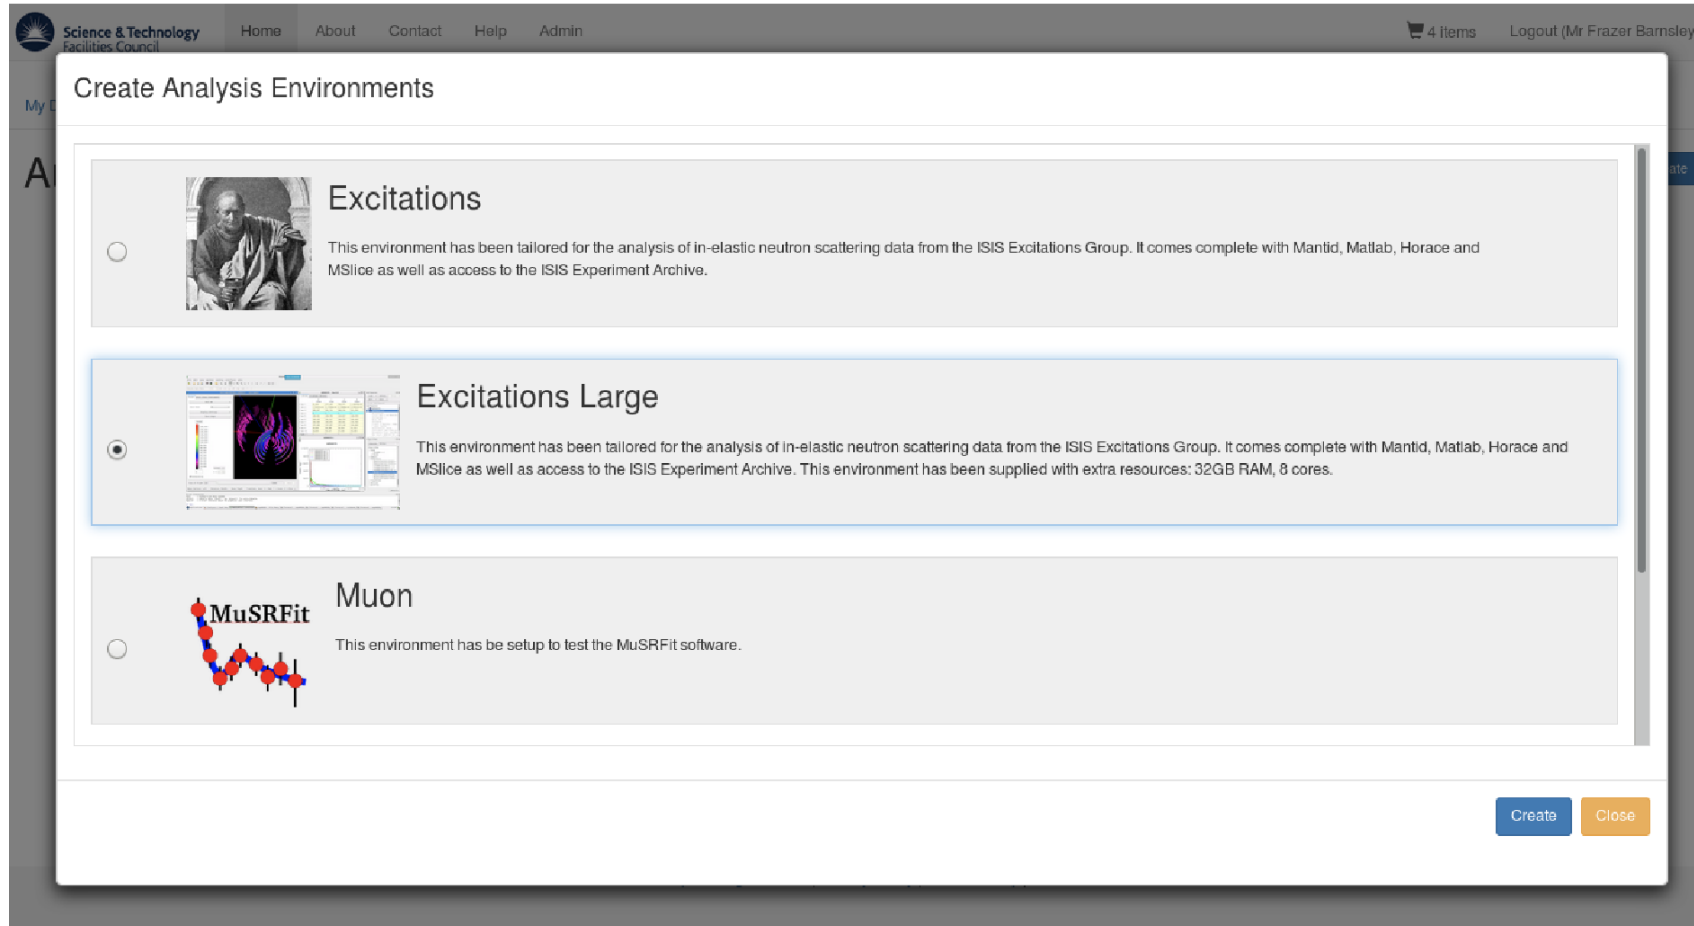
# DAaaS

- DAaaS stands for Data Analysis as a Service.
  - The DAaaS group (~6FTE) to support this.
- The main users are the 3 major facilities at RAL:
  - Diamond, ISIS and CLF
- ~10,000 scientists a year with data sets ranging from MB to many TB.




DAaaS is a cloud-based data analysis platform, which enables scientists to access a virtual environment powered by the STFC Cloud via a web browser. With DAaaS, scientists can have everything they need for their data analysis, from raw data to software tools, sharing their analysis and working together. We are providing scientists with all the tools they need to deliver their research in an efficient way all in one convenient place.

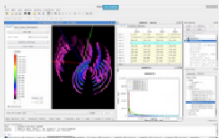
# DAaaS user experience

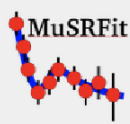


Science & Technology Facilities Council Home About Contact Help Admin 4 Items Logout (Mr Frazer Barnsley)

## Create Analysis Environments

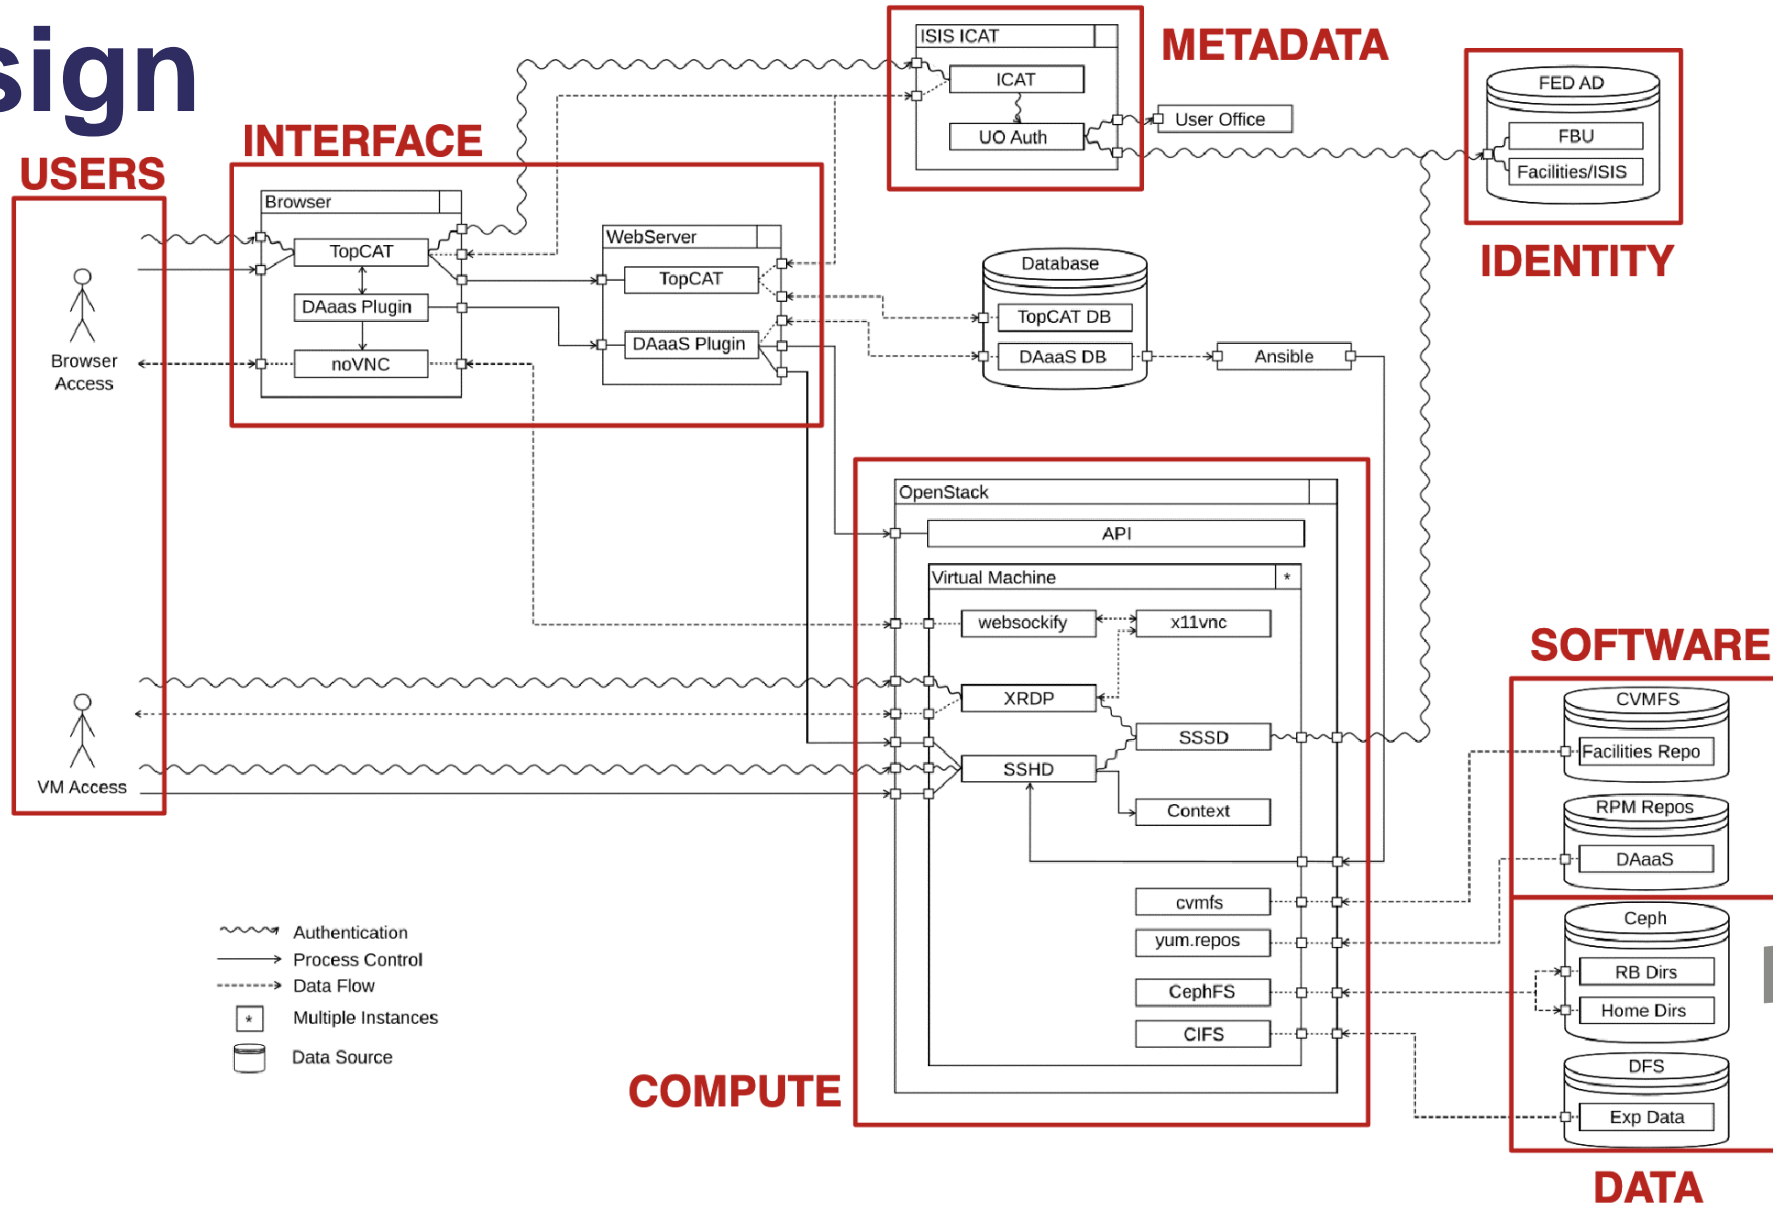
 **Excitations**  
This environment has been tailored for the analysis of in-elastic neutron scattering data from the ISIS Excitations Group. It comes complete with Mantid, Matlab, Horace and MSlice as well as access to the ISIS Experiment Archive.

 **Excitations Large**  
This environment has been tailored for the analysis of in-elastic neutron scattering data from the ISIS Excitations Group. It comes complete with Mantid, Matlab, Horace and MSlice as well as access to the ISIS Experiment Archive. This environment has been supplied with extra resources: 32GB RAM, 8 cores.

 **Muon**  
This environment has been setup to test the MuSRFit software.

Create Close

# Design





Science and  
Technology  
Facilities Council

# Questions?