




Science and  
Technology  
Facilities Council

Scientific Computing

# Design considerations for an environmentally sustainable Datacentre

Dr Peter Oliver, Scientific Computing, STFC  
[peter.oliver@stfc.ac.uk](mailto:peter.oliver@stfc.ac.uk)



# Agenda

**1 Strategic Case**

**2 STFC Environmental Strategy**

**3 Build and Operation**

**4 Summary**



# Agenda

**1 Strategic Case**

2 STFC Environmental Strategy

3 Build and Operation

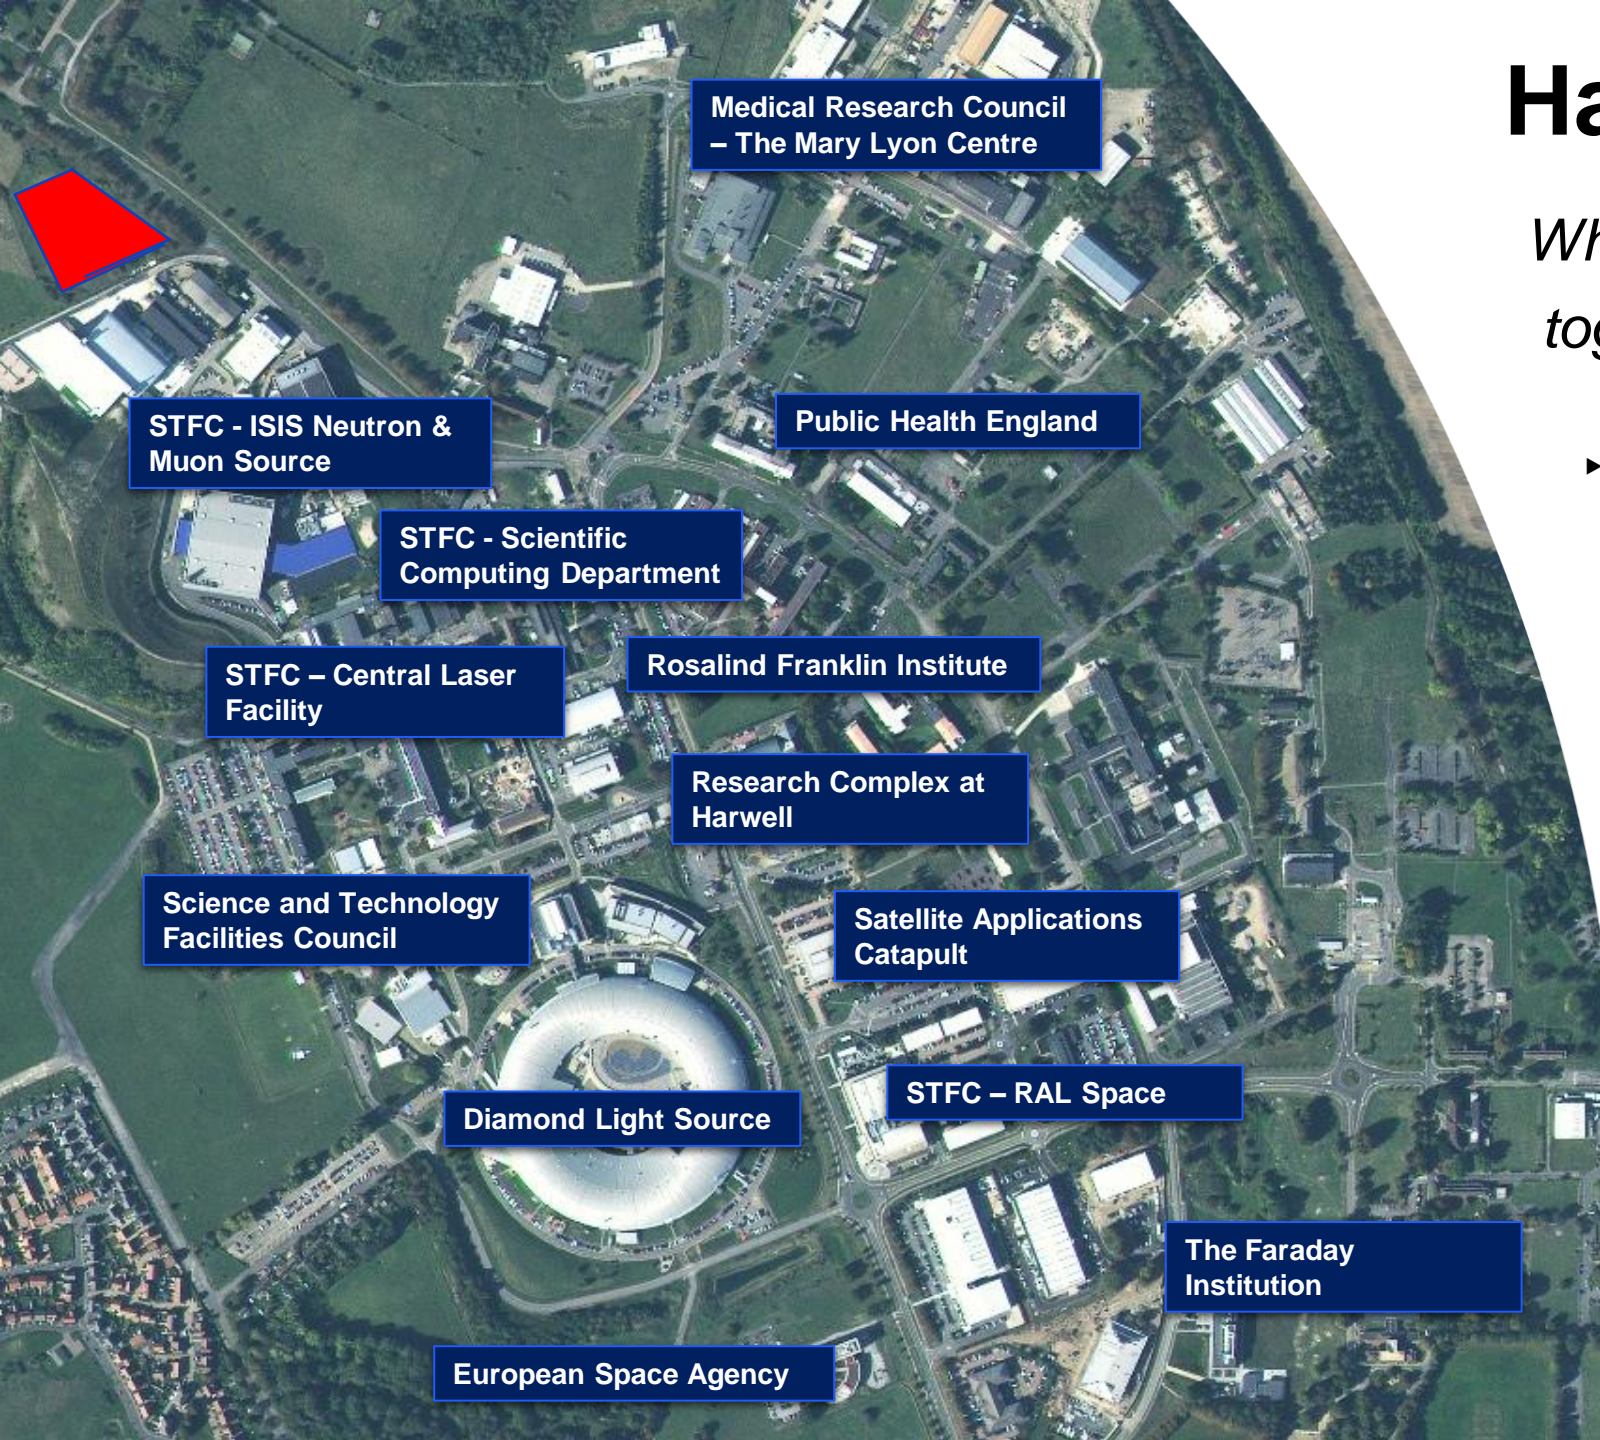
4 Summary



# Harwell Campus

*Where science and industry comes together to tackle global problems*

- ▶ £2+bn national infrastructure that can study viruses to galaxies
- ▶ 6,000 people designing and operating facilities
- ▶ Supporting research at 70+ UK universities
- ▶ Working with scientists in 60+ countries



**HARWELL**

# Research Computing Centre (RCC): Requirements Capture


- Major cross-council computational and storage requirements in future years
- Data rates doubling every 6-12 months.



Activity	Purpose	Interested RC
Ada Lovelace Centre (ALC)	Provides software, data services and skills to exploit data from large scale national facilities including Diamond, ISIS and CLF/EPAC	STFC/EPSRC/ BBSRC/MRC
Extreme Photonics Applications Centre (EPAC)	It will bring together world-leading interdisciplinary expertise to develop and apply novel, laser based, non-conventional accelerators and particle sources which have unique properties	STFC
IRIS	A co-operative consortium of STFC science users and compute providers, to deliver High Throughput Computing (CPU and GPU), data management and storage capability to achieve the science goals of the National Facilities (including ISIS, Diamond, CLF), of PPAN projects and instruments (including SKA, LSST, LIGO,LHC, DUNE), and CCFE	STFC
National Quantum Computing Centre (NQCC)	100+ qubit NISQ demonstrator hardware platform(s), Quantum software, algorithm & applications development High performance, scalable qubit technology development, Roadmap and architecture towards fault-tolerant general purpose quantum computing	EPSRC
Rosalind Franklin Institute (RFI)	A new national research centre at the interface between the engineering, physical, and life sciences	EPSRC
Diamond Light Source (DLS)	Diamond Light Source is the UK's national synchrotron science facility	STFC Wellcome
Diamond Light Source II (DLS II)	Diamond-II upgrade, a co-ordinated programme of development that combines a major machine upgrade with new instruments and complementary improvements to optics, detectors, sample environment and delivery capabilities, and computing, as well as integrated and correlative methods	STFC Wellcome
Electron Bio-Imaging Centre (eBIC)	eBIC provides scientists with state-of-the-art experimental equipment and expertise in the field of cryo-electron microscopy	STFC/BBSRC MRC/Wellcome
JASMIN	JASMIN is a globally-unique data analysis facility. It provides storage and compute facilities to enable data-intensive environmental science.	NERC
CCP4	Collaborative Computational Project for Macromolecular X-Ray Crystallography	BBSRC
CCP-EM	The Collaborative Computational Project for electron cryo-microscopy (CCP-EM) supports users and developers in biological EM.	MRC
Mary Lyon Centre	A national facility providing world-class expertise, tools and resources to generate and characterise genetically altered mouse models for use	MRC

# Science Drivers

Major new cross council Digital Research Infrastructure (DRI) required to fully exploit science infrastructure



# RCC: Cross Council Scientific Research

- RCC will need to host a variety of equipment such as
  - Large scale data storage
  - High Throughput Computing (HTC)
  - High Performance Data Analytics (HPDA)
  - High Performance Computing (HPC) and
  - Specialist computing hardware, for example Machine Learning (ML)
- This encompasses a broad range of power requirements from <math><10\text{kW}/\text{rack}</math> to .
- RCC is the datacentre only and NOT not the compute & data resources
  - Computing/Data equipment purchases are driven by major scientific programmes/projects/facilities
- Analysis of requirements indicates the need for 11MW of computing load by 2030 and 20MW by 2035.
  - Phased delivery, with 6.6MW capability by 2026, option to install more capacity
  - Capacity in excess of 6.6MW would follow an additional review of requirements and funding request
  - Meet critical requirement for DLS, STFC, NERC by 2025/6 and DLS II from 2028
- Part of a UKRI eco-system of datacentres
  - Complementing Hartree Centre and national capabilities such as ARCHER 2 and Exascale.



# Agenda

1 Strategic Case

2 STFC Environmental Strategy

3 Build and Operation


4 Summary





# Environmental Sustainability

- Environmental Sustainability is a critical Strategic aim across UKRI and STFC



# Agenda

1 Strategic Case

2 STFC Environmental Strategy

3 Build and Operation

4 Summary



# What does environmental sustainability mean for a datacentre build?

- The RCC is a new datacentre and can adopt international standards for the build.
- BREEAM
  - The world's leading science-based suite of validation and certification systems for a sustainable built environment.
  - External validation of the buildings design
- The current BREEAM assessment is broken down into 10 categories and 5 standards.


- Energy
- Health & Wellbeing
- Innovation
- Land Use
- Materials
- Management
- Pollution
- Transport
- Waste
- Water



- The RCC is aiming for BREEAM Excellent but has the potential for outstanding

# BREEAM - Ecology


- **Ecology Survey** “A large portion of the site comprises lowland calcareous grassland priority habitat. Any loss of this habitat should be limited where possible, with any unavoidable loss to be compensated for through enhancement of retained areas. “
- RCC will keep existing trees where possible




- Site boundary
- Other neutral grassland
- Calcareous grassland
- Buildings
- Hardstanding
- Scattered scrub
- Ruderal/ Ephemeral
- Blackthorn scrub
- Street trees
- Fence

# BREEAM - Ecology

- **Ecology Survey** “A large portion of the site comprises lowland calcareous grassland priority habitat. Any loss of this habitat should be limited where possible, with any unavoidable loss to be compensated for through enhancement of retained areas. “
- RCC will keep existing trees where possible



- Site boundary
- Other neutral grassland
- Calcareous grassland
- Buildings
- Hardstanding
- Scattered scrub
- Ruderal/ Ephemeral
- Blackthorn scrub
- Street trees
- Fence



# What does environmental sustainability mean for a datacentre? In Operation

- The RCC will be able to support 6.6MW of Computing Load.





A typical UK house uses 241kWhrs per month



RCC 4,752,000kWhrs per month  
**~20,000** houses

# What does environmental sustainability mean for a datacentre? In Operation

- Only ~38% of the UK's electricity is generated from renewables



- As we have only 1 electricity grid. The electricity you use comes from all sources even if on a “green tariff”. We need to consider CO<sub>2</sub> emissions.

# What does environmental sustainability mean for a datacentre? In Operation

- The RCC needs to be as power efficient as possible.
- **Power usage effectiveness (PUE)** is a ratio that describes how efficiently a computer datacentre uses energy; specifically, how much energy is used by the computing equipment (in contrast to cooling and other overhead that supports the equipment).
- Typical datacentres have a PUE of ~1.5
  - (as does the existing R89 Datacentre at RAL)
- This means that for every 1MW of computing power you use 0.5MW for cooling and other overheads
- Increasing efficiency and reducing PUE clearly has benefits on power usage!
- The RCC is targeting a PUE of 1.1



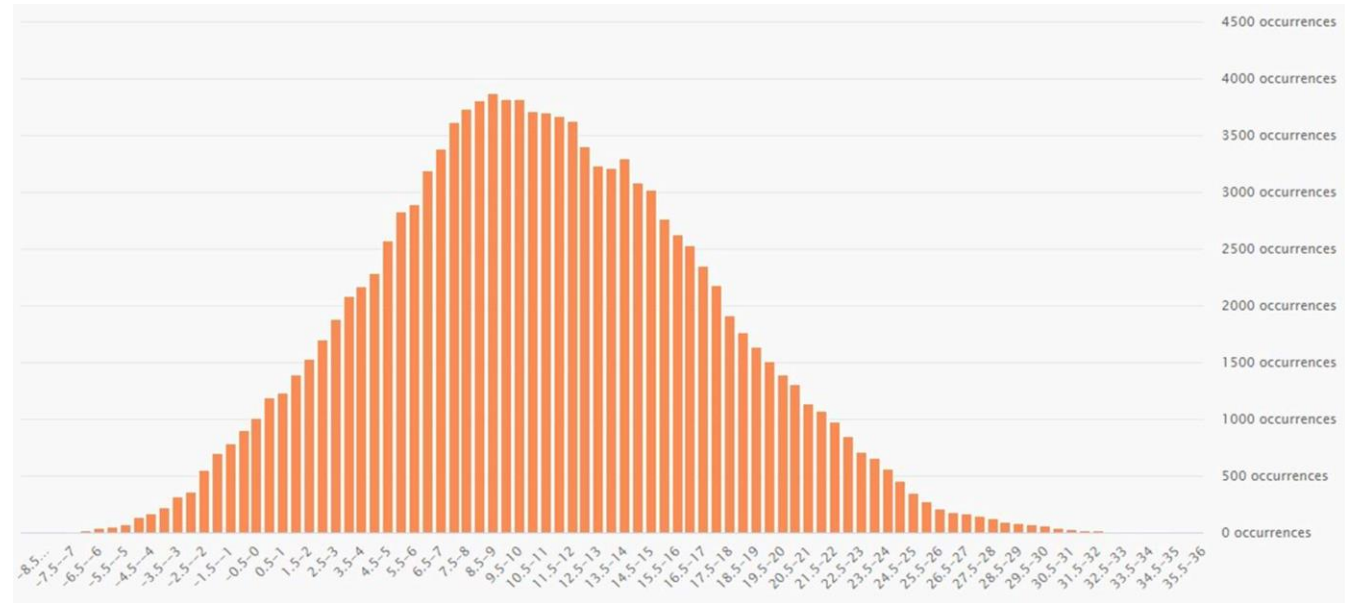
$$PUE = \frac{\text{Total Datacentre Power Usage}}{\text{IT Equipment Power Usage}}$$



# How to achieve a PUE of 1.1?

## Understanding the Weather

Temp °C	# Occurrences
30 – 31	77
31 – 32	35
32 – 33	16
33 – 34	16
34 – 35	9
35 – 36	2



Hourly ambient air dry bulb temperature for Harwell 2008-2021



Climate change will affect these so need to plan for the future for a datacentre with a 20 year life

# How to achieve a PUE of 1.1?

Understanding the Technology

Rear Door Heat Exchangers 100kW/rack

Type	Summary
Dry Air Coolers	Free Cooling, efficient average PUE, high peak PUE as needs chillers
Adiabatic cooling	Free Cooling, efficient average PUE, uses spray water, maintenance issues
Cooling Towers	Free Cooling efficient PUE, Uses LOTS of water, high maintenance
Hybrid Coolers	Free Cooling, efficient average and peak PUE, contained water, acceptable water usage as anticipated need low.



# PUE, CO<sub>2</sub> and Power Costs (2028)

Power costs 37.9p/kWh (based on Crown Commercial Services estimates)

kgCO<sub>2</sub>e/kWh 0.167 (UK Gov Green book)

Computing Load 6.6MW

PUE	Total MW	£m Electricity	tCO <sub>2</sub> e 1000's	# Houses 1000's
1	6.6	21.9	10	20
1.1	7.26	24.1	11	22
1.2	7.92	26.3	12	24
1.3	8.58	28.5	13	26
1.4	9.24	30.7	14	28
1.5	9.9	32.9	15	30
<b>Diff 1.5-1.1</b>	<b>2.6</b>	<b>8.8</b>	<b>4</b>	<b>8</b>

**Environmental Sustainability is a double win, less CO<sub>2</sub> and smaller Electricity Bill!**

## Additional activities to increase efficiency of the RCC

There will be a large area for Solar power generation building on STFC investment solar panels across the RAL and Daresbury. The existing investment is estimated to deliver 3,450MWh of energy each year at RAL and 1,200MWh at Daresbury annually.




There are great examples of where the waste heat from datacentres has been used for domestic district heating, agriculture and other processes. STFC is exploring if the waste heat can be used for new campus buildings.

# An Efficient Datacentre is not the end of the story.

The previous calculations have assumed all the racks and servers are performing “useful work”

However.....

- Need to optimise usage as an idle resource is not doing useful work
- Optimise applications
- That 10+ year old code could be re-engineered and perform more efficiently.
- Newer libraries
- Are you using the application in the best way?
- New approaches such as ML could drastically cut down time to solution (but training costs could be high)





# Summary

- Build and Operation matter
- Using standards such as BREEAM for the build minimises the environmental impact of the build
- Landscaping to minimise impact and enhance ecology and bio-diversity
- Cooling technology choices matter
- At a PUE of 1.1 can save up to £8.8m electricity compared to PUE 1.5!
- A low PUE **decreases tCO2e** and **decreases the electricity** costs

**BREEAM**<sup>®</sup>  
**EXCELLENT**



Scientific Computing





Science and  
Technology  
Facilities Council

Scientific Computing

A series of thin, light blue lines radiate from the top and left edges of the dark blue rectangular area, creating a sense of motion or data flow. These lines vary in length and angle, some pointing towards the center of the rectangle.

# Questions?



Science and  
Technology  
Facilities Council

Scientific Computing

# Thank you

[scd.stfc.ac.uk](http://scd.stfc.ac.uk)

 [@SciComp\\_STFC](https://twitter.com/SciComp_STFC)