UKRI DRI
Net Zero
by 2040

Measuring Carbon

NetZero and the IRISCAST Project

Audit of Carbon Costs

Net Zero DRI Scoping Project

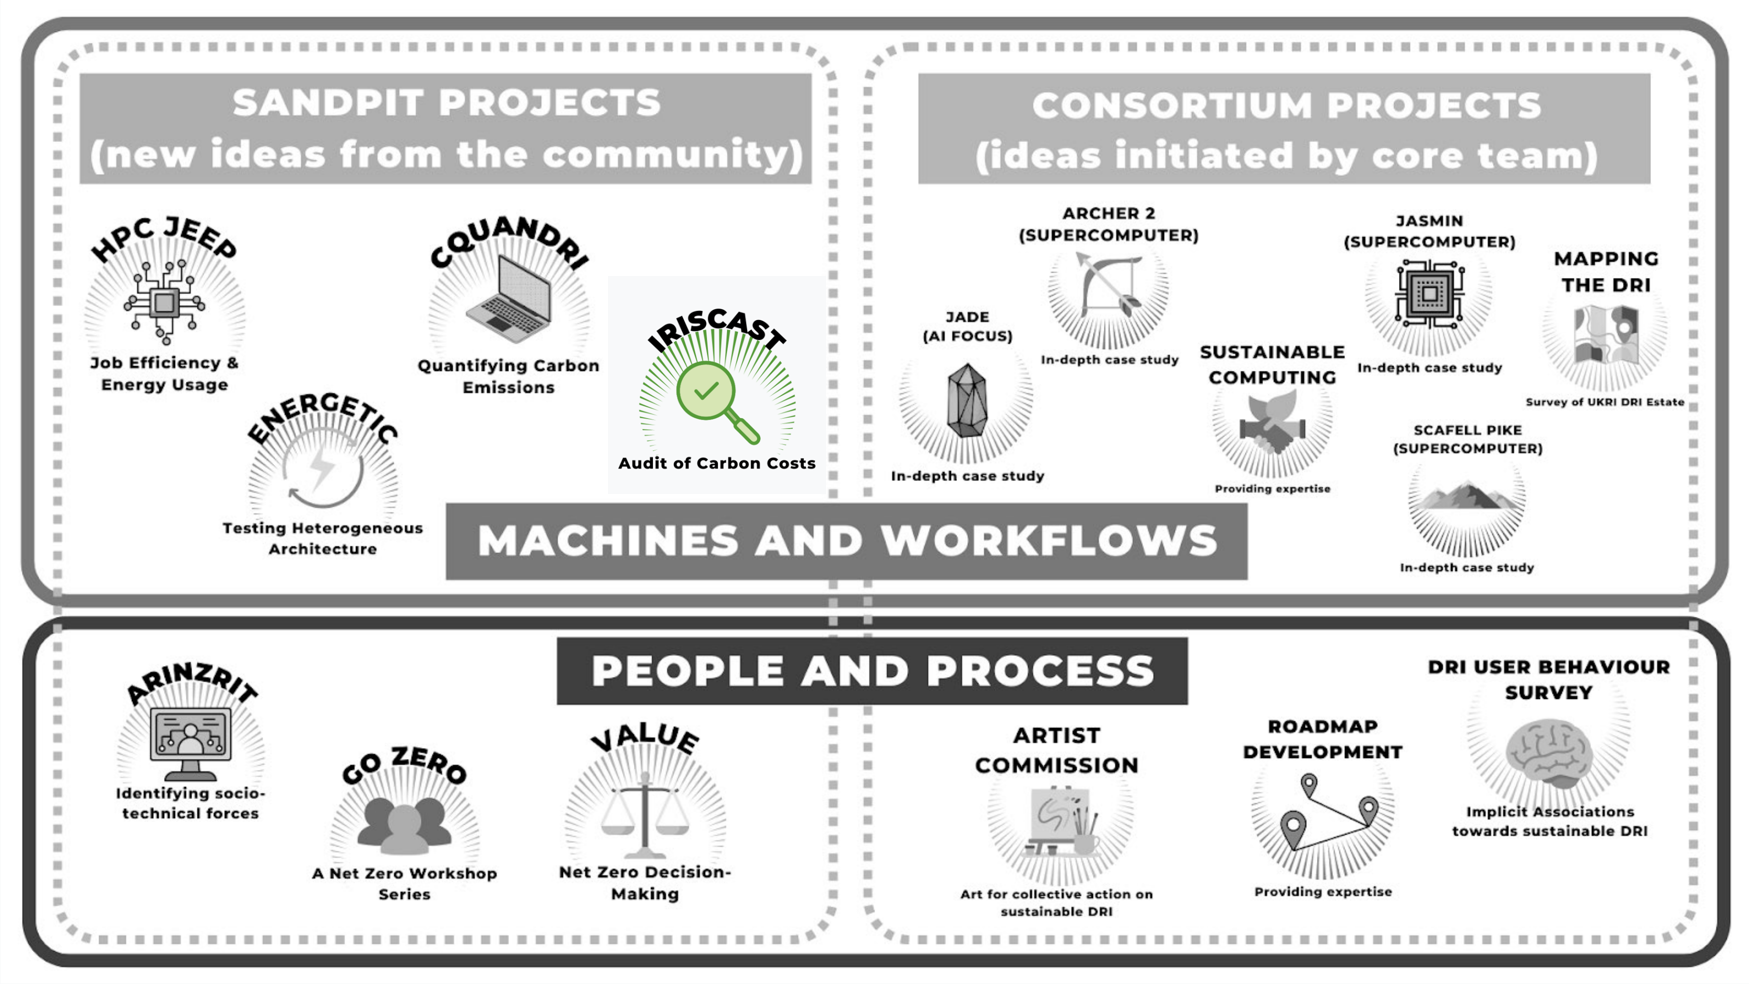

https://net-zero-dri.ceda.ac.uk/

Good robust decisions need good robust information

Speaker: Dr Alex Owen

Queen Mary
University of London

Queen Mary University of London
UNIVERSITY OF CAMBRIDGE
epcc

Measuring Carbon

NetZero and the IRISCAST Project

Audit of Carbon Costs

Jon Hays (QMUL)
Nic Walton (Cambridge)
Adrian Jackson (Edinburgh)
Alison Packer (STFC)

Alex Owen (QMUL)
Alex Ogden (Cambridge)
Anish Mudaraddi (STFC)

Dan Traynor (QMUL)
Derek Ross (STFC)
Alexander Dibbo (STFC)
Jon Roddom (STFC)
Martin Summers (STFC)
Jacob Ward (STFC)
Dan Whitehouse (Imperial)
Alastair Basden (Durham)

UKRI Science and Technology Facilities Council | Scientific Computing

6 Month Project Funded within UKRI Net Zero Scoping Project

UKRI DRI
Net Zero
by 2040

DiRAC

iris

GridPP

archer2

IRISCAST
Audit of Carbon Costs

DIGITAL

RESEARCH

INFRASTRUCTURE

Heterogeneous

Evolved

Distributed

Self-Organising

eInfrastructure for Research and Innovation for STFC

IRIS is a cooperative community brining together (mainly) STFC computing interests

Formed bottom up by science communities and compute providers

Works closely with STFC but run by the community

IRIS Science Director is Prof J Hays who is also IRISCAST Project PI

https://www.iris.ac.uk/

What is IRISCAST?

IRISCAST is the IRIS Carbon Audit SnapshoT

Audit of Carbon Costs

24 Hour snapshot across multiple 'IRIS Facilities'

The Challenge

Estimate carbon costs for scientific computing across a broad heterogeneous landscape

Identifying the key drivers for carbon costs

Identifying the hurdles and barriers

Communicating the carbon costs to drive change

Working coherently across different communities

The Project

Work together coherently across different facilities with different remits, tooling, and capabilities.

Learn by doing!

Document the gaps, the barriers and the issues, drive requirements for future work and decision making

Communicate across our communities and build a foundation for future action

Carbon Model

$$C_t^p = C_a^p + C_e^p$$

Carbon Cost (C) for a period (p) is sum of active carbon (a) and embedded carbon (e)

$$C_a^p = CM_e^p \left(E_{nodes}^p + E_{network}^p + E_{cooling}^p + E_{power}^p + E_{facility}^p \right)$$

Carbon Intensity of Power (Grid)

Measure Energy Usage (E)
Obtain Carbon Grid Intensity (CM)

$$C_e^p = \sum_{1}^{\text{nodes}} \sum_{t=0}^{p} \frac{C_{\text{enode}}}{L_{\text{node}}} + \sum_{1}^{\text{networks}} \sum_{t=0}^{p} \frac{C_{\text{enetwork}}}{L_{\text{network}}} + \sum_{1}^{\text{facility items}} \sum_{t=0}^{p} \frac{C_{\text{efacility}}}{L_{\text{facility}}}$$

Inventory of Equipment
Obtain Embedded Carbon
Estimate Lifetime (L)

IRISCAST
Audit of Carbon Costs

24 Hour Snapshot at Six Facilities

Learning By Doing

Summary Inventories

Node Model	Quantity
Dell PowerEdge R640	118
Mellanox SN2410	4
APC APDU9953	12

Facility Inventory at QMUL

	Specification		
Model	CPU	RAM	Quantity
Dell PowerEdge C6320	x2 Intel Xeon CPU E5-2690 v4 @ 2.60GHz	256 GB DDR4-2400MHz ECC	60

Facility Inventory at Cambridge

Node Model	Quantity
PowerEdge R410	68
PowerEdge R430	60
PowerEdge R440	15
PowerEdge R6525	30
ProLiant SL2x170z G6	24
SYS-6028TP-HTR	12
X9DRT	24
Unknown (Generic Server)	8

Facility Inventory at Imperial

	Node Specification		
Model	CPU	RAM	Quantity
Dell PowerEdge C6420	x2 Intel Xeon Gold 5120 CPUs	512 GB	452
Dell PowerEdge C6525	x2 AMD EPYC 7H12	1024 GB	360

Facility Inventory at Durham

Node Specification		
CPU	RAM	Quantity
AMD Epyc 7502	256GB	246
Intel Gold 6126	192GB	164
Intel E5-2650v4	128GB	201
Intel E5-2650v3	128GB	88
Network Switches		-

Facility Inventory at STFC SCARF

	Node Specification		
Model	CPU	RAM	Quantity
Dell C6420	Intel Xeon 4108	96GB	96
Dell C6525	AMD Epyc7452	512	138
Supermicro	AMD Epyc7452	512	238
Supermicro	Intel 6130	384GB	74
Dell	Various	Various	10
GPU Nodes	Various	Various	94
FPGA Node	Intel 6148	192GB	1
Control Plane	Various	Various	12
Storage Nodes	Various	Various	105
Network Switches	Various	Various	-

Facility Inventory at STFC CLOUD

Facilities

Cambridge IRIS HPC/Cloud
STFC SCD Cloud
STFC SCARF
QMUL GridPP Tier 2
Imperial GridPP Tier 2
DiRAC (Durham)

Build a Community

IRISCAST
Audit of Carbon Costs

Energy Vs Power

For active carbon we need to know about ENERGY usage

IRISCAST
Audit of Carbon Costs

Time Stamped Energy usage is more robust than instantaneous power

APC PDU's SNMP query	✓	`PowerNet-MIB::rPDU2DeviceStatusEnergy.1`	**Hoctowatthour (100 wh) units**
APC port n SNMP query	✓	`PowerNet-MIB::rPDU2OutletMeteredStatusEnergy.n`	
IPMI query (freeipmi)	✓	`ipmi-oem dell get-power-consumption-data`	
IPMI query (ipmitool)	?	`ipmitool sensor list`	**Probably Power Measurements Energy is better!**

PDU Vs IPMI

IPMI approx 20% low cf PDU
Except at QMUL where 1.5%
(APDU9953)

More questions than answers:
Check your calibrations!

System	Facility (kWh)	PDU (kWh)	IPMI (kWh)	TurboStat (kWh)	No. of Nodes
QMUL	1299	1299	1279	1214	118
CAM	261	-	261	-	59
DUR	8154	8154	6267	-	876
STFC CLOUD	3831	-	3831	-	721
STFC SCARF	4271	4271	3292	-	571
IMP	944	-	944	-	117
Total	18760				

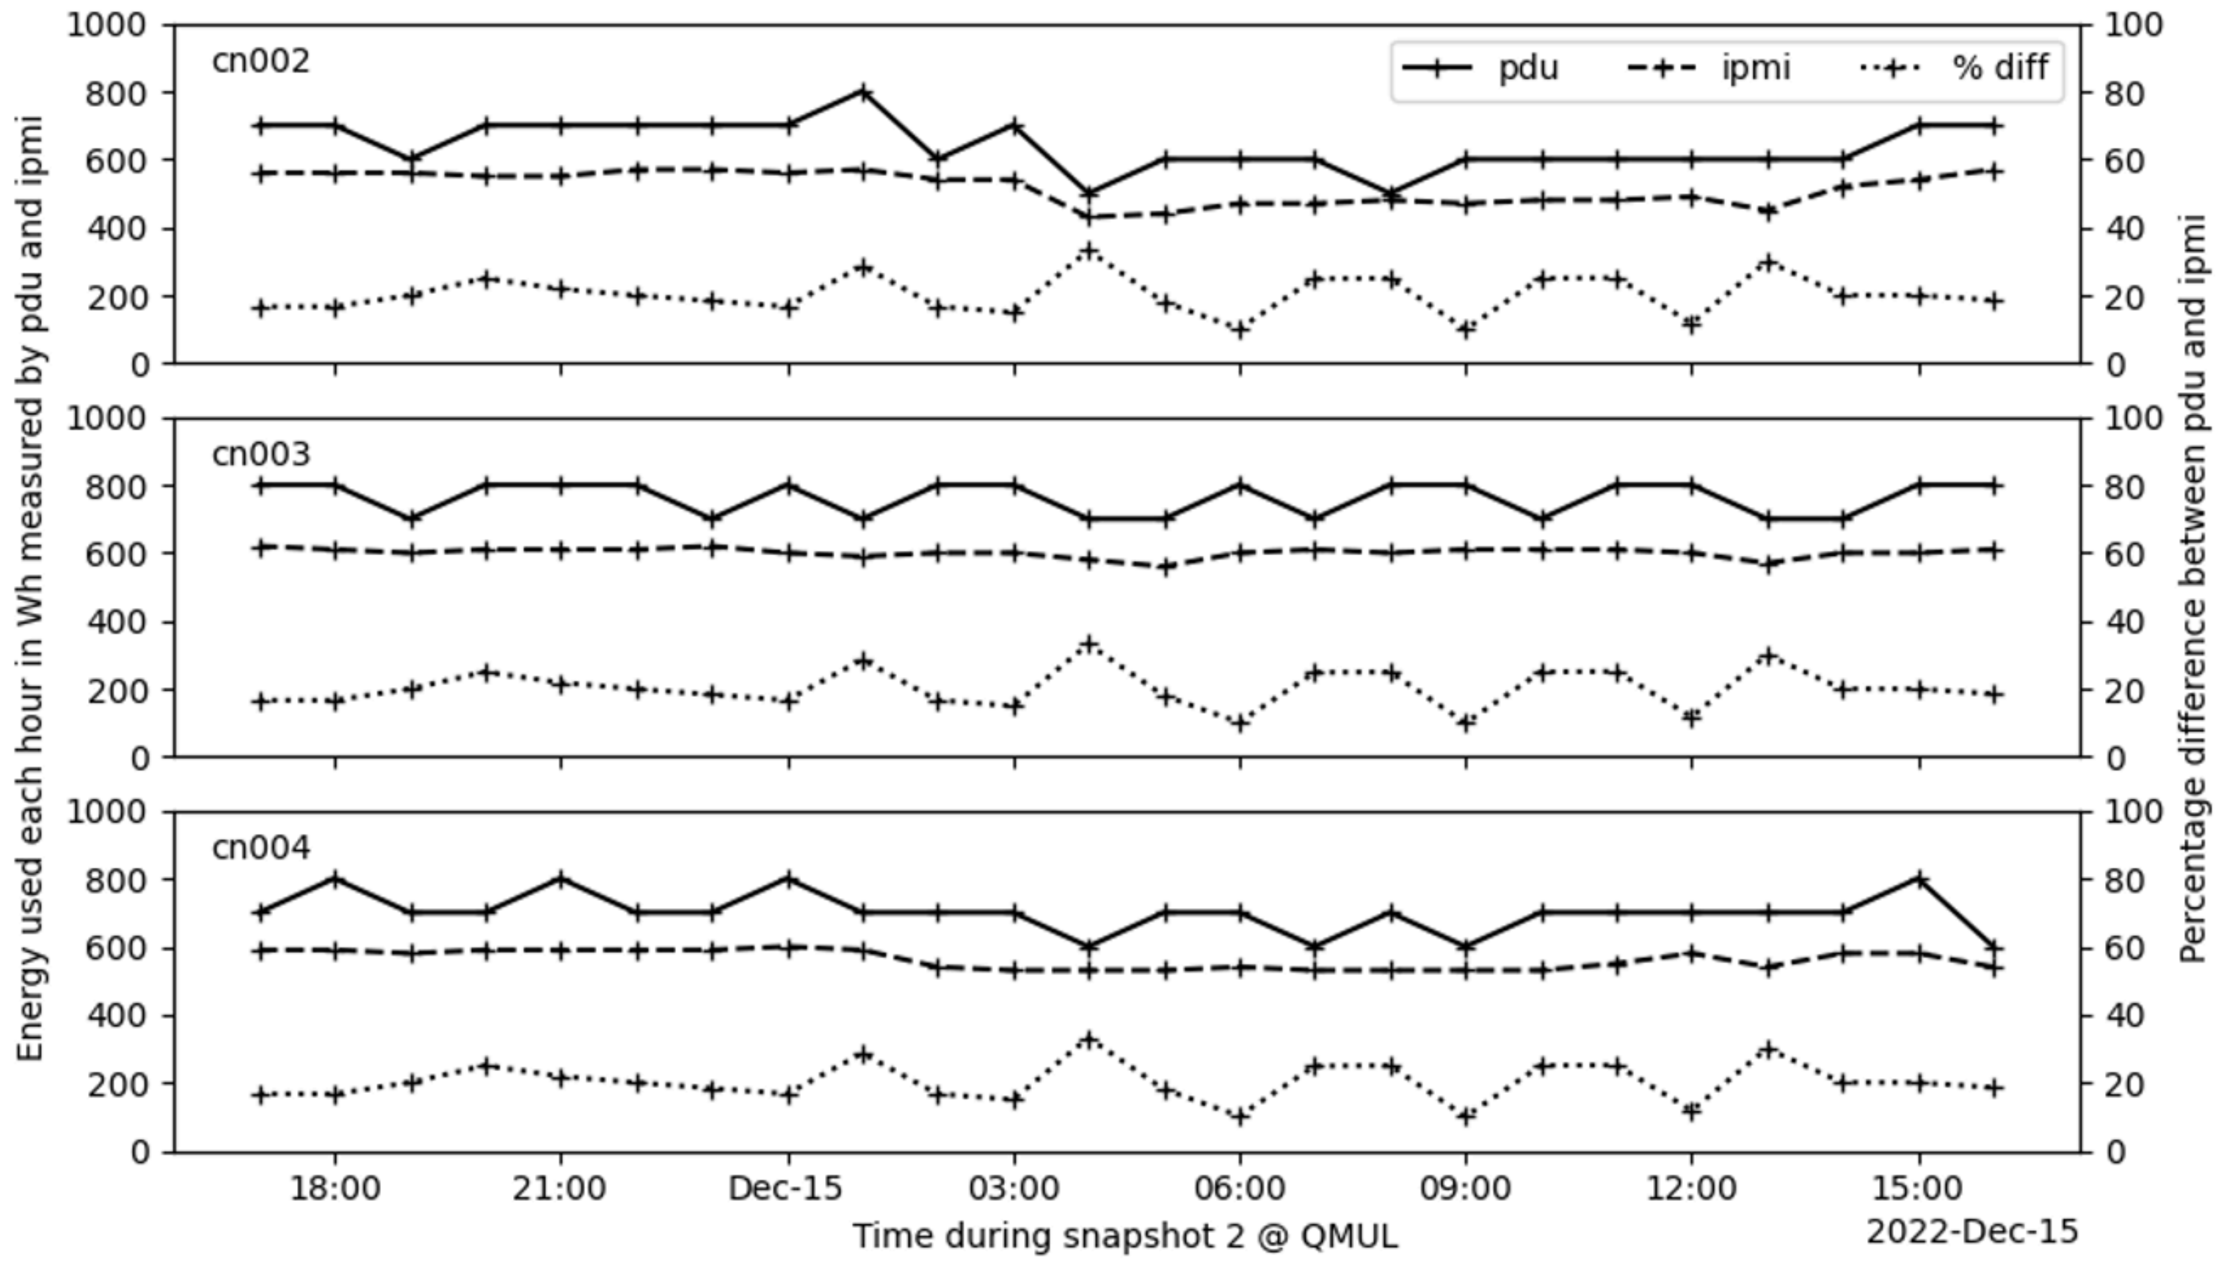

20% difference between APC AP8459WW port measurements and IPMI measurements

Re test at QMUL
with AP8459WW
per port PDU
20% difference!

IRISCAST
Audit of Carbon Costs

Tools at each Level

Facility Level: Cooling energy usage/PUE generally poorly known

Facility	Enclosure Level			Node Level		
Name	Device	Protocol	Tool	Device	Protocol	Tool
QMUL	PDU	SNMP	Net-SNMP	BMC	IPMI	free-ipmi
Cambridge	-	-	-	BMC	Redfish	Prometheus
Durham DiRAC	PDU	SSH	SSH	BMC	IPMI	unknown
STFC Cloud	PDU	SNMP	LibreNMS	BMC	IPMI	ipmitool
STFC SCARF	PDU	SNMP	LibreNMS	BMC	IPMI	ipmitool
Imperial	-	-	-	BMC	IPMI	ipmitool

Comparison of predominant data collection methods at IRISCAST sites. Notably using intelligent Power Distribution Units (PDUs) and Baseboard Management Controllers (BMCs).

Precision of ipmitool output identified as a problem by Durham. They may have a patch!

Job Level: SLURM queues can report Job Energy usage

Turbostat and other RAPL tools?

IRISCAST
Audit of Carbon Costs

Carbon Model

Model a range of scenarios

Measuring computer energy usage is the easy bit.

Cooling energy usage/PUE less well known.

Computer embedded carbon figures hard to find.

Other equipment embedded carbon figures even more hard to find.

IRISCAST

Audit of Carbon Costs

Factor	Scenario Low	Scenario Medium	Scenario High
Carbon Intensity (gCO_2/kWh)	50	175	300
PUE	1.1	1.3	1.6
Server Embodied Carbon ($KgCO_2$)	400	-	1100
Server Lifespan (years)	3	5	7

		Total carbon footprint estimate ($kgCO_2$) (Percentage active carbon)		
Server embodied carbon	**Server lifespan**	**PUE Low Carbon Intensity Low**	**PUE Medium Carbon Intensity Medium**	**PUE High Carbon Intensity High**
Low	**3**	1950 (55%)	5293 (83%)	10186 (91%)
	5	1600 (67%)	4943 (89%)	9836 (95%)
	7	**1449** (74%)	4792 (92%)	9685 **(96%)**
High	**3**	3483 **(31%)**	6826 (65%)	**11719** (79%)
	5	2519 (42%)	5862 (75%)	10755 (86%)
	7	2106 (51%)	5449 (81%)	10342 (90%)

IRISCAST 24 hour snapshot roughly 1-4 people on 12 hour return Jet

Potential to reduce carbon emissions by an order of magnitude!

IRISCAST Proposes

High Level Feedback

Carbon Equivalent per month

Low Level Feedback

Figure of merit per Job → Carbon Equivalent per Job

Audit of Carbon Costs

1. Future DRI procurement to include a score based on embedded carbon costs and equipment energy usage.
2. New computer hardware to include energy measurement capability such as IPMI (or per port PDUs) and require the supplier to provide best estimates of embedded carbon costs.
3. Measure energy used by cooling infrastructure and the computing infrastructure.
4. Facilities to keep an inventory of equipment including embedded carbon cost and idle power draw.
5. Monthly (or other periodic) reporting of carbon usage by facilities based on 3 and 4 above. Roll into standard grant reporting regime.
6. Collect per job (or VM) energy usage by using tools like Slurm (correctly configured). Combined this with embedded carbon from inventory and electricity carbon intensity to feedback job carbon cost to the end user to drive improvements in user code and workflow.
7. Identify user communities and the authors of community codebases so that useful feedback can be given to them to drive the development of more efficient code and workflows.

UKRI DRI
Net Zero
by 2040

Measuring Carbon

NetZero and the IRISCAST Project

Audit of Carbon Costs

Good robust decisions need good robust information

Speaker: Dr Alex Owen

Queen Mary
University of London