NetDRIVE Workshops

Queen Mary University London

February 18/19, 2025

Martin Juckes, Sarah Sparrow, Alex Owen, Graeme Smith

Contents

Meeting page: https://indico.ph.qmul.ac.uk/e/NetDRIVE-2025-02

Cite as: Juckes, M., Sparrow, S., Smith, G., & Owen, A. (2025). NetDRIVE Workshops Queen Mary University London February 18/19, 2025. Zenodo. https://doi.org/10.5281/zenodo.14999697

Overview

On 18 -19 February 2025, a meeting was held in London entitled NetDRIVE: Priorities and Ethics of Research Computing Investments.

Thie meeting marked the initiation of NetDRIVE Stage 2 and aimed to address the following questions:

- How can we maintain competitive computing capability when we have 10-fold increase per decade in the price of a competitive machine?
- How do we manage ethical issues around a national compute resource together with associated data?
- How do we develop the new skills needed both for rapidly advancing technology and evolving sustainability targets?

The meeting brought together users and stakeholders from across the digital research community and featured plenary presentations, as well as breakout sessions that enabled attendees to form groups to formulate ideas to address questions above. Remote attendees could also join the plenary sessions and formed an online breakout group to address the questions. The full briefing document on the three questions is preserved at appendix 1.

The sessions resulted in several ideas and actions for NetDRIVE as it builds its vision for a sustainable future, as well as recommendations for UK Research and Innovation (UKRI), UK universities and the research computing community in general.

The meeting featured presentations from Professor Jon Hays, Head of the Particle Physics Research Centre at Queen Mary University of London; and from GreenDIGIT, a project set up to address the problem of reducing the environmental impact of Research Infrastructures from the generic conceptual point of view by addressing all major factors that define and influence the RI environmental and climate impact. There was also the announcement of a funding call for community activities.

The meeting was hosted by Queen Mary University of London at the historic Barts Pathology Museum. The museum was built in 1879 for the purpose of housing anatomical specimens to be used by those studying medicine and is still in use by QMUL's medical students today.

The Network for sustainable Digital Research Infrastructure Vision and Expertise (NetDRIVE) is developing a forum for managers, software engineers, academics and others to come together to build a common vision for a sustainable future and to incite a transition to sustainable working practices in the Digital Research Infrastructure communities by 2040 or sooner.

Workshop Outputs

These recommendations flow from the meeting. In some cases, there may need to be further consultation and refinement before implementation; delegates at this initial workshop did not have the time to explore practical details of all suggestions.

Quantitative standards which define measurable properties which can be reported with a value and an uncertainty are separated from qualitative standards which might describe a process or workflow. The former fall primarily in the technical themes and the latter in the behavioural themes.

1. Training, Human Capacity, and Skills

a. Mandate training element for PhDs to cover efficient use of DRI/green computing/optimisation. Consolidate training (i.e. set aside funding for training). Efficiency should be part of basic coding training, especially in all science courses where computing is or is becoming important.
b. Identify and fund centres of excellence around technologies; improve accessibility of platforms, codes, and support; invest in looking after code; targeted Research Software Engineer support.
c. [WG ethics] Help researchers to explore different techniques of getting to desired solution (implement through sustainability element of proposals).
d. Compute providers to identify sustainability expert in each team and consider making it a distinct role.
e. Universities to embed sustainability skills into courses and seek accreditation; sustainability skills, such as software efficiency, will be needed for research and industry.
f. NetDRIVE to survey technical career pathways with a view to re-implementing best practice in research computing career pathways.
g. Encourage transfer of knowledge while we have known sources (before they leave for private sector).

2. Measures, Metrics, and Quantitative Standards

a. Funding bodies should take measures of carbon/environmental impact into account at the stage of funding decisions; clear evaluation of CO2 footprint in proposals and business cases; engage with and implement guidelines on metrics being developed in NetDRIVE. Standardised metrics gathering, accounting and feedback mechanisms should be implemented at UKRI funded facilities and services. Estimates based on upper likely bounds should be used for services which cannot provide evidence-based information. Ensure transparency in carbon accounting.

b. Consult and learn from (i) French National Research Agency (ANR)1 experience and their dept for ecological transitions (Lannelongue , 2025); (ii) Welcome Trust Environmental Sustainability Funding Policy2; (iii) Green Software foundation3 for individuals to develop a standard approach to assessing carbon footprint of proposed work of planned infrastructure.
c. [Software builders] To ensure users have information about footprint. Large codes can be inefficient especially with inexperienced users.
d. Develop metrics to demonstrate sustainable use of resources.
e. [WG metrics] Information on units and process of metrics. Transparent explanation of processes of estimation and boundaries.
f. [snapshot] NetDRIVE to commission survey of impact tracking for research benefits and best practice in carbon impact tracking (especially to look at the benefits of research).
    - NetDRIVE to use above to help conduct some carbon impact studies so we have some 'good news' stories.
g. NetDRIVE to formulate guidance on metrics and the tools available to gather them (with regards CO_2).
h. ACTION: Providers to measure cooling energy usage as well as compute/storage energy usage.
i. NetDRIVE to seek out training best practice to help train the trainers and then train users.
j. ACTION: NetDRIVE to identify barriers to measuring carbon footprint

3. Assessing, Testing, Qualitative Standards

a. UKRI to evaluate how effective feedback mechanisms are at changing behaviour. Qualitative methods such as asking PI's how they feel feedback has affected their behaviour/approach to carbon footprint might be more useful than a quantitative approach.
b. [CA snapshot] Collect evidence on impact of RSE teams. What are the cultural issues stopping RSEs from looking at energy efficiency? What determines balance of functionality vs. Efficiency?
c. [CA best practise] Document existing green computing courses (e.g. Manchester, Oxford, various French Universities).
d. [WG metrics] Improve awareness and methodology of recording baselines and success in shifting the dial (reductions from baseline).
e. Share knowledge of cost-efficient code profilers (given high cost of commercial profilers); Improve access to tools which support efficient software engineering; Access to tools to review code; Support for actioning information that reviews provided so that users can learn lessons from metrics which indicate inefficient code. Require users to explain what they have done in response to code review.

f. Create and adopt standards to peer review how sustainable the compute capability really is.
g. ACTION: UKRI to clearly communicate new sustainability guidelines/policy to users with a clear 'what' and 'why'.
h. Figure out how to "account" for carbon, dealing with different uncertainties and assumptions in different scopes, and with double counting inherent in scope 3 emissions.
    - Carbon accounting system to default to "bad" values where providers do not provide data, to motivate providers to gather real metrics.
    - ACTION: NetDRIVE to commission prototype CO2 accounting and feedback system

4. Advocacy, Communication and Culture Change

a. Encourage multi-disciplinary collaboration, e.g. require funding for RSE to support software/compute resources. Action to review grants on that basis
b. ACTION: UKRI to understand and argue on behalf of research community concerning the need for an advanced DRI to solve research problems; the “DRI” should mean people and software as well as hardware. DRI is not just a supercomputer and research are not measured in exaflops.
c. [best practice activity] Need a place to share good practice and support peer recognition. For example, a new interdisciplinary journal for practical carbon efficiency: can publish as proceedings/look for support from e.g. Copernicus. Perhaps use ArXiv.
d. Leverage national research computing events by offering NetDRIVE sessions at these events.
e. Encourage & support publication ("software corridor" in journals).
f. Demonstrate competitive capability through diversity of resources.
g. A list of research computing resources in the UK: setup a website summarising the systems as a wiki style so it can be edited in the future and make it interactive.

5. Process, Evaluation

a. Government/UKRI to announce clear budgets with timelines.
b. NetDRIVE to use case studies and evaluation of best practice to establish guidance on decommissioning of HPC equipment; full life-cycle plan should be in place when funding is approved.
c. Enable researchers to balance and make decision about what is worth doing; facilitation to link researchers together.

Appendix 1: Workshop Briefing

Breakout Aims

We have three breakout sessions scheduled over two days. The first should be used to collect and clarify ideas. Some views expressed within the group may be contradictory: to aim here is to record these areas of contradictory views rather clearly than trying to resolve them. In the 2nd breakout the main areas of agreement and contention should be identified in order to begin organising the range of views expressed in the first breakout.

The final breakout should be used to set out recommendations and guidance for UKRI and the community.

Each group should seek to address the three questions listed below and assign discussion points to the four project themes (also listed below). Moderators will try to ensure that all themes are covered.

Three Questions

Q1: **How can we maintain competitive computing capability when we have net zero challenges and a 10-fold increase per decade in the price of a top-end machine?**

Context: Performance per dollar has increased approximately 100-fold per decade since 1940; performance of the top machine has gone up 1000-fold per decade. Claiming progress significantly above past performance trends, Cerebras have announced the capability to build a 250-exaflop machine[1], but the (undisclosed) price will probably be in the range of $50billion. This hyperinflation is reflected in the 3-fold price increase of the machine topping the TOP 500 list from $200m for Summit in 2019 to $600m for El Capitan in 2024.

- ★ What is a realistic ambition for the scale of our research computing resource? Should this be expressed in terms of international ranking or scale of budget?

Q2: **How do we manage ethical issues around culture change and resource usage for our national compute resource and associated data?**

Context: It is widely accepted that net zero targets cannot be achieved in any meaningful way without behaviour change.

- ★ If behaviour change means, for instance, more preparation before gaining access to large compute resources, how should we ensure that this does not distort fairness of access?
- ★ How do we maintain accessibility while enhancing awareness of negative impacts of resource use? Do we need to ration usage or educate users?

Q3: **How do we develop the new skills needed both for rapidly advancing technology and evolving sustainability targets?**

Context: The commercial software sector is expanding at a rate which academic organisations cannot match. For example, Nvidia claims more than 4 million developers working on their systems through 40,000 partner companies[2]. On the other hand, the Independent Review of The Future of Compute: Final report and recommendations (2022) recognises the severe implications for national security of any gaps in our national skill base.

- ★ Is there a substantial difference between the type of new skills needed now and the skills which were new 10 years ago? Do we need a different approach to skills or is it just a matter of expanding and updating existing courses?
- ★ Do we need different approaches to management and different organisational structures?
- ★ What is motivating early career software engineers and computational scientists today?

Four NetDRIVE Themes

Green Software Engineering

This theme covers a wide spectrum of digitised instructions and knowledge systems, whether a code used or developed by a researcher, or the environment within which such codes run (e.g. schedulers on multi-user machines or use of ChatGPT as a tool). It covers the whole lifecycle of development of software, from gathering requirements, through design, implementation and deployment, and the science of doing so in an efficient manner. Traditionally, Software Engineering has focussed on functionality, reliability, extensibility and how well the software product performs in terms of time, hardware resources and people to maintain it. Green software engineering shifts the focus to sustainable use of resources which can be achieved both through efficient programming and an inclusive approach to project management and development.

Target audience: Research software engineers, computational scientists,

Topics: Skills and Career Pathways, data workflows, exploiting/taming artificial intelligence, exploiting novel architectures

Objective: Develop and embed the paradigm of green software engineering in the UKRI research culture.

Systematic Change

"*Sustainability - the critical application of creative thinking and technology to secure our future on this small planet*" (Richard Rogers, 1995)1

Increase in computational efficiency over recent decades have been offset by increases in capacity, resulting in a net increase in the carbon footprint for computation. Such gains in efficiency in computation are essential, but to convert efficiencies into emissions reductions, avoid the rebound effect, and remain aligned with the balanced pathway to net zero, we need a paradigm shift in the evaluation of computation and the use of computing in research. We need continuous assessment and focus on the mission of achieving sustainability as well as active measures to counter the risk of enhanced demand negating efficiency gains. Working with peers and suppliers to understand the mutual benefits of a low carbon supply chain is essential. UKRI needs to support the Climate Change Committee's balanced pathway for achieving national sustainability by delivering high levels of innovation while restraining resource usage.

Target audience: Team and institute leaders, strategic planners, policy makers

Topics: System level view, delivering value, block rebound, culture shift, EDI, RRI, open science and data,

Objective: Ensure that the UKRI DRI achieves a clear leadership role in the net zero transition.

Machine Room and Hardware

Amazing gains in computational power and capacity are being delivered through advances in technology which also drive increases in complexity of hardware and associated supply chains. We need to bring sustainability to the forefront of the management of our infrastructure. Full and authoritative life cycle analysis of infrastructure will be needed, including effective measurement and management of impacts during the use phase, clear contracts and conditionalities to develop a low carbon supply chain, and clear end-of-life planning for hardware.

Target audience: Infrastructure funders, machine room operators, infrastructure policy,

Topics: Procurement, lifecycle analysis, sustainable capacity, HPC, cloud

Objective: Deliver platforms for sustainable research computing.

Galvanising Individual Action

Behaviour change is difficult to guide but crucial in the pursuit of DRI Net Zero. All staff (from student to CEO) should be mandated and empowered to take proportionate action to drive change and reduce the environmental impact of their work. Discussions among colleagues, proven to be effective in fostering positive changes, should be encouraged. work must start now with commitment appropriate to the climate

emergency while recognising the need for adjustments and learning from experience as we seek a pathway to sustainability.

Topics: Empowerment, motivation, creative thinking.

Target audience: UKRI staff using or managing the DRI.

Objective: Build a sense of individual and community engagement and ownership of the transformation

Appendix 2: Venue, Agenda, Presentations

The meeting was held at the Pathology Museum, St. Bartholomew's Hospital.

Day 1

Session 1.1: 13:00 – 14:50

- [20] Welcome (Alex) and Introduction (Martin/Sarah)
- [30] Sustainable Computing (Jonathan Hays)
- [60] Breakout 1 – Sharing ideas and concerns

Session 1.2: 15:10 – 17:00

- Feedback from Breakout Session 1
- [30] GreenDigit (Catalin Condurache)
- [60] Breakout 2 - Consolidating and balancing ideas

Day 2

8:30 Coffee

Session 2.1: 09:00 – 11:00

- [5] Welcome (Alex)
- [25] Feedback from Breakout Session 2
- [30] Flexible time
- [60] Breakout 3 - Recommendations and Conclusions

Session 2.2: 11:20 – 13:05

- [50] Feedback from Breakout Session 3 and Plenary Discussion
- [50] NetDRIVE Funding Opportunity
- [5] Close

Appendix 3: Delegates

- Tahmina Ajmal University of Bedfordshire
- Alastair Basden DiRAC / Durham University

- Anjali Bhatt STFC (UKRI)/IRIS
- Sarah Campbell NERC
- Fatima Chami CEDA STFC
- Catalin Condurache EGI Foundation
- Robert Currie University of Edinburgh
- Tom Deakin University of Bristol
- Dawn Geatches Innovate UK Business Connect
- Sukhpal Singh Gill Queen Mary University of London
- Tom Griffin UKRI STFC - Scientific Computing
- Jonathan Hays Queen Mary
- Jessica Huntley STFC
- Adrian Jackson EPCC, The University of Edinburgh
- Philip Jackson Science and Technology Facilities Council
- Martin Juckes IRIS IAM
- Marcus Keil UCL
- Loïc Lannelongue University of Cambridge, UK
- Daohai Li Queen Mary
- Agnieszka Marczuk-Dupim
- Laurents Marker NCAS
- Laura McGuire Edge Hill University
- Sara Moverley JISC
- Alex Owen Queen Mary
- Kirsty Pringle University of Edinburgh
- James Richings University of Edinburgh
- Michael Rudgyard Alces Flight
- Chiranjibee Sangroula Ulster Business School Alumni
- Andrew Sansum IRIS
- Graeme Smith University of Oxford
- Lorna Smith EPCC, The University of Edinburgh
- Sarah Sparrow University of Oxford
- Harry Swift STFC
- Joseph Thacker STFC-UKRI
- Ilian Todorov UKRI STFC
- Sam Tygier STFC
- Werner van Niekerk UCL
- Michele Weiland EPCC
- Dan Whitehouse Imperial College
- Mark Wilkinson STFC DiRAC HPC Facility
- Jeremy Yates STFC DiRAC

- Jo Zhong LSBU/UCL