Sustainability in Digital Research Infrastructure

UKRI Net Zero DRI Scoping Project final technical report

Updated 1 August 2023

Author list:

Martin Juckes[1], Michael Bane[2], Jennifer Bulpett[1], Katie Cartmell[1], Miranda MacFarlane[3],

Molly MacRae[1], Alex Owen[4], Charlotte Pascoe[1], Poppy Townsend[1].

1. Science and Technology Facilities Council
2. Manchester Metropolitan University
3. King's College London
4. Queen Mary University of London

If you have any questions about how to use the report, you can contact the CEDA helpdesk: support@ceda.ac.uk.

Preface

This report provides a full account of evidence collected in the UKRI Digital Research Infrastructure (DRI) Net Zero Scoping project. It combines detailed technical analysis with a literature survey and results from community and stakeholder engagement activities.

The target of achieving net zero emissions by 2040 is extremely challenging, but the project has encountered a strong determination to meet that challenge in all areas of the DRI and the research community which depends on it.

This technical report, which provides a detailed account of outcomes and recommendations from project activities, will be complemented by an overview document[1] which will provide the key conclusions and recommendations in a more accessible narrative.

Martin Juckes, June 2023.

Acknowledgements

We gratefully acknowledge funding from the UKRI Digital Research Programme (NE/W00713). We also acknowledge support and feedback from the Steering Committee, the Scientific Advisory Board, the Sandpit Review Panel and the DRI Mapping Expert Advisory Group, all of whom provided valuable time and expertise to guide the project.

Special thanks go to our project partners, this report would not have been possible without your hard work and enthusiasm. Thank you to anyone else who contributed to our project in any way. There were many operational, communications and events staff behind the scenes who greatly helped us with tasks such as; sharing and advertising our events, attending or contributing to events/meetings/creative workshops, providing general feedback and advice, and so on. A full list is given in Appendix 4.

Reference this report: Juckes et al., 2023: Sustainability in Digital Research Infrastructure UKRI Net Zero DRI Scoping Project final technical report. https://doi.org/10.5281/zenodo.8199984

Centre for Environmental Data Analysis
RAL Space
Harwell Campus
Email: support@ceda.ac.uk

This work is funded by the UKRI Digital Research Programme on grant NERC (NE/W007134/1). The project website is net-zero-dri.ceda.ac.uk/.

[1] The overview document will be available at this location: https://doi.org/10.5281/zenodo.8203117
This is a reserve DOI, it has not been published yet (as of 1 August 2023) but it will be available soon.

Executive Summary

The United Nations Secretary General, Antonio Guterres, recently warned of the “climate time bomb” and urged developed countries reach net zero emissions by 2040 as part of “a survival guide for humanity” (and indeed all life on the planet)[2]. These are strong words and strong action is required immediately. Funded by the UK Government, UK Research and Innovation (UKRI) has a duty to ensure the emissions it is responsible for are net zero by 2040 or preferably much sooner. At the time when the UKRI target of 2040 was set it was intended to demonstrate leadership by being a decade in advance of the global target and may need revision, so the emphasis is increasingly on “sooner” than 2040.

UKRI is committed to creating a state-of-the art national digital research infrastructure (DRI) that will seamlessly connect researchers, policy makers and innovators to the computers, data, tools, techniques and skills that underpin the most ambitious and creative research. UKRI and the UKRI DRI are well positioned to play a leading role in the national and global transition to a sustainable society through their ability to draw on expertise across the full range of the technical, social, and economic challenges that we face. The UKRI DRI can provide leadership firstly by demonstrating and delivering on a commitment to achieving net zero and secondly as an agent of change driving the technical innovation and behaviour change which is needed to pave the way to net zero.

Escaping carbon dependency will require a reassessment of existing habits, values, and behaviours across the research community and wider society. UKRI is committed to becoming net zero by 2040 or sooner. The key areas contributing to DRI emissions are the purchase of electricity, the manufacture of hardware (including the footprint of extracting raw materials), emissions associated with the use of the UKRI DRI, such as the footprint of laptops, and, last but not least, the impact of the UKRI DRI on societal emissions through research and innovation outcomes. Following best practice, we do not treat different categories as equivalent but do include targets for all major contributions.

Delivery of a net zero DRI will require action on both operational emissions, dominated by the power supply, and the hardware supply chain. With the nation facing constraints on national power generation capacity in the coming decades, the high electricity demand of the exascale machines demanded by the national research programme will need to be well justified not only in terms of scientific creativity but also in terms of delivering certifiable impact accelerating the transition to a sustainable society.

To support net zero targets, and given the need and value of DRI, UKRI awarded £1.86 million to the Net Zero DRI Scoping project (hereafter "the Project") to investigate how UKRI can achieve net zero digital research infrastructure. The Project has collected a comprehensive range of evidence based on a literature survey, expert-led investigations across a range of focused projects, and stakeholder engagement. More than 20 institutions and over 40 researchers contributed to this work.This technical report provides a detailed assessment of the evidence and recommendations formed by the Project and sets out the key initial steps that we believe will allow the UKRI DRI to play a leading role in navigating an equitable transition to a sustainable future.

Technical evidence is collected into a **toolkit**, providing a subject-based synthesis which will be useful to practitioners seeking guidance on specific topics, and a **roadmap** organised in three delivery pathways. The toolkit and roadmap are complemented by a journey metaphor

---

2
https://www.reuters.com/business/environment/un-chief-urges-faster-shift-net-zero-after-report-highlights-climate-threat-2023-03-20/

which captures the need for a community-wide paradigm shift and a baseline which quantifies our point of departure.

In **the toolkit** evidence-backed specific recommendations are gathered together under six strategic themes:

1. **Mission Focus**: continuous assessment and focus on the mission of achieving sustainability; active measures to counter the risk of enhanced demand negating efficiency gains.
2. **Recognition of shared responsibility**: mandate and empower all staff (from student to CEO) to take proportionate action to drive change and reduce the environmental impact of their work; community building; encourage discussion among colleagues and learn from others to foster positive changes in behaviour.
3. **Action-based-research**: work must start now with commitment appropriate to the climate emergency while recognising that there will be a need for regular checks and adjustments; focus on progress not perfection; small steps; learn from experience.
4. **Work with peers and suppliers**: through contracts, conditionalities, and understanding mutual benefits, to develop a low carbon supply chain [essential in the longer term]
5. **Build and Share Knowledge**: providing leadership, support and advice for business cases and large procurements feeding into reporting; central hub for information and institutional knowledge [also likely to create short term results]
6. **Green Software Engineering**: creating a body of expertise around green software engineering, providing training, developing tools, metrics, expert assessment, and standards to transform current approaches to writing code, and supporting codes running in data centres, such that GSE becomes the norm rather than an optional extra.

The toolkit sets out the "what" can be done, whereas the roadmap is setting out "how" the UKRI DRI can implement the recommendations by 2040 or sooner. **The roadmap** is organised into three delivery pathways:

- **UKRI policy and governance**, Creating a policy framework which can deliver the steps needed to achieve the net zero ambition.
- **Delivery partnership**, in which funders and facility leads or service providers work together to implement the fundamental changes required.
- **Competitive funding**, which develops the necessary capabilities and tools drawing on the creativity, diversity and strength in depth of the UK academic community.

The journey metaphor captures the sense of transition, paradigm shift, and positive aspirations for a sustainable future. During the Project there has been a very positive and enthusiastic reception to ideas and working together to inspire action. There is certainly not a single action that brings us to net zero, but the consensus has been that the sooner we collectively start implementing changes (however large or small) to reduce carbon emissions, the more we can help ourselves to meet net zero by 2040, or sooner.

Understanding the journey is also key to understanding the complex choices needed to justify the expenditure of scarce resources on the research and innovation which is moving society forward to a better future. Driving innovation and change is the core role of UKRI and the UKRI DRI. The Climate Change Committee (CCC) makes it clear that the avoidance (or limitation) of damaging and disruptive adjustments on the journey to net zero is best achieved through a combination of technical innovation and behaviour change. This national picture is reflected in the DRI roadmap proposed in this report.

The scale of challenges presented by the climate crisis can overwhelm us, but it is also creating opportunities. The collective focus created by a shared national and international goal can bring people from many different backgrounds together and generate new inspiration and creativity. It has been a privilege and a joy to see this happening in our Project.

The report defines reference values for a 2002 base year and associated targets for a range of emission categories, along with discussion of the very large uncertainties.

Table of contents

1 Introduction and Context

This report sets out the conclusions of the UKRI Net Zero Digital Research Infrastructure (DRI) Scoping project providing UKRI and their community with an outline roadmap for achieving net zero DRI by 2040 or sooner.

This report aims to provide recommendations for actions required to achieve a net zero UKRI-DRI by 2040 or sooner. It also offers guidance and advice to UKRI DRI stakeholders; many of the actions needed to deliver real change in the carbon footprint are outside the direct control of UKRI, so dealing with the areas of indirect control will be vital.

This report is broken down into the following sections:

- This section provides some introduction and context, with Section 1.3 defining Digital Research Infrastructure (DRI)
- Section 2 contains a synthesis of the recommendations (as illustrated in Figure 1) delivered by the project collected together under:
    - A roadmap with three delivery pathways supported by actions
    - A toolkit with six strategic themes underpinned by 180 evidence based recommendations
- The evidence behind the recommendations is presented in Section 3, with subsections for each of the project activities delivering evidence.

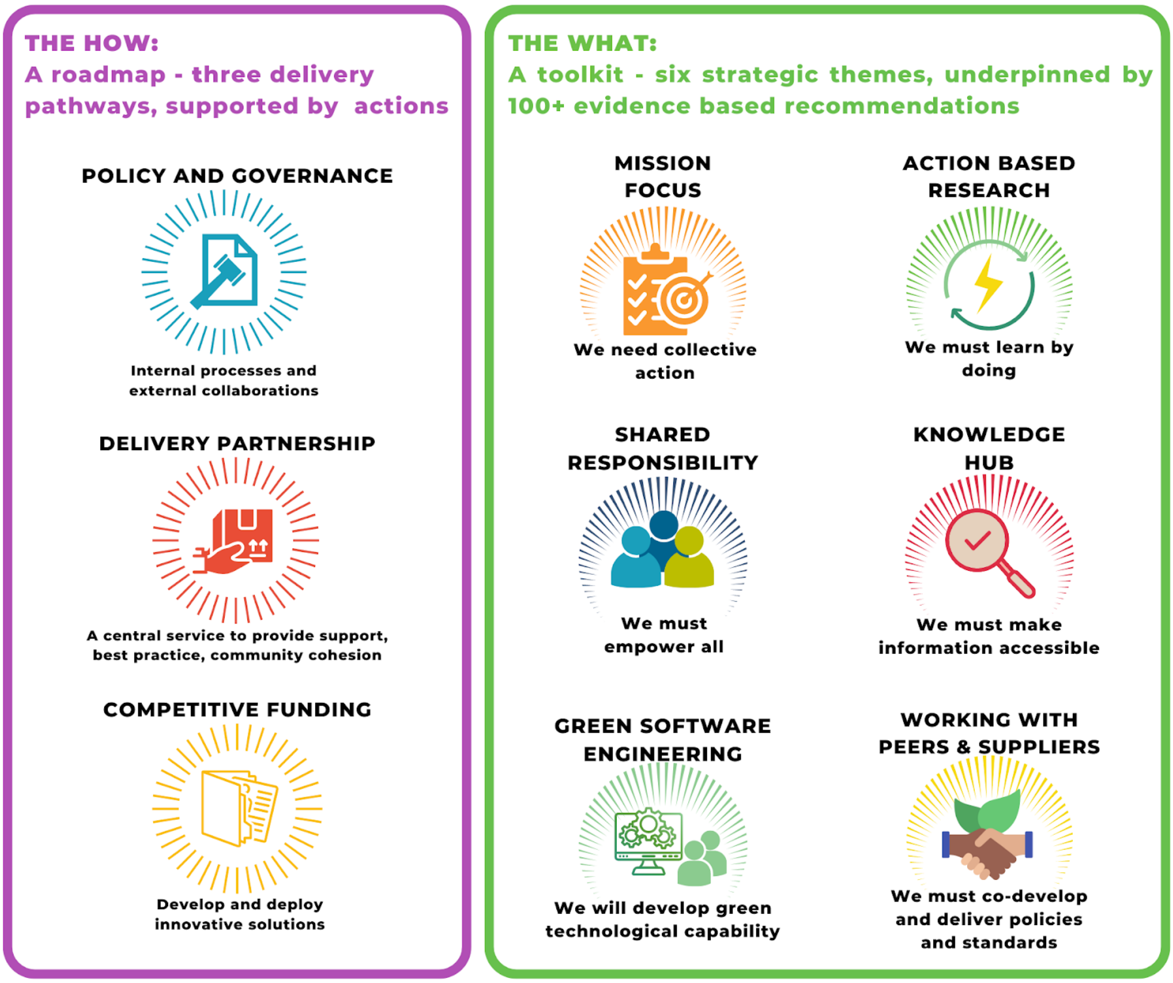

Figure 1: This report synthesises evidence based recommendations produced by technical studies and community consultation. These are organised through a dual approach comprising a toolkit and a roadmap. Full details are in the Synthesis section below.

The figures, tables and boxes in this document are all numbered to relate to the section of document they can be found in - if one section has more than one figure/table/box, additional letters are added. We have done our best to describe all figures within the text to ensure accessibility, however if you need the information in an alternative format please let us know.

1.1 About the project

The UKRI Net Zero Digital Research Infrastructure (DRI) Scoping project was initiated to provide recommendations which would bring the UKRI Digital Research Infrastructure (UKRI DRI) into alignment with the institutional and national ambition to achieve net zero by 2040 or earlier. This strategic objective is outlined by the UKRI Environmental Sustainability Strategy, published in 2020.

This technical report is supported by a Summary report - which is less technical and provides a high-level overview of this document.

The project has three key objectives to:

- Collect evidence to inform UKRI digital research infrastructure investment decisions.
- Provide recommendations for UKRI and their community with an outline roadmap for achieving carbon neutrality across all digital research infrastructure by 2040 or sooner.
- Enable UKRI to play a positive and leading role in the national and global transition to a sustainable economy.

Key details:

- £1.86 million project funded by UKRI, administered by NERC
- The project has a 19 month timescale running from September 2021 to June 2023
- It has been delivered by the project team based within the Centre for Environmental Data Analysis (CEDA) and the National Centre for Atmospheric Science (NCAS) with ~40 researchers from 20 different institutions collaborating as partners across 16 distinct projects

From a functional perspective, evidence has been gathered through four streams of work (see Box 1.1).

Box 1.1: Evidence gathering activity
1. **Literature review**: a survey of current ideas in academic and technical literature 2. **Sandpit projects**: new ideas proposed by the community. Funded activities via a 'sandpit' process, designed to stimulate research collaborations in service of the Project, took place over three days of workshops and generated a range of community initiated projects to address areas of interest highlighted by the Project team. 3. **Consortium projects**: commissioned activities targeting specific problems and/or challenges identified by the core Project team, and which required expert redress. 4. **Community and consensus building activities**: conducted in parallel with evidence gathering projects to ensure that key stakeholders were active

participants in developing and shaping recommendations.

From a thematic perspective, the project activity has been organised in two broad subject areas:

(1) Machines and Workflows: the hardware and software infrastructure that sits at the centre of the DRI.

(2) People and Processes: the expert staff, the systems and institutions that frame their work, and the scientific user community delivering the UKRI programme of research and innovation.

These two subject areas match the twin pillars of the CCC Balanced Pathway: high innovation and behaviour change. The machines and workflows represent both the leading edge of the innovation which is needed to achieve net zero and a high carbon-intensity activity which needs to be brought into line with the net zero target. People and processes are inextricable from the DRI and are essential to extracting maximum value from physical infrastructure and hardware.

1.2 Situational Context

Net zero presents us with **many opportunities** to lead a transformation to a fairer and more equitable society. There are also **many challenges** associated with the technical barriers to achieving net zero and in realising the organisational transformations required. In addition, defining net zero at the national and institutional level is contentious. Therefore, whilst the importance of figures and metrics to support positive change is not understated, net zero is not only considered to be a figurative destination. We consider net zero to be a societal challenge which requires sustained collaboration and cooperation across sectors to be achieved. In our pursuit of transformational change to meet the net zero challenge we also have an opportunity to establish a fairer, more equitable society. A bold vision for net zero therefore encompasses a just social transformation alongside technical solutions.

1.2.1 National Net Zero target

The UK government's plans for net zero greenhouse gas (GHG) emissions by 2050 came into effect in June 2019, updating existing legislation under the Climate Change Act 2008. The preferable approach outlined by the UK government advisory body, the Climate Change Committee (CCC), is the balanced pathway to net zero[3]. The CCC advocates a leadership-driven approach in which the UK, in line with the Paris Agreement, moves to net zero faster than less developed nations.

For the CCC the concept of net zero is based on Nationally Determined Contributions (NDCs) as defined in the Paris Agreement (in which reporting responsibility is limited to emissions occurring within national boundaries). The CCC, however, extended the scope of their recommendations to include aviation and shipping in recognition of the critical role played by these sectors.

The balanced pathway requires high levels of innovation and moderate behaviour change to reduce emissions without significant economic disruption. Power supply for existing activities

[3] Established under the Climate Change Act 2008, the Climate Change Committee provides advice to the UK and devolved governments, and reports to Parliament on progress towards reducing greenhouse gas emissions, and preparing for and adapting to the impact of climate change.

is level, with uplift in power generation devoted to activities which displace carbon intensive fuels (Climate Change Committee, The UK's path to net zero, 2020).

At the time of writing, the latest report from the CCC on progress in adapting to climate change (March 2023) finds that no sector is adequately prepared to meet present and future climate risks facing the UK across "cities, communities, infrastructure, economy and ecosystems"[4]. Progress towards climate change mitigation and adaptation must be galvanised by collective pursuit of the net zero target, through both behaviour change and innovation:

> "*Without urgent, effective, and equitable mitigation and adaptation actions, climate change increasingly threatens ecosystems, biodiversity, and the livelihoods, health and wellbeing of current and future generations.*"
>
> (IPCC, 2023)[5]

In March 2023, alongside the release of the Synthesis Report of the IPCC Sixth Assessment Report, and in response to its findings, the UN has urged leaders of developed countries to bring net zero targets forward by a decade to 2040[6].

1.2.2 National Landscape: Digital Strategies

The UK Digital Strategy (DCMS, 2022) sets out the national context for the digital sector, with a focus on the positive role that it can play in enabling the transition to net zero. This vision for the sector encompasses the National Data Strategy (2020), which highlights the interdependency of skills, data management and data availability with a digital infrastructure that is able to support efficient, secure and accessible data sharing. That these themes emerge in our own project recommendations is indicative of the blurred boundary between commercial and research interests with respect to digital infrastructures. It follows that net zero computing in research has interdependency with commercial and industry practices. The National AI Strategy (2021) also falls under the digital strategy scope, and outlines a high-growth agenda for AI driven technologies and associated infrastructure. AI technology offers new capabilities for data analysis and problem solving, and is positioned as a key innovation which can be leveraged in service of the net zero target for emissions reductions[7].

The "Independent Review of The Future of Compute: Final report and recommendations" (Department of Science, Innovation and Technology, March 2023)[8] sees a strong digital infrastructure as a necessary element of the national transition to net zero. The report notes that "investing in compute will bring wide-ranging benefits" but that "the UK is falling behind on compute and the government will need to take substantive action if it is to achieve its ambitions". That report and review also sets the scene for a future in which digital infrastructure both forms the basis for research, industry and government activity, and also offers a means of multi-sector and multi-national collaboration.

---

4 https://www.theccc.org.uk/publication/progress-in-adapting-to-climate-change-2023-report-to-parliament/#recommendations-to-government

5 SYNTHESIS REPORT OF THE IPCC SIXTH ASSESSMENT REPORT (AR6) - Summary for Policymakers.

6 UN Secretary-General António Guterres address, 20 March 2023: https://press.un.org/en/2023/sgsm21730.doc.htm

7 p45, National AI Strategy: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1020402/National_AI_Strategy_-_PDF_version.pdf

8 https://www.gov.uk/government/publications/future-of-compute-review/the-future-of-compute-report-of-the-review-of-independent-panel-of-experts

If these strategic ambitions are to be supported and facilitated by the research community, a digital research infrastructure that is fit for purpose must be central to the UK's national ambition for digital growth. This is the ambition outlined by directives for growth and expansion of AI capacity and exascale compute in the Future of Compute report. These are accompanied by a recommendation to "manage the sustainability of building and using compute through planning, procurement and innovation". As a cornerstone within the national landscape, environmentally responsible computation must be considered a key agenda item for UKRI with respect to all aspects of the digital research infrastructure.

1.2.3 UKRI and its digital research infrastructure

The UKRI Digital Research Infrastructure (UKRI-DRI) supports a huge range of research, and occupies a key role within the national landscape. It includes high performance computing to support extreme simulations, multiple mid-range computing facilities supporting local and specialist communities, and institutional machine rooms supporting research departments and teams. Wilkinson et al. (2020) set out a case for a capacity at the high performance end founded on research requirements, discussing key applications in the realms of astronomy, molecular physics, climate, extreme weather, the solid Earth, computational biology, medicine, digital humanities and mathematics. The increasing digitisation of research practice in subject domains not previously associated with the DRI, and its implications for net zero is amongst the findings set out in the evidence generated by this project (Section 3).

UKRI define their DRI to include[9]:

- large scale compute facilities, including high-throughput, high-performance, and cloud computing
- data storage facilities, repositories, stewardship and security
- software and shared code libraries
- mechanisms for access, such as networks and user authentication systems
- people: the users, and the experts who develop and maintain these powerful resources.

This definition includes both large scale compute facilities such as data centres, super computers and clusters (linking multi-site facilities to generate, analyse or otherwise utilise large volumes of data), and small server rooms supporting institutional activities. User devices, such as laptops, are not part of the digital research infrastructure, but are part of the scoping study because they make up a significant element of the Scope 3 emissions of the infrastructure. DRI includes both the hardware, the facilities to host the hardware, the personnel who use and support the hardware and the associated activities of those personnel required in order to use the hardware. The carbon footprint of DRI thus includes Scopes 1, 2 and 3 for each of these, as discussed in Section 1.3.

The operation of digital infrastructure is not carbon neutral, and the carbon emissions associated with its expansion and use are significant and growing (Hopper et al. 2020*; Royal Society 2022*).

The UK currently has 15 computers in the Top500 list, 5 of which are in the UKRI DRI. In the High Performance Conjugate Gradients (HPCG)[10] ranking, which is often considered a better guide for performance of user applications, only 3 in the HPCG ranking list (many sites have not submitted this benchmark yet). At the outset of this project, no overview of the UKRI-DRI facilities existed (see §3.4.3 for details of the DRI Mapping database created in this project).

9
https://www.ukri.org/what-we-offer/creating-world-class-research-and-innovation-infrastructure/digital-research-infrastructure/

10 https://www.hpcg-benchmark.org/

However, even without a detailed inventory of DRI assets, it is evident that the use of DRI is growing. The reasons for the increasingly widespread use of digital tools in research are manifold, and include "... advances in science and increased adoption of machine learning techniques in many research domains, including those not previously considered to be computationally intensive" (Section 3.3.1). This trend superseded the Covid-19 pandemic, and has been driven largely by breakthroughs in information technology, most notably (with respect to the National Digital Strategy and UKRI) advances in AI and Big Data (Kraus et al, 2021).

A slightly different perspective is given by Bichsel (2012), reporting results from an EDUCAUSE survey of US institutions showing that demand for digital research computing is driven in part by a competitive element: "having research computing resources makes an institution more competitive in recruiting and retaining faculty with research computing needs".

National and institutional agendas for digital solutions to a variety of problems may also actively drive greater demand for DRI: BBSRC (2020) shows that the proportion of 'data-rich' submissions to the BBSRC Responsive mode funding calls has risen from just over 20% in 2006 to over 50% in 2018. The NERC Digital Strategy (NERC 2022*) announces an ambition to use "digital technologies to transform our understanding and management of the natural world".

A broader trend, linking technological progress with economic growth, reductions in pollution and improvements in health outcomes in Organisation for Economic Co-operation and Development (OECD) member countries lends support to the notion that technology leads to population level benefit (Nghiem et al. 2021), although these authors do not, unfortunately, include greenhouse gases in their analysis). Although further exploration of this question is beyond the direct scope of this project, it is drawn out to some extent with respect to the extension of values and responsibilities associated with digital research practices and infrastructure development (Sections 3.3.7 and 3.5.3.3).

Existing UKRI strategic investment plans for the DRI support the national ambition to position the UK as a 'Science and Technology superpower' (Department of Science, Innovation and Technology, March 2023). However, without rigorously applying the principle of environmental sustainability to digital infrastructure and associated research practices, there is a risk that the expansion of digital innovation will undermine the ambition for a net zero DRI by 2040.

The Science Based Targets initiative (SBTi), which provides guidance for organisations developing net zero targets, recommends that organisations take steps to address all emissions associated directly or indirectly with their operations. For the UKRI DRI this implies taking responsibility for supply chain emissions of the computers that we purchase.

1.2.4 Stakeholders of the Project and the UKRI DRI

Stakeholders fall into several distinct groups, and include: manufacturers and commercial entities providing hardware, commercial and research software engineers (RSEs), funding bodies, research institutions and higher education institutions (HEIs), policy makers, and researchers.

UKRI interacts with stakeholders directly through policy and governance operations, and the procurement and funding of services, equipment and research. As a major funder of research and innovation, UKRI is also able to leverage its position to influence stakeholders indirectly through leadership and advocacy activity.

1.3 Carbon Context

1.3.1 Understanding Scopes and Responsibilities

In order to capture the carbon emissions associated with the UKRI-DRI consideration is given to all three scopes as outlined by the Greenhouse Gas Protocol (Figure 1.3.A), which include emissions both directly and indirectly controlled by UKRI. We treat the UKRI DRI as a virtual organisation, and consequently consider electricity usage in a UKRI-funded university to be equivalent to electricity usage in a UKRI facility. This provides an aggregate scope for reporting emissions, which goes some way to capturing an actionable view of the overall impact and capacity of UKRI DRI (Sections 1.3.3 and 1.3.5). However, responsibility for emissions across the DRI is often shared, leading to considerable uncertainty in numerical allocation of emissions.

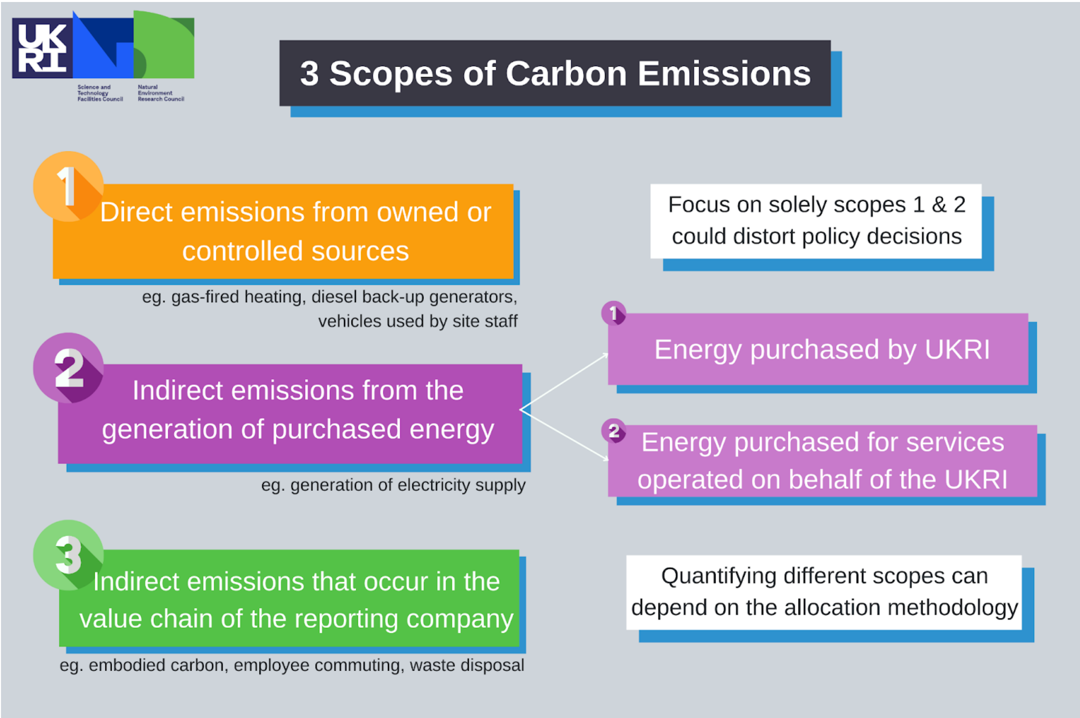

Figure 1.3.A: The 3 scopes of carbon emissions defined by the Greenhouse Gas Protocol form the basis of national and international reporting requirements (GHG protocol, 2004, 2011, 2013, 2019). When carbon emissions cannot be attributed uniquely to specific products of activities the allocation approach can be used, although results depend on methodological choices (e.g. Ernst and Young 2020*).

Our interim report (see Section 3.2) outlined findings which confirm that the largest footprint uncertainties are in the supply chain and in the impact of the UKRI DRI on society[11]. These impacts are felt acutely in locations where land use for mining activity displaces communities and generates high volumes of chemical waste in addition to carbon emissions. Furthermore the energy requirements of large scale DRI, such as data centres, have the potential to displace electricity supply which would otherwise serve much needed housing

[11] The supply chain and the societal impact are two elements of the GHG Protocol Scope 3 emissions. The UKRI DRI supports many projects which have a beneficial effect on societal emissions, but we don't have access to quantitative information or even clear guidance on how such impacts could be quantified.

developments[12]. Both of these impacts are included within the Scope 3 emissions of the Greenhouse Gas Protocol. Treating local emissions as fully equivalent to emissions associated with mining activities at the other end of the supply chain creates many intractable problems[13]. However, ignoring those remote emissions is clearly incompatible with the high ambition of UKRI. UKRI has a leading role in the UK to ensure the net zero ambitions encourage an equitable and inclusive transition that benefits all of society[14]. UKRI must be open and transparent in how it is responsible in ensuring changes are beneficial to all, cover all three scopes, and do not negatively impact others.

1.3.2 International Initiatives and Targets

There is a growing recognition of the urgency of taking action, and a wide range of national and international activities are already making significant progress. The Science Based Targets Initiative (SBTi) is rapidly gaining recognition as a focus for setting consensus and has recently set the first globally recognised corporate net-zero standard. For example, the Alliance for Sustainability Leadership in Education (EAUC) is bringing the SBTi principles into the UK university sector to make sustainability a core component of running a good organisation, and not as an add-on or luxury.

The SBTi standard recommends organisations set targets relative to a base year, with specific targets for all significant scopes and categories.There are strong indications that focus on ad hoc metrics can be counterproductive and we should be keeping focus on the target of achieving net zero carbon balance and going beyond that to stabilise the climate (e.g. Barron et al., 2021).

1.3.3 An overview of the DRI and its carbon footprint

Figure 1.3.3.A shows a comparison of large scale DRI facilities (e.g. supercomputers and data centres) compared to smaller scale devices such as laptops and university servers. These are roughly equivalent to each other - and the carbon estimates for each are approximately equal to an aircraft flight around the entire world 79 times.

The key areas of emissions associated with the DRI include the purchase and generation of electricity (Scope 2, Figure 1.3.A), termed ‘active carbon’, the manufacture of hardware and the mining of minerals to provide raw materials for manufacture (Scope 3, Figure 1.3.A), termed ‘embodied carbon’ (also known as ‘embedded carbon’).

Electricity is the largest single element of the carbon footprint for most UKRI DRI services and facilities, and a number of credible actions to address emissions are outlined in our interim report (though they are not without complexity as mentioned within Interim Report (see Section 3.2).

By contrast, robust estimates of embodied carbon are very hard to arrive at. Some manufacturers provide reports, but the smallprint generally indicates that these are provided to show the relative amplitude of major components of the footprint and should not be used for overall carbon budget calculations. Using the Malmodin et al. (2018) estimate the total

---

[12] An illustrative media article detailing the compromises necessitated by new data centre developments in West London, UK: https://thedeveloper.live/opinion/opinion/energy-shortage-how-data-centres-are-blocking-housebuilding

[13] Responsibility for the supply chain emissions is shared across many stakeholders.

[14] UKRI needs to avoid the negative associations such as those expressed by Action Aid (2020): 'The term “net zero” is used by the world’s biggest polluters and governments as a façade to evade responsibility and disguise their inaction or harmful action on climate change'.

carbon footprint of the ICT manufacturing as 730 Mt CO2-equivalent and a Global ICT market report[15] estimate of sales value at 4,900 billion USD, gives an embodied carbon intensity by sales value (bulk allocation based on global ICT market in 2018) of 0.15 kg / USD (this approach can be referred to as the "economic input-output approach"[16]). Major suppliers of hardware are making significant commitments such as 100% renewable energy supply by 2030 or net zero by 2050[17] (see Sections 1.3.4 and 3.4.5 below for some context around the carbon figures).

The market price associated with services offering carbon dioxide removal from the atmosphere ranges from £5 per tonne for ecosystem storage (e.g. tree planting) to £900 per tonne for deep geological storage. The social cost, estimated at £10,000 per tonne, is even higher (Interim Report, Table 2.5). As might be expected, the quality of the offerings at the lower end of the price scale is reduced, with little or no guarantee that aspirations to remove CO2 from the atmosphere will be met.

The operation of the DRI is also contingent on the actions and movement of people (as outlined by 2.3.1 People and Machines), and therefore emissions associated with travel (flying in particular), usage and maintenance also contribute to its overall GHG emissions footprint. Lastly, the positive impact of the DRI on societal emissions through research outcomes has not been overlooked; it is considered a significant point of tension which requires resolution through the application of shared principles driving progress towards both national net zero targets and a net zero DRI (discussed further in Sections 3.3.7 and 3.5.3.3).

---

15 https://www.researchgate.net/publication/324167731_Information_and_Communication_Technology_ICT_Industry_in_the_Fourth_Industrial_Revolution_Prospects_and_Challenges_for_Workers_in_Asia-Pacific/figures?lo=1

16 Cambridge University methodology statement: https://www.environment.admin.cam.ac.uk/files/uoc_methodology_statement.pdf

17 https://csrreportbuilder.intel.com/pdfbuilder/pdfs/CSR-2021-22-Full-Report.pdf, https://esg.tsmc.com/download/file/TSMC_TCFD_Report_E.pdf

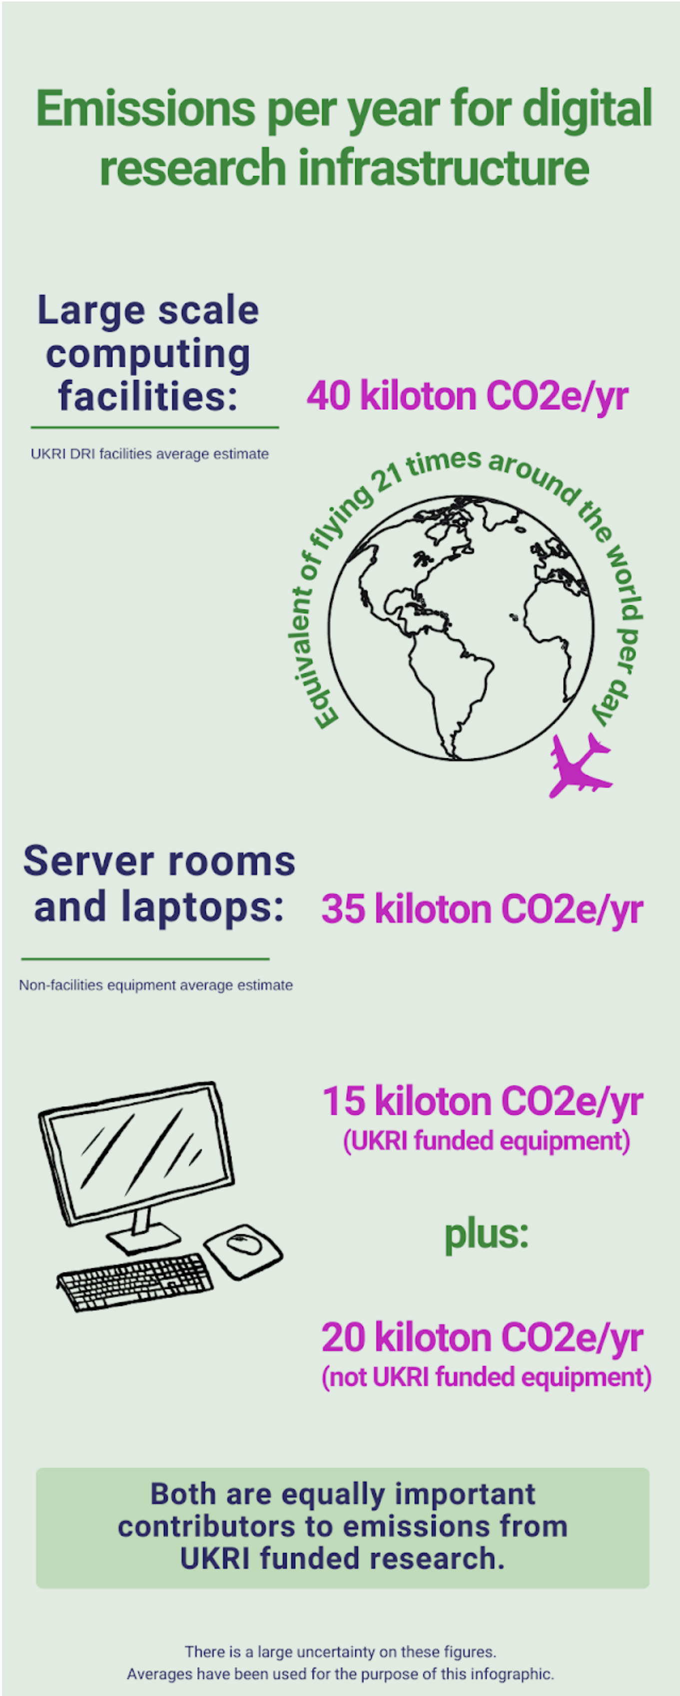

Figure 1.3.3.A: comparison of estimated carbon emissions per year for different aspects of the UKRI DRI ecosystem.

1.3.4 Use Phase and Embodied Carbon - the IRISCAST Case Study

The IRISCAST project (Hays et al., 2023) investigated the sources of carbon emissions from the digital research infrastructures at the six IRIS institutions over a 24 hour snapshot period in November 2022. The six institutions were Queen Mary University of London, University of Cambridge, Durham University, STFC Cloud, STFC SCARF and Imperial College London.

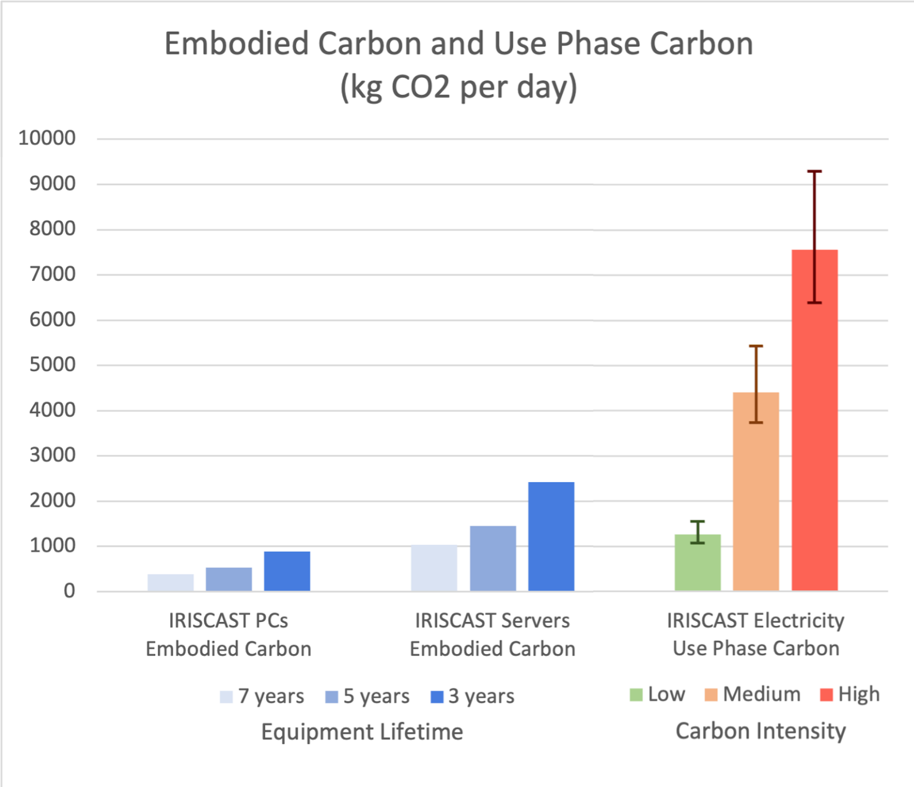

Figure 1.3.4.A: IRISCAST 24 hour snapshot of embodied carbon and active carbon.
Blue bars: Embodied carbon estimates for all 2398 IRIS nodes from six institutions. Embodied carbon assumptions are for typical PC characteristics and server characteristics. Blue shading indicates the embodied carbon contribution to the 24 hour snapshot given different equipment lifetimes.
Green/Amber/Red bars: Use phase (active) carbon estimates for all 2398 IRIS nodes for low, medium and high carbon intensities, the height of the bars are representative of medium Power Usage Effectiveness (PUE), the "error" bars indicate the active carbon range for each carbon intensity from low to high PUE.

1.3.4.1 Use Phase Carbon

Use phase (or active) carbon refers to carbon emissions from the active use of DRI, that is, from the electricity supply. The carbon intensity of the electricity supply in the UK shows a lot of variation over time and regions. UK average carbon intensity ranged from 60 to 300 gCO_2 /kWh in November 2022. Therefore, to prevent calculations from being overly affected by the specific time chosen for the snapshot, carbon intensity values were chosen to be representative of low, medium and high carbon intensities (50, 175 and 300 gCO_2/kWh respectively).

Use phase carbon emission estimates due to the electricity consumption of IRIS equipment:

- Low Carbon Intensity: 969 kg CO_2,
- Medium Carbon Intensity: 3391.5 kg CO_2,
- High Carbon Intensity: 5813.8 kg CO_2.

Not all of the electricity consumed by a DRI is used to power computing equipment, the operation of a DRI facility also requires extra energy for power distribution and cooling. The additional energy required for operation is represented by a facilities' Power Usage Effectiveness (PUE) metric. Values representative of low, medium and high PUE used by IRISCAST were 1.1 (10% extra energy required), 1.3 (30% extra) and 1.6 (60% extra). The range of PUE values used by IRISCAST are in broad agreement with PUE information gathered in the DRI mapping survey (Stephens et al., 2023).

Use phase carbon estimates when energy required to operate the facilities is included:

- Low Carbon Intensity and Low PUE: 1066 kg CO2
- Medium Carbon Intensity and Medium PUE: 4409 kg CO2
- High Carbon Intensity and High PUE: 9302 kg CO2.

1.3.4.2 Embodied Carbon

The estimates of embodied carbon made by IRISCAST used values for a Fujitsu PC[18] and a typical Dell rack server[19]. The individual carbon estimates were scaled to be representative of all 2398 nodes. The contribution to carbon emissions for a given 24 hour period is dependent on the lifetime of the equipment, the IRISCAST estimates assumed lifetimes of 3, 5 and 7 years.

1.3.5 The Base Year and Targets

The Science Based Targets initiative (SBTi)[20] provides guidance on setting ambitious corporate climate action targets. They advise that targets should be set relative to a base year and adjustments made only when specific criteria are met, such as a significant change in the inventory. This report follows this advice and Table 1.3.5.A lists those areas which are within scope of the "net zero" ambitions discussed throughout this report. The base year is set as 2022 (the year of the project during which measurements and estimates were made).

Scope & Category	**Description**	**Comments**
Scope 1	On-site emissions	On-site emissions will be a small component of the total facility footprint.
Scope 2	Power supply	The power supply dominates the operational carbon footprint.
Scope 3, Category 1 & 3	Procurement of hardware	It is clear that the embodied carbon footprint of HPC hardware is large, but reliable quantification is not possible at this stage.
Scope 3, Category 6	Academic Travel	Academic travel is a major source of emissions for the academic sector as a whole, but the allocation of academic travel to DRI

[18] Fujitsu ESPRIMO P9010 Desktop PC
https://www.fujitsu.com/global/documents/about/environment/Life%20cycle%20analyses%20of%20Fujitsu%20Desktop%20ESPRIMO%20P9010%20June%202021.pdf
[19] Dell Typical Rack Server
https://i.dell.com/sites/content/corporate/corp-comm/en/Documents/dell-server-carbon-footprint-whitepaper.pdf
[20] https://sciencebasedtargets.org/

		operations is unclear.
Scope 3, Category 10 & 11 (direct)	User devices	Scope 3 emissions include those arising from user activities which are integral parts of service access, such as obtaining and running devices to access the service.
Scope 3, Category 10 & 11 (indirect)	Impact	As above, but looking at the impact of DRI usage. Listed here as a separate element to emphasise that there is no equivalence between indirect impacts, which may be carbon negative, and positive impacts in other scopes and categories.

Table 1.3.5.A : The dominant elements of the DRI footprint set in the context of the GHG Protocol emission scopes and categories.

Scope & Category	Baseline	Targets(*)
Scope 1	Not available	Reduce to zero by 2040;
Scope 2	20 to 60 kiloton per year	50% reduction by 2030; by 2040 to a level which can be balanced by certified removal of carbon from the atmosphere;
Scope 3, Category 1 & 3	3.4 kiloton per year (extremely high uncertainty)	Reduce as fast as possible to a level which can be balanced by certified removal of carbon from the atmosphere; Develop tools to enable accurate quantification.
Scope 3, Category 6	Not available	
Scope 3, Category 10 & 11 (direct)	25 kiloton per year (extremely high uncertainty)	
Scope 3, Category 10 & 11 (indirect)	Not known	Seek to quantify.

Table 1.3.5.B : Baseline estimates for dominant elements of the DRI footprint and emissions reduction targets, using location based carbon intensity of power. Estimates using the market-based carbon intensity will be considerably lower as many host institutions use power supply contracts claiming zero or near-zero emissions. Baseline estimates are described in more detail in Table 1.3.5.A.

**: Any actions to reduce emissions in one scope must also be audited to ensure that they are delivering a net reduction across all scopes rather than displacing emissions into a different scope.*

There are substantial uncertainties around the UKRI DRI baseline emissions, not least because of uncertainty in the scope of the inventory (there is large uncertainty in the nature of small and medium scale facilities). Following the 2020 guidance from SBTi and ITU[21], we set distinct targets for different components of the carbon budget (Table 1.3.5.A and B).

Nevertheless, the total carbon footprint of the UKRI DRI is tiny in comparison with the IT sector as a whole, and the whole sector is a small component of global emissions. However, the breadth and excellence of the academic community in UKRI which can tackle the

---

[21] Guidance for ICT companies setting Science Based Targets https://sciencebasedtargets.org/resources/legacy/2020/04/GSMA_IP_SBT-report_WEB-SINGLE.pdf and Greenhouse gas emissions trajectories for the information and communication technology sector compatible with the UNFCCC Paris Agreement https://www.itu.int/rec/T-REC-L.1470-202001-I/en

technical and societal challenges, acting as an agent of change, provides UKRI with the opportunity and obligation to take a leadership role (cf. Radinger-Peer and Stoeglehner 2013).

1.4 Positive Transformation on the Journey to Net Zero

> *"The challenge we face today is to break with a system that treats technology and finance as a route to short term profit, rather than as a means to social and environmental ends; the urban planner who drives a motorway through the middle of a city merely to advance the single minded goal of mobility employs technology for the wrong ends. So too do those developers and architects who design the largest, cheapest, and most profitable building with no regard for a city's public life or environment. I am wild about science, not about science run wild. And that thought I suppose is at the heart of my concept of sustainability - the critical application of creative thinking and technology to secure our future on this small planet"*
>
> *Richard Rogers, Reith Lectures, Sustainable City - The Culture of Cities, 12 Feb, 1995.*[22]

The science that informed the commitment to reach net zero and the intertwined biodiversity targets gives us an unprecedented opportunity to plan the world we give to future generations. There are grave risks and challenges, but also huge opportunities. This report sets out initial steps in areas of consensus building, organisational structures, and delivering change.

The opportunities will not fall into our hands easily. There is a massive challenge and much work to do, but this falls well within the core UKRI remit of using research to transform society. As noted by Prof. Dame Leyser (UKRI Chief Executive, 2023) in the quote below[23], we need to combine clear path-setting priorities from the centre with the best in creative thinking from across the UK research community.

> *"As with so much in research and innovation, success requires an ability to support creative, free ranging, open, bottom-up activity, and to connect it dynamically to top-down priorities, ensuring both current and future needs are met."*

By 2040, it is likely that the term 'net zero' will mean something different to what it currently does. The current narrative, solutions, goals and terminology will have changed. These are all dynamic concepts shaped by the social, economic and climate contexts of the time, and have already changed greatly in the last 10 years. We must be adaptable, progressing through iterative and open processes, in order to keep up with these inevitable changes. Whilst our scoping project has been tasked with helping UKRI reach its current net zero goals, we believe that UKRI must use its influence and leadership capability to go above and beyond in tackling the climate emergency our society faces. We must embed transformational change into every aspect of our research practices and processes. We all have a shared responsibility to deliver personal and collective change.

UKRI has a significant role in funding and supporting research and innovation in service of net zero targets across sectors and industries. It therefore has a reciprocal responsibility and leadership capability for adopting research and innovation to minimise the impact of research activity. The DRI is uniquely placed as an increasingly cross-council, cross-sector

---

[22] https://www.bbc.co.uk/sounds/play/p00gxnzz

[23] The quote refers to UKRI as a whole but the phrase very clearly captures the need for both a directed and a creative element in the transition to Net Zero. www.ukri.org/blog/uk-research-and-innovation-a-critical-national-asset

resource which is facilitating solutions-based research into environmental sustainability. There is a distinction between the applications of DRI for research into sustainable solutions, and the environmental impacts associated with the DRI itself. Expanding usage and demand is driven by national political and investment agendas, in which UKRI is a key stakeholder. This tension must be acknowledged, and proactively managed in order to avoid increasing environmental impacts of DRI in the short, medium and long term. To date, increases in computational efficiency have been offset by increases in affluence (Bol, Pirson and Dekimpe, 2021), resulting in a net increase in the carbon footprint for computation (Interim Report, Section 3). Efficiencies in computation are essential to reducing emissions under the CCC balanced pathway. If we wish 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.

The many complexities of governance, shared responsibilities, technology, supply chains, economic, political and social contexts should not divert us from simple truths:

- UKRI has both an immense intellectual capacity and a mission to transform society.
- Society is going through a communication and information revolution powered by digital technology: UKRI can exploit its leading role in digitally enabled research to demonstrate how the information revolution can fit with net zero targets.

There is a great opportunity for UKRI here to deploy the full range of knowledge across the research sector to understand.

> *The science is unequivocal: a global increase of 1.5° C above the pre-industrial average and the continued loss of biodiversity risk catastrophic harm to health that will be impossible to reverse.*
> *…*
> *Targets are easy to set and hard to achieve. They are yet to be matched with credible short- and longer-term plans to accelerate cleaner technologies and transform societies. Atwoli et al. (2021)*

1.4.1 The Journey to a Net Zero DRI by 2040

A central conclusion of this research is arguably the most obvious: it is unlikely that adequate progress to net zero digital research infrastructure will occur without a clear plan to ensure it does so. However, we do not anticipate that detailed plans will all be developed centrally. Plans to achieve net zero will be developed at various levels within the overall UK research community; some as DRI-specific plans and others as plans within larger efforts to accomplish net zero at the level of an entire institution. There will, however, be commonalities between them. At all levels, a journey is being undertaken, which entails three distinct phases:

- Setting out: preparing, route planning and allocating associated responsibilities.
- The journey: undertaking the transition to net zero.
- Arriving: successful accomplishment of the transition to continuing operations on a sustainable, zero-carbon basis.

The toolkit and roadmap are complemented by this journey metaphor which captures the need for a community-wide paradigm shift and a base year which provides a reference point for our departure.

The journey metaphor captures the sense of transition, paradigm shift, and positive aspirations for a sustainable future.

1.4.2 Stages on the journey to net zero

The transition to net zero presents considerable challenges. We need to be careful to distinguish between transition activities which are a necessary part of the journey and sustainable activities which we aim to establish by 2040 or sooner (see Figure 1.4.2.A).

On the journey itself we are going through a process of transition and reducing carbon footprints. During the journey we will need an advanced research infrastructure to drive the innovation which is needed to achieve net zero. Building and operating this infrastructure will, while we are on the journey and short of our net zero destination, inevitably incur a carbon footprint. In this stage, these investments and their footprints must be justified by robust evidence that the infrastructure is moving us towards our destination. At the final stage, arriving at net zero, we have achieved sustainable infrastructure. At this point we will understand and have the means to ensure that we can practically meet net zero, on a dynamic basis and responsively to changes in research etc. This will liberate us from a range of constraints and open up huge possibilities. At this stage we will no longer need to consider justification of a carbon footprint and will be able to focus on other societal challenges.

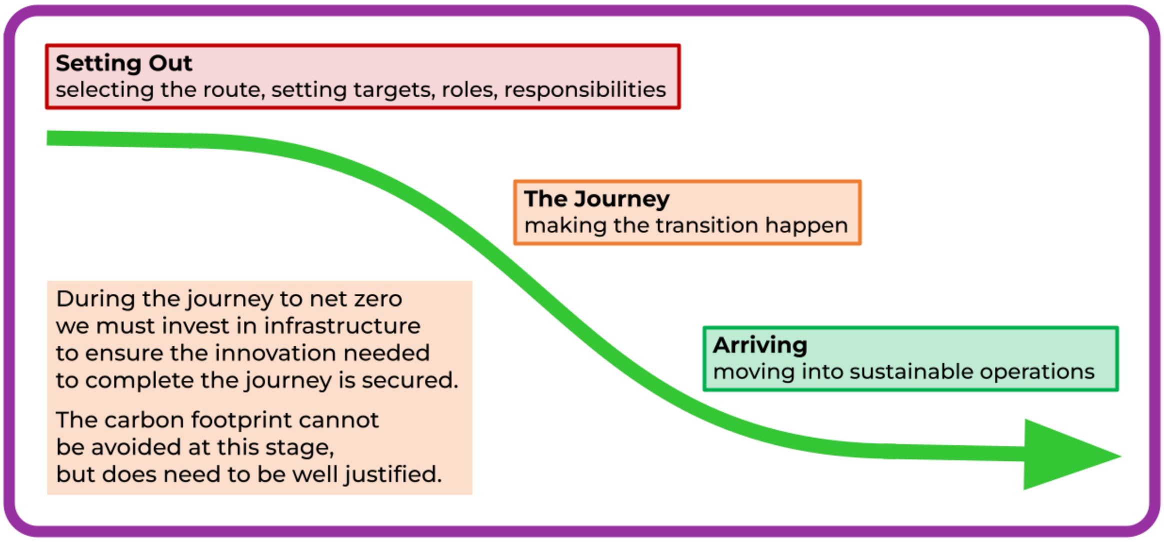

Figure 1.4.2.A: The journey to net zero starts with a "setting out" phase characterised by a lot of discussion and exploratory work, followed by a period of transition, and then arrival at net zero, which is itself a stage on a longer journey to stabilising the climate.

1.4.3 Vision for 2040

In planning our pathway to a sustainable DRI in 2040 it is useful to have a vision of the future that we want to build, a vision which captures the ambition for the sustainable world of the future. It is hard to improve on the closing words of Richard Rogers' quote (see Section 1.4 above): "*the critical application of creative thinking and technology to secure our future on this small planet*".

The journey to net zero computing requires transformational change which can deliver huge benefits as part of the broader transformation to a new generation of digitally enabled research. The transition to sustainable research computing must follow and support the national transition to a sustainable economy, as outlined by the CCC balanced pathway. The evidence generated by this scoping project provides the skeleton for a framework to support a vision for net zero computation that is fair and inclusive; there is precedent for including social justice within the operation of research practices, and research enabled by a net zero DRI should be no different. We have also sought to ensure that our recommendations are resilient to scenarios as yet undefined; the timeframes for new technological capabilities are unfolding with uncertainty, and have both potentially positive and negative bearings on the pathway to net zero.

Given that priorities for investment and evaluation are choices, which are dynamic and subject to institutional and government influence, it is vital that a shared set of principles, values and responsibilities guide responsive action and adaptation in the face of change. The rationale for a broad assessment of the environmental and social impacts of the DRI is echoed in the aspirations for an expanded, interconnected and advanced digital infrastructure in the UK; digital technologies enable connections and impacts on local, national and international scales between diverse communities of practice and focus. When we consider the impact of DRI, an equally wide definition of geographies and communities must be considered across all scopes. This requires cooperation, transparency and adaptability. If these ways of working can be considered solutions to challenges associated with net zero, they reflect approaches being developed to address fiscal, social, and ethical dimensions of AI and Big Data usage[24].

The changes required for environmentally sustainable computing are often synergic and overlapping with effective data management practices, efficient resource usage, and Open science. Our aspiration is that a net zero DRI reached under these conditions will leverage wider benefits to the research community and wider society, beyond the net zero target.

The sustainable research environment of 2040 cannot be designed or specified today, so we must combine strategic long term objectives with specific short term actions which can be completed and evaluated within months or a few years.

Box 1.4.3.A: A vision for UKRI DRI 2040
Facilities have a five-star sustainability status, with everything from the tea bags in the staff canteens to the racks of servers in the data centres covered by a comprehensive life-cycle analysis. Virtual and augmented realities transform our interactions with data and with each other, reshaping our notions of space and time and shattering existing barriers to understanding. Experts provide a resource of digital excellence supporting a transformed national economy. The UKRI DRI reputation for environmental excellence and its leading role in promoting productivity through Open Science policies and workflows attracts leading researchers from all over the world.

[24] A critical field guide for working with machine learning datasets. Part of the 'Knowing Machines' research project, developing critical methodologies and tools for understanding, analysing, and investigating training datasets, and studying their role in the construction of "ground truth" for machine learning: https://knowingmachines.org/critical-field-guide#sec-6

2 Synthesis - what do we know now?

Section 3 below will set out the project results from a project perspective, summarising what was achieved by each activity within the project. Here, in Section 2, that material is synthesised and restructured in order to make it more accessible for those who wish to use the results. There are four subsections, the first, Section 2.1, sets out results in the form of a toolkit with a subject-based synthesis which will be useful to practitioners seeking guidance on specific topics. Secondly, in Section 2.2 a roadmap provides a synthesis in terms of actions and target dates. Finally, Section 2.3 presents principles, knowledge gaps, tensions and highlights.

Both the toolkit and the roadmap are derived from 180 detailed recommendations, as illustrated in Figure 2.A). Individual recommendations are indicated by a unique reference, in the format "[*number*]", which can be used for citation purposes. The evidence is presented in Section 3 below, and a full list of recommendations is available in Appendix 3 or available on Zenodo in a spreadsheet format, which allows you to sort and interact with the data.

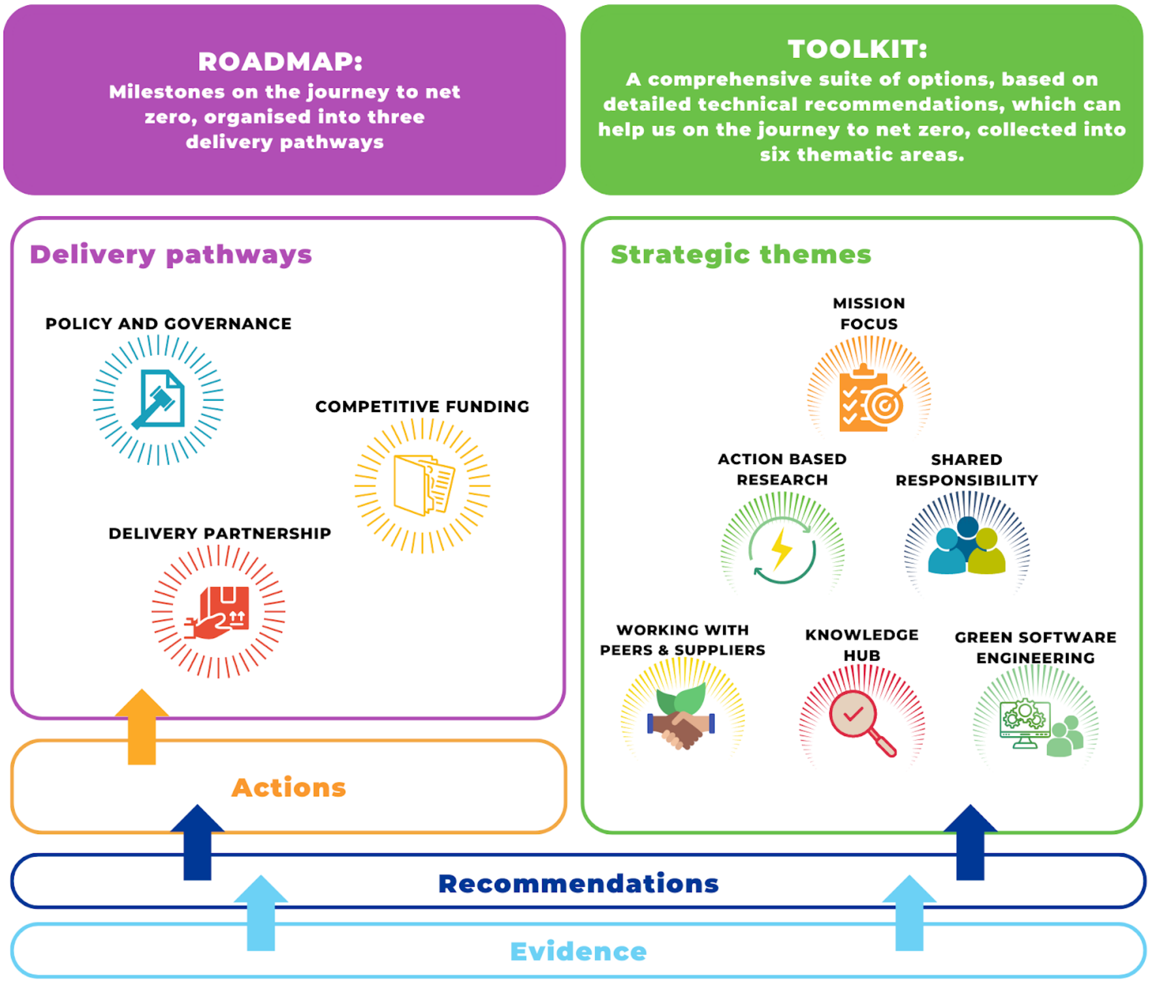

Figure 2.A: Net zero delivery framework. The roadmap and the toolkit are both based on the recommendations derived from evidence. For explanation of the elements of the toolkit and the roadmap see sub-sections 2.1 and 2.2 respectively.

The toolkit sets out the "what" can be done, whereas the roadmap is setting out "how" the UKRI DRI can implement the recommendations by 2040 or sooner. Decision makers and funders can use the roadmap to shape future work.

2.1 Toolkit - what can be done?

In this Section we discuss how the toolkit sets out "what" can be done. Whereas the roadmap is discussed in Section 2.2 as to "how" the UKRI DRI can implement the recommendations by 2040 or sooner. Within the toolkit synthesis, there is a discussion of a range of options for managing and reducing carbon footprint. DRI staff and users can use the toolkit to support their own attempts to reduce their carbon footprints.

The toolkit is gathered together in six strategic themes: mission focus, shared responsibility, action-based-research, working with peers and suppliers, support and advice, and underpinning it all, green software engineering (see Box 2.1.A). Each theme cites specific individual recommendations from evidence collected by the Project. Individual recommendations are indicated by a unique reference, in the format [number], which can be used for citation purposes. A full list of the specific recommendations can be found in Appendix 3.

Box 2.1.A: Six Strategic Themes that make up the toolkit
1. **Mission Focus**: continuous assessment and focus on the mission of achieving sustainability; active measures to counter the risk of enhanced demand negating efficiency gains. 2. **Recognition of shared responsibility**: mandate and empower all staff (from student to CEO) to take proportionate action to drive change and reduce the environmental impact of their work; community building; encourage discussion among colleagues and learn from others to foster positive changes in behaviour. 3. **Action-based-research**: work must start now with commitment appropriate to the climate emergency while recognising that there will be a need for regular checks and adjustments; focus on progress not perfection; small steps; learn from experience. 4. **Work with peers and suppliers**: through contracts, conditionalities, and understanding mutual benefits, to develop a low carbon supply chain [essential in the longer term] 5. **Build and Share Knowledge**: providing leadership, support and advice for business cases and large procurements feeding into reporting; central hub for information and institutional knowledge [also likely to create short term results] 6. **Green Software Engineering**: creating a body of expertise around green software engineering, providing training, developing tools, metrics, expert assessment, and standards to transform current approaches to writing code, and supporting codes running in data centres, such that GSE becomes the norm rather than an optional extra.

We examine each of these in turn, covering a range of options for managing and reducing carbon footprint that DRI staff and users can use to reduce carbon emissions.

2.1.1 Mission Focus

"Continuous assessment and focus on the mission of achieving sustainability; active measures to counter the risk of enhanced demand negating efficiency gains."

The national net zero strategy foresees a transition to zero carbon electricity supply being achieved through a range of factors including demand constraint. There will be some increased capacity available to offset emissions in other areas, primarily to power electric transport and replace the existing fleet of fossil-fueled cars.

The UKRI mission of transforming society cannot be achieved within a constrained energy supply to facilities. Therefore consistency with the national mission can only come by clearly positioning any necessary expansion in facility energy consumption in the context of the vision outlined by the CCC Balanced Pathway: for high levels of innovation driving a smooth transition to sustainability.

The use of the term "climate emergency" has become so widespread that we often forget the implications. UKRI scientists are regularly reporting new findings that underline the urgency of taking significant and transformative action. There is a parallel emergency in biodiversity. It is not a competition, we must deal with both. The national strategy requires both enhanced innovation and restrained use of energy. There is a strong will at all levels of the organisation to address this emergency, but action is blocked by lack of time, knowledge and tools.

The net zero target has a clear objective meaning when applied to global emissions, but precise interpretation becomes ambiguous and controversial when applied to organisations because of fundamental ambiguities about allocation (Interim Report, §2.6) and compensation measures through offsetting and capture (Interim Report, §3.9). On deeper analysis, it becomes clear that the objective of demonstrating leadership is a safer guiding principle than any specific quantitative metric.

UNESCO (2021) notes the urgency of "fostering equitable access to scientific information" in order to create resilience and the capability to respond to global emergencies.

The 'net zero' objective is often treated as an objective and physically measurable target, but there is an increasing body of literature treating it as an emerging social norm (see discussion in Section 1.2 above). This brings risks, challenges and opportunities.

Commercial hyperscale data centres can provide significant advantages of scale, and many providers have well-developed reporting on emissions and sustainability. However, quantitative comparison between different institutions remains challenging because of a lack of standards to ensure consistent interpretation of reported data.

2.1.1.1 Review of detailed recommendations

Net zero is much more than a technical challenge, it requires, in less than two decades, a fraction of a career, a transformation of the way that we work and interact with each other. Such transformations are far from unusual in a research community dedicated to transforming society for the better, but a strong mission focus is needed to ensure that this transformation stays on track to achieve sustainability. UKRI should use its capacity in social sciences, arts and humanities, and in economics, to understand the range of societal views, the avenues of consensus which open-up potential for accelerating transition and the emerging (or exploding) discords which can block or reverse change.

The community is looking to UKRI to provide leadership which needs to be backed by effective messaging, a clear narrative which maintains the momentum needed for transformational change. Cross-community engagement supporting grassroots initiatives, backed by continuity of expertise and innovative capacity, will be an essential element of the transformation [25][49][53][116][118][42][119][22][98].

There is a need to embed the net zero ambition in all organisational layers of UKRI and across all research councils. Policies, backed by evidence, need to be put in place to ensure consistency in reporting. Transparency and consistency in reporting is critical, even when precision in terms of numerical carbon footprints is not achievable, including reporting of carbon impact of research activities enabled by the DRI. Emissions reduction pathways should be established for all facilities. As with reporting, consistency and transparency will be critical. As both the operation and use of the DRI is deeply interconnected with other elements of the UKRI estates and communities, policies need to be integrated broadly. [109][99][82][23][27][71][21].

UKRI needs to take an exhaustive approach to reducing emissions because there are many unanswered questions about the scalability, affordability and sustainability of proposed mechanisms for compensation of residual emissions by processes which actively extract greenhouse gases from the atmosphere (Interim Report, (Juckes et al., 2022)).

A centralised resource is needed to gather best practices and promote positive change. A review process is needed to ensure that well intentioned measures do not trigger so-called rebound effects which boost demand instead of reducing resource use and counter risks of missed opportunities or policies diverging from the mission. The mission focus element of the toolkit will enable robust management of the rebound effect and strategic risks through policy constraints and restraints combined with an agile delivery partnership (see Figure 2.1.1.A). [105][7].

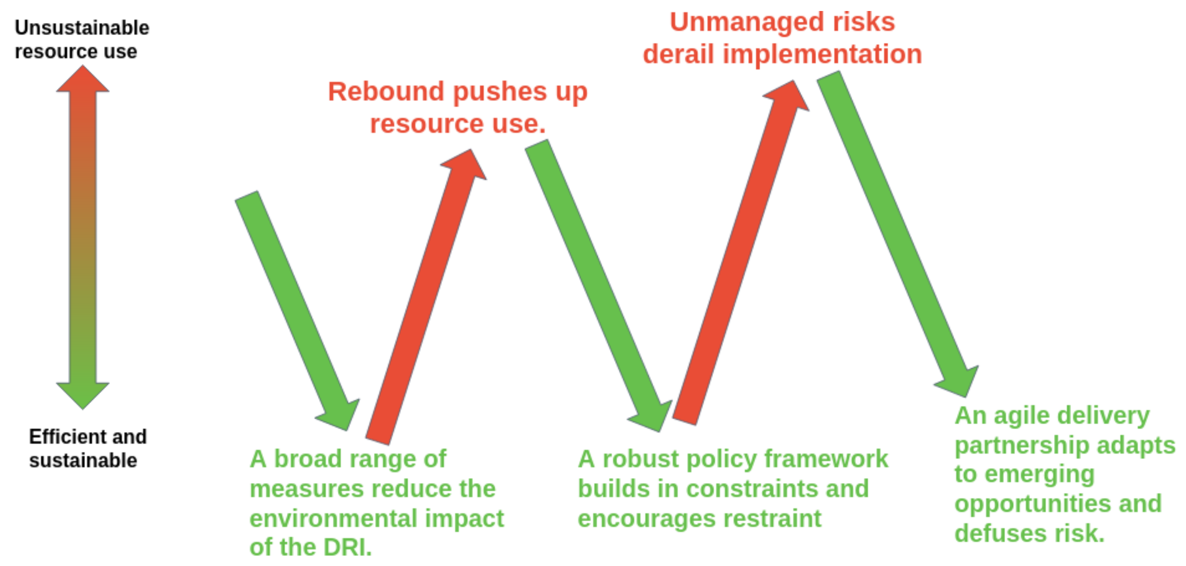

Figure 2.1.1.A : The mission focus element of the toolkit will enable robust management of the rebound effect and strategic risks through policy constraints and restraints combined with an agile delivery partnership.

Open Science and FAIR[25] data should be advanced to ensure efficient and effective use of the DRI. At the same time there needs to be proactive engagement in new technology and attention to continuous improvement in order to exploit all opportunities for greater sustainability. Policies need to cover not only the operation of machine rooms but also the sustainability of staff and user activities required or supported by the DRI such as travel and purchase of connecting devices. [148][95][18]

The DRI serves a large and growing research community. Consistent and efficient user engagement with the service needs to be backed by appropriate support and requirements at the grant approval stage. Environmental impact statements should be mandated with clear carbon and sustainability assessments.

[25] *Findability*, *Accessibility*, *Interoperability, and Reuse*

Investment decisions and award of capital funds are just as important as grant funding and should follow consistent policies.

Procurement of hardware and of electricity are two major pathways to emissions and hence also major leverage points for emissions reduction. HPC procurement and maintenance must reflect carbon efficiency priorities. Facilities need to develop a green energy supply, exploiting on-site generation, power purchase agreements promoting renewable generation investment, and off-grid energy supply.

People and behaviour will make or break the transition to sustainability. Change should be encouraged through awareness raising and incentives. Support is needed for sustainable ways of working, e.g. through personal net zero action plans, monitoring of personal resource use, enhanced low-carbon conference support and through encouraging staff conversations about sustainability.

2.1.1.3 Acting on many time scales using many levers

The most cost-effective approaches to net zero typically involve greater degrees of cooperation and lower degrees of potential direct influence. For instance, it makes little physical sense to buy equipment manufactured in a factory driven by coal-power and then compensate for the emissions through expensive carbon sequestration technologies (Figure 2.1.1.B).

Figure 2.1.1.B: Different implementation pathways categorised in terms of low to high potential for direct decision making (left) can map onto a range of lever-points ranging from immediate measures to multi-year processes.

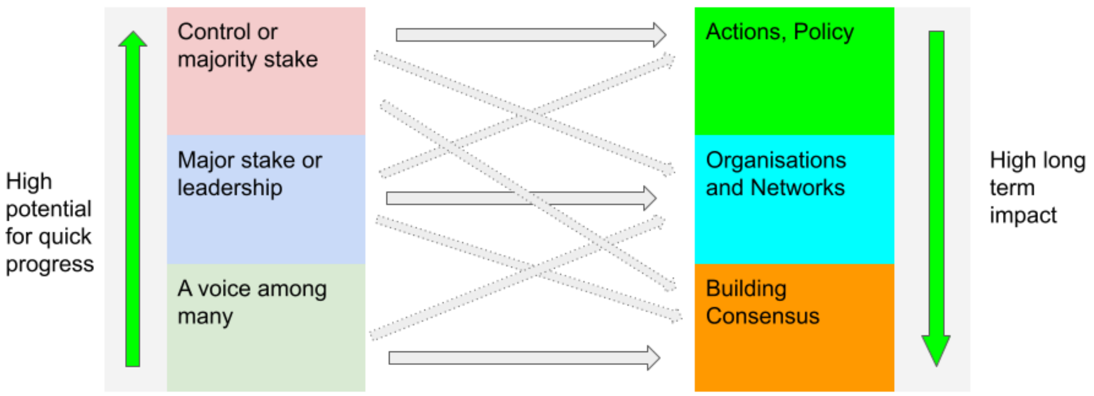

2.1.2 Recognise shared responsibility

“Mandate and empower all staff (from student to CEO) to take proportionate action to drive change and reduce the environmental impact of their work; community building; encourage discussion among colleagues and learn from others to foster positive changes in behaviour.”

Shared responsibility means each person, irrelevant of position in an organisation, has a responsibility to take action to reach net zero. This responsibility should be recognised by each person and by each entity within the UKRI ecosystem. The responsibility extends to understanding that there is a range of appropriate actions that can be taken, and that the end-goal is “net zero” emissions, thus some expenditure of carbon emissions during research in order to save a larger amount of emissions elsewhere is admissible.

Evidence collected (see Section 3) identified the need for UKRI to raise awareness and understand and encourage behaviour change with engagement by all Research Councils, Furthermore, it identified the need for the embedding of net zero principles at all layers of UKRI, with active engagement with each Research Council to ensure inclusion of their carbon footprint within a DRI database (see recommendations: [2][23][25][38][49][99][105][109]).

These could be achieved by developing & sharing best practices, by Net Zero Action Plans, and by a centralised resource, particularly regarding design and running of data centres and offering of incentives. Various evidence recommended community building and engagement, including encouragement of discussions across disciplines to come up with creative solutions around net zero DRI.Other recommendations include support for code writing collaborations and mechanisms to reduce data duplication and replication, and recommending use of the FAIR approach[26] (see recommendations: [3] [24] [31] [47][49] [50] [53] [54] [94] [105] [116] [129] [141-144]).

The ARCHER2 case study (Smith et al., 2023) recommended "good citizen" behaviour [79], whilst regular surveys of employees to help target future education campaigns was a key recommendation arising from the User Behaviour Survey [132] (McGuire, 2023). Both are complimented by our Interim Report's recommendation that DRI investment decisions should be informed by views of the research community [1]. Likewise, individuals and institutions might support carbon reductions via use of virtual conferences, by encouraging and participating in a modal shift to low carbon commuting options, and reviewing working from home policies [118-121]. The latter point requires consideration of the complex balance between reducing emissions associated with travel versus the potentially lower efficiencies of home offices compared to corporate shared spaces.

Grant applications should include an assessment of their carbon impact [44,47,83,122,154, 165,166,179] or perhaps a full environmental impact statement [21,154]. This impact will need to be assessed, perhaps akin to an Ethics Study of all research proposals. As discussed by CI-UK Keynote Speaker, Prof Michèle Weiland, there needs to be a common understanding of the balance between the negative impact of using resources to conduct research and the positive impact which research delivers. UKRI supports a huge range of work that is accelerating the transition to net zero, from improving the understanding of the climate system through a variety of models, to enhancing the efficiency of renewable energy with digital twins of wind turbines[27]. This theme was picked up by the ARCHER2 case study to recommend that the evidence of societal/economic impact is weighted against the carbon impact of the proposed project [83]. The Training and Standards workshop also produced a recommendation for mandatory carbon reporting in funding calls [156], and the ARINZIT sandpit project (Friday et al., 2023) further recommends that sufficient resources are made available to publicise and implement Net Zero policies [27].

In terms of infrastructure, facility managers were identified with specific responsibilities, particularly in terms of reporting and providing support for users to achieve reductions. A key means of empowering individuals to make environmentally conscientious decisions is to report energy data to users, grant holders and the wider research community [[36,37, 56, 59, 69, 70, 88]. Various projects suggested facilities should be obliged to improve performance (including energy) measurement systems [74], perhaps via monthly facility reports [65], with UKRI reporting collective annual figures [103]. Such collective reporting should take proactive steps to include facilities where UKRI may have only partial ownership [112]. The HPC-Jeep sandpit project (Turner and Basden, 2022) goes further with a direct recommendation for energy-based charging for use of DRI facilities [57].

[26] https://www.nature.com/articles/sdata201618
[27] https://gtr.ukri.org/projects?ref=EP%2FX022765%2F1

In terms of tackling embodied emissions current procurement processes need to be reviewed and amended to encourage energy efficient solutions [43], involving a wider range of input (e.g. from hardware experts) [123]. This is necessary to support a common energy reporting standard or perhaps standards for delivery and exploitation of services/software as recommended by the ENERGETIC project (Bane et al. 2023) and Training and Standards workshop [37,147]. UKRI grant calls have a role to play, since research practice guidelines and stipulations can encourage the use of current specialised low carbon data storage/compute facilities such as JASMIN/ARCHER2 [131].

Clearly, net zero cannot be achieved without consideration of active carbon emissions associated with energy supply. The ARCHER2 case study calls for use of green energy [78]. The findings from our literature review supported recommendations for procuring 100% off-grid renewable electricity with clauses for renewable investment [9], and the elimination of on-site use of fossil fuels [16].

Similarly, optimising use of current facilities across the DRI ecosystem could be used to undertake computational research at centres with lowest carbon Intensity (supporting lower operational carbon budgets) and reduce need to expand new DRI infrastructure (so not increasing embodied emissions of the UKRI DRI estate) [26]. It is recommended that DRI site inventories include both operational & embodied emissions, which specifically record the idle power draw of each data centre [64]. In order to support steps taken to reduce this several projects recommended estimating and investigating energy use and carbon intensity to track and reduce DRI's emissions [32,73, 76]. The ARINZRIT sandpit project further recommends investigating mechanisms for the use of excess heat generated by DRI facilities (specifically data centres) within their local communities, in order to reduce energy waste and improve carbon efficiency [30].

A principle of shared responsibility is that everybody pulls together to do what they can within the scope of their remit and capability. This is vital in the challenge to meet the 2040 target date for UKRI DRI to be net zero. Recognising where community building is already taking place and supporting these groups alongside new initiatives is a key recommendation arising from the Interim Report (section 3.2). UKRI should ensure the continuity of net zero expertise and activities through actively supporting communities of practice and investing in research in this space to grow new best practice [22]. Further to which, whilst the roadmap presented in Section 2.3 sets out actions to be delivered by UKRI, successful implementation of the roadmap actions assumes collective engagement and enactment of change.

The community building element of the Project has been highly productive (Section 3.5). Through a range of activities the Project has brought interested parties representing diverse backgrounds together, which has generated lively and creative discussion. **Community building supports all areas of the net zero challenge, creating energy and drive to address other recommendations.** Crucially, the community is not just those who run and use the current DRI facilities, but people from across the UKRI community who have an interest in a productive national research programme in a world in which digital services play an increasingly central role.

As with risk assessment, environmental impact assessment must be embodied in every-day work practices at a level which is appropriate and proportionate to the task at hand, whether this is submitting a computational task to a DRI facility or planning a major infrastructure investment. Any environmental impact assessment for DRI use and development must be consistent with environmental impact assessment in other UKRI activities, such as major investments on physical infrastructure or planning academic travel.

2.1.3 Action-based Research

> *"Work must start now with commitment appropriate to the climate emergency while*

recognising that there will be a need for regular checks and adjustments; focus on progress not perfection; small steps; learn from experience."

On the 1st May 2019 the UK Parliament declared an environment and climate emergency. In such circumstances, and in the context of this report, it is essential that work on decarbonising the UKRI DRI starts now with a resolve to: focus on progress not perfection by making small steps and learning from experience; recognise this will be an iterative process where regular checks and adjustments will likely be needed.

The user survey, described in Section 3.4.4, found that different individuals will engage in different ways with the DRI net zero work. Thus, to ensure that this variation - in use patterns and attitudes towards net zero by users - is accounted for, it is recommended that a range of behaviour change techniques and strategies are employed [133]. Follow-up surveys may be required to measure effectiveness and track changes in attitudes so that adjustments can be made [132].

Pressure of daily work will prevent engagement with net zero efforts so it is recommended that sustainability practices are built into workplace policies to enable sustainable behaviours to be integrated into daily work rather than being perceived as an additional burden [117]. This is supported by the ARINZRIT recommendation (Friday et al., 2023) to formalise net zero research incentives to reshape academic practice, promoting research which truly embeds a sustainable approach to DRI [25, 31]. For example: by assessing DRI's full lifecycle carbon cost in peer-review processes of funding applications and during project execution and review meetings, and by offering best paper awards for delivering results with minimal environmental impact. This again takes multiple approaches and focuses on embedding the change into daily research work life.

Moving to full lifecycle analysis overnight is not practicable. Therefore the VALUE sandpit project (Boulton, 2023) makes the recommendation to initially use energy estimates and average grid carbon intensity to plan and track DRI use emissions until more detailed systems are in place [73], using information by new systems to monitor energy usage by facilities, servers and users (IRISCAST (Hays et al., 2023) and DRI Mapping (Stephens, Kayumbi & Lambert, 2023) recommendations [64-67] *and* [102-104,107,112]).

The GoZero and the DRI Mapping projects recommend capturing the current state of UKRI DRI [48,96]. IRISCAST recommends gathering inventories of equipment at sites and idle power draw measurements [64]. Combining these recommendations would result in a database of key information which is periodically updated using automated data feeds [66-67]. Work to start building such a database and data feeds should start now. This dataset, after a number of iterations to learn from any mistakes, can then be used in the medium term to make more detailed carbon cost calculations. Such an audit could also investigate if consolidating small DRI resources makes sense from a carbon cost perspective. Idle power consumption was also considered by CQUANDRI (Schien et al., 2023) which recommends various techniques to significantly reduce idle power draw and thus operational carbon emissions [32-34].

Some further investigation and research is recommended into the significance, with regards carbon emissions of: end user devices (laptops/desktops etc); the academic network; and research use of commercial public cloud providers [97]. Getting accurate comparable carbon costs for these items may be difficult or impossible and care must be taken when evaluating such data. Further investigation into time shifting DRI workloads to periods with a lower carbon intensity is also recommended as is an investigation into the significance of data triage to store data into less carbon intensive tape archives where appropriate [32, 60].

Further consideration of extending the lifetime of DRI equipment to get the most out of the embodied carbon cost is also recommended [15].

There are a number of other actions that can be taken now. At the infrastructure level it is recommended that DRI sites test themselves against the PRACE best practice recommendations [84]. Sites can also consider delaying decommissioning of kit to extend the lifetime of DRI equipment to get the most out of the embodied carbon cost [15]. For Laboratory infrastructure adoption of the LEAF (Laboratory Efficiency Assessment Framework) is recommended [134]. It is also recommended that sites share knowledge and best practice of heat reuse systems at DRI resources with the aim of making heat reuse into a standard practice across UKRI [102].

At the user level feedback of energy use or carbon costs of user jobs will help drive user behaviour change and it is recommended that DRI resources start to do this where possible [90]. Additionally options for energy-based or carbon cost charging or allocation should be explored for current and future HPC services [57]: again to drive user behaviour change.

These recommendations collectively show that much can be achieved by getting started now, and much more can be achieved by learning and adjusting as progress is made.

2.1.4 Working with peers and suppliers

“Through contracts, conditionalities, and understanding mutual benefits, to develop a low carbon supply chain”

Significant aspects of the reduction of the scope 2 emissions from power supply and the scope 3 emissions associated with purchased goods (the so called embodied carbon) and the ongoing value chain of the UKRI DRI are outside the direct control of UKRI. These emissions can be represented as a capacity factor multiplied by a carbon intensity factor. The capacity factor can be mitigated by making more efficient use of power and hardware, but we cannot achieve net zero without addressing the carbon intensity factor without engaging with suppliers.

Investment decisions need to exploit community expertise, and steps need to be taken to minimise barriers to adopting new technology and to ensure that enhanced efficiency does not just lead to expanded capacity [1][5][7][123]. To avoid introducing burdensome requirements, reporting details need to be appropriate to the scale of activity [97]. There is a need for both long-term planning and contractual arrangements and flexibility to combine funding streams (e.g. ensure there is no accounting barrier to investing in efficiency) [9][12][89]. Improved standards and guidelines, as well as embedded green principles, are needed to ensure consistency in reporting across the supply chain [158, 174, 176].

There is a need to work with peers and suppliers to develop informative categories of facilities and services and recognised accreditation for practitioners [100, 152].

There is a need for a policy on overall power consumption, including both the DRI and the facilities it supports, to align with the national balanced pathway [159].

Strong relationships, both contractual and outside contracts, with suppliers are needed to enable accurate analysis of life-cycle costs. Specific measures include provision of reliable information on embodied carbon and capabilities to inform users of power usage to facilitate efficient usage, including measures of idle power draw [13][14][15][43][55][61][64][75][76].

Influence over scope 2 emissions, from generation of electrical power, can be yet further removed because power is often not procured directly by UKRI but by the DRI hosting organisations. Scope 3 emissions from the procurement of servers and from data transfers outside the data centres are hard to quantify precisely.

Procurement rules are set by broader institutional policies and constrained by competitive forces from the commercial world. It is clear that better quantification of the carbon footprint is important, but also that effective measures to reduce the footprint can be taken on the information available now.

The fact that the UKRI DRI cannot directly control the crowd activities does not mean that there are no opportunities for leadership. If anything, the reverse is true: in confronting net zero it is leadership in the crowd that is particularly important. Providing leadership in the crowd is business-as-usual for UK scientists. It was, for instance, a British scientist, John Houghton, who led the first three IPCC reports which laid the foundations for the global adoption of a net zero target.

2.1.5 Knowledge hub

“Providing leadership, support and advice for business cases and large procurements feeding into reporting; central hub for information and institutional knowledge [also likely to create short term results].”

Information needs to be accessible and authoritative. Within the knowledge hub theme we collect recommendations that contribute to the creation and dissemination of a body of knowledge which will support the UKRI community in the journey to net zero. The knowledge needs to be delivered not only to the facility managers and users, but also to the funders and policy makers to ensure that appropriate funding is available.

The CCC (UK Climate Change Committee) balanced pathway requires a high level of innovation. This will require advanced research facilities to support a world leading research and innovation programme. Where power requirements of new facilities exceed the baseline and thus depart from the CCC balanced pathway assumption of level power draw from the national grid, these new facilities must have a clear and explicit justification for the expanding power draw.

Training of staff at all levels is needed, both to increase awareness and understanding of implications of climate change and net zero and to provide technical competence to deliver change. Training needs to be backed by an active programme of learning and discovery. The roadmap to net zero will pass through unexplored territory and training material will need to be regularly updated with lessons learned from exploratory pathfinder projects at UKRI and elsewhere.

The various projects have identified training and support for stakeholders, for sharing information and best practices (and in investing to formulate improved best practices), and to support how individuals and institutions can find answers to their questions regarding Net Zero. By stakeholders we consider end users, researchers, ECRs and students, and also facility managers and UKRI policy makers. Specific recommendations include training for all stakeholders on all aspects of energy [86], energy efficiency and carbon emissions (i.e. some form of Carbon Literacy[28]) [108]; and Net Zero (e.g. building upon the Green Software Practitioner[29]) [146]; and on wider sustainability matters. Several projects including ENERGETIC (Bane et al. 2023) advocated training up on expert matters [39, 140], such as those covered within “Green Software Engineering”. This is echoed in further recommendations arising from the Training and Standards workshop and User Behaviour Survey (McGuire, 2023) (advising mandatory training within UKRI on why it is important to change practices in support of Net Zero [149] and incorporating sustainability practices into workplaces to facilitate more sustainable behaviours [117] respectively). Training should be

[28] https://carbonliteracy.com/

[29] https://training.linuxfoundation.org/training/green-software-for-practitioners-lfc131/

extended to PhD students and Early Career Researchers (ECRs) working with specialised low carbon storage/compute facilities such as JASMIN/ARCHER2 [130].

Aligned with these training recommendations, is the requirement to raise awareness at various levels.The ENERGETIC project highlighted the need to raise awareness that optimising code can lead to lower energy consumption (and that one can optimise for lower energy consumption)[38]. The User Behaviour Survey noted that more information is required to encourage stakeholders to collaborate when writing code such that Green Software Engineers can advise accordingly [129]. This takes the traditional HPC idea of co-design (for speed) and re-aligns for co-design to minimise carbon emissions. This design of code should incorporate consideration of deployment and make use of new/alternative technologies (hardware) and techniques & tools (software) that most appropriately reduces emissions [6, 35]. In order to maximise such savings guidance and training on the sustainable use of new technologies is needed. Support and advice are prerequisite to implementing Green Software Engineering across the DRI. This should be resourced and accessible to users [151], facility managers [106, 145], and the wider research community [125,127]. Specifically, the Sustainable Computing project recommends the adoption of the Laboratory Efficiency Assessment Framework[30] (LEAF) principles into DRI in order to reduce environmental impact [134].

The measurement (or estimate & prediction) and reporting of carbon used needs to be standardised (see Green Software Engineering) and best practice developed [48, 85].

Embodied emissions are also vitally important for everybody to understand. The User Behaviour Survey had a balance of recommendations from documenting how equipment upgrades could be more efficient (operationally) to the re-use or recycling of unwanted equipment and creating case studies (or similar) to highlight how reuse has helped others and the challenges of recycling.

The User Behaviour Survey expanded the recommended information, support and advice given to all stakeholders to help raise awareness of all environmental impacts within the workplace, and where to access training and support to reduce this, with management taking a structured approach (e.g. appraisals) to ensure adoption of positive practices [128].

More information is needed about the extent of the DRI and the DRI value chain (the value chain includes computers which are not majority funded by UKRI but which are used to support existing user ways of working). From public websites, EPSRC tier-2 facilities and large STFC HPC facilities have a combined core count of close to 300,000, which is just below 50% of the core count of the ARHER2 HPC. Additional significant facilities funded through MRC and BBSRC have not been well characterised, but are likely to be comparable in scale to STFC and EPSRC investments, so we can conclude that overall capacity in tier-2 scale facilities is close to the ARCHER2 core count. Section 3.4.5 estimates a total of between 100 and 300 GWh per year. Energy consumption of non-facility equipment is estimated to be slightly smaller, but the supply chain footprint of non-facility equipment (which is included in DRI scope 3 emissions) dominates.

2.1.6 Green Software Engineering

“Creating a body of expertise, providing training, developing tools, metrics, expert assessment, and standards.”

Computing requires software, and the use of computers creates carbon. The term “software” 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. from schedulers on multi-user machines to use of ChatGPT as a tool). The term “Software

[30] https://www.ucl.ac.uk/sustainable/leaf-laboratory-efficiency-assessment-framework

Engineering" typically 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.

Many projects have made technical recommendations. Some of these may not precisely fit the above discussion of Green Software Engineering, and even more clearly do not fit any of the other Six Strategic Themes. For example, recommendations relating to procurement could be considered under "Working with Peers & Suppliers" but is also covered within this section on Green Software Engineering since options over deployment of software should consider embodied emissions (as well as operational emissions) and thus lifetime of hardware (or purchase of more efficient architectures).

During the UKRI Net Zero DRI Scoping Project, various projects have made direct recommendations around UKRI providing expert support in Software Engineering regarding designing, writing, developing, deploying and maintaining software with a primary focus on reducing energy consumption [6,17, 29, 39, 45, 86, 108; 125,127,130, 137-140, 147]. Further projects have recommended tasks for users (namely, users to obtain and publish energy data) and/or for facility managers that require some training and support in making them happen [42, 52, 56, 60, 67, 68, 70, 72, 80, 85, 107]. That is, software should be designed, developed and implemented to limit energy consumption, and the users of software packages (including trained machine learning models, climate models etc) empowered to use such software with least environmental impact.

These recommendations can (and should) be met by UKRI investing in Green Software Engineering. This should include several approaches from applying current best practices (e.g. via a team of "Green Software Engineers" (GSEs)) and investment in the community to develop further best practices. It is recognised there are currently some, but limited numbers, of people in the UK with experience and expertise in writing green software. It is recommended to invest in GSEs and an initial GSE pool of talent would be a central UKRI team of people with relevant skills/experience to provide support and training in how to make both software greener and the use of software greener. These GSEs could help with selected codes (for example the software codes with high resource use identified in the HPC-JEEP report (Turner and Basden, 2022) which will give immediate carbon savings. The GSEs would write these up as Case Studies and provide training, such that further people gain relevant skills who then support and train further people (and recursively until all researchers, and users/managers of facilities have acquired such knowledge). The pool of GSEs would grow over time and members would be available to UKRI researchers and facility managers.

The learning developed as part of Green Software Engineering expertise needs to be embedded in all software engineering training and practices. This falls within UKRI's remit regarding provision of CPD, advocating for best practice requirements - influencing teaching of undergraduate and postgraduate students - and strategic investment directives for centres of learning and research. Moreover expertise needs to be available to continue development of this area of work across disciplines (in particular to support bespoke solutions for subject domains where code writing is not a core skill associated with the discipline as it currently stands e.g. arts & humanities, economics and social sciences)

The Scoping Project has identified key areas, and these would be part of the skillset of the pool of GSEs. Namely, (i) an initial step to reducing energy is the ability to measure the energy consumption; (ii) embodied emissions are a substantial element of the total emissions of computing infrastructure; (iii) various technologies and techniques can be employed to reduce energy consumption of a given simulation; (iv) where and when software is run affects the carbon footprint; (v) understanding options for not running software or running software tools in a less carbon intensive manner; (vi) sharing best practices with a shared knowledge repository; (vii) continual improvement in knowledge and skills (particularly as technologies/techniques develop and standardised approaches emerge). The ENERGETIC sandpit project specifically recommended UKRI investment in accelerator technologies (e.g. Field Field-programmable gate arrays (FPGA) and Graphical Processing Units (GPUs)) and their support [35].

The GSE team would, as a whole, have expertise and experience across this range and work together with users to quickly identify and then implement the most appropriate approaches that would lead to greatest carbon savings for each unique case. As noted above, this team would share knowledge including new Case Studies and an important aspect of the support to the client is to train them in order to be able to take such further interventions themselves in future.

Many recommendations focus on how facilities could be run in support of Green Software Engineering. As noted above, obtaining energy data is required in order to evaluate inventions to support maximum energy savings per user, per code, per facility and thus across all of UKRI. To this end, appropriate instrumentation is vital [80], covering cooling and other supporting infrastructure [63], and such data needs exposure (in a standard manner) at user level [37, 104, 150, 153]. Data per job should also be accessible to facilities managers and the public, and should support appropriate charging mechanisms to reflect the true cost of consumed resources [85, 107]. Facilities should provide mechanisms to support the use of lower amounts of energy for a given computational/analysis and to use the least carbon Intensive option [51]. Specifically, an appropriate "green scheduler" could cap run-time behaviour and re-distribute work to data centres with (at the given time) lower Carbon Intensity, reschedule to another time (when Carbon Intensity is lower) or potentially automate use of more efficient architectures [32, 33, 60, 72]. Such a scheduler does not yet exist so should be something UKRI invests in making happen.

Facilities can also help reduce emissions by taking steps to reduce idling power and the default frequency of CPUs and GPUs (whilst being aware that in some cases lower frequency may result in higher energy-to-solution) [34, 77, 81]. Facilities can maximise the use of embodied emissions by moving to high utilisation of resources (which may mean powering down some nodes) and consideration of the hardware lifecycle from procurement to grave]. Namely, extending lifetime (and re-use) of servers [15], maximising utilisation [82, 135], buying the most appropriate hardware such as accelerators [35], and avoiding a "rebound" effect during procurement whereby efficiency gains become offset by higher usage [7]. Further discussion of procurement is under Section 2.1.2 above (Recognise Shared Responsibility).

Many of the recommendations have, by necessity, had some focus on given DRI sites. There is also a carbon footprint due to use of the academic networks (e.g. JANET) and this needs further investigation [114]. Similarly many researchers make some use of (public) cloud computing, not just for email and storage (via their institutions) but also for research purposes (e.g. running simulations and analyses). There needs to be an assessment of the carbon footprint of UKRI's use of public cloud and how this can be reduced. UKRI should also determine whether (or rather in which cases) there is an environmental advantage of using cloud to other solutions [115].

It is important to note that UKRI does not (and should not) approach Green Software Engineering solely by itself. For example, UKRI could sign up to and then work with the Green Software Foundation to share & develop expertise and to publicise effective steps to reducing carbon footprint. A key aspect of experienced GSEs is to know which approach will be most appropriate (save most carbon) in the client's unique case. As the Green Software Foundation (GSF)[31] note: "green software aims to reduce the carbon emissions associated

[31] https://greensoftware.foundation/

with software. As such, green software is software that emits the least amount of carbon possible. This can be achieved by making software more energy efficient, using less hardware, or letting software do more when the electricity is clean and less when it's dirty."[32]

Over and above having GSEs available to support researchers and facility managers in current best practices, is the urgent need to discover further efficiency savings.This requires investment into research across UKRI remit, from EPSRC's new technologies/techniques to ESRC research regarding reduction of barriers for paradigm shift to consider energy (rather than say time or money) as key priority. Specific recommendations for research and investment are given throughout this report.

[32] https://stateof.greensoftware.foundation/ (accessed 1st April, 2023)

2.2 Roadmap - how can we do it?

In this Section we discuss the roadmap and "how" the UKRI DRI can implement the recommendations by 2040 or sooner.

The roadmap details proposed actions to be delivered by 2024, 2028, 2030 and 2040. The short-term (by 2024 and 2028) deliverables reflect foundational activity urgently required to support and enable mid-term actions (to be delivered by 2030 and 2040) which are required for a net zero DRI.

Actions have been mapped onto three delivery pathways (Table 2.2.A), which are aligned with the UKRI delivery areas identified in the 2020 UKRI Environmental Strategy. Firstly, in the UKRI Direct Operations delivery area, we have the policy and governance pathway which sets out steps needed to create a policy framework which can deliver the steps needed to achieve the net zero ambition. Secondly, there is a delivery partnership pathway (in the lead procured activity area), in which funders and facility leads or service providers work together to implement the fundamental changes required. Lastly, competitive funding (the "funded" delivery area) develops the necessary capabilities and tools drawing on the creativity, diversity and strength in depth of the UK academic community.

Table 2.2.A: Describing the three delivery pathways: distinct but interdependent mechanisms for affecting progress towards a net zero DRI, which correspond to UKRI delivery areas cited in the UKRI Environmental Sustainability Strategy.

Delivery Pathway	**Description**	**UKRI delivery area**
Policy and Governance	Creating a policy framework which can deliver the steps needed to achieve the net zero ambition.	UKRI Direct Operations
Delivery Partnership	Funders and facility leads or service providers work together to implement the fundamental changes required.	Lead procured
Competitive Funding	Develops the necessary capabilities and tools drawing on the creativity, diversity and strength in depth of the UK academic community	Funded

The proposed actions presented in Table 2.2.B reflect areas of contingent activity, reached through an iterative analysis of all the recommendations which arose from evidence gathering activities (see Section 3). The majority of actions in the roadmap aim to reduce carbon emissions within scopes 2 and 3, associated with activity under the influence of UKRI operations.

There is an interplay between activities in the three pathways. For instance, there is a need to develop monitoring capabilities and community standards through funded projects before effective policies can be designed and implemented. This applies both to technology standards and to standards or norms of behaviour.

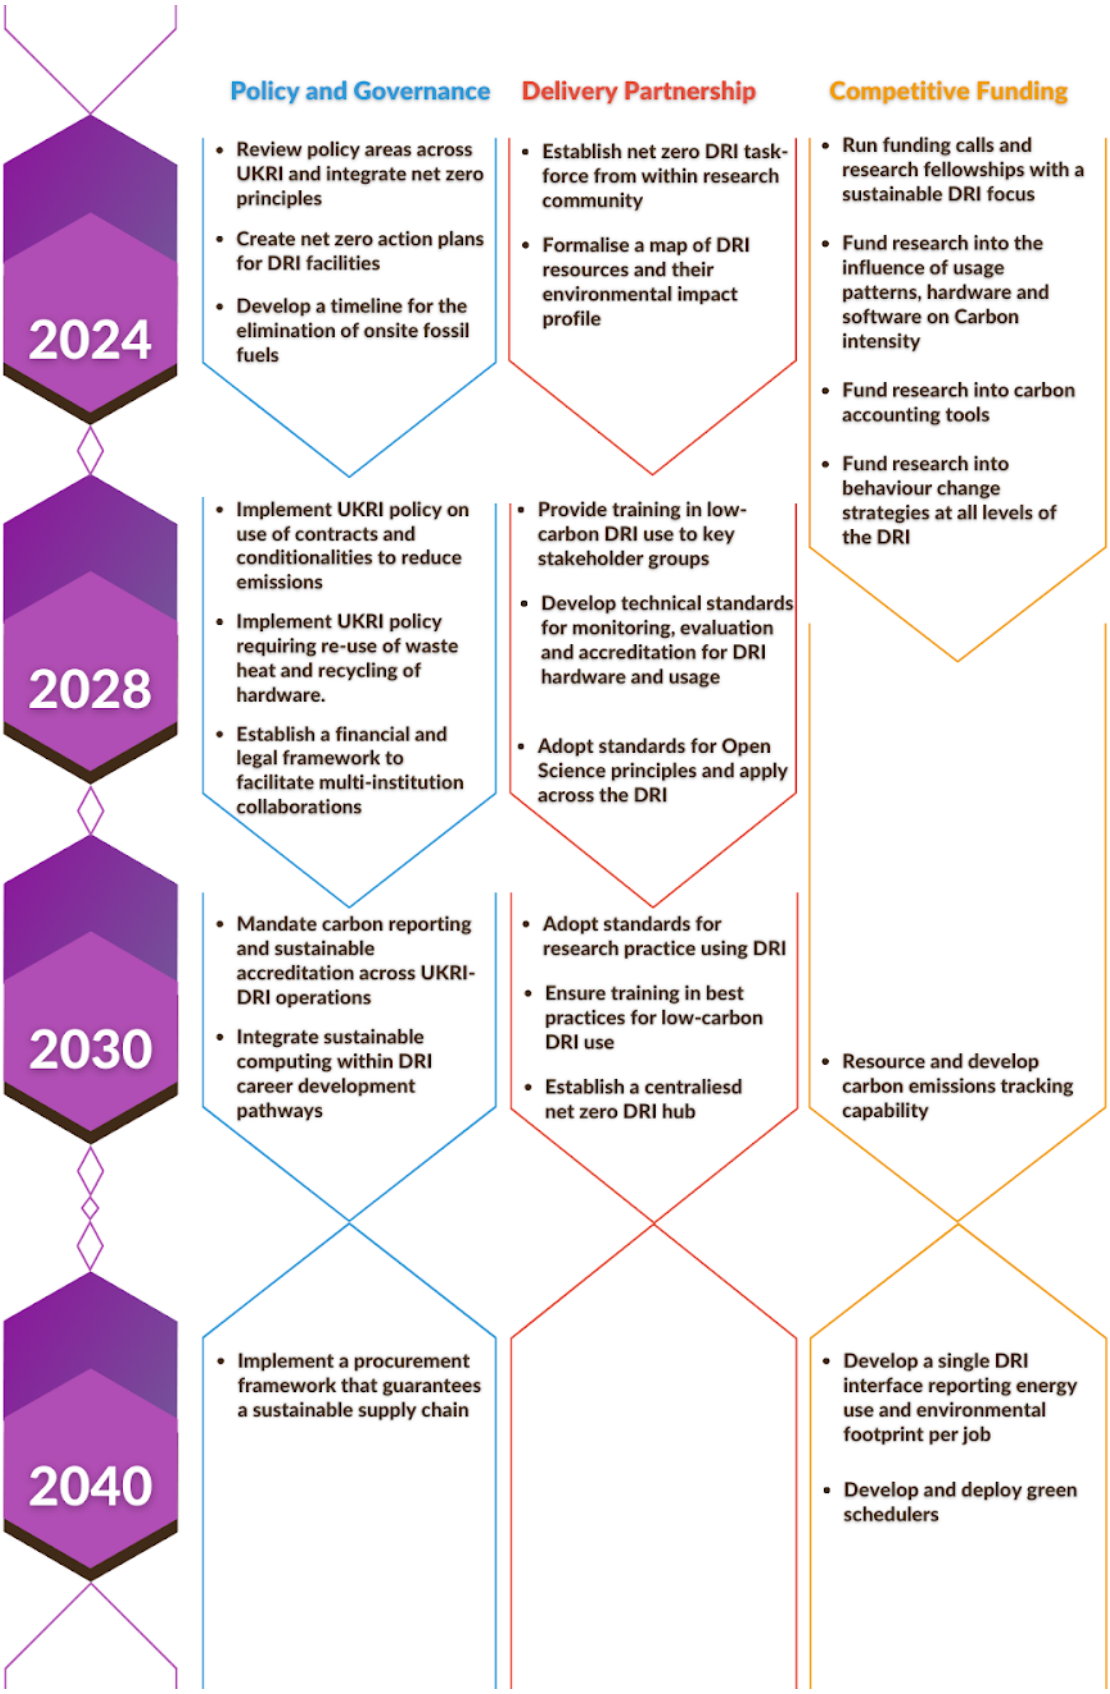

Figure 2.2.A: Graphic illustrating the delivery of roadmap actions along parallel delivery pathways. The year tracks the due date for the delivery of actions, which are interdependent.

Table 2.2.B: *Roadmap actions proposed to reach net zero DRI by 2040 or earlier. Within each time period actions are not indicated in a fixed order. See [Appendix 3](#) for the full list of detailed recommendations.*

Action Number (*)	Actioned by (Year)	Delivery Pathway	What	Outcomes and impacts
1	2024	Policy and Governance	Review policy areas across UKRI and integrate net zero principles	Net zero ambition integrated into all areas of policy resulting in consistent and effective action
2	2024	Policy and Governance	Create net zero action plans for DRI facilities, considering active and embodied carbon, and including decommissioning	Consistent and transparent approach enables rapid transition without disrupting competitive and collaborative working
3	2024	Policy and Governance	Develop a timeline for the elimination of onsite fossil fuels	Clarity about deadlines will result in timely procurement decisions and minimise disruption
4	2024	Delivery Partnership	Establish net zero DRI task-force from within research community to provide guidance and support on policy development and implementation	A task-force with the capacity to maintain current and authoritative information will enable consistent decisions taking advantage of best knowledge available.
5	2024	Delivery Partnership	Create a map of DRI resources and their environmental impact profile	A clear understanding of current footprints, UKRI DRI stakeholders will be able to make effective decisions about priorities for decarbonisation
6	2024	Competitive Funding	Run responsive-mode funding calls and research fellowships focused on sustainable DRI	Creation of a robust body of expertise around green software engineering will enable optimal use of resources.
7	2024	Competitive Funding	Fund research into carbon impact of heterogeneous hardware and associated software combinations and usage patterns	Clear understanding of carbon impact related to heterogeneous hardware/software combinations and use patterns

8	2024	Competitive Funding	Fund research into carbon accounting tools and methodologies	A robust and effective framework for monitoring carbon impacts
9	2024	Competitive Funding	Fund research into behaviour change strategies at all levels of the DRI	A clear understanding of the measures that UKRI must take to facilitate change to more sustainable working patterns
10	2028	Policy and Governance	Implement UKRI policy on use of contracts and conditionalities to reduce emissions in scopes 2 and 3 via procurement contracts and power-purchase agreements	Establish clear and legally binding commitments from suppliers that will take us to net zero in scope 3 emissions
11	2028	Policy and Governance	Implement UKRI policy requiring re-use of waste heat and recycling of hardware.	Minimisation of waste ensures that the net zero strategy is also aligned with other sustainability goals.
12	2028	Policy and Governance	Establish a financial and legal framework to facilitate multi-institution collaborations to support sharing of resources and efficient use of capacity	Greater flexibility leads to enhanced collaboration among UKRI funded institutions.
13	2028	Delivery Partnership	Provide training in low-carbon DRI use to early career and established researchers, research software engineers, facilities managers etc	Habits of efficiency adopted early by new generations of researchers
14	2028	Delivery Partnership	Develop technical standards for monitoring, evaluation and accreditation of DRI usage	Consistent reporting on energy and carbon footprints becomes possible
15	2028	Delivery Partnership	Develop standards for Open Science principles	Open Science becomes common in the UKRI research community, leading to greater research efficiency
16	2030	Policy and Governance	Shift from voluntary to mandatory environmental sustainability reporting and accreditation across: research grant funding, UKRI majority funded facilities, and procurement	A comprehensive approach to reporting gives greater clarity about areas of progress and areas where greater efforts are needed
17	2030	Policy and Governance	Integrate sustainable computing within DRI career development pathways	A workforce motivated to enhance sustainability leads to rapid progress in elimination of waste

18	2030	Delivery Partnership	Adopt standards for research practice using DRI	Open Science enables greater efficiency in scientific workflows
19	2030	Delivery Partnership	Ensure training in best practices for low-carbon use accompanies software, hardware and facility development	When new hardware arrives, the user community is well prepared to make full use of it
20	2030	Delivery Partnership	Establish a central net zero DRI hub to provide support and advice to UKRI research communities and ensure visibility of low-carbon initiatives within DRI	All operators and users will have easy access to consistent and up-to-date information.
21	2030	Competitive Funding	Resource and develop carbon emissions tracking capability, making the connection between user activity and carbon impact	Establishing the capability at a facility or service level will enable design of a system wide approach and start the process of gathering data systematically
22	2040	Policy and Governance	Implement a procurement framework that guarantees a sustainable supply chain; Prioritise sustainability in DRI design and investment decisions and develop DRI resources in partnership with the research communities	Supply chain element of scope 3 emissions reduced to near zero.
23	2040	Competitive Funding	Develop a single DRI interface reporting energy use and environmental footprint per job	Clear and accessible metrics facilitate detection of inefficiencies by users and operators of facilities
24	2040	Competitive Funding	Develop and deploy green schedulers to enable optimal distribution of computational workloads and carbon efficient energy usage across facilities	Reduction in resources wasted by running on inappropriate hardware or underused facilities.

2.3 Principles, Accounts, Issues and Highlights

2.3.1 Emerging Principles to guide the Net Zero DRI journey

The following principles set out fundamental propositions and values aligned broadly with four principles of ethics (beneficence, nonmaleficence, autonomy, and justice):

- UKRI and the UKRI DRI can provide leadership in the national and global transition to net zero as a world class organisation dedicated to transforming tomorrow.
- Transparency and clarity in reporting DRI carbon footprints need to be enhanced at all levels.
- There is growing urgency, driven by the escalating climate crisis and the enthusiasm of the research community for positive change. Action to reduce carbon footprints needs to be taken in parallel with efforts to enhance clarity about the scale of our footprints. Responsibility is distributed, but not diluted, across the organisation.
- Equity and justice are central to the ethos of UKRI and to the net zero ambition.

2.3.2 Carbon Accounts

2.3.2.1 Examples

This section presents a few examples of carbon footprints to give the later discussion some context. Table 2.3.2.A below provides examples of the power consumption and emissions of both facilities and devices. This includes active carbon associated with electricity usage, and embodied carbon associated with the supply chain. Table 2.3.2.B provides examples of the emissions footprint for various activities including uses of DRI (running a search on an internet browser, training a model using climate data, publishing a scientific paper etc).

Item	Annual Power	Annual Scope 2 emissions (kg of CO2e)	Estimated Supply chain emissions	Annualised Supply Chain Emissions (kg of CO2e)	Comment
Laptop	26kWh	8.3	321	80	Based on Dell Latitude E6530, 4 year lifetime. Note, this is for the default system configuration.
JASMIN	1.5GWh	318,000	Not applicable	50,000	Based on UK grid-average carbon equivalent intensity of electricity generation (212g/kWh, BEIS 2021) and sector-typical factor for scope 3 (16%).
IRIS (Low Estimate)	6.4GWh	389,090	-	139,795 (includes a contribution for the	Derived from IRISCAST 24 hour audit of 6 DRI resources with: low estimates of PUE & carbon intensity; low supply chain

				embodied cost of the data centre building)	emissions; 7 year computer lifetime; and 60 year building lifetime. (see Section 3.3.6, Hays et al. 2023)
IRIS (High Estimate)	9.2GWh	3,395,230	-	882,205 (includes a contribution for the embodied cost of the data centre building)	Derived from IRISCAST 24 hour audit of 6 DRI resources with: high estimates of PUE & carbon intensity; high supply chain emissions; 3 year computer lifetime; and 60 year building lifetime. (see Section 3.3.6, Hays et al. 2023)
ARCHER2	20GWh	4,240,000	Unknown		Based on UK grid-average carbon equivalent intensity of electricity generation (212g/kWh, BEIS 2021). This is significantly lower if the local grid carbon intensity or supplier carbon intensity is used. (Section 3.3.5, Turner and Basden, 2022)
UKRI DRI facilities total (low estimate)	100GWh	20,000,000			Based on the sustainable computing report (Section 3.4.5, Vanderbauwhede, 2023), adjusted to use BEIS 2021 carbon intensity of electricity.
UKRI DRI facilities total (high estimate)	300GWh	60,000,000			
Non-facilities, UKRI funded equipment (low estimate	20GWh	6,000,000			
Non-facilities, UKRI funded equipment (high estimate)	70GWh	25,000,000			
Non-facilities, not UKRI funded equipmen	10GWh	5,000,000			

t (low estimate)					
Non-facilities, not UKRI funded equipment (high estimate)	100GWh	36,000,000			

Table 2.3.2.A: Carbon footprint of sample facilities and devices.

Item	Footprint	Comment
Search (service)	0.2g	Single Google search query, ignoring user devices
Search (activity)	7g	User activity involving several queries.
One hour streamed video	16 g	Much higher values are widely cited but are not consistent with overall power draw of data centres.
Regular Latte	400 g	Assuming dairy milk and air-freighted coffee.
Simulating one year with a climate model (low resolution)	1 kg	With around 10^6 mesh points.
Scientific Publication	5.4 kg	
One year of growth by a tree (absorption).	-46kg	Based on 7.3 t/ha/a for poplars at 156 trees per hectare and a lifetime of 26 years (Cannell, 1999).
Simulating one year with a climate model (high resolution)	600 kg	With around 10^8 mesh points.
PhD	21000 kg	

Table 2.3.2.B: Carbon footprints of a range of activities. Further details in the Interim Report (Juckes et al., 2022).

2.3.2.2 Literature Survey

The quality of information available is steadily improving. For instance DELL now provides a detailed life-cycle assessment for some of their servers[33]. For the Poweredge R6515, R7515, R6525, R7525 rack servers they estimate 64-71% of emissions come from the use phase (based on grid-average carbon intensity for the USA and EU in 2020) and the majority of the remainder comes from manufacture. The manufacturing is, in turn, dominated by integrated circuit production, mainly the SSD. The R7515 has an assessed embodied carbon footprint of 1,338 kg CO2e [100 year] and a list price of £3,0026, working out at 0.44 kg per GBP. DELL conclude that energy consumption, waste and emissions of the manufacture of the SSD NAND flash far outweigh the regular metallurgical or plastic production processes seen with other components such as the server chassis[34].

For many devices a simplified LCA using the PAIA approach is used, e.g. the Latitude 7320 Detachable[35] has an estimated embodied carbon of 60kg and a price tag of £1929, giving 0.03 kg per GBP.

Pirsen et al. (2023) provide a comprehensive review of the integrated circuit (IC) supply chain landscape. They aggregate results from a wide range of sources on the basis of IC area and conclude that carbon footprint has decreased only modestly, by 17% over 30 years. This represents a substantial reduction in footprint per transistor as the density of transistors has increased rapidly. However, increase in computational demand has led to a slight increase in IC area produced (a massive increase in the number of transistors).

As hardware is powered with increasing proportions of renewable energy[36] emissions from operational energy consumption (use phase carbon) will reduce relative to emissions from hardware manufacturing. Hardware manufacturing emission will eventually dominate the carbon footprint. For user devices the embodied carbon is generally greater than the carbon emissions during use (Jattke et al. 2020, Clement et al, 2020)

Chip manufacturing accounts for most of the carbon output attributable to hardware systems. As capability increases from one device generation to another, a rising percentage of hardware life-cycle emissions will come from manufacture (Gupta et al., 2022).

Using renewable energy to power chip fabrication facilities will reduce the carbon emissions from hardware manufacturing. However, even under optimistic renewable-energy projections, manufacturing will continue to represent a large portion of hardware-life-cycle carbon footprints (Gupta et al., 2022). On a positive note, suppliers are working to reduce carbon intensity of production, with targets such as 100% renewable electricity supply to factories by 2030 and new zero operation by 2050 aiming for 100% renewables to power factories by 2030[37].

---

33 https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/full-lca-of-dell-severs-r6515-r7515-r6525-r7525.pdf

34 https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/lca-poweredge-r6515-r7515-r6525-r7525.pdf

35 https://www.delltechnologies.com/asset/en-us/products/laptops-and-2-in-1s/technical-support/latitude-7320-detachable-pcf-datasheet.pdf

36 Green Book supplementary guidance https://www.gov.uk/government/publications/valuation-of-energy-use-and-greenhouse-gas-emissions-for-appraisal

37 Intel Corporate Responsibility Report, 2021-22. https://csrreportbuilder.intel.com/pdfbuilder/pdfs/CSR-2021-22-Full-Report.pdf, TSMC Task Force on Climate-related Financial Disclosures (TCFD) Report, 2020. https://esg.tsmc.com/download/file/TSMC_TCFD_Report_E.pdf

Semiconductor manufacturers use a variety of high global warming potential (GWP) fluorinated gases to etch the intricate circuitry patterns on silicon wafers and to rapidly clean chemical vapour deposition (CVD) tool chambers[38]. In 2018 the Taiwan Semiconductor Manufacturing Company (TSMC), the world's largest semiconductor foundry which supplies chips to Apple, reported that whilst 63% of emissions from chip manufacture came from the consumption of energy, nearly 30% of emissions were due to high GWP perfluorocarbons (PFCs), chemicals and gases[39].

Etching and cleaning processes can also be achieved with compounds that have less potent global warming potential and the industry is investigating such options. The switching of these agents must be established during the design phase of the manufacturing plant. It can take up to 18 months to achieve reliable chip manufacture so changing to different compounds mid-manufacture is not commercially viable[40]. Therefore the set-up stage of a chip fabrication facility represents a discreet window of opportunity to implement emission reduction strategies.

2.3.3 Knowledge gaps,ambiguities, and constraints

Key barriers to net zero DRI include:

- **Ambiguity of the net zero concept at smaller scales**

  One person's scope 3 emissions (e.g. supply chain) are another person's scope 2 emissions (e.g. energy), so there can be a danger of double counting. Whilst net zero makes sense on a global scale, carbon emissions, and therefore net zero, can be difficult to quantify at an institutional level.

  There is a further ethical dimension to net zero. Institutional net zero policies should not block the wider societal pathway to net zero, for instance by consuming a disproportionate quantity of limited renewable energy resources.

- **Limitations around the definition and measurement of scope 3 emissions** (those emissions arising through the activity of suppliers, manufacturers and researchers associated with the DRI and beyond the direct influence of UKRI)
- **Limitations around measuring embodied carbon in hardware and facilities construction processes** *(related to the point above)*
- **Lack of available carbon accounting tools for the accurate evaluation of carbon emission**
- **Displacing (renewable) energy supplies from local community assets**
- The UKRI DRI must operate with consideration of local impacts regarding power consumption, storage and supply, as well as responding to the needs of the research community
- **Situational ambiguity** refers to situations where there is uncertainty or confusion about what is expected or required of individuals and can lead to a "fog of enactment"[41]

---

[38] https://www.epa.gov/f-gas-partnership-programs/semiconductor-industry
[39] TSMC Corporate Social Responsibility Report, 2018 https://esg.tsmc.com/book/2019/en/files/downloads/TSMC_CSR2018Highlights_E.pdf
[40] The computer chip industry has a dirty climate secret https://www.theguardian.com/environment/2021/sep/18/semiconductor-silicon-chips-carbon-footprint-climate
[41] https://www.thegreenwebfoundation.org/publications/report-fog-of-enactment/

2.3.4 Tensions between competing objectives

There will be tensions between objectives that could be seen as opposing or competing. In many cases it is not that one set of objectives can or should be chosen in exclusion of other objectives. Rather, for example, we need to find a balance. This is not a modern problem, and is illustrated in Figure 2.3.4.A where it is necessary to steer (via regular reviews) between the extreme of unfettered innovation (Scylla) or of excessive restraint (Charybdis).

This section explores the key tensions discovered during the Scoping Project.

Figure 2.3.4.A: Steering a passage between the rock of Scylla (inadequate steps and the to cut emissions leading to a crash) whirlpool of Charybdis (constraints on innovations leading to a tragic downward spiral). (John Doyle, National Portrait Gallery)

2.3.4.1 Expanding Energy Consumption in the DRI

- **Expanding energy consumption in the DRI (in line with increased adoption of digital methods and developments within computational research)**
- The Climate Change Committee report on pathways to Net Zero highlights the importance of restrained or constrained energy consumption. They recommend a balanced pathway which involves some restraint in energy usage but at the same time requires innovation to power the transition.This balance creates a tension between the need for innovation and restraint on energy use which could dampen innovation. Neither the extreme of unfettered innovation (Scylla in Figure 2.3.4.A) or of excessive restraint (Charybdis) provides a satisfactory outcome. To steer between the two we need to have regular reviews of the course.
- Some contributors have taken it for granted that this constraint will be passed down to the UKRI and the UKRI DRI while others assume that the UKRI DRI is a special case for which, by virtue of the large societal benefits it delivers, should be allowed an exemption from this constraint.
- This reflects a broader tension between strategic objectives and environmental sustainability. It is essential that policy frameworks act to reduce/minimise tensions across UKRI operations with the principle of environmental sustainability and the ambition for net zero within the DRI.

2.3.4.2 Requirements of stakeholders at various scales across the DRI

There are tensions or conflicts between the requirements of stakeholders at various scales across the DRI:

- HPC-Jeep (Turner and Basden, 2022) finding about the CPU intensity of different code bases suggests that at the facility level, such as ARCHER2, managers could schedule CPU intensive software (the kind that needs additional cooling) during periods where the electricity supply from the grid is powered with a greater proportion of renewable energy (less carbon intensive). Conversely less CPU intensive software could be preferentially scheduled during periods of high carbon intensity electricity supply.
- GO Zero (Manika et al. 2023) suggestion of a demand-push policy for the allocation of carbon budgets along with financial budgets to research projects by grant awarding bodies.
- There is potential conflict between the desire to minimise the carbon expenditure on DRI use at the individual project level and decisions made about scheduling jobs to minimise energy consumption at the collective facility level.
- For instance, a climate modeller might wish their simulations to be run during periods of low carbon intensity energy supply to make best use of their carbon budget allocation. Whereas a facility manager might prefer to run climate model software during periods of high carbon intensity because the input output restrictions of the climate model software requires less energy draw for the CPU.

2.3.4.3 Avoiding Air Travel versus Maintaining Communication

Academic air travel is not a major element of the UKRI DRI carbon footprint, but it is a substantial element of the footprint of the research sector as a whole. There are significant challenges to addressing the footprint of air travel, and strong variation in opinions. Efforts to decarbonise the DRI will be undermined if the values and ambitions are not applied equally to air travel.

Braun and Rödder (2021*) to refer to the problem as “elephant in the sky”. Reyes-Garcia (2022) identifies the need for normative standards for travel policies. They propose a net-zero compliant policy which specifies no flights for trips of less than 6-weeks duration. A policy of this form would clearly leave a gap in the academic experience that many have come to expect and depend on, but there could also be benefits. Crumley-Effinger and Torres-Olave (2021) discuss the hypermobile elite and the exclusion of less mobile researchers. They highlight the role of virtual meetings in increasing accessibility and discuss the responsibility of delegates.

On the other hand, one UNFCCC report[42] considers air travel emissions as being unavoidable.

The actual amount of travel by researchers appears to be determined more by social and structural reasons than by attitudes to climate change. The reasons given for travel include a belief in superior quality of face-to-face interactions (Whitmarsh et al. 2020, 2021). Tseng et al. (2022) consider that “it is generally accepted that air travel and face-to-face meetings are absolutely irreplaceable elements of a successful academic career”. They analyse the problem using the Stephenson et al (2010) energy cultures framework, which comprises three elements: cognitive norms (e.g. beliefs, understandings), material culture (e.g. technologies, building form) and energy practices (e.g. activities, processes).

A pilot for reducing carbon emissions for conference travel (Skiles et al. 2021) suggested organising a virtual conference, but having hubs set up in various locations locally/regionally

[42] https://unfccc.int/sites/default/files/resource/CNN%20Guidelines.pdf Climate Neutral Now - Guidelines for Participation (UNFCCC, 2021).

and using them as networking sites where people could attend the virtual conference, but still have face-to-face networking opportunities with other local researchers. The authors write ‘This conference within a conference approach allowed for reduced cost and travel, increased local and regional networking and created an international conference’. See also Pasek et al. (2022) for further discussion.

Alfonso et al. 2023 conclude that Herculean efforts will be needed to achieve sustainable air travel by 2050. Recent progress in electric flight suggests that commercial electric flight, in the form of one or two prototype aircraft with capacity to carry 2-4 passengers tens of kilometres, may commence within a year or two (for example, there may be electric air-taxis at the Paris olympics in 2024[43]). This will be a major milestone, but is clearly some way from scaling to sustainable provision of millions of long distance flights.

Optimists continue to seek novel solutions, such as converting seaweed to aviation fuel[44], but hard constraints imposed by the laws of physics and the finite capacity of our planet point to a need to plan for reduced access to air travel.

At the same time, investment in enhanced digital communication technology and culture can mitigate the drivers for academic air travel.

The ways forward are thus:

1. Promote and enhance pathways for networking and career development which do not rely on frequent air travel, for example by exploiting less-frequent journeys of longer duration and improved online communication.
2. Watch and monitor the advances in low-carbon flight and adapt the 2040 Net Zero roadmap to reflect emerging capabilities.
3. Develop a policy for conference organisation and attendance which incorporates the values and ambition of the Net Zero DRI policy on sustainability and accessibility.

2.3.5 Highlights

We need to start taking action. Within this scoping project substantial gains in DRI computational efficiency have already been achieved, there are more easy gains waiting to be discovered.

Mitigation and resilience is central both to reliable operation in a volatile future and to managing our resource usage effectively.

Net zero is not a milestone or a target, it is the society that we want to build.

There is a great enthusiasm for action in the community, though also some reservations about the pace of change.

The transformation to a sustainable society is a positive aspiration, we need to grasp the opportunity.

---

[43] https://www.barrons.com/news/flying-taxis-star-at-paris-air-show-next-stop-the-olympics-b0839d8f (accessed 2023-06-21).

[44] https://www.bbc.com/future/article/20230323-climate-change-how-seaweed-could-power-planes . Note, however, that seaweed is also being investigated as a potential source of biomass for carbon offsetting and other applications, and tends to grow in areas which are important for global food supply.

3 Evidence

This section of the report contains a summary of all the evidence collected through project activities and provides links to more detailed technical reports which have been published by partners during the project.

The evidence gathering process sought to address the heterogeneity of the physical DRI, and of the user communities accessing this resource. Four types of activities were undertaken: A literature review surveyed current ideas in academic and technical literature; funded and commissioned activities, termed ‘sandpit’ and consortium projects respectively; and stakeholder engagement activities, which were conducted in parallel with evidence gathering projects to ensure that key stakeholders were active participants in developing and shaping recommendations.

A summary of all activities undertaken is provided in Table 3.A. The activities broadly fell into two themes: machines and workflows, and people and process. Figure 3.A shows how the sandpit and consortium projects fit into these themes.

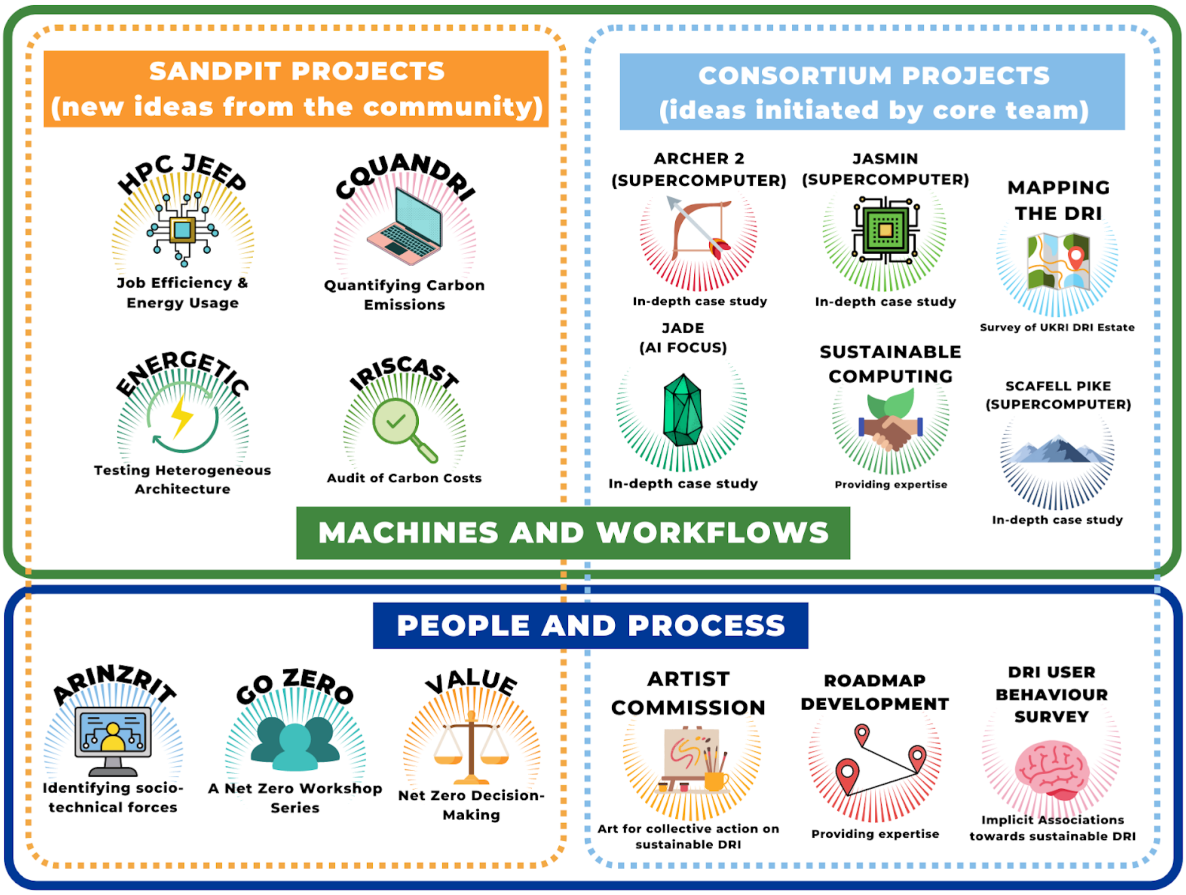

Figure 3.A: Sandpit and consortium projects broadly fell into two themes: machines and workflows, and people and process. The literature review and community and consensus building activities covered both themes.

Table 3.A: Summary of Evidence Gathering Activity. Links to project reports are embedded or can be found within the relevant evidence sub-section below.

Evidence Gathering Activity	Name

3.2 Literature Review	Interim Report (Juckes et al., 2022)	
3.3 Sandpit projects	ARINZRIT (Friday et al., 2023)	
	CQUANDRI (Schien et al., 2023)	
	ENERGETIC (Bane et al., 2023)	
	Go Zero (Manika et al., 2023)	
	HPC-Jeep (Turner et al., 2022)	
	IrisCast (Hays et al., 2023)	
	Value (Boulton, 2023)	
3.4 Consortium projects	Facility Case-studies	ARCHER2 (Smith et al., 2023)
		JASMIN (Lambert, Stephens & Kayumbi, 2023)
		JADE2
		Scafell Pike
	Exploiting batteries to reduce carbon intensity	
	DRI Mapping Database (Stephens, Kayumbi & Lambert, 2023)	
	User behaviour survey (McGuire, 2023)	
	Sustainable Computing (Vanderbauwhede, 2023)	
	Artist Commission	
3.5 Community and Consensus Building Activities	Training and Standards Workshop	
	Procurement Workshop	
	Values and Responsibilities Workshop	
	Early Career Researchers Workshops	
	Community Engagement Workshop	

3.1 Recommendations Table

The recommendations that have emerged from all evidence gathering activities are collated in a recommendations table in the Appendix 3. This includes a full list of the detailed specific recommendations, including citable reference numbers.

The table maps each recommendation to the toolkit theme and roadmap action which it informs. You can also find all recommendations in an interactive spreadsheet document found on Zenodo here: https://doi.org/10.5281/zenodo.8199893

3.2 Literature Review (Interim Report)

The interim report, published in August 2022, contains an extensive literature survey. More recent literature which brings extra value is cited in the synthesis section above.

The interim report is available on the Zenodo platform: "Complexity, Challenges and Opportunities for Carbon Neutral Digital Research" (https://doi.org/10.5281/zenodo.7016951).

3.3 Sandpit projects

What follows is a summary of the aims, methodology and findings of each sandpit project.

3.3.1 Applying Responsible Innovation to the Net Zero Research Infrastructure Transformation (ARINZRIT)

As UK research increasingly relies on DRI, which in turn drives up environmental costs, the ARINZRIT sandpit project (Friday et al., 2023) aimed to investigate the social and organisational factors that shape, and are shaped by its use and provision. The consideration of technical factors was left to other sandpit projects. By understanding the broader landscape of policy, practice and engagement with DRI, this sandpit project aimed to 'map the landscape' of how policies and practices interplay and affect growth in DRI. This section summarises the ARINZRIT final report.

To explore these wider socio-cultural influences affecting DRI provision and use, 25 interviews were conducted across stakeholder groups of users, providers, commissioners, and senior management relating to DRI. The interviewees spanned multiple disciplines, institutions, and research councils and, whilst not fully representative of the sector, capture a variety of experiences that point towards some of the complexities and relationships to be considered in moving towards Net Zero ambitions for DRI. These experiences also demonstrate that interventions should not just be technical.

The ARINZRIT interviews yielded the following key findings.

DRI use is increasing due to advances in science and increased adoption of machine learning techniques in many research domains, including those not previously considered to be computationally intensive. In parallel, the social sciences are developing new approaches using digitised artefacts, and examining social life as it occurs in online spaces. Concurrently students, both undergraduate and postgraduate, are increasingly expected to use DRI as part of their studies, further embedding DRI use into academic tools and practices.

> *"... Students read papers, right, that's part of the work. So, they investigate the field, and [...] the majority of those papers are results from like, big models, [...] they want to compete with that, and publish in a certain conference or journal..."*[45]

Increasing demand drives DRI growth as computation use and the provision of more services creates new opportunities for innovation and knowledge creation. Nationally, demand for services outstrips supply, resulting in DRI expansion and upgrade in line with UKRI funding cycles. At a HEI level, provision is shaped by the need to spend budgets in funding cycles as well as researchers' desire for ownership and control of 'own' facilities (e.g., local workstations, compute clusters). There is currently little to limit this growth in DRI, or its use. Targeted funding and continued embedding of DRI and computationally intensive methods in all disciplines may even be accelerating this.

[45] This and following quotes in this subsection come from the ARINZRIT user survey.

> *"... It's more a case of the more we add, the more it would be used: there's basically more demand than supply. ..."*

> *"...We are generally slightly underprovided I think, and therefore people do feel a little bit restricted as to what they can do..."*

DRI use can be inefficient. Poor software development practice such as lack of code optimisation or codebases poorly attuned to modern hardware lead to inefficiencies. Policy, practice and culture may lead researchers to make inefficient use of DRI for example by using resources at higher priorities/speed than optimal and encouraging replication of computation to meet reviewer expectations. Unnecessary computation where results already exist or where poorly formulated or redundant research questions are posed is also clearly inefficient. Together these factors ultimately decrease DRI's efficiency and increase hidden environmental costs, without resulting in an increase in the knowledge created.

DRI provision can be unsustainable in the sense that some pockets if DRI are under utilised and others are oversubscribed. This inhomogeneous provision can be attributed to new DRI being provided too soon (wasting hardware) or too late (creating inefficiency), and by constraining procurement protocols such as ringfenced clusters for health applications. Provisioning for peak demand periods (e.g., student term times, academic paper deadlines) can lead to overprovisioning as does underutilisation of one DRI resource which may result in the creation of a duplicate resource elsewhere. DRI is underutilised in other ways: for example the absence of waste heat reuse due to capital cost, logistical, organisational or estate related reasons. DRI procurement can also be driven by funding availability with short deadlines to spend the money which can result in allocation of funding to projects that can complete quickly, possibly increasing inhomogeneity, rather than to longer term efforts to rebalance the inhomogeneity of DRI resources.

> "... If there's duplication of platforms or resources, then that's wasteful. If there are inefficiencies, because there aren't ways of joining up different platforms, that's wasteful. ..."

> "... And Estates never hooked up those pipes like it was in the design spec. It's really frustrating. ..."

The growing trend to FAIR data has implications for Net Zero DRI. Data which meet the principles of findability, accessibility, interoperability, and reusability are known as FAIR data. This easier findability and greater potential for reuse of data could avoid unnecessary duplication of DRI usage. However, the process of FAIRification of data requires an ecosystem of supporting tools and services that in turn require DRI resources, such as long-term data preservation, machine processing of data, training and upskilling. The balance between benefits and costs is currently poorly understood: a variety of influencing factors interact together to drive the generation and use of FAIR data, and the balance in terms of net zero between effort put in and benefits derived is still unclear particularly when considering long term factors.

Academic practice shapes unnecessary and inefficient use and provision of DRI through expectations by reviewers and academic communities on the quantity of DRI use implicated in producing work acceptable for publication; local academic norms and 'lock in' to old possibly inefficient software; and a desire for unrestricted use and control leading to the provision and inefficient use of extra local DRI. However, care needs to be taken to balance against innovation in academic practice using DRI that can be said to increase the quantity and quality of research.

> "...the fundamental thing for me, is that academia is still culturally, very much rooted in this sort of 19th century independent scholar, slightly artisanal approach to research. And that's increasingly at odds with the reality of how work is done on the ground, but it's still the culture. ..."

There is a lack of awareness and training toward achieving Net Zero DRI, with its environmental impact being hard to find or contextualise, and therefore, largely unknown among researchers (including RSEs). Where awareness does exist, Net Zero DRI is a low priority and actions are usually grounded on informal knowledge that does not always lead to sustainable, evidence-based change. Time and cost are frequently used as proxies for carbon, but these are inadequate. Users say that better information would help to inform their decision-making on how to use DRI facilities.

> "... So for me, the speed of the run is a proxy for how energy efficient it is. And that might not always have been true, and it probably isn't going to be true going forward. But I don't know, at the moment, I have no way of measuring that. As far as I know, there are no tools that I can access to do that. ..."

There is no clear ownership, oversight, or resource for progressing towards Net Zero DRI, resulting in little overt or cultural pressure to consider or embed Net Zero in DRI activities. Users of DRI individually have relatively little agency in addressing Net Zero and are – generally – embedded in a 'publish or perish' culture, making it critical that research funders and institutions work collectively to establish practice for DRI that is Net Zero compliant. Where cost is incurred, this is often decoupled from the researchers at the point of use, and funders, practitioners and academic communities are not typically holding each other to account for DRI's environmental impacts. Resultantly, there is no obvious ownership, human or financial resource set aside to address this issue. If Net Zero DRI proves more expensive or complex than the default DRI, or there are no resources provided to address the cost and complexity, it is unlikely to be addressed.

> "...there really is zero external pressure. You know, the university when they made the Net Zero announcement. Pretty much nothing changed..."

> "...declaring climate emergencies, but they're not actually putting any money into delivering those things...."

3.3.2 Carbon Quandry with Digital Research Infrastructure (CarbonQuanDRI): Decarbonising UKRI Electricity Use – A Case Study of GridPP

The CarbonQuanDRI sandpit project (Schien et al., 2023) aimed to produce a system dynamics model to evaluate interventions for decarbonising parts of the DRI while taking into account the geographical and temporal variations in the electricity supply carbon intensity. This work was set in the specific context of the Tier-1 GridPP services, operated by the Science and Technology Facilities Council (STFC) at the Rutherford Appleton Laboratory (RAL). This section summarises the CarbonQuanDRI final report.

A run-time power performance analysis of representative production workload at the GridPP Tier-1 facility was undertaken, and a systems dynamics model developed. A new quantitative model of locational marginal carbon intensity was also developed. This was then applied to evaluate some intervention scenarios.

Analysis of the run-time power performance data that was collected at the GridP Tier-1 showed that 6 hourly time resolution on Power Distribution Unit (PDU) power measurements was insufficient for statistically mapping variation in electricity consumption, hence carbon costs, to variations in individual services run by the GridPP Tier-1.

This led to a series of curated workloads being run on two similar servers that had additional power monitoring measuring at 5 second intervals comparing an Intel CPU and an AMD CPU. It was found that the AMD server was more energy efficient than the Intel server used in this experiment, however it is not clear that this can be generalised. Both servers had similar job execution times when using the same number of cores or executing the same number of jobs. Interestingly increasing the number of cpu-threads used reduced the runtime and overall energy consumption per thread, despite the higher power draw. It was found that triggering jobs sequentially or in parallel yields the same energy consumption per job. The energy consumed as a baseline by servers running no useful work is non-negligible. The level of utilisation of DRI alone is therefore not enough to indicate energy efficiency, since the energy consumed even at idle may be high unless managed.

The system dynamics model was used to evaluate a scenario whereby the GridPP Tier-1 over-provisions hardware such that all computing jobs can be paused for an hour at peak local marginal carbon emissions, but get the same amount of computing work done each day as in the standard non-over-provisioned case. Essentially time shifting the work to use less polluting energy. Initial, unverified, modelling indicates that for a range of options there may be significant emission savings based on the local average and marginal emissions factors identified.

CarbonQuanDRI highlights that it is important to use the right Carbon Intensity Metric when considering interventions in terms of carbon cost in an electricity consuming system such as DRI. The grid average carbon intensity gives the electricity carbon intensity averaged across the country. If the UK’s electricity grid is modelled as 29 interconnected regions, as in the 29 bus model used by CarbonQuanDRI, then a Locational Average Emission Factor can be calculated in essentially the same manner as the system average carbon intensity, but region by region (and then accounting for any importing or exporting of energy from/to other regions).

This model was used to calculate the system average carbon intensity for 2021, as well as the locational average and marginal emissions at the location of STFC’s RAL laboratory. The results show that this site has generally slightly higher than average emissions, though during periods of very low demand its emissions are less than average. The locational marginal emissions follow the carbon intensity of different generating technologies, and are often much higher or lower than the average figures. The system average carbon intensity is a simple, perhaps simplistic way to evaluate things at a strategic planning level. The Locational Average Emission Factor is perhaps more appropriate at this strategic planning level, especially when considering the geographical location of DRI resources. For more tactical and operational decisions then a new measure of locational marginal carbon intensity (also known as locational marginal emission factor) is probably more suited especially when deciding things minute by minute or hour by hour.

3.3.3 Energy-aware heterogeneous computing at scale (ENERGETIC)

There is a natural heterogeneity between UKRI DRI resources due to variations in design choices as well as a heterogeneity designed into individual DRI resources due to the use of GPU and FPGA accelerator technologies. The ENERGETIC sandpit project (Bane et al. 2023) set out to investigate whether the use of accelerators would give significant energy savings compared to CPU-only computations, and to determine which classes of code are most suited to accelerator technologies. This section summarises the ENERGETIC final report.

To underpin this work ENERGETIC also investigated methods for measuring energy use in a standardised, portable manner irrespective of architectural platform. The project expected to identify challenges and barriers to energy monitoring at the application level on shared facilities and so from the outset aimed to host a workshop, colocated with CIUK 2022, involving key stakeholders with the intention of disseminating the projects experiences, gaining a better understanding of the current landscape of energy monitoring in compute-focused DRI, and discussing a framework for comparable energy measurements across devices and systems.

Approaches to measuring energy usage on CPU, GPU and FPGA architectures were investigated by a literature review and practical “hands-on” experimentation. Additionally the stakeholder workshop of around 25 attendees discussed the matter of fair energy usage measurements across architectures. This informed the choice of energy monitoring used in the experimental phase of this project.

Four representative benchmark algorithms were chosen, which each had ports to CPU,GPU and FPGA architectures, so that comparative energy-to-solution measurements could be made to find which architecture is most efficient for each benchmark. The list had originally been longer but lack of off-the-shelf working FPGA ports restricted the available choice. The remaining four benchmark codes were insufficient for meaningful mapping to typical work-loads within UKRI DRI thus no energy saving estimates for switching codes from CPU’s to accelerators could be calculated.

Measurements were made on: the ExCalibur FPGA testbed hosted by EPCC; the Myriad cluster at UCL; and on a Linux Workstation at Newcastle University.

The benchmarks were run on a range of AMD and Intel processors, NVIDIA GPUs and Alveo FPGAs.

The ENERGETIC project found that there are a plethora of energy measurement software tools available, and in development, that are predominantly compute-architecture specific. A small selection of these software tools were compared with direct physical measurement of energy usage and in general were found to provide slight overestimates of the energy usage. A generic software methodology for measuring energy that could be used universally across the UKRI DRI would be somewhat challenging given the wide variety of hardware available employed by various DRI resources and would need some sort of calibration or validation.

In spite of disparate seed discussion topics a set of common themes appeared across all three workshop discussion groups.

There was broad agreement that standards for energy measurement across platforms are highly desirable. Since there are various valid choices concerning the specific implementation of energy measurement, it is very important that the chosen methodology is clearly documented. It is often not clear exactly what is being measured by energy measurement tools, or how. An ideal situation would be that all UKRI DRI systems had the same (or equivalent) tools available to consistently compare energy usage across systems. This would need to be mandated in procurement processes and would require support from hardware vendors and integrators, as energy measurement needs to be supported in both the hardware and operating system as well as by application-level tools. Even system administrators may not have access to the information required to monitor energy usage at the level of a particular job, or even a single node.

It was broadly agreed across discussion groups that better tools are needed. There are high-quality tools which support specific architectures, or even multiple architectures, but only one at a time. The user is responsible for implementing whole-node energy measurements (or as close to it as the available tools allow). Integration with the job scheduling tool Slurm (as on ARCHER2) creates a very user-friendly experience, as it requires no additional effort from the user and makes the data available asynchronously within 24 hours of job completion, but does lack transparency about how the data is collected. There was no known or suggested standard way to include the use of shared resources such as the network and storage in the energy usage.

The final theme which occurred across discussion groups was the need for education on all aspects of energy efficient computing. Users who wish to make their own scientific computing 'greener' may be unsure how to make appropriate measurements of their energy usage (in part due to the complexity of tools), how to interpret the data they may gather, or what to do about it once they have the data. One suggestion was the development of 'Green Software Engineers (GSEs)'. Given the finding of the wider ENERGETIC project (Bane et al. 2023) that even on heterogeneous systems parallel efficiency is still a good proxy for energy efficiency, the skills of a GSE would not be dissimilar from those of Research Software Engineers (RSEs) or HPC developers, but there is room for specific training on energy measurement and reduction, and the impact of hardware choices.

The benchmark studies showed that the use of GPU or FPGA accelerators were in general beneficial in terms of energy to solution and time to solution, often significantly so, particularly for large problem sizes: although the gains & optimal choice of accelerator vary depending on the specifics of the algorithm.

The overall greatest improvement in energy efficiency was found when using the FPGA to calculate a 2D Fast Fourier Transform, where there was also a significant improvement in runtime, showing the potential for FPGAs to be both time and energy efficient when programmed well.

GPUs are well known to be efficient at dense matrix algebra and the findings reinforce that showing a nearly 40-fold improvement in the runtime while improving the energy efficiency more than two-fold.

The performance on the FPGA, both in time and energy to solution, was highly sensitive to the configuration/optimisation of the available software port. Thus, whilst there are energy savings observed by use of FPGA compared to either CPU or GPU, further savings could be expected given further optimisation of the FPGA ports for the given FPGA cards used. It is likely that CPU and GPU ports have had more time and community engagement in code optimisation than for FPGA ports.

There is clearly scope for accelerators, both GPU and FPGA, to significantly reduce the energy-to-solution for many computational problems. More skill is required to port code from CPU to GPU and further still to port and optimise on more specialist FPGA accelerators. Most current DRI accelerated HPC systems contain only GPUs which are significantly simpler to program using current development stacks than FPGAs. Currently access to specialist GPU and FPGA programming skills are not readily available to researchers from within UKRI.

3.3.4 Giving Voice to, and Empowering Stakeholders of UKRI DRI: A Net Zero Workshop Series (Go Zero)

The Go Zero sandpit project (Manika et al. 2023) acknowledged that much of the knowledge and many of the skills needed to solve net zero challenges are already in the community but the challenge is to enable mechanisms for sharing those skills and enabling behaviour change. Go Zero led a series of three participatory workshops with the aim to connect and enthuse the DRI community, empowering stakeholders to take ownership of the path to net zero. This section summarises the Go Zero final report.

Having obtained ethical approval, snowball sampling techniques were used to seek expressions of interest from 60 people across the UKRI and JISC[46] DRI communities. A set of incentives were offered to encourage participation. Approximately half of those approached expressed an interest in participating via the first of four online surveys. Each 4 hour online workshop was designed to include three expert speakers to help frame the workshop discussion. There was some overlap in the roles of expert speaker and participant. The Go Zero team numbered 7 and at least 6 from the team attended each workshop. The

[46] https://www.jisc.ac.uk/

total attendance of the sequential workshops naturally decreased due to demands on participants' diaries.

The first workshop had a total attendance of 36 and addressed ideas around where does the energy consumption come from and how can it be made more efficient. The second workshop had 28 attendees and addressed ideas around what to do with an unused carbon budget allocation (“fossil fuel left over”), how to reduce unavoidable carbon emissions (“limit unavoided emissions”) and how a “green scheduler” might change DRI users behaviour. The third and final workshop was attended by 21 individuals and considered what key actions will need to be taken to meet net zero targets in the context of UKRI DRI: ranging from personal action plans through to strategic actions and potential recommendations.

Fourteen participants attended all the sequential workshops. Following each workshop an online survey using likert scale items was used to evaluate the workshop and to track changes in participants’ views regarding: their comfort engaging with the UKRI DRI community; their perceived level of net zero importance for UKRI and for the their job role; their perceived knowledge on the topic of net zero DRI; and their intentions and ability to take action to meet net zero targets.

The workshop findings are summarised below.

Data triage and the correct choice of data storage medium, be that solid-state disk, conventional hard disk or magnetic tape has an impact on the carbon efficiency of data storage due to the nature of the embodied and operational carbon emissions for each media type.

Web front end database applications have potential for carbon optimisation by reviewing portal architecture and improving the efficiency of database queries and through the judicial use of caching.

To improve the efficiency of high performance computer (HPC) systems some additional restrictions on HPC user behaviour may need to be imposed, for example time to solution may be extended to schedule jobs to run when the energy grid has a lower carbon intensity. It may be more carbon efficient to use a centralised computing resource than a local one.

User behaviour change is difficult to achieve but crucial in the pursuit of DRI Net Zero as it is users who will need to make improvements to how code is produced, reviewed, tested, deployed and run to make processing more efficient. This will require training users and communities of users in the importance of good practice with regards to carbon efficiency. Users will also need feedback from energy/carbon monitoring systems in terms they can understand to help drive a user culture to shift away from running code as fast as possible towards understanding of the need for efficiency. If meaningful feedback and monitoring tools are made available to users then people will likely innovate and compete to reduce their carbon impact within a green computing community. Well defined communication activities that will create synergies between actions, giving communities tools and the motivation needed to push actions will also drive user engagement making the big changes feasible.

A green computing campaign would potentially be an effective method to highlight the carbon impact of computing to encourage all researchers to think before deploying their codes, and re-using existing data rather than generating new datasets to reduce carbon emissions.

Improving research software engineering (RSE) skills is needed to evolve software development to increase job efficiency by both improving code, development operations and approaches for running code.

Electricity procurement frameworks currently do not include carbon impact details which will be needed in future to make informed decisions regarding carbon efficiency.

The operational carbon emissions of an HPC system must be considered during procurement alongside the embodied carbon emissions of the equipment and the total carbon cost of associated cooling and auxiliary equipment. This consideration of lifetime embodied carbon costs may result in decisions to increase the lifetime of existing HPC equipment in order to increase embodied carbon efficiency, possibly through optimised maintenance regimes.

Funding councils could ask research proposals to include a consideration of their carbon impact.

Monitoring and assessment of DRI usage may lead to and then regulate a policy of rationing compute resources.

Policy and standards to allow a common language and a common accounting basis are needed, a “global criteria for carbon intensity” if you will. This will enable comparative metrics such as emissions per research paper or emissions per DRI resource or facility to be established. This enables transparency and community coherence but would also drive allocation of accountability and responsibility. These common terminologies, standards, frameworks, and education & training resources are also required, to bring the community to a collective understanding at all levels, which will be important to make large changes feasible.

The allocation of carbon budgets along with financial budgets to research projects by grant awarding bodies could act as a demand-push policy. Predicting the carbon consumption of a project might not be as straightforward as financial budgets, especially as the carbon intensity of the UK electricity grid is constantly changing. UKRI will need to develop tools to allow applicants a transparent and clear carbon content assessment of their proposals and projects to underpin this.

There could be systems to reward work that reduces carbon footprints, and these systems should be clear at the time of awarding funding, for example by allowing grant holders to roll over unused financial budgets where there is a commitment to use such funds to further reduce carbon emissions.

A database for green computing could be set up to act as a repository of detailed documentation about what has been done to increase energy efficiency in past projects. Such a resource could prove a valuable starting point for DRI stakeholders looking to optimise new codes. Similarly to an ethics review or animal use review in research applications, consulting this database for green computing could be made a mandatory part of the grant application process to show that applicants are familiar with established measures to decrease the footprint of their intended computations.

The development of a “green scheduler” could enable time shifting of electricity usage to times of low carbon intensity. With the incorporation of both operational and embodied carbon cost data into the scheduler then carbon accounting per job, project or user could be implemented and carbon budgets could even be enforced at the scheduler level with the introduction of carbon quotas or caps. Making such a scheduler a community-driven development would increase DRI stakeholder engagement and drive trust in the system compared to a top-down imposed solution.

Trying to reduce demand for new hardware and stretching lifetimes of existing hardware should reduce the impact of embodied carbon costs. As the electricity grid reduces its reliance on carbon based generation the operational carbon costs will reduce and the embodied carbon cost will increase in significance. This will need adjustment to institutions’ procurement and maintenance policies and associated IT and business policies to clearly define this aim.

Some of the changes needed to achieve Net Zero must be community-led, with groups coming together to plan to make short-term savings. Individuals and communities need the time and space to explore what is possible, as this can drive the less motivated and build capacity and momentum.

A big part of the UKRI DRI drive and motivation must be targeted funding for Net Zero projects, as those are already driving behavioural change in parts of the community, proving that this incentive is effective. In the long term it is critical to embed an environmental impact carbon footprint assessment on every research proposal with clear criteria, mandatory reporting, and auditing of the entire process.

The development of individual action plans in the final workshop found that the 5W's technique of asking What, Who, When, Where and Why helped the participants structure their plans in an actionable way which importantly identified who would be responsible for an action and why it would work along with potential barriers or obstacles that would need to be overcome. These action plans highlighted the need for implementing carbon-aware strategies for hardware replacement, implementing power monitoring infrastructure, presenting power consumption data or carbon usage estimates to users, and using RSE to review codes and workflows to improve carbon efficiency.

On the face of it, the pre and post workshop survey data for all participants seems to indicate that there may have been an increase in participants' knowledge and intention to take action as the workshops progressed. However, a rigorous statistical evaluation of the data from the 8 participants that completed all four surveys found that the sequential workshops were effective in advancing knowledge, but not in advancing intentions to take action. Go Zero asserts that this proves that sequential workshops as a community engagement mechanism can help DRI stakeholders build and advance knowledge which later in turn could potentially influence their intentions to take action to meet net zero targets. However, Go Zero caution against assuming that knowledge always leads to behaviour change.

3.3.5 Energy Usage on ARCHER2 and the DiRAC COSMA HPC services (HPC-JEEP)

The HPC-JEEP sandpit project (Turner and Basden, 2022) investigated what level of energy information could be extracted from current and historical per-job data from the national ARCHER2 and DiRAC services; and how this data can be analysed and synthesised to provide the information required for funders and researchers to allow them to make informed decisions about how to manage HPC resources to extract the maximum research benefit in the most emissions efficient way. This section summarises the HPC-JEEP final report.

Grant allocations of national supercomputing and HPC services such as ARCHER2 and DiRAC have traditionally be awarded in units of compute time (e.g. node-hours) with users not informed, nor paying much attention to their energy use; and services not generally reporting breakdowns of energy use other than for total electricity charges. However ARCHER2 and DiRAC COSMA services already collect power and energy measurements for operational reasons. That information along with existing job accounting data means there is scope to reprocess this into useful energy metrics to enable informed decisions on the path to Net Zero.

ARCHER2 is a UK national supercomputing service funded by UKRI and hosted by the University of Edinburgh's EPCC. The ARCHER2 system is a liquid-cooled HPE Cray EX with 5,860 compute nodes interconnected by 768 HPE Cray Slingshot 10 switches backed with 12 PB of Storage. Jobs are scheduled on ARCHER2 using the Slurm Workload Manager.

Nominal power draw estimates show that the ARCHER2 compute nodes comprise 83% of the total energy usage, the interconnect fabric 14%, cooling distribution units 2% and storage 1%. Significantly this excludes external plant such as cooling infrastructure. Idle compute node power draw is around 40% of the 600W peak.

Baseboard Management Controllers (BMCs) on the ARCHER2 compute nodes provide energy data which is read by the Slurm “pm_counters” energy plugin. The plugin aggregates data from across all relevant nodes used by a job and stores the overall node energy usage for a job in the Slurm accounting database. It should be noted that other system energy usage is ignored on the basis that the node energy usage is 83% of the system usage.

A python program was written to query the energy and other job accounting information from the Slurm accounting database. Existing mapping of executable names to software packages allow categorisation of job types by software and research area.

The DiRAC Memory Intensive system, COSMA is a UK national HPC facility funded by UKRI hosted by the Institute for Computational Cosmology at Durham University. The current cluster COSMA8, comprises 360 Dell compute nodes interconnected by 44 HDR 200 Gbps InfiniBand switches backed by 6.5 PB Storage.

In-service energy usage measurements show that the COSMA8 compute nodes comprise 80% of the total energy usage, the storage 9%, interconnect fabric 7% and cooling distribution units 0.5% with 3% remaining due to ancillary systems. Significantly this excludes external plant such as cooling infrastructure. Idle compute node power draw is around 17% of the 900W peak. Jobs are scheduled on COSMA8 using the Slurm Workload Manager.

Energy used by a job is calculated by querying the cumulative energy counter provided by the nodes baseboard management controller (BMC). The software “ipmitool” can be used to read out this counter

Baseboard Management Controllers (BMCs) on the COSMA8 compute nodes provide a cumulative energy counter. This counter is read at the start and end of a job using the ipmitool software tool and the difference provides the energy used by a job. The precision with which the standard ipmitool reports the cumulative energy counter is insufficiently precise so COSMA8 uses a modified version of ipmitool to improve this precision. Additionally due to a firmware bug the BMC over-reports the energy consumption by 50% so, until this is fixed by the vendor a correction factor is applied to the raw data.

ARCHER2 data for an 8 month period starting in December 2021 showed that over 90% of the total energy use was due to the top ten energy consuming research areas. Further analysis of these ten areas shows that as expected, there is a close correspondence between node hour use and energy use - the more compute resources you use, the more energy you use. However, there are some cases where the energy use is more or less than would be expected based on the node hour use. In particular: Climate Science has a significantly lower energy use (approx -20%) than node hour use; Combustion has a significantly higher energy use (approx +10%) than node hour use; and Astrophysics and Cosmology have a significantly lower energy use (approx -20%) than node hour use.

Considering that for Climate Science >90% of usage is attributed to the software “UK Met Office Unified Model” and that for Combustion ~70% of usage is attributed to SENGA and ~20% to OpenFOAM, these three software packages were analysed further. Further analysis of Astrophysics and Cosmology software packages was not possible as those are predominantly user compiled code.

Code profiling would be required for a full analysis but with knowledge of the code and previous research HPC-JEEP suggest that the “UK Met Office Unified Model” is probably I/O bound (waiting for access to storage) which will allow nodes to lower CPU energy consumption while waiting for I/O operations to complete. This would explain the lower “energy density” that was found for the “UK Met Office Unified Model” and hence for Climate Science research.

Similar first order analysis of the SENGA and OpenFOAM software leads to the suggestion that these codes are both memory bound, a condition known to let CPU’s make use of turbo-boost capabilities which increase power draw, but perhaps without significant improvement in performance. This would explain the higher “energy densities” found for these 2 software packages and hence for Combustion research. This also leads to the suggestion that for applications with these characteristics, it may be possible to reduce the energy use without impacting performance.

DiRAC Cosma8 energy usage data for Q2 2022 revealed a wide range of average power usage per node: from 210 W to 860 W. It is suggested the lowest “energy density” jobs come from codes not being used correctly. For example submitting a single core job to a full node, leaving much of the node idle. Similarly it is suggested that the highest “energy density” jobs are probably CPU intensive or MPI processes which spin or poll while waiting for data (rather than being interrupt-based). At both extremes there may be quick wins by educating users on how to correctly submit jobs and correctly set MPI flags for proper use of interrupts.

The compute nodes in the ARCHER2 and DiRAC COSMA8 clusters each contributed over 80% of the power consumption, excluding cooling plant, for the relevant clusters, this can reasonably be generalised to most clusters. It is therefore important to understand why the idle power draw of compute nodes on ARCHER2 (200 to 250 W) is so high compared to the idle power draw on COSMA8 (150 W): work is ongoing to understand this.

HPC-JEEP have found that by developing some analysis tools raw data from the Slurm scheduler and IPMI power monitoring can be used to to produce useful energy usage analyses and, for COSMA energy usage summaries can be emailed quarterly to users and project PI’s to raise awareness and provided information about their net-zero impacts and inform possible future allocation estimates.

We can already envisage how such data could potentially be used by service operators to make decisions on how services could be provided. For example:

HPC-JEEP set out three examples of how these new insights into energy usage could potentially be used by service operators to make decisions on how services could be provided. The examples are:

- Research areas/applications that are energy dense per unit of time (e.g. Combustion research on ARCHER2) could potentially be preferentially scheduled to run when more free cooling is available (e.g. overnight or during the colder months of the year) to reduce the cooling overheads for the service during warmer periods.
- Regional carbon intensity forecasts for the energy supply to large scale HPC facilities could be used to reduce emissions by preferentially scheduling high energy density usage during periods of forecast low carbon intensity (e.g. when a high proportion of electricity comes from wind power) and preferentially scheduling low energy density usage when the forecast carbon intensity is higher.
- During extreme heat events where cooling large HPC systems can become difficult, use of the service (or a portion of the service) could be limited to those software/research areas that are less energy dense. This would preserve the utilisation and impact of the services without requiring nodes to be taken out of service.

It is also clear that a better understanding of why different applications have different energy density profiles and what the performance variability is when the energy draw is capped is required to make operational decisions and to provide useful advice to researchers on how they can be energy-efficient in their use of HPC systems.

In the HPC-JEEP analysis of energy usage it was found that the cpu node energy consumption was the overriding factor, however the analysis excluded external equipment such as UPS’s and cooling equipment. A full analysis of energy usage that includes external equipment is clearly needed to see if these are a significant factor or not. It may be that these factors do not add significant overheads or are only of significant interest for particular restricted time periods (e.g. periods of high external air temperatures), but they do need to be measured and analysed for their impact to be understood.

3.3.6 IRIS[47]-Carbon-Audit-SnapshoT (IRISCAST)

Taking a learning by doing approach the IRISCAST sandpit project (Hays et al. 2023) set out to conduct a carbon audit across a number of heterogeneous UKRI DRI resources taking account of the embodied carbon and operational carbon contributions via a carbon model. This section summarises the IRISCAST final report.

IRISCAST set out to identify where the carbon cost sits across a large multi-site infrastructure including the embodied carbon costs of both infrastructure (such as computers and data storage devices) and housing (such as data centre buildings and equipment). A carbon audit of a number of DRI resources was to be conducted simultaneously across sites during three consecutive 8 hour snapshot periods. They aimed to: identify factors that influence carbon cost; find relevant, useful, and practical metrics; and use these metrics as inputs to a carbon model that can be used to inform decision making. IRISCAST sought to capture insights and lessons learned to feed into future larger scale projects towards carbon cost modelling, reporting, and NetZero.

From the outset IRISCAST realised that they needed to engage the community to conduct the audit and then to interpret and communicate the findings, and so two workshops were undertaken.

A stretch goal was to start to contribute to the development of dashboard tools to allow easier assessment of carbon costs.

IRISCAST reached out to a number of potential partners and hosted an in person workshop at QMUL in September 2022 to discuss and plan the technical aspects of the carbon audit. A set of data collection objectives was agreed. In broad terms each site was tasked with collecting an inventory of equipment at the facility, enclosure and node level. Each site was also tasked to collect available power usage information at the facility, enclosure, and node level. Available job or payload level information was also requested. Due to the short time frames and the volunteer technical effort the data collection was limited to information that was already readily available, or that could be obtained within the modest resources and effort available. The technical collection and delivery of the data was then delegated to sites with sites deciding exactly which equipment was in scope for the audit.

In parallel a carbon model was developed to turn the raw data into KgCO2e (Kilogram Carbon-dioxide equivelant).

3.3.6.1 IRISCAST Audit

Six UKRI DRI resources contributed to the IRISCAST audit.

The QMUL GridPP T2 cluster deemed four racks in scope for audit. 4x Mellanox SN2410 switches; 12x APC APDU9953 intelligent Power Distribution Units (PDUs); and 118x Dell PowerEdge R640 compute nodes were evenly distributed across the 4 racks.

SNMP was used to collect cumulative energy readings from the PDUs thus measuring power into the racks. Baseboard Management Controllers (BMCs) were queried via FreeIPMI to gather cumulative energy consumption for each compute node. Direct measurement of switch energy consumption was not possible. At the sub-node level turbostat was used to gather energy consumption of CPU and RAM using RAPL facilities. Job level information was obtained from the Slurm scheduler logs. Slurm was not configured

[47] eInfrastructure for Research and Innovation for STFC https://www.iris.ac.uk/

to collect job energy information. However, there was no capability to measure energy usage at the facility level; neither the air conditioning electricity feed nor the server room electricity feed had any metering available. Data was logged and an Inventory of equipment was collated.

The Imperial GridPP T2 cluster is housed in a shared data centre outside of the cluster administrators’ control and so there was no capability to make measurements at the facility level. PDU measurements were not routinely captured and were unavailable for the snapshot periods. In scope for audit were 241 nodes broadly comprising seven models of hardware, having been procured in batches over a number of years from Dell and HPE among other vendors. A python script used IPMItool to gather instantaneous power usage from each node's BMC and logged that with operating system load average and CPU utilisation. An inventory of equipment was collated.

Cambridge University's Research Computing Services hosts 60 IRIS funded Dell PowerEdge C6320 compute nodes which were in scope for audit. Prometheus Redfish was used to gather power usage data from each node's BMC. CPU and RAM usage data was also captured using Prometheus Node Exporter along with CPU usage and Idle time. An inventory of equipment was collated. Facility level energy usage including cooling was not available.

The DiRAC COSMA7 and COSMA8 clusters hosted at Durham University participated in the IRISCAST audit. COSMA8 was also part of the HPC-JEEP sandpit project described earlier (Turner and Basden, 2022). Electricity feeds to each rack were monitored and distributed by PDU’s. Measurements of power or current were retrieved over ssh. Node level data was collected from BMCs via the IPMI protocol from the 452 Dell PowerEdge C6420 COSMA7 nodes and 360 Dell PowerEdge C6526 COSMA8 nodes. The COSMA clusters are managed by a Slurm job scheduler which was configured to collect energy usage per-job. Job level information including this per job energy usage was extracted from the Slurm accounting database. Facility level energy usage including cooling was not available. An inventory of equipment was collated.

The Scientific Computing Application Resource for Facilities (SCARF) cluster run STFC at the Rutherford Appleton Laboratory participated in the IRISCAST audit. The 27 racks that house SCARF are powered by 75 PDUs of varying models: APC AP8653, APC AP7953 and APC AP8953. These were queried over SNMP. Mellanox Network switches reported energy usage over SNMP while Dell and NETGEAR switches did not have this facility.

IPMItool was used to make instantaneous power readings from BMCs on the 571 SCARF nodes of 11 different Supermicro or Dell models purchased between 2015 and 2021. This heterogeneity highlighted that IPMI is not the same on all models and different models/BMC’s provide different energy or power readings. The instantaneous power reading from the IPMItool was used as it is the most consistent across all node types. Job level information was queried from the Slurm scheduler accounting database. Slurm on the SCARF cluster is not configured to collect job energy usage. Facility level energy usage including cooling was not available. An inventory of equipment was collated.

The STFC Cloud hosted at RAL is a private OpenStack cloud service for STFC and IRIS users to create and run virtual machines (VMs). The STFC Cloud participated in the IRISCAST audit and is co-located with the SCARF cluster described immediately above. Facility level energy usage including cooling was not available.

The 30 racks that house the STFC Cloud are powered by 93 APC AP8953 PDU’s. These were queried over SNMP. The 31 Mellanox switches also provided energy usage data over SNMP while the 41 NETGEAR switches did not. IPMItool was used to query the instantaneous power from the 582 hypervisor nodes, 103 storage nodes and 21 auxiliary nodes. At the payload level which for a cloud is the VM level IRISCAST could not find a tool to monitor individual VM energy usage. PowerTop seemed promising but can only really be used on laptops. Prometheus should be investigated in future. An inventory of equipment was collated.

The Opensearch dashboard system at RAL was used to collect the data for IRISCAST. The STFC Cloud data was already there via the data pipeline briefly described above while other data was imported from site supplied CSV, XML and json files as needed. Additional data analysis was conducted in Python using pandas dataframes.

3.3.6.2 IRISCAST Carbon Model

IRISCAST went on to define a carbon model to calculate the total carbon cost (C_t^p) attributable to operating a DRI resource during a specific time period. This is in essence simply the sum of the active/operational (scope 1&2) carbon cost (C_a^p) for that period plus the embodied (scope 3) carbon cost (C_e^p) apportioned to that period.

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

The complexity comes when formulating what is included in C_a^p and C_e^p and how they are derived. For active carbon costs, C_a^p, the difficulty for this metric is deciding and defining which resources are included in the DRI, apportioning the percentage of resources shared by the DRI and other infrastructure, and defining the scope of resources. The computer nodes involved in a specific DRI are relatively straightforward to identify as is the network equipment, although a little more care may be needed in identifying the network demarcation point. The wider campus network and internet were deemed out of scope.

Facility carbon costs are also included in C_a^p where these facilities are identified as: Cooling systems for the DRI resources and buildings hosting the DRI resources; power distribution units and transformers supplying DRI resources and infrastructure; uninterruptible power supply (UPS) resources supporting DRI systems; Facility electricity usage, such as lighting, fire and security systems, as well as other ancillary systems within the data centre/building hosting the DRI resources. Clearly carbon costs of shared resources in this list must be apportioned to DRI and non-DRI uses as appropriate for that facility and adding some uncertainty to the result.

Putting this together gives:

$$C_a^p = \sum_{1}^{nodes} \sum_{1}^{p} C_{anode} + \sum_{1}^{networks} \sum_{1}^{p} C_{anetwork} + \sum_{1}^{facility\ items} \sum_{1}^{p} C_{afacilities}$$

For all of these carbon calculations we can consider the following form:

$$C_{ax}^p = E_x^p \times CM_e^p$$

where C_{ax}^p is the active energy carbon impact of item x, E_x^p is the energy used by item x during the period p, and CM_e^p is a factor to convert the energy used into carbon equivalent units derived from the electricity/energy supply mix used by item x for period p.

This allows a simplification such that

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

Where E_{nodes}^p is the total energy used by all the nodes in the period, $E_{network}^p$ the total energy used by DRI network in the period, and $E_{cooling}^p$, E_{power}^p and $E_{facility}^p$ are the total apportioned energies for cooling, power distribution and other facility support in the period.

The embodied carbon costs require the consideration of the carbon emitted in creating the resources in the first place and can be similarly represented as:

$$C_e^p = \sum_{1}^{nodes}\sum_{1}^{p} C_{enode} + \sum_{1}^{networks}\sum_{1}^{p} C_{enetwork} + \sum_{1}^{facility\ items}\sum_{1}^{p} C_{efacilities}$$

For embodied carbon, C_{enode} is the carbon emitted in creating, delivering, and installing a given node, likewise for network components and for the facility components discussed earlier. Embodied carbon is not dependent on the operational time period under consideration as the actual emissions happened at the time of creation of the resource. However, it is sensible to apportion those embodied carbon emissions to the period under consideration in relation to the expected lifetime of the resource. For instance a computer with an expected lifetime of 5 years would have 20% of its embodied carbon emissions apportioned to each year of its life. However a building with an expected 20 year lifetime would only have 5% of its embodied carbon cost apportioned to each year. For shared resources an additional apportionment to each shared use should be made.

3.3.6.3 IRISCAST Findings

A community workshop hosted in January 2023 by the Institute of Astronomy in Cambridge, allowed IRISCAST to disseminate their work and findings and to learn further lessons from discussion with the community.

IRISCAST's approach of, on the whole, delegating measurement techniques to sites meant a range of techniques were used and similarities and differences between sites could be noted and many lessons could be learned.

It was found that the energy usage of cooling and other facility items was poorly understood at sites. Estimates of a range of power usage effectiveness (PUE) were used to evaluate the IRISCAST carbon model but clearly sites may wish to add an electricity meter to cooling installations to better understand this.

At the rack level PDU's queried over SNMP was the predominant method of electricity monitoring. At the node level BMC's queried over IPMI was the predominant method of electricity monitoring. However, some sites measured cumulative energy usage and others measured instantaneous power. Cumulative energy usage is a more appropriate and robust method of data collection for carbon audits while instantaneous power may be more suited to operational monitoring. Where PDU and IPMI/BMC data were both available it was seen that PDU readings were between 1.5% and 20% higher than the IPMI data. Some further work would be needed to fully understand this variation and to understand the accuracy and precision of the readings.

The three 8-hour snapshots showed energy usage was fairly throughout the 24 hour period, with some variation between sites but no overall trend emerging. STFC Cloud data allowed comparison between storage and compute node energy usage showing that in total the storage nodes used roughly half the energy of the hypervisor nodes during the snapshots.

Slurm was a popular choice of scheduler at the IRISCAST sites, however only the Durham clusters had Slurm configured to collect job-level energy usage. Clearly configuring schedulers to collect this information where possible is desirable.

IRISCAST found some evidence that the CPU energy usage was the main driving factor behind node energy usage. The snapshot nature of the IRISCAST methodology combined with the relatively high usage levels of the DRI resources measured meant that node idle power could not be calculated from the data gathered. It was also evident that a good understanding of idle power draw is important along with resource utilisation measures as the idle power emissions should really be apportioned to the jobs that do run on a DRI resource.

The aim of IRISCAST was not to provide highly accurate carbon audits of a small number of DRI resources on a particular day, but to establish the methods and processes to do that in future and to understand how changes or variations to the system will affect the carbon costs. A range of three Electricity supply carbon Intensities were chosen (low, medium and high) based on the variation in the UK Grid carbon intensities over the month of November 2022. Direct energy measurements of nodes and where available network equipment were used but as facility cooling data measurements were not available three scenarios of low,medium and high PUE were chosen to allow estimates of the total active carbon cost to be calculated.

The embodied carbon costs were evaluated based on the inventories of equipment collected from each site. Manufacturer datasheets were consulted where available. The availability of this information from the vendors in a usable form being identified as a difficulty. Having reviewed a number of server datasheets two scenarios of embodied server footprints were chosen and these were apportioned over each of five server lifetime scenarios of 3 to 7 years.

No site specific data for building embodied carbon was available so by estimating the floor area occupied by each DRI resource's racks and taking a typical per floor area embodied carbon cost from the literature and taking a 60 year building lifespan an estimate for apportioned building embodied carbon could be made. Due to the long lifetimes the building model showed that the building embodied carbon had a low impact on the DRI carbon cost.

Combining the above partial models, the complete carbon model for each of the scenarios showed that there was an order of magnitude difference in carbon emissions between the worst and best scenarios, which means there is scope for improving existing DRI carbon costs. This and the hitherto rapid rate of change in Information Technology with equipment renewal cycles of typically 5 years means that there is the potential for rapid improvements across UK DRI.

IRISCAST found that volunteer sites have been keen to get engaged and further their own net zero journey.

Sites did not have a good understanding of the energy-usage/PUE of their cooling systems nor those systems' embodied carbon costs. The same can be said of many of the shared facility items that underpin the DRI resources.

Gathering time-stamped cumulative energy readings eases the calculation of accurate carbon costs, and they are robust in the face of missing readings, compared to instantaneous power readings.

embodied carbon costs of compute equipment is hard to obtain and often generic in nature not relating to exact computer specifications. It was found to be more difficult to get carbon cost figures for networking equipment.

The data gathered in the IRISCAST snapshots was insufficiently detailed at the job level to be able to distinguish different kinds of jobs and hence to compare different computer carbon efficiencies or to investigate how jobs running on the same note affect each others carbon efficiency. A synthetic workload study rather than a snapshot study might be more appropriate to investigate these issues in future.

As a DRI is essential to conducting research UKRI will wish to maximise the research output for every tonne of carbon emissions. Funding bodies already have procedures for allocating financial resources to projects based on research merit. It is clear that it is these procedures that must be extended to additionally allocate a carbon budget to projects also based on research merit.

High level feedback, using as far as possible absolute measures, will be needed to inform these decisions which could be accomplished by periodically obtaining: reports of cumulative energy usage of DRI and associated cooling; reports of an embodied carbon inventory of the DRI equipment (including measured idle power draw and expected lifetime) and best estimates of building embodied carbon; and reports of DRI facility usage and idle figures. Including this carbon reporting into standard grant reporting procedures would seem an efficient way to achieve this high level feedback to funding bodies. This reporting underpinned by a dashboard – an IRIS Carbon Dashboard for example – would give high level information to decision makers at DRI level and funding body level information to allocate carbon resources efficiently and effectively.

However, the above would do little to help the end user researcher or research community to drive carbon reduction in their workflows and code bases. This is where low-level feedback can make an impact. As IRISCAST has shown, few sites are setup for such low-level feedback. The energy reporting in the Slurm workload manager would be an easy first step in providing low-level feedback and so sites using Slurm should consider configuring that.

IRISCAST measurements have shown there is likely, at least initially, to be a difference between the sum of the carbon costs of the low-level feedback and that of the high-level feedback. It is envisaged that these feedback mechanisms would develop and mature over time and discrepancies between high- and low-level feedback would diminish.

If data collection is to be rolled out across the UKRI DRI estate it will be important to define exactly what data should be collected and the format in which it should be stored and transmitted. Such data collection should include: Equipment inventories including embodied carbon of equipment, idle power draw, expected lifetime; Time stamped cumulative energy usage of the DRI resource and separately the supporting facility with any percentage apportionment also noted.

The IRISCAST carbon model indicates that the building embodied carbon costs are likely to be a small contribution to the total while active carbon costs of the equipment are likely to be the largest contribution thus the area with the highest potential for climate impact. This indicates that continuing the process of decarbonising the electrical supply will significantly reduce DRI climate impacts.

The carbon model approach used by IRISCAST can not account for or compare alternative ways of achieving the same research: it should be remembered that often computer modelling replaces many polluting physical experiments, each with their own carbon cost.

As the energy grid carbon intensity decreases so does the active carbon component of the model making the embodied component more significant: leading to the conclusion one should extend the lifetime of DRI resources to make the most of that embodied carbon sunk cost. However, this effect must be balanced against procuring newer more energy efficient computers which can do more useful work.

The total modelled carbon emissions of DRI sites involved in the IRISCAST audit during the snapshot period were said to be equivalent to flying between 1 and 4 people (depending on the scenario used) on a return flight to a destination 12 hours away. Clearly this means that the total UKRI DRI, of which IRISCAST was a small part, is not an insignificant carbon cost but the carbon modelling shows that carbon costs can vary by factors of ~10, hence there is significant opportunity for the carbon footprint of DRIs to be reduced.

3.3.7 Value and Net Zero Decision Making (VALUE)

The VALUE sandpit project (Boulton, 2023) aimed to support the overall NetZero DRI roadmap by investigating the relationship between the emissions a DRI occupying task emits versus the ‘value’ that the task adds to UKRI. Understanding this will allow UKRI to make informed decisions extracting maximum ‘value’ for minimum emissions. This section summarises the VALUE final report.

The VALUE project had three objectives. Firstly to make an assessment of the current UKRI approach to DRI allocation decision making. Secondly to review how other organisations approach ‘emissions-vs.-value’ comparisons. Thirdly, to make a recommendation to UKRI of a framework to approach ‘emissions-vs.-value’ comparisons for NetZero DRI use.

To assess the current UKRI DRI allocation approach an interview was conducted with a member of the ‘UKRI Net Zero Project Scoping Team (Ag Stephens – CEDA Head of Partnerships) to gain an understanding of how UKRI currently operate their DRI, and where decision making structures do or do not exist.

A literature review of ‘value’ assessment /quantification methods within organisations and a review of comparison techniques was undertaken. Additionally the broader system of constraint optimisation within data-centres was also investigated in literature to ensure these comparisons result in positive action towards an emissions objective.

The knowledge and wisdom gleaned from this interview and the literature review were then combined to form a recommendation to UKRI about ‘emissions-vs.-value’ supported by a suggested framework to ensure Net Zero commitments are met.

The interview to gather information about the current UKRI DRI allocation system found the following.

UKRI is made up of 12 research councils. The DRI ranges from HPCs to university server rooms and everything in between. Due to the vast and distributed nature of this infrastructure, an exact profile of the DRI is not currently known to UKRI. A project is underway to perform a ‘stock take’ of DRI and confirm what is ‘in scope’ for this project. This includes a questionnaire to data centre managers about the energy use profiles of their DRI.

UKRI does not have any overarching policies / systems in place for assigning whether or not a request is allowed to use a specific piece of DRI. It is assumed that this decision making process is organised at the local level by the local infrastructure manager. Existing constraints (e.g., financial, time of use, and infrastructure capacity constraints) are currently managed at a local level in a variety of ways.

There is no formal system to track the ‘value’ that any use of DRI generates. Forthcoming questionnaire aims to find if ‘scientific throughput’ is being tracked / monitored at data centre levels. For ‘higher tiered’ resources such as HPCs, users must already go through an approval process where they are asked to predict energy consumption of their request. However, after acceptance there is little monitoring of actual infrastructure use.

Overall the DRI is very heterogeneous with no generic cloud communication capability between sites. There is the possibility that cloud-based requests will be implemented in the future but this is not likely to occur anytime soon.

This currently decentralised system of infrastructure control seems to follow the usual, and logical, approach to resource constrained decision making. Being that systems are only put in to place to monitor and scrutinise resource allocation when unscrutinised use of the infrastructure approaches a limit (such as infrastructure capacity) or impacts on a budget (such as a financial budget for electricity use). E.g., we pay little attention to resource use limits unless we have to.

Historically there has been no constraint on the emissions that each request for DRI use generates and so there has been no need to track or make decisions around emissions allocation to DRI use. HPCs have always had financial budgets (for capital costs and running costs) associated with them and management structures have evolved to optimise within these limits. In a world of NetZero computing (and the pathway towards it), the carbon emissions associated with DRI use will need to be monitored and limited in line with climate safety pathways.

To progress towards a NetZero DRI, UKRI will need to implement a system to monitor and manage their infrastructure use and the energy it consumes.

A review of the literature regarding 'value' concluded the following.

The term 'value' is used to describe the benefit that a task generates to an organisation while 'task' is used to refer to any job or process that occupies a piece of DRI and/or consumes electricity through DRI. Traditionally for profit companies track 'value' via the balance sheet.The tracking of 'value' in a non-profit organisation like UKRI is a more complex process.

There are a variety of value tracking techniques and comparison methods available to an organisation. These broadly fall into two camps: Direct quantification of a task's benefits - which is resource intensive in time and money; and Proxy systems which can allow for comparison without direct quantification and are less resource intensive.

Specific performance measurement systems (and their associated comparisons) should be catered to the organisation's requirements. Detailed direct quantification is a time consuming way to estimate value. Less resource intensive methods should be used where possible. The VALUE report details seven direct comparison methods which avoid direct quantification of value which in this summary are simply listed: Scoring, Relative comparison, User justification of benefits, User self-determined scoring, Willingness to pay, Monetization, Contingent Valuation.

Decisions to reduce carbon emissions to Net Zero clearly need to be made at a range of levels: strategic, tactical and operational.

To effect operational level decisions aligned to strategic aims the VALUE project suggests that a Smart Carbon Scheduler be deployed across DRI sites. Job schedulers are a well-established framework to achieve certain digital resource allocation objectives and ensure operation within constraints. Historically these aims have been performance based or cost based. However recently there has been growing attention around developing schedulers to prioritise a reduction in carbon emissions due to infrastructure usage. The VALUE report looks at a handful of carbon schedulers that have been discussed in the literature. These broadly fit into two camps: time shifting demand and geographical shifting of demand. Due to the heterogeneous nature of the UKRI DRI in general it is not possible for tasks to migrate from one DRI resource to another. This rules out further consideration of geographical shifting of demand in the short to medium term. Time shifting demand however may be possible. This allows demand to be shifted to times of low carbon intensity for example when the wind blows. The key unifying driver behind the design of the four time shifting scheduler designs reviewed is that there is a need for data centres' electricity demand to be as flexible as possible to respond to a renewable based grid. This driver will become more important as renewable penetration in local grids increases and carbon emission budgets tighten.

The UK energy grid will inevitably deploy large scale grid energy storage to help match the supply and demand but this currently has huge financial and environmental costs. Therefore, temporal shifting of electricity demand will be important in a low emission future.

Appreciating that the UKRI research councils are somewhat independent in their decision making, the VALUE project finds that only flexible and adaptable solutions and recommendations make sense in this context. To help value vs. emission comparisons to be made the VALUE project outlines a smart carbon scheduler system to optimise the operation of equipment within a NetZero pathway. The key inputs are as follows. An electricity grid carbon intensity forecast, itself informed in part by the energy infrastructure mix, current demand and weather. A Carbon Emission tolerance or budget itself informed by the UKRI DRI NetZero roadmap and the carbon emissions of tasks already run in that budgeting period. Last but not least the tasks themselves which include predicted energy/carbon usage and some kind of priority or value information. In summary, the Smart Carbon Scheduler itself will compare and prioritise tasks into queues. This prioritisation should avoid detailed direct quantification of value but use a less resource intensive proxy as discussed earlier. The jobs in the queues can then be allocated to run on infrastructure using a 'run criteria' which aims to optimise carbon efficiency. The level of complexity of this 'run criteria' may vary from simply running jobs when the grid carbon intensity is below a threshold, to multiple thresholds based on job priority, to a complex rule which tries to optimise 'value' based on a dynamic forecast of grid carbon intensity.

Additionally it is suggested that an energy use register or carbon use register is maintained. This can be used for strategic and tactical decision making and for carbon budgeting. For the majority of facilities simple estimates / measurements of energy use could be paired with grid information to give quick and reasonably accurate snapshots of DRI emission profiles. More detailed tracking of emissions can then be focused on the highest emitters and those that need larger reductions as carbon budgets tighten. Such a 'scrutiny' or focused approach would help track such emissions in an efficient way and ensure decarbonisation effort is focussed in the areas where it is most needed.

To be reasonable, solutions will need to be modular and flexible to fit each of the individual research councils and to be adopted by DRI resources. The quantification of 'value' by performance measurement systems should be catered to an organisation's specific operation and that large effort requirements associated with detailed direct quantification should be minimised. Task schedulers were identified as an overarching system to implement DRI emissions reduction through value comparison. UKRI will need a system that monitors the emissions from DRI use, and alters it in line with the diminishing carbon allowance required for a NetZero pathway. It is suggested that an energy use register will help with this as would a flexible smart carbon scheduler system at each DRI facility to: ensure each facility operates in line with NetZero pathway; schedule tasks / allocate resources as efficiently as possible; and inform emission reduction decisions. Detailed design work would be needed to turn this high-level, modular framework into an initial implementation. Pursuit of perfectly detailed emissions tracking should not be allowed to delay initial design of downstream projects such as a smart carbon scheduler which would be flexible enough to incorporate detailed emission tracking as it becomes available.

3.4 Consortium activities

What follows is a summary of the aims, methodology and findings of each consortium project.

3.4.1 Facility Case-studies

The case studies of exemplar facilities at national level gave opportunities to examine the net zero practices and issues of centralised resources at this scale. Two sets of facilities and two distinct approaches were taken: one concentrating on ARCHER2 at the University of Edinburgh (Smith et al., 2023), and the other dealing with STFC-operated facilities JASMIN and Scafell Pike/JADE 2 (Lambert, Stephens & Kayumbi, 2023).

These three facilities (two of them within the same centre and so treated together) were chosen for the interview-based case studies on the grounds of their close connections with the Net Zero project team and their representation of different types of national-level infrastructures.

This section summarises the ARCHER2 final report and JASMIN and Scafell Pike/JADE 2 final report.

ARCHER2 (https://www.archer2.ac.uk/) is the UK's National HPC service, funded by UKRI (via EPSRC and NERC) and operated by EPCC, the University of Edinburgh, at the Advanced Computing Facility. It offers a general-purpose service supporting a wide range of research communities and software with a corresponding wide range of technical requirements and use cases.

JASMIN (https://jasmin.ac.uk/), operated by STFC on behalf of NERC, is the UK's data analysis facility for environmental science, providing both compute and storage facilities for data-intensive research. It is a highly versatile infrastructure supporting over 1500 users exploring a wide range of topics in environmental science and engaged in a variety of workflows. Offerings include development and running of analysis codes, long-term data curation services, high-throughput data transfers, hosting services for environmental science, and delivering training.

The two facilities known as Scafell Pike and JADE 2 are both part of the service offered by the Hartree Centre (https://www.hartree.stfc.ac.uk/), which is devoted to the exploitation of supercomputing, data science and artificial intelligence technologies for the benefit of UK business and industry. Scafell Pike is a HPC facility that supports modelling and simulation for engineering, computational chemistry, physics and some life sciences work, while JADE 2 is a GPU cluster with heavy AI application usage. Both are managed by STFC's Digital Infrastructure department.

Key findings

(a) ARCHER2

The ARCHER2 case study (Smith et al., 2023) produced a series of 13 recommendations under the headings Societal value; Best Practice; Power and Energy Efficiency; Energy Consumption at the User Level; and Commissioning and Decommissioning. These recommendations are directed at a variety of stakeholders including service providers, data centre managers, vendors and end-users (researchers).

From the study of ARCHER2 itself, some important findings emerged. It was observed that the idle power draw is a large fraction of the peak power draw, leading to the recommendation that high utilisation is vital for efficient use of resources. High-quality, reliable instrumentation at all levels (per-job, hardware, power supply, cooling) is seen as very important for the enhancement of energy efficiency. During the period of the study, some improvements in ARCHER2's power efficiency were already made: the power consumption was reduced in May 2022 from around 3.22 MW to 2.94 MW, and again in December 2022 to 2.53 MW, an overall reduction in power usage of more than 20%. A final very important finding came from experiments with processor frequency: changing the default processor frequency from 2.25 to 2.0 GHz reduced energy consumption substantially while having relatively little impact on performance on benchmarks.

(b) Scafell Pike/JADE 2

Emerging from the interviews conducted with DRI managers, a number of findings and recommendations were produced that appear to be of general relevance and validity to Net Zero and DRI, in the following areas.

Power monitoring: Best practice should be shared among DRI managers, and those responsible for design and procurement,

Management of energy contracts: Negotiation of contracts, at whatever level it takes place, should take account of the patterns of usage at DRI level with a view to opportunities for optimisation.

User practices and behaviours: The possible scope and implementation of user-related monitoring should be examined and disseminated.

Local sources of energy and use of heat: Any practice and experience should be shared on local sources of energy and heat reuse

Opportunities in technological developments: DRI managers should be encouraged to identify technological opportunities with implications for Net Zero attainment

Principles of procurement: Principles and priorities of procurement should be re-examined to assess their alignment with Net Zero goals, within the necessary constraints.

Knowledge base on data centre design: A central shared knowledge base on data centre design and best practice should be developed and sustained

Aims

(a) ARCHER2

As part of the wider Net Zero DRI Scoping project, the ARCHER2 case study investigated the emissions associated with ARCHER2 and how to improve the energy efficiency of the service.

(b) Scafell Pike/JADE 2

The aim of the Scafell Pike/JADE 2 case study was to explore the current concerns and future expectations of the facility managers with respect to Net Zero, and to draw conclusions and recommendations of general relevance and validity to Net Zero and DRI. As per the project plan, another goal was to inform the design of the database of UKRI DRI, and there was a strong interaction between the two lines of work.

Methodology

(a) ARCHER2

ARCHER2 was compared with other large HPC systems in terms of best practice for HPC data centres and power efficiency. Issues such as future power requirements, reuse of waste heat and instrumentation were considered. The power draw by component was examined, as well as the estate on which the facility is housed (the Advanced Computing Facility), which is relevant for cooling. Investigations were made of the effect of changing processor frequency on ARCHER2 in terms of energy consumption and performance.

(b) Scafell Pike/JADE 2

The method adopted for the Scafell Pike/JADE 2 case study comprised six steps.

1. **Selection of DRIs**
2. **Preparation of interview scripts.** A systematic approach was taken with four groups of questions focussing on ‘What is known?’, ‘What is done?’, ‘What is the external environment?’ and ‘User communities’.

3. **Conduct of interviews.** The interviews (four were held altogether) followed the script and explored the issues. Notes taken at the time provided the raw material for the next step of the process.
4. **Reports on interviews**
5. **Analysis of interviews.** The chosen approach was SWOT analysis (Strengths, Weaknesses, Opportunities and Threats) applied to the structured notes of the interviews.
6. **Conclusions and recommendations**

3.4.2 Exploiting batteries to reduce carbon intensity of power usage

The cost of electricity, in both monetary and carbon terms, varies throughout a 24-hour period with changing demand and the availability of different forms of power generation. It is therefore feasible, and in all likelihood desirable, to use energy storage technologies to exploit this fact; storing energy at periods of over-supply where costs and carbon intensity are low, and then discharging this energy, and therefore refraining from drawing power from the national grid, when demand outstrips supply as costs and carbon intensity peak. A study was therefore conducted to look into the feasibility of applying such a scheme to DRI, specifically on the JASMIN supercomputer based at the Rutherford Appleton Laboratory in Oxfordshire.

Two open source datasets were used to evaluate this: the Octopus Energy agile tariff for an approximation of wholesale energy cost, in units of £/kWh; and carbon-intensity data from the UK's National Grid, in units of gCO2/kWh. These were averaged over a 2-year period into a 24-hour half-hourly averaged profile, Figure 3.4.2A. They were found to be covarying, with the two peaks and two troughs in roughly the same location in the 24-hour cycle.

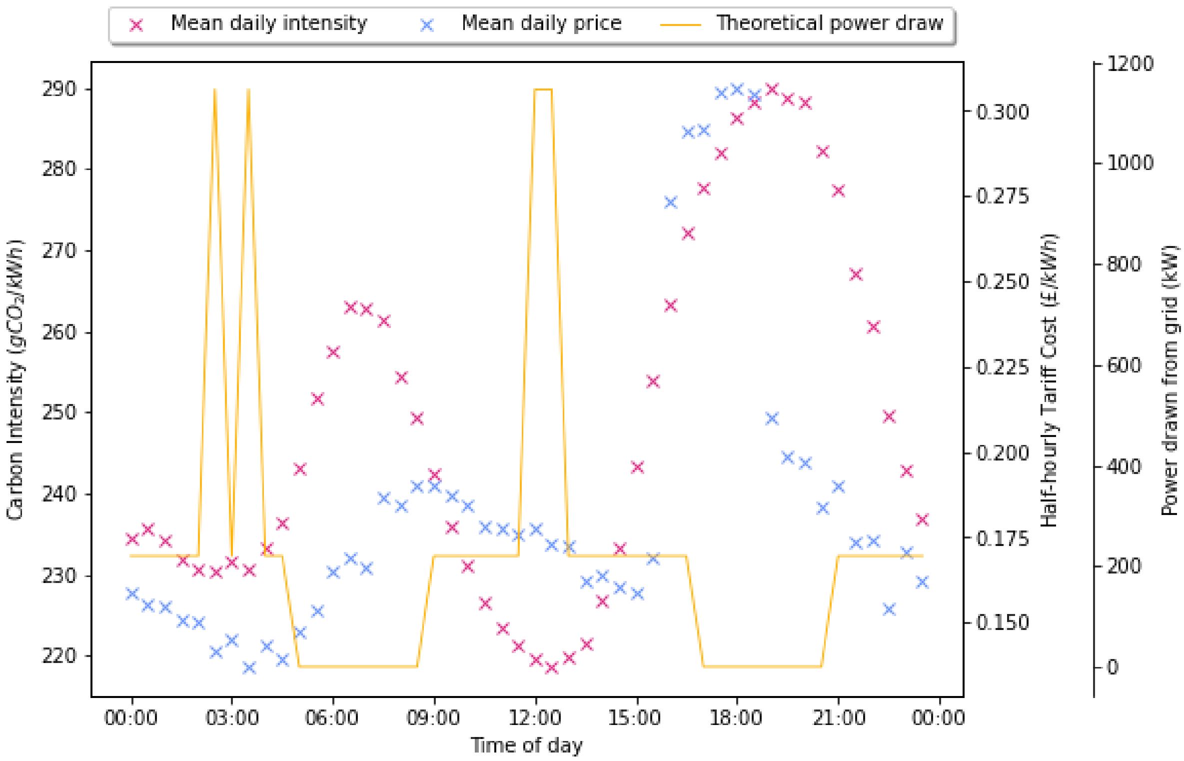

Figure 3.4.2A: *Plot of the 24-hour half-hourly averaged profile of carbon intensity (magenta) as obtained from national grid carbon intensity API, along with the half-hourly averaged profile of monetary cost as obtained from the Octopus Agile tariff. Overlaid in yellow is the theoretical power consumption curve for the JASMIN system for the 24-hour period, with the*

charge and discharge of a theoretical 880kWh battery timed at the points of lowest and highest carbon intensity respectively. The plotted figure shows a 100% efficiency of charge and discharge, but a range of efficiencies were modelled.

The peak was found to be roughly 4-hours wide, so a battery would have to be selected which could supply the power requirements of JASMIN for that 4-hour period. At time of writing, JASMIN consistently draws approximately 220kW of power from the grid, requiring a battery of 880kWh capacity to sustain all of this draw for the 4-hour period. Commercial batteries were found to be available at this size at a price of approximately £150/kWh, totalling a cost of around £110,000. With this information, the application of the battery's charge-discharge cycle to both the monetary and carbon cost profiles could be modelled, i.e. have them draw 1100kWh for two 1-hour periods at the two lowest carbon intensity points, and having them draw no power over the two 4-hour peaks of carbon-intensity. Power was drawn for only 1 hour as this is the highest 'safe' recharge time quoted for most commercial grade batteries, but higher recharge times could be used, if desired, albeit to the detriment of battery operational lifespan.

Upon applying the charge-discharge cycle it was found that, over a 24 period, with a battery charge-discharge efficiency of 95%, the total cost was reduced by 8.5%, and the total carbon cost was reduced by 4%. If we assume the use of a Lithium-ion battery, the initial capital outlay would be repaid in as little as 4-years. Lithium-ion batteries boast an estimated 10 year lifespan given proper use and maintenance, so there could potentially be a 6 year period of 10% cost reduction, resulting in potential savings of £170,000. This analysis can also be applied to the embodied carbon of the battery, i.e. the cost of manufacturing and transportation etc. Estimates for this vary significantly, with the value depending on many assumptions but principally the energy mix at location of manufacture. An average value of 150 kgCO2/kWh was taken, resulting in a total battery carbon cost of 132,000 kgCO2 equivalent and a payback time of 6 years.

Other battery technologies could feasibly be used, such as Vanadium redox flow batteries (VFRBs). These typically, at the time of writing, have approximately double the monetary cost per kWh of Lithium batteries, resulting in a longer payback period – in this case 8 years. They purportedly have approximately the same carbon costs per kWh as Lithium-ion batteries. However, they also have better environmental credentials when evaluated more holistically as they have both a substantially longer lifespan, with some estimates being upwards of 20 years, and better potential for recycling, which could reduce their lifetime carbon emissions significantly.

In conclusion, this study found that it is indeed feasible to produce monetary and carbon cost savings by better managing DRI power consumption using a large-scale battery as energy storage. The cost of the battery could be repaid in monetary terms in 4 years and in carbon terms in 6 years. There are caveats however as only a limited study was carried out. It would be worth considering how larger or smaller batteries with different discharge cycles could change this outcome. Additionally, the method by which the charges and discharges were timed could be made more dynamic to result in greater savings, i.e. using a longer time series and a moving window of charge/discharge. Finally there are probably other inefficiencies and intricacies to take into account when drawing and storing power at the hundreds of kW scale that were not taken into account. Nevertheless, the study clearly shows that the idea is feasible. With further reductions expected in the cost of batteries and as grids become more decarbonised - resulting in the double benefit of lower embodied carbon in battery production and lower carbon intensity at charge time - the benefits are also likely to become far greater in the coming years.

3.4.3 DRI mapping database overview

Introduction and Aims

The UKRI digital research infrastructure (DRI) represents a vast and diverse range of systems, people and activities that is inherently difficult to measure. The UKRI Net Zero Digital Research Infrastructure Scoping Project took on the task of gathering a representative sample of DRI facilities with the aim of understanding what the DRI looks like, along with a range of properties that relate to its wider carbon footprint. It is expected that future activities within this remit will require a database of all significant DRI assets in terms of their technical specification, energy usage, efficiency measures and user communities. Capturing a timeline of this information will allow analysts to compare different systems in order to identify opportunities for positive change.

The work was carried out by a team of Science and Technology Facilities Council (STFC) staff from the Scientific Computing and RAL Space Departments. This section summarises the DRI mapping final report (Stephens, Kayumbi & Lambert, 2023).

Methodology

The first phase of the work was to compile a list of DRI facilities. This work was done by interrogating the UKRI InfraPortal records and knowledge of project partners, to identify relevant facilities and their contact points. After an in-depth process, the final count of facilities to survey was 123.

The next stage involved developing a comprehensive questionnaire to be emailed out to each of the facility managers/contacts. A questionnaire was developed, in conjunction with the wider team in the scoping project, with sections focussed on (1) Contact and facility info, (2) Funding and community, (3) General characteristics, (4) Technical details, (5) Energy supply, (6) Energy usage, (7) Environmental considerations and constraints, (8) User engagement and (9) Other information and feedback. Since we were aware that the DRI is highly diverse, it was likely that some questions were not applicable, so they were phrased in a manner that would allow respondents to provide simple answers, or to elaborate on details or relevance. A number of case studies were carried out in the project, the information from these also fed into the design of the questions.

A Google Form was chosen as the technology to deliver the survey, as it was free and simple to use, and the results could be easily converted to CSV or Excel formats. The questionnaire was tested on a small sample of colleagues who run DRI facilities, and their feedback was used in developing the final version. The survey was carried out between November and January 2023, with recipients being contacted via e-mail, and in-person where they were known to the team.

An Expert Advisory Group (EAG) was created to help the project team decide on the most appropriate way of analysing and sharing the results of the survey. The key recommendations from the EAG were that (1) an internal database would be analysed for the main results, (2) a public dataset would be generated with anonymised and categorised responses, (3) the public dataset would be published alongside metadata and guidance on how to use it, including an interactive Jupyter Notebook demonstrating how to read and visualise the dataset.

The responses were exported to a CSV file and processed to improve the quality of the dataset. The transformation of the original responses through to the public dataset included the following steps: anonymisation, exclusion of irrelevant fields/records, mapping of answers to categories or numerical bins, handling of missing values, outlier checking, correlation checking and summarisation of free-text responses.

Results, Key Findings and Recommendations

Of the 123 facilities contacted, 51 responses were received, and 7 of those were excluded as not being appropriate or relevant to the DRI dataset. The results therefore consisted of 44 valid records.

The results were reviewed, analysed and are presented in the report in a series of pie charts. The key findings were extracted and mapped to a set of 20 recommendations. The findings and recommendations can be grouped into the following topics:

- Defining the DRI
- Engaging the community
- Disseminating knowledge and best practice
- Creating a database of DRI information
- Innovating and supporting the community
- Net Zero policy and strategy
- Undertaking further research

Within the scope of **defining the DRI**, it is imperative that a consistent and reliable picture of the physical DRI is found, either from a system-wide survey or other methods such as sampling and modelling. Given the variability in the DRI, it is important to categorise the DRI facilities by a defined set of metrics. Additionally, to avoid missing parts of the UKRI carbon footprint, we must ensure partially UKRI-funded resources are appropriately captured.

Regarding engagement of the community, the level of engagement with each component of the DRI should reflect its size and relative impact. It was identified that some UKRI DRI facilities are already collecting data relevant to their carbon footprint and are keen to engage. UKRI should work with the contacts made within the DRI Mapping process to explore the most effective ways of maintaining and evolving community engagement to build the DRI database required in future. Incentives, such as funding and support, should be considered to foster engagement. It is also important that direct contact is made with each Research Council in future engagement to ensure full coverage of the DRI.

In terms of **disseminating knowledge and best practice,** a centralised service or hub should be created to gather best practice and promote positive change. it is important to share the experiences of those parts of the DRI community that have committed to using renewable energy, and to consider whether UKRI should consider a policy of purchasing electricity from renewable sources across the entire DRI. Additionally, some respondents are re-using waste heat, so their knowledge and experience should be shared with other facilities. Since a variety of energy usage monitoring currently takes place, support should be given to facility managers to enable them to develop and deploy systems for monitoring and reporting energy usage. Information, **support and training is essential to bring about the change required to meet the Net Zero target**, this needs to focus on understanding and approaches for more sustainable practices (at all levels).

In relation to the **creation of a database of DRI information**, it is important to create standards for recording and reporting energy usage and efficiency. A public database of key information about the UKRI DRI should be collected and updated routinely. The measurement of annual energy usage should be a fundamental measure recorded in the database.

Regarding **innovation and support**, a fair approach should be developed for recording, monitoring and reporting energy usage information at multiple levels (e.g. per job, user or project). Once a comprehensive dataset exists, then UKRI should investigate whether larger facilities bring about more efficient use of DRI, and if so, make arrangements for greater centralisation of infrastructure and services.

In terms of **Net Zero policy and strategy**, there has been a great deal of good will and interest in making changes towards a lower carbon footprint from both users and managers of DRI facilities. Net Zero needs to become an explicit and significant concern for all in the UKRI community. It should be embodied in all layers of the organisation, and there should be suitable recognition, acknowledgement, and reward for positive action towards sustainability.

Regarding **further research**, the relative contribution of three areas requires more investigation: personal devices (e.g. laptops, tablets, and phones), digital networks (i.e. the physical networking infrastructure) and use of public cloud computing and storage. These were considered outside the scope of the DRI Mapping Survey and have not been estimated in terms of how they compare to the DRI facilities. UKRI needs more data on the impact and trends of energy usage related to each of these.

Conclusion

Mapping the UKRI Digital Research Infrastructure was a complex and involved task. There is no clear delineation of what the DRI is, and there is great variety in terms of the size, purpose, function and usage of the components of the DRI. Whilst some major parts of the DRI are already engaged with the Net Zero agenda, others are lacking information and training to be able to assess their carbon footprint. The work outlined in this report demonstrates a pathfinder process that is essential for UKRI to accurately monitor and improve its energy usage and progress towards the Net Zero target. The generation of a list of the significant DRI Facilities is a pre-requisite to capturing the current situation. Furthermore, establishing a methodology for measuring or modelling the long tail of small facilities and personal compute is also required. The results and inferences within this work highlight the value of collecting and sharing information about the DRI to draw out best practice and opportunities for knowledge-exchange. Building a more comprehensive database of this information, and capturing temporal updates, will be essential for UKRI to reach Net Zero by 2040.

3.4.4 User survey

The aim of this research was to investigate the motivating/enabling factors and barriers to sustainable behaviour, as well as willingness to change, in several work-related domains, namely, general work-related behaviour, digital infrastructure usage, research, data storage and code writing.

Implicit attitudes of participants to low carbon choices were also measured. Implicit attitudes are often a better predictor of behaviour, particularly those behaviours that are subject to social desirability such as sustainable behaviour, low carbon choices and racial bias. Implicit attitudes were measured using a carbon Implicit Association Test (IAT).

This section summarises the User Behaviour Survey Report (McGuire, 2023) which can be found at User Behaviour Survey final report.

3.4.4.1 Survey Respondents

284 participants took part in this study (male n=123, female n=150, non-binary n=2, prefer not to say n=9), they were recruited via university newsletters, emails, as well as through social media.All respondents were service users in some capacity - be it small (e.g., use of desktop applications, email, SharePoint etc.), medium (e.g., research workflows), or large (e.g., major use of parallel processing infrastructure). Respondents worked in a variety of research sectors including Arts and Humanities, Environmental Science, Medical Research, Higher Education etc.

Respondents were asked to select which age category they belonged to and what stage in their career they were at, their responses are reported in table 3.4.3.A below.

<table>
<tr><th>Age range</th><th>Percentage of Respondents</th><th>Career stage</th></tr>
<tr><td>18-24</td><td>3.9%</td><td rowspan="8">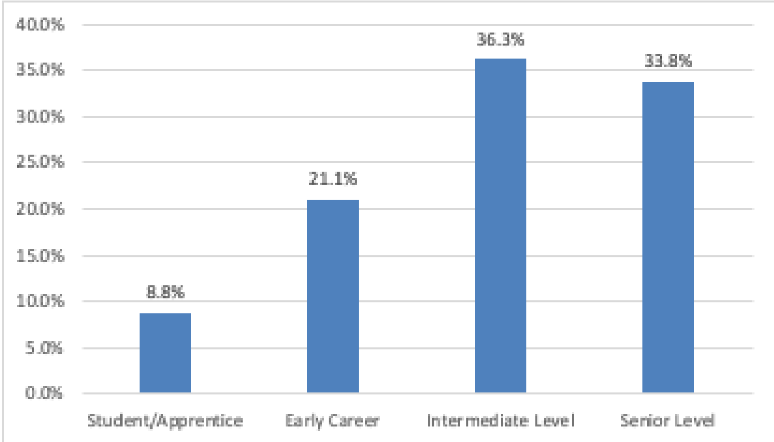
</td></tr>
<tr><td>25-34</td><td>20.8%</td></tr>
<tr><td>35-44</td><td>28.9%</td></tr>
<tr><td>45-54</td><td>23.9%</td></tr>
<tr><td>55-64</td><td>18.3%</td></tr>
<tr><td>65-74</td><td>2.1%</td></tr>
<tr><td>75 and over</td><td>0.35%</td></tr>
<tr><td>Prefer not to say</td><td>1.76%</td></tr>
</table>

Table 3.4.3.A Respondents were asked to select which age category they belonged to and what stage in their career they were at

3.4.4.2 Questionnaire

Participants were presented with twelve statements grouped under four categories: general work-related behaviour, technology/equipment, research, and data/code/large storage. E.g. “I do my best to reduce the amount of energy my research requires”. For each statement participants were asked to report on a 5-point scale how frequently they carry out these behaviours ‘Always’, ‘Often’, ‘Sometimes’, ‘Rarely’, or ‘Never’.

Participants who selected ‘Always’, ‘Often’ or ‘Sometimes’ were presented with a list of motivators/enablers and were invited to select those which were the best fit to their own reasoning/behaviour. E.g. “The environment is important to me and I want my work to reflect my personal values”.

Participants who selected ‘Rarely’, or ‘Never’ were presented with a list of barriers and were invited to select those which were the best fit to their own reasoning/behaviour. E.g. “I wouldn’t know how to reduce the environmental impact of my research”.

Alternatively, participants could select ‘Other’ and use the free response box to type in their own words about what motivates/enables them to carry out, or what prevents them from carrying out the specific action.

In addition to this, participants who selected ‘Rarely’ or ‘Never’ were asked if they would be willing to change their behaviour on a three-point scale: ‘Yes’, ‘No’, ‘Maybe’. Participants were then encouraged to use a free response text box to provide additional information, or justification.

The statements presented to the participants are listed in table 3.4.3.B below alongside the number of participants who responded and their most frequent (modal) response.

General Work-Related Behaviour Statements	**Number of responses**	**Modal response**

1	In general, I look to reduce my environmental impact in my workplace	284	Often 43.3%
2	Carbon emissions take precedence when I decide if travelling to a conference is worthwhile	256	Sometimes 32.4%
3	Carbon emissions take precedence when I decide if travelling to my usual place of work is worthwhile	255	Often 23.5%
Technology and Equipment Statements		**Number of responses**	**Modal response**
4	I feel excited about upgrading to newer digital technology if it is more energy efficient	272	Sometimes 31.3%
5	I look to save energy when running equipment	281	Often 39.9%
6	I consider the environmental impact of equipment upgrades	268	Often 28.4%
Research Statements		**Number of responses**	**Modal response**
7	I do my best to reduce the amount of energy my research requires	234	Often 36.3%
8	I think about the link between storing data and its environmental impact	272	Rarely 26.8%
9	I make my research data available following community standards so it can be reused more easily (e.g. utilising open access data repositories)	202	Always 34.2%
Data Storage and Code		**Number of responses**	**Modal response**
10	I look for efficiencies in the code I write so it uses less energy	141	Sometimes 24.1%
11	I seek to collaborate when writing code	133	Often 30.8%
12	I run software on national or regional computing resources e.g. JASMIN or ARCHER2	129	Never 56.6%

Table 3.4.3.B *The statements presented to the participants are listed alongside the number of participants who responded and their most frequent (modal) response. In the user behaviour survey report the full response distribution for each statement is provided including a disambiguation of responses by career stage.*

In addition participants were asked about their ability to calculate the carbon footprint of everyday activities as well as the carbon footprint of work-related activities including research, data storage, and digital technology use. Over half the respondents (56.4%) reported that they did find it difficult to assess the carbon footprint of their everyday activities. The responses were even more stark with regards to the carbon footprint of work-related activities with 79.8% of respondents reporting that they found it difficult.

3.4.4.3 Measures of Attitudes and Segmentation Analysis

Explicit Measure

Participants were asked to complete a 5-point Likert scale and were asked to select which statement best described them from 1- ‘I strongly prefer products with a high carbon footprint to a low carbon footprint’, 3, ‘I like high carbon and low carbon footprints equally’, 5 ‘I strongly prefer products with a low carbon footprint to a high carbon footprint’.

This was completed by 281 people. The mean Likert score for this sample was 4.5 which is between a ‘Moderate preference for low carbon’, and a ‘Strong preference for low carbon’

Implicit Measure

Implicit attitudes were measured using a computerised categorisation task called an Implicit Association Test or IAT. The IAT measured the speed of association between different paired concepts: ‘High carbon’ and ‘low carbon’ represented by images, and ‘good’ and ‘bad’ concepts represented by words.

The images that were used to represent high carbon included plastic bottled water, heavily packaged fruit, and single use carrier bags. The images used to represent low carbon included tap water, loose (non-packaged) fresh fruit, and a reusable bag (see Figure 4.3.3.a).

The words used to represent good and bad concepts were:

> GOOD: Happy, Lovely, Nice, Wonderful, Superb, Marvellous, Excellent, and Great.
> BAD: Sad, Unpleasant, Nast, Horrible, Awful, Terrible, Evil, and Appalling.

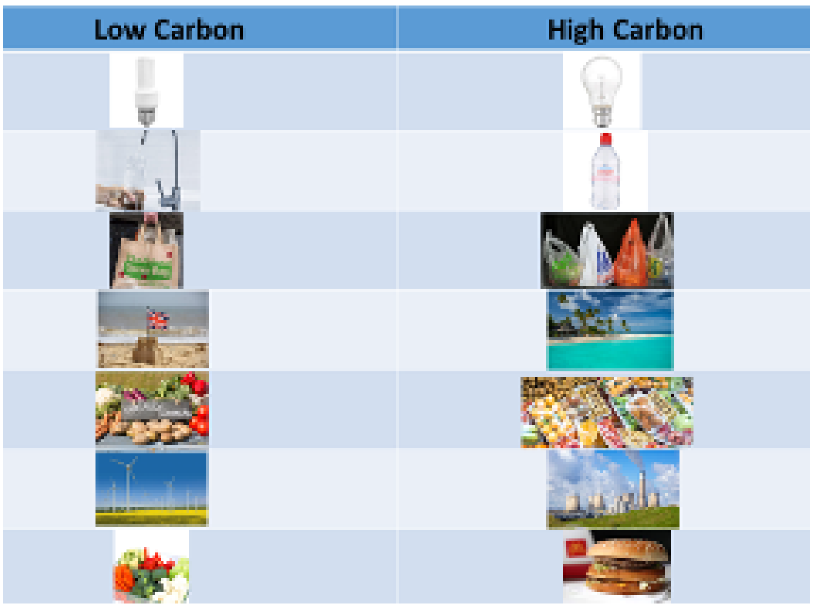

Figure 4.3.3.a: Images used in the Implicit Association Test

The implicit Association Test was completed by 231 participants. Scores ranged from -0.4 (‘Slight implicit preference for high carbon’) to +1.6 (‘Strong preference for low carbon’).

Segmentation Analysis

A segmentation analysis was used to separate participants into four categories (figure 4.3.3.b) based on their explicit scores (measured using the Likert scale) and implicit scores (measured using the Implicit Association Test). Due to the distribution of scores on the explicit and implicit measures the analysis focused on two of the segments – 'True Greens' (those who report a very strong pro-low carbon attitude and have a very strong pro-low implicit attitude to low carbon) and 'Surface Greens' (those who report a strong explicit attitude to carbon, but have a low implicit score).

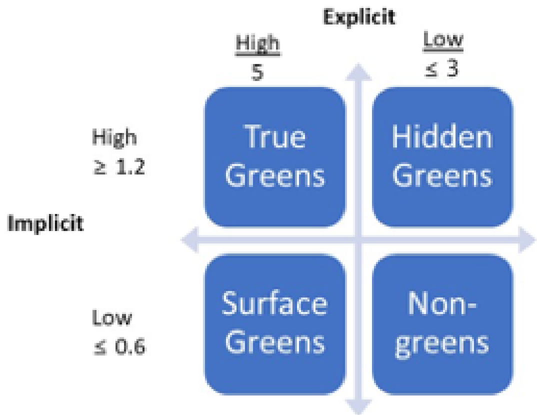

Figure 4.3.3.b: Segmentation analysis categorised by high/low explicit scores and high/low implicit scores

A qualitative analysis was conducted to determine if there were any differences in how respondents report perceived barriers to sustainability.

'True Greens' demonstrated a strong concept of personal responsibility in making change, for example, when asked about data storage and its environmental impact, one respondent wrote 'I am always willing to learn and listen to sustainability issues'.

'Surface Greens' tend to downplay their personal responsibility in overcoming barriers, for example, when responding to the statement 'I do my best to reduce the amount of energy my research requires', one respondent wrote 'We are NOT being encouraged to work in an environmentally conscious manner.

The comments made by 'True Greens' highlight their self-efficacy (that they feel they can make a difference), and response efficacy (that recommended action steps will make a difference). On the other hand, Surface Greens downplay the level of self-efficacy and response efficacy as demonstrated by the comment "Our volumes are small, so any impact of a choice would be insignificant' when responding to the statement 'I consider the environmental impact of equipment upgrades'.

3.4.4.4 Key Findings

1. The first conclusion from this research is very positive. Respondents did want to do something about their environmental impact around the workplace and they felt very comfortable about raising environmental issues with their teammates, thus demonstrating that networks within the workplace are an important facilitator to pro-environmental behaviour and should be encouraged, rather than utilising a top-down approach. However, the biggest barrier preventing respondents from looking to reduce their environmental impact in their workplace was the pressure of work – particularly time pressure and workload pressure, as well as other issues taking a higher priority.

2. When it came to specific work-related behaviour, respondents were keen to reduce their emissions from travelling to conferences. One major motivator/ enabler here was that people benefit from the time saved when attending a conference virtually. However, many of

the respondents felt that face-to-face networking is more effective. This was particularly true of early career researchers. People are willing to consider the environmental impact of conference travel in the future – but only if online conferencing improved somewhat, and if web-conference software could facilitate more effective networking strategies.

3. When it came to travelling to work, the respondents were very mindful of the carbon footprint of travelling to work each day and many enjoy walking or cycling because it keeps them fit as well as lowering their carbon footprint. However, many would much prefer to work from home whenever possible, but not all institutions allow this. Another barrier was that the lower carbon option - be it walking or cycling, getting the bus or train – is either not an option due to distance as it would take too long for them to get to work, or that public transport is too unreliable. Financial cost was also a barrier – with public transport costs being much more expensive than driving to work.

4. One obvious way of lowering the carbon footprint in the digital industries is to use more energy efficient technology and equipment. Many of the respondents were excited about this possibility – especially because more energy efficient technology would have added financial benefits. This was particularly true for those who did a lot of their work from home. A major barrier, however, was that not everyone feels that they have the power within the institution to make decisions when upgrading equipment. Some people were also cautious about whether old equipment may go to waste which would have a major negative impact on the environment. Respondents reported that they would also like more evidence that upgrading is preferable to extending the life of existing technology.

5. When it came to saving energy when running equipment, the results were a little surprising - although 66.2% selected ‘Always’ or ‘Often’ in response to the statement ‘I look to save energy when running equipment', only 26.3% said that they would ‘Always’ do this. In terms of the behaviours they engaged in 44.1% turn the equipment off - either because that is what they always do, or they do this specifically for good environmental reasons. But that number should be approaching 100% and many of the respondents did admit to leaving equipment on standby even though they are aware that they should be turning it off. People are willing to change in this regard, but they would like more information about how to save energy and best practices that they should follow. Other respondents need convincing that it is important. So, there is clearly more work to be done here in terms of encouragement and/or reminding people to turn equipment off.

6. When it comes to considering the environmental impact of equipment upgrades one would imagine that it is one of those transitions where the environmental impact of the equipment is at the forefront of people’s minds, yet only 23.5% of respondents said that they would ‘Always’ consider this during the upgrading process. And when they did so, they often gave a reason that they liked the idea that the equipment they currently use could benefit other people. But again, this is an issue of control and responsibility and the major barrier here preventing people from considering the environmental impact of equipment upgrades was that they did not have much input into this procedure. One obvious implication of this might be to increase people’s input into the procurement process.

7. When it comes to people’s own research, which is often very personal and can be reflective of the researcher/ research team, it was interesting that only 16.2% said that they ‘Always’ try to reduce the amount of energy their research requires. The main motivator seemed to be that the environment is important to them, and they want their work to reflect their personal values. But a major barrier was that they simply did not know how to go about reducing the amount of energy their research requires, and they do not believe that small changes regarding their own research will make a positive impact on the environment. So, there is clearly a need here for more information and education about its importance as well as how to go about doing this.

8. One striking conclusion from the responses to the statement ‘I think about the link between storing data and its environmental impact’ the most common response was that people rarely consider this issue at all, and this was particularly true of early career researchers, students, and apprentices. Some respondents did, and they reported that their main motivator was that it made their work more efficient when they cleaned up their data files. The main barrier, however, was that they simply are not aware of the environmental impact of storing data. So again, this is a signal for better information, particularly for student/apprentices and early career researchers so they can better understand and make more informed choices and how to be more environmentally friendly in this regard.

9. In terms of whether people make their research available following community standards, what is interesting was that approximately a third said that they ‘Always’ did this, but almost a third said that they ‘Sometimes’, ‘Rarely’, or ‘Never’ engaged in this process. One of the motivators for employing this clearly environmentally friendly behaviour was that it makes sense to avoid the duplication of effort, and that it enhances the impact of their own work when other researchers use their data. However, some respondents were concerned about sharing data that contains sensitive information, so this made the process somewhat more problematic because confidentiality needs to be considered. Others had tried to do this, but found the process too complicated.

10. When it came to the issue of whether respondents actively looked for efficiencies in the code they write to use less energy, interestingly, less than 20% said that they ‘Always’ did this. But although some people are motivated to do this to detect inefficiencies, many people had never considered this before or they simply did not know how to find information about the efficiency of their code. So, there is clearly a need for training here.

11. When it came to collaborating when writing code, just 15.8% of respondents said that they ‘Always’ did this. The most common motivator was that a fresh pair of eyes are more likely to find bugs. However, the major barrier here was that respondents felt that they did not have anyone who they could ask. So again, perhaps more guidance about possible collaborations would be very useful feature to consider. There was also a lack of awareness of how collaborating could potentially save energy - so training in this regard would also be instrumental.

12. In terms of specialised computing facilities where the infrastructure has been designed to generate direct environmental consequences (e.g. JASMIN or ARCHER2) it was somewhat surprising that the most common response was that this set of respondents had ‘Never’ considered using such facilities. Interestingly 100% of student/apprentices selected ‘Never’ in response to this statement. Those who would consider using national or regional computing resources said that it would allow them to analyse larger datasets and run larger models. But what was interesting was that respondents did not see any additional efficiencies when using this, so in their words, the environmental consequences did not override the obvious efficiency savings when running on local computing resources. In addition to this, some respondents were not aware of these facilities, but they said that if they did know how to use them, they most certainly would.

13. Just over half of the sample reported that they find it difficult to assess the carbon footprint of their everyday activities and 79.8% of the sample reported that they find it difficult to assess the carbon footprint of their work. So, there is clearly a need for information in this regard, as well as training as to how people can assess the carbon footprint of their work, and everyday activities.

14. In terms of reported attitudes to low carbon, it was welcome to see how positive these attitudes were. Yet we know that reported attitudes have limited implications for everyday behaviour, decision-making and actions.

15. Importantly, the implicit attitudes were also extremely positive with much of the sample holding a strong positive implicit attitude to low carbon. This combination of high reported attitudes to low carbon, with high implicit attitudes to low carbon sets up a very useful combination for pro-environmental behaviours. But different approaches to encouraging pro-environmental behaviour is needed when targeting people within the different segments and this might be a critical issue in this domain.

16. The vast majority of respondents fell into the 'True Green' category and the point about True Greens is that their attitudes (both explicit and implicit) point in the right direction. If people were not engaging in pro-environmental behaviours all of the time indicated by the choice of 'Always' in response, then this was to do with either cost issues or information issues, which clearly have to be worked on.

3.4.5 Sustainable computing

This consortium activity looked to answer the following questions:

1. Relative importance of emissions resulting from UKRI funded research but not directly attributable to facilities. In particular, we considered (1) emissions from computing equipment that is not part of UKRI managed facilities and (2) data movement to and from facilities and storage.

a. Based on estimates of the typical hardware usage amongst UKRI researchers and the number of researchers paid by UKRI (58,000), we conclude that desktop machines, laptops, local servers and small clusters contributions to UKRI emissions (including embodied carbon) are of the same order as the UKRI-managed facilities, and that therefore it is important to reduce the emissions from this aspect of UKRI research. (see https://zenodo.org/record/8072018 for more details about the methodology and data used to obtain these estimates).

In summary, we estimate the following for compute equipment used by UKRI funded researchers:

- UKRI facilities: energy consumption 100 GWh/y to 300 GWh/y; carbon footprint 20 kton CO_2e/y to 60 kton CO_2e/y
- Non-facilities, UKRI funded equipment: energy consumption 20 GWh/y and 70 GWh/y; carbon footprint 6 kton CO_2e/y to 25 kton CO_2e/y
- Non-facilities, not UKRI funded equipment: energy consumption 10 GWh/y to 100 GWh/y; carbon footprint 5 kton CO_2e/y to 36 kton CO_2e/y

b. Data movement to and from facilities and storage does not directly contribute to UKRI emissions. This is because the network infrastructure (Jisc's Janet network) is effectively always on and the energy consumption does not depend on the bandwidth utilised (https://zenodo.org/record/7778575).

However, because the demand for bandwidth from UKRI research grows continuously, Janet is effectively upgraded continuously. From our estimate, energy consumption of Janet is around 50 GWh/y. This corresponds to about 10,000 ton CO2e/y. (see https://zenodo.org/record/7778536#.ZCL5QxXMKWY for details on the estimate)

2. Impact of the software on emissions, i.e. how much can be saved by using different programming languages, programming frameworks, compilers and runtime systems, and more widespread use of energy-efficient accelerators, as well as adoption of better software engineering practices.

The main conclusion from this work (see https://zenodo.org/record/7709401#.ZCLxURXMKWY and https://zenodo.org/record/7709483#.ZCL5tBXMKWY for more details) is that there are very considerable savings to be made in many ways in the process of development and deployment of research software. Some of these gains are policy based, e.g. to avoid wasting computing time by occupying machines for access; others are concerned with reducing the need to re-run experiments (code review, testing); others with proper design of experiments to minimise energy expenditure to obtain required results; and finally code efficiency, in terms of choice of programming language, program optimisation and compilation optimisation.

We propose to train researchers in basic practices (e.g. software testing and review, DoE design) that will greatly improve the overall energy efficiency of their digital experiments, and provide researchers with access to green software engineers to on the one hand help them with their practice and on the other hand improve the energy efficiency of their code.

3. What technologies could make a dramatic difference in 20 years (horizon scanning).

This part of the project is mainly based on the very comprehensive (over a thousand references) report on the topic, the IEEE International Roadmap for Devices and Systems (IRDS) 2021 Edition, Beyond CMOS. The report looked at many aspects of post-CMOS technologies but there is not mention at all of the key issues related to sustainability of future computing devices (see https://zenodo.org/record/7778432#.ZCLy7RXMKWY for more details).

- Post CMOS: CMOS scaling is projected to stop by 2030. There is as yet no clear candidate to replace CMOS. Considering the many potential technologies, we conclude that the emissions from manufacturing of post-CMOS ICs are likely to be of the same order as today for at least 2040.
- Storage: As the carbon footprint of SDD is much higher than that of RAM, I have focused on the former. In conclusion, at least until 2040, advances in data storage will not lead to a reduction of emissions from data storage, and likely even lead to an increase per unit of data stored.
- Beyond-CMOS computational technologies:

The ITRS report discusses a number of viable beyond-CMOS technologies but most of these are likely to take the shape of specialised accelerators rather than mainstream CPUs. All of them are still in very early stages, but several of them are particularly promising in terms of energy consumption, in particular analog computing approaches and adiabatic computing which could lead to accelerators for matrix-vector multiplication, linear equation solvers and differential equations solvers.

However, based on the state of the art, we conclude that none of the beyond-CMOS computational technologies can help with getting to net zero by 2030. It will likely take until 2050 before some of them become mainstream, in particular because the most promising technologies in terms of sustainability (low embodied carbon, low energy consumption) will also require radically different programming models, which is likely to slow down adoption in similar fashion as seen by the advent of multicore and accelerator technologies

3.4.6 Art Commission

Art to inspire collective action on sustainable digital research infrastructure

Our roadmap towards carbon neutral digital research infrastructure will involve UK researchers from across disciplines and take an integrated and coordinated approach. To meet the ambitious sustainability target - of reaching net zero digital research infrastructure by 2040 - a broad transformation will be required. The arts have a role in this transformation as a medium through which people can consider their values, shift attitudes, reflect on issues, and find connection.

In September 2022, an open funding call for £25,000 was made available by the project to inspire collective action towards sustainable digital research infrastructure. Proposals from artists, arts collectives, or teams of makers were encouraged. A panel of experts was convened to review applications. In November 2022, Artist Paul Millhouse-Smith was commissioned. The commission ran between November 2022 until June 2023.

Background

This piece of work came about following conversations between the core team and members of the Arts and Humanities Research Council (AHRC) community. Two initial meetings were set up with experts from the AHRC community. These included informal discussions about new ideas and opportunities, new perspectives about challenges/barriers for DRI in AHRC, and ways of framing/sharing the scoping project to this particular community.

As a project team, we found these conversations particularly valuable because it challenged our perceptions and changed how we were planning to do some of the work. It also led to a new idea: to commission an artist, arts collectives, or teams of makers to work with the project as part of our consortium.

We are very grateful to the consulted experts as it resulted in an inspiring, interesting and unique piece of work that we (as a core team) would not have considered doing without their input. We are also grateful for their support as champions who have shared our project with their communities.

About the commission

Paul Millhouse-Smith is a multi-disciplinary artist and technologist who explores our relationship with the past and new technologies.

For this project Paul set out to explore how the decisions we make everyday affect and shape our future. The resulting installation comprises physical and virtual artworks, inspired by conversations between Paul and the research community, digital research infrastructure facility managers and the core project team, across a series of creative workshops, and 1-2-1 meetings.

Paul's work encourages people to look at the challenge of net zero for digital research infrastructure from a fresh perspective and inspire meaningful change. It is designed to stimulate engaging conversations across the research community and beyond.

Figure 3.4.6A: Paul Millhouse-Smith standing in front of the 3D ceramic printer that will be used to create his physical vases.

Creative workshops

Four creative workshops were held online and attended by people from across the UK digital research infrastructure community in January and February 2023. They were designed to engage diverse groups of people with the project, and to enable them to think about digital research infrastructure and the net zero challenge in different ways (whilst learning about others perspectives too).

Paul invited participants to engage in creative making sessions and captured conversations related to the critical decisions around the work we (the research community) do. These conversations took place both within workshops and via separate 1-2-1 meetings. Offline creative activities were also provided for those unable to attend the workshops.

Workshop participants explored what they hold important about research, digital infrastructure, and the world we live in - in the context of our rapidly changing climate and what we are set to lose. This was done across a combination of mark-making exercises in which participants were asked to respond to audio, verbal and visual provocations. A subsequent exercise required them to consider what they held sacred, in relation to digital research infrastructure and the net zero transition. Using participants' own ‘found’ (household) objects to expand on notions of value and importance, this half of the workshop facilitated diverse and lively conversations.

Out of a total of 41 registrations the workshops were attended by 22 members of the research community attended a creative workshop (19 people cancelled in advance or were no shows).

Many participants verbalised their enjoyment during the workshops, others were visibly emotional whilst discussing their perspectives. This quote, from one of the workshop attendees summarises their experience:

> “*Thanks for actually getting me to do art again. I think as a researcher we kind of get wrapped up in the research and forget to take time out for things like this - and it can really help to get us excited about our research fields again and see them from a different perspective!*”

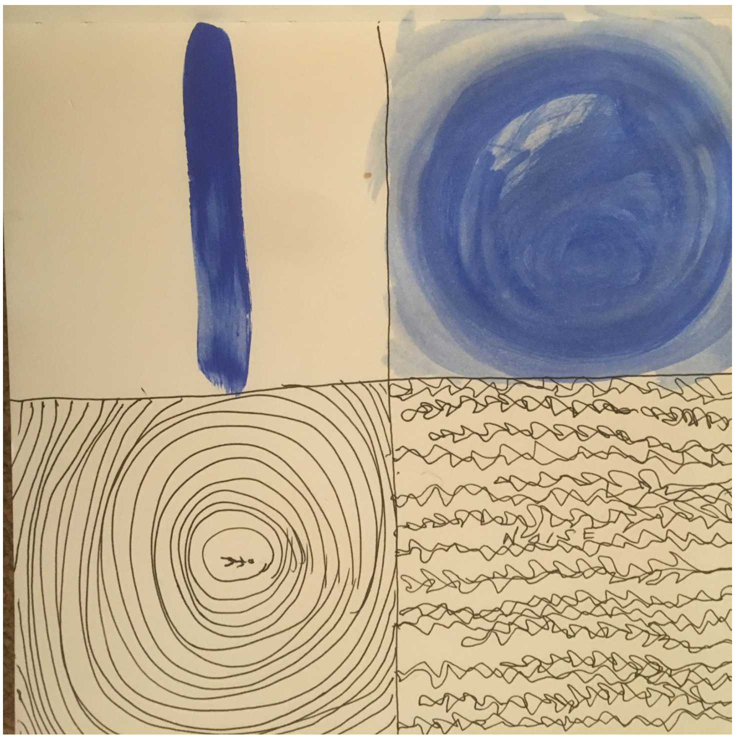

Figure 3.4.6B: Image drawn by an attendee of one of the creative workshops.

Some key points/themes discussed during the workshops:

Perspectives and values

Things appear differently to different people - even when we're all given the same information. Everyone has different ways of interpreting what is important to them, based on their values and differing backgrounds, contexts and experiences. This makes people behave in different ways - and respond to the same 'call to action' in completely different ways.

Recommendation [141-142]: UKRI must use a variety of different inclusive methods of engaging communities and communicating about challenges, related to climate change/net zero/digital research infrastructure, to influence diverse groups of people to work together and enable collective action.

Stress, creativity and digital collaboration

Many participants noted their enjoyment at the physicality and messiness of the exercises. They also enjoyed discussing their perspectives about the climate emergency with a group of people they wouldn't usually interact with - something that was possible due to hosting the workshops fully online.

Tackling the climate emergency is both stressful and challenging. In times of stress (e.g. covid lockdowns), many people turn to creativity and craft. We also rapidly learn how to overcome challenges (e.g. using video calls and working at home). Collaborating and encouraging creativity digitally can be a means for working together and improving how we do things.

Recommendation [143 - 144]: UKRI should enable mechanisms for researchers from across disciplines to work together creatively. This will encourage sharing of perspectives, discussion of solutions and how to implement changes for digital research infrastructure and wider research practices.

Physical and virtual artwork pieces

Drawings, images, and texts from the creative workshops (and the offline exercises) have been transformed into narrative reliefs. Inspired by some of the world's earliest examples of narrative art and information sharing via pottery, Paul created six one-metre-high vases using cutting-edge 3D ceramic printing technology.

These ceramic cylinders - and vessels of 'data' - conjure up a sense of monolithic forms. They are decorated in a way that reflects the digital research infrastructure transformation and the perspectives of the workshop participants.

As one of the oldest human inventions, the practice of pottery has developed alongside civilization. Ceramics is an ancient and enduring art-form, which is now being revolutionised by the latest printing technologies and digital design practices.

Clay, as a sustainable material, has the ability to change and take on different forms, to last, or to be broken down and re-made. Using clay for this commission is a metaphor for lasting change - and the change that is required for reaching net zero digital research infrastructure.

A 3D playable virtual environment has also been designed by Paul, to be experienced through a web browser, virtual reality headset, or film projections. Hologram-like versions of the vases are displayed inside a re-imagined digital infrastructure environment, for virtual visitors to explore and interact with. A way for people to experience the artwork in their own time and place. The online installation can be freely viewed.

The physical vases and virtual reality space form an immersive installation and visitor experience. A one day exhibition occurred on Thursday 11 May at the Advanced Research Centre in Glasgow. This exhibition was alongside the cross-UKRI stakeholder workshop (see Section 5.4.4.4). A selection of images from this event can be found below.

The in-person and online installations stimulate further thinking and conversations on climate change, digital infrastructure, and the future of UK research.

Figure 3.4.6C: in total there were 7 ceramic vases each based on conversations from the creative workshops. One of the vases provided an interactive element for attendees where you could make your mark in the clay. The vase in the foreground of the image is based on a butterfly wing - importance of nature was a very common discussion during the workshops.

Figure 3.4.6D: the exhibition room during the private viewing for attendees of the cross-UKRI stakeholder engagement workshop.

Figure 3.4.6E: this ceramic is based on a sound recording taken from the machine room of a UKRI funded digital research infrastructure (JASMIN).

Conclusions

For the workshop participants, and for members of the UK digital research infrastructure community more broadly, the art commission offers the opportunity to observe what can be created by bringing different disciplines together for collective action. It has been particularly valuable in understanding different perspectives, challenges and knowledge (or lack of it) about digital research infrastructure. We found it inspiring to hear the passion, emotion and interest from a range of people - many of whom we would not have engaged with if this work had not been undertaken.

We have several future aspirations for this piece of work, including:

- Using the physical exhibition to inspire conversations at events - both within the digital research infrastructure community and more broadly (e.g. at public events like the STFC Daresbury and Harwell open weeks).
    - There could also be the option for running creative workshops alongside a physical exhibition.
- Encourage other UKRI funded projects to consider working outside of your comfort zone and engage with others who may have different perspectives to you.

This art commission has been a brilliant example of how working across UKRI sectors can bring together diverse groups and stimulate conversations on climate change, digital infrastructure, and the future of UK research. We found the process particularly valuable in terms of helping us to shape the framing of the wider projects ‘big picture’ messages - to ensure they are reflective and representative of the communities we have engaged with.

The assets created could be re-used by others to stimulate further thinking and conversations (subject to resource availability). Please contact Poppy Townsend (via support@ceda.ac.uk) if you are interested in hosting a future exhibition or are interested in finding out more about the process behind completing this work.

3.5 Community and consensus building activities

3.5.1 Community building

Throughout the project, the core team has built a community of interested parties. This has been achieved via a range of different methods.

Mailing list

A project mailing list was gathered and built on since the start of the project (November 2021). The mailing list exceeds over 300 contacts. The mailing list contact ranges from UKRI cross council members, academia, industry and general public. The mailing list was managed by the project's communication manager and communication support resources. For cross council engagement, the project team used the UKRI staff directory to initially identify communication support. Once identified, emails were sent to ask for communication support (level of support and communication in each internal council). In addition, the stakeholders were mapped into categories of engagement, therefore depending on the stakeholder categorisation, the communication method was specific for each group.

Mailing list information included various information that either raised awareness, engaged with the audience or provided information. For example, project surveys, social media updates (project website, twitter, linkedin), webinar information and funding opportunities.

Stakeholder Engagement

The project created a stakeholder engagement and communications plan to; establish methods of communication, identify and map stakeholders, to enable effective communication and engagement. The project team identified the project stakeholders (internal and external) and assessed their interest and influence. Each stakeholder was mapped into one of four categories; Category A (High Influence, High Interest), Category B (High Influence, Low Interest), Category C (Low Influence, High Interest) and Category D (Low Influence, Low Interest). Mapping identification allowed the project team to define the level of the relationship, identify specific engagement strategies, understand communication delivery method and frequency and mitigate potential risks. The project team shared the stakeholder engagement and communication plan with all stakeholders for review. Feedback was incorporated and version control applied. Stakeholder engagement and communication plan was regularly reviewed throughout the project lifecycle and updated where required.

Meetings with relevant community members or organisations

The project team effectively engaged with a variety of stakeholders through multiple channels of communication and engagement. For example, regular webinars presented by the project team, posting on social media, updating the project website, presenting at events and conferences. In addition, the project team assembled project boards with the relevant community members for strategic direction, guidance and decisions. Through project board members and the growing community, additional contact details were shared. Contact details ranged from internal cross council connections and external connections.

Project board meetings consists of:

Steering committee members - Engineering and Physical Sciences Research Council, Medical Research Council, Met Office, National Centre for Atmospheric Science, University of Edinburgh, University of Glasgow,

Scientific Advisory Board - Cambridge University, National Centre for Atmospheric Science, Natural Environment Research Council, Imperial College London, University of College London, University of Glasgow

Head office meeting - National Centre for Atmospheric Science, Natural Environment Research Council, Science Technology and Facilities Council

Sandpit Projects

Brunel University, Heriot Watt University Manchester Metropolitan University, National Oceanography Centre, University of Bristol, University of Cambridge University of College London, University of Durham, University of Edinburgh, University of Lancaster, University of Newcastle, University of Oxford, University of Reading, Queens Mary College of London, Science Technology and Facilities Council

Consortium

Edge Hill University, King's College London, National Centre for Atmospheric Science, Scientific Computing Department, University of College London, University of Edinburgh, University of Glasgow

Additional stakeholders

Ada lovelace Institute, Cancer Org, Catapult, DARE UK, Decarbonize Org, Darvis, Digital humanities Climate Coalition, FCB Studios, Health Data Research Innovate UK, Health Foundation, Intel UK, High End Compute, Microsoft, NC3R, Novopower, Tech UK, Turing, University of Exeter, University of Leicester, University of Manchester, University of Sheffield, University of York, UKRI, UKRI Environmental Sustainability Team, Wellcome,

3.5.2 Information sharing events and resources

These events and resources were created by the core project team in order to aid community understanding about the project throughout its development.

Webinars, Sandpit and Consortium Events

Different types of webinars were held throughout the duration of the project. In the initial phase of the project, monthly webinars were held to help with community building and engagement. Webinars and drop in sessions were held from May to December (2022) for the sandpit and consortium projects, to provide project progress and help with any questions. A project partner slack channel was also created to facilitate team building and help with project synergies.

Findings and Recommendations Webinars held in January 2023.

The first webinar held was on technical and operational challenges, this included a range of projects presentations (7 presenters), with a total of 60 attendees. The second webinar held was on community and organisational challenges, this included a range of project presentations (5 presenters), with a total of 44 attendees. Both webinars can be found on YouTube, webinar 1 can be found here and webinar 2 can be found here. The slides for both webinars can be found, webinar 1 can be found here and webinar 2 can be found here.

Conferences

Computing Insight UK Event 2nd of December 2022 - Project team presented at the event and hosted a talk panel, with speakers and questions.

PV Conference 2nd to 4th of May 2023 - Project team to present a Net Zero DRI Project poster. Conference detail were shared with all project partners, encouraging collaboration and aiding community building.

Stakeholder Engagement Sessions

Project team held numerous stakeholder engagement sessions to help build the community, list is as followed;

- Values & Responsibilities Workshop - Online, Friday 24 March 2023
- Early Career Researchers Workshop - Online, Tuesday 28 March 2023
- Community Engagement Workshop - Online, Thursday 30 March 2023
- Early Career Researchers Workshop - London, Tuesday 4 April 2023
- Scenarios and Roadmap meeting - Online, 26th April 2023
- Cross UKRI stakeholder engagement workshop - hybrid, Thursday 11 May 2023

Digital engagement

The project engaged with the target audience across multiple different channels, list is as followed:

- Project website, the link can be found here - UKRI Net Zero Digital Research Infrastructure Scoping Project | net-zero-dri - UKRI Net Zero Digital Research Infrastructure Scoping Project (ceda.ac.uk)
- Twitter Account, the account can be found here - @cedanews
- Email account, the email address is here - support@ceda.ac.uk

3.5.3 Evidence gathering events

These events were hosted by the core project team in order to gather additional evidence not covered by other parts of the project. These were close invite events for key stakeholders or domain specific experts.

3.5.3.1 Training and standards workshop

Our interim report (Juckes et al., 2022) highlighted a gap around existing training and standards to support environmentally responsible research using the DRI. In addition to which, in-depth case studies of HPC facilities (i.e. JASMIN (Lambert, Stephens & Kayumbi, 2023)) further highlighted the need for training at the user level (see Section 3.4.3). In order to determine what kind of training and standards might be developed and implemented in response to these findings, the project convened an online workshop on Tuesday 24 January 2023 with invited experts to focus on the topic. 14 attendees represented UKRI research council staff, project partners and external organisations including Health Data

Research UK (10 invitees were unable to attend, but contributed to the discussion outcomes and consensus priority actions).

The discussions at this workshop (see Table 3.5.3.A) were indicative of support for UKRI investment in a hub or centre of excellence for environmentally responsible use of digital research infrastructure (DRI). The agreed ambition of this service should be to:

- Act as a centralised repository of 'green' resources, training opportunities and expertise
- Ensure that resources are visible and accessible to a range of stakeholders across UKRI research councils and the wider research ecosystem
- Coordinate reporting and monitoring (as part of developing meaningful standards and metrics)

Where existing training and/or professional development provision already exists, it would be pragmatic to integrate environmentally responsible practice in software design, data management and facilities management (see Table 3.5.3.A, Actions A, B, G, H).

In addition, funding must be leveraged both to support access to training and continuing professional development, as well as incentivise environmentally responsible research design and monitoring by embedding related criteria in research proposals and evaluation (see Table 3.5.3.A, Actions L-M).

<u>Design & outcomes</u>

The workshop focused on the types and value of training and standards with respect to various levels and operations of the DRI.

Ahead of the workshop contributors were asked to respond to two prompts using an online whiteboard (Miro), which formed the basis for live discussion:

1. What behaviours/activities might have a significant impact on the UKRI DRI carbon footprint, and which of those were likely to have the largest impact or greatest significance at scale.
2.
    a. Which of those activities or behaviours can be significantly addressed by advances in training and/or standards?
    b. Additional considerations included whether training is specialist or general, one-off or repeated; and regarding the scope of standards (narrow yet impactful or lower-impact but scalable).

Four break-out groups each discussed one of the emergent topics from the online whiteboard:

1. User behaviour, Application and Software Development
2. System and Facility Management, and Policy
3. Data Lifecycle management, Data sharing and Open Compute
4. Funding incentives, Assessment criteria and Strategic Policy

Each group was facilitated by a member of the project team or project partners.

All the suggested actions arising from discussions are listed in Table 3.5.3.A The wording of actions were refined by the project team, and subsequently circulated to all invitees for final comment and agreement. They were asked to select and rank three actions, generating three high priority actions through consensus:

1. Train individual users in good practice, such as code testing and optimisation, reflecting the latest knowledge and tools (Action B)
2. Train facility managers to monitor and evaluate energy/carbon usage at both user and system level (Action A)
3. Develop/agree standards for metrics reporting to enable consistent measurement and monitoring of carbon usage, across platforms, sectors and institutions (Action F)

These priority actions are interdependent. They reflect the importance of ensuring that users of the DRI at all levels are equipped with the skills and knowledge necessary to enact environmentally sustainable practices across the UKRI research community (priority actions 1-2). They also emphasise the need for standards to be established to enable consistent reporting and evaluation of progress (priority action 3). This is essential to enabling joined-up progress towards emissions reductions across all levels of the research ecosystem.

Table 3.5.3.A: Actions arising from group discussions held during the Training and Standards workshop. Citable references (as listed in Appendix 3) are included as [N].

Discussion	**Suggested Actions**
User behaviour, Application & Software Development	A) Train facility managers to monitor and evaluate energy/carbon usage at both user and system level [145] B) Train individual users in good practice, such as code testing and optimisation, reflecting the latest knowledge and tools [146] C) Train dedicated teams of Research Software Engineers (RSEs) to provide cross-sector support and optimise scientific code for deployment [147] D) Develop Open Science and FAIR data standards and train researchers in them to maximise good practice, efficiency, data sharing, discoverability and reuse [148] E) Create mandatory training within UKRI on WHY it is important to change working practices to achieve Net Zero [149]
System and Facility Management, and Policy	F) Develop/Agree standards for metrics and reporting to enable consistent measurement and monitoring of carbon usage, across platforms, sectors and institutions [150]
Data Lifecycle management, Data sharing and Open Compute	G) Provide training to individuals on carbon-efficient data management practices [151] H) Extend the scope of Data Management Plans to include the environmental impact of the complete data lifecycle [152]

	I) Develop accreditation for training in environmentally sustainable data management as part of ongoing professional development [153] J) Develop standards for the delivery and exploitation of big data through carbon-efficient services and software. E.g. server-side subsetting to reduce data transfer and storage loads [154]
Funding incentives, Assessment criteria and Strategic Policy	K) Mandate the inclusion of an environmental impact statement, along with mitigating actions, within applications for research funding [155] L) Include a budget within funding calls to support researchers in engaging with Net Zero goals, e.g. general training for scientists or access to specialist expertise. [156] M) Mandate carbon monitoring and reporting in funding calls [157]

3.5.3.2 Procurement workshops

Two workshops were held on funding of digital research infrastructure procurement, its operation and use in UKRI with stakeholders from UKRI including NERC, MRC, STFC and a number of research facilities. These were held on Thursday 22 September 2022 and Monday 6 February 2023.

In the first of these we discussed best practice, understanding aspirations, describing barriers, understanding how we work with others and understanding and managing the rebound effect.

The focus of the second workshop was to review issues and recommendations, and then to discuss and agree an outline consensus on procurement recommendations.

The recommendations below provide a clear and stable long-term vision aimed at providing consistency and clarity in the face of rapidly evolving opportunities and threats, while examples provide specific changes that can be implemented no to deliver immediate progress. Citable references (as listed in Appendix 3) are included as [N].

1. Standards and Guidance for Spending Decisions

Spending decisions are made in many contexts and often in many layers. Many aspects of the carbon budget are baked in by decisions high in the decision tree. Responsibility is shared but proportionate to the influence at each stage. Significant change will depend on action being taken at the higher levels, but this may depend on or require behaviour change from a wider user community.

RECOMMENDATION 1.1 [157]: All spending decisions must include a proportionate assessment of their impact on the UKRI carbon budget and on the implementation of the Net Zero policy. In most cases staff do not have the tools and knowledge at hand, so work on training (recommendation 3.1) and standards (recommendation 1.2).

- **Example 1.1.1**: UKRI is considering an allocation of £500k to refresh DRI facilities which will be allocated through an Announcement of Opportunity (AO). The specification of the call should ensure that the expectations are aligned with Net Zero policy. Bids must include an appropriate level of consideration of carbon footprint in an environmental sustainability assessment, and guidance on preparation of the environmental sustainability assessment must be included in the call. Care must be taken to ensure that call specification does not impose constraints on environmentally sustainable decision making by bidders.

- **Example 1.1.2**: A facility manager is procuring equipment to maintain services. They must provide an appropriate life cycle assessment of the equipment.

RECOMMENDATION 1.2 [158]: Standards need to be developed to ensure that environmental sustainability assessments are made robustly and efficiently. There are many existing standards, but there are also cases for which new standards need to be set or developed.

- **Example 1.2.1**: When information and communication technology (ICT) equipment is being purchased an estimate based on bulk allocation of carbon footprint by price applied to the global market should be used, using a figure provided by UKRI and updated annually. The current value, for embodied carbon, is 0.15 kg CO_2e per GBP of ICT investment (see Section 1.3.3 above). Additional figures providing specific information can be provided if available. The bulk allocation of carbon footprint by price can be used at all levels of the decision hierarchy, from Treasury downwards.

- **Example 1.2.2**: Life cycle assessments need to be created with an appropriate level of detail taking into account the limitations of the information that is available. For the use phase of equipment, the environmental impact of power supply must consider both a central estimate of usage based on the Climate Change Committee balanced pathway and a risk assessment looking at the consequences of potential departures from the pathway or changes in policy.

RECOMMENDATION 1.3 [159]: UKRI must have a policy on overall power consumption of facilities which is aligned with the Climate Change Committee balanced pathway to net zero. Although the overall consumption barely registers on the scale of the national sectoral analysis considered by the CCC it is important that UKRI should provide leadership in explaining how their investment decisions align with the CCC recommendations, including the recommendation that power consumption for existing activities needs to be held constant or reduced in order to enable a timely transition to renewable power.

- **Example 1.3.1**: To maintain or enhance their place as a world leading scientific community, UK physicists need access to a new generation of facilities which will result in a substantial increase in power draw. Capping the power draw would limit the UK to facilities which would be considered at best as second rate in comparison with international competitors. In order to deliver the level of innovation required by the UK society, economy, and by the assumption of high innovation which is implicit in the CCC balanced pathway, facilities which expand the power draw are needed. UKRI must develop a policy which sets out both the rationale for an expanded power draw and the steps that need to be taken by facility planners to ensure consistency with the UKRI policy.

2. Pitfalls and Unintended Consequences

RECOMMENDATION 2.1 [160]: UKRI must take a proactive approach to ensuring that net zero policy does not disrupt research programmes and that prioritisation of low carbon investment and purchasing options does not have a disproportionate negative impact.

- **Example 2.1.1**: Where zoom meetings replace in-person meetings we lose the direct experience of seeing our colleagues in their working environment and opportunities for informal and unscripted interactions which often seed creative thinking. UKRI must ensure that both their directly employed and their funded staff have adequate opportunities for informal and unscripted interactions with colleagues and stakeholders.

- **Example 2.1.2**: Procurement policy should be co-developed with vendors, establishing opportunities and risks, to ensure that policy is fit for purpose.

- **Example 2.1.3**: A question is asked about the overall approach, who is ultimately responsible? As with health and safety, every employee has responsibility for their area of influence, but overall responsibility should be held by the CEO.

RECOMMENDATION 2.2 [161]: UKRI must ensure that steps taken to reduce environmental impact do not end up having the opposite effect through feedbacks such as the rebound, or Jevons, effect.

- **Example 2.2.1**: UKRI should monitor the impact of spending decisions and changes to policies and processes that govern spending decisions. Monitoring should be done through both quantitative metrics and expert assessment, with provision for open discussion and anonymous feedback, to ensure that effective communication about positive and negative consequences.

3. Training and Learning

RECOMMENDATION 3.1 [162]: Training of staff at all levels is needed, both to increase awareness and understanding of the implications of climate change and the net zero policy, and to provide technical competence to deliver change. Training needs to be backed by an active programme of learning and discovery. The roadmap to net zero will pass through unexplored territory and training material will need to be regularly updated with lessons learned from exploratory pathfinder projects at UKRI and elsewhere.

- **Example 3.1.1**: Staff involved in making significant DRI procurement decisions need to be provided with training and support for collective development of guidance on appropriate assessment methodologies in the context of their work.

- **Example 3.1.2**: Staff preparing business cases for major DRI investments should be provided with support to enable effective integration of environmental sustainability planning into investment and operation plans.

- **Example 3.1.3**: Investments in new facilities should include adequate provision for training of operators and users of the facilities to ensure maximum gain from the financial and carbon investments.

3.5.3.3 Values and responsibilities workshop

In order to effectively implement change to achieve net zero computing, a shared set of values and distributed responsibilities are needed. Responsible research practice is organised according to sets of values and/or principles which are reflected in mandatory requirements (e.g. data governance, ethics risk assessments) and funding criteria such as fiscal efficiency and derivative benefit of research outcomes. The need to reduce carbon emissions associated with DRI can come into tension with existing values/principles guiding research practices at all levels of the ecosystem. 27 experts were invited to participate in this online workshop on Friday 24 March 2023. Those attending represented a range of stakeholders across the research ecosystem including: UKRI research councils, UK funding bodies (Wellcome and Cancer Research UK), research institutions (Sanger and Ada Lovelace), researchers and project partners. Representatives from external organisations experienced in effecting sector wide change (NC3Rs[48] and Arts Council England[49]) were also invited in order to offer insight into affecting change at the level of communities of practice.

Key messages arising from this workshop aim to address the practical need to integrate environmentally sustainable use of DRI across research design, evaluation and implementation.

They have been framed as suggested actions. Citable references (as listed in Appendix 3) are included as [N]:

A. leverage institutional power: demonstrate action at the institutional level, with visibility across the research community in order to ensure cooperative organisational and individual responsibility and inspire positive change [163]
B. make information about the relative benefits of different actions clear and readily accessible to the research community, in order to empower decision making by individual researchers and groups [164]
C. include environmental sustainability within funding assessment and award processes, so that it is planned into the project from the outset, and evaluated as part of funding applications [165]
D. request that projects estimate their carbon footprint, even crude ones[50] [166]
E. provide/develop a rating of host (research) institutes by the sustainability of their operations and projects [167]

These actions have been organised roughly with respect to contingencies. An example is readily made regarding project's needing to report on their carbon footprint (action D). In order to ensure a fair burden of responsibility, this will require host institutes to have average CO2 emissions for their activities, e.g. compute, project meetings, travel, etc. This needs to be available to researchers before they can comply with any reporting requirements. The actions here lend support to every strategic theme in the toolkit (Section 2.1), with a particular emphasis around supporting communities of practice through sharing of knowledge, resources, and emerging best practices.

<u>Design and outcomes</u>

---

[48] NC3Rs is the UK national centre for the replacement, refinement and reduction of animals in research: https://www.nc3rs.org.uk/

[49] Arts Council England (ACE) is the UK national development agency for creativity and culture, which has pioneered the inclusion of environmental impact reporting in the international landscape since 2012: https://www.artscouncil.org.uk/blog/taking-action-environmental-responsibility

[50] "The proposal would say: we're planning to run 1,000,000 CPU hours of compute, and host 3 in-person project meetings, each with 10 people travelling from the UK, 5 from the US, so the total CO2 would be XYZ.

The workshop included presentations on the DRI user survey results (see Section 3.4.3) the sustainability of data-driven health research, and two case studies. One from NC3Rs on the adoption of methods and practices to replace, refine and reduce animals used in research, and another on the safety transition analogy. Invited attendees were selected on the basis of their expertise with respect to (i) research concerning ethics and values associated with environmental sustainability of digital (research) methodologies and infrastructures, (ii) funding assessment and associated decision making processes and (iii) affecting change within research or other communities of practice.

The workshop aimed to:

- determine which values and/or factors pose the most significant challenge for implementing environmental sustainability during decision-making processes and/or practices associated with use of DRI (including Big Data, AI, and Algorithms)
- highlight existing case studies of affecting change where conflicting values have been aligned/resolved
- identify existing mechanisms that could be useful in supporting decision making that integrates environmental sustainability principles (e.g. ethical frameworks, carbon reduction guidelines)

The outcomes which follow reflect one or more of the aims above.

First, factors which presently guide decision making in research design, implementation and evaluation were highlighted according to those which are presently afforded high priority. These include: cost, quality, overarching strategic direction or policy of research institutions and funders, benefit or impact, ethical research practice, and novelty or 'innovativeness' (Figure 3.5.3.A).

Figure 3.5.3.A: Factors identified as guiding and influencing decision making in research design, implementation and evaluation, generated as part of a pre-workshop activity remotely using Miro board.

A second set of insights into considerations for integrating equality, diversity and inclusivity (EDI) with the principle of environmental sustainability was also captured as part of the same pre-workshop activity. These are highlighted in Figure 3.5.3.B, and reflect the factors which could or should inform how environmentally sustainable decisions relating to DRI use and investment are made. By including evaluation of harms or potential impacts to biodiversity, human and non-human, health across geographies and communities beyond those targeted by the research focus (*"but also in terms of those who will be affected by mitigation inaction"*, Figure 3.5.3.B), the net zero DRI objective can be viewed through an anti-colonial and equitable lens. Given that environmental sustainability is the principle driving net zero targets at national and institutional levels, these factors reflect a wider scope for capturing the impacts of DRI, beyond GHG emissions.

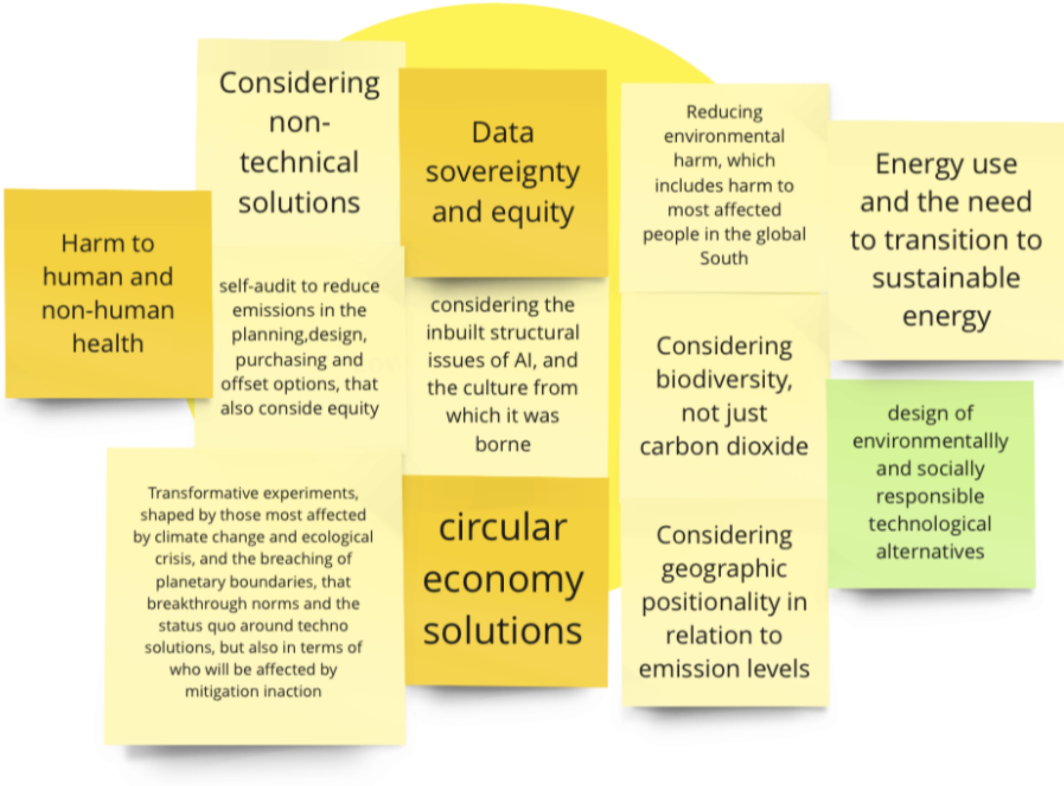

Figure 3.5.3.B: *factors which require consideration in order to integrate EDI with sustainability across digital research practices. These inputs were generated as part of a pre-workshop activity remotely using Miro board.*

Factors which were considered to be in tension or aligned with the principle of environmental sustainability of the DRI were drawn out of the second pre-workshop activity (Figure 3.5.3.C). Contributors were asked to 'list the factors that align with and/or come into tension with the principle of environmental sustainability in decision making processes associated with the use of DRI (including Big Data, AI, and Algorithms)'. Factors considered to be in alignment include: responsive policies (e.g. efficient facilities usage, monitoring energy consumption), technical expertise (enabling 'green' use of software and hardware), continued professional development (CPD; to support energy efficient procedural change), data sovereignty and FAIR data principles (i.e. data that is Findable, Accessible, Interoperable, Reusable[51]). Factors in tension with environmental sustainability are framed according to associated increases in demand for DRI and usage. For example, advances in technological capability can increase efficiency, but this is often framed as a gain of capacity,

[51] https://www.go-fair.org/fair-principles/

rather than a reduction in emissions, frequently leading to net increases in use (the rebound effect). Similarly the interdependence of software and hardware can preference replacement over maintenance of existing hardware, where new programmes are incompatible with existing infrastructure. These factors can be exacerbated by strategic objectives which prioritise scientific advancement and impact of research outputs without consideration towards the environmental impact.

These factors accompanied discussion in break-out groups. All groups focused on the relevance of existing and suggested frameworks or guidelines to address conflict between principles and to drive more environmentally sustainable use of DRI. Discussion also generated examples of existing frameworks, guidelines, and standards which might be exploited to support the net zero agenda (Table 3.5.3.B). These are suggested examples, and further investigation is required in order to determine actual utility. However, the exercise serves to highlight an existing body of knowledge and tools which share a common aim in their application; to ensure environmentally responsible practice across procurement and operational practices.

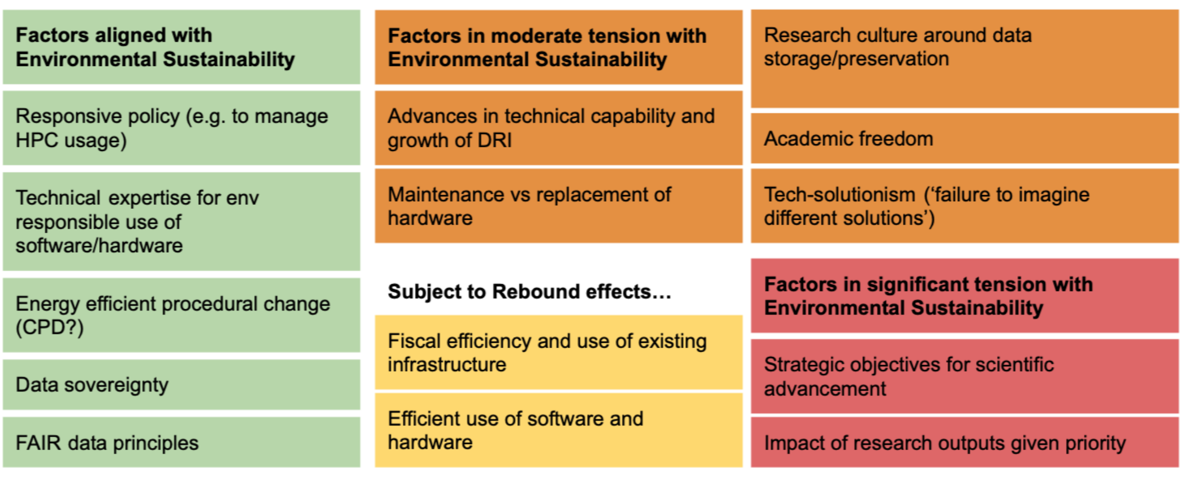

Figure 3.5.3.C: *The factors outlined here emerged from inputs generated as part of a pre-workshop activity remotely using Miro board. They were grouped according to whether contributors indicated that they aligned or conflicted with the principle of environmental sustainability by the project team in advance of the workshop, and shared with attendees via presentation and a follow-up opportunity for comment.*

Table 3.5.3.B: *existing frameworks, guidelines or standards shared as part of breakout discussions. Participants were asked to consider available tools relevant to driving environmentally sustainable decisions, which have relevance to the net zero transition of the DRI.*

Framework, guideline or Standard	Target audience	Aim	Link
AA1000 Stakeholder Engagement Standard	Businesses, investors, governments, and other entities	To integrate effective stakeholder engagement with strategic operations	https://www.accountability.org/standards/
B Corp Certification	Businesses	To monitor and evaluate social and environmental impact	https://www.bcorporation.net/en-us/certification

Global Reporting Initiative (GRI) standards	Businesses, and other entities	To create global standards for organisations to report their (environmental) impacts	https://www.globalreporting.org/standards/
SBTi's Net-Zero Standard	Private Corporations	To enable organisations to set science-based emissions reduction targets	https://sciencebasedtargets.org/net-zero
GHG Protocols	Businesses, governments, and other entities	To create global standards and frameworks to measure and manage greenhouse gas (GHG) emissions	https://ghgprotocol.org/standards
Sussex Sustainability Research Programme (SSRP)	International, national and local stakeholders (various)	Research projects on SDGs, to understand trade-offs and co-benefits	https://www.sussex.ac.uk/research/centres/sussex-sustainability-research-programme/about
Procurement Law			https://bills.parliament.uk/bills/3159

3.5.4 Feedback and consensus building events

These events were hosted by the core project team towards the end of the project to ensure that feedback was gathered about the project outputs and consensus was reached (where possible).

These events generated three key recommendations:

- The community are looking to UKRI to lead the way and go further than simply reaching net zero. UKRI must create, use and share open, adaptable, iterative processes and resources about net zero DRI that can be critiqued and contributed to by the research community. UKRI should ensure processes and funding routes are flexible and able to evolve to match the changing requirements - whilst embedding environmental sustainability into every aspect of UKRI.
- UKRI must provide clear messages and narrative to empower the research community to make the necessary changes needed for net zero DRI. Providing support, examples of best practice and relevant success stories are essential. These should be provided via a central service in an open access way. These resources could be relevant to wider net zero agenda, and should not be solely limited to the DRI.
- Cross-disciplinary community building is essential for continuing the momentum needed for transformational change. UKRI should fund cross-disciplinary activities that bring together members of different research areas to continue community building work and support (many currently unfunded) grass-roots activities.

3.5.4.1 Early career researchers workshops

The proposed changes associated with this project's recommendations will implicate researchers at every career stage, and across varied fields of research. It was therefore considered pertinent to engage directly with early career researchers (ECRs), defined as within 10 years of completing an undergraduate degree or equivalent.

Two workshops were held, one online on Tuesday 28 March 2023, and one in person at the Science Gallery London on Tuesday 4 April 2023. The workshops sought to discuss the implications, changes and barriers associated with the project's strategic themes, as well as how ECRs can influence change at institutional levels.

The workshops consisted of a mixture of presentations and discussion sessions. Discussions were fruitful, with many attendees positively engaged in the work.

The online workshop had 52 people register from a range of research disciplines, 12 people attended on the day. The in-person workshop had 30 people register to attend, 6 people attended on the day. Whilst these attendance rates were disappointing compared to the high interest shown via registration, the discussions were still incredibly valuable to the project team and the participants of the workshops.

Figure 3.5.4.A*: a word cloud created by attendees at the online early careers workshop showing the diverse ways they use digital research infrastructure.*

Based on discussions, three areas for recommendations emerged (citable references (as listed in Appendix 3) are included as [N]):

1. **Ensure collective responsibility for Net Zero DRI [168]**
    - move towards an empowered and equitable research community, bringing top-down and grass-roots actions together
    - develop a shared and accessible body of knowledge to support collective action, avoiding silos of practice

2. **Resource green tools, including metrics, guidelines and continued professional development (CPD) [169]**
    - ensuring resource availability to address capability gaps between researchers in High Income Country and Low and Lower-middle Income Country settings

3. **Embed environmental sustainability across all research activity [170-171]**
   - integrate environmental sustainability within graduate and post-graduate courses to support workforce demand for ‘green skills’
   - ensure that academic career progression is supported via alternatives to flying, and improvements to virtual networking capability

The discussion and themes captured demonstrably overlap with the project’s strategic themes, and emphasise the importance of a holistic and coherent approach to net zero computing which supports engagement and inclusion of ECRs. Discussion at the events also helped emphasise the importance of framing around the narrative given in this technical report. The team have included changes to this report’s introduction based on the views of early career researchers.

<u>Design and outcomes</u>

These workshops aimed to engage ECRs with the project, set in the context of UKRI Environmental Sustainability Strategy, and the growing application of Digital Research Infrastructure (DRI) across subject domains.

The ECRs attending represented a range of academic disciplines, and were invited via public online platforms (Twitter, LinkedIn), departmental mailing lists for social and medical sciences, and informatics research at King’s College London, networks accessed via UKRI research councils and communications from the Science Gallery London.

During break out discussions attendees were asked to consider some of the strategic themes (indicated in bold) arising from the project’s evidence base so far:

1. **Shared responsibility from personal to collective**
2. **Working with others - from peers to suppliers**
3. **Developing new knowledge and continuous improvement**
4. **Leadership and guidance for the UK research community**
5. Mission focus: appropriate response to the climate emergency
6. Action-based-research / learning-by-doing / just-get-on-with-it

These were chosen by the project team on the basis that they held most relevance to ECRs, and ensured focused discussion.

Across two break-out discussions groups were asked to think about some of the recommendations, and asked to discuss their present ‘perspectives and values on the digital research infrastructure’ and their ‘future aspirations for environmental sustainability’ in relation to the recommendation. These topics were echoed in discussions held within the Community Engagement workshop (Section 3.5.4.2). Comments and suggestions were captured, and subsequently grouped by the project team to produce cross-cutting themes (summarised in Table 3.5.4.A).

Table 3.5.4.A: *Summary of themes arising from early careers researchers during break-out discussion considering the strategic themes according to their present perspectives and future aspirations for a net zero DRI*

Present perspectives	Future aspirations
Theme 1: Ensure collective responsibility for Net Zero DRI	An empowered and equitable research community

<table>
<tr><td></td><td>A shared & accessible body of knowledge</td></tr>
<tr><td>Theme 2: Resource green tools, metrics, guidelines and CPD</td><td>Resourcing goes beyond green tools, metrics, guidelines and CPD…</td></tr>
<tr><td colspan="2">Theme 3: Embed environmental sustainability within:</td></tr>
<tr><td>Academic career progression</td><td rowspan="2">All levels of research ecosystem</td></tr>
<tr><td>Research design, assessment and implementation</td></tr>
</table>

3.5.4.2 Community engagement workshop

A community engagement workshop was hosted online, by the project team, on Thursday 30 March 2023 to engage with UKRI funded researchers about the project's strategic themes. The main aim of the workshop was to obtain feedback from a cross-section of the UKRI academic community, with balanced participation achieved by a random sample from UKRI Principal Investigators.

The purpose of the workshop was to:

- share initial findings, from the UKRI Net Zero DRI Scoping project, with researchers funded by UKRI
- capture how the digital research infrastructure is used across research councils
- generate insights into possible challenges re: implementation of environmentally sustainable recommendations for the digital research infrastructure
- generate discussion about our project recommendations and identify gaps for further exploration
- obtain feedback from a cross-section of the UKRI academic community, with balanced participation achieved by a random sample from UKRI Principal Investigators

Over 150 researchers were invited to attend, chosen via a balanced cross-section of UKRI researchers funded by each research council in 2020-21. Despite regular invites sent to the researchers, only 20 people registered for the event, and 8 people attended.

Research areas (and interests) represented by attendees included statistics (design and analysis of experiments), software engineering (approaches to support food security and a circular economy), engineering for sustainable development (investigating a circular economy in sub-saharan africa), genetics (gene expression and regulation) and biomedical engineering (assistive technologies and disability).

Key themes

The attendees were asked to consider the same questions and strategic themes as described in the early careers workshops (Section 4.4.4.1). Discussion initiated around the perspectives and values on the DRI amongst the researchers present generated suggested actions which were grouped by (Citable references (as listed in Appendix 3) are included as [N]):

A. Challenging assumptions re: tech solutionism and interrogate scope of benefit vs risk [172]
B. Improving interoperability and accessibility of infrastructure [173]
C. Producing guidelines for procurement and best practice [174]
D. Investing in green resources, training and skills [175]
E. Embedding green principles in funding processes [176]
F. Prioritising environmental sustainability and raise awareness [177]

A subsequent discussion focusing on a future, environmentally responsible DRI resulted in three further groups of actions, deemed necessary to ensure an equitable, diverse and inclusive future:

G. Implement learning from applicable case studies to support environmentally responsible use of DRI (e.g. UKRI leadership in reduction, replacement and removal of animals use in research, Wellcome Trust leadership in promoting public engagement with research**)** [178]
H. Incentivise environmentally sustainable research practice [179]
I. Advocate for sharing of resources including infrastructure (facilities) and data [180]

3.5.4.3 Scenarios and roadmap meeting

The 'Scenarios and Roadmap' meeting took place on Wednesday 26 April 2023. The meeting was divided between discussion of a draft (version 1) of the roadmap, which was presented to invitees one week ahead of the meeting, and scenarios influencing the net zero ambition.

The purpose of the meeting was articulated by the following aims to:

- share our roadmap to a net zero DRI and identify any barriers to effective implementation of recommendations
- discuss how best to frame an on-going review process for the implementation of the roadmap
- discuss scenarios influencing the net zero ambition for DRI (and more broadly UKRI's strategic objective for net zero across operations) by 2040 or earlier. In particular we would like to discuss the **Climate Change Committee (CCC) balanced pathway to net zero.**

Attendees were invited on the basis of expertise with respect to DRI facilities management, policy research and implementation, and strategic guidance. The discussion generated was intended to provide feedback to the project team with respect to the above aims. Key areas of discussion concerned (i) mitigating against the rebound effect, (ii) ownership and responsibility of proposed actions/the net zero agenda within the DRI (iii) mitigating against tensions between net zero and existing policy areas within UKRI (iv) priority areas for action (v) 'good citizen' behaviour of data centres and/or large scale DRI facilities (i.e. ARCHER2 back-up generators have been able to provide energy to the grid during periods of high demand over winter), and (vi) evaluating research reliant on the DRI with respect to carbon impact.

On the latter point, it was noted that Carbon should be considered a limited resource, in the same way as financial resources. In the financial case frameworks and processes already tension different kinds of research in order to support investment decisions, and it is therefore pragmatic to consider Carbon under a principle of constraint[52]. This is consistent with recommendations arising from the ARINZRIT and VALUE sandpit projects (sections 3.3.1 and 3.3.7 respectively), to develop mechanisms for allocating carbon budgets to areas of research/or for considering when use of DRI is justified in light of the societal challenge of net zero (balancing the projected outcome or ambition of research with the means used to achieve it - the impact of research practice).

Integrating environmental sustainability into existing 'cost-benefit' analysis processes in UKRI investment decisions was agreed to be central to addressing the challenge of net zero. To this end, steps are already underway as part of the UKRI Environmental Sustainability Strategy. Consideration of the DRI requires specific tools - some of which are unavailable/not yet developed. However the overarching message from projects (Sections 3.2-5) is that there is not time to wait until we have a perfect solution. This is a sentiment echoed within the Action based Research and Mission Focus areas of action (Sections 2.2.1 and 2.2.3).

3.5.4.4 Cross UKRI stakeholder engagement workshop

A 3 hour workshop called 'Recommendations and implementation for a net zero DRI - cross-UKRI stakeholder workshop' was hosted by the project team on Thursday 11 May. It was run both online and in-person to encourage attendance in the most accessible way. The in-person location was at the Advanced Research Centre in Glasgow alongside the commissioned art exhibition (see Section 3.4.6).

The meeting was closed invite only, with attendees representing a mixture of funders, infrastructure managers and experts, project partners and the core project team. The workshop provided attendees with an opportunity to discuss the findings, recommendations and proposed roadmap produced by the Project.

Table 3.5.4.4.A: showing numbers of invited stakeholders and how they attended

Attendance type	**Attendance numbers on the day**
In person	9
Remote	28
Unavailable	33
Total invited	**70**

The purpose of the meeting was to:

[52] The UKRI infrastructure fund already has quite stringent sustainability criteria to demonstrate sustainability going forward.

- share our recommendations and gather expert input as to how these would work for a range of stakeholders
- allow attendees to consider what they need, in order to act upon the recommendations within their work, their teams and organisations
- discuss what needs to come next to start to address the recommendations and build on the Project's outcomes

Outcomes of the event:

- shared understanding and agreement gained amongst stakeholders that:
  - the project recommendations are fit for purpose and use
  - the proposed roadmap and next steps are suitable

The project team presented some introduction and background to the project, including scope and the process of collecting evidence and recommendations. Key messages and a proposed roadmap (draft version 6) were also discussed. The presentations were followed by discussion with attendees, focussing on the following questions:

- Are the project recommendations fit for purpose and use?
- Are the proposed roadmap and next steps suitable?

Discussion was lively and positive with many questions and comments raised by a variety of attendees. The main discussion points focussed around:

- Presenting positive side vs. avoiding “excuses”
- Target audience
- Presenting big-hitting actions clearly
- Theory of change
- How do we bring out the value of the forum and the leadership opportunity

This feedback has been considered by the project team and appropriate action taken (e.g. addressed in this report). Overall, the information presented to the stakeholders was well received. There was a good shared understanding and general agreement gained amongst stakeholders that the project recommendations are fit for purpose and use, and the proposed roadmap and next steps are suitable.

3.6 Lessons learnt

Lessons learnt were captured throughout the project lifecycle. The project team created a lessons learnt spreadsheet, which captured the following; project team lessons learnt throughout the life cycle of the project, sandpit projects feedback after the sandpit events were hosted, stakeholder engagement feedback from all workshops held and external feedback from project partners. The project used two format types and questions for capturing lessons learnt. The project team and external feedback was collated using RAL Space Project Management Office lessons learnt template. Sandpit events and Stakeholder engagement feedback was collated via both survey questions (using a google form) and eventbrite.

Project team will hold two project reviews, first meeting will be held in April 2023 and a final meeting to be held in June 2023. The project team will summarise and highlight what worked well, what did not work well and what could be improved. This information will be shared internally and externally.

Appendix 1. Glossary and Abbreviations

A1.1 Acronyms

AI - Artificial Intelligence The ability of a machine to display human-like abilities such as learning, reasoning, and creativity.

AMD - Advanced Micro Devices, Inc. CPU & GPU & FPGA manufacturer.

ARCHER2 - UK National Supercomputing Service provided by UKRI, EPCC, HPE Cray and the University of Edinburgh. https://www.archer2.ac.uk

BBSCR - Biotechnology and Biological Sciences Research Council, a UKRI council. https://www.ukri.org/councils/bbsrc/

BMC - Baseboard Management Controller, a microcontroller separate from the main CPU to monitor and manage a computer and its components, with the help of sensors. A baseboard management controller is embedded within the main circuit board of the device or computer and can help a single administrator to monitor a large number of servers or devices remotely.

CCC - The UK Climate Change Committee, The Climate Change Committee (CCC) is an independent, statutory body established under the Climate Change Act 2008. It's purpose is to advise the UK and devolved governments on emissions targets and to report to Parliament on progress made in reducing greenhouse gas emissions and preparing for and adapting to the impacts of climate change. https://www.theccc.org.uk

CIUK - Computing Insight UK, A UKRI conference aimed at people using computational and data facilities, including UKRI DRI. https://www.scd.stfc.ac.uk/Pages/CIUK2022.aspx

COSMA - COSmology MAchine, The DiRAC Memory Intensive System, COSMA is a UK national HPC facility funded by UKRI hosted by the Institute for Computational Cosmology at Durham University. https://www.dur.ac.uk/icc/cosma/

CPU - Central Processing Unit, Also known as "processor" (or "CPU processor"). The CPU is the computer component responsible for interpreting and executing most of the commands from the computer's other hardware and software. Today's CPUs typically have many (e.g. 64) "cores".

DCMS - Department For Culture Media and Sport, UK governmental department. https://www.gov.uk/government/organisations/department-for-culture-media-and-sport

DiRAC - Distributed Research Utilising Advanced Computing, Academic led provision of distributed High Performance Computing services to the STFC theory community. https://www.dur.ac.uk/icc/cosma/dirac/

DRI - Digital Research Infrastructure, the building blocks of DRI include: large scale compute facilities including high-throughput, high-performance, and cloud computing; data storage facilities, repositories, stewardship and security; software and shared code libraries; mechanisms for access, such as networks and user authentication systems; people, the users and the experts who develop and maintain these powerful resources.

EPCC - Edinburgh Parallel Computing Centre, Supercomputing centre based at the University of Edinburgh. https://www.epcc.ed.ac.uk/

EPSRC - Engineering and Physical Sciences Research Council, A UKRI council. https://www.ukri.org/councils/epsrc/

ExCALIBUR - Exascale Computing ALgorithms & Infrastructures Benefiting UK Research, UK research programme that delivers high-performance simulation software for the highest-priority fields in UK research.https://excalibur.ac.uk/

FAIR - Findable, Accessible, Interoperable, Reusable, Guiding Principles for scientific data management and stewardship. The FAIR Guiding Principles for scientific data management and stewardship | Scientific Data (nature.com)

FPGA - Field Programmable Gate Arrays, Semiconductor devices based around a matrix of configurable logic blocks (CLBs) connected via programmable interconnects. FPGAs can be reprogrammed to desired application or functionality after manufacturing, allowing potentially higher efficiencies than using general purpose CPU architectures.

GPU - Graphics/Graphical Processing Unit Hardware (typically at the end of an PCI-e bus) that supports throughput of very large (1000s) of lightweight threads to overcome context switching. A GPU has many (1000s) of lightweight GPU Cores that are very good at (e.g.) vector arithmetic but GPU gives poor performance on codes with lots of logical switching.

GridPP - Grid for Particle Physics, a collaboration of particle physicists and computer scientists from the UK and CERN that manage a distributed computing grid across the UK. The collaboration oversees a major computing facility called the Tier1 at the Rutherford Appleton Laboratory (RAL) along with the four Tier 2 organisations of ScotGrid, NorthGrid, SouthGrid and LondonGrid (formerly LT2). https://iopscience.iop.org/article/10.1088/0954-3899/32/1/N01

GSE - Green Software Engineering, Green Software Engineering relates to the use and development of software such that its carbon footprint is reduced. This encompasses programmers, users, facilities' operational staff and covers amendments to writing, running and scheduling code and to how data centres are run. A Green Software Engineer is an experienced practitioner of Green Software Engineering.

HEI - Higher Education Institution, A college, university, or other provider that offers and delivers higher education, leading to the award of an academic degree.

HPC - High Performance Computing, (also see "large scale computing" LSC)

HTC - High Throughput Computing, the use of many computing resources over long periods of time to accomplish a computational task

IAT - Implicit Association Test, Measures the speed of association between different paired concepts.

IPCC - Intergovernmental Panel on Climate Change, United Nations body for assessing the science related to climate change. IPCC — Intergovernmental Panel on Climate Change

IPMI - Intelligent Platform Management Interface, A standard for management and monitoring a computer using a BMC

IPMI tool - Intelligent Platform Management Interface tool, ipmitool is a utility for managing and configuring devices that support IPMI. https://github.com/ipmitool/ipmitool

LSC - Large Scale Computing, see "Large-scale computing: the case for greater UK coordination", https://www.gov.uk/government/publications/large-scale-computing-the-case-for-greater-uk-coordination

MMU - Manchester Metropolitan University, A UK HEI

NCAS - National Centre for Atmospheric Science, NERC funded research centre that undertakes research on atmospheric sciences. https://ncas.ac.uk/

NERC - Natural Environment Research Council, a UKRI council. https://www.ukri.org/councils/nerc/

OECD - Organisation for Economic Co-operation and Development, Inter-governmental organisation. https://www.oecd.org/about/

PDU - Power Distribution Unit, A device fitted with multiple outputs designed to distribute electric power, especially to racks of computers and networking equipment located within a data centre.

PUE - Power Usage Effectiveness, Ratio of total power to run a data centre (e.g. for compute and for cooling) divided by the power solely for the compute. A ratio of 1.0 is ideal.

QMUL - Queen Mary University of London, A UK HEI. https://www.qmul.ac.uk/

RAL - Rutherford Appleton Laboratory, STFC funded research laboratory. Rutherford Appleton Laboratory – UKRI

RAM - Random-Access Memory, Temporary storage component (short-term memory) of a computer.

RAPL - Running Average Power Limit, Intel interface to query a CPU for its power usage and performance measurements.

RSE - Research Software Engineer/Engineering, A software engineer or the use of software engineering principles in an academic research context

SCARF - Scientific Computing Application Resource for Facilities, A DRI resource hosted at the Rutherford Appleton Laboratory. https://www.scd.stfc.ac.uk/Pages/SCARF.aspx

Slurm - Simple Linux Utility for Resource Management, The Slurm Workload Manager, formerly known as Simple Linux Utility for Resource Management, or simply Slurm, is a free and open-source job scheduler for Linux and Unix-like kernels, used by many of the world's supercomputers and computer clusters. https://slurm.schedmd.com/overview.html

SNMP - Simple Network Management Protocol, protocol for monitoring the health and welfare of network equipment. https://net-snmp.sourceforge.io/

STFC - Science and Technology Facilities Council, A UKRI council. https://www.ukri.org/councils/stfc/

UCL - University College London, A UK HEI. https://www.ucl.ac.uk/

UKRI - United Kingdom Research and Innovation, A non-departmental public body sponsored by the Department for Science, Innovation and Technology. https://www.ukri.org/

UNESCO - United Nations Educational, Scientific and Cultural Organization, A United Nations agency promoting world peace and security. https://www.unesco.org/en

UPS - Uninterruptible Power Supply, A device to automatically switch to a backup power supply during a mains power cut so that attached equipment is not affected by the power cut. Typically used to protect computer equipment in data centres.

A1.2 Definitions

Accelerator Technology - Microprocessors that are capable of accelerating certain workloads. Examples of processing accelerators are GPUs (Graphical Processing Units) and FPGAs (Field Programmable Gate Arrays).

Active Carbon - In the context of this report, Active Carbon is synonymous with Operational Carbon, the carbon released from the ongoing operation of an infrastructure.

Benchmarking - Benchmarking is usually associated with assessing performance characteristics of computer hardware, for example, the floating point operation performance of a CPU, but there are circumstances when the technique is also applicable to software. Software benchmarks are, for example, run against compilers or database management systems (DBMS). Benchmarks provide a method of comparing the performance of various subsystems across different chip/system architectures.

Carbon Intensity - An emission rate of a given pollutant (equivalent carbon dioxide in this report, unless specified otherwise) relative to the amount of a specific activity, or an industrial production process; for example kilograms of carbon dioxide released per megajoule of energy produced.

Carbon Intensity of Power by Location - Also known as "location based carbon intensity", is the average emissions of the grid that you draw power from. For the UK this is the UK grid emissions factor published by the UK Government.
https://www.zevero.earth/post/location-vs-market-based-carbon-reporting

Carbon Intensity of Power by Sub-Grid - This approach to defining carbon intensity reflects the different mixes of generation in different parts of the national grid.

Carbon Intensity of Power by Market - Also known as the "market based carbon intensity", this reflects the emissions associated with the generating company or companies from which you purchase power. The customary market-based approach does not reflect the carbon impact of the power delivery by the national grid. UK Government guidance requires organisations to use the location-based approach but allows market-based metrics to be reported in parallel and suggest that "Organisations using a market-based figure may want to consider adding narrative information on whether their contractual arrangements cause additional renewable electricity generation".
https://www.zevero.earth/post/location-vs-market-based-carbon-reporting

Carbon Neutral - Carbon equivalent emissions are balanced by carbon sequestration.

ChatGPT - ChatGPT is an artificial-intelligence (AI) chatbot. It is a member of the generative pre-trained transformer (GPT) family of large language models, pre-trained on large datasets of unlabelled text and able to generate novel human-like text.

Cluster - An HPC cluster typically comprises several nodes connected by a network. A node comprises one or more sockets. Each socket holds a processor (or "CPU") and each processor may have 1 or more cores (also known as "CPU cores").

Code Optimisation - Code Optimisation is any method of code modification (directly or indirectly e.g. by a compiler) to improve code quality and efficiency.

Compute Node - When several computers are connected together in a network (forming an HPC cluster), each of the computers is referred to as a node in the network. An HPC **cluster** typically comprises several nodes connected by a network. A node comprises one or more **sockets**. Each socket holds a processor (or "**CPU**") and each processor may have 1 or more cores (also known as "**CPU cores**").

Core - An HPC cluster typically comprises several nodes connected by a network. A node comprises one or more sockets. Each socket holds a processor (or "CPU") and each processor may have 1 or more cores (also known as "CPU cores").

Delivery Areas - Three elements of UKRI programme delivery: "UKRI Direct Operations", "Lead procured", and "funded".

Delivery Pathways - Three elements of the roadmap: "Policy and Governance", "Delivery Partnership", "Competitive Funding". Each of the three delivery pathways is associated with one of the UKRI Delivery Areas.

Embodied/Embedded Carbon Embodied Carbon is the amount of carbon pollution emitted during the creation and disposal of a device (sometimes referred to as “Embedded Carbon”).

Enclosure - In the context of DRI, enclosures are the structures used to house compute nodes. An example of an enclosure is a vented cabinet known as a rack.

Green Software - Software that is written to emit the least amount of carbon possible.

Green Software Engineering - Green Software Engineering relates to the use and development of software such that its carbon footprint is reduced. This encompasses programmers, users, facilities’ operational staff and covers amendments to writing, running and scheduling code and to how data centres are run. A Green Software Engineer is an experienced practitioner of Green Software Engineering.

Idle Power Draw - The idle power draw of a computer is the energy consumed when it is on but not executing code.

Intel - Intel Corporation, CPU & GPU & FPGA manufacturer

Life Cycle Analysis - Compilation and evaluation of the inputs, outputs and potential environmental impacts of a product system throughout its life cycle. An LCA identifies the main contributors (materials, energy sources, etc.) to key environmental impacts throughout the product’s entire life cycle.

Machine Learning - The use and development of computer systems that are able to learn and adapt without following explicit instructions, by using algorithms and statistical models to analyse and draw inferences from patterns in data.

Marginal Emissions - Marginal emissions are the emissions that would come online if new load were added, they are nearly always greater than the average emissions.

Net Zero - Reaching carbon neutrality via carbon reductions through behaviour change and technological innovation, with carbon sequestration used to balance remaining emissions. In the words of the Science Based Targets initiative net-zero means: “Setting corporate net-zero targets aligned with meeting societal climate goals means: (a) reducing scope 1, 2 and 3 emissions to zero or a residual level consistent with reaching net-zero emissions at the global or sector level in eligible 1.5°C scenarios or sector pathways and (b) neutralising any residual emissions at the net-zero target date – and any GHG emissions released into the atmosphere thereafter”, where “neutralising” means: “Measures that companies take to remove carbon from the atmosphere and permanently store it to counterbalance the impact of emissions that remain unabated."

NVIDIA - Nvidia Corporation, GPU manufacturer.

Open Science - An approach to the scientific process that focuses on spreading knowledge as soon as it is available using digital and collaborative technology.

Operational Carbon - The carbon released from the ongoing operation of an infrastructure. Sources include lighting, power, heating, ventilation, cooling etc.

Port Code - Porting is the process of adapting code for the purpose of executing it in a different operating environment, such as a different CPU chip set or moving between CPU/GPU/FPGA configurations.

Sandpit Project - A method for generating cross-disciplinary, projects where academics and practitioners come together for a short time to create new projects around a given theme. These are pitched to a panel of experts for feedback. Successful projects receive funding.

Socket (processor) - An HPC cluster typically comprises several nodes connected by a network. A node comprises one or more sockets. Each socket holds a processor (or "CPU") and each processor may have 1 or more cores (also known as "CPU cores").

Sustainability - The quality of causing little or no damage to the environment and therefore able to continue for a long time.

Switch (network) - A network switch connects users, applications, and equipment across a network so that they can communicate with one another and share resources.

Tier (computing) - UKRI operates national (tier 1) and regional (tier 2) compute capability with universities providing local (tier 3) systems. International pre-exascale and exascale systems are referred to as tier 0.

Transformational Change - The type of organisational change that completely reshapes an organisation. It occurs in response to, or in anticipation of, significant changes in an organisation's environment or technology.

Appendix 2. Bibliography

Action Aid (2020): NOT ZERO: How 'net zero' targets disguise climate inaction. https://actionaid.org/publications/2020/not-zero-how-net-zero-targets-disguise-climate-inaction

Frederico Afonso, Martin Sohst, Carlos M.A. Diogo, Simão S. Rodrigues, Ana Ferreira, Inês Ribeiro, Ricardo Marques, Francisco F.C. Rego, Abdolrasoul Sohouli, Joana Portugal-Pereira, Hugo Policarpo, Bruno Soares, Bruna Ferreira, Edgar C. Fernandes, Fernando Lau, Afzal Suleman, Strategies towards a more sustainable aviation: A systematic review, Progress in Aerospace Sciences, Volume 137, 2023, 100878, ISSN 0376-0421, https://doi.org/10.1016/j.paerosci.2022.100878 .

Bane, Michael K., Brown, Oliver, Bhowmik, Deepayan, Quinn, Jamie, Ali, Teymoor, Quinn, Jamie, & Stansby, David. (2023). ENERGETIC: Final Report. Zenodo. https://doi.org/10.5281/zenodo.7692272

Barron, Alexander R. et al. (2021). “Carbon neutrality should not be the end goal: Lessons for institutional climate action from U.S. higher education”. In: One Earth86-4.9, pp. 1248–1258. issn: 2590-3322. DOI: https://doi.org/10.1016/j.oneear.2021.08.014

BBSRC (Nov. 2020). Review of Data-Intensive Bioscience. Tech. rep. Authoritative. url:

https://www.ukri.org/wp-content/uploads/2020/11/BBSRC-201120-ReviewOfDataIntensiveBioscience.pdf

BEIS (2022). Greenhouse gas reporting: conversion factors 2021. Tech. rep. Authoritative. url: https://www.gov.uk/government/publications/greenhouse-gas-reporting-conversion-factors-2021 (visited on 06/08/2022)

Bichsel, Jacqueline (2012). Research computing: The enabling role of information technology. Tech. rep. Opinion. url: https://library.educause.edu/resources/2012/11/research-computing-the-enabling-role-of-information-technology (visited on 05/16/2022)

Braun (2021) https://www.wiso.uni-hamburg.de/en/forschung/forschungszentren/css/working-paper-series/wp3-roedderbraun-2021-academicairtravel.pdf

Bol, Pirson and Dekimpe (2021) Moore's Law and ICT Innovation in the Anthropocene, Design, Automation & Test in Europe Conference & Exhibition (DATE), Grenoble, France, pp. 19-24, doi: 10.23919/DATE51398.2021.9474110.

Cannell, MGR (1999). “Growing trees to sequester carbon in the UK: answers to some common questions”. In: Forestry 72.3, pp. 237–247.

Louis-Philippe P.-V.P. Clément, Quentin E.S. Jacquemotte, Lorenz M. Hilty, Sources of variation in life cycle assessments of smartphones and tablet computers, Environmental Impact Assessment Review, Volume 84, 2020, 106416, ISSN 0195-9255, https://doi.org/10.1016/j.eiar.2020.106416.

Crumley-Effinger, Max and Blanca Torres-Olave (2021). “Kicking the Habit: Rethinking Academic Hypermobility in the Anthropocene”. English. In: JOURNAL OF INTERNATIONAL STUDENTS 11.1, SI, pp. 86–106. issn: 2162-3104

DCMS (2022). UK Digital Strategy. https://www.gov.uk/government/publications/uks-digital-strategy/uk-digital-strategy

DELL (2022). Carbon Footprint of a Dell Latitude 6530. Tech. rep. Commercial. url: https://i.dell.com/sites/doccontent/corporate/corp-comm/en/Documents/WP-E6530.pdf

DSIT (2020). National Data Strategy. https://www.gov.uk/government/publications/uk-national-data-strategy/national-data-strategy

Ernst and Young (2020). Product Carbon Footprinting – a study on methodologies and initiatives: Final Report. Tech. rep. Moderated. url: https://ec.europa.eu/environment/eussd/smgp/pdf/ProductCarbonFootprintstudy.pdf

ExCALIBUR Design & Development Working Groups (June 2021). Research Software Engineer Knowledge Integration Landscape Review. Tech. rep. Moderated. DOI: 10.5281/zenodo.4986062.

Friday, Adrian, Preist, Chris, Bird, Caroline, Lord, Carolynne, Widdicks, Kelly, Kayumbi, Gabin, Lambert, Simon, & Jackson, Adrian. (2023). Learning from the Big Picture: Applying Responsible Innovation to the Net Zero Research Infrastructure Transformation (ARINZRIT). Zenodo. https://doi.org/10.5281/zenodo.7966424

Greenhouse GHGs Protocol (partnership). “The Greenhouse Gas Protocol”. url: https://ghgprotocol.org (visited on 07/29/2022) (cit. on pp. 2, 29)

Gupta, U. *et al*., "Chasing Carbon: The Elusive Environmental Footprint of Computing," in *IEEE Micro*, vol. 42, no. 4, pp. 37-47, 1 July-Aug. 2022, doi: 10.1109/MM.2022.3163226. https://doi.org/10.1109/MM.2022.3163226

Hays, Jonathan, Walton, Nicola, Jackson, Adrian, Mudaraddi, A., Packer, A., & Owen, R.A. (2023). IRISCAST: Final Report. Zenodo. https://doi.org/10.5281/zenodo.7692451

Hopper, Andy et al. (Dec. 2020). Digital technology and the planet: Harnessing computing to achieve net zero. Tech. rep. Authoritative. url:https://royalsociety.org/-/media/policy/projects/digital-technology-and-the-planet/digital-technology-and-the-planet-report.pdf (visited on 05/16/2022)

Jattke, Marleen et al. (Sept. 2020). “Environmental Implications of Service Life Extension of Mobile Devices”. In: Electronics Gores Green 2020+ - The Story of Daisy, Alexa and Greta: Proceedings. Berlin, Germany: Fraunhofer Institute for Reliability and Microintegration, pp. 163–170. isbn: 978-3-8396-1659-8. url: https://online.electronicsgoesgreen.org/wp-content/uploads/2020/10/ProceedingsEGG2020v2.pdf (visited on 08/17/2022)

Juckes, Martin, Pascoe, Charlotte, Woodward, Lucy, Vanderbauwhede, Wim, & Weiland, Michele. (2022). Interim Report: Complexity, Challenges and Opportunities for Carbon Neutral Digital Research. Zenodo. https://doi.org/10.5281/zenodo.7016952

Kraus, Sascha et al (2021) Digital transformation in healthcare: Analyzing the current state-of-research, Journal of Business Research, 123: 557-567, DOI: https://doi.org/10.1016/j.jbusres.2020.10.030

Anne Pasek, Dani Admiss, Tega Brain, Angela YT Chan, Maya Chowdhury, Laura Clarke, Julie Freeman, Louise Hargreaves, Ellie Harrison, Lara Houston, Marija Bozinovska Jones, Adi Kunstman, Jane Lawson, Gretta Louw, Martino Morandi, Padmini Ray Murray, Inge Paneels, Irini Papadimitriou, Anne Pasek, Radek Przedpelski, Caroline Sinders, Rebecca Sinker, Brian Sutherland, Marloes de Valk, Chris Wright, Digital Decarbonisation Consensus & Conjectures, Sept. 2022. https://sunlightdoesntneedapipeline.com/2022/08/17/digital-decarbonisation-consensus-conjectures/

Lambert, Simon, Stephens, Ag, & Kayumbi, Gabin. (2023). Net Zero DRI Case Studies Report. Zenodo. https://doi.org/10.5281/zenodo.7828052

Manika, Danae, McCarroll, Niall, Buck, Justin, Lopez, Alvaro Lorenzo, Geatches, Dawn, Ahrens, Florian, Krumdieck, Susan, Keshtkar, Hossein, Fayyad, Nadine, & Hassanloo, Hamidreza. (2023). Giving Voice to, and Empowering Stakeholders of UKRI DRI: A Net Zero Workshop Series (Go Zero). Zenodo. https://doi.org/10.5281/zenodo.7692290

McGuire, Laura. (2023). User Behaviour Survey Final Report: The motivators/enablers, and barriers to sustainable Digital Research Infrastructure use. Zenodo. https://doi.org/10.5281/zenodo.7827919

Pirson, T., T. P. Delhaye, A. G. Pip, G. Le Brun, J. -P. Raskin and D. Bol, "The Environmental Footprint of IC Production: Review, Analysis, and Lessons From Historical Trends," in IEEE Transactions on Semiconductor Manufacturing, vol. 36, no. 1, pp. 56-67, Feb. 2023, doi: 10.1109/TSM.2022.3228311.

Radinger-Peer, Verena and Gernot Stoeglehner (Apr. 2013). “Universities as change agents for sustainability – framing the role of knowledge transfer and generation in regional development processes”. In: Journal of Cleaner Production 44, pp. 85–95. DOI:10.1016/j.jclepro.2012.12.003

Reyes-Garc´ıa, Victoria et al. (2022). “Decarbonizing the academic sector: Lessons from an international research project”. In: Journal of Cleaner Production 368, p. 133174.

Tseng, S.H.Y., Higham, J., Lee, C. (2022). Academic Air Travel Cultures: A Framework for Reducing Academic Flying. In: Bjørkdahl, K., Franco Duharte, A.S. (eds) Academic Flying and the Means of Communication. Palgrave Macmillan, Singapore. https://doi.org/10.1007/978-981-16-4911-0_13

Schien, Daniel, Zilberman, Noa, Greenwood, David, Shabajaee, Paul, & Dewhurst, Alastair. (2023). Preliminary Release: CarbonQuanDRI: Decarbonising UKRI Electricity Use – A Case Study of GridPP Tier-1. Zenodo. https://doi.org/10.5281/zenodo.7692133

Stephens, Ag, Kayumbi, Gabin, & Lambert, Simon. (2023). Report on the DRI Mapping Survey 2023 (UKRI Net Zero DRI Scoping Project) (1.0). Zenodo. https://doi.org/10.5281/zenodo.7805988

Turner, Andrew, & Basden, Alastair. (2022). HPC-JEEP: Energy Usage on ARCHER2 and the DiRAC COSMA HPC services (v1.1). Zenodo. https://doi.org/10.5281/zenodo.7137390

Smith, Lorna, Simpson, Alan, Moran, Laura, & Turner, Andy. (2023). ARCHER2 Net Zero Case Study. Zenodo. https://doi.org/10.5281/zenodo.7788498

UNESCO (2021): Recommendation on Open Science https://unesdoc.unesco.org/ark:/48223/pf0000379949.locale=en

Vanderbauwhede, Wim. (2023). Energy consumption of facility and non-facility computing equipment used by UKRI researchers. Zenodo. https://doi.org/10.5281/zenodo.8072018

Whitmarsh Lorraine, Capstick Stuart, Moore Isabelle, Köhler Jana, Le Quéré Corinne, Use of aviation by climate change researchers: Structural influences, personal attitudes, and information provision, Global Environmental Change, Volume 65, 2020, 102184, ISSN 0959-3780, https://doi.org/10.1016/j.gloenvcha.2020.102184

Whitmarsh Lorraine, Capstick Stuart, Moore Isabelle, Köhler Jana, Le Quéré Corinne, Corrigendum to "Use of aviation by climate change researchers: Structural influences, personal attitudes, and information provision" [65 (2020) 102184], Global Environmental Change, Volume 67, 2021, 102202, ISSN 0959-3780, https://doi.org/10.1016/j.gloenvcha.2020.102202

Wilkinson, Mark et al. (Oct. 2020). UK Research & Innovation: Science case for UK Supercomputing. Tech. rep. Opinion. url: https://excalibur.ac.uk/wp-content/uploads/2021/07/UKRI-Supercomputing-Science-Case.pdf

Appendix 3. Methodology

Literature Survey

An exploratory approach has been taken to develop the literature survey. The topic areas defined in the project work plan, as modified on the basis of feedback by the assessment panel and Steering Committee, have been taken as the basis for literature searches. Key results and themes are drawn out from the search results.

Objective analysis of keywords using the van Eck and Waltman (VOSviewer) application and search results from Web of Science is also used to obtain an overview of topics addressed within subject areas and to analyse clusters of activity.

Drafts and Review

- Zero order draft : limited review
- First order draft : limited review
- Draft Final: open review

Grey Literature

Where possible, this report uses peer reviewed literature. In many areas, however, important information is published through grey literature. This includes authoritative reports from, for instance, the Climate Change Committee or the International Energy Authority, reports by research consortia or consultancies, and even news articles.

We follow the IPCC (e.g. Editorial 2011) in accepting grey literature but taking care to use it with caution and also to be clear about the nature of the information being used. In particular, press releases can give a reliable indication of the ambition and aspiration of organisations, and that is an important indicator of the direction of development in a fast moving industry, but should not be used as the basis of detailed technical analysis. Some organisations, such as the IEA, have rigorous systems for gathering information and their reports can be considered to be as reliable as peer- reviewed literature.

In this report, grey literature will be assigned one of the following categories:

• Authoritative: produced by an authoritative organisation with transparent procedures or containing sufficient evidence of an internal review process (e.g. through community engagement);

• Authoritative Opinion: produced by an individual or individuals placed in a position of trust in an organisation with a reputation for academic independence.

• Moderated : a paper which is produced by a project or organisational unit which has some form of publication process.

• Commercial: produced by a commercial organisation.

• Opinion: treated as representing the opinion of the authors.

The category assigned will be displayed in the bibliography, as in the following example, which is an editorial from the Nature Climate Change journal:

> *Editorial (Aug. 2011). "Evolving the IPCC". In: Nature Climate Change. Authoritative Opinion. url:https:/ /www.nature.com/articles/nclimate1189.pdf?origin=ppub.*

In the text, grey literature is distinguished from peer reviewed literature by a star.

Bibliography Review

One or more technical reviewers will be engaged to perform a systematic review of all sources cited in the bibliography, checking the categorisation of grey literature, the consistency of attributed information with the source and validity of links provided.

It has not been possible to complete this secondary check for the interim report, but a more systematic review will be put in place for the final review.

Appendix 3. Full Recommendations Listing

Four tables present all 180 recommendations arising across the different evidence gathering activities conducted as part of the Project. You can also find all recommendations in an interactive spreadsheet document found here: https://doi.org/10.5281/zenodo.8199893

Each recommendation table includes: a citable reference, the source report (e.g. which evidence gathering activity it relates to), a title, statement (brief description), toolkit theme, and which roadmap action it relates to.

A3.1 Literature Review

See Section 3.2 for discussion.

Citable Ref #	Source Report	Recommendation title	Statement	Theme(s)	Roadmap Action
1	Interim report	1. DRI investment decisions to take into account the views of the community.	The UKRI needs to ensure that investment decisions are backed by a deep understanding of the views of the research community.	**Working with peers and suppliers**	**4**
2	Interim report	2. Leverage soft capacity to accelerate change through consensus.	UKRI should use its capacity in social sciences, arts and humanities, and in economics, to understand the range of societal views, the avenues of consensus which open-up potential for accelerating transition and the emerging (or exploding) discords which can block or reverse change.	**Mission Focus**	**4**
3	Interim report	3. Establish a focal point to enable	Create a focal point which can bring together the strands of activities across	**Mission Focus**	**20**

		coherent Net Zero delivery.	the research sector to enable a coherent approach to delivering the Net Zero UKRI DRI.		
4	Interim report	4. Invest in DRI throughput capacity challenges.	The UKRI DRI needs to invest in the development of capability to deliver effective scientific throughput in a rapidly evolving digital landscape.	**Mission Focus**	**6**
5	Interim report	5. Blending the discipline categorisations of infrastructure and research	Steps must be taken to minimise barriers to the adoption of potential efficiencies arising from new technologies in hardware and software, e.g. ensuring adequate cross over between experts of science and experts of technology.	**Working with peers and suppliers**	**4**
6	Interim report	6. Match new technologies with information that is reliable and relevant, and appropriate training and tools.	Ensure the introduction of new technologies is matched by appropriate resources for training and expert user support, particularly training scientists in use of new technology.	**Green software engineering** Knowledge Hub	**19**
7	Interim report	7. Avoid the rebound effect at the procurement stage	Ensure the procurement framework enables the conversion of efficiency gains into carbon savings rather than simply resulting in higher usage.	**Working with peers and suppliers** Mission Focus	**10**
8	Interim report	8. 100% off-grid renewable electricity supply.	Best practice for an individual institution is to adopt 100% off-grid renewable electricity supply, but there are limitations of scale and location.	**Mission Focus**	**1**
9	Interim report	9. Power Purchase Agreements with renewable investment.	Adoption of multi-year power purchase agreements with renewable investment clauses.	**Working with Peers and Suppliers**	**10**
10	Interim report	10. On-site grid-scale battery storage	Construction of grid-scale battery storage matched to the institutional power demands.	**Mission Focus**	**3**
11	Interim report	11. On-site renewable power and batteries.	Building on-site renewable generation and grid-connected power storage to mitigate load on the national grid.	**Mission Focus**	**3**
12	Interim report	12. Procurement that balances investment in energy efficiency against investment in infrastructure.	Institutions making purchases on behalf of UKRI must be empowered to balance investments in efficiency against investments in energy intensive infrastructure.	**Working with peers and suppliers**	**10**

13	Interim report	13. Procurement contracts with sustainability clauses.	Add sustainability clauses in procurement contracts.	**Working with peers and suppliers**	**10**
14	Interim report	14. Procurement: supply line relationships.	Build relationships along the supply line to work on mutually beneficial solutions.	**Working with peers and suppliers**	**10**
15	Interim report	15. Extend life of and re-use equipment.	Look at the whole life cycle of equipment and opportunities to extend the life and reuse potential of equipment.	**Working with Peers and Suppliers** Action-based research	**11**
16	Interim report	16. Fossil fuel free on-site energy generation.	Eliminating on-site use of fossil fuels will require a clear timeline. This should be developed by 2025.	**Mission Focus** Action-based research	**1**
17	Interim report	17. Digital collaboration tools.	The UKRI DRI should facilitate and promote digital collaboration tools and awareness to reduce carbon intensity and enhance access to the research programmes.	**Green software engineering**	**3**
18	Interim report	18. Continuous improvement of measures to avoid emissions.	UKRI needs to be exhaustive in exploring what can be avoided before, during, and after taking steps to deal with unavoided emissions.	**Mission focus**	**6**
19	Interim report	19. Sustainability of carbon sequestration.	Given uncertainty in the scalability of biochar and other carbon removal innovations, the UKRI need to couple investments with research into their sustainability.	**Mission focus**	**6**
20	Interim report	20. Carbon offsetting action should be linked to long-term guarantees.	UKRI should ensure that any offsetting investments are linked to guarantees of institutional continuity, e.g. through a trust.	**Mission Focus**	**22**
21	Interim report	21. Impact of research on environmental sustainability.	There is currently no requirement for researchers to consider that their research could have a negative impact on sustainability. The existential crises that face us in climate and biodiversity need to be reflected in every grant application as a key element of ethical and societal responsibility.	**Mission Focus** Recognise shared responsibility	**16**
22	Interim report	22. Continuity of expertise and innovative capacity	The UKRI DRI must ensure continuity of activities which can assess best practice and deliver guidance to all	**Mission Focus** Recognise	**6**

		to realise the net zero roadmap.	those involved in funding, procuring, operating and using digital research infrastructure.	shared responsibility	

A3.2 Sandpit Projects

See Section 3.3 for discussion.

Citable Ref #	Source Report	Recommendation title	Statement	Theme(s)	Roadmap Action
23	ARINZIT	1. Sector-wide Net Zero DRI policies	Make informed Net Zero DRI policies from transparent evidence of its environmental impacts, involving sector-wide policy to ensure all research institutions share DRI environmental data with UKRI which follow a consistent carbon calculation method and consider DRI's full lifecycle impact.	**Recognise shared responsibility**	**12**
24	ARINZIT	2. Establish and promote sector-wide FAIR data and code protocols	Establish and promote sector-wide FAIR data and code protocols to maximise visibility and re-use of existing data and code, and minimise duplicate or unnecessary processing and storage	**Green software engineering** Working with peers and suppliers	**15**
25	ARINZIT	3. Formalise Net Zero research incentives to reshape academic practice	Formalise Net Zero research incentives to reshape academic practice, promoting research which truly embeds a sustainable approach to DRI (e.g., by assessing DRI's full lifecycle in peer-review processes and funding applications and calls, during project execution and review, offering best paper awards for delivering results with minimal environmental impact).	**Mission Focus** Action-based research	**1**
26	ARINZIT	4. Flexible sharing of sector-wide DRI	Support flexible sharing of sector wide DRI for researchers to utilise available computational resources when required, avoiding underutilised DRI and the expansion of new and unnecessary DRI elsewhere.	**Knowledge Hub** Recognise shared responsibility	**12**
27	ARINZIT	5. Resourcing for	Publicise and resource mandatory	**Mission**	**3**

		mandatory Net Zero policies	Net Zero and climate emergency policies so that low carbon options are the default choice by ensuring appropriate funding and regulations are compatible with addressing the ambitions and cost of Net Zero	**Focus**	
28	ARINZIT	6. Clear processes in funding decisions for DRI use	Establish clear decision processes in funding applications for whether DRI use is required, based on consistent processes for defining type and use of DRI and methods for determining its environmental impact	**Mission Focus** Working with peers and suppliers	**3**
29	ARINZIT	7. Training and support for sustainable DRI use, and inclusion of RSEs in funding	Offer training to researchers on sustainable DRI use and better software engineering practices to ensure best choice and use of appropriate DRI hardware e.g., via specialist ‘Research Software Engineers’ (RSEs), costed into or shared across projects, and supported beyond their lifetimes to avoid inefficient use of DRI and DRI duplication.	**Green software engineering**	**13**
30	ARINZIT	8. Integration of DRI as part of wider estate infrastructure	Recognise DRI role as part of the wider infrastructure and embed in institutional policy and practice, ensuring valuable outputs (e.g., heat) are integrated into institutions’ estates and beyond (i.e., local, regional, national) to maximise value and avoid waste	**Mission Focus** Recognise shared responsibility	**11**
31	ARINZIT	9. Establish and follow best practice in sustainable use of DRI by researchers	Ensure researchers follow best practice in the sustainable use of DRI whilst recognising the need to advance knowledge, e.g., by reusing DRI, data and code where possible, ensuring new code is optimised, embedding FAIR data practices, and considering whether the proposed research or new DRI is really required.	**Recognise shared responsibility** Knowledge Hub	**16**
32	CQUANDRI	1. Investigate the Carbon Benefits from increasing Flexibility in Demand	Monitor the maturity of Locational Marginal Emission Factors (LMEF). When they are robust, evaluate the whole life cycle net benefit from increasing flexibility in demand. This	**Green software engineering** Action-based research	**7**

			can take many forms (Battery, overcapacity, throttling capacity). Some of these options imply procurement of additional hardware with it's associated embodied impact and cost. Any modelling should take this into account.		
33	CQUANDRI	2. Develop carbon-aware resource allocation policies	Develop a resource allocation policy that allows users to use resources in a more carbon-efficient way (e.g. longer runtime for less carbon emissions) and allow for "spot" workloads that can be deallocated when carbon intensity is high or when PUE increases	**Knowledge Hub** Action-based research	**7**
34	CQUANDRI	3. Reduce the idle power consumption of equipment	The idle power consumption of equipment constitutes a significant element in overall energy consumption (and carbon emissions), that is not affected by better software design. Reducing this idle power by various techniques can significantly reduce emissions.	**Green software engineering** Action-based research	**6**
35	ENERGETIC	1. Invest in accelerators according to energy savings	UKRI should invest in accelerator hardware and associated training & development in order to reduce energy to solution of HPC workloads. The accelerators should include GPU and FPGA technologies and further benchmarking efforts funded to improve understanding of differences in energy efficiency across different classes of algorithms.	**Action-based research**	**10**
36	ENERGETIC	2. Provision of required energy consumption data	UKRI compute-focused DRI administrators endeavour to provide a common interface or framework to abstract underlying monitoring mechanisms and expose energy consumption data to users.	**Recognise shared responsibility**	**21**
37	ENERGETIC	3. Commitment to reporting energy consumption data to users	UKRI commits to requiring user-level energy monitoring in a specified format to be available on all UKRI-funded DRI, and incorporates this into the funding and procurement processes.	**Green software engineering** Recognise shared responsibility	**16**

38	ENERGETIC	4. Raising awareness	UKRI invests in raising awareness of the need to optimise for least energy consumption	**Recognise shared responsibility**	**4**
39	ENERGETIC	5. Support for greener compute via upskilling and Green Software Engineers	UKRI invests in upskilling researchers regarding (env.) sustainable software principles & best practises, via Green Software Engineers (GSEs) and encouraging researchers and system admin to undertake Green Software training (e.g. https://trainingportal.linuxfoundation.org/learn/course/green-software-for-practitioners-lfc131)	**Green software engineering**	**6**
40	GO ZERO	1. Data Storage	Employ data 'triage' techniques e.g. data compression, filtering multiple instance of the same data; choose physical hardware (e.g. disks and tapes) for storage rather than virtual due to lower energy use	**Action-based research**	**14**
41	GO ZERO	2. Management of online portals	Improving searching methods to access data faster by using specific software, e.g. mongoDB, or by using other methods such as labeled data; making available search history of past queries and their efficiencies; return to accessing physical libraries.	**Green software engineering**	**14**
42	GO ZERO	3. Participatory resource monitoring	Design and implementation of energy and carbon monitoring system to feedback how using DRI affects emissions and energy footprints. The goal is to encourage "sustainable behaviour" by design.	**Green software engineering**	**23**
43	GO ZERO	4. Procurement Process	Current procurement process leaves little room to account for energy efficiency in hardware and software.	**Working with Peers and Suppliers**	**10**
44	GO ZERO	5. Changing Funding Body Policies and Monitoring	Pressure could be applied from funding bodies to request an assessment of carbon-impact for each proposal and to perform life cycle analysis. To do this there’s a requirement for the developments of tools and information to help people include these data in proposals. Funding councils could ask proposals to consider efficiency and	**Working with peers and suppliers**	**3**

			plan for reducing footprint.		
45	GO ZERO	6. Develop RSE skill base	Training for RSEs to develop efficient code	**Green software engineering**	**6**
46	GO ZERO	7. Carbon efficient HPC procurement and maintenance	Align the procurement, replacement and maintenance of HPC with Net Zero policy	**Mission Focus** Working with peers and suppliers	**11**
47	GO ZERO	8. Include carbon costs within the assessment process and planning of research grants	Addition of the carbon assessment of the use of HPC where applicable in the research bidding process	**Recognise shared responsibility**	**3**
48	GO ZERO	9. Develop best practice for measurement of carbon use on UKRI HPC	Measurements and estimates of carbon used on HPC need to be based on consistent best practices	**Action-based research**	**9**
49	GO ZERO	10. Encouraging behaviour change through raising awareness and offering incentives	Currently, if they even consider the carbon cost of their computations, individual users 'decide' if and how they change their own practices to reduce their energy consumption.	**Recognise shared responsibility** Knowledge Hub	**4**
50	GO ZERO	11. Data capture, storage and access	Reduce the need or drive to reproduce data that already exists by encouraging communities to adopt more openness and data-sharing practices	**Recognise shared responsibility**	**14**
51	GO ZERO	12. Cap and Share/ Green Scheduler	Setting a cap and share system that caps the emissions at a climate safe amount and that manages the usage of the DRI accordingly. To meet the needs of fulfilling computation, the behaviour will adapt to the fluctuations of renewable energies to 1) make the most use of them and 2) to reduce the need of fossil fuels. It has yet to be researched how this system could look like and how it can be implemented in DRI activities.	**Green software engineering**	**16**
52	GO ZERO	13. Auditing	Doing environmental and energy audits to map out the current state of the DRI system in focus.	**Green software engineering**	**5**
53	GO ZERO	14. Net Zero Action	Actors in the DRI system are to	**Mission**	**2**

		Plans	design net zero action plans that focus on community-led initiatives that are achievable and that realise the fossil-fuel downshift.	**Focus** Recognise shared responsibility	
54	GO ZERO	15. Community building	Help to share knowledge and best practices on getting to NetZero	**Knowledge Hub** Working with peers and suppliers	**4**
55	HPC-JEEP	1. Embodied emissions data	Information on embodied emissions of the technology offered should form part of any procurement of an HPC system. This may also require updating framework purchasing agreements	**Working with peers and suppliers**	**10**
56	HPC-JEEP	2. Energy use data	Energy collection functionality integrated with the scheduling software and available to users at an individual compute node level (and ideally, at the level of other components: e.g. switches, file systems) should be a mandatory requirement for all HPC services	**Green software engineering** Recognise shared responsibility	**16**
57	HPC-JEEP	3. Energy-based charging	Options for energy-based charging should be explored for current and future HPC services	**Green software engineering** Action-based research	**6**
58	HPC-JEEP	4. Procurement benchmarks	Procurement benchmarks for HPC services should be evaluated on energy efficiency metrics as well as traditional performance metrics. This may also require updating framework purchasing agreements	**Working with peers and suppliers**	**10**
59	HPC-JEEP	5. Report energy/emissions use to users	Services should report energy and emissions (embodied+energy-based) back to users and project leaders and set the numbers in context (e.g. compared to typical home energy use or emissions from a typical journey)	**Green software engineering** Knowledge Hub	**16**
60	HPC-JEEP	6. Explore energy/emissions aware scheduling	If per-job energy and emissions data is available there is the possibility of linking the job scheduler to forecasts of weather (which influence cooling	**Green software engineering** Action-based	**24**

			overheads) and/or emissions (as forecast on the grid) to schedule work to, for example, to minimise cooling overheads or to minimise emissions from energy use.	research	
61	IRISCAST	(1) Procurement Scoring includes energy usage and carbon cost	We recommend that for future DRI procurement rounds that embodied carbon costs and equipment energy usage start to be included as part of procurement evaluation scoring.	**Working with peers and suppliers**	**10**
62	IRISCAST	(2) Procure equipment to measure node energy usage, Require hardware embodied carbon costs in procurement documents	As a further step it is recommended that computer hardware be specified to include energy measurement capability such as IPMI (or per port PDUs) and also to require the supplier to provide details or best estimates of embodied carbon costs.	**Working with peers and suppliers**	**11**
63	IRISCAST	(3) Direct energy use measurements of power to compute clusters and to cooling	DRI sites will need to ensure that they have energy measurements in place for cooling infrastructure as well as for the energy supply to the computing infrastructure.	**Green software engineering**	**11**
64	IRISCAST	(4) Equipment Inventory to include embodied carbon and idle power draw	Further, DRI sites will need to keep an embodied carbon inventory of equipment that is in service, and it is recommended that a measure of the idle power draw is added to these inventories.	**Working with peers and suppliers** Knowledge Hub Action-based research	**5**
65	IRISCAST	(5a) "Monthly" Reporting of carbon cost and usage/idle statistics	With such an embodied carbon inventory, including idle power draw, and with such energy measurements, monthly reporting of carbon costs at sites should be instituted, along with reports of usage/idle time of compute clusters allowing the carbon cost of the idle time to be deduced. Reporting should be rolled into standard grant reporting as much as possible.	**Recognise shared responsibility** Action-based research	**16**
66	IRISCAST	(5b) "Monthly" Reporting merged into standard grant reporting	**With such an embodied carbon inventory, including idle power draw, and with such energy measurements, monthly reporting**	**Recognise shared responsibility** Action-based	**16**

		procedures.	**of carbon costs at sites should be instituted, along with reports of usage/idle time of compute clusters allowing the carbon cost of the idle time to be deduced. Reporting should be rolled into standard grant reporting as much as possible.**	research	
67	IRISCAST	(6a) Collect per payload energy usage	It is recommended that sites are encouraged to configure Slurm and other tools to collect per job (or in cloud environments per virtual machine) energy usage. This information will need to be combined with the inventory information on embodied carbon and with information on the energy mix to obtain the carbon cost of a particular job. A mechanism to feed this back to the end user will need to be developed to drive improvements in user workflow and some codebases.	**Green software engineering** Recognise shared responsibility Action-based research	**16**
68	IRISCAST	(6b) Feedback of payload energy usage	It is recommended that sites are encouraged to configure Slurm and other tools to collect per job (or in cloud environments per virtual machine) energy usage. This information will need to be combined with the inventory information on embodied carbon and with information on the energy mix to obtain the carbon cost of a particular job. A mechanism to feed this back to the end user will need to be developed to drive improvements in user workflow and some codebases.	**Green software engineering** Recognise shared responsibility	**16**
69	IRISCAST	(7) Identification of user communities and codebases	A way to identify user communities and the authors of community codebases so that useful feedback can be given to them to drive more efficient code and workflows will also need to be developed.	**Recognise shared responsibility** Working with peers and suppliers	**4**
70	VALUE	1. Set up an 'energy use register'	It is recommended that an energy use register is developed and kept updated to track carbon emissions across all the infrastructure that is	**Mission Focus** Working with peers and	**23**

			considered 'in scope' for this project. Different levels of 'scrutiny' should be applied to tracking the energy use of different facilities in the register. For the majority of facilities simple estimates / measurements of energy use could be paired with grid information to give quick and reasonably accurate snapshots of DRI emission profiles. More detailed tracking of emissions can then be focused on the highest emitters and those that need larger reductions as carbon budgets tighten.	suppliers	
71	VALUE	2. Set emissions reduction pathways for all councils / facilities	See Section 6.2.4 of report	**Mission Focus** Recognise shared responsibility	**2**
72	VALUE	3. Set up a system to ensure emissions reductions occur.	The recommended 'efficient' system would be using a task scheduler as outlined in Section 6.2 of the report.	**Green software engineering** Mission Focus	**24**
73	VALUE	4. Use estimations of energy use + grid carbon intensity to track DRI use emissions	Rough estimates of predicted task energy use can be paired with grid carbon intensity forecasts and compared retrospectively against actual power consumption of the facility. Detailed emissions tracking work can then be developed where it would be most beneficial (e.g., facilities with the highest 'scrutiny' required).	**Green software engineering** Action-based research	**7**
74	VALUE	5. Require facilities / councils to implement performance measurement systems	Sections 5.2.1 & 6.2.2 of report	**Mission Focus**	**16**

A3.3 Consortium Projects

See Section 3.4 for discussion.

Citable	Source	Recommendation	Statement	Theme(s)	Roadmap

Ref #	Report	title			Action
75	Case Study - ARCHER2	1. Information on carbon impact of manufacturing and delivery	For future procurements, vendors should be required to provide high quality information on the carbon impact of the manufacturing and delivery of the hardware.	**Working with peers and suppliers**	**10**
76	Case Study - ARCHER2	2.Environmentally aware building of data centres	Data centre construction should use environmentally aware builders and should begin to make estimates of the cost of carbon construction.	**Working with peers and suppliers**	**10**
77	Case Study - ARCHER2	3. Reduction in idling power	Vendors should be required to ensure that idling power is reduced significantly.	**Green software engineering** Working with peers and suppliers	**16**
78	Case Study - ARCHER2	4. Green energy supply	Data centres should be required to use green energy.	**Mission Focus** Working with peers and suppliers	**10**
79	Case Study - ARCHER2	5. Good-citizen behaviour by data centres	Data centres are custodians of significant power and should therefore actively engage in good-citizen behaviour. For example, in times of potential nationwide power shortages, Data Centres could reduce power utilisation or donate generated power to the grid.	**Recognise shared responsibility** Working with peers and suppliers	**12**
80	Case Study - ARCHER2	6. Instrumentation is vital	Instrumentation is vital. This includes, for example, accurate per job data, all hardware components, data centre power supply and all cooling elements.	**Green software engineering** Action-based research	**7**
81	Case Study - ARCHER2	7.Changes to CPU frequency	Some, but not all, codes may benefit from a reduction in CPU frequency. Correct decisions require good information (e.g., instrumentation).	**Green software engineering** Action-based research	**7**
82	Case Study - ARCHER2	8. High utilisation required	High utilisation is vital for efficient use of resources as, for example, idle power utilisation is high and embodied carbon is a significant fraction of the lifetime emissions.	**Green software engineering** Recognise shared responsibility	**16**
83	Case Study - ARCHER2	9. Evidence of societal and/or economic value of a service	Service Providers should evidence the societal and/or economic value of their service, to be evidenced against the carbon impact of their service.	**Recognise shared responsibility**	**10**

84	Case Study - ARCHER2	10. Test against best practice recommendations e.g. from PRACE.	Based on the PRACE infrastructure workshops reports, a set of best practice recommendations have been made for all data centres. Data centres should test their services against these recommendations, which include: planning for future power requirements; efficient water cooling; re-use of waste heat; broad instrumentation; energy efficiency; and energy grid-friendly scheduling.	**Mission Focus** Knowledge Hub Action-based research	**14**
85	Case Study - ARCHER2	11. Report energy use at the job level. Charging mechanisms involving energy use.	To guide users towards making better energy/carbon decisions, Service Providers should consider reporting to users the energy use of their jobs and implementing charging mechanisms that include a component of energy use.	**Green software engineering** Knowledge Hub	**16**
86	Case Study - ARCHER2	12. Training for users on energy efficiency and carbon emissions	Better training should be provided for users on energy efficiency and carbon emissions. This could potentially build on the work of the Green Software Practitioner open-source training: https://learn.greensoftware.foundation/	**Knowledge Hub** Recognise shared responsibility	**13**
87	Case Study - ARCHER2	13. Positive environmental impact from users	Researchers should be encouraged to report their positive environmental impact.	**Recognise shared responsibility**	**4**
88	Case study - JASMIN, JADE2, Scafell Pike	1. Appropriate power monitoring	Best practice should be shared among DRI managers, and those responsible for design and procurement, on suitable levels of monitoring power usage and how this may be effectively deployed in practice.	**Recognise shared responsibility**	**4**
89	Case study - JASMIN, JADE2, Scafell Pike	2. Disconnect from sources of energy	A recommendation to devolve management of energy contracts to individual DRI managers would be impractical and undesirable. However the negotiation of contracts, at whatever level it takes place, should take account of the patterns of usage at DRI level with a view to opportunities for optimisation.	**Working with peers and suppliers**	**11**
90	Case study - JASMIN,	3. User behaviour and energy use	It is not necessarily the job of DRI managers to interact with users or to	**Action-based research**	**7**

	JADE2, Scafell Pike		have in-depth awareness of user practices. However, monitoring of energy use related in some way to user activity would complement the general usage monitoring. The possible scope and implementation of such monitoring should be examined and disseminated.	Working with peers and suppliers	
91	Case study - JASMIN, JADE2, Scafell Pike	4. Local energy and reuse	Any practice and experience should be shared on local sources of energy and heat reuse, leading to a knowledge base on costs and opportunities.	**Knowledge Hub** Working with peers and suppliers	**4**
92	Case study - JASMIN, JADE2, Scafell Pike	5. Seeking technological opportunities	DRI managers should be encouraged to identify technological opportunities with implications for Net Zero attainment; job roles should include this explicitly.	**Mission focus** Green Software Engineering	**17**
93	Case study - JASMIN, JADE2, Scafell Pike	6. Procurement priorities	Principles and priorities of procurement should be re-examined to assess their alignment with Net Zero goals, within the necessary constraints.	**Mission focus**	**22**
94	Case study - JASMIN, JADE2, Scafell Pike	7. Data centre design knowledge base	A central shared knowledge base on data centre design and best practice should be developed, with input from DRI managers, and sustained into the future with mechanisms to influence actual practice.	**Mission Focus** Knowledge Hub	**20**
95	Exploiting Batteries	1.Large scale battery for DRI	Feasibility study to investigate practicality of using large-scale batteries to daily draw and store energy from the gird during hours of low carbon intensity to power a local DRI during hours of higher carbon intensity.	**Mission Focus,** Action-based research	**3**
96	DRI Mapping	1. Define the physical DRI	Create a consistent and reliable estimate of the physical DRI. This might be generated from a system-wide survey and/or a model based on sampling example facilities.	**Action-based research**	**5**
97	DRI Mapping	2. The level of accounting and engagement should match the size and impact of each DRI	Future engagement with DRI contacts should take into account the relative size and impact of different facilities. The amount of detail required from each component of the DRI should be	**Working with peers and suppliers**	**10**

		component	proportional to the potential impact of making policy and system changes. NOTE: This process should be managed carefully to avoid missing opportunities where a change to a relatively small system could be significant if applied across the DRI (e.g. a change to all UKRI laptops).		
98	DRI Mapping	3.Engage with interested parties and incentivise early engagement	Work with the contacts made within the DRI Mapping process to explore the most effective ways of maintaining and evolving community engagement to build the DRI database required in future. Consider incentives for improving engagement, such as: · Funding mechanisms for early adopters · Community champions – to reach out via their own networks · Support (i.e. tools/best practice advice/training) for early adopters · Introduce a voluntary reporting system, such as that provided by the Laboratory Environmental Assessment Framework (LEAF)	**Mission Focus**	**4**
99	DRI Mapping	4. Engage with all Research Councils	Contact should be made with each Research Council to ensure that the main DRI footprint of each is appropriately captured in a future DRI database.	**Recognise shared responsibility**	**5**
100	DRI Mapping	5. Categorise the DRI facilities by a defined set of metrics	Building on the results here, a set of categories should be defined to describe different types of DRI facility. A specific set of metrics and information should be agreed to in relation to each facility type.	**Working with peers and suppliers**	**14**
101	DRI Mapping	6. Share knowledge and experience of green energy adoption within the community	Some parts of the community have demonstrated that purchasing renewable energy can be combined with financial and scientific success. UKRI should consider a policy of purchasing electricity from renewable sources across the entire DRI. NOTES: · Experience and knowledge should	**Mission Focus** Working with peers and suppliers	**4**

			be shared in relation to on-site renewable energy generation. · Facility managers should be invited into the energy purchase discussions to improve the connection between high-level policies and practices at the facility level		
102	DRI Mapping	7. Share knowledge and experience of heat re-use in data centres	Engage with respondents that have heat re-use systems in place. Look for opportunities for knowledge exchange to make heat re-use into standard practice for UKRI.	**Knowledge Hub** Action-based research	**4**
103	DRI Mapping	8. User "annual energy usage" as fundamental measure in DRI database	The initial design for a database of DRI facilities should include annual energy usage.	**Action-based research** Knowledge Hub	**5**
104	DRI Mapping	9. Create standards for recording and reporting energy usage and efficiency	Agree standard metrics for recording and sharing efficiency and energy usage information. These metrics need to be developed in conjunction with appropriate tooling and hardware/software specifications that are accessible to all system managers. These standards should incorporate the GHG Protocol Emission Scopes.	**Action-based research**	**14**
105	DRI Mapping	10. Create a centralised resource to gather best practice and promote positive change	Provide a service, such as a centralised hub/group/knowledge-base, to gather best practice and share the outcomes/impacts of positive changes.	**Mission Focus** Knowledge Hub	**20**
106	DRI Mapping	11. Support facility managers in monitoring and reporting energy metrics	Provide guidance and support for facility managers to enable them to develop and deploy systems for monitoring and reporting energy usage.	**Knowledge Hub** Recognise shared responsibility	**14**
107	DRI Mapping	12. Develop a fair system for capturing energy usage at multiple levels.	Develop a fair approach for recording, monitoring and reporting energy usage information at appropriate levels (e.g. per job, per user, per group, per research theme). NOTE: It was pointed out that there may be unforeseen negative consequences of making energy	**Action-based research** Recognise shared responsibility	**23**

			usage information public (or within an institution) – it could be used to shame a high-energy user. A solution must be carefully managed to avoid such outcomes.		
108	DRI Mapping	13. Resource Net Zero and sustainability training	Information, support and training is essential to bring about the change required to meet the Net Zero target. This needs to be made a funding and resourcing priority so that: 1. The knowledge and human resources exist to provide the information/training. 2. The need to undertake training on sustainability issues (such as "green software engineering") is a funded and supported activity across all levels including decision-makers, PIs, system managers and end-users.	**Green software engineering** Knowledge Hub	**13**
109	DRI Mapping	14. Embed Net Zero in all layers of UKRI	Net Zero needs to become an explicit and significant concern of every single DRI user and stakeholder. A system of recognition, acknowledgement and reward is required to incentivise and encourage positive action on sustainability from all staff and users.	**Mission Focus**	**3**
110	DRI Mapping	15. Create a public database of DRI information that is updated regularly	A public database of key information about the UKRI DRI should be collected and updated routinely. Where possible, automated data feeds should be used to improve the data coverage and quality.	**Knowledge Hub**	**5**
111	DRI Mapping	16. Examine economies of scale: should more compute be centralised?	Gather a comprehensive dataset on the efficiency of facilities compared with their overall energy usage (e.g. their size). If there is a significant correlation: investigate options for moving more processing into centralised services to achieve higher overall efficiency.	**Green software engineering**	**7**
112	DRI Mapping	17. Ensure partially UKRI-funded resources are appropriately captured	Future information gathering regarding the carbon-footprint of UKRI-funded resources should ensure that some large facilities are not ignored because UKRI only funds	**Action-based research** Recognise shared responsibility	**5**

			a minority share. If they are significant in size, then this should be captured in a comprehensive future database.		
113	DRI Mapping	18. Analyse the contribution and impact of personal devices within the DRI	A detailed audit or analysis should be undertaken to understand what percentage of the overall UKRI DRI relates to personal devices such as desktop machines, laptops, tablets and phones.	**Green software engineering**	**7**
114	DRI Mapping	19. Analyse the contribution and impact of digital networks within the DRI	Carry out an assessment of the percentage contribution of the academic network to the overall carbon footprint of the DRI. If it is significant, then investigate how that footprint could be reduced.	**Green software engineering**	**7**
115	DRI Mapping	20. Analyse the contribution and impact of public cloud computing within the DRI	Carry out an assessment of the percentage contribution of the use of the public cloud to the overall carbon footprint of the DRI. Consider whether there might be environmental advantages of using the cloud for certain use cases.	**Green software engineering**	**7**
116	User Behaviour Survey	1. Encourage conversations about sustainability.	Conversations with teammates and working groups about sustainability should be encouraged, but advice about sustainable behaviour in the workplace should not come from the managerial level.	**Recognise shared responsibility**	**4**
117	User Behaviour survey	2. Re-strategise workplace practice to allow for more sustainable behaviours.	Build sustainability practices within workplace policies to enable an ease for sustainable behaviours without impacting upon employees' time.	**Mission Focus,** Action-based research	**16**
118	User Behaviour Survey	3. Encourage virtual conference attendance wherever possible	UKRI should encourage virtual conference attendance wherever possible and/or incentivise low carbon travel options in funding applications.	**Mission Focus**	**3**
119	User Behaviour Survey	4. Develop more effective networking strategies for web-conferencing	UKRI need to develop more effective networking strategies for online web-conferencing if we are to cut emissions through air travel.	**Mission Focus**	**7**
120	User Behaviour Survey	5. Encourage a low carbon commute to work	Incentivise the low carbon commute e.g. benefits for those using a car share scheme e.g. priority parking for those who are car sharing. Work with local bicycle shops to come up with	**Recognise shared responsibility**	**4**

			payment plans/discounts for purchases of bikes. Work with local councils to get discounted public transport for employees.		
121	User Behaviour survey	6. Rethink all policies about working from home given the obvious carbon/environmental implications	Workplaces should allow some element of working from home. By allowing employees the option of working from home it is reducing carbon emissions from the daily commute	**Recognise shared responsibility**	**1**
122	User Behaviour survey	7. UKRI/ all funders to include a 'Carbon Assessment' for all grant applications	In many funding applications and ethics applications academics are required to include a 'Risk Assessment' for their research and reviewers can determine whether their research is deemed 'low risk', 'medium risk', or 'high risk'. This procedure could be adapted to assess the carbon footprint of planned research projects by using a 'Carbon Assessment' tool. Categories would need to be carefully selected and include items like data storage, data sharing, equipment usage, conference travel etc. and a weighting would need to be assigned for each of the different categories. For example, if in the funding application a researcher was requesting financial support for attending a conference abroad and they were flying, this would be categorised as 'high carbon'. If on the other hand the researcher was requesting support for conference fees but not requesting travel money because they are attending a conference virtually, this would be deemed as 'low carbon'. Similarly, travelling to a conference within the U.K using public transport would be considered 'medium carbon'. Conference travel would have a high weighting due to the carbon emissions of transport, compared to equipment usage, for example. The application could then	**Recognise shared responsibility**	**2**

			be assigned an overall ‘Carbon Assessment’ score (totalling all of the categories with the weighting calculated). Funders will then be able to assess the carbon footprint of the research project and feed back to the researchers and advise as to how to reduce it (if necessary).		
123	User Behaviour survey	8. Allow people more input into the procurement process	Give people more input into the decision-making when upgrading to more energy efficient equipment. In most institutions only allocated staff make purchasing decisions about new equipment, and these upgrades may not necessarily be more energy efficient. If people who are working with specialised equipment can argue a case for a more energy efficient upgrade.	**Working with peers and suppliers**	**4**
124	User Behaviour Survey	9. Demonstrate how old equipment has benefited others	Provide evidence to employees as to how old equipment has benefited other people, for example, if old equipment has been donated to schools in harder to reach areas, if equipment is helping people in third world countries etc. and utilise an easy recycling/reconditioning scheme for obsolete technology to encourage less waste.	**Knowledge Hub**	**4**
125	User Behaviour Survey	10. More training regarding saving energy when running equipment	A policy for best practices for saving energy when running equipment is needed in all departments with clear instructions for all technology. In addition to this - more information about the carbon implications of equipment left on standby is needed. Reminders around the workplace should be placed in visible locations to remind people to turn equipment off when it is not in use.	**Knowledge Hub**	**1**
126	User Behaviour Survey	11. Clear communication of the cost/carbon benefit of equipment upgrades	Provide staff with clear documentation as to how specific equipment upgrades are more energy efficient.	**Knowledge Hub** Working with peers and suppliers	**4**
127	User	12. Training for best	People need to be made aware of the	**Green software**	**4**

	Behaviour Survey	practice for reducing the amount of energy all research requires.	environmental impact of their research and how to reduce it. Develop information campaigns to raise levels of response efficacy that small changes can make a difference to the carbon emissions it uses.	**engineering** Knowledge Hub	
128	User Behaviour Survey	13. Information regarding the carbon emissions of activities around the workplace with training for best practice for reducing this.	People need to be made aware of the environmental impact of their work – including various activities around the workplace. Information needs to be provided as well as training to increase levels of response efficacy and clear instructions provided as to how to go about reducing carbon emissions. Again, a six-month appraisal may help to provide continued encouragement for people to be more conscious of the environmental impact of their work.	**Knowledge Hub** Recognise shared responsibility	**4**
129	User Behaviour Survey	14. Provide more information and encouragement to facilitate the workforce to collaborate when writing code	Make the workforce more aware of the energy implications in writing code and to encourage collaboration.	**Knowledge Hub** Green software engineering	**4**
130	User Behaviour Survey	15. Provide training to PhD students and ECRs to work with specialised low carbon data storage and computing facilities (e.g. JASMIN and ARCHER2).	UKRI need to include expert training for their funded PhD students and ECRs to utilise specialised supercomputing centres and encourage the use of large data storage facilities such as JASMIN and ARCHER2 to ensure students and researchers are aware of the capabilities (in terms of energy efficiency) of these supercomputing services	**Green software engineering** Knowledge Hub	**13**
131	User Behaviour survey	16. Incorporate/encourage the use of specialised low carbon data storage and computing facilities in UKRI policies	UKRI need to incorporate the use of specialised low carbon data storage and computing facilities (e.g. JASMIN and ARCHER2) into their policies and funding calls to facilitate energy efficient computing. It may be necessary to include training for those who do not know how to use such facilities.	**Knowledge Hub**	**10**

132	User Behaviour Survey	17. Regular profiling of the workforce on attitudinal measures using both self-reports and implicit associative measures	Regularly test employees on their explicit attitudes to carbon and their implicit attitudes to carbon. This would be useful in identifying what some of the main obstacles are in guiding the workforce in more sustainable practices.	**Action-based research**	**4**
133	User Behaviour Survey	18. Adopt different approaches to behaviour change strategies.	Behaviour change strategies will need to be different to encourage behaviour change for individuals falling within the different segments of high/low explicit attitudes and high/low implicit attitudes, for example, those with lower implicit attitudes will need a different approach to raise their level of response efficacy and self-efficacy.	**Action-based research**	**9**
134	Sustainable Computing	1. General recommendation	Our general recommendation is that institutions should adopt the LEAF (Laboratory Efficiency Assessment Framework) for Digital Infrastructure.	**Green software engineering** Working with peers and suppliers Action-based research	**2**
135	Sustainable Computing	2. Minimising full lifecycle emissions of institutional facilities	Monitoring of facility load and coupling to replacement cycle; Adopting access policies that incentives efficient time use of facilities and prevent squatting	**Mission Focus**	**11**
136	Sustainable Computing	3. Optimising code and reducing coding mistakes	Make the use of a workflow implementing best practice for testing and a transparent code review process mandatory for deploying on controlled facilities.	**Green software engineering**	**16**
137	Sustainable Computing	4. Providing training on Software Engineering best practices	• Principles of Green Software Design • Code review: focus on energy efficiency of architecture, algorithms and implementation • Testing: unit, integration and acceptance testing • Use of build systems and revision control systems, in particular Continuous Integration workflows • Compilation optimisation: use of compiler options to optimise energy efficiency • Design of Experiments: ensure that	**Green software engineering**	**13**

			the experiment produces the required results with minimal energy expenditure		
138	Sustainable Computing	5. Offering researchers expert support	Offering researchers expert support in the form of either trained Research Software Engineers or volunteer researchers trained to be experts.	**Green software engineering** Knowledge Hub	**4**
139	Sustainable Computing	6. Ensure adequate training of experts	Ensure adequate training of experts, specifically: • Code review for energy efficiency • Green Software Engineering practices • Code energy efficiency evaluation and optimisation • Compilation optimisation for energy efficiency • Design of Experiments optimisation for energy efficiency	**Green software engineering**	**6**
140	Sustainable Computing	7. Invest in research and development of better tools	• UKRI should encourage research into novel approaches to improve energy efficiency of scientific software. • UKRI should provide specific funding schemes to develop proof-of-concept tools for improving energy efficiency into software products usable by non-expert researchers	**Green software engineering** Action-based research	**9**
141	Art commission	1. Diversity in communications channels	Employ specialist communications and engagement expert/s whose sole responsibility is to work with (sit within?) the environmental sustainability group to ensure messages are clear and engaging for a diverse audience. Their main responsibility is to influence diverse groups of people to work together and enable collective action for net zero DRI.	**Recognise shared responsibility**	**4**
142	Art commission	1. Diversity in communications channels	The specialist communications expert should oversee continuation of community engagement about net zero DRI, provision of support, examples of best practice and relevant success stories (e.g. from existing work being done by UKRI	**Recognise shared responsibility**	**4**

			researchers, and also the outputs from the competitive funding calls).		
143	Art commission	2. Enabling creative partnerships	fund a series of creative workshops and physical art exhibition that uses the existing art commission work to inspire conversations about net zero and digital research infrastructure.	**Recognise shared responsibility**	**4**
144	Art commission	2. Enabling creative partnerships	fund projects that enable cross-UKRI researchers to work together creatively to share perspectives and discuss solutions about how to implement changes for net zero dri.	**Recognise shared responsibility**	**6**

A3.4 Community and Consensus Building Activity

See Section 3.5 for discussion.

Citable Ref #	**Source Report**	**Recommendation title**	**Statement**	**Theme(s)**	**Roadmap Action**
145	Training and Standards	1. Train facility managers to monitor and evaluate carbon usage	Train facility managers to monitor and evaluate energy/carbon usage at both user and system level	**Knowledge Hub** Recognise shared responsibility	**13**
146	Training and Standards	2. Train individual users in [environmentally responsible] good practice	Train individual users in good practice, such as code testing and optimisation, reflecting the latest knowledge and tools	**Knowledge Hub** Recognise shared responsibility	**13**
147	Training and Standards	3. Train RSEs to provide cross-sector support to optimise code for deployment (re: carbon efficiency)	Train dedicated teams of Research Software Engineers (RSEs) to provide cross-sector support and optimise scientific code for deployment	**Green software engineering**	**13**
148	Training and Standards	4. Develop Open Science and FAIR data standards	Develop Open Science and FAIR data standards and train researchers in them to maximise good practice, efficiency, data sharing, discoverability and reuse	**Green software engineering** Working with peers and suppliers	**15**
149	Training and Standards	5. Mandate training on import of working practices to acheve Net Zero DRI	Create mandatory training within UKRI on WHY it is important to change working practices to achieve Net Zero	**Knowledge Hub**	**16**
150	Training and	6. Co-develop	Develop/Agree standards for	**Action-based**	**14**

	Standards	standards for metrics and reporting methodology	metrics and reporting to enable consistent measurement and monitoring of carbon usage, across platforms, sectors and institutions	**research**	
151	Training and Standards	7. Train individuals on carbon-efficient data management practices	Provide training to individuals on carbon-efficient data management practices	**Knowledge Hub** Recognise shared responsibility	**13**
152	Training and Standards	9. Develop accreditation for training in environmentally sustainable data management as part of on-going professional development	Develop accreditation for training in environmentally sustainable data management as part of ongoing professional development	**Working with peers and suppliers** Green software engineering	**13**
153	Training and Standards	10. Develop standards for the delivery and explotation of carbon efficient services and software	Develop standards for the delivery and exploitation of big data through carbon-efficient services and software. E.g. server-side subsetting to reduce data transfer and storage loads	**Green software engineering** Working with peers and suppliers	**14**
154	Training and Standards	11. Mandate the inclusion of an environmental impact statement	Mandate the inclusion of an environmental impact statement, along with mitigating actions, within applications for research funding	**Recognise shared responsibility**	**3**
155	Training and Standards	12. Include a budget within funding calls to support researchers in engaging with Net Zero goals	Include a budget within funding calls to support researchers in engaging with Net Zero goals, e.g. general training for scientists or access to specialist expertise.	**Green software engineering** Working with peers and suppliers	**1**
156	Training and Standards	13. Mandate carbon monitoring and reporting in funding calls	Mandate carbon monitoring and reporting in funding calls	**Recognise shared responsibility**	**3**
157	Procurement	1.1 All spending decisions must include a proportionate assessment of their impact on the UKRI carbon budget and on the implementation of	All spending decisions must include a proportionate assessment of their impact on the UKRI carbon budget and on the implementation of the Net Zero policy.	**Mission Focus**	**10**

		the Net Zero policy.			
158	Procurement	1.2 Standards need to be developed to ensure that environmental sustainability assessments are made robustly and efficiently.	Standards need to be developed to ensure that environmental sustainability assessments are made robustly and efficiently. There are many existing standards, but there are also cases for which new standards need to be set or developed.	**Working with Peers and Suppliers**	**14**
159	Procurement	1.3 UKRI must have a policy on overall power consumption of facilities which is aligned with the Climate Change Committee balanced pathway to net zero.	UKRI must have a policy on overall power consumption of facilities which is aligned with the Climate Change Committee balanced pathway to net zero. Although the overall consumption barely registers on the scale of the national sectoral analysis considered by the CCC it is important that UKRI should provide leadership in explaining how their investment decisions align with the CCC recommendations, including the recommendation that power consumption for existing activities needs to be held constant or reduced in order to enable a timely transition to renewable power.	**Working with Peers and Suppliers**	**10**
160	Procurement	2.1 UKRI must take a proactive approach to ensuring that net zero policy does not disrupt research programmes and that prioritisation of low carbon investment and purchasing options does not have a disproportionate negative impact.	UKRI must take a proactive approach to ensuring that net zero policy does not disrupt research programmes and that prioritisation of low carbon investment and purchasing options does not have a disproportionate negative impact.	**Mission Focus**	**4**
161	Procurement	2.2 UKRI must ensure that steps taken to reduce environmental impact do not end up having the opposite effect through feedback such as the rebound, or Jevons	UKRI must ensure that steps taken to reduce environmental impact do not end up having the opposite effect through feedback effects such as the rebound, or Jevons effect.	**Mission Focus**	**22**

		effect.			
162	Procurement	3.1 Training of staff at all levels is needed, both to increase awareness and understanding of the implications of climate change and the net zero policy, and to provide technical competence to deliver change	Training of staff at all levels is needed, both to increase awareness and understanding of the implications of climate change and the net zero policy, and to provide technical competence to deliver change. Training needs to be backed by an active programme of learning and discovery. The roadmap to net zero will pass through unexplored territory and training material will need to be regularly updated with lessons learned from exploratory pathfinder projects at UKRI and elsewhere.	**Recognise shared responsibility**	**13**
163	Values and Responsibilities workshop	A. leverage institutional power: demonstrate action at the institutional level, with visibility across the research community in order to ensure cooperative organisational and individual responsibility and inspire positive change	See section 3.5.3.3	**Mission Focus**	**1**
164	Values and Responsibilities workshop	B. make information about the relative benefits of different actions clear and readily accessible to the research community, in order to empower decision making by individual researchers and groups	See section 3.5.3.3	**Knowledge Hub**	**4**
165	Values and Responsibilities workshop	C. include environmental sustainability within funding assessment and award processes, so that it is planned	See section 3.5.3.3	**Recognise shared responsibility**	**12**

		into the project from the outset, and evaluated as part of funding applications			
166	Values and Responsibilities workshop	D. request that projects estimate their carbon footprint, even crude ones	See section 3.5.3.3	**Recognise shared responsibility**	**12**
167	Values and Responsibilities workshop	E. provide/develop a rating of host (research) institutes by the sustainability of their operations and projects	See section 3.5.3.3	**Recognise shared responsibility**	**5**
168	Early Career Researchers workshop	1. Ensure collective responsibility for Net Zero DRI	(i) move towards an empowered and equitable research community, bringing top-down and grass-roots actions together (ii) develop a shared and accessible body of knowledge to support collective action, avoiding silos of practice	**Recognise Shared responsibility**	**4**
169	Early Career Researchers workshop	2. Resource green tools, including metrics, guidelines and continued professional development	Resource green tools, including metrics, guidelines and continued professional development, ensuring resource availability to address capability gaps between researchers in High Income Country and Low and Lower-middle Income Country settings	**Green software engineering**	**8**
170	Early Career Researchers workshop	3. Embed environmental sustainability within research design, practice, and assessment	Embed environmental sustainability within research design, practice, and assessment	**Mission focus**	**1**
171	Early Career Researchers workshop	3. Embed environmental sustainability within academic career progressions pathways	Ensure that academic career progression is supported via alternatives to flying, and improvements to virtual networking capability; integrate environmental sustainability within graduate and post-graduate courses to support workforce demand for ‘green skills’	**Mission focus**	**17**
172	Community Engagement	A. Challenging assumptions re: tech	See section 3.5.4.2	**Green software**	**4**

	workshop	solutionism and interrogate scope of benefit vs risk		**engineering**	
173	Community Engagement workshop	B. Improving interoperability and accessibility of infrastructure	See section 3.5.4.2	**Action-based research**	**6**
174	Community Engagement workshop	C. Producing guidelines for procurement and best practice	See section 3.5.4.2	**Working with peers and suppliers**	**14**
175	Community Engagement workshop	D. Investing in green resources, training and skills	See section 3.5.4.2	**Green software engineering**	**6**
176	Community Engagement workshop	E. Embedding green principles in funding processes	See section 3.5.4.2	**Working with Peers and Suppliers**	**12**
177	Community Engagement workshop	F. Prioritising environmental sustainability and raise awareness	See section 3.5.4.2	**Mission Focus**	**22**
178	Community Engagement workshop	G. Implement learning from applicable case studies to support environmentally responsible use of DRI (e.g. UKRI leadership in reduction, replacement and removal of animals use in research, Wellcome Trust leadership in promoting public engagement with research)	See section 3.5.4.2	**Knowledge Hub**	**4**
179	Community Engagement workshop	H. Incentivise environmentally sustainable research practice	See section 3.5.4.2	**Recognise shared responsibility**	**12**
180	Community Engagement workshop	I. Advocate for sharing of resources including infrastructure (facilities) and data	See section 3.5.4.2	**Recognise shared responsibility**	**12**

Appendix 4. Named contributors to the scoping project

The table below shows 94 named individuals (alphabetical by surname) who significantly contributed to the project.

Thank you to anyone else who contributed to our project in any way. There were many operational, communications and events staff behind the scenes who greatly helped us with tasks such as; sharing and advertising our events, attending or contributing to events/meetings/creative workshops, providing general feedback and advice, etc.

This project has been a very collaborative effort - and would not have been possible without everyone's involvement.

First Name	Surname	Organisation	Involvement in project
Florian	Ahrens	Heriot-Watt University	Project partner - sandpit
Burak	Akyol	University of Bristol	Project partner - sandpit
Anna	Angus-Smyth	Natural Environment Research Council	Funder
Alberto	Arribas Herranz	European Environment Agency Scientific Committee (previously Microsoft)	Scientific Advisory Board
Richard	Bailey	Engineering and Physical Sciences Research Council	Steering Committee
Michael	Bane	Manchester Metropolitan University	Project partner - sandpit, Technical report author
Alastair	Basden	Durham University	Project partner - sandpit
Nick	Beard	National Centre for Atmospheric Science	Project partner - consortium
Deepayan	Bhowmik	Newcastle University	Project partner - sandpit
Caroline	Bird	University of Bristol	Project partner - sandpit
Mary Ethna	Black	National Centre for Atmospheric Science and University St. Andrews	Chair of Scientific Advisory Board, Steering Committee, Sandpit Review Panel
Jack	Boulton	Heriot-Watt University	Project partner - sandpit
Oliver	Brown	EPCC, University of Edinburgh	Project partner - sandpit
Justin James Henry	Buck	National Oceanography Centre	Project partner - sandpit
Jennifer	Bulpett	Centre for Environmental Data Analysis	Core project team member
Matt	Burrows	The University of Reading / Engaging Environments	Critical friend in Policy and Communications

Katie	Cartmell	Centre for Environmental Data Analysis	Core project team member
Tao-Tao	Chang	Arts and Humanities Research Council	Art Commission Review Panel
Ruth	Chaplin	National Centre for Atmospheric Science	Project partner - consortium
Neil	Chue Hong	EPCC, University of Edinburgh	Steering Committee/Sandpit Review Panel/DRI Mapping Expert Advisory Group
Charlotte	Clarke	Natural Environment Research Council	Funder
Jose Alejandro	Coronado Arciniegas	University College London	Scientific Advisory Board, Project partner - consortium
Alastair	Dewhurst	Science and Technology Facilities Council	Project partner - sandpit
Damu	Ding	University of Oxford	Project partner - sandpit
Adrian	Friday	Lancaster University	Project partner - sandpit
Emma	Fryer	TechUK	Steering Committee
Dawn	Geatches	Science and Technology Facilities Council	Project partner - sandpit
Alyssa	Gilbert	Imperial College London	Scientific Advisory Board
Glenn	Greed	UK Met Office	Steering Committee
David	Greenwood	Newcastle University	Project partner - sandpit
Jonathan	Hays	Queen Mary University of London	Project partner - sandpit, DRI Mapping Expert Advisory Group
Xinpeng	Hong	University of Oxford	Project partner - sandpit
Adrian	Jackson	EPCC, University of Edinburgh	Project partner - sandpit
Sophie	Janacek	UK Research and Innovation DRI Programme	General advisor
Catherine	Jones	Science and Technology Facilities Council	Sandpit Review Panel
Martin	Juckes	Centre for Environmental Data Analysis	Core project team member
Rainer	Kattel	University College London	Project partner - consortium
Gabin	Kayumbi	Science and Technology Facilities Council	Project partner - sandpit and consortium
Susan	Krumdieck	Heriot-Watt University	Project partner - sandpit
Luca	Kuhn Von Burgsdorff	University College London	Scientific Advisory Board, Project partner - consortium

Simon	Lambert	Science and Technology Facilities Council	DRI Mapping Expert Advisory Group, Project partner - sandpit
Loïc	Lannelongue	University of Cambridge	DRI Mapping Expert Advisory Group
Ian	Leslie	University of Cambridge	Scientific Advisory Board
Carolynne	Lord	Lancaster University	Project partner - sandpit
Alvaro	Lorenzo Lopez	National Oceanographic Centre	Project partner - sandpit
Miranda	MacFarlane	Kings College London	Project partner - consortium, Technical report author
Molly	MacRae	Centre for Environmental Data Analysis	Core project team member
Danae	Manika	Brunel University London	Project partner - sandpit
Mariana	Mazzucato	University College London	Scientific Advisory Board, Project partner - consortium
Niall	McCarroll	University of Reading	Project partner - sandpit
Fergus	McDonald	Health Data Research UK	Sandpit Review Panel
Laura	McGuire	Edge Hill University	Project partner - consortium
Chris	Michaels	The National Gallery, London	Art Commission Review Panel
Paul	Millhouse-Smith	Freelance artist	Project partner - consortium
Stephen	Mobbs	National Centre for Atmospheric Science	Project Investigator
Lorenza	Monaco	University College London	Project partner - consortium
Sophie	Mosselmans	Summer student at the Centre for Environmental Data Analysis	Core project team member
Anish	Mudaraddi	Science and Technology Facilities Council	Project partner - sandpit
Erinma	Ochu	UWE Bristol	Art Commission Review Panel
Peter	Oliver	Science and Technology Facilities Council	DRI Mapping Expert Advisory Group
Alex	Ogden	Institute of Astronomy, University of Cambridge	Project partner - sandpit
Alex	Owen	Queen Mary University of London	Project partner - sandpit, Technical report author
Alison	Packer	Science and Technology Facilities Council	Project partner - sandpit
Alison	Pamment	Centre for Environmental Data Analysis	Project partner - consortium
Charlotte	Pascoe	Centre for Environmental Data Analysis	Core project team member

Jess	Phillips	Natural Environment Research Council	Funder
Chris	Preist	University of Bristol	Project partner - sandpit
Jamie	Quinn	University College London	Project partner - sandpit
Stefan	Reis	UK Centre for Ecology & Hydrology & NERC Science Committee	Steering Committee
Harriett	Richardson	National Centre for Atmospheric Science	Art Commission Review Panel, Critical friend in Communications
Daniel	Schien	University of Bristol	Project partner - sandpit
Marian	Scott	University of Glasgow	Steering Committee
Paul	Shabajee	University of Bristol	Project partner - sandpit
Andrea	Sharpe	Natural Environment Research Council	Funder
Emily	Shuckburgh	University of Cambridge	Scientific Advisory Board
Alan	Simpson	EPCC, University of Edinburgh	Project partner - consortium
Lorna	Smith	EPCC, University of Edinburgh	Project partner - consortium
Adam	Staines	Medical Research Council and Biotechnology and Biological Sciences Research Council	Steering Committee
Ag	Stephens	Centre for Environmental Data Analysis	Core project team member
Ole	Stubben	University of Glasgow	Project partner - consortium
Alexander	Taylor	National Centre for Atmospheric Science	Project partner - consortium
Sean	Tonkin	University College London	Project partner - consortium
Poppy	Townsend	Centre for Environmental Data Analysis	Core project team member
Andy	Turner	EPCC, University of Edinburgh	Project partner - sandpit
Wim	Vanderbauwhede	University of Glasgow	Scientific Advisory Board, Sandpit Review Panel, Project partner - consortium
Graham	Waddell	UK Research and Innovation (environmental sustainability team)	General advisor, Steering Committee, DRI Mapping Expert Advisory Group
Nicholas	Walton	Institute of Astronomy, University of Cambridge	Project partner - sandpit
Michele	Weiland	EPCC, University of Edinburgh	Sandpit Review Panel

Daniel	Whitehouse	Imperial College London	Project partner - sandpit
Kelly	Widdicks	Lancaster University	Project partner - sandpit
Lucy	Woodward	Summer student at the Centre for Environmental Data Analysis	Core project team member
Adam	Young	Tech UK	Steering Committee
Claire	Young	University of Bristol	Project partner - sandpit
Noa	Zilberman	University of Oxford	Project partner - sandpit