Developing a Sustainable Institutional Research Computing Culture

Jeremy Cohen[1]
Advanced Research Fellow, Department of Computing
Director of Research Software Engineering Strategy

Hannah Scott[1], Diego Alonso Alvarez[1], Michael Bearpark[1], Christopher Cave-Ayland[1], Rae Goddard[2], Emily Lumley[1], Katerina Michalickova[1] and Andrew Turner[3]

[1] Imperial College London, [2] Paraphrase.studio, [3] EPCC, University of Edinburgh

SC4RC26: Sustainability Conference for Responsible Research Computing
Wednesday 6th May 2026, Geneva, Switzerland

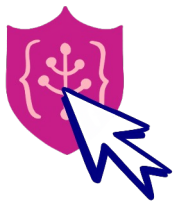

IMPERIAL

Developing a research culture focused on sustainability

Outline

- **Developing a research culture focused on sustainability**
  - Motivation and opportunity
  - Challenges for research
- **Understanding your research and computing infrastructure**
- **The 3 pillars of a sustainable research computing culture**
- **Final thoughts / conclusions**

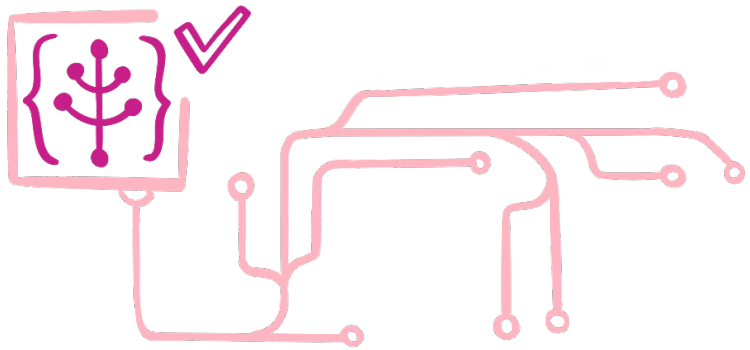

Community | **Education** | **Advocacy**

Developing a research culture focused on sustainability

Motivation

- Research computing now a core part of the research lifecycle across almost all domains
- Driven by AI/ML/LLMs, demand for resource capacity is growing rapidly
  - McKinsey predicted (April 2025) global data centre capacity demand to almost triple by 2030 – **70%** of increase due to AI workloads (with caveats)[1]
  - Planned new data centres to ”add around **6.2GW** of IT power capacity by 2030” [2]
- How much of this demand is the result of inefficient/poorly architected code, inefficient model training methods or unnecessary use of resource-intensive tooling?

[1] **The cost of compute: A $7 trillion race to scale data centers. McKinsey Quarterly, 28th April 2025.** https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers

[2] **C. Hahn. The UK’s data centre boom. Oxford Economics, 8th Dec 2025.** https://www.oxfordeconomics.com/resource/the-uks-data-centre-boom-growth-trends-drivers-and-the-rising-power-challenge/

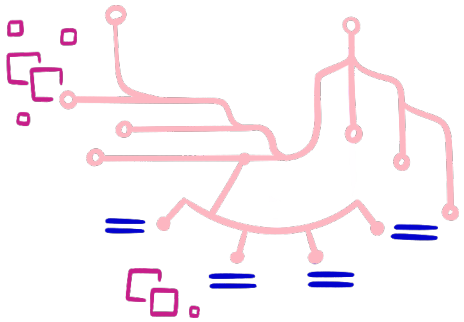

Developing a research culture focused on sustainability

Motivation

- Increased demand balanced by increasingly efficient hardware and models[3]
- But expected cumulative 2025-2035 emissions significantly higher than previously suggested, representing 0.9-3.4% (34-123million tonnes) of projected UK CO_2 emissions over this period[3]
- Large-scale research computing infrastructure increasingly operated partly/fully by renewable energy (e.g. ARCHER2[4]) – but scope 3 emissions can still be significant – important to "do more with less" (resources)

[3] **UK Compute Roadmap, Compute Evidence Annex: changes corrected, 23 April 2026. DSIT/UKRI Available at:** https://www.gov.uk/government/publications/uk-compute-roadmap/compute-evidence-annex-changes-corrected-23-april-2026

[4] **Environmental Sustainability, ARCHER2. EPCC. Available** at: https://www.archer2.ac.uk/community/sustainability/

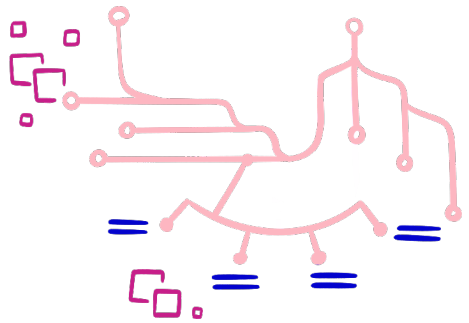

Challenges: Writing efficient, optimised code requires time and specialist skills

- Significant proportion of researchers write code;
  A significant proportion of developers are self-taught
  - As far back as 2014, an SSI survey found **56%** of researchers at UK Russell Group universities write software, with over a fifth of them having no software engineering training[5]
  - HackerRank's 2018 Developer Skills Report surveyed over **39,400** developers - almost **74%** are in-part self-taught[6]
- Researchers increasingly vibe coding / leveraging AI-assisted coding to some extent
- Lack of software training makes critically assessing AI outputs challenging!

[5] **S. Hettrick. It's impossible to conduct research without software, say 7 out of 10 UK researchers. SSI blog. Dec 2014.** https://www.software.ac.uk/blog/its-impossible-conduct-research-without-software-say-7-out-10-uk-researchers

[6] **2018 Developer Skills Report. HackerRank.** https://www.hackerrank.com/research/developer-skills/2018

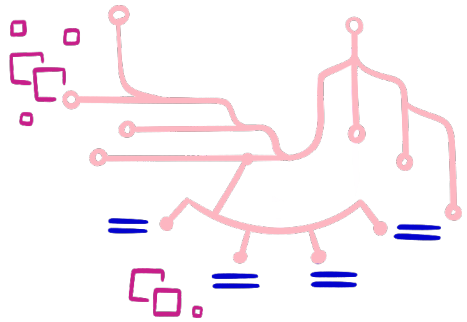

Challenges: Researchers' approach to software development

- Growing complexity of large-scale digital research = new skills to learn
- Focus often on results – need to balance with building recognition of the value of optimisation
- Domain researchers who code often have different priorities to software engineers:

Focus on results rather than code architecture/structure/efficiency	Sometimes long-running, inefficient code is quicker than learning code optimisation skills! (institutional resources often free at point of use)	Code often made open – may be adopted by others, growing impact of inefficiencies

- Sustainability isn't yet the focus – need to improve support for researchers – build skills into developing competency frameworks such as DIRECT?

Challenges: Motivating change in a complex, multidisciplinary environment

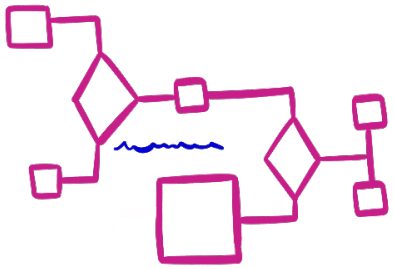

- *Can't answer the question about how much resource is wasted by inefficiencies in code and use of compute resources*
- **But we can encourage our research community to think about their own personal resource/infrastructure use and develop sustainability awareness through basic training**
- *Can't (and shouldn't) prevent researchers using AI tools for their work*
- **But we can educate them on the how these tools work, associated costs and what other skills to learn to support use of AI - make informed decisions about when/how to use AI**

Understanding your research and computing infrastructure

Understanding your research

How much energy is your use of LLMs for software development consuming?

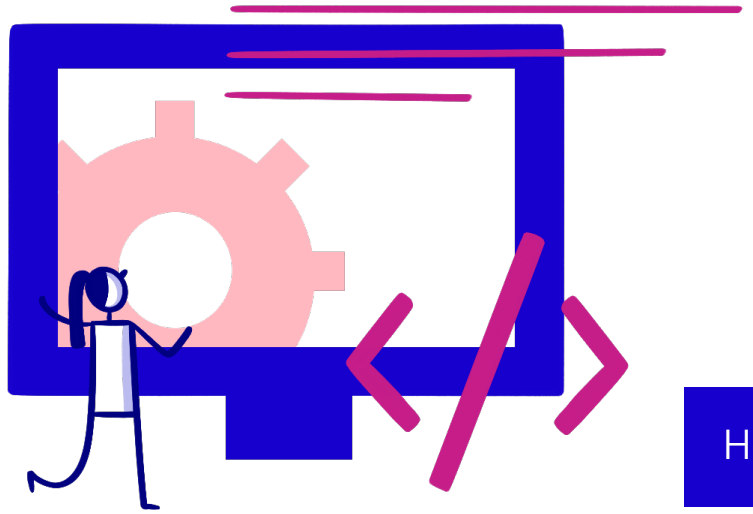

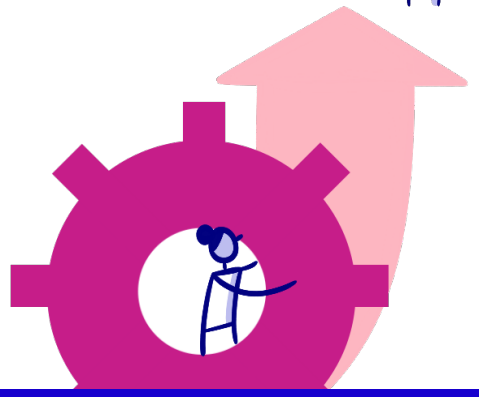

How efficient is your code?

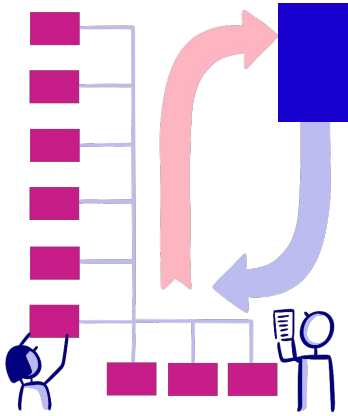

Are you efficiently managing your data?

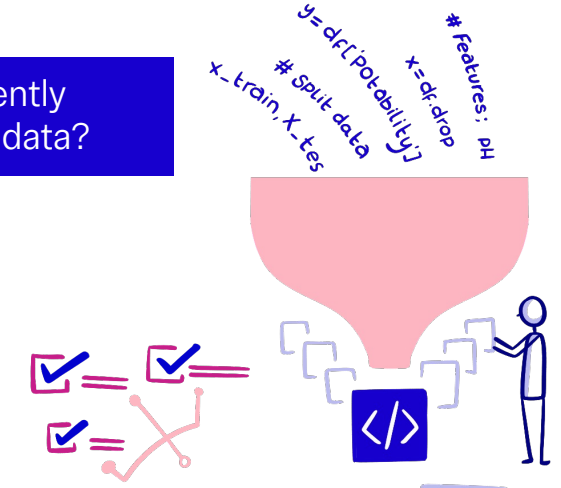

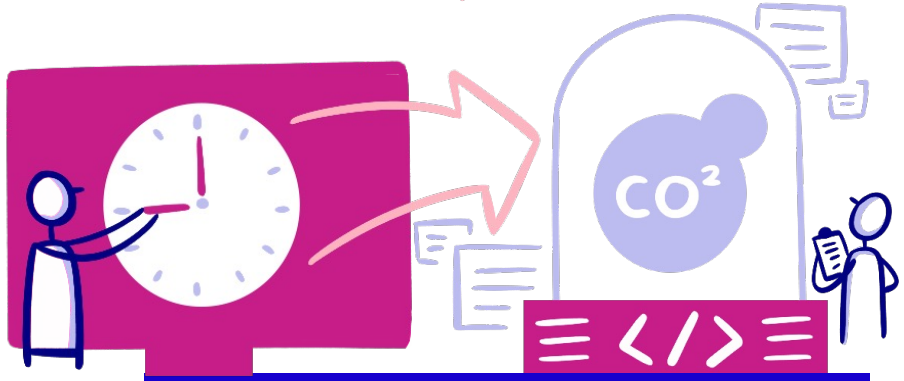

Do you know how efficient your computing infrastructure is?

Understanding your local infrastructure

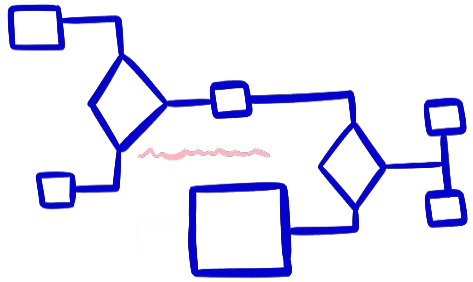

- Can’t leave it to regional/national infrastructure operators to deliver skills and awareness around sustainability/efficiency
- We need a local “people and skills” infrastructure to support our computing use and underpin efficient use of computing infrastructure
- Local infrastructure is also increasingly complex and heterogeneous
- Example:
  - Imperial College London’s Research Computing Service provides High Performance Computing and Research Data infrastructure + Research Software Engineering skills, support and training.
  - HPC infrastructure supports 6,000+ registered users – mostly research students/researchers/academics

The HEX Platforms – Imperial College London

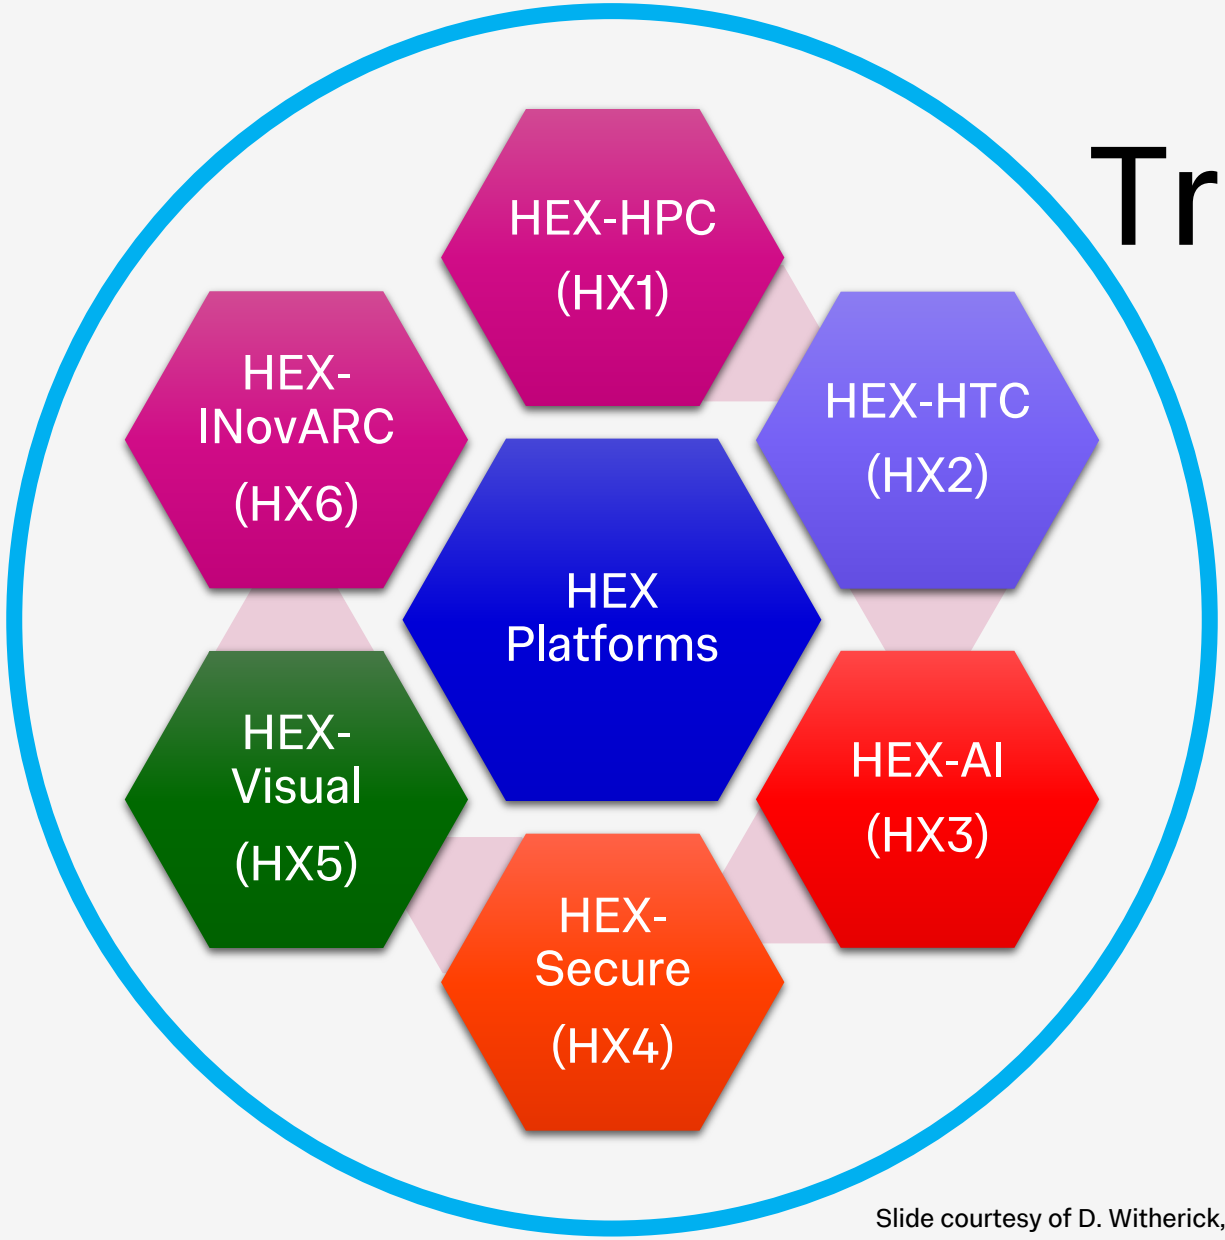

Training

Slide courtesy of D. Witherick, Research Computing Service, Imperial College London

HX2
HEX-HTC

Cluster Specification

High throughput facility to eventually replace CX3	
27,000+ cores of Intel's 6th Generation of Xeon Scalable server processors (Granite Rapids)	
96 NVidia Hopper-based H200 (141GB) GPUs	
2PB of all-flash local disk using the Weka data-storage platform	
800 Gb ethernet core	200 Gb ethernet to every compute node 800 Gb (4 x 200 Gb) to every GPU node
RHEL9.6 with Slurm workload scheduler	

Slide courtesy of D. Witherick, Research Computing Service, Imperial College London

Understanding local infrastructure

- To make effective, sustainable and efficient use of such local infrastructure, we need more than just training
- Need to bring the research community and institutional leadership with us on our journey to sustainable research computing
- Embed core principles of sustainability and efficiency into our computational research alongside quality, reproducibility and sustainability

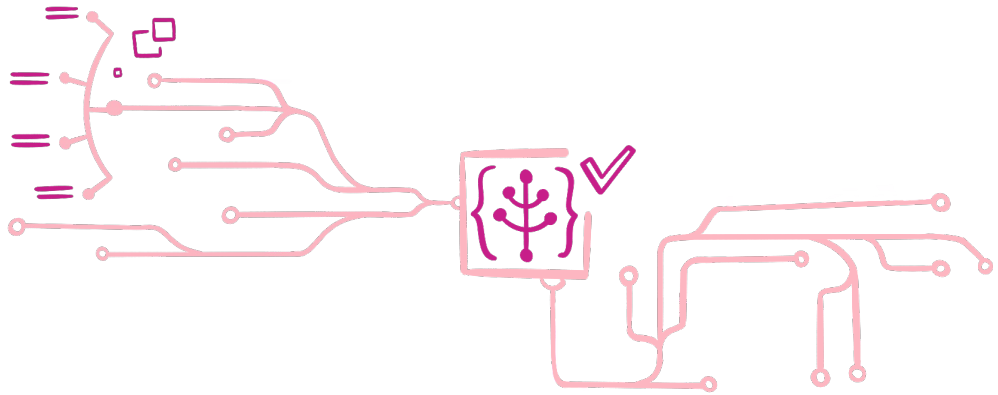

A Sustainable Research Computing Culture

The 3 pillars of a sustainable research computing culture

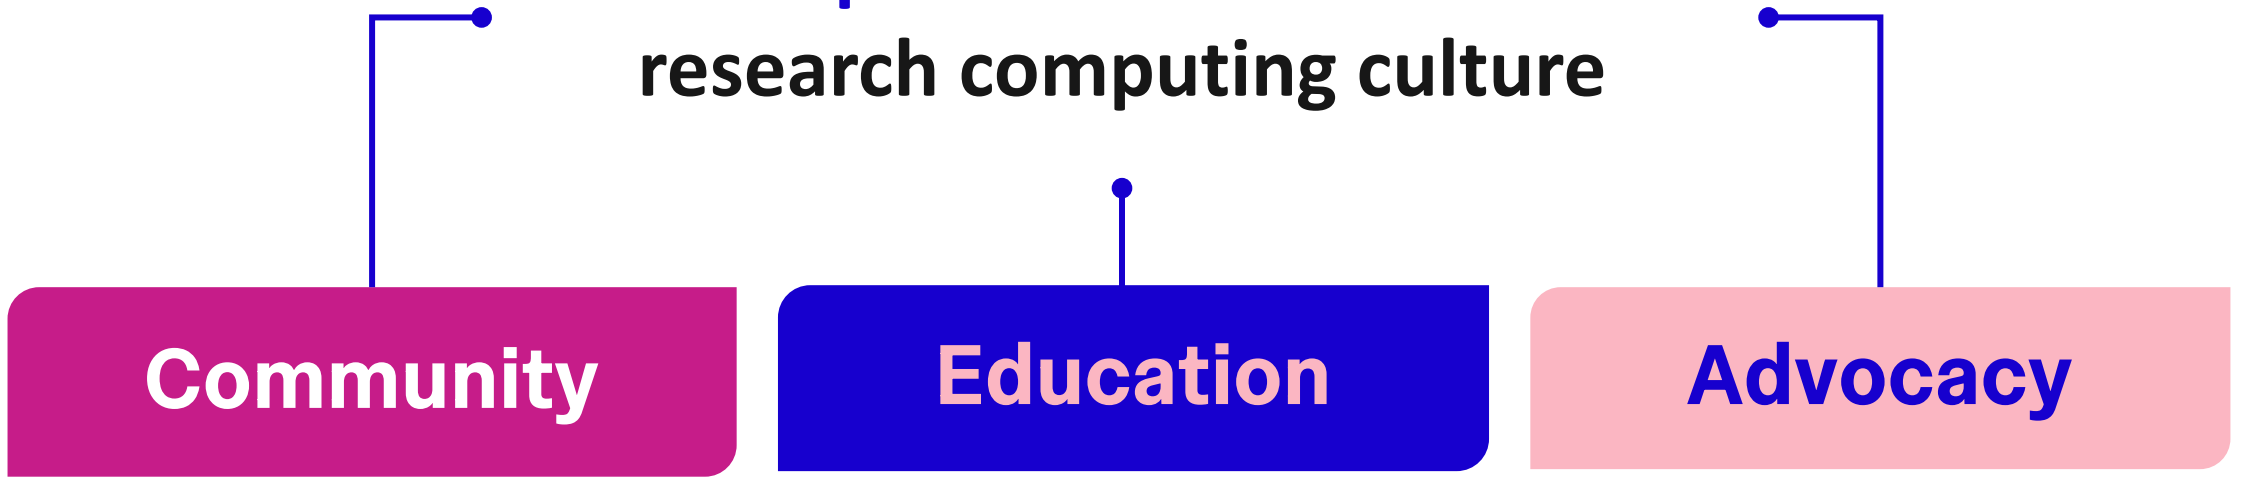

The 3 pillars of a sustainable research computing culture

Community

- **Collaboration is at the core of what researchers do**

  - Leverage existing research communities to raise awareness of sustainable research computing practices

  - Share knowledge, ideas, benefits, opportunities across local groups/communities

- **Green DiSC: Digital Sustainability Certification**

  - Beneficial for institutions, resource operators/managers, research departments

  - Important community building activity – get buy-in from community as part of certification process

The 3 pillars of a sustainable research computing culture

Community

- **Events, hackathons, clinics**
  - Build on community interest and engagement to share knowledge, skills, best practices through seminars
  - Hackathons – green coding hackathons (ecathons?)
- **Build connections with other local/regional communities**
  - Share green/sustainable computing practices
  - Motivate a wider shift towards more efficient institutional computing models
- **Consider a local Champions scheme**
  - Ref: NetDRIVE Sustainable Research Computing Champions[7]
  - Example model: STEP-UP Digital Research Technical Champions

[7] **Call for Sustainable Research Computing Champions. NetDRIVE Project.** https://uknetdrive.org/netdrive-calls/call-for-champions-for-transformational-change

STEP-UP Digital Research Technical Champions

- **PhD students providing digital skills knowledge and support**

  - Sustainable computing can be an important element of this

  - Run through the London region STEP-UP project – 4 Central London university partners

  - Champions focus on work within their home department – engage with the rest of the cohort through regular catch-ups

  - Paid for their time, flexibility to work on their own ideas to best support their local community

STEP-UP
https://step-up.ac.uk

The 3 pillars of a sustainable research computing culture

Education

Development of a **“Digital Research Sustainability”** course

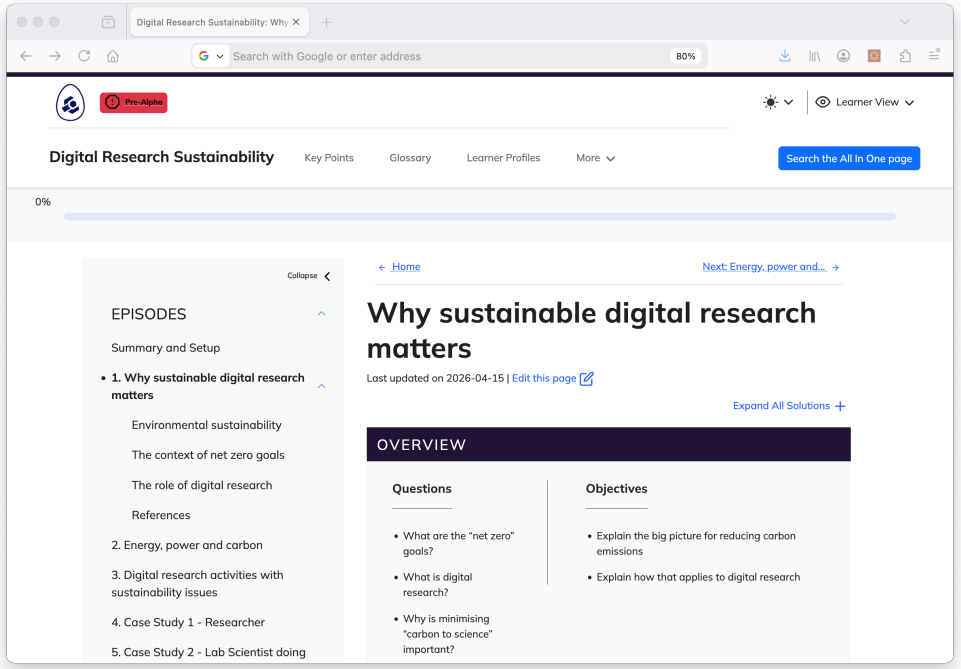

The 3 pillars of a sustainable research computing culture

Education

- Development of a **"Digital Research Sustainability"** course
  - Initially motivated by our Green DiSC certification activities
  - Development led by Imperial's RSE team - supported by Research England research culture funding via Imperial
  - Will be made available under an open licence for others to use/adapt
- A combination of general research sustainability background/skills and case studies
  - Help researchers to understand how to estimate different aspects of emissions in a number of scenarios

The 3 pillars of a sustainable research computing culture

Education

- **"Digital Research Sustainability"** course – target audience

  - Not just researchers

  - Technical staff – digital Research Technical Professionals (dRTPs)

  - Professional staff interested in understanding environment impact of computing

- Delivery plans

  - Early stage pilots from May 2026, plans for ongoing regular delivery

  - Exploring delivery partnerships with e.g Imperial's Early Career Researcher Institute (ECRI) and Research Computing Service

  - Looking at potential for core material to support any possible future staff skills training in sustainability

The 3 pillars of a sustainable research computing culture

Advocacy

We can build a community and train people – how do we bring this together to support meaningful change?

- **Engagement**
  - Highlight the benefits of sustainable, responsible research computing use
- **Easy wins**
  - Straightforward shifts in behaviour individuals can make – everyone can contribute
- **Incentives**
  - **Prizes** – e.g. link in with institutional awards schemes
  - **Competitions** – e.g. improve efficiency of existing open source tools
  - Such incentives also underpin advocacy – raise awareness, amplify message

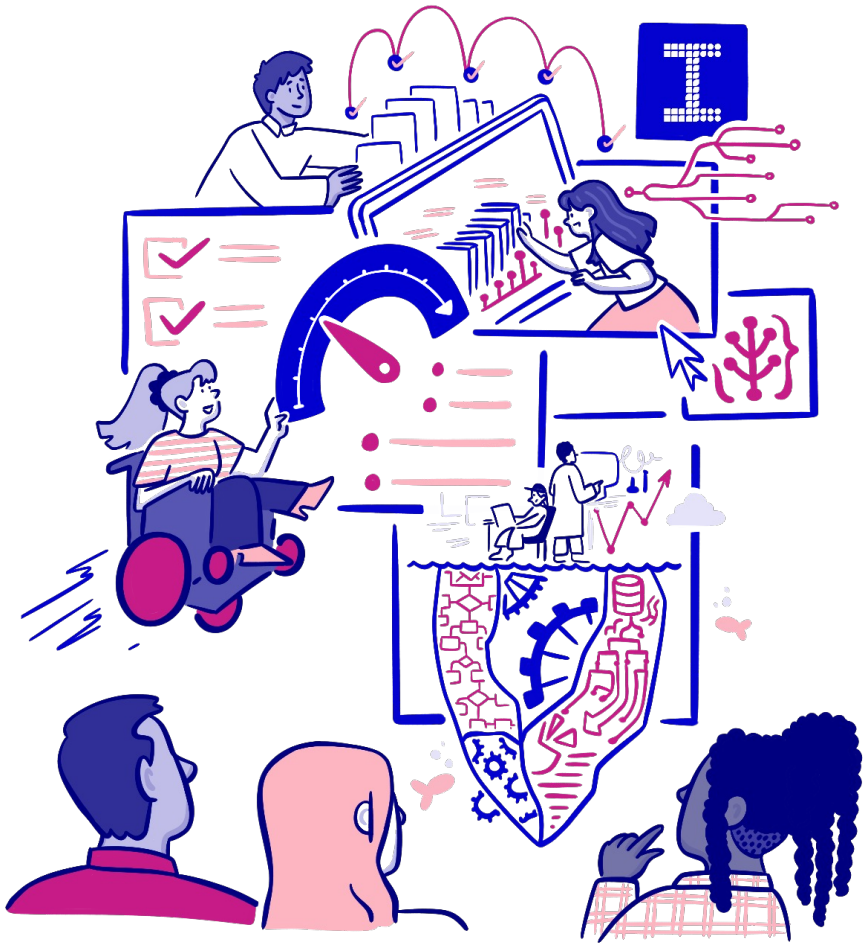

The 3 pillars of a sustainable research computing culture

Advocacy

Engagement

- Build support from central IT leadership + institutional leaders
- Consider finding an institutional “champion”
- Lower electricity costs – contribute to sustainability goals
- Deliver talks to research departments – build on community
- Sustainable research computing advisory group/committee?

- **All part of getting the message out to the institutional community**

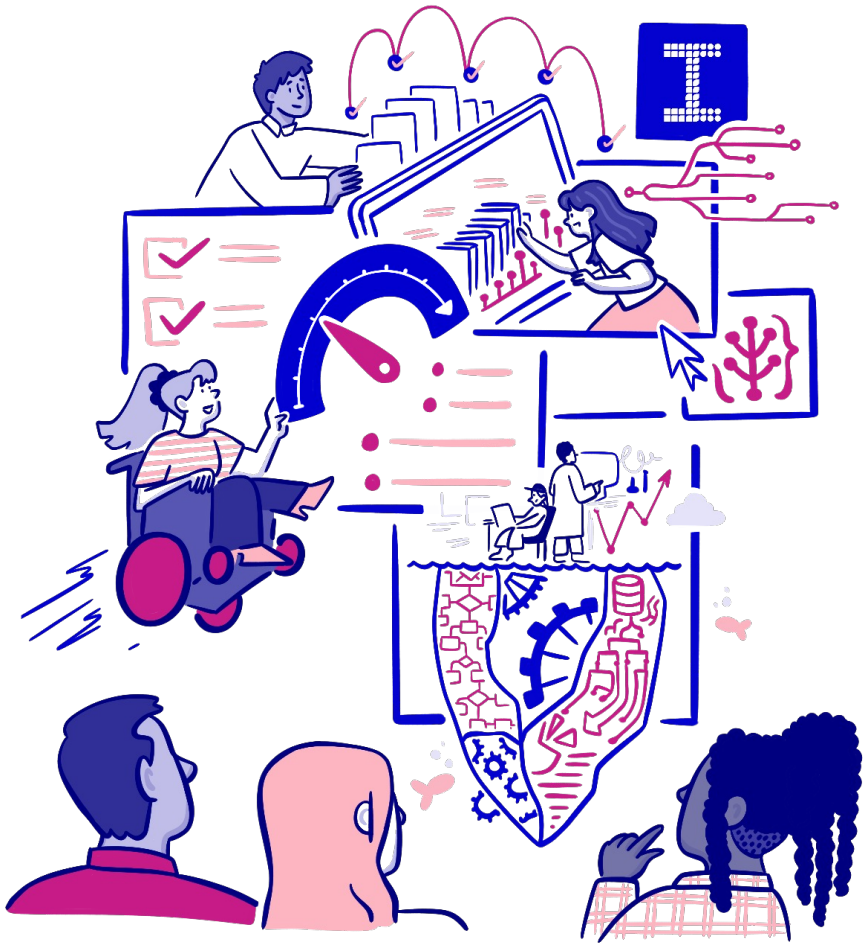

The 3 pillars of a sustainable research computing culture

Advocacy

Easy wins

- Everyone has a stake in supporting sustainable computing use
  - Researchers: use of large-scale HPC infrastructure
  - Professional services staff: SaaS tools, LLMs, data storage
- Offer straightforward, actionable guidance applicable to all
  - E.g. Imperial's **"7 steps to ICT sustainability"**
- Looking at large-format artwork to inspire / raise awareness

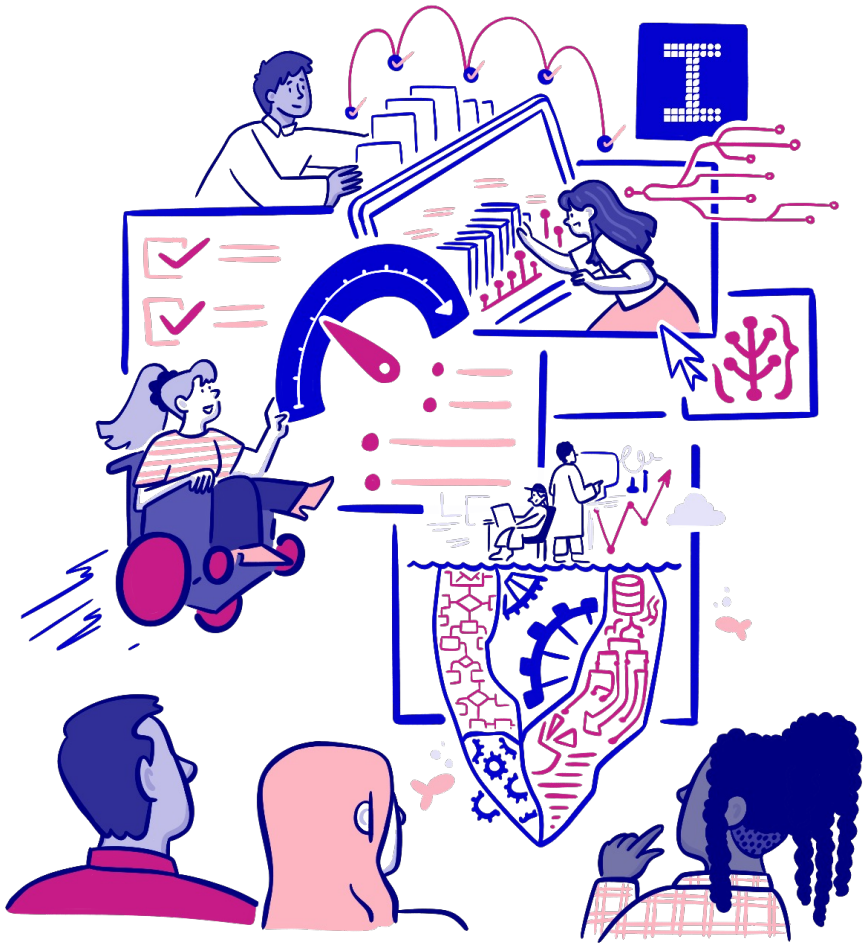

Final thoughts

Sustainable research computing use for developers

- Researchers / RSEs writing code have a particular challenge
  - Complex code, models, datasets
  - Profiling and optimising code is an advanced topic
  - Energy cost to debugging/testing without necessary skills
- Increasing introductory training available
  - E.g. the “Profiling & Optimisation (Python)[8]” course developed by Robert Chisholm (University of Sheffield)
- Increasing opportunities to build on/learn from existing work
- Open source scientific codes/AI models increasingly being optimised with a sustainability/efficiency focus.

[8] **R. Chisholm and J. Migenda. Performance Profiling & Optimisation (Python):** https://carpentries-incubator.github.io/pando-python/ DOI: https://doi.org/10.5281/zenodo.16902755

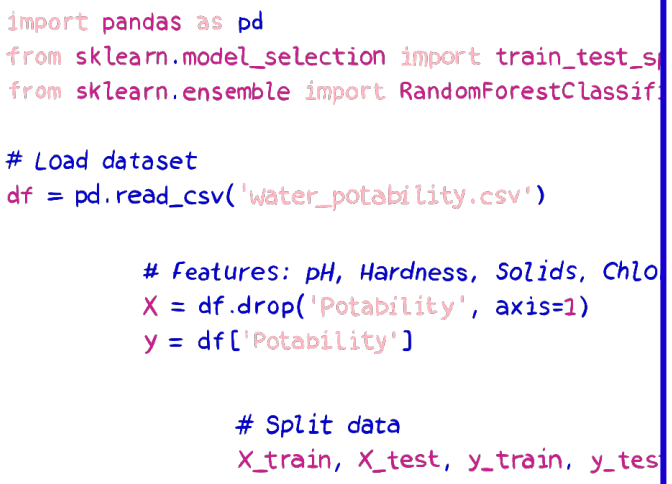

Final thoughts

Bringing it all together

A sustainable institutional research computing culture

- Sustainability is a key element of modern institutional strategies
- Research computing/IT (especially HPC) is a key contributor to an organisation's emissions
- More efficient use isn't possible without the support of staff and students
- Build an environment with an embedded culture that includes thinking about environmental/impact of research computing

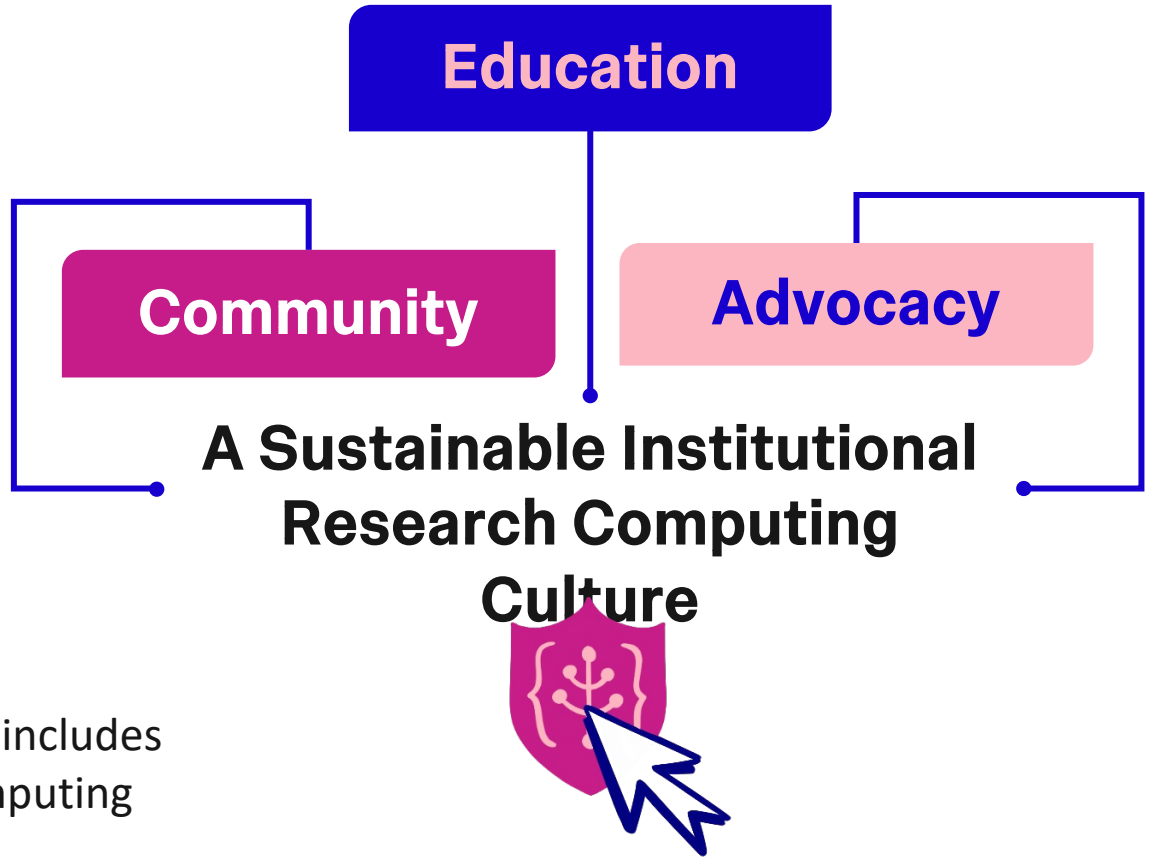

Final thoughts

Conclusions

Rapid, ongoing shift to digital research + AI/LLMs

- Need to think much more about the environmental cost of our computing use

Needs a coordinated institution-wide approach

- Build a culture that includes thinking about the impact of computing use – not just researchers – everyone has a stake and can contribute

Support this by building a sustainable institutional research computing culture

- **Community:** Engage researchers and technical staff, develop a community
- **Education:** Upskill our community via a digital research sustainability course
- **Advocacy:** Work with key stakeholders to raise awareness and build institutional support

Thank you
Questions?

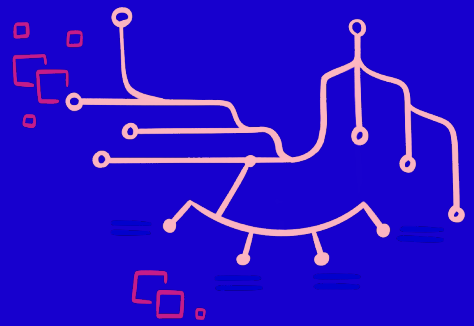

Jeremy Cohen:
jeremy.cohen@imperial.ac.uk

Hannah Scott:
h.scott@imperial.ac.uk

The project team would like to thank **Imperial College London and Research England** for funding this work using **Research England research culture funding** via Imperial. **JC** acknowledges additional support from **UKRI-EPSRC** under the **STEP-UP** project **(EP/Y530608/1). At** acknowledges support for input on this project from NetDRIVE (https://uknetdrive.org). Illustrations by **Rae Goddard**

IMPERIAL