




Science and
Technology
Facilities Council

Scientific Computing

Towards a Federated Accounting Framework for Sustainable Research Computing: A Landscape Review of the UK's Digital Research Infrastructure

*Or: a tour of what we measure, what we think we measure,
and what we quietly ignore.*

Dr.rer.nat. Deniza Chekrygina
Federating Services Group
Scientific Computing, STFC UKRI



iris

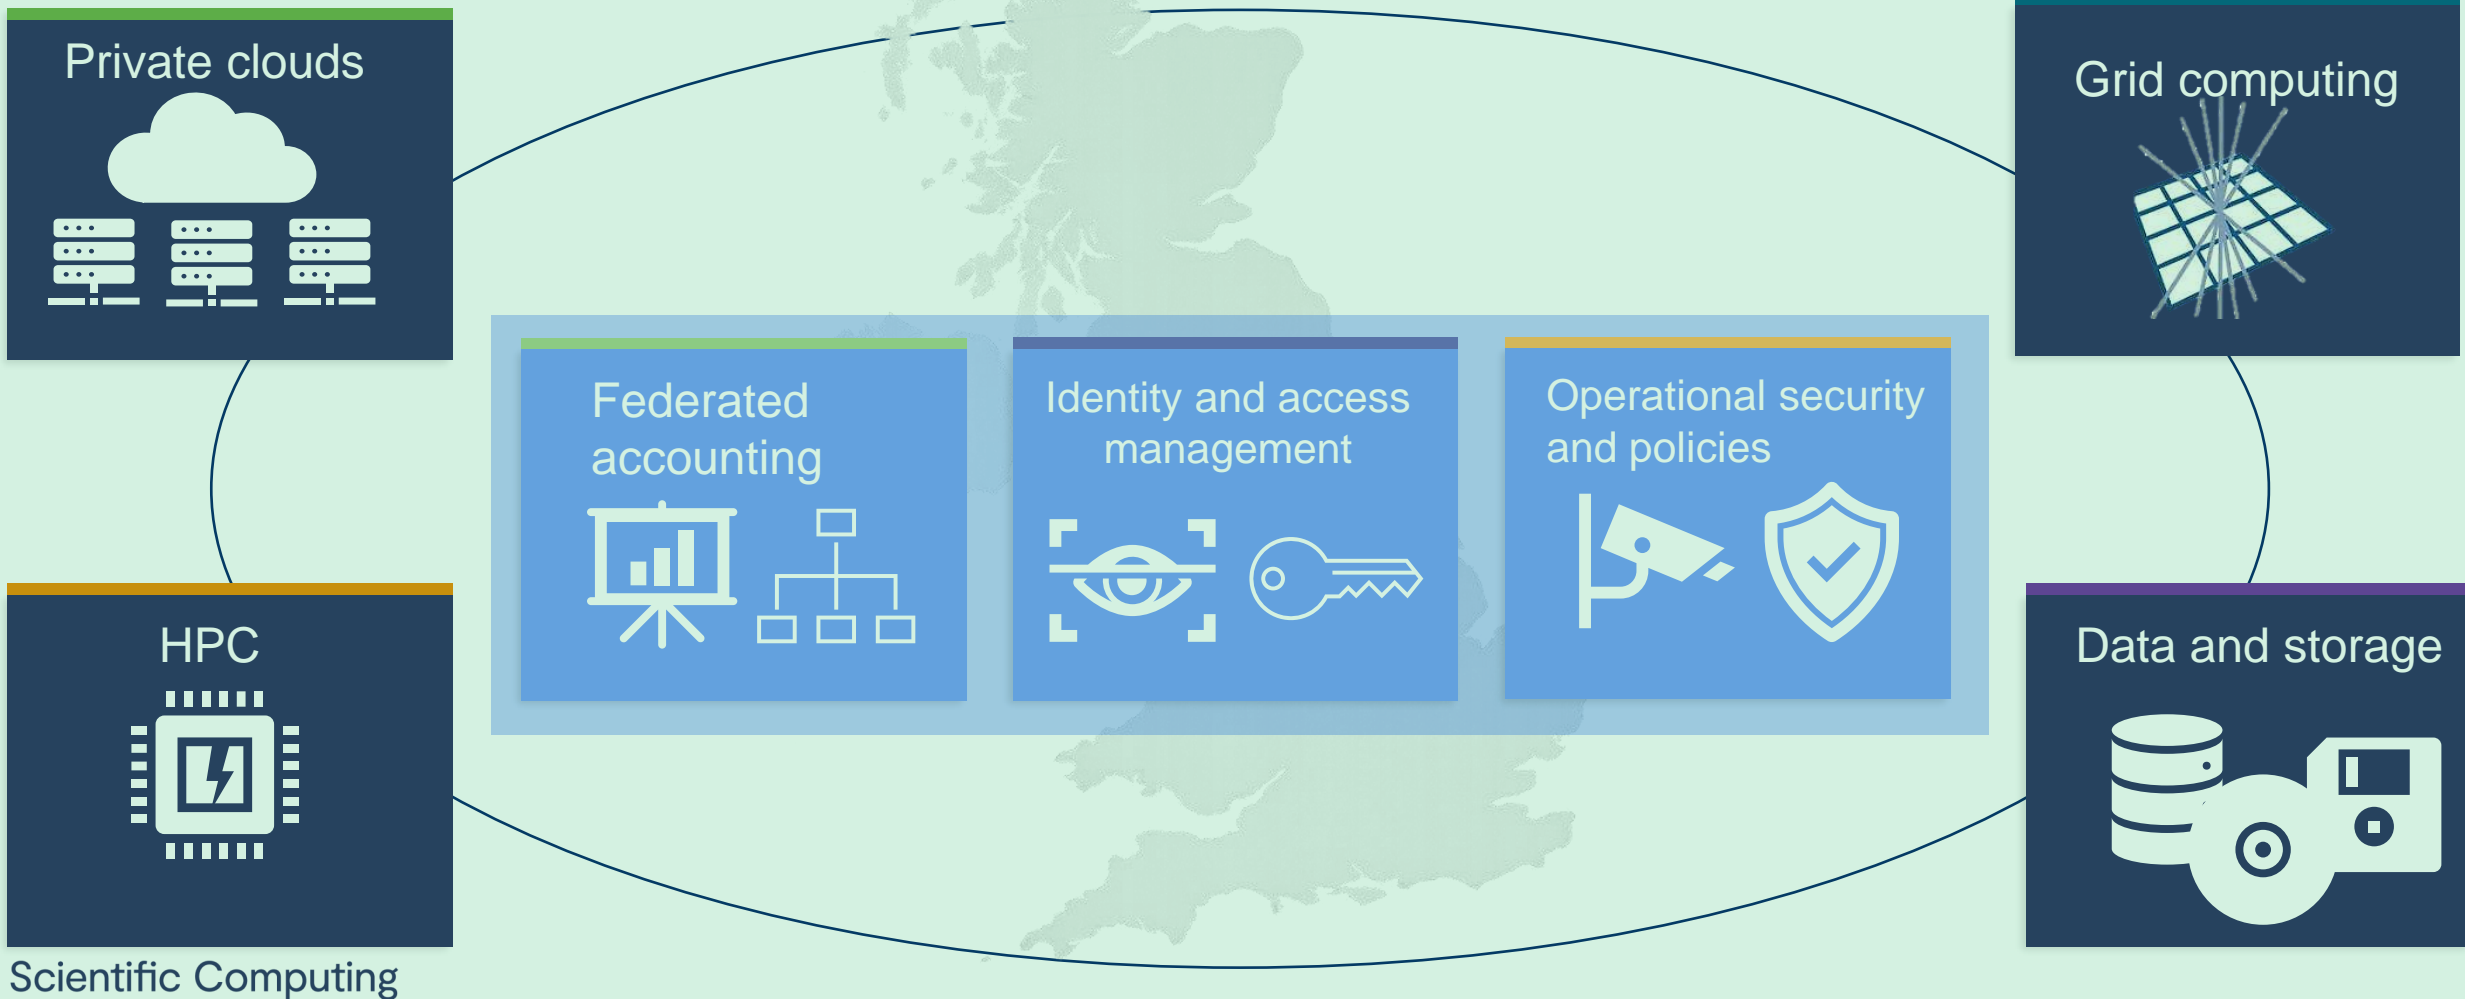
Introduction

The world-class research infrastructure was built across the country. Then, each site quietly got on with tracking it in their own way. And never quite got around to comparing notes.

Meanwhile, 140 countries pledged net-zero. The International Telecommunication Union called for a 45% cut in technology sector emissions by 2030. UKRI committed to Net Zero by 2040. All of these require one thing first: energy and carbon data. Which was never consistently collected.

So somebody had to go and ask the awkward questions. That somebody was me.


IRIS is a **cooperative community** creating digital research infrastructure to support STFC science, facilities, and beyond...



Private clouds



Grid computing



Federated accounting




Identity and access management




Operational security and policies



Data and storage



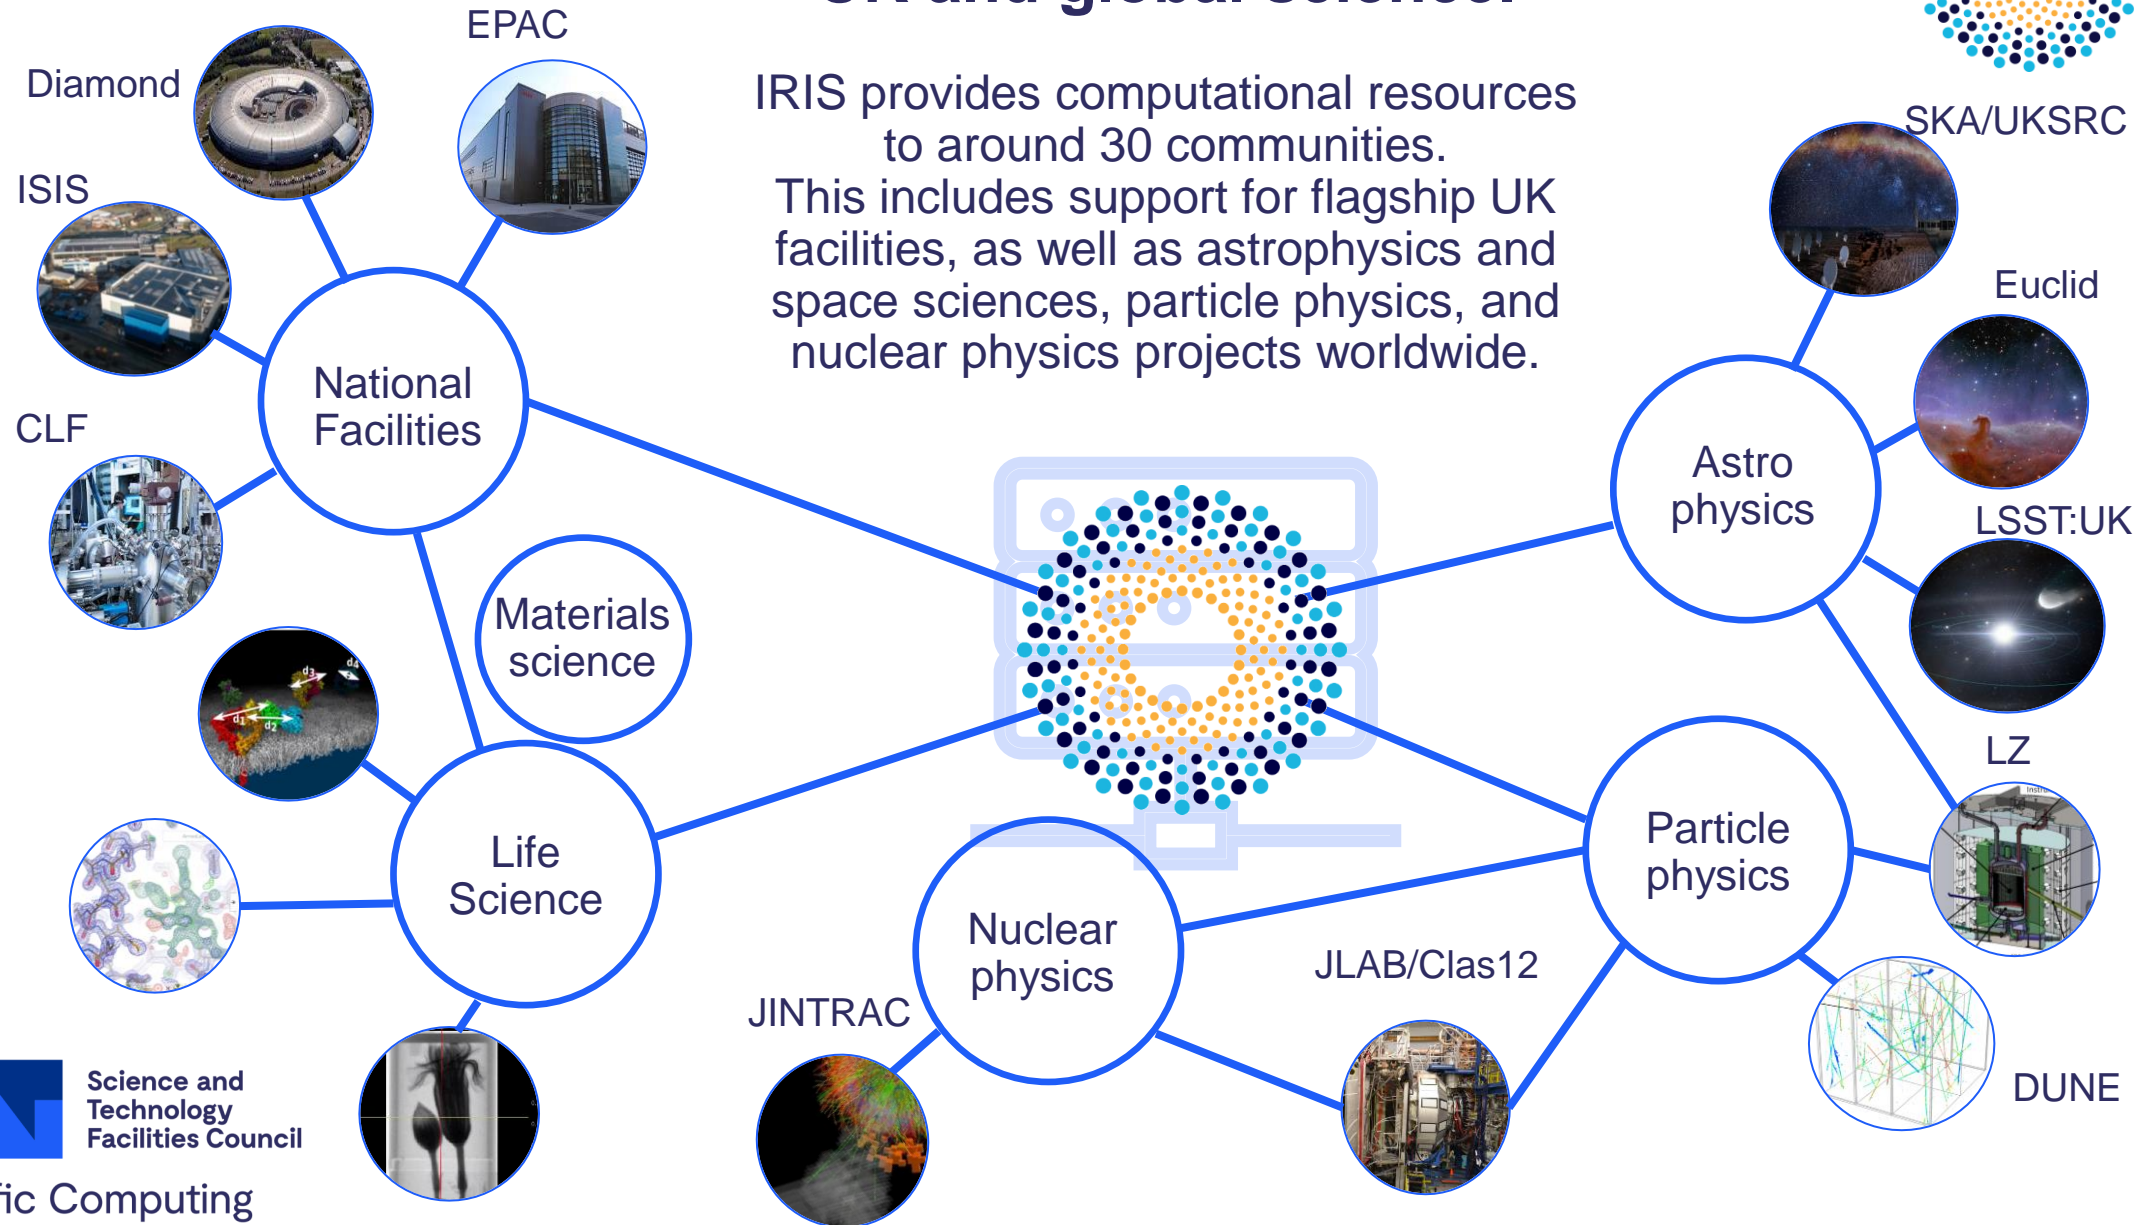
HPC



Scientific Computing

IRIS has a broad impact on UK and global science.

IRIS provides computational resources to around 30 communities. This includes support for flagship UK facilities, as well as astrophysics and space sciences, particle physics, and nuclear physics projects worldwide.



Back to the landscape review of the hardware use accounting...

The summer of landscaping

Nothing says "summer" like emailing 80+ people about "hardware use accounting" during their annual leave.

80+

people
contacted

many of whom were
on a beach at the time

38

in-depth
interviews

Interrogations
about the community and their
requirements

29

questionnaire
responses

13 questions covering
everything from quota enforcement
to carbon (I was made to cut 33
more questions. I'm still not over it.)

Those who responded

National HPC

ARCHER2 / EPCC · DiRAC
(Durham,
Leicester, EPCC) · Baskerville ·
Supercomputing Wales ·
NI-HPC / Kelvin2 · N8/Bede

Grid & HTC & AI compute

GridPP · STFC / APEL · IRIS ·
Isambard / BriCS

Cloud & Data Platforms

JASMIN ·
Hartree Centre ·
EIDF · DAFNI

Trusted Research Environments

DARE UK · EPCC Safe Haven ·
HIC Tayside ·
East of England SDE ·
Yorks & Humber SDE ·
East Midlands SDE · ONS

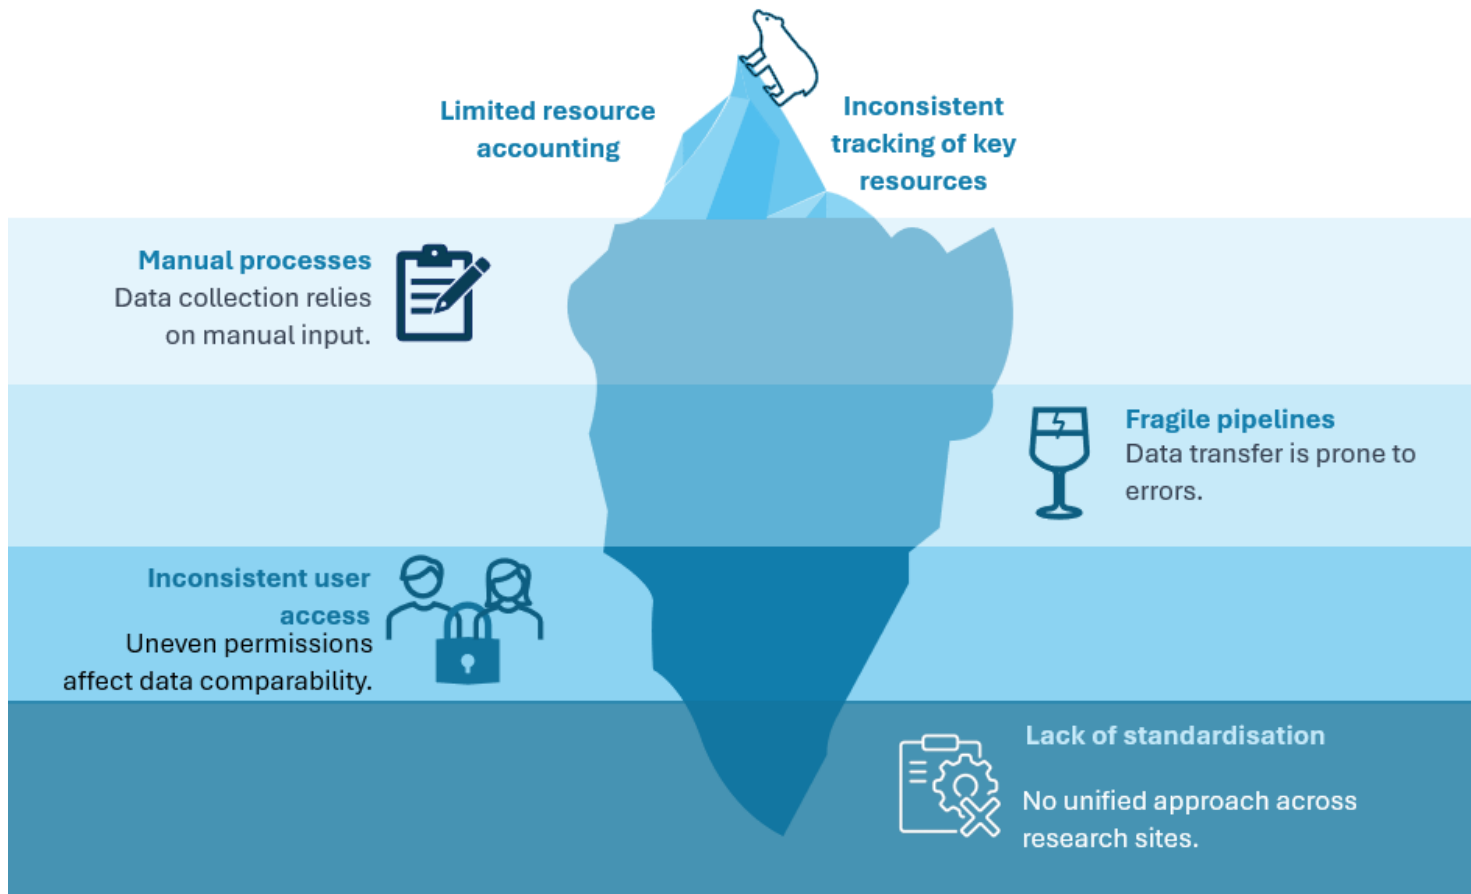
Research Communities

Alan Turing Institute ·
UKCP / HEC ·
BIOFAIR · NBI ·
Archaeology Data Service ·
UK Data Service ·
AHRC / RICHeS · CoSTAR

Governance & International

UKRI DRI · JISC ·
EOSC · ICHEC ·
Freiburg / AUDITOR

What was found



Maturity of tracking across the DRI

CPU time

Established

Well tracked. The one thing we nailed.

GPU usage

Emerging

Some sites track it. Often summaries, sometimes just allocations.

Storage

Emerging

Usually manual. Little granularity.

Energy

Early

A handful of sites. Rarely per-job.

Carbon

Nascent

Somewhere between 'aspirational' and 'imaginary'. Nobody expects it.

Behind the 'nascent' label

Question3: Do you currently see any energy or carbon-related data linked to your compute usage, and would this information affect your decisions if it were available?

29

responded to questions in one form or another

8

clearly confirmed collecting energy data

1

provides **regularly** energy and carbon data per job to users

Granularity unknown



Of the 8 who confirmed collecting energy data, several could not specify whether it was per job or per rack. The precision of what exists is itself unclear.

Trust in the data is questioned



Several respondents raised doubts about whether the system for energy readings is reliable. And whether it is possible to have a reliable attribution.

Collecting ≠ sharing



Some collect energy data internally for reporting. Almost none of it reaches users. Visibility at the point of compute use is nearly absent.

At least half avoided the second part of the questions.

What 'nascent' is blocking

5 respondents said accessible energy and carbon data would change how they run their sites or how users behave.

User behaviour

Some researchers said they would configure jobs differently : choosing lower-carbon time windows, adjusting resource requests, or selecting sites, if they could see the environmental cost.

Charging models

Sites said energy and carbon data could inform how allocations are priced, moving toward sustainability-linked pricing rather than compute-time alone.

Site policy

Service design, procurement decisions, and workload scheduling would all be informed by better data. Some sites are already planning for this but they're waiting for guidance and the metrics.

The main barrier

Shared infrastructure

Hard to isolate one job's energy draw from shared hardware

Attribution limits

Scheduler data may not reflect real consumption accurately

No carbon conversion

Energy data exists in places , while carbon calculations rarely do

No shared standard

Every site measuring differently means nothing compares

What everyone agreed on (for once)

Remarkably, from ARCHER2 Supercomputer to the Archaeology Data Service, these four themes came up every time.



Use-Centric/Role based design

Easy access to data , beginner-friendly interfaces, alerts, and role-based dashboards.

Interfaces a PI can use without filing a support ticket and waiting three weeks.



Accountability & Transparency

Tracking access, linking usage to funding, and standardised reporting.

Clear attribution to projects and grants. Consistent reporting. No more 'it depends which site you ask.'



Sustainability

Tracking carbon and energy usage, and efficiency of compute resources.

Energy and carbon metrics linked to jobs. Not just because policy says so, people genuinely want to know.



Interoperability

Unified platforms, shared accounting models, and UK-wide standards.

And make it connected - Stop every site reinventing the same wheel.

The Goldilocks question

Too little? Too much? Or just right?

Do nothing (much)



Keep calm and carry on.
Minor dashboard tweaks.
Low cost. Gaps persist.
Policy alignment: nil.

Too cold

Encourage alignment



Sites adopt shared tools
voluntarily.
Hopeful. Gradual.
Falls apart when the
enthusiasm fades.

Getting warmer

Federate nationally



Shared schemas, APIs,
dashboards. Sites keep
their own tools but meet
a national baseline.
Delivered in phases. Will
provide necessary balance.

Just right

Centralise everything



One centrally operated
platform replaces all local
systems. Maximum
consistency.
Maximum resistance.
Good luck with that.

Too hot

How the framework would work (hopefully)

Enhanced system connectivity and security - *Federated login without five separate accounts*



Collaborative decision-making process - *A governance model someone actually designed*

Guidelines ensuring data comparability - *Apples to apples, not SLURM to spreadsheet*

Adaptability to local research needs - *Room for local tools (yes, even the questionable ones)*

Governed and funded

UKRI-led Advisory Board for vision and standards · Working Group for the actual work

Estimated 3–4 FTE over 2–3 years, assuming we build on what exists, not start from a blank page

What needs to happen

Start

- Advisory Board and Working Group
- National schema for hardware accounting
- Prototype central federated portal
- Pilots at 3–5 diverse sites
- Shared sustainability and computing KPI framework

Stop

- Bespoke systems with no interoperability
- Ignoring GPU/storage gaps in reporting
- Ignoring carbon and energy accounting
- Using non-standard formats
- Hoping it'll sort itself out. It won't. It never does.

Change

- Accounting → foundational infrastructure
- Spreadsheets → API-first dashboards
- CPU-only → full-stack metrics
- One-size views → role-based dashboards

Fund

- Shared collectors, dashboards, schema
- Onboarding packs and training for sites
- KPI tracking infrastructure
- Usability testing (yes, really, with actual users)

PEAR Accounting Service

The **P**ortal for **E**nhanced **A**ccounting **R**ecords (PEAR) will serve as a production quality replacement to the APEL system currently used by EGI and CERN WLCG. PEAR will:

- Cover resource usage data for all main Hardware elements (CPU, GPU, storage, network)
- Be extendable with environmental related metrics (PUE, CO2 intensity, Water measures, Benchmarks, etc)
- Be flexible and modular to support a wide range of communities.

The work to deliver a new accounting system is funded by the EU through the 3-year long **ENSURE*** project, which starts in September 2026.

As well as producing a new platform, this project will develop carbon and sustainability benchmark tooling, to enable consistent reporting across sites, iterating on work done by the Universities of Glasgow, Manchester, along with other international partners.

The design of PEAR will be strongly shaped by the recommendations set out in the UK Landscape Review.

PEAR Partners

- STFC & CERN are leading on the repository and portal/dashboards side of development
- Universities of Glasgow and Manchester will be working on defining and capturing sustainability metrics
- University of Freiburg will lead on the client-side development, delivered through improvements and enhancements to the AUDITOR tool.
- EGI will be coordinating the larger ENSURE project, which is made up of 20 partners.



universität freiburg



Scientific Computing



Scientific Computing

What needs to happen

Start

- Advisory Board and Working Group
- National schema for hardware accounting
- Prototype central federated portal
- Pilots at 3–5 diverse sites
- Shared sustainability and computing KPI framework

Stop

- Bespoke systems with no interoperability
- Ignoring GPU/storage gaps in reporting
- Ignoring carbon and energy accounting
- Using non-standard formats
- Hoping it'll sort itself out. It won't. It never does.

Change

- Accounting → foundational infrastructure
- Spreadsheets → API-first dashboards
- CPU-only → full-stack metrics
- One-size views → role-based dashboards

Fund

- Shared collectors, dashboards, schema
- Onboarding packs and training for sites
- KPI tracking infrastructure
- Usability testing (yes, really, with actual users)

