Federated

IRIS Digital Research Infrastructure: an update

J. Hays, IRIS Science Director

IRIS Collaboration Meeting

December 2024

iris eInfrastructure for Research and Innovation for STFC

IRIS is a cooperative community bringing together STFC computing interests

Formed bottom up by science communities and compute providers

Works closely with STFC but run by the community

IRIS Activities

Build infrastructure

Build community

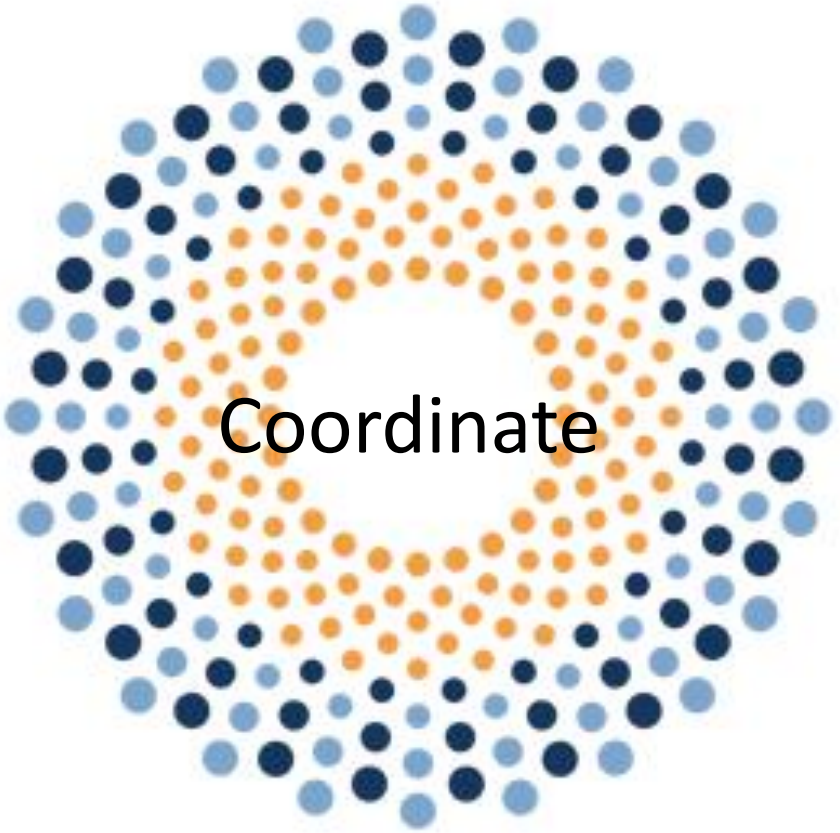

Invest in computing hardware

Support training

Support digital asset creation

Making the case for investment

IRIS Structure

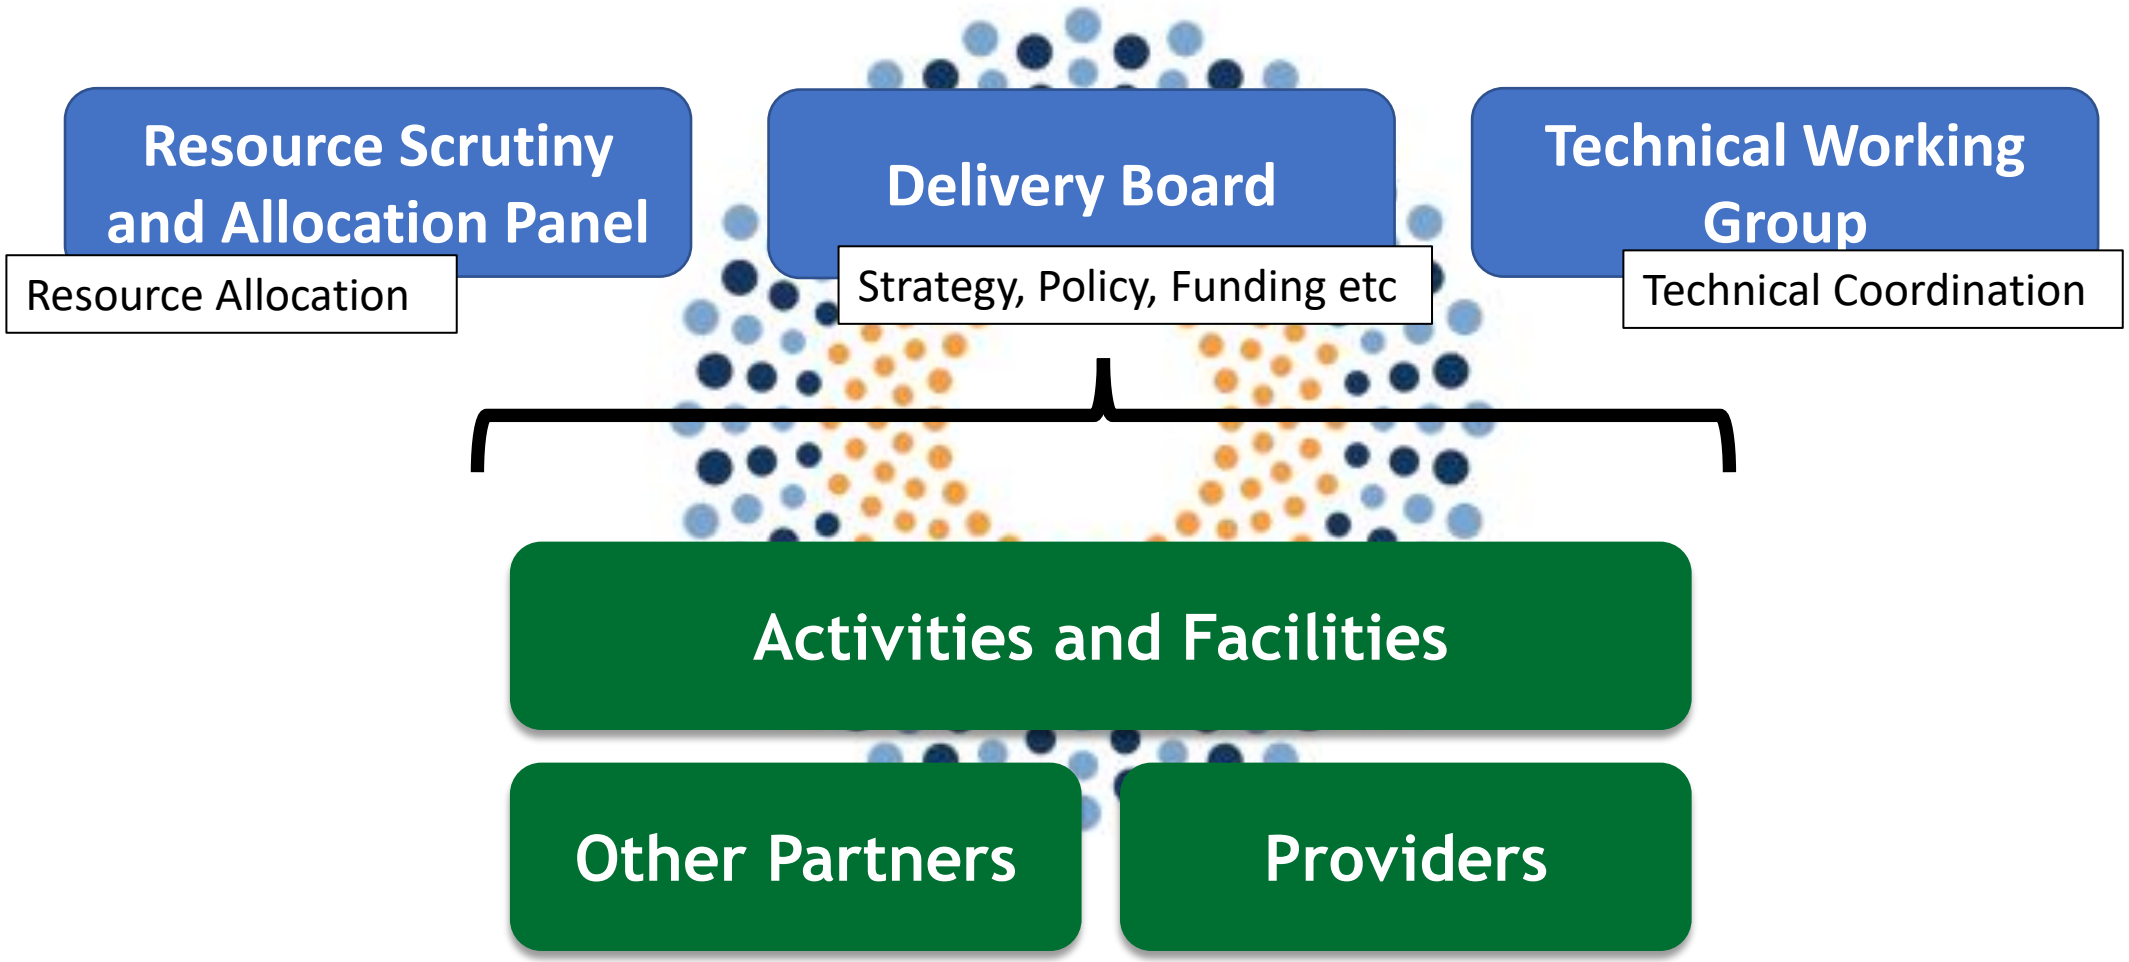

Funding

- 2017:
  - £1.5m Capital : £1m hardware, 0.5m digital assets
- 2018:
  - £16m (4M / year over 4 years, 11M Hardware, 5M Digital Assets for National Facilities
- 2020:
  - Capital injections ~ £6m : 3M capital for IRIS, 2M for LSST (Rubin Observatory), 1M for SKA
- 2021:
  - UKRI Digital Research Infrastructure funding ~ £2m
- 2022:
  - STFC Capital funding £2.4m – anticipated recurrent over 3 years
  - UKRI DRI funding – £3.5m additional funding in 22/23 + 23/24
  - Additional capital supporting Astronomy programme £2.3m
- 2023:
  - STFC Capital funding £2.4m
  - UKRI DRI Phase 1b from 22/23 award - £400k
  - UKRI DRI Phase 2 - £1.4m (£400k 23/24, £1m 24/25)
- 2024:
  - STFC Capital funding £2.4m
  - UKRI DRI Phase 2 - £1.2m

IRIS and its Community

- Technical Working Group
  - Weekly meetings
    - Technical discussions – presentations, showcases etc
    - Operations reports
  - Regular ad-hoc workshops as and when needed
    - Wide range of topics from identity management, cloud computing, GPU usage, machine learning etc
- Targeted at the IRIS community but open to all

Digital Asset Creation

Digital Assets

Essentially software (or similar) products (capitalizable) – with a defined end point or delivery of the asset

Additional track of funding “Other DA” programme available to full IRIS Community

Often infrastructure related but can also be closer to science

UKRI DRI funding also allows some non-capital projects
- can be used for R&D as well as asset production

Digital Asset Creation

Digital Assets

Identity management - IRIS-IAM
- Working towards single-sign on
- Closely aligned with LSST and SKA work
- Integrated with the DiRAC Federation Project

Radio astronomy on-boarding
- eMERLIN, ALMA, and others
- Documentation, tutorials, community engagement
- Evaluation for IRIS integration

Algorithm Development
- Efficient random numbers on FPGAs
- Support for FAST-HEP toolkit

Data management
- RUCIO (LHC) from development for non LHC usage – multi-VO support, token based authentication support, etc (LSST, DUNE, SKA)
- High performance data transfers for SKA – using FTS and RUCIO (leveraging hardware investment by ExCALIBUR)

Core infrastructure
- Security policy
- Information Security Management – supporting STFC SOC
- Accounting development – building on work at Tier 1
- VMDIRAC – workflow management building on DIRAC (LHCb) for non-LHC
- Tools and services portfolio

What IRIS is not…

IRIS is not a turn-key computing solution

IRIS currently has no direct resource funding for user support – some community support as “best effort”

IRIS can provide physical resources for groups who can make use of it –

Activities may need their own:

- Software frameworks
- Support staff

IRIS might not have the kind of hardware needed for every activity

Key Challenges : Funding

- Funding
  - Continues to be a key challenge
  - Received all anticipated funding this year (~£3.6m)
    - While an excellent outcome in a tough year - this is insufficient to supply all requirements for 2025/26 based on RSAP forward look
  - Approval to spend came very late this year
  - No solid information on funding for next year
  - Funding model very much based on flat-cash + extras
    - RSAP provides good quality information on future requirements and is fed into scenario planning and provided “up-the-chain”
  - Successful in exploiting pooled hardware provision but need to continue to look for efficiencies

Key Challenges : Planning

- Forward planning
  - Continues to be a key challenge
  - Driven mainly by funding uncertainties
    - Causes extra work in scenario planning
    - Increases delivery risk with respect to March 31st deadline
  - Reduced coherence across STFC
    - Risk of shortfalls when projects are approved without funding for their compute
    - “Right-sizing” of projects essential
  - Scope creep
    - Risk of shortfalls when projects wish to expand scope of work

Key Challenges : Provider Support

- Providers paid in capital grants
  - Supports purchase of hardware
  - Associated MOU outlines hardware provision
  - No explicit payment for electricity
  - No explicit payment for staff support
- Current model:
  - 4 year life-span for IRIS allocation purposes
    - assuming typical 5 year life-span of equipment
  - 15% additional capital for hardware to be used at providers discretion (within some restrictions)
  - Amounts to around 35% additional funding in-lieu of direct funding
- Key Challenge (apart from overall funding level)
  - Non-equipment funding paid at 80% FEC
    - FEC model of funding doesn't fit well with multi-institute science where the science funding (overhead carrying) is separate from infrastructure funding (capital, non-overhead carrying).

Key Challenges : Tensioning

- Tensioning
  - When resources insufficient some projects will not get their allocations
  - Science is not tensioned by IRIS
    - Science funding decisions made by STFC through its boards and panels
    - No double jeopardy
  - IRIS continues a policy of balancing shortfalls at each facility across all projects
    - While maintaining a pragmatic approach of dialogue with all affected projects to find the best outcomes
  - No mechanism exists to handle large shortfalls
    - This necessarily needs to involve relevant STFC boards/panels etc

Key Challenges : Environmental Sustainability

- Carbon budgets are coming
  - Need to build on previous projects
    - UKRI DRI NetZero Scoping project – IRISCAST, CQUANDRI, etc
  - Need to work together with existing partners
    - SCD, GridPP, DiRAC, etc
  - How do scope-3 (embedded / supply chain) emissions fit into this?
- Need accounting rules, standards, and approaches to be defined
  - Who does the carbon cost get charged to?
  - How to apportion costs across different groups?
  - How to meet UKRI/STFC and HMG requirements?
  - Reporting?
- Need technology solutions to collect data
  - Existing accounting on CPU time is good 0th order proxy but need to do better

Forward look

Governance changes

- SEAG is gone to be replaced by CAB
  - Details are still being worked out
- Stronger engagement with Science Boards
  - Need to create new processes

Technical Review Panel

- Standing agreement with STFC Astronomy to provide technical review for UKSRC and LSST/Rubin projects for capital spending
- Formalising process for this – draft documents currently under consultation
- Plan is to enable this to be used by other projects and facilities to provide technical review and recommendations on request (non-decision making) that may link to CAB, SB, etc.

Forward look

- Digital Assets / DRI Projects
  - Tricky funding situation at present
  - Need more targeted approach to support delivery plan
    - Working groups?
    - Build on work done by partners and associated projects (DiRAC, GridPP, AIRR-FED, etc.)
    - Continue to seed new initiatives
  - Potential focus areas:
    - Federation services
    - Distributed data
    - Multi-site cloud integration
    - Environmental sustainability
    - … many more potential areas

Summary

IRIS is a federated cooperative community bringing together STFC (and beyond) computing interests

Bringing groups together
Building Infrastructure
Supporting Science
Lobbying for the future

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