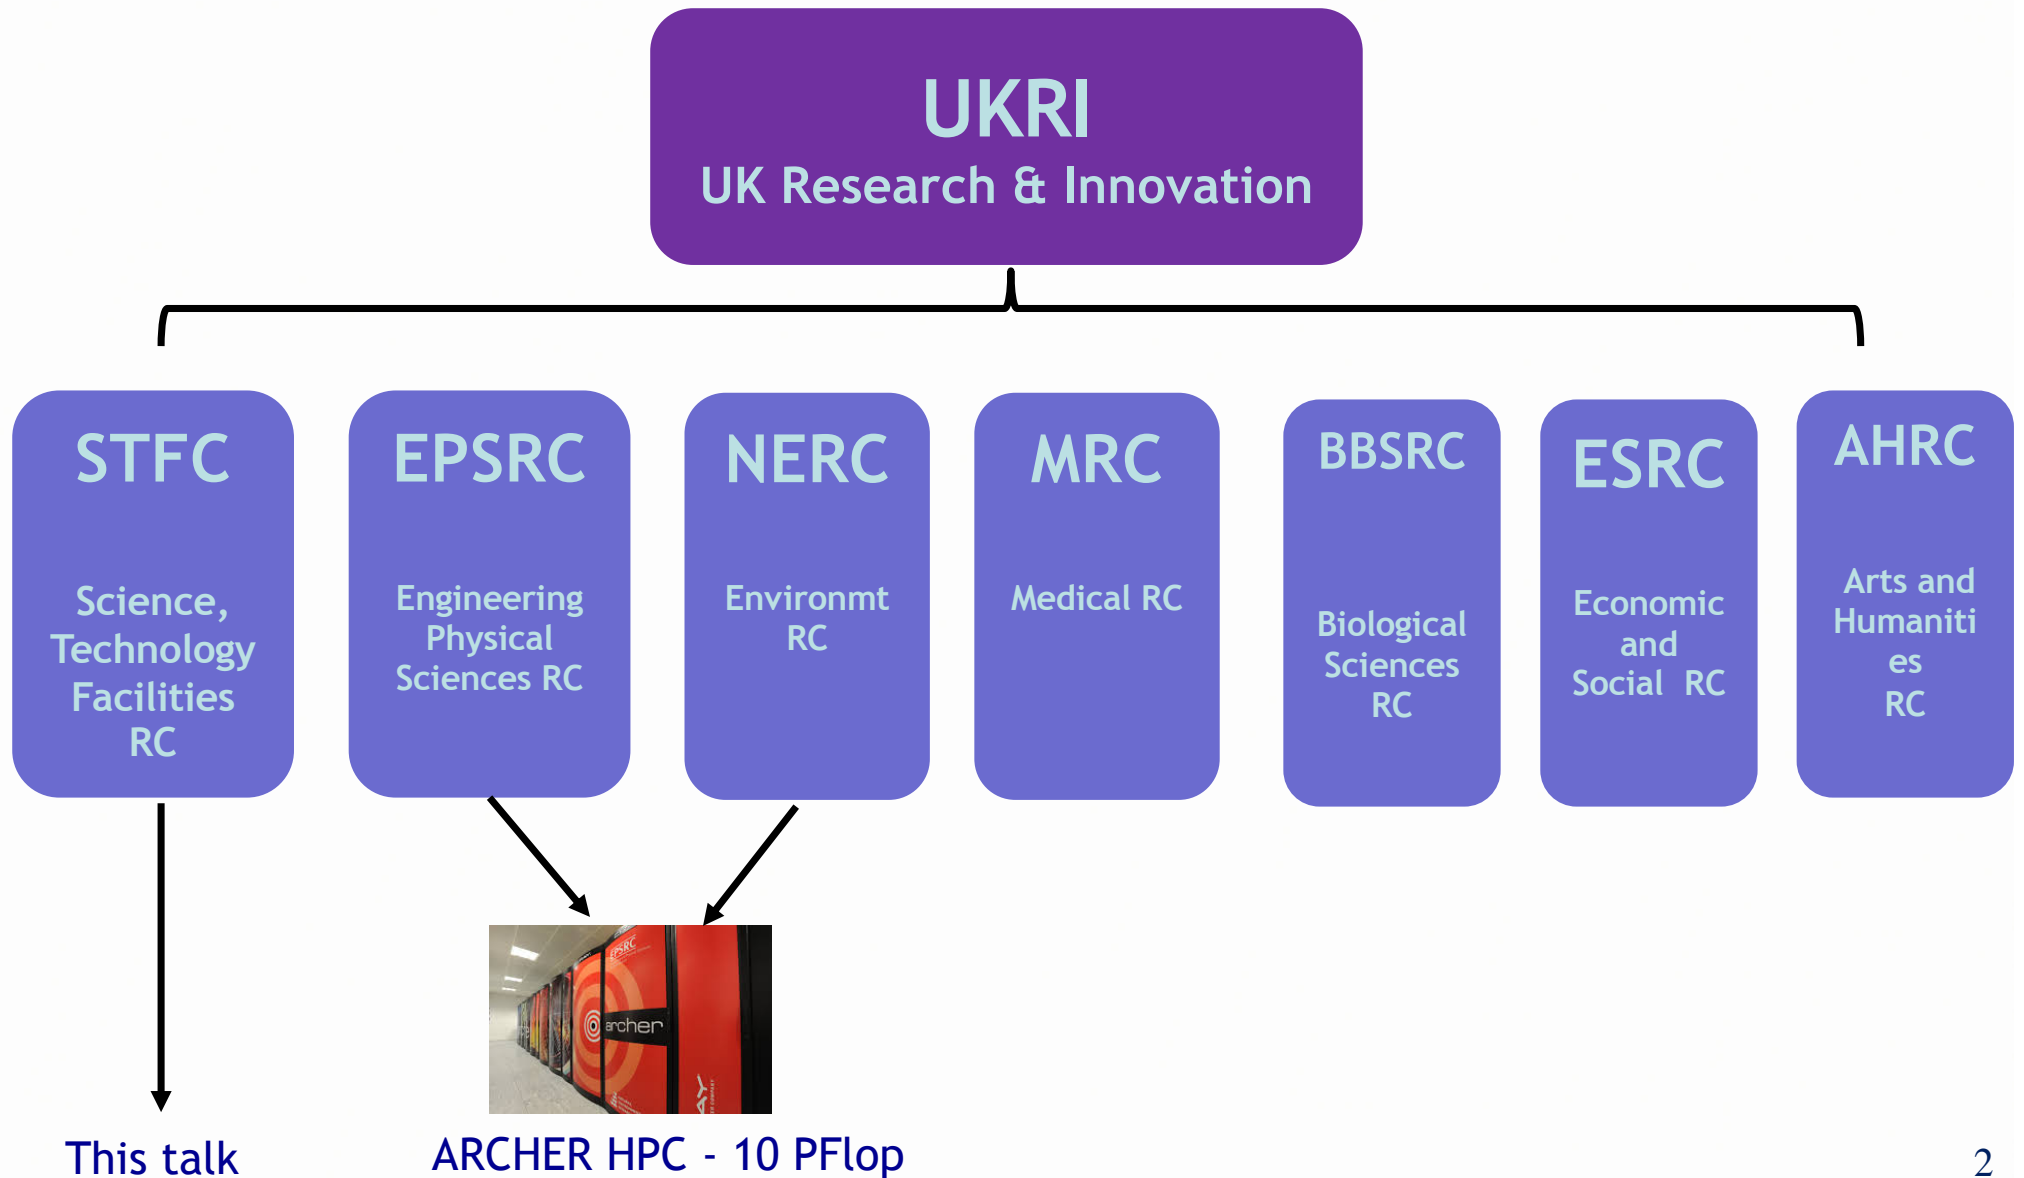


IRIS (UK)

eInfrastructure for Research
and Innovation for STFC

WLCG/HSF Meeting
22-Mar-2019
JLAB/ Virginia

Pete Clarke
University of Edinburgh
On behalf of IRIS Community



STFC

National
Laboratories


Large scale Facilities

(Science)
Programmes





National Facilities:

Diamond Light Source
ISIS (neutrons)
Central Laser Facility
RAL Space



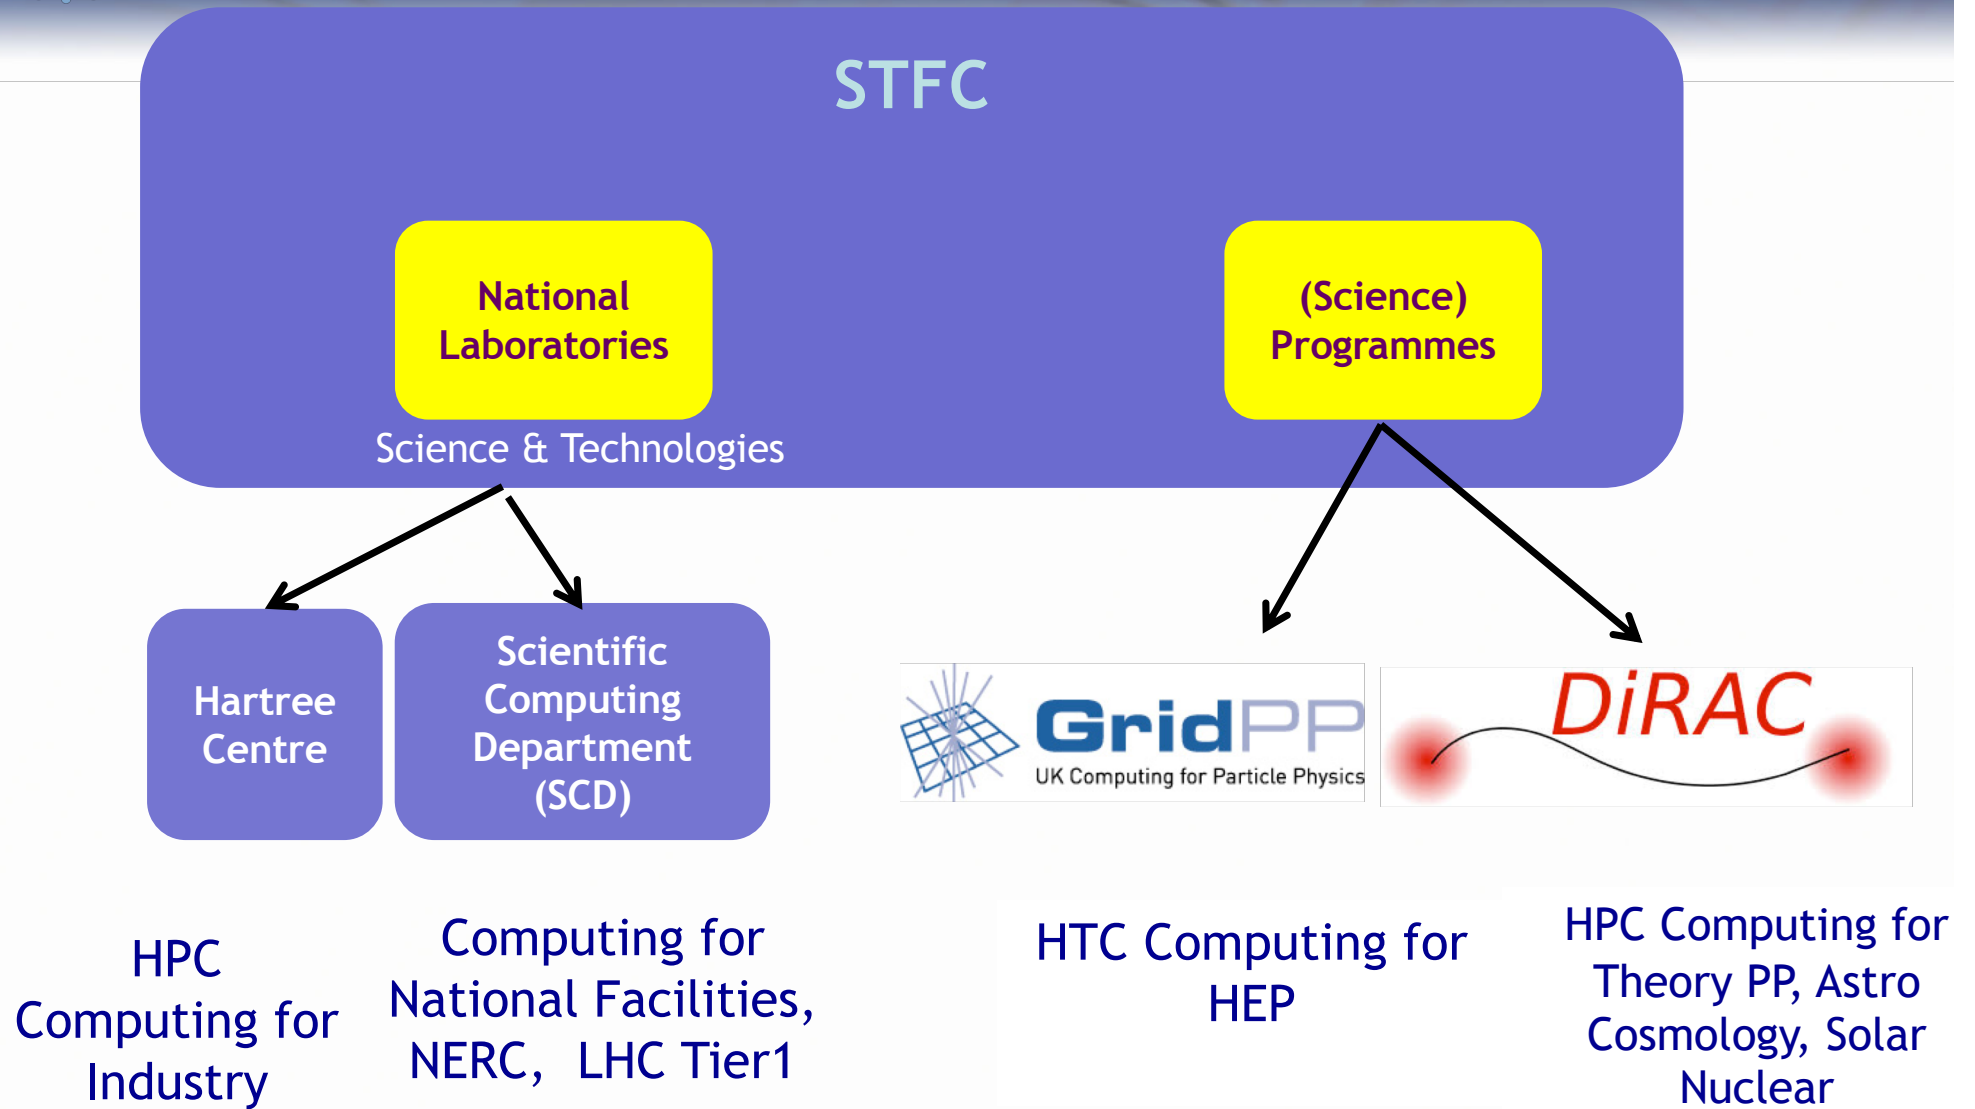
Supports programmes

HEP,
Astronomy,
Astro-Particle,
Nuclear





Computing Provision from STFC





Computing Provision from STFC


STFC

National
Laboratories

(Science)
Programmes

Hartree
Centre

Scientific
Computing
Department
(SCD)



GridPP
UK Computing for Particle Physics

DiRAC

HPC
Computing for
Industry

Computing for
National Facilities,
NERC, LHC Tier1

HTC Computing for
HEP

HPC Computing for
Theory PP, Astro
Cosmology, Solar
Nuclear

IRIS : a mutually consenting coordination body



IRIS: a Community Initiative

- **A Community Initiative to bring STFC computing interests together**
 - Formed bottom up by the science communities and compute providers
- **It is an association of peer interests:**
 - Particle Physics: GridPP: LHC + other HEP
 - Astro: LOFAR, LSST, EUCLID, SKA,
 - Astro-particle: LZ, Advanced-LIGO, CTA
 - STFC Scientific Computing Dept (SCD)
 - Nuclear Physics (growing demand)
 - DiRAC HPC Facility
 - Diamond Light Source
 - ISIS Neutron Source
 - Central Laser Facility
 - Hartree HPC Centre
 - CCFE (Culham Fusion Centre)



[GridPP central to IRIS]

As we are in an HEP meeting.....

- GridPP was of course a major part of starting this
- GridPP de-facto centralises all computing provision for HEP in UK
 - Physical Hardware
 - People to run WLCG services
 - ~ 10% of WLCG
- GridPP has provided resources to “others” since a long time
 - ~ 10% of its resources
- GridPP staff help other communities
 - on best efforts basis







IRIS -what does it do (1): Coordinate

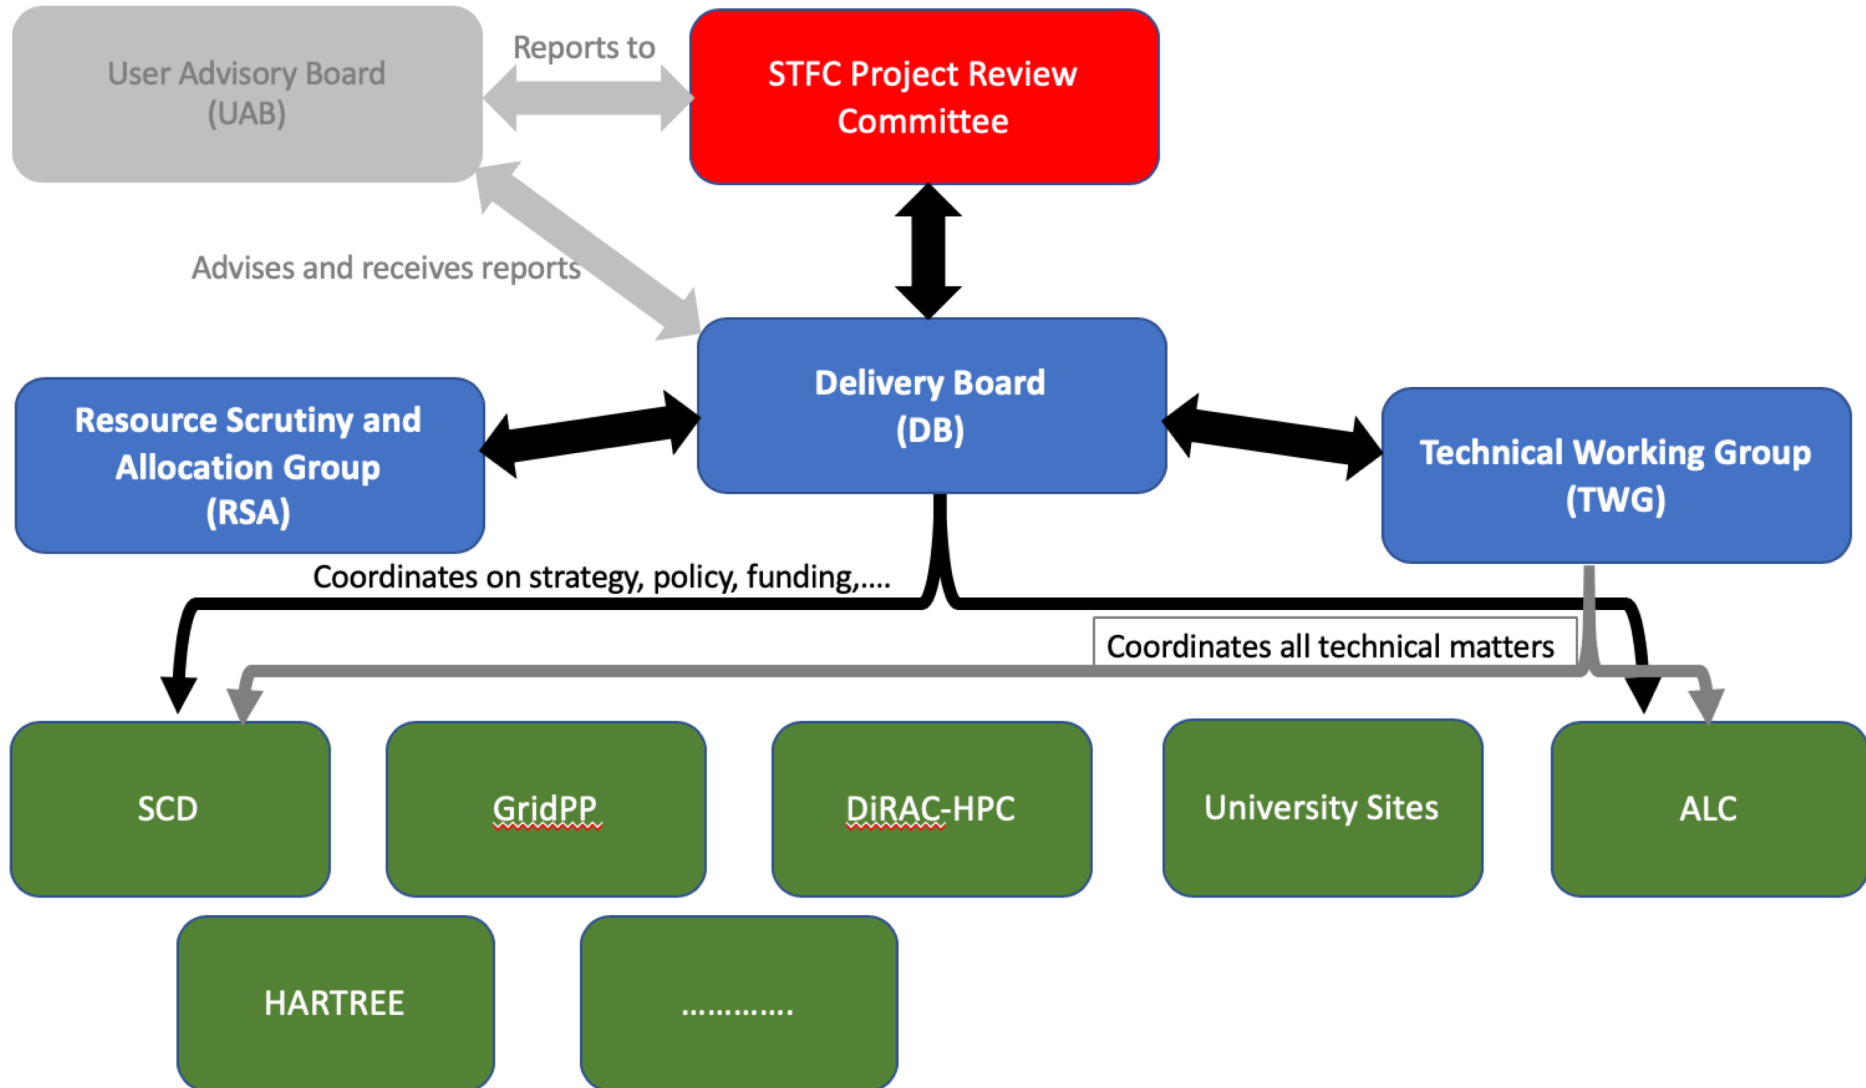


Owned by



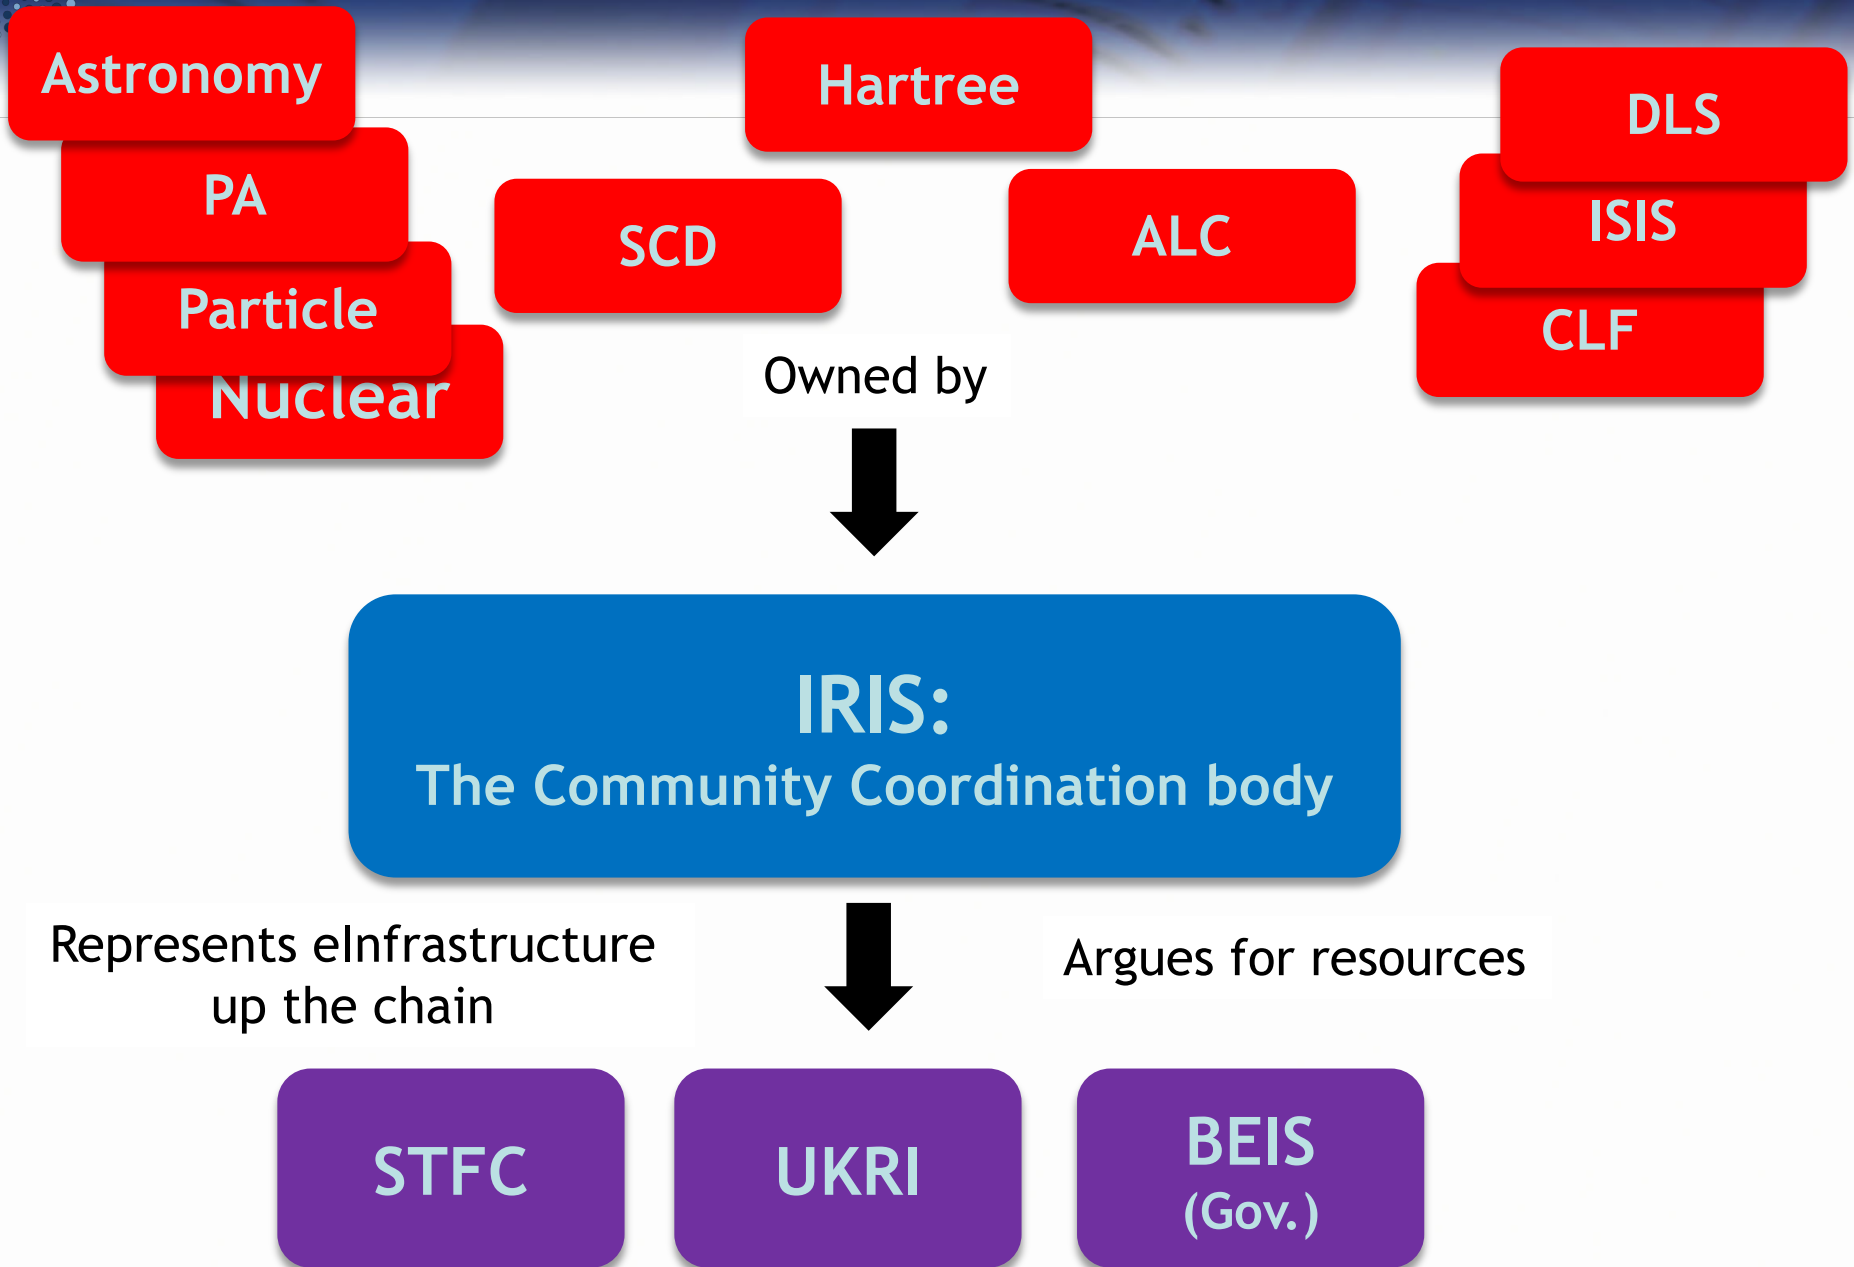
Coordinates Interests & Shares Resources







What does it do (2): argue for resources





What does it do (3): Spends money with actual “providers”

IRIS

Provides ££££ for additional CPU and storage resources via

SCD

GridPP

DiRAC

Runs hardware & supports communities where possible

DLS

ISIS

LZ

LIGO

CTA

CLF

**All
Particle**

SKA

LSST

???

Note: IRIS is not funded consistently - capital only - no staff for support




IRIS : “what has IRIS ever done for us”



Established spirit of consenting cooperation


- This does not come by top down edict
- This comes by establishing trust slowly
- By helping people on their own terms
- Can never be taken for granted - so needs constant work



iris

Obtained significant funds & lobbied for more

- 2017:
 - £1.5M Capital for hardware
- 2018:
 - £16M (4M p.a. for 4 years)
 - 11.5M Capital for Hardware
 - 5.5M for RSE staff (Research Software Engineers) **but NOT for HEP**
- 2018:
 - Case presented to UKRI for long term investment of ~ £20M p.a.
 - Lots of “Roadmaps” written
 - Cloud
 - Data&HTC
 - HPC
 - Networking
 - Software and People



iris

Obtained significant funds & lobbied for more

- 2017:
 - £1.5M Capital for hardware
- 2018:
 - £16M (4M p.a. for 4 years)
 - 11.5M Capital for Hardware
 - 5.5M for RSE staff (Research Software Engineers)
- 2018:
 - Case presented to UKRI for long term investment
 - Lots of “Roadmaps” written
 - Cloud
 - Data&HTC
 - HPC
 - Networking
 - Software and People



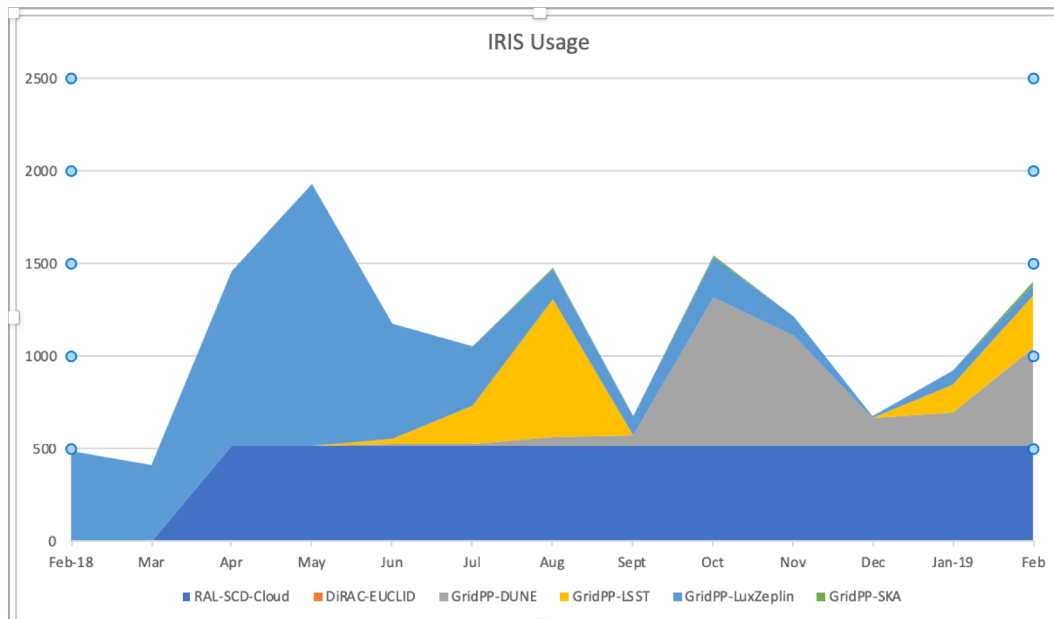
for HEP

p.a.



Deployed resources for new communities

- **Deployed physical resource:**
 - **2018: 8000 cores + 7 PB disk + tape**
 - **2019 14000 cores + 14 PB disk + tape**



Includes:

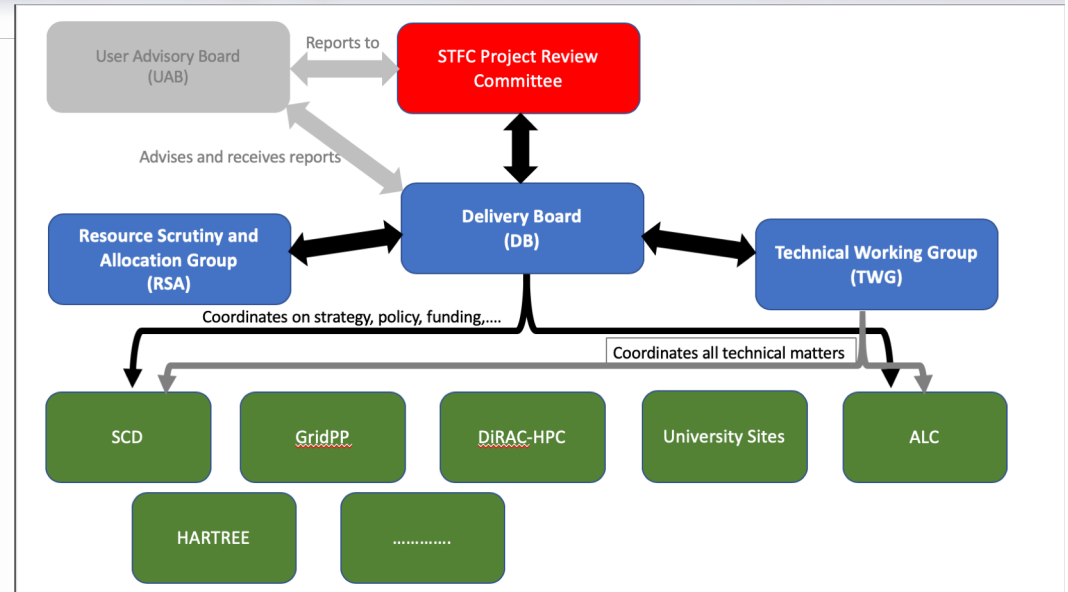
- LSST
- SKA
- EUCLID (Satelite)
- Lux Zeplin
- DUNE

- **Invested in some software development**
 - **Deployment of OpenStack federation/deployment tools**
 - **Developing Rucio for multi-community use**
 - **Developing DIRAC for multi community use**
 - **Advancing AAI across STFC**



Set up Resource Scrutiny & Allocation Group

- **RSAP**
 - Modelled on RRB/CRSG [& DiRAC RAC]
- Partners provide annual resource requirements doc
 - Justify resources
 - Not science



This means the “co-operative” is so far trusted to disburse its own funds for benefit of its partners using internal due diligence



- All user communities are different
 - Sociology
 - Tradition
 - Support for computing (or lack thereof)
 - Scale of their computing challenge
 - Stage of their computing challenge
- One cannot/should not (aggressively) tell other communities that they should do it your way.
 - Particularly don't say we have "solved all this so you should use it"
 - It is irrelevant whether this is true or not → Other Communities have to come to problems by themselves
- So you have to help them do what they feel they need to do "at this time"

- Most other communities (in UK at least) do not, and will never have, a distributed computing problem
 - So they probably don't need WLCG technology
- But some do: SKA
 - SKA SRC (Science Regional Centre) people know they have a lot of data and a distributed computing challenge.
 - They are working with CERN, and working with GridPP in UK
 - In UK happy with Grid+DIRAC so far
 - Looking at Rucio
- Other examples:
 - Case study : EUCLID
 - Case study : LSST

Enabled by local expertise
and good will in IRIS sites



- Many other communities need > 2 Gb/core