

Co-design of HPC services: the examples of DiRAC & ExCALIBUR

Mark Wilkinson

Director, STFC DiRAC HPC Facility

N8 CIR Community Day, 13 June 2022

DiRAC-3 Phase 1

DiRAC

Memory Intensive "COSMA8" (Durham)



DELL EMC

- 360 TB RAM
- Large-scale cosmological simulations



Atos

Extreme Scaling "Tursa" (Edinburgh)

- GPU-based system
- Large lattice-QCD simulations

Data Intensive "DlaL" (Leicester)



Hewlett Packard Enterprise

- Heterogeneous architecture for complex simulation and modelling workflows



DELL EMC

Data Intensive "CSD3" (Cambridge)

- Heterogeneous architecture for complex simulation and modelling workflows

Project Office (UCL)

Applying the scientific method to HPC service design

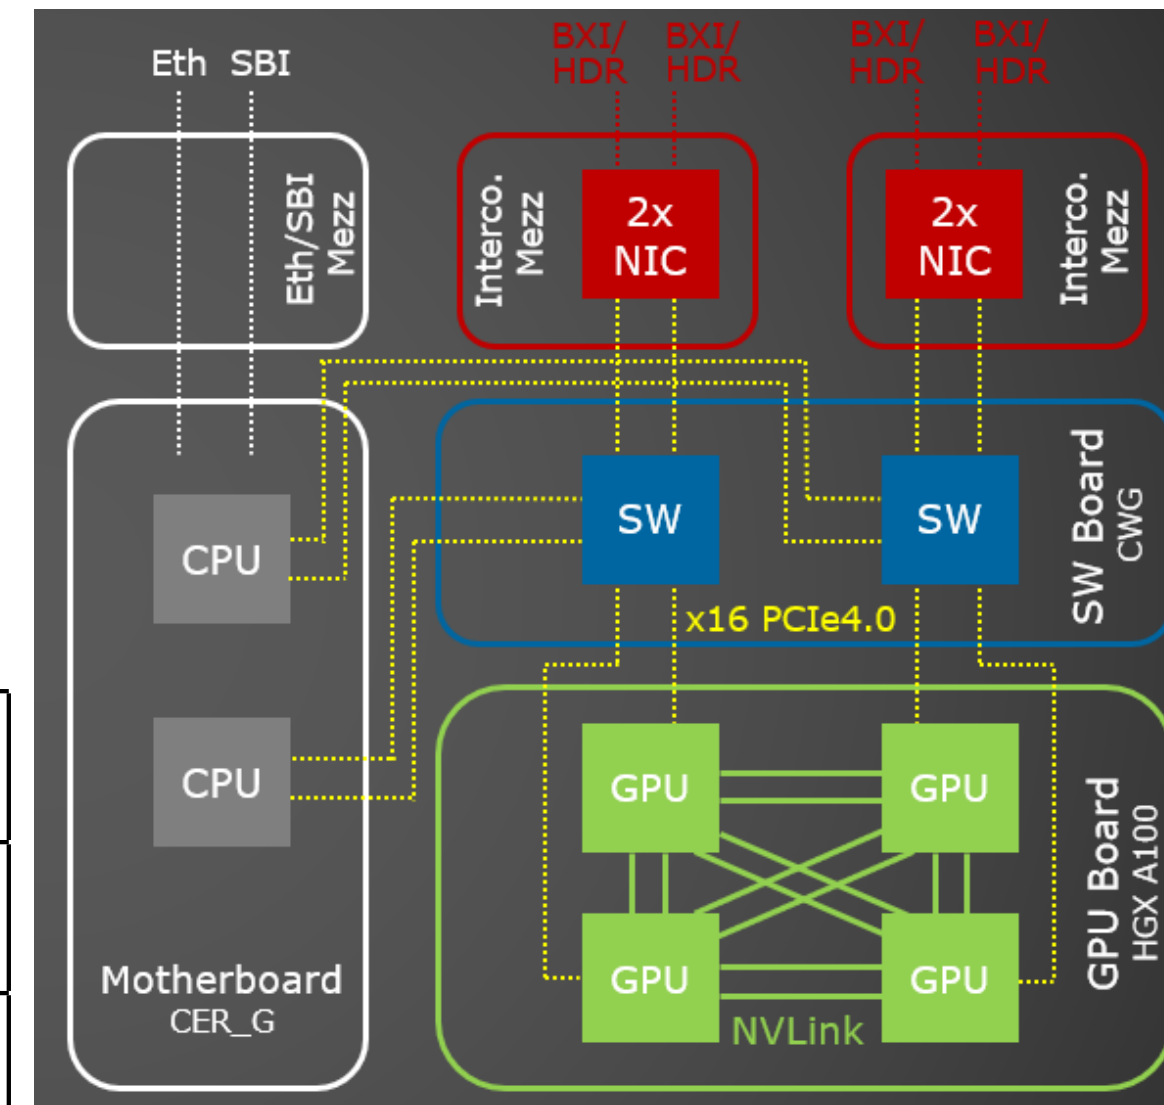


What did we expect from Tursa?

Tursa:

- 112 ATOS Sequana XH2000 nodes
- 4x Nvidia A100-40 GPUs per node
- 2x AMD Rome 7302 CPUs
- 4x HDR-200 infiniband per node

nodes	Performance (Tflop/s)
1	9.2
16 measured	5.3 per node
16 committed	5.83 per node
Network committed	>160 GB/s



Performance of Grid code - Boyle et al.

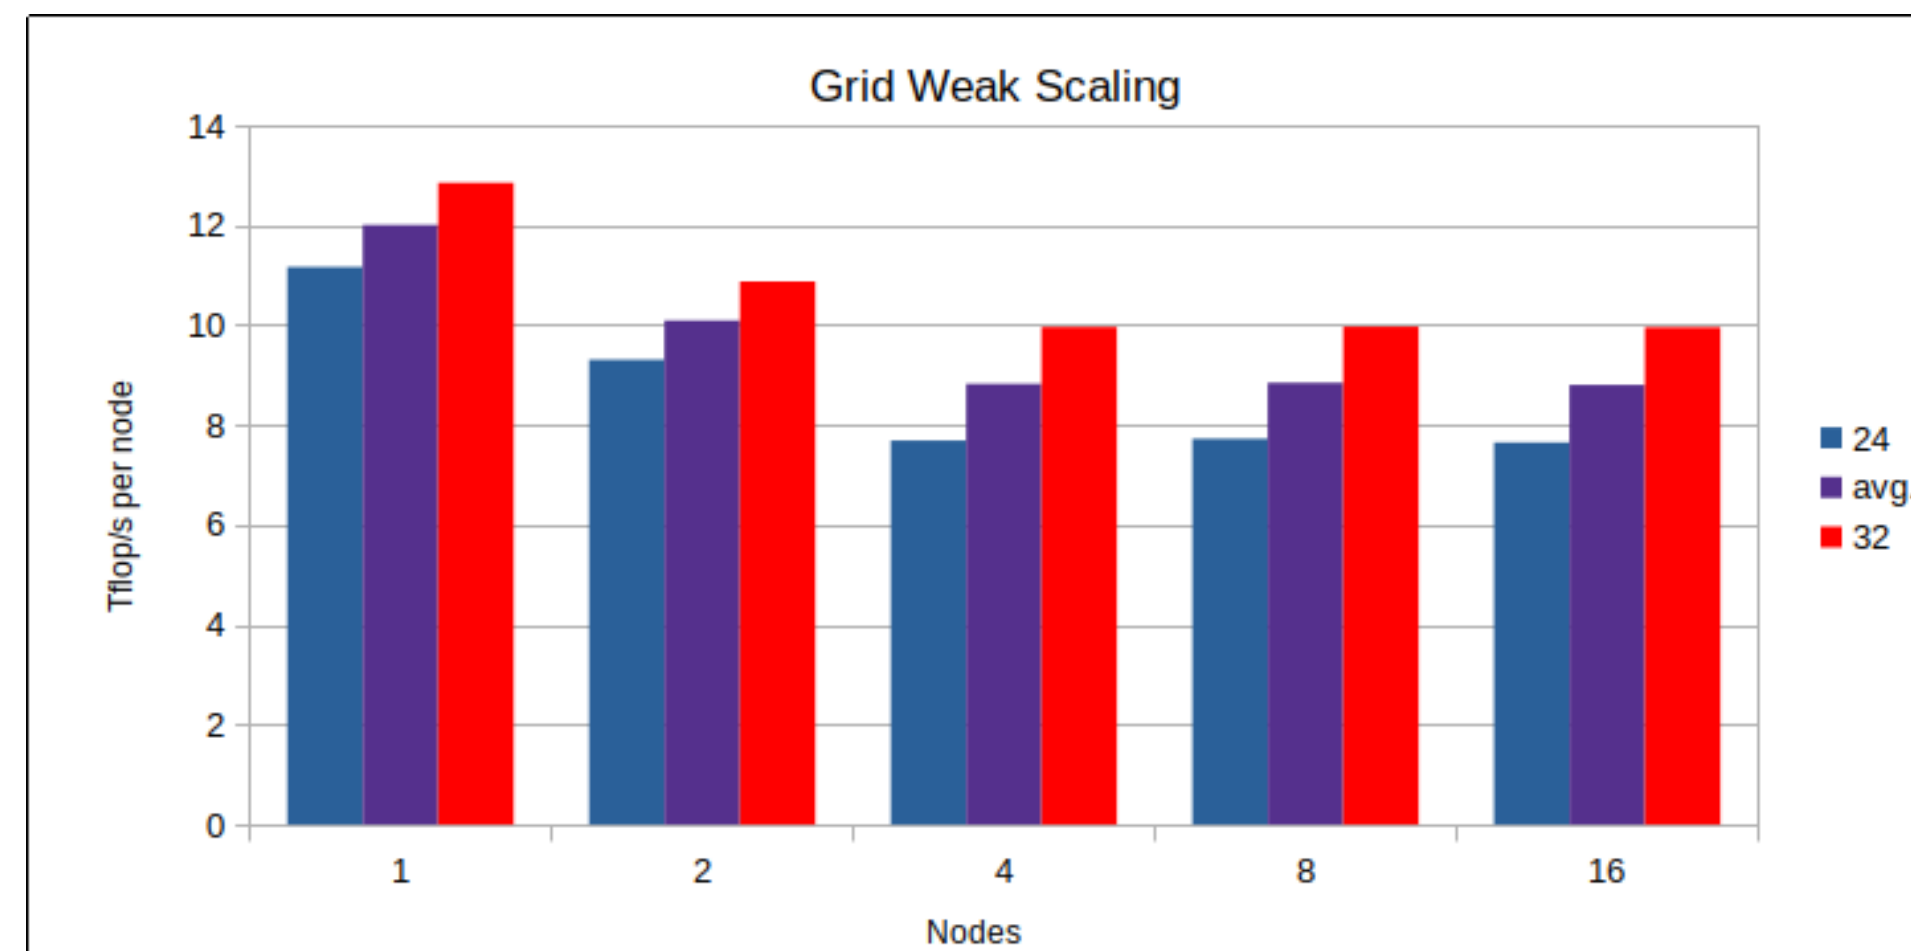


How does this performance compare?

- Tesseract 512 nodes \times 150 Gflop/s per node = 76.8 Tflop/s
- Tursa 16 nodes \times 5.3 Tflop/s per node = 84.8 Tflop/s

Commissioning Outcomes: Benefits of co-design

Stage	1 node	% inc.	16 nodes	% inc.	speed up 512 tess
Measured	9.2	-	5.3	-	1.1
Committed	9.2	-	5.83	10%	1.22
Acceptance	9.65	5%	6.15	16%	1.28
Commissioning	12	30%	8.8	66%	1.83
Peak	12.9	40%	9.9	87%	2.06



Tesseract : Tursa

- $1468 \times 150\text{Gflop/s} = 220\text{Tflop/s}$: $112 \times 9.9\text{Tflop/s} = 1109\text{Tflop/s}$

UKRI Supercomputing Ecosystem - principles

Goal: ensure all researchers have access to sufficient provision of the most appropriate hardware for their science.

Scope and design

- Research-driven
 - requirements defined by peer-reviewed science cases.
- Assessed based on scientific and industrial productivity
- Sustainably funded, including power, ResOps, RSE and algorithm development

Delivery

- Community-driven to ensure efficiency and productivity
- Value for money - evidence-based decision-making
- Co-designed with industry partners
- Supported by investments in people and skills

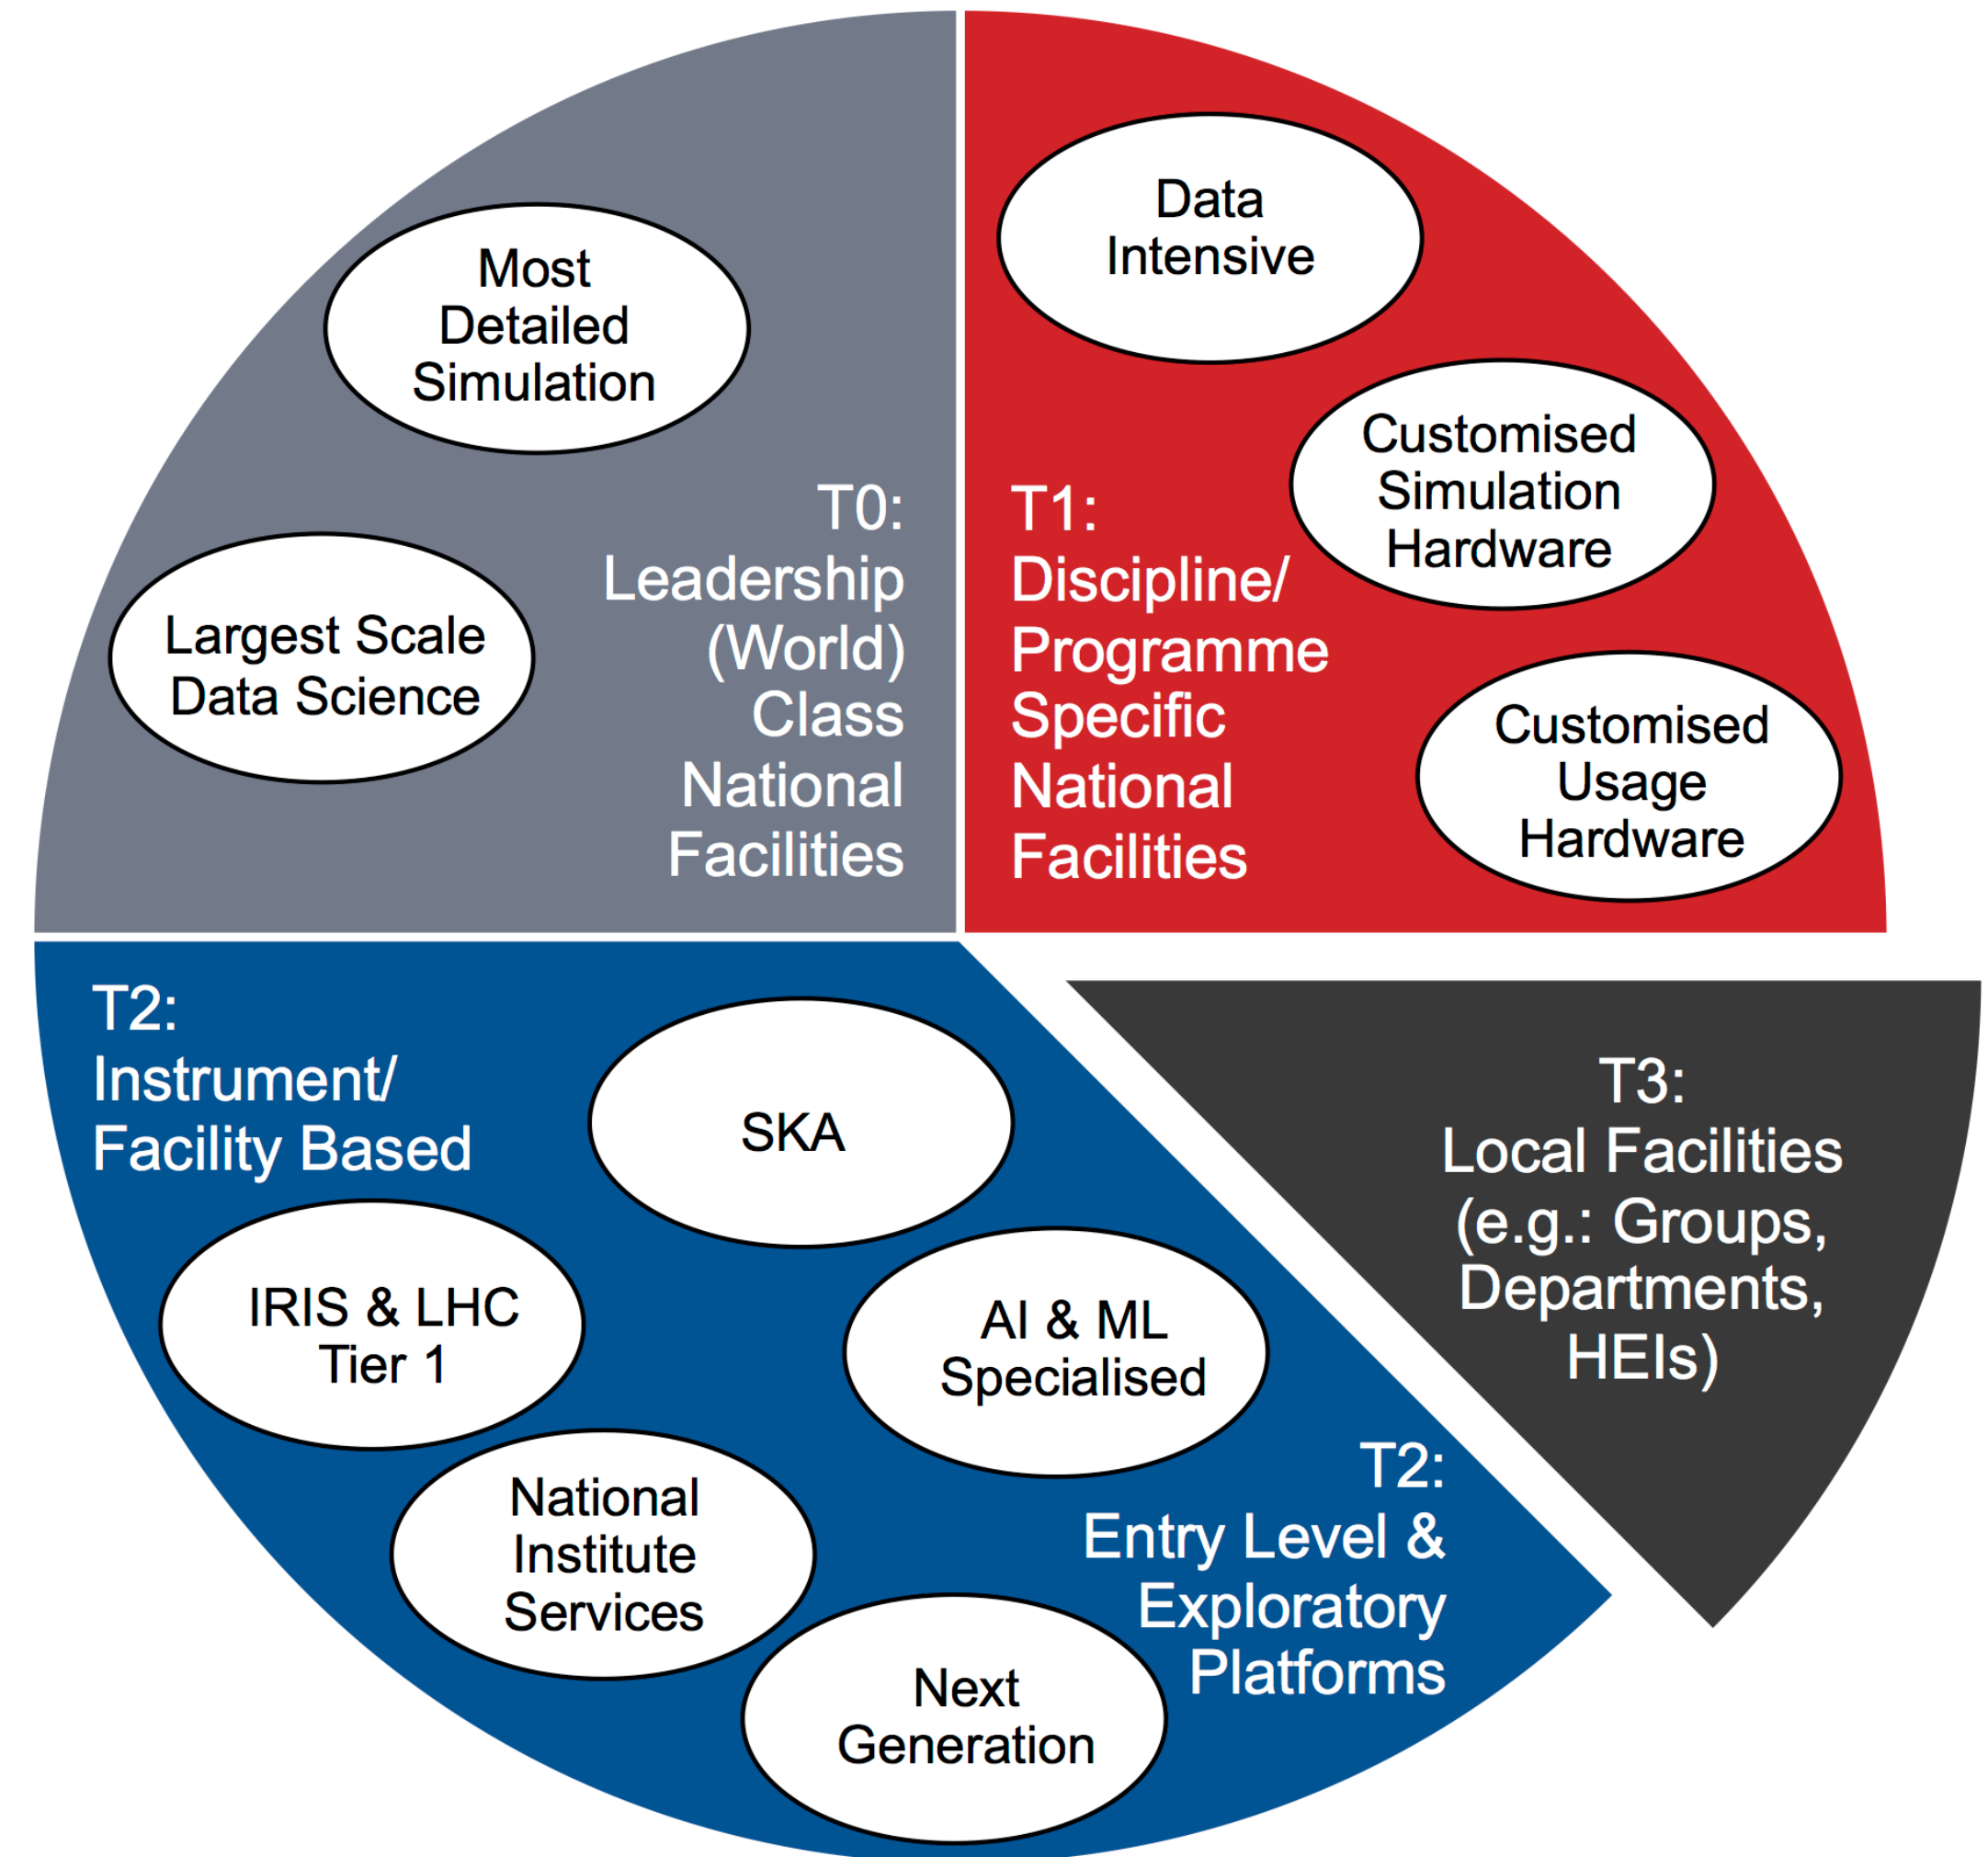
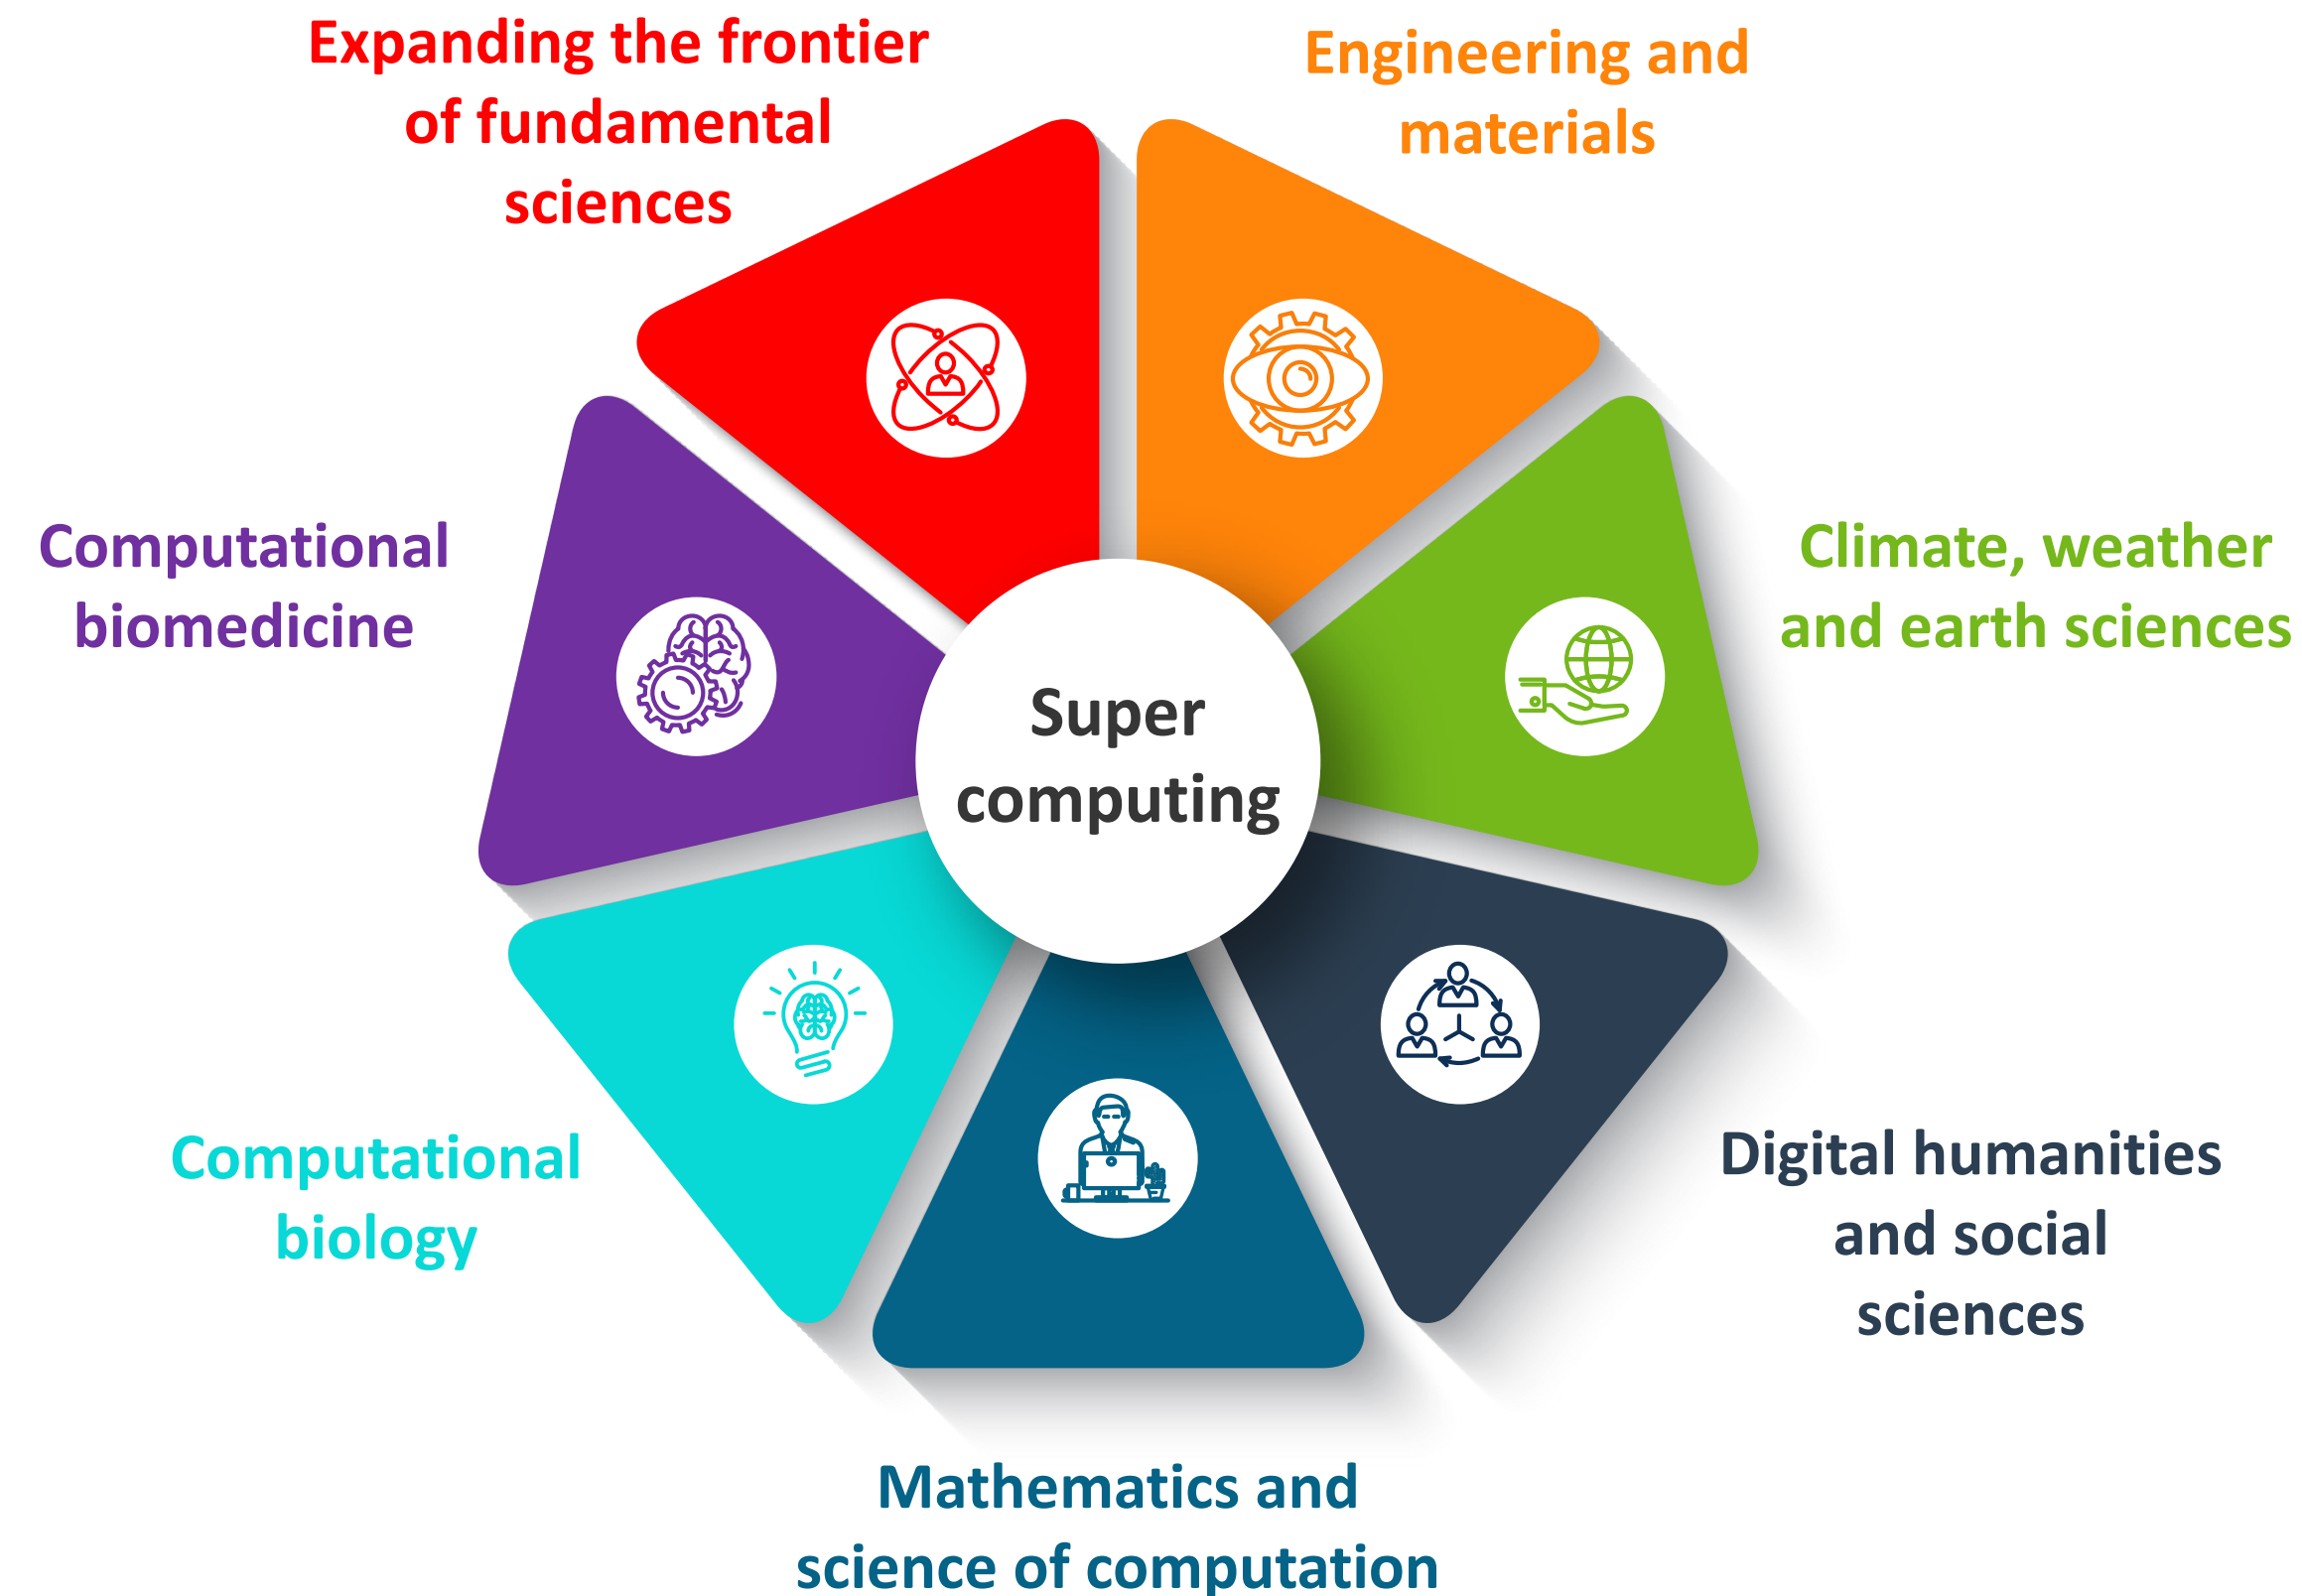


Image credit: Bryan Lawrence & Clare Jenner

UKRI supercomputing science case



Editors: Wilkinson (Chair),
Birney, Chakraborty, Coveney,
Ford, Johnson, Lawrence,
Parsons, Prescott, de Roure,
Sijacki, Trigg, Wingate

- Supercomputing science case assembled with inputs from across UKRI research communities
- Science drivers for the Tier-0 (leading to exascale) and Tier-1 (petascale)
- Internationally peer-reviewed in EPSRC-managed peer review process
- First step in the design process for UKRI supercomputing services
- Noted convergence and cross-fertilisation of simulation and AI workflows
- Available online: <https://excalibur.ac.uk/ukri-science-case-for-uk-supercomputing/>

ExCALIBUR High Level Overview

- **Exascale Computing Algorithms and Infrastructures Benefiting UK Research**
 - £45.7M from the Strategic Priorities Fund (SPF)
 - Led by UKRI and the Met Office with UKAEA
 - The UK's 5 year Exascale programme
 - Primary focus is on software and algorithms
 - 10% of budget allocated to testbeds exploring novel Hardware and Enabling Software (H&ES)




Photo credit: UK Atomic Energy Authority (UKAEA)

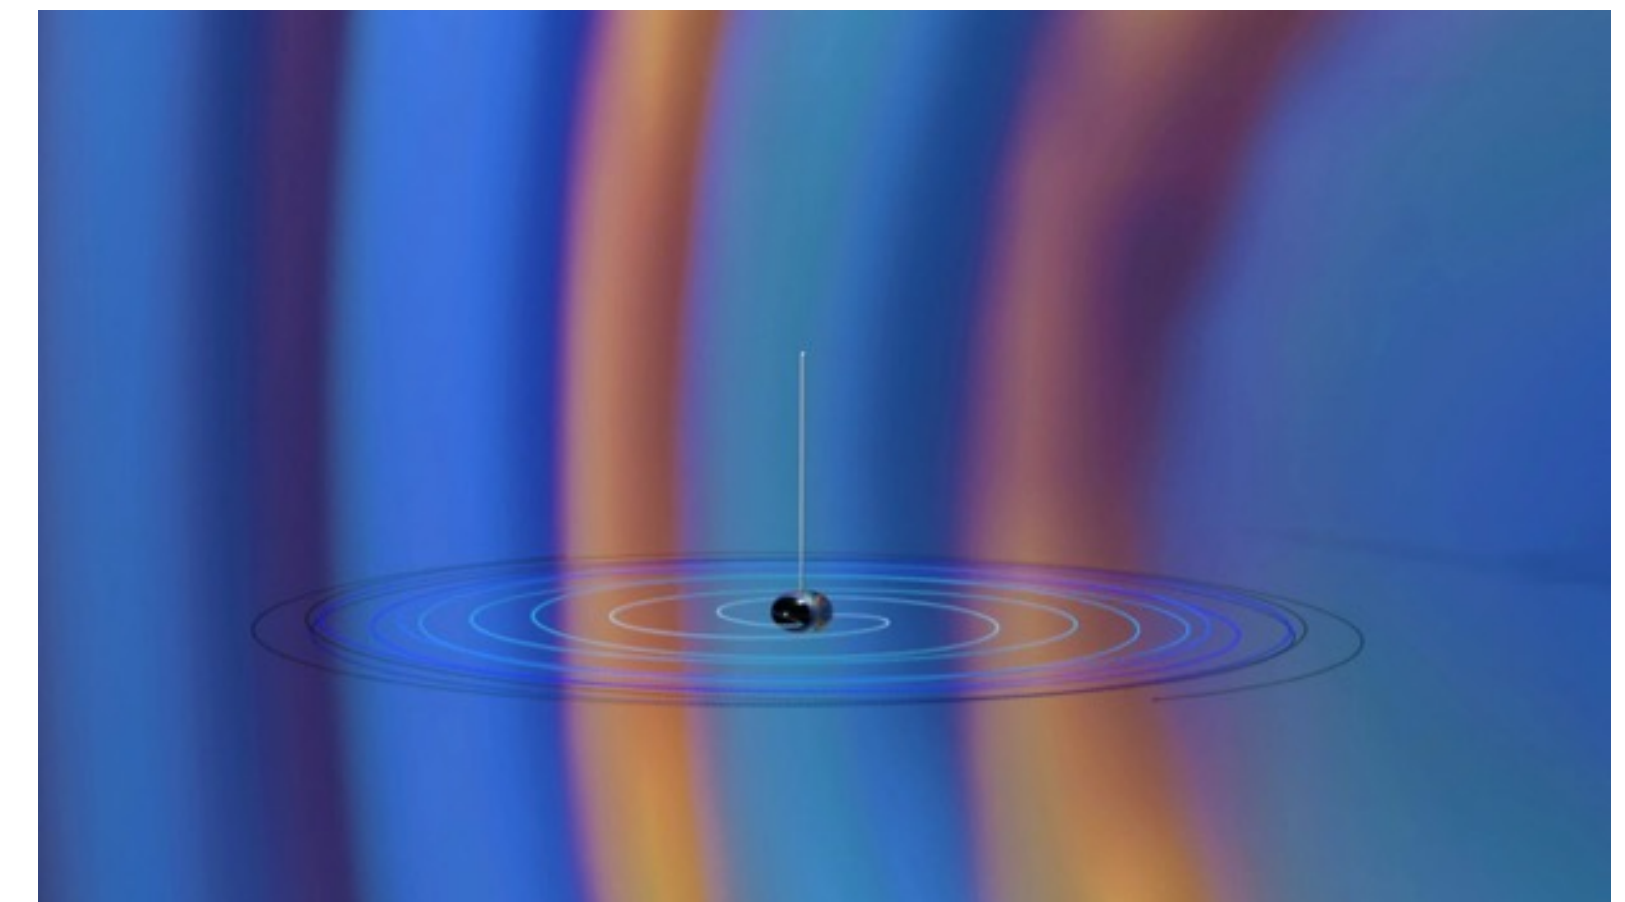
ExCALIBUR H&ES Key Ingredients

Hardware and enabling software:

- CPUs (ARM, AMD)
- Accelerators (GPUs, FPGAs, 'AI' etc)
- Interconnects (BlueField-1/2, Rockport)
- Filesystems (DAOS etc)
- Tooling (Compilers, debuggers, profilers etc)
- Benchmarking



Includes new paradigms!




- Durham hosts three test-beds: Bluefield; AMD GPU; Rockport

Towards UK Exascale

- Two activities underway:
 - UK Exascale supercomputer project for 2025
 - Design Study work on pre-Exascale pathfinders
- Lessons from ECP, EuroHPC, Fugaku, DiRAC-3, ARCHER2 *et al*:
 - Co-design systems with technology partners, aiming for cloud presentation
 - Design and benchmark against science community use cases
 - Embrace the art of the possible – pragmatic solutions required around power and cooling
 - “Exaflop system or Exascale ecosystem?” Discuss.



```
33 var self = getObj(data, id),
34     parents = [];
35 if(self){
36     parents.push(self);
37     if(self.pId){//If pId is not @
38         parents = parents.concat
39             (getSelfAndParents(data,
40              self.pId));
41     }
42 }
43 //Through the id to find themselves and all
44 //its paren function getSelfAndParentsId(data, id)
45 function getSelfAndParentsId(data, id) {
46     var self = getObj(data, id),
47         ids = [];
48     if(self){//In the presence of self
49         ids.push(id);
50         if(self.pId){//If pId
```



Summary

- The DiRAC HPC facility supports UK theory research in astrophysics, particle physics, cosmology and nuclear physics.
- DiRAC services are defined based on the workflows they support
- Hardware and software co-design delivers more productive HPC services

(Better systems) + (Better software) = Better research

- ExCALIBUR Hardware & Enabling Software programme is providing access to test-beds for all UKRI researchers

dirac.ac.uk

excalibur.ac.uk



DiRAC - HPC



@DiRAC_HPC