The UKRI supercomputing ecosystem 2020-30

Mark Wilkinson
Director, DiRAC HPC Facility

Talk based on the *draft* white paper:

UKRI National Supercomputing Roadmap 2019-30

Editor: Mark Wilkinson

UKRI Supercomputing Working Group members:
Rob Akers, Paul Alexander, Peter Boyle, Robert Davey, Tom Griffin, Clare Jenner, Alison Kennedy, Leigh Lapworth, Bryan Lawrence, Richard Lawrence, Simon McIntosh-Smith, Mark Parsons, Andrew Sansum, Paul Selwood, Arthur Trew, Mark Wilkinson, Jeremy Yates, Serena Viti.

Supercomputing in research

- Direct numerical simulation and modelling is core research activity
  - HPC systems are main scientific instruments for theory research
- Computational requirements of models increasing due to
  - Increased resolution: running models with existing physics at finer scales.
  - Increased complexity: introducing new physics into models to reflect progress in theoretical understanding; often needed to match resolution
  - Coupling of models: multi-physics, multi-scale modelling;
  - Quantification of modelling uncertainty using large ensembles of simulations to provide robust statistics
- Constant process of refining and re-defining our tools
- Growing requirement for simulations and modelling concurrent with observations so that models evolve in line with data acquired
  - Observational/experimental facilities need access to significant local computing capabilities as well as option to burst out to the larger, national facilities.

Supercomputing and AI – complementarity not competition

- Increased cross-fertilisation of ideas between “HPC” and “AI” communities:
  - Hardware requirements between simulations and AI are converging
  - AI algorithm optimisation using approaches developed for HPC
  - AI use in analysis of simulation data
  - AI-driven “computational steering” to increase efficiency of simulations
  - AI replacing parameterised “sub-grid” models

- Explainable AI – significant development work needed to
  - increase value of AI as a scientific research tool
  - reduce “black-box” application of AI

Current UK supercomputing provision

Provider	Activity	Location(s)
Met Office	Weather forecasting and climate modelling	Exeter
ARCHER	Engineering, Physical Sciences, Natural Environment	Edinburgh Parallel Computing Centre (EPCC)
DiRAC	PPAN Theory	Cambridge, Durham, EPCC, Leicester
Hartree Centre	Industry and Commerce	Daresbury Laboratory (DL)
EPSRC Tier 2 services	Engineering, Physical Sciences	Cambridge, HPC Midlands+ (at Loughborough), Bristol, EPCC, Oxford (at DL), UCL (JISC Southern Data Centre), N8 (at Durham), Queen's Belfast
Research Data Facility	EPSRC, NERC, PPAN	EPCC
JASMIN (NERC)	Natural Environment	Rutherford Appleton Laboratory (RAL)
Earlham Institute	Life Sciences	Norwich

UKRI National Supercomputing 2019-30 White Paper (Draft)

UK supercomputing investment has fallen behind…..

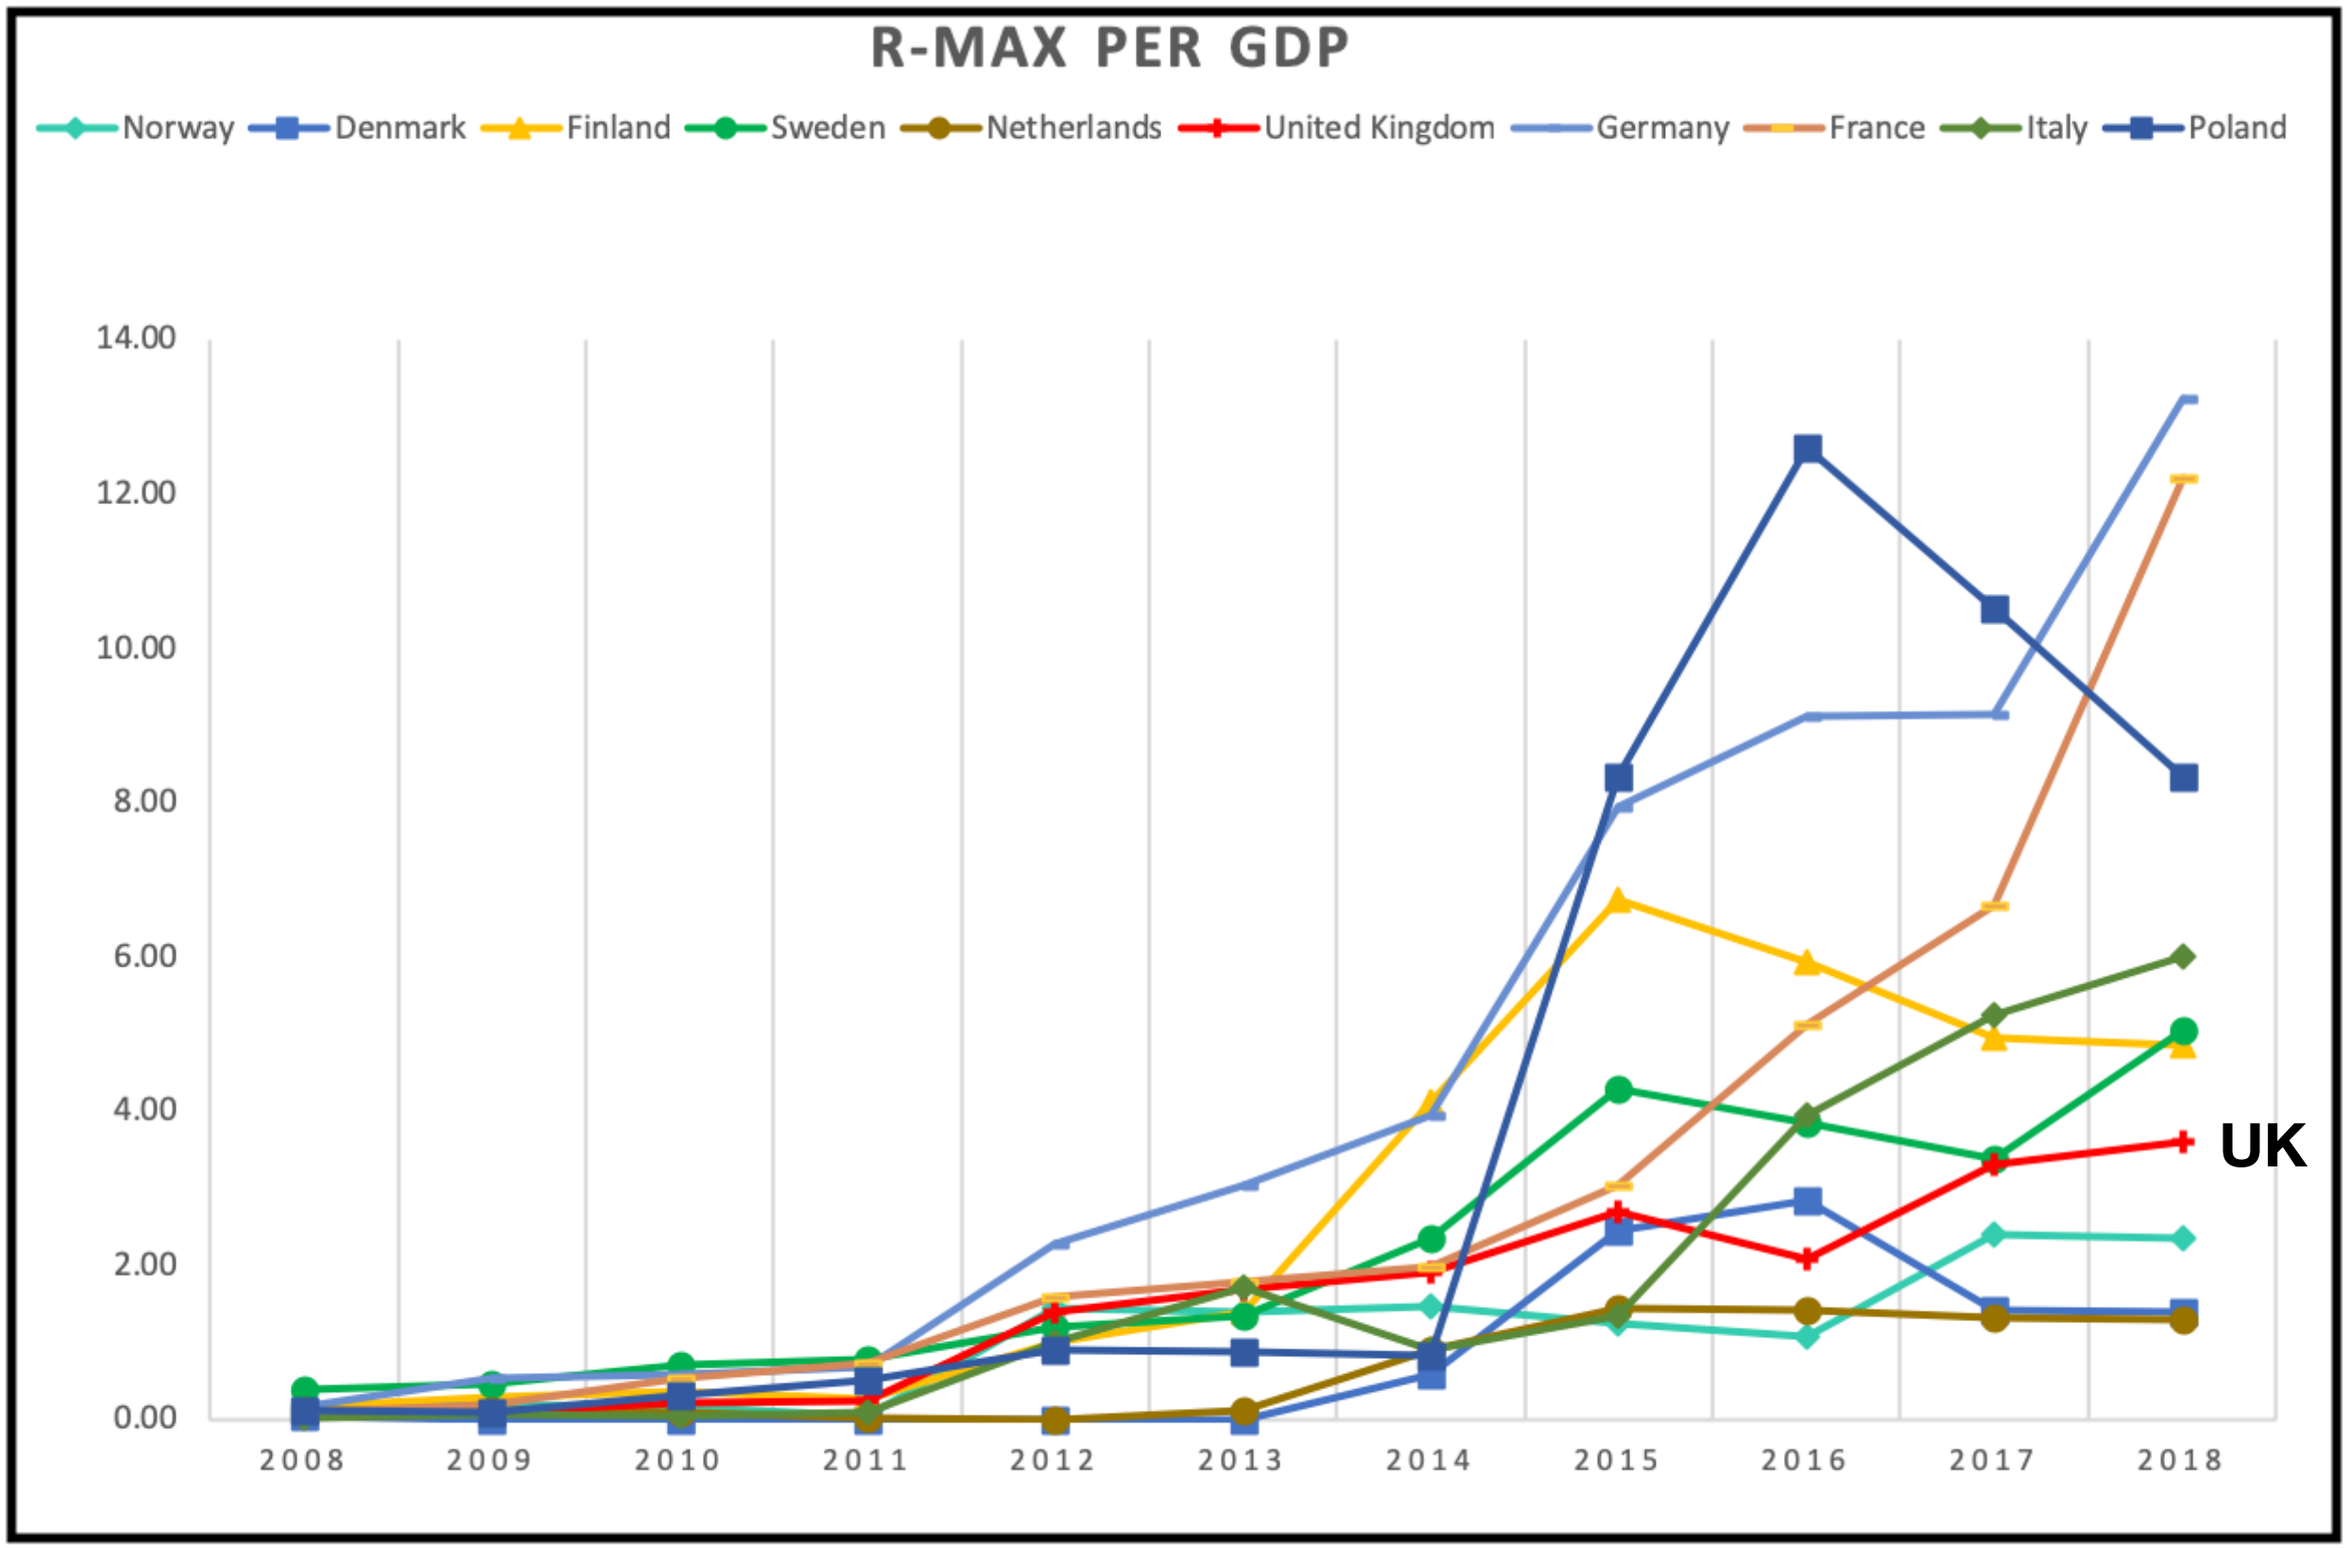

UKRI National Supercomputing 2019-30 White Paper (Draft)

International comparisons

Country	Peak Pflop/s	System Name	No of systems above 5Pflops (Nov 2019)
USA	201	DOE ORNL Summit	20
China	125	Sunway TaihuLight	2
UK	***22-35***	***DiRAC-3 (est.)***	***-***
UK	***28***	***ARCHER 2 (2020; est.)***	***-***
Switzerland	27	Piz Daint	1
Germany	27	SuperMUC	5
Korea	26	Nurion	1
Japan	25	Oakforest-PACS	4
Italy	19	Marconi	1
Taiwan	15	Taiwania 2	1
Spain	10	Mare Nostrum	1
Canada	7.5	Beluga	1
Saudi Arabia	7.2	Shaheen II	1
Australia	7	Gadi Phase 1	1
UK	**4.7**	**DiRAC-2.5y (combined)**	**0**
UK	**4.3**	**Scafell Pike**	**0**
UK	**2.6**	**ARCHER**	**0**

Vision

“Our aim is for the UK to have a supercomputing infrastructure that:

- enables world-leading research and innovation for both academia and industry
- enables and supports cross-disciplinary activity;
- keeps pace with technology developments;
- is well-coordinated to realise the most effective use of the combined resources that we have.”

UKRI National Supercomputing 2019-30 White Paper (Draft)

UKRI Supercomputing Ecosystem - principles

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
- Co-designed with industry partners
- Supported by investments in people and skills

UKRI National Supercomputing 2019-30 White Paper (Draft)

UKRI Supercomputing Ecosystem - requirements

1. A range of well-coordinated supercomputing services
   - appropriate diversity of scales and architectures to support research programme of UKRI and UK industry.
2. Sustainable long-term funding for UK supercomputing to facilitate strategic planning and maximise value for money
   - Include funding for staff and power
3. Support for software engineering work and algorithm development
   - required to exploit new systems and maximise science output.
4. Deploy an exaflop system, delivering genuinely exascale science
   - “exascale” refers to both compute (ExaFlops) and data (ExaBytes).
5. Align with the HPC vendor technology roadmaps and support activities leading to effective exploitation of emerging technology
   - to seed a competitive advantage to both the scientific community and the UK economy.

UKRI National Supercomputing 2019-30 White Paper (Draft)

UKRI Supercomputing ecosystem

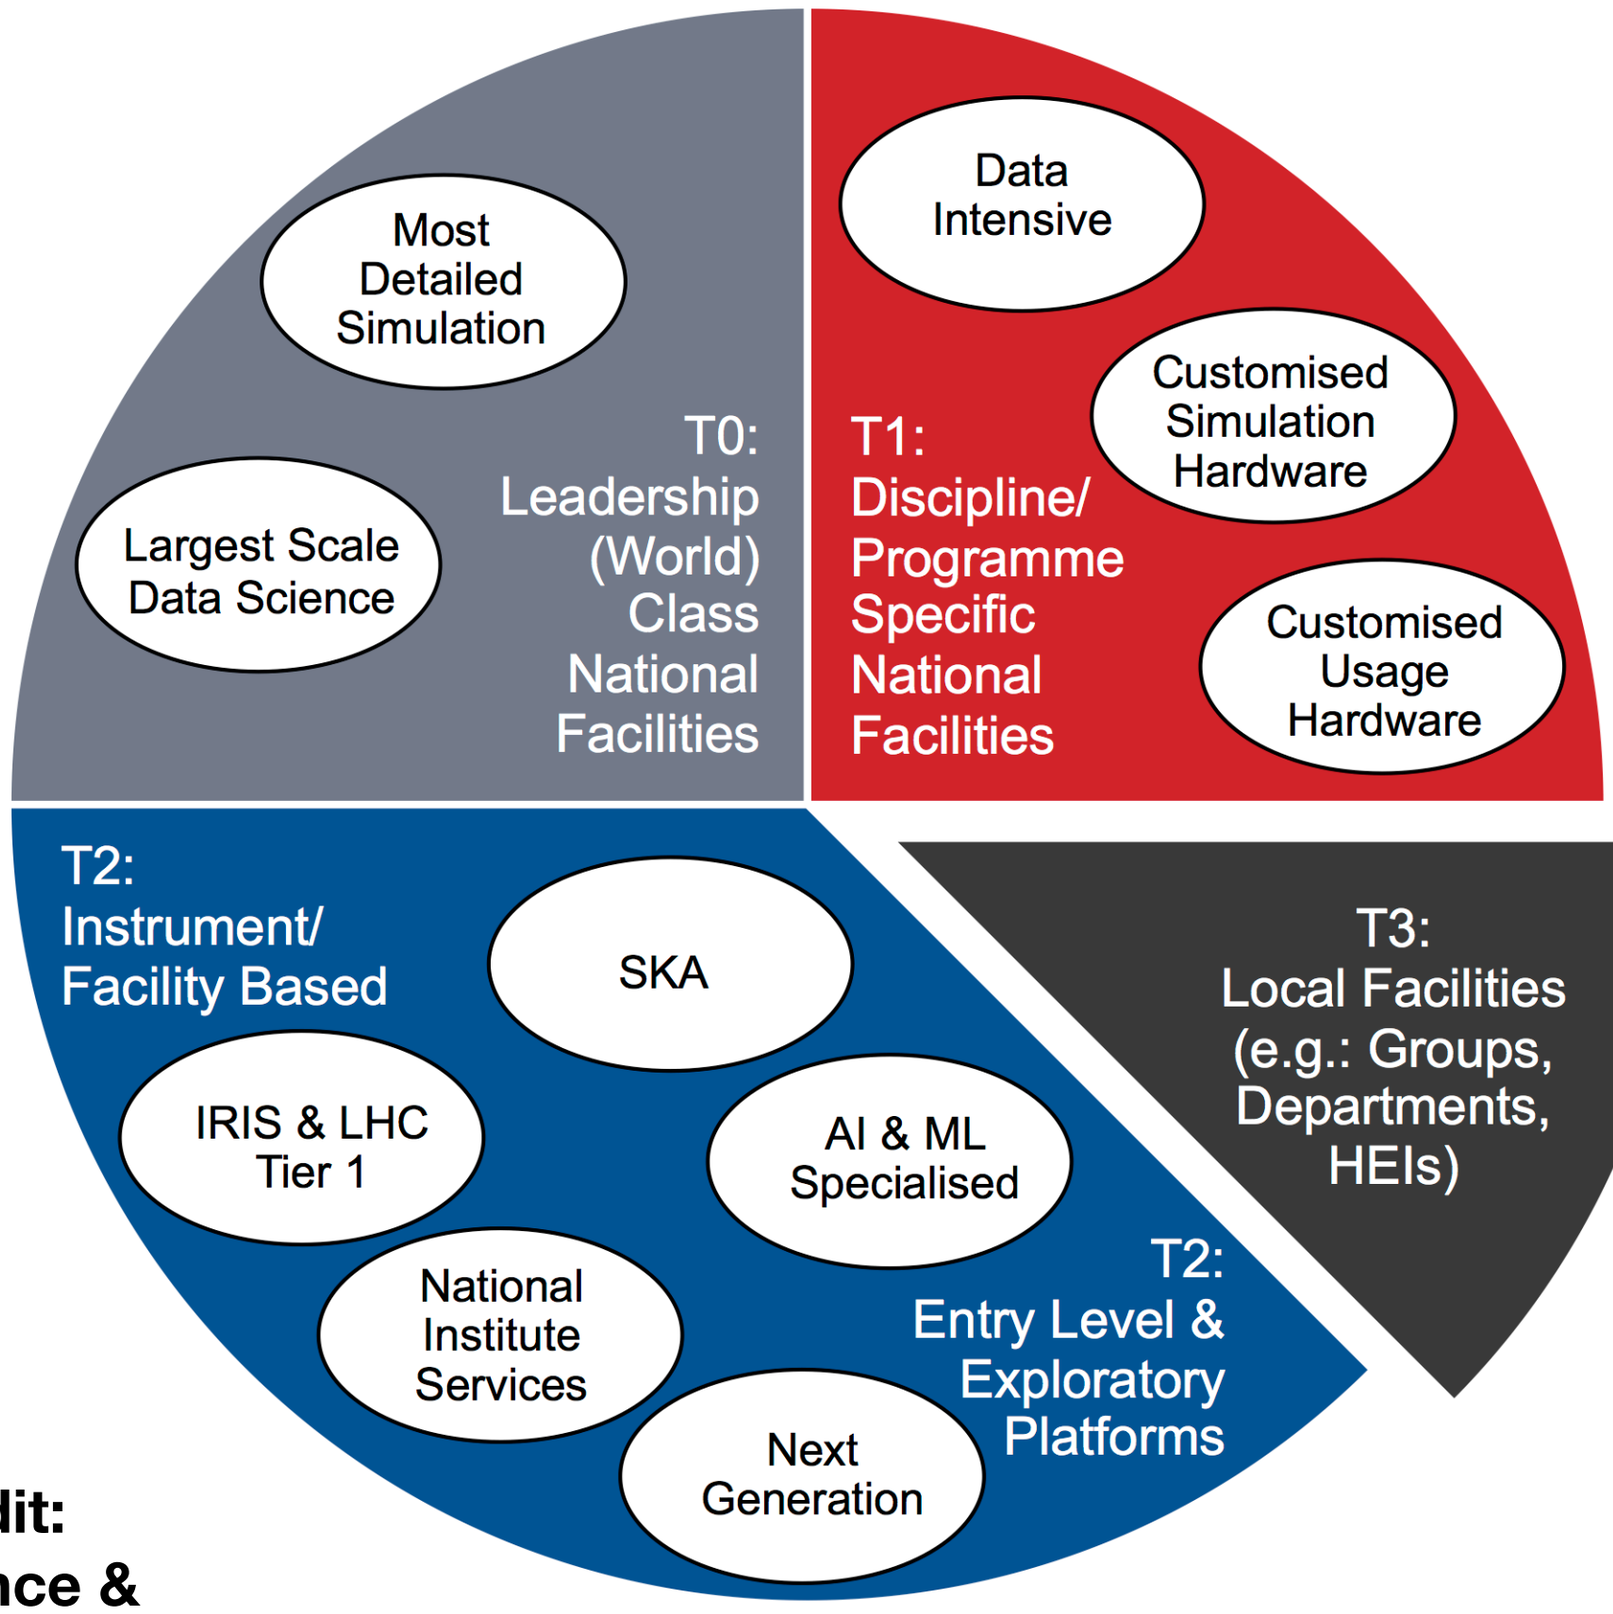

Image credit: Bryan Lawrence & Clare Jenner

UKRI National Supercomputing 2019-30 White Paper (Draft)

UKRI Supercomputing roadmap

	Baseline							
National Tier 0 Service (Extreme Scaling)		**National Tier 0 Service** (55PF) **Pathfinder**		**National Tier 0 Service** (200-300PF) **Pre-exascale**				**National Tier 0 Service** (1EF)
National Tier 1 Services (Extreme Scaling Memory Intensive Data Intensive)	**Scafell Pike** **Archer** **Met Office** **DiRAC2.5y** **Jasmin**	**Jasmin** **Archer 2** (25PF) **DiRAC3** (20-30PF)	**Jasmin** **Met Office** **Hartree** (10PF)	**Hartree** (10PF) **Jasmin** **Archer 3** (50PF) **DiRAC3b** (100PF)	**Jasmin**	**Jasmin**	**National Tier 1 Services** (300PF) **Archer** **DiRAC** **Jasmin**	**Jasmin**
National Tier 2 Services (Memory Intensive Data Intensive)	**Earlham**	**National Tier 2 Services**				**National Tier 2 Services** (10PF) **Earlham** **IRIS**		
	2018/19	**2019/20**	**2020/21**	**2021/22**	**2022/23**	**2023/24**	**2024/25**	**2025/26**

UKRI National Supercomputing 2019-30 White Paper (Draft)

UKRI e-Infrastructure

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

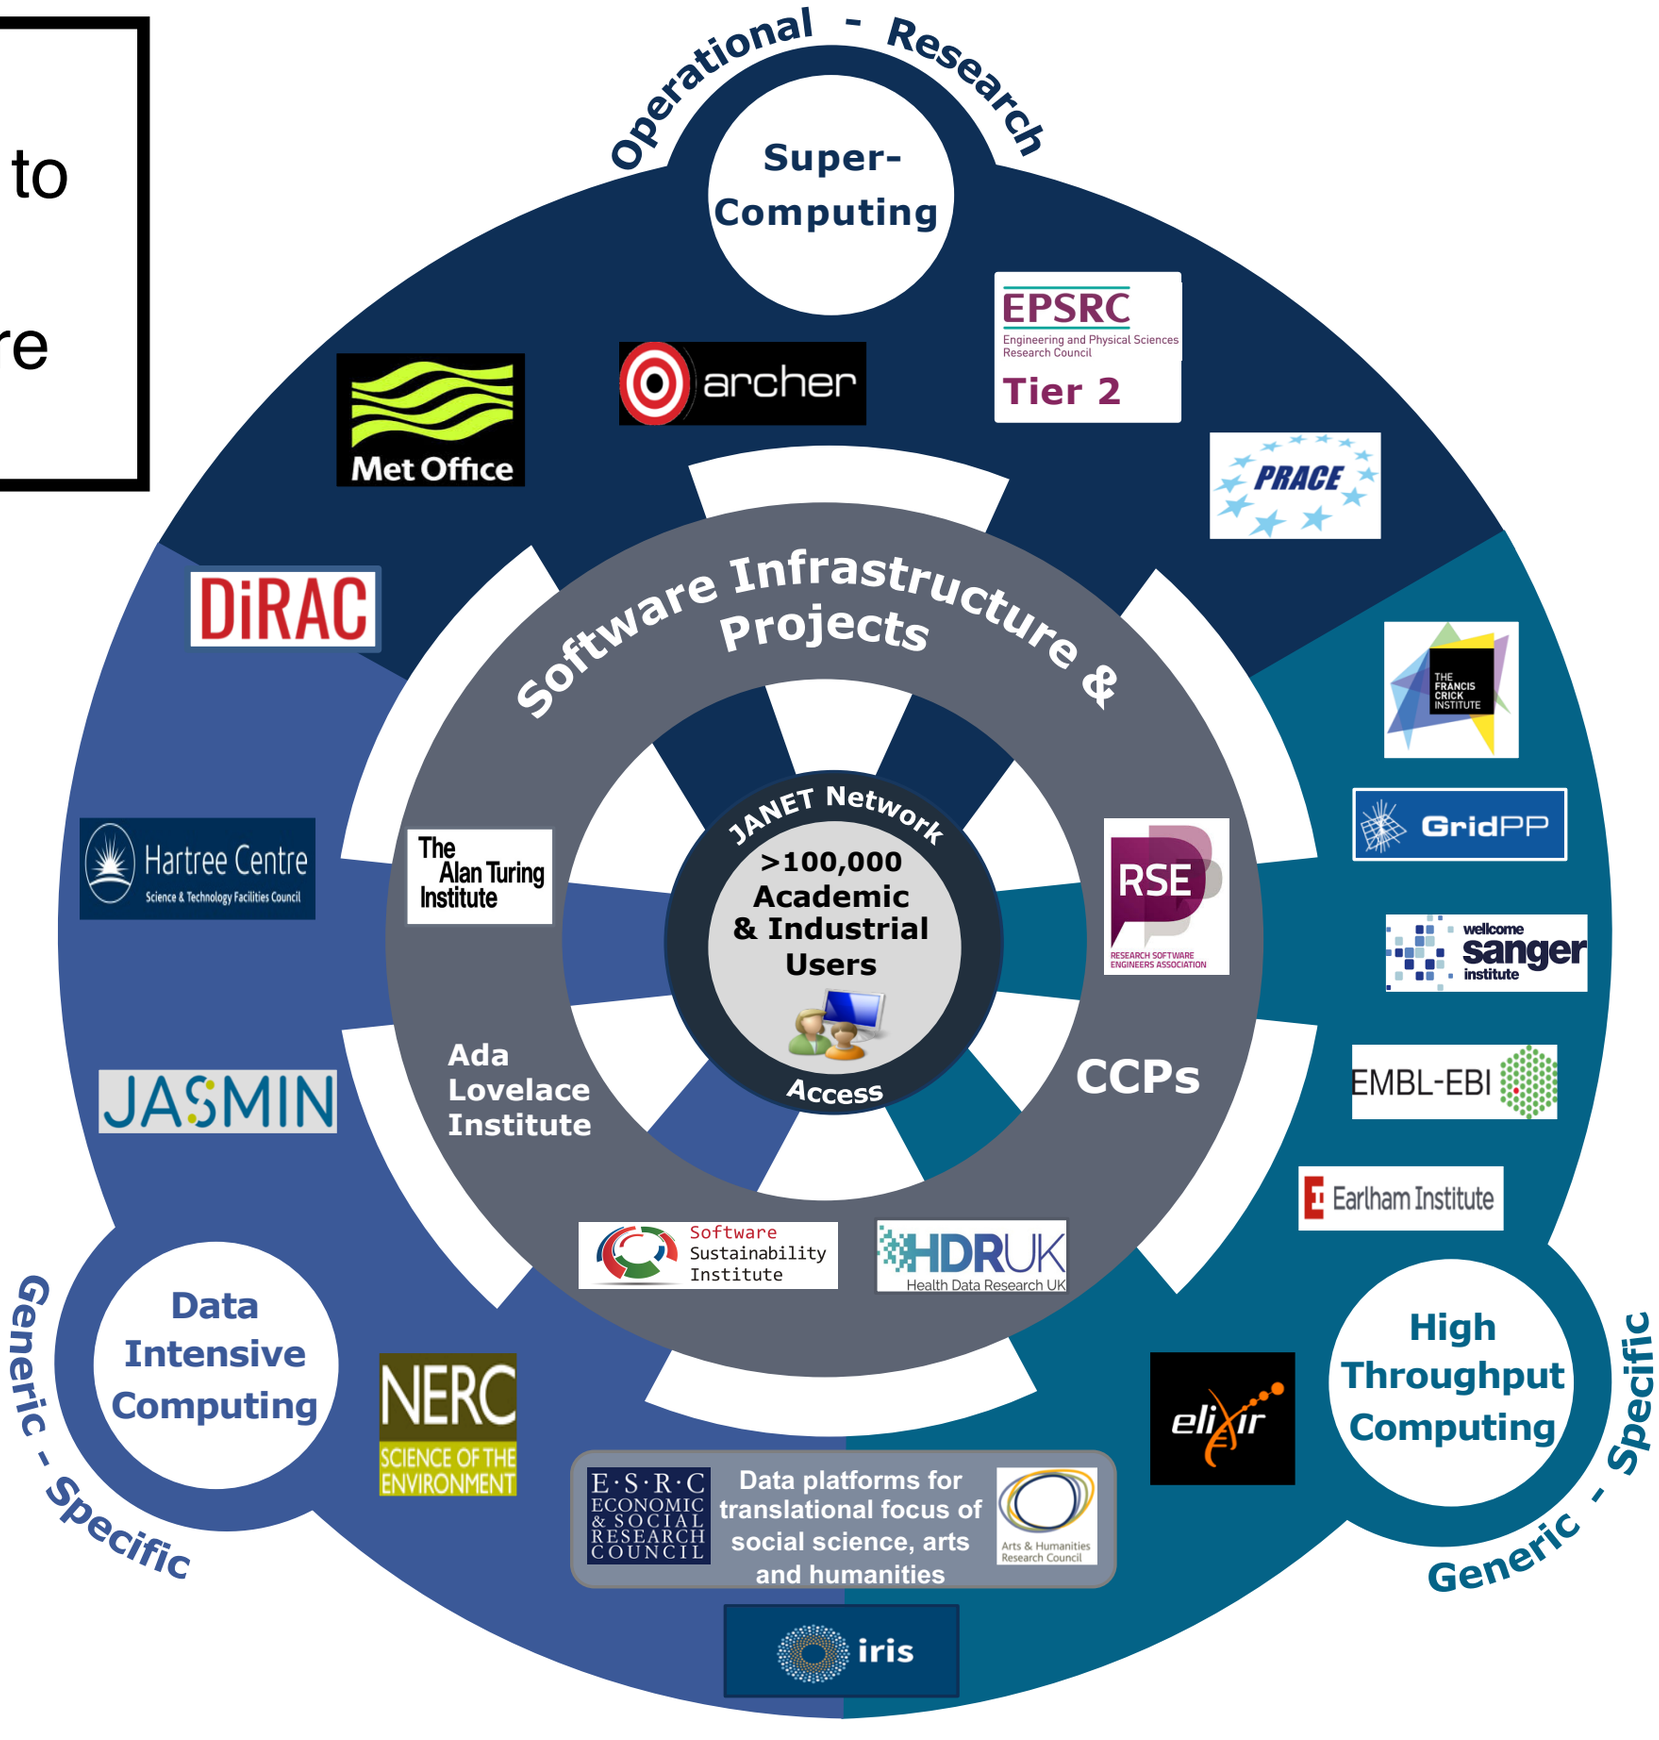

Image credit: Clare Jenner

Supercomputing White Paper - recommendations

0) Establish a UKRI Supercomputing ecosystem which is:

- research-driven - requirements defined by science case
- sustainably funded, including OPEX, RSE and algorithm development
- Delivers at all scales, including exascale

1) Assemble a peer-reviewed “UK Supercomputing” Science Case;

2) Establish a national technology horizon-scanning programme;

3) Actively engage with research communities in the development of the roadmap and ecosystem;

4) Establish a joint technology development programme to develop algorithms, hardware and software;

- Include comp. scientists, comp. mathematicians, ICT researchers

5) Establish a large-scale, securely-funded, long-term programme of software re-engineering;

6) Explore opportunities to increase coordination between existing supercomputing services across UKRI

7) Develop a high-profile frontier call (akin to INCITE) to support exceptionally ambitious research (including software development).

ExCALIBUR

- Recommendation 2 of supercomputing white paper:
  - "Establish a national technology horizon-scanning programme"
- Recently announced initiative led by UK Met Office, UK Atomic Energy Agency and UKRI to develop software for exascale
- DiRAC is managing ExCALIBUR "Hardware & Enabling Software" project on behalf of UKRI community:
  - Small-scale deployments of a range of processors, network switches, memory, etc.
  - Supporting code benchmarking and testing
  - Annual £1M investment to support addition of new hardware
  - Leverage vendor co-funding and support
  - Access available to all UKRI projects for testing purposes
  - Enable RSE teams and researchers to target future hardware

UKRI supercomputing science case

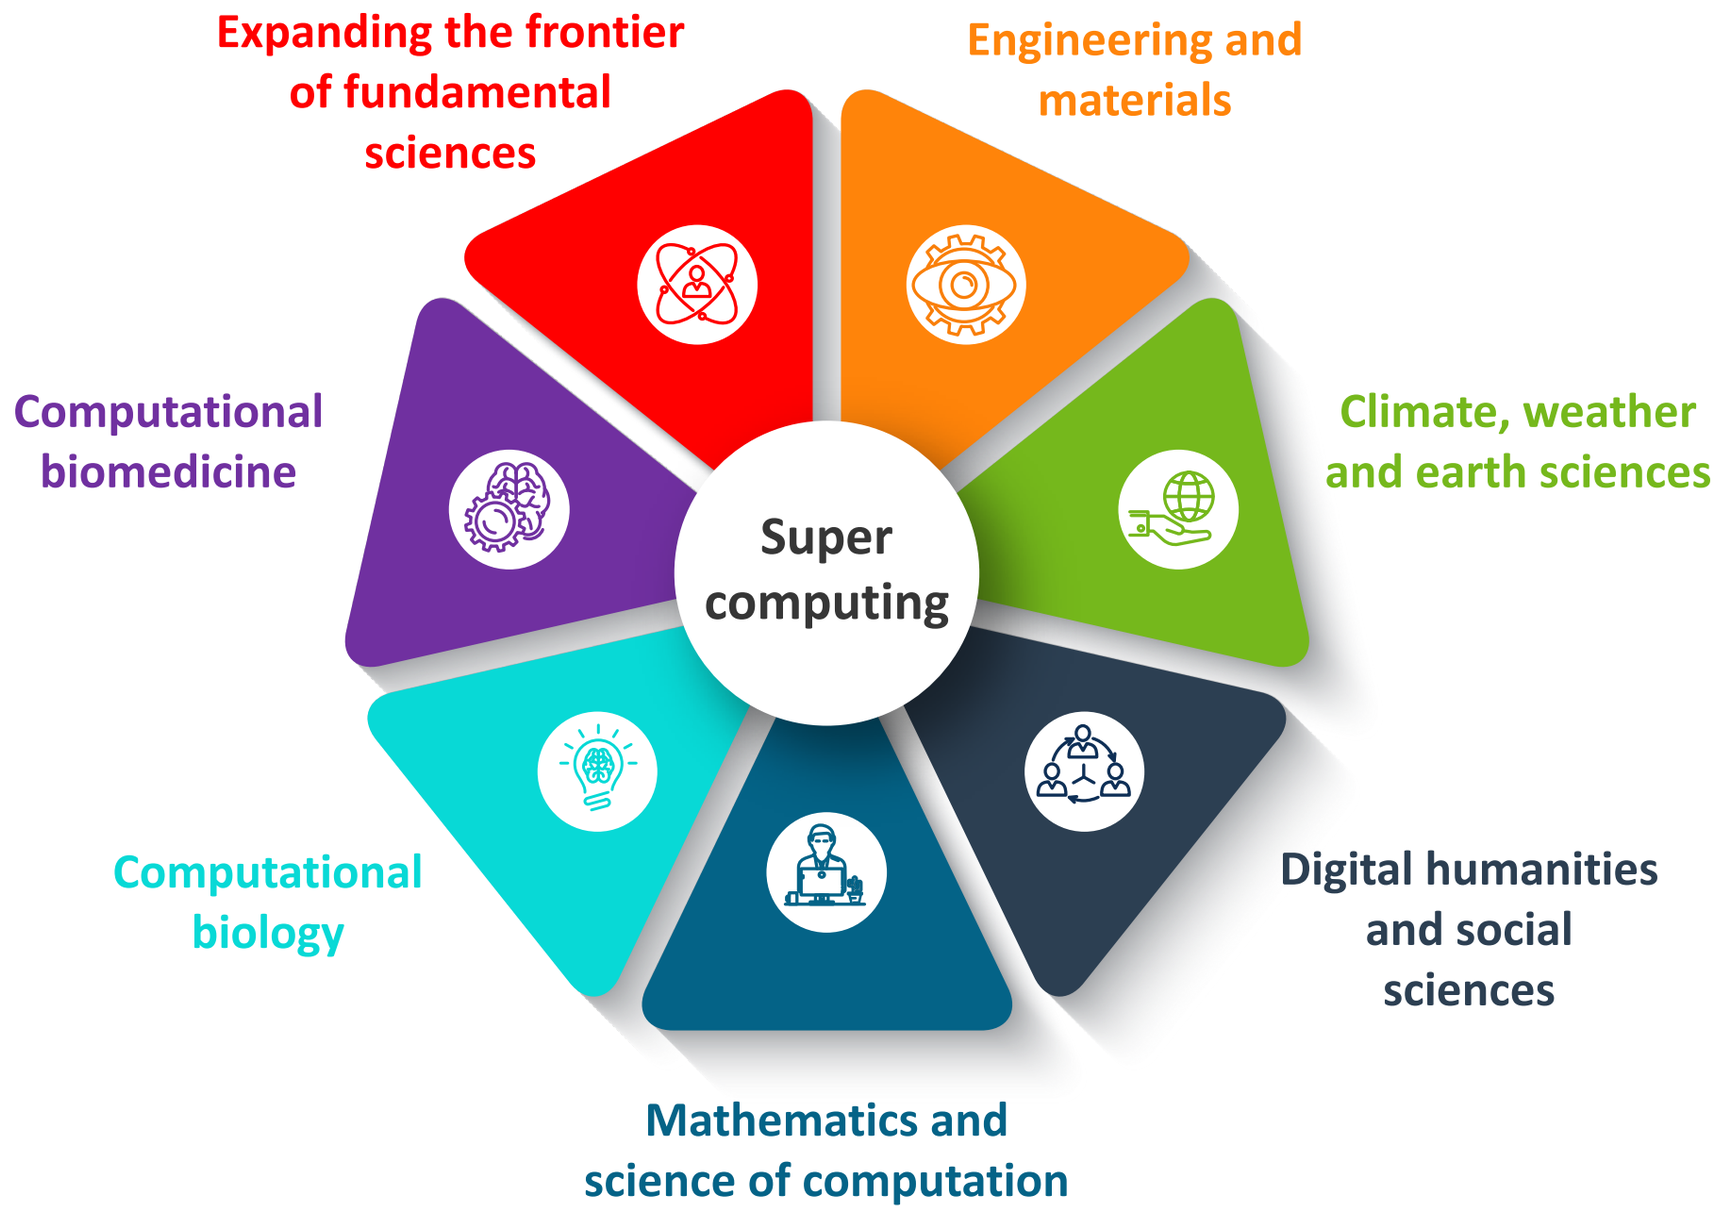

- Supercomputing science case assembled with inputs from across UKRI research communities
- Science drivers for the Tier-0 (leading to exascale) and Tier-1 (petascale)
- Currently being peer-reviewed
- Will be used to support business cases associated with the roadmap

Mathematics and science of computation

Key challenges which must be addressed to ensure productivity of new supercomputing systems:

1. Develop the new area of Mathematics at scale
2. Performance modelling and next generation benchmarking
3. Composable languages and tools across supercomputing applications
4. Working with industry to design the next generation of supercomputing systems
5. Next generation development cycle
6. New Research Excellence Framework (REF) Unit of Assessment for Computational Science.

Industry access to UKRI Supercomputing ecosystem

At least three categories of collaborative research projects:

1. Access to academic systems to develop ideas (“selling cycles”);
2. Understand scaling
    - exploring what large and/or novel systems can deliver;
3. New research in collaboration with academics
    - developing new models/simulations or methods of exploiting data

Conclusions

- Draft UKRI Supercomputing white paper presents an ambitious vision for the next decade
  - interdependent with other parts of NeI roadmap
- Development of draft white papers and roadmap was genuine cross-community effort
- All aspects of roadmap are research-driven
  - science cases and benchmarks underpin all planning
- Significant increase in investment is required to maintain UK leadership and competitiveness
- Collaboration between UKRI communities and industry partners is essential to deliver the roadmap
- Watch for opportunities and engage!