THE UK EXASCALE PROJECT

Professor Mark Parsons

EPCC Director
Associate Dean for e-Research

epcc

Exascale supercomputing and the UK

- Since mid-2018 it has been clear the UK needed an Exascale supercomputing strategy
- UK has not joined EuroHPC and will not be eligible to take part in EuroHPC projects or access calls once Horizon 2020 ends
- UK has been involved in EuroHPC negotiations
  - UK Government Representative from Business, Energy & Industrial Strategy (BEIS) Department attended meetings with Technical Advisor from December 2017
- Strategic case made in August 2018 to undertake a scoping study to develop an Exascale Outline Business Case for UK Government

Exascale Project Working Group

- BEIS and UKRI asked me to set up an small Exascale Project Working Group to develop an Exascale Outline Business Case
- EPWG members
  - Mark Parsons – EPCC (Chair)
  - Susan Morrell / Richard Gunn – UKRI/EPSRC
  - Alison Kennedy – UKRI/STFC
  - Leigh Lapworth – Rolls Royce plc
  - Paul Selwood – UK Meteorological Office
  - Mark Wilkinson – DiRAC (Chair of Science Case WG)
- Supported by
  - Nick Trigg (STFC), Chris Johnson (EPCC), Ritchie Somerville (consultant) and Mark Graham (consultant)

What is a Green Book Business Case?

- **The Green Book: Central government guidance on appraisal and evaluation (2018)** sets out its preferred approach to how to develop a rationale for intervention
- Critical to the **options appraisal** are a clear understanding and articulation of why change is needed through **a statement of rationale for intervention**
- The complementary guidance **Guide to Developing the Project Business Case (2018)** is also used.
- The case is 25 pages body plus 15 pages annexes

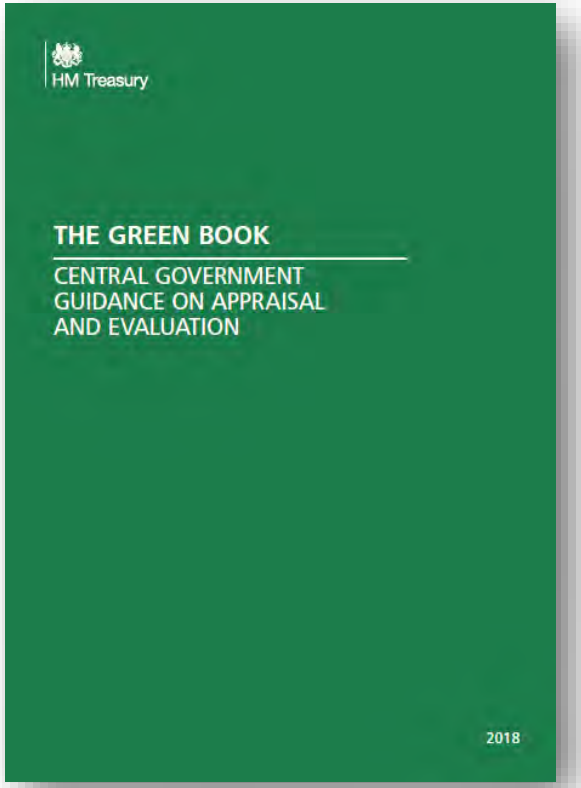

(132 pages)

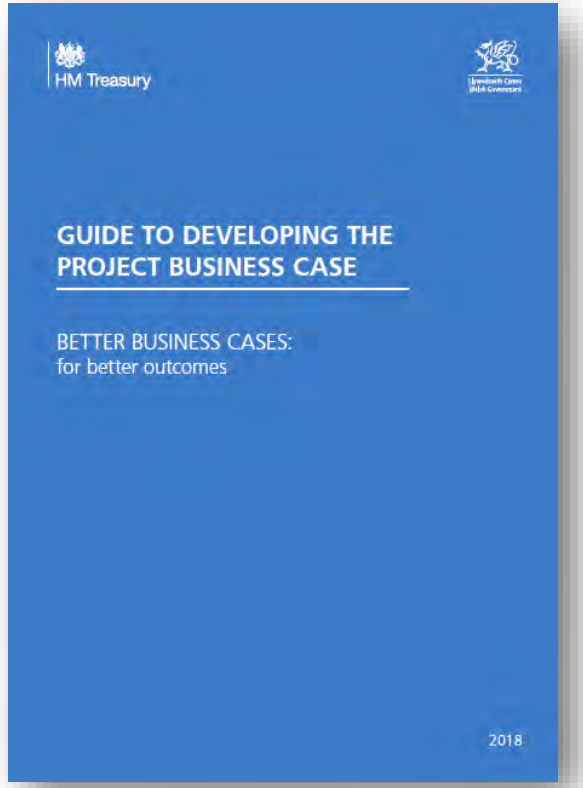

(134 pages)

epcc

Four key questions

1. Do we need an Exascale capability?
2. What will an Exascale supercomputer look like?
3. What will it cost and how will it be funded?
4. How will it be hosted?

Making the case – International Context

- China 2020 / 2021
  - Little known at present – updated CPU plus accelerator as per Sunway
- Japan 2020
  - Fugaku system – based on A64FX ARM processor
- USA 2021
  - Frontier system – based on AMD EPYC CPU + AMD GPU
  - Aurora system – based on Intel A21 CPU + Intel PV GPU
- Europe 2020 / 2023
  - Pre-Exascale hosting sites chosen – Finland / Spain / Italy – late 2020
  - Exascale systems based on EPI – late 2023

Making the case – Overall Context

Leadership Class Systems	• Internationally relevant • Most detailed simulations and largest data problems
Discipline/Programme Specific Systems	• Dedicated to specific disciplines / programmes • Some targetted custom hardware • Specific usage patterns
Entry level and Instrument/Facility Specific Systems	• Smaller scale systems for smaller tasks • Often linked to specific research facilities • Exploration of new computing technologies
Local Systems	• University level systems • Departmental or research group level systems • Small company systems

Making the case – Science Case

- UKRI have been developing an e-Infrastructure Roadmap for the past 18 months
  - The overall roadmap has been supported by a series of White Papers – one of these was on supercomputing
- In parallel with development of Outline Business Case, Mark Wilkinson has led development of Supercomputing Science Case
  - Focus on supercomputing as a whole – not just Exascale
  - Case spans the sciences
  - Focuses on how a 10X increase and 100X increase (Exascale) on today's resource will transform scientific discovery
- Currently being internationally peer reviewed

Making the case – Science Case themes

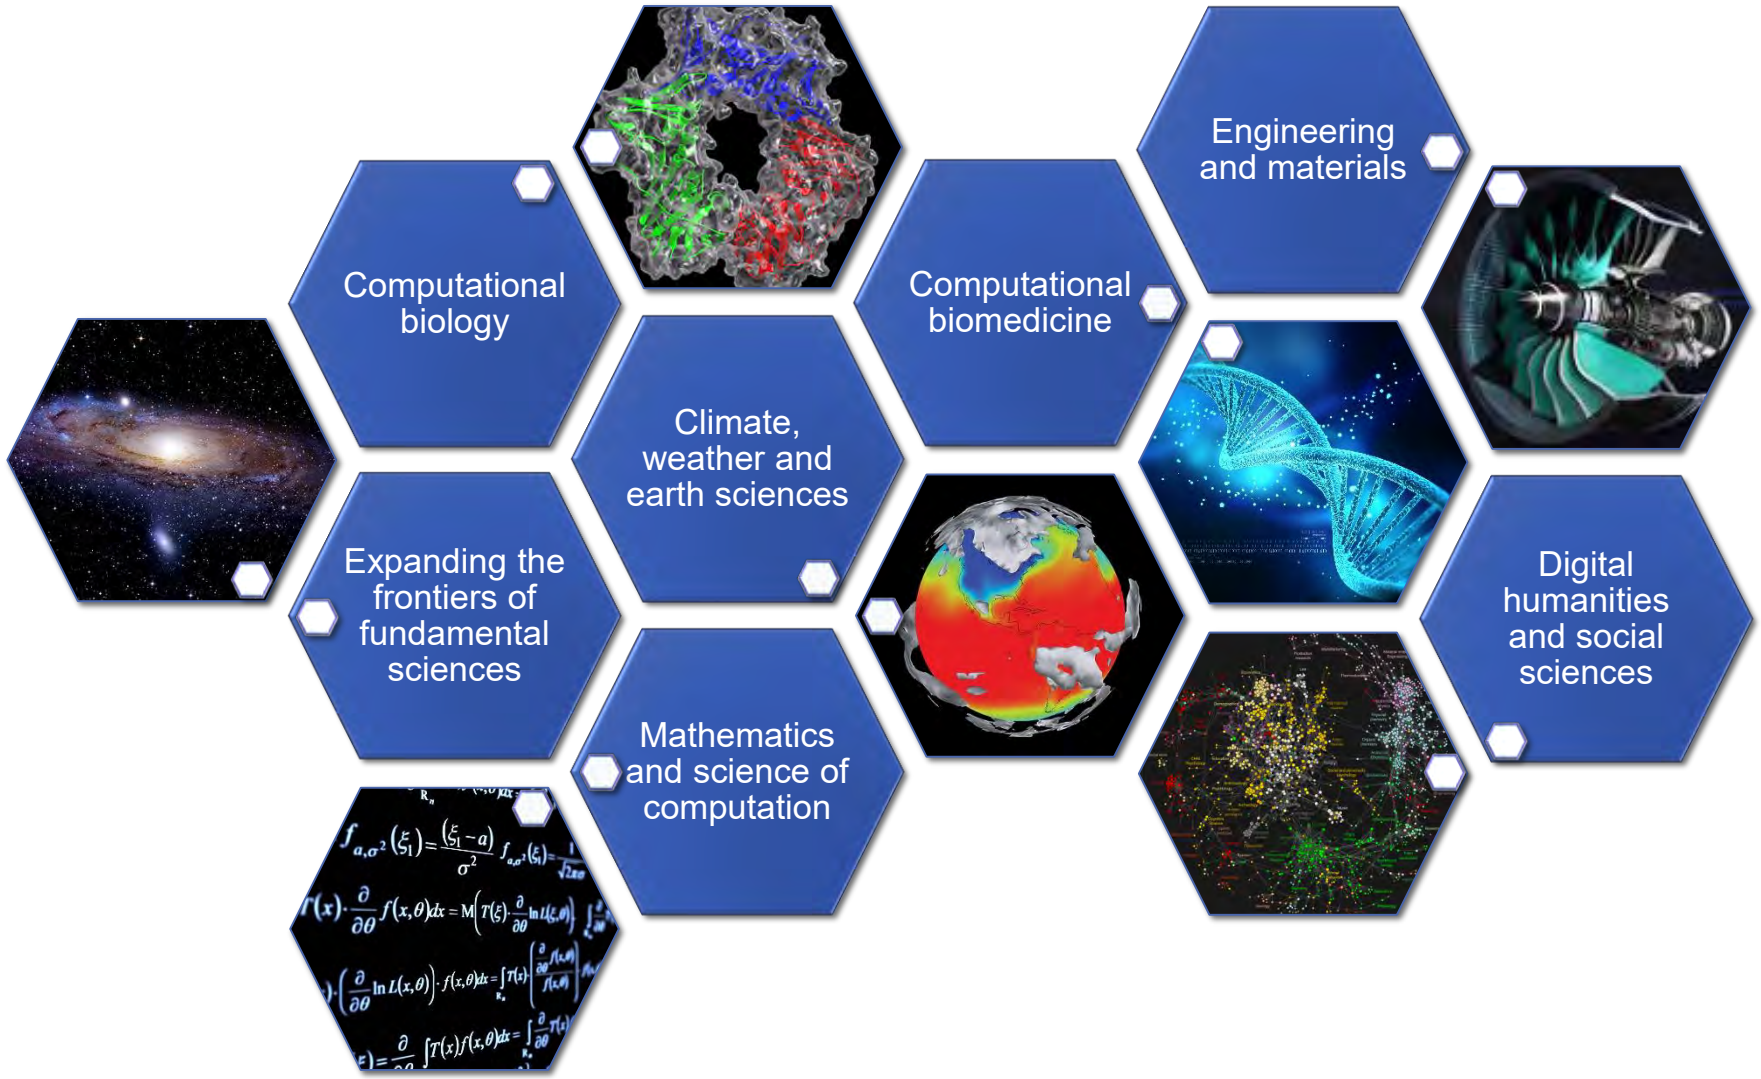

Making the case – Specific Requirements from Government

- System should support both traditional Modelling & Simulation and Artificial Intelligence / Deep Learning applications
  - Technology choices may be impacted by this
  - But future technologies blur the distinction
- System should support both scientific user communities and industry users
  - A greater focus is proposed with regard to industry use for research
  - Pay-per-use production access will be supported
  - Specific support for SMEs
- System should be operational by Q4 2022 / Q1 2023

Technology – Building an Exascale system today

- USA's SUMMIT system is the world's fastest computer today according to Linpack
- 2,414,592 CPU core
- R_{peak} = 201 Petaflop
- Power consumption
- To reach the Exasca
  - 12 million CPU cores
  - 65 Megawatts
- … very high levels of parallelism

… but new technologies are coming

Technology – two routes to the Exascale

- Cores-only route
  - Favoured by Japan
  - Evidence this approach leads to longer lifetimes of systems
  - Some hardware support for AI
  - Larger power requirements and physical dimensions
- Cores plus accelerator route
  - Favoured by the USA
  - Traditional multi-core processors coupled to accelerator
  - Sweet spot seems to 10 cores per GPU – pushes towards 1 socket + 4 GPU blades at the Exascale

Technology – Japanese versus American model

- Japanese model has attractions but more difficult to sell to Government
- Lower peak performance – but much longer science lifetime
- Look at HPCG:

June 2019 HPCG Results

Rank	Site	Computer	Cores	HPL Rmax (Pflop/s)	TOP500 Rank	HPCG (Pflop/s)	Fraction of Peak
1	DOE/SC/ORNL USA	**Summit** – AC922, IBM POWER9 22C 3.07GHz, dual-rail Mellanox EDR Infiniband, NVIDIA Volta V100 IBM	2,414,592	148.600	1	2.926	1.5%
2	DOE/NNSA/LLNL USA	**Sierra** – S922LC, Power9 22C 3.1GHz, Mellanox EDR, NVIDIA Tesla V100 IBM / NVIDIA / Mellanox	1,572,480	94.640	2	1.796	1.4%
3	Riken Center for Computational Science Japan	**K computer** – , SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu	705,024	10.510	20	0.603	5.3%
4	DOE/NNSA/LANL/SNL USA	**Trinity** – Cray XC40, Intel Xeon E5-2698 v3 16C 2.3GHz, Aries, Intel Xeon Phi 7250 68C 1.4GHz Cray	979,072	20.159	7	0.546	1.3%

epcc

Technology – key insights from NDA briefings

- High Bandwidth Memory is coming
  - Most Exascale blades will have HBM
  - Many designs have no DRAM at all
- We have a three-way competition in the processor world
  - Intel v. AMD v. Arm
- GPUs are getting ever more powerful
  - The market is also going to broaden
- Cabinet energy densities are rocketing
  - Today's 80-100KW cabinets will be eclipsed by cabinets at 300KW+
  - Density of blades is a key battleground

Technology – How parallel? How large?

- For a cores only approach it should be possible to get to 1 Exaflop theoretical peak with 4-5 million cores
- A cores plus GPU approach will reduce the number of cores but overall parallelism will increase as GPUs have much higher parallelism
- 1 Exaflop power requirements range from 20MW to 160MW
- Size of systems is highly dependent on density of blades and cabinet design
  - Number of cabinets ranges from circa. 60 to over 800!
- Key metric is always usefulness for research (science / industry)

Funding – What will an Exaflop system cost?

- Need to consider:
  1) Hardware costs including maintenance
  2) Accommodation costs
  3) Power costs
  4) Software and support programme costs

  } £700m – £1.2bn

- Also need to consider lifetime of system
  - Outline Business Case considers a 5 year and 8 year scenario
- Also need to consider that an Exascale system cannot do it all – continued (enhanced) investment in the rest of the e-Infrastructure ecosystem will be required

Hosting – Advanced Computing Facility

- The ‘ACF’ will host the WCDI
- Opened 2005
- Purpose built, secure, world class data centre
- Houses wide variety of leading-edge systems
- Major expansion in 2013
  - 8 MW power upgrade
  - 550m^2 machine room + 850m^2 plant room
- New £20m expansion started
  - Further machine room and another 24MW

H y

- T
- C
- I
- I
- I
- I

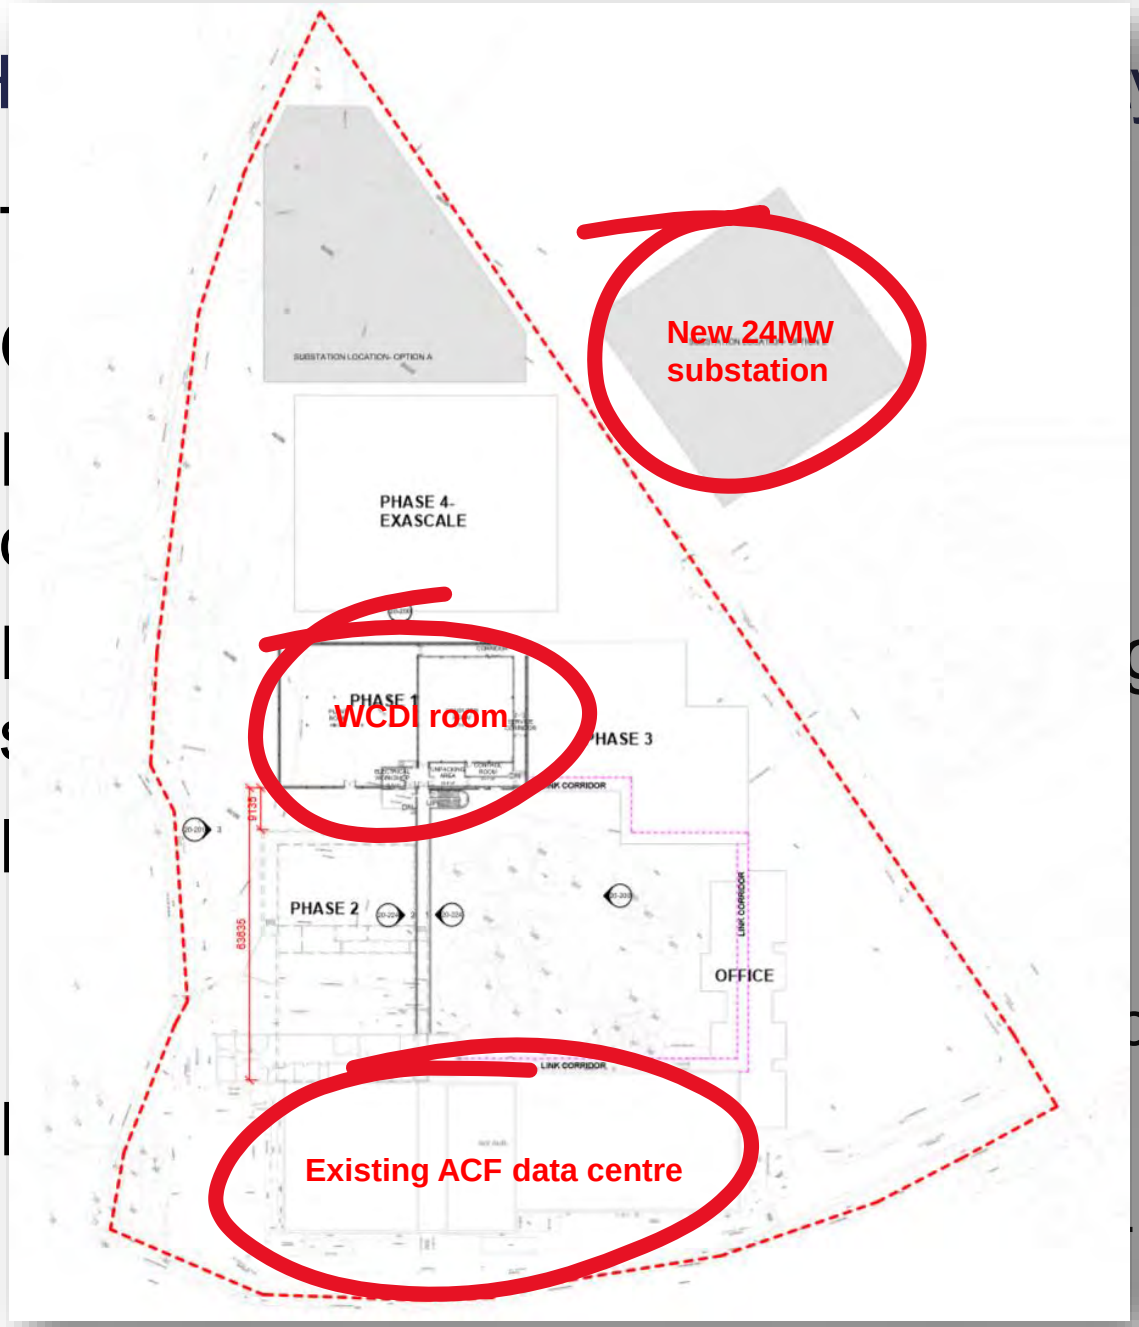

ge

om

MW

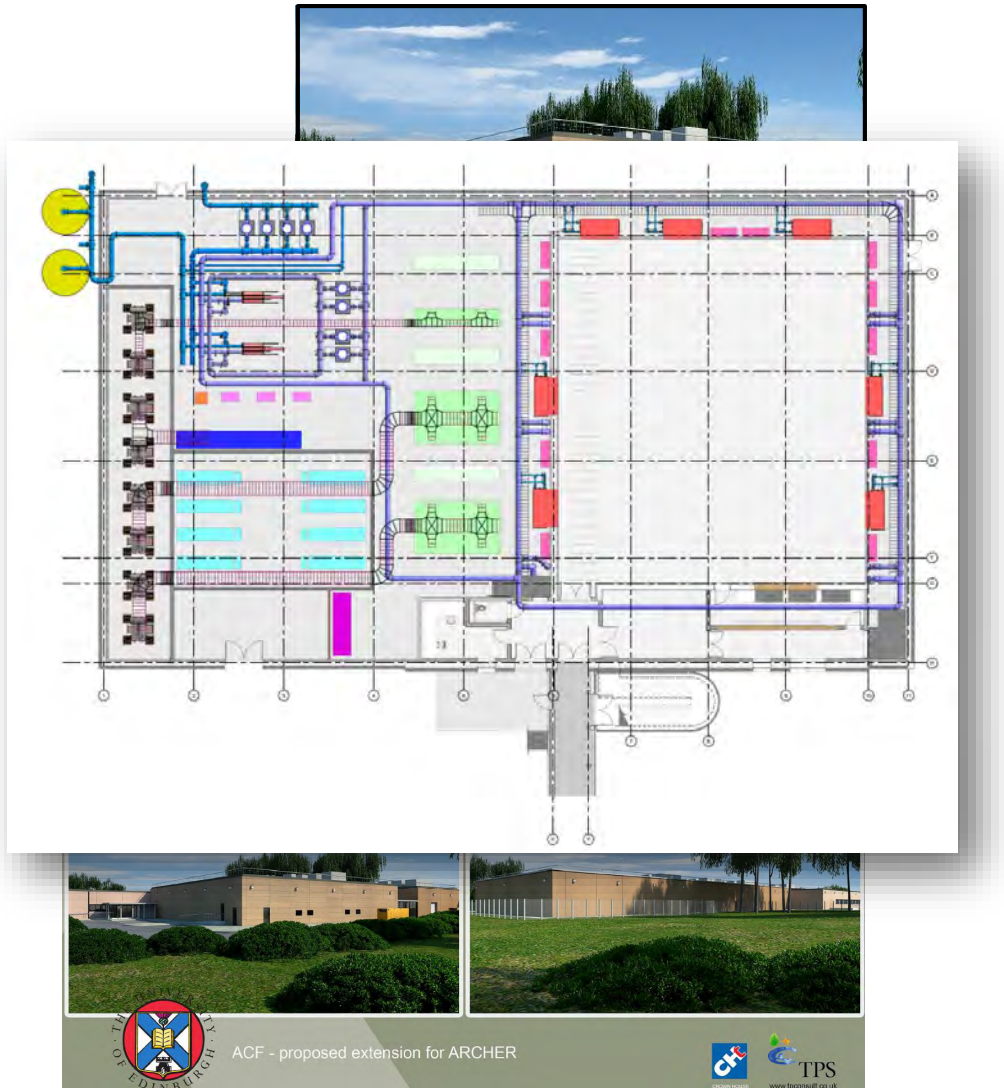

epcc

100% Renewable Energy

- The University of Edinburgh is part of the Scottish Public Procurement contract for electricity
- We choose the 100% renewable energy option

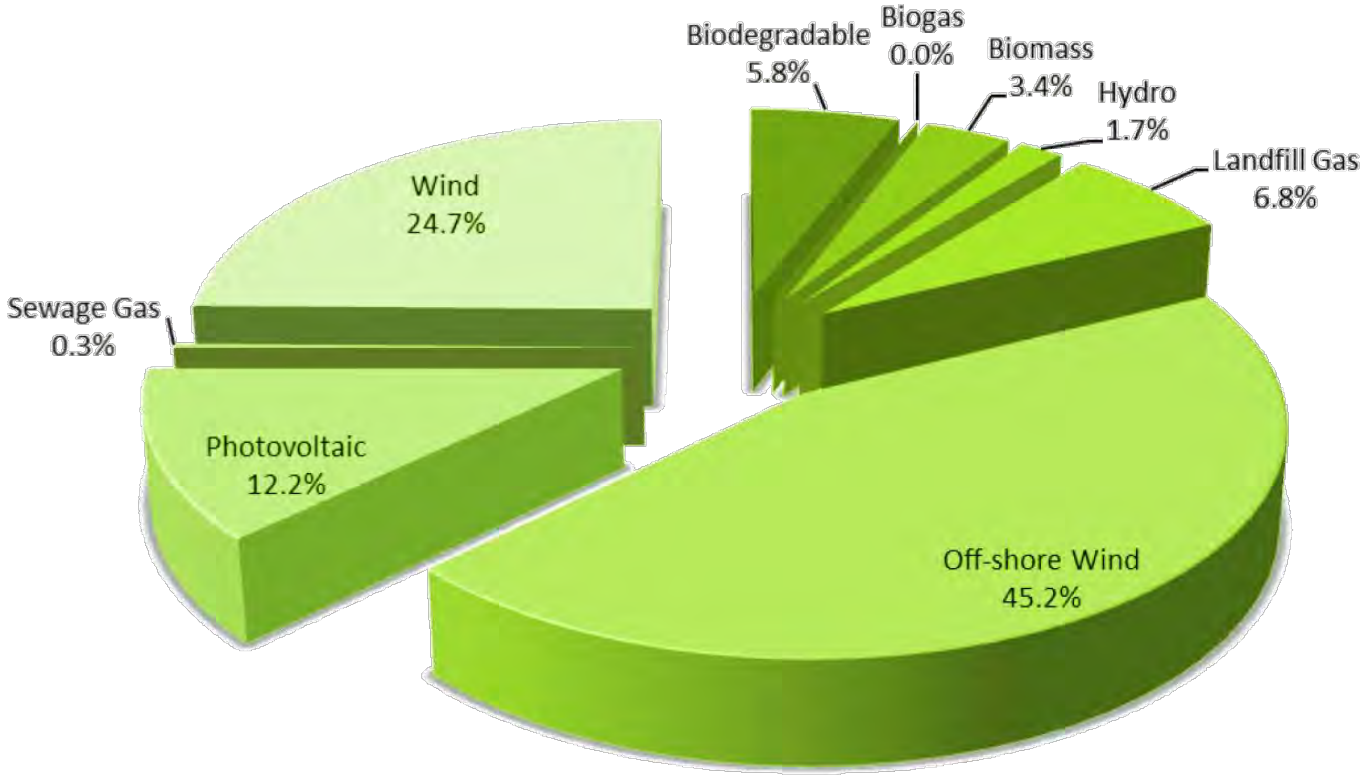

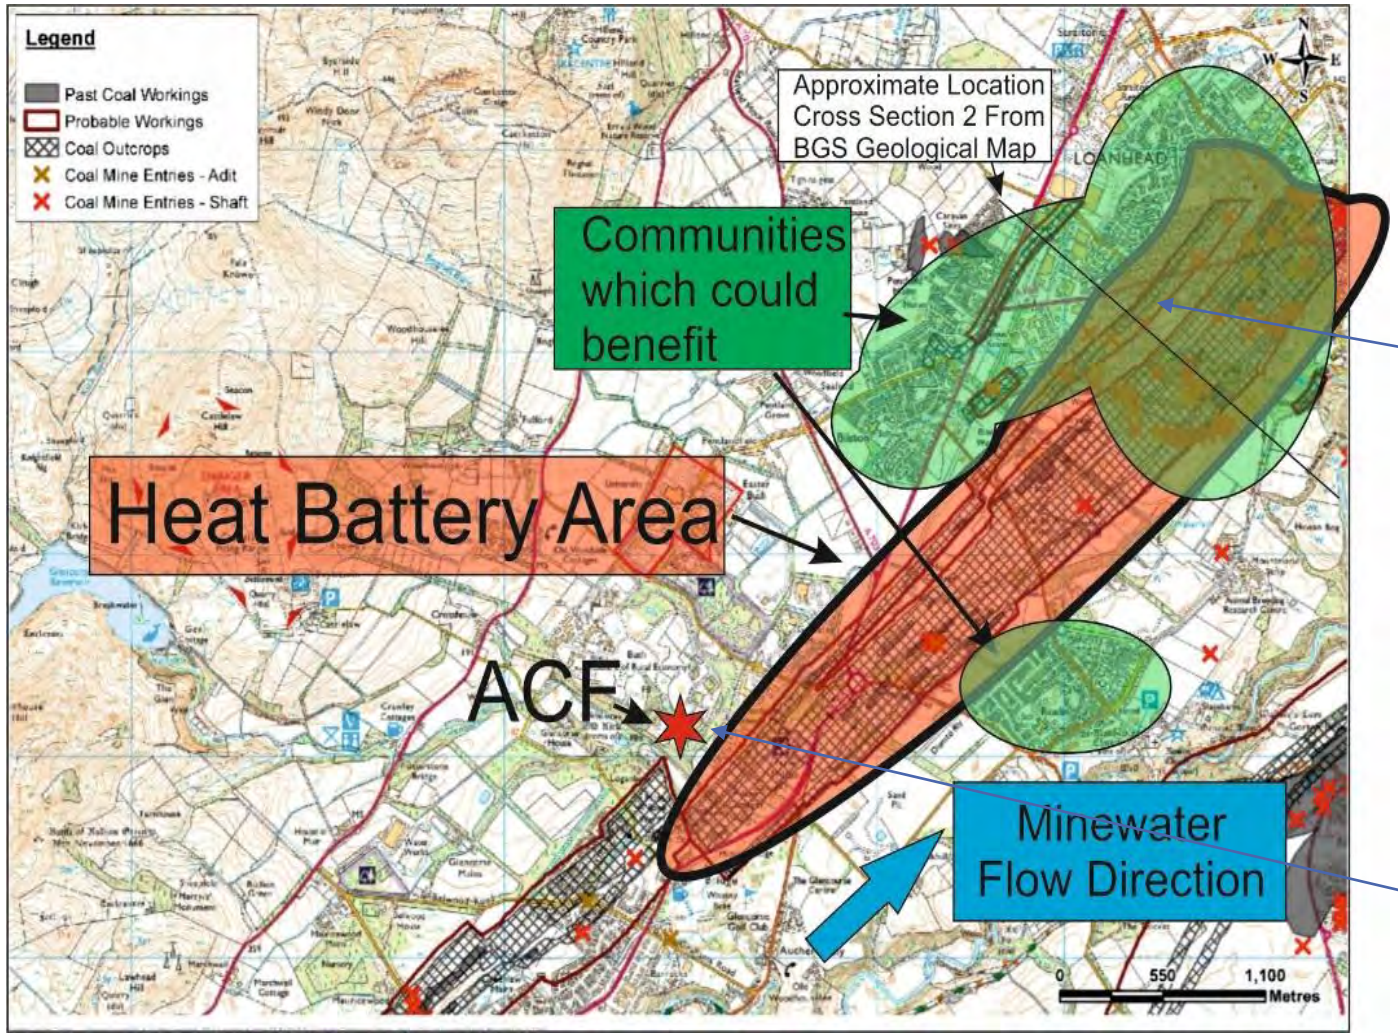

Bilston Glen - Shaft Details:
No. 1 shaft (NT 2714 6510)
7.3m diameter, 751m deep
No. 2 shaft (NT 2725 6507)
6.1m diameter, 733m deep
https://canmore.org.uk/site/72676/bilston-glen-colliery

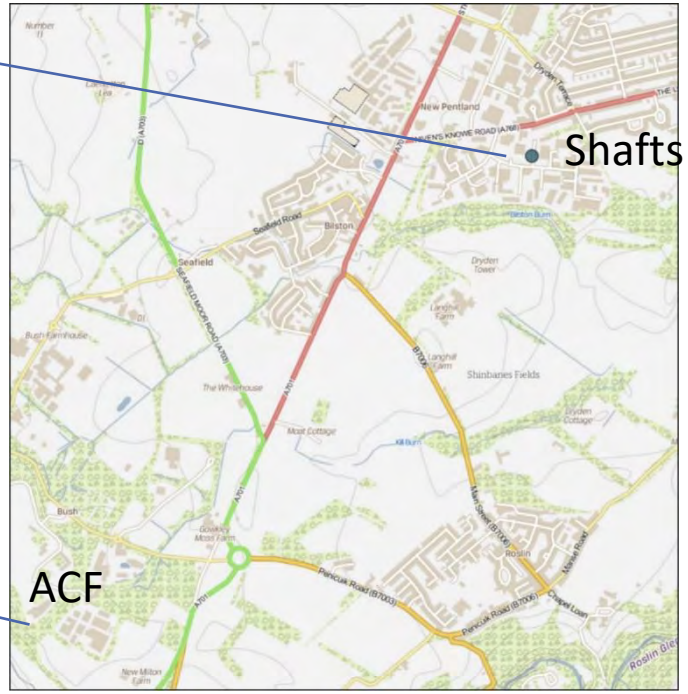

One of the National Coal Board's (NCB) most successful superpit developments, and designed to go much deeper than neighbouring mines into the Midlothian coalfield basin, exploiting the limestone coals, with an intended output of 1 million tons per annum. Commenced 1963, closed 1988.

Bilston Glen Colliery, 670m, 15.0C, Minewater
Monktonhall, 866m, 25.5C, Rock
Lady Victoria, 768m, 18C, Minewater

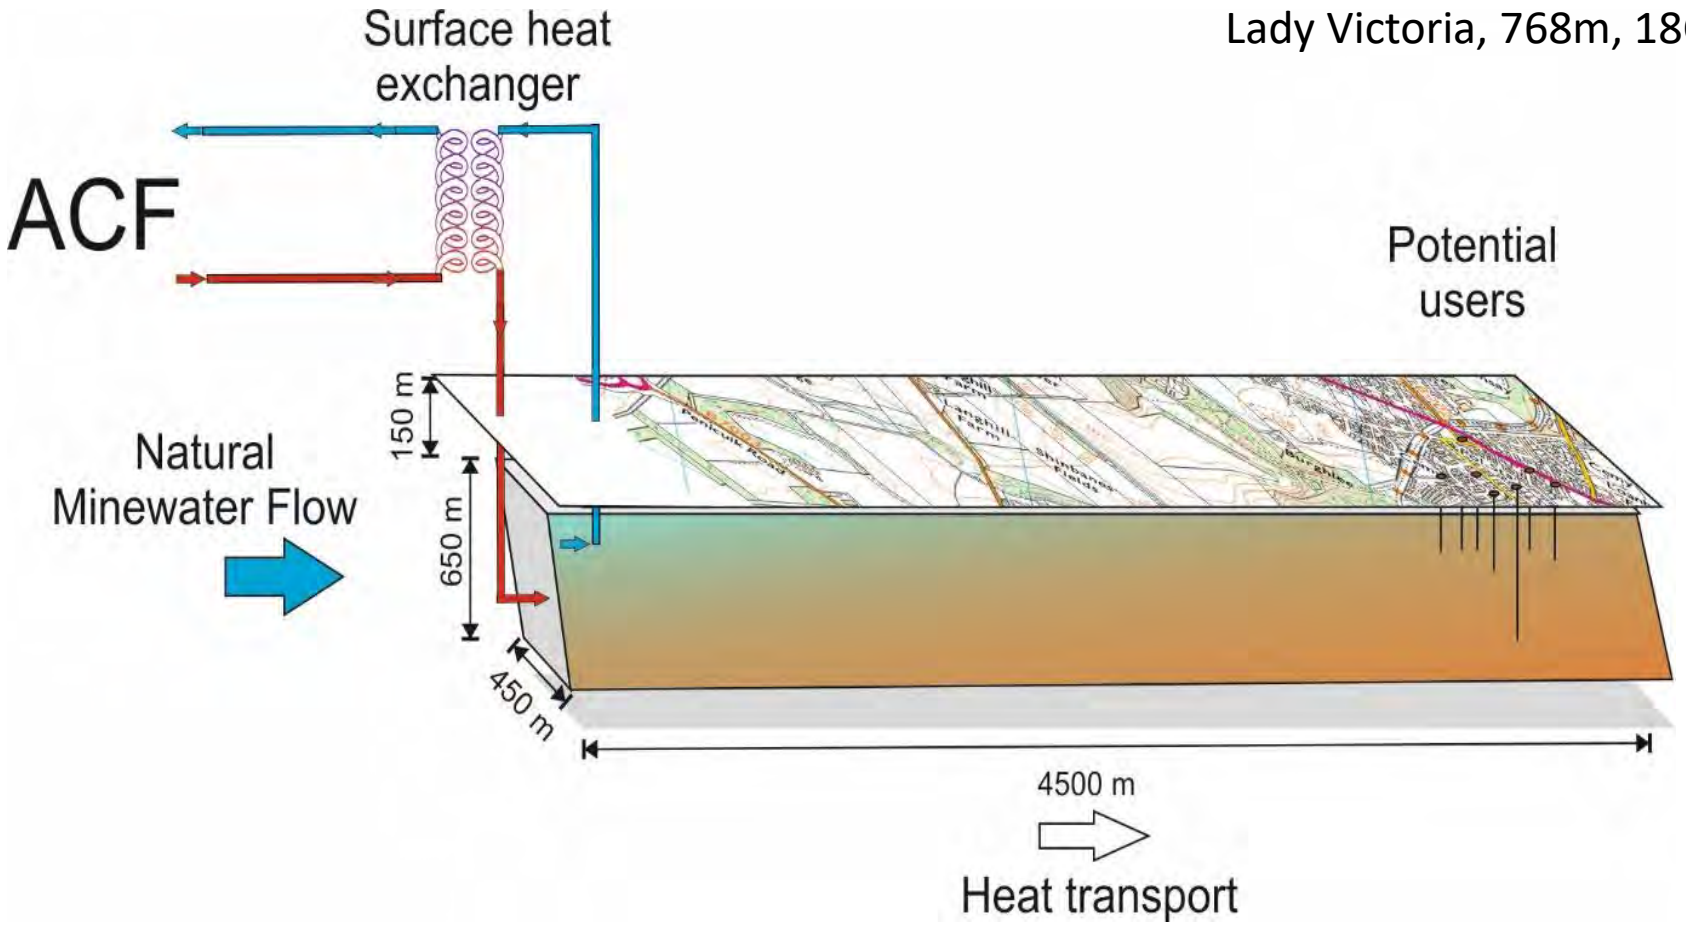

Summary

- The UK is taking the advent of the Exascale era seriously
- An Outline Business Case for a UK Exascale system has been prepared
- Aligned with the larger e-Infrastructure road mapping activity
- Activities have already started with regard to the hosting environment due to long lead times