



Early Experience from Running the UK's National AI Supercomputer Service

Prof. Simon McIntosh-Smith, Director & Founder,
Bristol Centre for Supercomputing (BriCS)



The Bristol Centre for Supercomputing: BriCS

- Established in **2024** to deliver **national AI and HPC services**
- We've grown rapidly to ~30 staff, drawn from HPC centres of excellence all over the world
- Based at the **University of Bristol**
- Delivering **over £300M of AI and HPC services** over the next 5 years
- Continuing to hire across all levels





DLC-CLEAN ENVIRONMENT. PLEASE AVOID CO

BriCS
Bristol Centre for Supercom

Department for
Science, Innovation
& Technology

NATIONAL
COMPUTERS
CENTRE

UNIVERSITY OF
BRISTOL

December 2025

Bristol Centre for Supercomputing systems

We've built **two new modular datacentres** and installed **three new HPE supercomputers** in less than 2 years (Q4 '23 – Q3 '25):

1. **Isambard 3** – 384 nodes of Grace-Grace CPUs, air cooled
2. **Isambard-AI phase 1** – 168 GH200 GPUs in EX2500, DLC
3. **Isambard-AI phase 2** – 5,280 GH200 in EX4000, DLC

All use HPE's Slingshot 11 high-speed network. Mixture of air and direct liquid cooled.

Isambard-AI phase 2



- **5,280 NVIDIA Grace-Hopper GPUs** in a new, 5MW HPE modular data centre (MDC) facility in Bristol, UK
 - **~21 ExaFLOP/s of 8-bit for AI, ~216 PFLOP/s 64-bit for HPC**
 - **PUE typically 1.06-1.08, powered entirely by local renewable electricity**
 - **Ranked #11 in the world** on the Top500 November 2025
- **Extremely rapid deployment**
 - Site constructed from scratch and supercomputer installed in ~12 months
 - ~2 months for bringup and acceptance



Department for
Science, Innovation,
& Technology

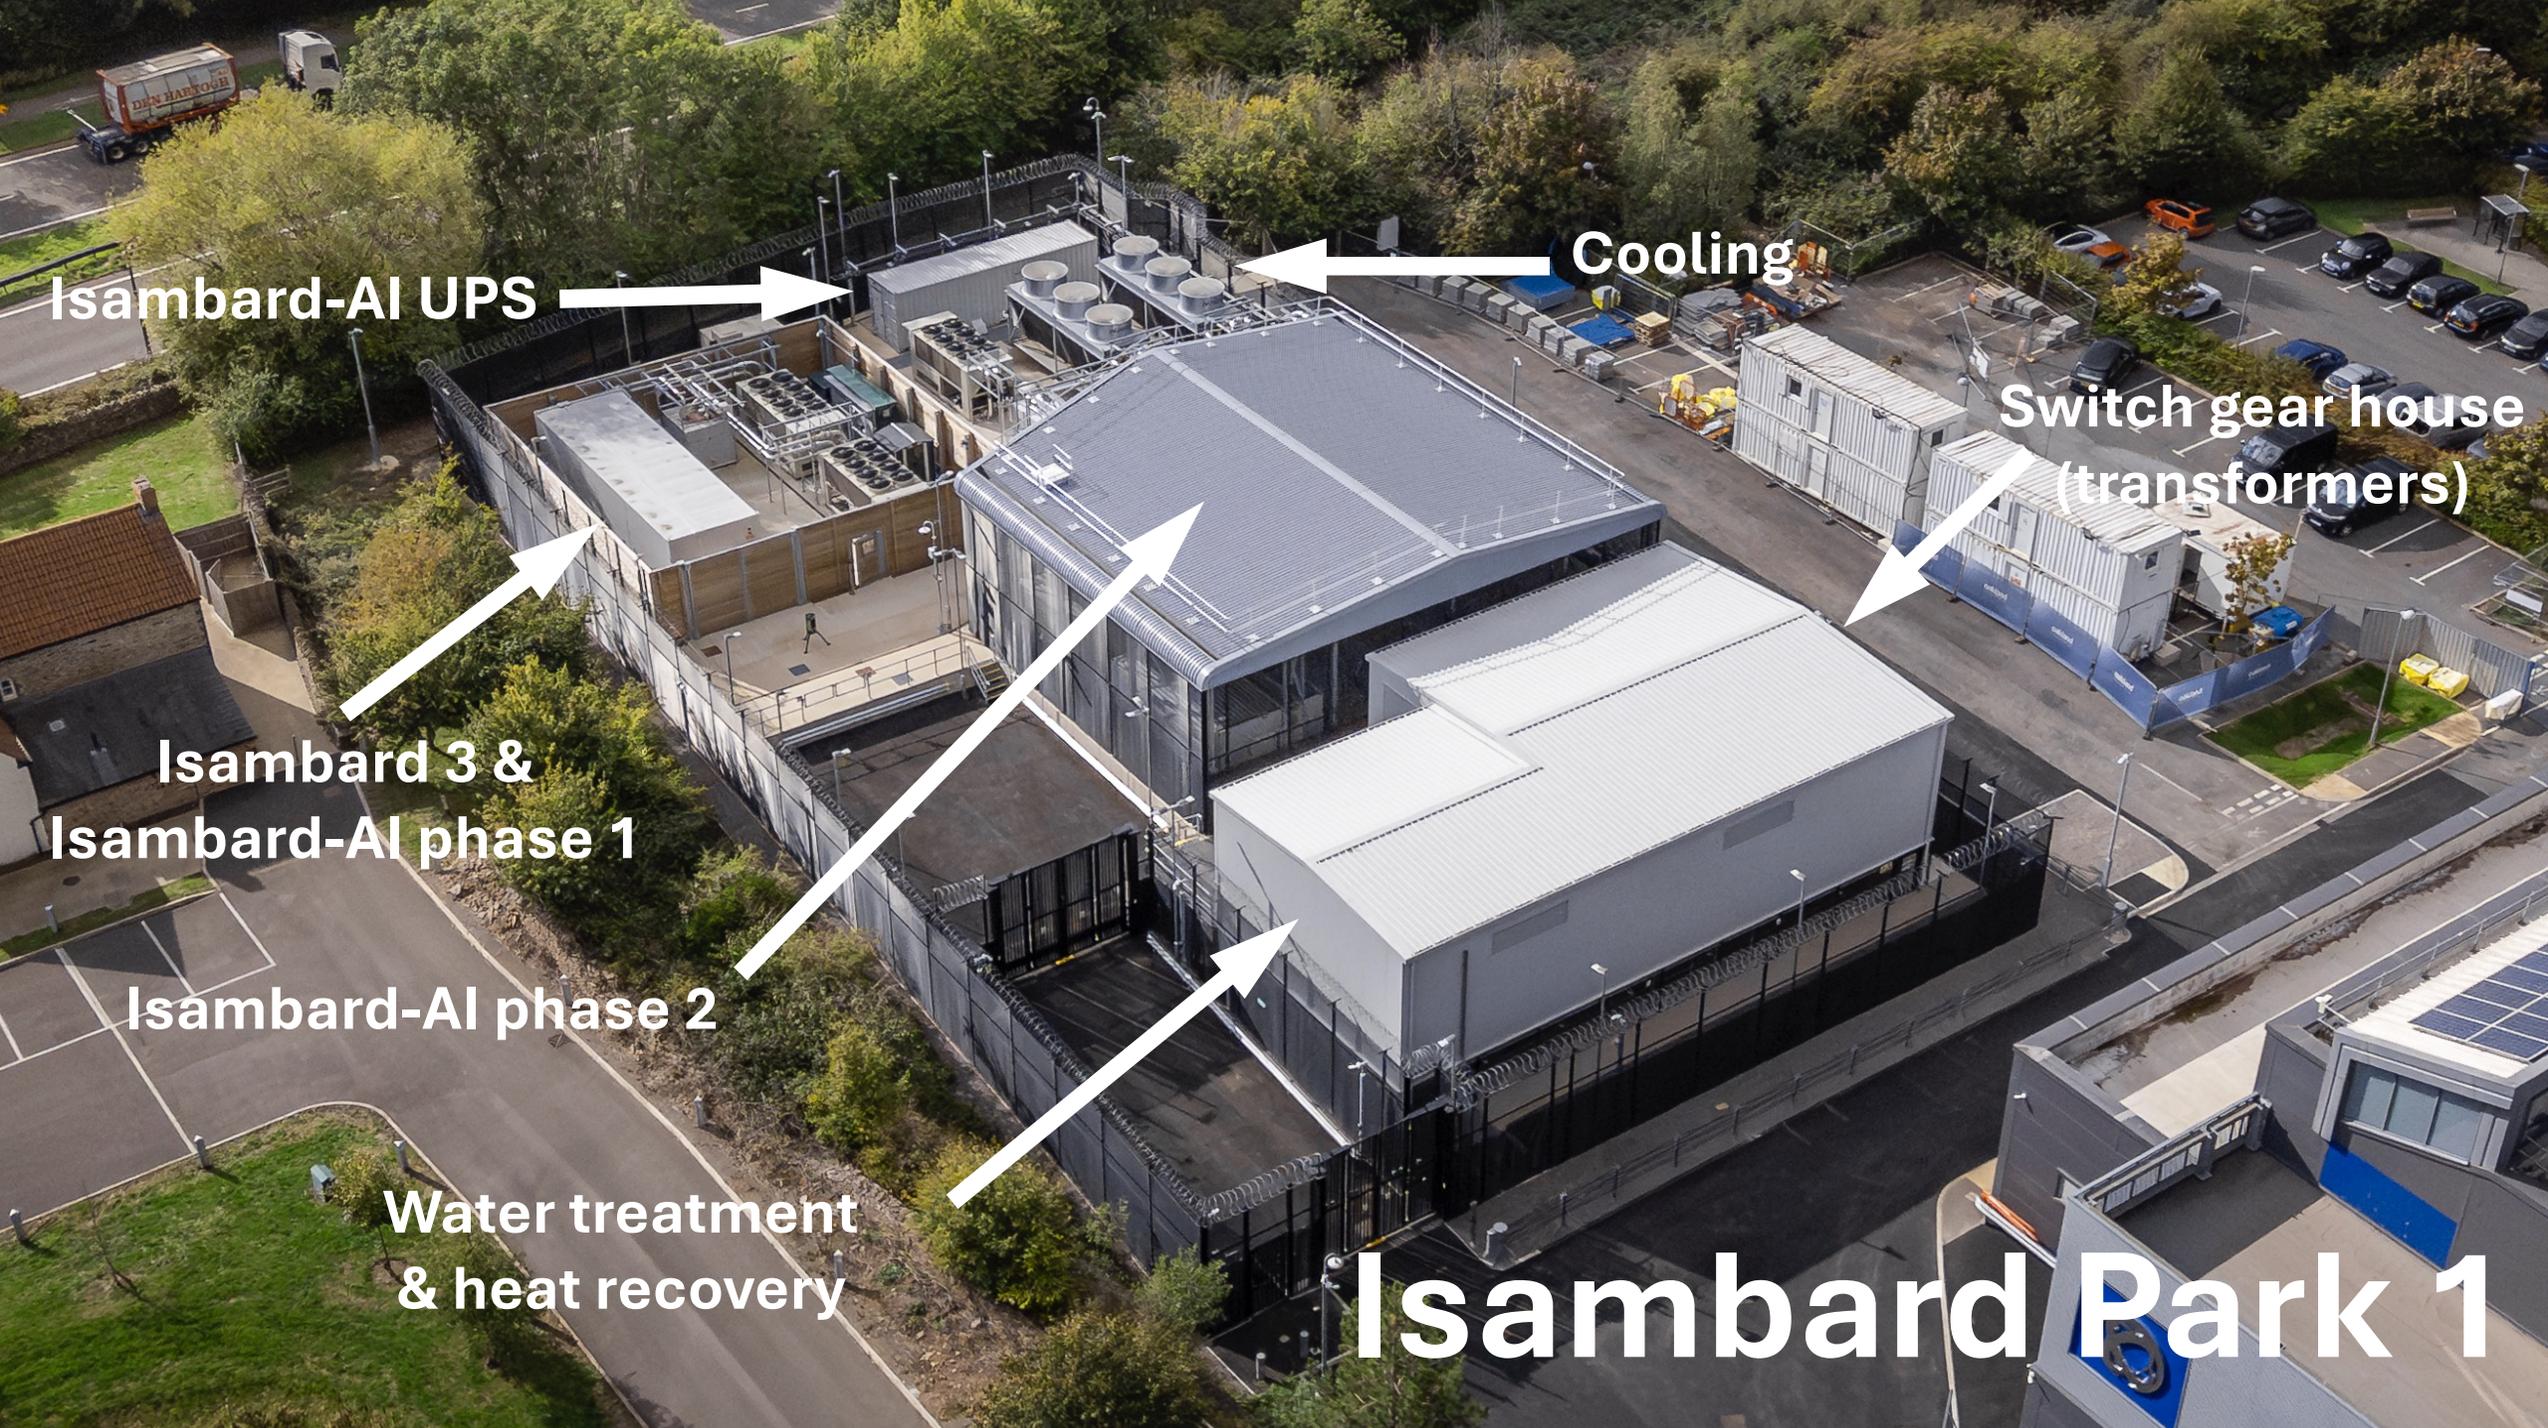
Isambard-AI site build: ground broken 3rd June 2024



Isambard-AI site build: ground broken 3rd June 2024



18-Feb-25 12:10



Isambard-AI UPS

Cooling

Switch gear house
(transformers)

Isambard 3 &
Isambard-AI phase 1

Isambard-AI phase 2

Water treatment
& heat recovery

Isambard Park 1



January 22nd 2026: HRH Prince William visiting Isambard-AI

Early experiences running the AI service

Grown to **1,600 users** and **600 projects** since August last year.

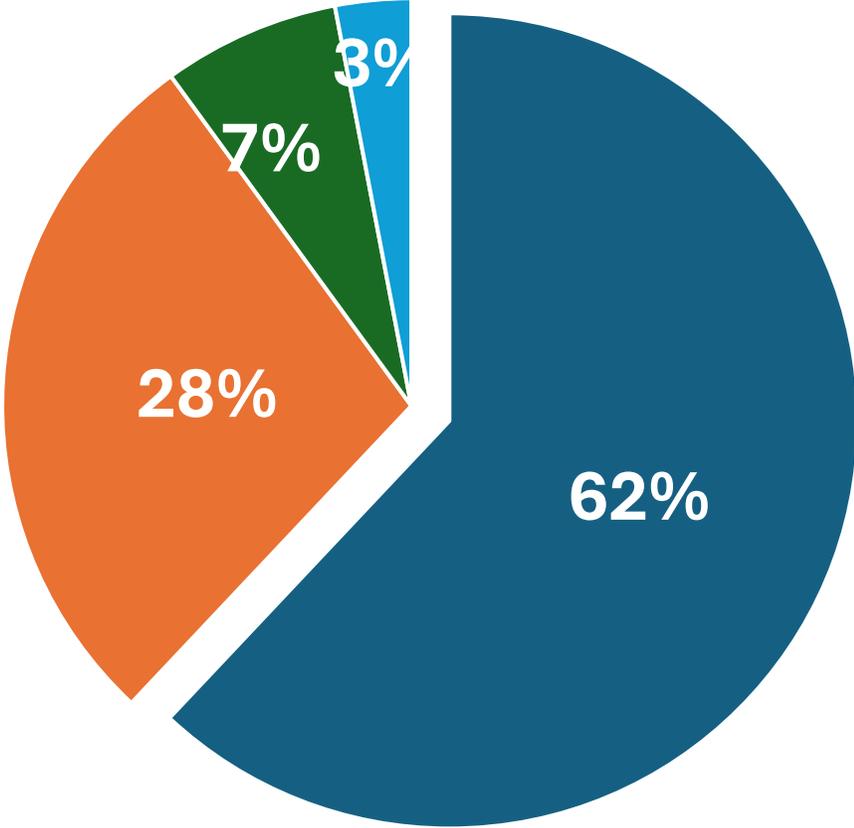
“AI first” users are not like HPC users.

- Lots of **interactive** users e.g. Jupyter Notebooks etc
- Preference for **Kubernetes** rather than SLURM
- Lots of use of standard packages containers
 - Most users aren't writing their own code
 - “AI for science” users are the exception

Usage across sectors

- Most usage by university-based research groups
- Government users next major user group
- Industry usage larger than expected this early in the ramp-up

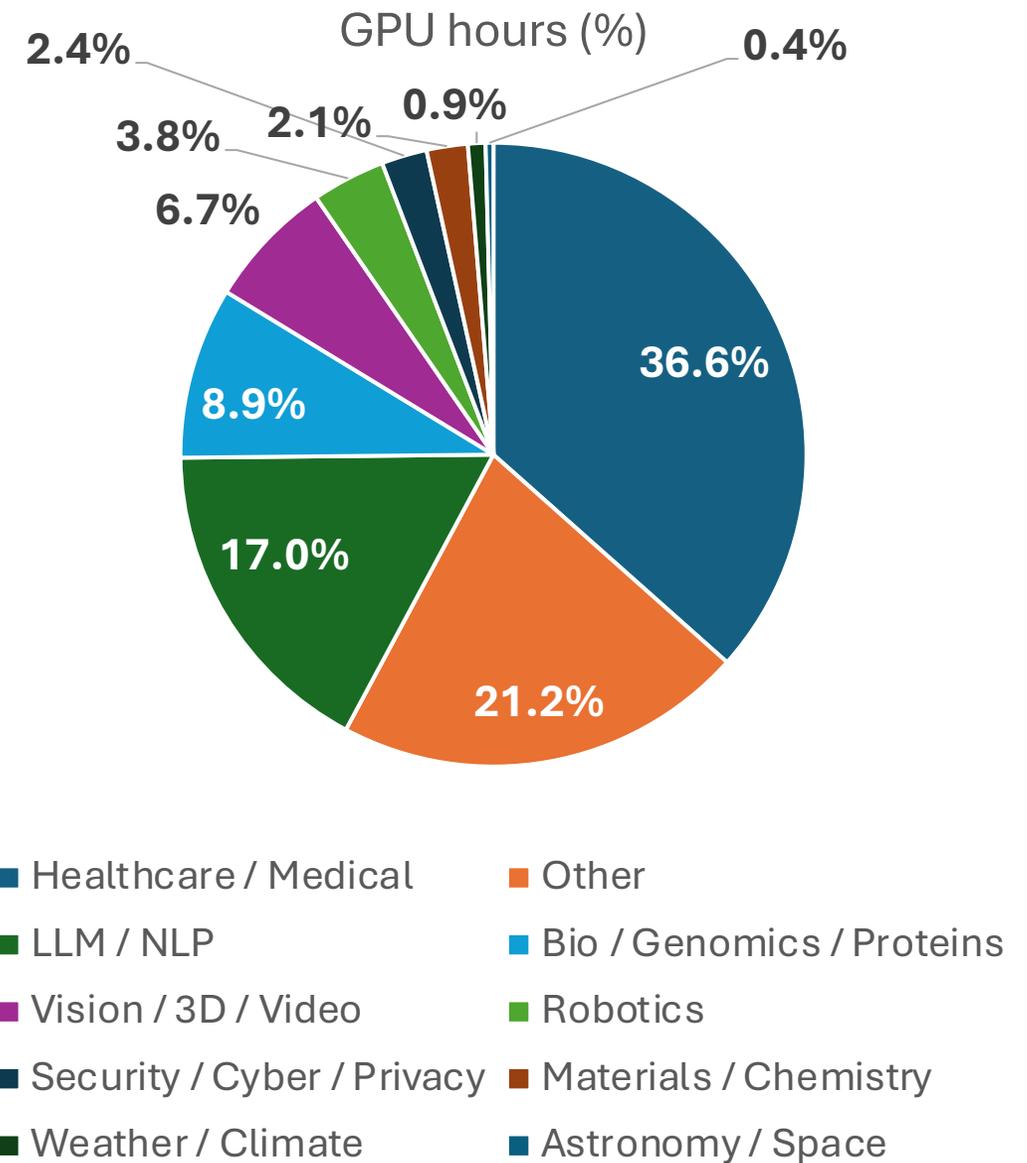
Usage by sector



- Academia
- Industry
- Public sector
- Other

Usage by theme

1. Healthcare / Medical $\approx 2,029,041$
2. Other (uncategorised large awards, e.g., some AISI / Sovereign systems work) $\approx 1,176,828$
3. LLM / NLP $\approx 944,313$
4. Bio/Genomics/Proteins $\approx 493,125$
5. Vision / 3D / Video $\approx 370,423$
6. Robotics $\approx 210,569$
7. Security/Cyber/Privacy $\approx 130,375$
8. Materials / Chemistry $\approx 117,923$
9. Weather / Climate $\approx 50,260$
10. Astronomy / Space $\approx 22,250$

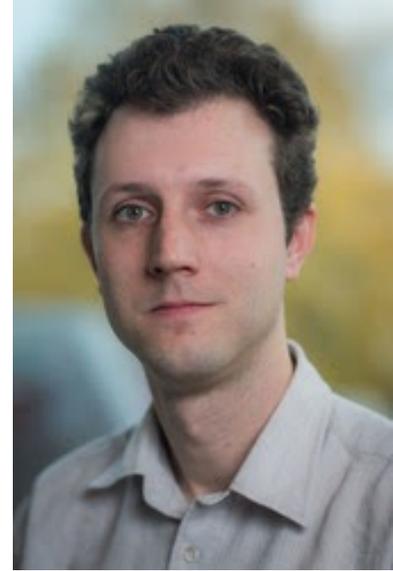


What does AI (Copilot) make of all this?

- **Foundation models everywhere** (vision, language, multimodal, protein/genomics, climate, simulation).
- **Generative simulation/world models** across robotics, fluids, weather/ocean, and materials.
- **Healthcare & life sciences**: imaging FMs, radiology/pathology harmonisation, variant effect prediction, protein design, antibody/TCR design.
- **Climate & environment**: hybrid AI-physics climate models, retrieval-augmented VLMs for EO, geophysical & ocean health FMs.
- **AI safety, robustness, and evaluation**: jailbreaks, routing stability for MoE, bias/fairness, trustworthy reasoning, data poisoning, monitoring/oversight.
- **Robotics/vision**: 3D/4D reconstruction, egocentric HOI, gesture generation, VLA models, digital twins.

Early AI science outputs from Isambard

A sovereign LLM for the UK – UK-LLM



- Creating an LLM based on open source, open data
- Trained across multiple languages used in the British Isles:
 - English, Scots, Welsh, Scottish Gaelic, Irish, etc
- Also trained on UK legal case law etc
- Led by Prof. Pontus Stenetorp at UCL, London
- Scaling up to thousands of GPUs on Isambard-AI is enabling the BritLLM team to scale their models to compete with large **open** LLMs seen today, e.g. LLM-jp at 172B and OLMO at 32B



Using OpenFold to improve healthcare



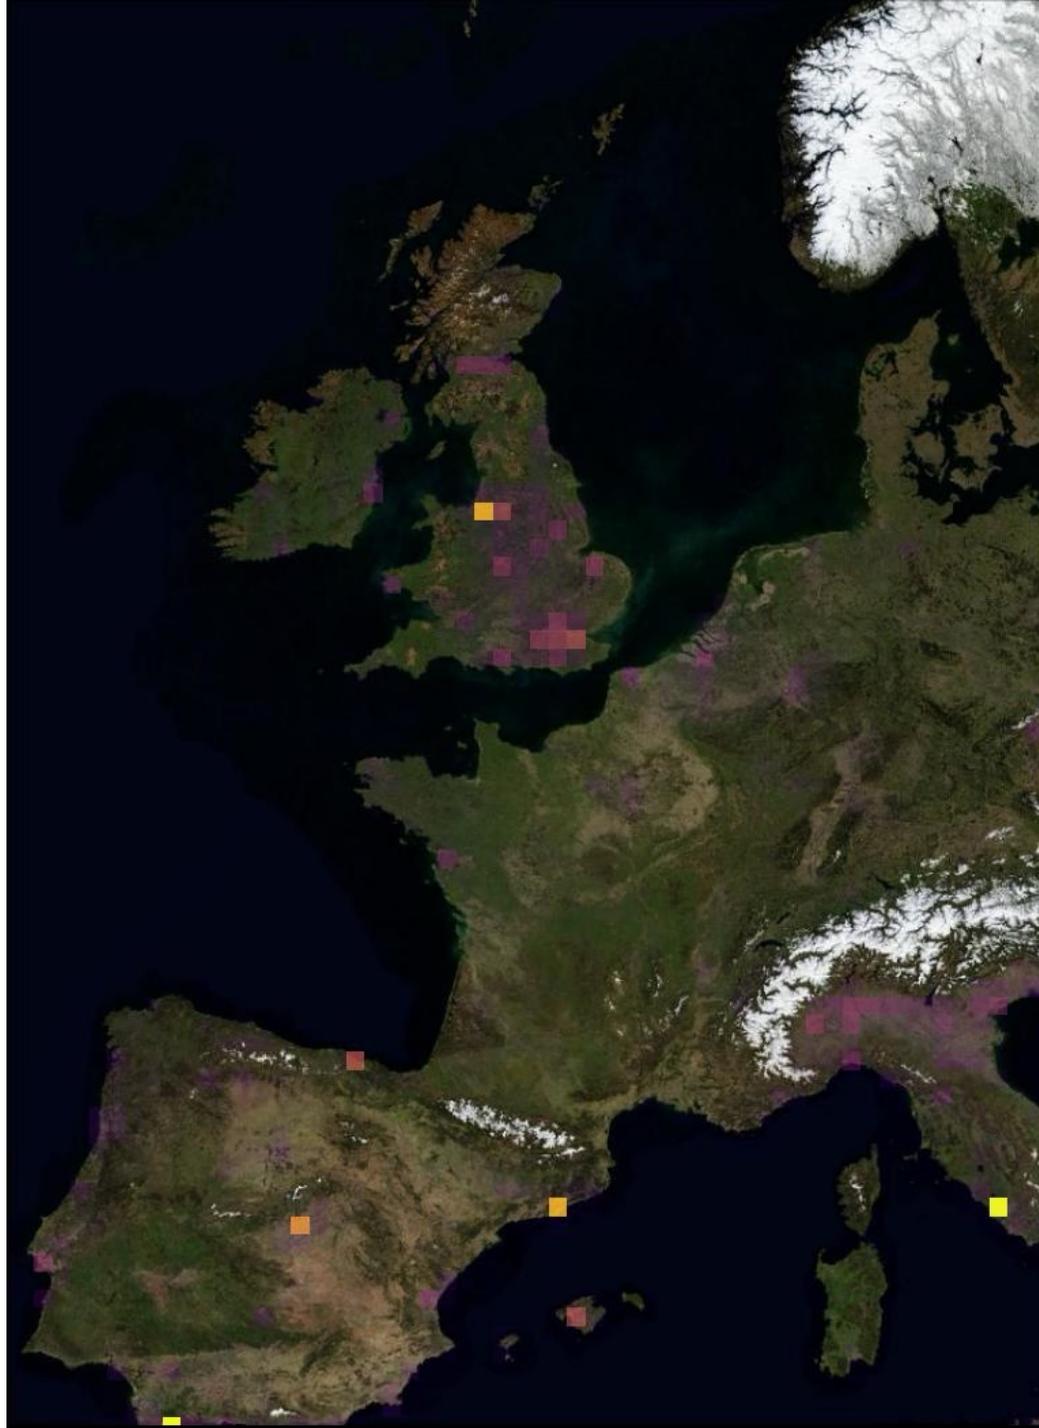
- Dr Jon Lees and Dr Danielle Paul at the University of Bristol are using OpenFold on Isambard-AI to understand how proteins interact within cells to cause disease
- The power of Isambard-AI allows them to start to reconstruct the cellular networks of proteins at unprecedented scales
- For example, they are studying a huge filamentous protein - the cardiac thin filament - a major component of heart muscle, to understand diseases such as hypertrophic cardiomyopathy, the most common inherited heart condition and leading cause of sudden adult death in the young

PolluGen: AI-based pollution forecasting



- A new high-resolution pollution dispersion model developed by the University of Manchester, using NVIDIA CorrDiff and Earth-2 Studio
- The model can use regional health and other data to help inform citizens and policymakers of air-quality impacts
- A pre-trained generative model is conditioned with low-resolution simulation data to infer a high-resolution output
- Led by Dr David Topping, University of Manchester







zebra: 9 ; conf: 0.89



zebra: 12 ; conf: 0.87



zebra: 10 ; conf: 0.88



zebra: 7 ; conf: 0.89



zebra: 1 ; conf: 0.89



zebra: 1 ; conf: 0.9



zebra: 8 ; conf: 0.89



zebra: 4 ; conf: 0.92



zebra: 13 ; conf: 0.87



zebra: 3 ; conf: 0.92



What's coming next in Bristol?



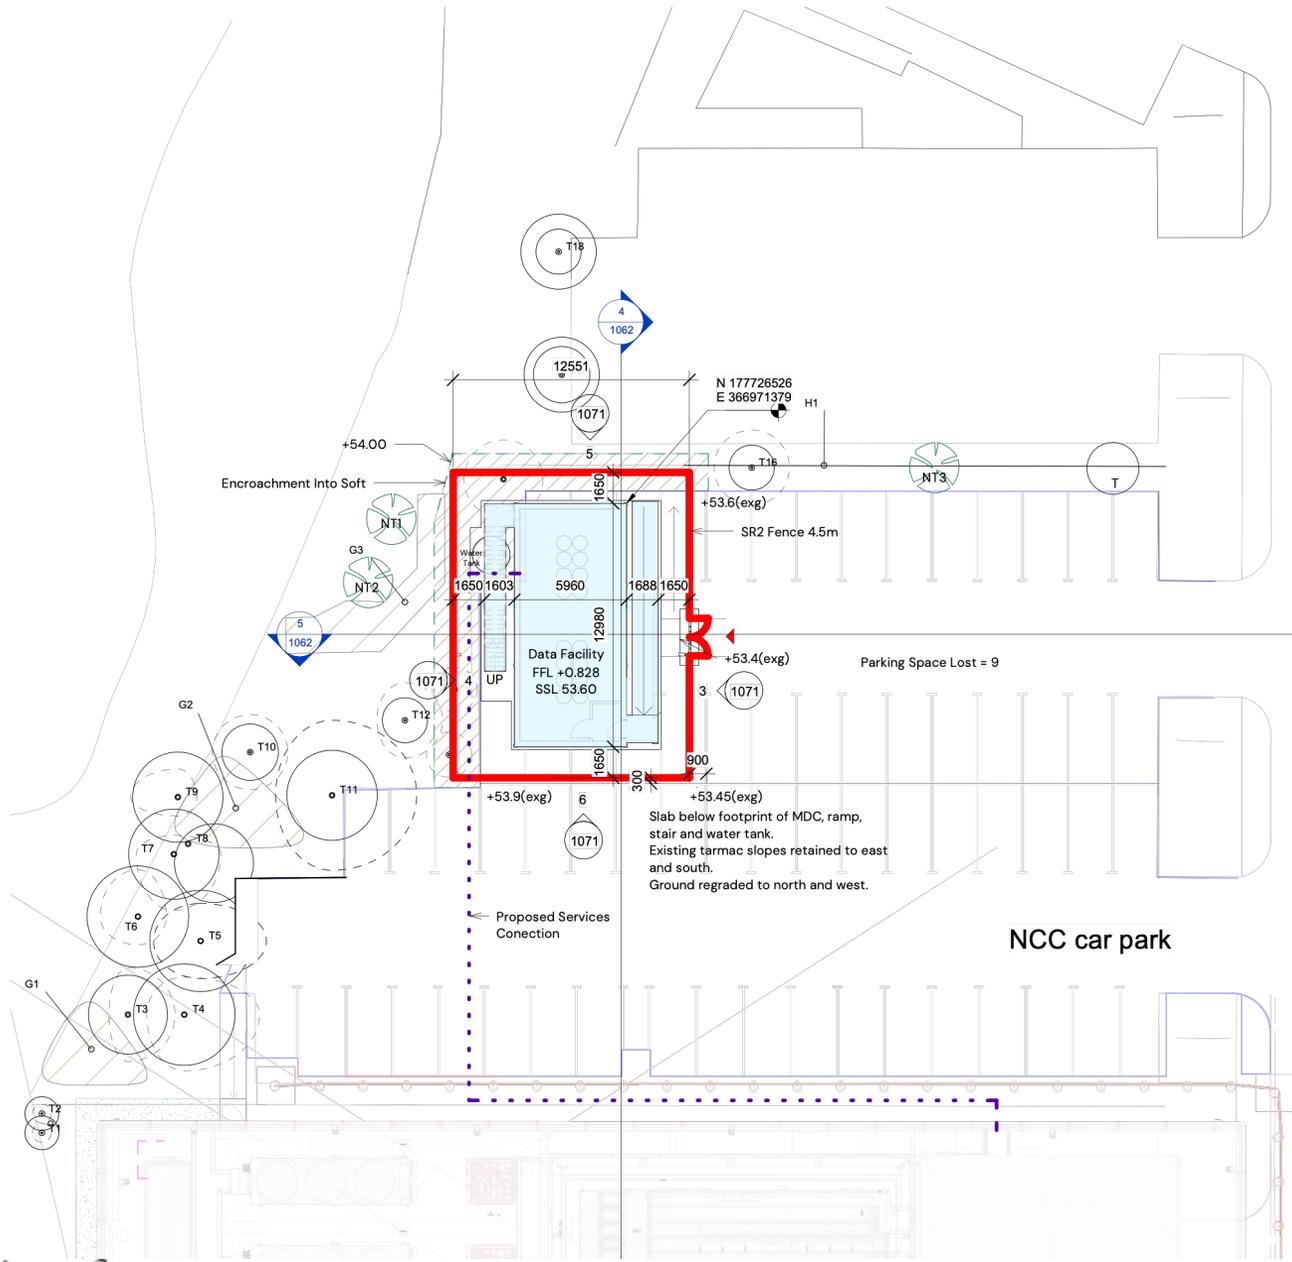
Isambard Park 2: 4-acres, 50+MW

**AI Data Facility
BRAID**

**Isambard Park 1:
0.5-acres, 8.6 MW**

Bristol AI Data Facility (BRAID)

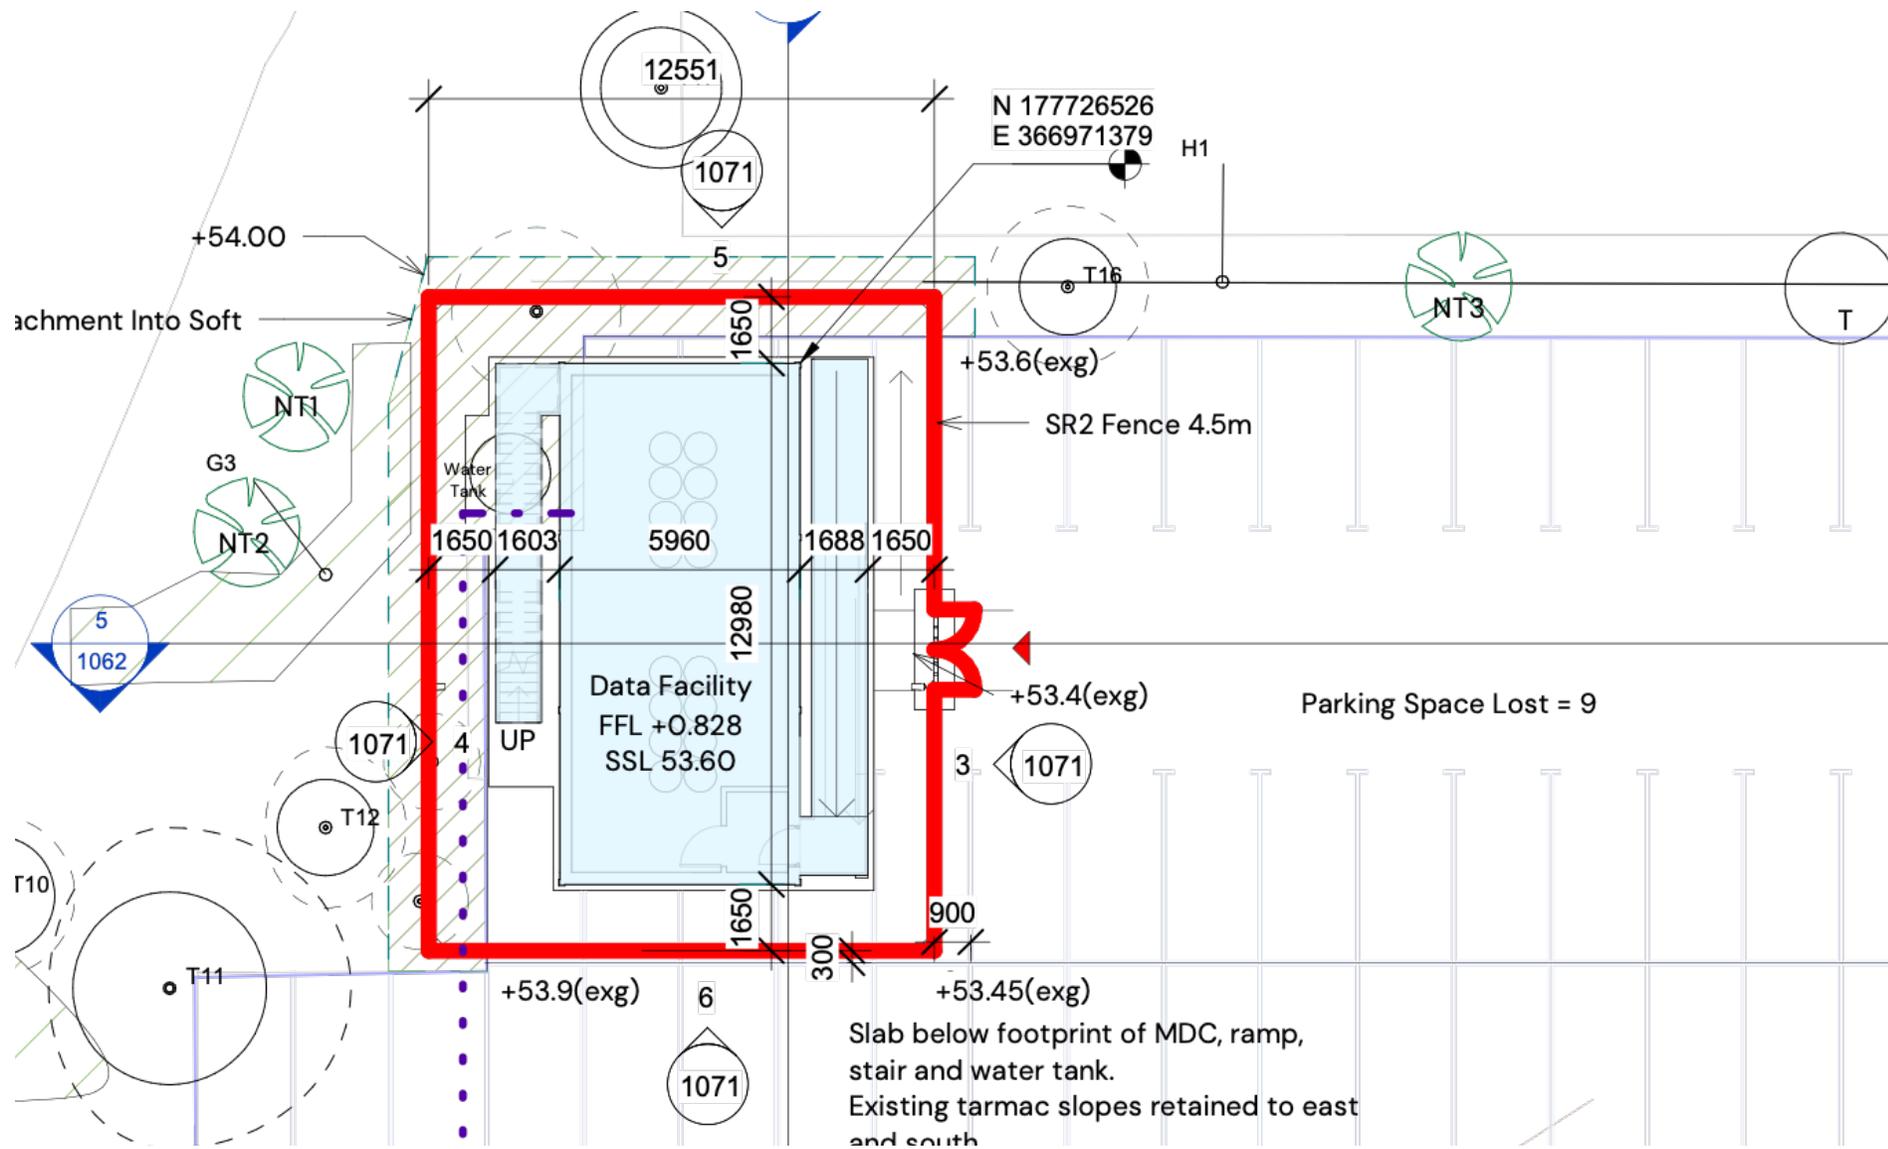
- A new, ~50 PByte AI data lake alongside Isambard-AI
- All solid-state for speed and flexibility
- Embedded Kubernetes cluster for processing
- Room to expand rapidly to ~150 PBytes
- New Dell MDC
- High bandwidth to local compute – hundreds of Gbps
- 100 → 400 Gbps into national network (JANET)
- Going live summer 2026



Proposed Plan Key

- Loss of parking 9
- Proposed Slab Level 53.60
- Cut a fill requirements to be confirmed
- Site boundary to be formed by 4.5m SR2 rated fence
- Proposed Site Area 174m2
- - - - - Proposed Services Connection
- GI/HI/TI Tree cluster, Hedgerows and Trees, refer to Arb' Survey
- NT1 Proposed New tree locations (contained within 1m3 tree pit).

National Composite Centre



Summary

- AI-focused services are **not** just rebadged HPC services
 - Interactive use and small jobs are important
 - Many users will happily scale to thousands of GPUs too
- AI4S users are a hybrid of AI and HPC
- AI users are coming from a broad range of research areas
- **Usage is changing rapidly** and **will continue to do so**
- **Sovereign AI** capability increasingly important around the world during periods of increasing geopolitical uncertainty

Acknowledgements

Isambard-AI is funded by the UK Government's Department of Science, Innovation and Technology (DSIT) via UKRI / STFC and EPSRC.

<https://www.linkedin.com/company/brics-uob/>

<https://www.bristol.ac.uk/supercomputing/>