

Co-design, deployment and operation of a Modular Data Centre (MDC) with air and direct-liquid cooled (DLC) supercomputers

BriCS: Sadaf R Alam, Ross Barnes, Naomi Harris, Dave Moore,
Simon McIntosh-Smith

HPE: Emma Akinyemi, Janine Over, Martin Podstata

CUG 2025

Bristol Centre for Supercomputing (BriCS)

- The GW4 (Bath, Bristol, Cardiff and Exeter) **Isambard project** initially set out to prove that a new **ARM-based** technology was relevant to supercomputing since 2016
- **BriCS** has been recently formed (Q3 2024) for managing Isambard Digital Research Infrastructure (DRI) projects within the Faculty of Engineering
 - Isambard-AI (a GH200 superchip cluster)
 - >£200M including additional Modular Data Centre, all cooling etc.
 - >£300M total investment over 5 years
 - Isambard 3 (a Grace superchip cluster)





Isambard-AI Phase 1 168 GH200
SS11 Cray HPE EX2500 & Isambard
3 384 Grace Superchip SS11

Isambard-AI Phase 2
5280 GH200, SS11 DLC
Cray HPE EX4000



🤖 Why MDC for Isambard supercomputers? See Isambard-AI: a leadership class supercomputer optimised specifically for Artificial Intelligence. <https://arxiv.org/abs/2410.11199> for details

ISC2025 (Green500)



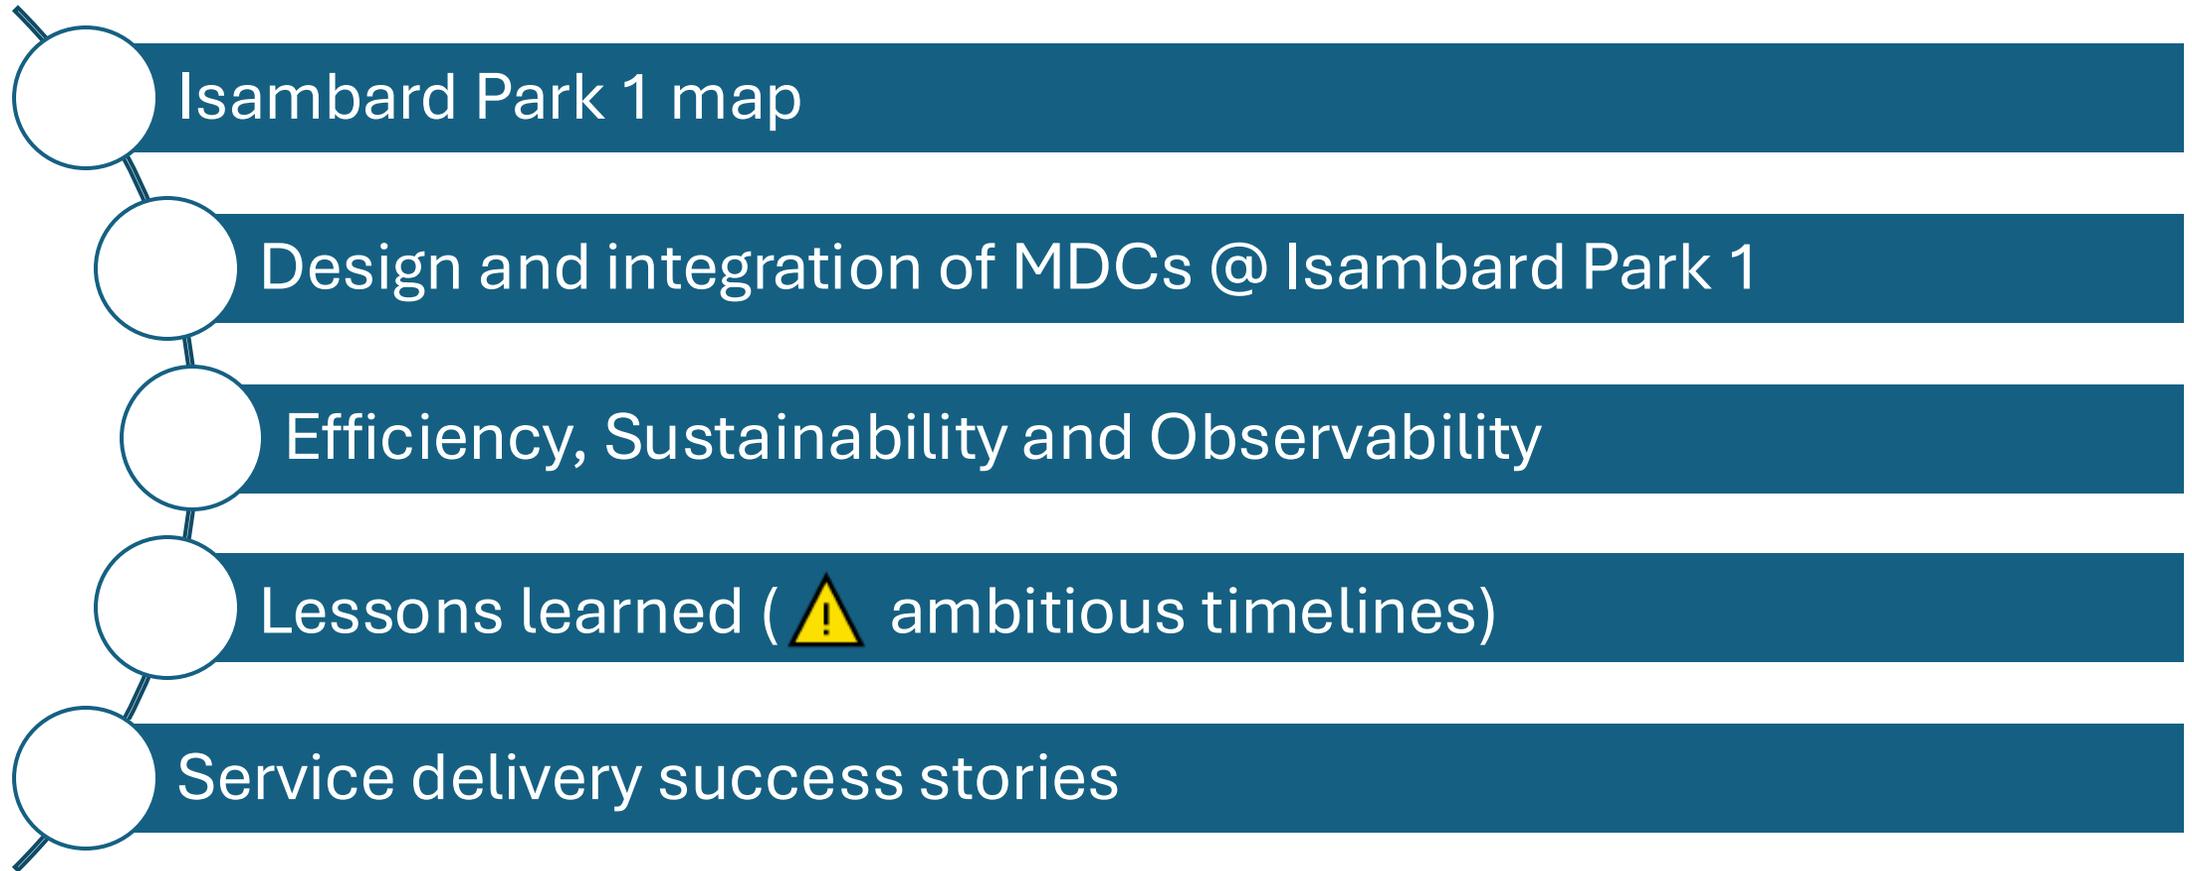
Bristol dominates uni cohort with double win at first AI awards

15 Sep 2024

The University of Bristol's leading role in artificial intelligence for higher education received public recognition with two wins in an inaugural awards ceremony for AI.



Outline



P2 cooling

P2 Power management unit (PMU)

P1 cooling

Air-cooled MDC with P1 EX2500 DLC and Isambard 3

NCC car park

P2 air-cooled MDC

P2 DLC data hall MDC

P2 switch house

POD Modules

Isambard 3 (Air cooled) and AI phase 1 (DLC cooled) =
POD DC20, operational since March'24

Isambard-AI phase 2 = DLC cooled and air-cooled
installed March'25

Fire detection and suppression

Biometric access control & CCTV

Power Modules

Isambard 3 and AI phase 1 (facility power 900 kW)

UPS backed by diesel generator

Isambard-AI phase 2 (facility power 4590 kW)

Power Management Unit (PMU)

UPS backed by diesel generator

Isambard Modular Data
Centre (MDC) template with
custom Data Centre
Infrastructure Manager
(DCIM) configuration

Cooling Modules

Isambard 3 and AI phase 1 = Freecooling chillers
(cooling capacity 2 x 430 kW), pump skid

Isambard-AI phase 2 = Freecooling chiller (cooling
capacity 430 kW) Adiabatic coolers (cooling capacity 2
x 2050 kW), pump skids

Heat recovery connection

Site integration

Isambard 3 and AI phase 1 = Utility Power available via
LV GRP enclosure

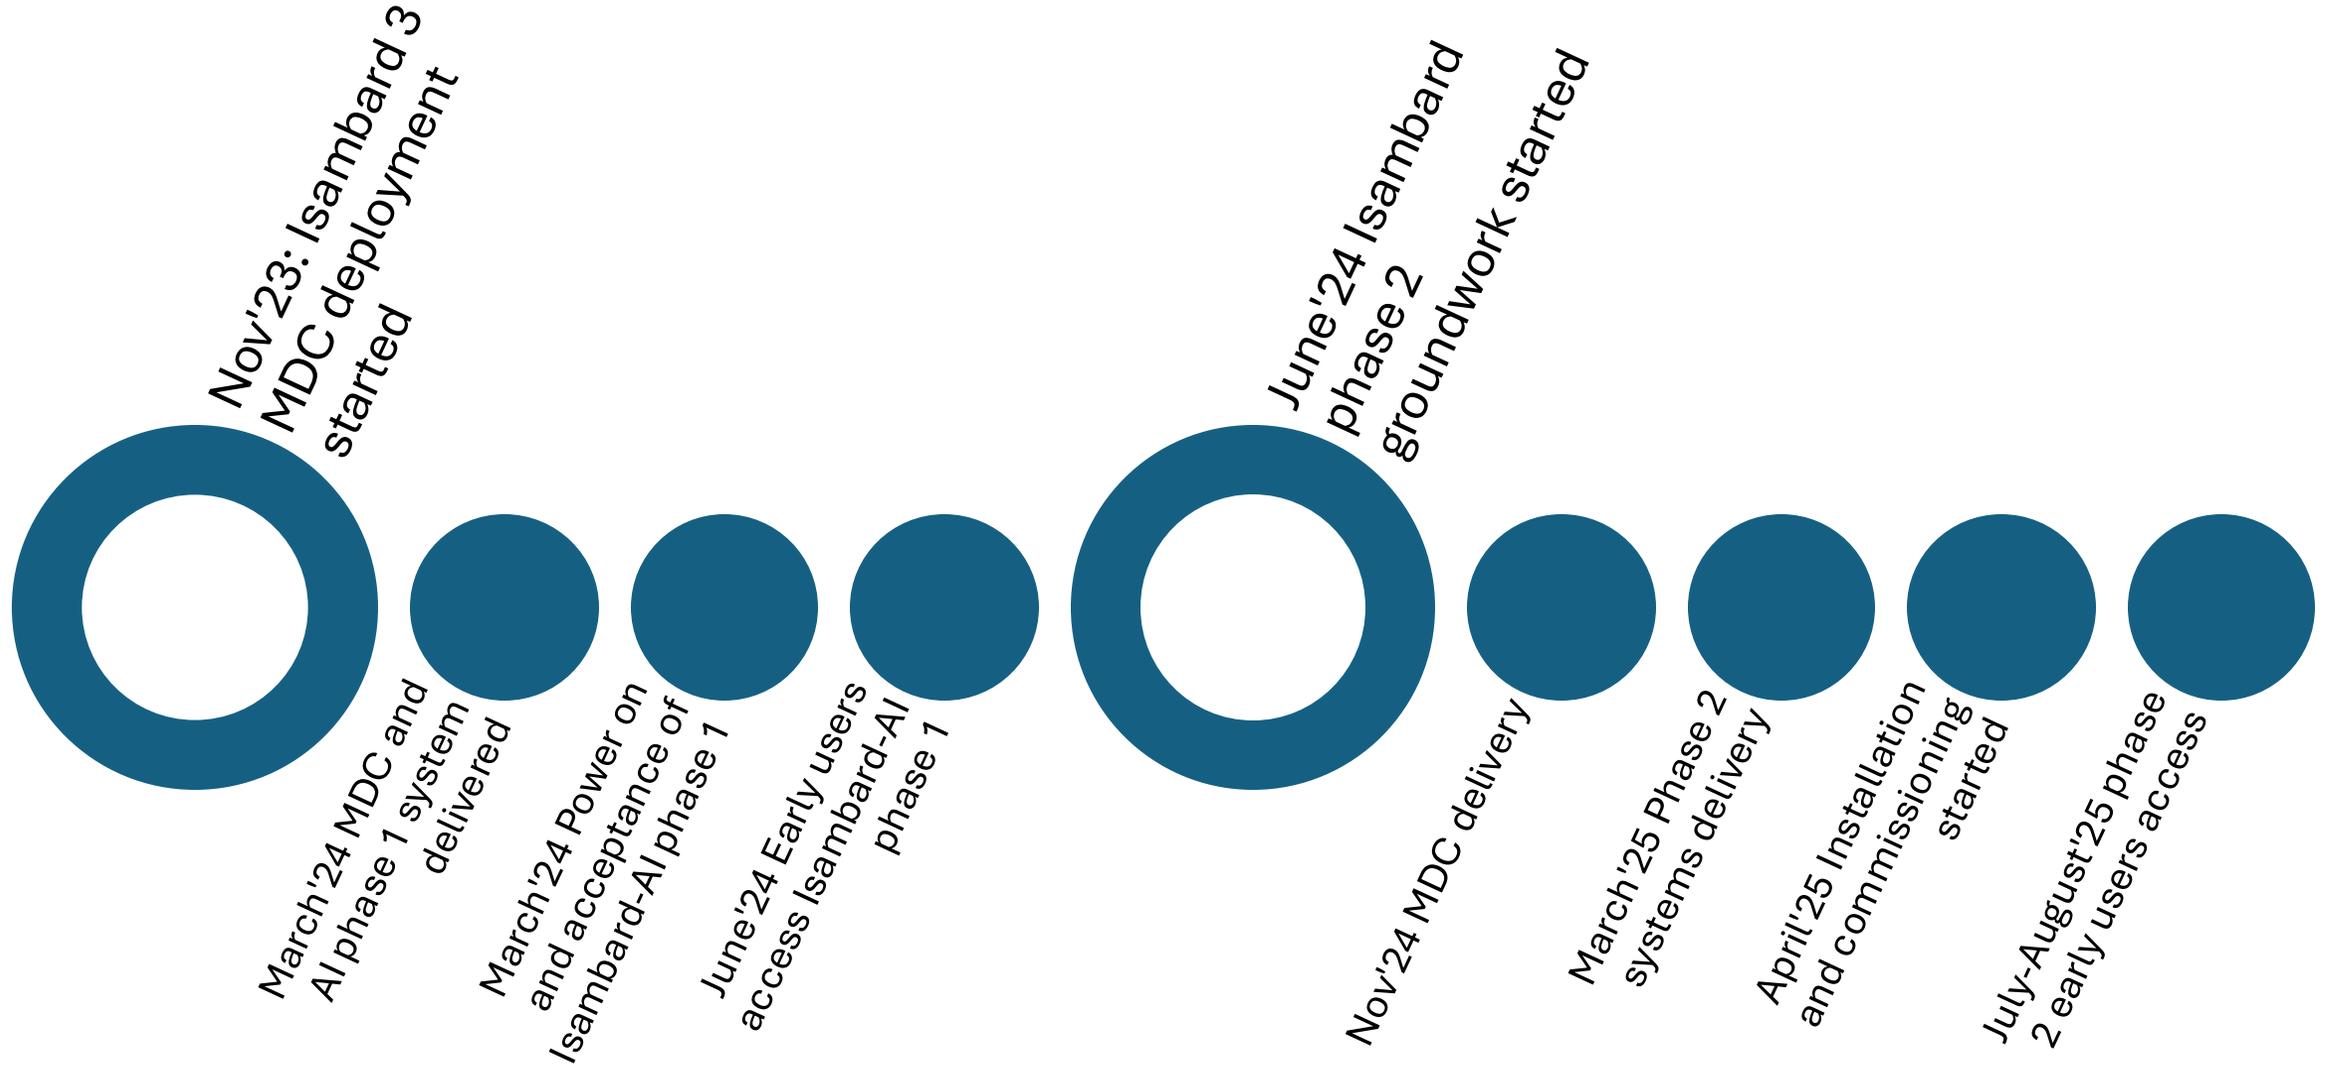
Isambard-AI phase 2 = Utility power available via MV
RMU, multiple MV/LV transformers and LV power
distribution panels in a dedicated switch house

Water treatment infrastructure

Heat recovery infrastructure enabling up to 4000 kW of
heat reuse.

	HPE Options	HPE provided	Site provided
<i>POD</i>	<ul style="list-style-type: none"> • For the POD IT module or All in on POD MDC • Outdoor or indoor • Standard products or customised • Scale-out options through spines and corridors 	<ul style="list-style-type: none"> • Outdoor, self-contained PODs for Isambard 3 & AI Phase 1 and AI phase 2 • Standard products or customised • Isambard-AI phase 2 is scaled out to fit 12 EX4000 cabinets plus connectivity with the air-cooled cabinets 	<ul style="list-style-type: none"> • Supporting pre-installation, deployment (in factory integration) and post-deployment support • Liaising with site contractors
<i>Power</i>	<ul style="list-style-type: none"> • MV and LV Switchgear boards • UPS options • Diesel generators • Power stabilisers • Transformers 	<ul style="list-style-type: none"> • LV power modules with UPS and Batteries to enable graceful shutdown of HPC kit • Diesel generators • Dedicated free cooling units • Fire detection • Downstream LV power distribution 	<ul style="list-style-type: none"> • Infrastructure for utility power connectivity including MV distribution transformers and upstream LV bus bars, etc. • Statutory provider upgrades
<i>Cooling</i>	<ul style="list-style-type: none"> • Chillers with free air cooling • Dry/hybrid coolers for DLC • Pump skids • Heat recovery infrastructure • On-site pipe works 	<ul style="list-style-type: none"> • Chillers with free air cooling • Dry/hybrid coolers for liquid cooling • Pump skids • On-site pipe works • Heat recovery interlinks 	<ul style="list-style-type: none"> • Supporting pre-installation, deployment and post-deployment support
<i>Site</i>	<ul style="list-style-type: none"> • Concrete slab • Network and cabling • Secure Cage • Fencing • Access controls • Security Systems 	<ul style="list-style-type: none"> • Networking in between PODs • Storage facility <div style="text-align: center; background-color: #f4a460; border-radius: 15px; padding: 10px; margin: 10px auto; width: fit-content;"> Flexibility and versatility </div>	<ul style="list-style-type: none"> • Concrete slab • Network and cabling • Secure Cage • Fencing • Access controls to PODs • Water treatment

Deployment Efficiency



Water Usage Efficiency

- Our projections are based on 100% IT load.
- Negligible when the ambient temperature is no higher than 19°C or 66.2°F.
- Wet mode activated for ambient temperature over 19°C
 - 3.67M litres of water across the ~11% of the hours in a year that are expected to be above 19°C in Bristol.
 - Flow rates 2.2m³ at 19-20°C to 7.7m³ at 35°C (2,200 litres up to 7,700 litres per hour once in wet mode).
 - 20,400  per year (30-50 gallons per bath per Gemini AI)
 - No water consumption ~90% of the time per year

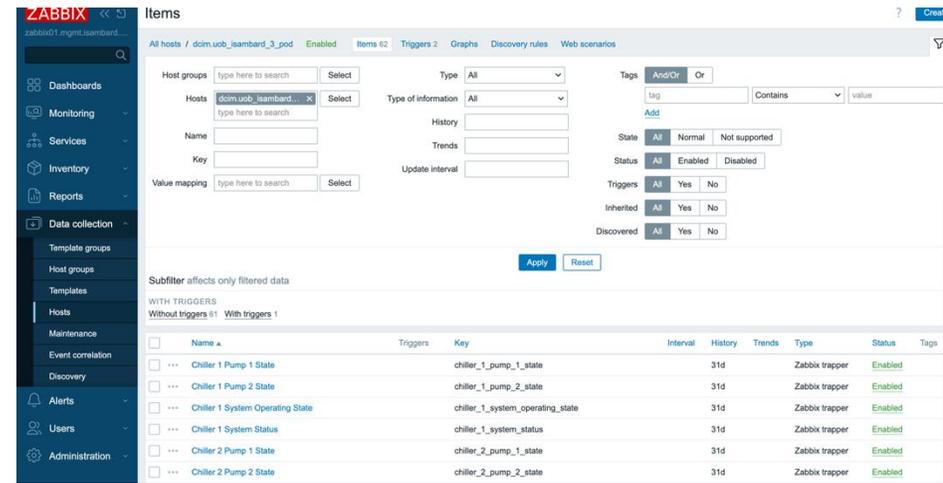
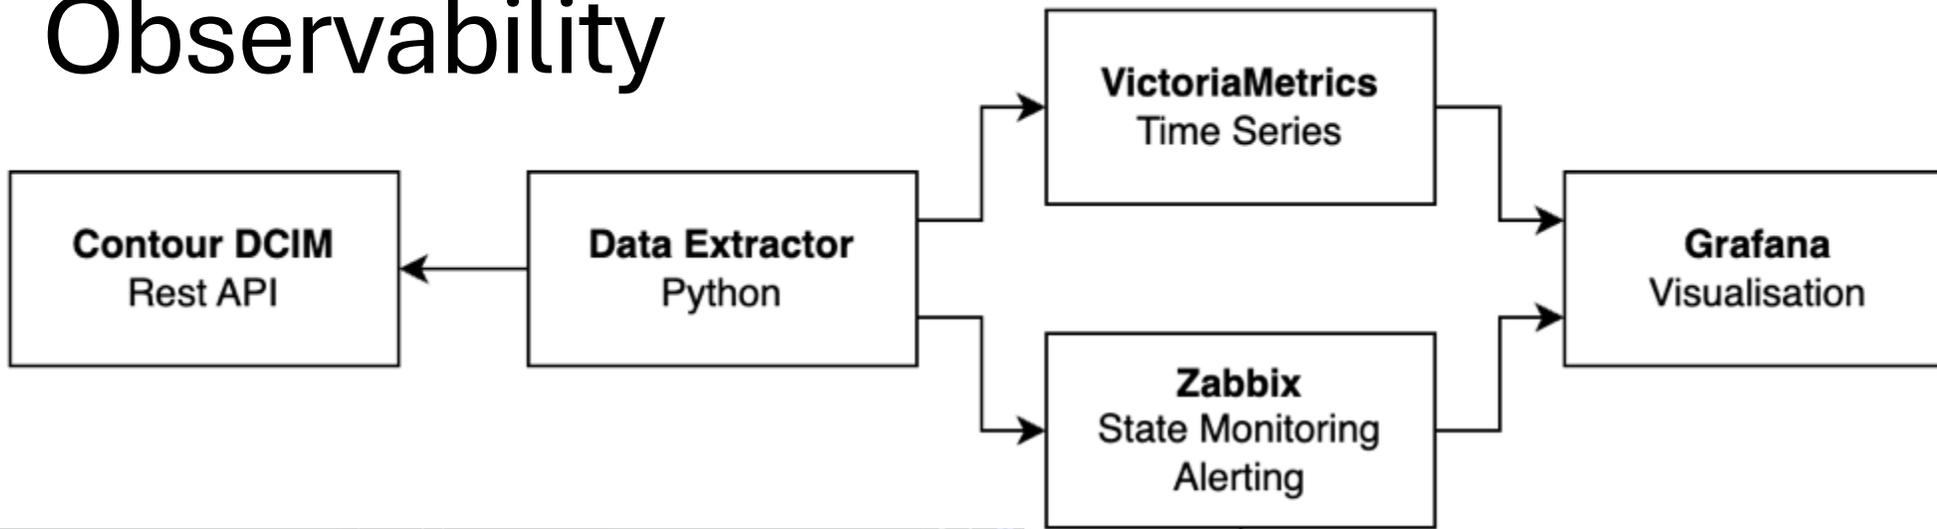
AI Overview

Hyperscale data centers, which can be larger than 10,000 square feet and house over 5,000 servers, can use a substantial amount of water, particularly for cooling their equipment. Water consumption varies depending on the cooling technology used, but evaporative cooling, a common method, can lead to daily water usage of between 1 and 5 million gallons (MGD) for a hyperscale data center. This equates to a yearly consumption of 365 million to 1.825 billion gallons, which is roughly equivalent to the water usage of 12,000 to 60,000 Americans.

Sustainability

- Energy Efficient High Performance Computing Working Group (EE HPC WG) <https://sites.google.com/lbl.gov/procurement/home>
 - Energy Efficiency Considerations for HPC Procurement Document: 2021 (~2022-2027)
 - Sustainability Chapter added to the Energy Efficiency Considerations for HPC Procurement Document: 2021 (version 2.0) https://drive.google.com/file/d/1nJYzxVmlivvG6R1fThophbA_N_fQM97a/view
- **Isambard-AI procurement fulfils**
 - **Advanced requirement** embodies water usage associated with production and water usage report.
 - **Innovative requirement** for providing an estimated measurement of embodied greenhouse gas emissions of the processes associated with the production according to some standard such as ISO 14067 as well as usage report
 - **Innovative requirement** for providing a product Life-Cycle Assessment (LCA) disclosure, including water consumption, pollution, and eco-toxicity. The disclosure should include a statement on the methodology used for the product LCA. There should be assurance that the product LCA is verifiable.
 - **Innovative requirement** for providing a supply chain responsibility report specific to the procured product, based on a documented disclosure framework

Observability



Lessons Learned (ambitious timelines)

- Site works uncertainties
 - Dependencies on several local entities and stakeholders
 - Never underestimate lawyers
- Just-in-Time findings
 - Introduction to detailed physical security plans (in a car park, not a secure facility)
 - Never underestimate insurance
- Co-design agility
 - Pros—optimising design for changing priorities
 - Cons—managing cost and program
 - Never underestimate highly motivated teams and partnerships



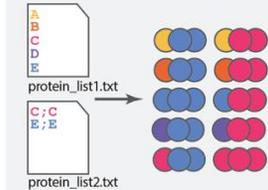
More Users, Projects & Success Stories

- 100+ projects, ~500 users
- UCL **BritLLM** (UK models specific for law, health, and finance applications, & local languages)
<https://llm.org.uk>, <https://arxiv.org/abs/2410.23956>
- AlphaFold and OpenFold for **Cardiac disease** (Understanding how gene mutations change protein complexes and cause inflammation)
- Sensitivity of models to backdoor attacks after **data poisoning** (How easy is it for malicious actors to insert poisoned training data for a universal jailbreak?)
 - BIAS-AI (**Bias** of diagnosis systems to skin colour)
 - Turing Institute fully **reproducible model training** including optimiser checkpointing (GPTNeoX)

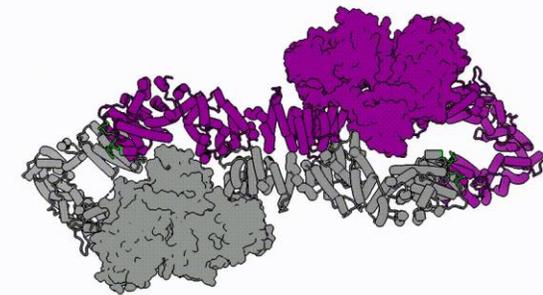
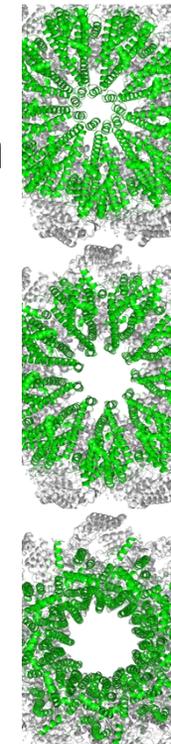
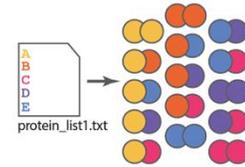
single input



multiple inputs



all versus all



NF1: morphing of inpainted and experimental models

AIRR Expression of Interest to use UK large scale compute

Overview

UK Research and Innovation (UKRI), on behalf of the Department for Science, Innovation and Technology (DSIT), invites researchers and innovators from across the UK to express their interest in accessing large-scale AI compute, in particular, the new AI Research Resource, the [Isambard-AI](#) and [Dawn compute services](#).

We are seeking expressions of interest from researchers and innovators who can demonstrate a clear need for AI compute and may be suitable for early access to the Isambard AI and Dawn services as part of their testing phase.

Closes 19 Dec 2025

Opened 14 Jan 2025

Contact

dri@ukri.org

Thank you

- Contact: brics-enquiries@bristol.ac.uk
- <https://docs.isambard.ac.uk/>



Acknowledgements

Isambard-AI is funded by the UK Government's Department of Science, Innovation and Technology (DSIT) via UKRI / STFC. The project has only been possible thanks to a very talented and extremely hardworking team of people across the University of Bristol, HPE, Contour, NVIDIA, the National Composites Centre, and our building contractors, Oakland.