How to rapidly deploy and deliver a highly sustainable national Exascale AI research resource?

Sadaf R Alam
University of Bristol
NHR PerfLab Seminar talk
July 30, 2024

Outline

- Introduction
- Story of Isambard-AI
  - A fast, sustainable, performance and accessible AI supercomputing platform
- Capabilities and potential for diverse research domains
- Early users and next steps

2	128	**Isambard-AI phase 1** - HPE Cray EX254n, NVIDIA Grace 72C 3.1GHz, NVIDIA GH200 Superchip, Slingshot-11, HPE University of Bristol United Kingdom	34,272	7.42	117	68.835

ISC, May 2024

University of BRISTOL

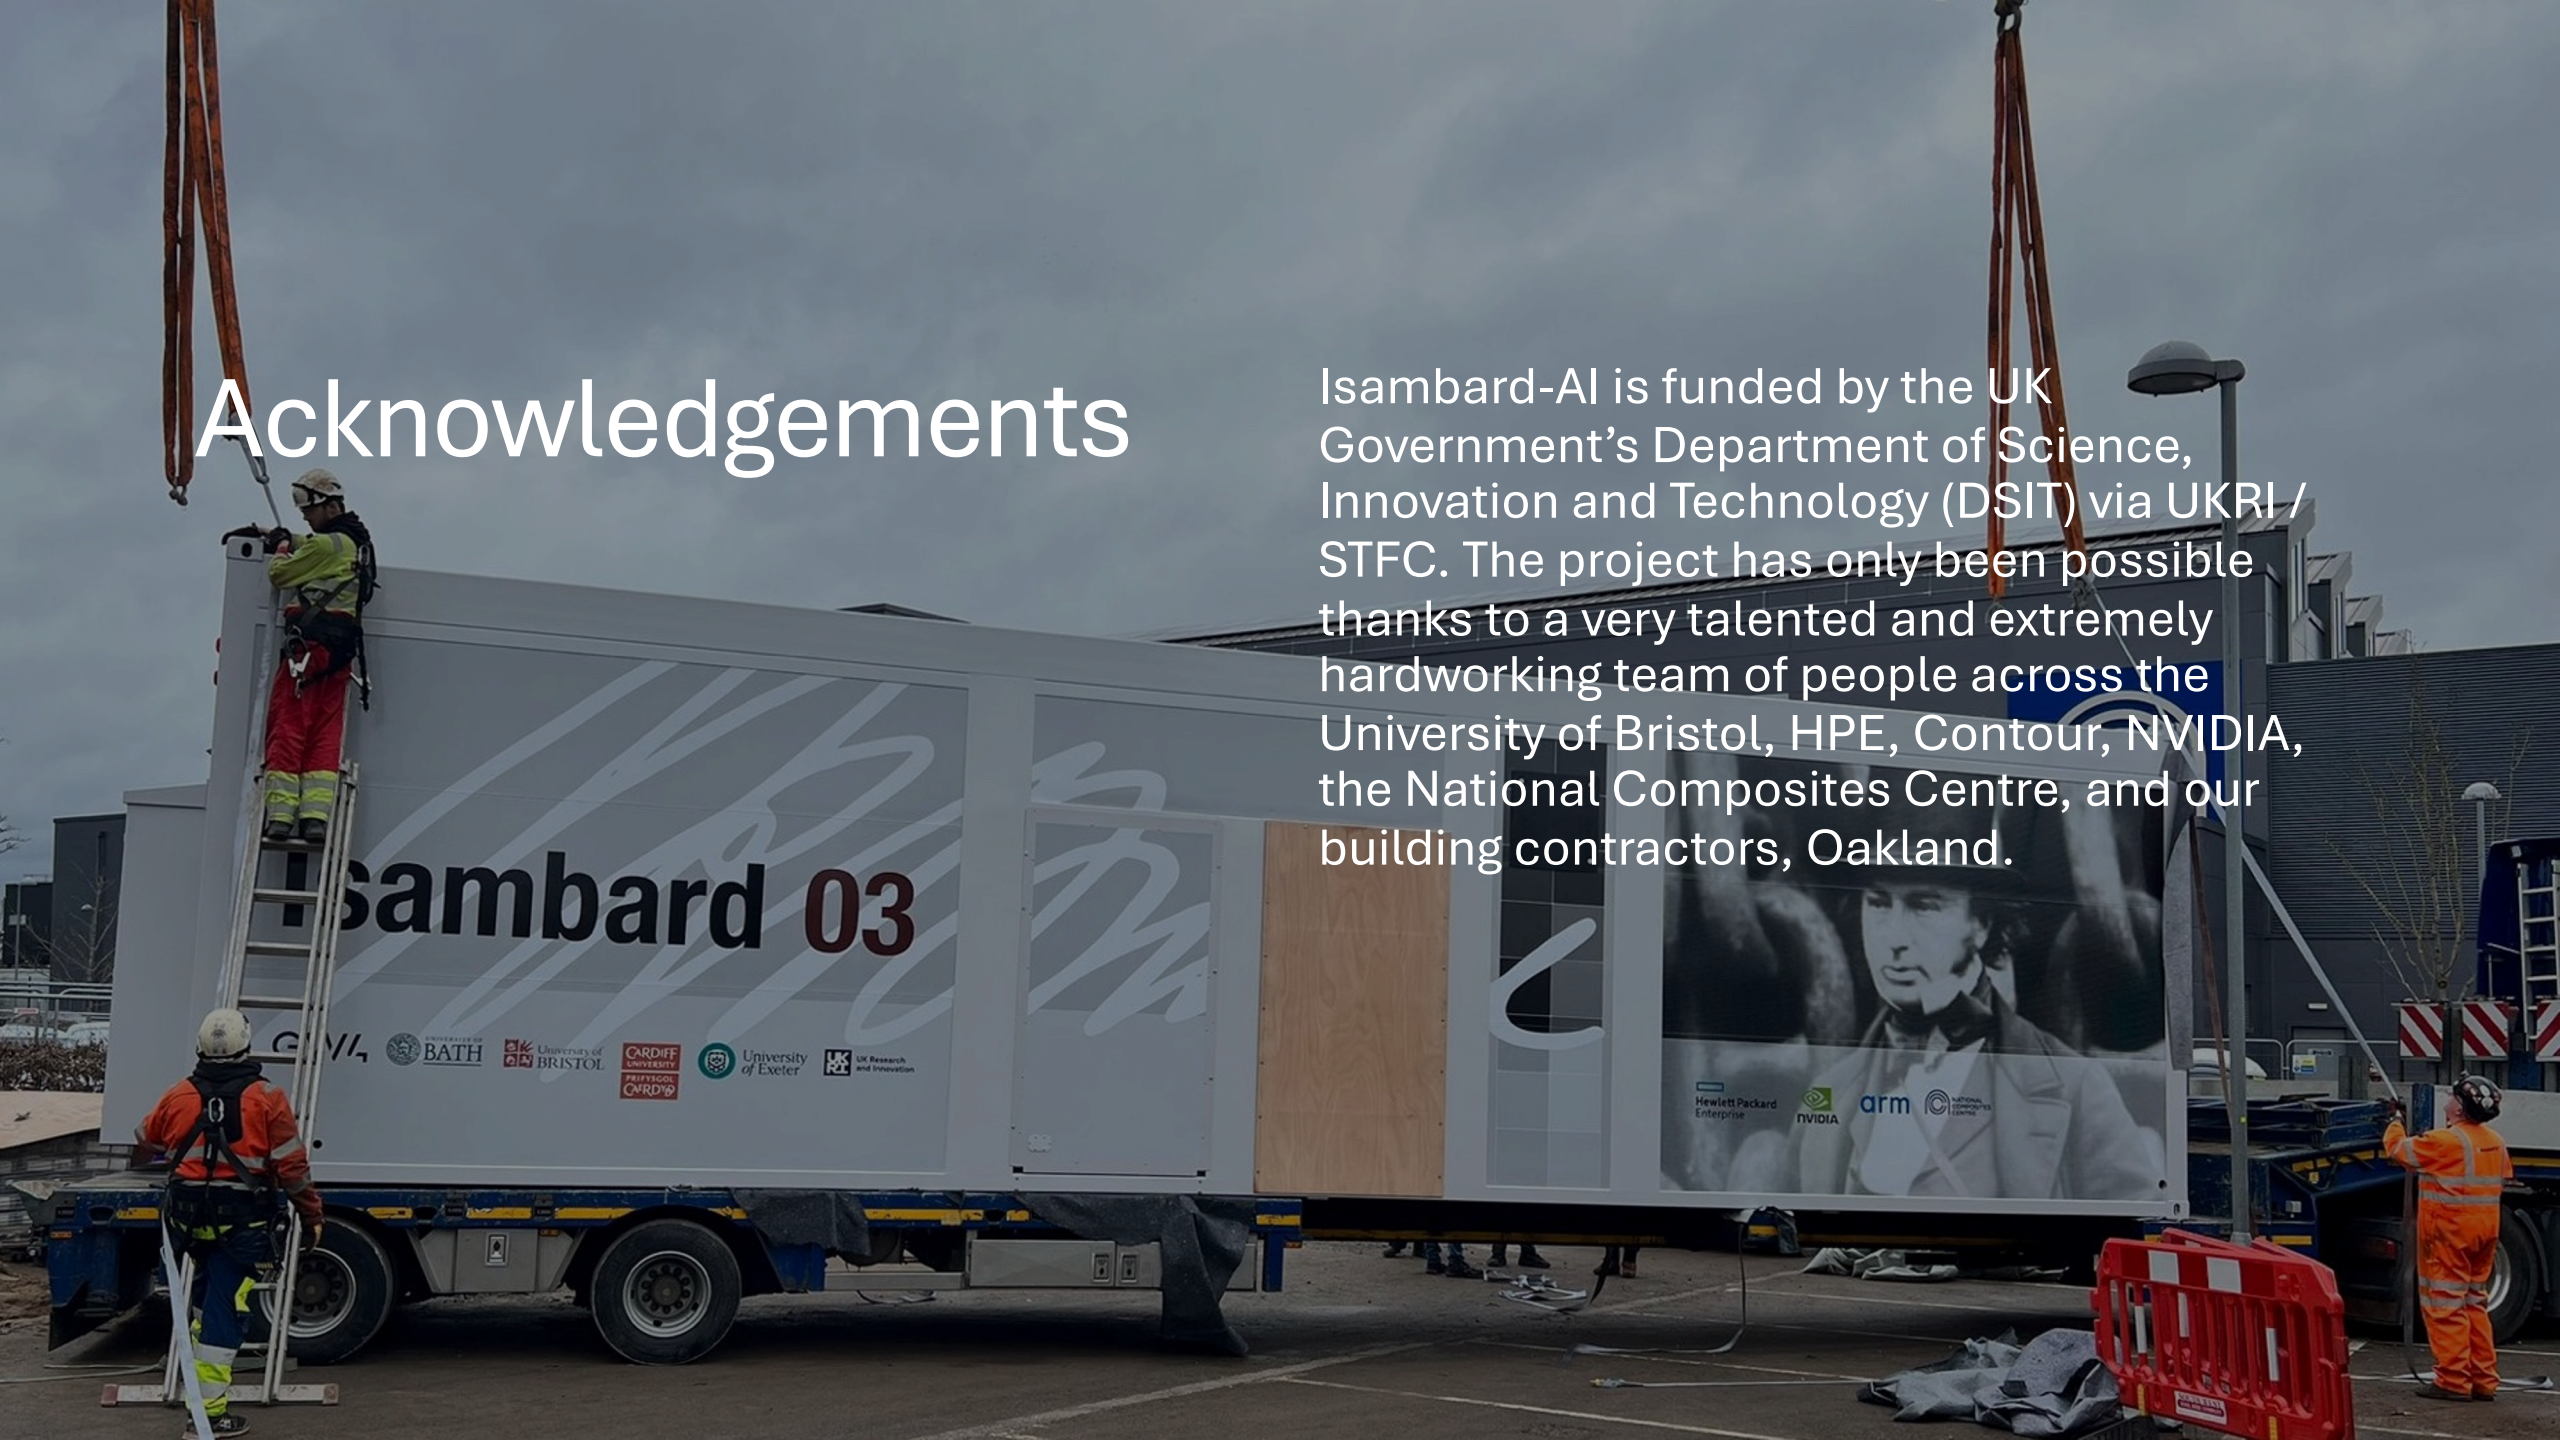

Acknowledgements

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.

Introduction

Who am I

- 2022-today 🇬🇧
  - University of Bristol-–Chief Technology Officer (CTO) for Isambard DRIs and Director of Advanced Computing Strategy
- 2009-2022 🇨🇭
  - Swiss National Supercomputing Centre—various roles including CTO, head of operations & computer scientist
- 2004-2009 🇺🇸
  - Oak Ridge National Laboratory (ORNL) Leadership Computing Facility liaison, computer scientist and postdoc
- 2000-2004 🇬🇧
  - University of Edinburgh PhD Researcher
- 1999-2000 🇬🇧
  - Analog Devices Software Engineer

https://insight.ieeeusa.org/articles/the-art-of-mentoring/

University of Bristol and Isambard Projects

- The University of Bristol is a red brick Russell Group research university in Bristol, England. It received its royal charter in 1909, although it can trace its roots to a Merchant Venturers' school founded in 1595 and University College, Bristol, which had been in existence since 1876 [Wikipedia]
- The GW4 (Bath, Bristol, Cardiff and Exeter) Isambard project initially set out to prove that a new Arm-based technology was relevant to supercomputing.
- Bristol School for Supercomputing (BriCS) has been recently formed for managing Isambard DRI projects within the Faculty of Engineering

University of BRISTOL

Story of Isambard-AI

Press release

Bristol set to host UK's most powerful supercomputer to turbocharge AI innovation

A new supercomputer set is to be built in Bristol, in a move to drive pioneering AI research and innovation in the UK.

- UK AI Research Resource dubbed Isambard-AI will be one of Europe's most powerful supercomputers
- new facility will serve as national resource for researchers and industry experts spearheading AI innovation and scientific discovery
- plans for the supercomputer backed by £900 million investment announced in March to transform UK's computing capacity

From: Department for Science, Innovation and Technology and The Rt Hon Michelle Donelan MP

Published 13 September 2023

A new supercomputer set to be one the most powerful in Europe is to be built in Bristol, in a move to drive pioneering AI research and innovation in the UK.

The UK government has confirmed the University of Bristol will host the new AI Research Resource (AIRR), which will serve as a national facility to help researchers maximise the potential of AI and support critical work into the potential and safe use of the technology.

The world-class AIRR cluster will vastly increase the UK's compute capacity – essential to achieving the UK's AI ambitions and securing its place as a world-leader in harnessing the rapidly developing technology. The cluster, which will be made up of thousands of state-of-the-art graphics processing units, or GPUs, will be able to train the large language models that are at the forefront of AI research and development today.

University of BRISTOL

https://www.gov.uk/government/news/bristol-set-to-host-uks-most-powerful-supercomputer-to-turbocharge-ai-innovation

Isambard-AI and UK's AI Safety Institute

UK Research and Innovation

Department for Science, Innovation, & Technology

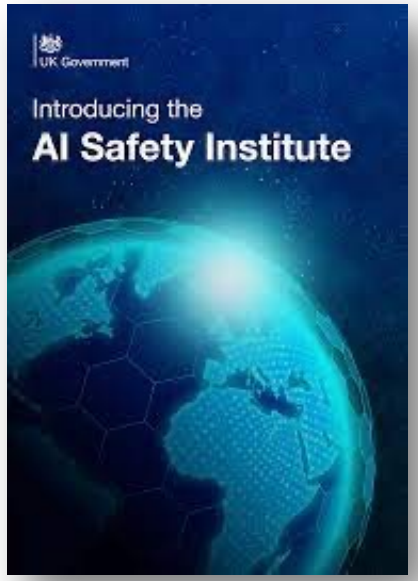

We must always remember that AI is not a natural phenomenon that is happening to us, but a product of human creation that we have the power to shape and direct. Accordingly, we are not waiting to react to its impacts but are choosing to be proactive in defining the trajectory of its development, to ensure public safety and human flourishing for years to come. This is why the UK is building the AI Safety Institute.

The Institute is the first state-backed organisation focused on advanced AI safety for the public interest. This is our contribution to addressing a shared challenge posed to all of humanity.

Running evaluations and advancing safety research will also depend on access to compute. The AIRR will integrate the recently announced Isambard-AI compute cluster at Bristol University, which will be one of the most powerful AI supercomputers in Europe.

Isambard Digital Research Infrastructures (DRIs)
Isambard-AI phase 2
5,280xNVIDIA GH200
12 HPE Cray EX Cabinets
Direct Liquid Cooled (DLC) computing & HPC networking
10.4 M AI TOPS (Trillion Operations per Second)
21 8-bit AI Exaflops

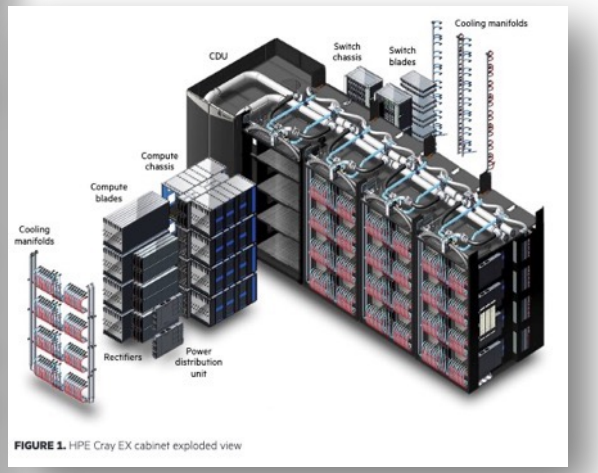

FIGURE 1. HPE Cray EX cabinet exploded view

https://www.gov.uk/government/publications/ai-safety-institute-overview/introducing-the-ai-safety-institute

University of BRISTOL

(ARM based Isambard 1, 2) Isambard 3 in 2023/4 (no data centre and Isambard PI Simon MS with a non-dedicated GW4 team)

Isambard-AI procurement

Isambard hiring started

"Who is Isambard?" running on Isambard-AI

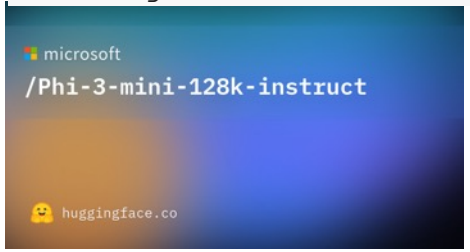

Service configuration and hardening for AI users

2016–July 2023 | Aug. 2023 | Sep.–Oct. 2023 | Nov. 2023 | Dec. 2023 | Mar. 2024 | May 2024 | July 2024

UK govt feedback on Bristol's Isambard-AI proposal

AI Safety Summit

Modular data centre (POD) and Isambard-AI phase 1 installed + a new team in ~3 months

Onboarding Early Users via co-design and technical readiness projects

University of BRISTOL

Isambard 1, 2 & 3 – Leading ARM for HPC since 2016 as a UK national tier-2 resource

- Isambard 1 and 2 hosted at the UK Met Office data centre
- Options considered:
  - Renting space in a DC—£££ plus not available for hundreds of KW DLC cabinets like Cray HPE XE
  - Building new-–time and ££££
- Solution—containersied data centre or MDC

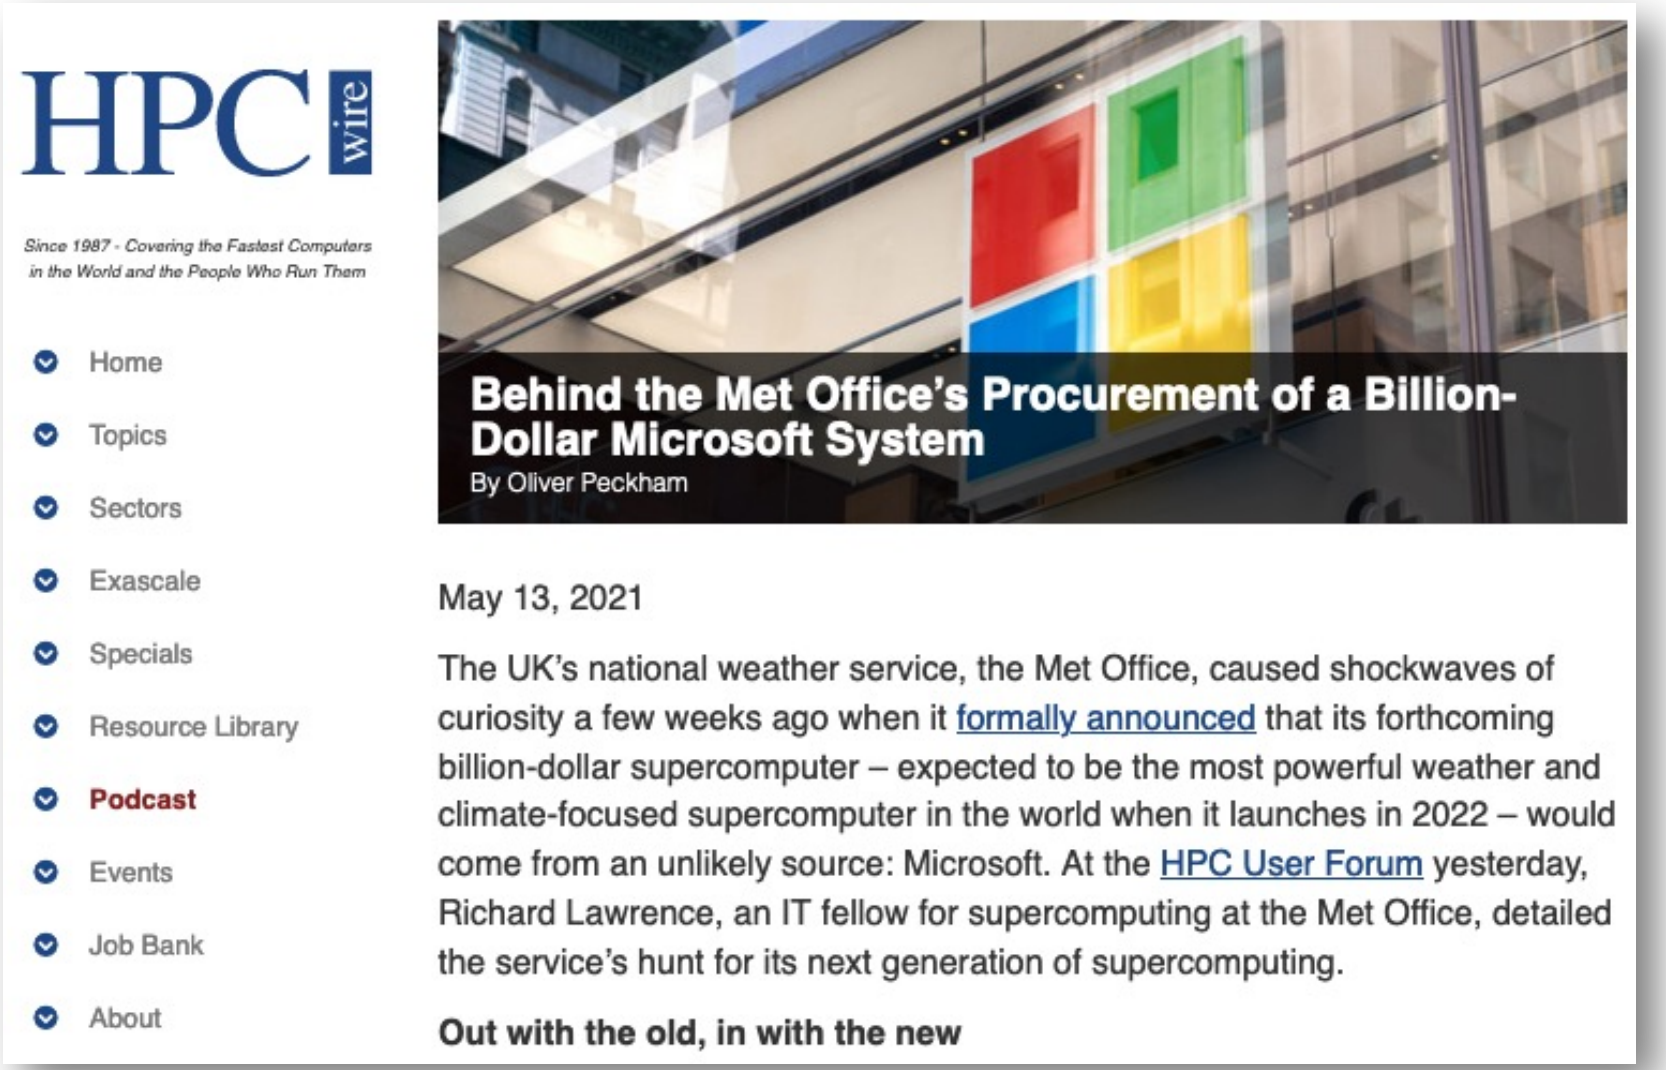

HPC wire

Since 1987 - Covering the Fastest Computers in the World and the People Who Run Them

- Home
- Topics
- Sectors
- Exascale
- Specials
- Resource Library
- Podcast
- Events
- Job Bank
- About

Behind the Met Office's Procurement of a Billion-Dollar Microsoft System

By Oliver Peckham

May 13, 2021

The UK's national weather service, the Met Office, caused shockwaves of curiosity a few weeks ago when it formally announced that its forthcoming billion-dollar supercomputer – expected to be the most powerful weather and climate-focused supercomputer in the world when it launches in 2022 – would come from an unlikely source: Microsoft. At the HPC User Forum yesterday, Richard Lawrence, an IT fellow for supercomputing at the Met Office, detailed the service's hunt for its next generation of supercomputing.

Out with the old, in with the new

University of BRISTOL

What is an MDC

- Modular, self-contained & agile
  - Described as Lego blocks—designed and tuned for functional and performance specifications e.g. high availability, security, etc.
  - Everything from all IT (compute, network & storage), power (UPS, batteries) and cooling can be included as self-contained units
- Efficient and flexible deployments
  - Typically built and commissioned in months
  - Accommodates different environmental conditions
    - On-site integration options
    - Offsite integration options
- Sustainable
  - Holistic, fine-grain telemetry via DCIM
  - Upgradable, refreshable, recyclable with a lifespan on 10-15 years

The Isambard Site (Isambard-AI phase 1 & 2, and Isambard 3)

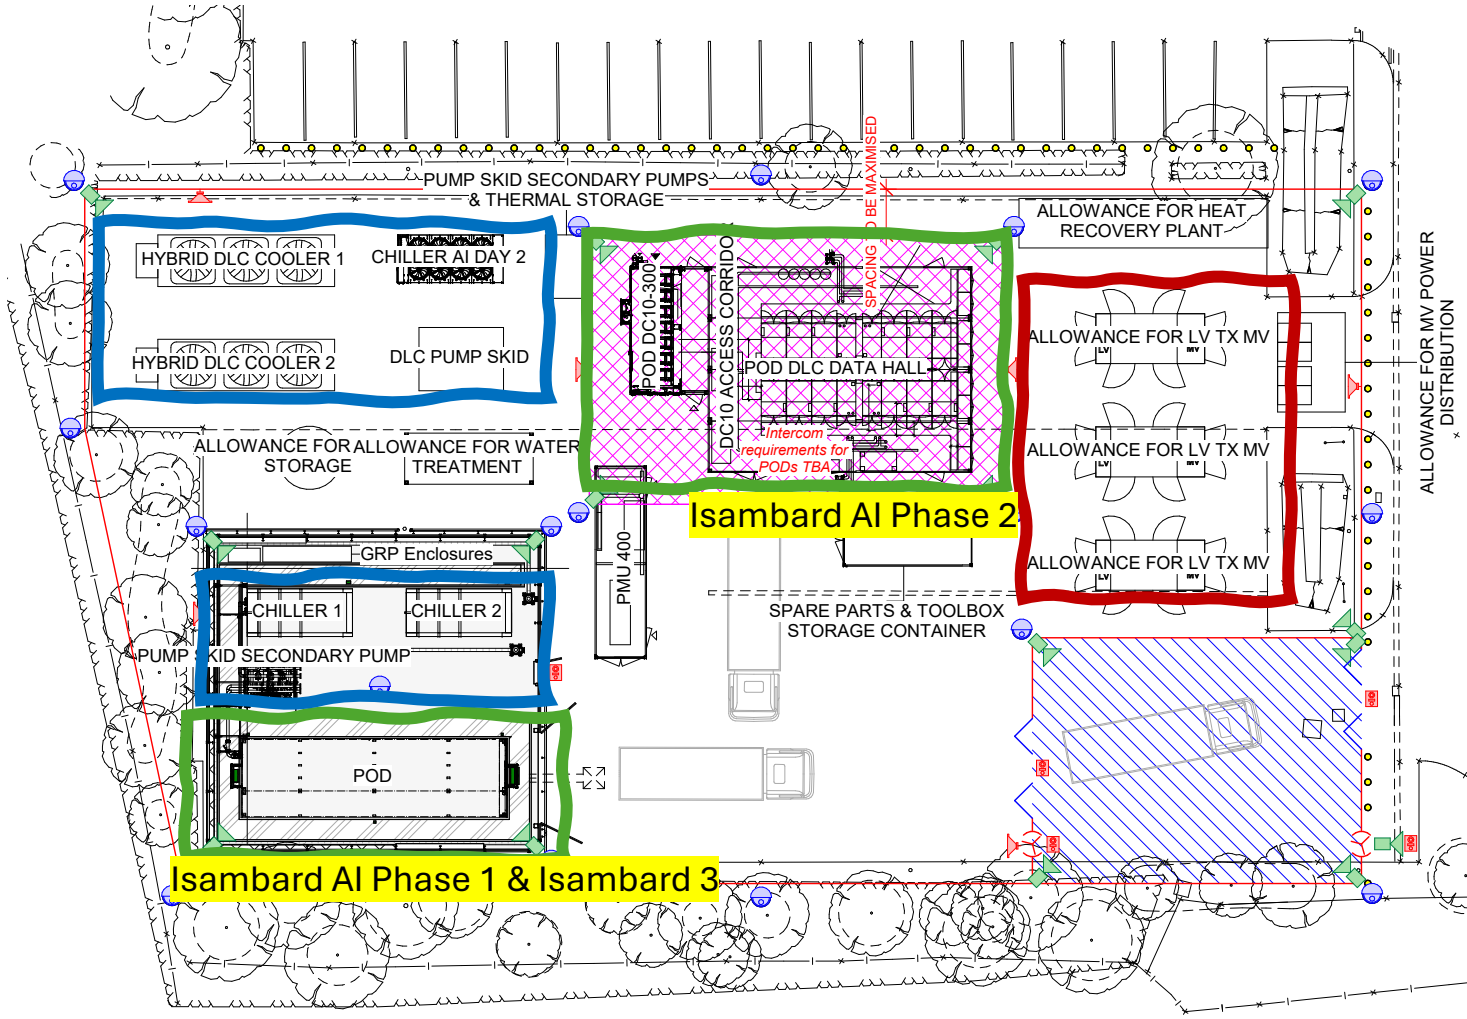

University of BRISTOL

Sustainability as a Key Design Principle

- Optimisation targets
  - PUE = Power Usage Effectiveness
    - Target <1.1
  - CUE = Carbon Usage Effectiveness
    - Non-fossil fuel sources
  - Plan for heat reuse for nearby buildings and local district heat circuit in future
- Aligning with university of Bristol Net Zero and sustainability targets for 2030
  - Categorising emissions
    - Scope 1 (~0%), 2 (90%) and 3 (10%)—based on an average UK data of 0.2123 kg CO2/kWh (IEA 2022 data)
  - Recycling 90% of components at the end of life in the UK

University of BRISTOL

Isambard Site – National Composite Centre (NCC) Facility in Bristol

- NCC—UK's Centre of Excellence for Composites Research and Development
  - Availability of power (~10 MW), networking and cooling
  - Heat reuse options
  - Co-location with industrial user community that has a digitalisation first approach

Physical Space Constraints at NCC

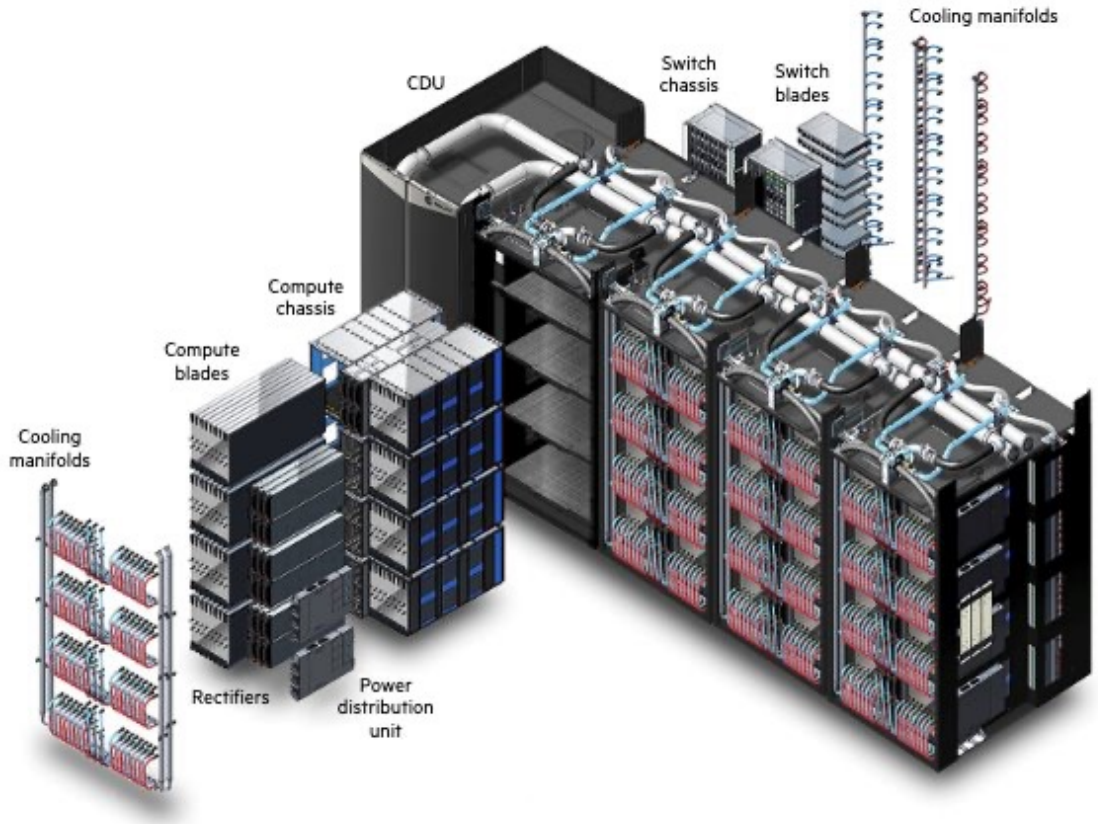

FIGURE 1. HPE Cray EX cabinet exploded view

5,280 GPUs in **12 EX4000 cabinets**

New Class of AI Supercomputer Connects 256 Grace Hopper Superchips Into Massive, 1-Exaflop, 144TB GPU for Giant Models Powering Generative AI, Recommender Systems, Data Processing

May 28, 2023

DGX SuperPOD with **256 GPUs in 16 cabinets** → **300 cabinets for 5300 GPUs ~30x increase in floor space**

University of BRISTOL

Isambard-AI vs Nvidia DGX SuperPOD

Isambard AI

- Space
  - 5280 GH200 in 12 HPE Cray EX compute cabinets
- Power
  - ~5MW inclusive of ecosystem
- TCO
  - CapEx ~£200M, TCO ~£300M over 5 years
- Software stack
  - Work in progress (see next slide)

DGX SuperPOD

- Space
  - 5,280 GH200 in ~330 compute cabinets
- Power
  - ~13 MW (40kW per H100 rack)
- TCO
  - CapEx >£400M, TCO >£600M over 5 years
- Software stack
  - Nvidia AI and ML optimised stack

Disclaimer: information is based on publicly available data and ChatGPT responses

University of BRISTOL

Convergence of Cloud and HPC on Isambard-AI

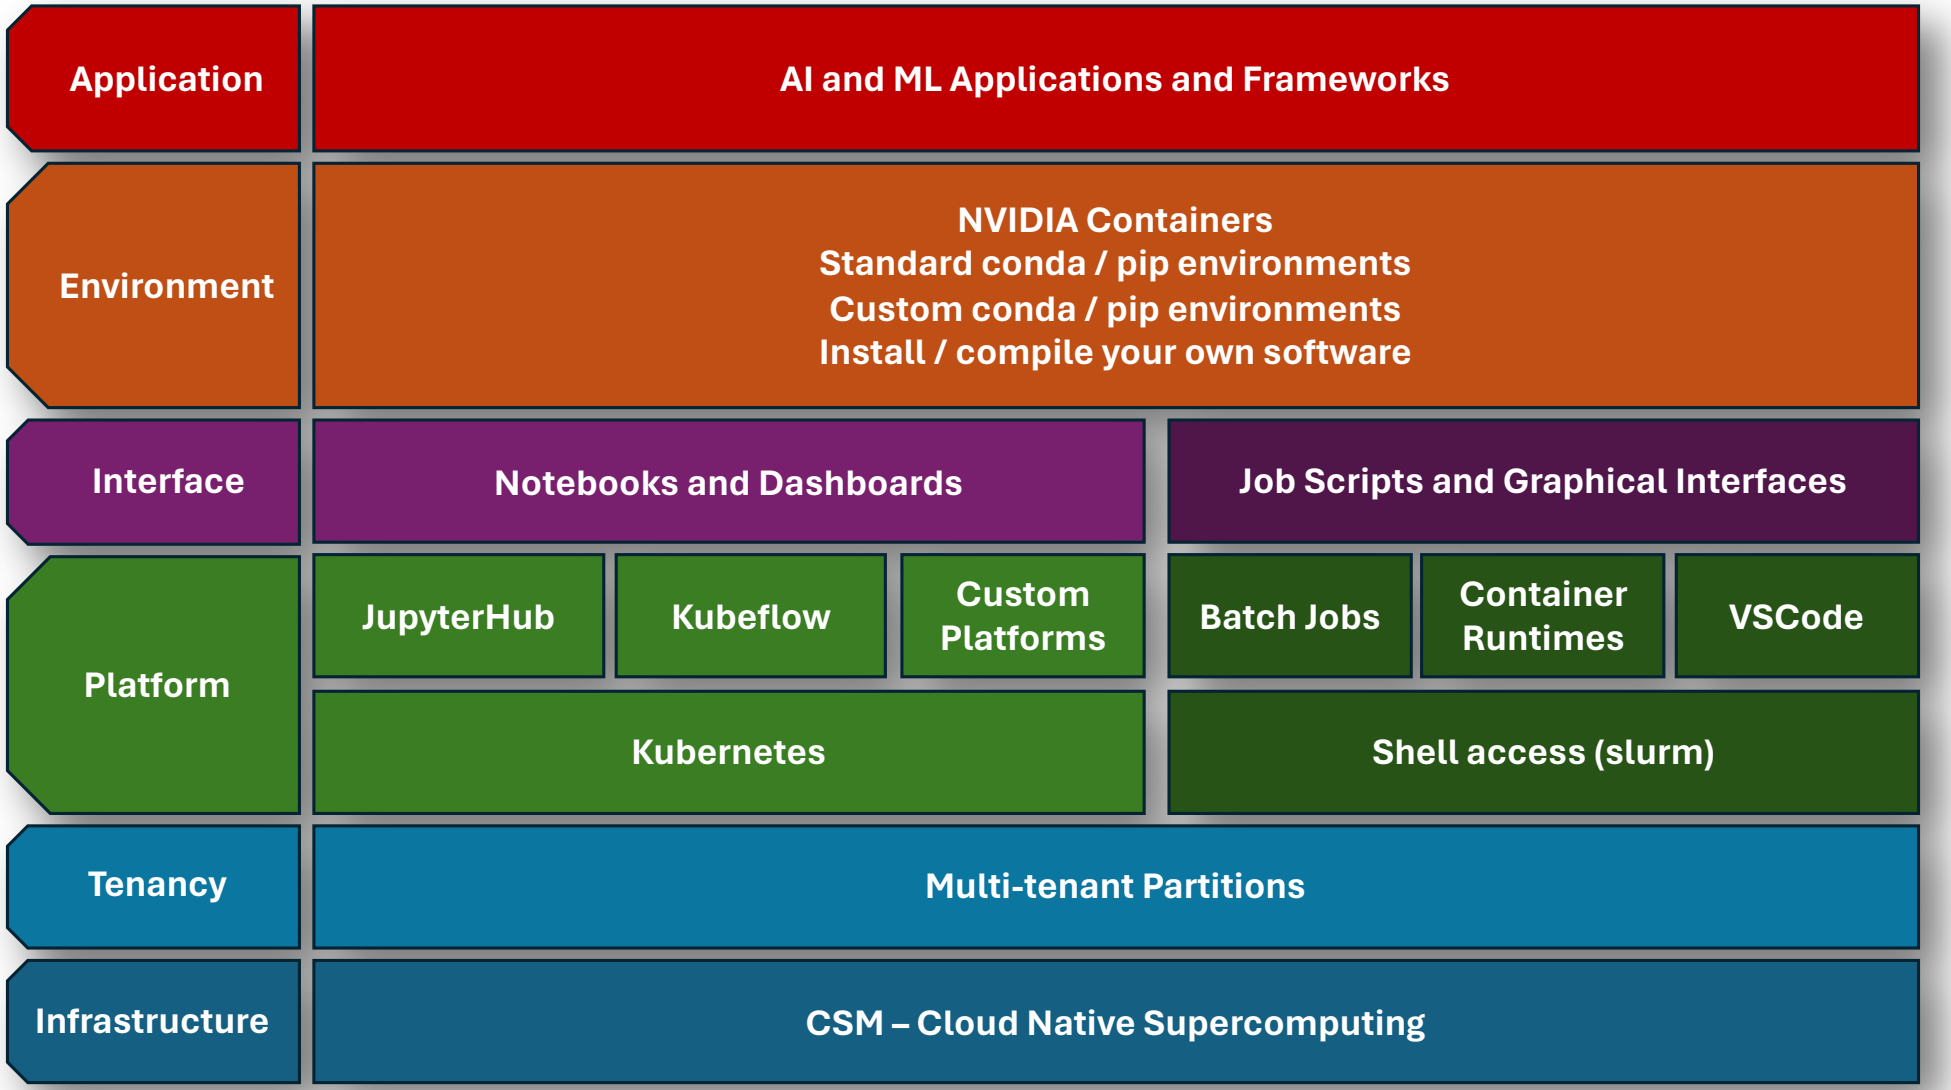

University of BRISTOL

HPE EX Series DLC and Nvidia GH 200

- HPE EX solution
  - Direct liquid cooling for high performance computing and networking
- 4-way Nvidia GH200 superchip
  - NVLink-C2C also only uses 1.3 picojoules/bit transferred—5x more energy efficient than PCIe Gen 5

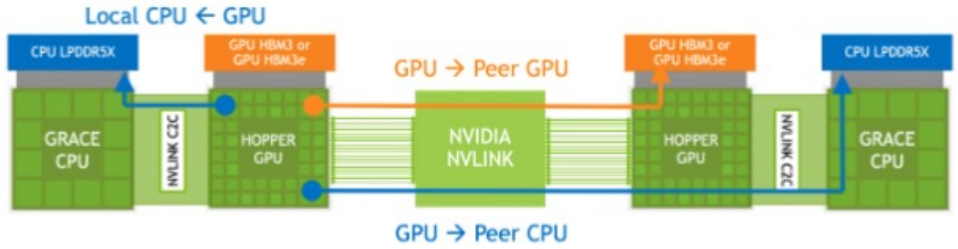

Figure 5. Memory Accesses across NVLink-connected Grace Hopper Superchips

Source: NVIDIA Grace Hopper Superchip Architecture Whitepaper

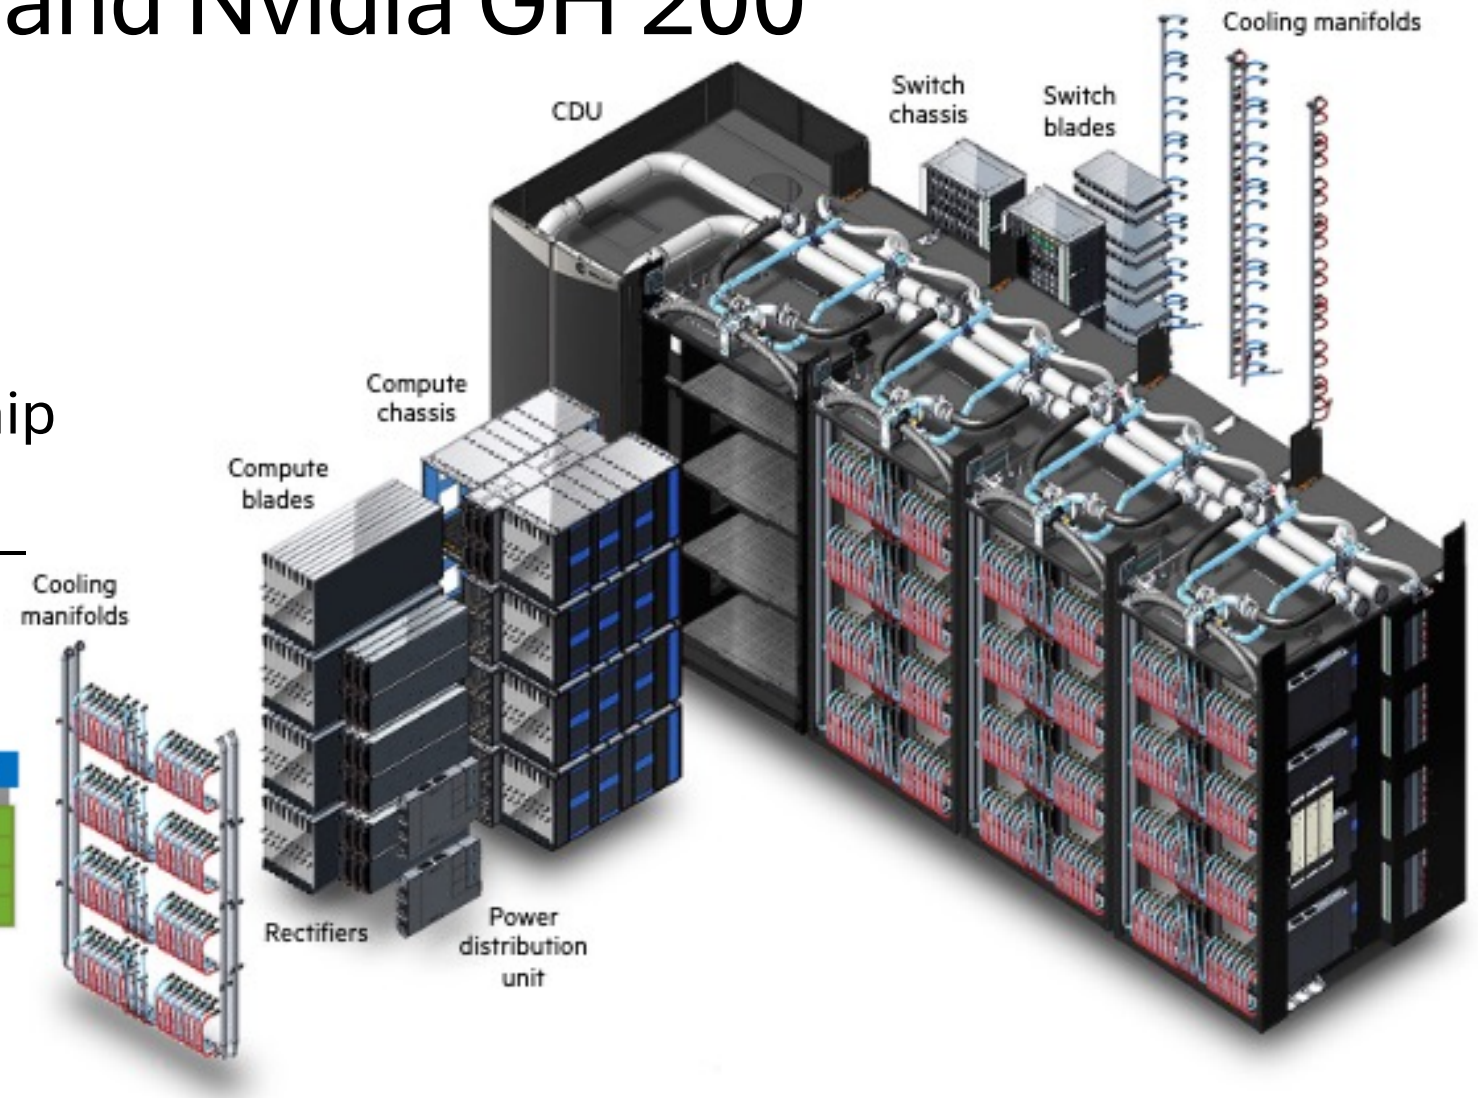

FIGURE 1. HPE Cray EX cabinet exploded view

Source: HPE CRAY EX Liquid-Cooled Cabinet for Large Scale Systems brochure

University of BRISTOL

Grace-Hopper Superchip & HPE EX Compute Blade

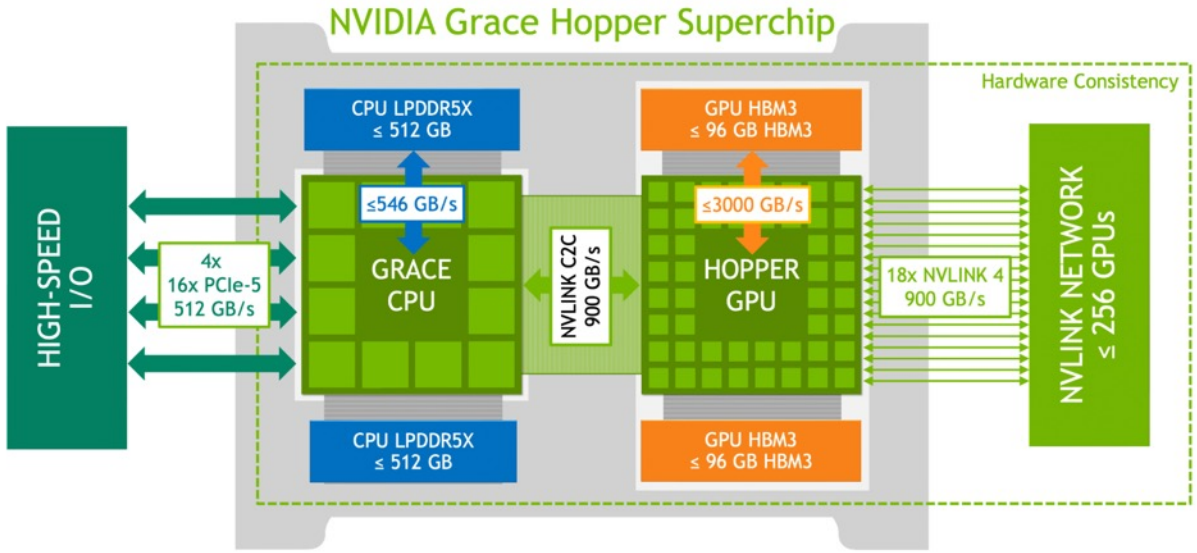

Source: NVIDIA Grace Hopper Superchip Architecture Whitepaper

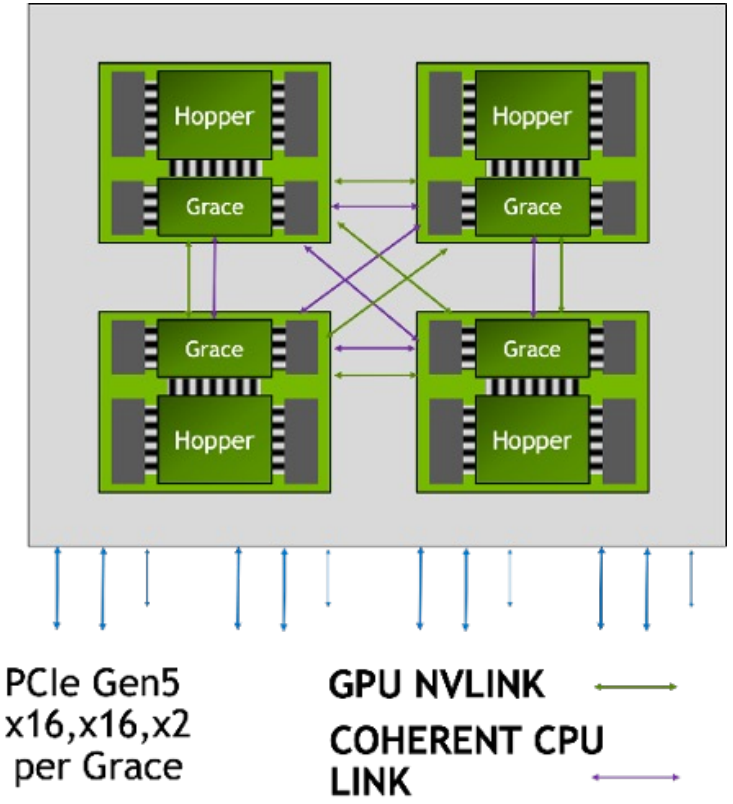

Source: HPE EX4000 Grace-Hopper blade

4 x Grace ARM CPUs
288 cores
512 GB Memory

4 x Hopper GPUs
~260 64-bit Tflops, ~16k 8-bit Tflops
384 GB High Bandwidth Memory

896 GB Memory Total
NVLink-C2C = 900 GB/s

Isambard AI node = 4 x GH200
Injection bandwidth = 4 x 200 Gbps

University of BRISTOL

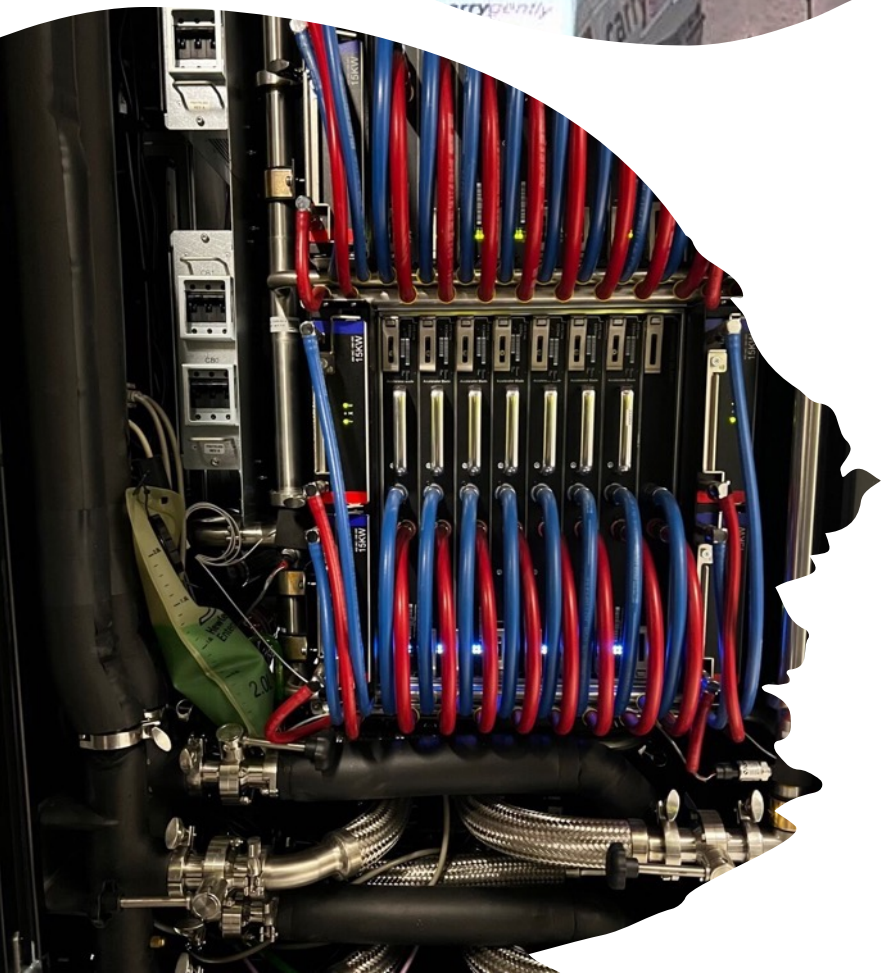

Isambard-AI Firsts

- First HPE Cray EX DLC inside an MDC/POD
- First deployment and power-up < 2 weeks
- First combined DLC (EX2500) for Grace-Hopper Superchip and air-cooled Proliant XD2000 Grace-Grace Superchip

Checklist

- Great team, collaborators and executive support from government, university, suppliers and contractors
- Clear vision, mission and goal orientation for delivering worldclass, federated digital research infrastructure
- Modular Data Centre (MDC) efficient, flexible and sustainable technology
- Proven technical solution (HPE Cray EX, SlingShot and Nvidia AI optimised GPUs)
- Mature AI and HPC software stacks from Nvidia and HPE

Capabilities and potential for diverse research domains

AI for Science

- AI is moving from computer science to computational science domains
  - AI-aided data collection and curation for scientific research
  - Learning meaningful representations of scientific data
  - AI-based generation of scientific hypotheses
  - AI-driven experimentation and simulation
- As size of data and size of models grows, capability of AI increases – leading to emergent capability
- Better understanding of how to "right-size" models to data / required capability

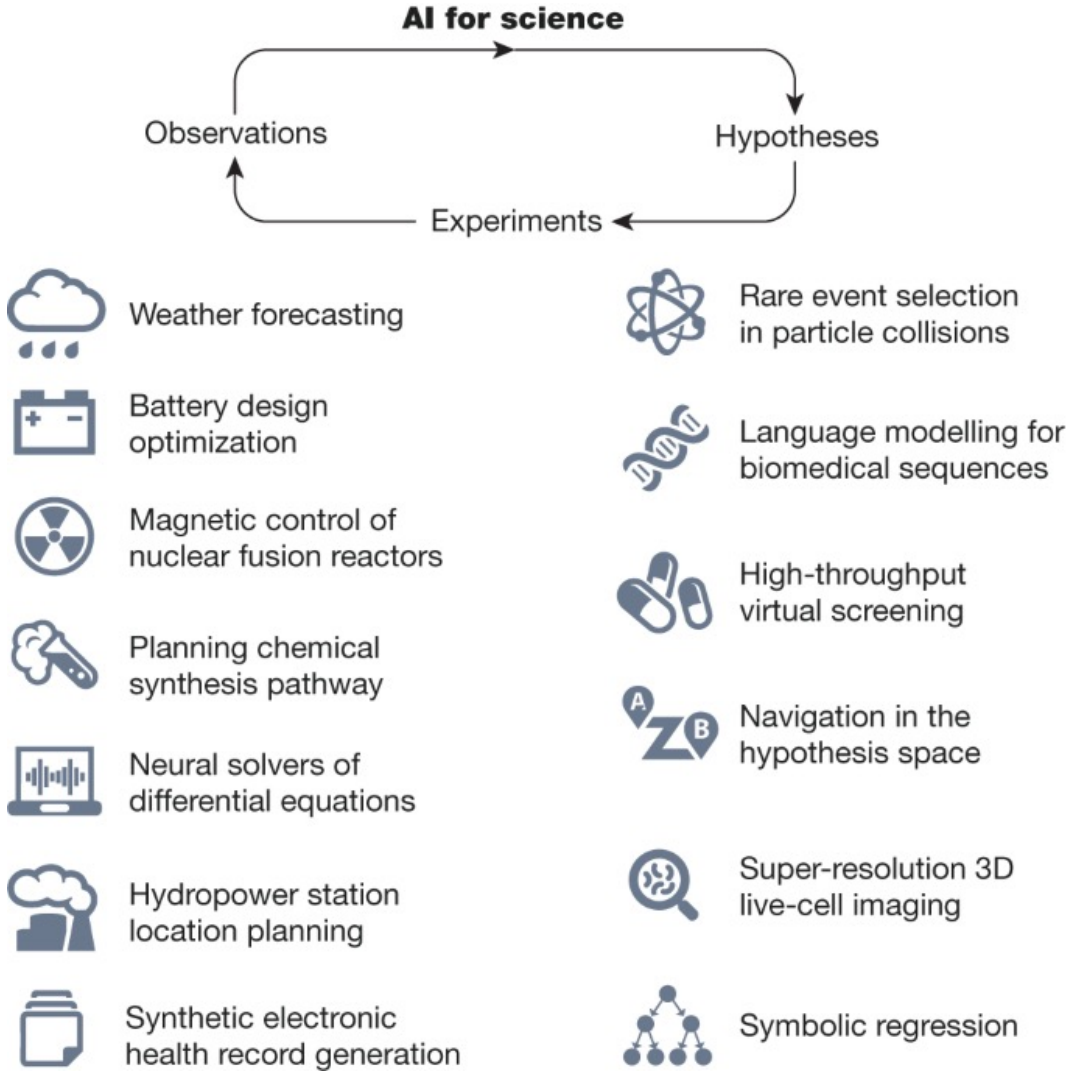

Wang, H., Fu, T., Du, Y. *et al.* Scientific discovery in the age of artificial intelligence. *Nature* **620**, 47–60 (2023).

University of BRISTOL

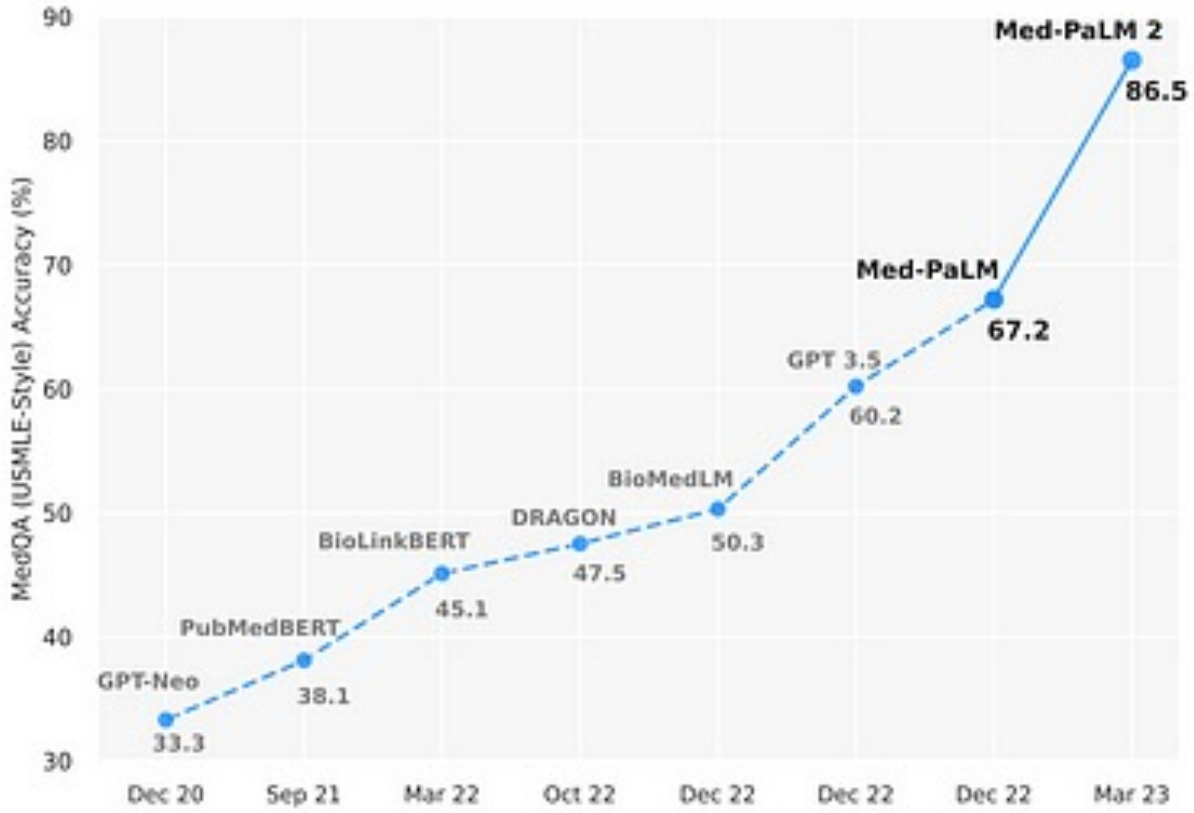

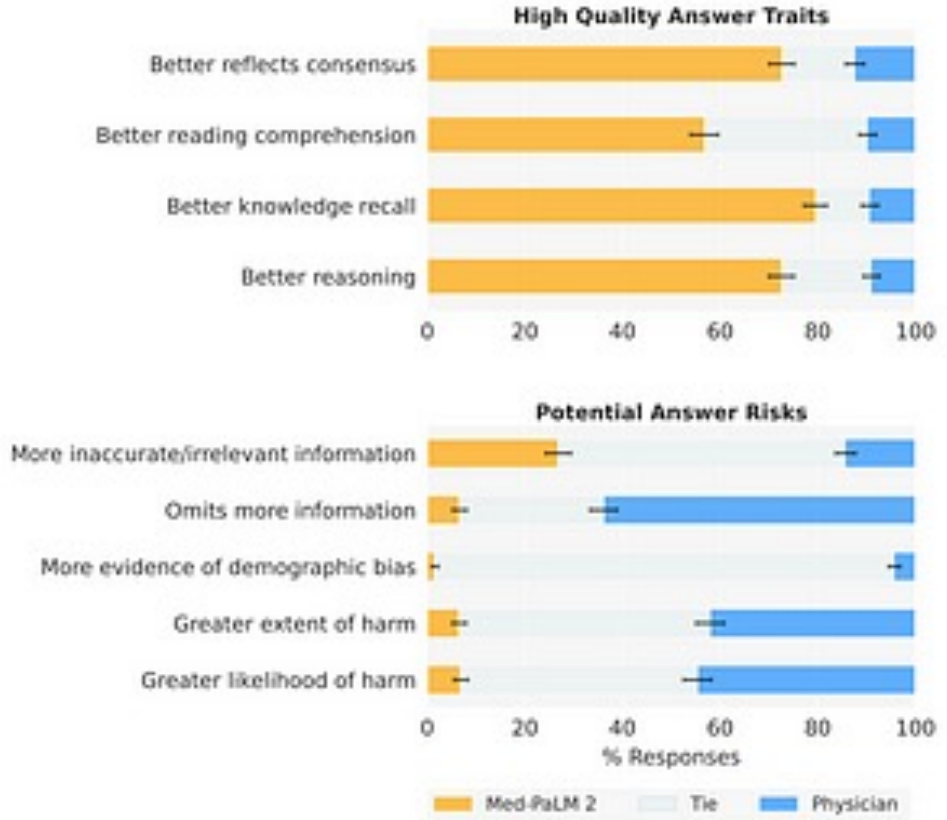

Left: Performance of various models on MedQA: Med-PaLM is the first AI model that passes the threshold; Med-PaLM 2 achieves significant leap; Right: Med-PaLM 2 outperforms human physicians in 8 of 9 aspects when answering 1066 consumer medical questions. The figure is copied from Singhal et al.

Source: https://medium.com/@AI_for_Science/ai-for-science-in-2023-a-community-primer-d2c2db37e9a7

Emergent Capabilities with increasing model size

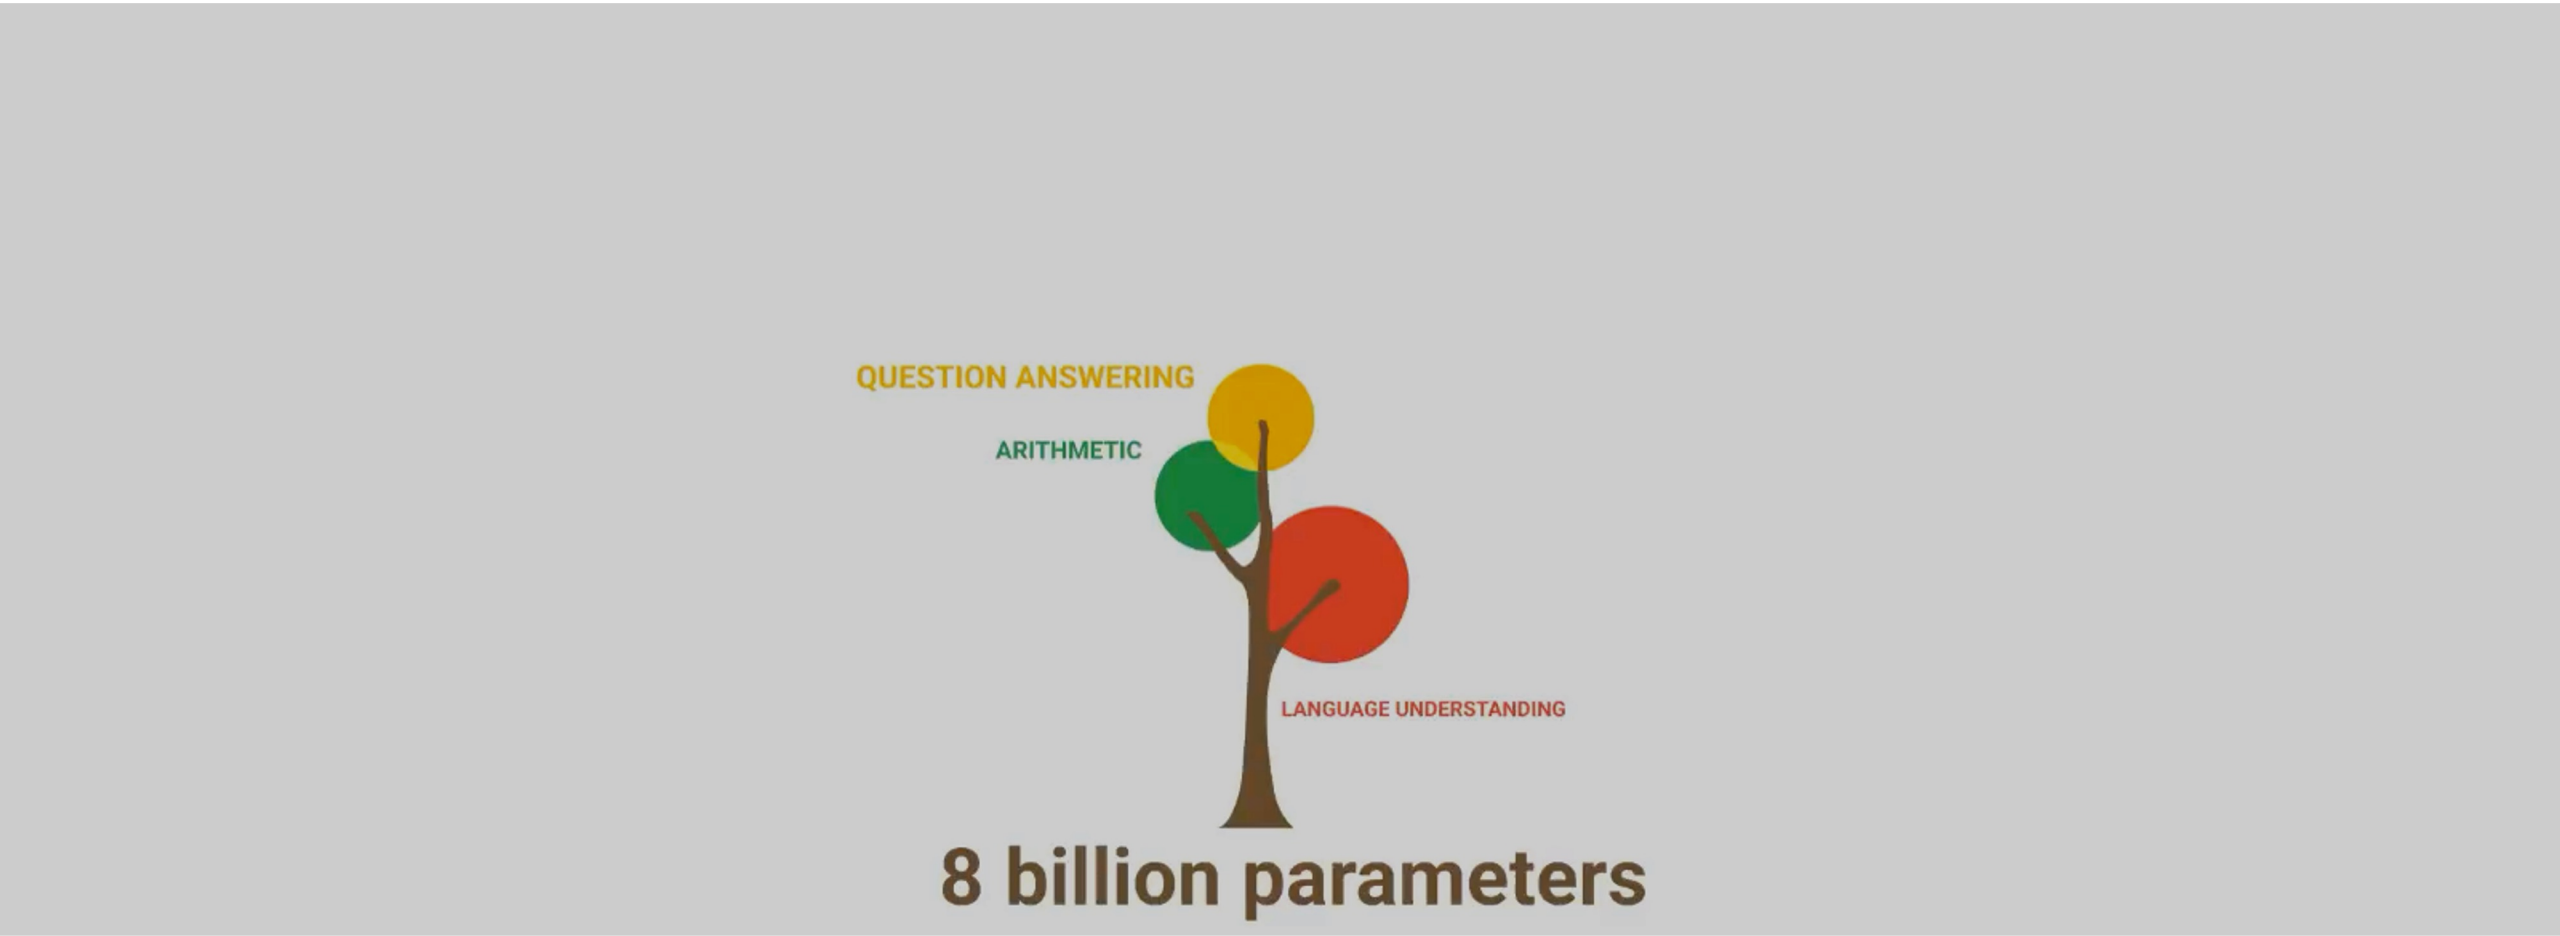

University of BRISTOL

https://ai.google/discover/palm2
https://the-decoder.com/google-palm-giant-language-ai-can-explain-jokes/

Model size increasing exponentially since 2018

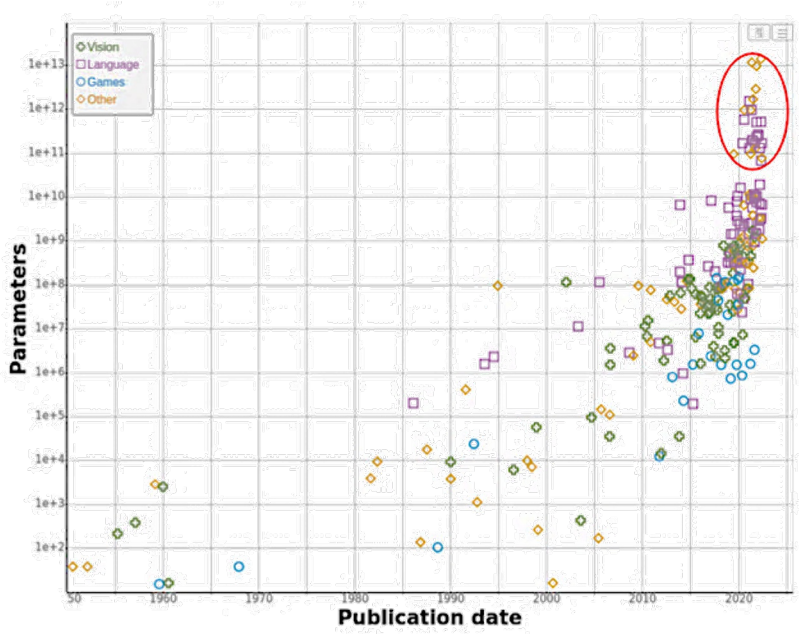

https://arxiv.org/pdf/2207.02852.pdf

”Open-source”

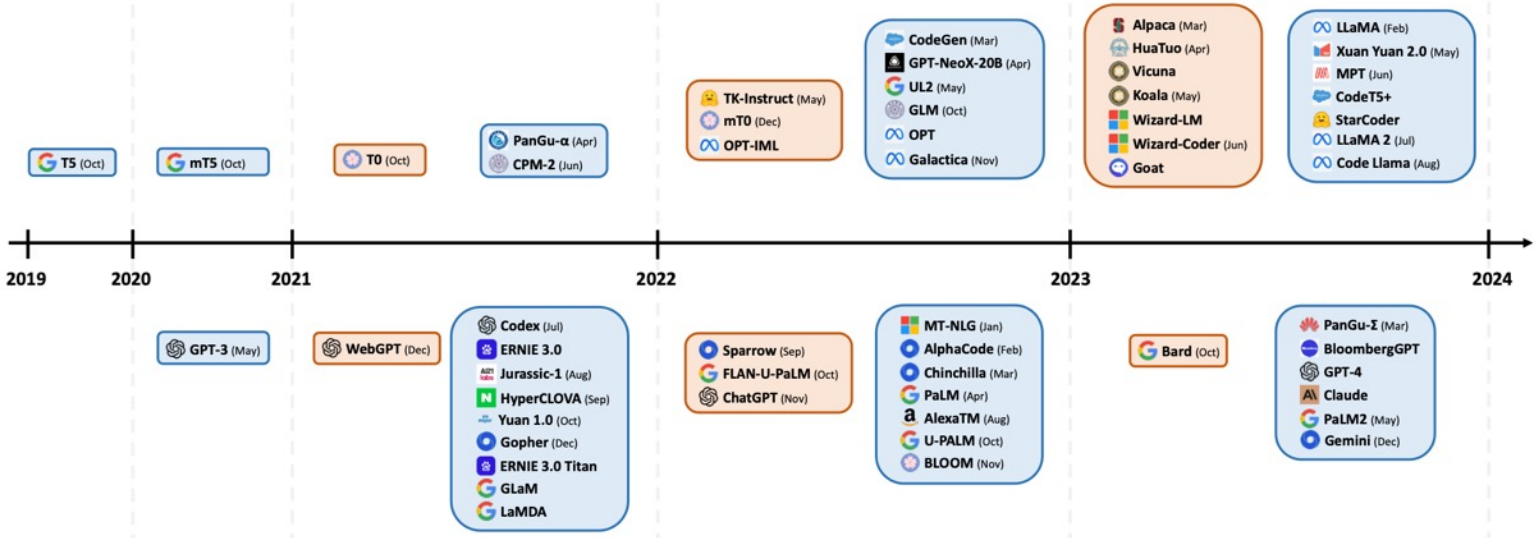

”Closed-source”

https://arxiv.org/pdf/2307.06435.pdf

University of BRISTOL

Scaling up – going beyond a laptop

NVIDIA GeForce RTX 4070
8 GB GPU memory
32 GB CPU memory
14 CPU cores
1 TB SSD disk
466 AI TOPS

More data!
Larger models!
More parameters!
Faster training!

Internationally competitive AI-powered research!

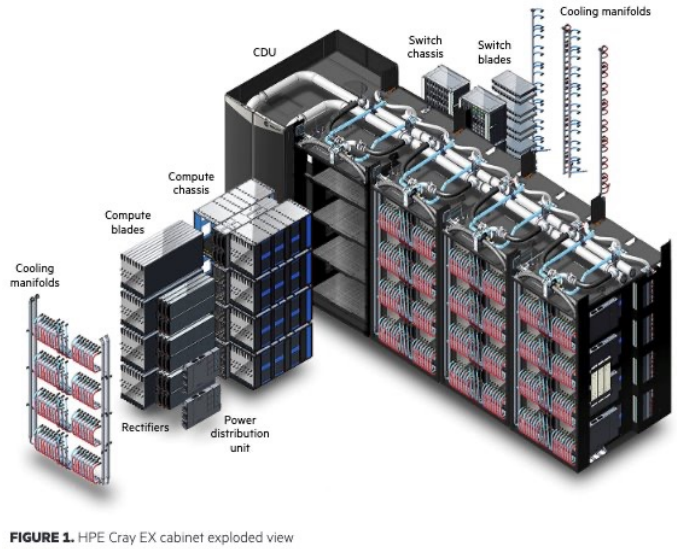

FIGURE 1. HPE Cray EX cabinet exploded view

5280xNVIDIA Grace Hopper GH200
384 GB GPU memory per node
512 GB CPU memory per node
288 CPU cores per node
~24,000 TB SSD disk
5280 x 1979 AI TOPS
10.4 M AI TOPS

University of BRISTOL

TOPS = measure of AI performance in trillions(!) of operations per second

Comparison to AWS—an AI user perspective

EC2 p5.48xlarge (8 x H100)
640 GB GPU memory
2,048 GB CPU memory
384 CPU cores
~28 TB SSD disk
$118 per hour
$2,832 per day

2 x HPE EX4000 Grace-Hopper blade

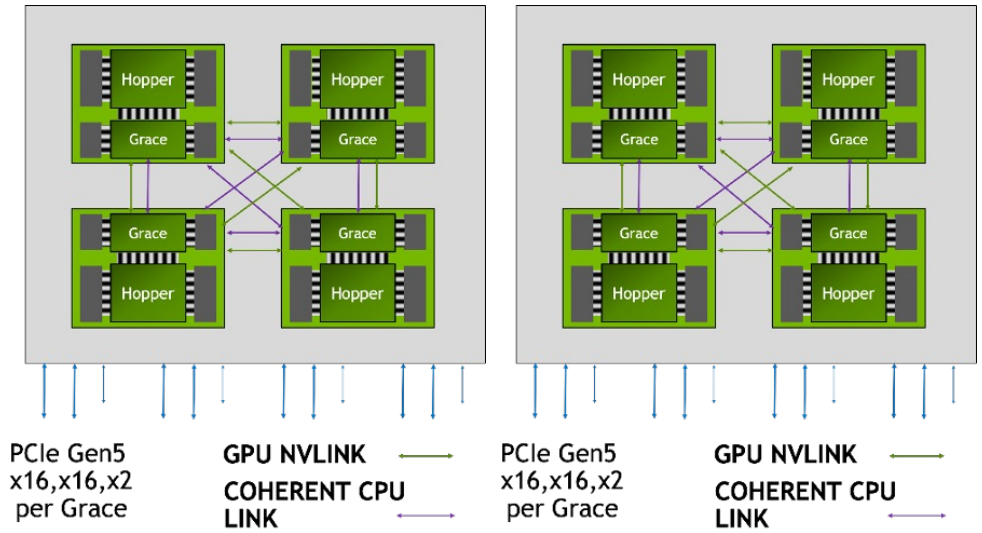

2 x Isambard-AI node (8 x GH200)
768 GB GPU memory
1,024 GB CPU memory
576 CPU cores
~24,000 TB SSD (shared)

Isambard-AI =~ 681 x p5.48xlarge =~ $80k per hour =~ $1.9m per day
…if you could get this many in one go
…(you would need a reservation – need a deal with AWS)

University of BRISTOL

Early Users & Next Steps

Examples of projects exploiting Isambard-AI

- Leveraging Alphafold for translational discovery research
- Performing a comprehensive data analysis of open-source skin cancer datasets to evaluate potential skin tone bias in the models developed from these dataset
- Machine Learning Models for Biomolecular Design
- BritLLM for producing training data, evaluation data, know-how, and open models aligned with UK interests
- Scaling Laws for Galaxy Images to investigate whether scaling laws hold outside of an ImageNet context, for images of galaxies
- AI-Powered Multi-Label Classification of Rare Great Ape Behaviours in Long-Tailed Video Data
- …

University of BRISTOL

AI Safety	1
Developing Generative AI	3
Developing General or Strong AI	1
Training / Retraining Large Lang...	5
Inference using Large Language...	6
Training Other Language / Text ...	1
Inference using Other Language...	1
Training Image / Video / Compu...	11
Inference using Image / Video / ...	10
Training Image Generation / Diff...	1
Inference using Image Generati...	1
Training time series / anomaly d...	1
Inference using time series / an...	1
Training general deep neural ne...	7
Inference using general deep ne...	6
Other	3

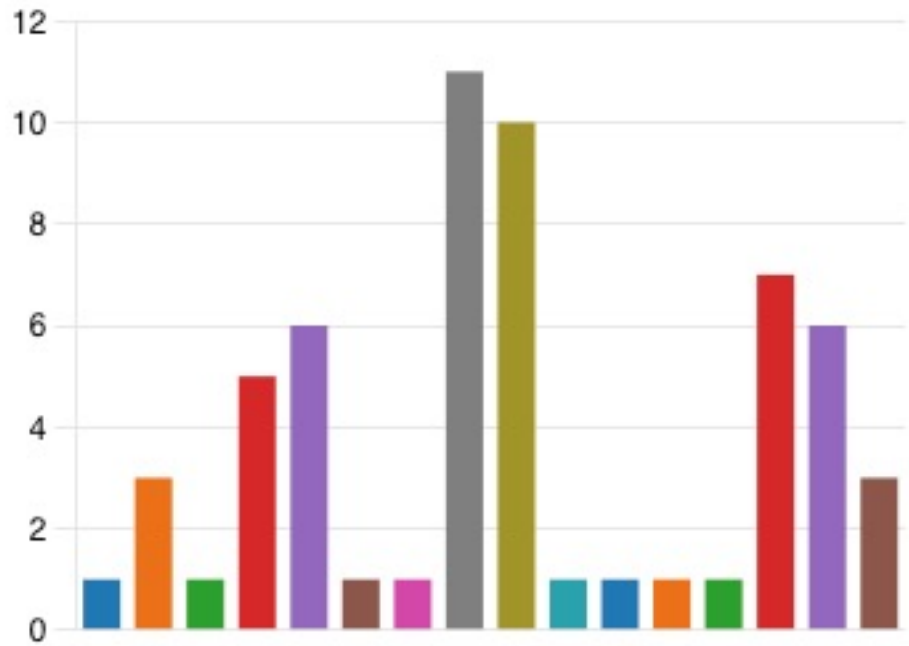

University of BRISTOL

PyTorch	17
PyTorch Lightning	5
transformers (Hugging Face)	8
Deepspeed	2
OpenCV	2
Keras	0
Tensorflow	3
JAX	4
scikit-learn	7
llama.cpp	1
flux.jl	0
Other	3

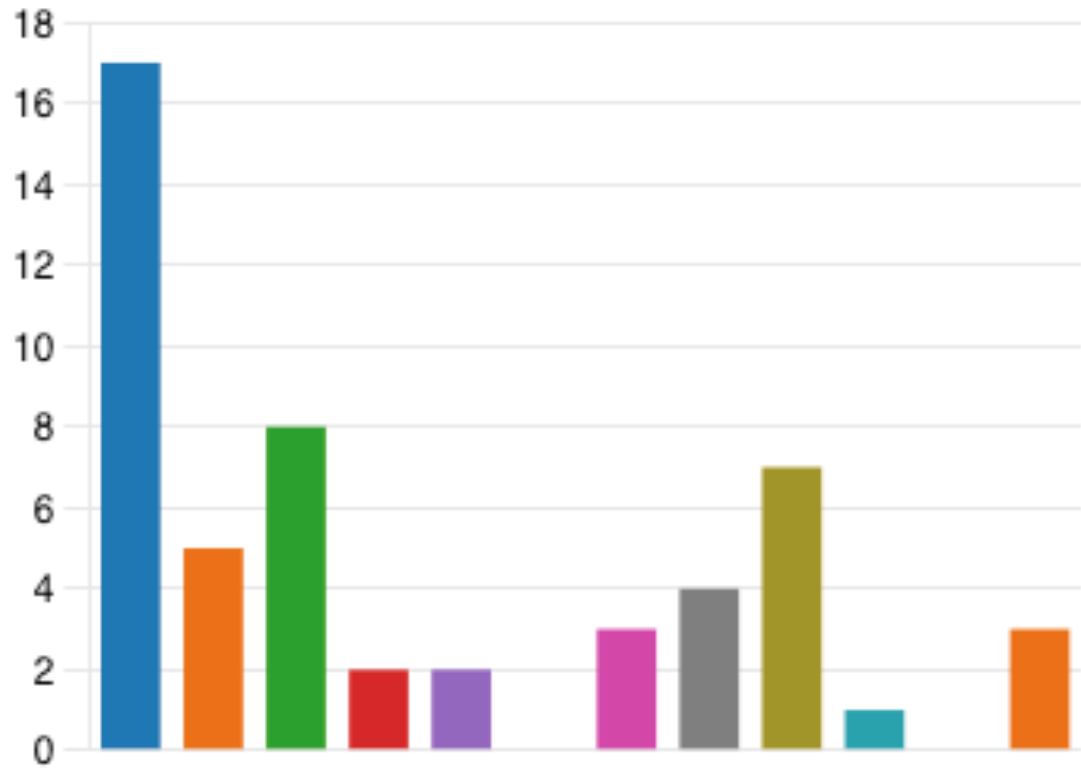

Evolving Software Stack for AI Users

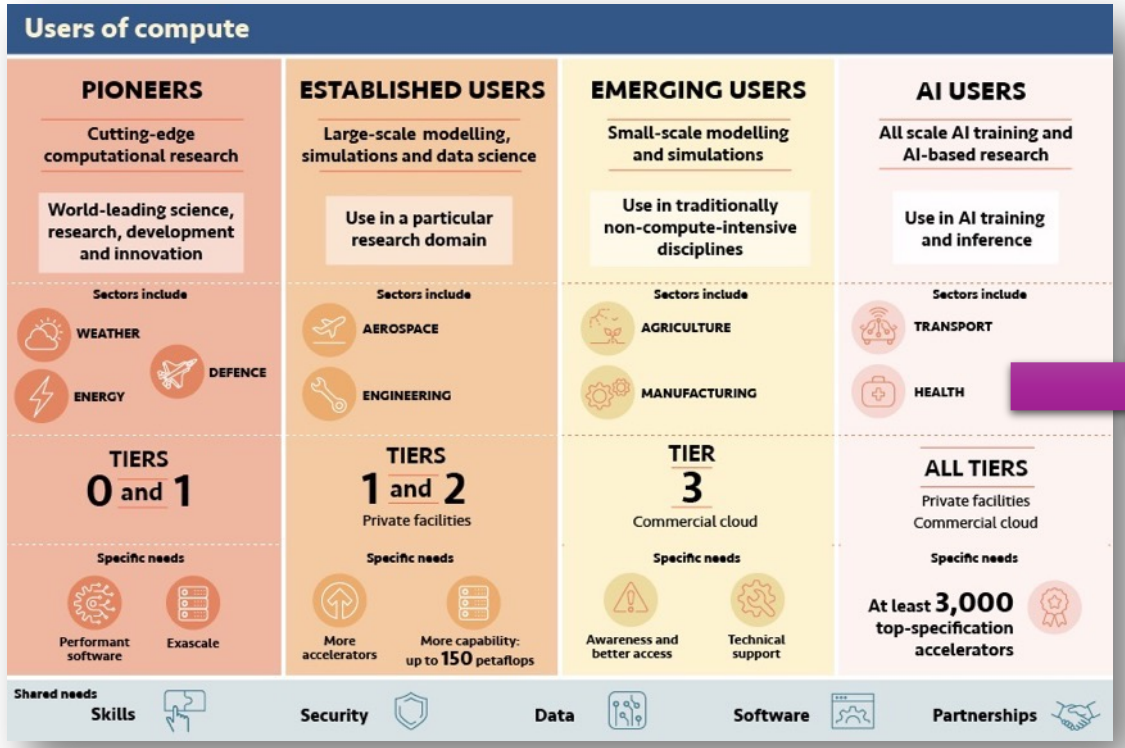

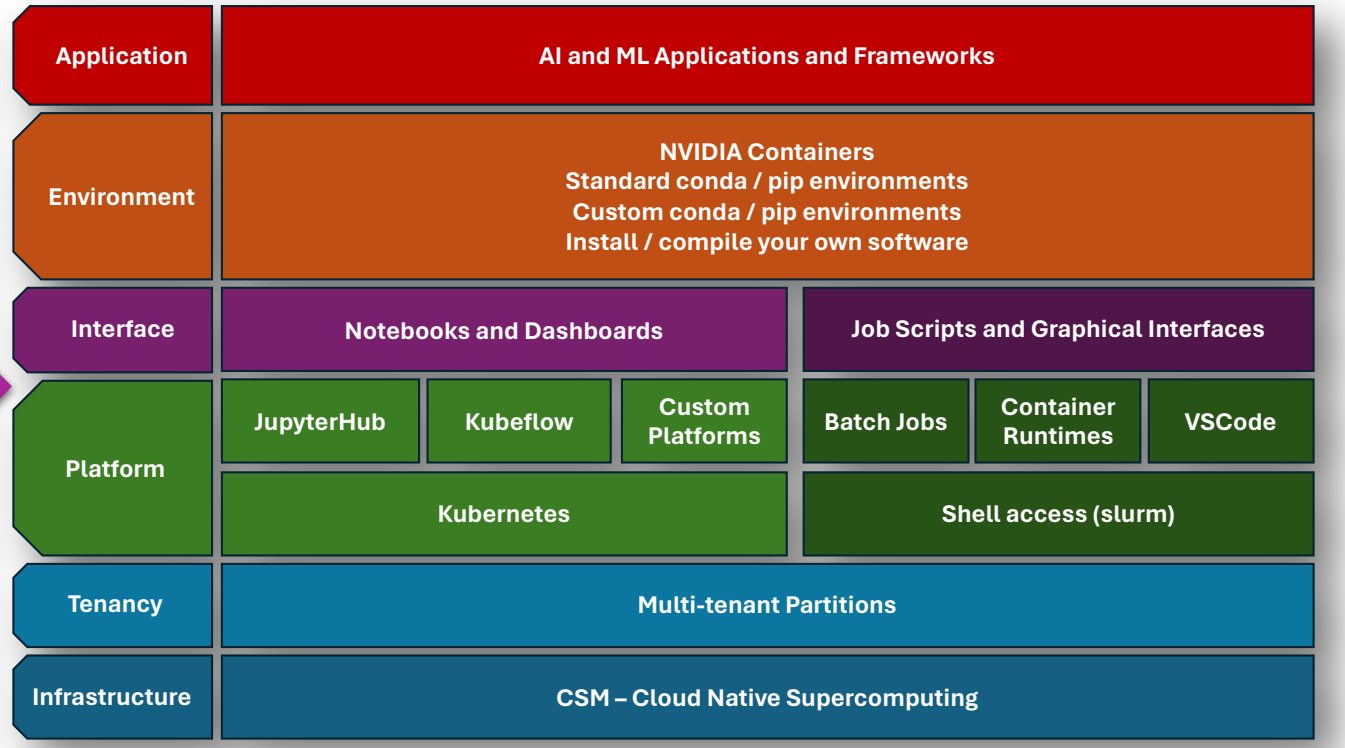

https://www.gov.uk/government/publications/future-of-compute-review/the-future-of-compute-report-of-the-review-of-independent-panel-of-expert

Converged Cloud and HPC software stack of Isambard-AI for diverse AI and ML platforms and hybrid workflows

University of BRISTOL

Scaling out in two phases

Phase 1 (~0.7 8-bit AI Exaflops)

Arrived in March 2024 – in Isambard 3 MDC
Piloting, on-boarding and staging services

1 x DLC EX2500 cabinet

21 blades (4-way Grace-Hopper)

42 nodes

168 GH superchips

12,096 Neoverse V2 Armv9 CPU cores

168 Hopper GPUs

21.5 TB CPU memory

16.1 TB high bandwidth GPU memory

37.6 TB total memory

AI high performance storage

~1 PB all-flash ClusterStor Lustre

Phase 2 (~21 8-bit AI Exaflops)

Arriving Autumn 2024 – new Isambard-AI MDC
Delivery of AI services

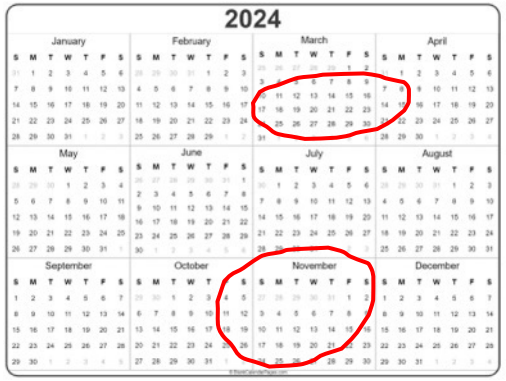

12 x DLC EX4000 cabinets

660 blades (4-way Grace-Hopper)

1,320 nodes

5,280 GH superchips

380,160 Neoverse V2 Armv9 CPU cores

5,280 Hopper GPUs

675 TB CPU memory

506 TB high bandwidth GPU memory

1.18 PB total memory

AI high performance storage

~27 PB all-flash storage!

(~20 PB Lustre, ~7 PB software defined VAST)

University of BRISTOL

Thank you

Stay tuned!

University of BRISTOL