A photograph of a server rack filled with multiple NVIDIA GPUs. The GPUs are gold-colored with the NVIDIA logo visible on each. The rack is dark, and the lighting highlights the GPUs.

JADE²

Tier 2 HPC

Wes Armour
JADE Day 2024
3rd October 2024



Welcome!

Quick introduction / update on JADE

More about today...



JADE 2

The Joint Academic Data Science Endeavour (JADE) is a national GPU facility dedicated to advancing world-leading research in machine learning and artificial intelligence research in the UK.

- 63x DGX V100 nodes (504 GPUs)
- DDN Ai400 flash storage (70TB)
- **1.1PB (Usable) spinning disk**
- InfiniBand EDR
- Hosted by the Hartree Centre



JADE 2 Consortium



JADE supports...

- World leading research groups across 20 UK universities.
- Over 600 active users.
- Providing over 3 million GPU/h to users every year.

When JADE 2 entered service, it was listed and the 16th most energy efficient supercomputer in the world.

GREEN



JADE 2 Funding model

Different to other Tier-2 services, JADE 2 has a unique funding model.

The CapEx for the machine was provided by EPSRC.

20% of machine time is given to the Hartree Centre for them to use or sell to commercial users to cover costs of running the machine.

Provided a low barrier to entry for partners.

JADE 2

End of Life...



JADE 2 has now reached end of life....

1st September – Stopped onboarding new users

1st November – Compute nodes will no longer accept Jobs

1st January – Machine will be physically decommissioned



JADE 2.5

Due to the age of JADE 2, continuing the service became prohibitively expensive.

As part of UKRIs interest in promoting a diverse computing landscape in the UK we secured funding to procure and run a technology pilot based on 24x AMD MI300 GPUs

JADE 2.5 Access

- Partner institutes can bid for time on the new platform. We expect to award about a month of time to each partner.
- We will have student “busy” weeks for students to access the machine.
- Those who use the machine will be asked to provide feedback on performance and usability.
- We hope to onboard first users in the next month.

JADE 2 Alternative...



Isambard-AI: Leadership AI computing for the UK



Prof Simon McIntosh-Smith (Director)
Bristol Centre for Supercomputing (BriCS)
University of Bristol

Isambard AI



What is Isambard-AI?

A new >£300M investment by UK Government in AI capability

Funding a **large new supercomputer optimized for AI** in a new, 5MW HPE modular data centre (MDC) facility in Bristol

- In Top 10 supercomputers in the world
- Had to be delivered extremely quickly – 12-18 months vs the 4-5 years usually required




What is Isambard-AI for?

- **To enable AI research in the UK:**
 - Training large language models
 - Large-scale inference
 - Foundational AI research
 - AI safety and understanding
 - Machine learning
 - Hybrid AI + simulation workflows
- Can run AI problems at a wide range of scales
 - From one GPU up to thousands



Status


- Isambard-AI phase 1 (168 Grace-Hopper GPUs) now running early access projects
- Calls to allocate for time on this system will be coming from UKRI/DSIT imminently
- Isambard-AI phase 2 (5,280 Grace-Hopper GPUs) will be commissioning in 2Q next year
- Calls for time on the full system will be issued next year



Example projects already running on Isambard-AI phase 1

- AI assisted drug design
- Wildlife conservation
- Heart health
- Tackling bias in AI skin cancer diagnosis tools
- Reducing memory footprint of transformers
- Improving large language model security

Contact us on brics-enquiries@Bristol.ac.uk from next week!



JADE Day Overview

[Home](#) [Programme](#) [Venue and More](#) [Bursaries](#) [Registration](#)



Welcome to JADE Day

Thursday 3rd October 2024

Worcester College, University of Oxford

JADE BACKGROUND & AGENDA

ABOUT JADE

Our first JADE Day is a national event for all JADE users across our 20 partner institutions.

JADE is owned by the University of Oxford with our service partners being the Hartree Centre, where JADE2, the UK's largest AI-focused supercomputer is hosted.

Based on NVIDIA's DGX servers, which have been specifically designed for high end, truly data intensive AI/ML workloads, the JADE service is comprised of 63 DGX servers, with 504 GPUs in total. Making JADE2 the de facto national AI compute service for UK HEI's. This technical and community driven focus makes the JADE service unique, not only in the Tier-2 landscape, but nationally too.



Thank you to **Lorjah Rajah** (JADE Consortium Manager)
for putting today together for us.

Please do provide feedback

JADE Day Agenda

Registration begins	09:30 - 10:00
Tea, coffee, pastries	09:30 - 10:00
Welcome & Introduction - Wes Armour, PI of JADE	10:00 - 10:15
Dave and Dave, Hartree Speakers	10:15 - 10:40
Break 1 - Tea, coffee	10:40 - 11:10
Steve Roberts, University of Oxford	11:10 - 11:35
Xiaowen Dong, University of Oxford	11:35 - 12:00
Lunch Break	12:00 - 13:00
David Llewellyn-Jones, Turing	13:00 - 13:25
Yishun Lu, University of Oxford	13:25 - 13:50
Breakout Session	13:50 - 14:50
Breakout Session Summary	14:50 - 15:00
Break 2 - Tea, Coffee	15:00 - 15:15
Sergio Perez, NVIDIA	15:15 - 15:50
Closing Remarks	15:50 - 16:00
Drinks reception / Networking Session & Poster Winner Announcement	16:00 - 18:00

JADE Day Breakout

One breakout session, 8 groups of about 10 people,
three topics...

Q1a. What did you have to change in your workflow to adapt to JADE?

Q1b. How would you have liked JADE to adapt to your workflow?

Q2a. What could your institution learn from JADE?

Q2b. What could JADE learn from your institution?

Q3a. What was most helpful in getting you started with JADE?

Q3b. What would you tell someone new to help them get started on JADE?

More to come later...

JADE Day Poster Prizes

1st Place - £100 Amazon voucher

2nd Place - £50 Amazon voucher

3rd Place - £25 Amazon voucher



Engineering and
Physical Sciences
Research Council





https://twitter.com/JADE2_HPC



<https://www.jade.ac.uk>

JADE²

Tier 2 HPC

