

AIRRFED: Federated access to compute (UK AI Research Resource Federation)

CIUK'24 Excalibur H&ES BO—Dec 6, 2024

Excerpts from AIRR Federation Demonstrator Project Report for the Delivery Board
(Oct 14, 2024)

By Sadaf Alam, Paul Calleja, John Taylor, Jeremy Yates, Martin Hamilton + thanks
to great work by teams from Bristol, Cambridge, Durham, Open Node & StackHPC
++ thanks to the Excalibur H&ES funding 2024-2025 +++ JISC campaigns & support



Bristol Centre for Supercomputing Documentation

Welcome

System Specifications

User Documentation >

Training >

Acknowledgements

Policies >

Service Status >

Getting Support

Welcome

This documentation site is the primary source of information about all Bristol Centre for Supercomputing (BriCS) facilities and services. Please use the site navigation menu to start browsing the site, or the quick links below to jump to key sections.

[Get started](#)

[Read the user documentation](#)

[Check service status](#)

[Get Support](#)

BriCS is home to:

[The Isambard-AI National AI Research Resource \(AIRR\)](#)

[The Isambard 3 Tier 2 Supercomputer](#)

Along with the above, the centre is responsible for its related services, such as the [BriCS Portal](#)

Around 60 projects and
300 users on
Isambard-AI phase 1
and Isambard 3 HPC in
< 6 months

Bristol Centre for Supercomputing

Choose your identity provider

University Login (MyAccessID)

Other Login (IdP of last resort)

BriCS (Administrators Only)

Having trouble logging in or are
unsure which identity provider to choose?

[Privacy Policy](#)

[Contact Us](#)

Details presented in SC24 S-HPC (Cybersecurity) workshop paper: “Federated Single Sign-On and Zero Trust Co-design for AI and HPC Digital Research Infrastructures”

<https://hpc.pnl.gov/S-HPC/2024/program.html>



Login with

Examples: University of Bologna, name@auth.gi

or



Login with eIDAS



Login with eduID Sweden



Login with

Bristol

University of Bristol
bris.ac.uk

City of Bristol College
cityofbristol.ac.uk

or



Login with eIDAS



Login with eduID Sweden



Bristol Centre for Supercomputing Documentation

Welcome

System Specifications

User Documentation ▼

Getting Started

Tutorials >

Guides >

Information >

Applications >

Training >

Acknowledgements

Policies >

Service Status >

Getting Support

Getting Started

This is a guide for anyone who has received an invitation to join an Isambard project.

 [Setup](#)

First time login through the portal.

 [Login](#)

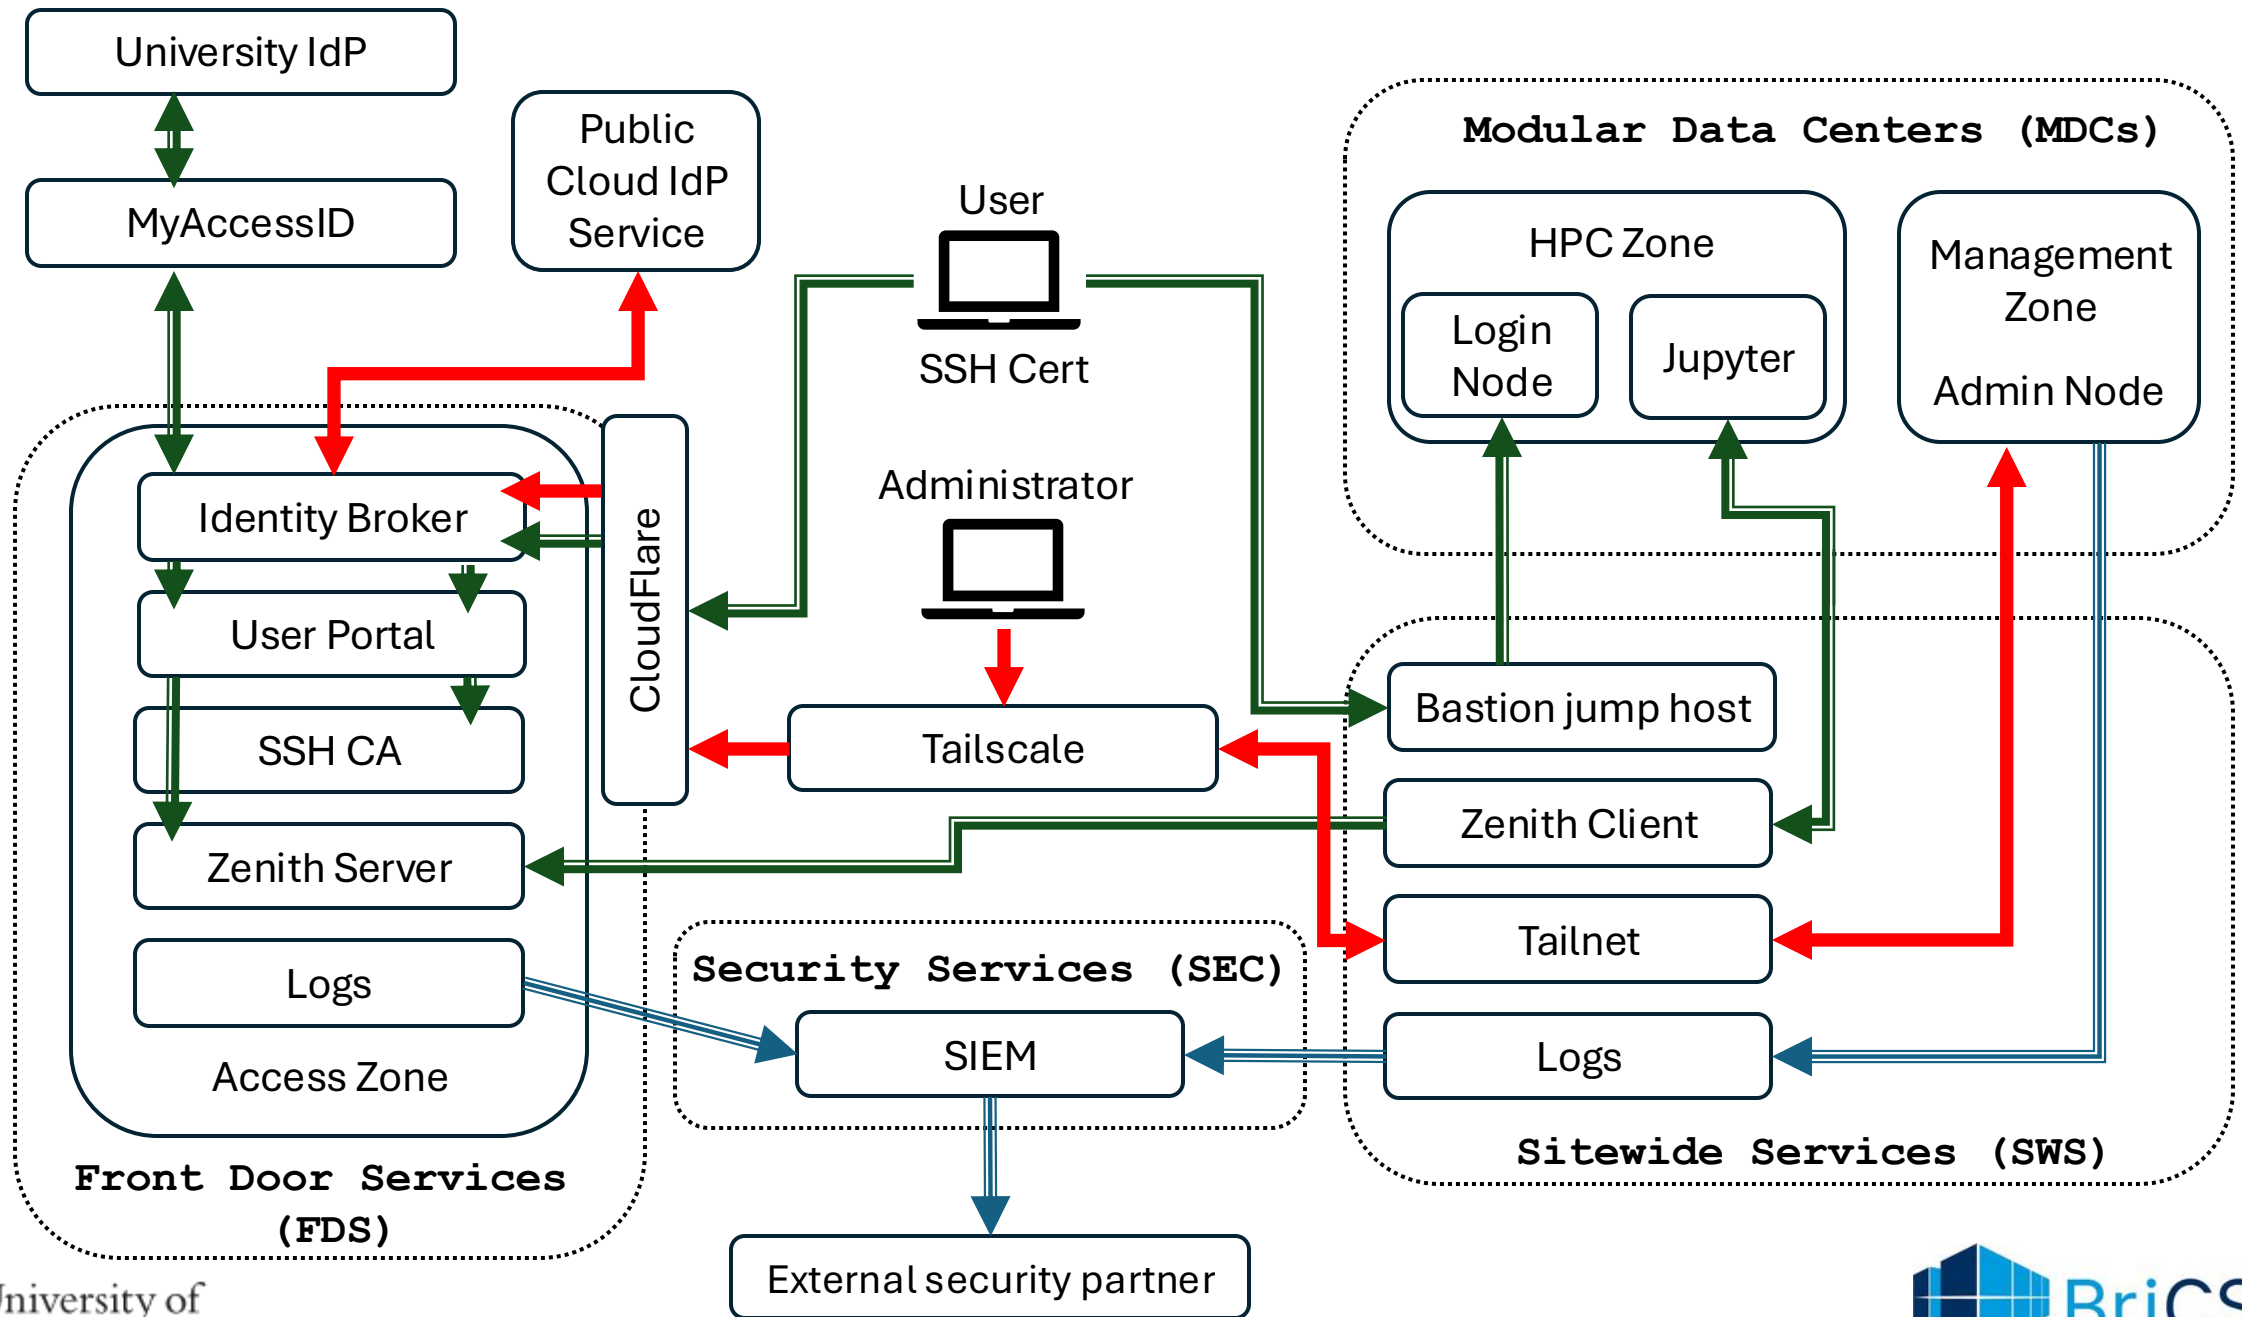
Logging into BriCS Supercomputers with SSH and `clifton`.

 [Manage project](#)

Manage your project on the portal (PIs only).

 [Get Support](#)

Opening support tickets on the helpdesk.




Selected Projects' Highlights (<6 months)

- 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

BriCS by CIUK25 (Isambard-AI)

Phase 1 (~0.7 8-bit AI Exaflops)

Arrived & accepted in March 2024
(Isambard 3 MDC)

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)

New Isambard-AI ~5MW MDC Delivery of AI services (s)

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

BriCS Isambard 3 Services CIUK'24

- 55,296 cores, 384 nodes, 6 racks of NVIDIA Grace 2x72 core @3.2GHz
- Slingshot 11 dragonfly network @ 200Gbps
- 2.0 PiByte of Cray/HPE ClusterStor E1000 Lustre storage
- Upgraded Multi Architecture Comparison System (MACS)
 - Added AMD Bergamo, Intel Sapphire Rapids + HBM, AMD GPUs...
 - Also integrated the newer parts of the MACS from Isambard 2
- Ready for national HPC allocation program (Jan'25)



DAWN



256 Dell PowerEdge
XE9640 nodes with DLC
to CPUs and GPUs



2x 48 cores Intel Xeon 4th
gen (Sapphire Rapids)
CPU



4x Intel Data Center Max
GPU 1550 (Ponte
Vecchio) GPUs



4x HDR 200 Gb/s per
node



~ 100 TFlops per node

DAWN

There are currently over one hundred early access projects running on Dawn, led by researchers from institutions across the UK

Early Access Performance

Very pleased with the BF16 results

- Distributed training works well
- Good scaling across nodes
- However, it is only using KYEV*

Results:

- Preliminary, with minimal tuning
- Favorable when using BF16
- Effective for AI workloads

* Know Your Environment Variables

IceNet & DAWN HPC

A probabilistic sea-ice forecasting system with real-world impact

Implementation

- High level libraries - facilitating the digital ecosystem using:
 - TensorFlow
 - Jupyter
 - Swift packaging
- Multiple approaches for delivery/integration of forecasts:
 - Forecast visualization dashboard
 - REST compatible API access
 - SDS outputs

Use-cases

PolarRoute

An eco-friendly marine vessel navigation system utilising IceNet sea ice forecasts as part of its input.

1. Multi Data Input
2. Climate Modeling
3. Variable Specifics
4. Route Planning

Journal of Adaptive Marine Operations Planning, 2020, 10(1)

Conservation applications

Used to forecast migratory behaviour of caribou and inform ship routes.

Rcc_018.ni

Unseen CT Scan

Finding

- A library designed for usage of agnostic AZ/ML backends
- Promotes code decomposition into generalised libraries, frees researchers from operational considerations
- Reproducible pipeline, easy porting to Open Stack cluster
- Previously run on AWS & JAGSML GPUs
- Observations on IBM with Intel GPUs
 - High data throughput with 32 data
 - Easy scaling with Swift
 - Training epoch time on test set < 10m

The Alan Turing Institute | British Antarctic Survey

Cambridge Solution

- **MyAccessID:** Federated Identity Provider
- **Keycloak:** Serves as the main broker for authentication
- **LDAP:** Stores user information for central account management.
- **Waldur:** Handles resource offering, project management, and integrates with SLURM. Provides an interface for Principal Investigators (PIs) to manage research projects and invite users
- **Open OnDemand:** Provides web-based HPC/AI access
- **SLURM Cluster:** Provides SSH Certificates based access to resources for job scheduling
- **K8s:** Providing K8s as a service

OpenStack



Work In Progress & Next steps for Cambridge

Work Completed:

- GitOps managed Waldur deployed
- User Onboarding and Offboarding flow defined
- Keycloak Integration
- Creation of test environment using OpenStack:
Providing Jupyter Notebooks for the users to use the DAWN resources through Open OnDemand web portal

Next steps:

- Integration of Atlassian Service desk with Waldur
- Connect MyAccessID with Keycloak
- Installing Waldur SLURM agent for reporting back the project resource usage to Waldur portal
- Present Kubernetes environment through Waldur
- Integration Test and Validation: Early user onboarding and feedback
- Providing K8s cluster



Progress Report

WP1 (IdP Thin Layer Federation)

- Thin federation layer in place for identity at Bristol; POC completed for Cambridge by StackHPC
- SSH CA demonstrator with federated IdP – in place at Bristol and in progress at Cambridge
- Cambridge integration of federation POC in progress
- Single sign-on with Zenith (a component of Azimuth – allows secure tunnelling) installed at Isambard-AI
- Cambridge Azimuth installation in progress with MyAccessID
- Prototype of (non-federated) IdP of last resort at Bristol – discussions ongoing at Cambridge and Durham as to what they will use.
- Bristol on-boarding in progress – tuning and hardening the MVP as well as improving the project/user experience
- Validation of slurm (batch) with SSO and ssh CA at Bristol and also via POC by StackHPC for Cambridge
- Jupyter notebook with SSO (stretch goal achieved) at Bristol and also via POC by StackHPC for Cambridge
- New Secure Access, Data Security and Cybersecurity work – no award of funds

WP2 (Data Transport using VAST and Other SSD systems)

- VAST SSD installations complete at Durham and Bristol; Cambridge TBD
- SSD resource access in progress (FTS3 demo subproject is needed)

Outreach

- Outreach via RSECon24
- CIUK and ExCALIBUR Wrap Meeting



We Ought to Fund Operationalisation for our Stakeholders and Users

- MyAccessID and IdP of last resort
 - Getting HEIs to publish R&S fields is an involved process (🙏 JISC)
 - Sometimes local security actions are blocking (for a few)
 - Non-academic users need to be catered for – Govt (high priority) – by an IdP of last resort (not federated, which is an issue)
- Hardening of demo airport (AIRR portal for national AI RR resources)
 - Accounting and authorisation capabilities included
 - Connecting to site and service specific portal (not centralised)
 - Bristol solution called OpenPortal being evaluated
- Several opportunities for federated data transfer platforms



Thank you

Please contact any member organization of the AIRRFED project if you are interested in participating—we are planning 3-4 outreach events.

Come and join us at 2 pm today for the Federation NetworkPlus breakout session