UNIVERSITY OF CAMBRIDGE

Building Federated AI Infrastructure While Preserving Site Autonomy

Federate the paperwork but keep the keys to your machine

Deepak Aggarwal, Research Computing Services

IRIS-TWG Meeting
8 September 2026

OVERVIEW

Today's Discussion

01 **The need, and the challenges**

Why a national compute resource is hard to run as one service across multiple independent sites

02 **The architecture for AIRR**

Federated identity, Allocator portal, site portal, OpenPortal connecting the portals

03 **Our approach, and where autonomy sits**

What we federated, what Cambridge kept, and how each challenge was answered

04 **It working in practice**

An allocation followed end to end

05 **The closed loop, and what it costs**

What allocators, the service provider and users each get

06 **What comes next**

Rollout to the NCRs, Dawn to Zenith

THE NEED

The UK AI Research Resource (AIRR)

£300M of investment in two of the UK's largest supercomputers, configured for AI workloads and serving research, industry and government.

THE TWO AIRR SYSTEMS AND THE CLOUD PILOT

Isambard-AI

UNIVERSITY OF BRISTOL

5448 Grace-Hopper GPUs

Dawn

UNIVERSITY OF CAMBRIDGE

1024 Intel Data Centre Max 1550 GPUs

Zenith follows as the next-generation system, with AMD Instinct MI355X GPUs

AIRR Cloud

STFC HARTREE CENTRE

AIRR Cloud Pilot launched in July 2026 to provide cloud credits for SMEs that can be used with one or more approved cloud service providers

The UK National Compute Resources (NCRs)

6

NCR sites

Separate from the AIRR systems above

THE NEED

Why We Are Building It This Way

AIRRFED — UK AI Research Resource Federation

EPSRC EXCALIBUR GRANT PROJECT, COMPLETED 31 MARCH 2025

A single pane of glass for secure federated **discovery, allocation, access and usage** of UK national AI and HPC resources, for both allocators and users.

> **A national resource has to look like one resource to the researcher and to the funder, while being delivered by independently operated facilities, each with its own policies, security posture, and hardware.**

FUNDERS NEED

To run calls, award compute resources across multiple sites, and see what was actually consumed without a quarterly spreadsheet from each site.

One cross-site view of awards and usage

RESEARCHERS NEED

One login with their own institutional credentials, and a working account within days of the award rather than weeks.

Institutional login, account in days

SITES NEED

To keep control of accounts, quotas, scheduling and security, and to reuse the operational processes they already run.

Local policy and control preserved

THE CHALLENGES

What We Are Tackling

01

Identity across institutions

Users arrive from hundreds of universities and companies, with no shared account namespace between them or with the sites.

02

Allocation decided nationally, enforced locally

Allocators awards resources; sites enforce them.

03

Onboarding at the pace of the award

A project approved today should not wait on account creation and access.

04

A clean split of responsibility

Allocators decide who gets access and on what terms. Sites provision against that decision — they do not want to adjudicate it.

05

Unlike sites, one catalogue

Different compute resources (hardware) have to be described as offerings a researcher can compare and apply for.

06

Support with no single owner

Whether extending a project, changing its membership, or closing it down, the researcher should have one place to go, not a different process for every site.

THE ARCHITECTURE FOR AIRR

The Whole Picture

ALLOCATOR
UKRI and DSIT

AIRRPortal

Calls, applications, assessment and awards (GPU-hours granted across AIRR sites)

Isambard-AI offering **Dawn offering** **AIRR Cloud offering**

↓ **project, allocation and team list as structured data via OpenPortal over encrypted connections, not an email**

SITE PORTAL
Cambridge RCS

Keycloak
Brokers identities via MyAccessID

Waldur
Holds projects, allocations and the site's offerings

Jira Service Desk
Tickets, operator approval and the audit trail

↓

INFRASTRUCTURE
Dawn, site-owned

LDAP
Local accounts, one per person

Compute-Slurm
Project accounts and GPU-hour limits

Storage
Project directories and quotas

↑ **Usage reported back up to site portal and allocator portal**

OpenPortal: Narrow-Permission Agents

OPENPORTAL

A distributed infrastructure management protocol. Portals and infrastructure communicate through peer-to-peer agents, and each agent handles only the operations it is permitted to perform.

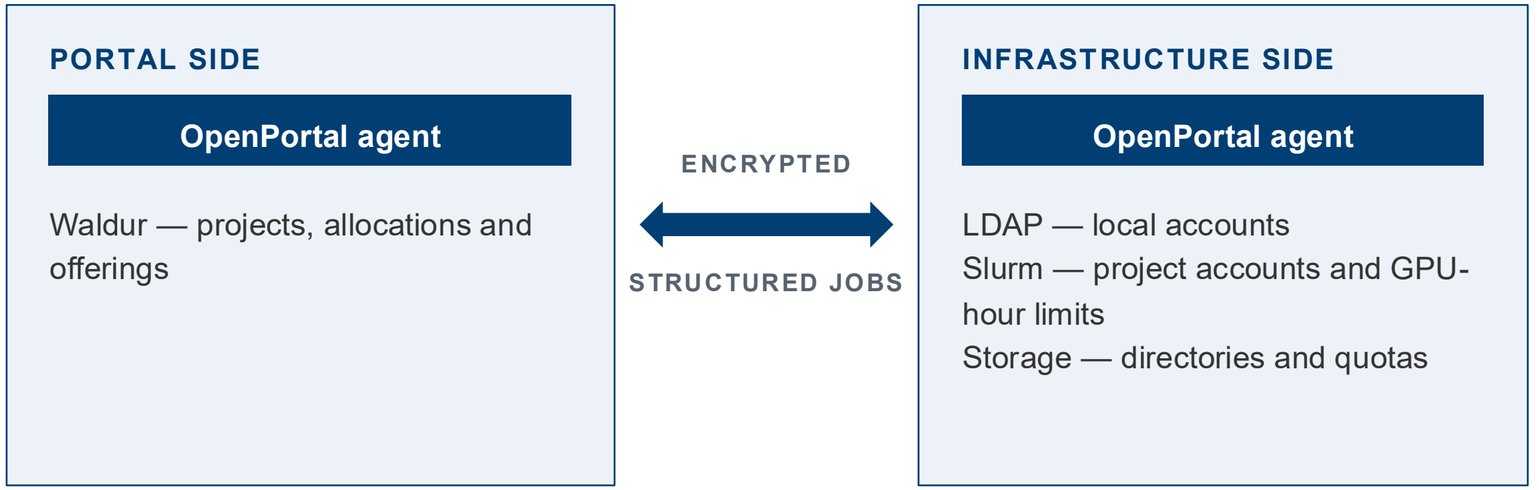

Agents at both ends coordinate every task; neither side holds credentials for the other.

WHAT THE AGENTS COORDINATE

- **Account creation**
- **Project management**
- **Filesystem provisioning**
- **Usage reporting**

Each agent is a small, statically compiled Rust executable, deployed at both the portal and the infrastructure.

THE ARCHITECTURE

Access Portal: Solution Components

Three concerns, kept separate: who you are, what you may use, and how you reach it.

AUTHENTICATION

Who you are

Home institution IdP

Reached through eduGAIN

MyAccessID

Issues a persistent federated identity

Keycloak

Identity broker at Cambridge, we never hold a password

Users without MyAccessID self-register with the Guest-IdP -> a local Keycloak account, the identity of last resort

AUTHORISATION

What you may use

Waldur

Allocations, monitoring and user management, across all service offerings

Jira Service Desk

Allocation and onboarding tickets, with the audit trail

PIs and users see their own allocations, and RCS manages every offering on one platform

ACCESS AND APPS

How you reach it

Open OnDemand

Browser access to HPC and AI resources

Kubernetes

Trusted Research Environments and inference as a service

SSH

Direct login and job submission, unchanged

THE ARCHITECTURE

Access Portal: Solution Implementation

Waldur deployment

Helm charts deployed by FluxCD into Kubernetes, with configuration held in Git so it is repeatable, auditable and self-healing.

MyAccessID and Keycloak

Separate realms for portal users and LDAP users, and a Guest-IdP realm behind both for people their institution cannot vouch for.

Waldur and Jira Service Desk

Allocation, project membership and onboarding raise tickets automatically; the ticket is the approval and the audit record.

Access to Open OnDemand

LDAP identities sign in through the brokered realm; notebooks, terminals and containers run against the same allocation as batch.

WALDUR GITOPS WORKFLOW

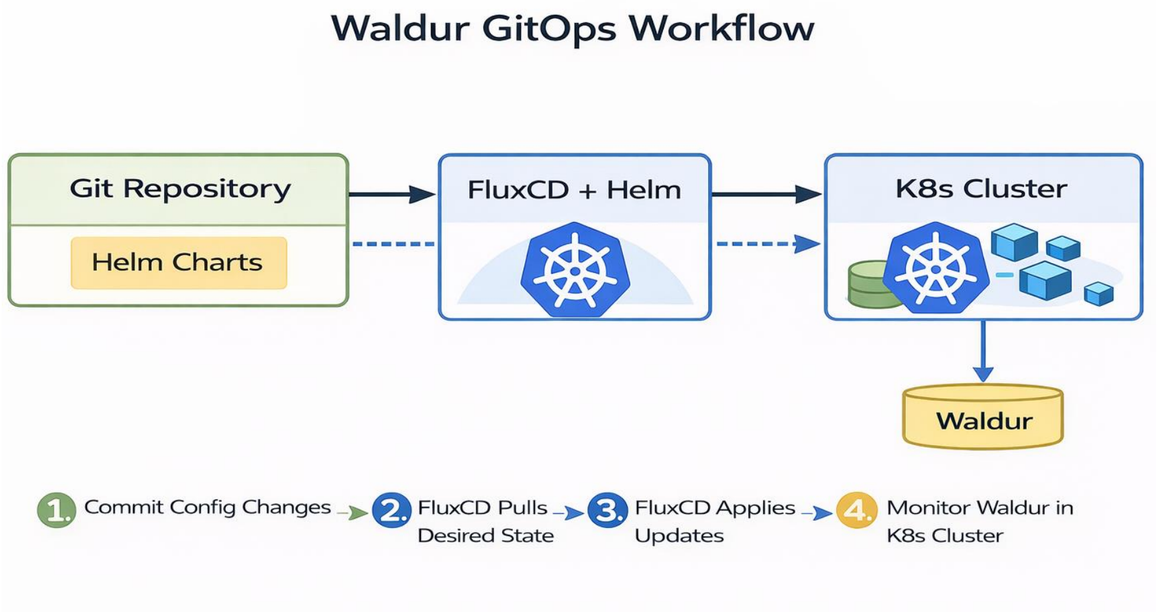

KEYCLOAK REALMS AND THE GUEST IdP

Realm 1 — portal users

Waldur
MyAccessID

Realm 2 — LDAP users

LDAP accounts
Open OnDemand and SSH
MyAccessID

Realm 3 — Guest-IdP

For users their institution cannot vouch for federally

THE ARCHITECTURE

Identity and Access Architecture for AIRR at Cambridge

MyAccessID	Guest-IdP	CSD3 LDAP
Home institution identities via eduGAIN	Self-registration: email as user id, plus password	The site's existing local accounts

↓ KEYCLOAK IDENTITY BROKER — THREE REALMS, THREE TRUST DECISIONS

REALM	WHO CAN LOG IN	BEHAVIOUR	WHAT IT UNLOCKS
`Realm1`	**MyAccessID users, or Guest-IdP users**	Open by design, so a newly invited user can reach their project before any local account exists	The AIRR Dawn Portal
`Realm2`	**Only users already in LDAP**	Login via MyAccessID or Guest-IdP is accepted only if the account exists locally; TOTP available	Open OnDemand, and SSH access
`Realm3`	**Anyone, by self-registration**	The Guest-IdP itself — identity of last resort for organisations not registered with MyAccessID, such as UKAEA and the Met Office	Feeds Realm1 and Realm2

OUR APPROACH

How We Chose to Build It

Reuse what already works operationally, and automate around it rather than through it.

Key Characteristics

- Platform federation via AIRRPortal
- Local operational control retained
- Human-in-the-loop with full audit trail
- Reuse existing RCS processes and tooling

OUR APPROACH

Where Autonomy Sits

FEDERATED ACROSS SITES

- User identity, via home institution credentials
- Project and allocation records
- Usage reporting and statistics
- The shape of the user interface
- Onboarding workflow and invitations

OWNED BY THE SITE

- Local accounts, LDAP and account policy
- Scheduler configuration and queue policy
- Filesystems, quotas and data placement
- Security posture, ISO 27001 scope, TREs
- Which services the site chooses to offer

Sites agree on interfaces, not on implementations. Cambridge kept its single-user-account policy and legacy CSD3 LDAP while joining the federation unchanged.

OUR APPROACH

How the Approach Answers Each Challenge

CHALLENGE	WHAT WE DID	WHAT IT CHANGED
Identity across institutions	MyAccessID brokered through Keycloak, plus an IdP of last resort for organisations outside eduGAIN	Researchers use their own university login. We never hold a password, and no account request is needed to reach the portal.
National allocation, local enforcement	Waldur holds the project and allocation record; identifiers map cleanly from portal to scheduler	An AIRRPortal award becomes a Slurm account and a Lustre quota with a traceable name, not a re-keyed request.
Onboarding at the pace of the award	Portal events raise Jira Service Desk tickets automatically, into the RCS processes that already exist	No new operational process to learn, and no email thread between the funder and the site team.
A clean split of responsibility	Access decisions stay in AIRRPortal; narrow-permission OpenPortal agents provision against them	The site never adjudicates who gets in, never hands over admin rights, and allocators see consumption live.
Unlike sites, one catalogue	Each site publishes its own offerings in Waldur, described in the same terms but on its own configuration	A researcher compares comparable things, and a site can add or retire an offering without a national change.
Support with no single owner	One RCS service desk queue for portal and cluster, plus the MyAccessID attribute checker for self-diagnosis	The user asks in one place, even when the fault sits with their own institution's identity provider.

WORKED EXAMPLE

Triggered from AIRRPortal

1. Triggered from AIRRPortal

Project Allocation

TEST-PROJ-001 – Test Project for AIRR (DAWN – Cambridge)

Start date: Dec 8, 2025 End date: Mar 22, 2026

Award: **10,000 GPU hours**

Team members (5)

LJ YF HJ YY YW

User Added to Project

New member:

HA hai@university.ac.uk

Role: Project member

Information passed

- Allocation (GPU-hours)
- User details (name, email)
- Project start / end dates

2. Cambridge Operations Workflow

All changes are handled through existing RCS support processes

JIRA **Automatic ticket creation**

Project Allocation Ticket

- Create / update SLURM project (group) account
- Allocate GPU-hours
- Configure storage allocation
- Set quotas and limits

RCS Operator

User Onboarding Ticket

- Create LDAP user account
- Add user to SLURM project
- Set project permissions
- Provide access to storage

RCS Operator

Audit & Traceability

- Full audit trail in JIRA
- Change approvals and records
- Alignment with RCS operational policies

3. Backend Systems – Dawn

LDAP

Username (dn)	Full name	Email	Groups	Status
dn-da61	Alice Smith	alice@example.org	airr-p77	Enabled
dn-da62	Bob Lee	bob@example.org	airr-p77	Enabled
dn-da63	Carol Chen	carol@example.org	airr-p88	Enabled
dn-da64	David Patel	david@example.org	airr-p88	Enabled

→ User identities created in LDAP

Slurm Accounting

Account	User	Resources
airr-p77-dawn-gpu	dn-da61	GPU partition, GPU hours
airr-p77-storage	dn-da61	Lustre storage quota
airr-p88-dawn-gpu	dn-da63	GPU partition, GPU hours
airr-p88-storage	dn-da63	Lustre storage quota
…	…	…

→ Project resources and user associations in Slurm

Filesystem (Lustre)

Project	Path	Size	Used	Quota	Status
airr-p77-storage	/lustre/airr-p77	5 TB	1.2 TB	5 TB	OK
airr-p88-storage	/lustre/airr-p88	5 TB	0.8 TB	5 TB	OK
…	…	…	…	…	…

→ Project & user directories created, quotas set on Lustre

WORKED EXAMPLE

End-to-End Cambridge Workflow

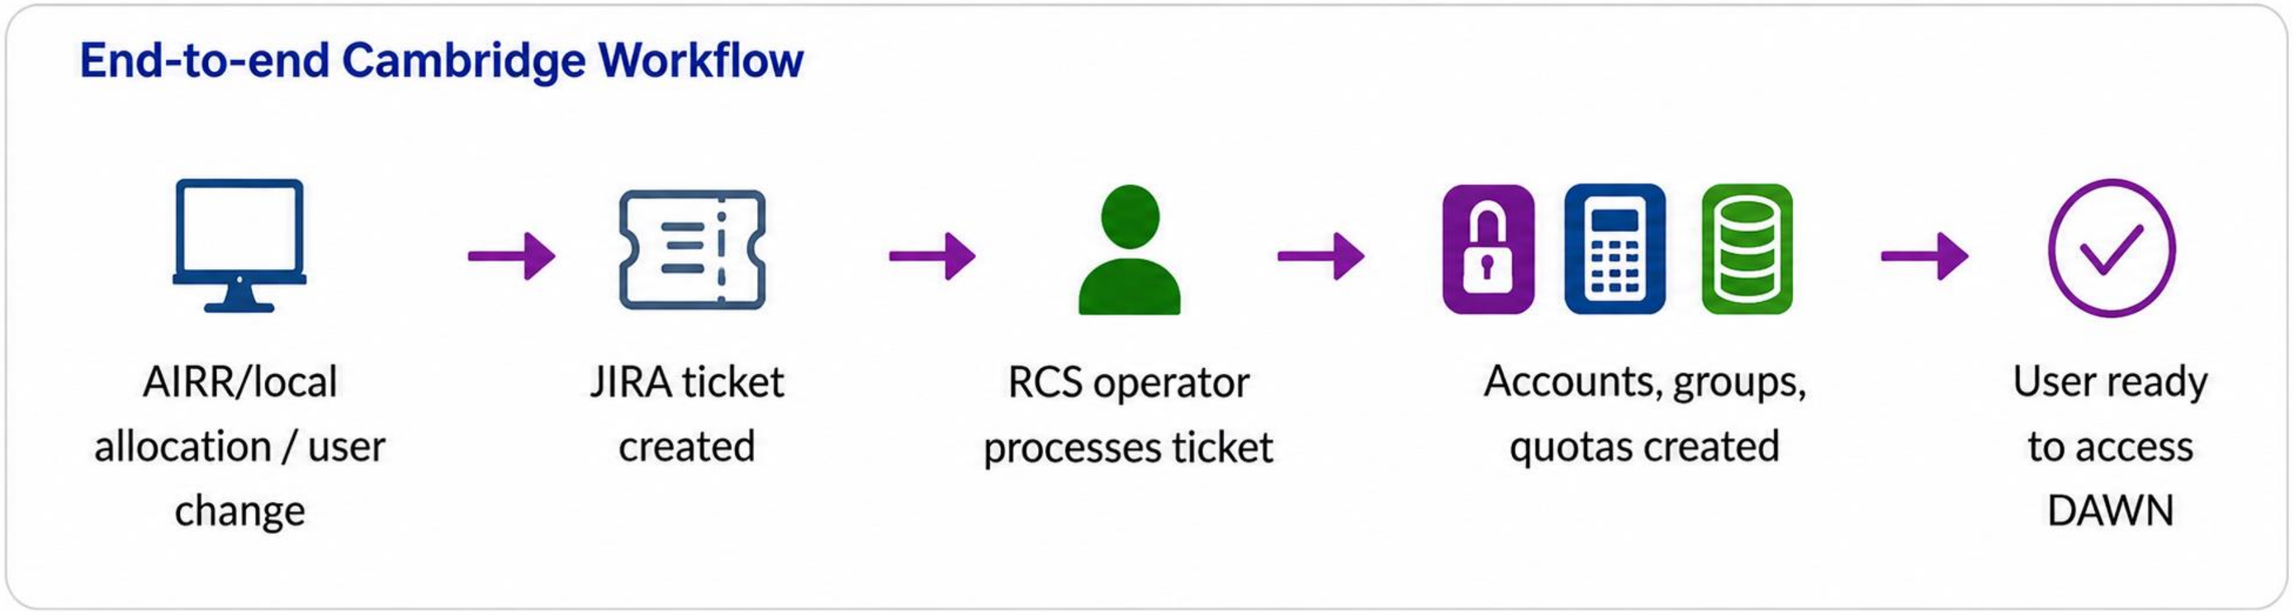

WORKED EXAMPLE

The User Journey

From institutional login to reported usage, without a single account request.

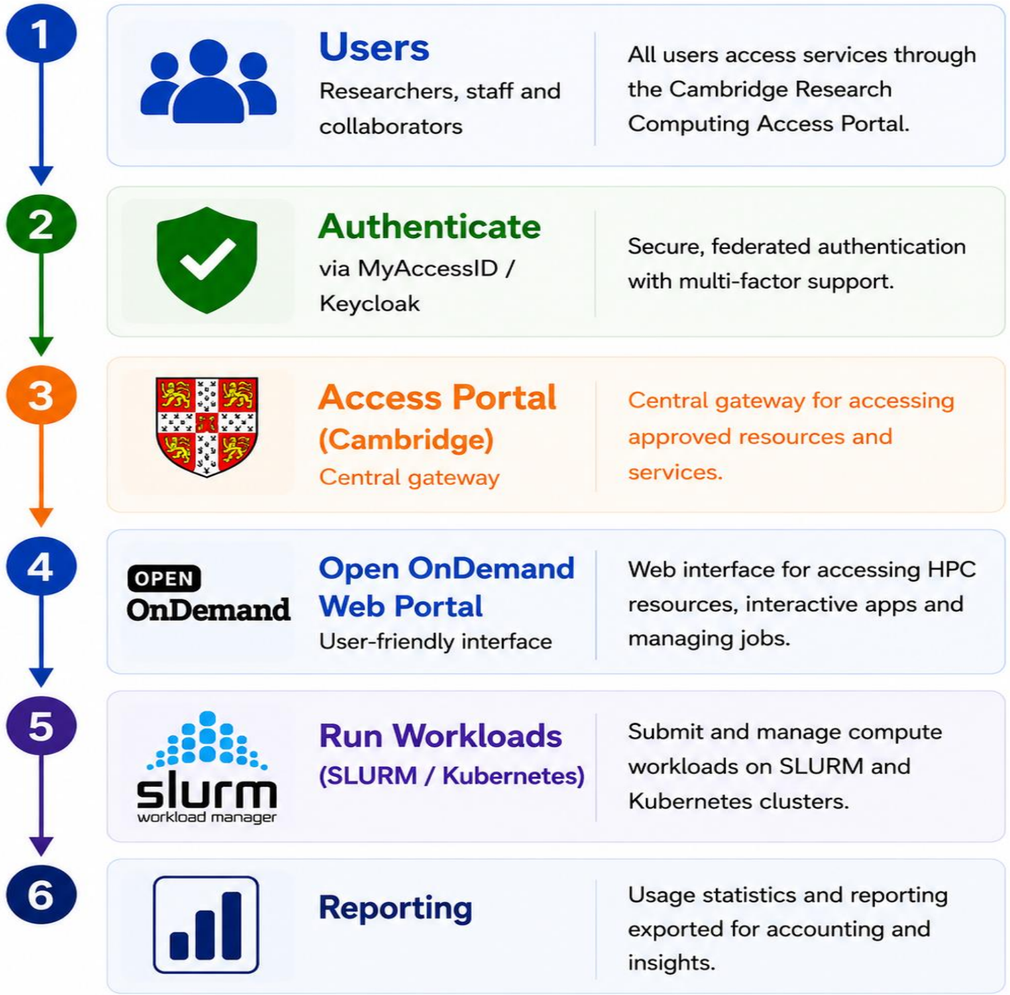

THE CLOSED LOOP

Allocators, Service Provider, Users

ALLOCATORS

UKRI and DSIT

Run calls, assess applications, award GPU-hours across both sites

allocation →

SERVICE PROVIDER

Cambridge RCS

Provisions accounts, projects, quotas and access through its own processes and systems

access →

USERS

PIs and researchers

Manage their own teams, run workloads, see what they have consumed

↑ **USAGE AND REPORTING CLOSES THE LOOP**

Consumption is reported from the cluster back to the site portal and up to AIRRPortal, so the next allocation decision is made on evidence rather than on a quarterly return.

Each party holds only what it should: the funder holds the decision, the site holds the infrastructure, the PI holds the team. Nothing depends on an email thread between them.

What Each Party Gets

Allocators

Run calls and award across both sites from one place

Redirect a request to the site with capacity

See live consumption instead of a quarterly report

An auditable record of every allocation decision

Service provider

Requests arrive as structured tickets, already complete

Existing RCS processes and tooling are reused, not replaced

Operational control of the cluster stays at the site

PIs manage their own teams, so admin load does not scale with users

Users

One login with home institution credentials

Working access in days rather than weeks

Notebooks, terminal and containers against the same allocation

Visible remaining budget, without asking anyone

Why This Matters

Better researcher experience

Simple, consistent access to resources

Faster onboarding

Streamlined processes and automation

Efficient resource utilisation

Unified visibility and allocation across sites

Stronger governance & compliance

Clear policies, auditability and accountability

Scale together as a federation

Share best practices and capabilities

RUNNING IT

Lessons from MyAccessID Integration

Registration takes time	Initial registration of the service took around 20 days. Plan it as lead time, not as a task in the sprint.
Redirect URLs are exact	Fully-qualified redirect URLs are required and wildcards are not supported — `https://example.org/*` is invalid. Every URL must be listed individually.
Mapping settings are critical	The integration itself was straightforward, but attribute mapping and sync mode settings are where the behaviour is decided.
Support is responsive	The MyAccessID support team has been quick to engage throughout, which materially changed how fast problems were resolved.

Access to the portal depends on the email address matching between the project invitation and the mail attribute released by the user's institution.

THE USER LOGS IN WITH

`abc1234@university.ac.uk`

THEIR IdP RELEASES

`firstname.lastname@college.university.ac.uk`

WHAT COMES NEXT

Rollout to the NCRs

The plan is to roll out the same site portal solution to the National Compute Resources. The architecture transfers; the local decisions do not.

TRANSFERS DIRECTLY

The Portal deployment

The MyAccessID and Keycloak brokering design, including the dual-realm split

OpenPortal agents for accounts, Slurm and filesystem provisioning

The Jira Service Desk automation and onboarding workflow

STAYS A LOCAL DECISION

How existing local accounts map to federated identities

Scheduler and reservation policy for the site's own workloads

Storage quotas, and data placement

Which service offerings appear in the site catalogue

WHAT COMES NEXT

Dawn to Zenith: The Platform Outlives the Hardware

TODAY

DAWN

1,024 Intel Data Centre Max GPUs, in service since November 2023 and on its way out

same portal →

NEXT

ZENITH

The next-generation Cambridge system, joining the same federation on the same interfaces

CARRIES ACROSS UNCHANGED

Federated identity, projects and allocation records

The site portal, its workflow and its support processes

OpenPortal agents, repointed at new clusters

CHANGES WITH THE HARDWARE

Accelerator architecture and the software stack around it

Scheduler partitions and queue policy

Capacity, and therefore what can be allocated

The site swaps the machine underneath not the way users get to it

Cambridge AIRR Team

Service Desk Team
Support and user operations

Babatunde Oyewo
DevOps Engineer

Eduardo Veiga
Delivery Manager

Stuart Rankin
Senior HPC Sysadmin

Deepak Aggarwal
Principal HPC Systems Manager

Simon Hartley
Senior AI RSE

K. Harrison
Senior AI RSE

Ryan Daniels
Principal AI RSE

Jonathan Steel
Keycloak Admin

UNIVERSITY OF CAMBRIDGE

THANK YOU

Sites agree on interfaces, not implementations which is what made the federation possible at all.

Deepak Aggarwal, Research Computing Services

da616@cam.ac.uk