FEDERATED IAM FOR COMPUTE SERVICES

National Federated Compute Services

Connect, coordinate, collaborate

epcc

IAM for NFCS

- IAM = Identity and Access management
- IAM is a subset of Authentication and Authorization
  - AA is very wide-ranging concern
  - IAM is more limited, Entities being managed are humans
  - Machine<->machine AA etc. out of scope
- What will a National Federated Compute Service need in terms of IAM?
  - UK research landscape already contains significant IAM and compute services
  - What are the gaps between what we have and what we need?

Federation

- A federation is where multiple distinct (and different) organizations work together
- Not (just) a technology problem:
  - Common Federation policies and agreements
  - Staff and procedures to support and operate the federation
  - Forums to evolve standards, technology and policies
- Still room for differences
  - Scope of federation activity is limited.

Approaches

Type	Rational	Organisation	Implementation
Distributed infrastructure (public cloud)	• Highly scalable • Fault tolerant • Economies of scale	• Single provider e.g. AWS • Distributed implementation	• Proprietary software • Accessed through single entry-point/API
Grid	• Fungible/interchangeable resource • Seamless/dynamic transfers of workloads	• Multiple providers • Strong governance to set technology/policy standards	• Highly standardised • Common software stack

Federation	• Diverse flexible ecosystem
	  • Different communities
	  • Different approaches
	• Cooperation on some cross-cutting concerns
		• Multiple providers
		• Governance for federated aspects only
			• Open standards
			• Defined interfaces
			  • Allow multiple implementations

Independent	• Happens by default	• Multiple providers • No explicit coordination	• Independent • Organic alignment only

Human/digital interfaces

- Primary mechanism for **people** to interact with digital resources is the web-browser.
  - Need to address federated web based IAM
- Compute services (particularly HPC) also commonly make use of SSH based command-line access.
  - Importance of this may be declining relative to web-based interfaces
  - SSH has its own authentication methods
    - Credential management, authorizations etc. for command line accounts now usually managed via web interfaces
    - Known mechanisms exist for mapping command-line -> web authentication if this proves necessary

UKAMF

- UK Access Management Federation
  - This is the existing UK research web-based IAM infrastructure
  - https://www.jisc.ac.uk/uk-federation
  - Mature service (over 20 years) Operated by JISC
  - Part of the international eduGAIN inter-federation
  - Highly valuable, especially for providing Single-Sign-On authentication
    - Over 95% of UK higher education institutions are part of UKAMF
    - Over 6000 international institutions are part of eduGAIN
  - Uses SAML
  - Authentication and attribute generation done by HEIs
    - Tied to employment and identities change when user changes job.
    - Typically, simple access rules e.g. **[staff|student]@*institution***

Intermediate IAM services

- Inevitably eduGAIN/UKAMF does not fully meet the needs of all research communities
  - Many communities run their own IAM services, layered underneath eduGAIN and providing additional functionality.
  - Typically implemented as a single “proxy” IAM service
  - The AARC project has developed a blueprint architecture https://aarc-community.org/architecture/ capturing recommended best-practice.
  - Some communities operate across multiple countries and multiple compute federations.
  - Use of OpenID Connect (OIDC) more common than SAML
- Examples
  - IRIS-IAM, DiRAC, Life science AAI, MyAccessID

Existing Service providers

- Existing compute service providers typically implement their own IAM layer
  - Groups/projects defined locally at each provider
    - Potentially annoying for research groups that use multiple facilities
    - More complex and error prone for group admins.
  - Local database of users
    - Requires housekeeping e.g. regular verification of user contact details.
    - Potentially annoying for users who use multiple services.
- Services tied to a particular community may share an intermediate IAM layer

Authentication requirements

- Federated services imply users may need to interact with multiple different systems run by different organisations.
- Users want to avoid overly frequent logins and multiple logins using different credentials/technologies.
  - Single-Sign-On (SSO)
    - SSO provides more than just common credentials. User sessions federate as well.
    - UKAMF/eduGAIN provides federated SSO where user authenticates using the login system of their home institution
      - Familiar, supported through local IT support, good for security incident response
      - No support for users without a home institution in the federation
  - Seamless cryptographic authentication
    - E.g. ssh-keys (ssh-key+MFA usually not seamless).
    - Less useful for non-technical users.
    - Need to collect and verify user ssh-keys

Access management requirements for compute services

- Consumable, exclusive-access resource
  - Resources used by one person/group are not available to anyone else.
  - Typically funded through research grant or equivalent
- Complex access rules
  - E.g. member of a “research project/group” **[manager|member]@group[/subgroup]**
  - Project membership is fluid and *controlled by academics*; may include students, RSEs, collaborators from multiple institutions.
    - Want canonical definition for the federation to make it easier to move groups between services.
    - Ideally be able to import definitions if already defined in community/service-provider IAM
  - Sub-groups may be desirable for large complex research programmes funded by a single grant.
  - Resource allocation and accounting may also reference these same groups
  - Groups may be funded via multiple grants (especially over time) so should be independent entities.

Identity requirement

- Some processes are handled by both people as well as machines
  - Especially in a loose federation
  - E.g. Incident management, user support
- Sometimes we need a unique “name” that is both usable by people and usable by machines
  - People can have multiple email addresses, and these can change over time.
  - UKAMF/eduGain identities are tied to a user's employment
  - ORCiD unique but not very human friendly
  - Community IAM often generates a community id (e.g. “dirac-id”)
- NFCS would also benefit from this.

Offline-access requirement

- For Compute services users are still relevant even when not logged in
  - local storage owned by users/groups.
  - long-running processes/virtual-machines etc.
- Users frequently have personal “accounts” that should be removed/locked if persons authorization status changes.
- Service operators may need to generate mailing list for a group
- Valuable to be able to evaluate up-to-date user attributes independently of user login sessions.

IAM requirements for UK federated compute

- Support federated authentication
  - Leverage UKAMF to allow institutional logins.
  - Add support for non institutional logins
- Shared group forming and management
  - Federate with existing/external community IAM systems to allow groups to be imported. Work needed to identify how best to do this.
  - Interface with allocation and accounting
- Community identities
- Offline access to attributes

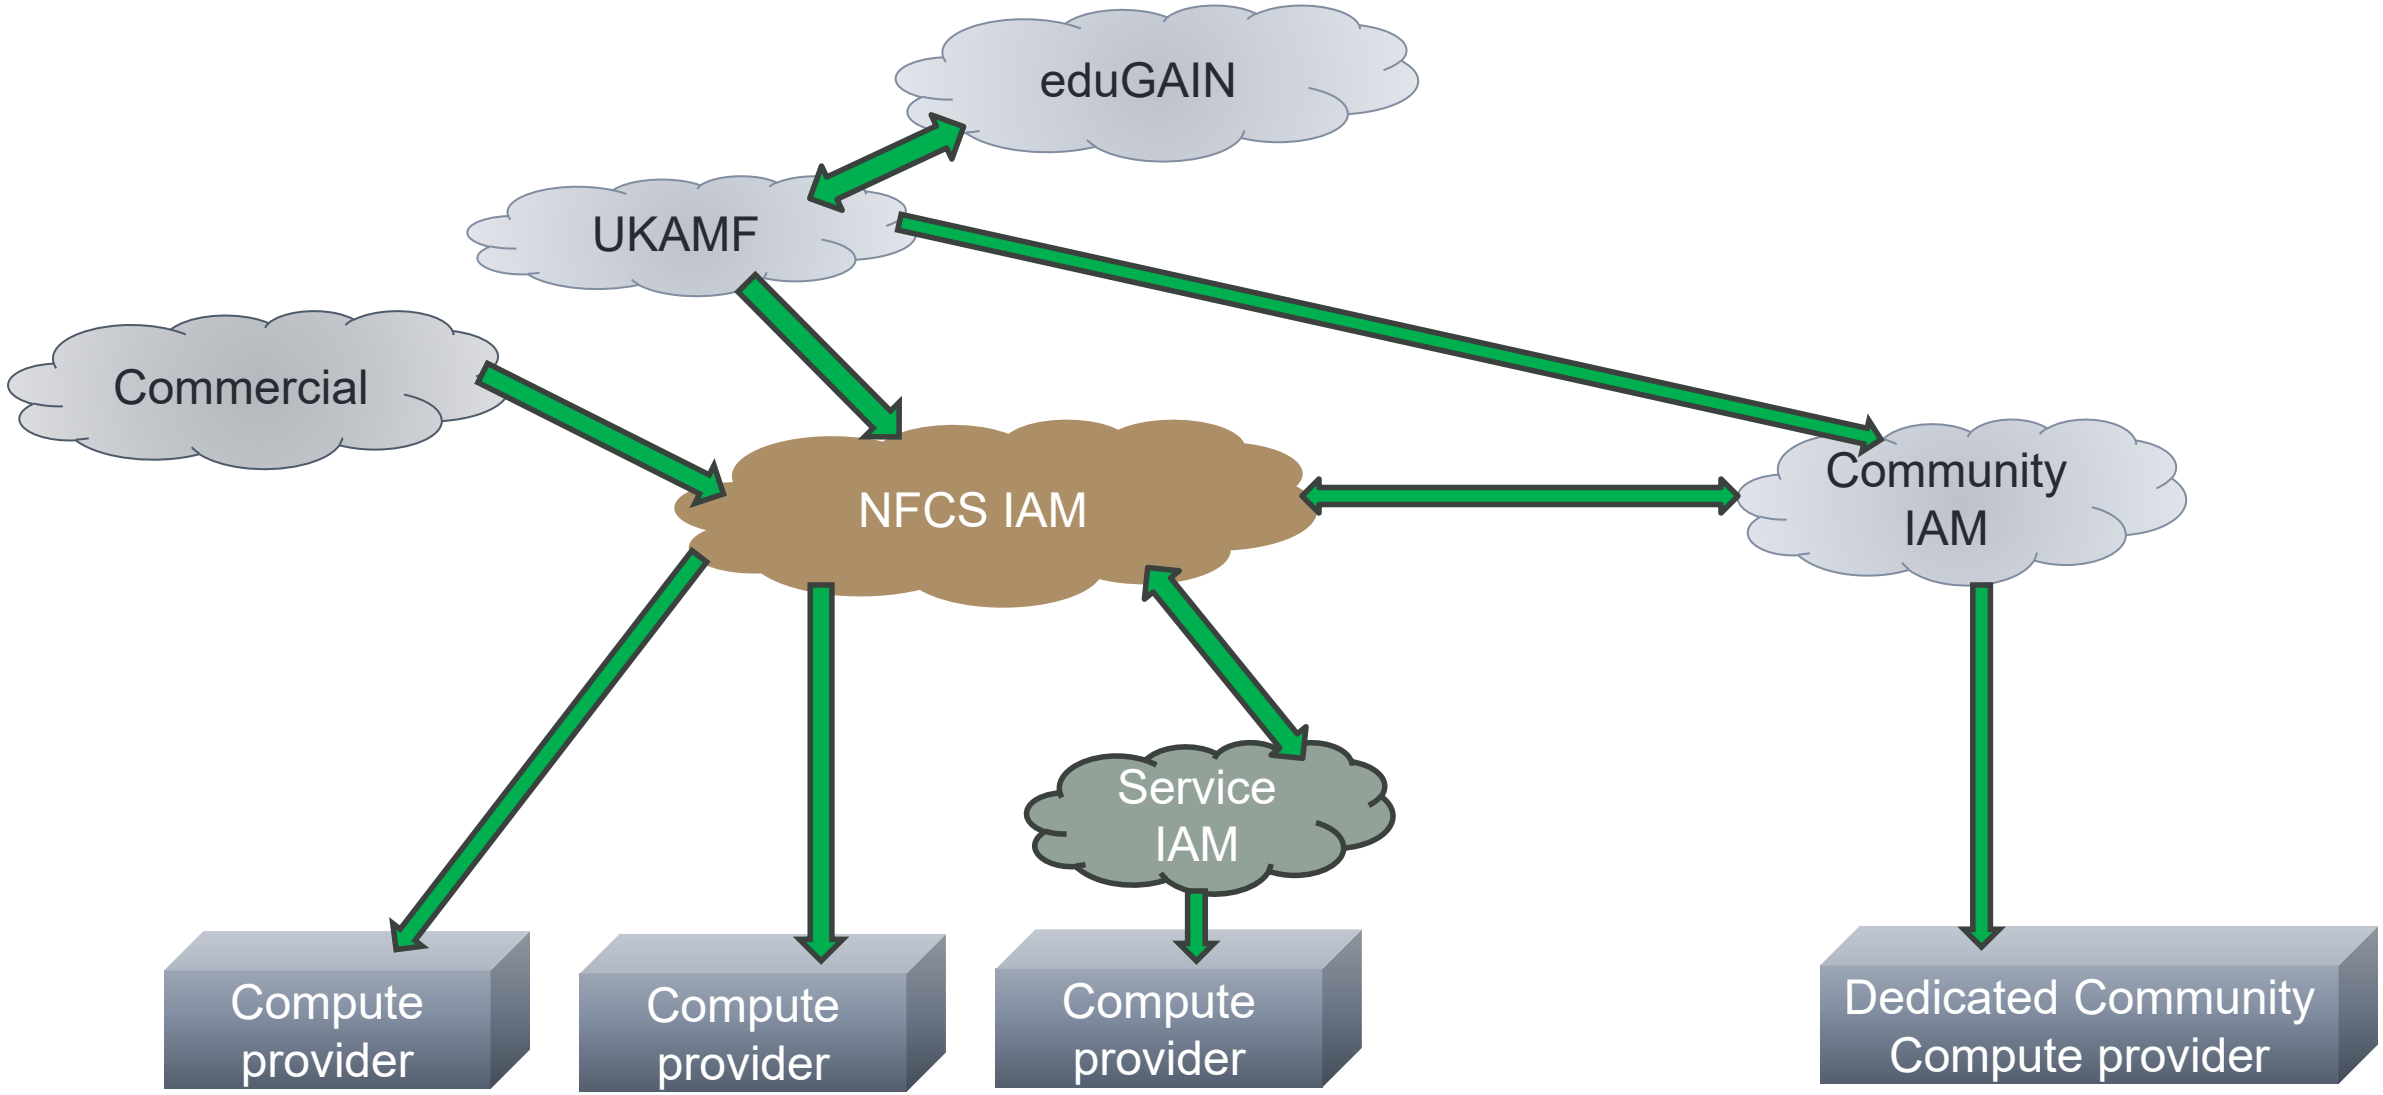

National Federated Compute Services
Connect, coordinate, collaborate

epcc

Strategy

- Requirements suggest NFCS needs its own AARC-compliant proxy
  - Much of the complexity is handled by the proxy
    - Compute services consume identities and groups from a single source
    - Unfortunately MyAccessID does not support group forming so dedicated NFCS service would be better
    - Existing implementations with group forming are available e.g. IRIS-IAM
  - Federation between proxies still being defined
    - Ongoing work for team operating proxy rather than entire federation
- Can be flexible about how compute service providers use group info
- Can leaverge groups/identities to build allocation and accounting systems.

Allocation and accounting