Federated AAI Adoption & Experience from MWA

Benjamin Oshrin

13 Nov 2019 • Utrecht

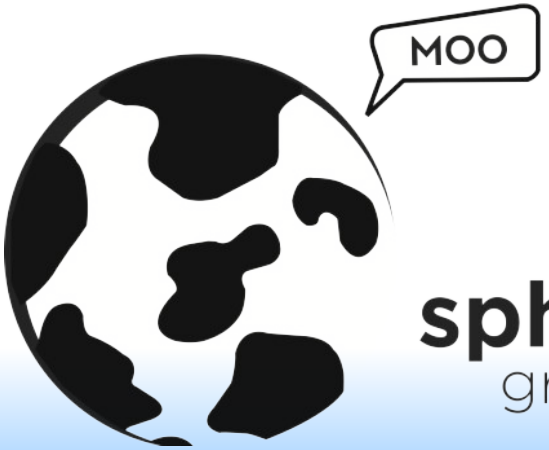

spherical cow group

MWA
MURCHISON WIDEFIELD ARRAY

Murchison Widefield Array (MWA)

- Located at the Murchison Radio-astronomy Observatory (MRO) on Wajarri Yamatji land
- Low frequency radio telescope (80 to 300 MHz)
- SKA precursor located at future site of the SKA low
- International collaboration led by Curtin University
- Science Operations began mid-2013
- 28 PB(!) of data archived at Pawsey Centre
- Growth: 3-8 PB / year depending on schedule & mode

Gurlgamarnu

MWA
MURCHISON
WIDEFIELD
ARRAY

Murchison Radio-Astronomy Observatory

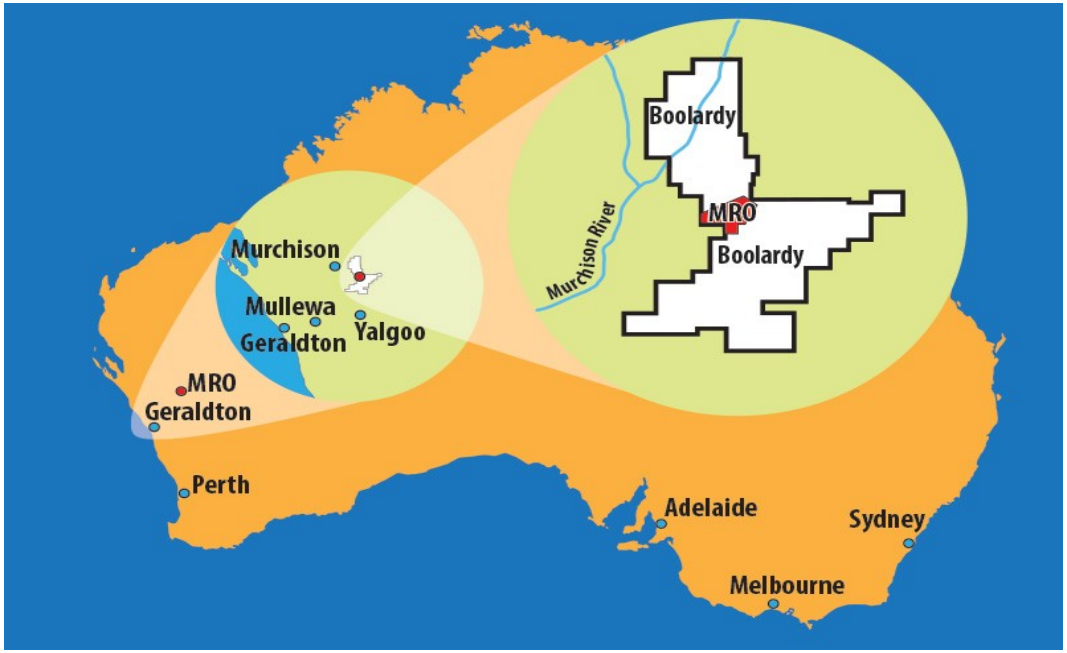

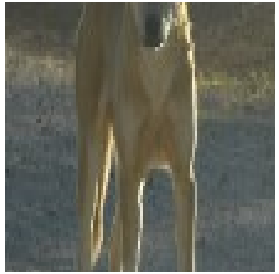

MWA
MURCHISON
WIDEFIELD
ARRAY

The Telescope

- 4096 dual-polarization dipole antennas
- 4 by 4 arrays or "tiles"
- 256 tiles (128 at a time)
- Dense core and 6 km longest baselines
- Several hundred square degrees at a resolution of several arcminutes

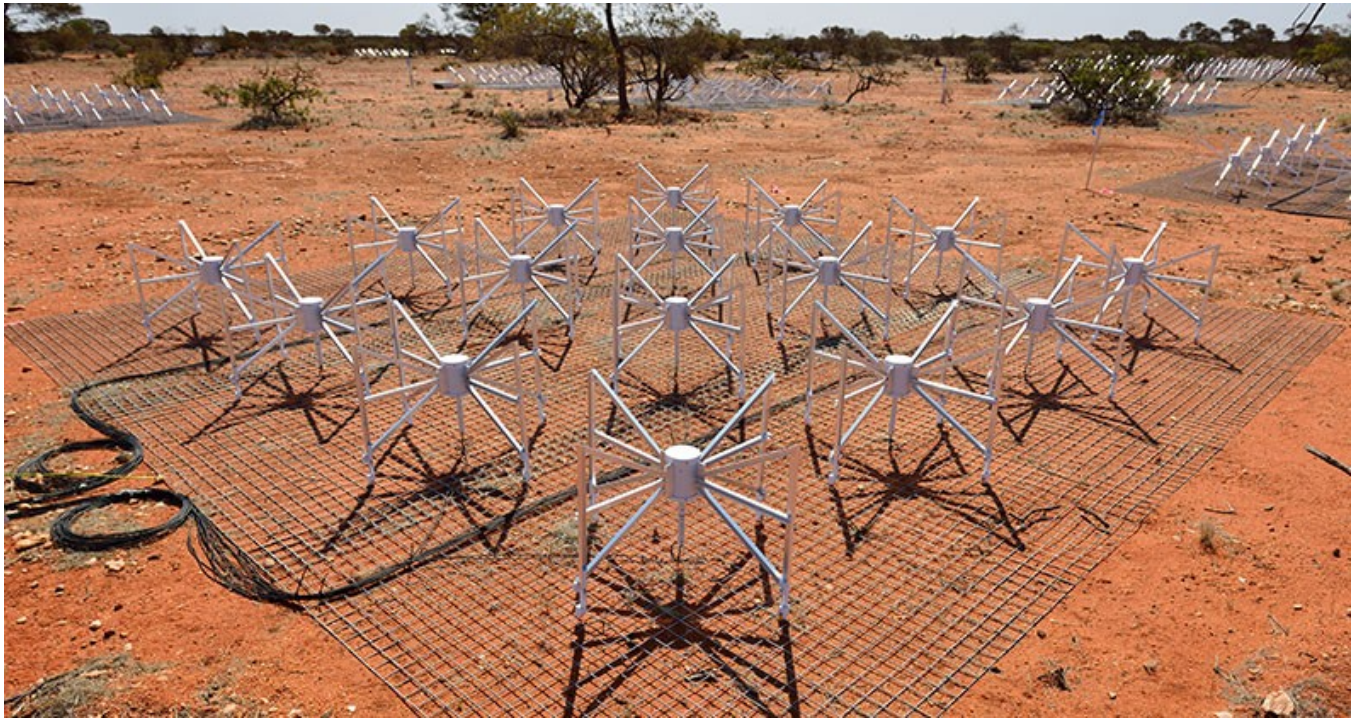

MWA
MURCHISON
WIDEFIELD
ARRAY

SCIENCE!

- Detection of light from the first objects in the universe (EoR)
- Galactic and extra-galactic surveys
- Fast signals e.g. FRBs, pulsars
- Solar / heliospheric studies
- Studying the atmosphere / ionosphere
- SETI
- Space junk tracking
- Testbed for SKA-low technologies

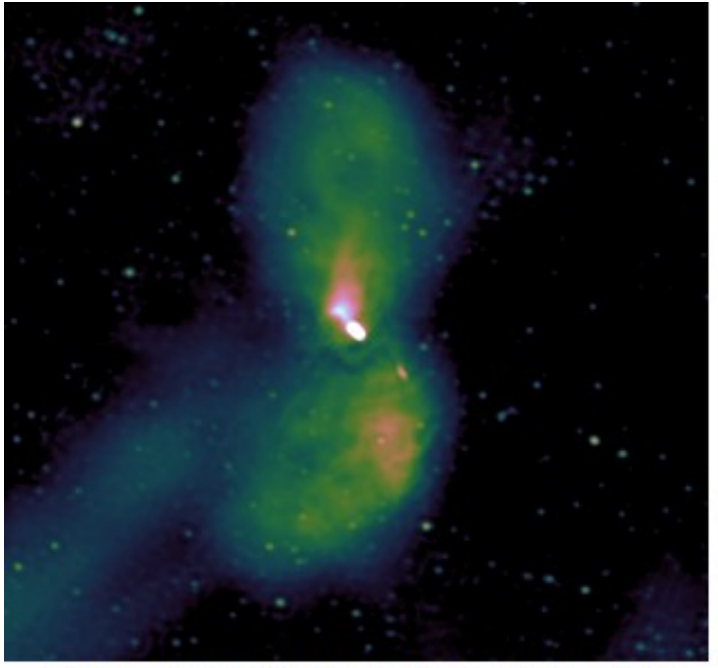

False colour image showing the nearby radio galaxy Cen A.
Credit: Randal Wayth and the MWA team.

MWA
MURCHISON
WIDEFIELD
ARRAY

An International Collaboration!

270+ members from:

- Australia
- Canada
- China
- India
- Japan
- Netherlands
- New Zealand
- South Africa
- Sweden
- UK
- USA

MWA Collaboration Services

- Wiki
  - TWiki
  - Owned by MIT, which left the Consortium
  - Contains a decade of semi-organized knowledge
  - Flat authorization model: a few admins, but otherwise no access management per domain

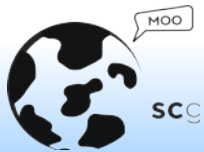

MWA Collaboration Services

- Mailing Lists
  - Mailman 2
  - Owned by MIT, which left the Consortium
  - Subscription management done manually by admin
  - Users want to send mail from multiple addresses
  - Users move institutions and subscribe with new address, but don’t/can’t unsubscribe old addresses

MWA Collaboration Services

- Source Code Repository
  - Gitolite
  - Hosted by Curtin
  - Authentication via SSH keys
    - Managed manually
- Data Access Portal
  - Domain specific application, implementing data access policies

MWA Growing Pains

- Management of membership as new institutions join the Collaboration, others leave
  - On/off-boarding, difficult, time-consuming, error-prone
  - No dedicated admin - Principal Scientist managing membership
- No public data access & limited controls on embargoed (non-public) data
- Aging / unsupported software
- Technical debt accumulated over a decade

MWA Federated Identity Management Approach

- SSO using user's own institution's credentials
  - Deploy "federation-aware" applications to replace aging systems
- Centralize user management
  - Single on/off-boarding processes (user driven)
  - Reduce admin burden on Principal Scientist, allow delegation
- Integrate with new MWA Data Portal: The MWA All Sky Virtual Observatory (ASVO)
- Solution: A federated approach with Australian Access Federation (AAF) & eduGAIN
  - Leverage Open Source software to provide core IdM services

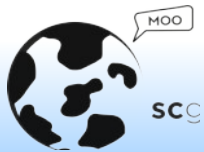

Federated Identity Management

Federated Identity
+
Identity Management

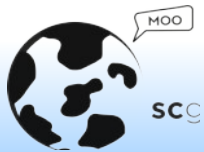

Federated Identity

- *Bring Your Own Identity*
- Institutional Credentials
- Increasingly, commercial / social credentials
  - (Google, Facebook, Weibo, etc)
- Federated Authentication with maybe some attributes

Evolution of Federated Identity

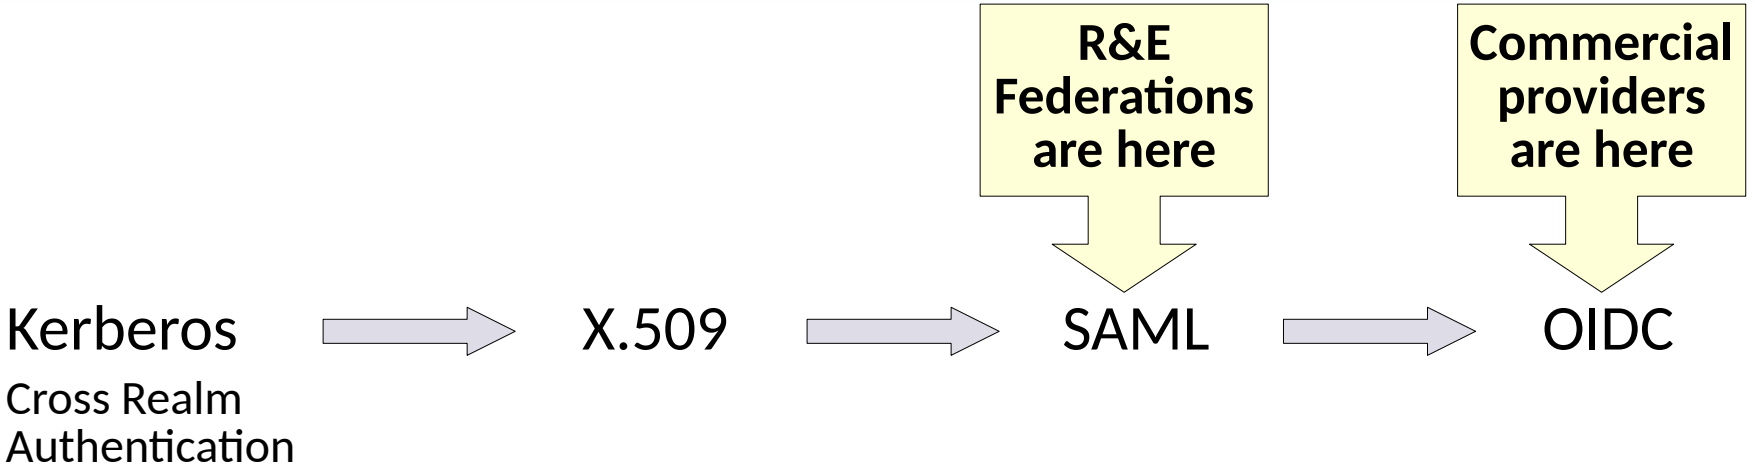

Federated Identity Complexities

- Identity Provider (IdP) Participation
  - IdP of Last Resort (United ID, Social Providers, AAF Virtual Home, DIY)
- IdP Attribute Release
- Level of Assurance / MFA
- User Experience / IdP Discovery
- Service Provider (SP) Integration
- Security Incident Reporting and Handling
- User Lifecycle Management
  - More of a local problem than a federation problem

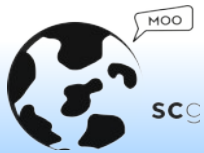

Federated Identity Management Community Work

- eduGAIN
- REFEDs Entity Categories
  - Research & Scholarship (“R&S”)
- REFEDs Security Incident Response Trust Framework for Federated Identity (“SIRTFI”)
  - Enables coordinated security incident response across federated organizations
  - Requirements
    - Self-assessment against SIRTFI criteria
    - Assert a security contact in metadata

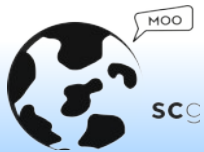

Federated Identity Management Community Work

- Federation Aware Software Development
  - Shibboleth (Shibboleth Consortium)
    - https://shibboleth.net
  - SimpleSAMLphp (UNINETT.no)
    - https://simplesamlphp.org
  - idpy: SATOSA, pyFF (Commons Conservancy)
    - https://idpy.org
  - COmanage (Internet2)
    - https://www.internet2.edu/comanage
  - Grouper (Internet2)
    - https://www.internet2.edu/grouper

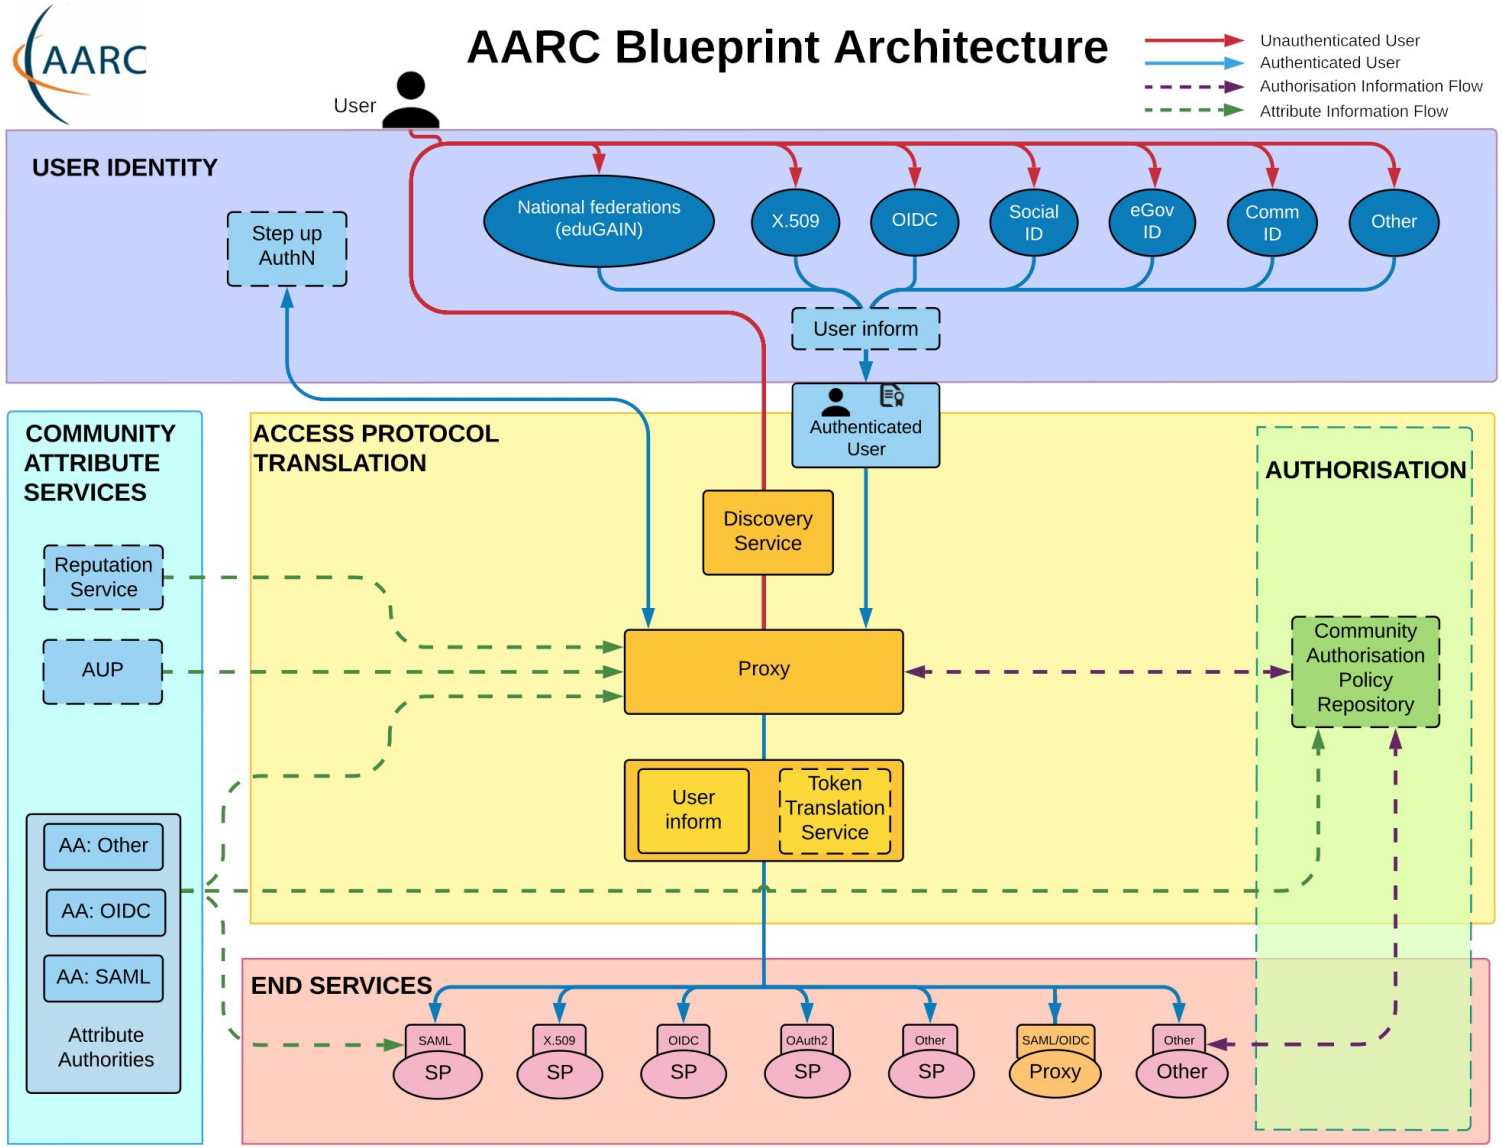
AARC Blueprint Architecture
Unauthenticated User
Authenticated User
Authorisation Information Flow
Attribute Information Flow
User
USER IDENTITY
Step up AuthN
National federations (eduGAIN)
X.509
OIDC
Social ID
eGov ID
Comm ID
Other
User inform
Authenticated User
COMMUNITY ATTRIBUTE SERVICES
ACCESS PROTOCOL TRANSLATION
AUTHORISATION
Discovery Service
Reputation Service
AUP
Proxy
Community Authorisation Policy Repository
User inform
Token Translation Service
AA: Other
AA: OIDC
AA: SAML
Attribute Authorities
END SERVICES
SAML SP
X.509 SP
OIDC SP
OAuth2 SP
Other SP
SAML/OIDC Proxy
Other Other

(Federated) Identity Management

- Enrollment and Lifecycle Management
  - Identifier assignment and management
  - Privilege/entitlement management
- Application Integration and Provisioning
- Automation
- Delegation / Self Service

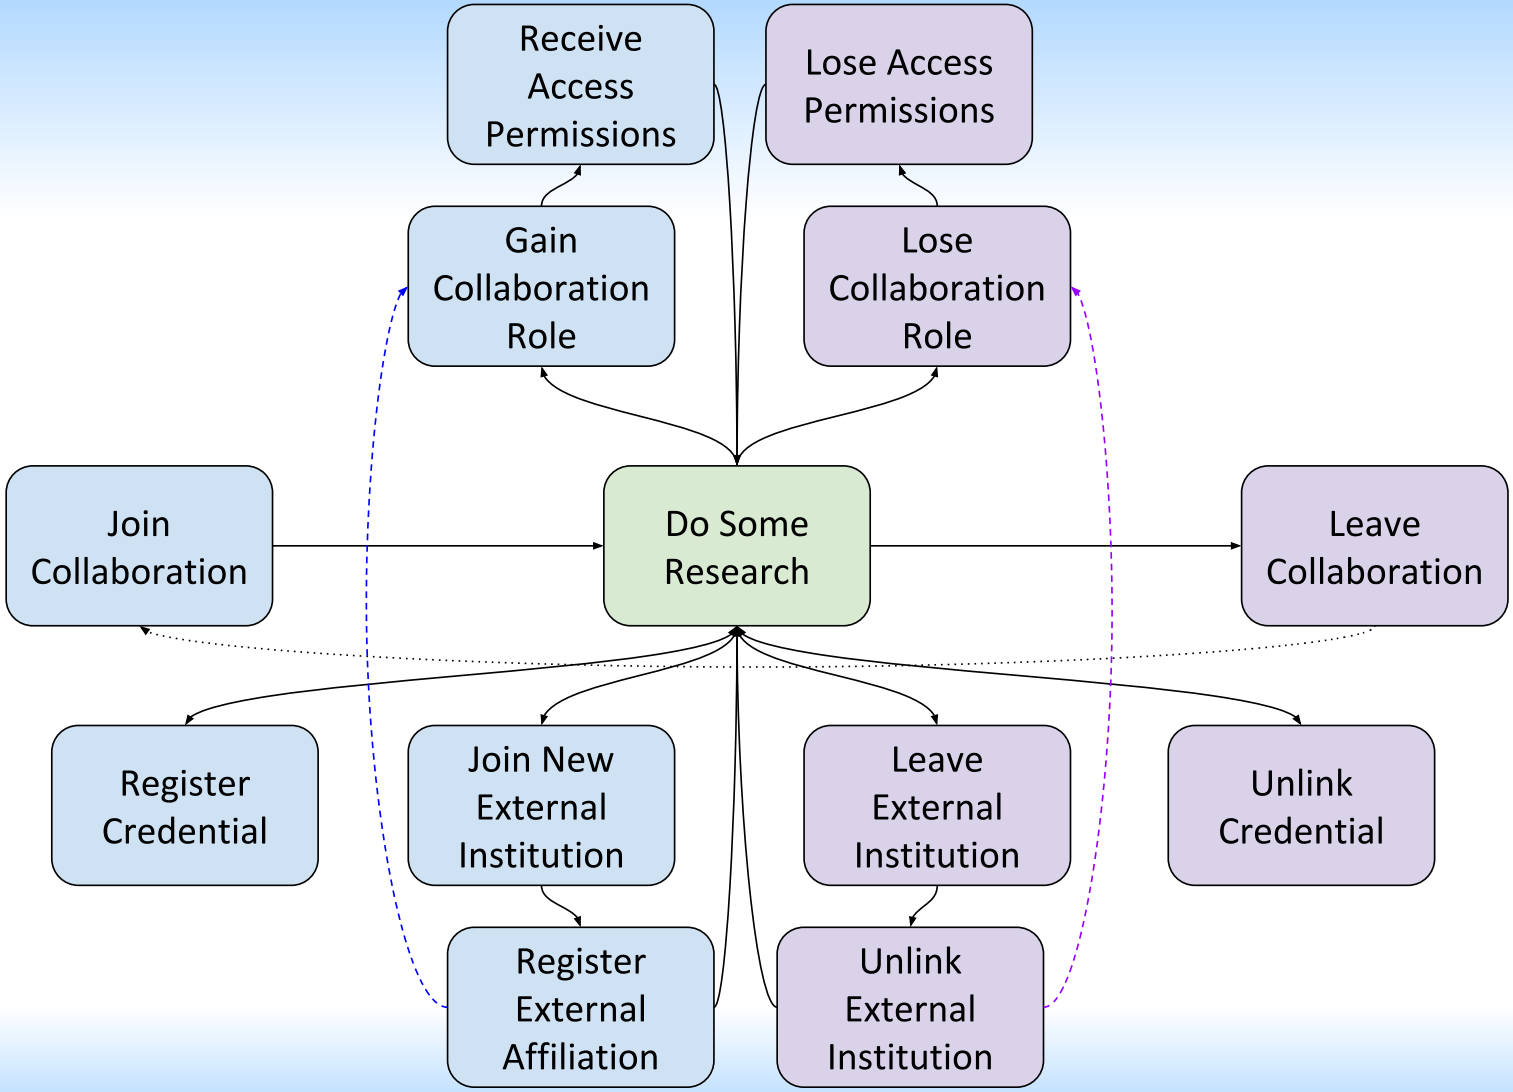
Receive Access Permissions
Lose Access Permissions
Gain Collaboration Role
Lose Collaboration Role
Join Collaboration
Do Some Research
Leave Collaboration
Register Credential
Join New External Institution
Leave External Institution
Unlink Credential
Register External Affiliation
Unlink External Institution

Federated Identity Management

Why Bother?

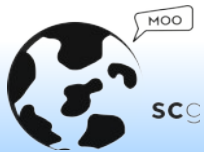

Why Bother?

- These problems need to be solved anyway, by investing a bit more up front in community-endorsed solutions and best practices, long term maintenance costs are lowered
- Anecdotal evidence suggests researchers prefer to use their institutional identity, when available
- A well designed Federated Identity Management approach speeds up onboarding/enrollment, and makes it more accurate
  - And the same for offboarding/expiration

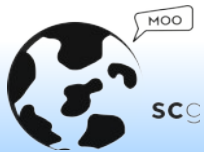

MWA's Federated Identity Management Approach

- Developed more or less concurrently with the AARC BPA
- COmanage Registry for user lifecycle (attribute + identifier) management
- SATOSA for IdP/SP SAML Proxy
- PyFF for SAML IdP Discovery
- eduGAIN for federated authn

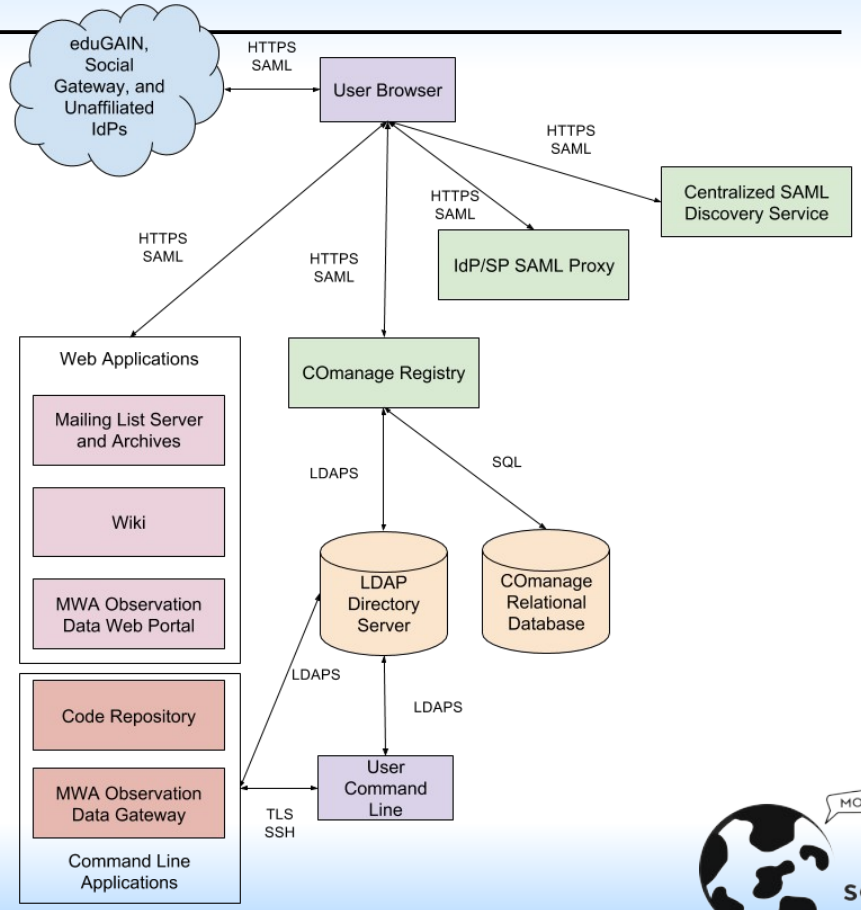

MWA Federated Identity Management

- Rolling out now to generally positive feedback
  - In particular from the administrators
- Proxy-centric architecture rapidly becoming the standard for VO Identity Management Architecture

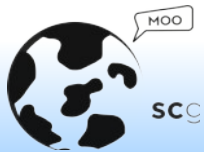

(fin)