HELMHOLTZ

RESEARCH FOR
GRAND CHALLENGES

HIFIS | HELMHOLTZ
FEDERATED
IT SERVICES

Helmholtz-AAI

Sander Apweiler (FZJ), Marcus Hardt (KIT), Uwe Jandt (DESY), Andreas Klotz (HZB)

October 2021

Outline

- Motivation
- Architecture
- Implementation
- Developments

Motivation

+

Overview

Historical records

- Helmholtz Data Federation (HDF) needed an AAI (back in 2017)
- Proof of Concept implementation of the AARC Blueprint Architecture
  - SP-IdP Proxy (in eduGain)
  - 4 Initial services (Nagios, OpenStack, dCache, WaTTS, ...)
  - OpenID Connect as a primary target
- Adaptation of the AARC Policy Development Kit
  - Security + Trust
  - Policy Compatibility with large infrastructures (WLCG, LIGO, XSEDE, ELIXIR, ...)

Then HIFIS started

I. Helmholtz AAI - Motivation

Slide courtesy of Uwe Jandt

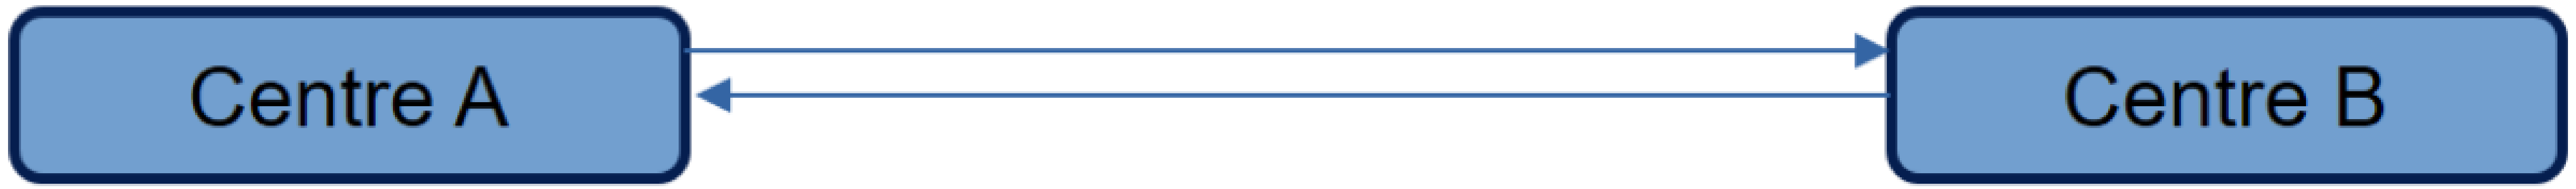

Conventional approach:
- 2 partners (1 service provider, 1 service user):
- separate contract for each service

- ➢ https://aai.helmholtz.de/howto
- ➢ https://aai.helmholtz.de/concept

I. Helmholtz AAI - Motivation

Slide courtesy of Uwe Jandt

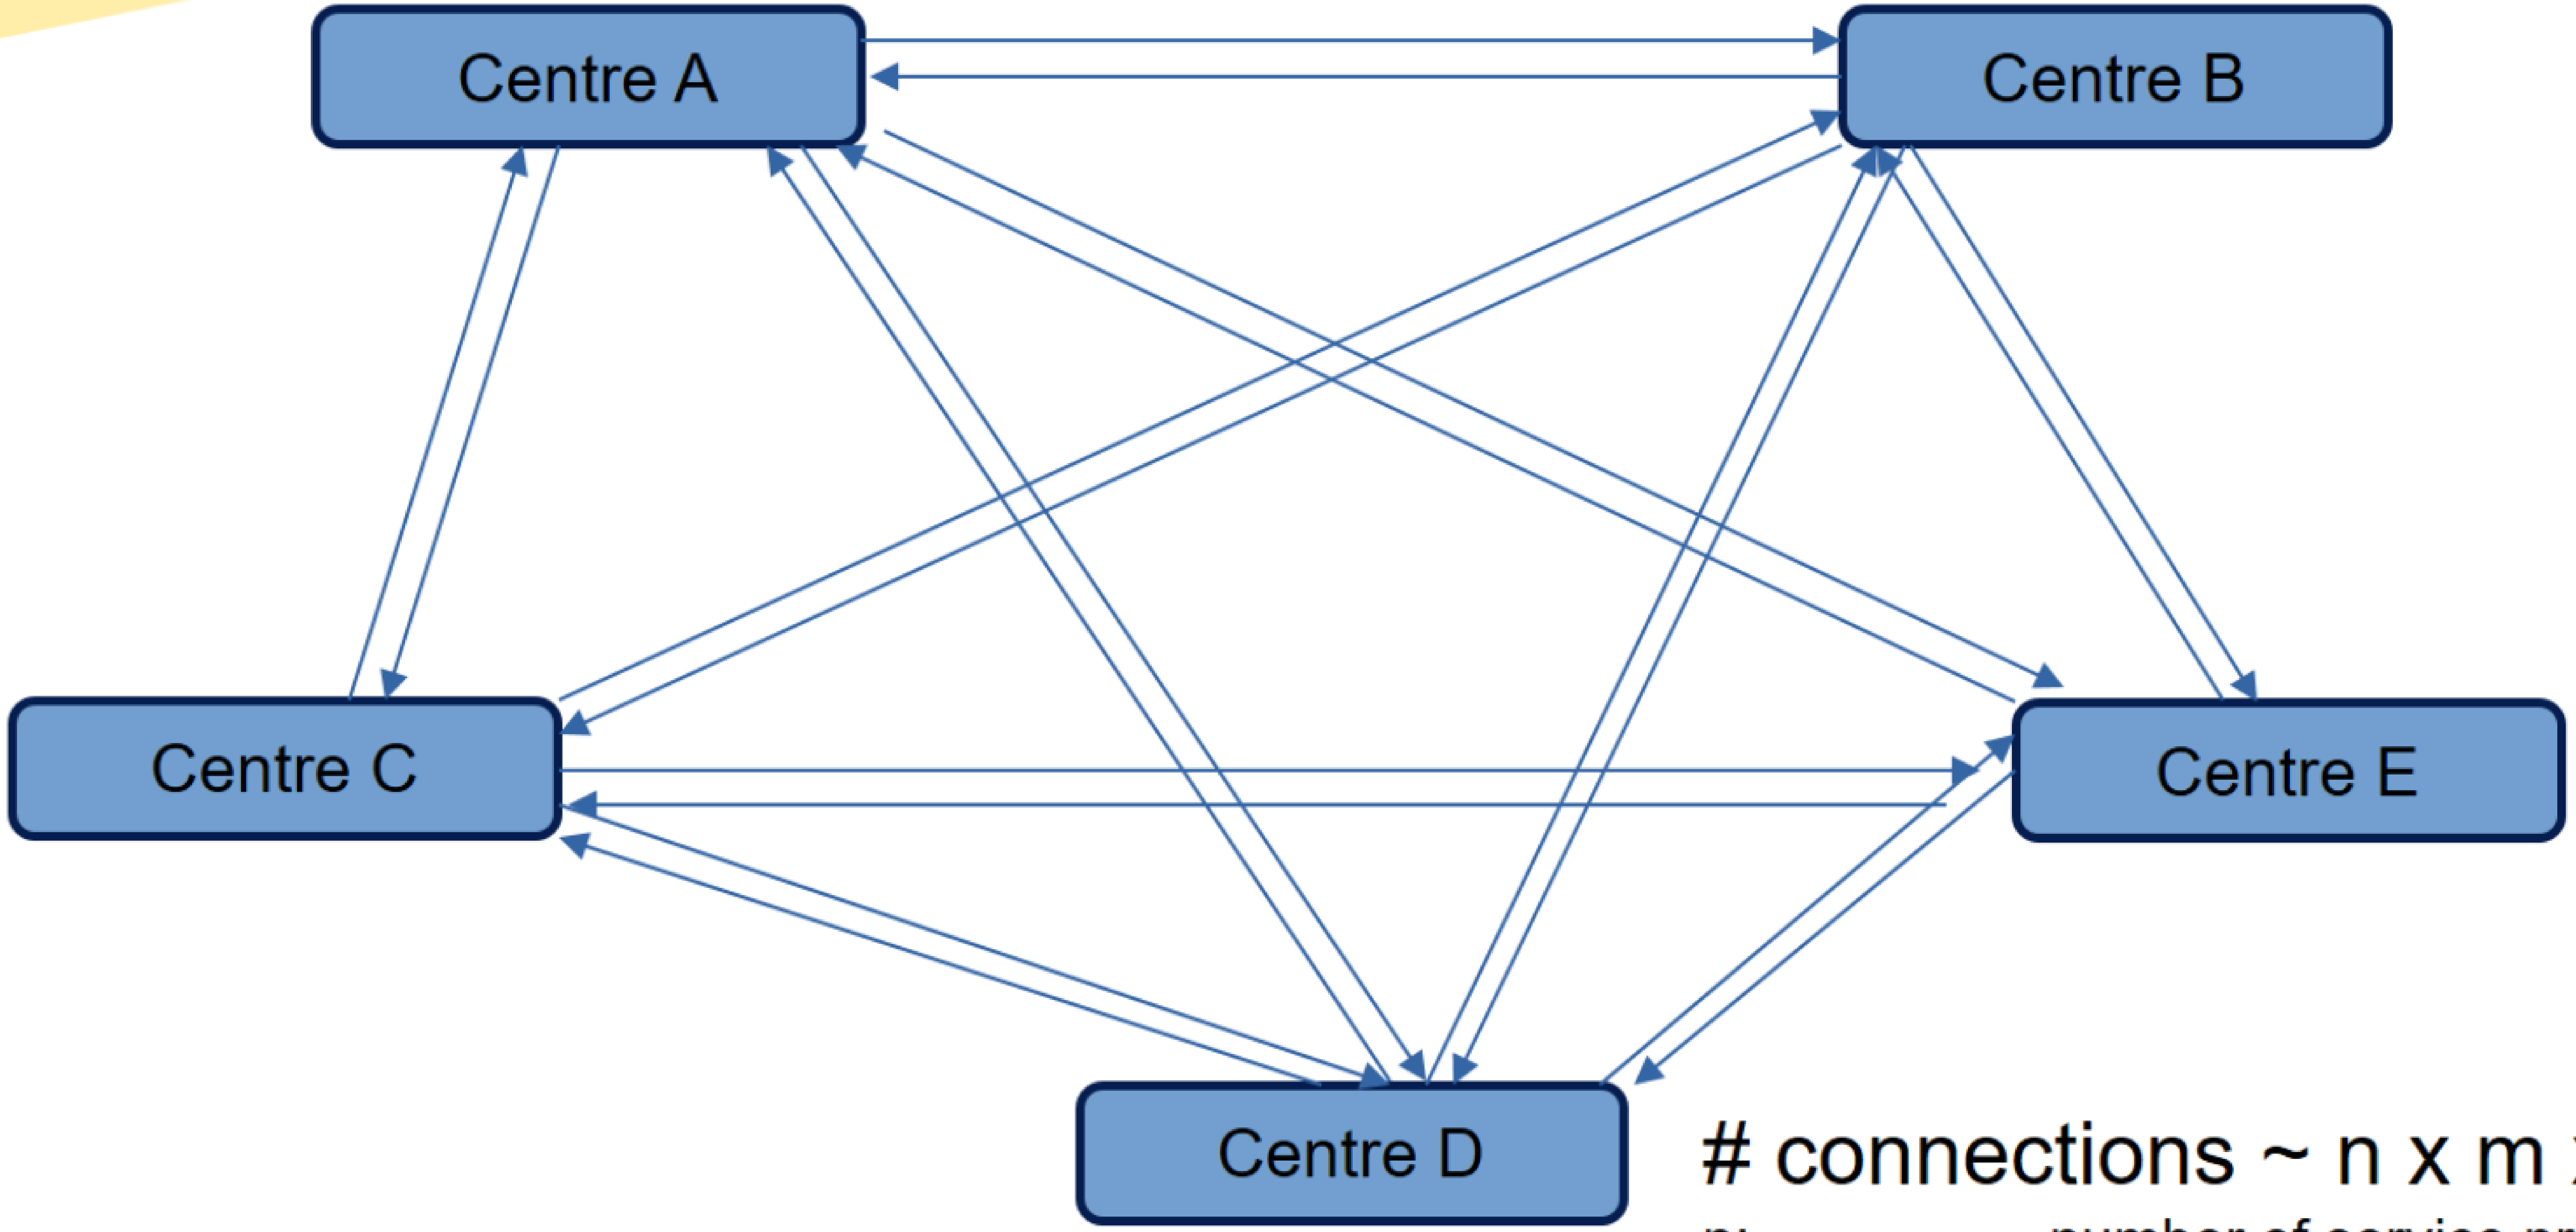

connections ~ n x m x s

n: number of service providing centres (>10)

m: numbers of using centres (19)

s: number of services

- https://aai.helmholtz.de/howto
- https://aai.helmholtz.de/concept

I. Helmholtz AAI - Solution

Slide courtesy of Uwe Jandt

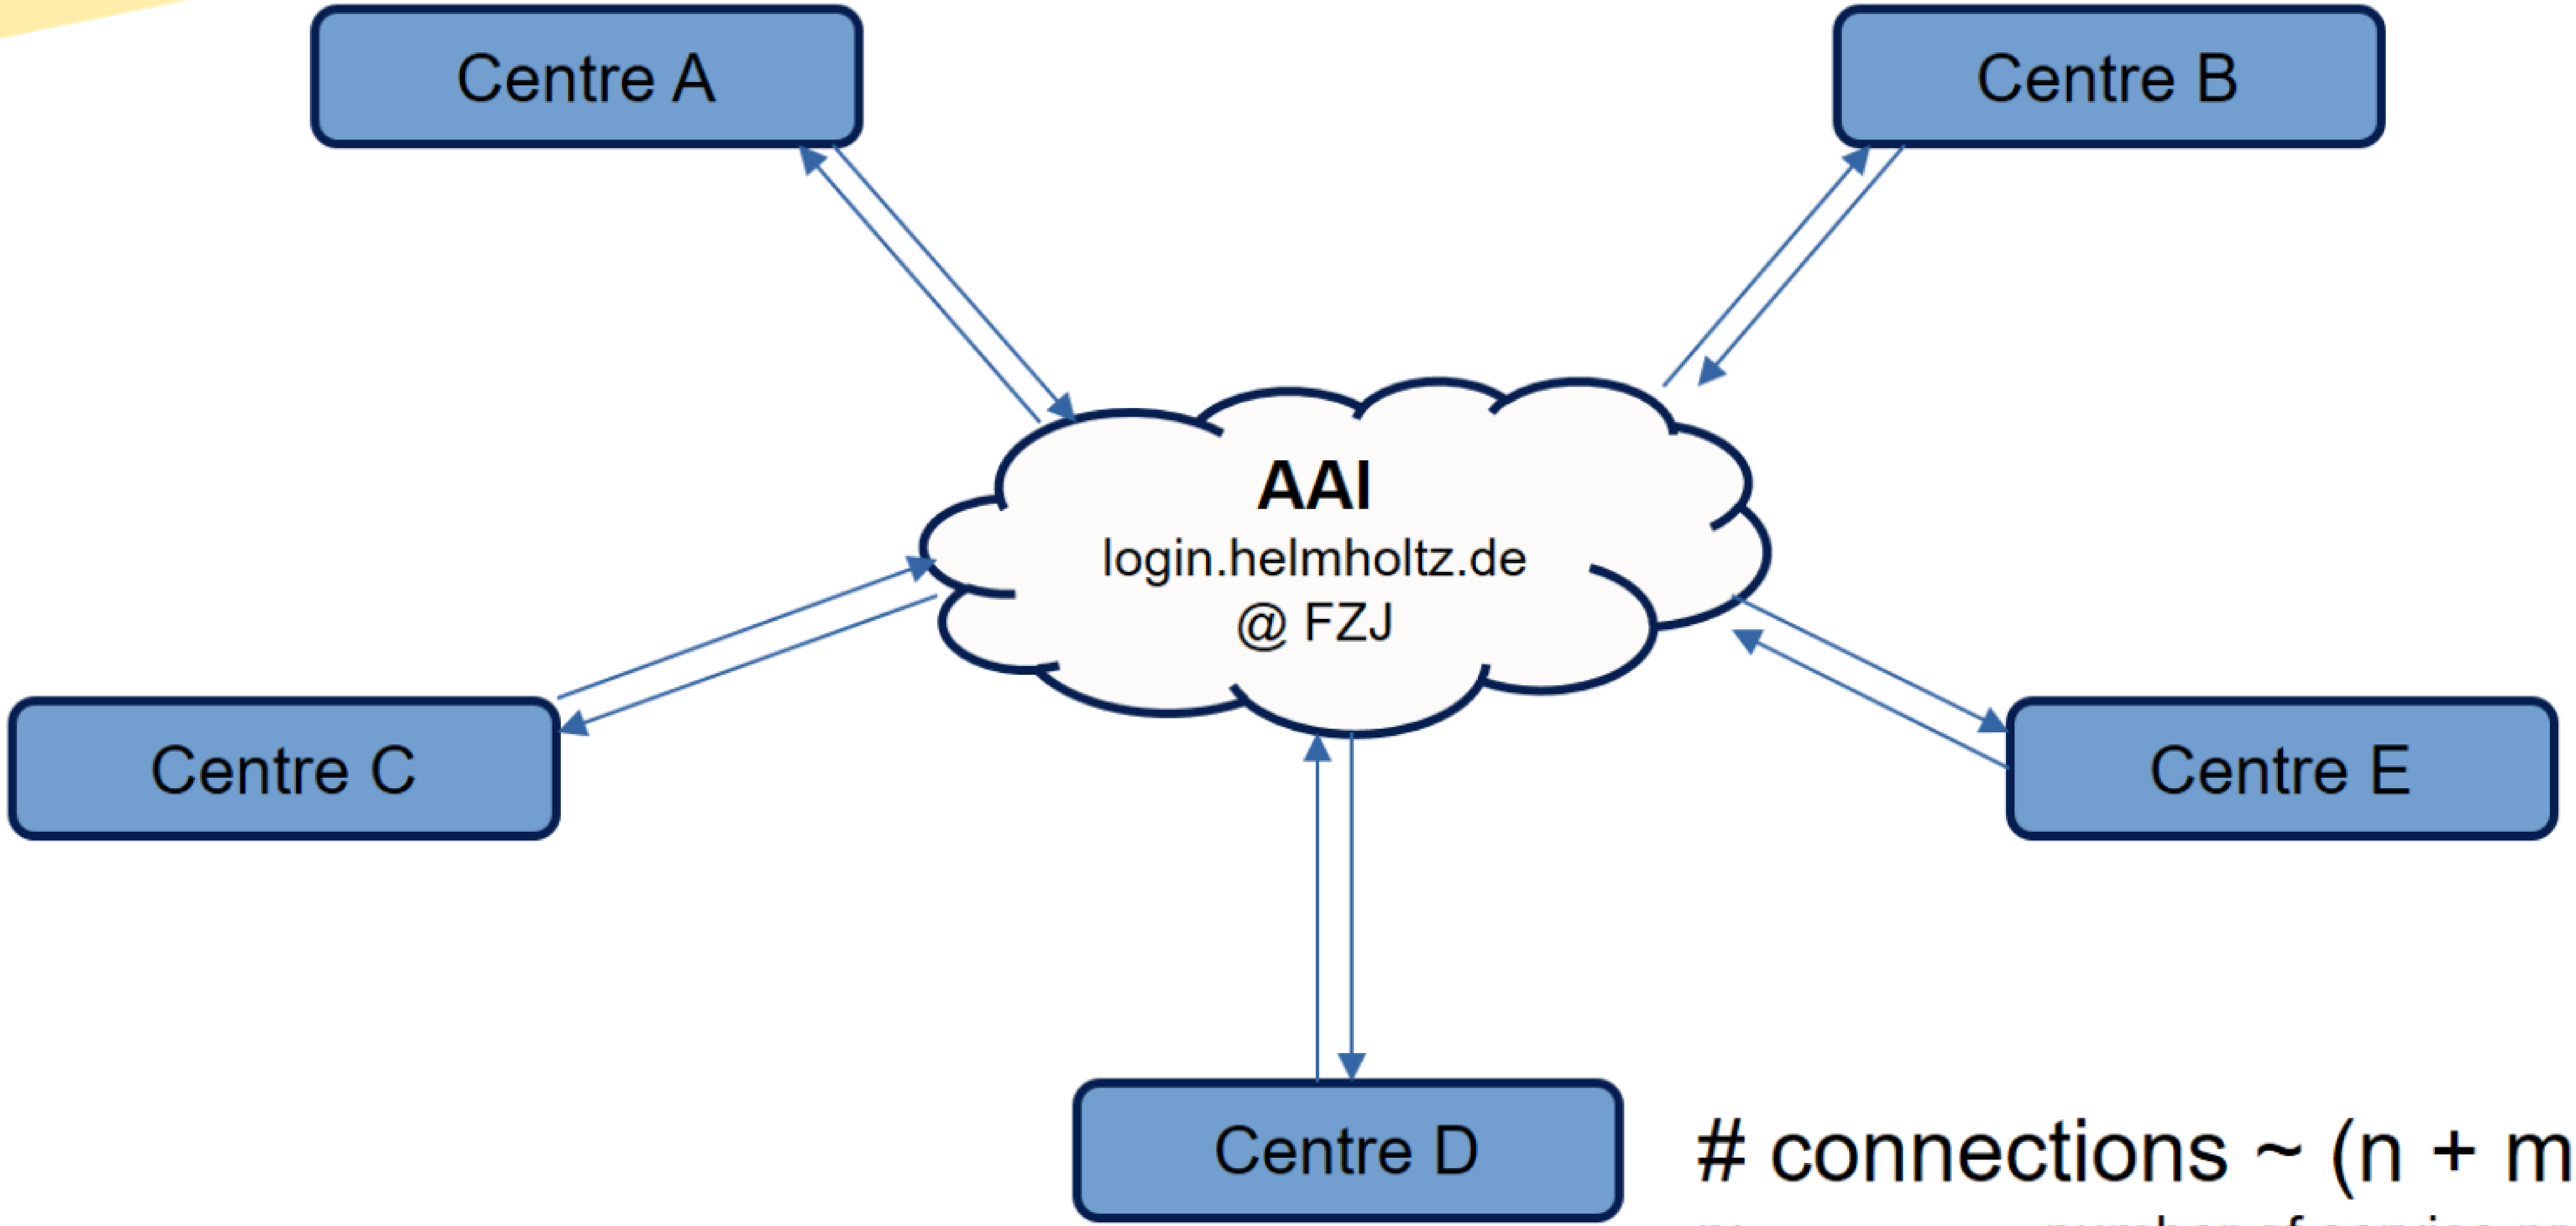

connections ~ (n + m)

n: number of service providing centres (>10)

m: numbers of using centres (19)

s: number of services

- ➢ https://aai.helmholtz.de/howto
- ➢ https://aai.helmholtz.de/concept

Helmholtz AAI: Goals

- Users can access **many** federated services
  - with the **one account** of their Home Organisation
- Seamless access
  - Enable Services for users at Helmholtz Centres and Partners
  - Enable researchers and guests to access Services
  - Very general approach => Don’t be limited by specific organisational structure
- Enable PIs to manage their own Virtual Organisations (VOs)
- Compatibility with the European Open Science Cloud (EOSC)
- Support for services beyond the browser
  - Delegation (Computing Jobs)
  - REST APIs
  - Shell access

Relation to DFN-AAI

- Overlap only in the acronym “AAI”
- Helmholtz AAI is *one* service inside DFN-AAI
  - it’s an SP-IdP-proxy
- Users come in via
  - **DFN-AAI**, eduGAIN,
  - And others: ORCID, Github, Google
  - Even “Homeless Users” **could** be supported

Föderationen: Basisinfrastruktur für BPAs

Slide courtesy of Wolfgang Pempe

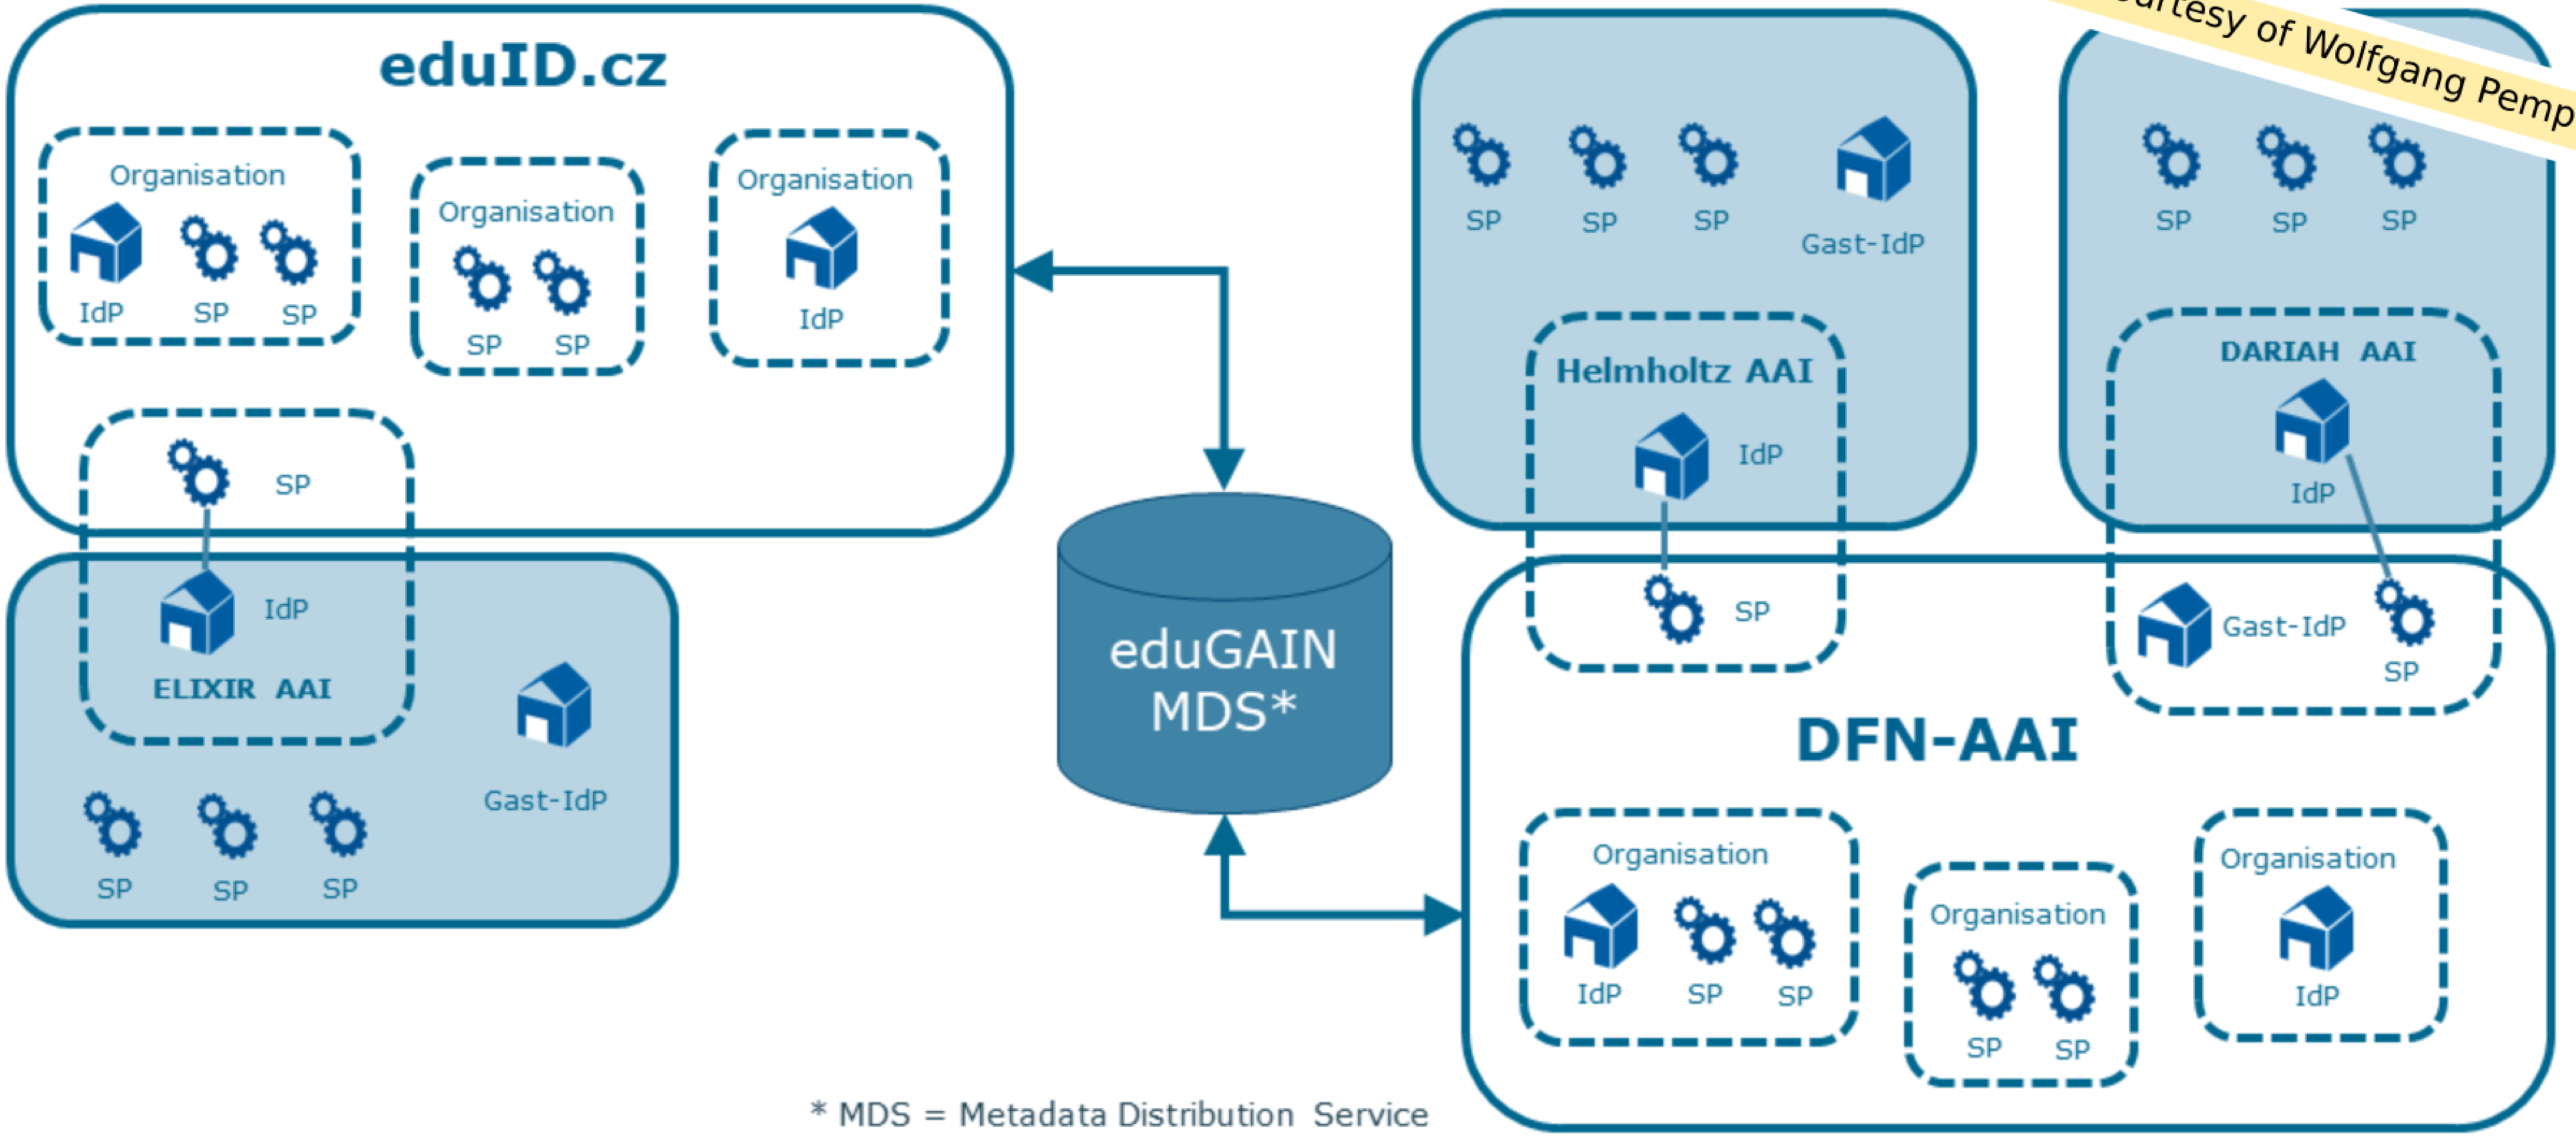

* MDS = Metadata Distribution Service

Architecture

AARC

- Authentication and Authorisation for Research Communities (AARC)
  - 25 Partners
  - 4 Years
- Mission
  - Analyse **existing** Architectures
  - Analyse **existing** Policies
  - => Give recommendations
    - 21 Final, ~15 more on the roadmap

AARC Results

- AARC Policy Development KIT
  - Fundamental policy templates for
    - Operating an infrastructure
    - Handling Incident Response
    - Manage Members
    - Requirements on Authentication
    - Risk Assessment
    - Data Protection
    - Privacy Policy
    - Service Operation
    - Acceptable Use Policy

  All policies designed to be GDPR compliant
- AARC Blueprint Architectures
  - Introduction of the “proxy” component

AARC Blueprint Architecture

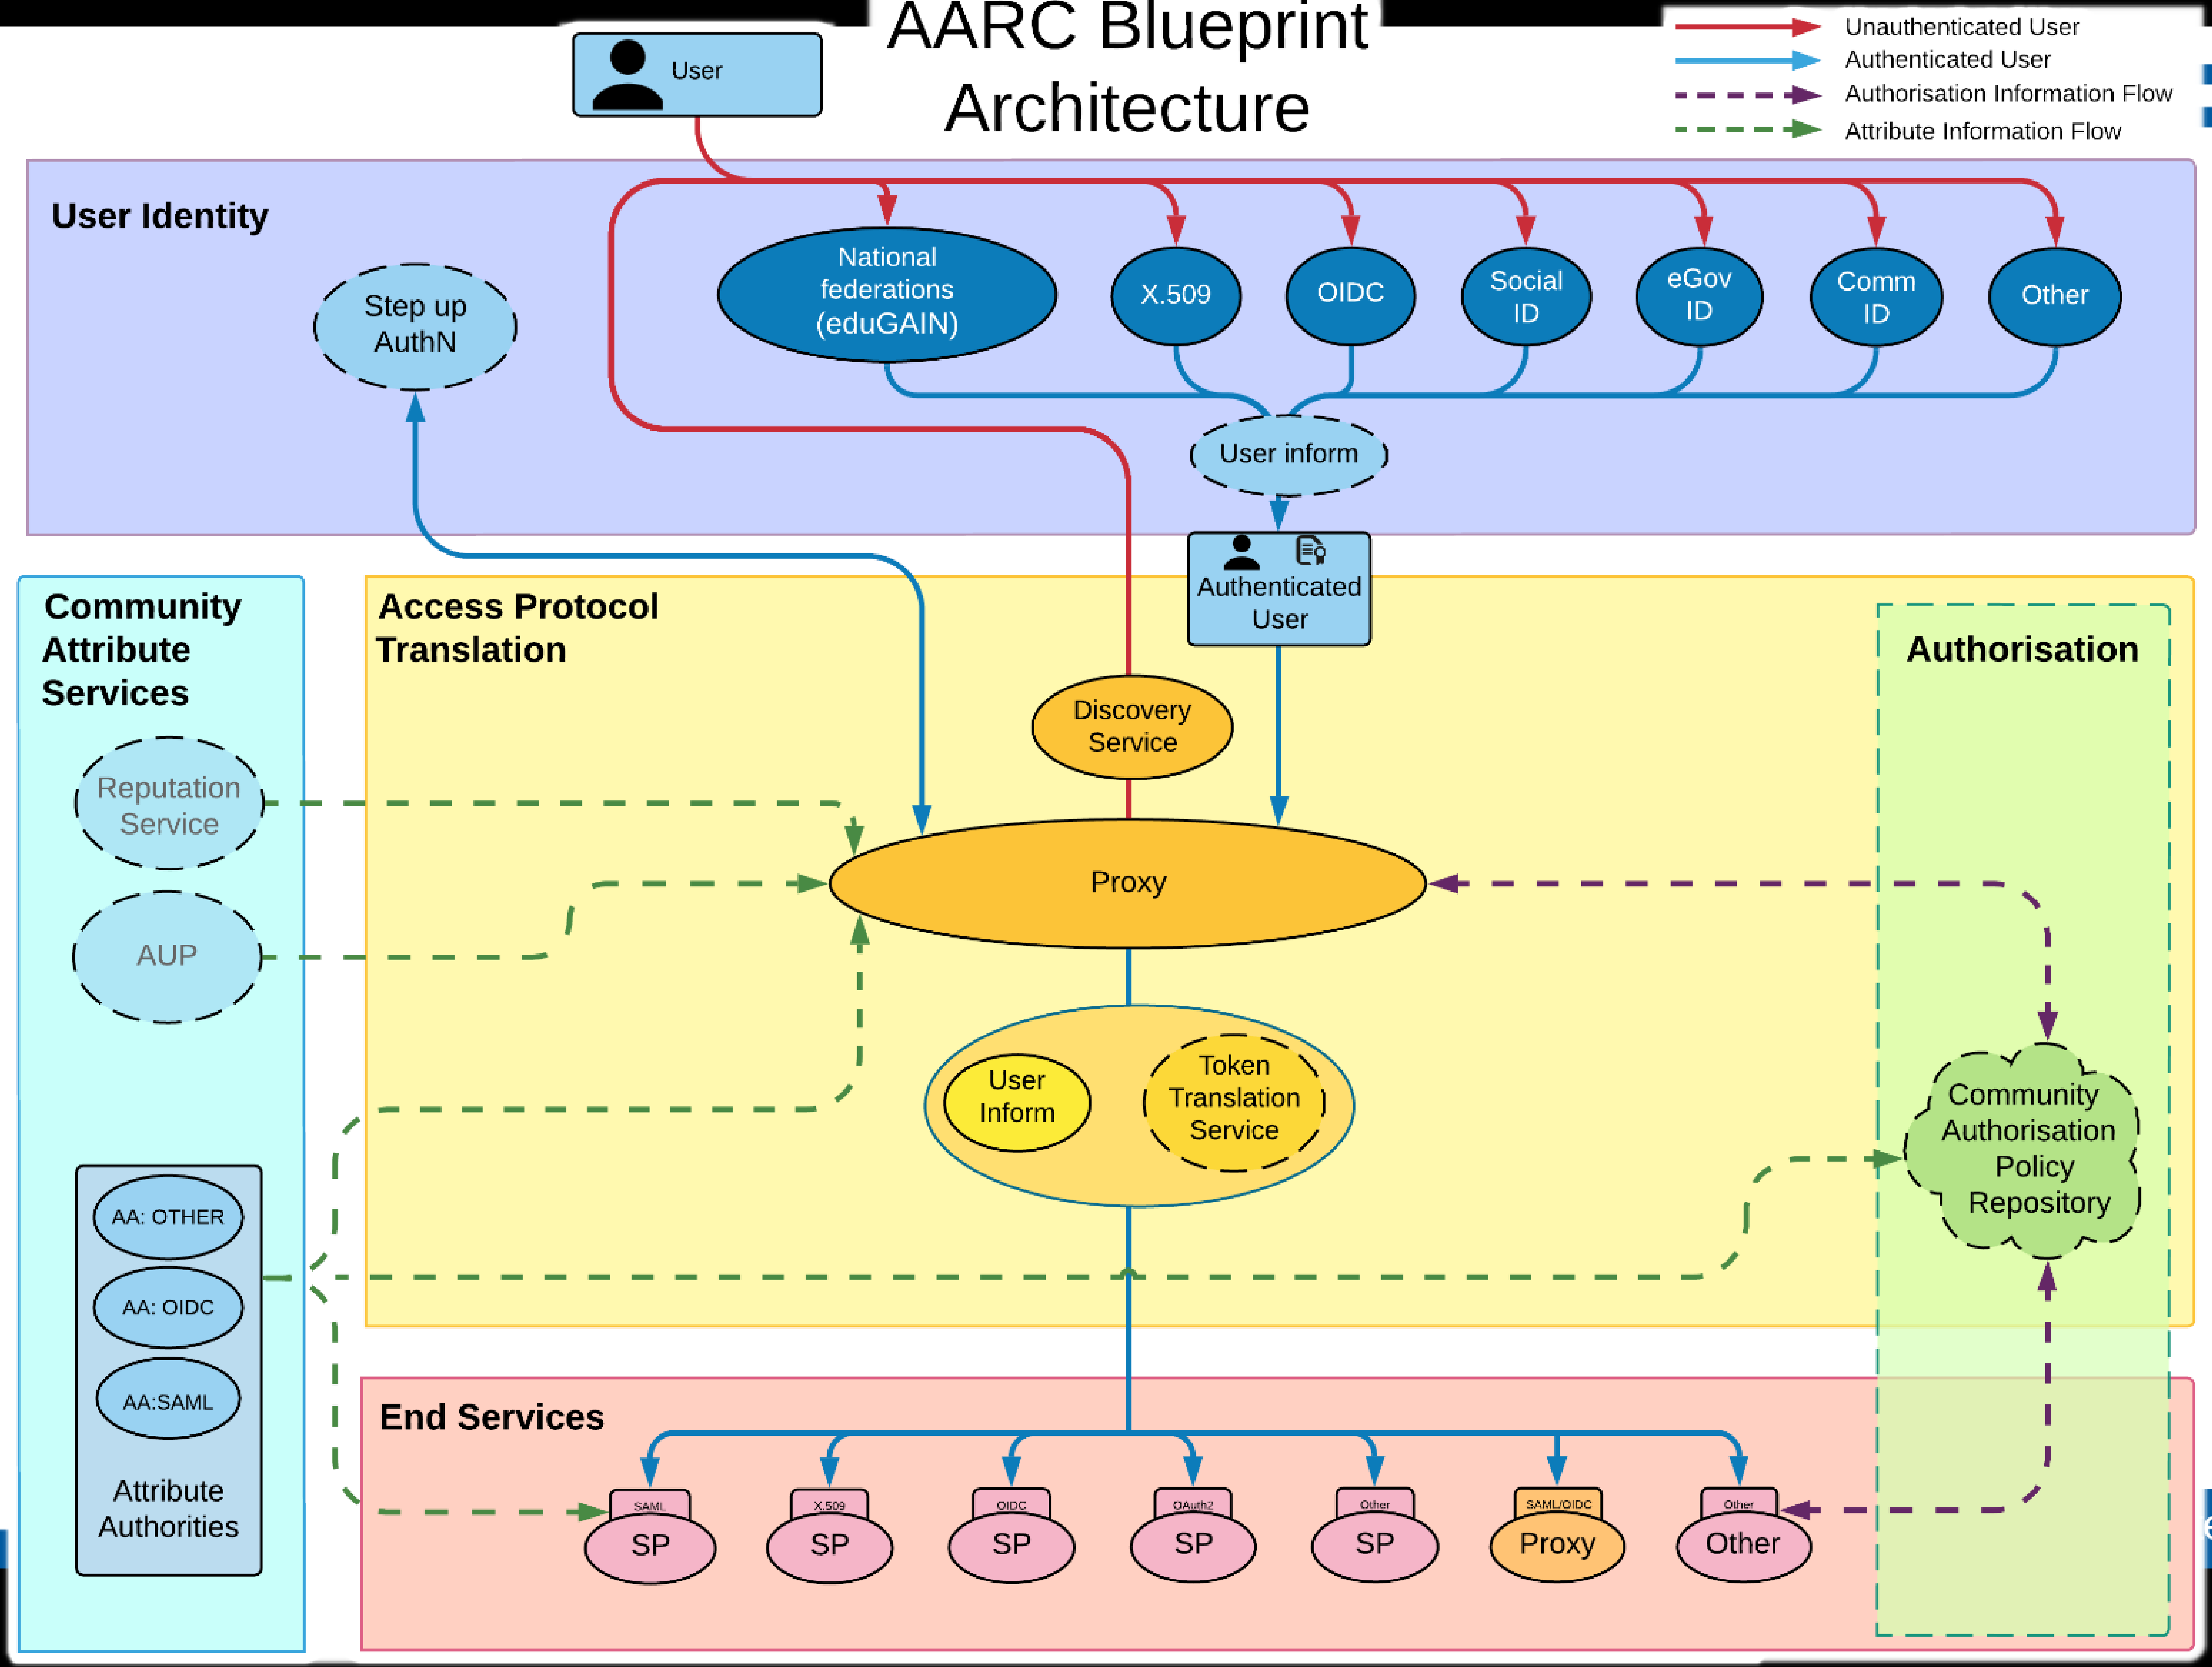

The “proxy” component

aka: “SP-IdP-Proxy”

- Scalability!
  - Stop every IdP needing to talk to every SP (`2800 * 1800`)
  - Reliable Attribute release
    - (Different nations, different IdPs, each with different schema)
- Third party **authorisation**
  - *“Community Attribute Services”*
- Enforcement of authorisation
- Protocol Translation
  - SAML, OIDC, X.509
- Stackable

Evolution of the BPA (~2019)

- Add differentiation:
  - Between **community** and **infrastructure**
  - Between “infrastructure run by a community” and “general e-Infrastructure”
  - Different types of Services
- Introduces the vocabulary used
- Full Document: https://zenodo.org/record/3672785/files/AARC-G045-AARC_BPA_2019-Final.pdf

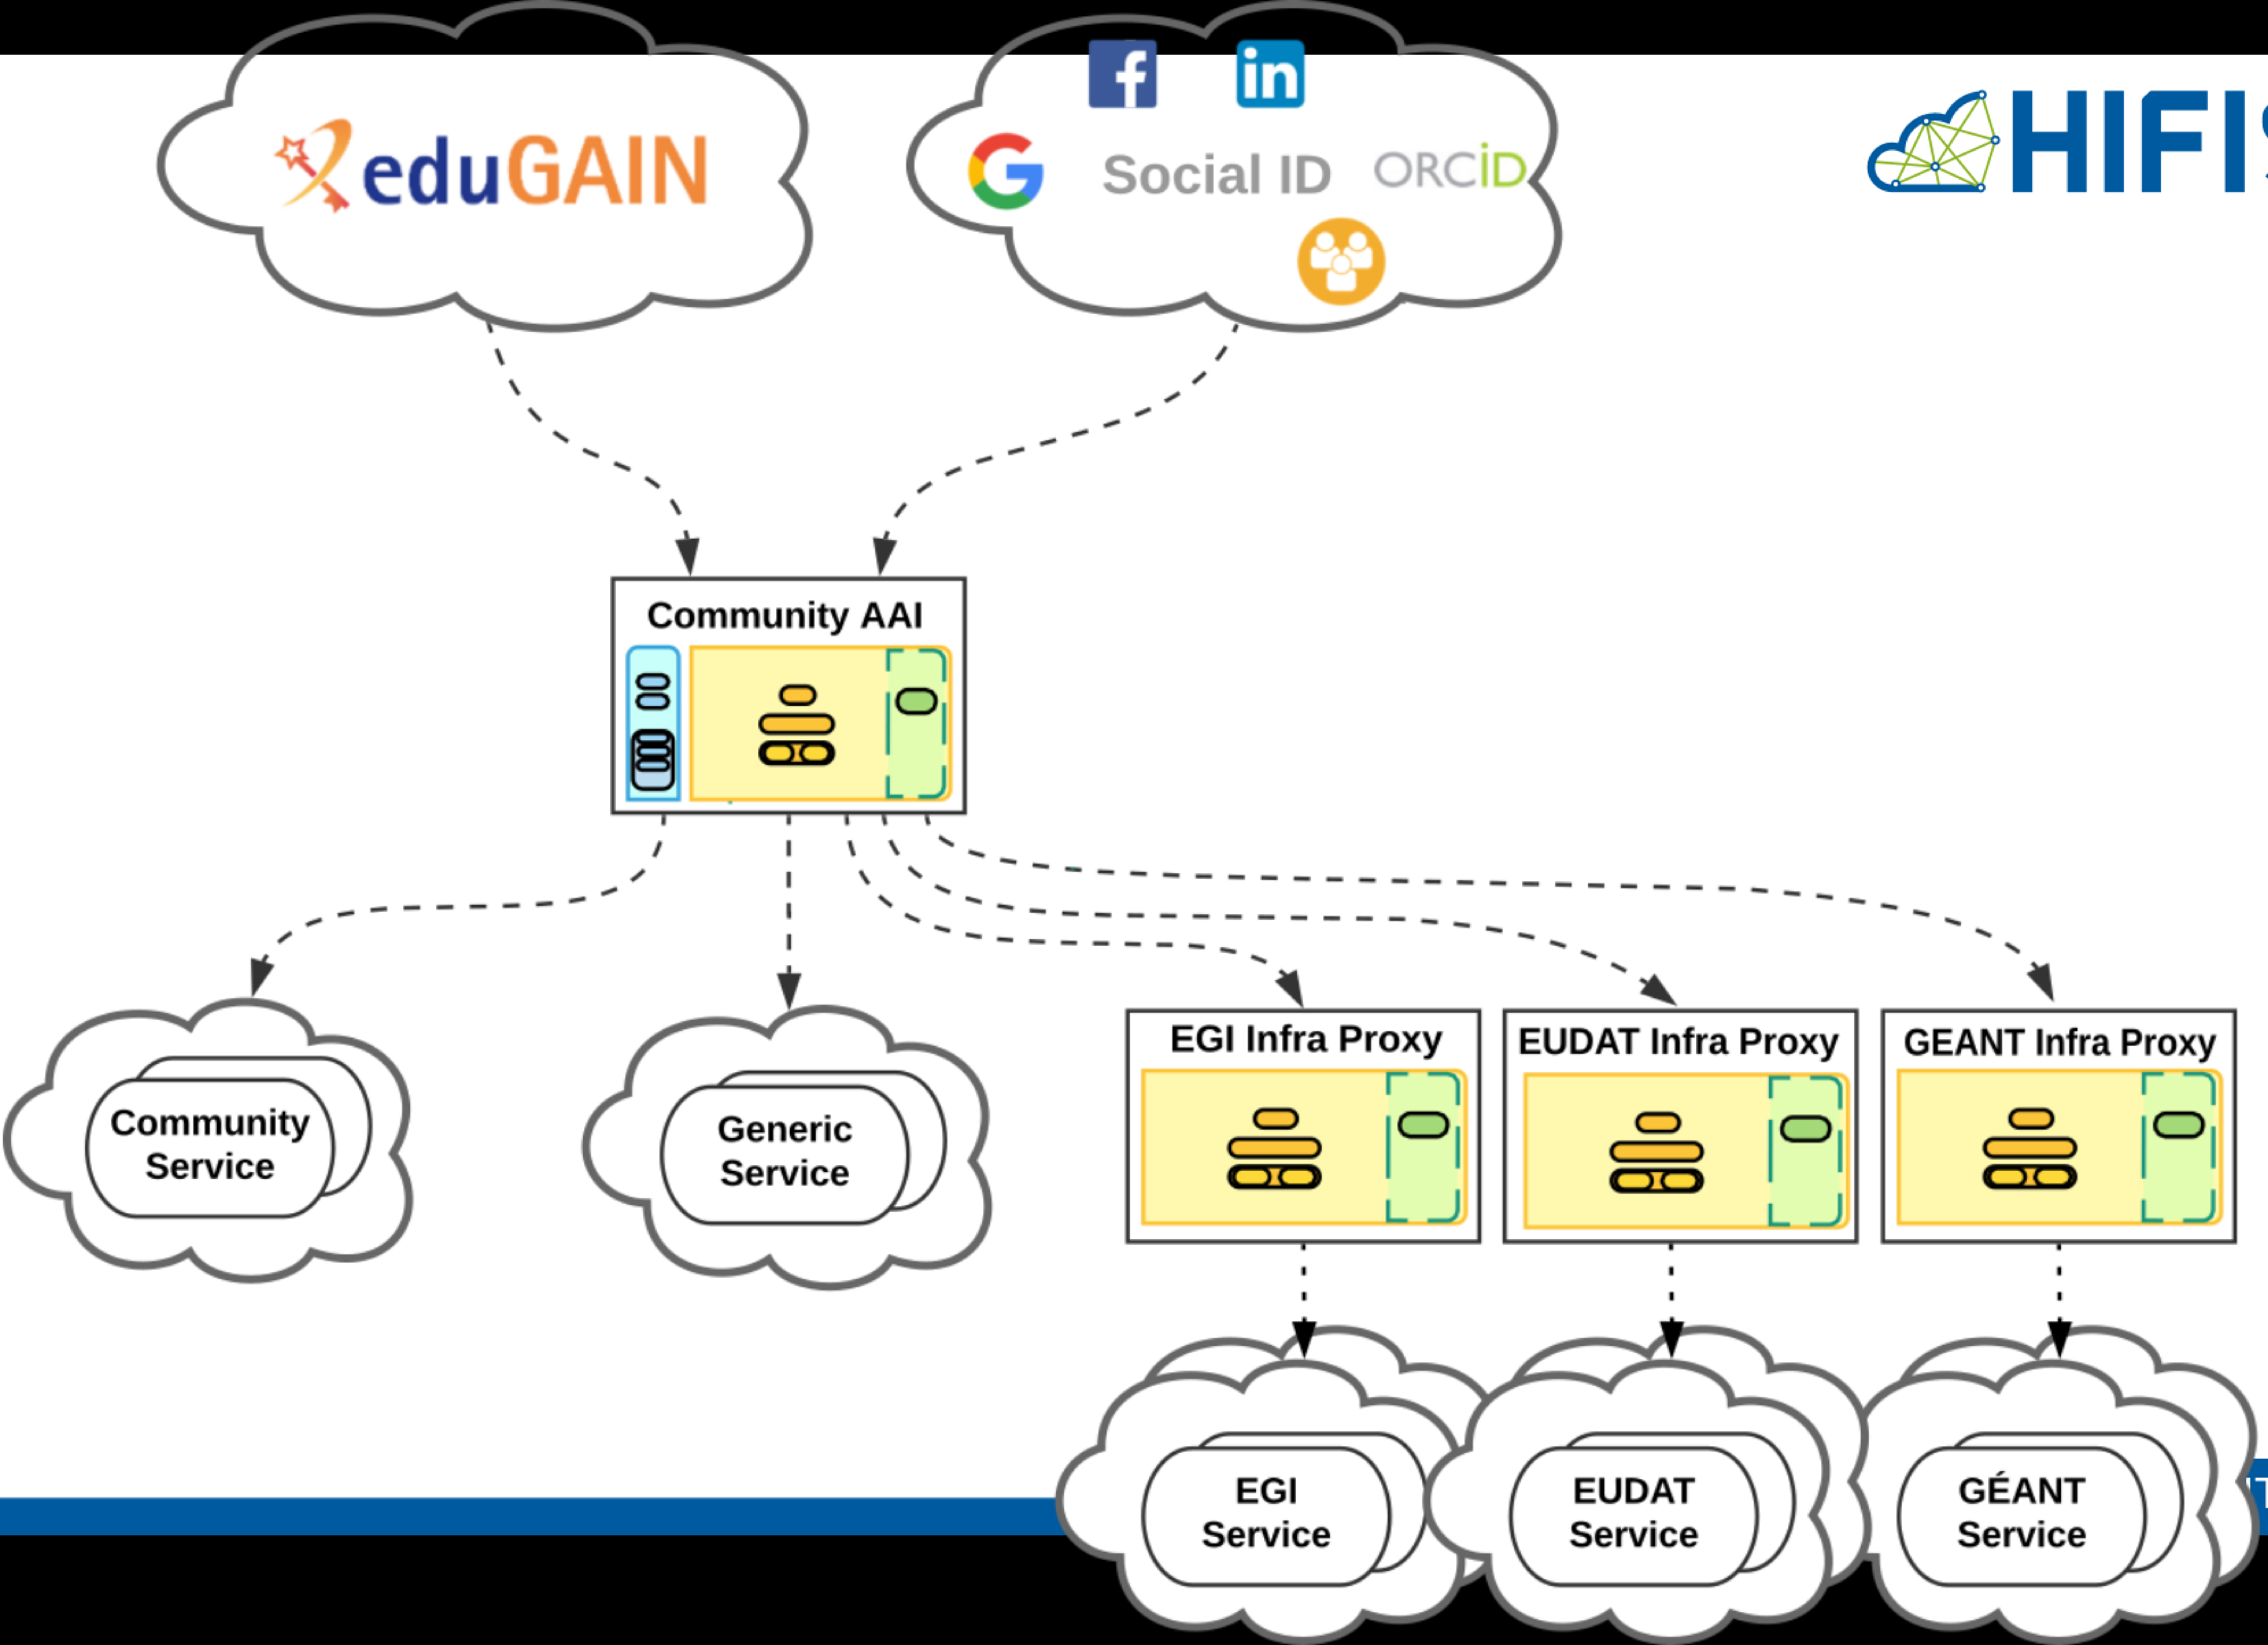
eduGAIN
Social ID
ORCiD
Community AAI
Community Service
Generic Service
EGI Infra Proxy
EUDAT Infra Proxy
GEANT Infra Proxy
EGI Service
EUDAT Service
GÉANT Service

HIFIS

Implementation

Helmholtz-AAI implements AARC BPA

Technical implementation

- Software: `unity` (also used for Eudat's b2access)
  - Production: https://login.helmholtz.de
  - Development: https://login-dev.helmholtz.de
- Self service Group Membership:
  - Principal Investigators can request a group
  - Manage their members

Helmholtz-AAI Features

- Well documented at https://aai.helmholtz.de
- Implements the Policy Development Kit
- Follows AARC recommendations (a lot of the `G0XY` documents)
  - <=> To use specific schemas for attributes and their content
  - HIFIS is an (observing) member of AEGIS
- Focus on OIDC
  - OpenID is not OpenID Connect
  - OpenID connect is defined be the OpenID Foundation
  - The OpenID protocol is deprecated
- Three levels of authorisation
  - Community based
  - Home Organisation Based
  - Assurance Based

Authorisation Management Based on Community

- **Virtual Organisation (VO)** approach
- Very similar: HPC compute projects
- **VO** Managers can administer community members
- Services can filter users by
  - VO Attributes
    - `climate -> ozone -> south-pole`
    - `cern -> cms -> admin`

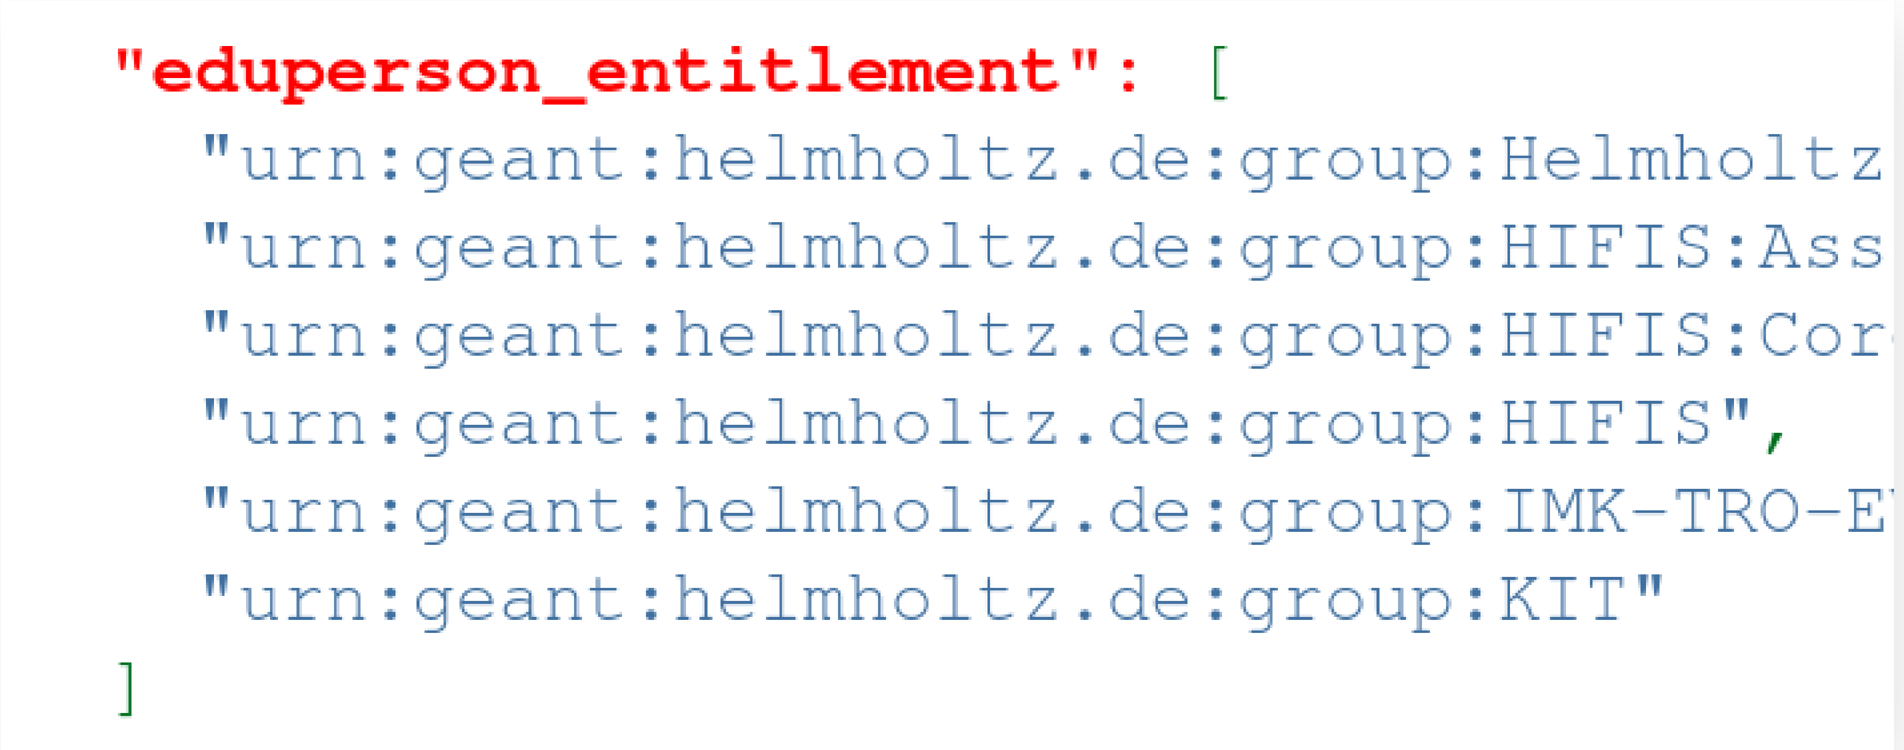

```
"eduperson_entitlement": [
  "urn:geant:helmholtz.de:group:Helmholtz
  "urn:geant:helmholtz.de:group:HIFIS:Ass
  "urn:geant:helmholtz.de:group:HIFIS:Cor
  "urn:geant:helmholtz.de:group:HIFIS",
  "urn:geant:helmholtz.de:group:IMK-TRO-E
  "urn:geant:helmholtz.de:group:KIT"
]
```

Authorisation Management Based on Origin

- **Home-IdP based** approach
- Home IdP can assert complementary information
- Services can filter users by
    - Home-Org asserted eligibility to use certain resources
    - Status: - Employee / Student / Guest

```
"eduperson_entitlement": [
  "http://bwidm.de/entitlement/bwLSDF-Syn
  "urn:mace:dir:entitlement:common-lib-te
]
```

Authorisation Management Based on Assurance

- Levels of Assurance: REFEDS Assurance Framework
  - Passport seen, Work-Contract available (Most academic Institutes)
  - Uniqueness of the identifier
  - Freshness of attributes
  - Verified Email Address (Social Media)
- Benefit
  - AAI can host “lesser-than-maximum” users
  - Scientists only need to upgrade their identity, if necessary to access service
  - Services can provide different levels of access

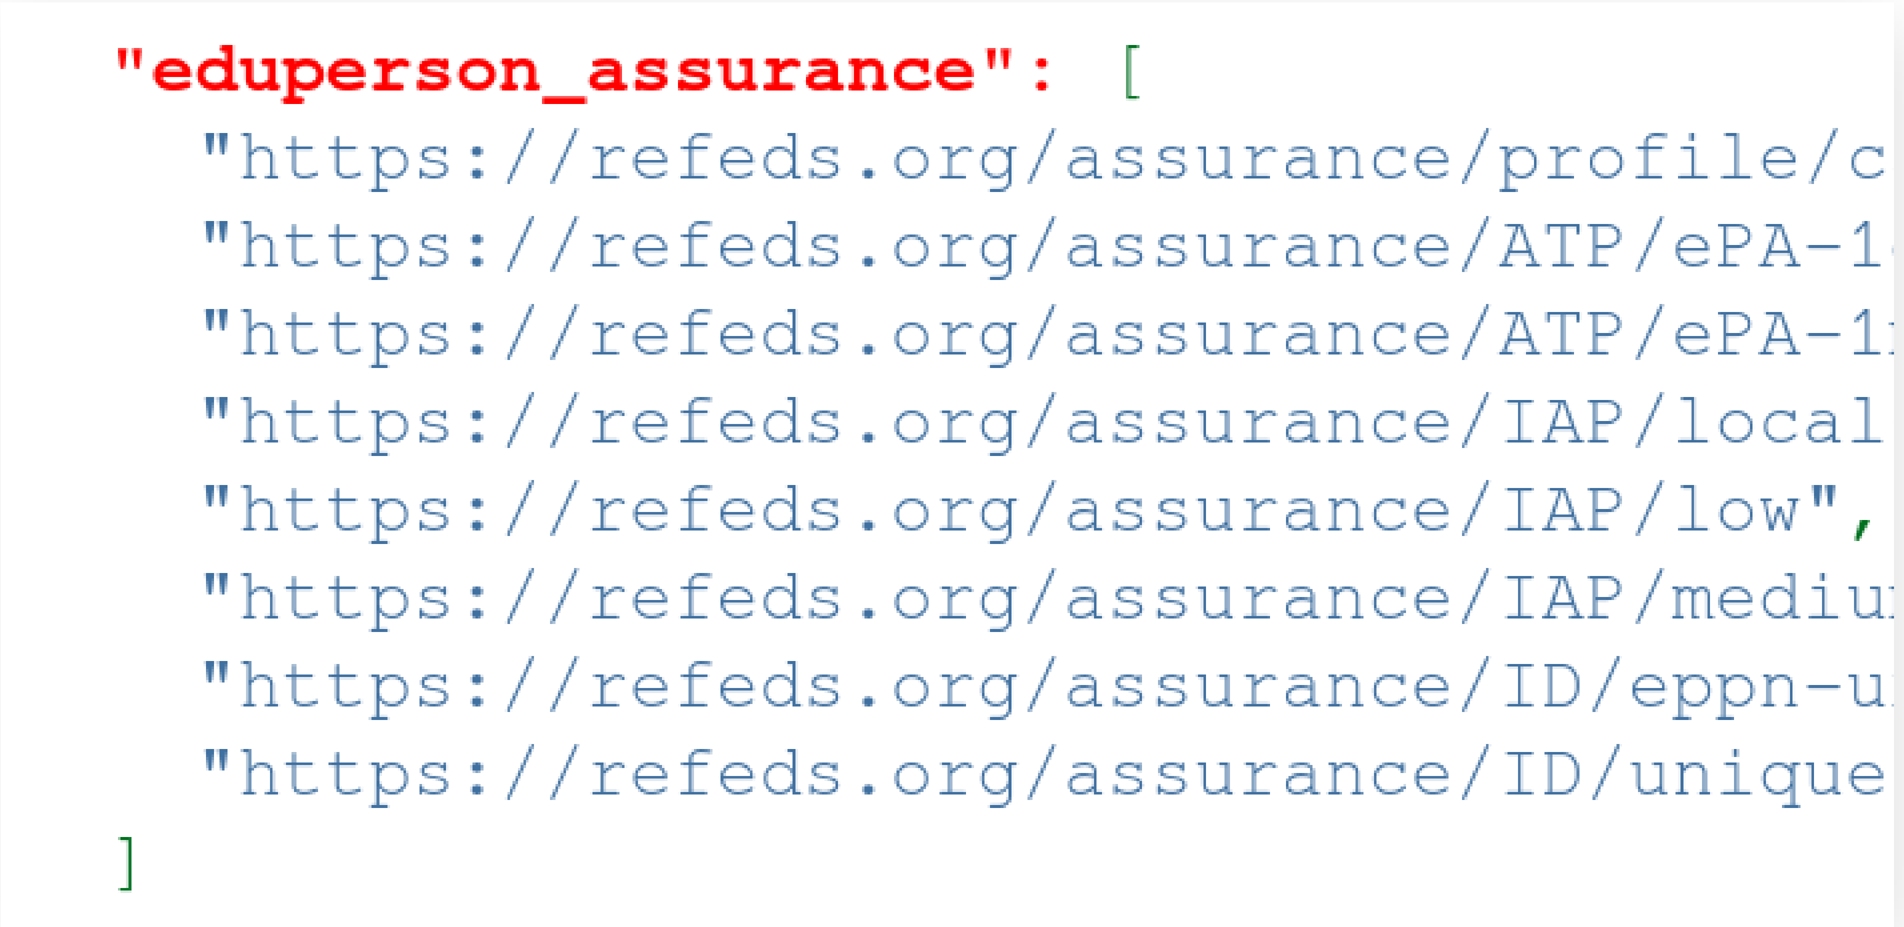

```
"eduperson_assurance": [
  "https://refeds.org/assurance/profile/c
  "https://refeds.org/assurance/ATP/ePA-1
  "https://refeds.org/assurance/ATP/ePA-1
  "https://refeds.org/assurance/IAP/local
  "https://refeds.org/assurance/IAP/low",
  "https://refeds.org/assurance/IAP/mediu
  "https://refeds.org/assurance/ID/eppn-u
  "https://refeds.org/assurance/ID/unique
]
```

Information available at services

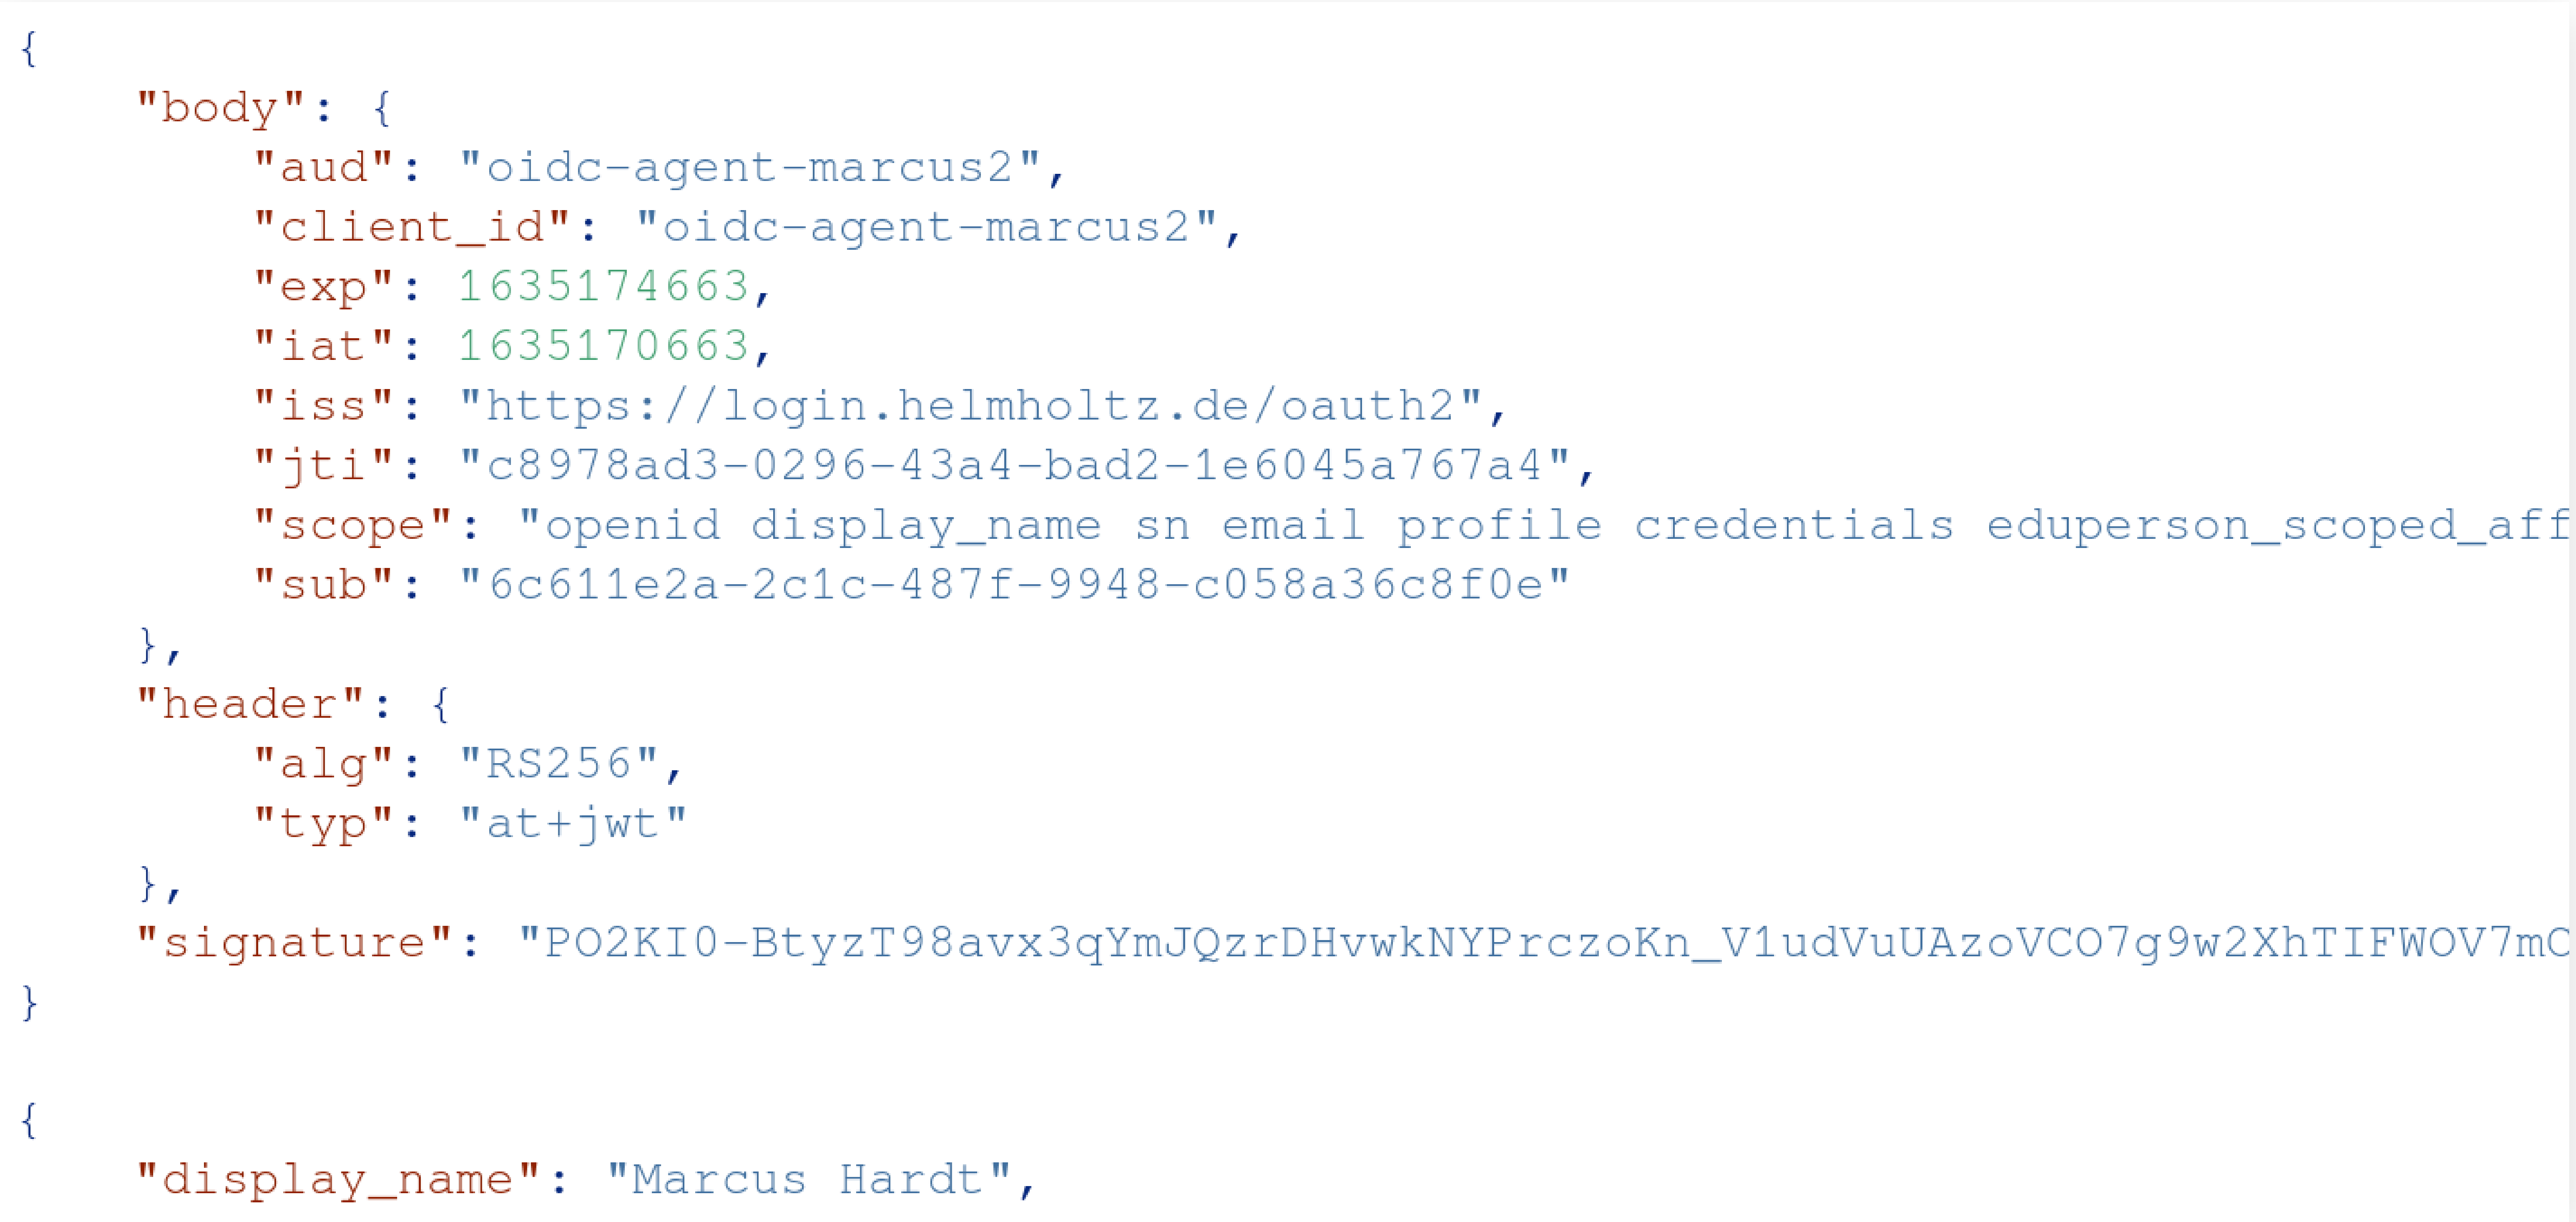

```
{
    "body": {
        "aud": "oidc-agent-marcus2",
        "client_id": "oidc-agent-marcus2",
        "exp": 1635174663,
        "iat": 1635170663,
        "iss": "https://login.helmholtz.de/oauth2",
        "jti": "c8978ad3-0296-43a4-bad2-1e6045a767a4",
        "scope": "openid display_name sn email profile credentials eduperson_scoped_aff
        "sub": "6c611e2a-2c1c-487f-9948-c058a36c8f0e"
    },
    "header": {
        "alg": "RS256",
        "typ": "at+jwt"
    },
    "signature": "PO2KI0-BtyzT98avx3qYmJQzrDHvwkNYPrczoKn_V1udVuUAzoVCO7g9w2XhTIFWOV7mC
}

{
    "display_name": "Marcus Hardt",
```

Connected Services

- Multi Protocol:
  - Identities: SAML, OpenID Connect, X.509
  - Services: OpenID Connect, SAML
- Integrated Services:
  - Helmholtz Federated IT services (HIFIS, hifis.net)
    - Drives development, documentation and service integration
    - cloud.helmholtz.de/services
  - Technially feasible: Rocketchat, Storage, Compute & more.
  - More pilot services at documentation pages
  - Exhaustive list: aai.helmholtz.de/services

AAI usage

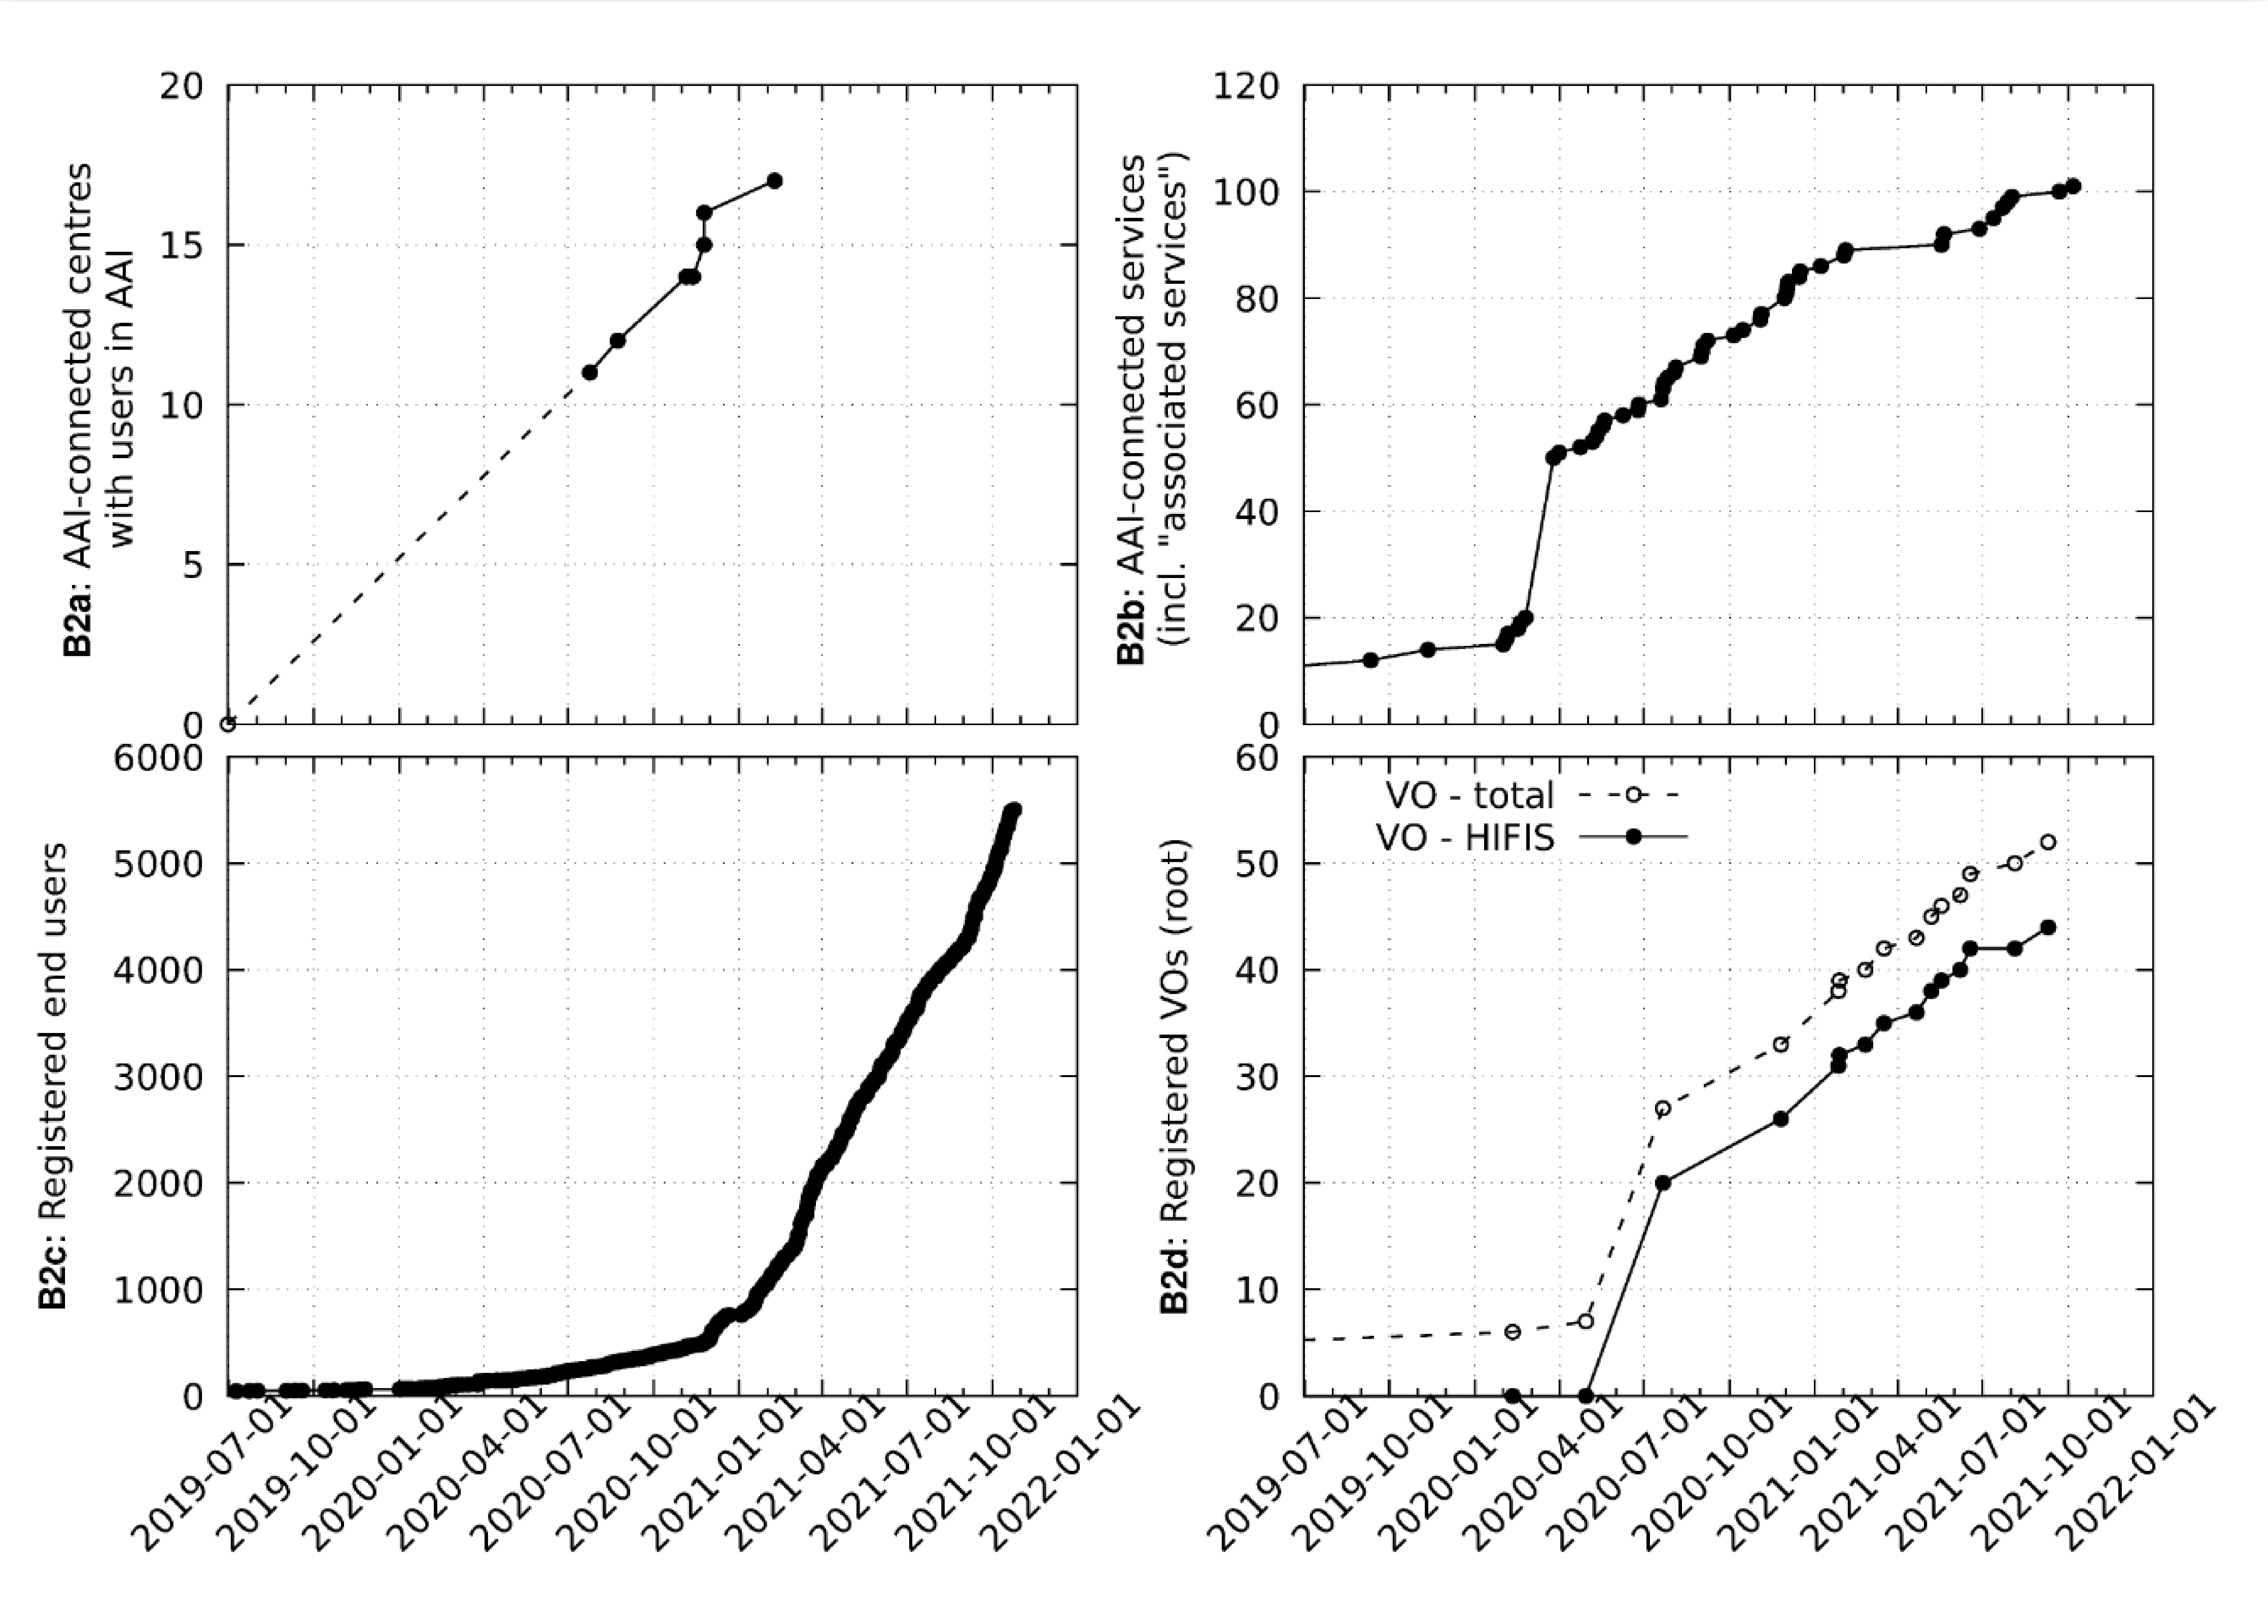
B2a: AAI-connected centres with users in AAI
0
5
10
15
20
B2b: AAI-connected services (incl. "associated services")
0
20
40
60
80
100
120
B2c: Registered end users
0
1000
2000
3000
4000
5000
6000
B2d: Registered VOs (root)
0
10
20
30
40
50
60
VO - total
VO - HIFIS
2019-07-01
2019-10-01
2020-01-01
2020-04-01
2020-07-01
2020-10-01
2021-01-01
2021-04-01
2021-07-01
2021-10-01
2022-01-01

AAI Developments in Helmholtz

Helmholtz Cloud Agent

- https://hifis.net/doc/service-integration/local-agent/
- First (exemplary) use-case: Nubes
  - Enable DESY cloud portal (cloud.helmholtz.de)
  - To use (Nubes)https://nubes.helmholtz-berlin.de resources at HZB
  - Challenge:
    - Exchange user provisioning information
    - Integrate local systems

oidc-agent

- Goal: Support OIDC on end user computers
- Initial goal: Unix commandline (linux + mac)
  - Handle all the issues with different OIDC Providers
  - Adequate security features
    - All sensitive information on disk is encrypted
    - Everything (sensitive) in RAM is obfuscated
    - Keep the user from stupid moves
  - Works just like ssh-agent
    - `oidc-agent, oidc-gen, oidc-add, oidc-token`
    - Including **agent forwarding** and **x-session integration**
  - Works well with many OIDC providers
    - Google, Eudat, eduTEAMS, EGI-Checkin, Elixir, Helmholtz-AAI, WLCG, Indigo IAM, KIT, Human Brain, ...
- New goal: Support for GUI environments (windows + mac + linux)

mytoken

- Mytokens are a new class of tokens
- Use case: Long running compute job
  - Longer than lifetime of Access Token

- **Mytoken Server**
  - Proxy for Refresh Tokens (RT)
  - Implemented as an extension of OIDC
- User flow:
  1. Create mytoken (MT)
  2. Use MT to obtain
     - Access Tokens (AT)
     - Other mytokens

```
[{"exp"         :1634300000,
  "nbf"         :1634400000,
  "geoip_allow":["DE"],
  "scope"       :"compute.create",
 },{
  "exp"         :1635300000,
  "nbf"         :1635400000,
  "geoip_allow":["DE", "FR", "NL"],
  "scope"       :"storage.write",
}]
```

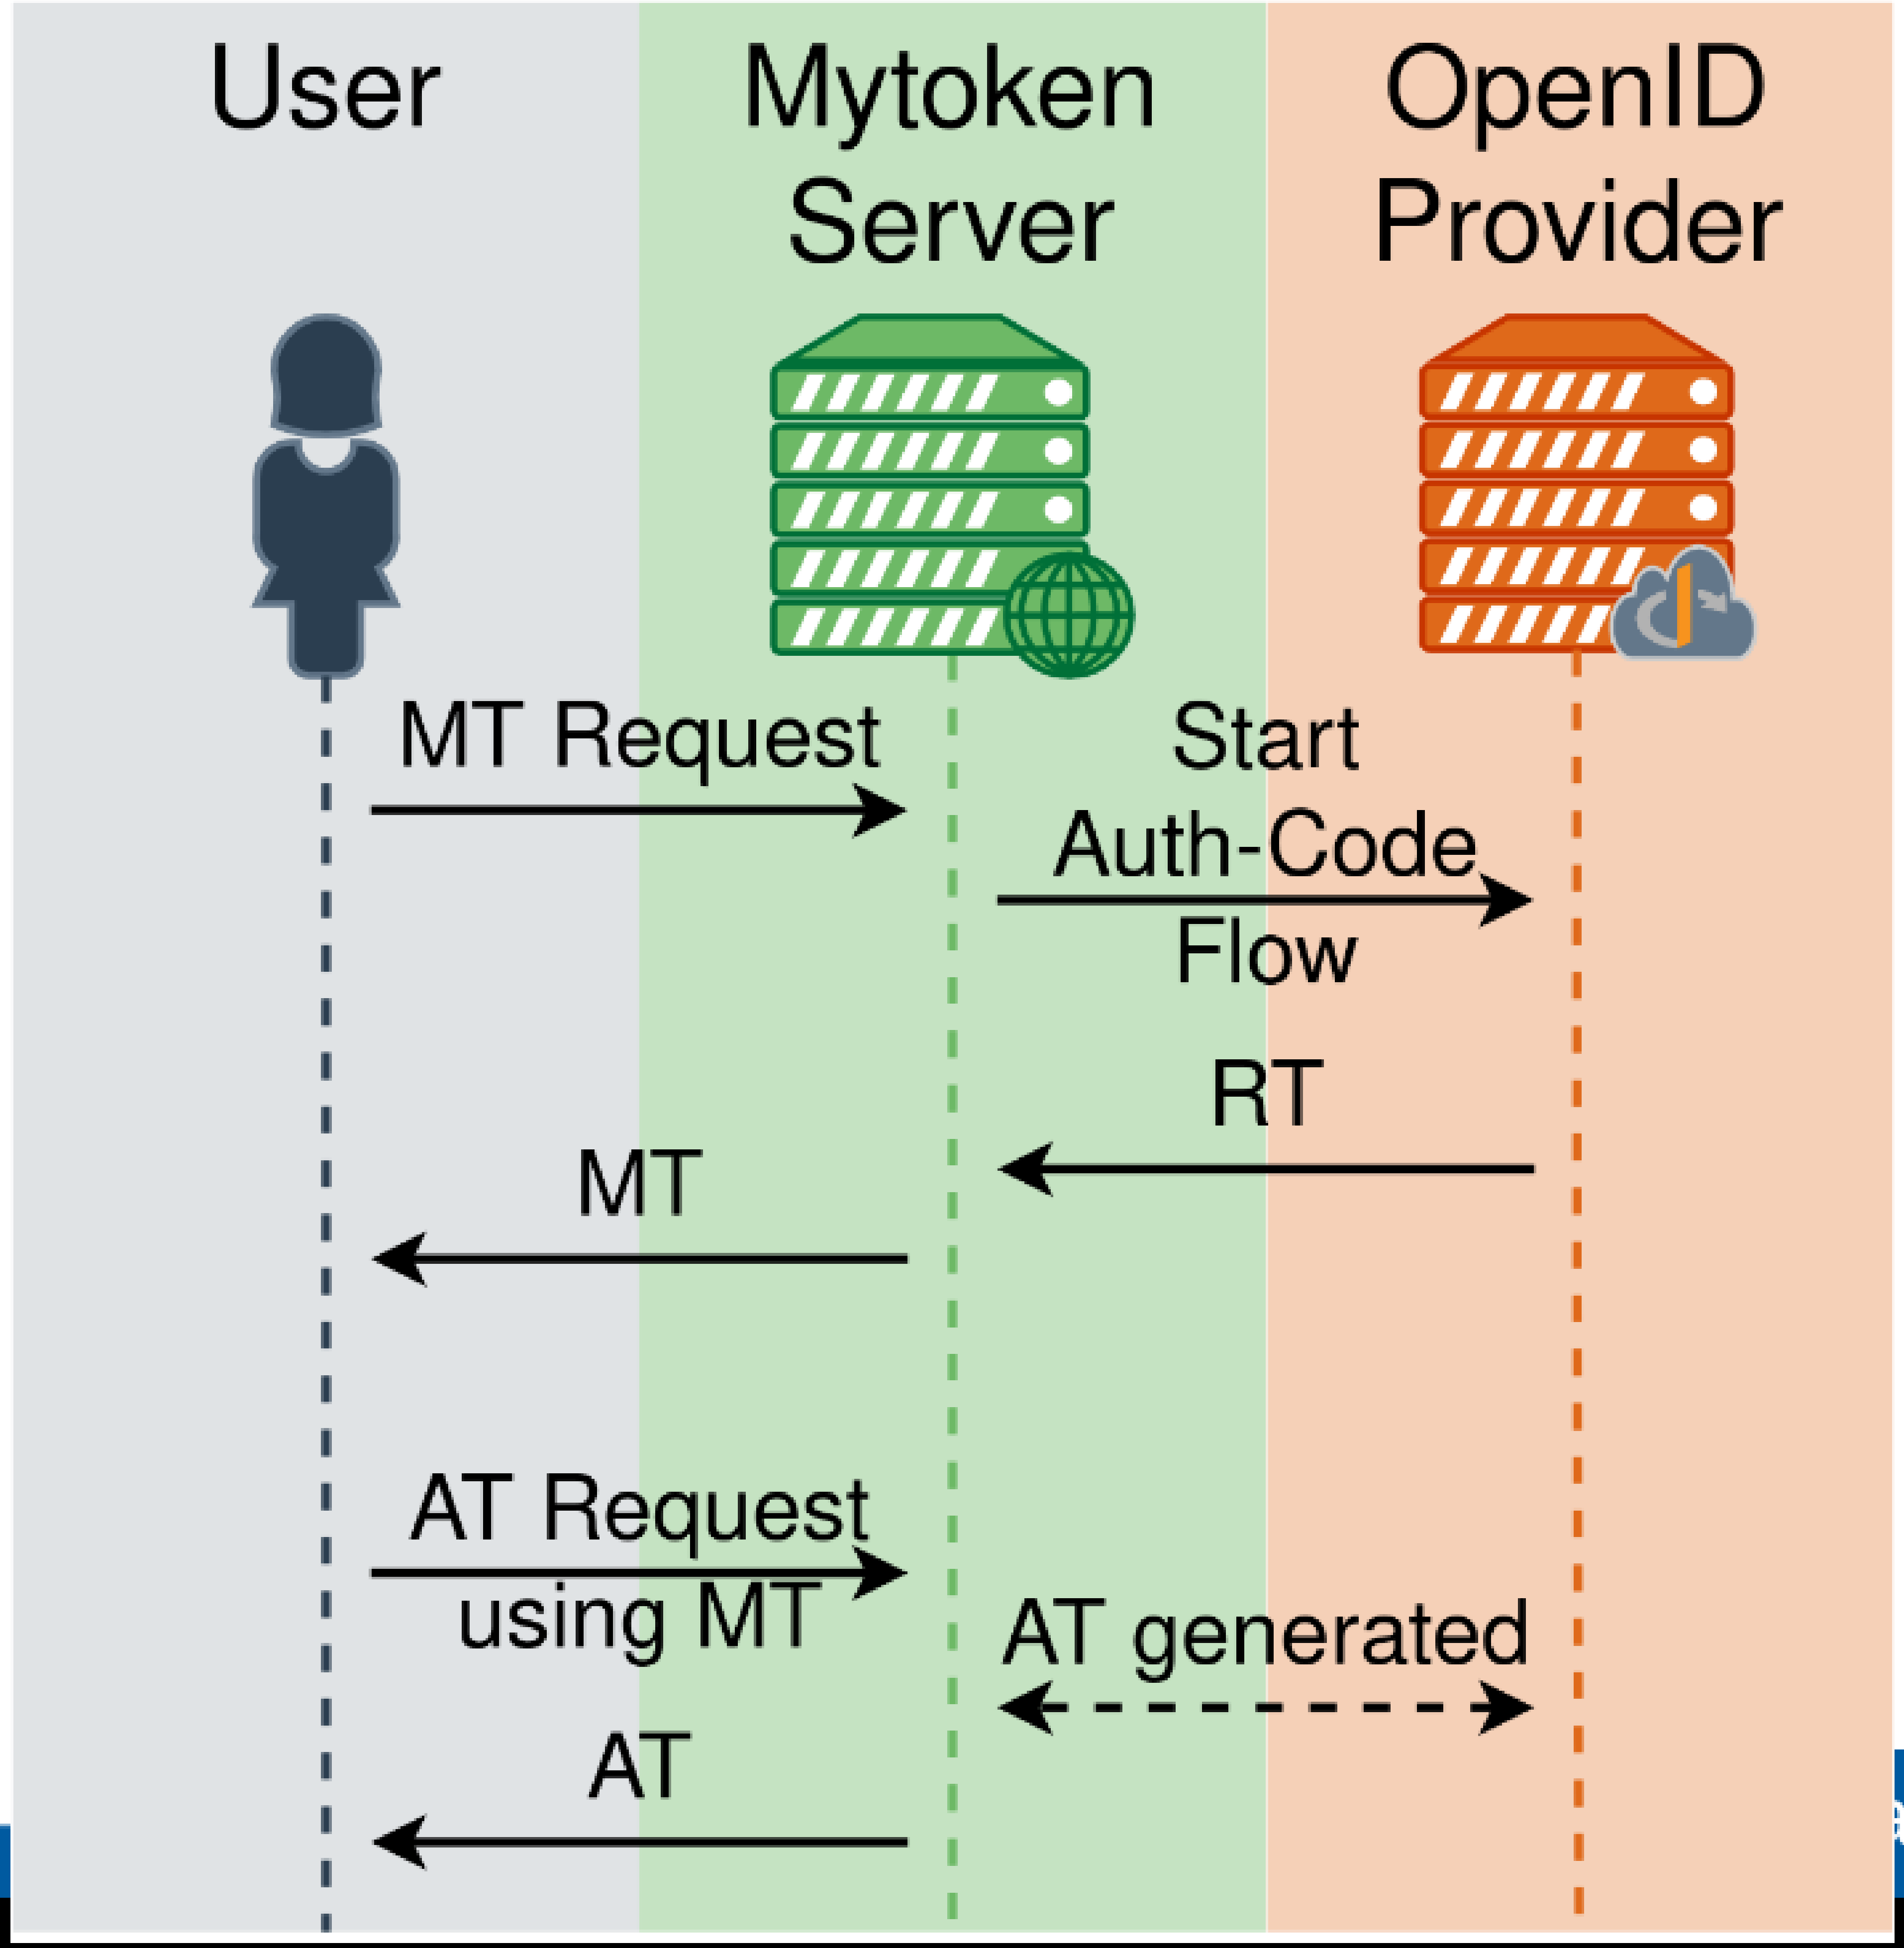

ssh-oidc

- Enable **ssh** via **federated identity** (OIDC)
  - without recompiling OpenSSH
  - with a clear authorisation concept
- Solution:
  - PAM module
  - Mapping Daemon
  - Client Wrapper
- Available for Linux
  - Mac and Windows in development (Putty, maybe: MobaXterm)
- Test it at https://ssh-oidc-demo.data.kit.edu

Outlook

- Integrate more services
  - Large Resources (HPC / Clusters)
- Spread the technology
- Interoperate with other Community AAIs
  - How to handle cross-community access?
  - How about OIDC-Federations?
- Contribute to AARC Guidelines:
  - IdP Hinting
  - SCIM and Deprovisioning
  - Expression of Entitlements
- Manage expectations, e.g. identity linking

More information

https://aai.helmholtz.de