

# Nikhef

 Maastricht University

How trust and identity enable  
global infrastructures for distributing computing

## In Infrastructures ... we Trust!

*ISGC 2025  
David Groep, March 2025*



# Collaborations: from small ...



Nikhef user room H1.37 – terminal stations in the early 1990's – image source: Nikhef

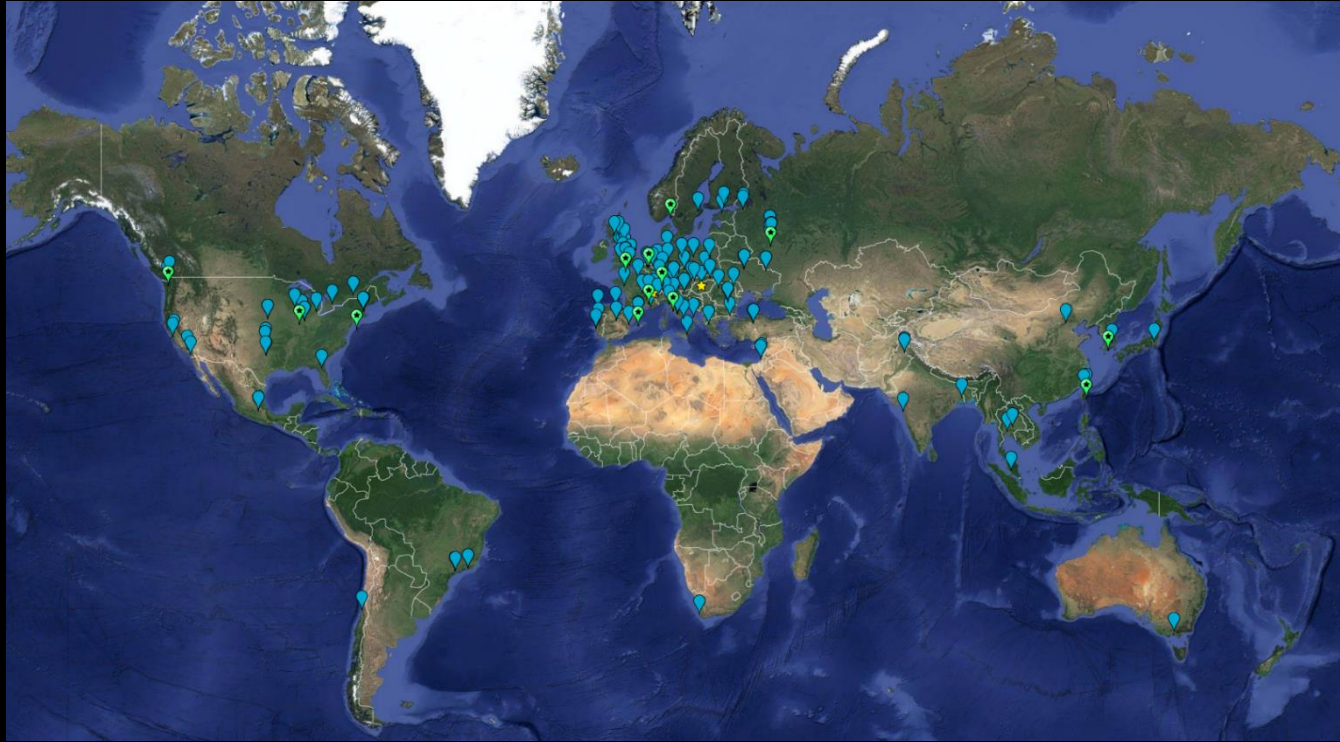
... to large collaborations (and shown here is a subset ...)



a small part of the CMS collaboration in 2017, photo credit CERN on behalf the CMS collaboration, CMS-PHO-PUBLIC-2017-004-3



# How many interactions? And just how many logins?



Worldwide LHC  
Computing Grid (~ 2024)  
~ 1.4 million CPU cores  
~ 1500 Petabyte  
disk + archival

170+ institutes  
42+ countries  
13 'Tier-1 sites'  
some multi-community:  
NL-T1 @ SURF & Nikhef

Earth background: Google Earth; Data and compute animation: STFC RAL for WLCG and EGI.eu; Data: <https://home.cern/science/computing/grid> ;  
LHC Computing Grid: [wlcg.web.cern.ch](http://wlcg.web.cern.ch), EGI: [www.egi.eu](http://www.egi.eu); ACCESS CI: <https://access-ci.org/>, NL-T1 and FuSE: [fuse-infra.nl](http://fuse-infra.nl), <https://www.surf.nl/en/research-it>



# Do we ask for ~ 12 000 x 170+ passwords for everyone?



Image "trust fall" by Barret Anspach (<https://www.flickr.com/photos/anspach/954545>) used under CC-BY-2.0 license

**Of course not!**

But 12 000 people is still a lot,  
*and many more than you would  
trust with your bank PIN ...*

Yet we have found mechanisms to  
collaborate beyond the canonical  
~150 people ("Dunbar's number")

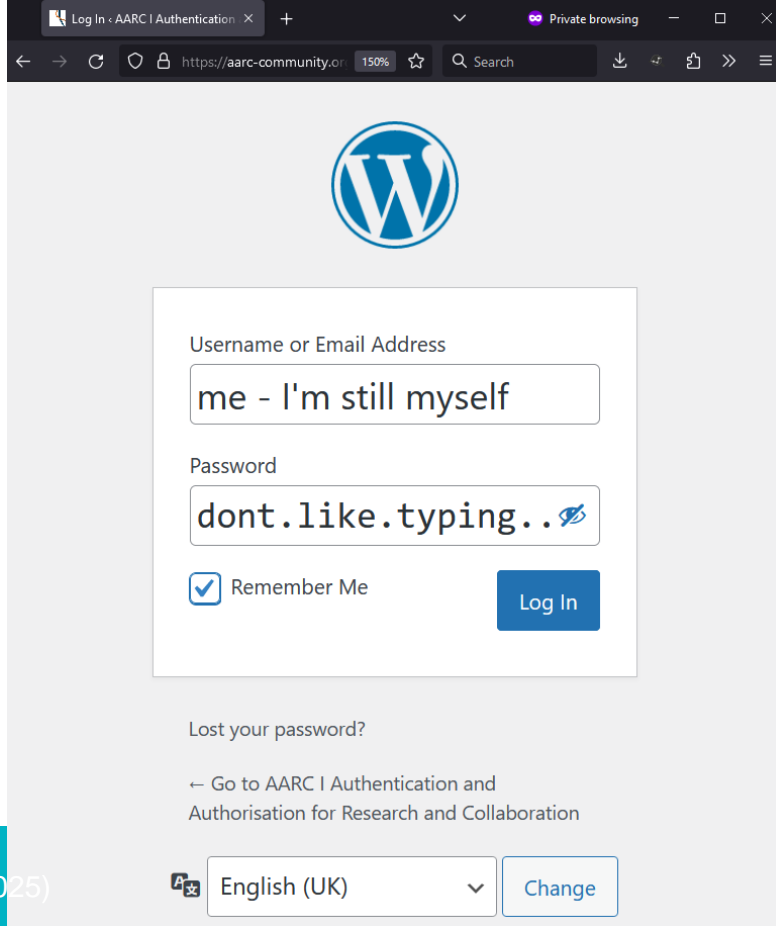
but what we built, may both be unique  
for our 'high-trust' research community  
... and be an example for others

# When you are asked to login again ...

## Authentication

demonstrating ‘you are you’

- ***authenticator***  
‘you’ remains same ‘you’
- ***vetted identity***  
‘you’ can be pseudonymous  
‘you’ can be a vetted person



# Self-asserted or 'pseudonymous' often not enough

*state of EU DataGrid and  
HEP computing in ~2000*



NATIONAAL INSTITUUT VOOR KERNFYSICA EN HOGE-ENERGIEFYSICA

Guest / students form (please)

1. This form is completed in connection with:
- work experience
- otherwise, visit

CERN/User Registration

**CERN COMPUTER CENTRE - US**

<http://cern.ch/it/documents/ComputerUsage/Comp>

To be returned to the User Registration box at the entrance of the computer room, completed by a user who requires a computer account, and is not yet registered in another group Department, and is not yet registered in another group

**To be completed by the User :**

It is **MANDATORY** to provide the following information, which will be treated confidentially and only be used for ensuring the correct supply name as registered by the Users' Office

FAMILY NAME(S) : .....

FIRST NAME(S) : .....

SEX [M] [F]      BIRTHDATE: Day ..... Month ..... Year .....

HOME INSTITUTE/FIRM: .....

NATIONALITY: .....\*CERN SUPERVISOR.....

\*CERN DEPARTMENT: . . . . \*CERN ID NUMBER (as on CERN card).....

**To be completed by the Group Administrator:**



**Fermilab**

**For Office Use Only**

ID:	Action:	ID Exp:	
Insurance:	Medical:	Safety:	
Computer:	Stkrm:	Family:	
NON-473:	Sensitive:	Verifier:	Date:

**Name:**


<b>SWIETZER</b>	<b>JOHN</b>	<b>JAMES</b>
Last	First	Middle

**University or Institution Name:**

<b>FLORIDA STATE UNIVERSITY</b>	<b>850-644-XXXX</b>
Telephone:	

**Experiment/Department:**

Exp. / Dept.	Spokesperson	Home Institution Contact	Contact Telephone
<b>D0</b>	<b>WOMERSLEY/WEERTS</b>	<b>SHARON HAGOPIAN</b>	<b>850-644-4777</b>





# Scaling credentials: per service per user

Many start with *credentials* dedicated to each service where you need access

- In a multi-organizational system becomes

$$a_{n_{\text{services}}} * a_{n_{\text{users}}}$$

- usually creates a strong link to authorization:

*different accounts for different roles,  
multiplying the number of credentials per user*

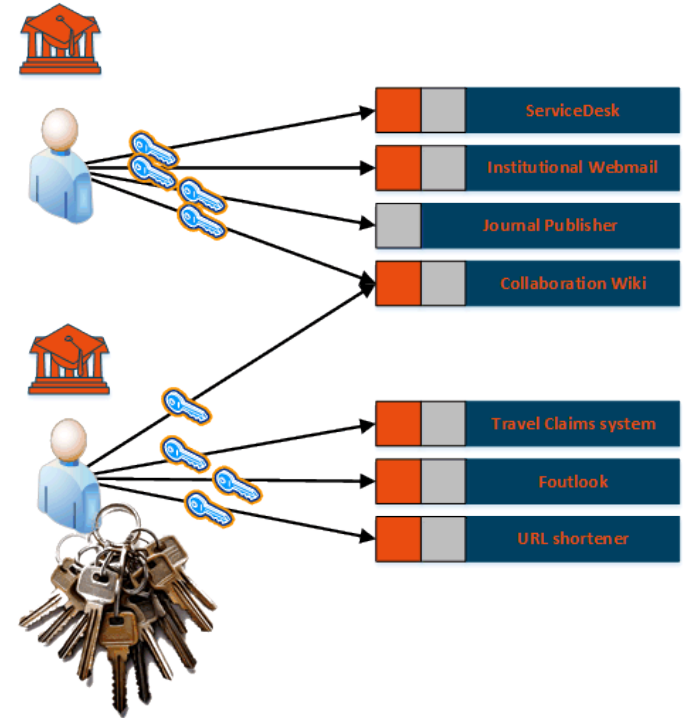


Image inspired by AARC NA2 training module "Authentication and Authorisation 101" – keychain image created by generative AI

# bilateral 'SSO': a single service, or a single identity source

#credentials required?


from previously

$$\mathcal{O}(n_{\text{services}}) * \mathcal{O}(n_{\text{users}})$$

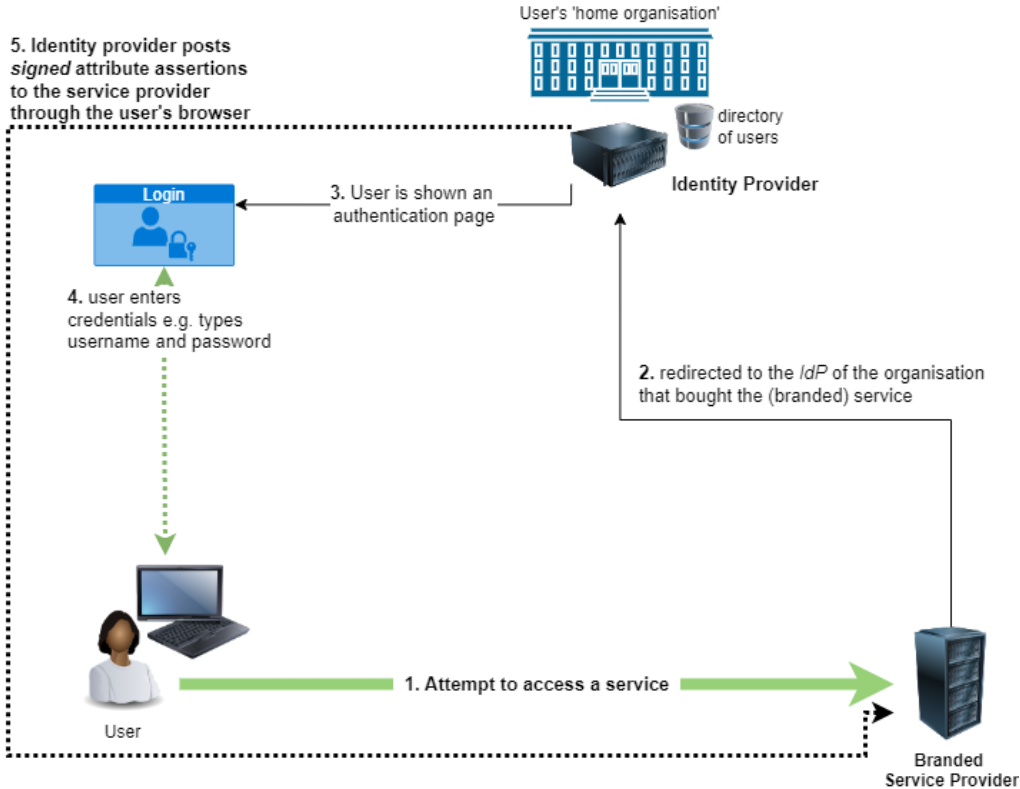
to

$$\mathcal{O}(n_{\text{users}}) + \mathcal{O}(n_{\text{services}} * n_{\text{home-orgs}})$$

*in first order at least*



# Single sign-on – why your browser keeps loading things



```
Extension: (SAML-tracer) - SAML-tracer — Mozilla Firefox
X Clear  || Pause  ↓ Autoscroll  ▾ Filter resources  0 Colorize  ⬇ Export  ⬇ Import

GET https://commute.nikhef.nl/
GET https://commute.nikhef.nl/favicon.ico
GET https://commute.nikhef.nl/commute/?auth=nikhef-ss0
GET https://sso.nikhef.nl/sso/saml2/idp/SSOService.php?SAMLRequest=fVJL SAML
GET https://sso.nikhef.nl/sso/module.php/nikhef/loginuserpass.php?AuthState=_9d4f7
GET https://sso.nikhef.nl/sso/module.php/consent/getconsent?StateId=_9d4f753ffc12c
GET https://sso.nikhef.nl/sso/resources/icons/favicon.ico
GET https://sso.nikhef.nl/sso/module.php/consent/getconsent?saveconsent=1&StateId
POST https://commute.nikhef.nl/simplesaml/module.php/saml/sp/saml2-ac3.php SAML
GET https://commute.nikhef.nl/commute/?auth=nikhef-ss0
GET https://commute.nikhef.nl/favicon.ico
```

## Glossary

'SAML' is the "Security Assertion Mark-up Language" an XML blob with information, usually digitally signed

```
HTTP Parameters SAML Summary
Version="2.0"
IssueInstant="2025-02-28T11:49:04Z"
<saml:Issuer>https://sso.nikhef.nl/sso/saml2/idp/metadata.php</
```

SAML-tracer plugin by Tim van Dijen (SSC-ICT) *et al.*  
<https://github.com/simplesamlphp/SAML-tracer>



# User-centric identity: 'I take my passport anywhere by myself'

Your 'home organisation' does not have to be in the loop ...



*user-centric* trust: you yourself hold a credential from a trusted third party and can use it *without having to ask 'home' each time*:

- Public Key Infrastructure client certificates (“X.509”)
- Verifiable credentials in wallets
- *and who remembers CardSpace?*

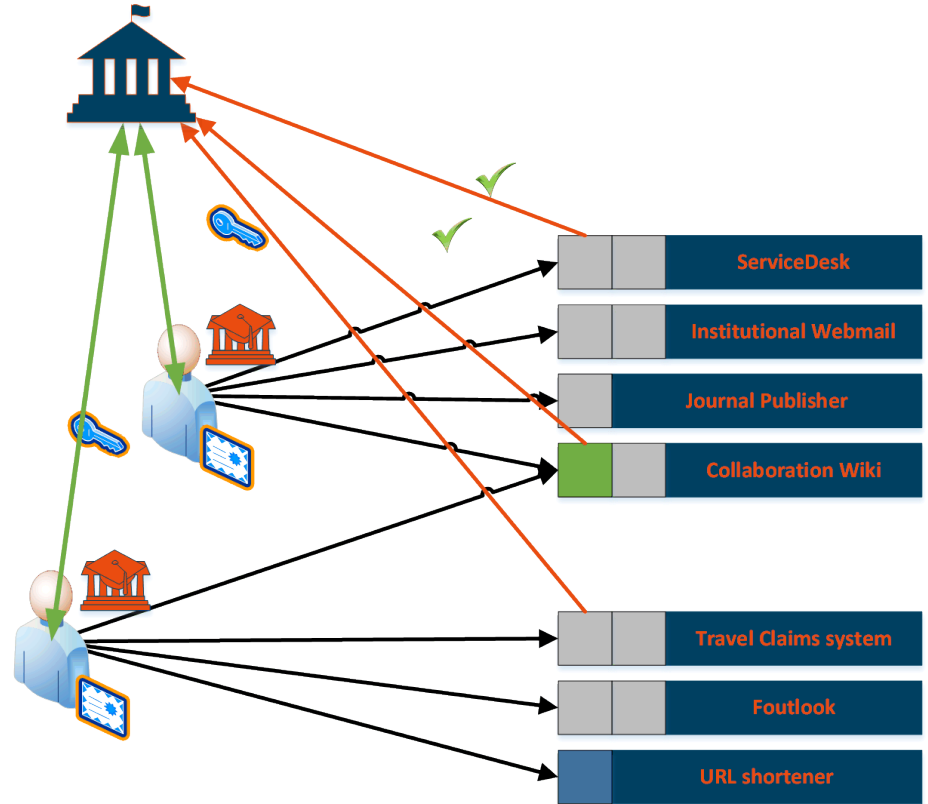
Passport image: cropped from original by Jon Tyson on Unsplash <https://unsplash.com/photos/Hid-yhommOg>

# User-centric AAI

A **trusted authority** giving the user a 'self-managed' credential, like a passport

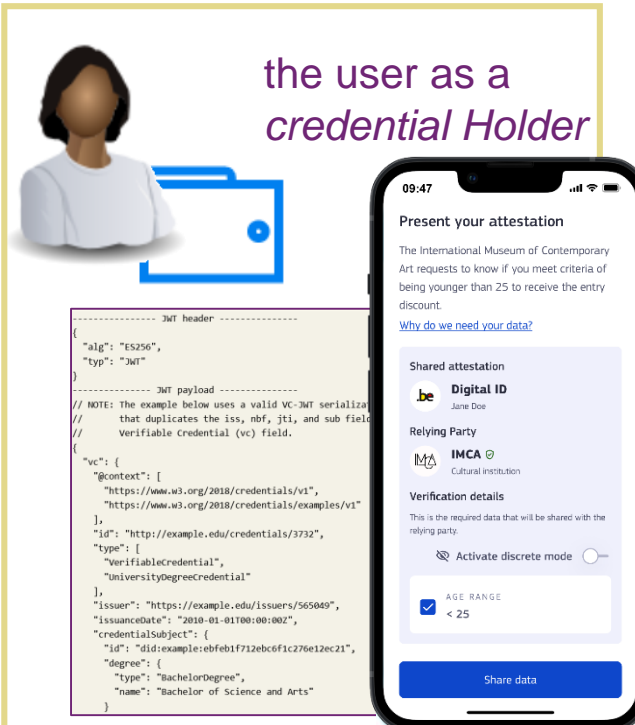
- a personal authentication digital certificate
- a verifiable credential in a wallet
- ...

verified (on-line and also offline)  
at the original trusted issuer  
or at an independent trusted verifier



# Identity wallets, held by the user, are another

the user as a *credential Holder*



```

----- JWT header -----
{
  "alg": "ES256",
  "typ": "JWT"
}
----- JWT payload -----
// NOTE: The example below uses a valid VC-JWT serialization
// that duplicates the iss, nbf, jti, and sub field
// of a verifiable credential (vc) field.
{
  "vc": {
    "@context": [
      "https://www.w3.org/2018/credentials/v1",
      "https://www.w3.org/2018/credentials/examples/v1"
    ],
    "id": "http://example.edu/credentials/3732",
    "type": [
      "VerifiableCredential",
      "UniversityDegreeCredential"
    ],
    "issuer": "https://example.edu/issuers/565049",
    "issuanceDate": "2019-01-01T00:00:00Z",
    "credentialSubject": {
      "id": "did:example:ebfeb1f712ebc6f1276e12ec21",
      "degree": {
        "type": "BachelorDegree",
        "name": "Bachelor of Science and Arts"
      }
    }
  }
}
    
```

Present your attestation

The International Museum of Contemporary Art requests to know if you meet criteria of being younger than 25 to receive the entry discount.

[Why do we need your data?](#)

Shared attestation

**be Digital ID**  
Jane Doe

Relying Party

**IMCA** Cultural institution

Verification details

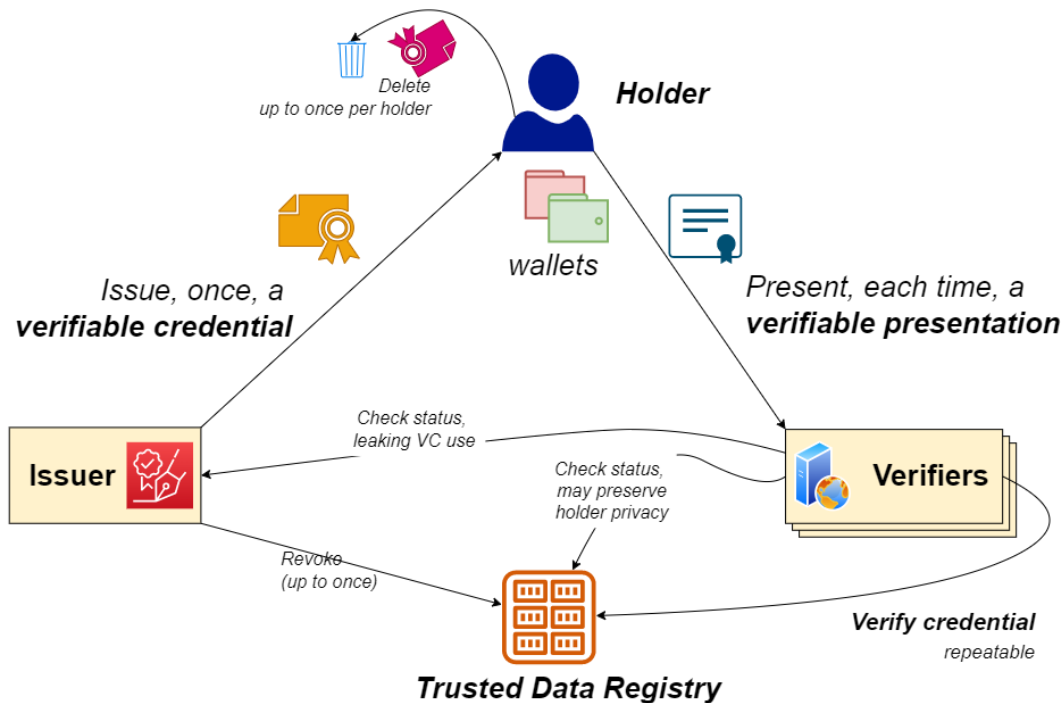
This is the required data that will be shared with the relying party.

Activate discrete mode

AGE RANGE

< 25

Share data



Flow diagram inspired by: Lifecycle Details (5.1), Verifiable Credentials Data Model v1.1, W3C Recommendation 03 March 2022, <https://www.w3.org/TR/vc-data-model/>  
 EU eID Wallet from [https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/european-digital-identity\\_en](https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/european-digital-identity_en)  
 An image: European Commission, at <https://ec.europa.eu/digital-building-blocks/sites/display/EUDIGITALIDENTITYWALLET/Security+and+Privacy>



# Can we scale better with an 'federated' Authentication and Authorisation Infrastructure ('AAI')



with one service provided to several organisations (universities)

*we will get to authorisation in a bit ...*

# Where are 'you' in the federated space – discovery!

The image illustrates the user journey from institution discovery to authentication. It shows the HARICA 'Choose Your Institution' page, which lists recent institutions: Nikhef (nikhef.nl), CERN (cern.ch), and Maastricht University. From this page, users can click on an institution to reach its specific login page. The 'Academic Login' button on the HARICA login page is highlighted in red. The authentication pages shown are for Nikhef (National Institute for Subatomic Physics), CERN Single Sign-On, and Maastricht University, each requiring a username and password.

An example cross-institutional service by HARICA, the GEANT TCS G5 provider, presenting a SeamlessAccess.org discovery page

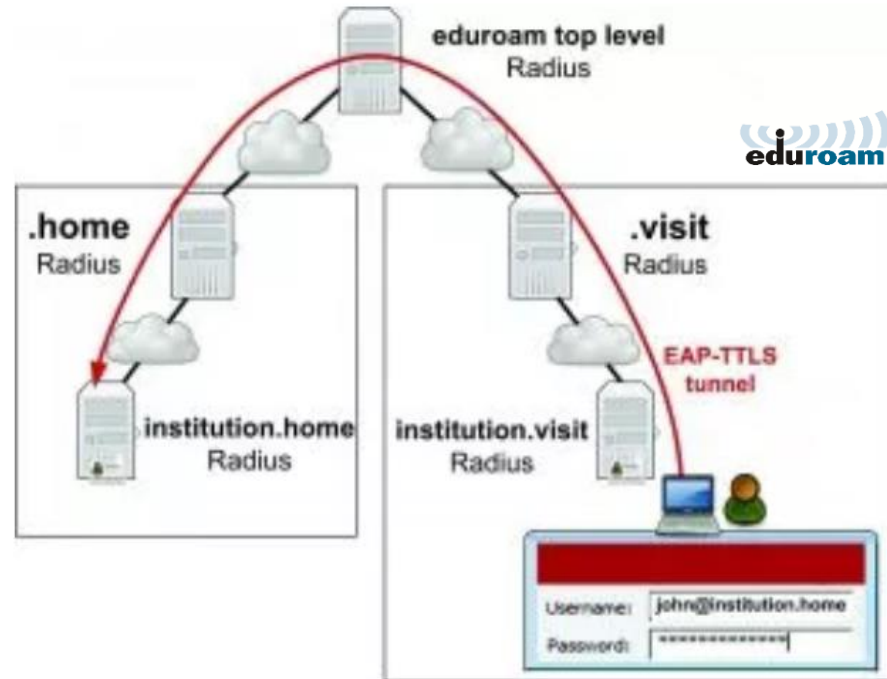
# The federation you most likely know ...

*service-specific* trust  
between organisations

hierarchical server path, based on  
a network-specific secure exchange

sending your credentials back  
to *only* your home institution

found via `<anon@domain.name>`



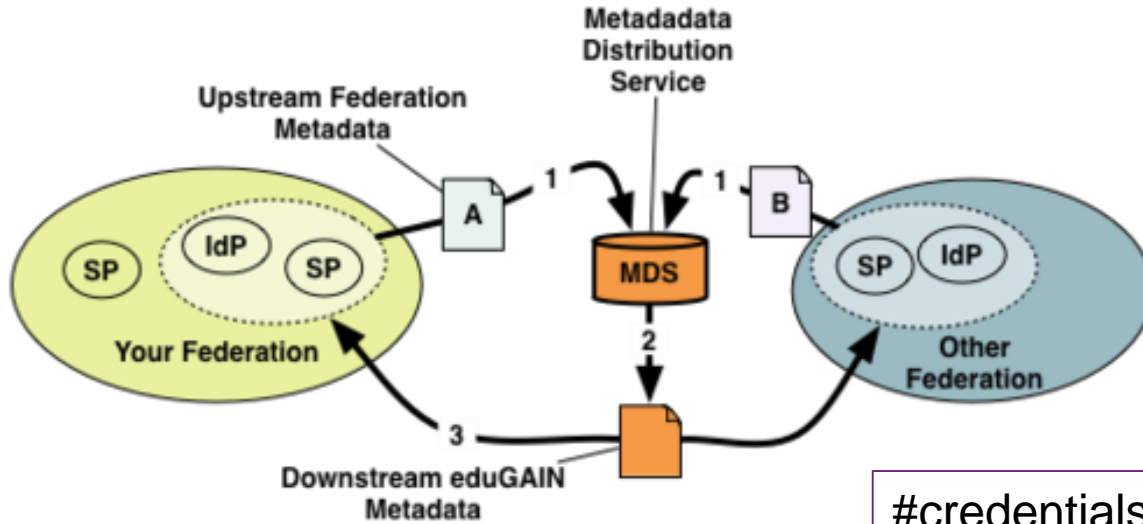
eduroam image from <https://eduroam.org/how/>, GEANT ; RADIUS: RC2865 <https://www.rfc-editor.org/rfc/rfc2865>; see also [freeradius.org](http://freeradius.org)



# We live in a federated world!



# Meta-data and trust in IdP-SP 'multi-lateral' federations




Listing of all entities (1872), 42 federations			
IdP	29 Mays University	EGCS	
1	Entity ID: <a href="https://kmlk.29mays.edu.tr/simplesaml/saml2/idp/metadata.php">https://kmlk.29mays.edu.tr/simplesaml/saml2/idp/metadata.php</a> Entity categories: Data Protection Code of Conduct v1, REFEDS Research and Scholarship Registrar: YETMIM Org: Istanbul 29 Mays University		<a href="#">Entity details</a>
IdP	A*STAR - Agency for Science, Technology and Research	EGCS	
2	Entity ID: <a href="https://a-star.singaren.net.sg/simplesaml/saml2/idp/metadata.php">https://a-star.singaren.net.sg/simplesaml/saml2/idp/metadata.php</a> Entity categories: REFEDS Research and Scholarship Registrar: Singapore Access Federation - SGAF Org: A*STAR - Agency for Science, Technology and Research		<a href="#">Entity details</a>
IdP	AAF Virtual Home	EGCS	
3	Entity ID: <a href="https://who.aaf.edu.au/idp/shibboleth">https://who.aaf.edu.au/idp/shibboleth</a> Entity categories: REFEDS Research and Scholarship SIRTFI: asserted Registrar: AAF Org: AAF Virtual Home		<a href="#">Entity details</a>
IdP	AARNet	EGCS	
4	Entity ID: <a href="https://shibboleth.aarnet.edu.au/idp/shibboleth">https://shibboleth.aarnet.edu.au/idp/shibboleth</a> Entity categories: REFEDS Research and Scholarship SIRTFI: asserted Registrar: AAF Org: Australian Academic and Research Network (AARNet)		<a href="#">Entity details</a>
IdP	ACCESS CI	EGCS	
5	Entity ID: <a href="https://access-ci.org/idp">https://access-ci.org/idp</a> Entity categories: REFEDS Research and Scholarship, <a href="http://id.incommon.org/category/registered-by-incommon">http://id.incommon.org/category/registered-by-incommon</a> SIRTFI: asserted Registrar: InCommon Org: National Center for Supercomputing Applications		<a href="#">Entity details</a>
IdP	ACOnet staff	EGCS	
6	Entity ID: <a href="https://idp.aco.net/idp/shibboleth">https://idp.aco.net/idp/shibboleth</a> Entity categories: Data Protection Code of Conduct v1, REFEDS Research and Scholarship Registrar: ACOnet Identity Federation Org: ACOnet staff		<a href="#">Entity details</a>
IdP	AIRCentre	EGCS	
	Entity ID: <a href="https://ido.aircentre.com/idp/shibboleth">https://ido.aircentre.com/idp/shibboleth</a>		

#credentials required?

$$\text{from } \mathcal{O}(n_{\text{users}}) + \mathcal{O}(n_{\text{services}} * n_{\text{home-orgs}})$$

$$\text{to } \sim \mathcal{O}(n_{\text{users}}) + \mathcal{O}(n_{\text{home-orgs}}) + \mathcal{O}(n_{\text{services}})$$

MDS meta-data flow: <https://wiki.geant.org/display/eduGAIN/Metadata+Flow+in+eduGAIN>  
 eduGAIN meta-data <https://mds.edugain.org/edugain-v2.xml> ; table excerpt from  
<https://technical.edugain.org/entities> showing only R&S IdPs, i.e. those supporting research ...



**78**

Identity Federations

**5100+**

Identity Providers

**3600+**

Service Providers

eduGAIN image: Davide Vagheti, GARR for GN\*-\*

# We progressed a lot since 2003 with identity federation

**NIKHEF**  
NATIONAAL INSTITUUT VOOR KERNFYSICA EN HOGE ENERGIEFYSICA

Guest / students form (please with a copy of your identity card)

CERN/User Registration Date: 01.03.2004

**CERN COMPUTER CENTRE - USER REGISTRATION FORM**  
<http://cern.ch/it/documents/ComputerUsage/CompAccountRegistrationForm-English.pdf>

To be returned to the User Registration box at the entrance of Building 513, after being completed by a user who requires a computer account in a Central Service provided by IT Department, and is not yet registered in another group or system or has already signed it before.

**To be completed by the User:**  
It is MANDATORY to provide the following information (except those with an \*). It will be treated confidentially and only be used for ensuring correct identification.  
Supply name as registered by the Users' Office or HR Division.  
FAMILY NAME(S): .....  
FIRST NAME(S): .....  
SEX [M] [F] BIRTHDATE: Day ..... Month ..... Year .....  
HOME INSTITUTE/FIRM: .....  
NATIONALITY: ..... \*CERN SUPERVISOR.....  
\*CERN DEPARTMENT: ..... \*CERN ID NUMBER (as on CERN card).....

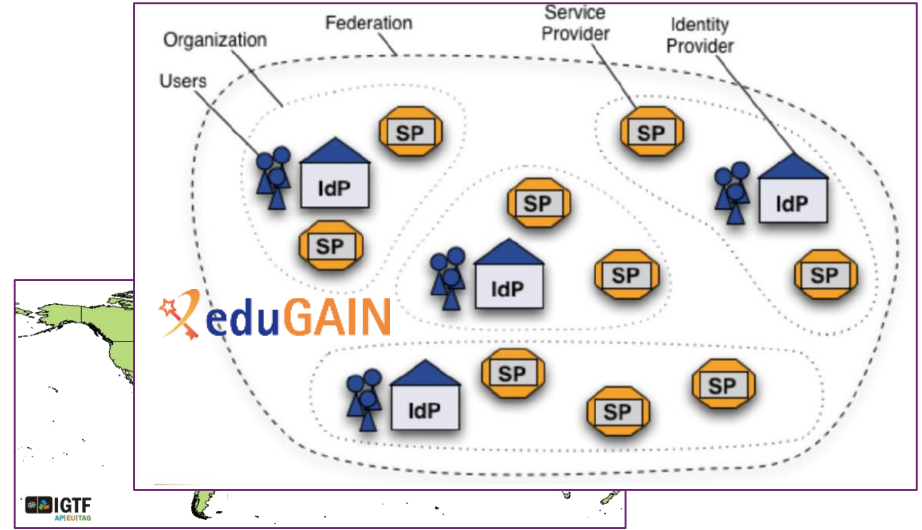
**To be completed by the Group Administrator:**

Exp / Dept.	Spokesperson	Home Institution Contact	Contact Telephone
D0	WOMERSLEY/WEERTS	SHARON HAGOPLAN	850-644-4777

SWITZER  
Last  
University or Institute  
FLORIDA STATE

Experiment/Department:  
Exp / Dept. Spokesperson Home Institution Contact Contact Telephone

4 5 6 7 8 9 0  
F1 F2 F3 F4 F5 F6  
CERN CLR ENT



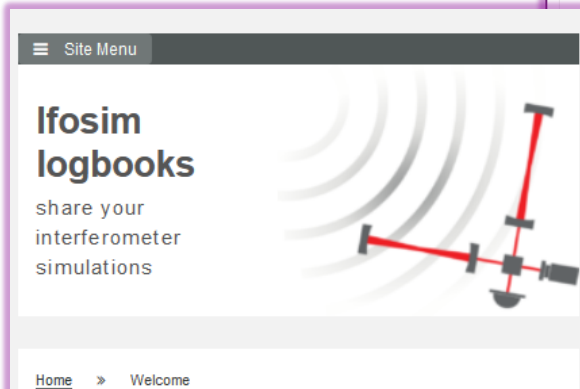
For eduGAIN federation the IdPs provide authentication from the home organisation, for the user-centric PKIX IGTF trust fabric, the CAs do. Then Service providers perform authorization, ... maybe using attributes provided by the IdP. But do they get them??

Right-hand image: Shibboleth IdP federation, Lukas Hammerle, SWITCH (CH), user-centric PKI credentials: Interoperable Global Trust Federation, <https://igtf.net/>

# Federated Success!

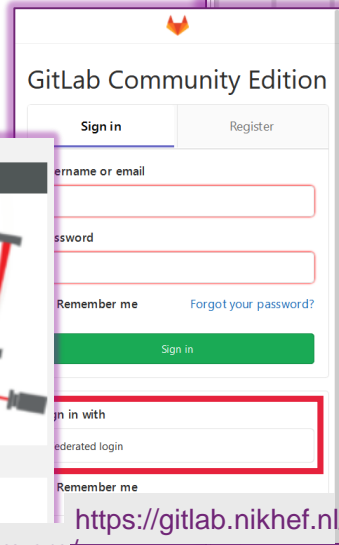
Login to GW's ifosim.org, to gitlab, or ... via the service proxy

*with any eduGAIN IdP for user authentication*



<https://logbooks.ifosim.org/>

ifosim federated AAI integration implementation by Mischa Sallé; per-country



<https://gitlab.nikhef.nl/>

<https://wayf.nikhef.nl/>

# Science infrastructures using our R&E 'federated access'

Large Hadron Collider

ALICE

ATLAS

LHCb

LHCb THCP

emso ERIC

EISCAT\_3D

OBSEA

EOSC-WeNMR portals


Marketplace

300+ Content providers

Spot On

WHISCY

3D-DART




Images: CERN <https://wlcg.web.cern.ch/>; HADDUCK, WeNMR, @Bonviniab <https://wenmr.science.uu.nl/>; Virgo, Pisa, IT; artist impression Einstein Telescope EMR region; EOSC portal in 2023, EGI catalogue <https://www.egi.eu/>

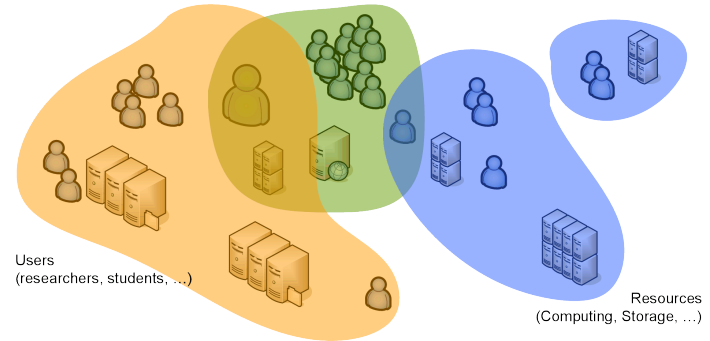


# They look similar, yet they are not ...

In the **Identity federation** picture, the source of authority is the *home organisation* via its IdP



In the **Community** picture, the source of authority is *the community itself*

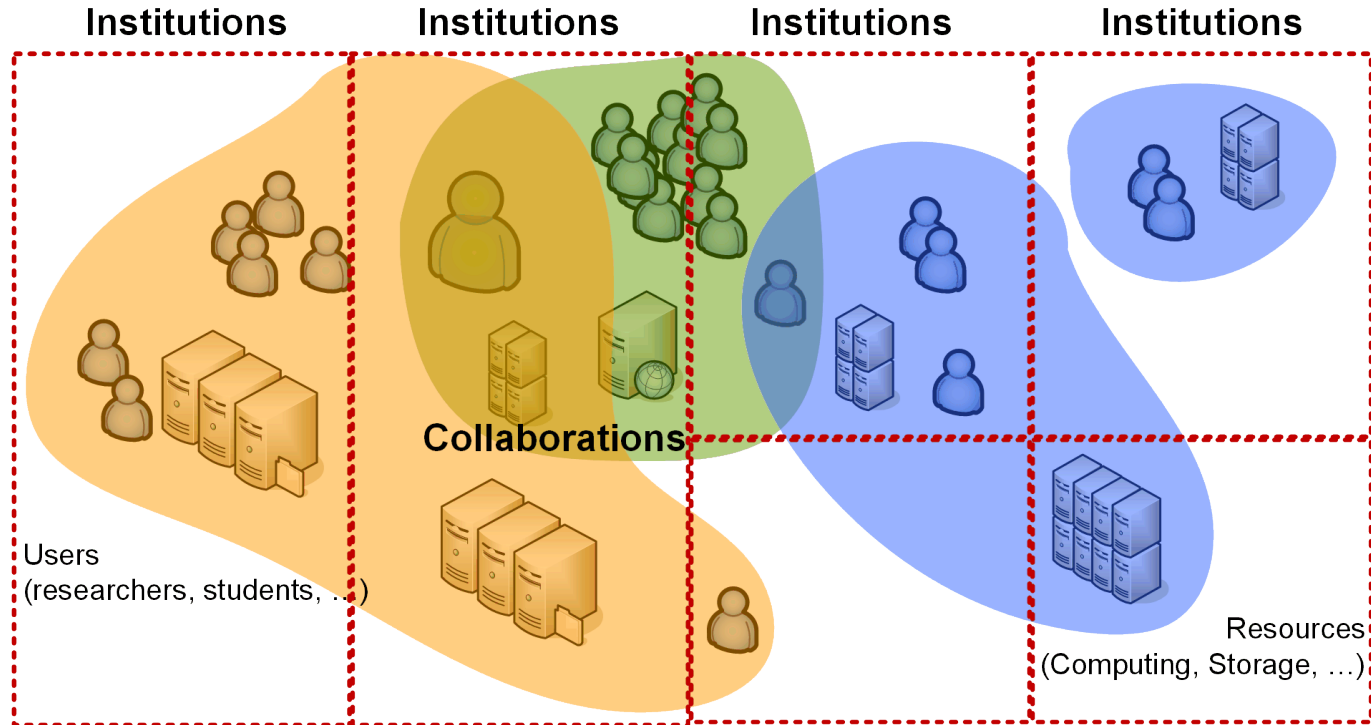


**the AuthN-AuthZ separation is fundamental**

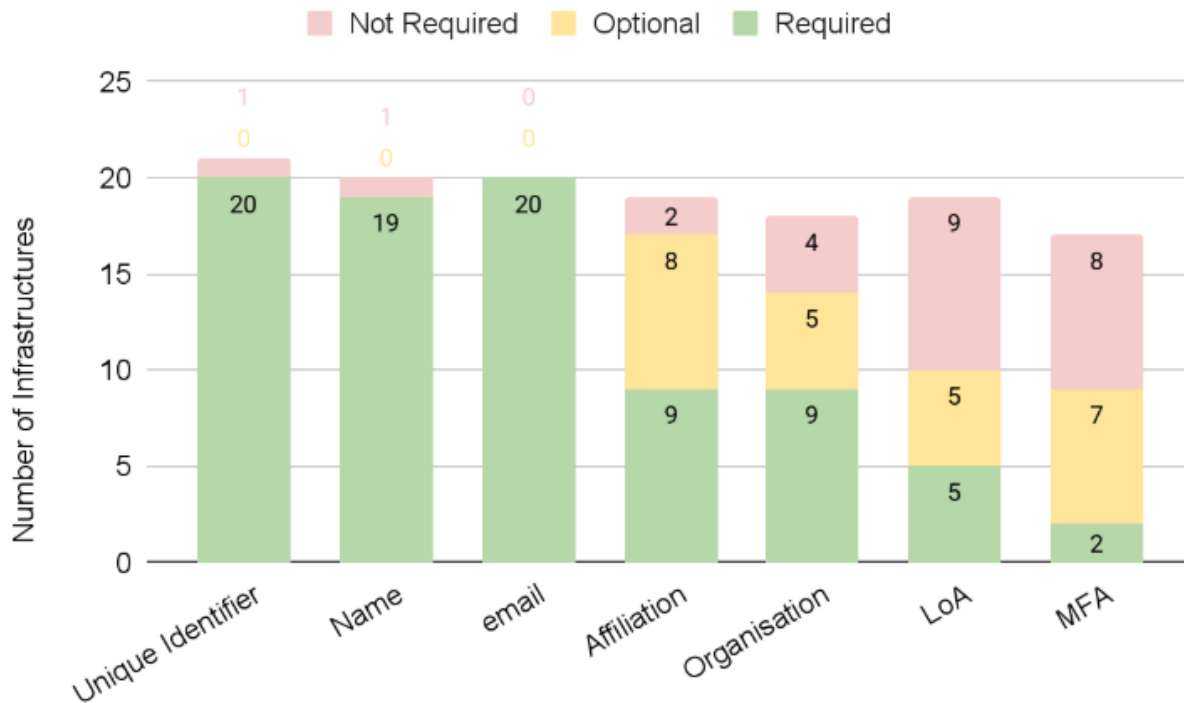
to the Federated (R&E) AAI, global IGTF PKI, VOMS, 'AARC BPA' AAI architecture ...

Right-hand image: Shibboleth IdP federation, Lukas Hammerle, SWITCH (CH)

# Since collaborations and institutions slice in different ways



# Research Infrastructures: what they *actually* need from ‘home’



## Glossary

Affiliation: what *type* of entity are you (student, faculty, alumnus, ...)

LoA: level of authentication assurance (like passport identity vetting and ‘freshness’ of data)

MFA: multi-factor authentication (password, 6-digit code, SMS, fingerprint)

Source: Marina Adomeit, Janos Mohasci, *et al.* AARC TREE Use-case collection and analysis (D3.2), 2025 (under review)

The one infra that did ‘not need a unique identifier’ actually stated: “<our infra> assigns own identifier upon registration” – so the unique identifier is *still* there!

# For starters: sharing good user identifiers is non-trivial 😞



of 6019 identity providers  
in 77 federations,  
only 1994 support R&S or Personalised access **33%**

## For IdP Operators

### What attributes should be released by an R&S IdP?

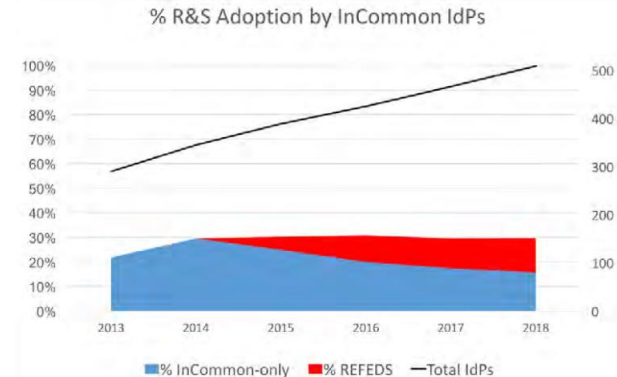
The Research & Scholarship specification defines a bundles of attributes that Identity Providers are encouraged to release to R&S services:

- personal identifiers: email address, person name, eduPersonPrincipalName
- pseudonymous identifier: eduPersonTargetedID
- affiliation: eduPersonScopedAffiliation

Category support is defined as follows:


An Identity Provider indicates support for the R&S Category by exhibiting the R&S entity attribute in its metadata. Such an Identity Provider **MUST**, for a significant subset of its user population, release all required attributes in

~ constant since 2018 😞

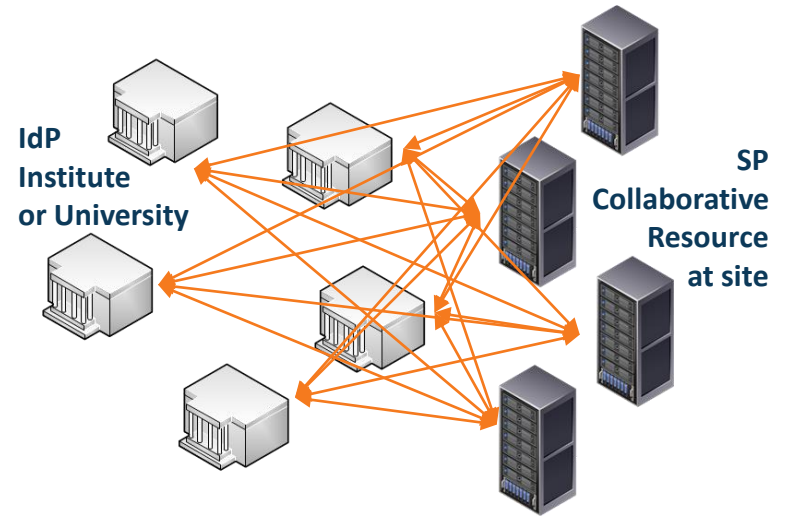


Graph: InCommon: Attributes-WG-Recommendations-May2018.pdf; Entity Category stats as per 2025-03-03, from <https://technical.edugain.org/entities>

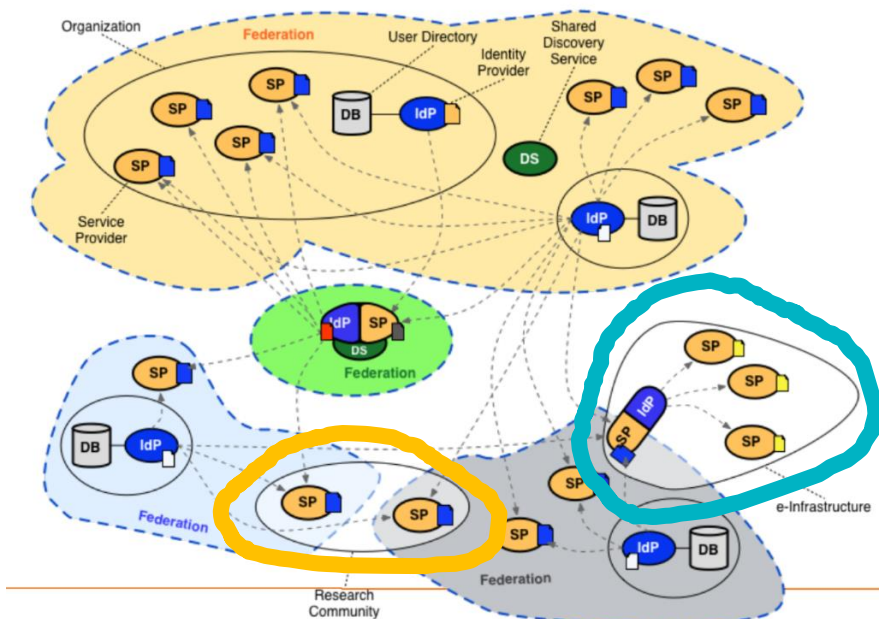
# A fundamental scaling issue remained unique to research



for identity and user data  
'n x m' agreements remain(ed)

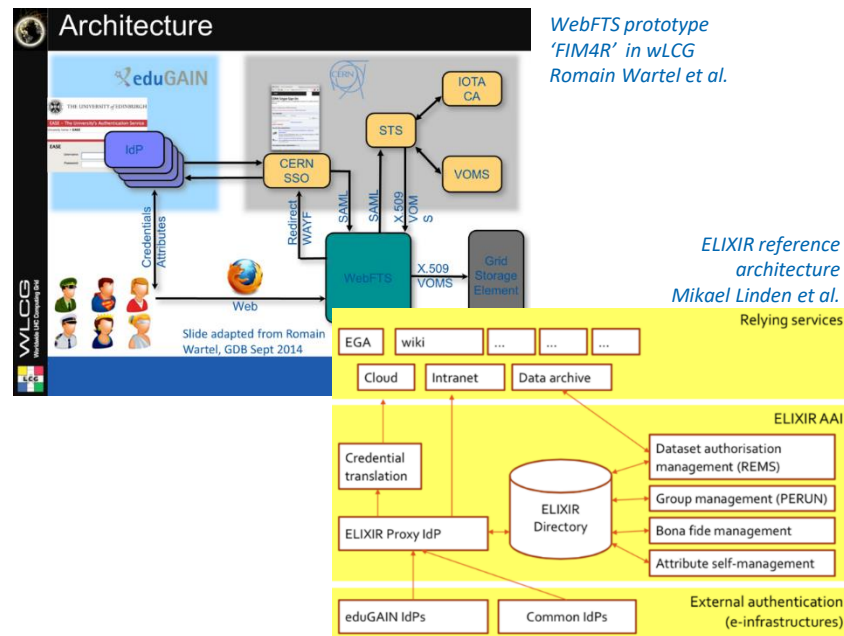


# Managing complexity: distributed diverse identity sources



they were composed of many services  
each of which had to manage federation complexity

Community images: Romain Wartel, CERN; Mikael Linden, CSC; Federation image (R): Lukas Hammerle, SWITCH




but most communities had started to invent  
their own 'proxy' model to abstract complexity



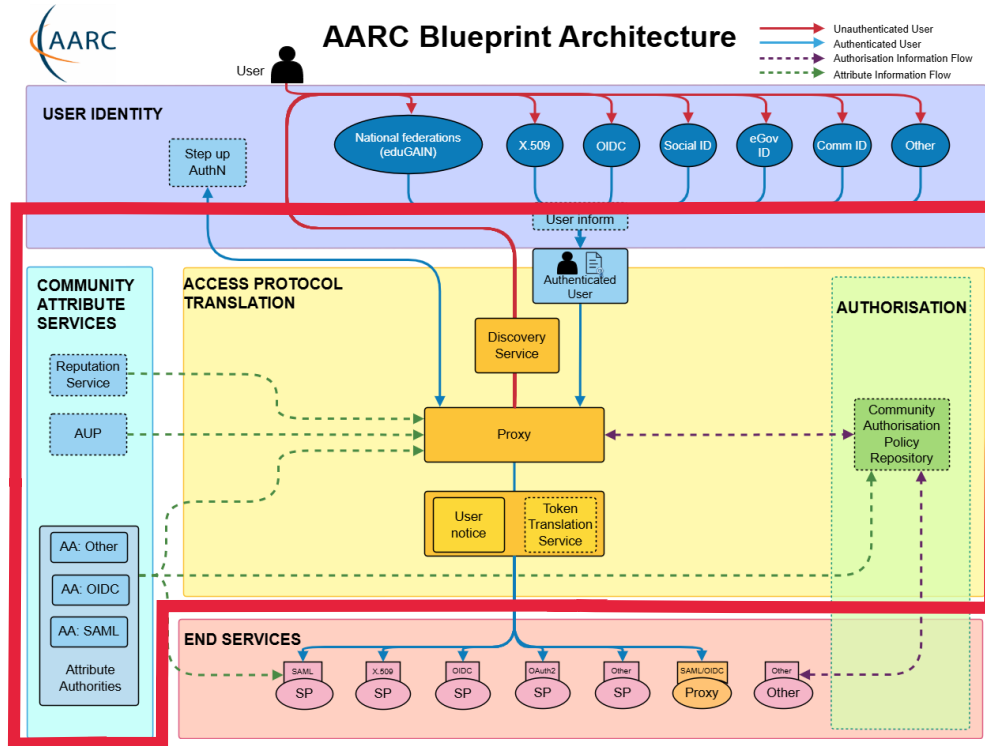
# The IdP-SP bridge

- Access services using **identities from users' Home Organizations**,
- but **hide complexity** of multiple IdPs, federations, and different technologies for authentication and authorisation
- **One persistent identity** across all the community's services through **account linking**
- **Access services based on role(s)** users have **in the collaboration.**
- For both **web** and **non-web** resources
- Integration of **guest identity solutions**
- **Support for stronger authentication assurance** mechanisms

*often known as proxy!*



# AARC Blueprint – making the bridge a first-class citizen



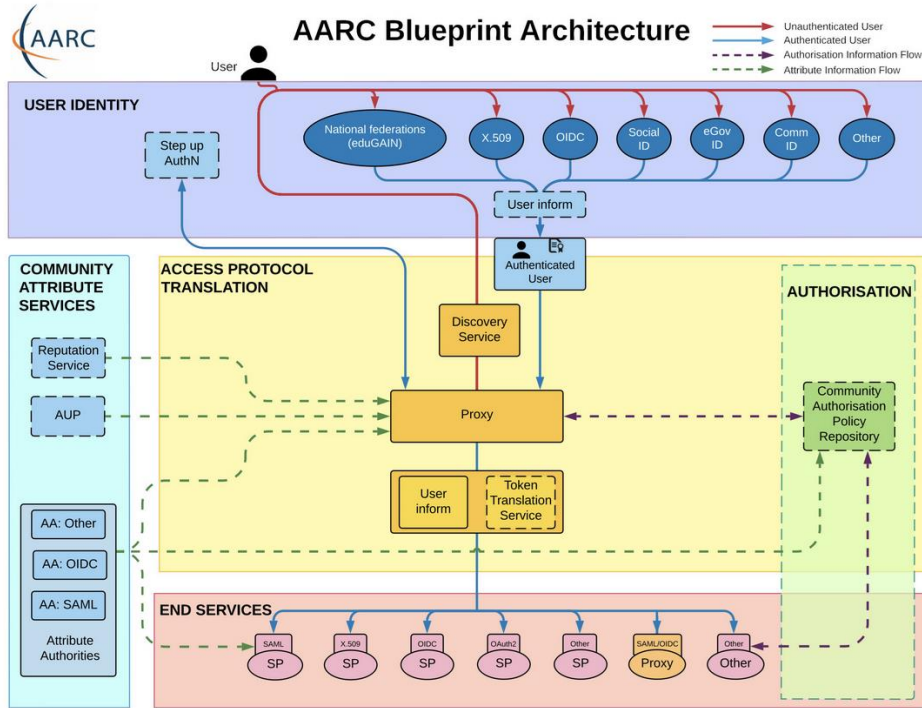
Manage users and access rights

with interoperable building blocks for 'AAI infrastructure' architects

that are

- technology-agnostic
- have multiple implementations
- come with policy templates & good practice guides

# More than just nice colours



<https://aarc-community.org/guidelines/>

Not sure how to begin with the AARC Blueprint Architecture? There are plenty of guidelines available but it can be a minefield at first. You probably want to start by designing the high level approach of your infrastructure based on the AARC Blueprint Architecture. There are several general topics you should consider, such as Data Protection (AARC-G042) and Federated Security Incident Response (AARC-ID51). Here you can find common questions matched to the relevant Blueprint Architecture component, along with links to guidelines that can help.

### Community Attribute Services:

- How should attributes from multiple sources be aggregated? AARC-G003
- How should I express the home institute of a user? AARC-G025
- How should I express the identifier of a user AARC-G026
- What are the best practices for running my Attribute Authorities securely? AARC-G071
- Which Acceptable Use Policy should I use to facilitate interoperability? AARC-ID44
- How should I infer the affiliation of a user? AARC-G057

### Authorisation:

- How should I manage authorisation information from multiple sources? AARC-G006
- How should group and role information be expressed to facilitate interoperability? AARC-G002
- How should resource capabilities be expressed? AARC-G027

### End Services:

- My service needs to act on behalf of the user – how should I handle credential delegation and impersonation? AARC-G005
- My services are not web based, how can I use identities from the proxy? AARC-G007
- How should Services hint which IP they would like users to use? AARC-G049
- Which Security practices should I follow? AARC-G014

### User Identity:

- How should I integrate Social Media Identity Providers? AARC-G008
- How should users link accounts, and how does that affect Assurance? AARC-G009
- How should services indicate that they would like users to authenticate with multifactor authentication, and how should my proxy forward that information? AARC-G029

### Assurance:

- How should assurance information of external identities be calculated? AARC-G031
- What can I say about assurance of identities from social media accounts? AARC-G041
- How is assurance impacted by account linking? AARC-G009
- How is assurance information be shared with other infrastructures? AARC-G021
- Which Assurance Profiles should I use, there are so many? AARC-ID50

### Access Protocol Translation:

- Which best practices should I follow for my Token Translation Services? AARC-G004
- How should I translate from Identity Federation information to X.509 certificates? AARC-G010

### Proxies:

- How can I ensure that my proxy is able to accurately claim that it supports best practices in Identity Federation? AARC-G015
- How should I express the home institute of a user? AARC-G025
- How should I express the identifier of a user AARC-G026
- How should I express assurance information for users when interacting with another proxy? AARC-G021
- How can my proxy simplify the discovery process for end-users? AARC-G061
- How can my proxy route the user to the correct discovery service? AARC-G062

What next? Are you looking for a kick start with your policies? Take a look at the **Policy Development Toolkit** which provides a set of templates:

Service	Infrastructure	Services (abide by)	Policy description	Example
Operations Security Policy	Infrastructure Management	Services (abide by)	This policy defines requirements for running a service within the infrastructure.	Google Doc
Acceptable Use Policy	Infrastructure Management (for baseline) & Research Communities (for community specific restrictions)	Users (abide by)	This is a template for the acceptable use policy that users must accept when using Research infrastructure. It is augmented by the Research...	Google Doc

Showing 1 to 9 of 9 entries Previous Next

# “Your attention to detail is appreciated”

## Even a simple challenge ...

“How to communicate affiliation of a user with the community”

## needs standards for interoperability!

- [AARC-G025 – Guidelines for expressing affiliation information](#)
- [AARC-G057 – Inferring and constructing voPersonExternalAffiliation](#)

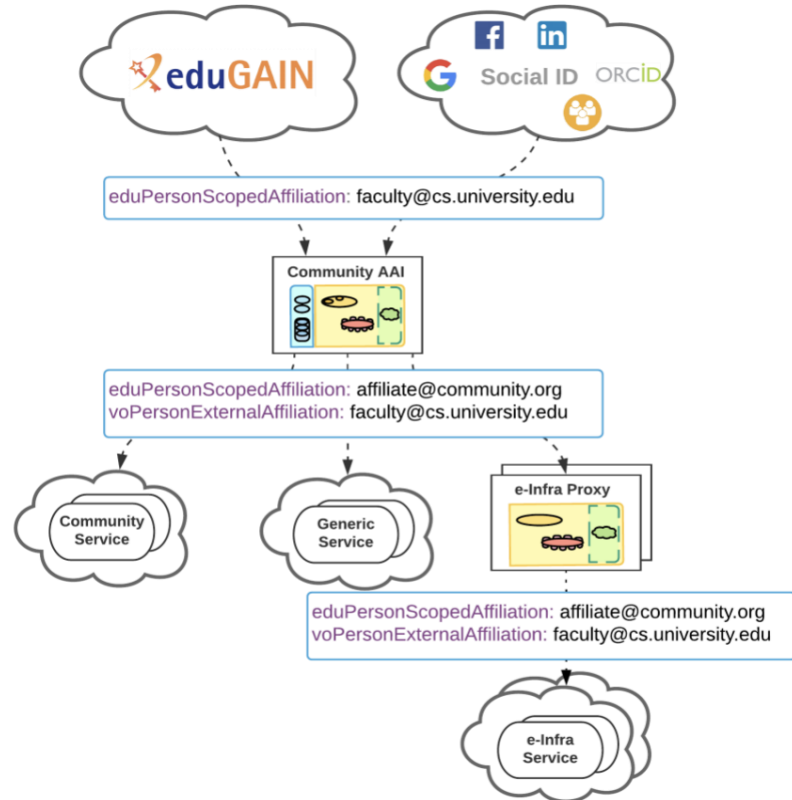
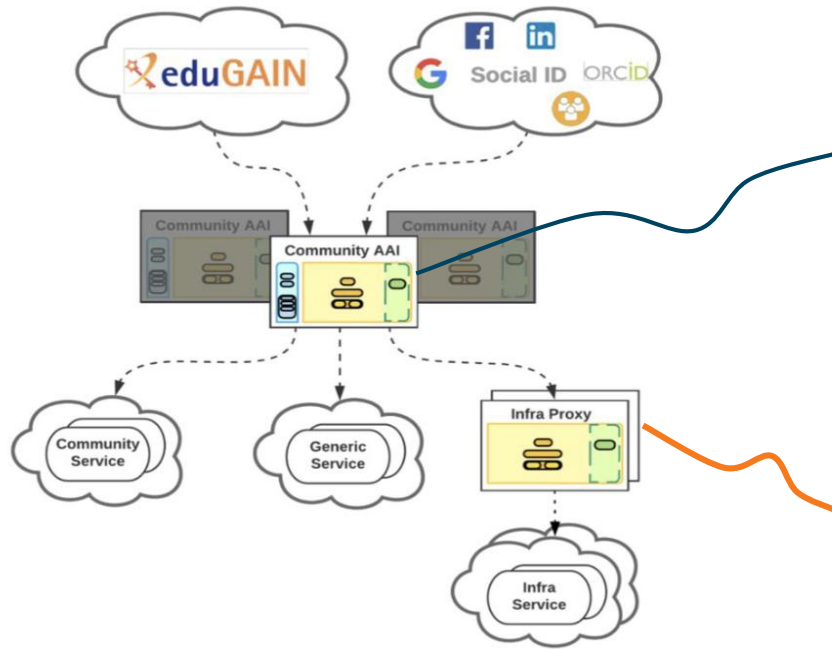


Image: Guideline AARC-G025 (AARC community); quote from the MoinMoin wiki software



# ... but one proxy is not enough in a research cloud



## Community AAI

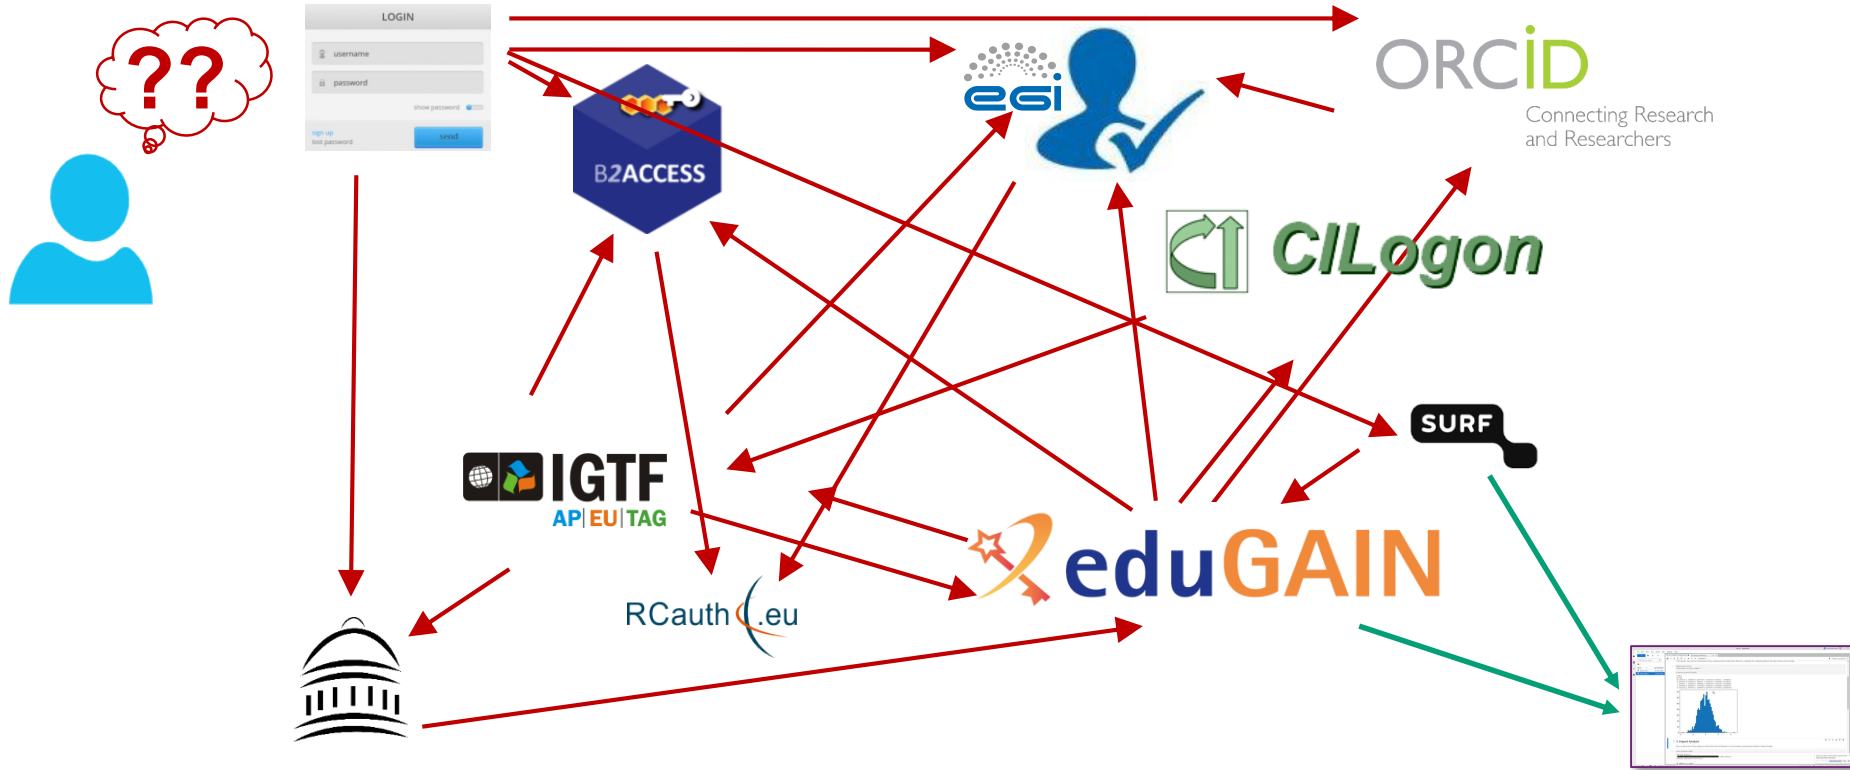
streamline researchers' access to services, both those provided by their own infrastructure as well as the services provided by shared infrastructures from other communities.

## Infrastructure Proxy

enables Infrastructures with large number of resources, to provide them through a single integration point, where the Infrastructure can maintain centrally all the relevant Policies and business logic for making available resources to multiple communities

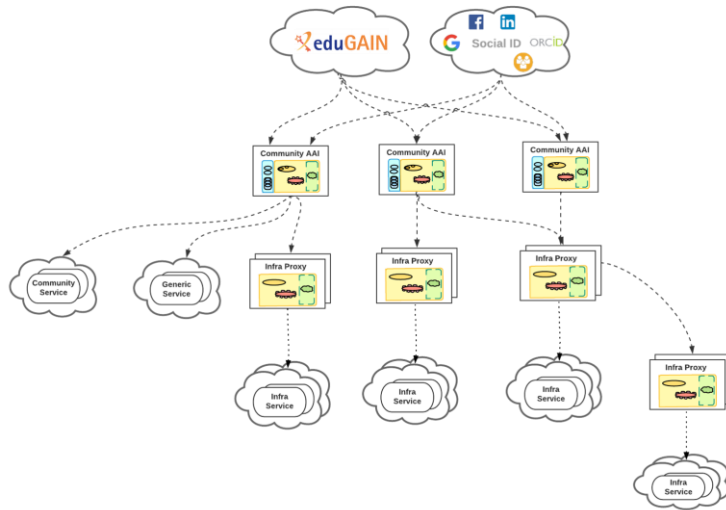


# Identity spaghetti: 1-loop, 2-loop and higher order diagrams




# We have seen many arrows before ... it needs federation!

Identity, community, infrastructure proxies and services form a **federation of proxies**



- bilateral registration  
*but then you have a scalability issue again*
- meta-data distribution of trust paths
  - **OpenID Federation**
  - **SAML meta-data**
- discovery and identity provider hinting



# European Open Science Cloud federation (2023 edition)

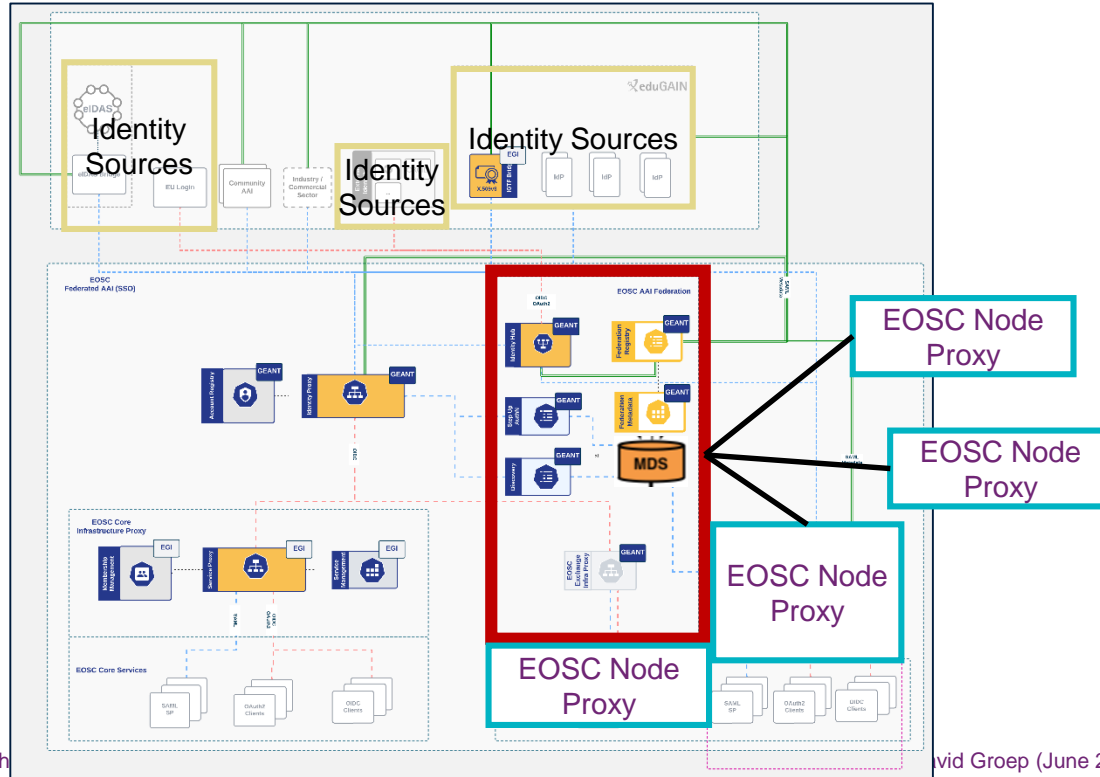
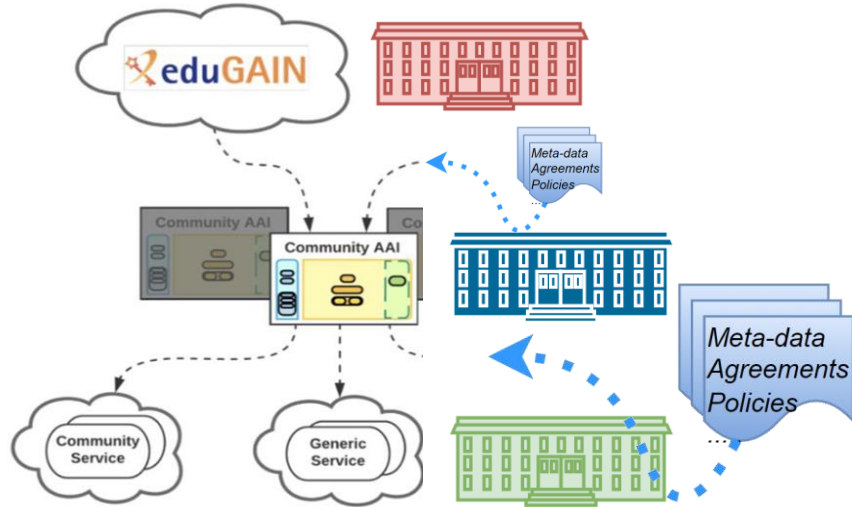


Image: EOSC AAI for the EOSC Core and Exchange

vid Groep (June 2023)

# And we need to 'decorate' the arrows with trust



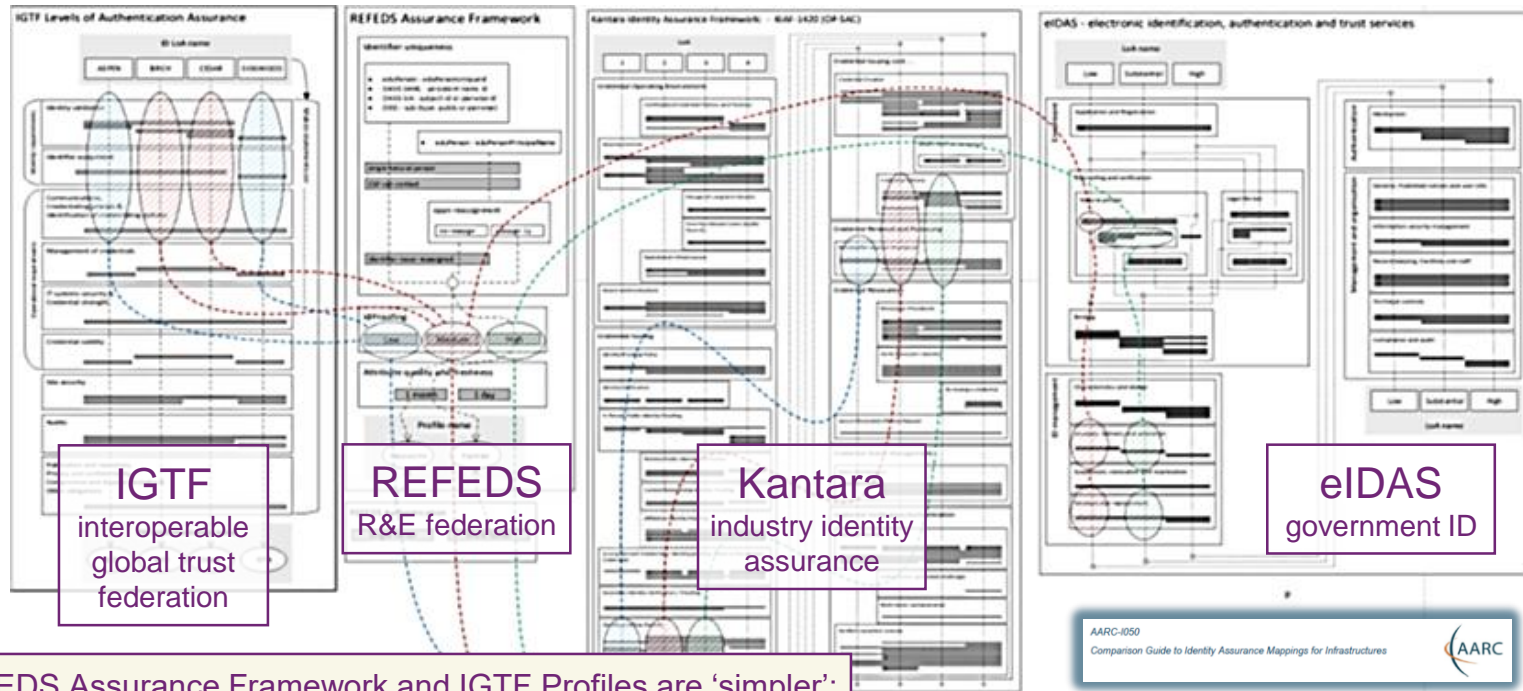
Each side of each arrow has *independent* parties

- we allow *them* to do part of the work we would otherwise do
- to make it easier and faster for users to perform their research
- but **we relinquish some control** beyond our organisation, our own policies, our own jurisdiction

**Why would we trust them to do that?**



# And even a simple ‘Who are you?’ is not always easy ...




REFEDS Assurance Framework and IGTF Profiles are ‘simpler’: academia is a higher-trust environment, leveraging self-assessed peer review

Source: <https://aarc-community.org/guidelines/aarc-i050>, Ian Neilson et al.



# Helping community and users: how much clicking through?



EUROPEAN COMMISSION  
DG COMMUNICATIONS NETWORKS, CONTENT AND TECHNOLOGY  
Directorate C - Enabling and Emerging Technologies  
Unit C.1 - High Performance Computing and Applications

## EOSC EU Node User Access Policy

Version 1.0

USER ACCESS POLICY

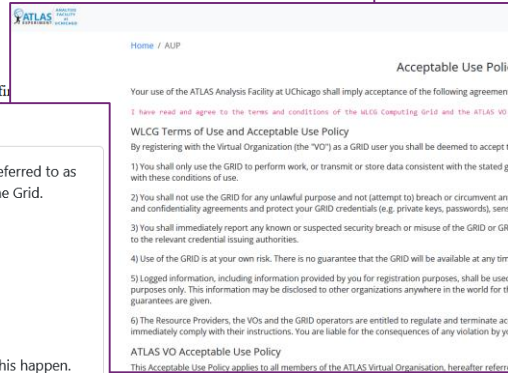
### 1. Purpose

This User Access Policy ("UAP") defines

### Acceptable Use Policy

text This Acceptable Use Policy applies to all members of Xenon Virtual Organisation, hereafter referred to as the VO, with reference to use of the European Grid Infrastructure (EGI), hereafter referred to as the Grid. The BiG Grid Executive Team owns and gives authority to this policy. Goal and description of the Xenon VO

The Xenon VO xenon.biggrid.nl is the incubator grid community for work on the international Xenon 1T and related experiments in the search for dark matter. Members of the VO will work to build, understand and analyse the detector and results related to the Xenon experiment and to "Monte-Carlo" studies that will be used to design, build and understand the detector, as well as work with the supporting computing infrastructure to make this happen. Members and Managers of the VO agree to be bound by the Grid Acceptable Usage Rules, VO Security Policy and other relevant Grid Policies, and to use the Grid only in the furtherance of the stated goal of the VO.



Home / AUP

### Acceptable Use Policy

Your use of the ATLAS Analysis Facility at UChicago shall imply acceptance of the following agreement:

I have read and agree to the terms and conditions of the WLCG Computing Grid and the ATLAS VO WLCG Terms of Use and Acceptable Use Policy

By registering with the Virtual Organization (the "VO") as a GRID user you shall be deemed to accept the following terms and conditions of use:

- 1) You shall only use the GRID to perform work, or transmit or store data consistent with the stated goals of the VO.
- 2) You shall not use the GRID for any unlawful purpose and not (attempt to) breach or circumvent any and confidentiality agreements and protect your GRID credentials (e.g. private keys, passwords), sensitive information, and other confidential information.
- 3) You shall immediately report any known or suspected security breach or misuse of the GRID or GRID resources to the relevant credential issuing authorities.
- 4) Use of the GRID is at your own risk. There is no guarantee that the GRID will be available at any time.
- 5) Logged information, including information provided by you for registration purposes, shall be used for purposes only. This information may be disclosed to other organizations anywhere in the world for the purposes stated in the policy.
- 6) The Resource Providers, the VOs and the GRID operators are entitled to regulate and terminate access to the GRID at any time. You are liable for the consequences of any violation by you.

ATLAS VO Acceptable Use Policy  
This Acceptable Use Policy applies to all members of the ATLAS Virtual Organisation, hereafter referred to as the VO.

## EGI Configuration Database Acceptable Use Policy and Conditions of Use (AUP)

This Acceptable Use Policy and Conditions of Use ("AUP") defines the rules and conditions that govern your access to and use of the EGI Configuration Database (EGIDB) (including transmission, processing, and storage of data) of the resources and services ("Services") as granted by the EGI Federation, and the Virtual Organisation to which you belong, for the purpose of meeting the goals of EGI, namely to deliver advanced computing services to support researchers, multinational projects and research infrastructures, and the goals of your Virtual Organisation or Research Community.

1. You shall only use the Services in a manner consistent with the purposes and limitations described above; you shall show consideration towards other users including by not causing harm to the Services; you have an obligation to collaborate in the resolution of issues arising from your use of the Services.
2. You shall only use the Services for lawful purposes and not breach, attempt to breach, nor circumvent administrative or security controls.
3. You shall respect intellectual property and confidentiality agreements.
4. You shall protect your access credentials (e.g. passwords, private keys or multi-factor tokens); no intentional sharing is permitted.
5. You shall keep your registered information correct and up to date.
6. You shall promptly report known or suspected security breaches, credential compromise, or misuse to the security contact stated below; and report any compromised credentials to the relevant issuing authorities.
7. Reliance on the Services shall only be to the extent specified by any applicable service level agreements listed below. Use without such agreements is at your own risk.
8. Your personal data will be processed in accordance with the privacy statements referenced below.
9. Your use of the Services may be restricted or suspended, for administrative, operational, or security reasons, without prior notice and without compensation.
10. If you violate these rules, you may be liable for the consequences, which may include your account being suspended and a report being made to your home organisation or to law enforcement.

The administrative contact for this AUP is: [operations@egi.eu](mailto:operations@egi.eu)

The security contact for this AUP is: [abuse@egi.eu](mailto:abuse@egi.eu)

The privacy notice is located at <https://gocdb-preprod.egi.eu/privacy.html>.

[Return to GOCDB homepage.](#)

# Good common practice: the WISE Baseline AUP



## Acceptable Use Policy and Conditions of Use

This Acceptable Use Policy and Conditions of Use ("AUP") defines the rules and conditions that govern your access to and use (including transmission, processing, and storage of data) of the resources and services ("Services") as granted by {community, agency, or infrastructure name} for the purpose of {describe the stated goals and policies governing the intended use}.

<To further define and limit what constitutes acceptable use, the community, agency, or infrastructure may optionally add additional information, rules or conditions, or references thereto, here or at the placeholder below. These additions must not conflict with the clauses 1-10 below, whose wording and numbering must not be changed.>

1. You shall only use the Services in a manner consistent with the purposes and limitations described above; you shall show consideration towards other users including by not causing harm to the Services; you have an obligation to collaborate in the resolution of issues arising from your use of the Services.
2. You shall only use the Services for lawful purposes and not breach, attempt to breach, nor circumvent administrative or security controls.
3. You shall respect intellectual property and confidentiality agreements.
4. You shall protect your access credentials (e.g. passwords, private keys or multi-factor tokens); no intentional sharing is permitted.
5. You shall keep your registered information correct and up to date.
6. You shall promptly report known or suspected security breaches, credential compromise, or misuse to the security contact stated below; and report any compromised credentials to the relevant issuing authorities.
7. Reliance on the Services shall only be to the extent specified by any applicable service level agreements listed below. Use without such agreements is at your own risk.
8. Your personal data will be processed in accordance with the privacy statements referenced below.
9. Your use of the Services may be restricted or suspended, for administrative, operational, or security reasons, without prior notice and without compensation.
10. If you violate these rules, you may be liable for the consequences, which may include your account being suspended and a report being made to your home organisation or to law enforcement.

<Insert additional numbered clauses here>

The administrative contact for this AUP is:

{email address for the community, agency, or infrastructure name}

The security contact for this AUP is:

{email address for the community, agency, or infrastructure security contact}

The privacy statements (e.g. Privacy Notices) are located at: {URL}

Applicable service level agreements are located at: <URLs>

## Purpose binding

ensure use is as intended for access grant

## Terms and Conditions

research data access conditions, permits, grant conditions

## WISE Baseline AUP


common 10 commandments that allow seamless cross-sectoral user movement

## Service level agreements

promises and recourse

## Privacy notice references



for access personal data policies



<https://wise-community.org/wise-baseline-aup/>



It all started here!



# AUP – The Taipei Accord

- (1) You may only perform work, or transmit or store data consistent with the activities and policies of the Virtual Organizations of which you are a member, and only on resources authorized for use by those Virtual Organizations.
- (2) You will not attempt to circumvent administrative or security controls on the use of resources. If you are informed that some aspect of your grid usage is creating a problem, you will adjust your usage and investigate ways to resolve the complaint.
- (3) You will immediately report any suspected compromise of your grid credentials or suspected misuse of grid resources to incident reporting locations specified by the Virtual Organization(s) affected and credential issuing authorities as specified in their agreements and policy statements.
- (4) You are aware that resource providers have the right to regulate access as they deem necessary for either operational or security-related reasons and that your use of the Grid is also bound by the rules and policies of the organizations through which you obtain access, e. g. your home institute, your national network and/or your internet service provider(s).

29 Apr 05      OSG Acceptable Use and Incident Response      8

Bob Cowles *Acceptable Use Policy and Security Incident Response Strategy in the Open Science Grid* – ISGC, 29 April 2005

# What about *unacceptable* use? “who dunnit” essential for incident response, but *what* have we just built?



So we have federation and single sign-on ...  
... but can we respond if something goes haywire?  
... can we share security incident information when needed?  
... timely and confidentially, protecting everyone’s reputation?

left: eduGAIN inter federation in 2025 (<https://technical.edugain.org/status>); logos on the right from the European e-Infrastructures and ESFRIs; center graphic: AARC

# 'Sirtfi' – what makes federated security different?




Organisations probably do 'something reasonable' for their own security ... but may not realise the implications for others

**Sirtfi** targets coordinated **response in a federated context**:


1. Enable **communication** and coordination in managing federated security incidents
2. Relevant **event data** is available to help collaborating incident responders.
3. **Security protections are applied** to federated transactions

Define capabilities for security incident response an IdP or SP **organisation can self-asserts** in federation meta-data

<https://refeds.org/sirtfi>



- [IR3] Notify security contacts of entities participating in Sirtfi when a security incident investigation suggests that those entities are involved in the incident. Notification should also follow the security procedures of any federations to which your organisation belongs.



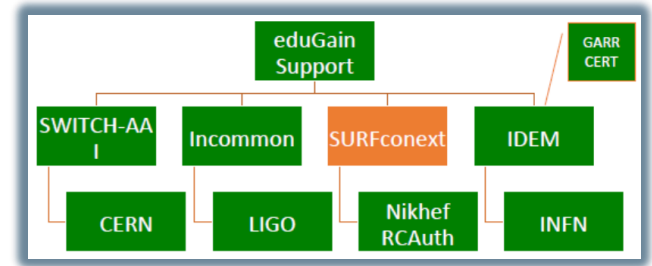
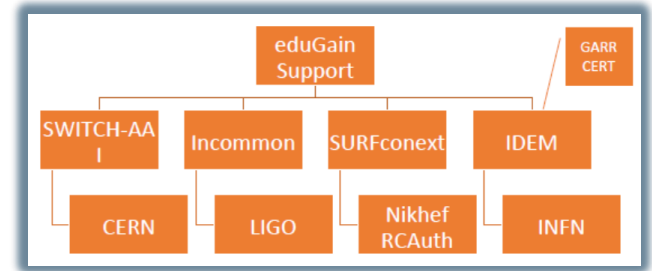
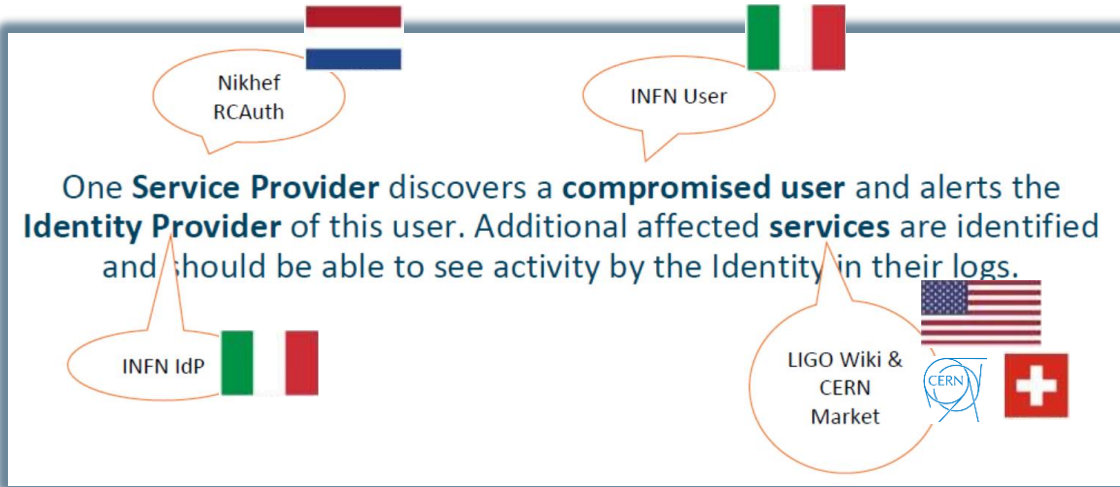
- Table of Contents
- Operational Security
  - Incident Response
  - Tracability
  - User Rules & Conditions





# A federated community security challenge

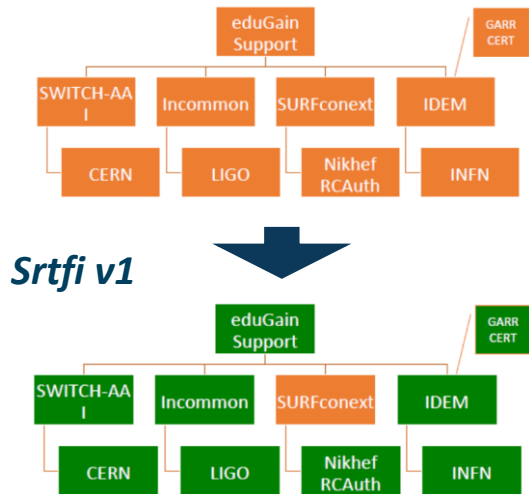
Can we coordinate our collective R&E response?  
'security challenges' based on the *Sirtfi* contact model



parties involved in response challenge

Report-outs see <https://wiki.geant.org/display/AARC/Sirtfi+Communications+Challenges%2C+AARC2-TNA3.1>

# Response across IdP-SP Proxies: the limits of Sirtfi version 1



Default Fed as proxy

**Fed-1**

No direct Coordinating team - Participant Communication




Support request situation

**Fed-2**

Direct Coordinating team - end entity Communication

Fed is always in CC

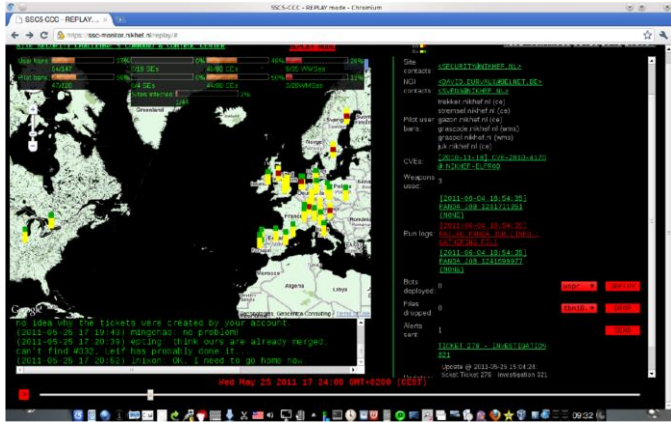


joint work with GN5 EnCo and eduGAIN CSIRT


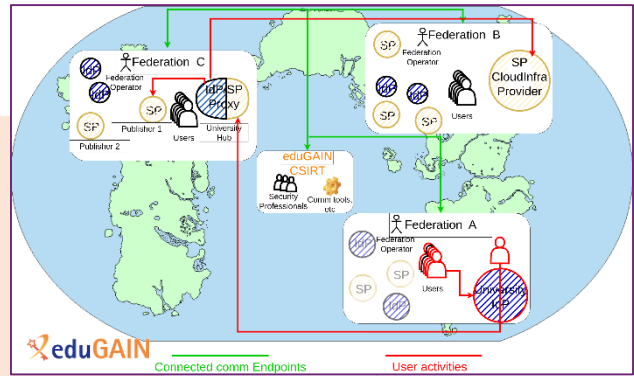




# Trust ... but verify



- Communication:
- Endpoints valid?
  - Form/Content OK ?
- Containment
- Ban "malicious" users
  - Find/Stop malicious processes
  - Find submission IP
- Forensics
- Basic Forensics on binary
  - Network traffic



Nikhef

## Nikhef CSIRT Traceability Challenge

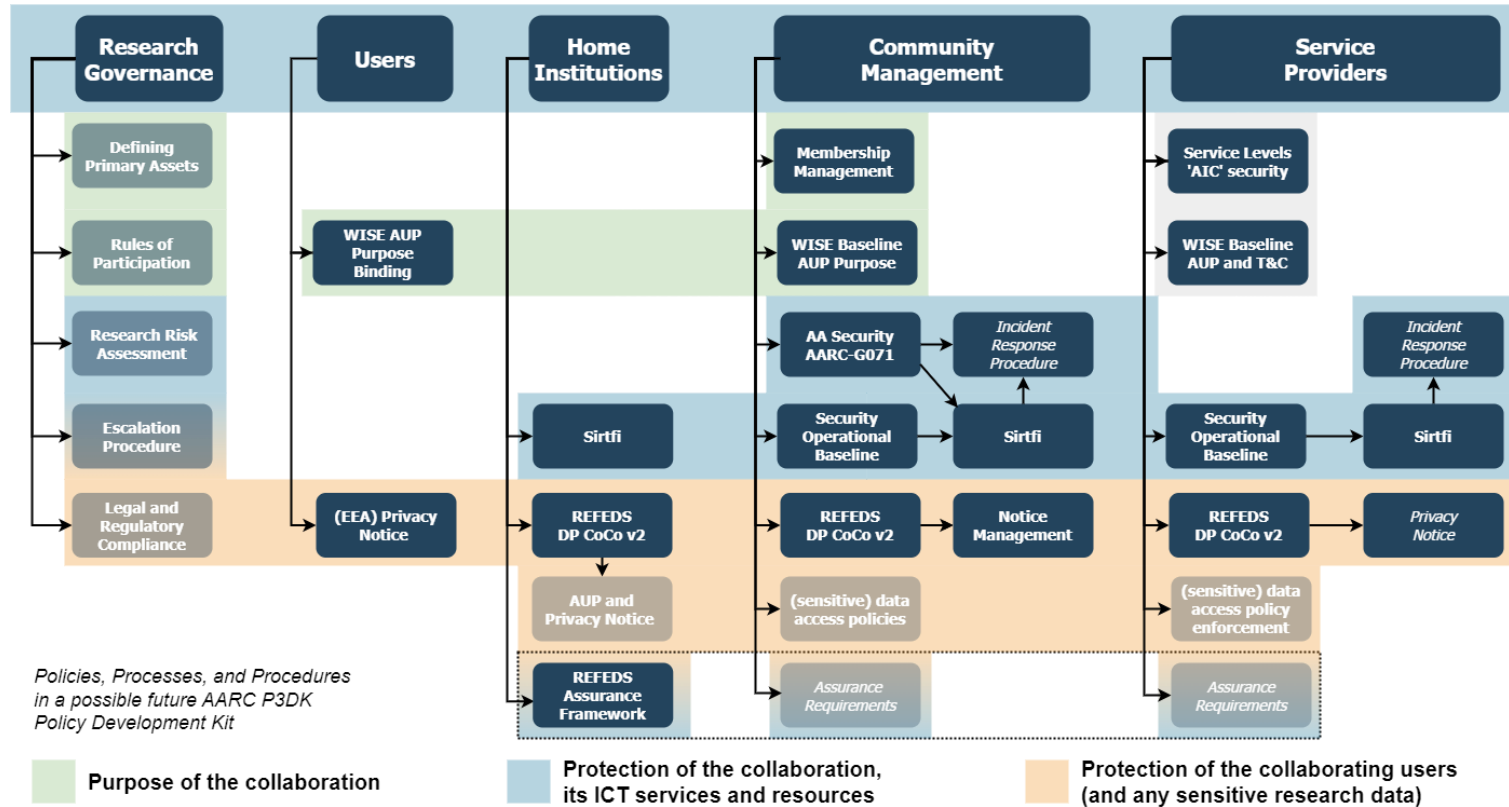
### introduction

Deze Traceability Challenge bestaat uit drie onderdelen, in (naar verwachting) oplopende moeilijkheidsgraad. Iedere challenge begint met een externe 'trigger' - aan het eind van dit document staan de hints en de goede (of in ieder geval: de 'gewenste') oplossing.

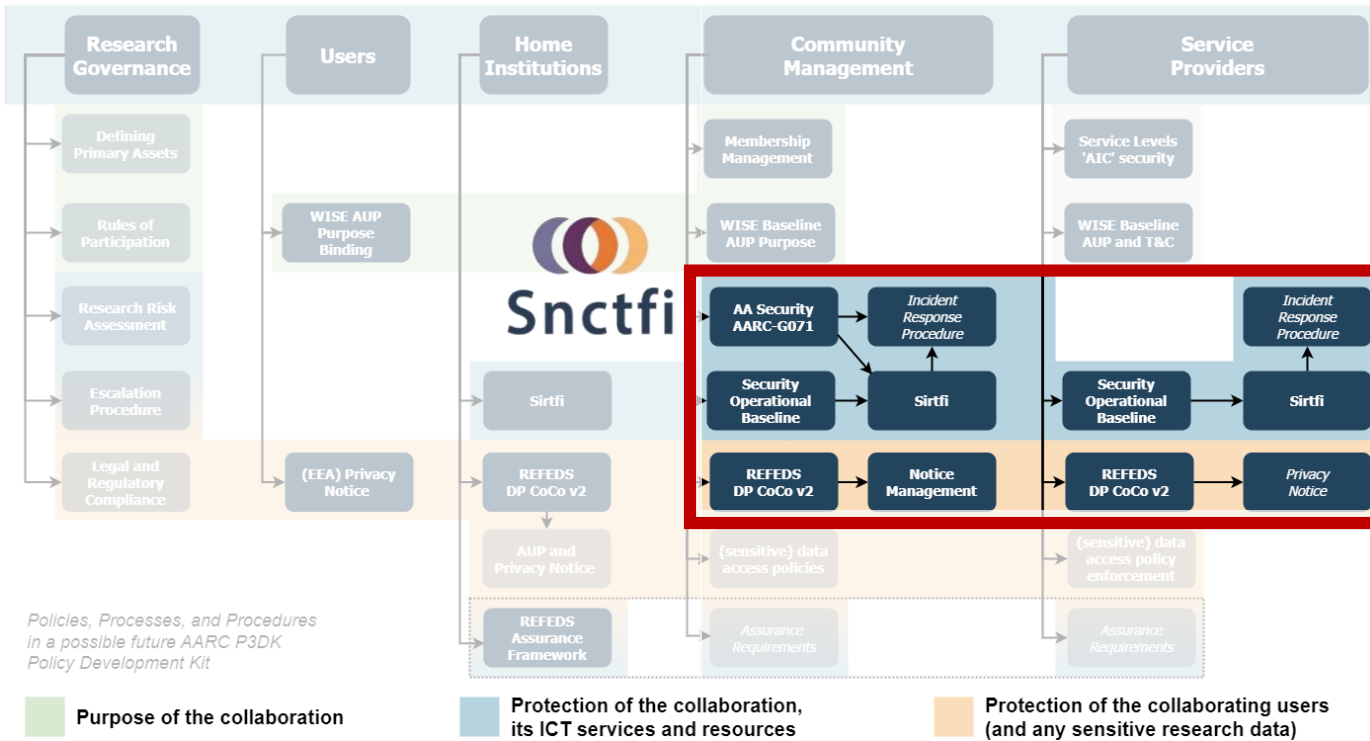
Veel plezier!



# Policy Development Kit: simplify by re-using good practice



# Providers manage complexity for research communities



communities sourcing 'well-operated' community platforms

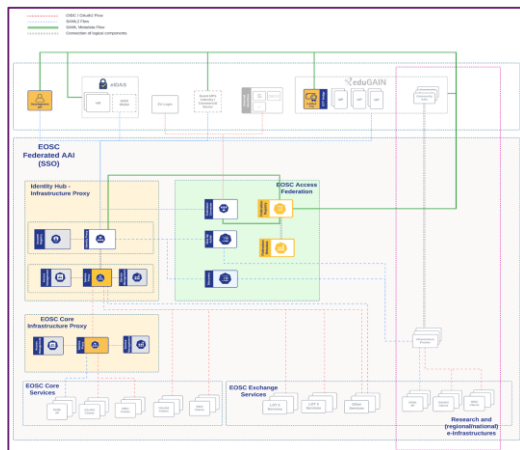


and a few more ...

through their scale gets federations to trust our AARC 'middle boxes'

# Enabling research: using the 'EOSC' with federated login

AARC compliant federation of 'national' and 'thematic' nodes in the European Open Science Cloud linked with other 'data spaces' and infrastructures



 EOSC | Federation



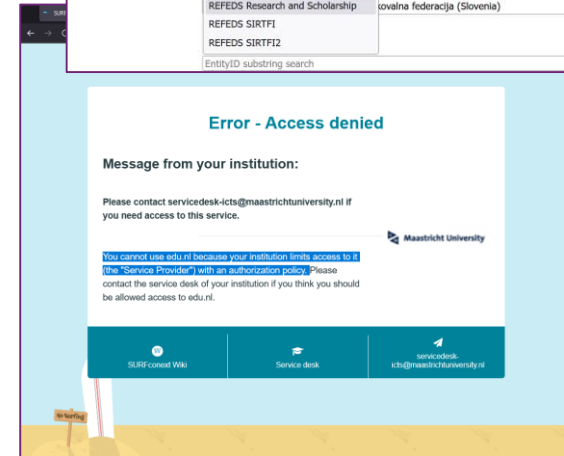
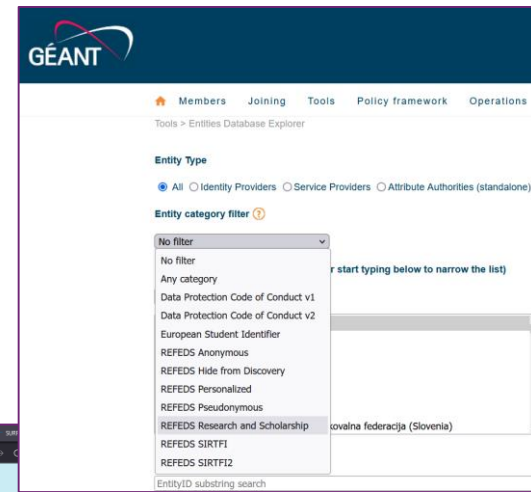
*The organisations invited to join the March 2025 kick-off workshop for the build-up phase of the EOSC Federation. All of the organisations are among the membership of the EOSC Association.*

<https://eosc.eu/eosc-about/building-the-eosc-federation/contributing-to-the-build-up-phase-of-the-eosc-federation/>; See also <https://wiki.geant.org/display/AARC/EOSC+AAI>

# And it needs everyone to work together

To scale trust in research infrastructures, we need to keep challenging ourselves ...

- *for eduGAIN: do we choose more trustworthiness and target baseline assurance, or more inclusiveness, but maybe less trust?*
- *for your university IT department: prioritize the primary mission of education and research, as both are now globally connected*
  - ‘we can use existing services from outside’
  - ‘we can contribute in collaborations in education and research’
  - ‘we teach our students to understand, study, and work with interconnected services and systems that are globally connected’*... rather than get stuck in an enterprise egg-shell approach?*
- *do our networks support a perimeter ‘fit for collaboration’?*



Images: <https://technical.edugain.org/entities>, Maastricht University blocking access to ... a privacy-friendly URI shortener ☺,

# Make and treat computing as the research instrument it is today – institutionally and globally



Institutional:  
Nikhef “Stoomboot”  
Analysis Facility



National Infrastructure  
SURF Snellius HPC



EuroHPC  
Joint Undertaking

as well as JP’s HPCI,  
US’s AccessCI, &c of course!

**There are today as much  
part of science  
as detectors are to physics  
*and: users should move  
seamlessly between tiers***

Photos: Nikhef NDPF, DelftBlue/TU Delft, SURF Data Repository, Snellius, SURF © DigitalReality; EuroHPC images: EuroHPC, LUMI Consortium, Jules Verne consortium



# And education labs are much like ad-hoc research collaboration

just slightly more organised than research ... I hope!



### SRAM API <sup>1.0</sup>

<https://sram.apispec.io>  
Documentation for the public APIs of SURF Res

#### Organisation All endpoints for organisations using an organisation API token

POST	/api/collaborations/v1	Post a new collaboration.	post_api_collaborations_v1	🔒
DELETE	/api/collaborations/v1/{co_identifier}	Delete collaboration.	delete_api_collaborations_v1_co_identifier	🔒
PUT	/api/collaborations/v1/{co_identifier}/members	Update collaboration membership.	put_api_collaborations_v1_co_identifier_members	🔒
DELETE	/api/collaborations/v1/{co_identifier}/members/{user_uid}	Delete collaboration membership.	delete_api_collaborations_v1_co_identifier_members_user_uid	🔒
PUT	/api/collaborations/v1/{co_identifier}/units	Update collaboration units.	put_api_collaborations_v1_co_identifier_units	🔒
PUT	/api/collaborations_services/v1/connect_collaboration_service	Connect service to collaboration.	put_api_collaborations_services_v1_connect_collaboration_service	🔒
PUT	/api/collaborations_services/v1/disconnect_collaboration_service	Disconnect service from collaboration.	put_api_collaborations_services_v1_disconnect_collaboration_service	🔒

Photo by [sunrise University](#) on [Unsplash](#); network diagram: FSE CSLab, Maastricht University; SRAM API: <https://sram.surf.nl/apidocs/>



# So: did we solve this inherently-cross-domain issue ... ?

Both Yes and No

Authentication and authorization 'AAI' infrastructures enable research every day

Building an *interoperable* system that enables multi-domain resource sharing remains a challenge

site map: WLCG sites 2021, visualization: CERN IT

# The AARC Blueprint – a very digestible architecture ... so




Photo credit: Marcus Hardt

# The AARC Blueprint – take a piece and feed collaboration!




Photo credit: Marcus Hardt



This work has also been co-supported by projects that have received funding from the European Union's Horizon programmes through grants AARC TREE and GEANT 5-2

Thanks to the AARC and global T&I Community, including folk from whom I re-used graphics in this overview. In random order: Licia Florio, Nicolas Liampotis, Christos Kanellopoulos, Marina Adomeit, Janos Mohacsi, Slavek Licehammer, Dave Kelsey, Ian Neilson, Marcus Hardt, Mischa Salle, Oscar Koeroo, Jouke Roorda, Davide Vagheti, Floris Fokkinga, Sven Gabriel, Hannah Short, Maarten Kremers, and many others!

# So let's digest all this ...



**SURF**

Nikhef



David Groep

davidg@nikhef.nl

<https://www.nikhef.nl/~davidg/presentations/>

 <https://orcid.org/0000-0003-1026-6606>

this work co-funded by and contributing to the Dutch National e-Infrastructure coordinated by SURF

