



Authentication and Authorisation for Research and Collaboration

Lessons Learned From AARC:

## Challenges to Pilot Components in Production E-Infrastructures

**Licia Florio** - GÉANT

**Christos Kanellopoulos** - GÉANT

**Paul van Dijk** - SURFnet

**David Groep** - Nikhef

Internet2 2017 Global Summit

April 25<sup>th</sup>, 2017

- 1. About the AARC project**
- 2. The AARC Blueprint Architecture**
- 3. AARC pilots + some pilot demos**
- 4. Sustainability for the AARC CILogin-like TTS Pilot + RCAuth demo**
- 5. Conclusions, lessons learned and looking ahead**

# Introduction – The AARC Project



## Authentication and Authorisation for Research and Collaboration



- Two-year EC-funded project
- 20 partners
  - NRENs, e-Infrastructure providers and Libraries as equal partners
- About 3M euro budget
- May 2015-2017 (2<sup>nd</sup> edition 2017-2019)
- <https://aarc-project.eu/>



- Offered a common ground for research communities, e-infrastructures and libraries to discuss challenges to adopt federated access
- Addressed some of the identified requirements
- Offered a funded and structured framework that made participation easier
- Defined models to help new research collaborations to implement interoperable AAs
- Tried to remain technology agnostic



## Why we run pilots?

We tested existing AAI components to assess to what extent they meet:

- Functional requirements
- Technical (AAI integration) requirements
- Required “readiness” levels




- Running pilots is inevitable to get a good sense of these aspects
- While running pilots, new clues and ideas arise
- Where possible improve components and ease deployability
- Improve visibility of useful AAI components for R&E


## Three types of pilots



To expand the reach of federated access (i.e. libraries)



To test technical and policy components in production infrastructures



To test cross infrastructure scenarios

# AARC: Analysis of User Communities and e-Infrastructure Providers



 05-10-2015  
**Deliverable DJRA1.1:**  
**Analysis of user community and service provider requirements**

Deliverable DJRA1.1  
Contractual Date: 31-03-2015  
Actual Date: 05-10-2015  
Grant Agreement No: 602665  
Task No: DJRA1.1  
Lead Partner: ESTEC  
Document Code: DJRA1.1  
Editors: Christos Kavalopoulos, Nicolas Lampiris, Mads van Oop, Peter Stratos

© DEKAT on behalf of the AARC project.  
The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 602665 (AARC).

**Abstract**  
This document, produced by AARC Task 1 "Analysis of user community requirements", identifies the requirements of user communities and service providers building upon the outcomes of previous activities such as the TERCOS AARC Study and the World e-identity survey. The requirements identified by these activities have been updated and enriched with user requirements that the team collected through a survey of user communities as well as a set of digital operators. These requirements are outlined here and will be provided as input for upcoming activities in AARC.

Attribute Release	Attribute Aggregation	User Friendliness	SP Friendliness
Credential translation	Persistent Unique Id	User Man. Information	Credential Delegation
Levels of Assurance	Guest users	Step-up AuthN	Best Practices
Community based AuthZ	Non-web-browser	Social & e-Gov IDs	Incident Response





Authentication and Authorisation for Research and Collaboration

# The AARC Blueprint Architecture

Christos Kanellopoulos - GÉANT



# International Research Collaborations



1. Users should be able to access the all services using the **credentials from their Home Organizations** when available.
2. Secure integration of **guest identity solutions** and **support for stronger authentication mechanisms** when needed.
3. **Access** to the various services should be granted **based on the role(s)** the users have **within the collaboration**.
4. Users should have one **persistent identity across all community services** when needed.
5. **Ease of use for users and service providers.** The complexity of multiple IdPs/Federations/Attribute Authorities/technologies should be hidden.



Attribute Release	Attribute Aggregation	User Friendliness	SP Friendliness
Persistent Unique Id	Credential translation	Credential Delegation	User Managed Information
Levels of Assurance	Guest users	Step-up AuthN	Best Practices
Community-based AuthZ	Non-web-browser	Social & e-Gov IDs	Incident Response



# Identity & Access Management for International Research Collaborations

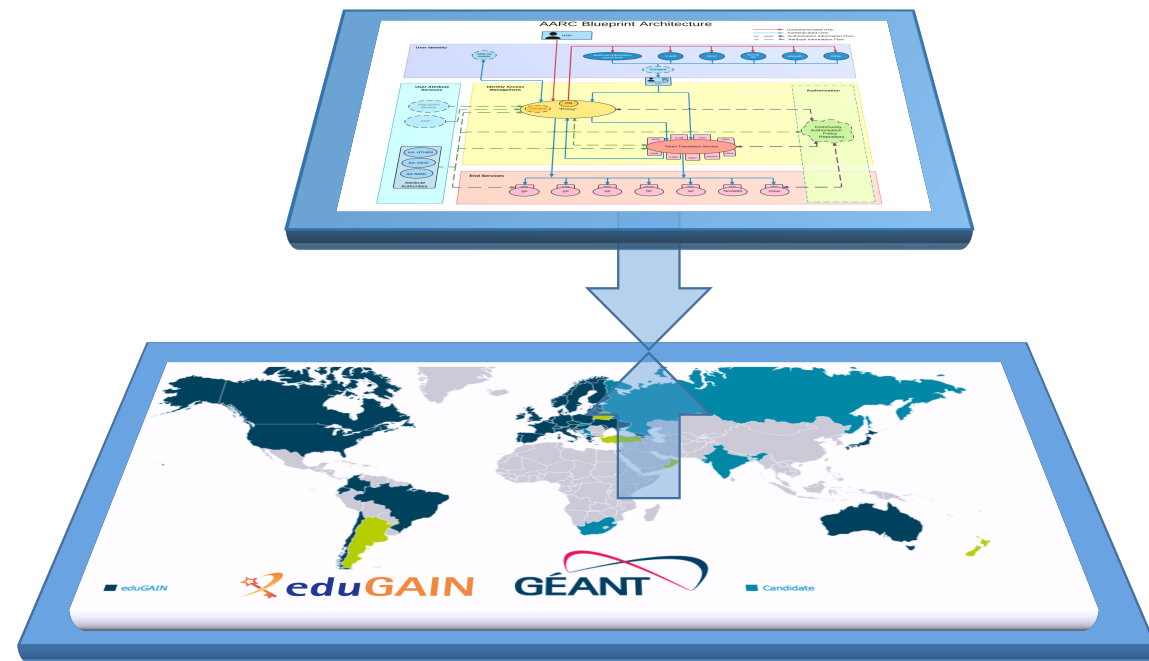


## A Blueprint Architecture for authentication and authorization

A set of architectural building blocks  
on top of eduGAIN

## eduGAIN and the Identity Federations

A solid foundation for federated  
access in Research & Education



# Identity & Access Management for International Research Collaborations

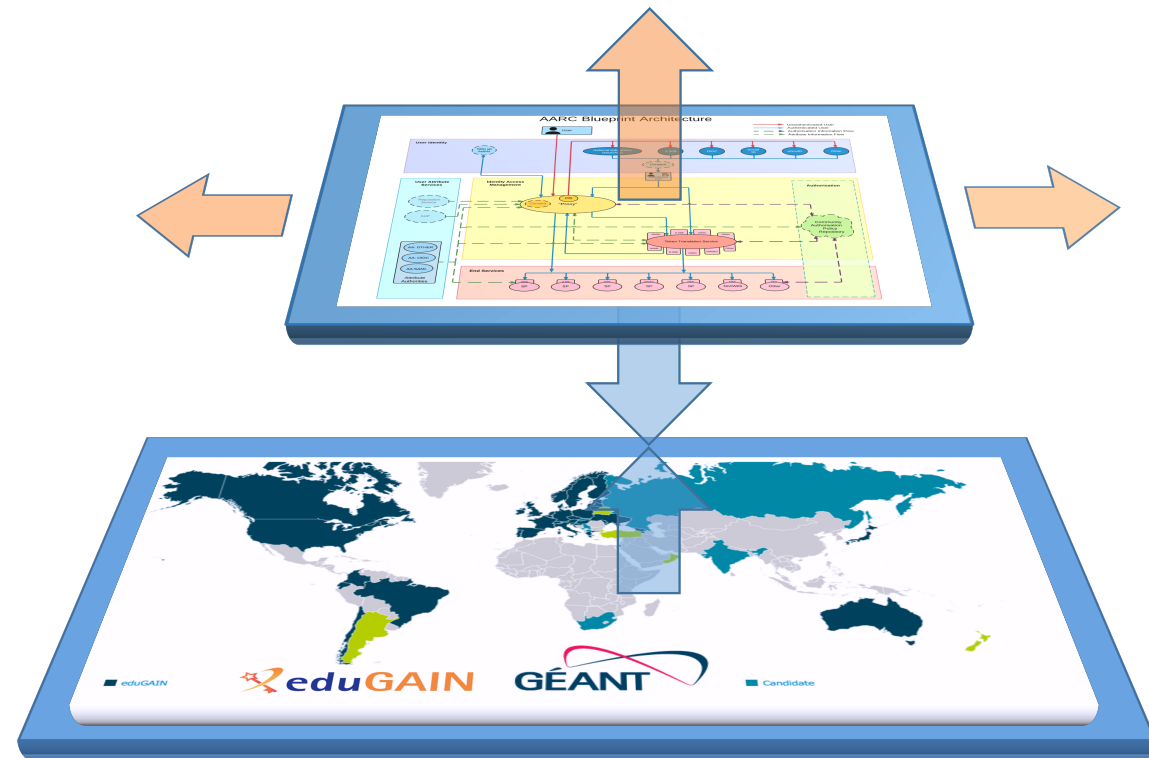


## A Blueprint Architecture for authentication and authorization

A set of architectural building blocks on top of eduGAIN

## eduGAIN and the Identity Federations

A solid foundation for federated access in Research & Education



# AARC Blueprint Architecture

## AARC Blueprint Architecture



- Unauthenticated User
- Authenticated User
- Authorisation Information Flow
- Attribute Information Flow

### User Community Requirements




05-10-2015  
**Deliverable DJRA1.1:**  
**Analysis of user community and service provider requirements**

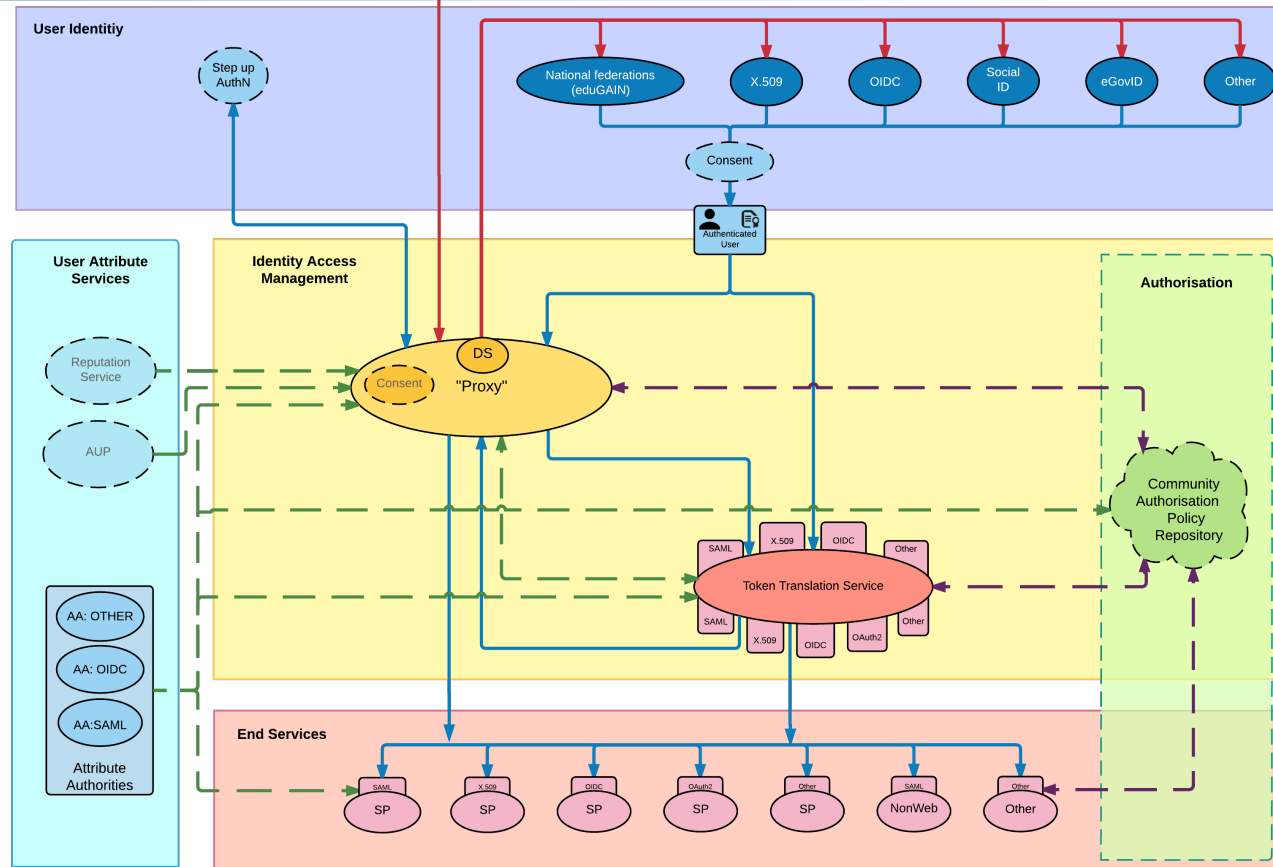
Deliverable DJRA1.1  
 Contractual Date: 31-08-2015  
 Actual Date: 05-10-2015  
 Grant Agreement No.: 603666  
 Work Package: JRA1  
 Task Item: JRA1.1  
 Lead Partner: EC.eu  
 Document Code: DJRA1.1  
 Editors: Christa Kanellopoulou, Nicolas Lampiris, Nels van Oij, Peter Stogias

© EC/ANT on behalf of the AARC project.  
 This research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 603666 (AARC).

AARC  
 This document, produced by JRA1 Task 1 "Analysis of user community requirements", identifies the requirements of user communities and services providers building upon the outcomes of previous activities such as the TRENDS AAA Study and the FRM11 workshop series. The requirements identified in these activities have been updated and enriched with new requirements that have been obtained through a survey of user communities as well as a set of targeted interviews. These requirements are analysed here and will be provided as input for upcoming activities in AARC.



<https://goo.gl/kSxENp>



<https://aarc-project.eu/blueprint-architecture/>

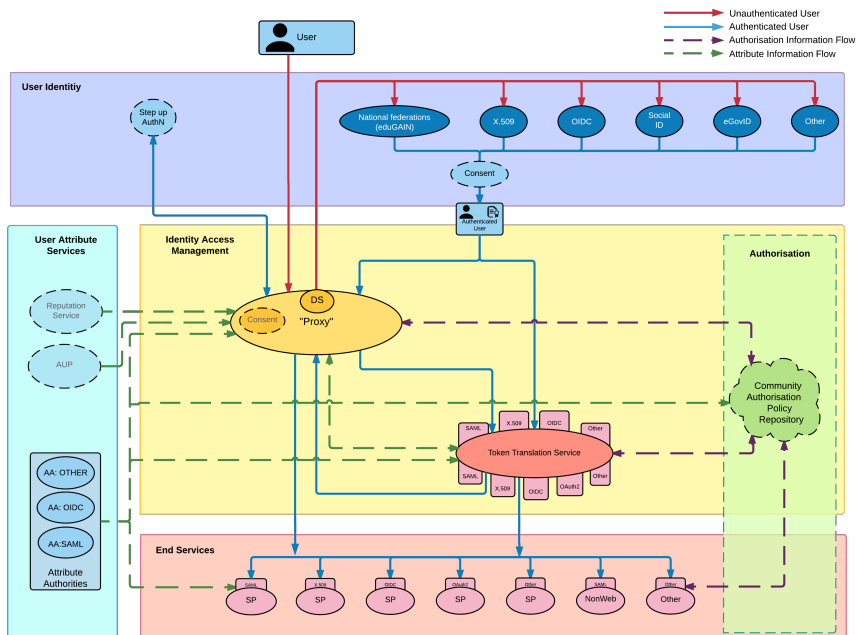
# AARC Blueprint Architecture



<https://aarc-project.eu/blueprint-architecture/>

## Guidelines and support documents

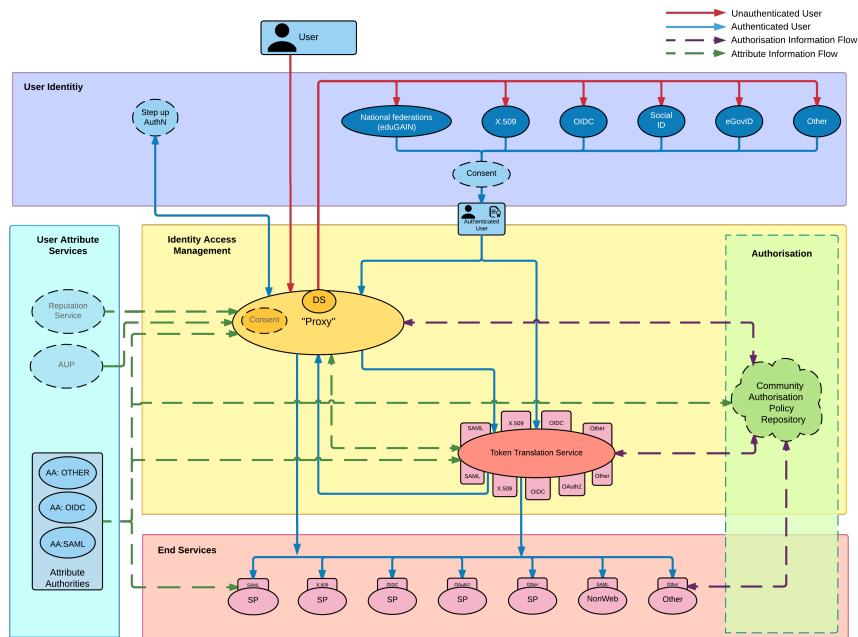
- Best practices for managing authorisation
- Expressing group membership and role information
- Scalable attribute aggregation
- Implementation of token TTS
- Credential delegation
- Non-web access
- Social media IdPs
- Use cases for account linking
- Use cases for LoA elevation via step-up authentication



# AARC Blueprint Architecture



<https://aarc-project.eu/workpackages/policy-harmonisation/>




## Policy recommendations & frameworks

- Security Incident Response Trust Framework for Federated Identity – Sirtfi
- Scalable Negotiator for a Community Trust Framework in Federated Infrastructures – Sncftfi
- Recommendations on Minimal Assurance Level Relevant for Low-risk Research Use Cases
- Differentiated LoA recommendations for policy and practices of identity and attribute providers
- Recommendations and template policies for the processing of personal data by participants in the pan-European AAI

# AARC Blueprint Architecture & Requirements





05-10-2015  
**Deliverable DJRA1.1:**  
**Analysis of user community and service provider requirements**

**Deliverable DJRA1.1**  
 Contractual Date: 31-08-2015  
 Actual Date: 05-10-2015  
 Grant Agreement No.: 653965  
 Work Package: JRA1  
 Task Item: JRA1.1  
 Lead Partner: EGI.eu  
 Document Code: DJRA1.1  
 Editors: Christos Kanellopoulos, Nicolas Liampotis, Niels van Dijk, Peter Solagna

© GEANT on behalf of the AARC project.  
 The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

**Abstract**  
 This document, produced by JRA1 Task 1 "Analysis of user community requirements", identifies the requirements of user communities and services providers building upon the outcomes of previous activities such as the TERENA AAA Study and the FIMAP workshop series. The requirements identified by these activities have been updated and enriched with new requirements that the team collected through a survey of user communities as well as a set of targeted interviews. These requirements are analysed here and will be provided as input for upcoming activities in AARC.

Attribute Release

Attribute Aggregation

User Friendliness

SP Friendliness

Persistent Unique Id

Credential translation

Credential Delegation

User Managed Inf.

Levels of Assurance

Guest users

Step-up AuthN

Best Practices

Community based AuthZ

Non-web-browser

Social & e-Gov IDs

Incident Response





Authentication and Authorisation for Research and Collaboration

## **AARC Pilots**

Paul van Dijk - SURFnet



## Pilots on the integrated R&E AAI



## The pilot approach in AARC

## Why we run pilots?

We tested existing AAI components to assess to what extent they meet:

- Functional requirements
- Technical (AAI integration) requirements
- Required “readiness” levels

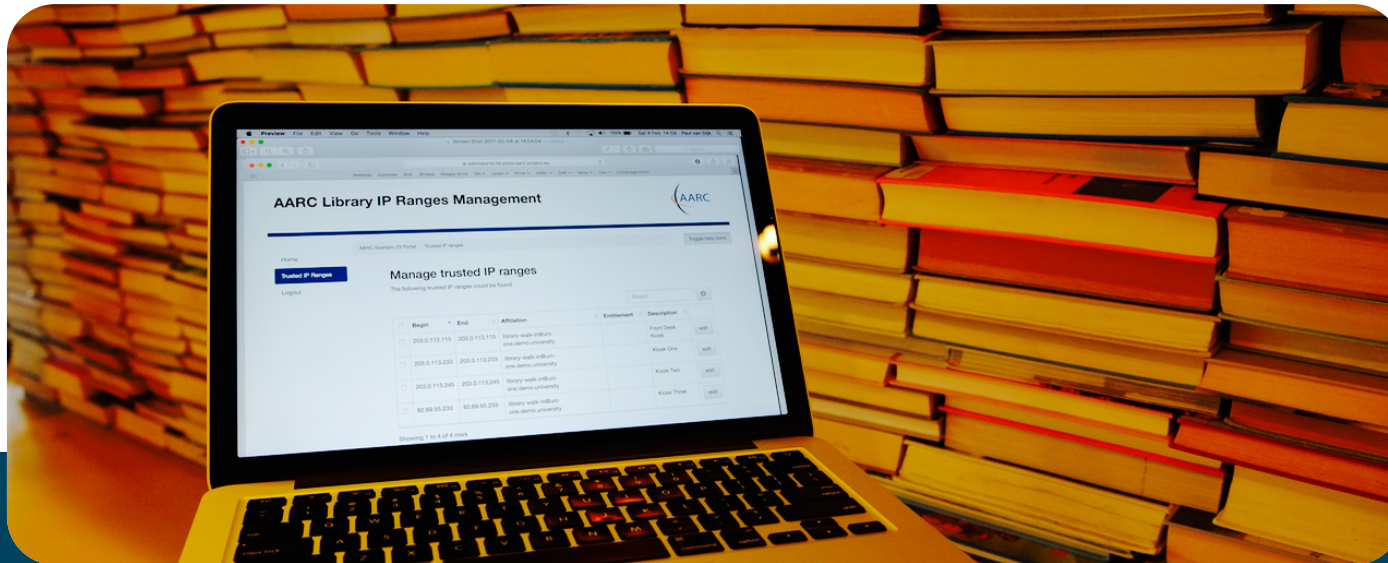


- Running pilots is inevitable to get a good sense of these aspects
- While running pilots, new clues and ideas arise
- Where possible improve components and ease deployability
- Improve visibility of useful AAI components for R&E

## Establish a common test-bed infrastructure

- A staging area for piloted services
- Technical platform delivered by 
- >20 VMs instantiated
- Using Ansible scripts for deployment
- SimpleSAMLphp DIY IdP available
- Online support by  staff





**Expand the reach of federated access (Libraries, external IdPs)**

## Library proxy pilot(s)

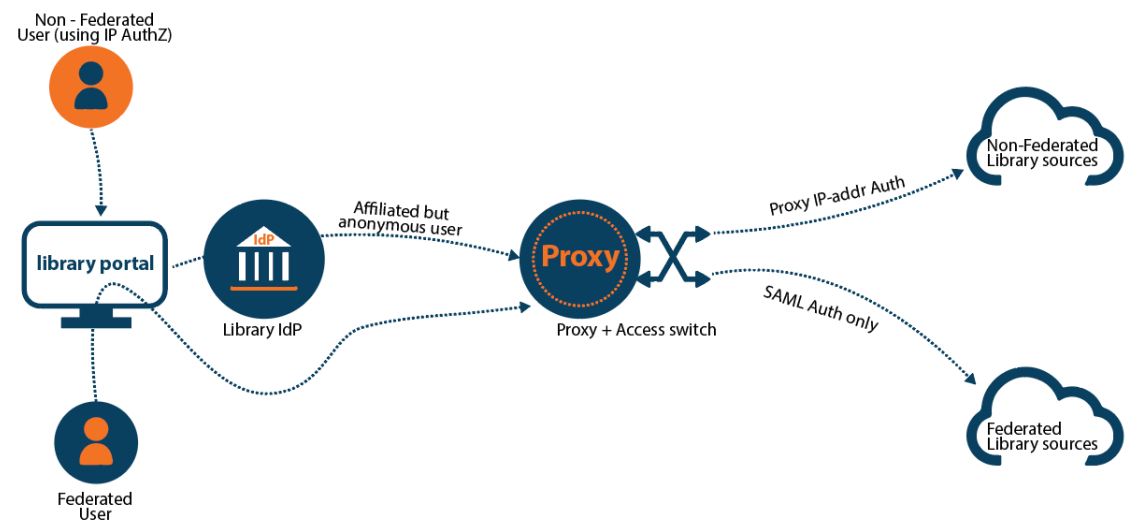
### Purpose

- Dealing with “walk by users” → Setting attributes based on IP-address
- Dealing with providers supporting IP-address access only → SP Proxy translating SAML to IP if needed

### Services/Components used

- Shibb add-on to filter on IP-address
- EZproxy with access switch mode
- Several library resources

[wiki.geant.org/x/a4qSAw](http://wiki.geant.org/x/a4qSAw)



## Social ID pilot

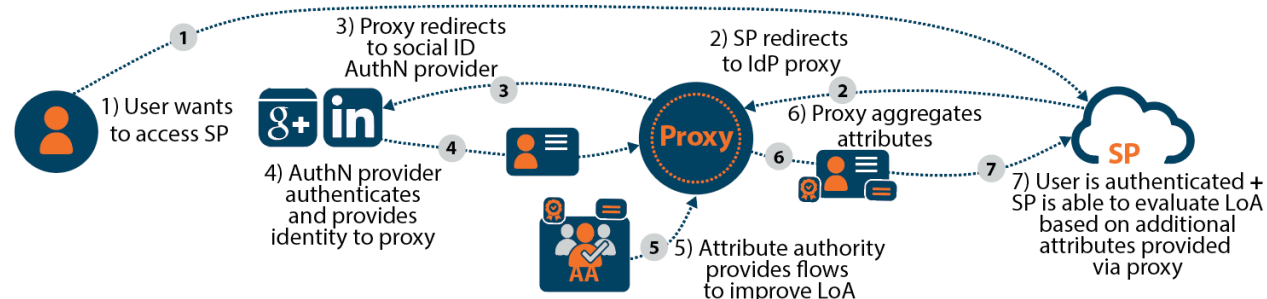
### Purpose

- Demonstrate possible mechanisms to include users with Social Identities
- Explore clues to enhance LoA of users

### Services/Components used

- Social ID providers (Google, ORCID, LI)
- COmanage AA
- SimpleSAMLphp proxy
- OpenStack Keystone SP
- Tested with EGI and AARC pilot community

[wiki.geant.org/x/ZlqSAw](http://wiki.geant.org/x/ZlqSAw)



BPA building blocks

Administrative domains





## Testing technical and policy components



# SAML – ORCID account linking pilot

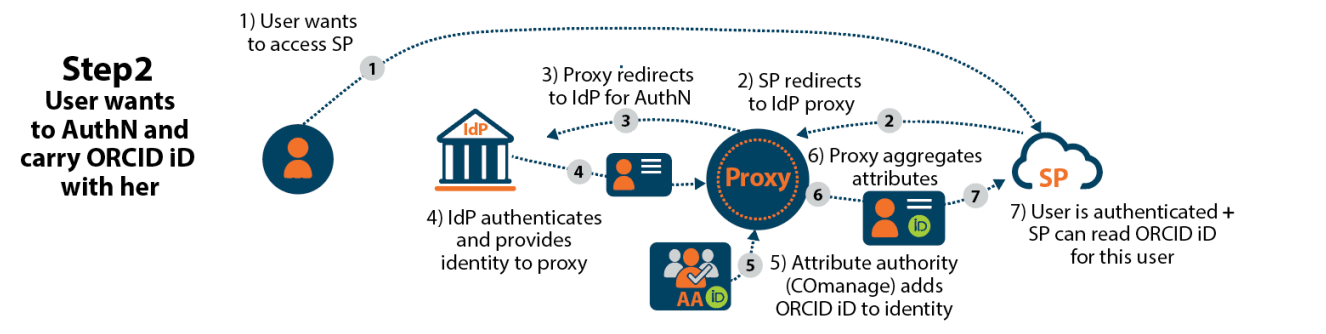

## Purpose

- ORCID provides persistent researcher centric IDs which are useful for use in collaboration services → include this ID in the assertion

## Services/Components used

- ORCID API - persistentID source
- COmanage – link ID to account
- Proxy – attribute aggregation
- Tested with the AARC community

[wiki.geant.org/x/WAH5Aw](http://wiki.geant.org/x/WAH5Aw)



# Attribute management & aggregation pilots



## Purpose

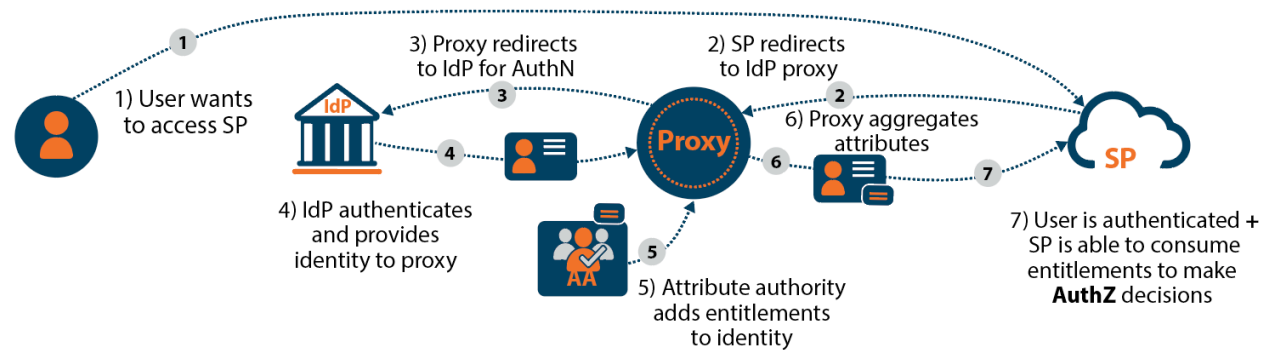
- Show how attributes from multiple AAs can be used for AuthZ in a fed. environment
- Delegate AuthZ decisions
- Minimize impact for SPs

## Services/Components used

- COmanage/PERUN AAs
- SimpleSAMLphp proxy
- OpenStack Horizon SP/BBMRI SPs

**EGI:** [wiki.geant.org/x/LAH5Aw](http://wiki.geant.org/x/LAH5Aw)

**BBMRI:** [wiki.geant.org/x/HgD5Aw](http://wiki.geant.org/x/HgD5Aw)



BPA building blocks

IdentityProvider

AttributeAuthority  
COmanage

Proxy  
simpleSAMLphp

TokenTranslation

ServiceProvider  
Just a service provider

Administrative domains

eduGAIN IdPs

E-infrastructure/Collab. Organizations  
ESI

eduGAIN SPs

# TTS: RCauth




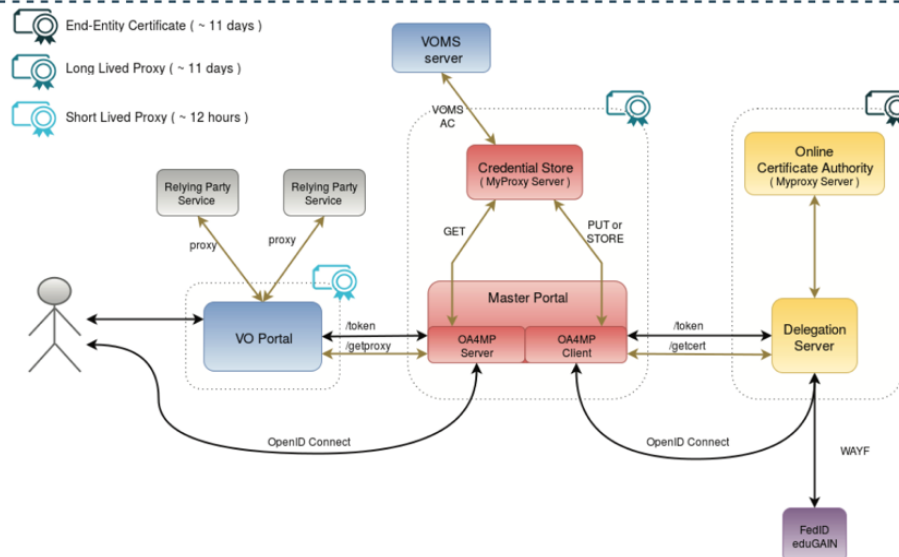
## Purpose

- Enable access to certificate based services for users with an institute account, generating certs on the fly
- Bridging eduGAIN & IGTF

## Services/Components used

- CILogon, adapted as RCAuth
- Several master portals
- Several science gateways
- SimpleSAMLphp
- VOMS Attribute Authority
- Tested with AARC community +...

[wiki.geant.org/x/yADaAw](http://wiki.geant.org/x/yADaAw)



# Bridging IGTF to eduGAIN



## Purpose

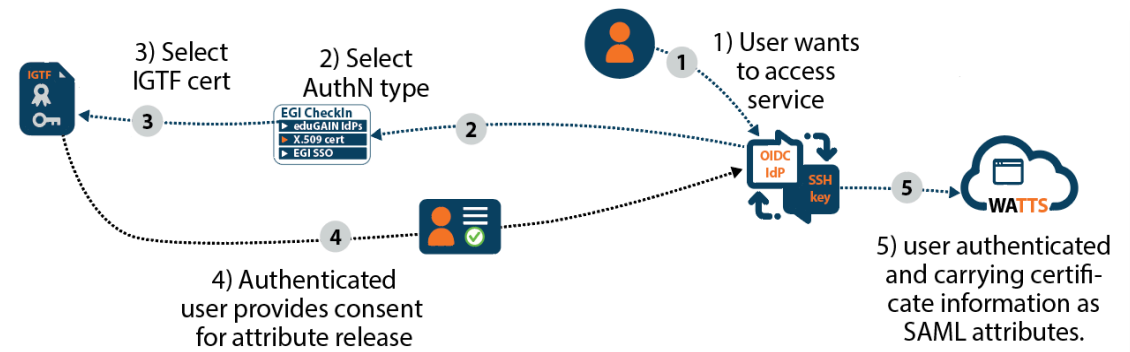
- Demonstrating how researchers can use X.509 certs to access eduGAIN services with substantial or higher LoA
- Not forcing them to use organization accounts



## Services/Components used

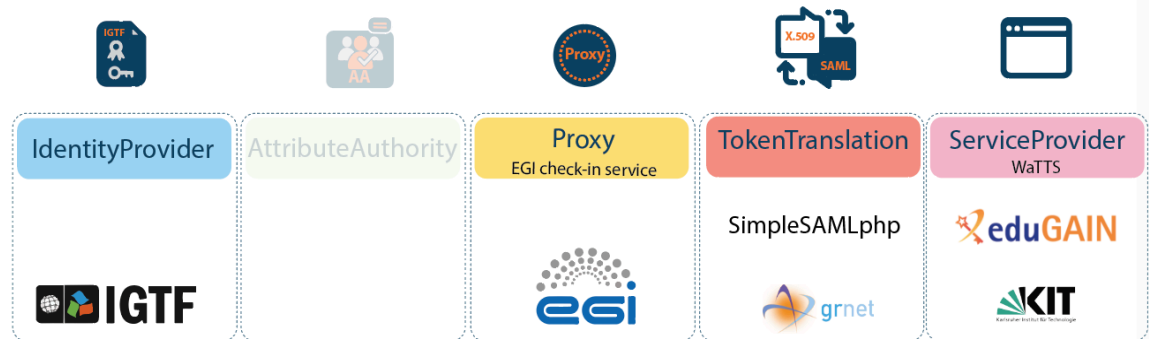
- SimpleSAMLphp add-on
- WaTTS one-stop-TTS-shop
- ~Okeanos infrastructure

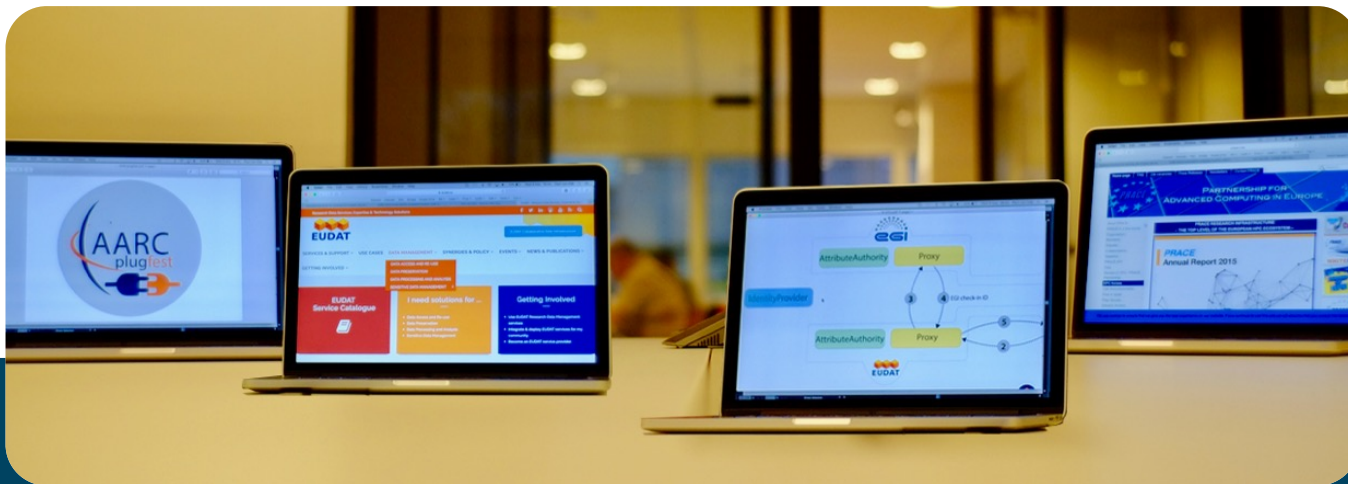
[wiki.geant.org/x/JoEKB](http://wiki.geant.org/x/JoEKB)



BPA building blocks

Administrative domains



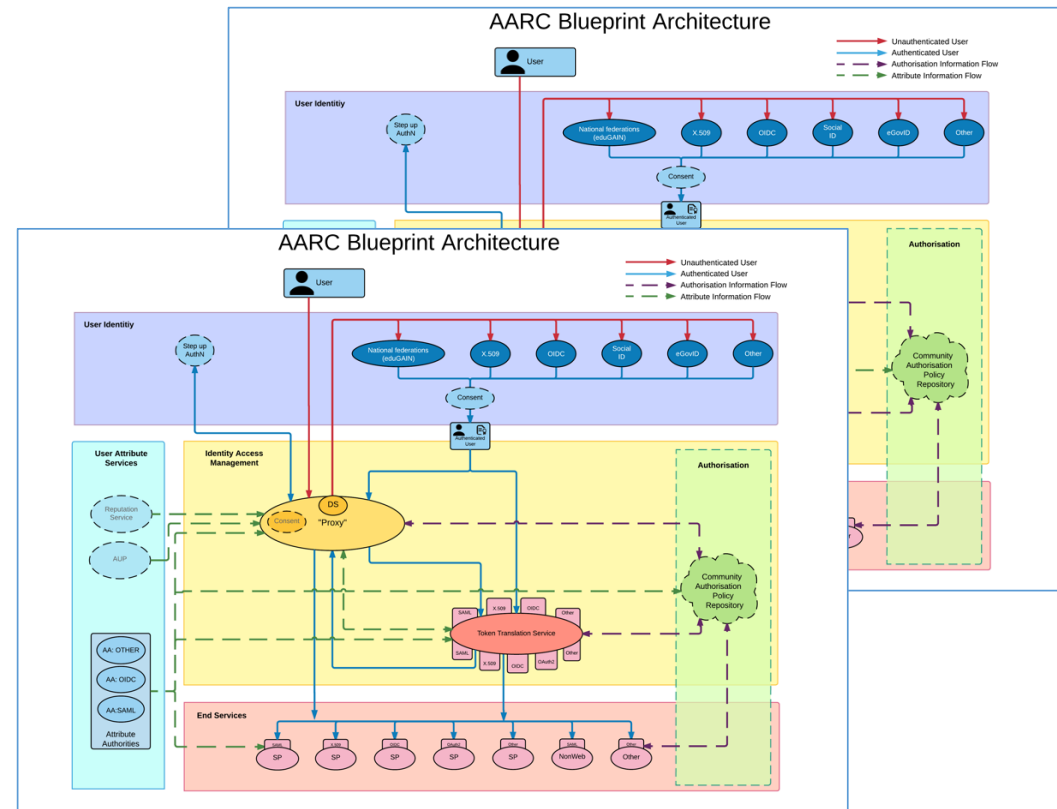


## Cross infrastructure pilots

# Cross-infrastructure pilots



AARC  
blueprint  
architecture



## Cross-infrastructure pilots, EUDAT - EGI

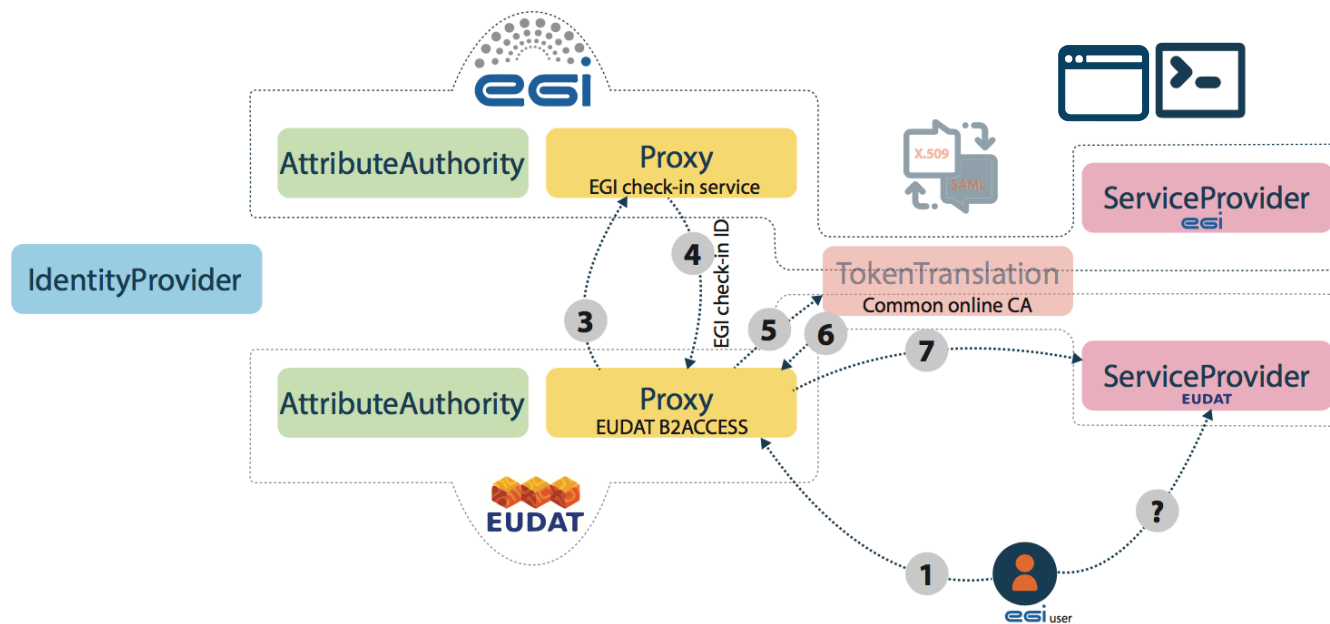
### Purpose

- Enable access to EUDAT services for users registered @EGI
- Bridging EUDAT and EGI infra

### Services

- EGI check-in service
- EUDAT B2ACCESS
- RCAuth (for non-web services)

[wiki.geant.org/soon](http://wiki.geant.org/soon) available





Authentication and Authorisation for Research and Collaboration

## Sustainability for the AARC CILogin-like TTS Pilot

*Distribution and models for the CILogon-like TTS Pilot  
for the European Open Science Cloud and  
the Dutch National e-Infrastructure coordinated by SURF*

**David Groep**

NA3 coordinator

Nikhef PDP (Advanced Computing Research) group



AARC I2GS session, EGI ENGAGE Conference

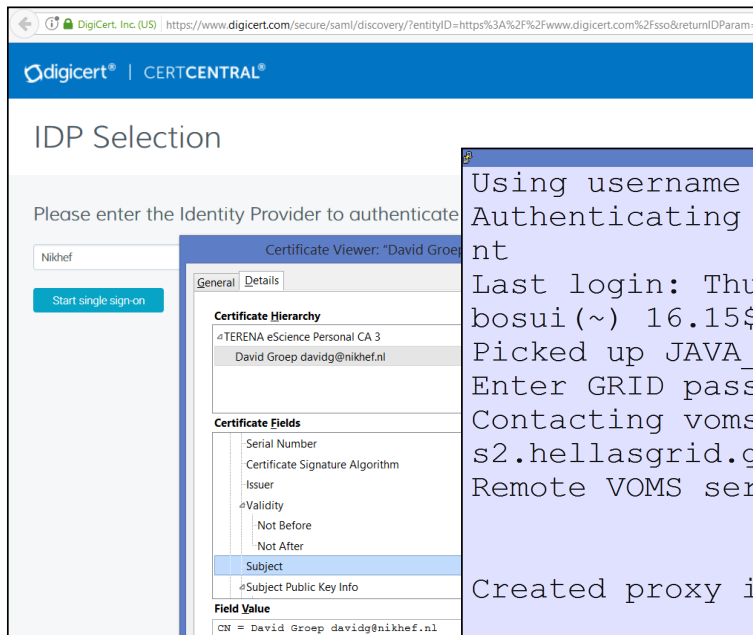
April, May 2017



# Seamless (eduGAIN) Access to (non-Web) Resources using PKIX?



## Traditional workflow – using a client-held credential



Works great *provided* the user understand the technology – and we may have found all users that know how to manage ☹️

```
Using username "davidg".
Authenticating with public key
nt
Last login: Thu Apr 13 17:43:46 2017 from 2a07:8500:120:e03b:
bosui(~) 16.15$ voms-proxy-init -voms dteam
Picked up JAVA_TOOL_OPTIONS: -Xmx512M
Enter GRID pass phrase for this identity:
Contacting voms2.hellasgrid.gr:15004 [/C=GR/O=HellasGrid/OU=h
s2.hellasgrid.gr] "dteam"...
Remote VOMS server contacted succesfully.

Created proxy in /tmp/x509up_u5917.

Your proxy is valid until Wed Apr 19 04:16:05 CEST 2017
bosui(~) 16.16$ █
```

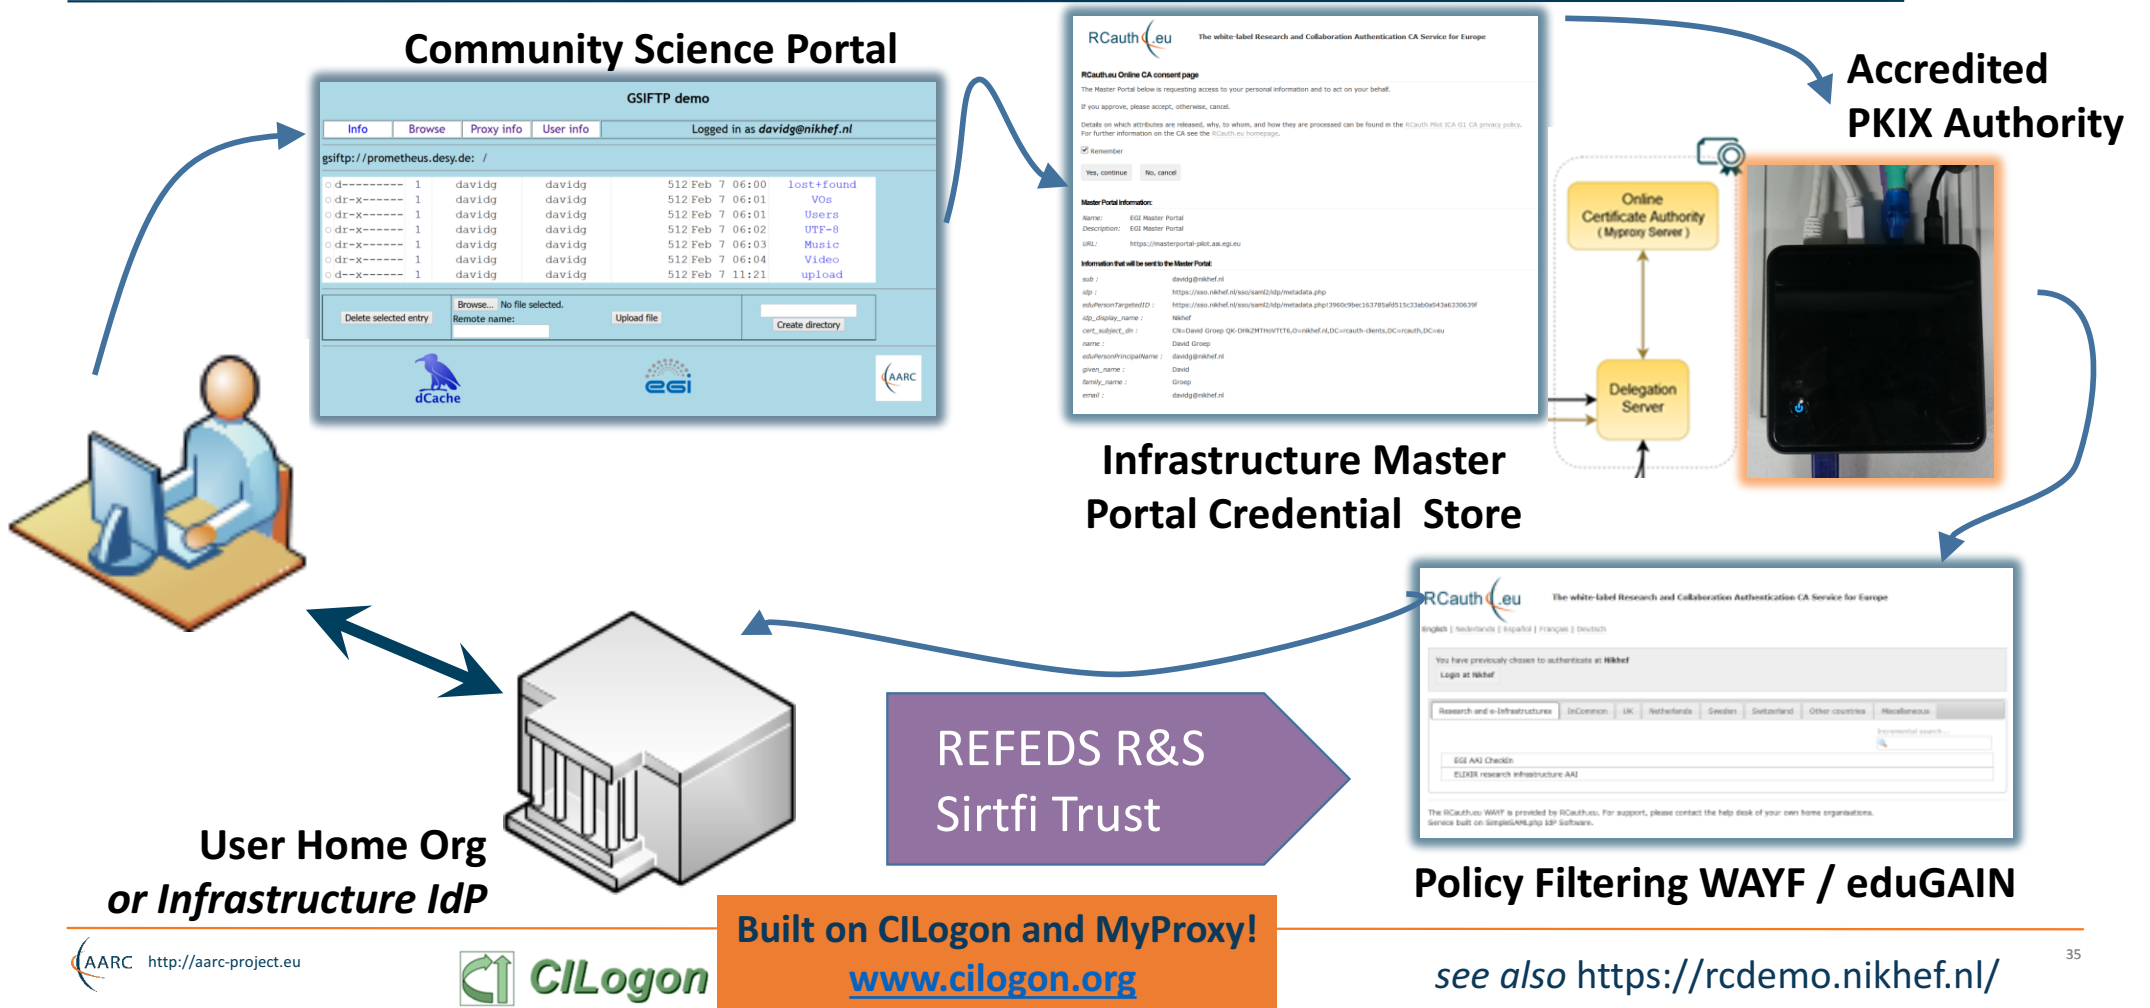
```
bosui(/user/davidg (davidg.emin)
bosui(~) 16.25$ gsissh sgmlhcb@kot.nikhef.nl -p 1975 'id -a && hostname -f'
uid=991(sgmlhcb) gid=2015(lhcbsgm) groups=2015(lhcbsgm)
kot.nikhef.nl
bosui(~) 16.25$
```

# Seamless (eduGAIN) Access via the CILogon-like TTS Pilot: aims




- **Ability to serve a large pan-European user base without national restrictions**
  - without having to rely on specific national participation exclusively for this service
  - serving the needs of cross-national user communities that have a large but sparsely distributed user base
- **Use existing resources and e-Infrastructure services**
  - without the needs for security model changes at the resource centre or national level
- **Allow integration of this system in science gateways and portals with minimal effort**
  - only light-weight industry-standard protocols, limit security expertise (and exposure)
- **Permit the use of the VOMS community membership service**
  - attributes for group and role management in attribute certificates
  - also for portals and science gateways access the e-Infrastructure
- **Concentrate service elements that require significant operational expertise**
  - not burden research communities with the need to care for security-sensitive service components
  - keep a secure credential management model
  - coordinate compliance and accreditation – and help meet EU privacy stuff in just one place to ease adoption
- *Optional elements: ability to obtain CLI tokens (via ssh agent or even U/P); implicit AuthZ*

# Flow for RCauth-like scenarios



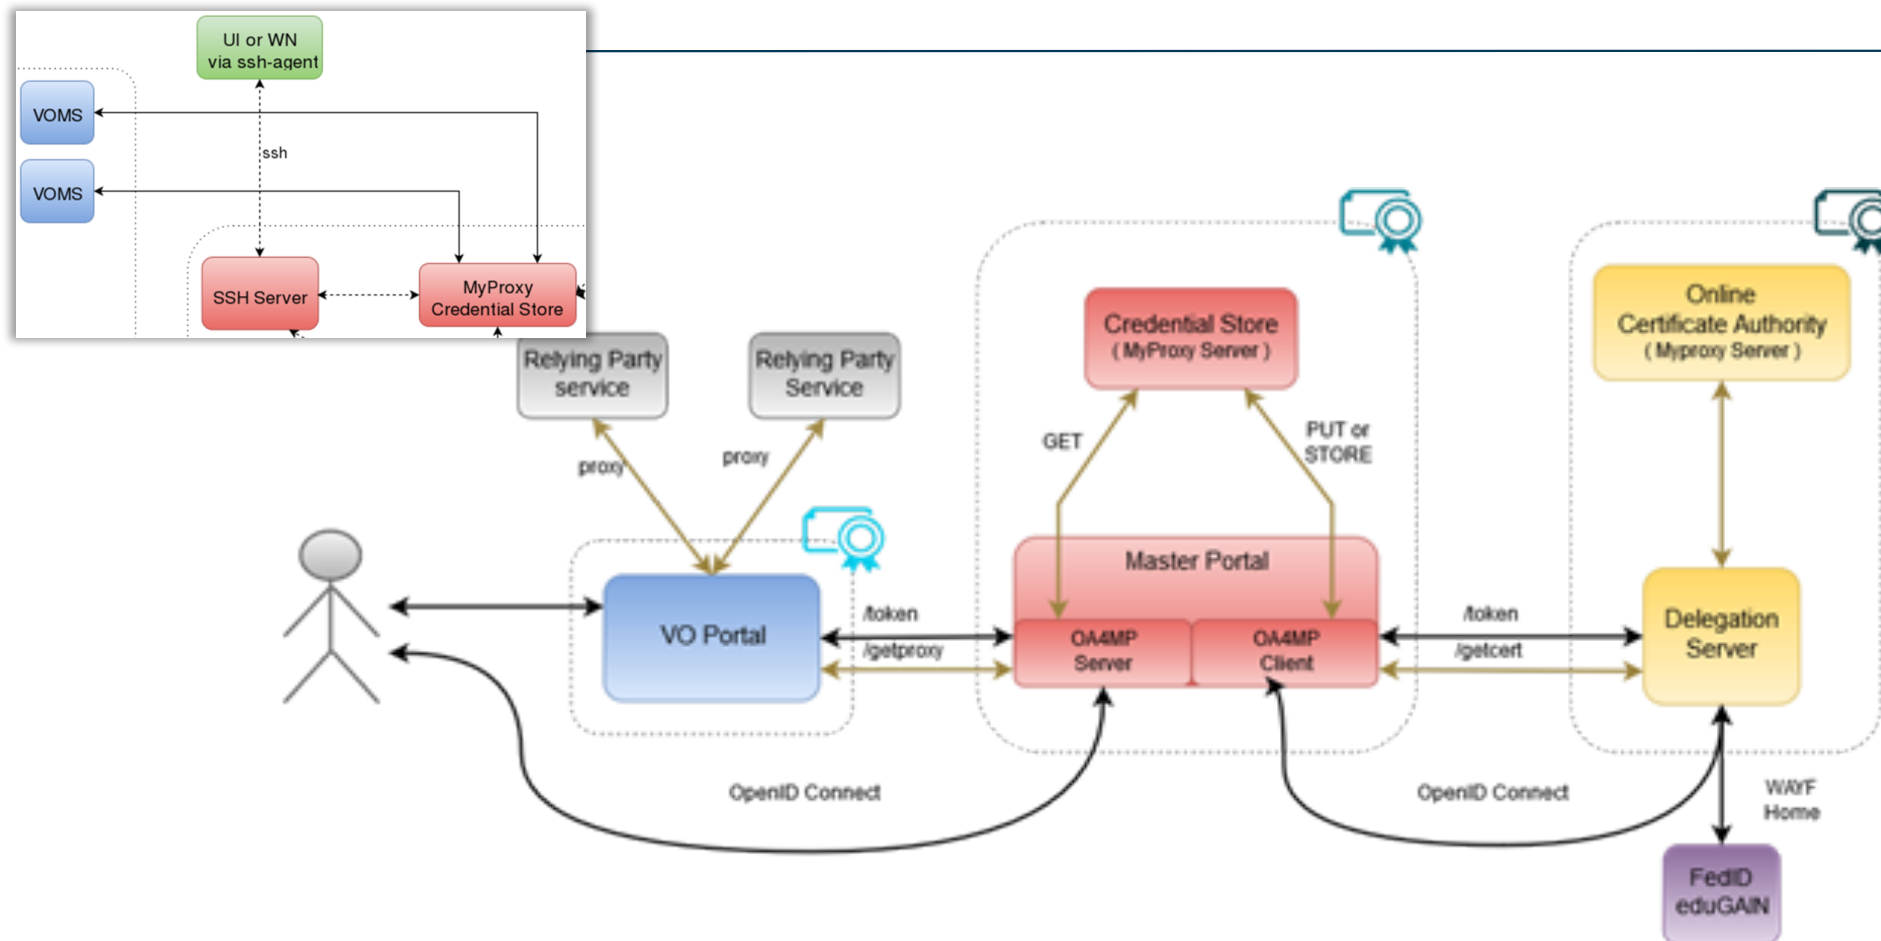
# DEMO



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

*you can do most things except access the Prometheus dCache pool right now  
that last elements needs your credential to be added to the permitted list*

# CILogon-like TTS Pilot - distributable elements



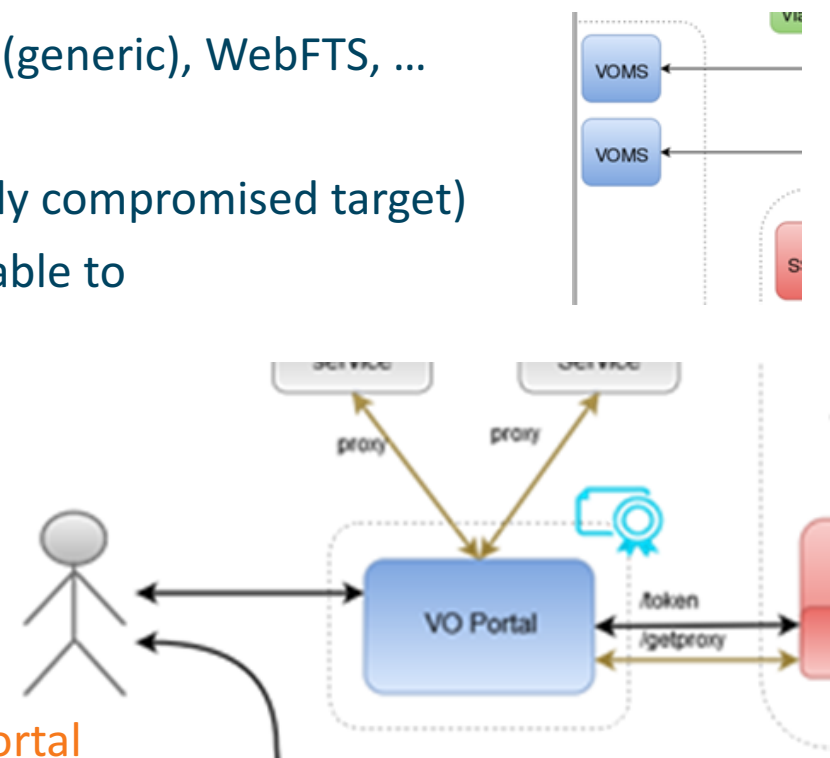
## Blue: VO services and resources

- Are around today, either self-managed or hosted, in most communities
- Science gateways, portals, e.g. HADDOCK, Galaxy, LifeRay (generic), WebFTS, ...
- Omnipresent (and has unfortunately proven to be an easily compromised target)
- Will have to *get credentials* from the MPs, but should be able to do so *only for authenticated users*
- Downtime will impact its' own users, but there will be many of these (same service by different sites?)

### Considerations:

Operated as today by the communities

Bound slightly stronger to the community via the Master Portal



## Red: the Master Portal (MP) and Credential Repository


- Needs to be a trusted element (to permit credential storage by the TTS)
- Requires some operational and security expertise (managed data centre, locked racks, access controls, ability to designate infrastructure for security operations, trained staff)
- Connects to (many) workflow-specific VO portals
- Connects to a single Delegation Service/TTS – and can give *IdP hints*
- Downtime of an MP disables resource access for connected VO portals

### Considerations:

One per Research or e-Infrastructure


Should be highly available (database sync)

Infrastructures can also build their own (e.g. in WaTTS, Unity, ...)



## Yellow: the Token Translator/OA4MP/CA service

- Needs security and policy expertise – and ability to maintain accreditation
- Needs operational and technical capabilities: hardware security modules, managed data centres, off-line and on-line secure areas, ROBAB-proof trained personnel, ability to designate infrastructure for security operations
- Connects to (a few, we hope) Master Portals (MPs) with explicit agreements *to take care of user credential protection and compliance*
- Connects (many, we hope scalably) federations, IdPs and (few) SP-IdP-Proxies
- *May have to present a WAYF, if the VO portal does not pass IdP entityID*



### Considerations:

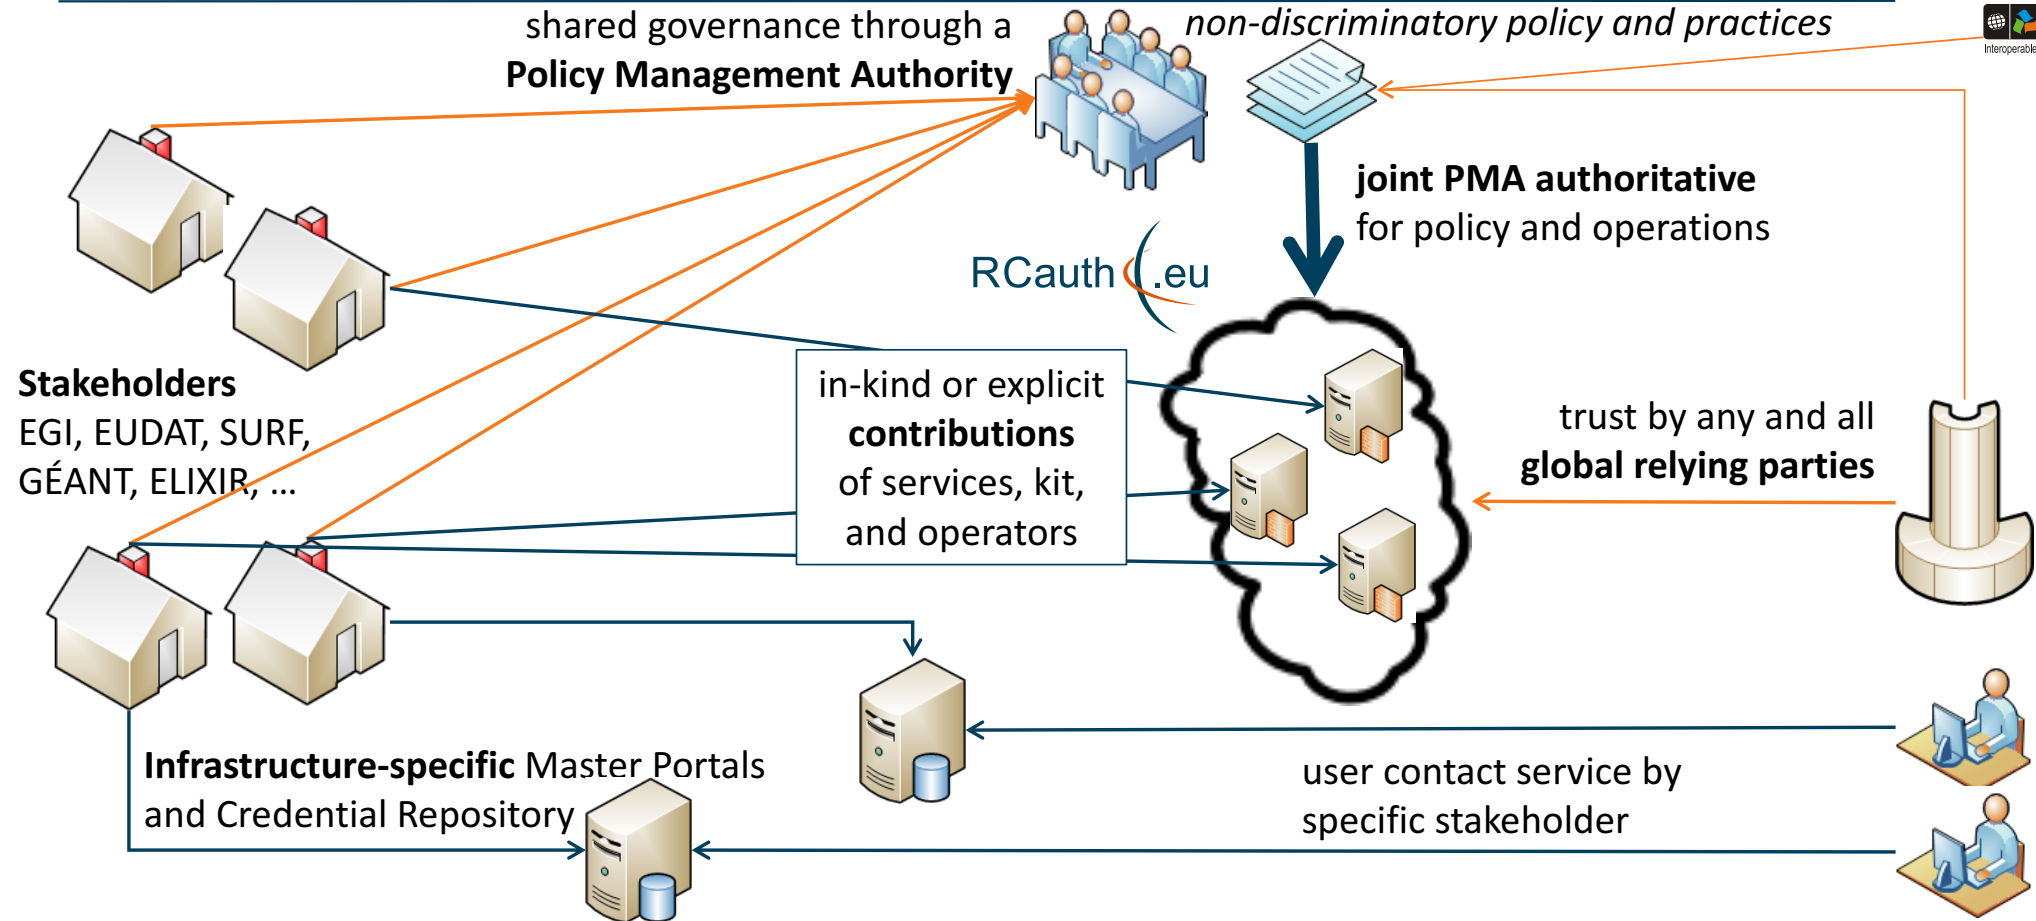
Trust and compliance, with IGTF accreditation

Single logical instance, with HA built in for production

Managed by a consortium: in Europe agreed by at least EGI, EUDAT, GÉANT, ELIXIR, and SURF



# Potential RCauth.eu management model



## Purple: connected federations and IdPs (proxies)

- In the generic case (conventional R&E federations) only limited control possible
- Infrastructure-managed IdPs will provide more specific capabilities, *e.g.*, uniqueness
- Connects to many services, of which the DS/TTS is just one
- Build on common technology (keep with SAML, no OIDC here)
- Shared policy compliance: REFEDS R&S, Sirtfi
- Negotiate only when needed (but a TTS must serve all users to prevent fragmentation!)
- Cope with heterogeneity (i.e.: use a 'filtering WAYF/entity filter proxy')

### Considerations:

eduGAIN registration, Sirtfi adoption, REFEDS R&S, filtering capability  
Needs a friendly registrar, but is otherwise 'just another SP'



# Sirtfi and R&S – policy needs for trust in identity providers and federations



- REFEDS and federation focused FAQ
- Definition of the global Security Contact meta-data profile for use in eduGAIN
- Namespace for Sirtfi Assurance at IANA
- Used in cyber ops roleplay exercises
- Promoted at I2TechX, FIM4R, Kantara, and TF-CSIRT
- Ingredient to the ‘CILogon pilot’  
*combination of*  
**REFEDS “Research and Scholarship”**  
*and*  
**Sirtfi v1.0**

The screenshot shows the SIRTFI website page. The header is red with 'SIRTFI' on the left, the URL 'https://refeds.org/SIRTFI' in the center, and 'REFEDS > SIRTFI' on the right. The main content area is white and contains the following text: 'The Security Incident Response Trust Framework for Federated Identity (Sirtfi) aims to enable the coordination of incident response across federated organisations. This assurance framework comprises a list of assertions which an organisation can attest in order to be declared Sirtfi compliant. Visit our [Wiki](#) to discover how your organisation can prepare itself for Federated Incident Response with Sirtfi.' Below this is a paragraph: 'REFEDS' [Sirtfi Working Group](#) has been active since 2014 and combines expertise in operational security and incident response policy from across the REFEDS community. Work to publish and implement the Sirtfi Trust Framework is supported by the [AARC Project](#).' At the bottom, there are three red icons in rounded squares: a group of people for 'Benefits', a document for 'Sirtfi v 1.0', and a question mark for 'FAQs'. Below each icon is a link: 'Why should I join? What are the [Benefits](#)?' for Benefits, 'View the [Sirtfi Framework](#)' for Sirtfi v 1.0, and 'Need [help](#)?' for FAQs.



Mar 17<sup>th</sup> 137 IdPs in eduGAIN that support Sirtfi

meets assurance requirements for RIs and EIs according to the IGTF “assured identifier trust”



Authentication and Authorisation for Research and Collaboration

## **Conclusions, Lessons Learned and looking ahead**



## Conclusions<sup>1</sup>

---

### Successfully...

- ✓ Deployed many different AAI solutions approx. 20
- ✓ Reused and glued together existing components
- ✓ Tested/discussed pilot results with communities
- ✓ Provided architecture and guidelines
- ✓ Provided software sources and deployment scripts

**Many results are being rolled out in production already in R&E infrastructures**

We've shown that

# Proxy = a key element (!)

in research collaboration use cases

E.g.:

- Library pilots → bridging SAML to IP-access
- E-infra pilots → collect, aggregate and forward AuthZ attributes, easy to digest by SPs
- TTS Pilots → bridging SAML to ssh/x.509 and v.v. while hiding complexity



## Conclusions<sup>3</sup>



Successfully **bridged** eduGAIN NREN (SAML) world to e-infrastructure (ssh, X.509) world



## Lessons Learned

---

- We needed this project to show the full **potential of AAI components for Research and Collaboration**
- Scoping and executing the **pilots** was a challenge
- Bring the communities together, speak the same language and increase mutual understanding. E.g. the **Blueprint Architecture** allowed to establish common understanding
- Engage with communities from the very beginning and have demos available to increase adoption
- Thanks to these results, AAI for research is on the radar, many e-infrastructures and research communities recognize the added value of AAI and are engaged now → **2<sup>nd</sup> edition of AARC**





# AARC 2<sup>nd</sup> edition – A wide range of research communities committed



More information: [aarc-project.eu](http://aarc-project.eu),  
Full list of pilots: [wiki.geant.org/x/RIOjAw](http://wiki.geant.org/x/RIOjAw)

A screenshot of the AARC website homepage as it appears in a web browser. The browser's address bar shows 'GÉANT Association'. The page header includes the AARC logo and the full name 'Authentication and Authorisation for Research and Collaboration'. A navigation menu below the header contains links for 'Welcome to AARC', 'Roadmap', 'Work Packages', 'Blog', 'AARC Infoshare', 'Documents', 'Meetings', 'Wiki', 'About', and 'Admin'. The main content area is divided into four columns, each with an icon, a title, a short description, and a call-to-action button. The columns are: 'Architecture' (orange cube icon), 'Training and Outreach' (green circle icon), 'Policy Harmonisation' (brown paper plane icon), and 'Pilots' (black smartphone icon).

**Architecture**  
AARC has designed an architecture to help e-infrastructures and research communities to enable secure, scalable and interoperable federated access to their resources.  
Discover AARC blueprint architecture

**Training and Outreach**  
AARC is producing different training modules and information packages on federated access and AARC results.  
Learn how AARC can help you

**Policy Harmonisation**  
AARC is working with research communities, e-infrastructures and identity federations to deliver a common policy framework for integrated AAI.  
Discover AARC best practices

**Pilots**  
AARC has built a testing environment to test technical and policy results based to address research communities requirements.  
View AARC results in the real world

Thank you  
Any Questions?



<http://aarc-project.eu/>



© GEANT on behalf of the AARC project.  
The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 653965 (AARC).