Authentication and Authorisation for Research and Collaboration

AAI Overview

Training for EPOS

Agenda

- Authentication and authorization processes
- Externalizing authentication (LDAP)
- Federated authentication (SAML and OIDC)
- Federations in R&D
- The inter-federation: eduGAIN

Access to applications

When you manage an application, you must guarantee it is **accessed only by people that have the right** to use it.

To achieve successfully this goal, two distinct processes need to be implemented:

- **authentication** of a user
- enforce of **authorization** rules for the user

Authentication (AuthN)

- Authentication is the **act of confirming the truth** of an attribute of a single piece of data or entity (the user of an application, for instance).
- Example (in the real world): authenticating the Mona Lisa.

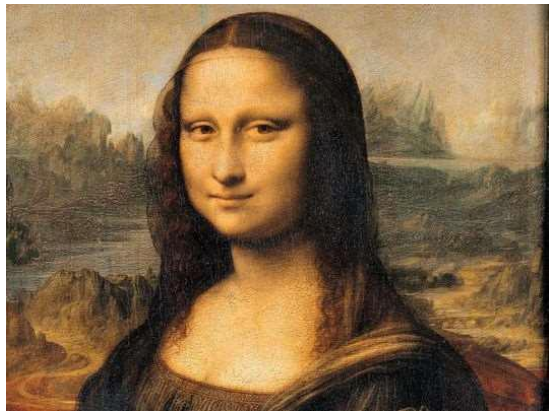


- In the digital world we tend to simplify the confirmation by using **username and password** *(the assumption is that password is known only by the intended user, so specifying the right password you're demonstrating you actually are who you pretend to be)*.

Authorization (AuthZ)

- Authorization is the function of **specifying access rights** to resources related to information security and computer security in general and to access control in particular.
  - More formally, ”to authorize” is to define an access policy.

- Example: going to a concert.

Authorization (AuthZ)

- In the digital world, defining the access rules user by user can be impractical.
- Authorization is often implemented with the so called **Role-Based Access Control** (RBAC)
  - **users are pooled** in groups based on their organizational role (e.g. payroll manager, project group A, …)
  - access **rights** are then **associated to roles**
- When a user access an application:
  - authenticates himself/herself
  - activate one or more roles (depending on the groups of belonging)
  - access application/services by leveraging RBAC authorization

Managing AuthN and AuthZ

As we have seen, an application to deal with authentication and authorization has to manage the following information:

1. **usernames** and associated **passwords**, *to authenticate users and verify they are who they pretend to be;*
2. institutional **roles**, *to describe the roles within the group or organization (used for RBAC);*
3. user **groups**, *to pool together users that have the same role in the organization (groups are associated to roles);*
4. **access policies**, *rules in the form of (role name -> access right) to describe which operation each role is entitled to perform and which not inside the application.*

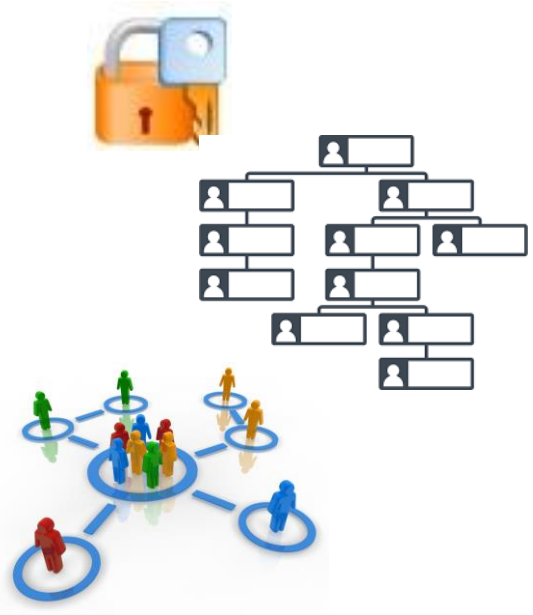

Externalizing authentication

For simplicity, and not to duplicate information, usually a **Directory** is used to collect username, password, roles and groups for the whole organization.

Directory services play an important role by allowing the sharing of information about users, systems, networks, services, and applications throughout the network.

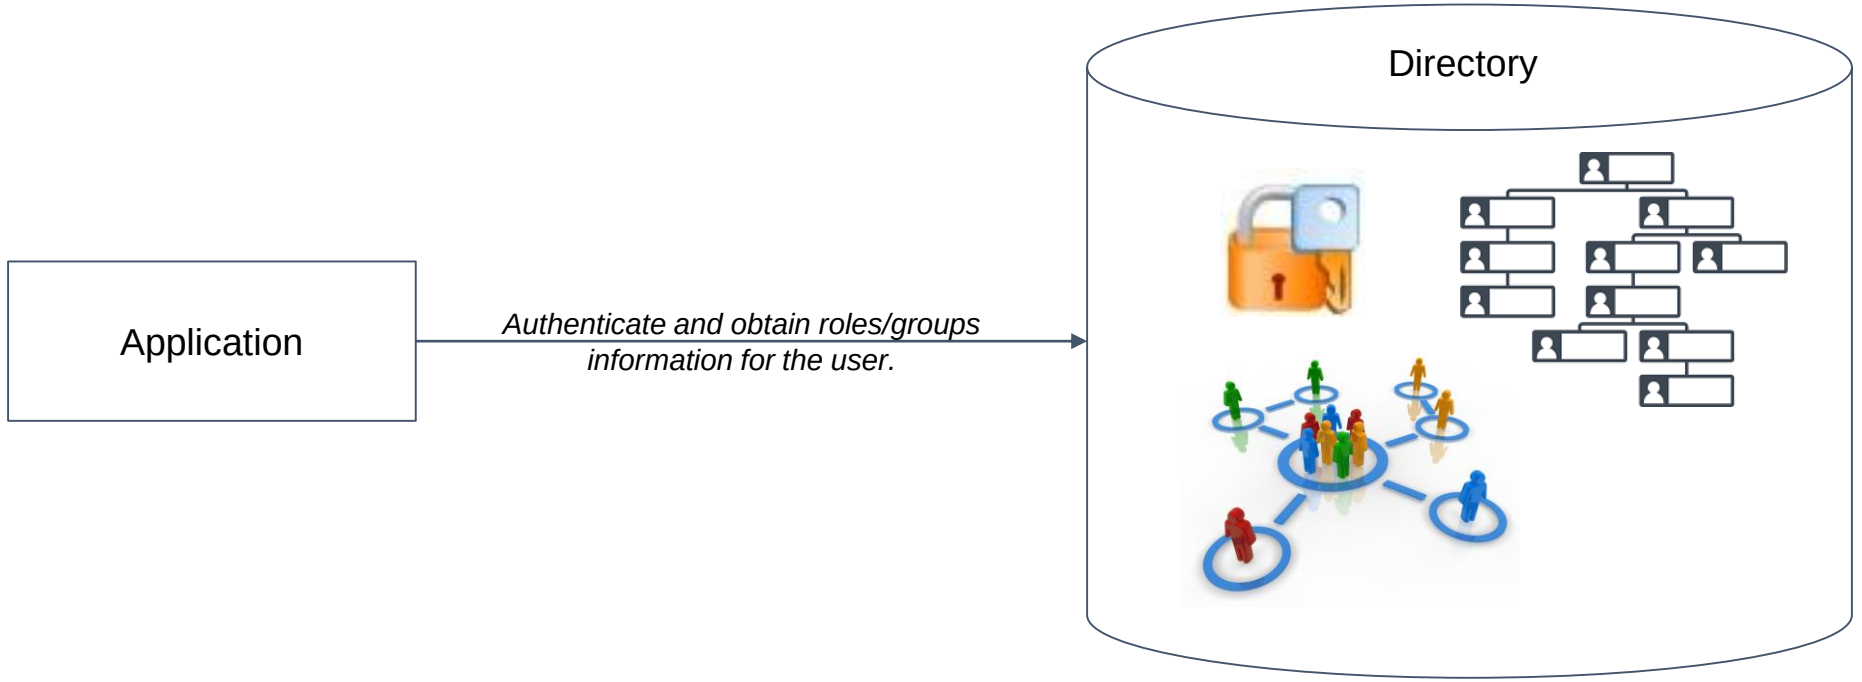

LDAP

The **Lightweight Directory Access Protocol** (LDAP) is an open, vendor-neutral, industry standard application protocol for accessing and maintaining distributed directory information services. LDAP is specified in a series of IETF RFCs.

Information in LDAP are structured as a **tree representing the organizational structure**.

Groups, people and users are then represented as nodes or leaves of the tree.

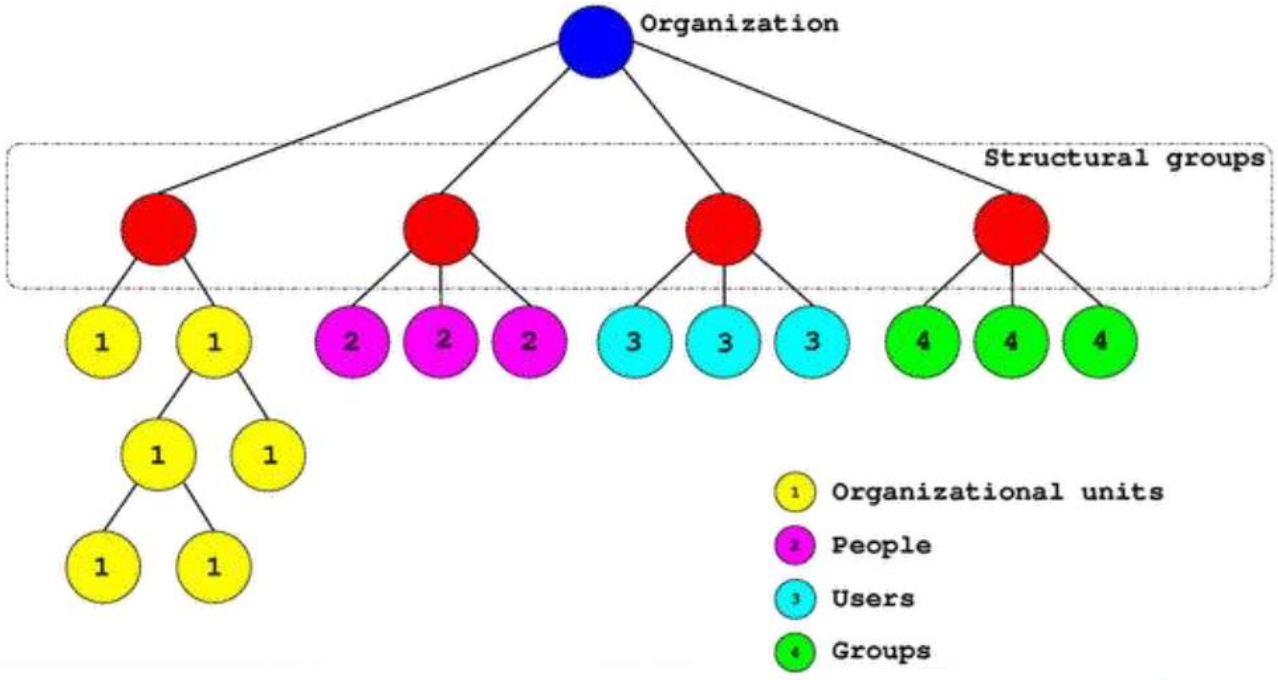

LDAP - Example

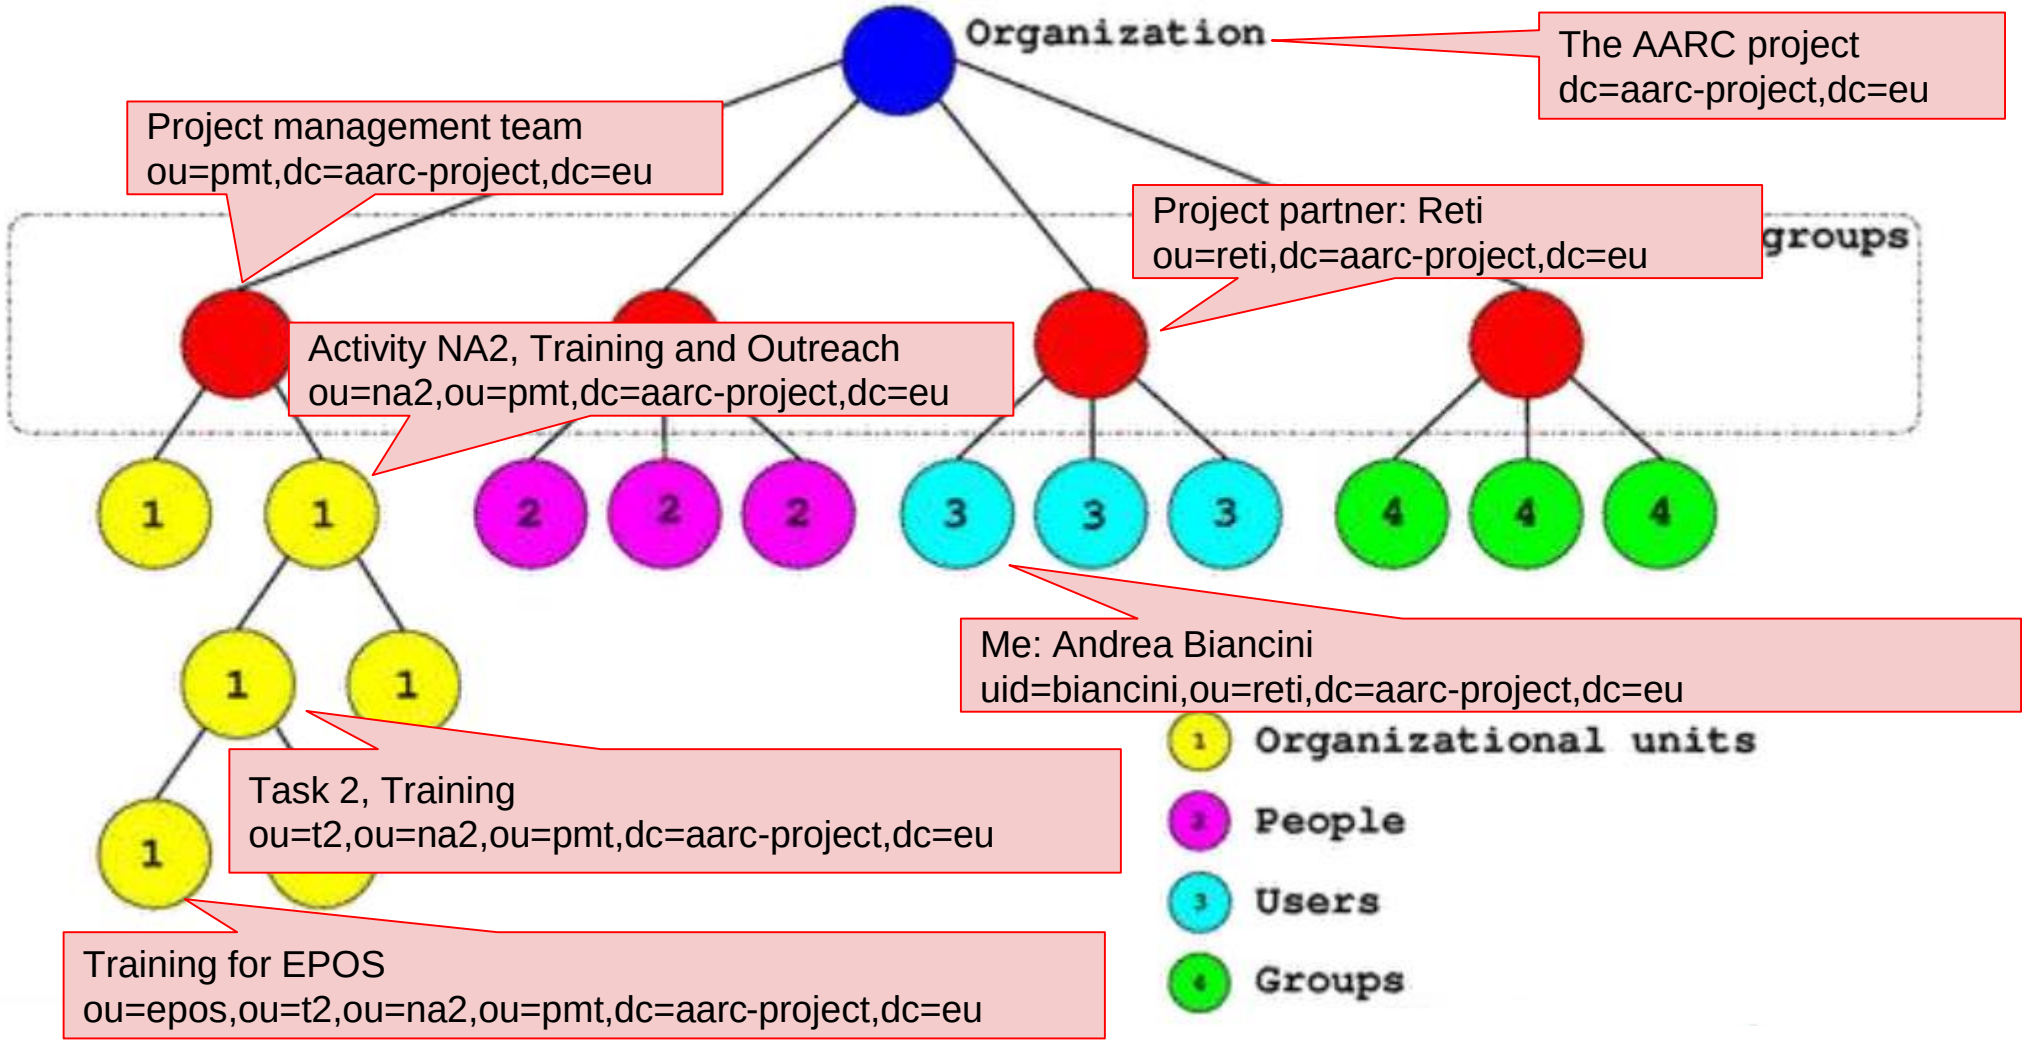

The AARC project
dc=aarc-project,dc=eu

Project management team
ou=pmt,dc=aarc-project,dc=eu

Project partner: Reti
ou=reti,dc=aarc-project,dc=eu

Activity NA2, Training and Outreach
ou=na2,ou=pmt,dc=aarc-project,dc=eu

Me: Andrea Biancini
uid=biancini,ou=reti,dc=aarc-project,dc=eu

Task 2, Training
ou=t2,ou=na2,ou=pmt,dc=aarc-project,dc=eu

Training for EPOS
ou=epos,ou=t2,ou=na2,ou=pmt,dc=aarc-project,dc=eu

Federated authentication

LDAP is widespread and quite always used to maintain authentication information of an organization. Most often, however, applications do have **users belonging to different organizations**: users that are authenticated, registered and associated to groups in LDAPs out of the domain of the organization servicing the application.

To permit these users to access the application, we use a **federated identity**:

- users can authenticate on different identity provider (IdP) services on the network;
- the different IdPs use similar protocols and user definitions so that applications can deal with users belonging to different organization in a similar manner.

Federated authentication

The objective of the AAI is, in a nutshell, to **simplify inter-organizational access** to resources. With a single login, for instance, a researcher can access applications at multiple organizations (universities or research institutions).

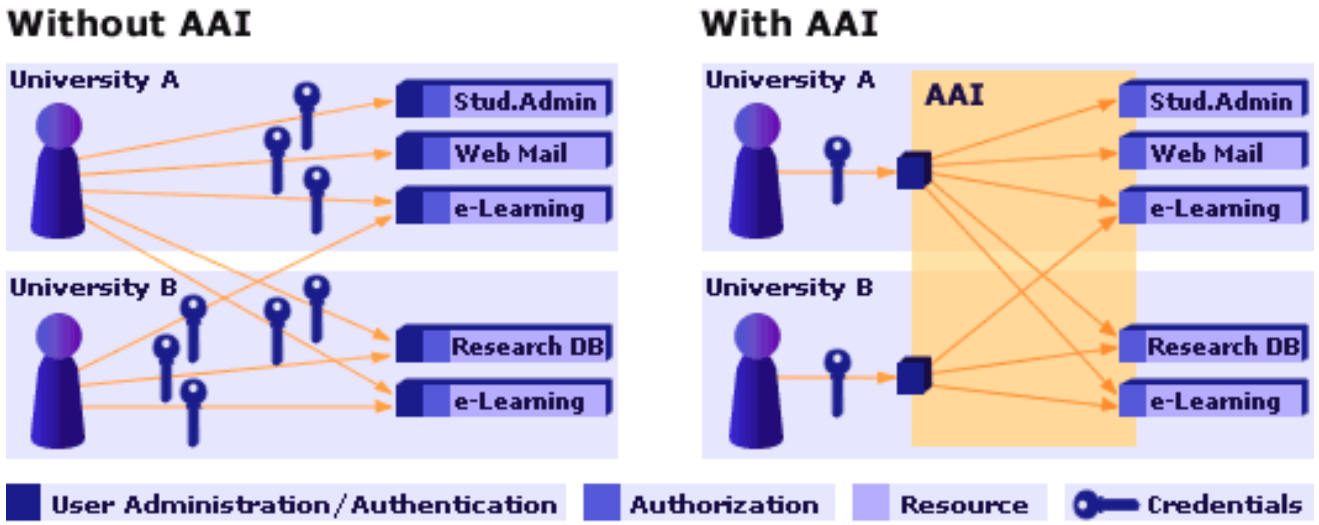

Benefits of federated authentication

- **A user registers only once** - namely with his/her so-called home organization to which the user is affiliated. This Home Organization is responsible for maintaining the user related information and provides the user with the credentials. Home Organizations can be institutions like universities, libraries, university hospitals etc.
- **Authentication** is always **carried out by** the **user's Home Organization**, which can also provide additional information about the user to the Resource upon Resource's request and user's consent.
- **All AAI-enabled Resources are available to a user** with a single set of credentials. At the same time, there is no need for Resource operators to register new users, because they get the required information directly from the user's Home Organization.
- An **access control decision** (authorization) is made by the Resource based on the retrieved information about the user.

What is a Federation

A federation is a **collection of organizations** that agree to interoperate under a certain rule set. Federations will usually define trusted roots, authorities and attributes, along with distribution of metadata representing this information.

In general each organization participating in a federation operates:

- one **Identity Provider** (IdP) for their users, and
- any number of **Service Providers** (SP) or applications.

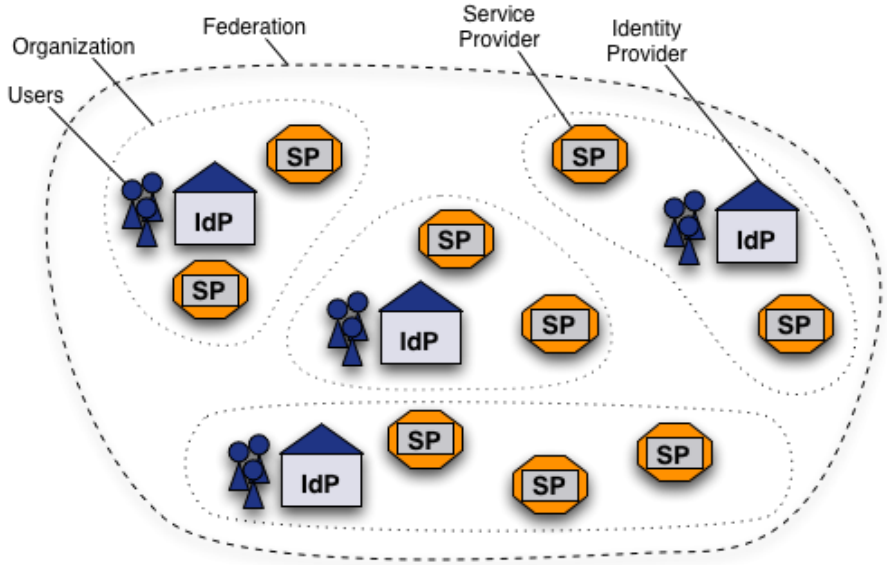

SAML

SAML (Simple Assertion Markup Language) is a **standard** that **facilitates the exchange of security information**. Developed by OASIS, SAML is an **XML-based framework**. SAML enables different organizations (with different security domains) to securely exchange authentication and authorization information.

To create a SAML infrastructure:

- an **IdP** must be **installed on top of each organization directory** to permit user authentication in the federation
- an **SP** must be **installed on top of each application** to consume authentication and authorization information obtained from the federation

OpenID Connect

OpenID Connect (OIDC) is a **standard** that **facilitates the exchange of security information**. Published by the OpenID Foundation, OIDC is a framework that uses **REST APIs** and **JSON** format. OIDC, as SAML, enables different organizations (with different security domains) to securely exchange authentication and authorization information.

To create an OIDC infrastructure:

- an **OP** must be **installed on top of each organization directory** to permit user authentication in the federation (similarly to IdP in SAML)
- an **RP** must be **installed on top of each application** to consume authentication and authorization information obtained from the federation (similarly to SP in SAML)

Why these two frameworks?

In our communities, we need to consider at least these two framework:

- **SAML** was implemented by different Research and Education Networks all over the world and has a very **strong legacy in R&E**;

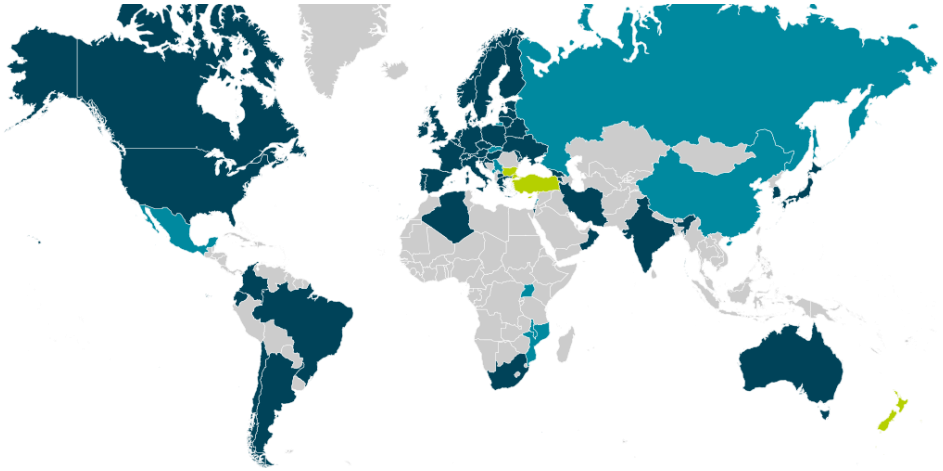

- **OIDC** is sustained by **many internet companies**.

Google KDDI Microsoft
NRI The Office of the National Coordinator for Health Information Technology ORACLE
PayPal PingIdentity Symantec.
verizon

Federations in R&D

A **federation operator** is an organisation that operates an identity federation. Operation typically includes at minimum:

- Collecting, processing and republishing federation metadata *(metadata permits to create a trust between IdPs and SPs to communicate securely)*
- Common policies and legal frameworks that all federation participants adhere to
- Guidelines and deployment instructions to operate services in the federation
- Helpdesk to assist with deploying services and debugging issues

Most academic federations are operated by the **national research and education network** (NREN). These organisations typically also operate the network connecting the universities and research organisations within a country.

The inter-federation

NRENs usually operates federation within a country.
To scale to a global level, R&E introduced the concept of **interfederation**.

Interfederation takes place if a **user from one federation accesses a service which is registered in another federation**. eduGAIN is the most known and largest academic Interfederation service to exchange trusted identity information across boundaries of (national) identity federations.

eduGAIN

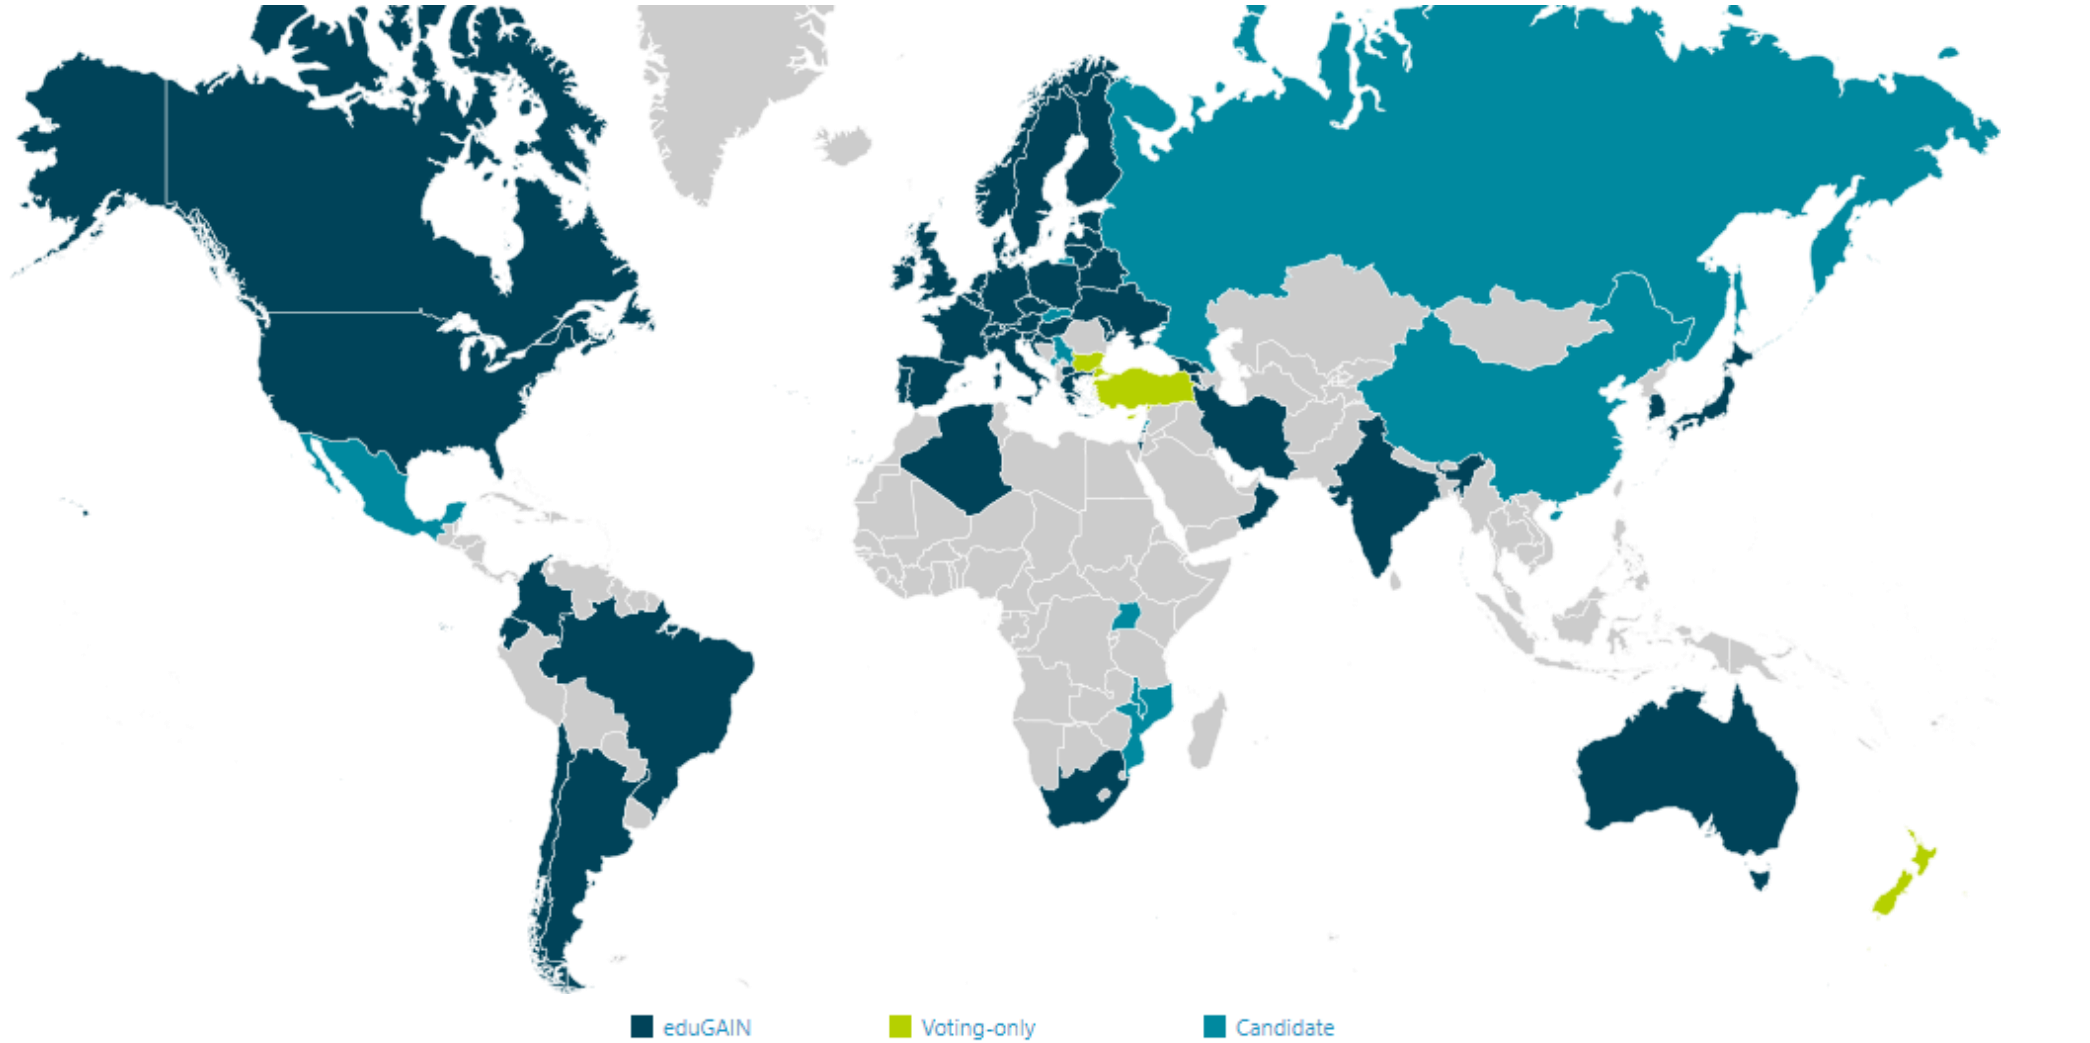

What we have learnt

- ★ What is authentication and what is authorization
- ★ Which are the most common ways of doing authentication and authorization
- ★ What is a directory for an organization
- ★ How the information used for AAI are stored in a directory
- ★ How to create a federation to create trust between different organizations
- ★ What is an inter-federation

Thank you

Any Questions?

https://aarc-project.eu

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