EGI Infrastructure and AAI solutions

Alessandro Paolini

Operations Officer, EGI Foundation

About me…

- master degree in physics
- worked in GRID related projects (EGEE, EGEE2, EGEE3, EGI-Inspire) for 11 years at INFN
  - Italian GRID infrastructure cohordination, operations activities, central services management, support to resource centres and users
- Since Nov 2015 at EGI Foundation as Operations Officer (EGI-Engage, AARC/AARC2 , EOSC-HUB)
  - Management of the daily operations of the EGI infrastructure
  - FItSM foundation and advanced levels
  - Piloting AAI implementations

www.egi.eu

EGI Federation (I)

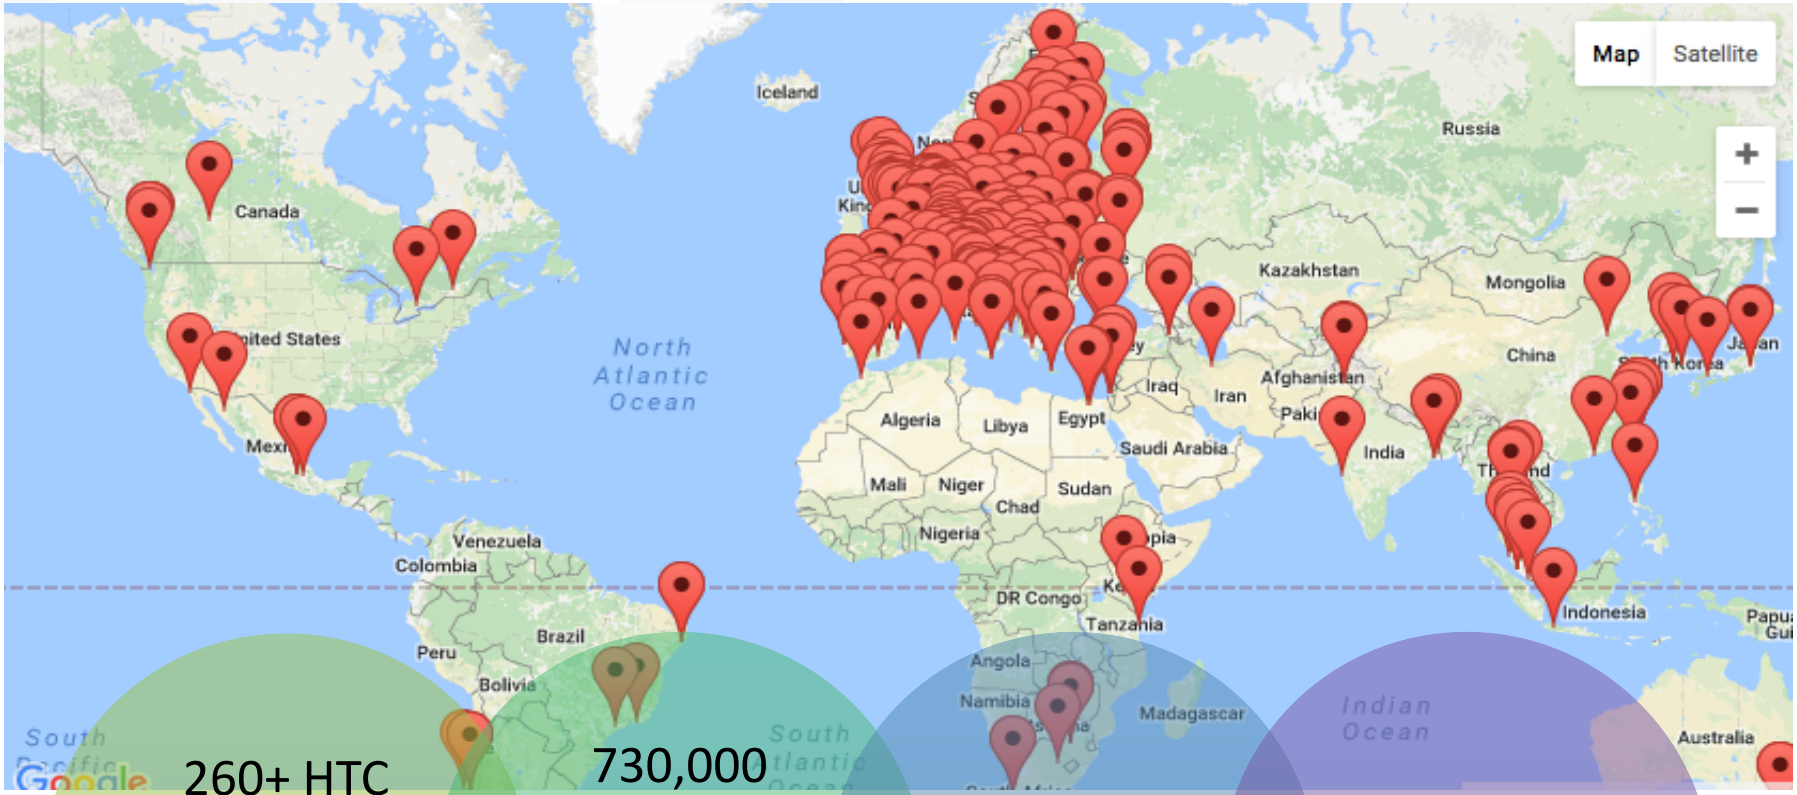

260+ HTC providers

(20 Cloud providers)

730,000 CPU cores

650 PB storage

1.7 Million jobs/day

2.8 Billion CPU hours/year

EGI Federation (II)

- Resource centres spread across 47 countries
  - United by a mission to support research activities
- EGI delivers services to support scientists, international projects and research infrastructures
- Integrates community, private and/or public infrastructures into a scalable data/computing platform for research.
- Uses federated identities, authentication and authorization
- Ensures interoperability of scientific applications and data across multiple providers bringing distributed computing to data
- EGI's mission is to create and deliver open solutions for science and research by federating digital capabilities, resources and expertise between communities and across national boundaries.

Research disciplines

Since 2010, EGI has been delivered unprecedented data analysis capabilities to more that tens of thousands of researchers from over hundreds virtual organisations covering many scientific disciplines

Physical Science

Medical and Health Sciences

Natural Sciences

Social Sciences

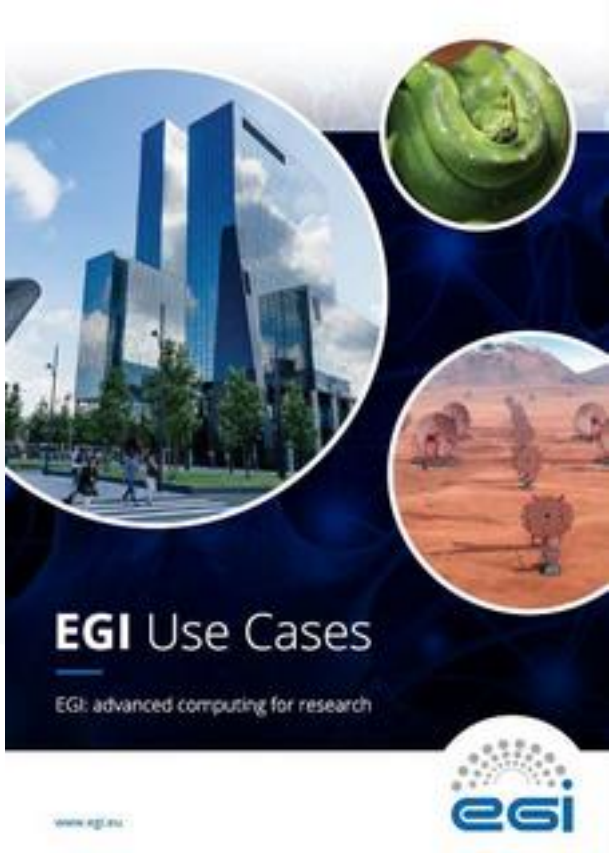

https://www.egi.eu/use-cases/research-stories/

Role for the EGI AAI services

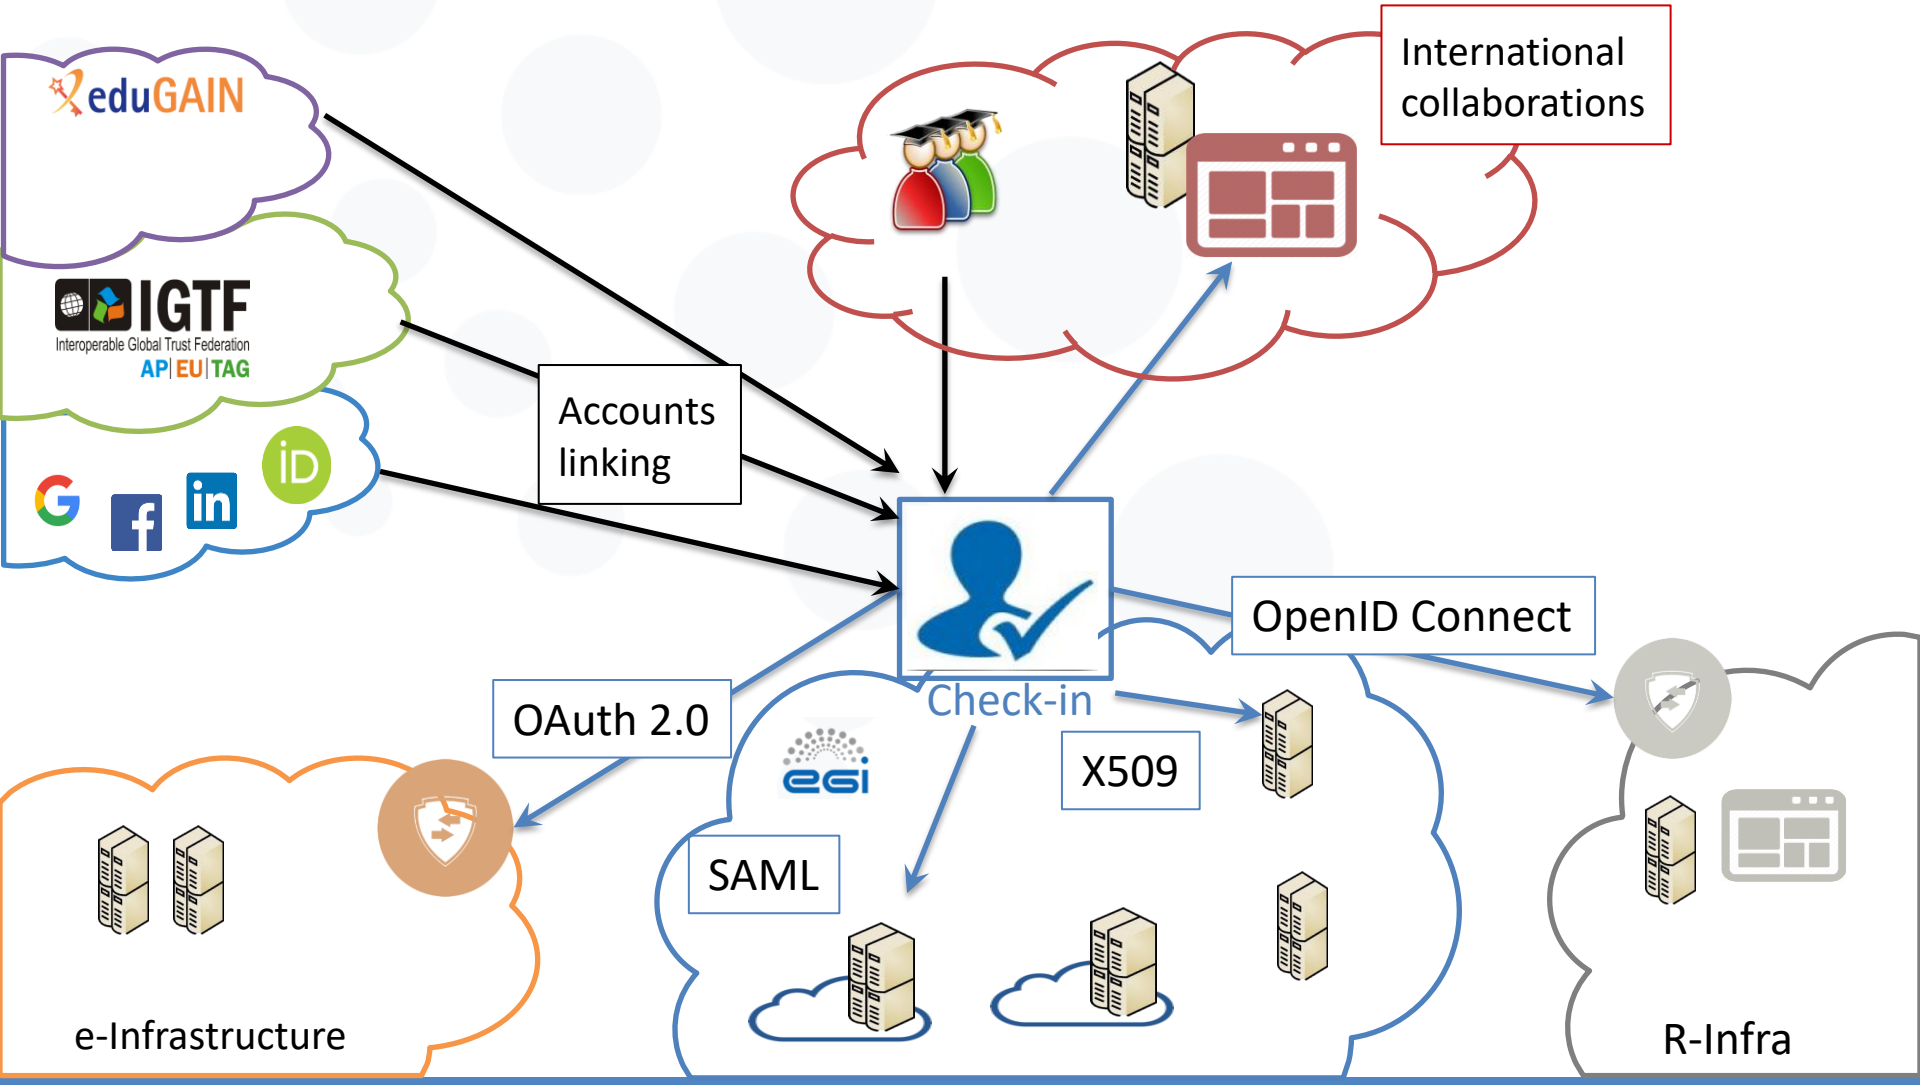

AAI Services: Check-in Overview

Check-in provides a reliable and interoperable AAI solution for the EGI service providers federation, and external service providers. It enables single sign-on to services through eduGAIN identity providers and other institutional or social media credentials

- Check-in development, lead by GRNET, has been supported by the EGI-Engage project. The EGI Council supports the long-term operations of the service.
- Check-in has been developed in close collaboration with the AARC project, and it implements the recommendations of the AARC Blueprint Architecture and Policy Framework (https://aarc-project.eu/architecture/)

ISO standards and policies

- In March 2017, EGI became the first European-wide publicly-funded e-infrastructure to be certified to ISO standards

- EGI has a mature and complete set of security policies and the processes to enforce them
  - Extended with Check-in specific policies:
    - ✓ Check-in acceptable usage policy
    - ✓ Check-in data protection policy
    - ✓ Agreement documents to integrate non-EGI and non-eduGAIN SPs and IdPs and maintain the compliance
- Check-in is supported by the EGI CSIRT and other security groups

Different entities with a common goal

- user communities, research infrastructures, federations, identity providers, and e-Infrastructures
  - sufficiently alike for sharing some common policy frameworks
- an open research commons benefits hugely if based on:
  - a set of a harmonized policy frameworks and best practices

Policy harmonisation in AARC project

- Activity focused on operational and security aspects and policies
  - delivers a set of recommendations and best practices to implement a scalable and cost-effective policy and operational framework for the integrated AAI
- Selected elements:
  - Assurance Level baseline (alongside a self-assessment tool)
  - Security Incident Response in federated environments
    - https://refeds.org/SIRTFI
  - Recommendations for Research and e-Infrastructures to Build Sustainable Services: http://go.egi.eu/cikhw
  - Scalable policy negotiation: adoption of 'entity categories' and the development of a policy framework for IdP-SP-proxies
    - "Scalable Negotiator for a Community Trust framework in Federated Infrastructures" or "Snctfi": http://go.egi.eu/hfymy
  - Protection of (mainly personal) data that is generated as a result of infrastructure use (e.g. in accounting)

Requirements for the minimal assurance profile (I)

In a federated AAI, the user's Home Organisation issuing and managing user's credentials determines the assurance level available for the user identity.

Currently, in research and education, there is no well-established assurance level framework. The assurance levels available depend on the policies and practices of the user's Home Organisation and the identity federation to which it belongs.

- requirements for a minimal assurance level which is still relevant for low-risk research use cases
- It is not expected that a Home Organisation must comply with these requirements for all of its user accounts. Instead, the Home Organisation must be able to tag the compliant accounts and logins

1. The accounts in the Home Organisations must each belong to a known individual person
   - Accounts must not be shared
   - need for a reliable audit trail
2. Persistent user identifiers (i.e., no re-assignment of user identifiers)
   - The identifier must have the property that it is never re-assigned, i.e. recycled to another person
   - Currently the identifier eduPersonPrincipalName is lacking this property

6 Requirements for the minimal assurance profile (II)

3. Documented identity vetting procedures (not necessarily face-to-face)
   - The Home Organisations must have a documented identity vetting process for its user accounts
4. Password authentication (with sufficient quality and some good practices)
   - For the low-risk research, authentication with passwords is sufficient
5. Departing user's eduPersonAffiliation must change promptly
   - within one month of the departure at most
6. Self-assessment (supported with specific guidelines)
   - Audits based on a complete and specific framework

Accounting and data protection

- The infrastructures need to be able to process data and meta-data about the users and their interaction with the systems
  - essential for accounting and for assigning use data to allocations
  - to be able to follow up on incidents in the infrastructure
- For providing recommendations, it is necessary to:
  - make an inventory of the relevant use cases
  - identify the types of data generated within the infrastructure as a result of its use
  - Identify the respective roles of the participants in the infrastructure with respect to data protection
  - take into account the wide diversity in laws and regulations related to personal data protection throughout Europe
    - Abide by EU General Data Protection Regulation

Accounting and data protection (II)

- Recommendations and template policies for the processing of personal data by participants in the pan-European AAI: http://go.egi.eu/jabrz
- EGI Policy on the Processing of Personal Data:
  - https://documents.egi.eu/document/2732
- Service providers need to provide a privacy policy explaining which kind of personal data are handled and why, and for how long they will be kept
  - Appointment letter with each partner stating roles and responsibilities (data controller/data processor)

Thank you for your attention.

Questions?

www.egi.eu

www.egi.eu

Backup slides

AARC Blueprint Architecture

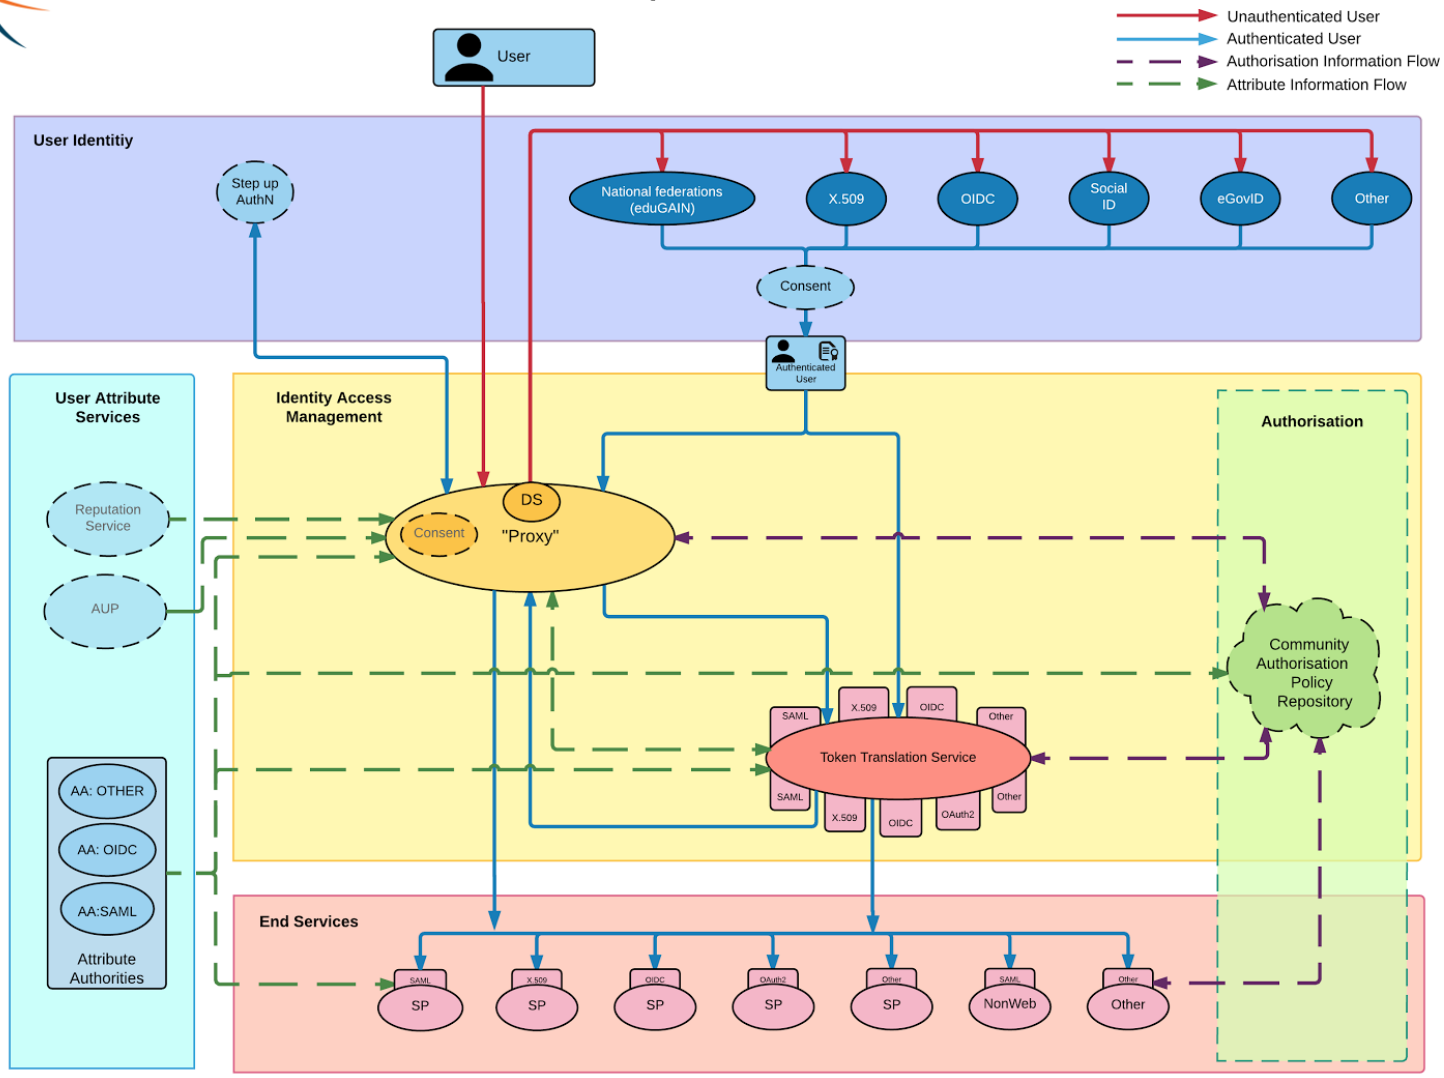

Deployment models

Check-in is offered in 3 deployment models:

- As a multi-tenant service running on EGI cloud infrastructure

- All the standard Check-in features
- Independent community management using COmanage or Perun
- Limited customisation of user-facing interfaces (e.g. community-specific themes for enrolment flows, group management)
- Limited customisation of AAI proxy behaviour

- As a dedicated service+ running on EGI cloud infrastructure
- As a dedicated service+ running on Community's cloud infrastructure

- Deployment of individual components
- Deployment of AAI platform as a whole
- Customisation of user-facing interfaces:
  - WAYF, enrolment, group membership UI
- Customisation of AAI proxy behaviour (e.g. attribute aggregation rules, service entitlements)
- Easy integration with the main Check-in instance, or other dedicated instances if necessary

www.egi.eu

Use cases

For communities operating their own AAI

Community's AAI connected to Check-in as an IdP Proxy to allow its users to access EGI services & resources

- ✓ Access EGI services without changing your authentication workflow

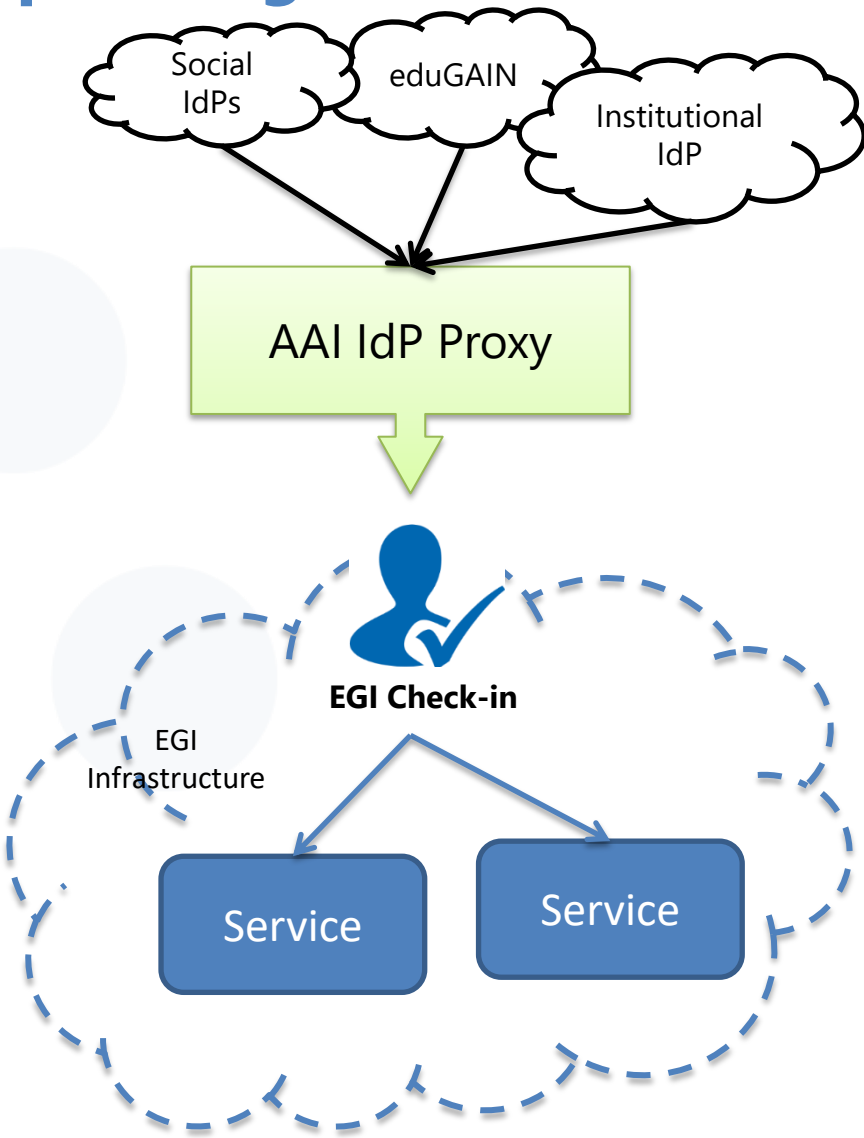

Examples: *ELIXIR Research Infrastructure - Check-in allows ELIXIR users to use their ELIXIR IDs to interact with relevant EGI services (Cloud, Configurations database, Applications on Demand*

For communities operating their own group management service

Community managing authorisation information about the users (VO/group memberships and roles) via their own group management service, which is connected to Check-in as an external attribute authority

- ✓ Check-in will handle the configuration of the IdPs and the aggregation of the attributes for the SPs
- ✓ No need to migrate the group management functionality to an EGI-specific attribute authority

Examples: *VOMS-managed VOs such as FedCloud and OPS, GOCDB attributes*

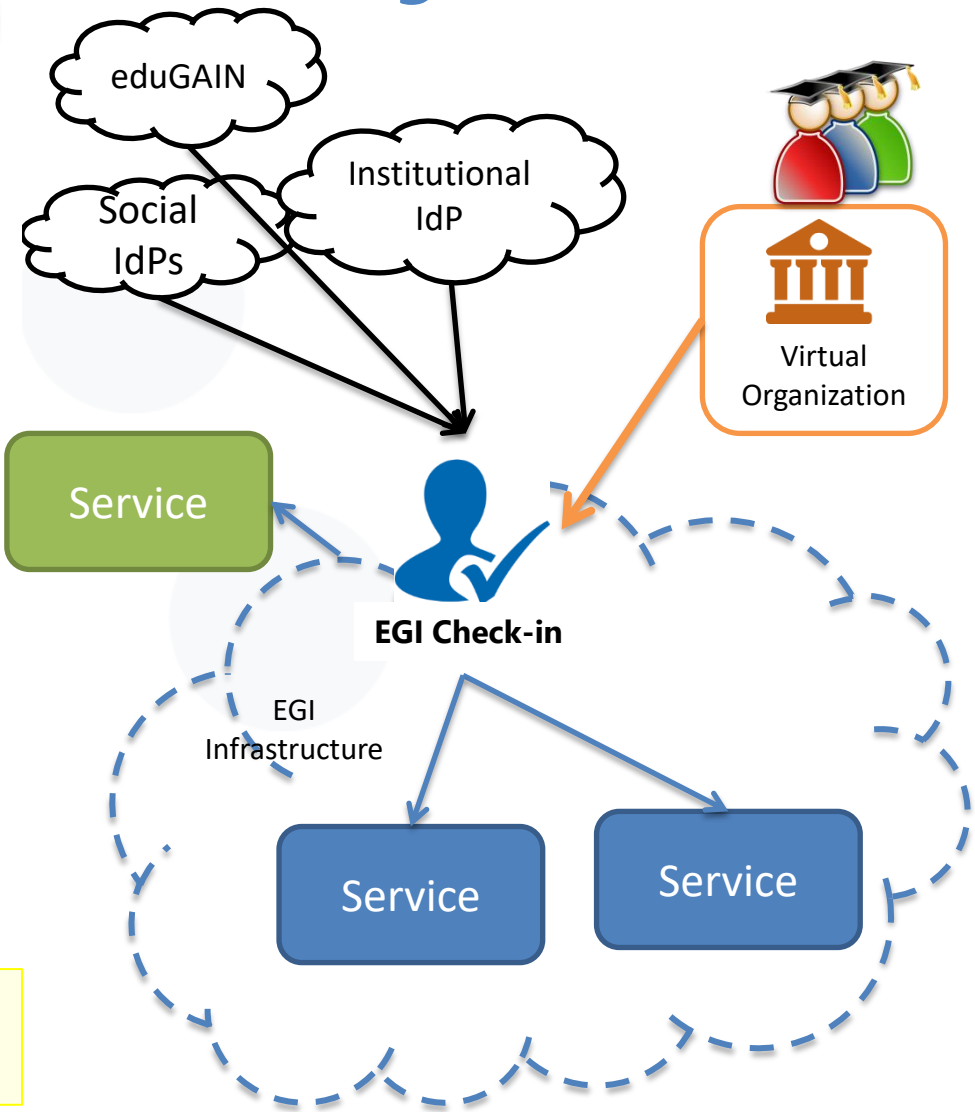

For communities in need of a ready-to-use group management solution

Communities that do not operate their own group management service can leverage the group management capabilities of the Check-in platform

- ✓ Ready-to-use solution
- ✓ Avoid overhead of deploying a dedicated group management service
- ✓ Support for multi-tenancy to allow authorised VO admins to manage the information about their users independently
- ✓ Easy connect to both EGI and non-EGI services

Examples: *Training and Long Tail of Science communities*

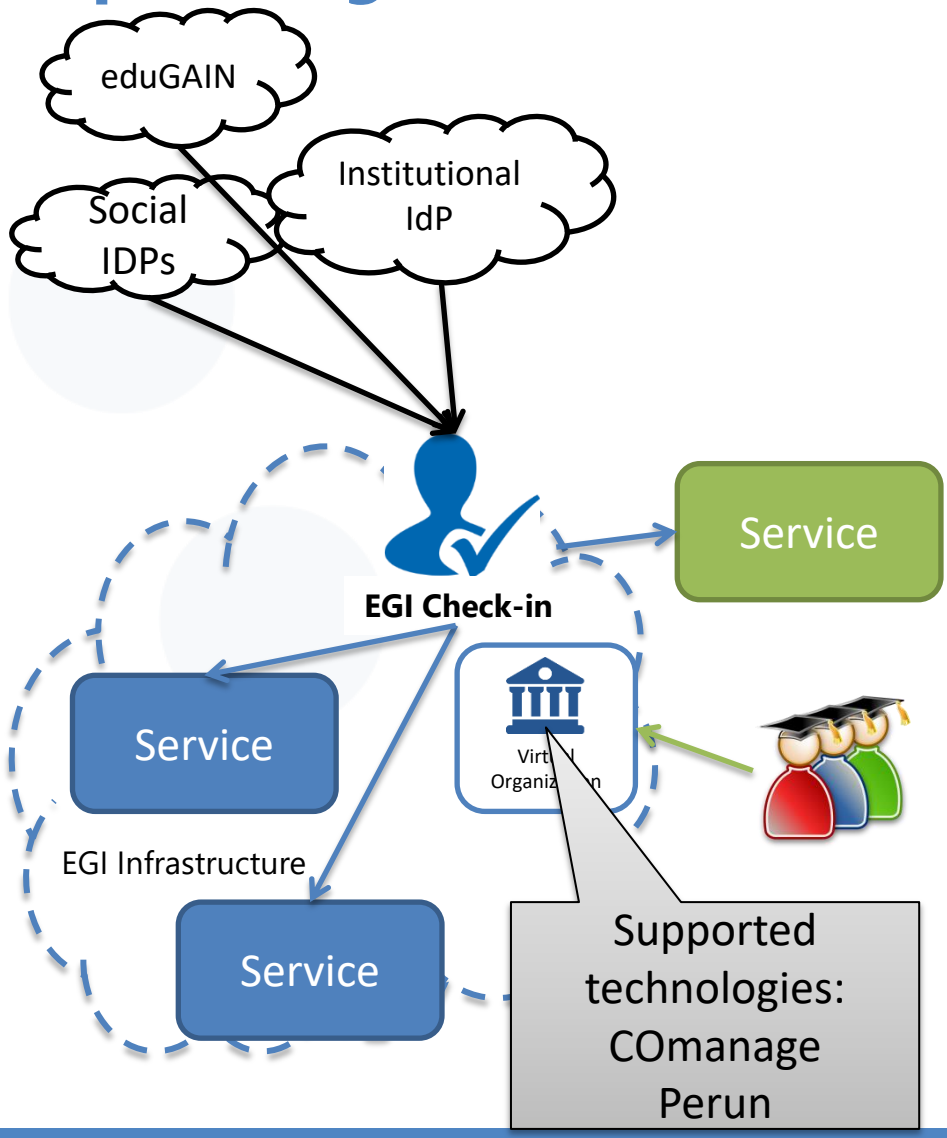

For service providers: AAI as a service

Check-in as an authentication proxy

- ✓ Enable login from institutional IdPs in eduGAIN and social media
- ✓ Minimal overhead for the service development
- ✓ All the other Check-in features are available for the SP: account linking, attribute aggregation, ..
- Prerequisites:
  - ✓ Service provider must accept EGI policies on data protection

Examples: *EDISON Community Portal*

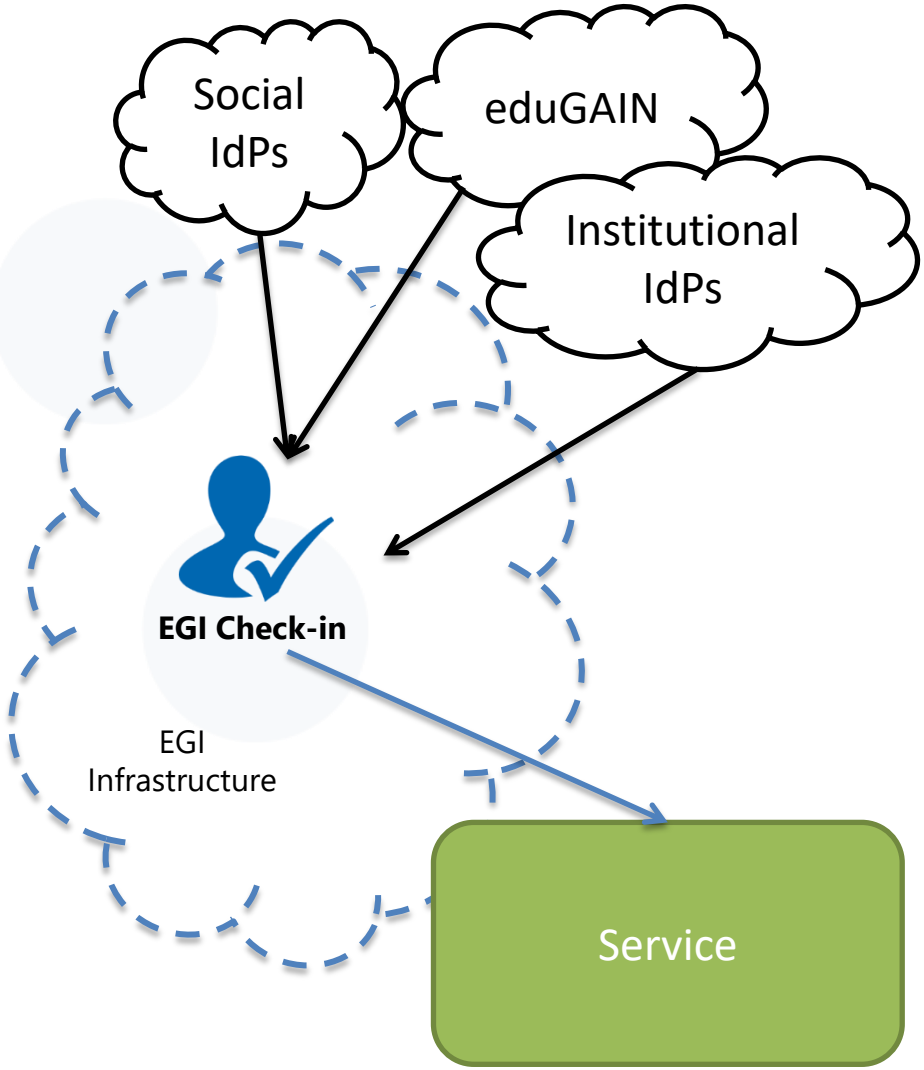

links

- www.egi.eu
- EGI policies: https://wiki.egi.eu/wiki/SPG:Documents
- https://aarc-project.eu/
- Differentiated LoA recommendations for policy and practices of identity and attribute providers, applicable to research use cases: http://go.egi.eu/wbrjj
- REFEDS Assurance framework: http://go.egi.eu/qxvzn