

Introduction the INDIGO IAM service

Andrea Ceccanti

INFN


andrea.ceccanti@cnafe.infn.it

JUNO Distributed Computing Meeting

July 3rd 2020



Brief intro to the WLCG AAI



WLCGG

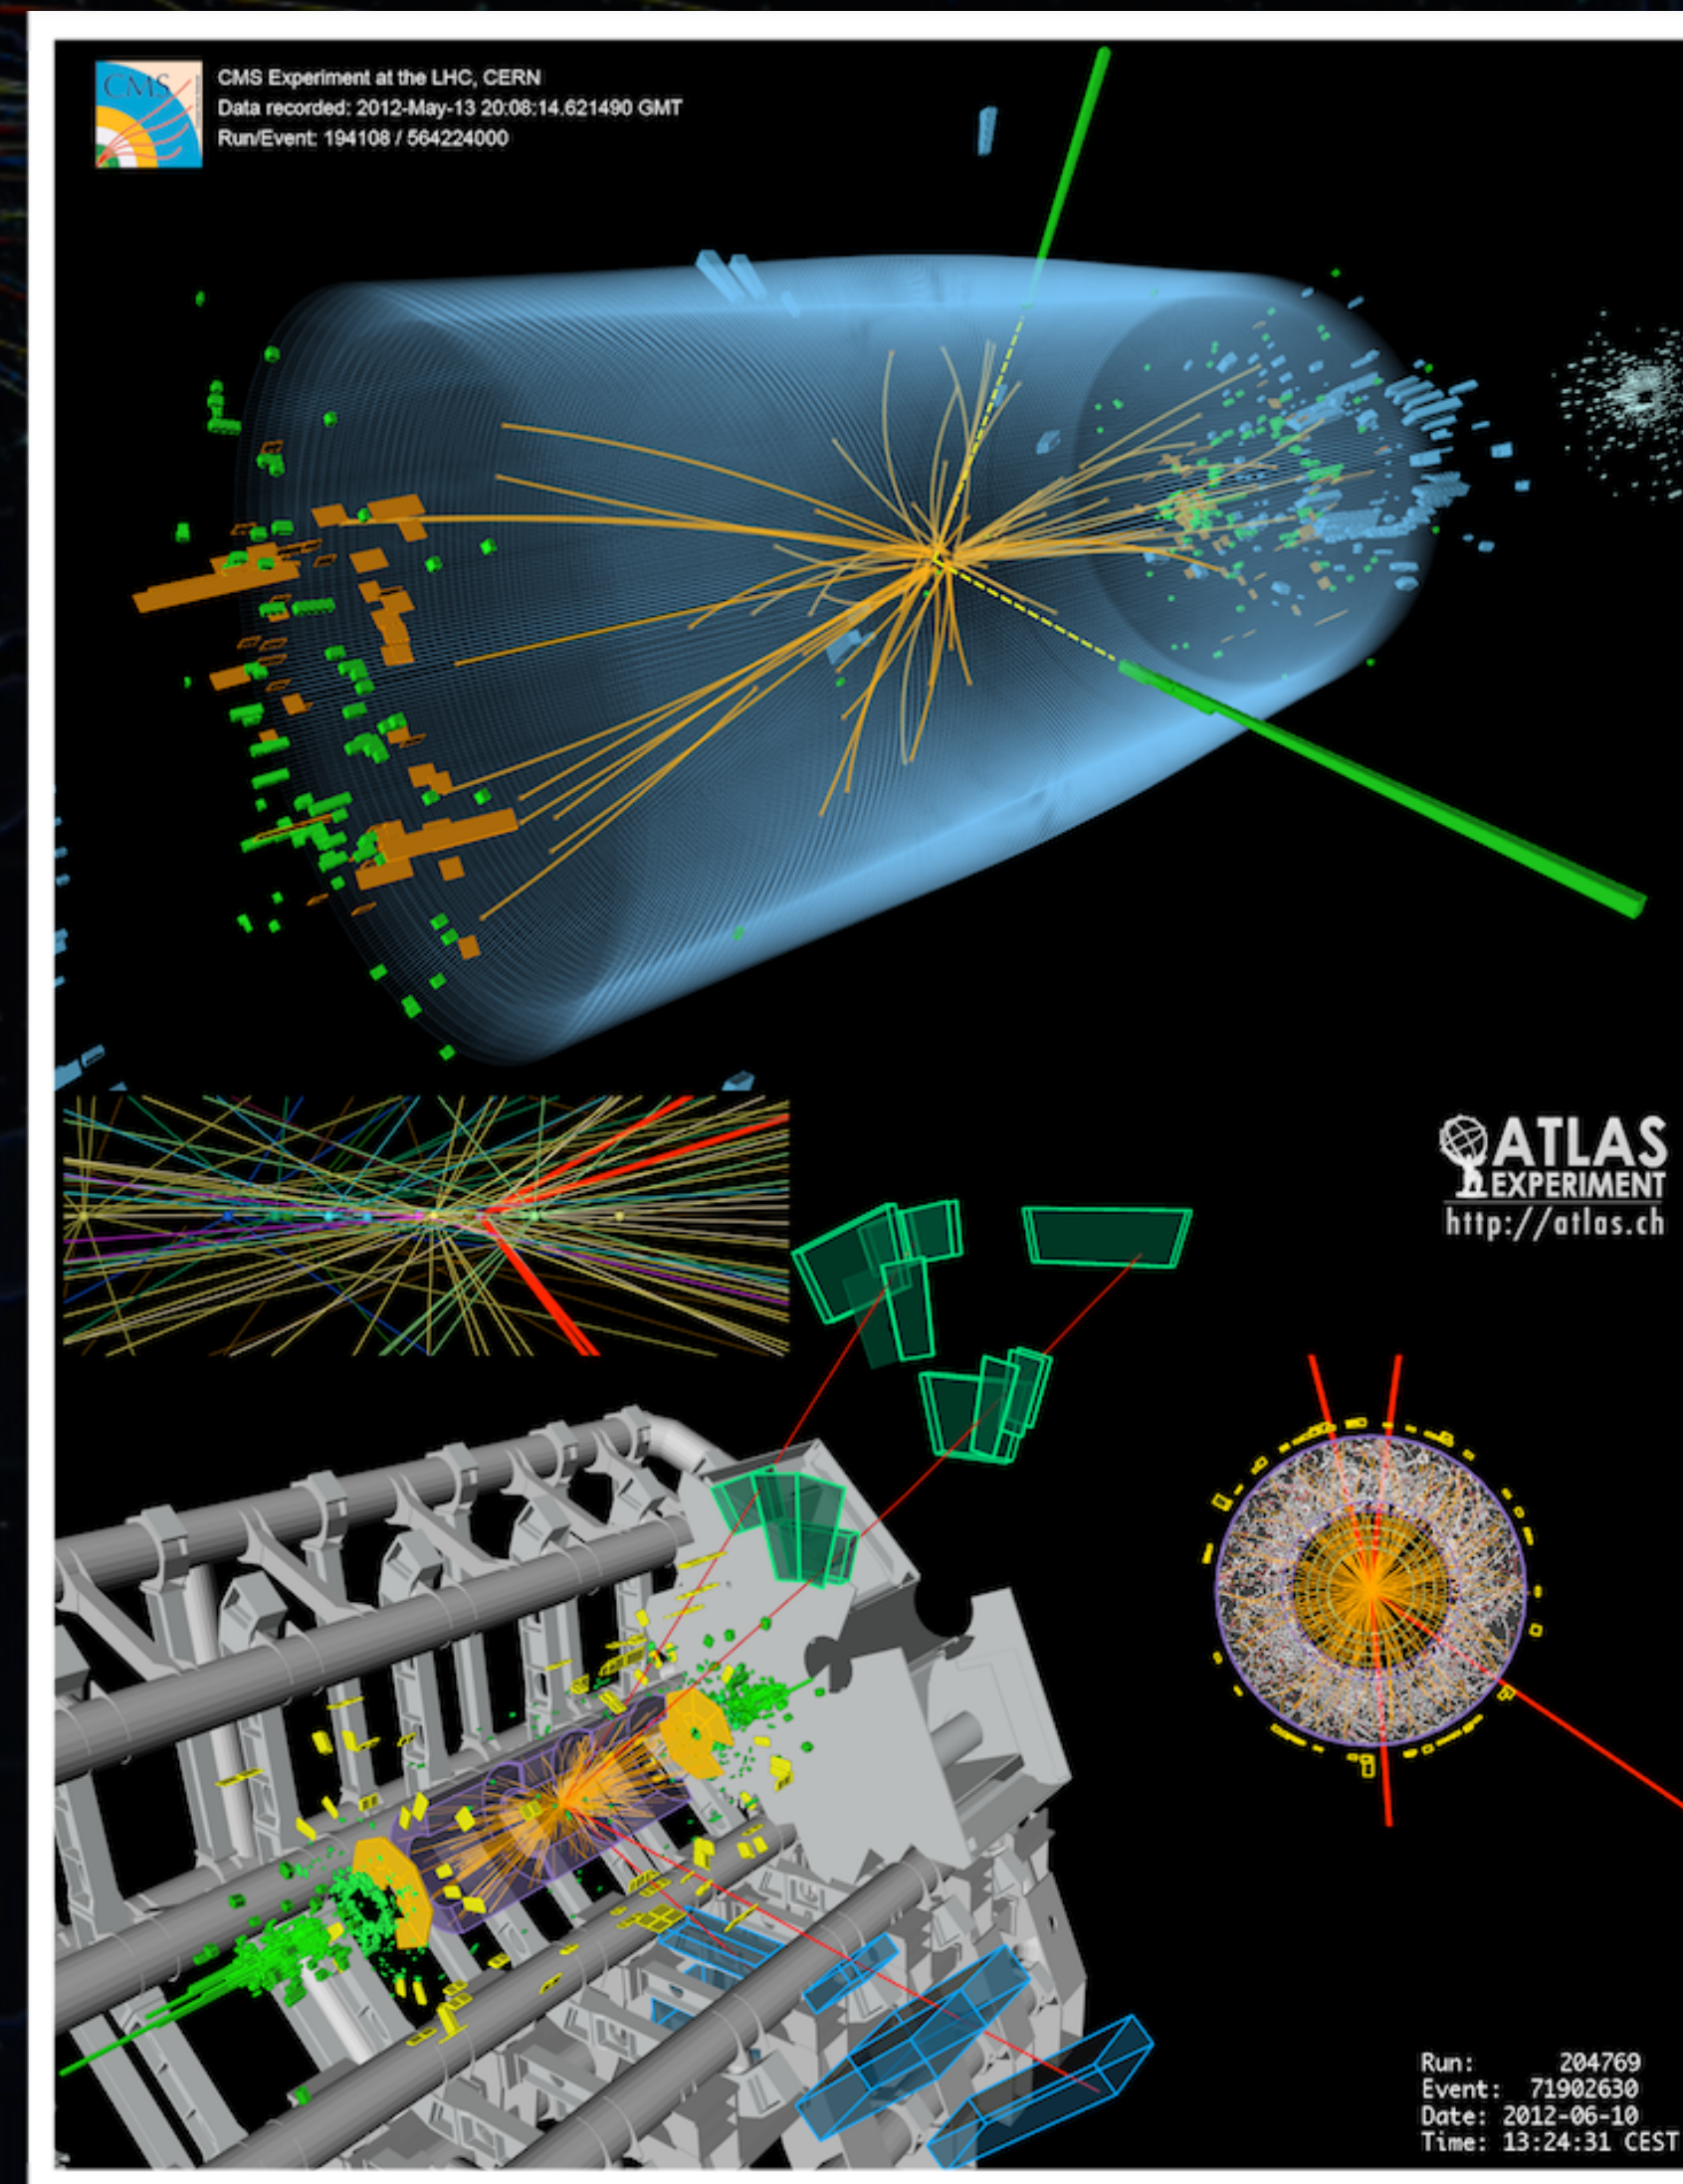
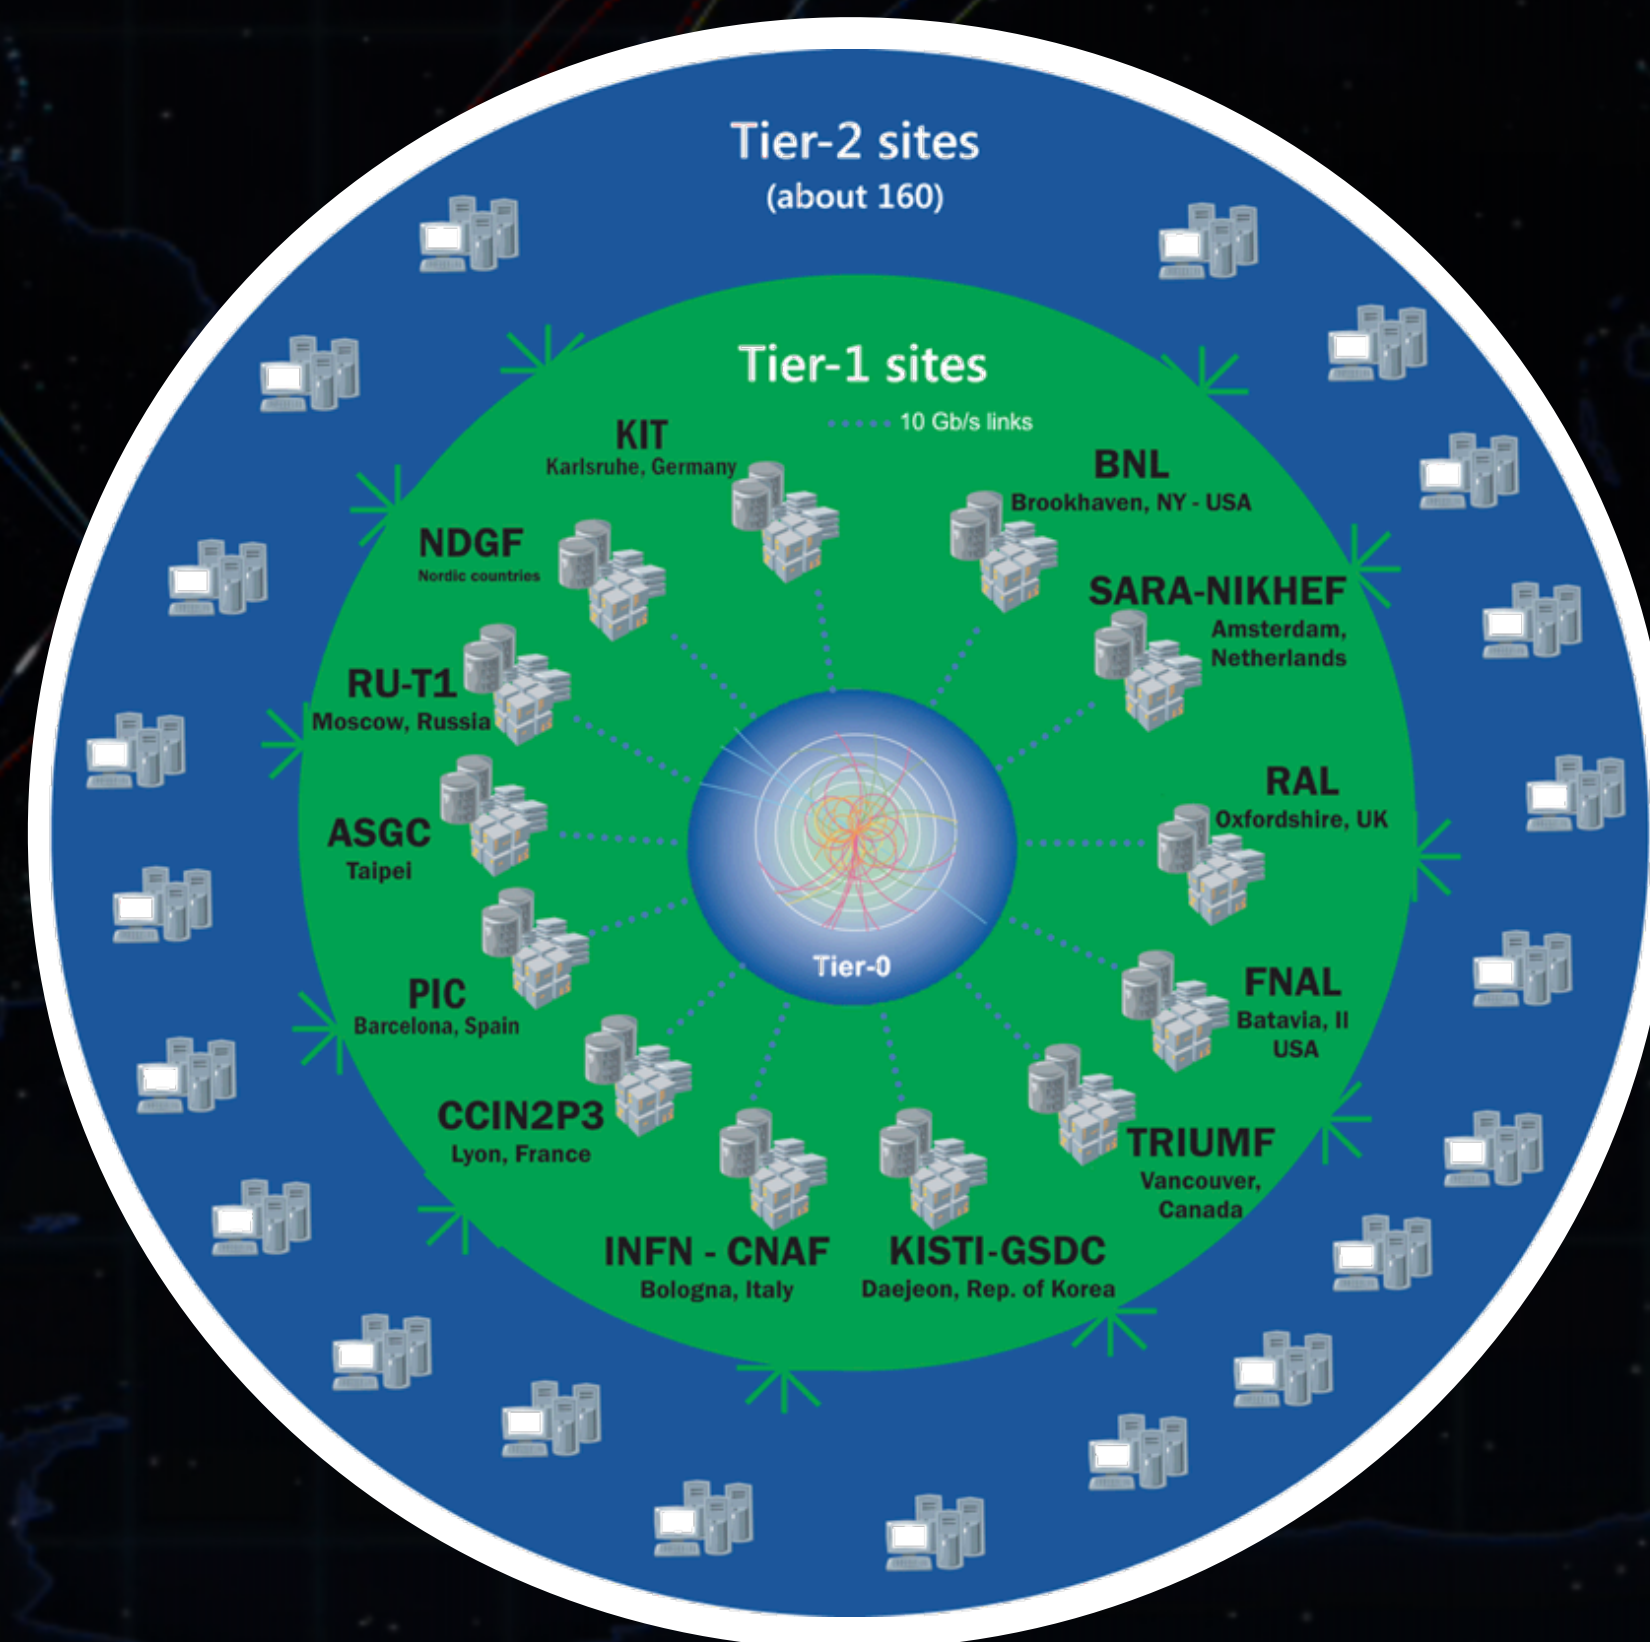
The Worldwide LHC Computing Grid (WLCGG), is a distributed computing infrastructure arranged in tiers – giving a community of over 12,000 physicists near real-time access to LHC data.

167 sites, 42 countries

~1M CPU cores

~1EB of storage

> 2 million jobs/day



Running jobs: 365118
Active CPU cores: 795836
Transfer rate: 18.35 GiB/sec

Virtual Organizations

WLCG brings together researchers from several institutions distributed all over the world.

Researchers typically collaborate in the context of scientific collaborations which represents **Virtual Organisations (VOs)**. **VOs**:

- group together users with a **common purpose**
- represents a **single integration point for infrastructure** providers (Grid sites, public clouds, etc...)
- provide **centralised management** of users enrolment and user lifecycle
- define their **authorization space** by organizing users in groups, assign them roles & other attributes

The WLCG AAI objectives (from 10K mt)

Controlled, secure sharing of computing and storage resources, provided across heterogeneous infrastructures (Grid, Cloud, HPC) in support of the WLCG experiments requirements

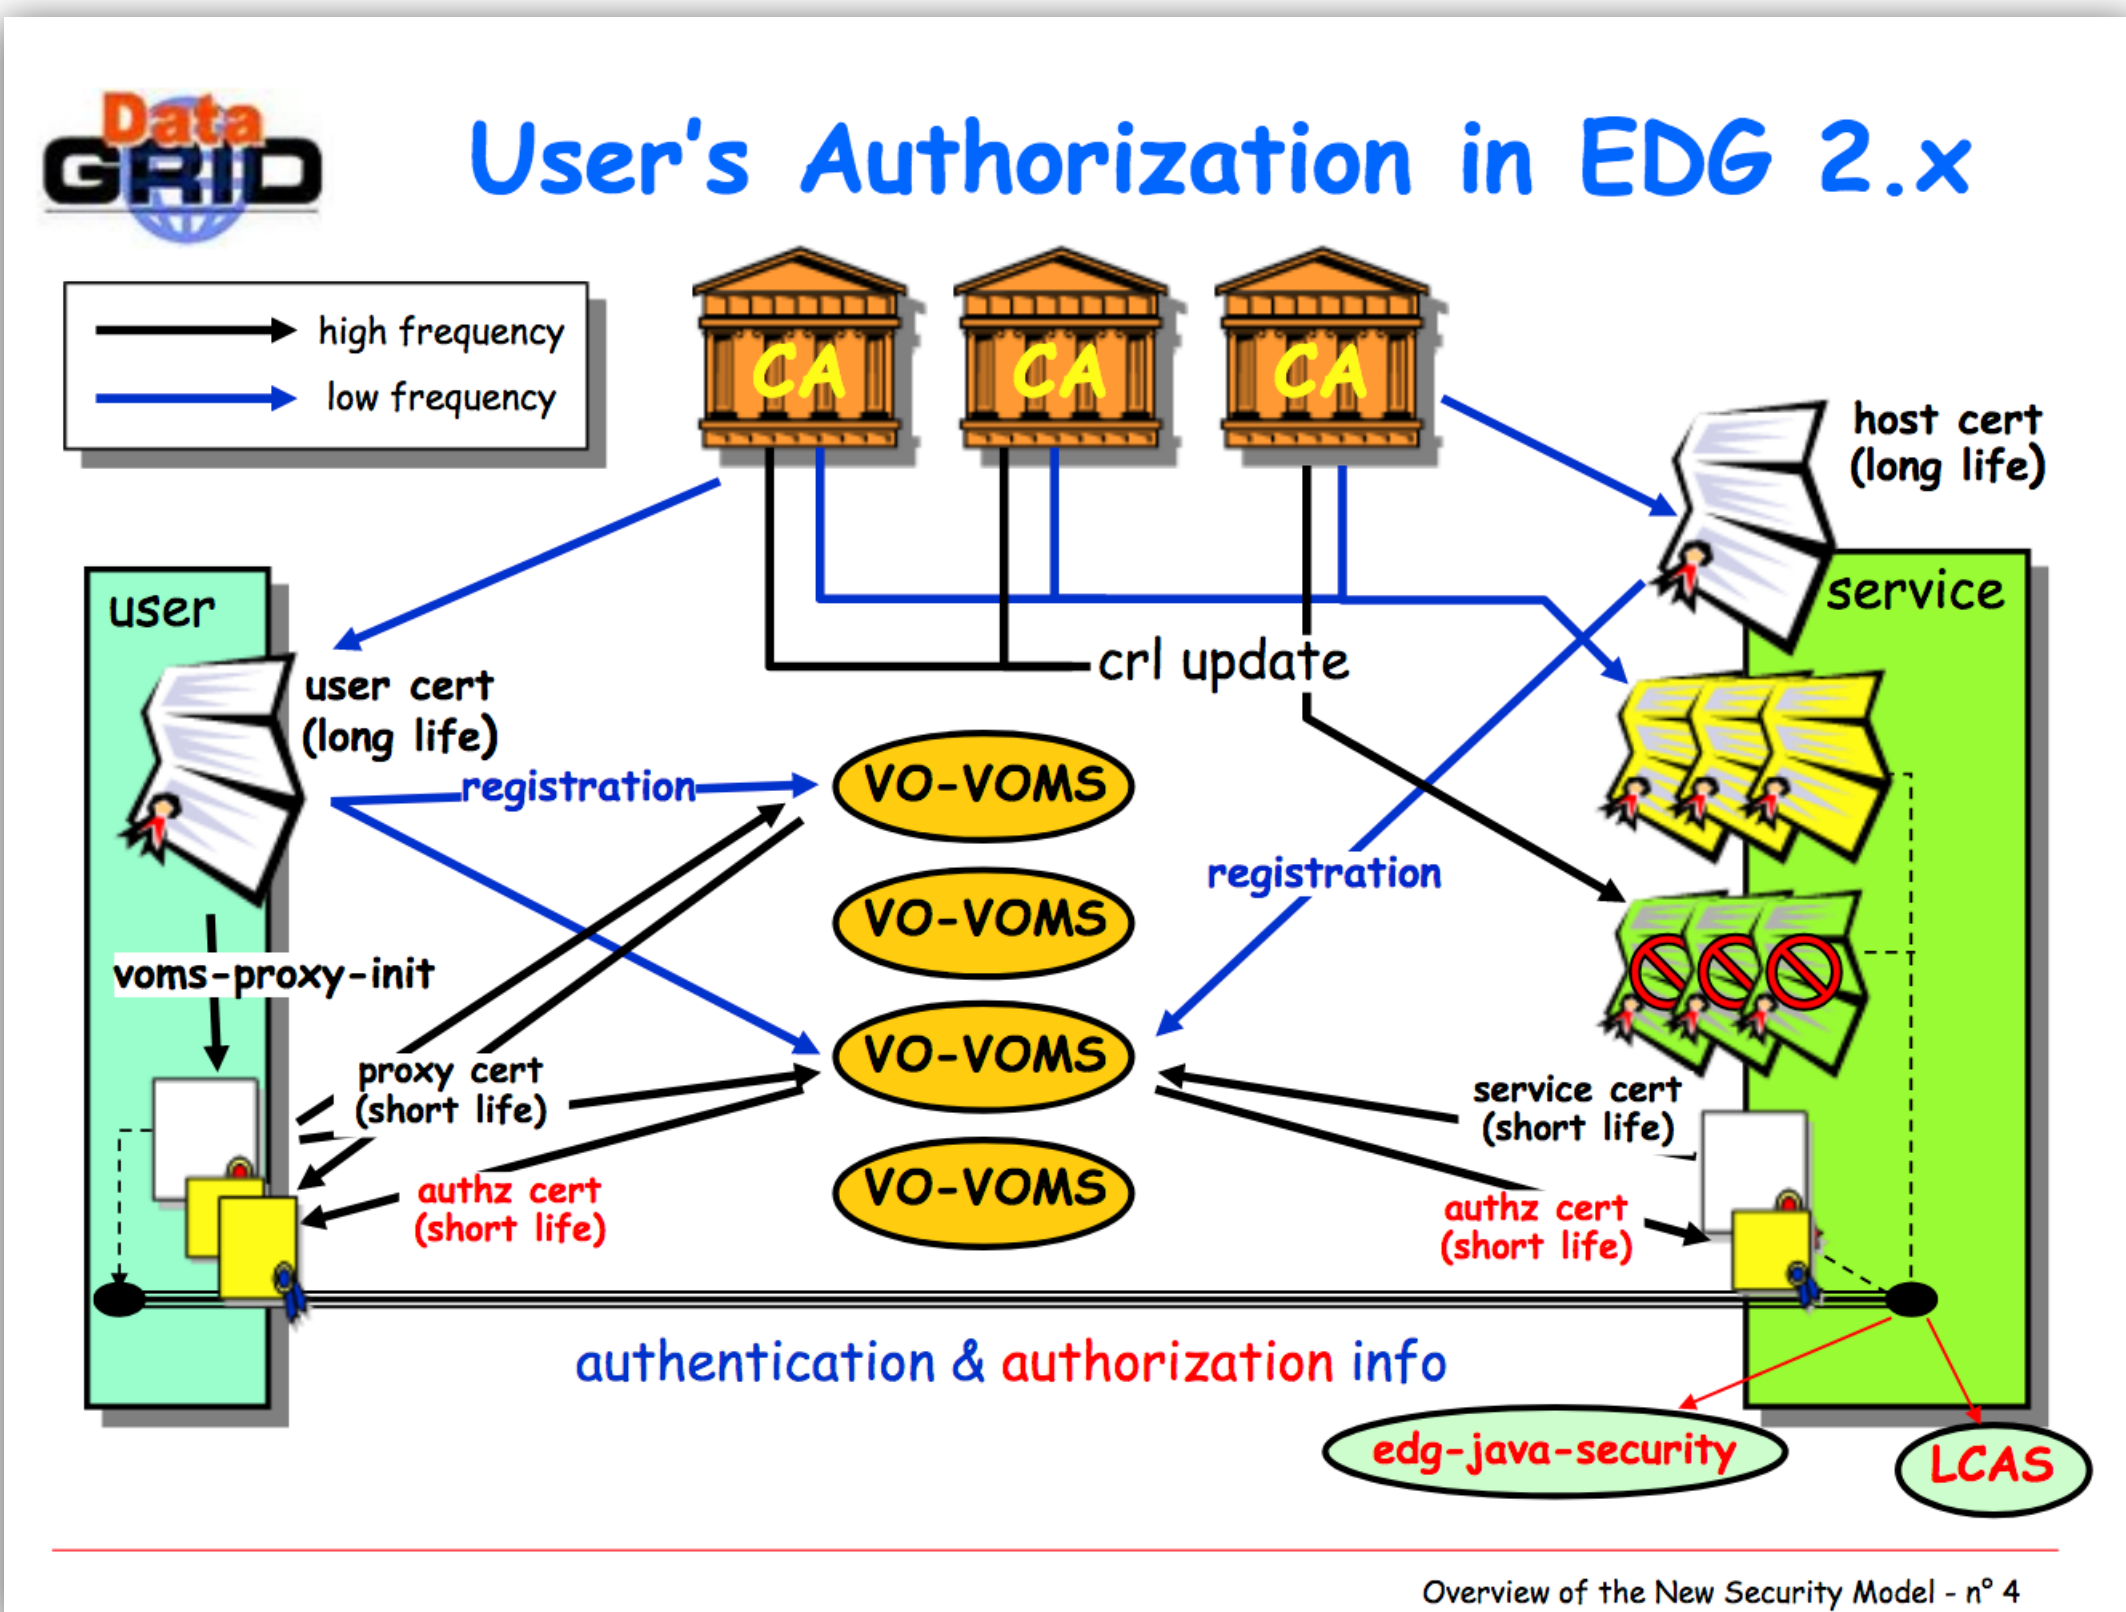
Consistent authentication and authorization across the infrastructure

VO-centric authorization model

Auditing, Traceability, Accounting

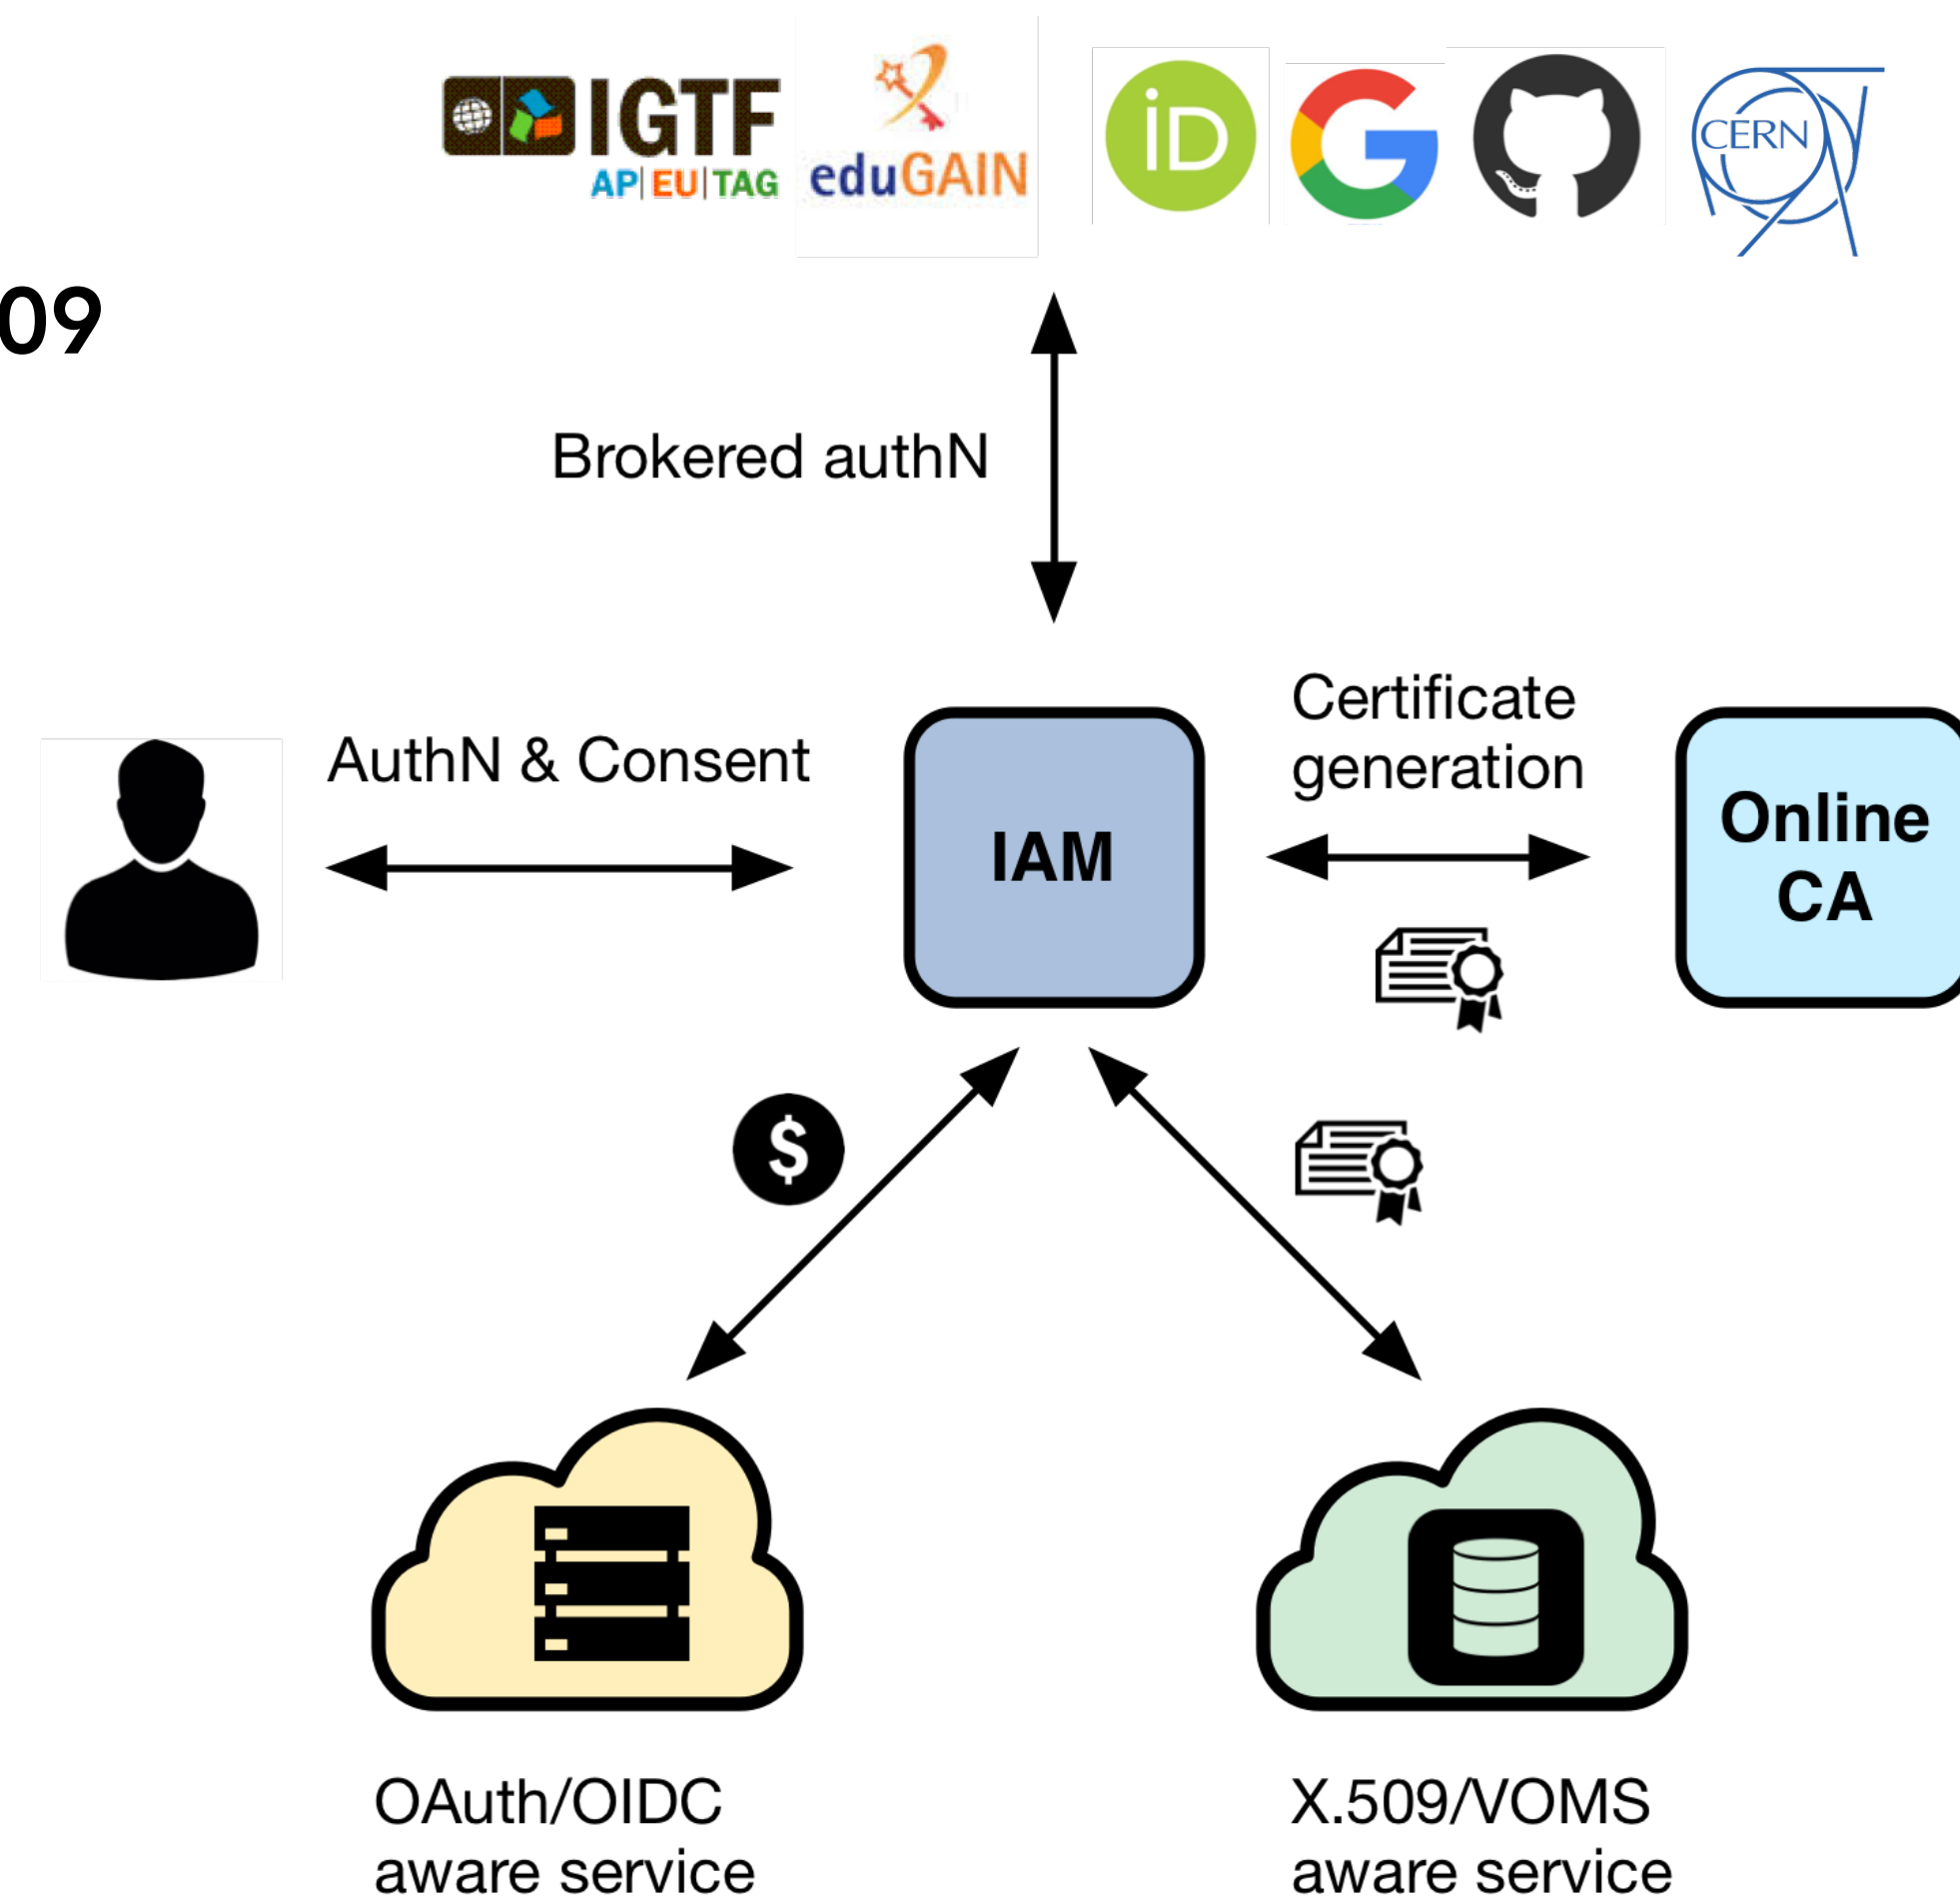

The WLCG AAI evolution

Current, X.509 based AAI



Future, token-based AAI

Move beyond X.509



Moving beyond X.509: main challenges

Authentication

- **Flexible**, able to accommodate various authentication mechanisms
 - X.509, username & password, EduGAIN, ...

Identity harmonization & account linking

- Harmonize multiple identities & credentials in a single account, providing a **persistent identifier**

Authorization

- **Orthogonal** to authentication, **attribute** or **capability-based**

Delegation

- Provide the ability for **services to act on behalf of users**
- Support for **long-running applications**

Provisioning

- Support provisioning/de-provisioning of identities to services/relying resources

Token translation

- Enable **integration with legacy services through controlled credential translation**

Moving beyond X.509: main challenges

Authentication

- **Flexible**, able to accommodate various authentication methods
 - X.509, username

Identity harmonization linking

- Harmonize multiple identities in a single account identifier

Authorization

- **Orthogonal** to authentication, **attribute** or **capability-based**

Delegation

- Provide the ability for **services to act on**

**Key challenge:
allow a gradual transition
to the new AAI!**

applications

provisioning of
resources

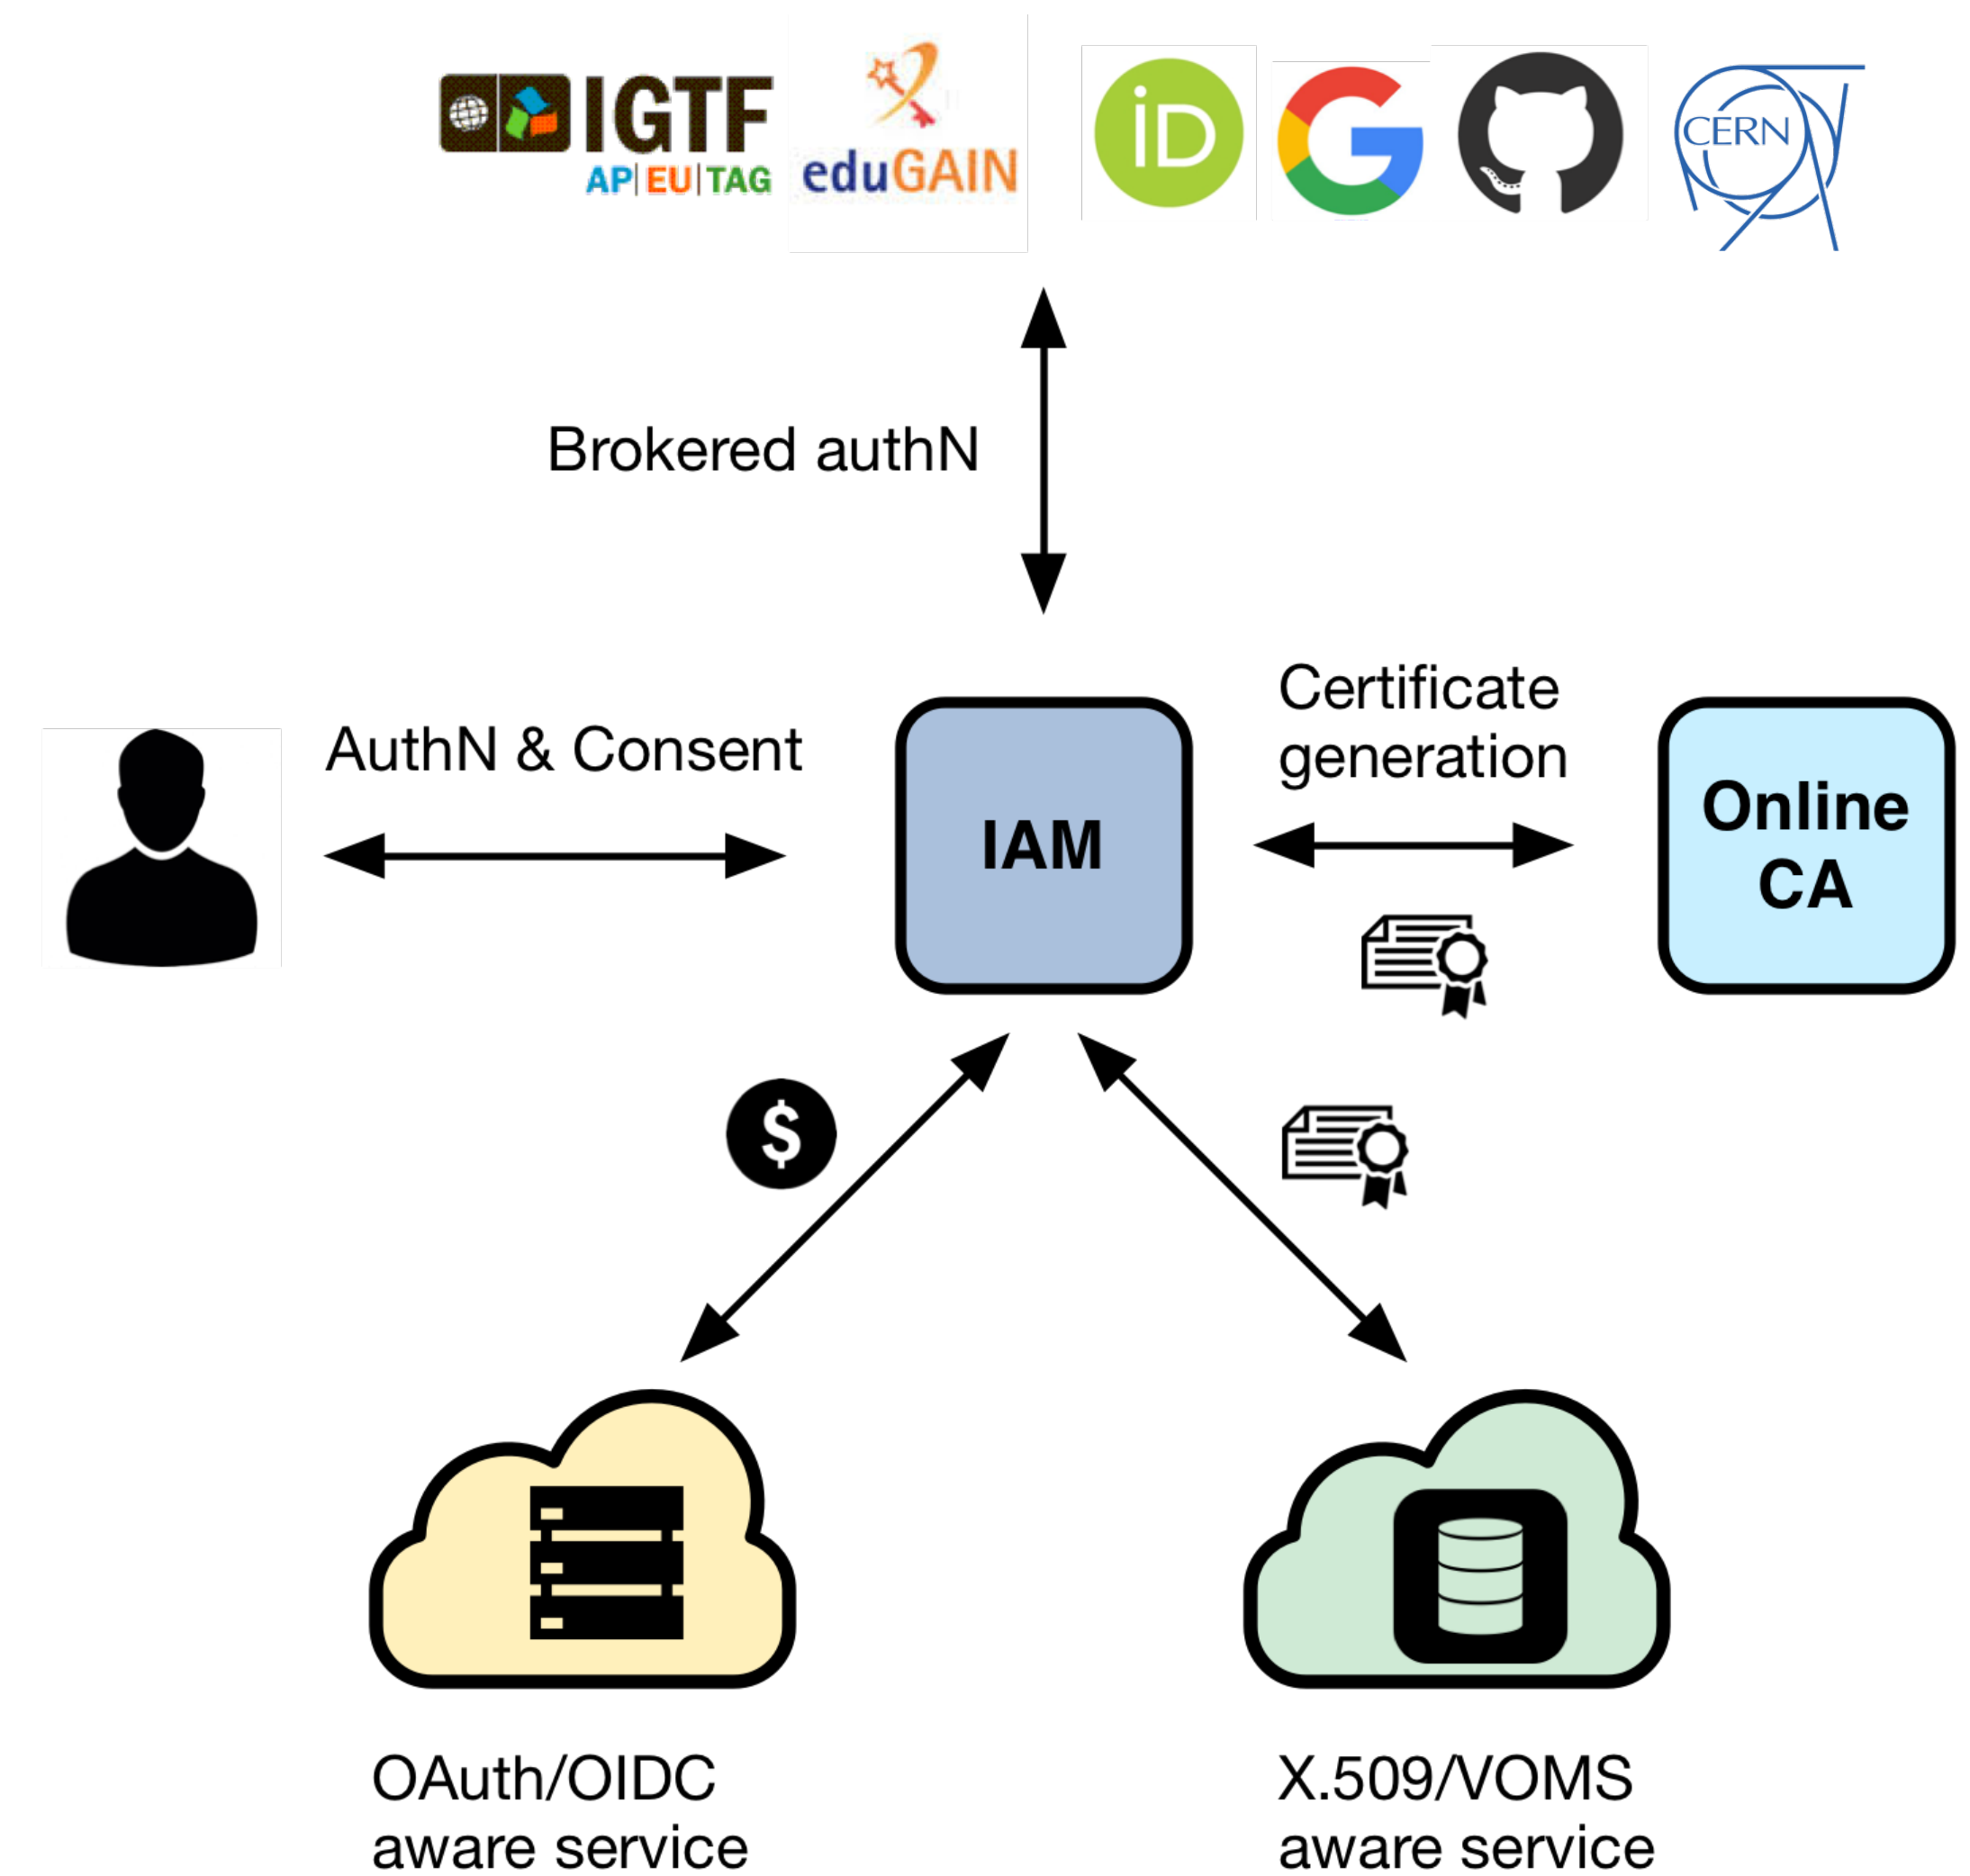
- Enable **integration with legacy services** through **controlled credential translation**

INDIGO IAM overview

INDIGO Identity and Access Management Service

A **VO-scoped** authentication and authorization service that

- supports **multiple authentication mechanisms**
- provides users with a **persistent, VO-scoped identifier**
- exposes **identity information, attributes and capabilities** to services via **JWT** tokens and standard **OAuth & OpenID Connect** protocols
- can integrate existing **VOMS**-aware services
- supports **Web** and **non-Web access, delegation** and **token renewal**
- supports **SCIM-compliant** provisioning APIs




INDIGO Identity and Access Management Service

Selected by the WLCG Management Board to be the core of the future, token-based WLCG AAI

- while ensuring backward compatibility with the existing infrastructure

Sustained by INFN for the foreseeable future, with current support from:



IAM key features

User enrolment & registration service

IAM currently supports two **enrolment flows**:

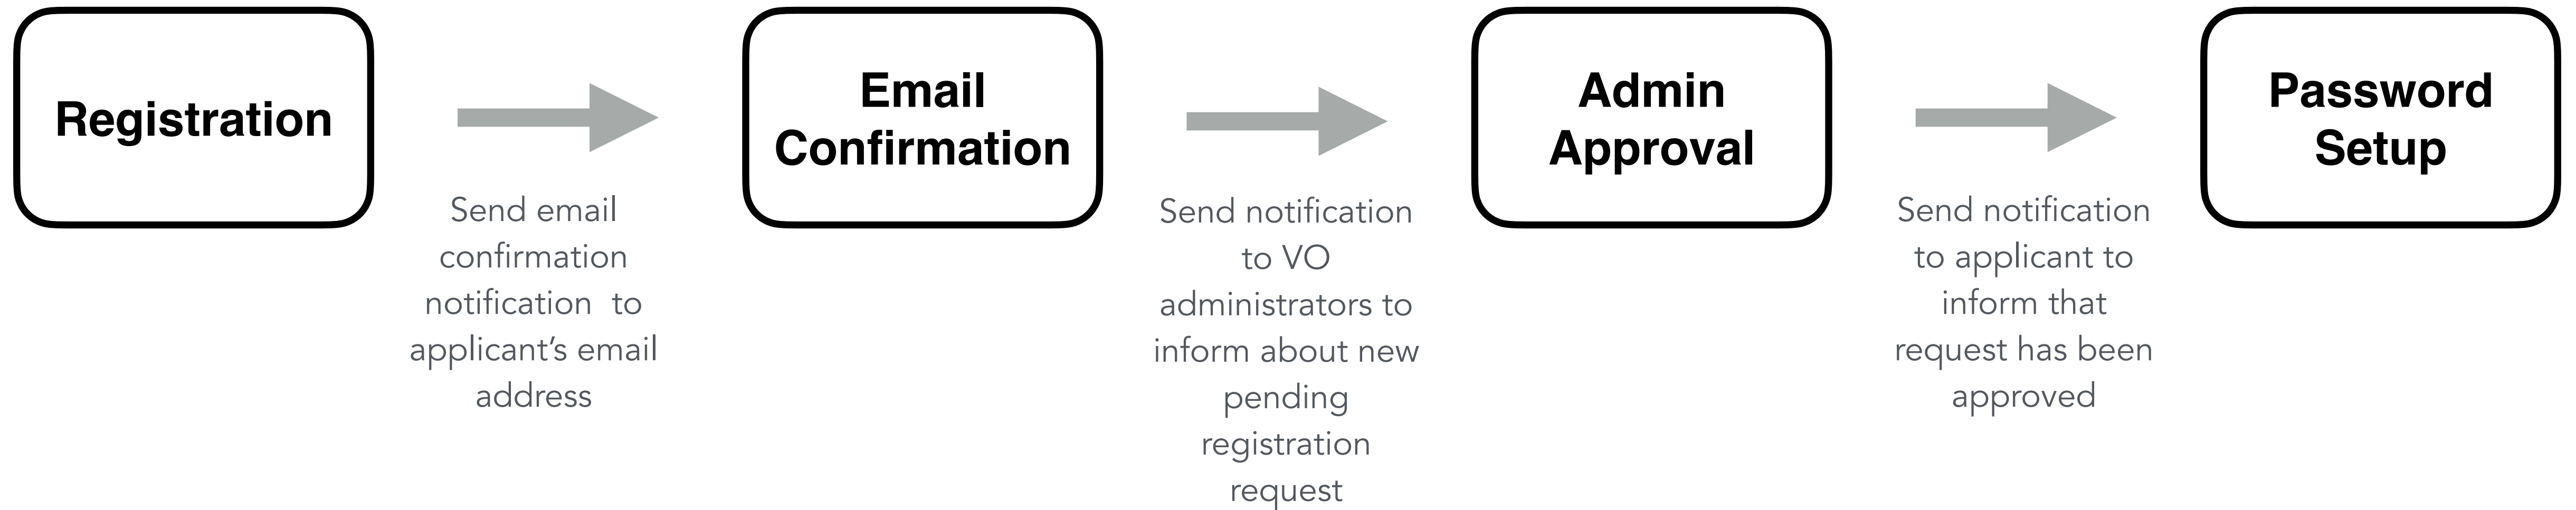
Admin-moderated flow

- The applicant fills basic registration information, accepts AUP, proves email ownership
- VO administrators are informed by email and can approve or reject incoming membership requests
- The applicant is informed via email of the administrator decision


Automatic-enrolment flow

- Users authenticated at **trusted, configurable** SAML IdPs are automatically on-boarded, without requiring administrator approval

IAM moderated enrolment flow



IAM moderated enrolment flow



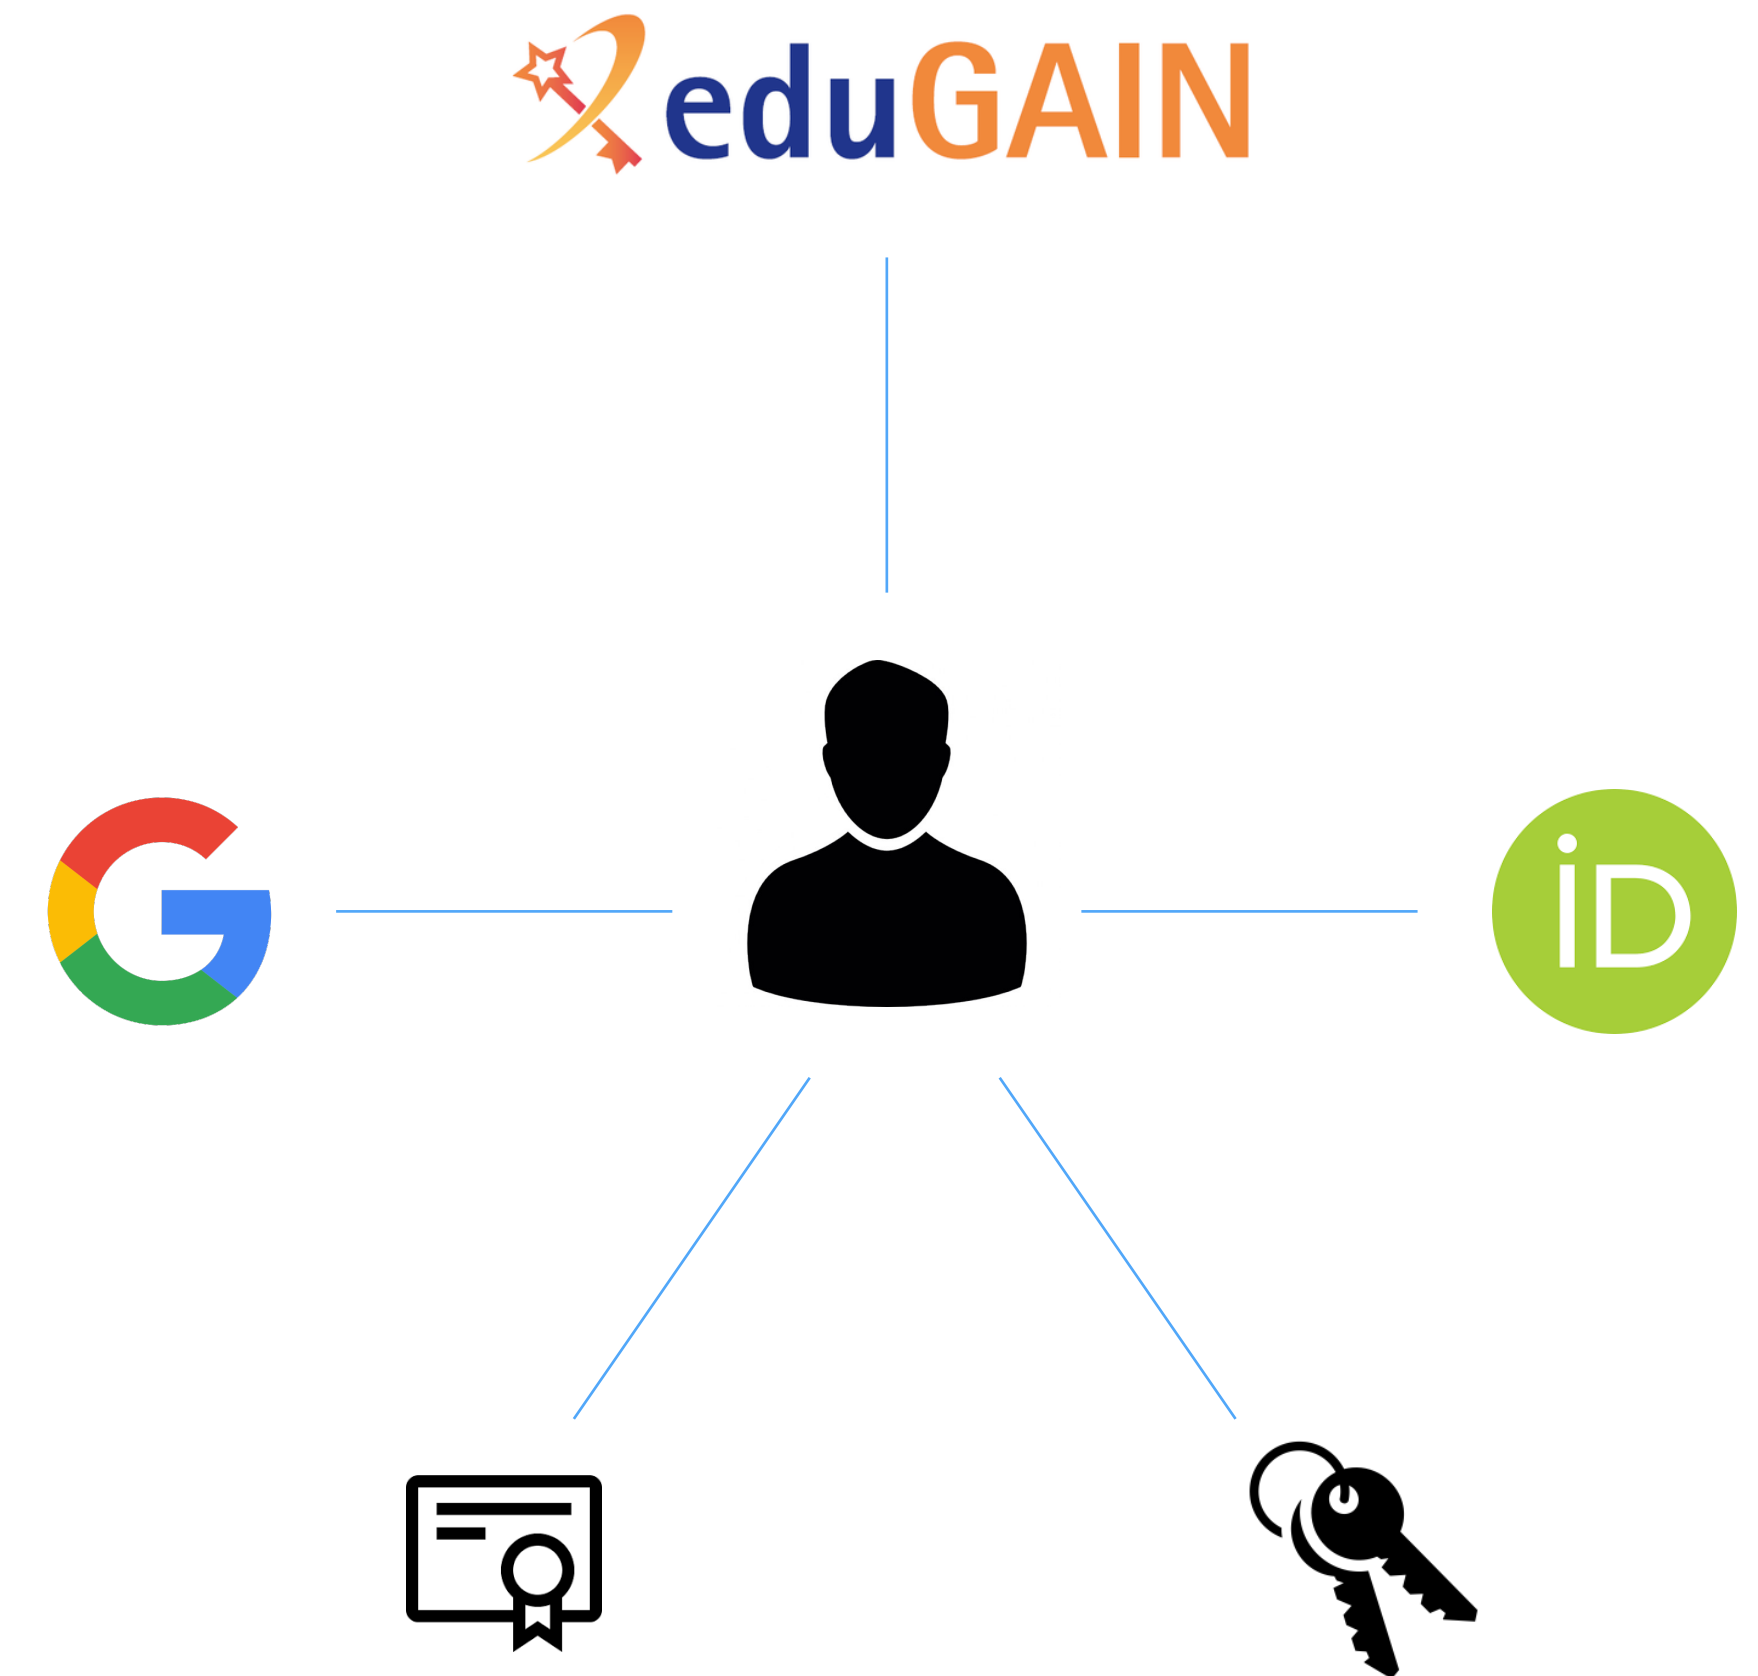
Flexible authentication & account linking

Authentication supported via

- **local username/password** credentials (created at registration time)
- **SAML** Home institution IdP (e.g., EduGAIN)
- **OpenID Connect** (Google, Microsoft, Paypal, ORCID)
- **X.509** certificates

Users can link any of the supported authentication credentials to their IAM account at registration time or later

To link an external credential/account, the user has to **prove** that he/she owns such account

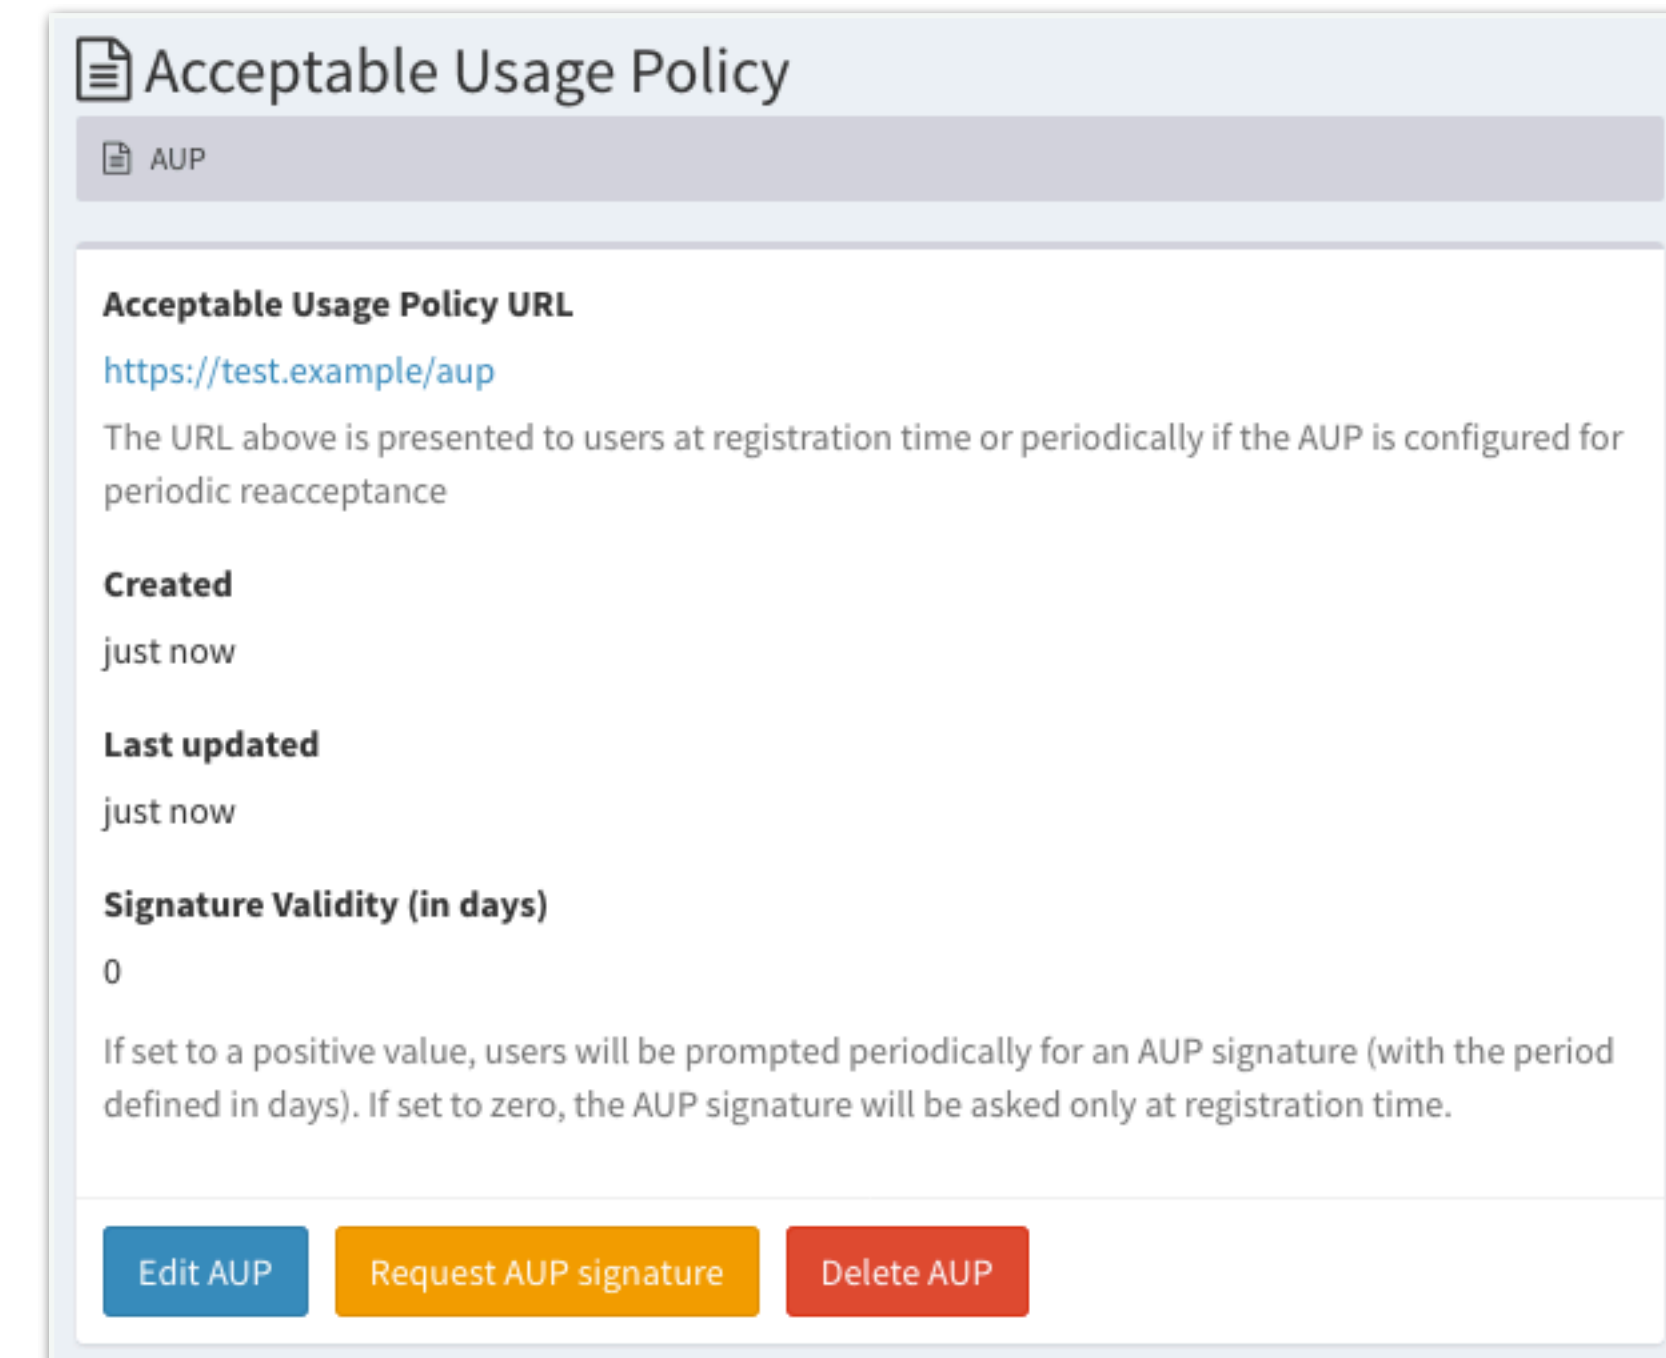


AUP enforcement support

AUP acceptance, if enabled, can be configured to be:

- requested once at user registration time
- periodically, with configurable period

User cannot login to the system (and as such be authenticated at authorized at services) unless the **AUP** has been accepted



The screenshot shows a web interface for managing an Acceptable Usage Policy (AUP). The title is "Acceptable Usage Policy" with a sub-header "AUP". The main content area displays the following information:

- Acceptable Usage Policy URL:** <https://test.example/aup>
The URL above is presented to users at registration time or periodically if the AUP is configured for periodic reacceptance
- Created:** just now
- Last updated:** just now
- Signature Validity (in days):** 0
If set to a positive value, users will be prompted periodically for an AUP signature (with the period defined in days). If set to zero, the AUP signature will be asked only at registration time.

At the bottom of the interface, there are three buttons: "Edit AUP" (blue), "Request AUP signature" (yellow), and "Delete AUP" (red).

SCIM provisioning APIs

IAM provides a RESTful API, based on the System for Cross-domain Identity Management (SCIM) standard, that can be used to access information in the IAM database

- users, groups, group memberships, etc...

The API can be used as an integration point towards external systems

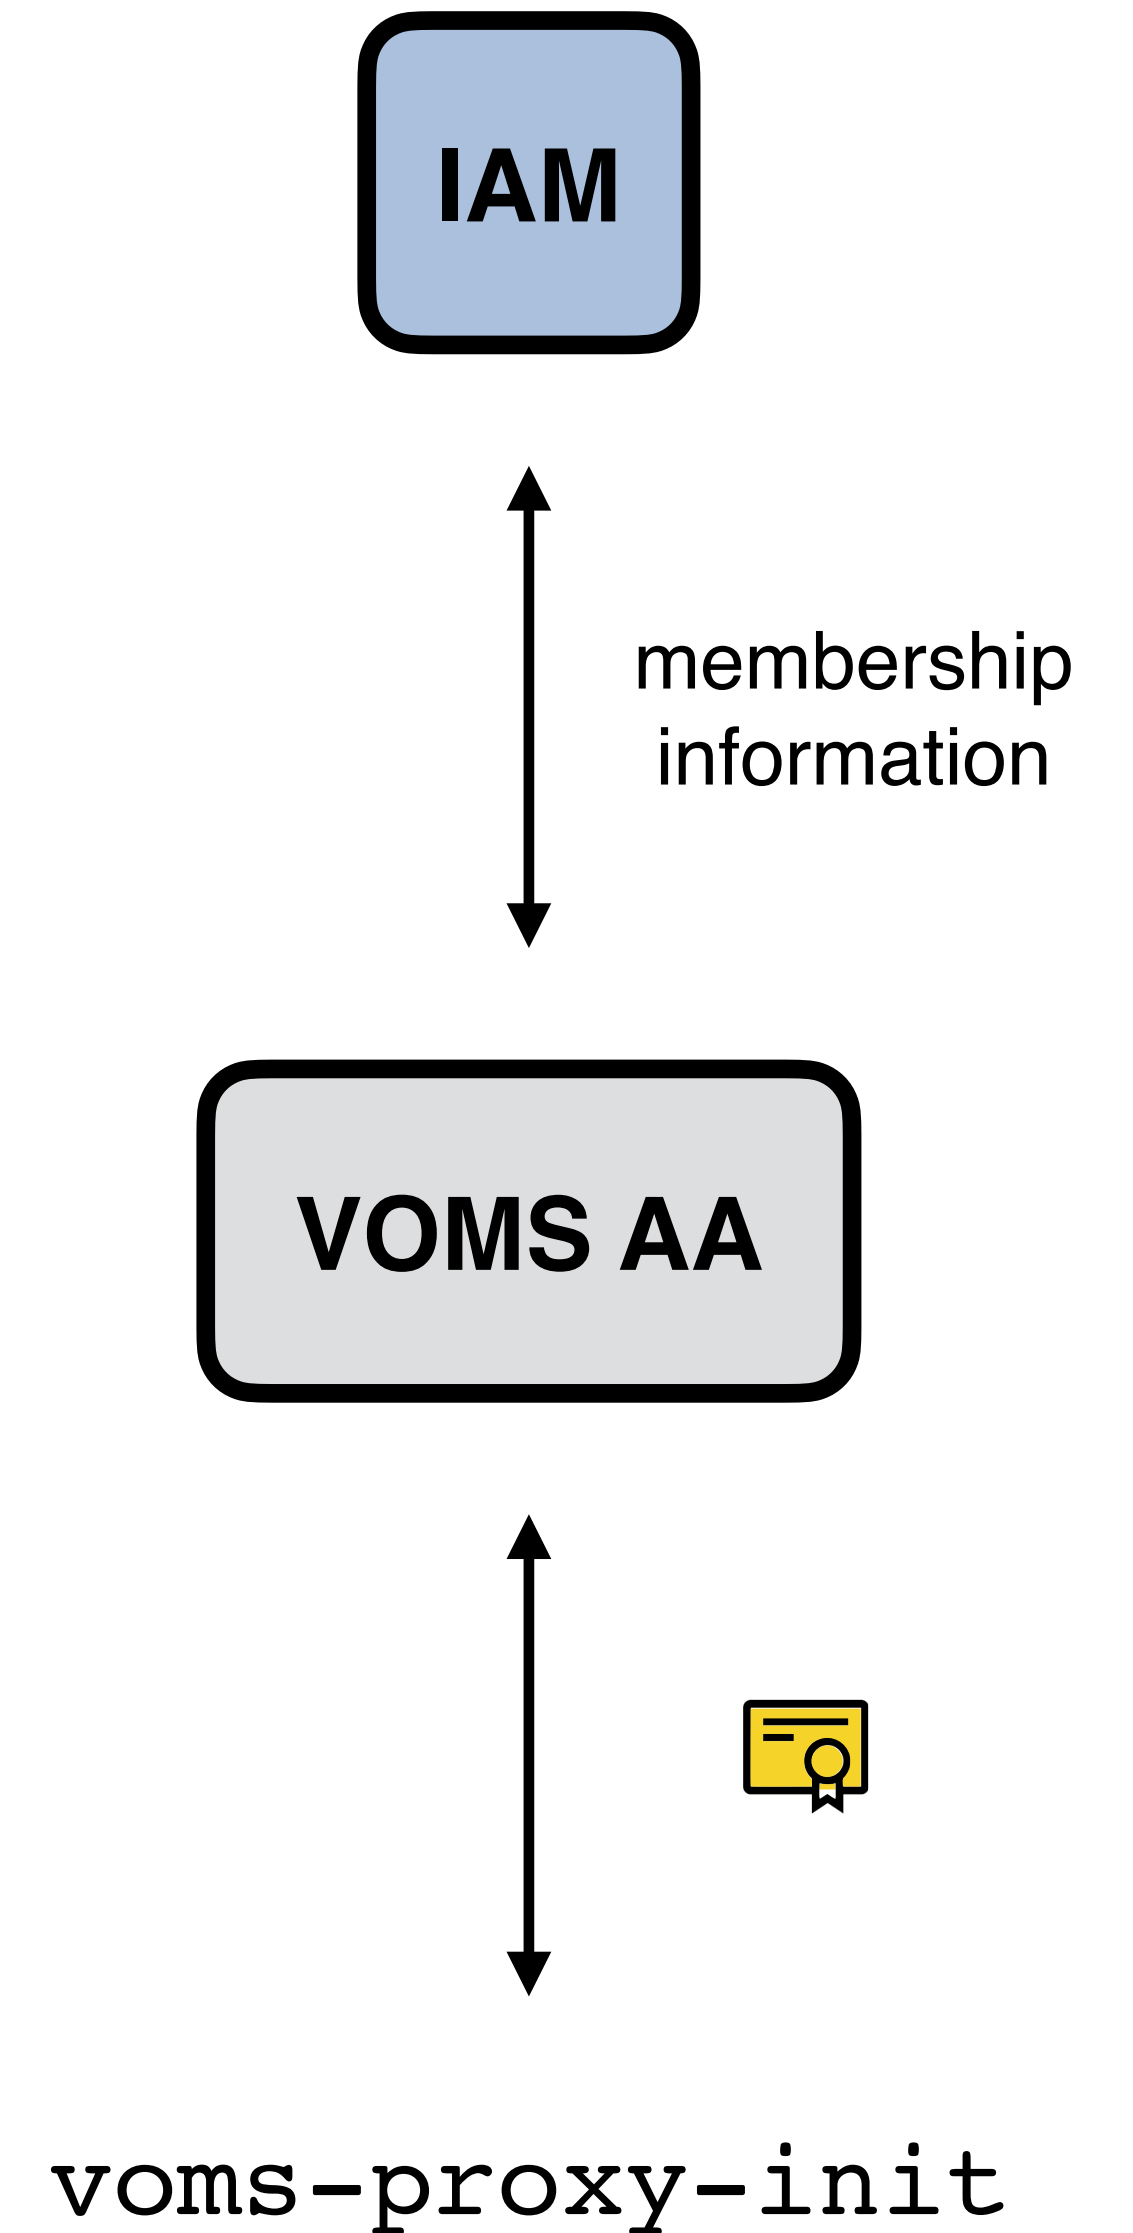
- Example:
 - The SCIM API is used in the integration with the HTCondor batch system to do account pre-provisioning based on IAM account information

VOMS provisioning

IAM includes a VOMS attribute authority micro-service that can encode IAM membership information in a **standard VOMS Attribute Certificate**

Proven compatibility with existing latest supported clients and Grid services

- e.g., data transfers in the ESCAPE data lake testbed rely on this



Easy integration with relying services

Standard OAuth/OpenID Connect enables **easy integration** with off-the-shelf services and libraries.

IAM has been successfully integrated with

- Openstack, Atlassian JIRA & Confluence, Moodle, Rocketchat, Grafana, Kubernetes, JupyterHub, **dCache**, **StoRM**, **XRootD (HTTP)**, **FTS**, **RUCIO**, **HTCondor**



Software Quality in IAM

Aim to have **~90% unit test coverage on all code:**

- now 33K LoC, 86,4% branch coverage, >1.2K tests

Open, **test-driven** development process

Static analysis tools

- [SonarCloud IAM page](#)

Multiple test suites

- **Unit tests**
- **Frontend test suite** (based on Selenium and Robot framework)
- **Deployment tests** (in CI)

Coverage

85.6% Coverage 818 Unit Tests — Coverage on New Code

Duplications

3.8% Duplications 72 Duplicated Blocks +0.0% Duplications

Size

24k 10

Add support to multiple OIDC providers #249

Open marcocaberletti wants to merge 2 commits into indigo-iam:develop from marcocaberletti:issue-229

Conversation 1 Commits 2 Checks 0 Files changed 35

marcocaberletti commented 14 days ago Member

This PR resolve issue #229.

- marcocaberletti Add support to multiple OIDC providers 34d63bd
- marcocaberletti requested review from andreaceccanti and enricovianello 14 days ago
- marcocaberletti added this to PRs ready for review in IAM next release 14 days ago
- New changes since you last viewed View changes
- Restore Link button caca09e

CnafSonarBot commented 14 days ago Collaborator

SonarQube analysis reported 1 issue

Note: The following issues were found on lines that were not modified in the pull request. Because these issues can't be reported as line comments, they are summarized here:

- OidcConfiguration.java#L97: Method has 10 parameters, which is greater than 7 authorized.

Add more commits by pushing to the issue-229 branch on marcocaberletti/iam.

Review requested Review has been requested on this pull request. It is not required to merge. Learn more. Show all reviewers

Deployment options

IAM as a service

- INFN provides IAM as a service to partner research communities. In this scenario, a dedicated IAM instance is deployed on the INFN infrastructure and configured according to the community needs. INFN takes care of keeping the service operational and up-to-date, while administrative control on the IAM instance is granted to the community.
- See <https://indigo-iam.github.io/docs/v/current/iam-aas/>

IAM on premise deployment

- IAM is an Apache-licensed identity solution, for which we provide **Docker images** on **Dockerhub** and RPMs and Deb packages
- See <https://indigo-iam.github.io/docs/v/current/admin-guide/>

Token-based AuthN/Z

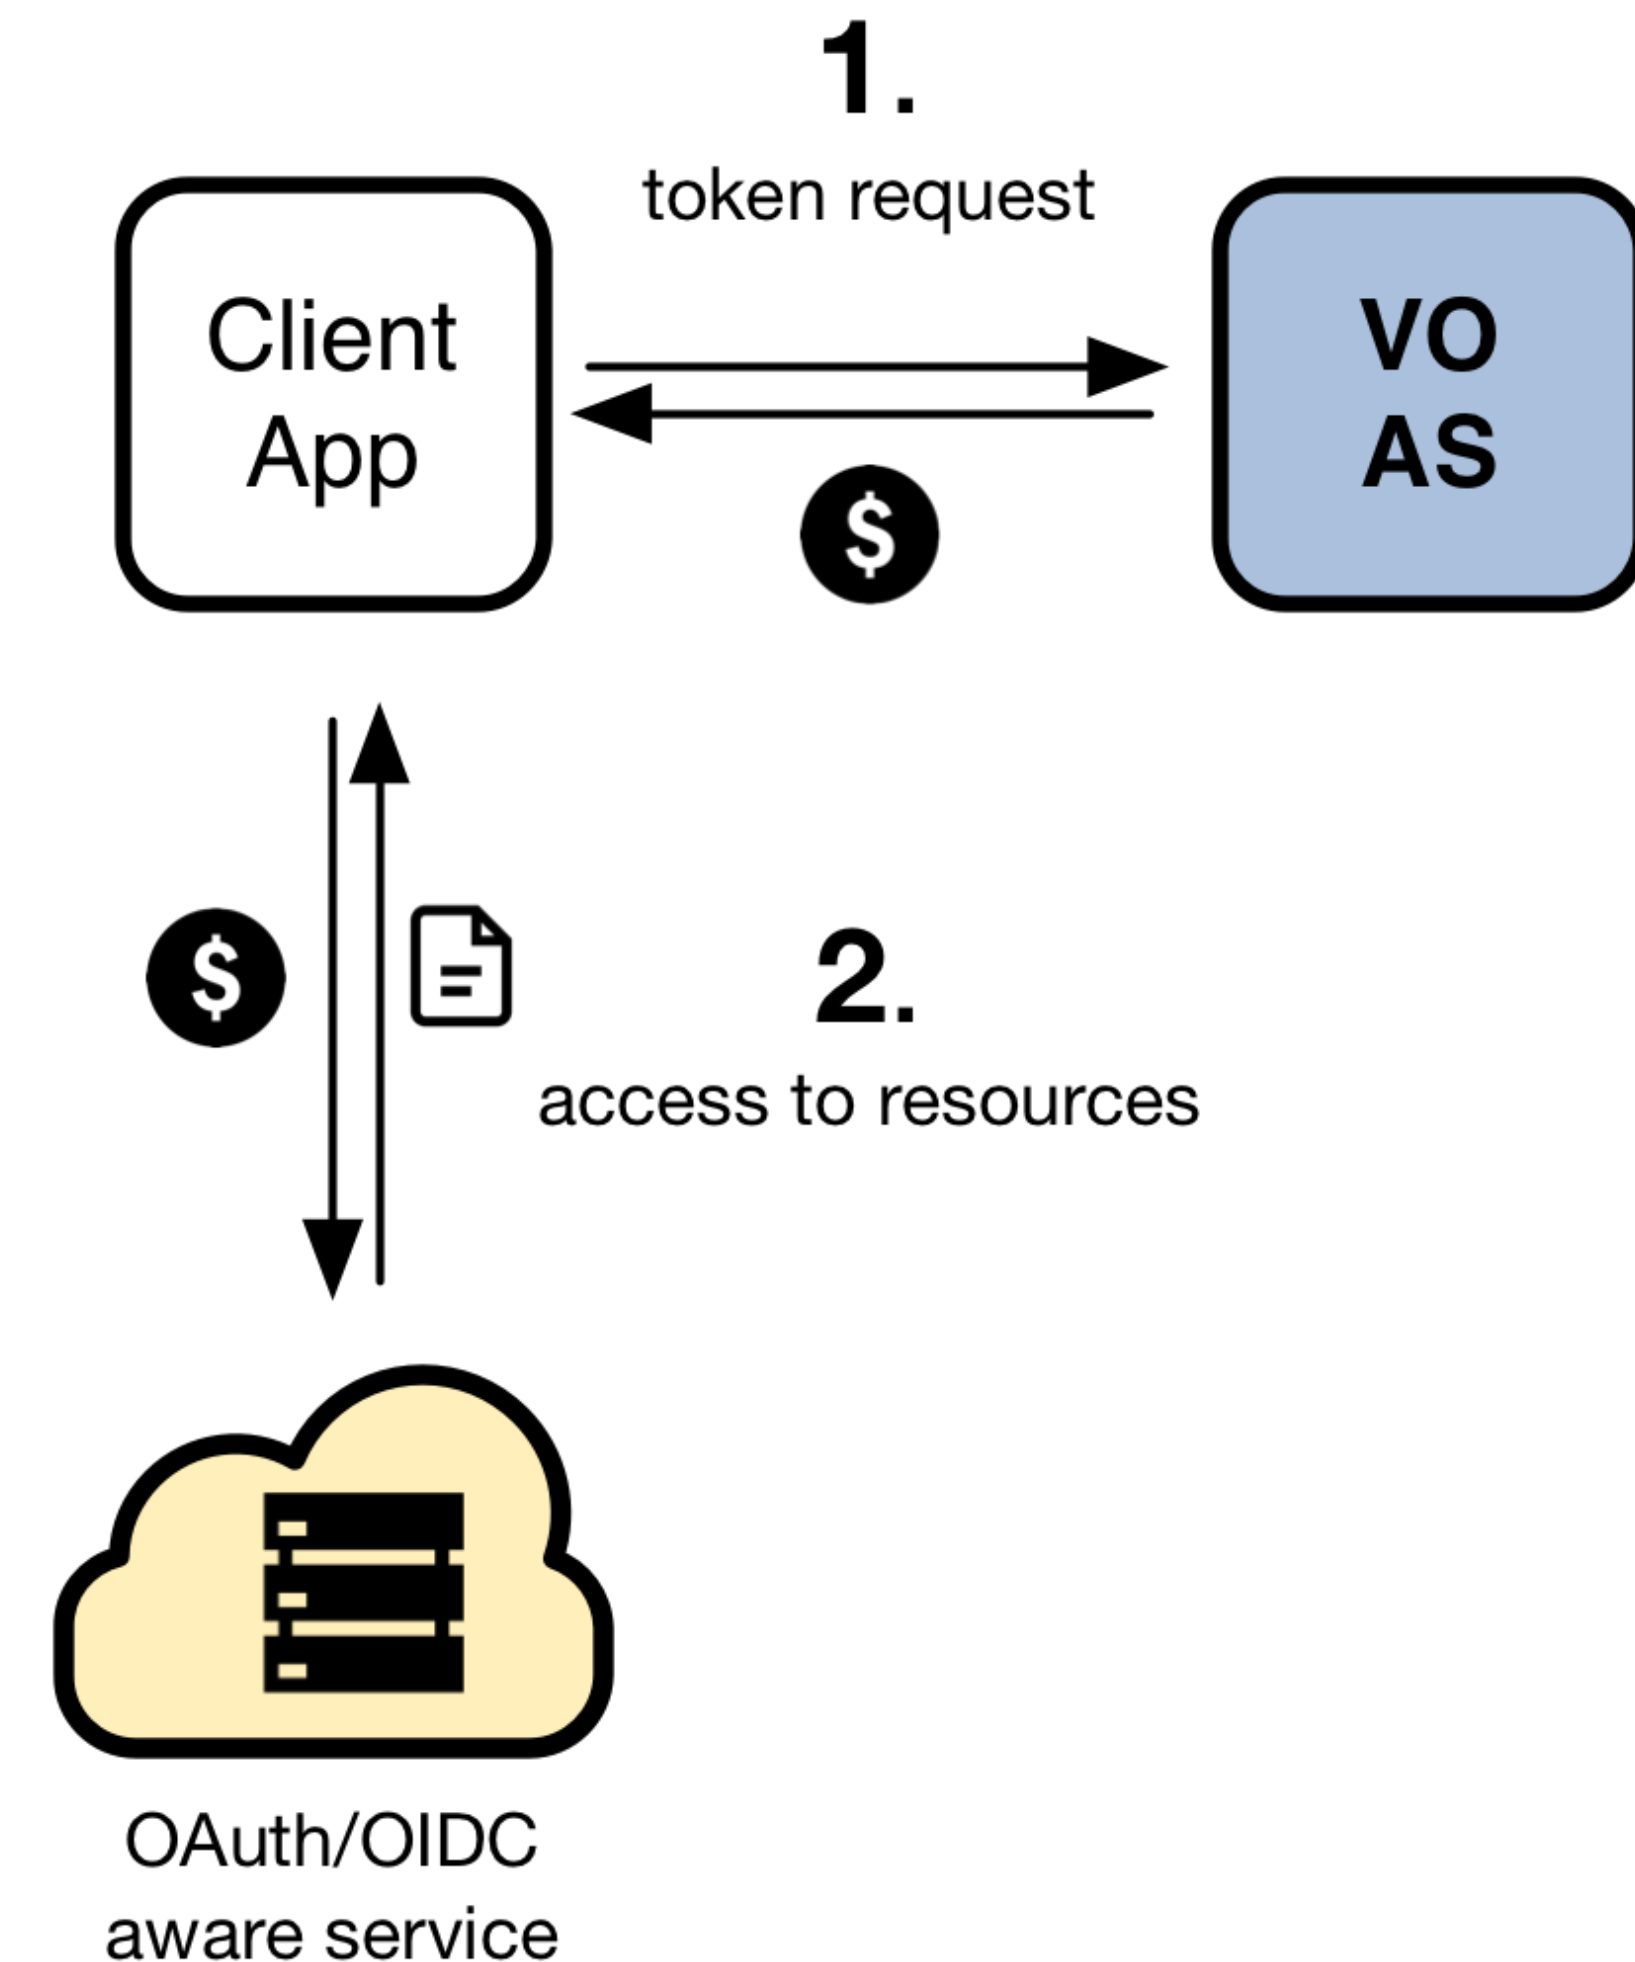
Token-based AuthN/Z from 10000 mt

In order to access resources/services, a **client application** needs an **access token**

The token is obtained from a **Virtual Organization** (which acts as an OAuth Authorization Server) using standard **OAuth/OpenID Connect** flows

Authorization is then **performed at the services** leveraging info extracted from the token:

- **Identity attributes:** e.g., **groups**
- **OAuth scopes:** capabilities linked to access tokens at token creation time



In practice...

The central authorization server (i.e., IAM) provides **attributes** that can be used for authorization at services, e.g.:

- groups/roles, e.g.: **analysis, production-manager**
- capabilities, e.g.: **storage.read:/, submit-job**

This information is exposed to services via **signed JWT tokens** and **via OAuth/OpenID Connect protocol message exchanges** (aka flows)

Services can then **grant or deny access** to resources based on this information.

Examples:

- allow read access on the **/analysis** folder to all members of the **analysis** group
- allow read access on the namespace to anyone with the capability **storage.read:/**

The WLCG Common JWT profile

How is **authentication** and **authorization** information encoded in **identity** and **access tokens**?

How is **trust** established between parties exchanging tokens?

What's the recommended **token lifetime**?

The screenshot shows a Zenodo record page for 'WLCG Common JWT Profiles'. The page includes a title, authors list, a technical note, and a preview of the document. The document text describes how WLCG users may use geographically distributed resources without X.509 credentials, mentioning OAuth2, OpenID Connect, and JSON Web Tokens. It also states that trust roots are established via OpenID Discovery or OAuth2 Authorization Server Metadata. The page features statistics: 136 views and 111 downloads. It is indexed in OpenAIRE. The publication date is September 25, 2019, and the DOI is 10.5281/zenodo.3460258. The license is Creative Commons Attribution 4.0 International.

Approach:

**rely on existing standards as much as possible,
extend only when needed**

Demo

Thanks!
Questions?

References

IAM @ GitHub: <https://github.com/indigo-iam/iam>

IAM documentation: <https://indigo-iam.github.io/docs>

WLCG Authorization WG: <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGAuthorizationWG>

IAM in action video: <https://www.youtube.com/watch?v=1rZlvJADOnY>

Contacts:

- andrea.ceccanti@cnaa.infn.it
- indigo-iam.slack.com