



Finanziato
dall'Unione europea
NextGenerationEU



Ministero
dell'Università
e della Ricerca



Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA



terabit

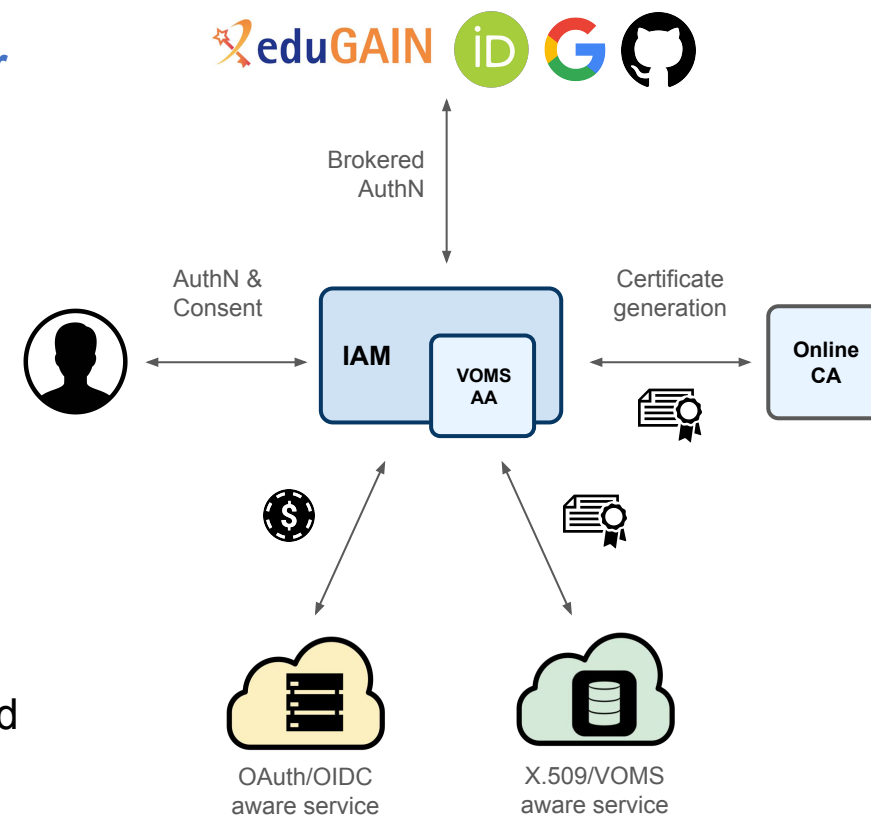
Authentication and Authorization with INDIGO IAM in the federation of computing resources

Roberta Miccoli - INFN CNAF

Work-Package Leaders Meeting | 15 - 16 ottobre 2024

INDIGO IAM in one slide

- Standard **OAuth2 Authorization Server** and **OpenID Connect Provider**
 - Easy integration with (web) applications
- **Java** application based on the **Spring Boot** framework
- **Multiple authentication mechanisms**
 - SAML, X.509, OpenID Connect, local users, etc.
- **Account linking**
- Moderated and automatic user enrollment
- Enforcement of **AUP acceptance**
- Management of Organization membership
- Issuance of **JWT** tokens and VOMS attribute certificates with **identity** and **membership information, attributes** and **capabilities**
- Typically deployed as a **Docker container**



INDIGO IAM in the computing federation

Role

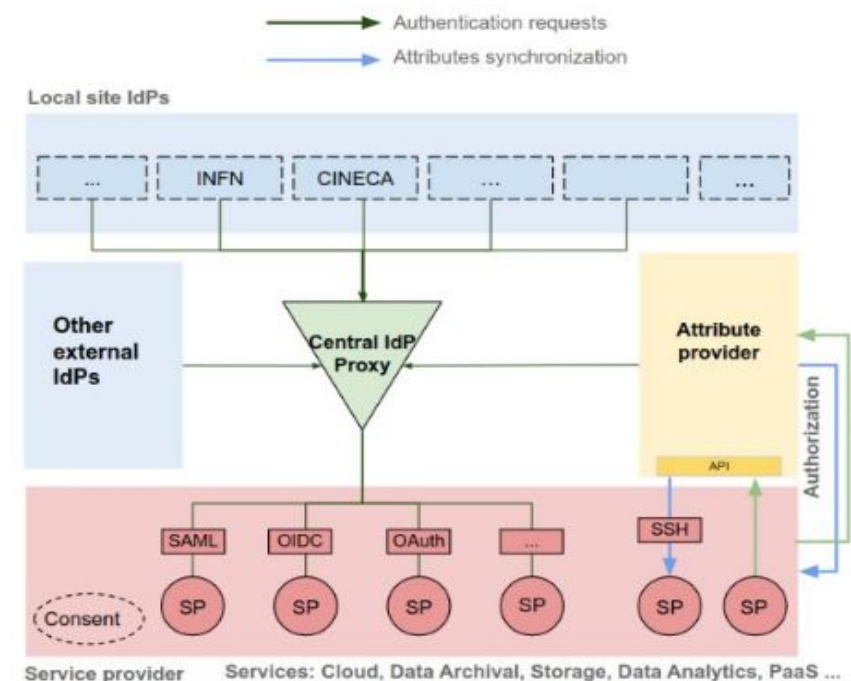
- Entrypoint for the computing federation, IAM acts as *Attribute Authority* and *IdP proxy* for the whole infrastructure

Objective

- Federate INFN and CINECA IdPs
 - CINECA and INFN users can register into INDIGO IAM through their IdP and be authorized to access the resources of both institutions

Advantages

- In-house development (mainly by INFN CNAF), born to satisfy the needs of scientific community
- Easy integration with third-party applications
- Backward-compatible with Grid-based authorization
- Support for capability-based authorization
- Allows the definition of policies for fine-tuned access privileges



PoC IAM: technologies

- The PoC IAM instance is deployed using **Docker Compose** on a Virtual Machine within the **INFN Cloud** infrastructure
 - deployed behind an **NGINX**
 - stores data in a **MySQL** database
- All the services belonging to the PoC infrastructure (the INDIGO PaaS orchestrator, RUCIO, etc.) that support authentication with the PoC IAM have been integrated by registering them as **clients** in IAM, exploiting the OpenID Connect technology

- **iam-be**: the main service (backend), <https://iam-poc-icsc.cloud.infn.it>
- **client**: an example of a client application, <https://iam-poc-icsc.cloud.infn.it/iam-test-client>
- **nginx-iam**: NGINX image used for TLS termination and reverse proxy which forwards requests to iam-be and client
- **voms-aa**: VOMS-AA microservice which releases VOMS proxies, <https://iam-poc-icsc.cloud.infn.it:15000>
- **nginx-voms**: NGINX reverse proxy which forwards requests to voms-aa (it differs by the nginx-iam service since it supports HTTPG)
- **db**: MySQL database used in read/write mode by INDIGO IAM and read mode by VOMS-AA
- **trust**: docker image for the Grid CA certificates, mounted in the `/etc/grid-security/certificates` path of the other services when needed

PoC IAM: state of the art

ICSC
Centro Nazionale di Ricerca in HPC,
Big Data and Quantum Computing

Welcome to **poc-icsc**

Sign in with your poc-icsc credentials

Username

Password

Sign in

Forgot your password?

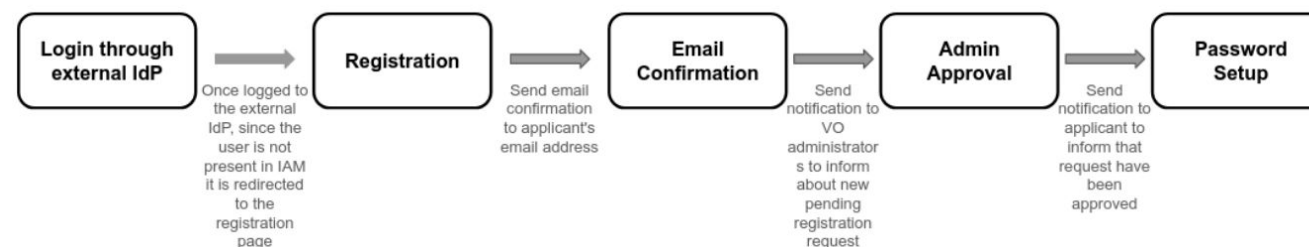
Or sign in with

CINECA

INFN
CCR - AAI

<https://iam-poc-icsc.cloud.infn.it/>

- Defined a **Virtual Organization** (VO), called *poc-icsc*
- Moderated user enrollment
 - it requires manual approval by IAM admins



- Authentication methods
 - external IdPs: **CINECA** dev instance of keycloak (*OIDC*) and **INFN AAI** (*SAML*)
 - X.509 certificates (if linked to the account)



PoC IAM: state of the art

ICSC
Centro Nazionale di Ricerca in HPC,
Big Data and Quantum Computing

Welcome to **poc-icsc**

Sign in with your poc-icsc credentials

Username

Password

Sign in

[Forgot your password?](#)

Or sign in with

CINECA

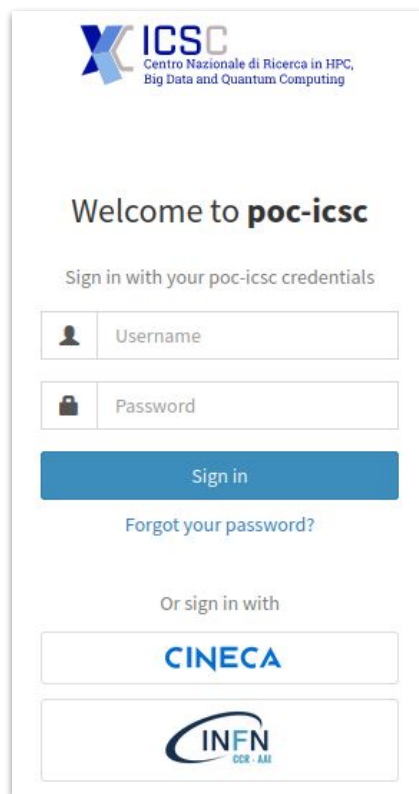
INFN
CCR - ASI

<https://iam-poc-icsc.cloud.infn.it/>

Attribute-based authorization

- Defined a set of **IAM groups** to enforce a more controlled access to federated resources
 - *poc-icsc/prod*: optional group (or VOMS role), necessary to submit third-party transfer jobs to FTS in the infrastructure, authenticating and authorizing with a proxy
 - *poc-icsc/admins/poc-icsc*: mapped to an OpenStack project on the federated Cloud providers used to instantiate services on the public network
 - *poc-icsc/priv-admins/poc-icsc*: not yet defined, will be mapped to an OpenStack project on the federated Cloud providers used to instantiate services on a private network

PoC IAM: state of the art

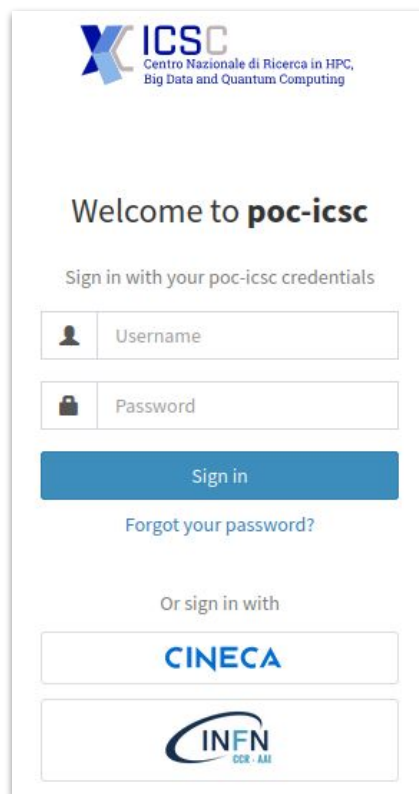


<https://iam-poc-icsc.cloud.infn.it/>

Scope-based authorization

- Defined **IAM scope policies** applied to each storage system for finer-grained read/write permissions in the federated namespace
 - read access to the entire namespace (/) is granted to users of the *poc-icsc* group
 - write access to the `/user/<iam-username>` namespace is granted to the user `<iam-username>`

PoC IAM: state of the art



<https://iam-poc-icsc.cloud.infn.it/>

The PoC IAM instance has been successfully integrated with

- PaaS orchestrator and its dashboard
- OAuth2 proxy services used by [interLink](#)
- RUCIO + FTS
- INFN Storage systems and INFN federated Cloud services involved in the PoC

PoC IAM integration in INFN Cloud



Welcome to the INFN Cloud Dashboard!

Please login, or register »

Compute Services

Scientific Community Customizations

Analytics

Data Services

Machine Learning

The dashboard features five main service categories, each represented by a cloud icon with a specific symbol: a server rack for Compute Services, gears for Scientific Community Customizations, a bar chart for Analytics, a database cylinder for Data Services, and a binary code pattern for Machine Learning.

<https://mycloud-poc-icsc.cloud.infn.it/>



PoC IAM integration in INFN Cloud

Welcome to the INFN Cloud

Please login, or register »

Scientific Community Customizations

Data Services

Analytics

Learning

<https://mycloud-poc-icsc.infn.it>

ICSC
Centro Nazionale di Ricerca in HPC,
Big Data and Quantum Computing

Welcome to **poc-icsc**

Sign in with your poc-icsc credentials

Username

Password

Sign in

Forgot your password?

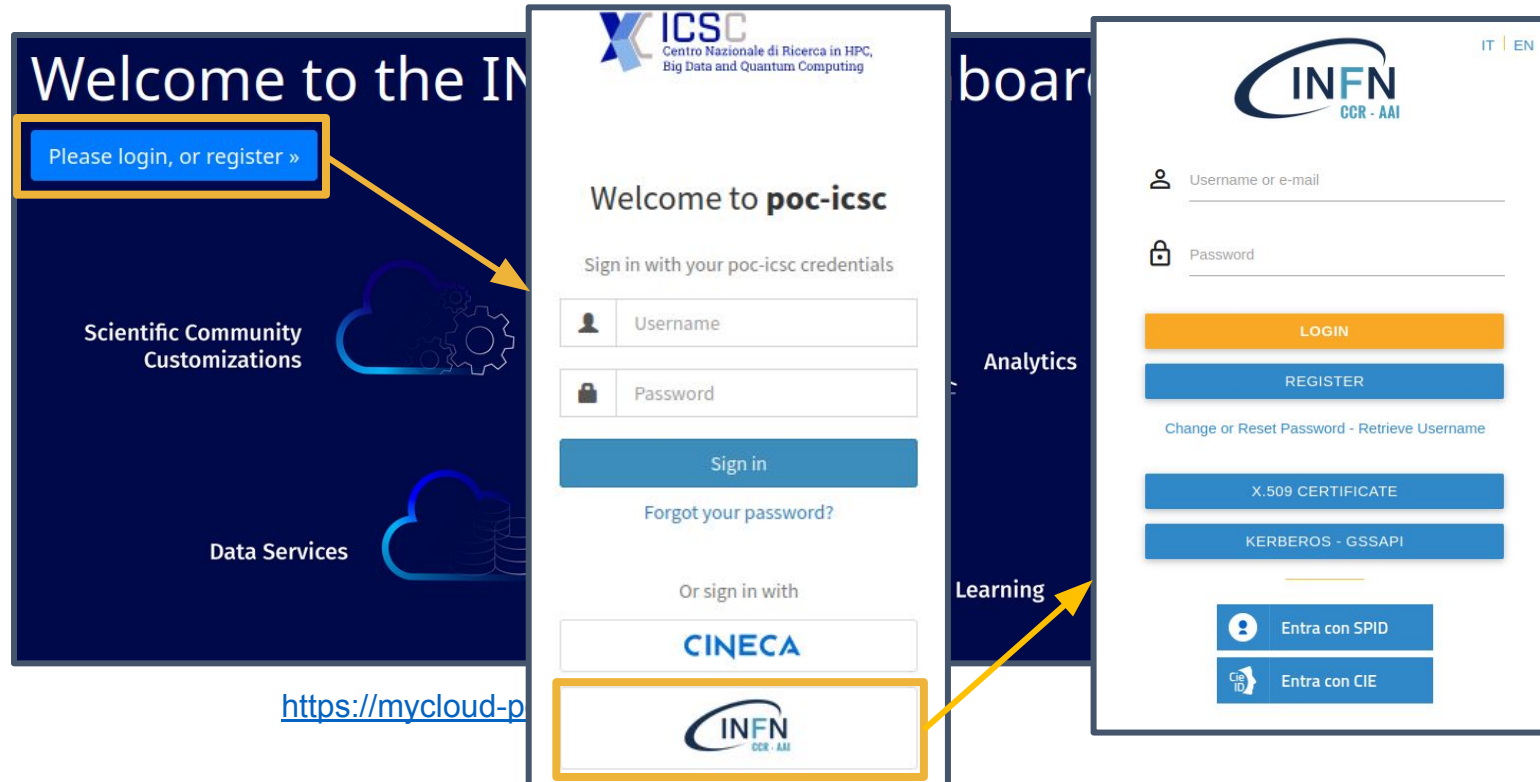
Or sign in with

CINECA

INFN
CCR - AM



PoC IAM integration in INFN Cloud





PoC IAM integration in INFN Cloud

Welcome to the INFN Cloud

Please login, or register »

Scientific Community Customizations

Data Services

<https://mycloud-poc-icsc.it>

ICSC
Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Welcome to poc-icsc

Sign in with your poc-icsc credentials

Username

Password

Sign in

Forgot your password?

Or sign in with

CINECA

INFN CCR - AAI

Analytics

Learning

INFN CCR - AAI

IT | EN

Username or e-mail

Password

LOGIN

REGISTRATION

Change or Reset Password

X.509 CERTIFICATES

KERBEROS

Entra con

Entra con

Dashboard

DEPLOYMENTS

ADVANCED

EXTERNAL LINKS

REPORT

CREATION COMPLETED 0

CREATION IN PROGRESS 0

CREATION FAILED 0

SERVICES

Search...

ON-DEMAND SERVICES

Virtual machine

Launch a compute node getting the IP and SSH credentials to access via ssh

INDIGO IAM as a Service

The on-demand deployment service for the INDIGO IAM provides a quick and easy way for organizations to deploy their own instance of the INDIGO IAM, which is an open-source Identity and Access Management system...

Settings

Help

Roberta Miccoli

The service catalogue is personalised for different groups of users, with the list of available services tailored based on group membership and roles

Current constraints

Identity Federation

- Due to existing certification policies and processes, a federated user must already be registered as a user at CINECA to access CINECA resources
- As a workaround for the PoC, all federated users must first be statically registered in CINECA LDAP and then registered in the PoC IAM instance using the same username from CINECA LDAP
- On the other hand, any federated CINECA user can transparently access both CINECA and INFN resources
- Ongoing discussions in ICSC spoke 0 and TeRABIT

Integration with CINECA

- Limited to a few processes and test endpoints (i.e., offloading part of the workflow of a Cloud application to HPC (CINECA) resources and data transfer between a test CINECA S3 endpoint and INFN resources)

Next steps

About PoC federation:

- Federate the test Cloud @CINECA with the PoC IAM instance and define an agreement between INFN and CINECA for resource access policies
- Apply security policies based on the ISO/IEC and ITU-T standards (e.g. **MFA**), and R&E community specifications (e.g. **Level of Assurance**)
- Define and apply access policies based on external IdPs authenticating through EduGAIN
- Enable automatic IAM user enrollment from trusted IdPs, thus allowing direct access to all ICSC/TeRABIT resources (to users coming from trusted IdPs)

About INDIGO IAM developments:

- Support **MFA**
- Support **OIDC Federation**
- Explore authorization with **Open Policy Agent** (OPA)

Exploring AuthZ with OPA

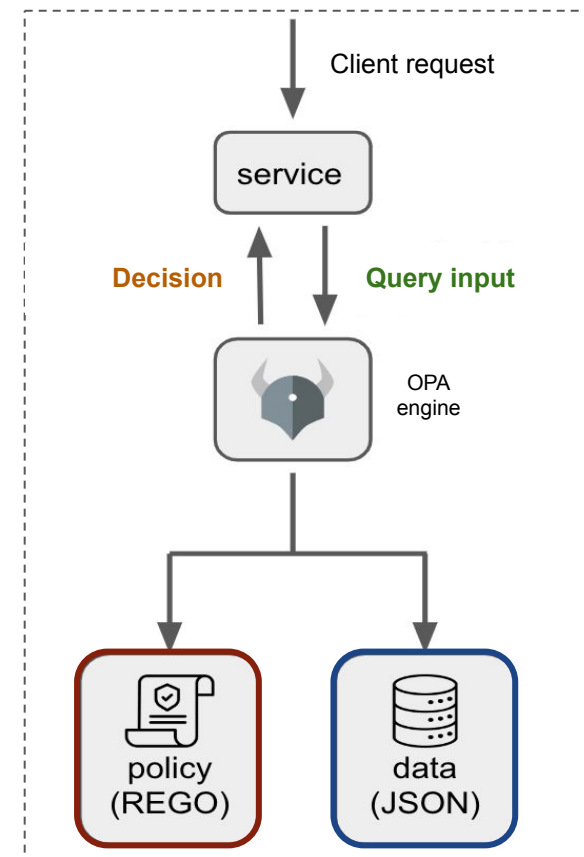
[Open Policy Agent](#) (OPA) is an open-source authorization engine based on a high-level declarative language (**Rego**) that allows the definition of policies as code

Rego is designed to express policies over complex hierarchical data structures

- policy authors can focus on what queries should return rather than how they should be executed
- Rego ensures high performance policy decisions, even with increasing number of rules

A service which needs to take policy decisions can **query** OPA with arbitrary structured data (JSON or YAML) as **input**

- OPA evaluates the query input against **policies** and optionally **data**
- OPA **decision** is not limited to a simple allow/deny answer, but can generate arbitrary structured data as output



Current usage of OPA for Grid and Cloud middleware

Integration with [StoRM WebDAV](#) service:

- OPA will replace the current StoRM WebDAV Policy decision Point (PdP) logic
 - supports both JWT tokens and X509 VOMS proxies

Integration with [StoRM tape REST API](#):

- OPA is used for JWT AuthN
- OPA is used for AuthZ in alignment with the same rules applied in StoRM WebDAV

Integration with INDIGO IAM:

- OPA is going to replace and evolve the IAM Scope Policy API
 - more readable policies
 - policies are also applied to clients to support the OAuth *client credentials* flow (not bound to a user)
 - backward compatible with current IAM scope policies syntax
- An OPA query took **~130 ms** to parse 10k policies, which in IAM reached the client timeout!



Finanziato
dall'Unione europea
NextGenerationEU



Ministero
dell'Università
e della Ricerca



Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA



Thanks for your
attention!





Useful references

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

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

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

OPA documentation: <https://www.openpolicyagent.org/docs/latest/>

OPA source code:

- StoRM Tape AuthN/Z: <https://baltig.infn.it/fagostin/storm-tape-authz>
- IAM OPA integration: <https://baltig.infn.it/fagostin/iam-opa-integration>

For general information:

- OAuth 2.0: <https://oauth.net/2/> and OAuth 2.1: <https://oauth.net/2.1/>
- OpenID Connect: <https://openid.net/connect/>

Contacts:

- iam-support@lists.infn.it



Discussion points

- Can we enhance INFN/eduGAIN users access to CINECA resources?
 - currently: all federated users must first be statically registered in CINECA LDAP and then registered in the PoC IAM instance using the same username from CINECA LDAP
 - this means that access token presented to CINECA resources must contain CINECA username
 - proposal: we could create a new IAM profile (*ICSC*) that customizes the access token payload, assuming that the INFN user has linked their CINECA account to their PoC IAM account
 - advantage: this decouples IAM username from CINECA username