RIA-653549

INDIGO – DataCloud

Authentication and Authorization in INDIGO DataCloud

Andrea Ceccanti
INFN
andrea.ceccanti@cnaf.infn.it

on behalf of the INDIGO AAI task force

INDIGO Datacloud

- An H2020 project approved in January 2015 in the EINFRA-1-2014 call
  - 11 M€
  - 30 Months (Apr. 2015 -> Sept. 2017)
- **Who:** 26 partners from 11 European countries
- **What:** develop an **open source** platform for computing and data targeted at **multi-disciplinary scientific communities**
- **Where:** provisioned over hybrid (public and private) e-infrastructures

INDIGO - DataCloud

INDIGO WP structure

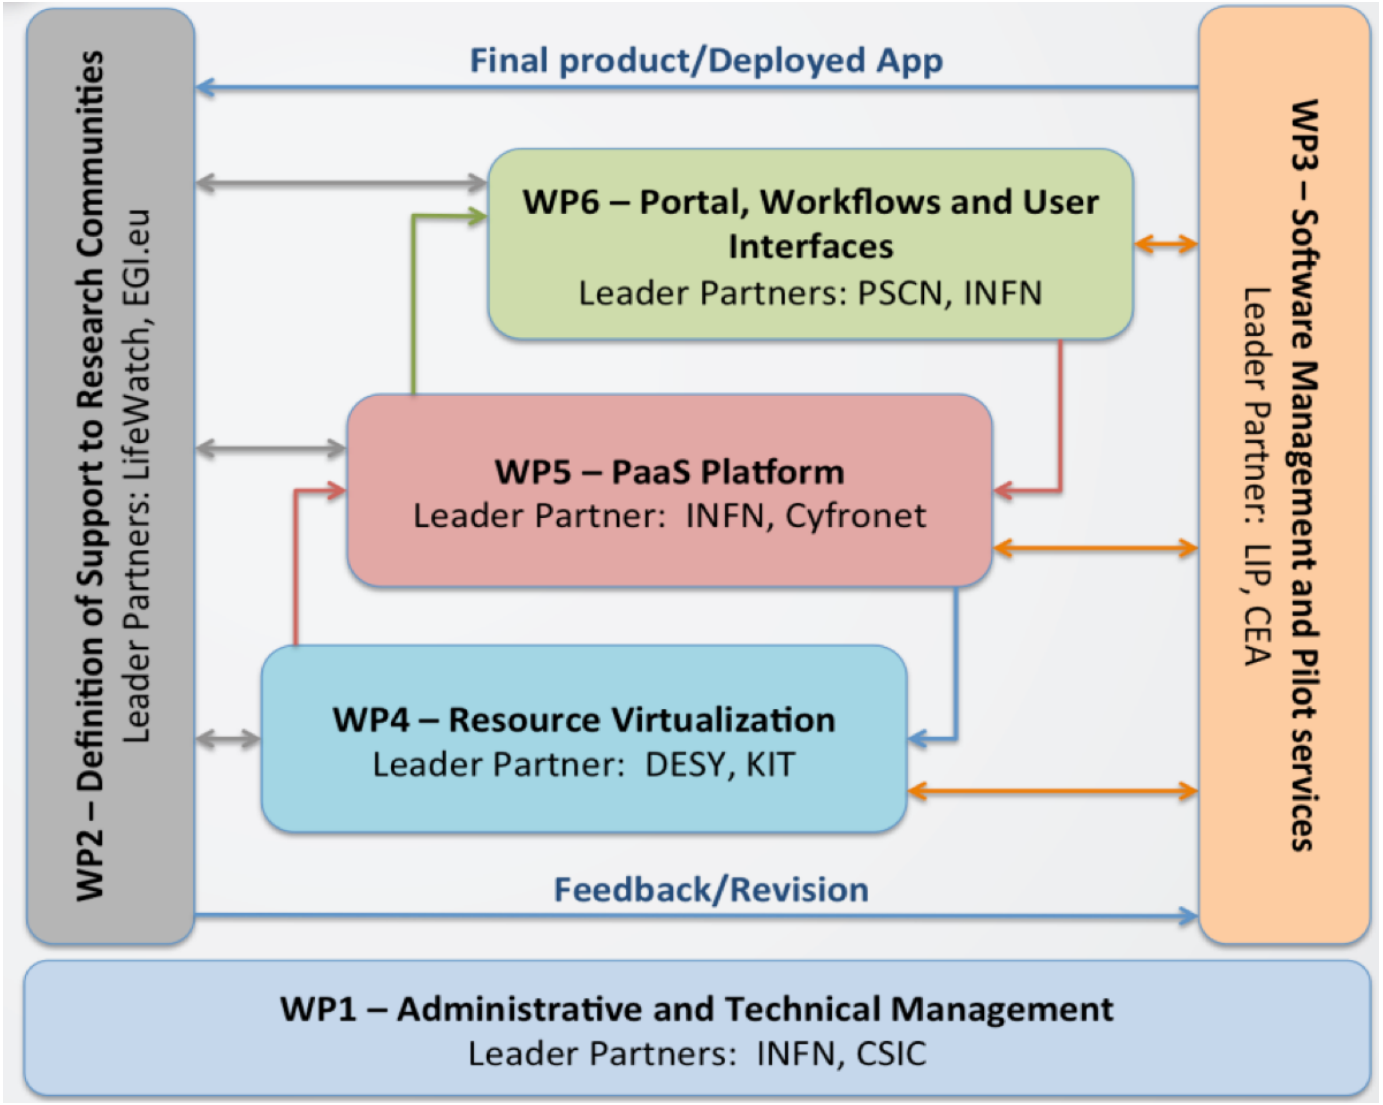

INDIGO Architecture

INDIGO - DataCloud

- Color codes
  - **Yellow:** implementation based on already available solution to be improved/ changed
  - **Orange**: New developments

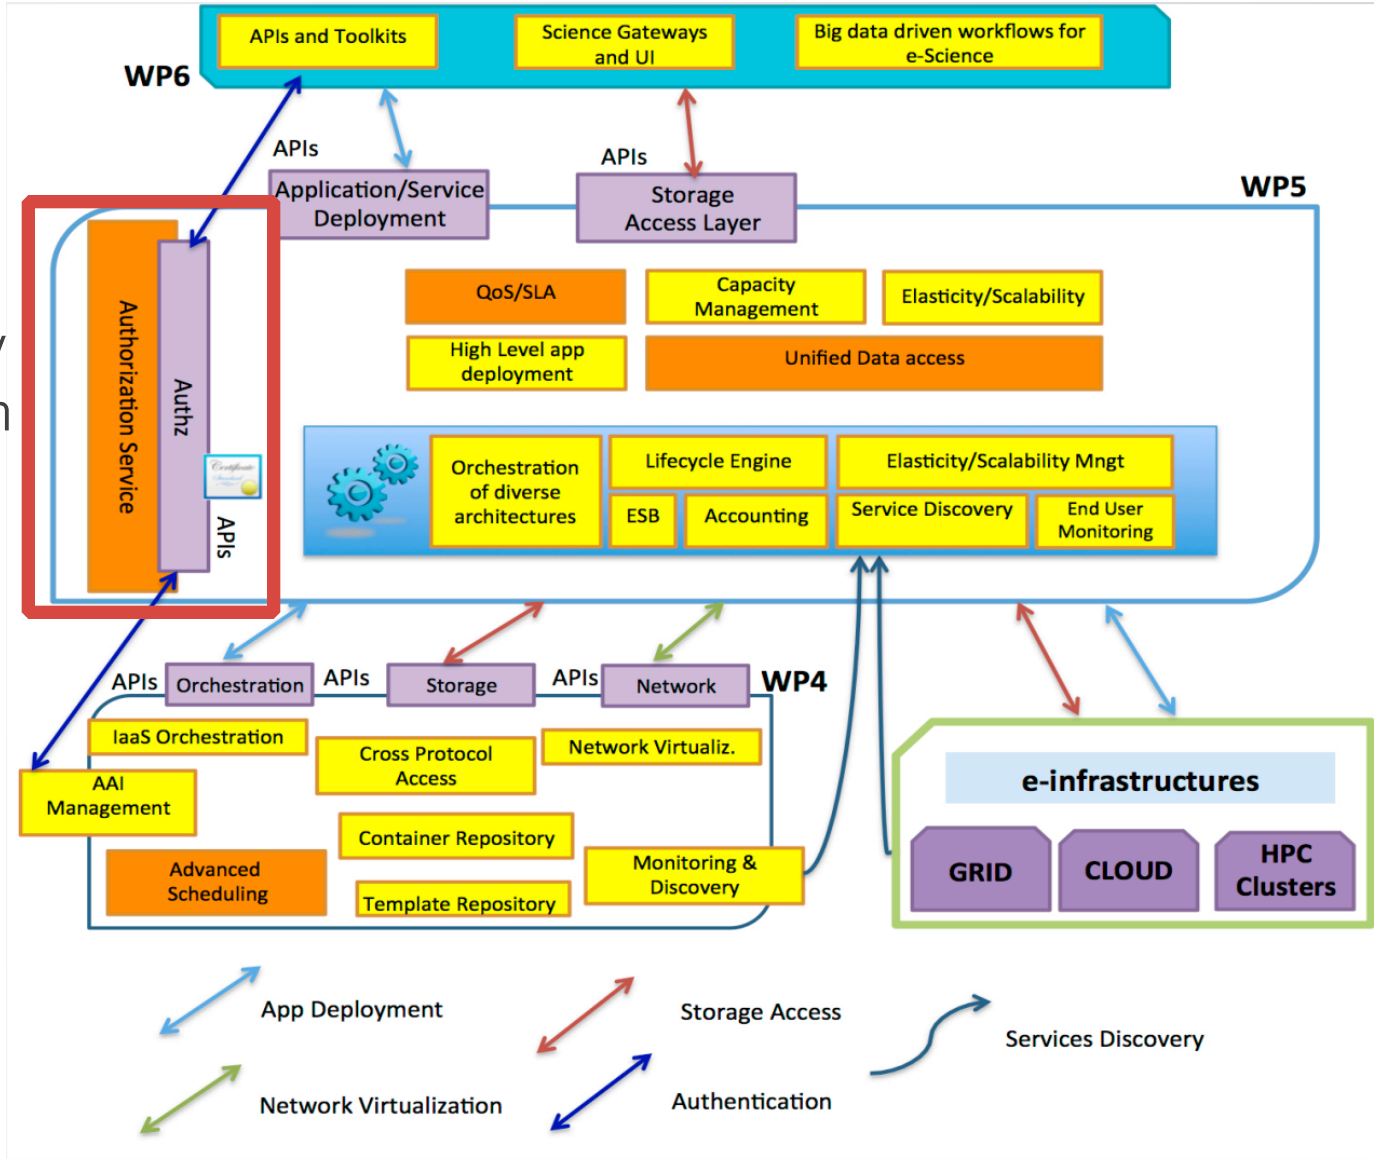

INDIGO - DataCloud

The INDIGO approach

- Based on Open Source solutions
  - widely supported by big communities
- whenever possible exploit general solutions instead of specific tools/services
  - or put effort in increasing the generality of tools developed in a given community
  - this will be important for sustainability of the architecture
- ensure that the framework offered to final users, as well as to developers, will have a **low learning curve**
  - existing software suites like ROOT, OCTAVE/MATLAB, MATHEMATICA or R-STUDIO, **will be supported** and offered in a transparent way

INDIGO - DataCloud

Scientific Portal “as a service”

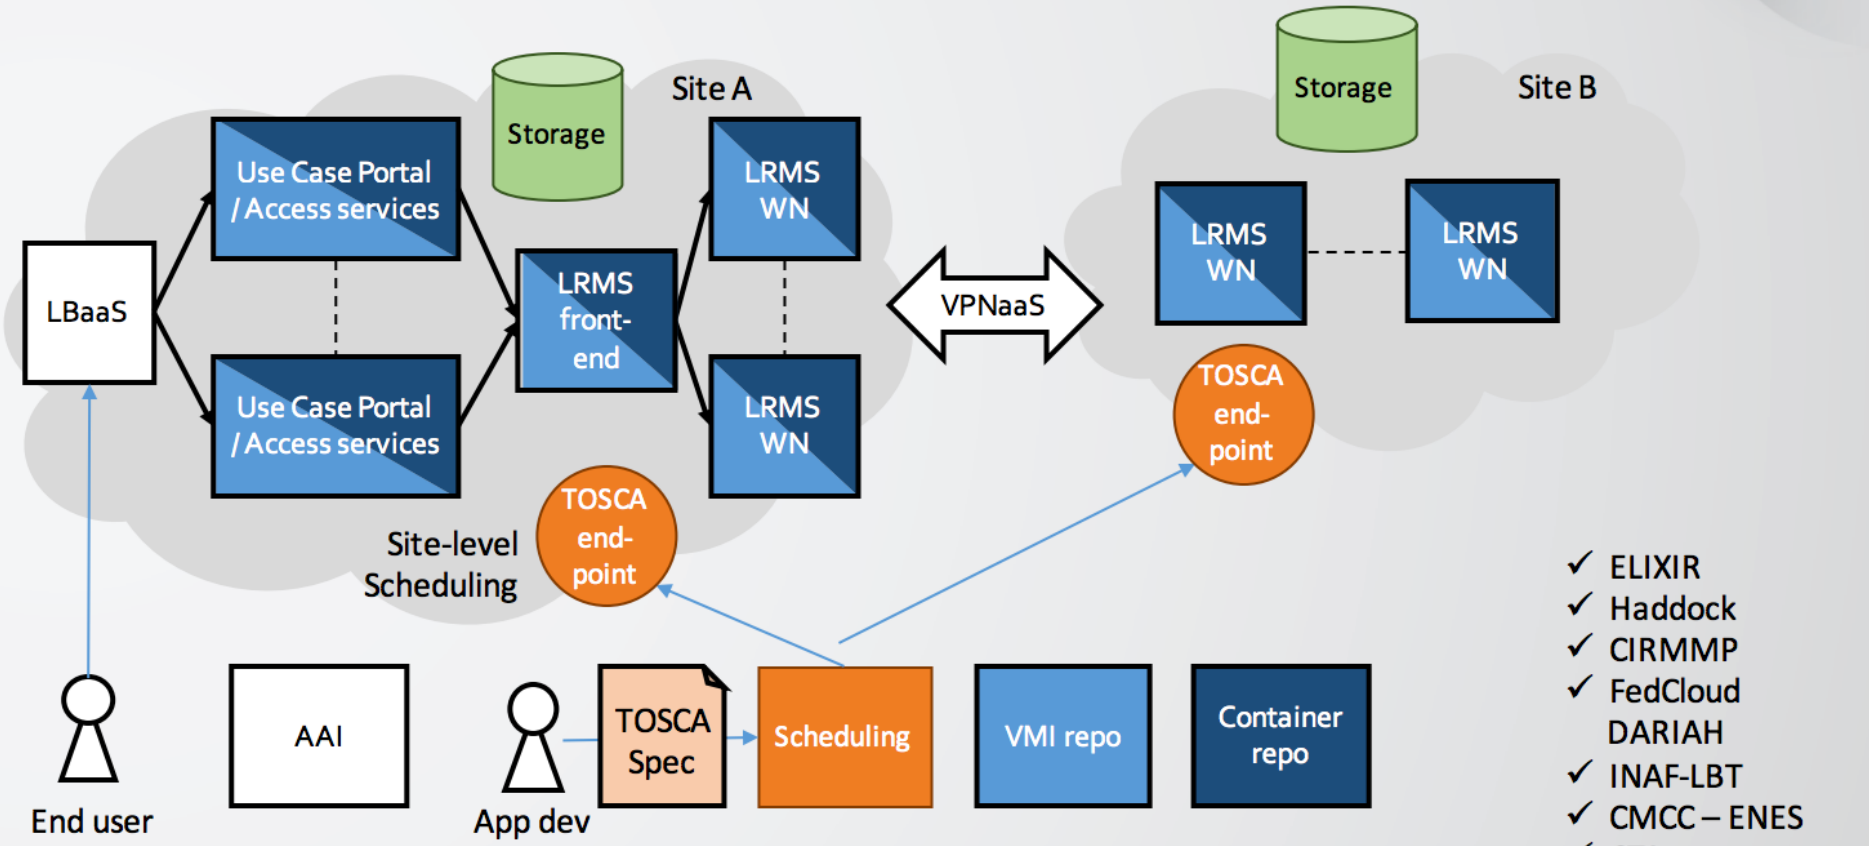

INDIGO - DataCloud

The AAI problem

- Heterogeneous infrastructures use heterogeneous authentication/authorization mechanisms
  - Hard to integrate resources from distributed infrastructures without common AAI ground
- Even where a single authentication technology is used, managing user and privileges on distributed resources in a **dynamic** and secure way is complex
- DCIs are not yet easily and securely accessible from common users
  - Federated identity support lacking or very limited

INDIGO - DataCloud

INDIGO AAI: approach

- How can we have common authN and auhtZ primitives that “just work” across several distributed infrastructures?
- Which tools should we provide to our users so that they have complete control on how authN and authZ is configured and performed on the resources (assembled from distributed providers) they will use for their research?
- How do we avoid reinventing the wheel? How do we exploit what is already available, leverage existing standards and ensure that what we develop is sustainable?

Authentication

Slide courtesy of Paul Millar

INDIGO - DataCloud

Identity layer challenges

- Support multiple AuthN mechanisms
  - SAML, OpenID-Connect, X.509
- Harmonise Identities
  - One INDIGO identity for multiple authN mechanisms
  - Persistent INDIGO identifier to services
- Support group membership and attributes
  - Linked to INDIGO identity and orthogonal to AuthN mechanism used

Identity in INDIGO

- The INDIGO identity layer speaks **OpenID-connect**
- The INDIGO Login Service is an OIDC provider
  - Authenticates users with supported AuthN mechanism
    - SAML, X.509, OIDC
  - Provides access to identity information through standard OIDC interfaces
- Can be seen as a first credential translation step

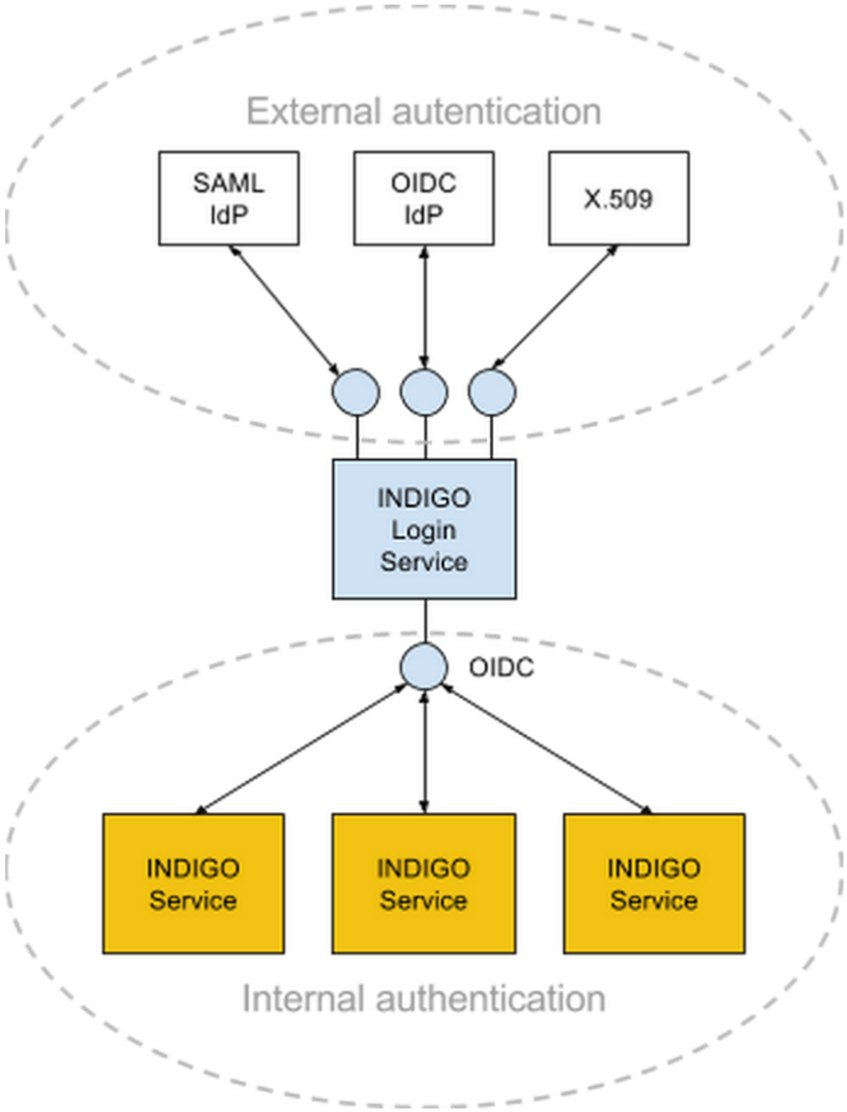

Why OpenID connect

- Standard and widely adopted in industry
  - Don’t reinvent the wheel
- Reduced client integration complexity
- Lots of things we need are covered and standardized
  - Dynamic Registration
  - Discovery
  - Token revocation
  - Session management
  - Distributed/Aggregated claims
- Friendly for mobile apps

INDIGO AuthN flow

Marcus

Access service

INDIGO Service

Marcus wants to access some service at INDIGO service

Home IdP

Indigo IAM

INDIGO AuthN flow

Marcus

Access service

INDIGO Service

INDIGO Services sees that Marcus is not authenticated, and redirects him to INDIGO IAM for authentication

Home IdP

Indigo IAM

INDIGO AuthN flow

Marcus

INDIGO Service

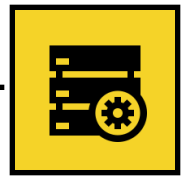

redirect to IAM
for AuthN

Home IdP

Indigo IAM

INDIGO AuthN flow

INDIGO - DataCloud

Marcus

INDIGO Service

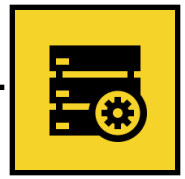

redirect to IAM
for AuthN

IAM lets Marcus choose how he wants to authenticate

Marcus chooses his Home IdP

Home IdP

Indigo IAM

INDIGO AuthN flow

Marcus

INDIGO Service

redirect to home
IdP for AuthN

Home IdP

Indigo IAM

INDIGO AuthN flow

INDIGO - DataCloud

Home IdP authenticates Marcus and sends back an AuthN assertion

INDIGO AuthN flow

Marcus

INDIGO Service

IAM validates assertion.
Marcus is now
authenticated at IAM.

Home IdP

Indigo IAM

INDIGO AuthN flow

Marcus

INDIGO Service

send back to IS
OIDC authz code

Home IdP

Indigo IAM

INDIGO AuthN flow

INDIGO - DataCloud

Marcus

INDIGO Service

exchange
authZ code
for OIDC ID-token
access token

Home IdP

Indigo IAM

INDIGO AuthN flow

INDIGO - DataCloud

Marcus

INDIGO Service

IS validates ID-Token.
Marcus is now
authenticated at IS

Home IdP

Indigo IAM

INDIGO AuthN flow

INDIGO - DataCloud

Marcus

INDIGO Service

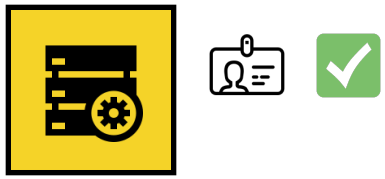

IS requests additional profile information about Marcus from IAM user info endpoint

Home IdP

Indigo IAM

INDIGO - DataCloud

Authorization

Slide courtesy of Paul Millar

INDIGO - DataCloud

Authorization challenges

- Support controlled delegation of privileges by design
- Provide the tools to support cross-organizational user and privilege management
  - Enrollment flows and group management
  - User information provisioning
- Provide tools to **dynamically** define, propagate, compose and enforce authorization policies based on identity attributes at various levels of the INDIGO stack
  - Uniform and consistent authZ over resources assembled from multiple, heterogeneous providers

AuthZ in INDIGO: OAuth 2

- A standard framework for delegated authorization to access HTTP protected resources
- Decouples AuthN from AuthZ
- Natural solution for delegated authorization in HTTP services

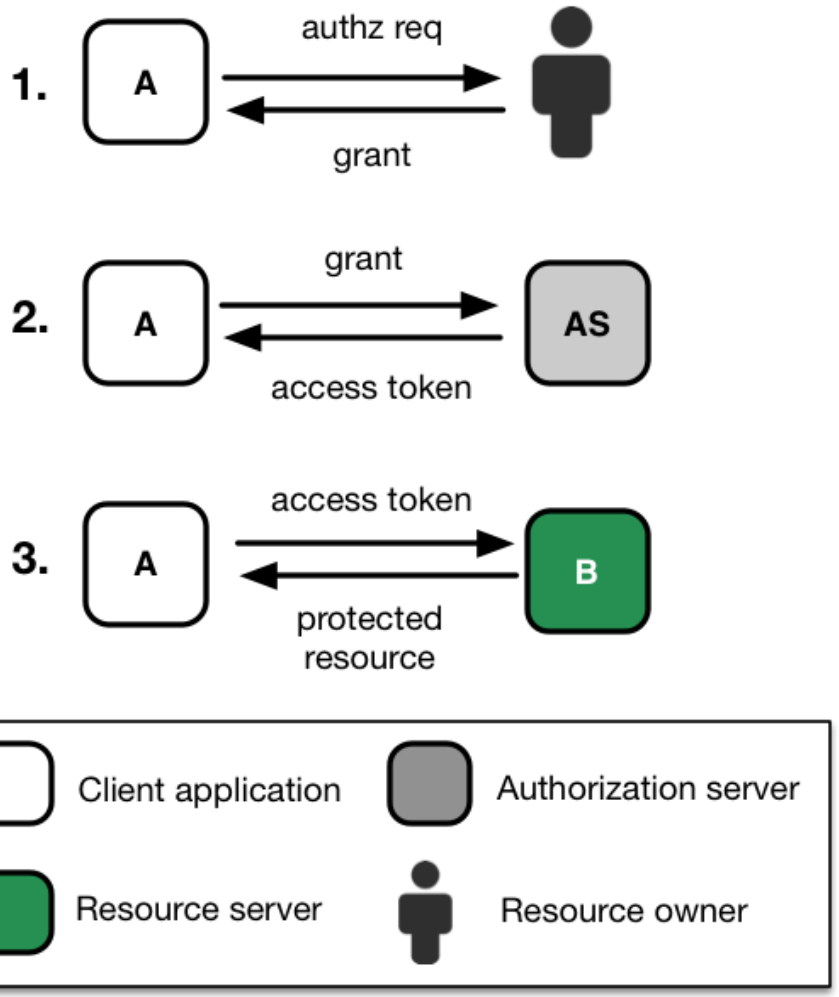

Authorization in INDIGO

- INDIGO services are modeled as APIs protected by an OAuth Authorization Service
- In order to access resources, a client needs an access token
- OAuth scopes used to
  - target the token to specific APIs
  - provide hints for local authZ
- Identity layer provides other attributes as base for AuthZ decisions

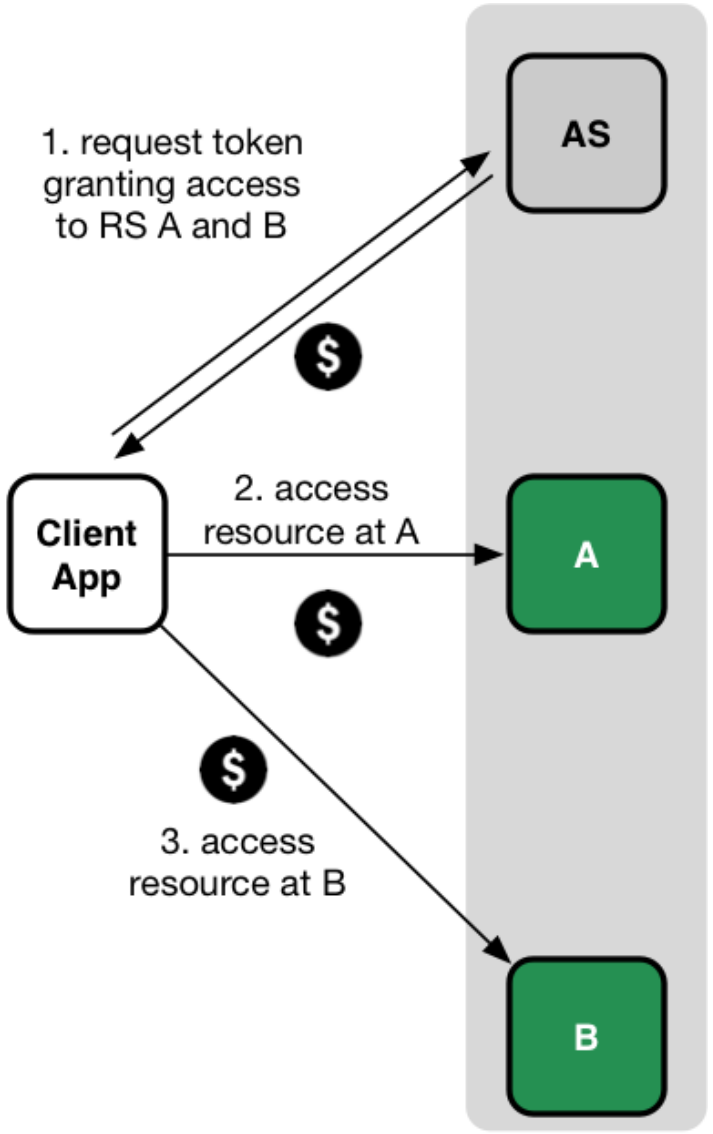

INDIGO - DataCloud

Delegation in INDIGO

- OAuth is **all** about delegating authority to a service to act on resources the user owns at another service
  - Scopes enable fine-grained controlled delegation of privileges
  - refresh tokens enable offline delegated access
- We plan to make this delegation model more flexible and secure for longer delegation chains by leveraging macaroons as OAuth tokens
  - macaroons are bearer tokens that can be further constrained along the delegation chain to limit their authority and the context of their applicability

INDIGO AAI architecture

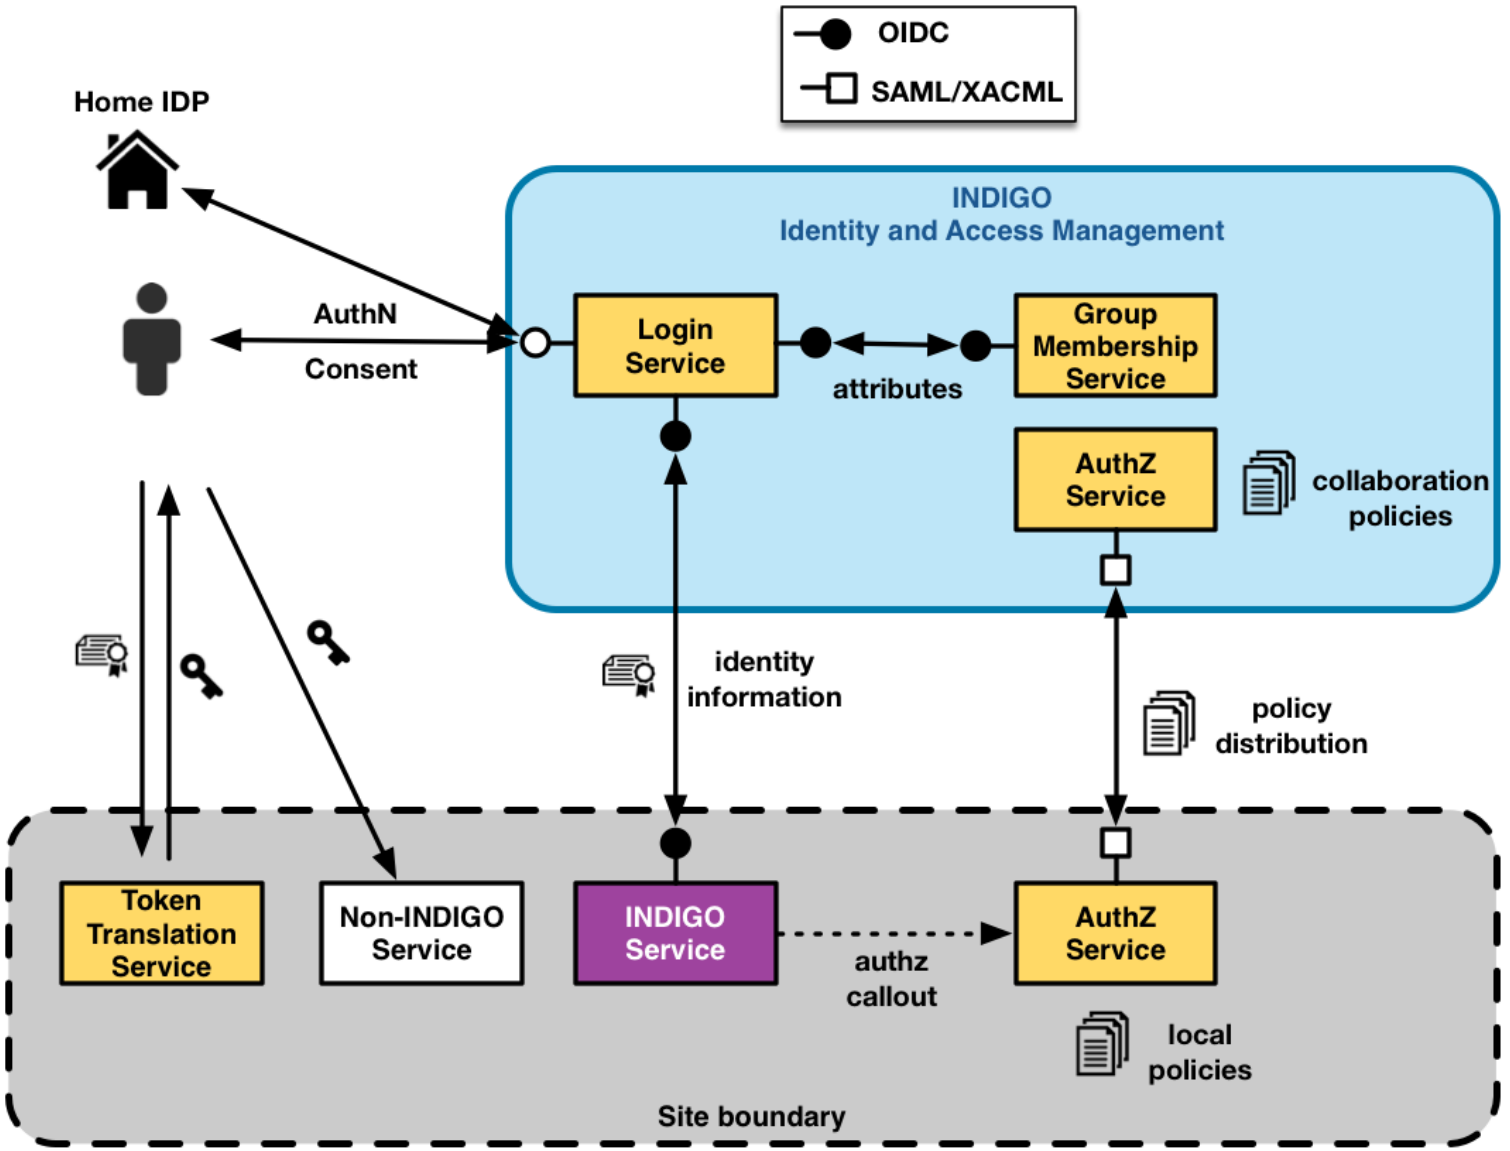

The IAM service

- Provides a unified view on identities and privileges on resources assembled from various providers
- Supports and integrates existing fed authN mechanisms
- Provides tools to define and manage enrollment flows for research communities

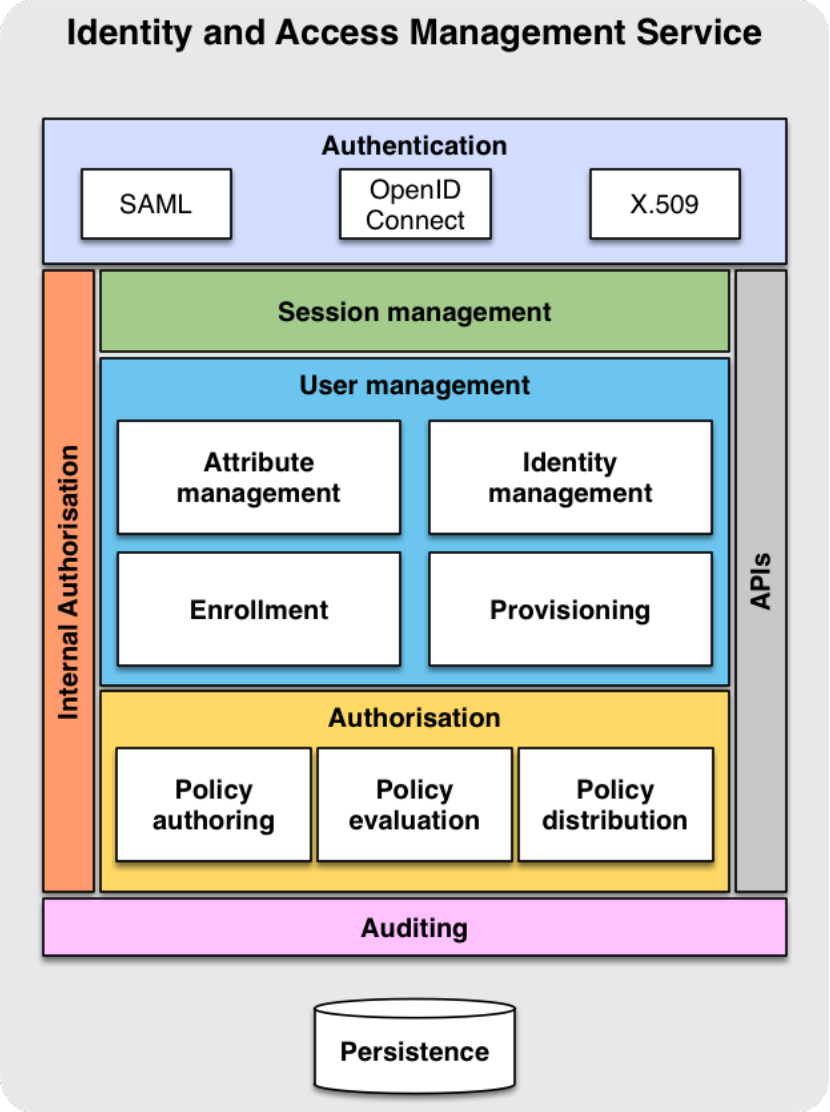

How Argus is affected

- Argus is the basis for the INDIGO IAM authorization component
- Planned Argus evolution work in INDIGO
  - Integrate support for INDIGO Identity in Argus
    - i.e. ability to define and enforce policies defined on attributes in the OIDC ID token

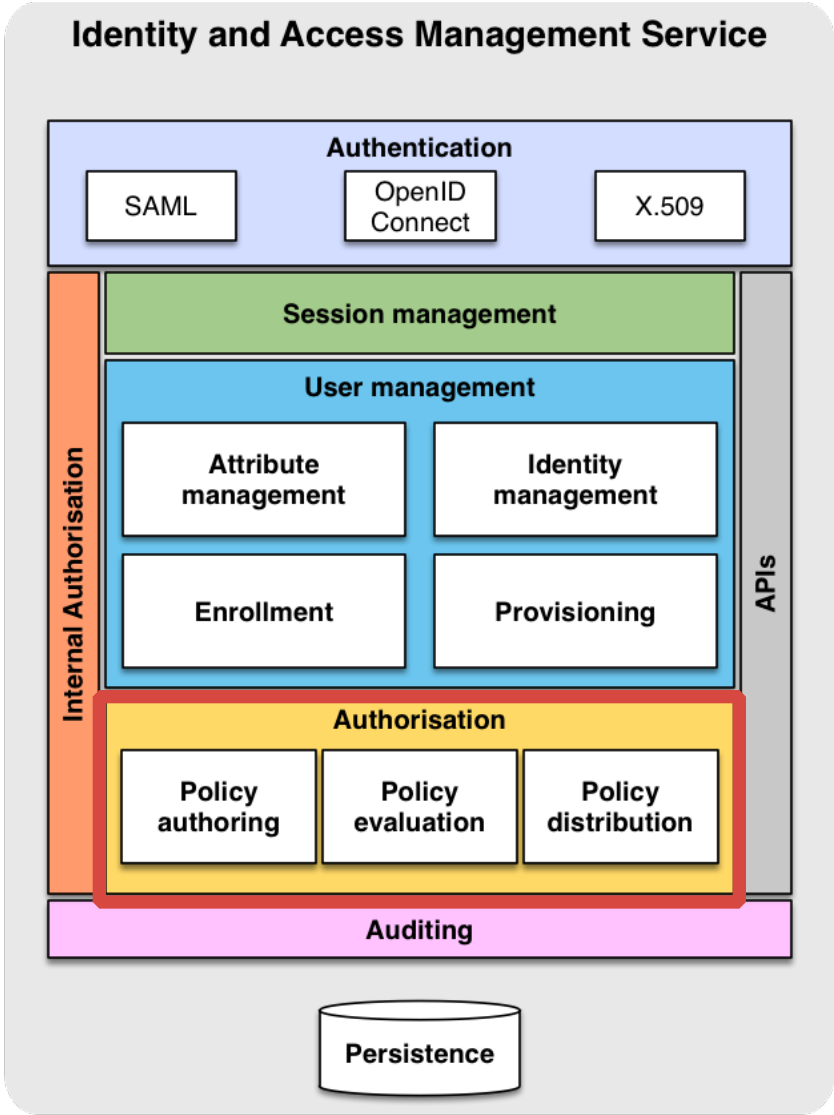

The Token translation Service

- INDIGO Token Translation Service
  - maps INDIGO identity & attributes to external service credentials
  - plugin-based architecture, will initially support translations to
    - ssh keypair
    - S3 keys
    - X.509 certificate

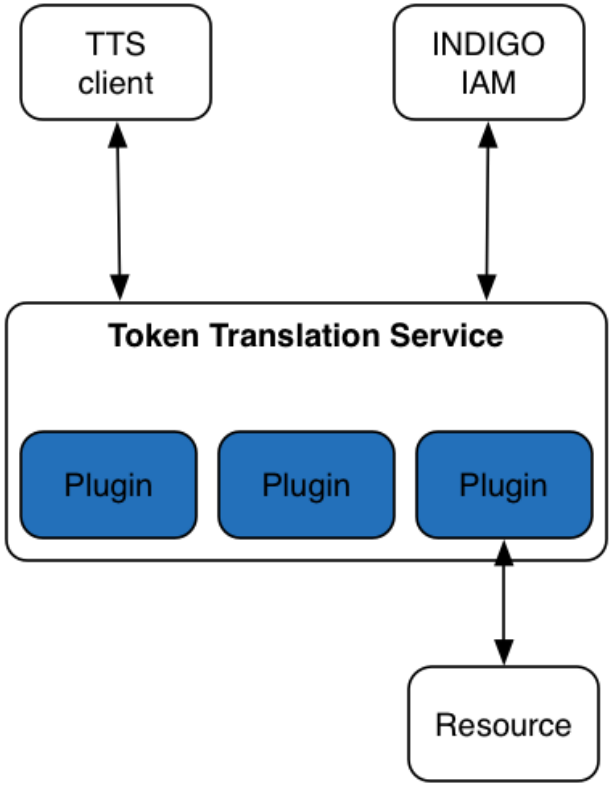

INDIGO - DataCloud

Timeline

- Architectural and design deliverables done
  - See https://www.indigo-datacloud.eu/documents-deliverables for all INDIGO deliverables
- Development activities started
- First official INDIGO release expected end of July, 2016
  - but we will start make available services as soon as they are ready enough to be tested

INDIGO - DataCloud

Thanks!

Questions?

indigo-aai-tf@lists.indigo-datacloud.eu