WLCG Transition from X.509 to Tokens: Status and Plans

CHEP 2023

Tom Dack
STFC, UKRI

On behalf of the

WLCG Authorisation Working Group

… five years ago

- Andrea Ceccanti presented the WLCG's initial plans for token migration

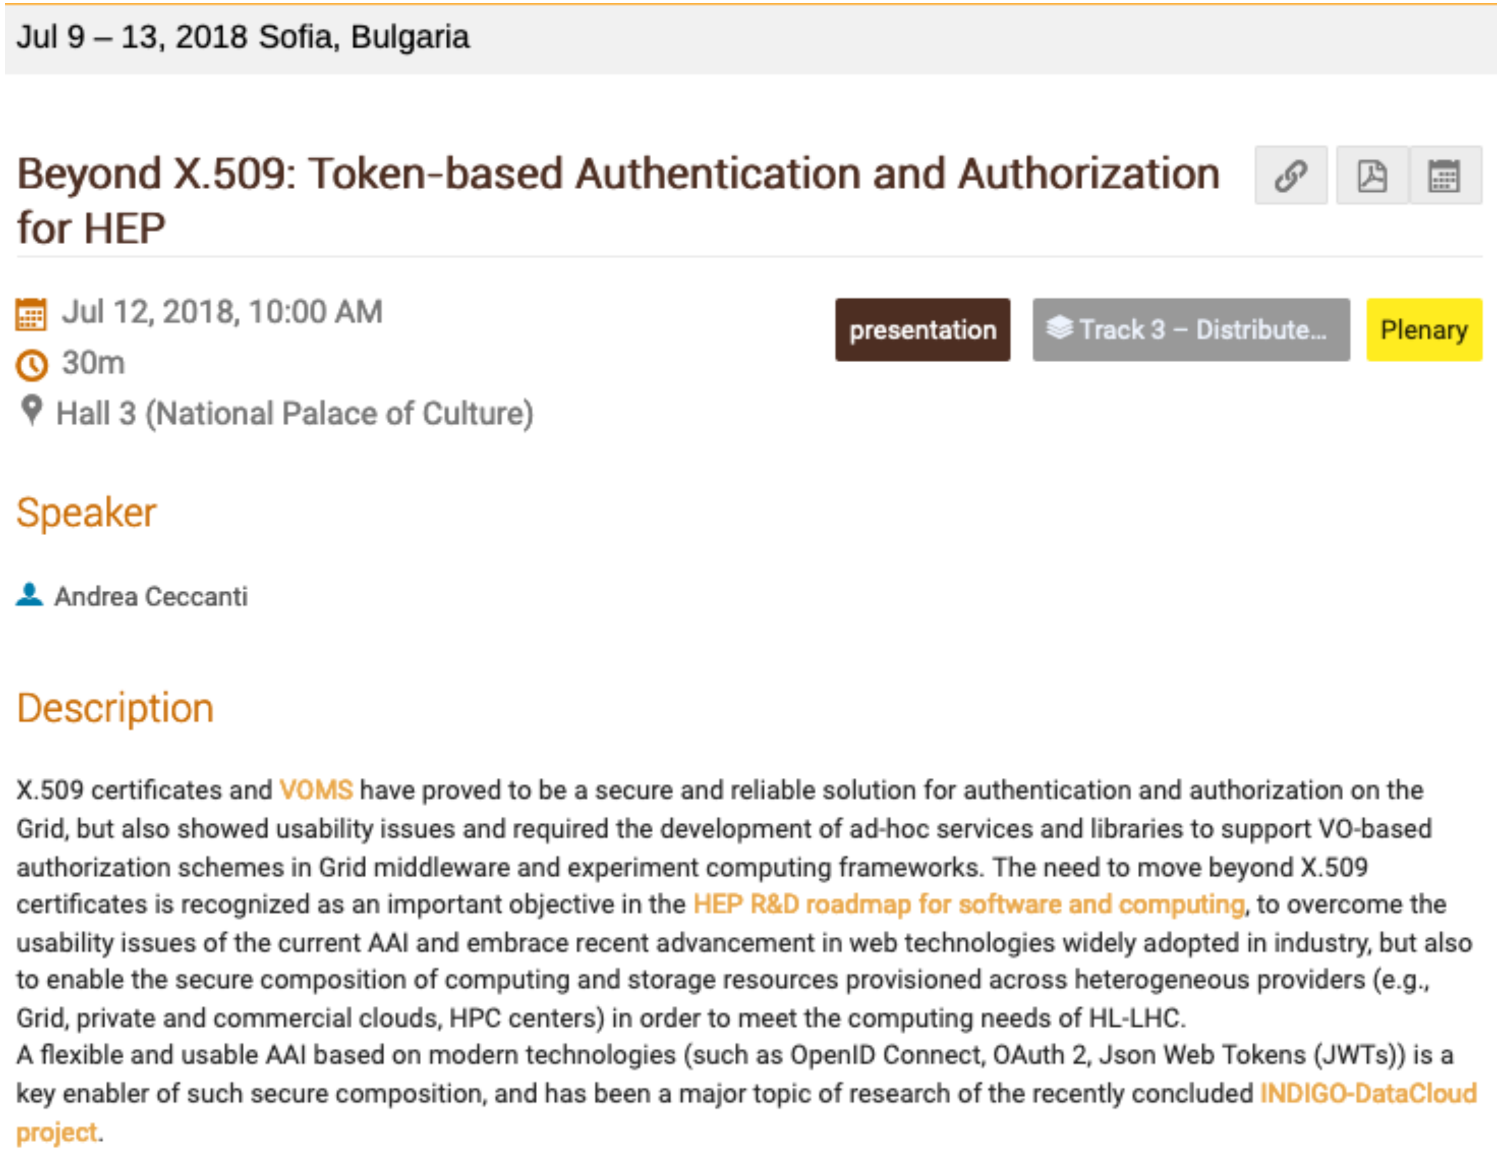

Jul 9 – 13, 2018 Sofia, Bulgaria

Beyond X.509: Token-based Authentication and Authorization for HEP

Jul 12, 2018, 10:00 AM

30m

Hall 3 (National Palace of Culture)

Speaker

Andrea Ceccanti

Description

X.509 certificates and VOMS have proved to be a secure and reliable solution for authentication and authorization on the Grid, but also showed usability issues and required the development of ad-hoc services and libraries to support VO-based authorization schemes in Grid middleware and experiment computing frameworks. The need to move beyond X.509 certificates is recognized as an important objective in the HEP R&D roadmap for software and computing, to overcome the usability issues of the current AAI and embrace recent advancement in web technologies widely adopted in industry, but also to enable the secure composition of computing and storage resources provisioned across heterogeneous providers (e.g., Grid, private and commercial clouds, HPC centers) in order to meet the computing needs of HL-LHC.
A flexible and usable AAI based on modern technologies (such as OpenID Connect, OAuth 2, Json Web Tokens (JWTs)) is a key enabler of such secure composition, and has been a major topic of research of the recently concluded INDIGO-DataCloud project.

Motivations for Change

- Industry Standards
  - Tokens are widely adopted in industry, leading to software which already supports tokens, and many standard libraries when integration is needed
- Ease of Use
  - X.509 certificates are difficult to handle for users, and unfamiliar to those entering
  - Prevalence of Token flows in industry/social applications makes the workflow easier to learn
- Flexible Authentication
  - Hard to integrate identity federations with VOMS, where a token infrastructure brings us steps closer towards this
- Finer Grained Access Control
  - Tokens provide direct routes for AuthZ relying on user attributes or capabilities
- Security Benefits
  - The usage of audiences and scopes allows restricting tokens to specific jobs or services with greater levels of flexibility

Token Based AAI for WLCG

- VO Scoped Token issuers, which can:
  - Support multiple authentication mechanisms
    - Currently supporting CERN SSO or X.509
  - Provide users with a persistent VO-Scoped identity
  - Expose identity information, attributes and capabilities to services via JWT tokens and standard OAuth & OpenID Connect protocols
  - Integrate with existing VOMS-aware services
  - Support both Web and non-Web access, delegation, and token renewal

- WLCG uses **INDIGO IAM** as its issuer

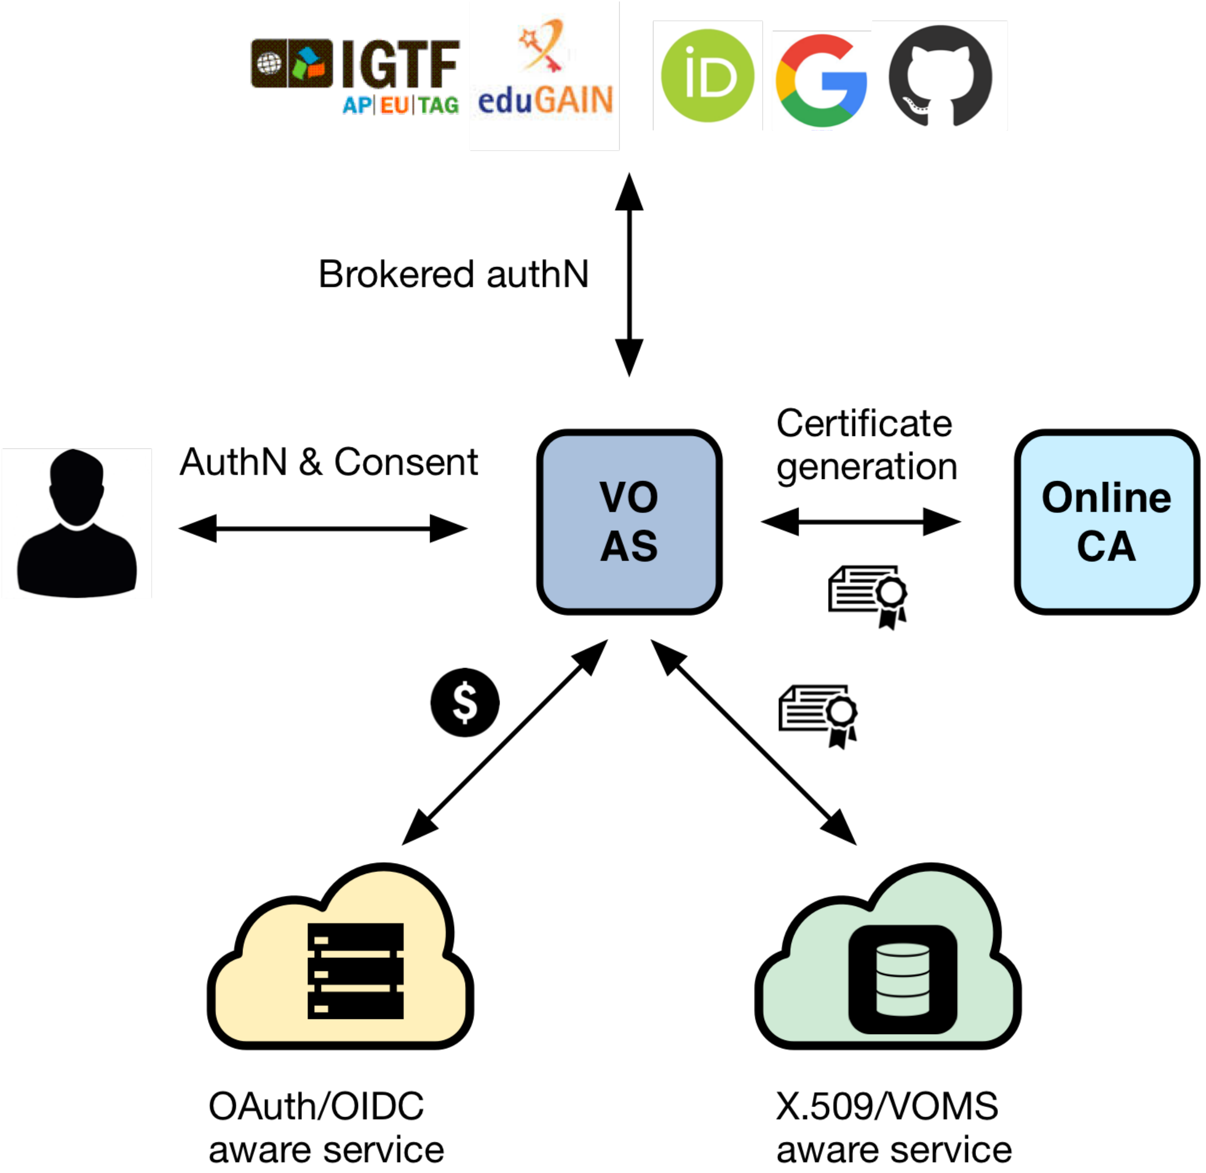

WLCG Token Infrastructure Design

Follows the AARC Blueprint: https://aarc-community.org/architecture/ but not all AEGIS recommendations

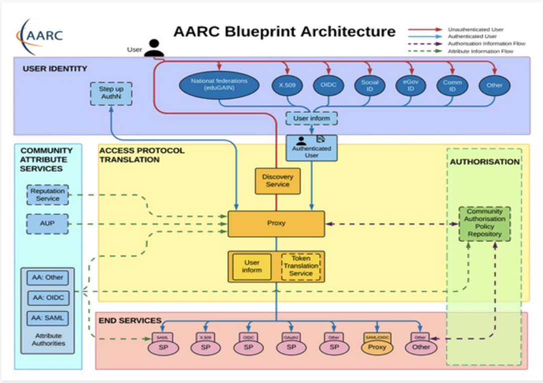
AARC
AARC Blueprint Architecture
USER IDENTITY
COMMUNITY ATTRIBUTE SERVICES
ACCESS PROTOCOL TRANSLATION
AUTHORISATION
END SERVICES
Discovery Service
Proxy
Token Translation Service
Community Authorisation Policy Repository

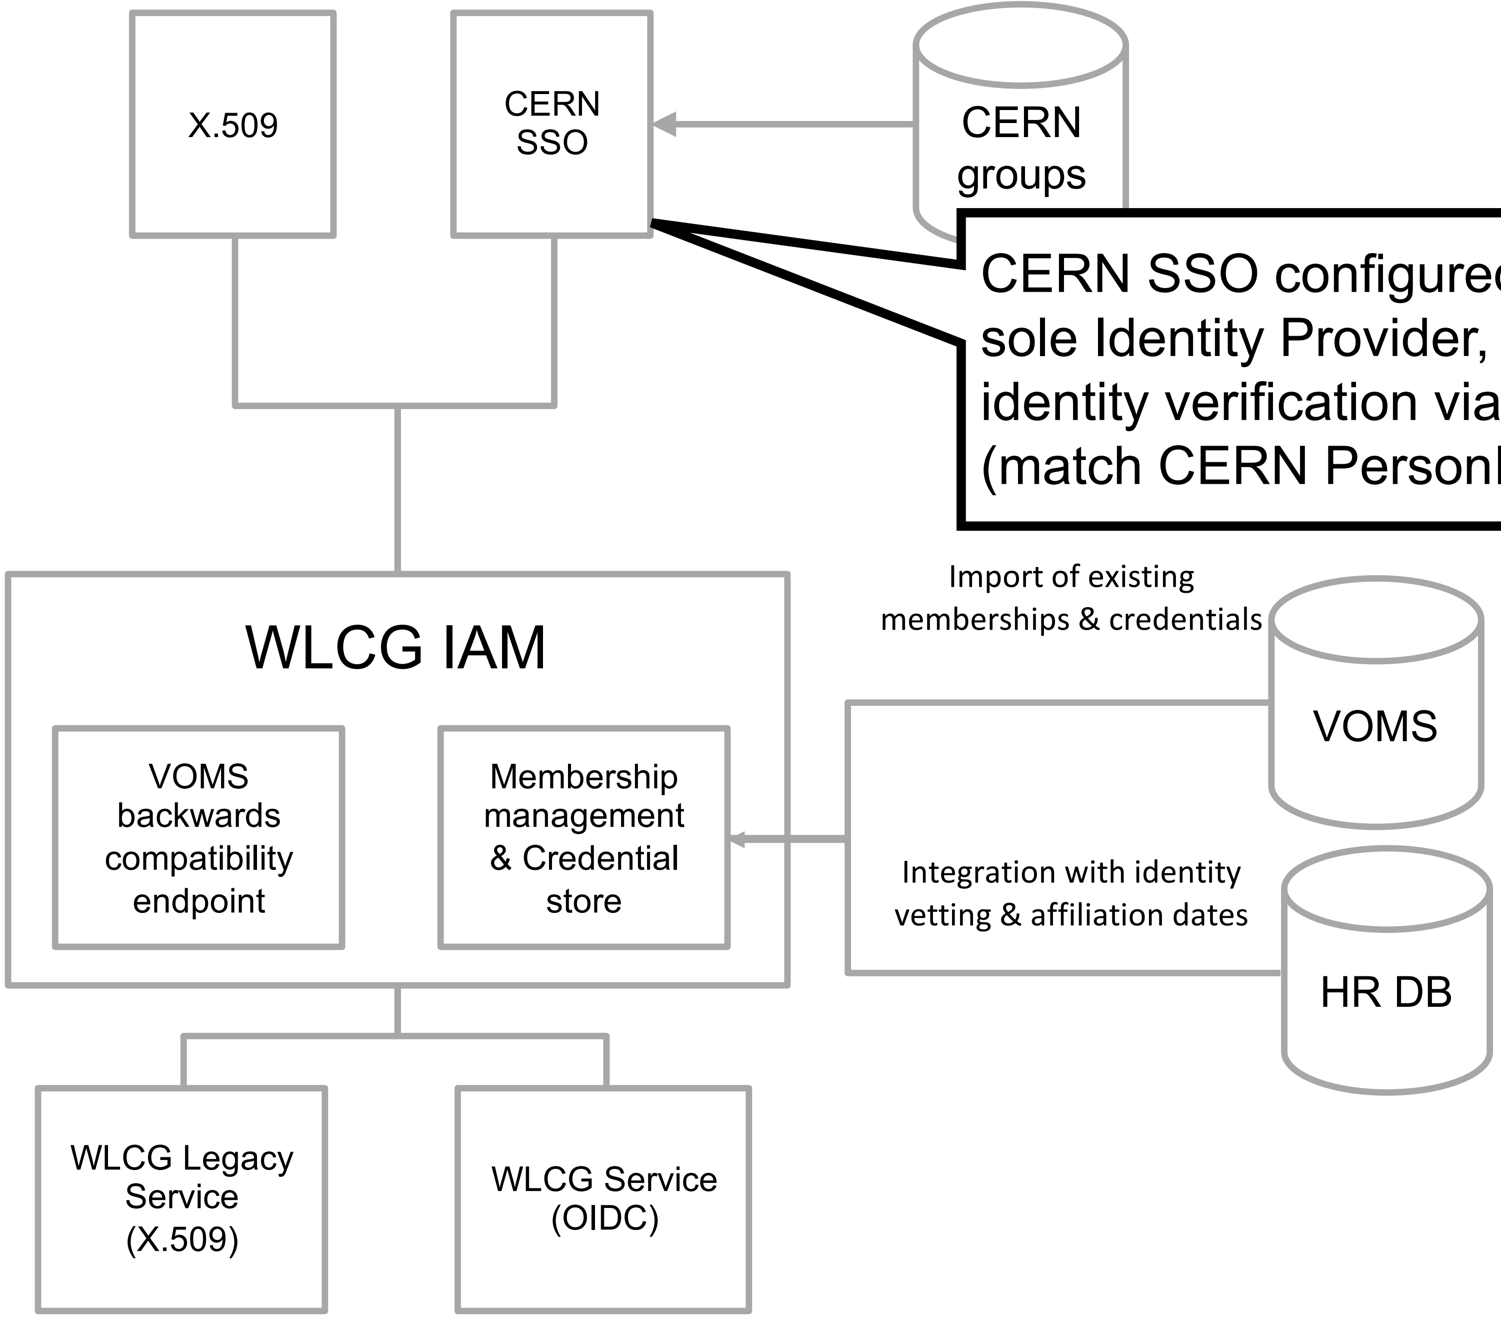
X.509
CERN SSO
CERN groups
WLCG IAM
VOMS backwards compatibility endpoint
Membership management & Credential store
Import of existing memberships & credentials
VOMS
Integration with identity vetting & affiliation dates
HR DB
WLCG Legacy Service (X.509)
WLCG Service (OIDC)

CERN SSO configured as sole Identity Provider, enables identity verification via HR DB (match CERN PersonID)

WLCG Token Infrastructure Design

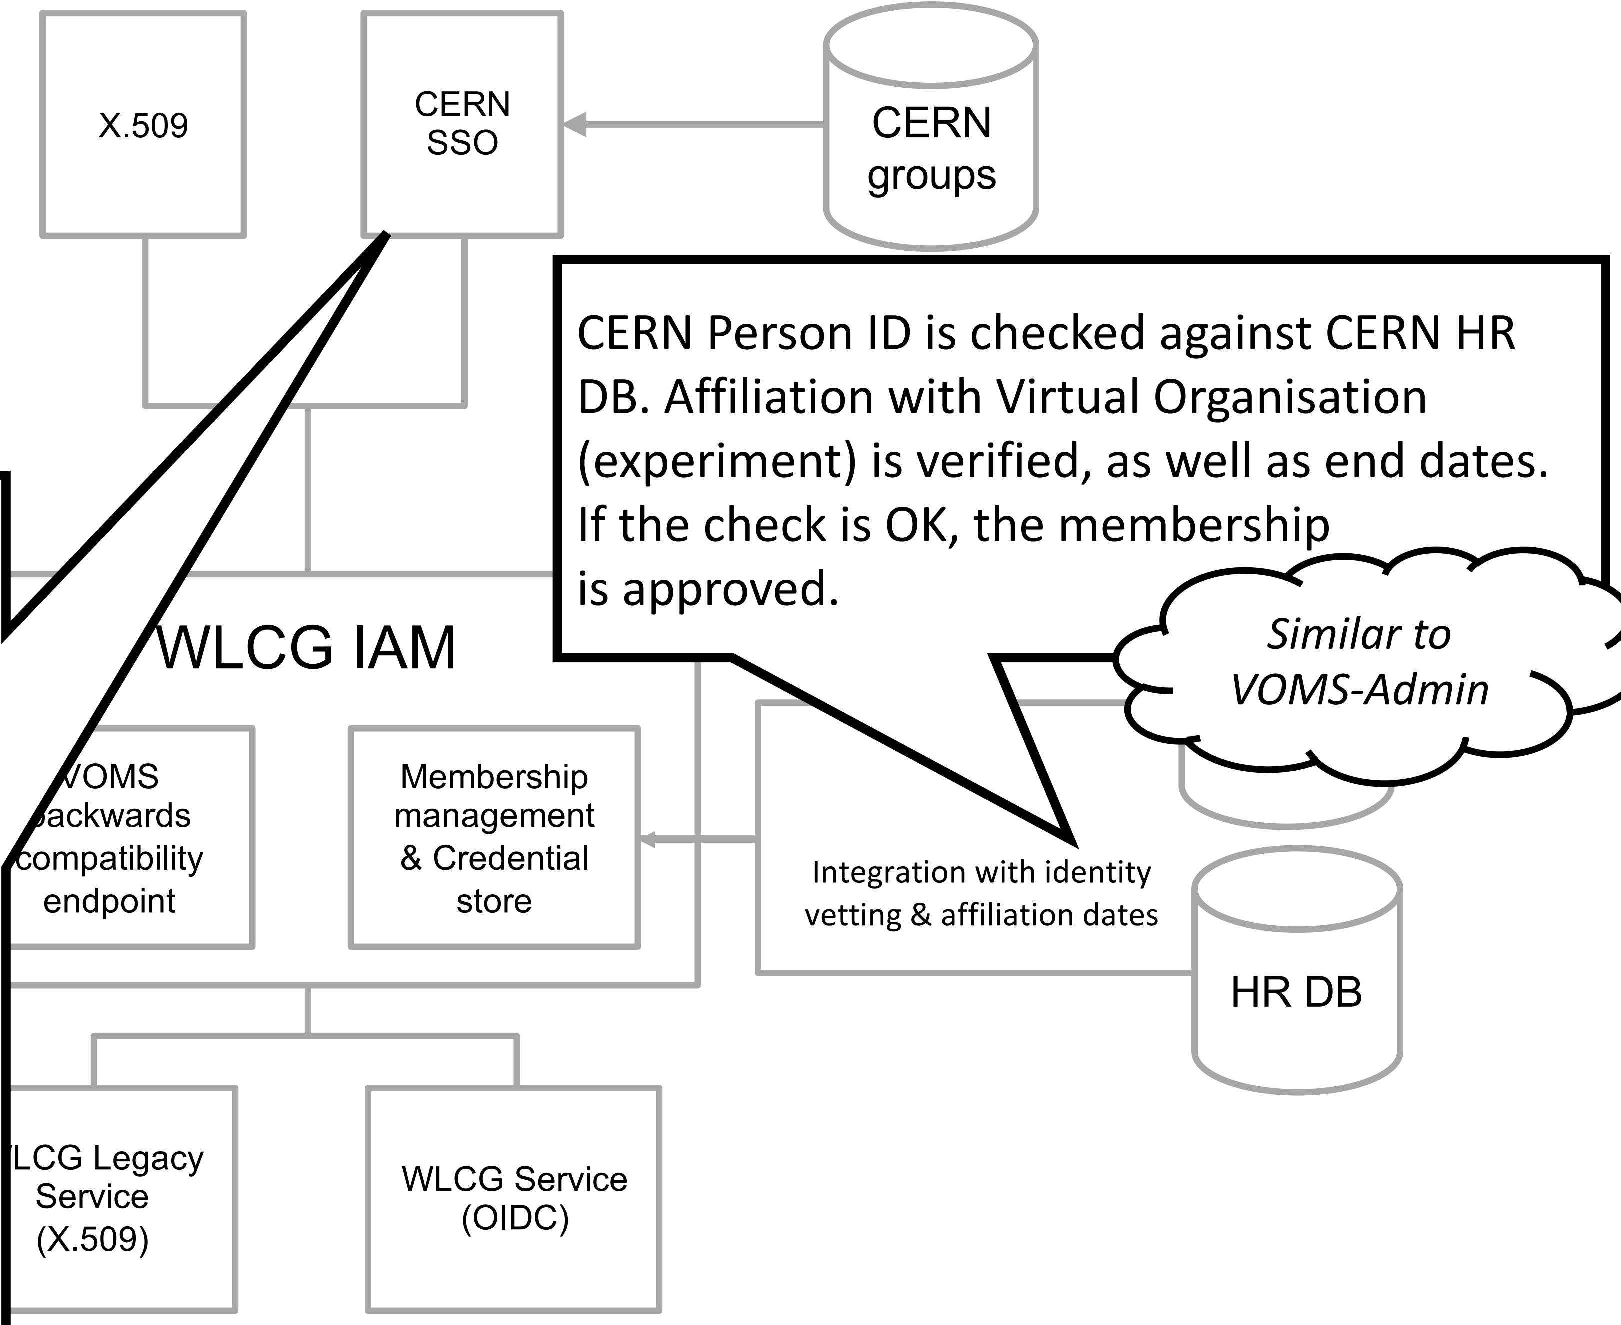
X.509
CERN SSO
CERN groups
WLCG IAM
VOMS backwards compatibility endpoint
Membership management & Credential store
Integration with identity vetting & affiliation dates
HR DB
LCG Legacy Service (X.509)
WLCG Service (OIDC)

CERN Person ID is checked against CERN HR DB. Affiliation with Virtual Organisation (experiment) is verified, as well as end dates. If the check is OK, the membership is approved.

Similar to VOMS-Admin

CERN SSO releases:

- Name,
- Email,
- CERN Person ID (indicates HR has performed ID check),
- CERN Kerberos Principal
- ...

Currently LHC experiment members can have CERN accounts, and so all access is via CERN SSO or Linked X.509, but aim is to work towards removing this need in future

WLCG Token Infrastructure Design

Outbound communication of user information to end services, which can be used to establish authorization.

VOMS proxies can still be used for as long as they are needed.

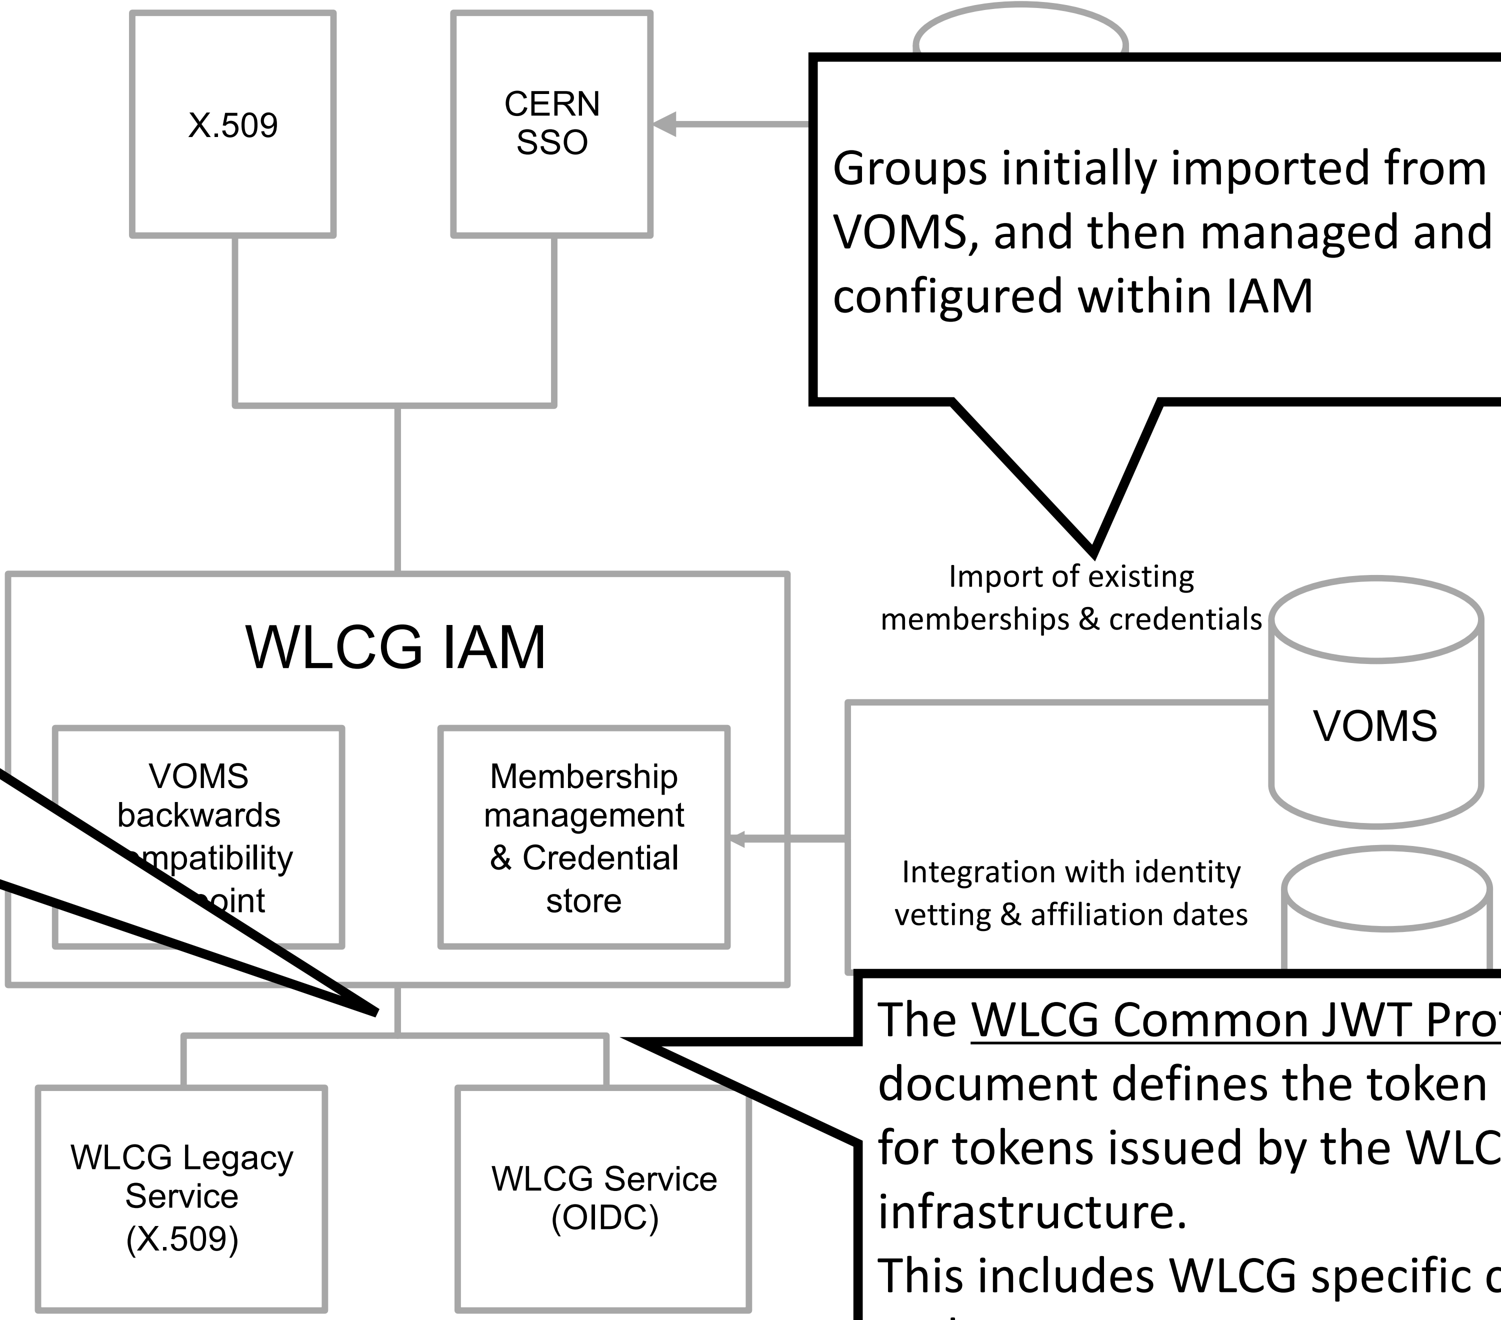

X.509

CERN SSO

WLCG IAM

VOMS backwards compatibility endpoint

Membership management & Credential store

Import of existing memberships & credentials

VOMS

Integration with identity vetting & affiliation dates

WLCG Legacy Service (X.509)

WLCG Service (OIDC)

Groups initially imported from VOMS, and then managed and configured within IAM

The WLCG Common JWT Profiles document defines the token profile for tokens issued by the WLCG infrastructure.
This includes WLCG specific claims and scopes.

Token uptake in WLCG and beyond

- IAM instances are in production for the four LHC experiments
  - Contents are automatically replicated *from* VOMS-Admin
    - Modulo known issues with workarounds
  - IAM VOMS endpoints have been in production use by ATLAS and CMS for over a year, alongside the legacy VOMS services

- ATLAS **must**, and CMS & SAM ETF **can** use tokens to *submit* pilot jobs to HTCondor CEs
  - These jobs still use X.509 VOMS proxies for **data management,** etc.
  - In particular the CEs on **OSG**, which only support tokens since May 2022
  - Still rely on X.509 for ARC-CEs

- For ALICE, LHCb, Belle-II, … these matters are currently WIP

- DUNE and Fermilab are progressing with tokens from *CILogon*
  - This supports the WLCG token profile, with DUNE using WLCG tokens
  - Hashicorp *Vault* is used to store refresh tokens, with *htgettoken* as a user interface

WLCG INDIGO IAM Deployments

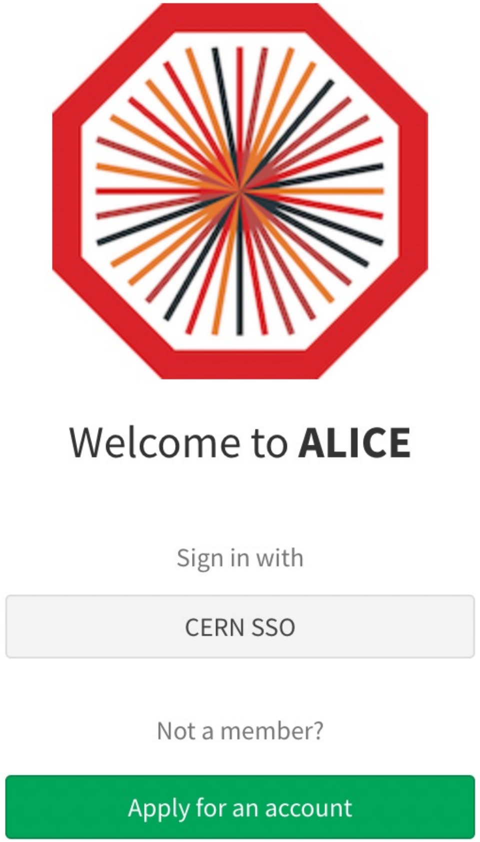
Welcome to **ALICE**

Sign in with

Not a member?

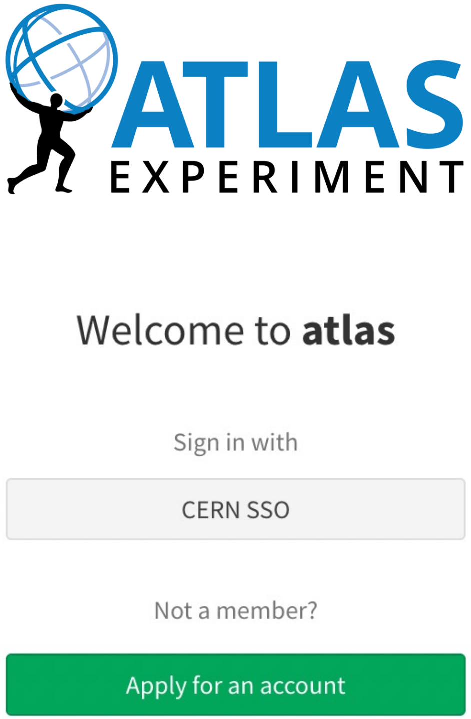

Welcome to **atlas**

Sign in with

Not a member?

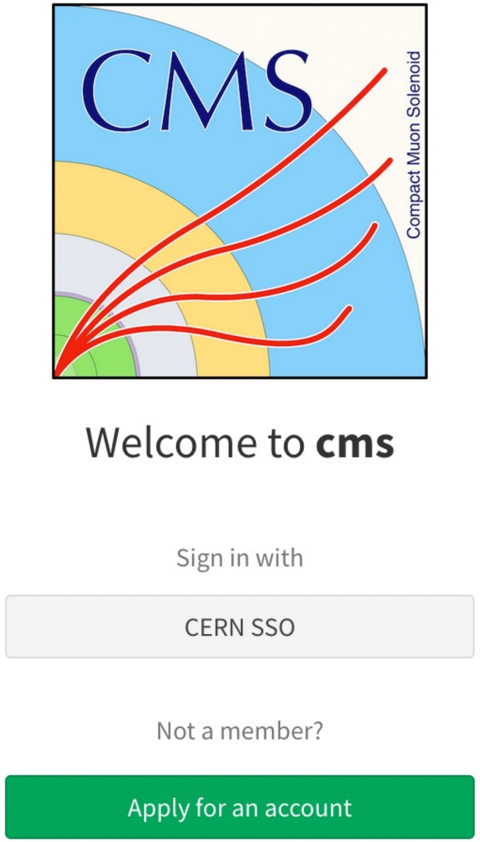

Welcome to **cms**

Sign in with

Not a member?

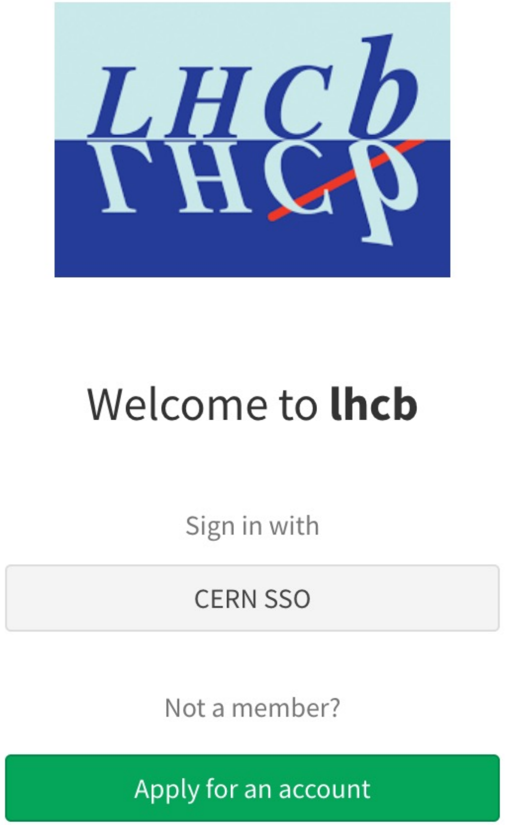

Welcome to **lhcb**

Sign in with

Not a member?

IAM Dashboard: https://<experiment>--auth.web.cern.ch
VOMS endpoint: https://voms-<experiment>-auth.app.cern.ch

WLCG INDIGO IAM Deployments

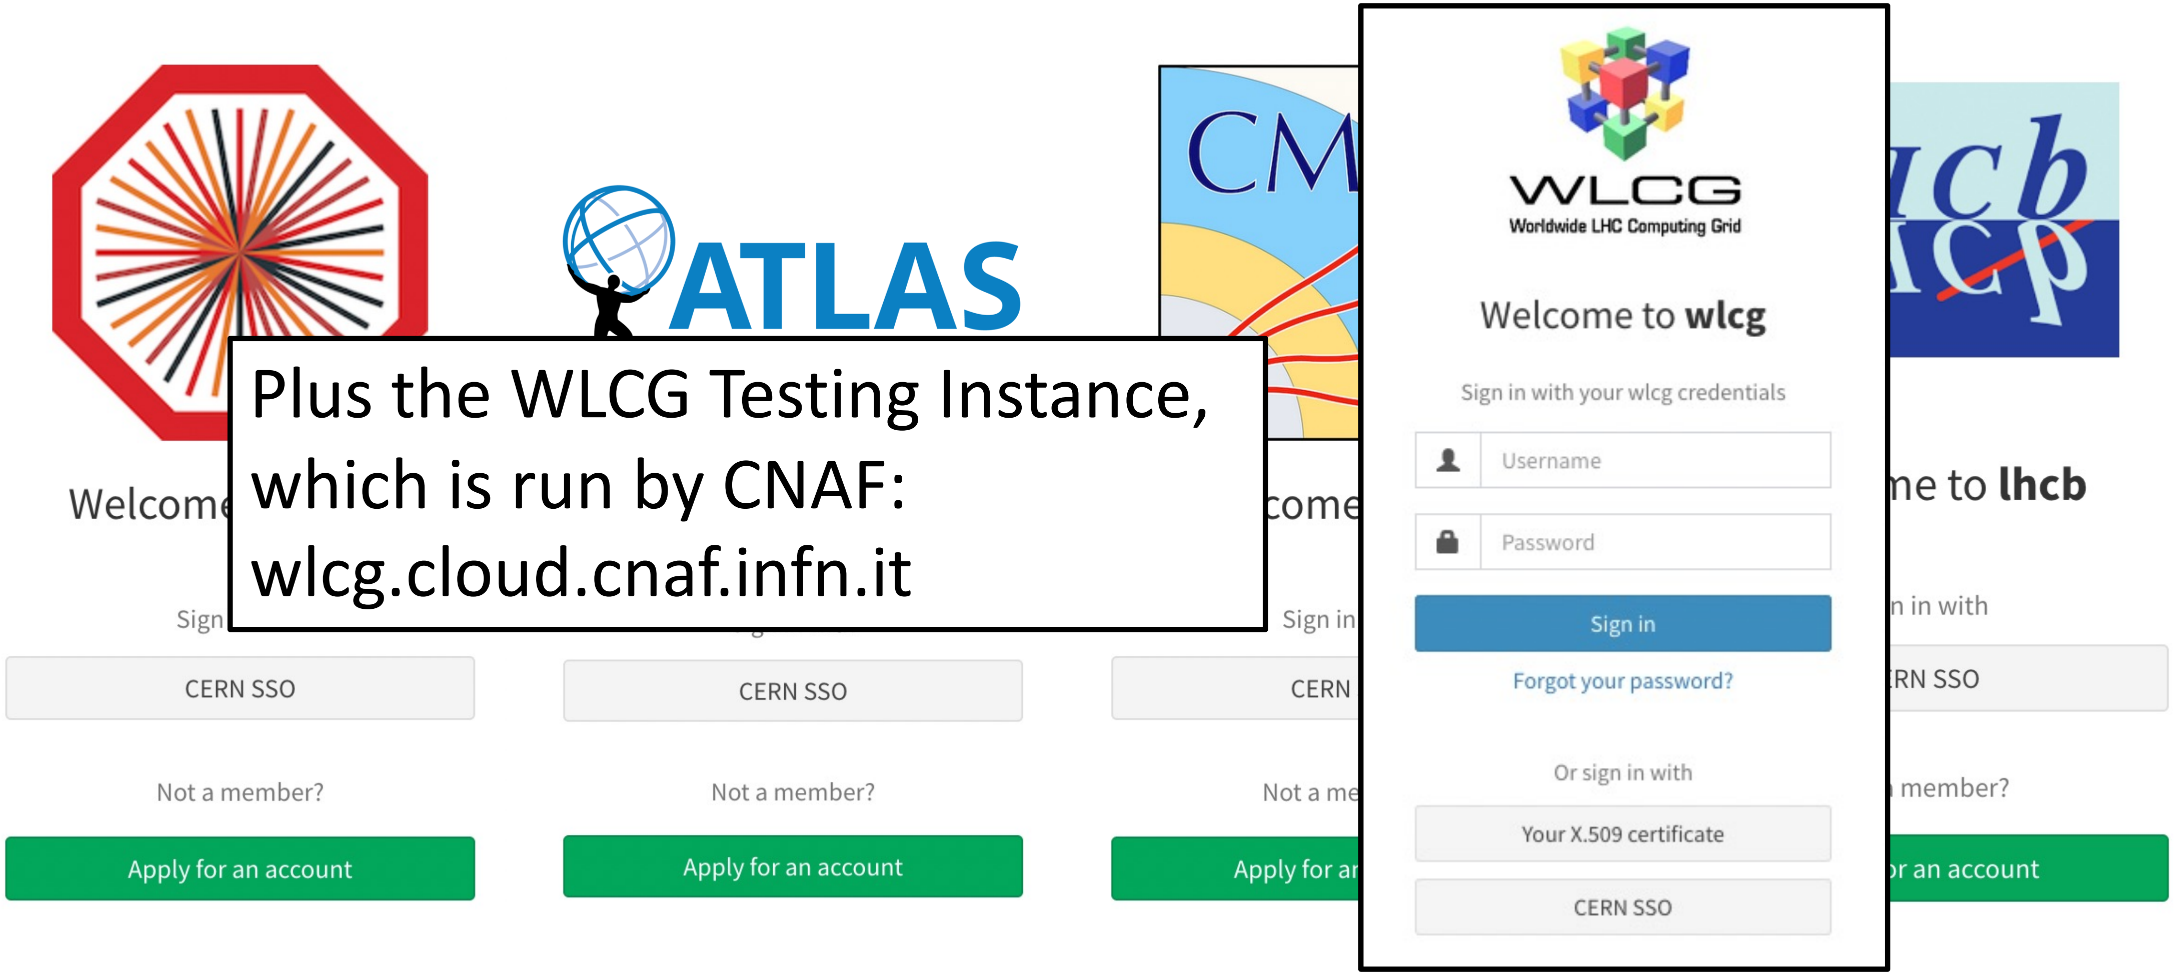

Welcome

Sign

CERN SSO

Not a member?

Apply for an account

CERN SSO

Not a member?

Apply for an account

come

Sign in

CERN

Not a me

Apply for ar

WLCG
Worldwide LHC Computing Grid

Welcome to **wlcg**

Sign in with your wlcg credentials

Username

Password

Sign in

Forgot your password?

Or sign in with

Your X.509 certificate

CERN SSO

me to **lhcb**

n in with

RN SSO

member?

r an account

Plus the WLCG Testing Instance, which is run by CNAF: wlcg.cloud.cnaf.infn.it

IAM Dashboard: https://<experiment>--auth.web.cern.ch
VOMS endpoint: https://voms-<experiment>-auth.app.cern.ch

Deployment Technical Details

- Deployed on CERN’s **Openshift** infrastructure
- IAM runs in **Docker** containers
- Configuration managed using CERN’s **gitlab**
- Logs sent to Elasticsearch
- Deployment managed by **Kubectl**
- Sectigo certificate for IAM dashboard
- CERN Grid Host Certificate for VOMS endpoint
    - CERN’s CP/CPS was updated to allow this with EUGridPMA approval

Token Transition Status

From the Token Transition Timeline:

- M.2 (Dec 2022) DIRAC versions supporting job submission tokens deployed for concerned VOs
  - LHCb have upgraded to v8.0 and validated job submission to HTCondor and ARC CEs with tokens
  - Currently planned for Belle-II in summer
- M.3 (Feb 2023) VOMS-Admin is switched off for one or more experiments
  - Supporting work underway, but pushed back to allow for further IAM development to improve on VO use-cases and concerns
- M.4 (Mar 2023) HTCondor installations at EGI sites have been upgraded to supported versions > 9.0.x
  - Work in progress, but expected date postponed.
- M.5 (Mar 2023) End of HTCondor support for GSI Auth (link)
  - Postponed to May
- M.6 (Mar 2023) Some storage endpoints provide support for tokens
  - Several CMS production storage services already pass token tests
  - A number of early adopters within ATLAS

For further details on M.4 and M.5, see the March GDB update presentation

Token Transition Status

From the Token Transition Timeline:

- M.7 (Feb 2024) Rucio / DIRAC / FTS have sufficient token support in released versions to perform DC24 using token authorization.
  - Currently WIP
- M.8 (Mar 2024) Sufficient storage endpoints support tokens to allow DC24 to be done using only tokens.
  - Having WLCG token support ready in time for DC24 is a major current emphasis
- M.9 (Mar 2025) Grid jobs use tokens for reading and stageout.
  - Requiring changes also *inside* the job pilots

Conclusions and outlook

- Though the token transition timeline has seen some delays, progress has been steady in many areas concerned

The main milestones at this time are about:

- The switch to HTCondor CE versions that no longer support GSI
  - Several scenarios to smooth the transition for legacy use cases
- Integration and deployment ready to run DC24 using tokens only
  - Aim for production-quality infrastructure at the majority of sites (preferably all T0/1 plus the big T2 sites)

- To be continued as we work towards:
  - ***M.10 (Mar 2026) Users no longer need X509 certificates!***
    - There is a longer gap between M.9 (Mar 25) and M.10 (Mar 26), so as to allow the development of scenarios and workflows to ensure a smooth *user* transition