Future Authorisation for WLCG

GDB

May 9th 2018

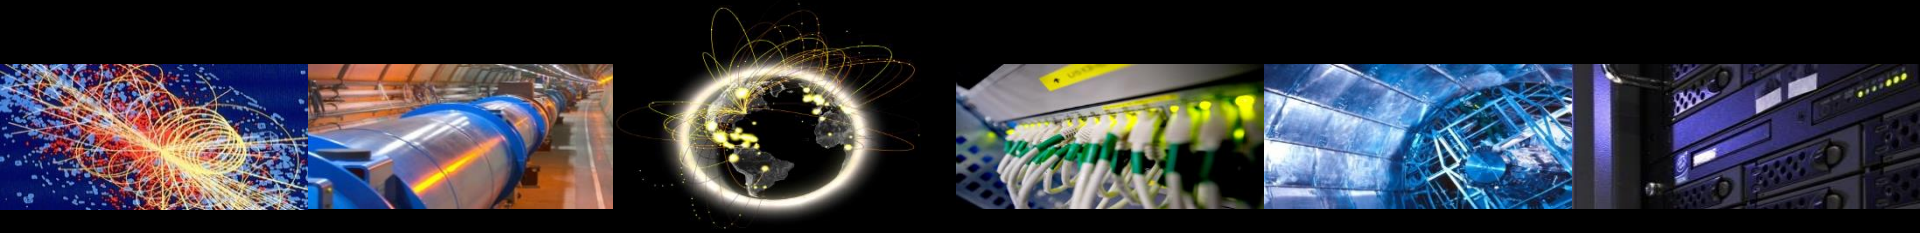

Agenda

- What is the WG Doing?
- Where do we need input?
- Next steps

What is the AuthZ WG doing?

Winter 2017

- WLCG AAI Requirements Gathering
- Identify Possible Solutions and launch WLCG AAI Pilots

Spring 2018

- Define authorisation token schema for use by infrastructures and VOs
- Monitor WLCG AAI Pilots

Autumn 2018

- Fully analyse operational impact
- Further liaison with VO management

Winter 2018

- Assess WLCG AAI Pilots

Why are we doing this?

- Evolving Identity Landscape
  - User-owned x.509 certificates -> federated identities
  - Current grid middleware does not support federated identities
  - How can we shield users from the complexities of X.509 certificate management ?
  - Token-based authorization widely adopted in commercial services and increasingly by R&E Infrastructures
- Data Protection
  - Tightening of data protection (GDPR) requires fine-grained user level access control, certain provisioning practices may need to be adjusted

Objective: Understand & meet the requirements of a future-looking AuthZ service for WLCG experiments

Where do we need input?

- Requirements
- Solution design
- Token Schema

Requirements

- VO Membership Management
  - Attributes? VO ID, ID of credential, Name, Email, Authorization
  - Support multiple federated credentials & their linkage
  - Active role selection
  - Token management achievable by the standard user
- Service Requirements
  - Attributes? Authorization plus traceability || Groups/Roles
  - Ease of implementation
  - Use standard approaches
  - Token integrity and validity verifiable
    - Without connecting to the issuer
  - For non-web, users should not have to manage identities in addition to their login session
- General
  - Support for fine grained suspension
  - Smooth transition from current X509-based to token-based AAI

Are these the correct requirements? What have we missed?

Solution Design

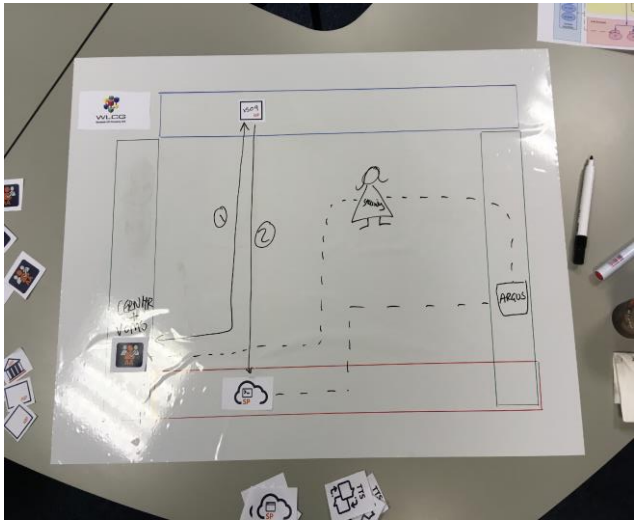

Current infrastructure allows access based on X509, including VOMS, CERN HR DB and Argus

Future infrastructure should support a range of credential types for users and services and provide a user friendly experience

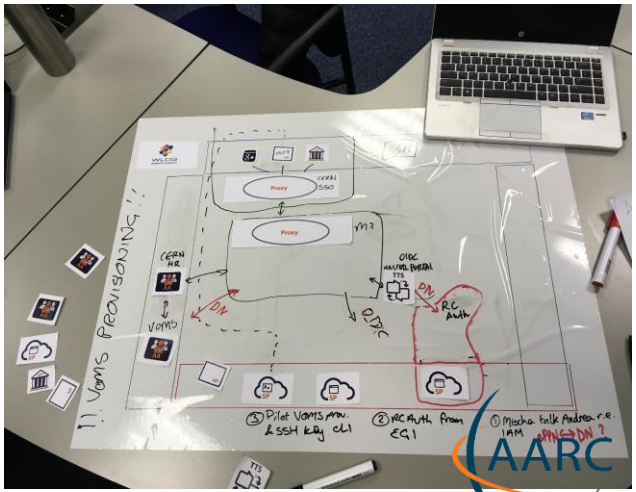

Solution Design

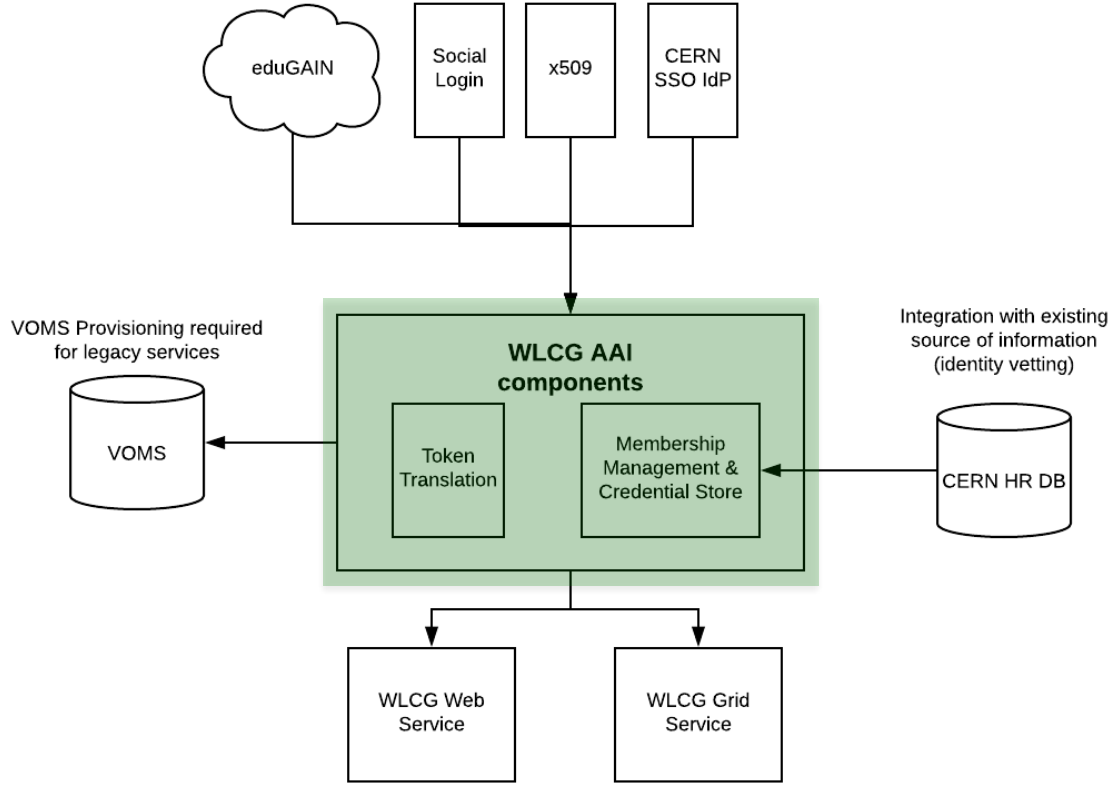

Solution Design

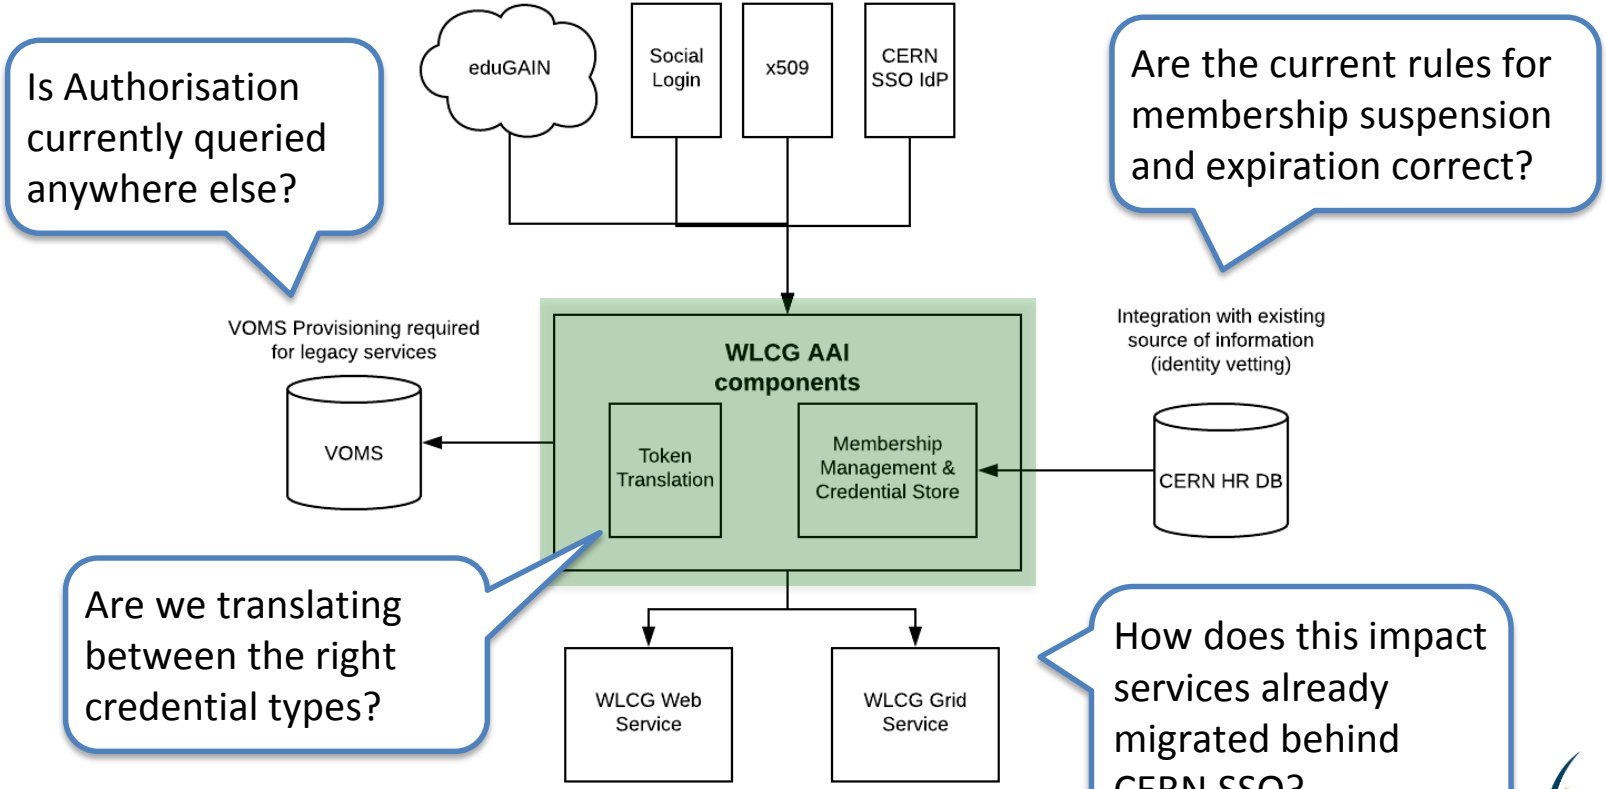

Solution Design

- Two solutions appear to meet the majority of requirements
  - EGI Check-in & COmanage
  - INDIGO IAM
- Additional integration required for
  - VOMS provisioning & lookup
  - CERN HR DB integration
  - AUP re-signing
- Input from VO managers will be required to assess these pilot implementations

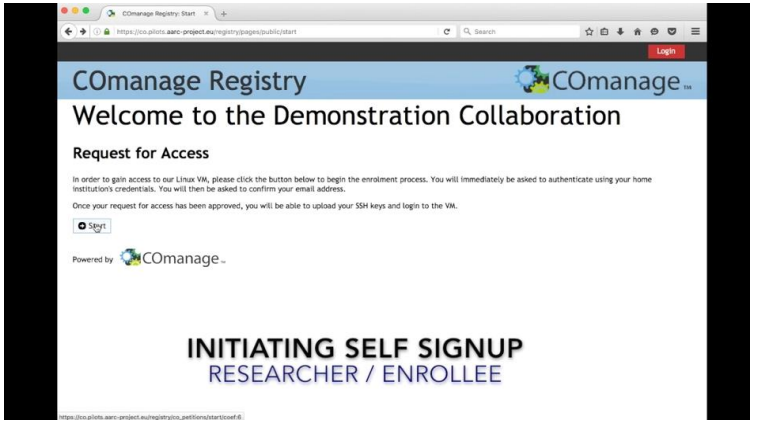

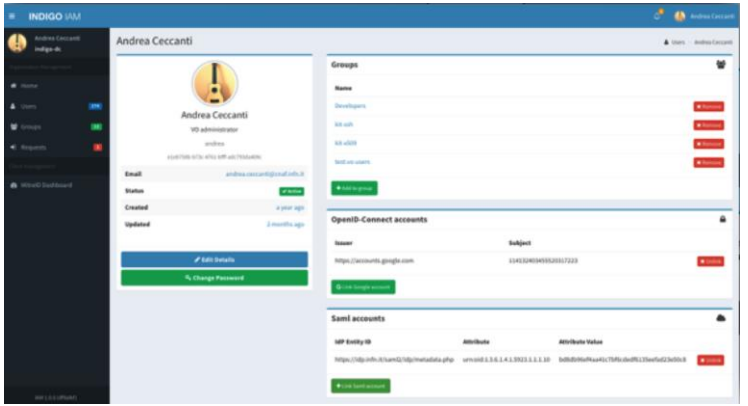

Token Schema

- Computing services are increasingly turning to token based authentication & authorization (OIDC, OAuth2…)
- Multiple infrastructure projects already using/supporting token based authorization but with diverging schemas or technologies
  - INDIGO IAM
  - EGI Check-in
  - SciTokens
  - dCache
  - ALICE tokens
- WLCG should define a token schema (or several!) to be issued by VOs and consumed by services

Token Schema

- Identity token
  - Proposal to stick closely to OIDC schema
  - Should groups/roles be included?
- Authorisation token
  - Do we need groups/roles at services?
  - Do we prefer capabilities?
  - Do we need both?
  - How should users be identified?
- Groups
  - Should groups be hierarchical?

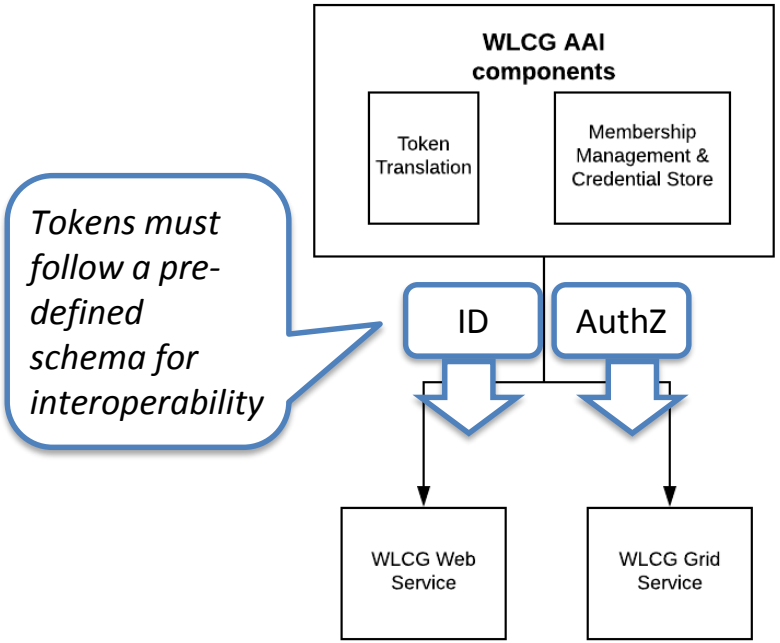

Next steps

- July Pre-GDB
- Aim to have a final drafts of the following to solicit feedback from stakeholders:
  - WLCG AuthZ Requirements
  - Token Schema
- Begin discussion on operational impact

How to participate

Participation in the WG is welcome!

E-group: project-lcg-authz@cern.ch

Twiki:
https://twiki.cern.ch/twiki/bin/view/LCG/WLCGAuthorizationWG

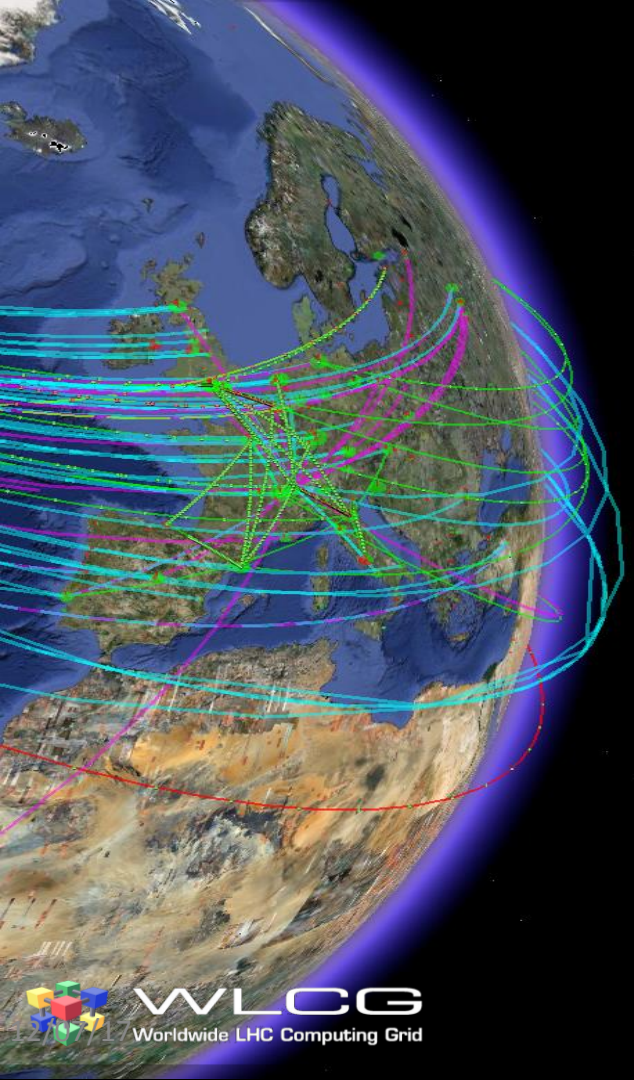

Questions?