Table of Contents

WLCG Issuer Deployment Architecture (DRAFT)..1

Proposal (Brian)...2

WLCG Issuer Deployment Architecture (DRAFT)

A few key decisions needed for the deployment of IAM instances for WLCG include:

* What is the content of the `iss` claim? * How many IAM instances should be run? Should there be a single multi-tenant instance for WLCG or multiple instances?

Proposal (Brian)

* The `iss` claims will be of the form:

- `https://cms.auth.cern.ch/`
- `https://atlas.auth.cern.ch/`
- `https://alice.auth.cern.ch/`
- `https://lhcb.auth.cern.ch/`

- These locations need not be the same as the token issuer (i.e., IAM) but are clear and memorable. It might be strategic to split the issuer string from the IAM instance hostname from the very beginning to help emphasize portability.

This approach requires that the metadata doc will be available at https://cms.auth.cern.ch/.well-known/openid-configuration

This can be easily implemented with the current IAM. Needs to be understood how things will work when IAM will be based on Keycloak.

* These will start as single-tenant instances of IAM. This decouples the VOs from having to share a single version -- allowing a "pathfinder" VO to proceed more quickly than the others. * We will start with CMS and stand up the IAM instance at `https://cms-iam.auth.cern.ch/`. * This instance will not be pre-populated with users - rather users will have to register via CERN SSO. The integration with the CERN HR database will allow us to determine whether the user is actually a CMS user.

This topic: LCG > WLCGIssuerDeploymentArchitecture
Topic revision: r3 - 2020-03-19 - BrianBockelman