Welcome to the ESCAPE AAI Webinar

April 02, 2020

https://indico.in2p3.fr/event/21072/

Please use this document to provide feedback/comments/suggestions!

Thanks!

Q1:

Is the SCIM standard relevant to integrating with external commercial services such as Active Directory ?

A1:
The SCIM APIs could be used for that, by implementing a middle layer that is responsible of syncing the contents of the AD with the IAM database, but there are simpler approaches to implement this integration.
Currently IAM does not natively support integration with AD/LDAP, but can be easily integrated with the deployment of a “glue” component that acts as “adapter”. We are already doing this for some use cases leveraging Keycloak for the integration.

In the future IAM will be based on Keycloak so it will have native AD/LDAP integration.

Q2:

What happens to a service that needs attributes from several VOs at the same time?
For instance, a data warehouse that hosts data from several VOs may provide a crossmatching feature.
In this case, the user must prove at the same time that he/she is a member of more than one group/VO at the same time.

A2:
Support for multi-tenancy depends on the software being integrated. Typically, libraries support integrating with multiple different VOs. At CNAF we integrate multiple VOs managed with IAM in storage, both using StoRM WebDAV or Apache with no problems.

Q3:

Can the email address of a registered account be changed/updated later?
For instance, if the user changes email provider or institution.

A3:
Yes.

Q4:

Can we ask for access through https://iam-escape.cloud.cnaf.infn.it/ now?

A4:
Yes.

Q5:

Where can we find instructions to deploy our own Indigo-based IAM service?

A5:
On the INDIGO IAM documentation page: https://indigo-iam.github.io/docs

Q6:

Can/How does IAM handle ssh keys? I.e. can it handle ssh key distribution to protected resources? +1

A6:

IAM supports storing SSH keys at the API level, but this functionality is not yet exposed by the IAM dashboard. Support in the dashboard will be provided soon. The idea is that then the keys linked to user profiles can be exposed to relying services via the SCIM provisioning APIs.

Q7:

Under what circumstances would you need to deploy your own instance, instead of using e.g. eduTEAMS?

A7:

IAM implements the concept of collaboration/Virtual Organization. It's a pretty self-contained alternative to EduTEAMS. Typically a collaboration will choose to use IAM or EduTEAMS depending on their requirements. In ESCAPE, IAM is the technology used for enabling access to the Datalake. Some functionality required to make the datalake work is not implemented in EduTEAMS. EduTEAMS, however, can be integrated with IAM as an external authentication provider for given collaboration, and we have plans to demonstrate this working integration soon.

Q8:

Can IAM be integrated with other IdP types such as ORCID, or an LDAP directory?

A8:

IAM can be easily integrated with SAML and OpenID Connect IdPs. An LDAP directory can be integrated currently using an intermediate service (Keycloak) that acts as an adapter. In the future IAM will integrate directly with LDAP directories.

Q9:

An update of the youtube IAM demo video (where necessary) would be greatly appreciated. Today's screen sharing was a bit problematic.

A9:
This is planned. The Webinar was meant to be recorded, but unfortunately there were technical issues.

Q10:

Can a service standing between the client and the storage/datalake request a token in an automated way (i.e. without password - as done with x509)? This is to avoid delegation.

A10:

Yes, there are ways to support this use case. But also delegation is greatly simplified with OAuth

Q11:

Is there already experience with integration of IAM and DIRAC?

A11:
No, this work still has to start. That said, there’s nothing really special in the DIRAC integration, and experience from the integration in storage elements and HTCondor will help.

Q12:
In slide 12, there is a key 'wlcg.groups' in the token claims. I guess this is a convention to partition the naming space for the groups and this convention is known and honored by all storage systems in the datalake.

If one wants to use IAM for a non-WLCG project, would that imply using the same naming convention (i.e. 'wlcg.myvo') or is it possible to add other prefix to the name space (e.g. 'astro.myvo'). In that case, I guess it would be needed to configure all the storage systems in the datalake to also support this new namespace convention.

Is this how it is supposed to work?

A12:

IAM can support multiple token profiles, i.e. ways to express information in tokens that are used for authentication/authorization in the infrastructure. In the slide the example shows a wlcg token. A different profile could be defined for other communities that is not “branded” as WLCG, and will be supported by IAM.