Fermilab Federated Identity Project

Mine Altunay

ICAC

February 9, 2022

Outline

- Quick recap from our last presentation.
  - What were our major problems back then and how we resolved them.
- New token-enabled scientific infrastructure.
- Experiments and service providers status.
- Transition planning.

Recap From Our Last Presentation

- At the time, we had the WLCG schema finalized. We were worried about reconciling the differences between SciToken and WLCG tokens and getting HTCondor to work with WLCG tokens as well as SciTokens.
  - This issue has been resolved.
  - We decided to adopt WLCG tokens to maintain our interoperability.
  - We requested HTCondor team to modify HTCondor so it now works with WLCG tokens as well as SciTokens.
  - WLCG Authorization Working Group is also working on reconciling the two schemas.
- CERN decided to use Indigo IAM. We were considering using CILogon as our token provider.
  - We decided to work with CILogon Token Issuer. CERN and our lab have different federation management systems, but we made sure they are both compatible.

Recap From Our Last Presentation

- At the time, we were very concerned about the necessary policy and procedures to extend federated access to users without Fermilab accounts. We discussed several contingency plans should we fall seriously behind OSG or WLCG schedules. We discussed identifying resources that we can federate right away in our contingency plans.
  - We did not fall back in our schedule. In fact, we are quite ahead of the WLCG token transition schedule.
  - We decided to focus on designing and implementing our scientific infrastructure first. We call this Phase 1.
  - We are currently in Phase 1, which is based on OAuth and JSON Web Tokens and it provides all the tools and software necessary for federated access to scientific resources.
  - Phase 2 will develop policy and procedures to extend federated access to users at our partner institutions.
  - We had a very large amount of challenging technical work to get our infrastructure token enabled. We did not want our technical work to be delayed while waiting for the policies to be developed.
  - We separated the policy work from the technical team, so Phase 1 and 2 could work in parallel.

New Token-enabled Scientific Infrastructure

- CILogon Token Issuer is our token provider. It comes with an LDAP server, which is used to store the user attributes and a list of privileges.
- To generate a token, CILogon Token Issuer retrieves a list of user privileges and attributes from the LDAP server and then puts them into the OAuth token.
- Ferry has been our central user management service for the scientific infrastructure for a long time. It has a list of VOs, user privileges, groups and so on. Ferry is programmed to periodically update the CILogon LDAP server.
- CILogon Token Issuer has been in production and can successfully generate tokens with correct privileges.

New Token-enabled Scientific Infrastructure

- We chose Vault as our new key repository and token renewal mechanism.
- Vault interacts directly with the CILogon Token Issuer and generates access tokens for the end user.
- We also created a new command line Vault client (called htgettoken) to fit our user needs. Htgettoken interacts with the end user and retrieves tokens from the Vault server.
- Vault also handles token renewals due to the short token lifetime.
- Our service providers integrate their clients with Vault client to handle tokens.

New Token-enabled Architecture

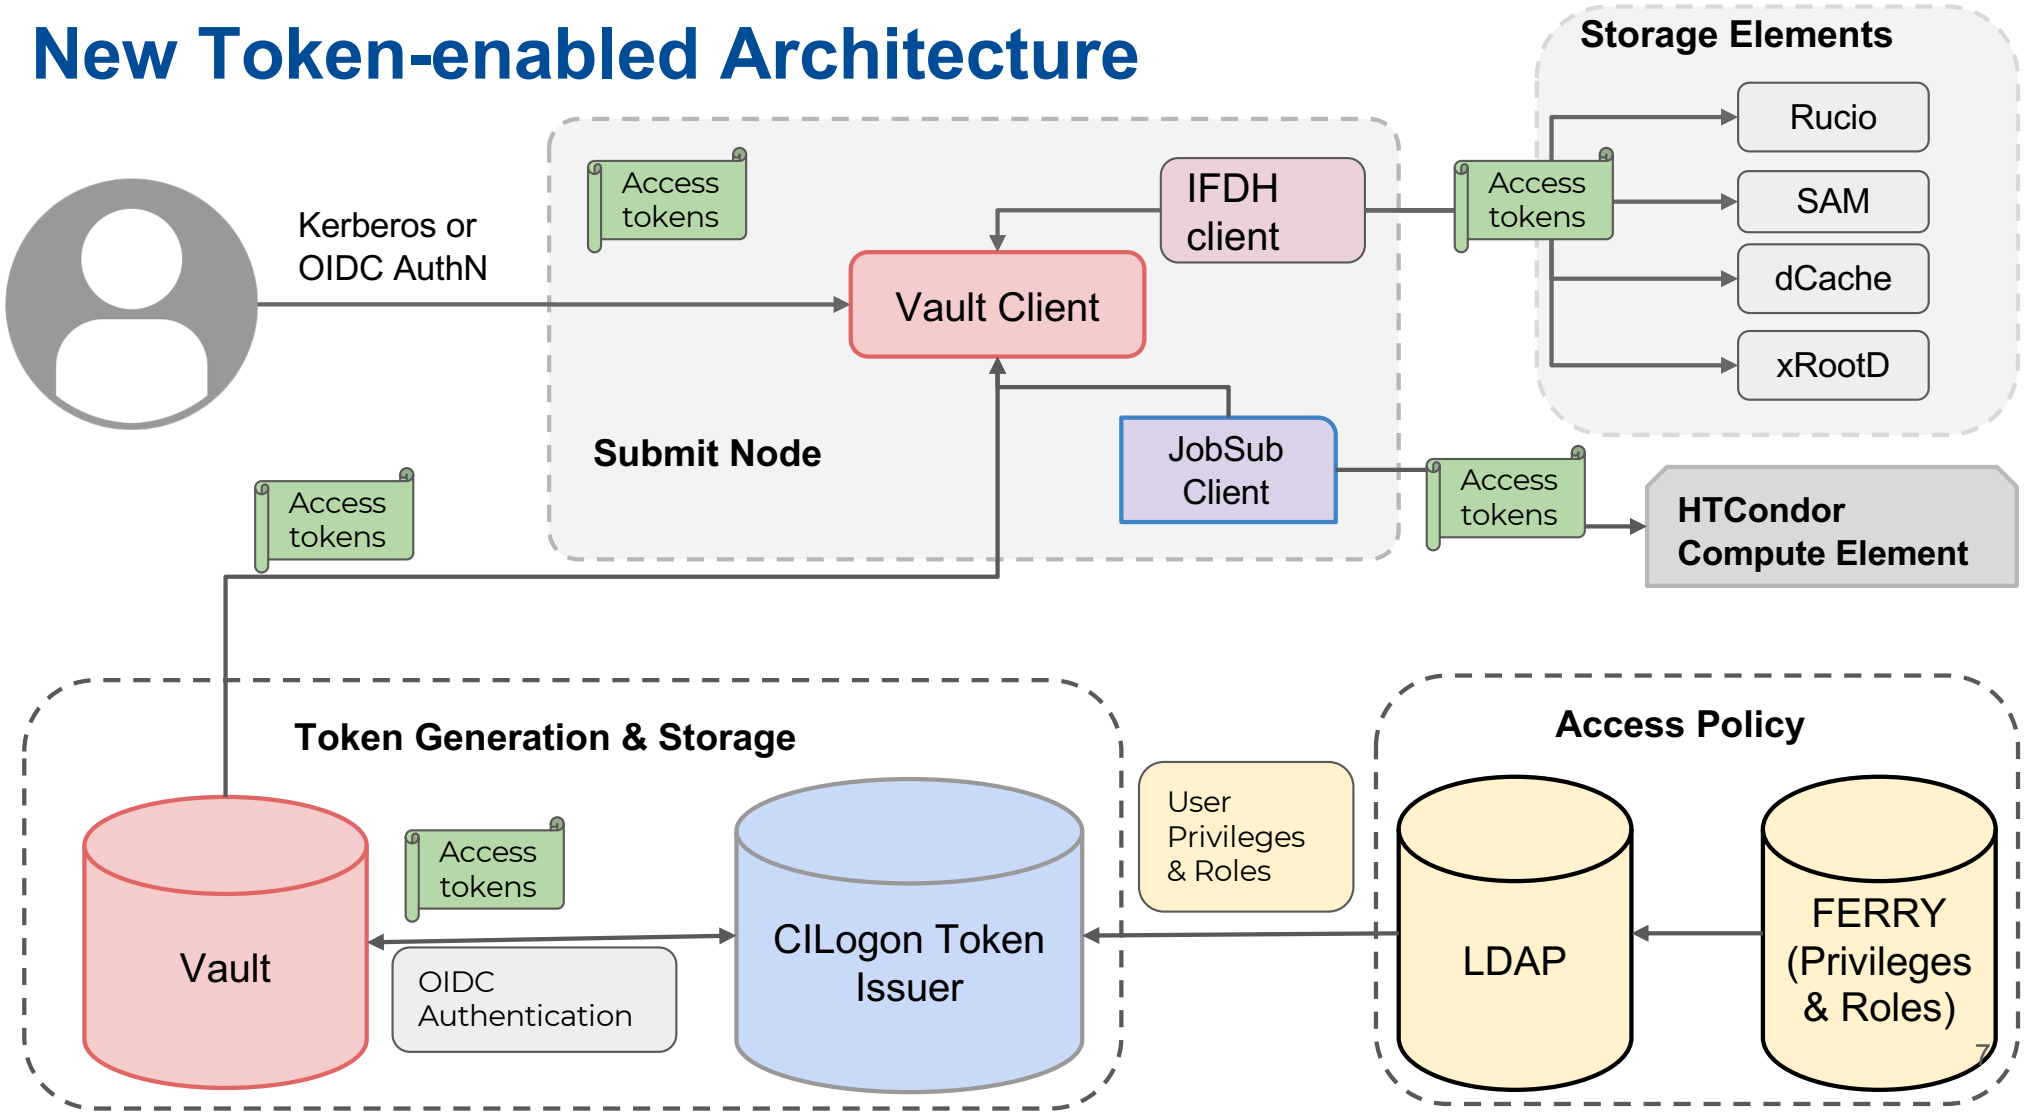

Token Renewals, Robot Tokens & Long Running Jobs

- Token lifetimes are only 3 hours, so no jobs can run without renewals.
- The renewed tokens are pushed to the long running jobs by HTCondor.
  - Our team helped to create two software components: Condor_submit and Condor_credmon_vault to renew and push tokens to the long running jobs.
- Automated processes need robot tokens, similar to robot certificates, which are continuously being renewed.
- Each time a new token is requested, Vault renews Access token (3-hour long) and Refresh token (1-month long). Long-lived Refresh tokens are stored in Vault repo and can be renewed indefinitely as long as a user is active at least once a month.
  - If a user is not active for a whole month, Refresh token will expire and user will have to authenticate as if it is the first time.
- Htgettoken supports robot tokens, which can be renewed automatically.
  - Uses kerberos robot credentials to authenticate to Vault.
  - Keeps generating robot access tokens periodically without allowing any expiration.

New Token-enabled Scientific Infrastructure

- All of the software development for the infrastructure is completed and has been in the testing process.
- We are moving our components into production.
  - CILogon Token Issuer and LDAP are already moved to production servers.
  - Ferry is in the process of moving.
  - However, we are not in production yet, since we have not completed our tests yet.

Service Providers Status

- Our service providers have been part of the architecture team meetings, so they kept up-to-date on our progress.
- HTCondor, Dcache, GlideinWMS, JobSub have been very active in this work. They completed software implementation and testing and are in the process of fixing discovered issues.
- We still have more services that need to be integrated with tokens. Now that our infrastructure work is completed, we can provide more resources to these services to help them become token enabled.

Experiments Status

- We started a task force dedicated to the experiments meeting regularly.
- 5 of our international VOs will get their own Token Issuers: DUNE, DES, SNBD, g-2, mu2e.
  - They already got their independent CILogon Token Issuer instances set up and working.
  - These 5 VOs can obtain tokens. They have completed basic job submission and Dcache testing successfully.
- The rest of our VOs will be served by the same Fermilab CILogon Token Issuer. They will not have a separate dedicated Token Issuer for themselves.

Transition Planning

- We are currently in the process of creating a transition plan.
- We plan to get all of our service providers token-enabled by the end of Spring 2022.
- In parallel to our work with service providers, we want to have experiments test token-enabled services. Services used by all experiments such as Dcache, JobSub and GlideinWMS have already been in the testing process.
- We plan to start transitioning our first VOs to tokens in the late summer. We will operate our infrastructure in a hybrid mode, where some VOs will use tokens and some will use certificates.

Transition Planning Challenges

- We are ahead of the WLCG schedule. WLCG does not plan to transition any VOs this summer yet.
- We are considering transitioning some of our VOs to tokens in late summer. DUNE is a likely candidate to go first.
- We need to understand how to transition an international VO while some parts of the VO services will remain with certificates.
  - We must make sure we do not break global VO infrastructure.
  - We remain in close contact with the experiment management and the WLCG.
- Hybrid Mode has unique challenges of not breaking any certificate-based system until we complete the transition.