Fermilab
U.S. DEPARTMENT OF ENERGY | Office of Science

Fermilab Federated Identity Project

Mine Altunay

OSG Token Transition Workshop

October 14, 2021

Federated Identity Project at Fermilab

- The goal of our project is to provide our scientific collaborators with federated access to our lab’s scientific infrastructure and resources.
    - Our collaborators will access the lab’s resources by using their identity credentials issued by our partner institutions.
- There are two phases of the project.
    - We are in Phase 1, where we design and build the architecture based on OAuth and JSON Web Tokens, and will only allow access to users with Fermilab accounts.
    - In Phase 2, we will develop policy and procedures to design how to extend the federated access to users at our partner institutions.

New Architecture

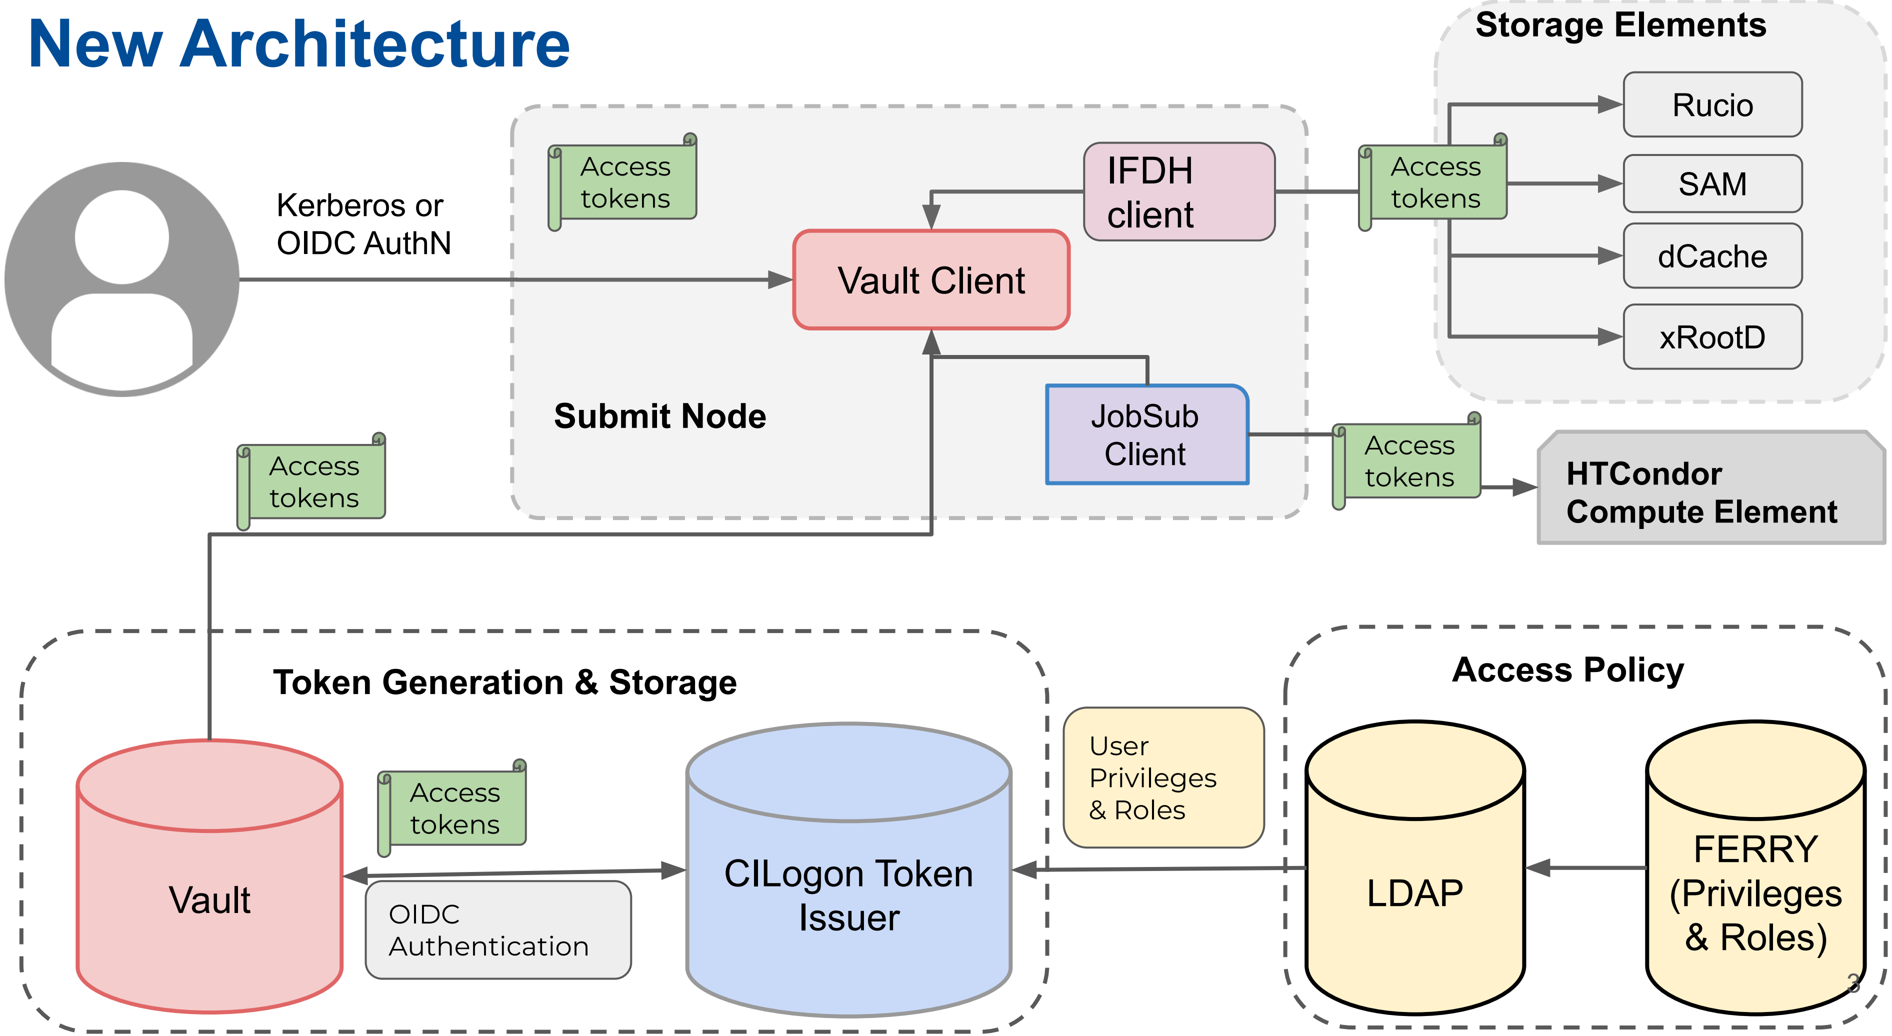

New Architecture

- Vault, CILogon and Ferry+LDAP are the main components of our architecture.
- CILogon is our Token Issuer.
- Vault is our token repository and also a client to CILogon Token Issuer to generate and receive tokens.
- Vault Client is at the submit node and retrieves generated tokens from the Vault repository. If there is no stored token, Vault generates a new one and sends it to the Vault Client.
- Ferry is the central database where we store each users’ privileges and roles (access policy). These privileges are sent to LDAP so that CILogon Token Issuer can read them programmatically and generate tokens accordingly.

Vault

- Needed a safe key repository to store the refresh keys.
    - Refresh keys can be renewed indefinitely.
    - A server based solution, similar to MyProxy, was preferable.
- Oauth2/OIDC is designed for web use, but almost all of our work is on the command line.
- Hashicorp Vault is a popular open source generic secret store server, that already supports OIDC and Kerberos
    - Needed some slight additions, submitted as PRs
    - Very flexible plugin architecture, REST/JSON API, and flexible access policies
    - Issues its own tokens for very flexible access to particular paths in its filesystem-like space
    - Needed configurator (htvault-config), and a new command line client (htgettoken) to control the flows

Initial Authentication

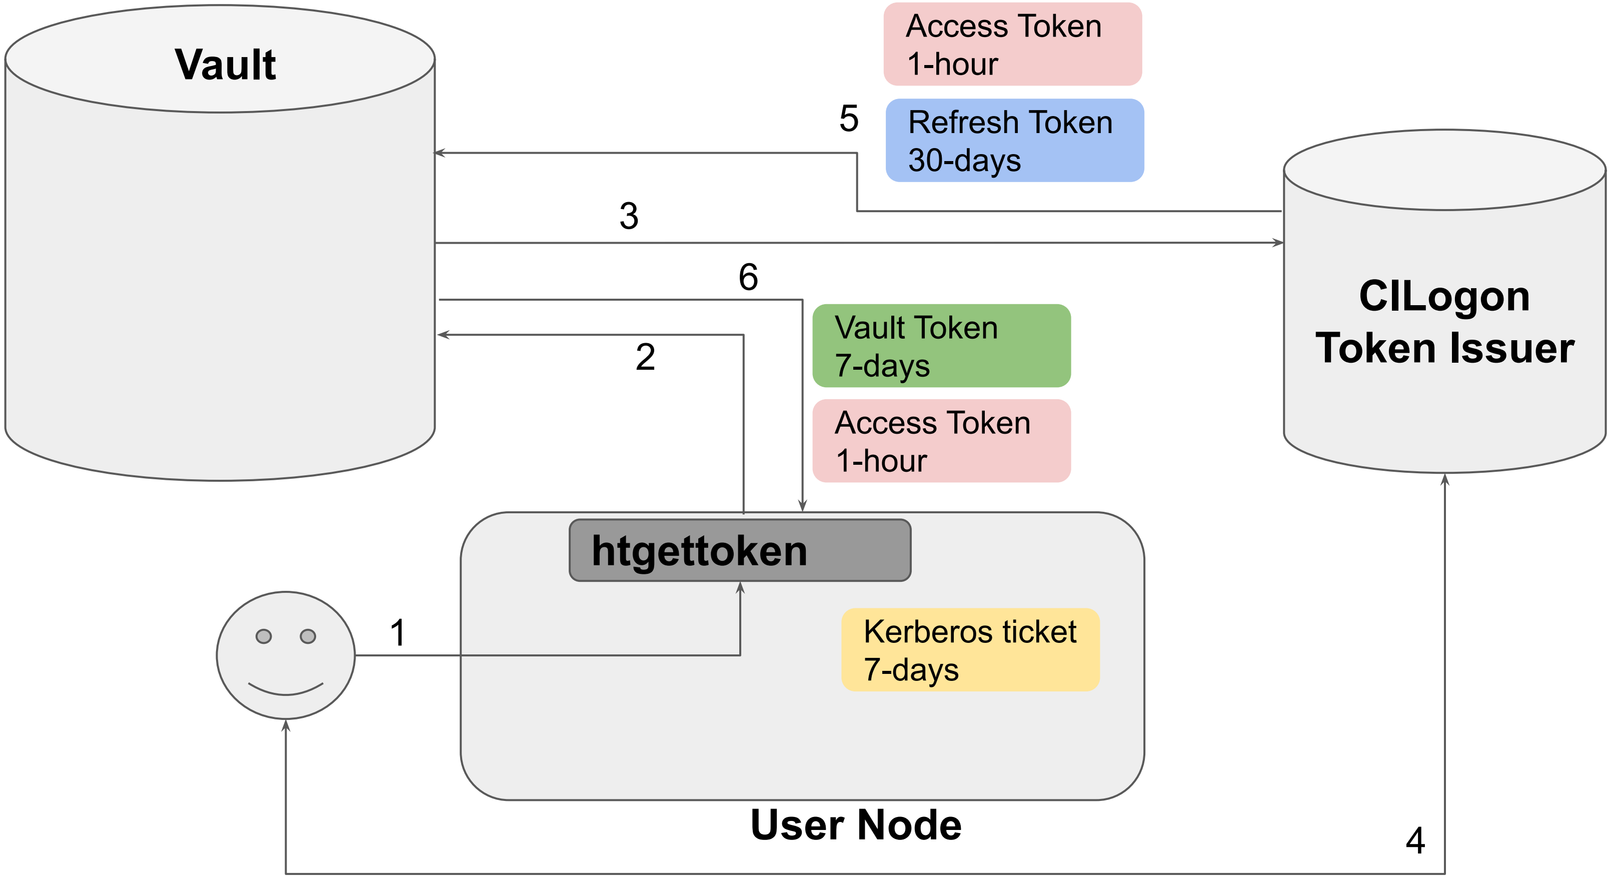

Token Renewals and Htgettoken

- Once a token is obtained, it can be renewed by the Vault token or by the Kerberos ticket.
  - Vault tokens expire after a week, after which Kerberos based renewal is necessary
- Each time a new token is requested, Access and Refresh tokens get renewed. Refresh tokens can be renewed indefinitely.
- If a user is not active for a month, Refresh token will expire and user will have to go through authentication flow as shown in previous slide.
- 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.

Token Renewal for Long Running Jobs, HTCondor Integration

- HTCondor uses long termed Vault tokens to create new access tokens and forward these to jobs.
- Condor_submit and Condor_credmon_vault are two components to achieve token renewal.
  - Condor_submit is configured to automatically invoke htgettoken as needed and store a long termed (4 weeks) Vault token in Credd.
  - Condor_credmon_vault uses the Vault token to get new short-lived access tokens and pushes them to the jobs.
- Will be in HTCondor’s release 9.0.6+ and 9.1.5+ and all OSG builds of 9.0+

Transition To Tokens

- All of our software development is completed except for patches.
- We are moving our components into production.
    - CILogon/LDAP is 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.
- We created a Token Task Force Group composed of experiment members, meeting biweekly
    - Invited experiments to obtain a token and conduct basic tests.

Experiment Feedback -- VO Mapping

- We have a number of VOs at Fermilab, and we cannot allocate an individual Token Issuer for each of them.
  - 5 of our international VOs will get their own Token Issuers.
  - The rest of our VOs will be served by the same Fermilab CILogon Token Issuer.
  - This causes an issue since the user mapping is based on the Token Issuer that indicates the experiment.
- We are currently working on the user mapping problem caused by having our experiments served by a single Token Issuer.
  - Dcache is developing a patch.
  - HTCondor will map based on Token Issuer and the Token subject field. For monitoring purposes, we will use ClassAds.
  - Holding a series of meetings with our other service providers
- Our 5 VOs with individual Token Issuers can obtain tokens and use the tokens to submit basic jobs and access data.
- We will do more in depth testing for the remaining VOs once the mapping problem is solved.

Token Lifetimes and Renewal Frequency

- Our access tokens are 1 hour long. Our experiments are concerned about the frequency of token renewal and the impact on the system load.
- Also concerned about network outages or other failures in the renewal system
  - Tokens with longer lifetimes have a better chance to finish their jobs.
- Our upper limit for access token lifetimes is 6 hours.
- We are currently discussing to increase access token lifetime to 6 hours.
  - We think experiments have valid points and decreasing the frequency of renewal would be useful.
  - Technically this is a simple change.

Confusion about Token Types and Renewal Process

- There is confusion about different types of tokens and credentials we use: Refresh, Access, Vault Tokens and Kerberos tickets.
    - How we use each token and when we use them.
- Some questions about why we had to switch from certificates.
- We provide 2 ways to renew Access tokens either via Vault tokens or with Kerberos.
    - Lots of questions about what happens when each token expires, which method to choose
    - Questions on what to do if the Refresh token expires.

Token Subject

- Currently, we use FNAL email addresses in the Token Subject field.
- CERN uses a random unique number for each user.
- We heard some requests that switching to a unique permanent number for each user may be easier than emails.
    - Arbitrary email addresses may be challenging for mapping.
- We do not have a common standard on the subject field value.
    - We will explore with WLCG whether all sites should use a common format

Multi Experiment Membership

- Some of our users have multiple experiment memberships. So, they will have multiple tokens for each VO.
- This caused some confusion about how different tokens can coexist.
- Htgettoken allows specifying the VO and the group information when requesting a token. Each token is stored separately in Vault and on the client side.
- Since this is different than our current model, it will take some time for our users to get used to it.

Experiment Transition

- Our earlier plan was to finish all experiments testing by January 2022.
- However, we found issues with user mapping for FNAL SubVOs. We are currently working on this issue.
    - Our testing schedule will be pushed back a few months.
- Our goal is to finish experiment testing by the end of Spring 2022.
- We will start creating a transition plan once all patches/fixes are in place and our architecture is finalized.
- We plan to start transitioning first VOs to tokens in the summer. We will operate our infrastructure in a hybrid mode, where some VOs will use tokens and some will use certificates.