2nd Meeting of ICAC
Federated Identity Status & Plans

Mine Altunay
October 15, 2019

Current Status

- Adoption of Federated Identity and Access Control is one of the top priorities for the lab
- Crucial for our success with DUNE, necessary for our lab to become the world's premier open physics laboratory
  - Allowing scientists transparent and seamless access to resources is prerequisite for our success
- We started a Federation Project that aims to bring federated access control to our laboratory infrastructure
  - Replacing certificates with federation tokens
  - Honoring tokens generated by our partner organizations and allowing scientists access to our resources based on these tokens

Federation Project

- Collaborating with WLCG and DUNE on Federation Project
- Organized and held a series of workshops at Fermilab last month
  - DUNE computing model workshop: https://indico.fnal.gov/event/21231/
  - WLCG SLATE security working group (container security): https://indico.fnal.gov/event/21485/
  - WLCG pre-GDB: https://indico.cern.ch/event/739896/
  - WLCG Grid Deployment Board https://indico.cern.ch/event/739882/
  - Federated Identity Management for Research https://indico.cern.ch/event/834658/
  - IRIS-HEP https://indico.cern.ch/event/840472/
- All of these workshops were very useful to build even stronger ties with our collaborators and make progress on technical challenges.

Status of US and European Efforts at Federation

- WLCG has voted to use Indigo IAM technology
  - A framework that can provide building blocks of federated access
  - Contributed by INFN
- SciToken project from the US also provides a federation framework
  - A token schema
  - Libraries for validating tokens and converting the token attributes into access rights
- CILogon project based at NCSA, Univ of Illinois, collaborates tightly with SciToken and provides the operational aspects of the federated access
  - Provides Token Issuers, Connectors, User Management services (COManage and LDAP), SAML Proxy and so on
  - It provides many services compliant with SciToken schema and libraries
- We collaborate with all our partners and try to build an interoperable solution

WLCG Authorization Working Group

- At our latest workshop, WLCG Authorization WG, which we are a part of, finalized the WLCG Access Token Schema
  - Big success by the WLCG Authorization WG
  - Schema ensures that all WLCG members can adhere to the same attributes and standards
- Fermilab recognizes that interoperability between all WLCG institutions are crucial for our success
  - We added some technical challenges to our todo list to ensure this
  - We will test sending jobs from Europe to US sites and US from to European sites
- SciToken schema adheres to the WLCG Token Schema; however, it does not implement the full WLCG schema
  - An important worry is the treatment of groups
  - SciToken does not allow expressing group memberships, rather expresses fine-grained access rights directly in a token

WLCG and SciToken Schemas

- SciToken tokens list the fine-grained access rights given to a user
- WLCG Tokens lists the groups that a user is member of
- WLCG schema allows both groups and access rights, so both approaches are fine
- But, we are concerned that if Fermilab sends a SciToken token to CERN without groups whether that token will have appropriate access
- Also, SciToken libraries do not process tokens with group lists
  - If Fermilab only uses SciToken libraries, then we may not fully process tokens coming from CERN
- We understand how important interoperability is
  - We will make sure US sites have appropriate libraries to process WLCG tokens
  - We will make sure Fermilab generated tokens can access European sites

User Tracing and Federation

- An important DOE requirement is to know which jobs belongs to which users
- Federation must provide technical means to easily associate jobs with its owners
- We will ensure that the solution we come up with must provide enough information about the end user that we can fully trace a job to the person who submitted the job.

DOE Foreign Visits and Assignments rules

- DOE requires all foreign users to go through an identity vetting process as part of Foreign Visits and Assignments (FVA) rules for all DOE Labs
- The vetting process can vary from less than a week to up to a few months
  - Even for users from the same country can have significantly different processing times
  - We do not have a good estimate over how long it should take for a particular user
- These rules also apply to our distributed infrastructure
  - Meaning users submitting remote science jobs will also have to go through this process
  - This means vetting hundreds of science users submitting analysis jobs
- This is a roadblock in front of “true” federation because:
  - Federation essentially means trusting another organization to do the vetting of the user and accepting the access tokens from that organization
  - FVA requires DOE Labs (such as Fermilab) to perform the user vetting independent of the home institution vetting

FVA and Federation

- A legitimate question: is there any benefit to Federation even with FVA rules
- Some small benefits are:
  - Once a user gone through the FVA, they may stop using Fermilab assigned credentials and switched to using their credentials from home organizations
  - This is a small benefit, but at least shields user from having to learn Fermilab passwords, kerberos tickets, tokens, etc.
  - But, the FVA process will need to be renewed as the user badge expires annually, so the benefit is actually rather small.
- We are still in the process of fully understanding and applying FVA rules to our laboratory
- One of the open questions is to understand the application of these rules to the distributed computing resources
  - Interactive access should fall under FVA
  - But, we are researching if people who only access through CEs and SEs should be able to follow old rules

FVA and Federation

- Given that our distributed computing resources are very well isolated form the rest of the laboratory, and
- We had no serious security incident in the past 10 year or so, we can conclude that our scientific distributed infrastructure does not increase our lab’s risk posture significantly
- However, applying FVA rules will significantly slow down the science and our international collaboration

FVA and Federation

- Another complication is the OSG's switch to SciTokens pretty soon
- We may fall behind the schedule due to implementation of FVA rules. We are working hard on these, but there is quite a lot to understand and implement.
- OSG time table is as follows
  - October 2019 OSG no longer carries OSG-specific patches for the GCT.3 All patches are upstreamed or retired.
  - January 2020 "GSI free" site demo. Show, at proof-of-concept / prototype level, all components without use of GCT.
  - November 2020: GlideinWMS sends last CMS production jobs using GSI
  - November 2020: Complete transition of production US ATLAS/CMS CEs to a version of HTCondor-CE that supports SciTokens
  - June 2021: Completely drop GSI support from production GlideinWMS factories

OSG Timetable and Federation

- It is a risk that Fermilab can seriously fall behind OSG's timeline
  - This means Fermilab will lose support for many software components
  - It can bring serious complications to US-CMS Tier2 and Tier3 resources since Fermilab Tier1 will be far behind their schedule
- CERN has a more relaxed schedule, but they are also starting to generate tokens in a test bed in the next few months
- Any more delay due to FVA can seriously hinder collaboration between many partner institutions
- We should develop contingencies when/if that happens
- Most of OSG software comes from NCSA/Scitoken project
  - They are handling the token switch for OSG
  - We are trying to understand how CILogon services can benefit us in future

Federation Future Steps

- To speed up the adoption of Federation, we can identify resources (e.g. web based resources) that may not be impacted by FVA and start working on them
- Increase our collaboration with all our partners (WLCG, CERN, SciToken, CILogon) and ensure interoperability
- Outsource services to projects such as CILogon as much as possible
- Develop contingency planning for serious delays from Fermilab
    - Ways to continue collaboration when CMS Tier2 and Tier3s along with OSG and CERN switch to federation
    - Ways to continue running jobs in a hybrid environment
    - Maintain software without external support