Tokens in ATLAS

Alessandro De Salvo

30/11/2022

Evolution from X509 to Tokens

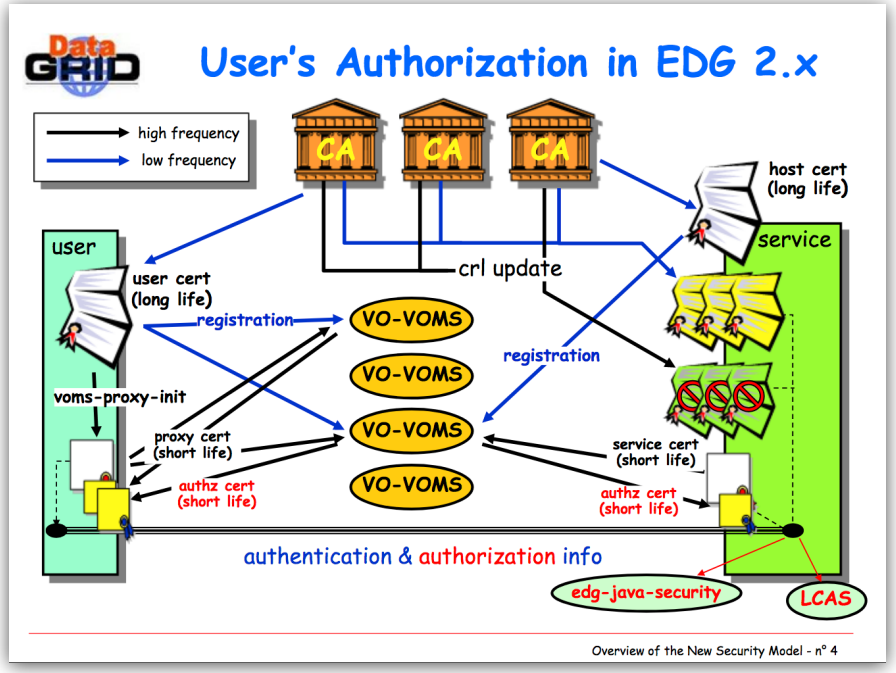

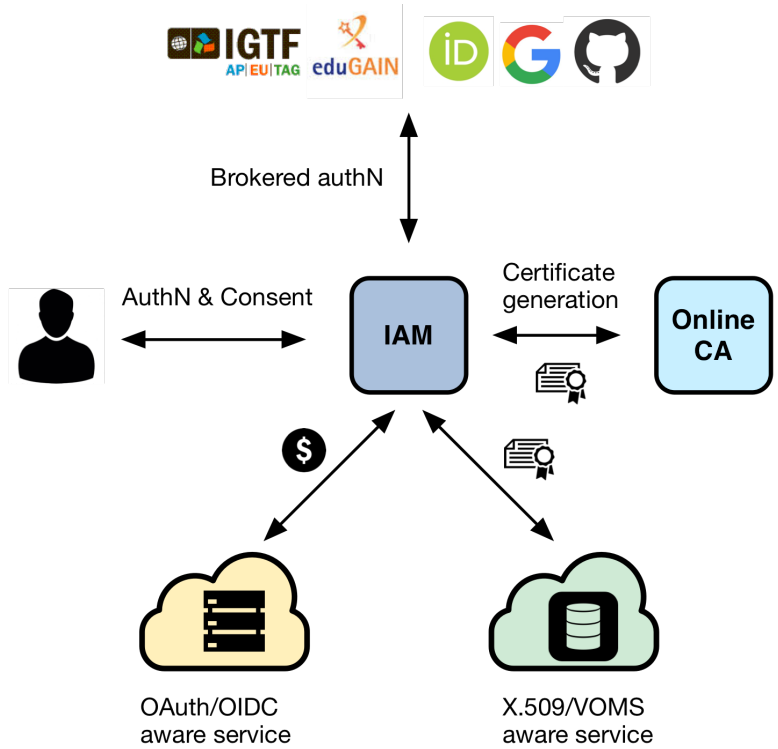

What is a Token

A. Ceccanti

To access computing and storage resources in the WLCG today you use a **VOMS proxy**, which provides information about **who you are**, **for which Virtual Organization (VO) you're acting** and **what you can do on the infrastructure** (i.e., VOMS groups and roles)

In the near future we will use **tokens,** which will provide more or less the same information

Tokens are obtained from a **VO token issuer** (e.g., IAM) using **OpenID Connect**

Tokens are **sent to services/resources following OAuth recommendations** (e.g., embedded in the header or an HTTP request)

Tokens are **self-contained**, i.e. their **integrity and validity can be verified locally** with no callback to the token issuer

WLCG Tokens

- JSON Web Token (JWT)
- Distributed over OAuth2.0 Protocol
- Contains identity and authorisation information from issuer (VO)
  - Groups and/or Capabilities
- Follows the WLCG Token Schema (https://zenodo.org/record/3460258)

A. Ceccanti

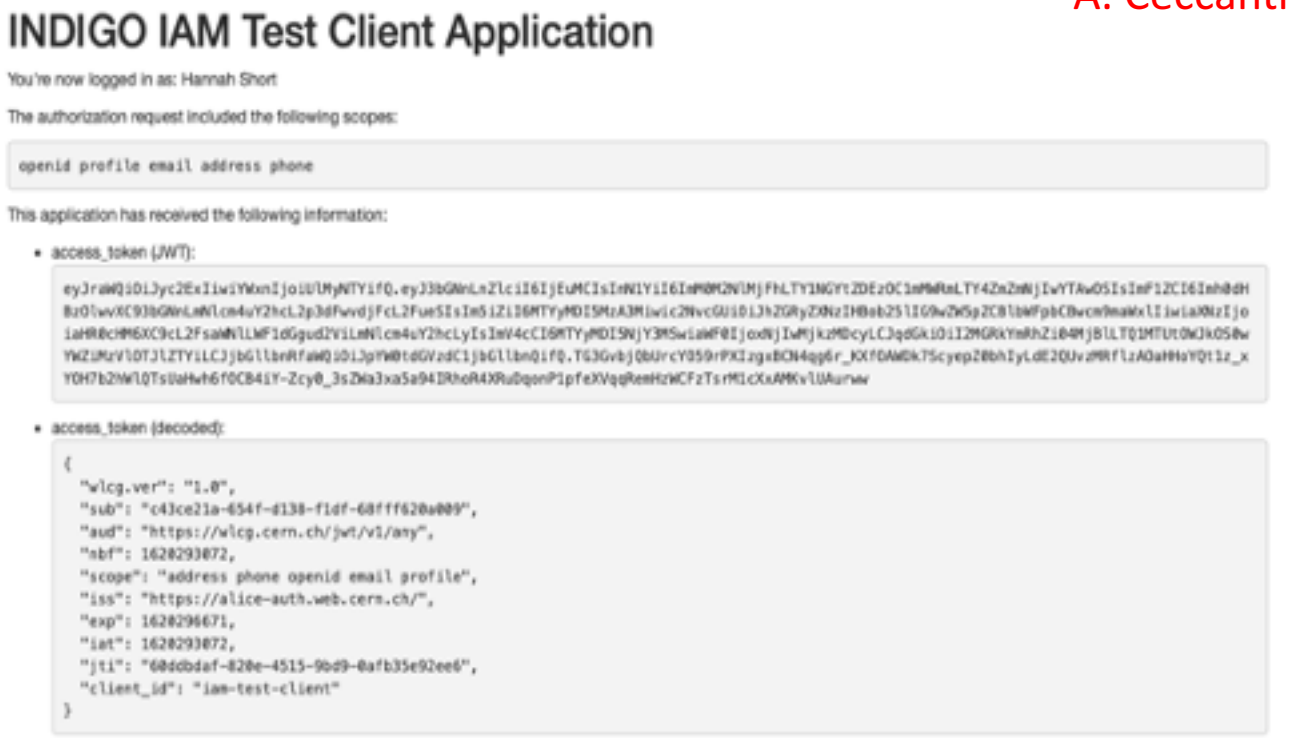

Example token from the IAM Test Client

The future of X509 and VOMS

A. Ceccanti

It will take **years** before we have migrated the whole infrastructure away from **user-managed** X.509 credentials

The transition will be **gradual**

During the transition we'll have **a mixed authN/Z model**: X.509/VOMS + tokens

Services at various level of the infrastructure **will have to understand both**

Token transition timeline

A. Ceccanti

Milestone ID	Date	Description	Dependencies	Teams
M.0	Feb 2021	Produce document with list of use cases for CMS VOMS-Admin API.	None	WLCG Ops
M.1	May 2021	WLCG baseline services include HTTP-TPC endpoints. Mind: tape services come later.	None	WLCG Ops, Storage providers
M.2	June 3/4 2021	WLCG hosts “CE and pilot factory hackathon”	None	Pilot framework providers
M.3	July 2021	Production IAM Instance(s) Available for at least 1 LHC experiment, likely CMS and possibly ATLAS	None	WLCG Ops, IAM, CERN IT
M.4	Oct 2021	Pilot job submissions may be performed with tokens.	M.3	Experiments, pilot framework providers, OSG/EGI, sites, Monitoring
M.5	Dec 2021	VOMS-Admin shutoff for CMS. IAM is sole authz provider for those (including for VOMS server)	M.3	WLCG Ops, CERN IT
M.6	Feb 2022	OSG ends support for the Grid Community Toolkit	M.1, M.4	OSG
M.7	Mar 2022	All storage services provide support for tokens	M.1	WLCG Ops, Storage providers
	?	All VOs shut off VOMS-Admin		
	Sept 2022	End of HTCondor support for GSI Auth (link)		
M.8	Oct 2022	Rucio transfers performed with token auth in production	M.7	Rucio, Experiments
M.9	Mar 2023	Experiments stageout & data reads performed via tokens.	M.7	Experiments
M.10	Mar 2024	X.509 client auth becomes optional.	M.9, M.8, M.4	Experiments

Indigo IAM

A. Ceccanti

An authentication and authorization service that

- supports **multiple authentication mechanisms**
- provides users with a **persistent, organization scoped** identifier
- exposes **identity information**, **attributes** and **capabilities** to services via **JWT** tokens and standard **OAuth & OpenID Connect** protocols
- can integrate existing **VOMS**-aware services
- supports **Web** and **non-Web access**, **delegation** and **token renewal**

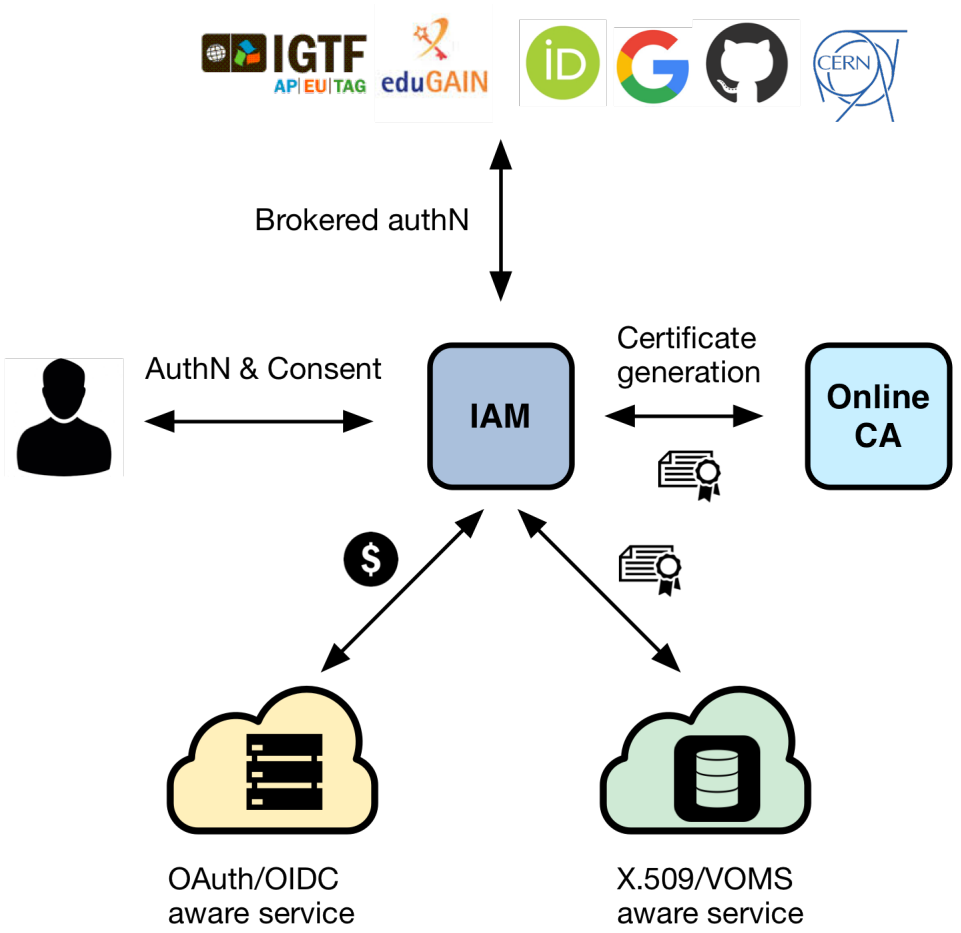

IAM enrollment workflow

A. Ceccanti

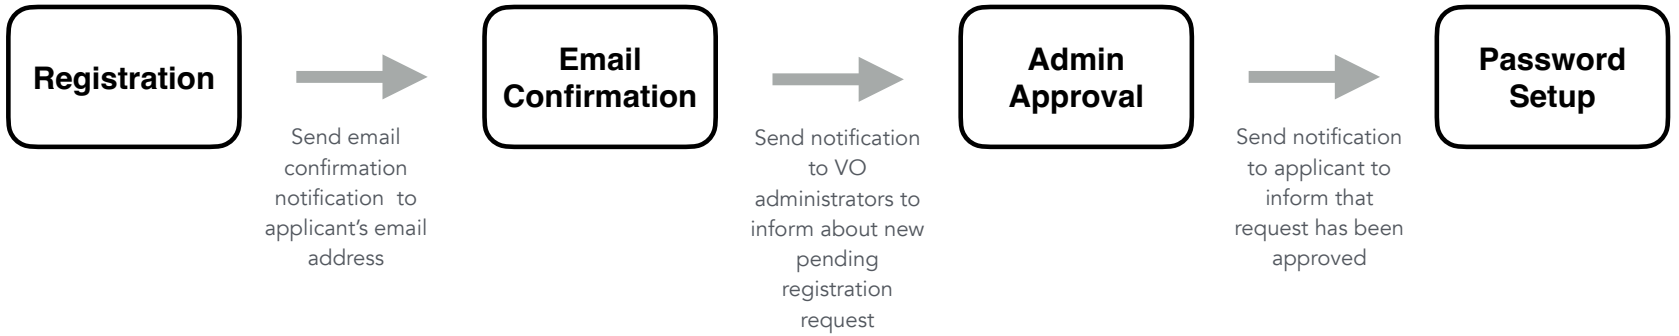

It's possible to set an **expiration time** for IAM accounts.

Once the account expires, login for the account is disabled.

IAM can be configured to remove expired accounts after a configurable grace period

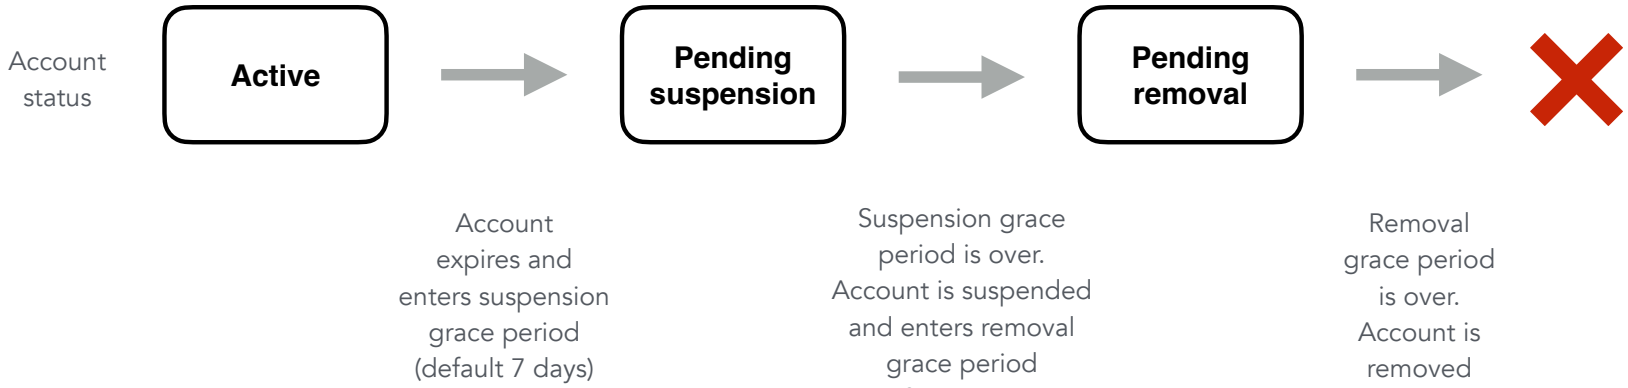

VOMS to IAM transition

A. Ceccanti

The transition from X.509 to tokens will take time so **IAM was designed to be backward-compatible with our existing infrastructure**

IAM provides a VOMS endpoint that **can issue VOMS credentials understood by existing clients and libraries**

- VOMS clients >= 2.0.16

The VOMS importer migration script has been developed to import users from VOMS to IAM

- users will **NOT** have to re-register in mass to IAM, and their IAM account will be automagically linked to their CERN account
- the script will keep IAM in sync with the VOMS instances until the VO registration process is migrated to IAM

At some point **IAM will be the only authoritative VOMS server for the infrastructure**

Legacy VOMS and ATLAS IAM VOMS

P. Vokac

- Starting in 2023 IAM service should have dedicated CERN person
  - Revive our original ATLAS IAM VOMS migration plan (slide 3) "plus one year"
    - Check remaining VOMS API usage and test new SCIM probes before the end of 2022
  - Follow other WLCG experiments IAM VOMS migration plans … may be already in Q1 2023(?)
    - Don't be the only experiment that fully rely on this new service
- Maintaining consistent state of two VOMS services not without problems
  - Each comes with their own "features" and difficult to understand behavior of service with two implementations on backend
  - Synchronization improved but delays may raise unnecessary questions
    - Some updates in "legacy" VOMS might not be correctly synchronized (destroy and re-create user account)
    - We need AUP signature time synchronization before IAM 1.8.x update
  - An increasing share of token usage
    - ~ 200Hz tokens for DC24
- List of technical issues that needs to be resolved before migration
  - Waiting for new IAM 1.8.x
    - Already tagged in git, but not yet deployed on CERN IAM instances
    - Most of the issues found during admin training should be fixed
  - VO Admins can test VO onboarding process with IAM
    - Require two different user certificates or guinea pig coleague
- Joining ATLAS VO documentation update once we migrate

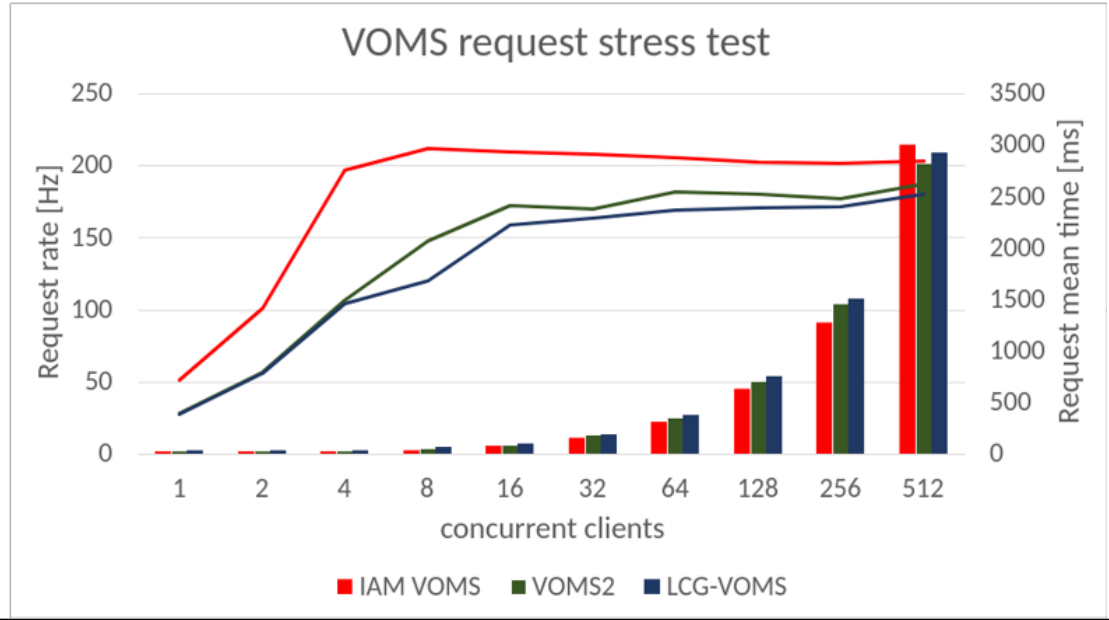

Rucio & Tokens in 2023

P. Vokac

- Extensively discussed by WLCG AuthZ group
  - Rough design described in Rucio token workflow evolution
  - Detail and ideas discussed by the WLGG AuthZ WG
  - Rucio become central point which enforce storage access
    - Users can't write to storage without contacting Rucio first
    - For reading we'll not be so strict
      - may even allow users to get tokens with storage.read:/ scope
- New developer started to improve token support in Rucio
  - Improve token support for data distribution with FTS (required by DC24)
  - Implement new schema for download / upload
  - Data removal
- Rely on token storage.*:/path scopes
  - Capability used to authz with storage
  - More secure, some security features can't be implemented with group authz

Tokens and job pilots/payloads in 2024

P. Vokac

- First we need storage & data distribution with tokens
  - Before pilot / payload can do something useful with tokens
  - Only storage.create:/path/to/file scope available to the job
- Tokens have short lifetime, avoid to distribute powerful refresh tokens
  - Easy to submit job with token, but this token can't be used by running job
  - Jobs may stay queued for a long time
  - Jobs runs for longer time that access token lifetime
- Initial discussion started during ARC/HTCondor-CE hackaton
  - Several topics discussed (tokens parsing, mapping, client side, security, testing)
  - Attempt to draw basic schema how deliver valid tokens to the job
    - Discussed just pull mode with Harvester
      - Establish communication channel with WFMS with secret passed with job
        - Use secret passed to the job to authenticate with WFMS
          - secret: "random string" registered in WFMS by Harvester
        - Uset to get right and valid access token
    - Can be extended for aCT + ARC-CE data management
      - Minor updates in ARC-CE REST
        - Passing tokens to ARC
        - Should be part of ARC-CE 7
      - Most of changes needs to be done on aCT side
        - Blind periodic token update on ARC-CE during whole job runtime(?)
  - Still need to decide which tokens will be passed to the job
    - Limit their power even more by employing features privided by MyToken / HashiCorp Vault
    - Analyze credential flow and evaluate consequences of stolen tokens

WLCG JWT token profile storage.create: Upload data. This includes renaming files if the destination file does not already exist. This capability includes the creation of directories and subdirectories at the specified path, and the creation of any non-existent directories required to create the path itself (note the server implementation MUST NOT automatically create directories for a client). This authorization DOES NOT permit overwriting or deletion of stored data. The driving use case for a separate storage.create scope is to enable stage-out of data from jobs on a worker node.

Only Rucio reaper daemon needs privileges to delete stored data (storage.modify scope in the token), no other component will ever have privileges to destroy any data

User interactions with tokens in 2025

P. Vokac

- Once we drop X.509 – credential management no more complicated
  - "we failed if current situation doesn't improve with transition to tokens"
    - Ideally credential management should be easier than `voms-proxy-init` once a day
  - Easy to achieve for web applications, tricky for CLI
    - Tokens comes with more granular privileges and shorter lifetime
    - User's have no idea which claims / scopes / audience should be in the token
    - CLI tools needs to be integrated with some local credential store
    - Several options based on HashiCorp Vault, MyToken, oidc-agent
      - FNAL integrated Vault with Kerberos
        - User in CLI can get tokens without providing any password
        - Still not integrated with CLI tools (user must call httokenget with right capability set)
- AF may start to use tokens (hidden from users) sooner
  - Is token support on data management side sufficient?
  - Does AF need also tokens for interaction with WFMS?

Tokens & ATLAS IAM

P. Vokac

- Just few for services with negligible load
  - HTCondor-CE job submission
  - JupiterHub tests, access to devel AMI version
    - Some details about registered client can be changed only by IAM admin
      - IAM admin – ATLAS VO Admins, me, CERN IAM service maintainers
- Most important use-cases – jobs and storage access – not yet covered
  - Rucio development, FTS improvements
  - Integration with Panda WFMS
- No common schema for WFMS
  - OSG/CMS – GlideIn WFMS (workshop)
  - LHCb – token comes with Dirac 8 (workshop)
  - We have to come with our own solution
    - Discuss ATLAS token credential flow with WLCG AuthZ
    - User interface is on WLCG AuthZ agenda
      - But no name assigned to this topic
- CE Hackathon in September 2022

IAM features

- Scope policies
- Token exchange policies
- JWT-based client auth

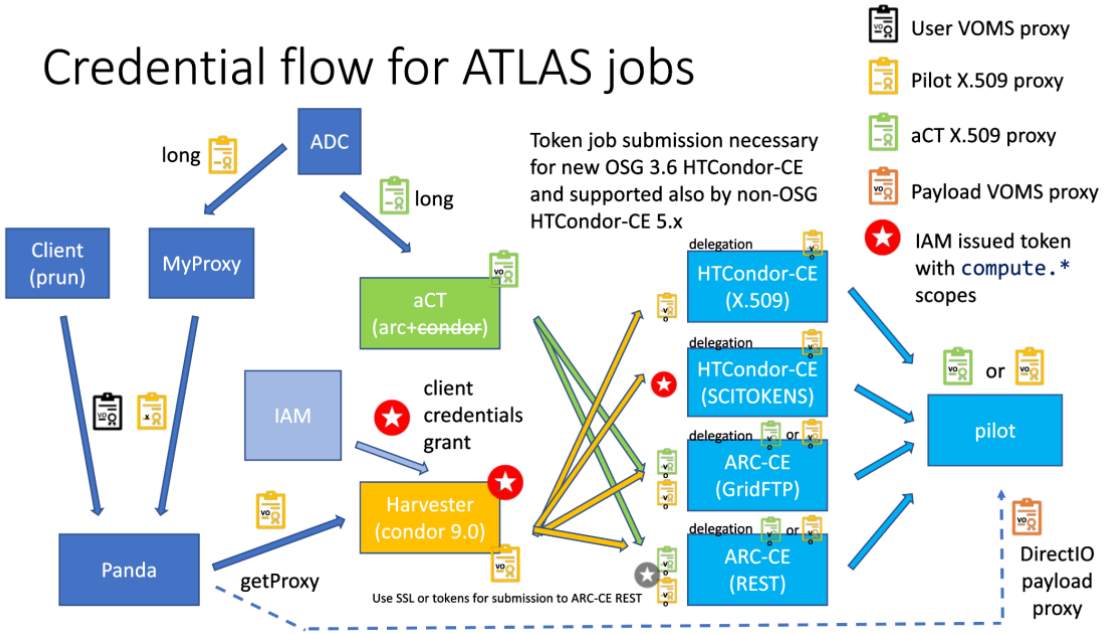

VO Managements and IAM

- VO managers are already instructed how to use IAM, but more training will be needed
- Person-power
  - Italy
    - Alessandro De Salvo
    - Elisabetta Vilucchi
    - Lorenzo Rinaldi
  - US
    - John De Stefano (soon to be replaced)
- No current time estimate for direct usage of IAM for registrations
  - For the moment the registration primary source is still the legacy VOMS-Admin, which is sinchronizwed every few hours to IAM
  - All services already instrumented to ask for VOMS proxies from IAM or directly tokens