

WLCG Token Transition Update

European HTCondor Workshop 2024

Brian Bockelman, 27 September 2024


Authorization, authentication, and credentials: A recap

Three important terms for this presentation:

- ▶ **Authorization:** Deciding whether an entity is permitted to perform an action.
- ▶ **Authentication:** Mapping an entity to an identifier.
 - ▶ Note: *Authentication* is often part of an authorization scheme.
- ▶ **Credential:** knowledge that establishes a fact (e.g., identity).
 - ▶ Not too far off from the 'credentials' the university provides: a diploma establishes the bearer has particular knowledge.
 - ▶ Classic example: a username/password is used as a credential to perform authentication.

Moving from identity mapping to capabilities

- ▶ Authorization on WLCG was always based on identity mapping*:
 - ▶ A request was authenticated to a global identity.
 - ▶ The global identity was mapped to a local identity.
 - ▶ The request was authorized if the local identity was authorized to perform the action.



<u>Scheme</u>	<u>Credentials</u>	<u>Authentication</u>	<u>Authorization</u>
Gmail login	Password, 2FA	Username	Access to your inbox
Building access	ID card	Identity in HR database	Elevators
International Travel	Passport	Identity according to US Government	Enter Switzerland
Baseball Game	Ticket	NONE!	Sit in section 4, seat 34B
Workshop	Zoom URL	NONE!	Attend this wonderful talk!

Identity Mapping for Distributed Computing

- ▶ Each compute center had the concept of user accounts – a login to the local site that's unique for each user.
 - ▶ Suppose we had 10,000 physicists and 60 sites. To establish a login for each person at each site, we would need to manage 600,000 usernames and passwords. **Ouch!**
- ▶ Instead, the WLCG utilized the idea of **identity mapping**.
 - ▶ Each user established one **global identity**.
 - ▶ A mechanism would be developed to map the global identity to a local identity.
 - ▶ The user is authorized to perform the operations of the local identity.
- ▶ What credential do we use for establishing an identity? The X.509 certificate!
 - ▶ Since you don't want to hand out your identity to others, the X.509 “proxy” serves as your “power of attorney”.

Group-based Mapping

- ▶ A conceptual leap occurred when sites realized they cared little about authorizing users but rather groups:
 - ▶ No need to know all 10,000 users and authorize them individually; rather, a site only cared that an identity was one of 5,000 members of a given experiment they supported
 - ▶ Example: every individual in the CMS experiment can store 1TB of data at my site.
 - ▶ Identities are only useful in determining whether two remote entities were the same individual.
 - ▶ As long as we know they belong to CMS, it is irrelevant if the username was “brian.bockelman” or “cms018”.
 - ▶ Sites still had a traceability requirement: need to determine all the actions “cms018” performed.

$A_N(\text{global ID}) \rightarrow \text{group}$

$A_Z(\text{group}) \rightarrow P_1, P_2, \dots, P_N$

Implementation: X.509 + VOMS

- ▶ Group mapping is implemented with the “Virtual Organization Management Services” (VOMS)
 - ▶ The VOMS server would sign a cryptographic extension to the user’s X.509 certificate or delegation.
 - ▶ Signing technology builds on X.509 attribute certificates.
 - ▶ This extension would assert group membership or roles.
 - ▶ A local mapfile would map the group / roles to a local identity. The rest proceeds as in identity mapping.
 - ▶ Note there’s no use of the user’s identity in this system!

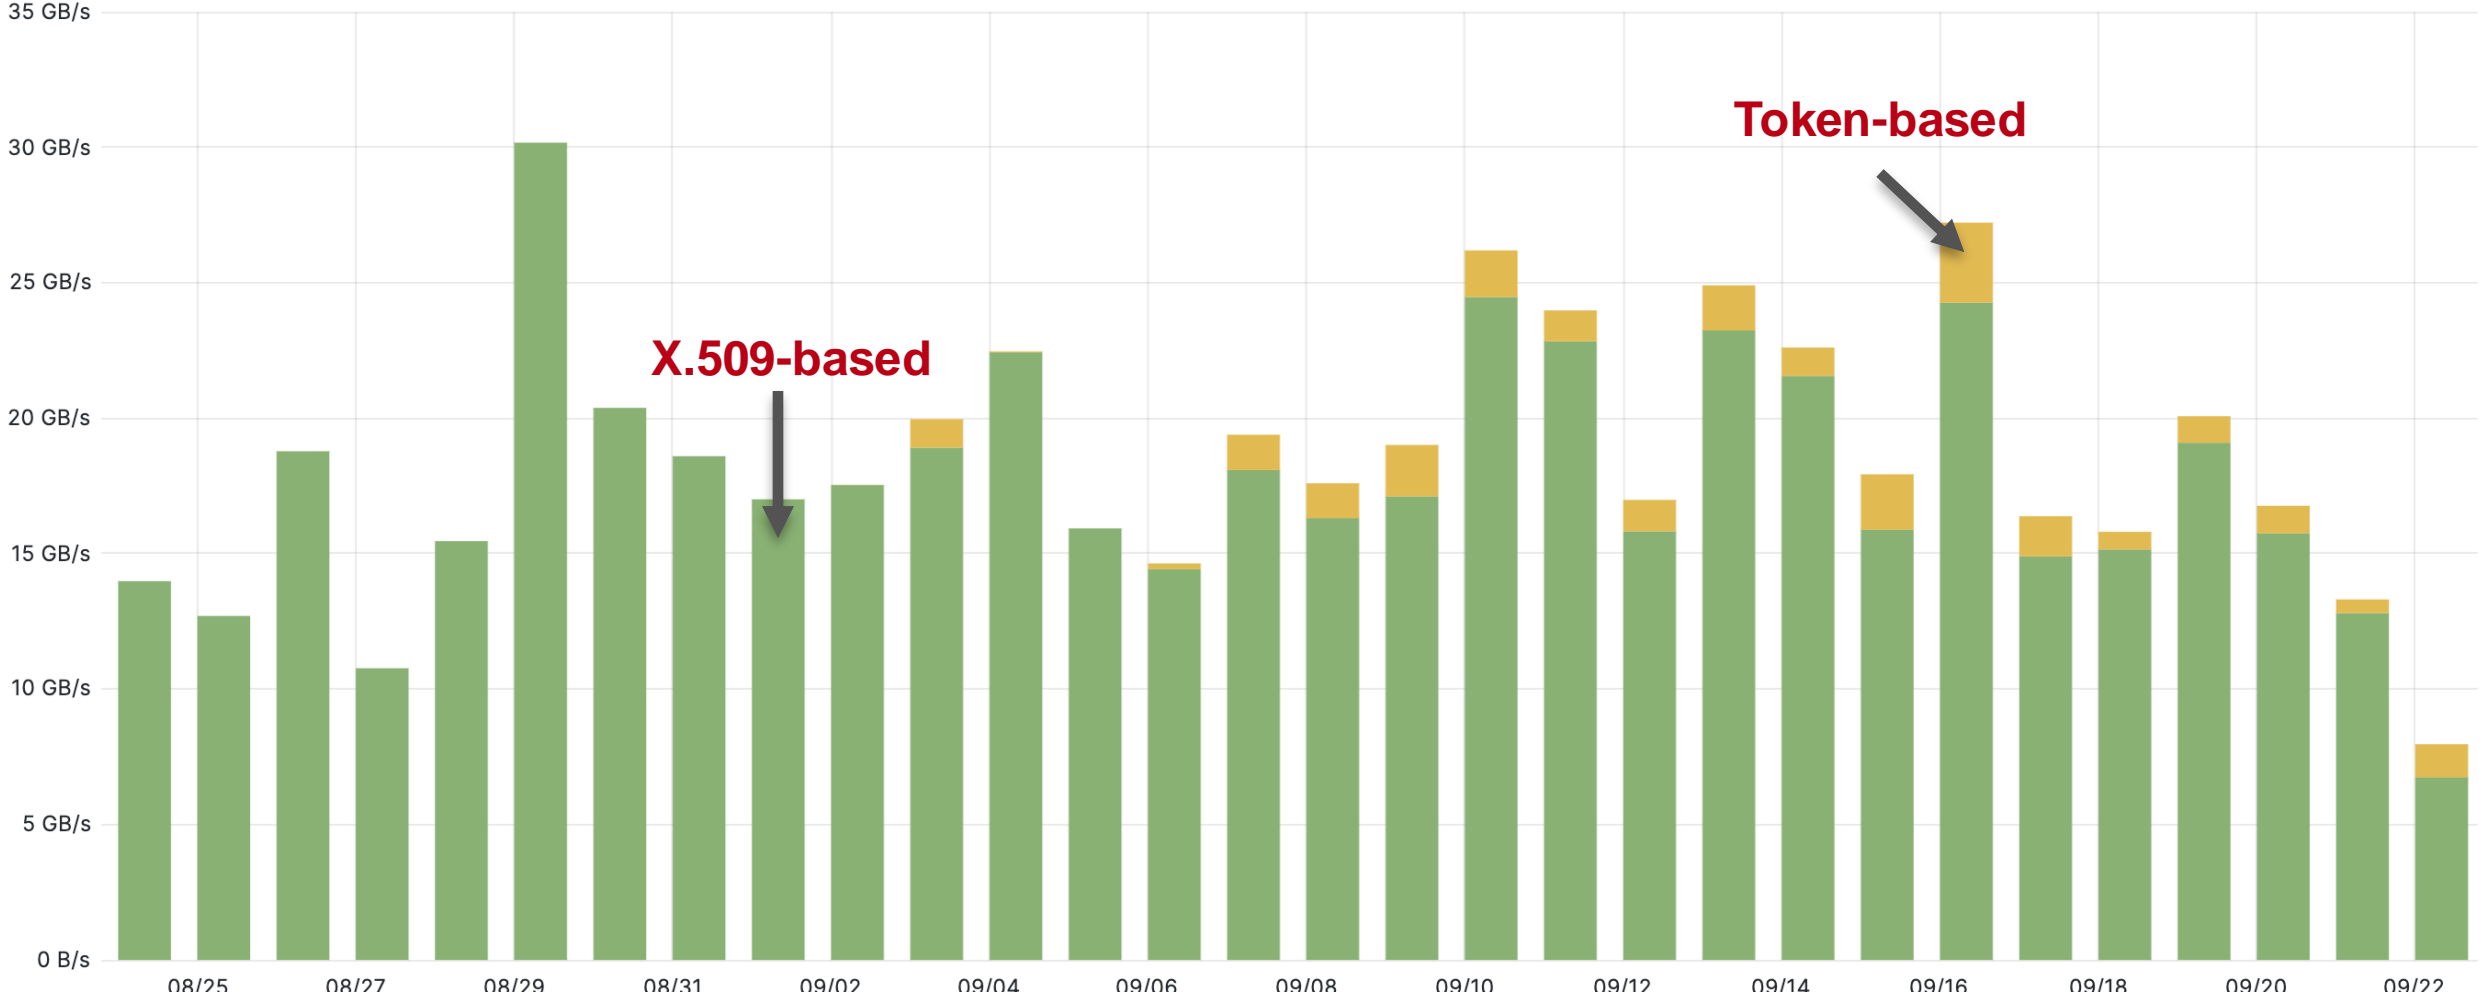
```
subject : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=bbockelm/CN=659869/CN=Brian Paul Bockelman/CN=1761602861
issuer  : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=bbockelm/CN=659869/CN=Brian Paul Bockelman
identity : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=bbockelm/CN=659869/CN=Brian Paul Bockelman
type    : RFC3820 compliant impersonation proxy
strength : 2048
path    : /tmp/x509up_u1221
timeleft : 11:59:54
key usage : Digital Signature, Key Encipherment
=== V0 cms extension information ===
V0      : cms
subject : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=bbockelm/CN=659869/CN=Brian Paul Bockelman
issuer  : /DC=ch/DC=cern/OU=computers/CN=voms2.cern.ch
attribute : /cms/Role=NULL/Capability=NULL
attribute : /cms/integration/Role=NULL/Capability=NULL
attribute : /cms/uscms/Role=NULL/Capability=NULL
timeleft : 11:59:54
uri     : voms2.cern.ch:15002
```

Human-friendly printout of the VOMS
extension

Overhauling Authorization for the WLCG

- ▶ Identities don't belong in the distributed authorization decisions (though needed for traceability).
- ▶ We want authorization to be based on the experiment's assertion, not the user identifier.
- ▶ We want to attenuate / reduce the credential's authorization as much as possible!
 - ▶ Principle of least authorization!
- ▶ Technology-wise, we wanted to get out of the Grid Security Infrastructure (GSI) business.
 - ▶ So did the Globus folks who wanted to push this in the first place!
- ▶ **Idea:**
 - ▶ Replace the X.509-based credential technology with more commonly-used JWTs (bearer tokens).
 - ▶ Replace identity-/group-mapping infrastructure with **capability-based**: credential encodes what you are allowed to do, not who you are!
- ▶ We took the idea and ran with it. For CMS:
 - ▶ Started developing the technology in ~2018 with the SciTokens project.
 - ▶ The WLCG profile (describing the authorization infrastructure) was published in 2020.
 - ▶ Converted the CEs to use tokens in 2021.
 - ▶ Majority of storage systems added support in 2023.
 - ▶ Peaked at ~50% data movement by volume used tokens for authorization in DC24.

CMS Production Transfers – last month



Tokens – more ways to reduce power!

- ▶ X.509 proxies were either “full power” or “limited” (with no clear definition of what “limited” met).
- ▶ Tokens have several dimensions to restrict the authorization:
 - ▶ **Time**: nbf (“not before”) / exp (“expiration”): Restricts when the token’s validity begins and ends.
 - ▶ **Scopes**: Positive statements of “what can be done” with the token.
 - ▶ Compute-related: compute.read, compute.{modify,create,cancel}. Authorizes interaction with CE.
 - ▶ Storage-related: storage.read, storage.{create,modify}, storage.stage. Authorizes interaction with storage endpoint. Can be further restricted by prefix (storage.read:/simulations/A/2024/09/26)
 - ▶ **Audience**: Who is the token for?
 - ▶ Prevents a stolen token from host A to be used to send jobs to host B.
- ▶ Design tradeoffs: which of these mechanisms to use?
- ▶ Finer grained = more secure
 - ▶ ... = more complexity = more load = more difficulty.

No free lunch!

What are the *downsides* of tokens?

- ▶ **Tyranny of choice:** multi-dimensional optimization problem versus one-bit. “Perfect is the enemy of the good”.
- ▶ Switch from proof of possession to **bearer**: Full token is sent over the wire, meaning receiver has full power of token.
 - ▶ This happens with GSI for GridFTP and the Globus CE (automatic delegation).
 - ▶ The problem is sloppy software: headers and signed URLs were never redacted in our community. Lots of logging of secrets / credentials to files.
 - ▶ **It can be done: need to catch up with the rest of the world.**
- ▶ More **reliance on token issuer**:
 - ▶ Tokens focus on much shorter lifetimes.
 - ▶ Finer grained = Larger cardinality of tokens.
 - ▶ Token issuers tend to keep created tokens in a database.
- ▶ **No standardized CLI client / client library.**
 - ▶ Would prefer we *don't* have a standalone “token-init”-style client but each CLI client (rucio, CRAB) use a common library to get the token it needs.

Who is your token issuer?

- ▶ In X.509 / VOMS world, there was approximately one single VOMS server implementation.
 - ▶ VOMS-Admin was never ported to RHEL8+.
 - ▶ VOMS C++ implementation went from April 2021 to July 2024 without a production release.
- ▶ For tokens, there's a (relative) plethora of options:
 - ▶ **Indigo IAM**: Developed by INFN, used by the LHC experiments and at RAL. All-in-one solution, including authentication, group-management, and a **VOMS attribute authority**.
 - ▶ **OA4MP**: Developed by NCSA, deployed by CILogon, used by OSG services and FNAL/JLab experiments. Scriptable interface, *only* issues tokens (integrates with existing identity and group management solutions).
 - ▶ **EGI-CheckIn**: Implements AARC profile, meant to be single-instance issuer.
 - ▶ **KeyCloak**: Community project (CNCF); supported by RedHat. Powers CERN SSO.
 - ▶ With modest tweaks, could generate tokens following WLCG profile.

Where are things going? CMS

- ▶ I'm not qualified to give an overview of every LHC experiment.
- ▶ Within my wheelhouse: CMS
 - ▶ SAM tests thoroughly probe site services (CE, storage) for token support.
 - ▶ Missing support: STORM ☹️
 - ▶ Will continuously expand list of sites using tokens with Rucio.
 - ▶ Web services now can use WLCG tokens
 - ▶ Job submission:
 - ▶ (Fall) Using “htgettoken”/Vault to distribute tokens to laptops / lxplus-type environments.
 - ▶ (Fall) HTCondor credmon to ensure valid token is always available in the job.
 - ▶ (Winter) Adapt workflow management tools to manage & use tokens.
 - ▶ (Fall/winter) “rucio upload”, “rucio download” with tokens.

Let HTCondor manage your tokens

- ▶ You want to use tokens in your job, right?
 - ▶ **Don't** put an access token in `transfer_input_files`. The token will expire!
 - ▶ **Don't** put a refresh token in your job and create a new access token from the EP. Refresh token is too powerful!
 - ▶ **Do** let the AP manage your credentials!
 - ▶ The job will always have fresh, valid tokens, periodically updated.
- ▶ There are three credmons:
 - ▶ OAuth2 authorization code flow: Traditional OAuth2 flow, requires running a web server at the AP.
 - ▶ Vault: Receives tokens from a Vault host; eases pain of OAuth2 flows + multiple APs.
 - ▶ Local issuer: Requires signing key locally, creates token on the fly.
- ▶ These are all mutually exclusive!
- ▶ Up next for credmon:
 - ▶ Merge all code bases together – have as many instances as you would like!
 - ▶ Add support for client credential flow: skips need for user interaction if AP is a trusted OAuth2 client.
- ▶ Still missing: device code flow ☹️

X.509, Redux

- ▶ X.509 is a credential type. How do we validate credentials?
 - ▶ **GSI** (Grid Security Infrastructure): A set of rules for validating credentials specific to the grid, including proxies, namespace verification. Implemented by the Grid Community Toolkit / Globus Toolkit.
 - ▶ TLS / **SSL** / Internet PKI: The more common, RFC-standardized set of rules for X.509 clients. No namespace verification, no custom transport layer. Implemented by OpenSSL.
- ▶ In HTCondor 23, we support TLS / SSL / Internet PKI.
 - ▶ Different protocol than prior “GSI” protocol
 - ▶ If VOMS support is enabled (23.5.x or later), then VOMS C++ library will be used to validate the client credentials. **Allows VOMS proxies.**
 - ▶ If you trust a library with fairly inactive development for your core security.

Want to learn how to enable SSL/VOMS authentication? How to authorize based on VOMS proxies?

See <https://htcondor-wiki.cs.wisc.edu/index.cgi/wiki?p=HowToUseProxiesWithSsl>

Outlook ahead

- ▶ The token transition is well on its way!
 - ▶ Not all experiments proceeding in lockstep but forward progress year-over-year.
 - ▶ Big use cases in production for big experiments (e.g., CMS data transfers).
- ▶ Services largely have support in place.
- ▶ “Scaling down” to smaller organizations is the current challenge, esp. with token issuers:
 - ▶ IAM is a good “all in one” choice if you were happy with VOMS-Admin.
 - ▶ OA4MP is a good option if you want to provide your own LDAP instead.
 - ▶ KeyCloak has lots of promise; looking for partners.
- ▶ Clients and client libraries are likely the next focus!

Questions?

FEARLESS SCIENCE

This project is supported by the National Science Foundation under Cooperative Agreements OAC-2030508. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.