WLCG Tokens and SciTokens in HTCSS

EU HTCondor Workshop 2022

Jaime Frey

Why Move to Tokens?

- HTC has used X.509 credentials for user authentication
  - Proposed by grid community with extensions
- Tokens are more widely supported
- Tokens allow a better security model

Tokens are Better Supported

- Grid community used X.509 with several extensions (e.g., proxies)
  - Never embraced by industry
- Tokens have been widely adopted
  - Standards: OAuth2, OpenId Connect, JSON Web Tokens
  - Many tools and language APIs

Tokens Allow a Better Model

- X.509 credentials are identity-based
  - Who you are
  - Passport
- Tokens can be capability-based
  - What you can do
  - Football ticket
- Token's power can be greatly curtailed

JSON Web Tokens (JWT)

- A set of key-value pairs…
- Signed by an issuer
- Some keys are standardized
  - iss: Token Issuer
  - exp: Expiration time
  - scope: List of authorizations
  - aud: Service token can be used at

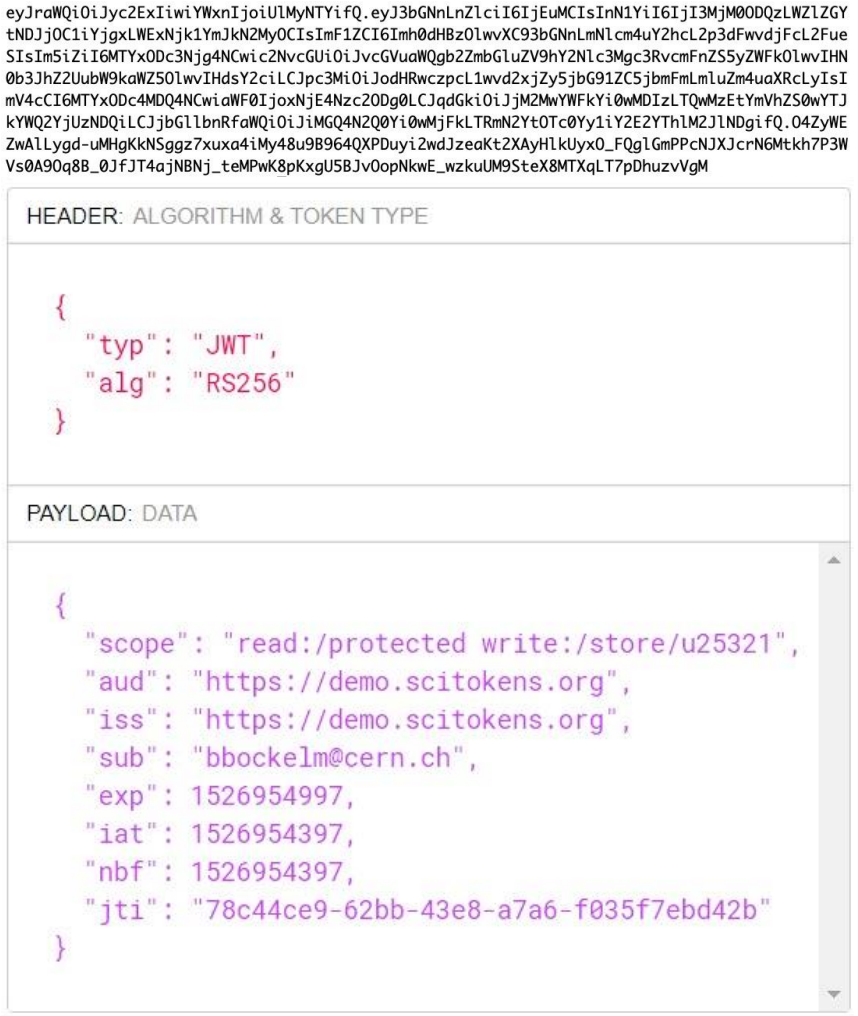

OAuth2 and OpenID Connect (OIDC)

- Standard frameworks for
  - User to authenticate with issuer and obtain a token
  - Service to validate a token presented by a user
- Widely used in industry
  - E.g. Google, Amazon, Facebook, Microsoft
- Lots of software and language support

SciTokens vs WCLG Tokens

- Both are based on OpenID Connect
  - WLCG tokens follow standard more strictly
- Both define file- and job-based authorizations
  - Additional authorization types are possible
- Format of scope names differs
  - SciTokens: read:/foo CONDOR:/READ
  - WLCG Tokens: storage.read:/foo compute.read
- HTCondor accepts both for job control

Token Discovery

- HTCondor tools look here for a token to use for authorization
  - `$BEARER_TOKEN`: value has token data
  - `$BEARER_TOKEN_FILE`: file has token data
  - `$XDG_RUNTIME_DIR/bt_u<id>`: file has token data
  - `/tmp/bt_u<id>`: file has token data
- First location with a valid token is used

Use With HTCondor-C(E)

- E.g., the factory use case
  - Using local AP to submit jobs to a remote AP
- Token must be in a file
- Add to your submit file
  - `scitokens_file = <filename>`
- Or if BEARER_TOKEN_FILE in environment, add this
  - `use_scitokens = true`

Authorizing SciTokens

- Specify in CERTIFICATE_MAPFILE
  - Which tokens/issuers to accept
  - Which identity (i.e. user account) to map them to
  - `SCITOKENS <issuer>,<subject> <username>`
- Accept specific token issuer and subject
  - `SCITOKENS https://demo.scitokens.org,jfrey jfrey`
- Accept all tokens from an issuer
  - `SCITOKENS https://demo.scitokens.org,.* jfrey`
- Set audience name of daemons
  - `SCITOKENS_SERVER_AUDIENCE = $(FULL_HOSTNAME)`

Play with SciTokens

- https://demo.scitokens.org
- Will issue any token you want
- Don’t use in production!
- Add these keys
  - "aud: "ANY”
  - (Or set config param SCITOKENS_SERVER_AUDIENCE)
  - "sub": "jfrey"
  - "scope": "condor:/READ condor:/WRITE"
- Configure AP to accept this issuer

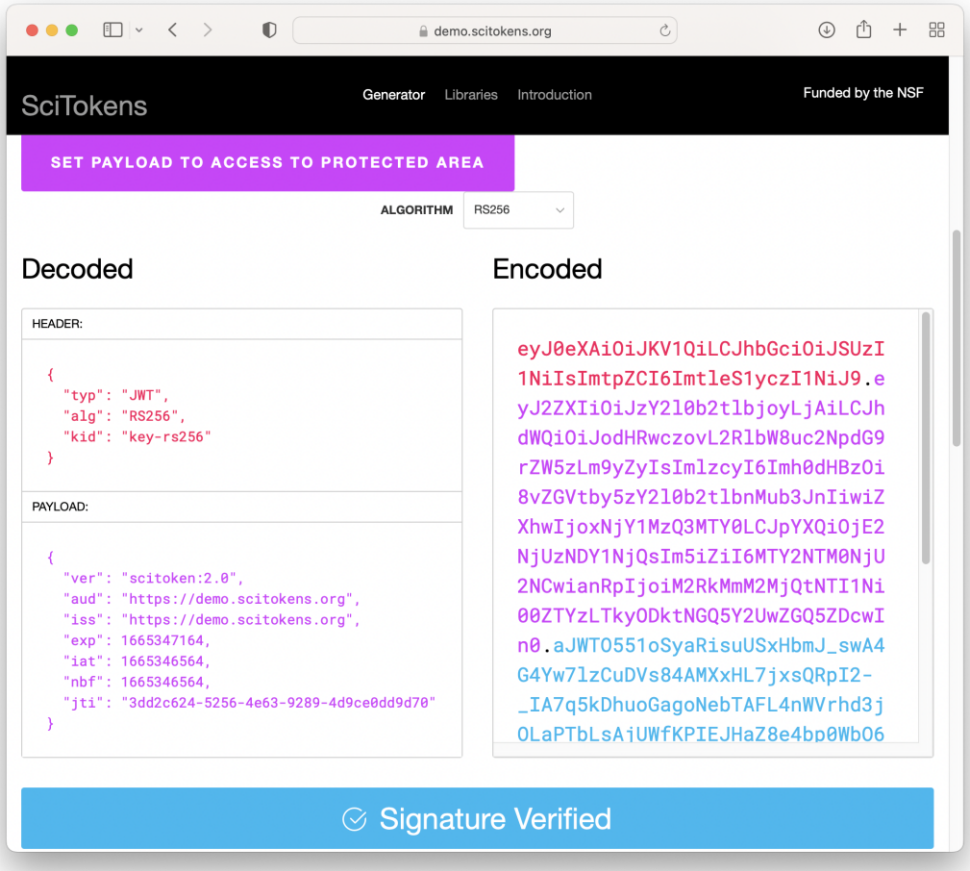

SciTokens vs IDTokens

- Both are JWTs and look similar
- SciTokens use asymmetric key
  - Anyone with issuer's public key can verify
  - Issuer must publish public key via https server
  - Suited for a VO accessing multiple services (including HTCondor pools)
- IDTokens use symmetric key
  - Issuer's private key required to verify
  - Suited for a single HTCondor pool
  - Doesn't use OAuth or OIDC

Tokens for the User Job

- Not exclusive to WLCG/SciTokens
- User can request tokens in submit file
  - use_oauth_services = box
  - box_oauth_permissions = read:/public # optional
  - box_oauth_resource = <resource> # optional
- condor_submit will do OAuth2 process to acquire tokens
  - Usually direct user to URL to authenticate with service
  - Tokens available to job under execute directory
- Supported services: box, gdrive, onedrive, dropbox, vault
  - Admin can add services

Coming Soon

- Currently, authorization based only on issuer and subject
  - Additional fields in token can't be considered
- Add callout mechanism
  - Allow admin to provide arbitrary logic for authorization and account mapping decisions
  - Akin to LCMAPS used with X.509 credentials
  - Easier to support other JWT-based tokens (e.g., EGI Check-In)
  - Coordinating with ARC CE for common interface

Coming Soon

- Improve token debugging
  - Allow admin to create fake token from any issuer
  - Only accepted by their CE and only for a brief time
  - Full software stack will accept token with no config changes
  - Admin doesn't need real token to debug auth problems

Coming Soon

- Improving SciTokens C++ library
  - Remove functions deprecated in OpenSSL 3
  - Add non-blocking interface
  - Fetching issuer's public key delayed by DNS problems

Thank You!

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

This work is supported by NSF under Cooperative Agreement OAC-2030508 as part of the PATh Project. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.