

Transfer Ecosystem Demonstrator

Brian Bockelman

<https://indico.cern.ch/event/696511/>

Goals for Today

- Outline ideas for HTTP-base transfer ecosystem.
 - Get (mostly) everyone on the same page
- Make sure everyone read the CHEP abstract submitted:
<https://docs.google.com/document/d/19ztslidwWI5XZtk1c-45qf5TWdkEgOXRZsjigV6xVlg/edit>
- Put together a map forward.
- Decide on a few community milestones.

HTTP Transfer Ecosystem

- **Very important:** This is not new. HTTP(S)-based third party copies have been occurring for >2 years.
 - Thanks to CERN IT and dCache teams for absolutely driving these efforts!
- Idea: Use (abuse?) a WebDAV extension, the COPY verb.
 - Eliminate the use of GridSite delegation so the “initial side” of the transfer does not receive the proxy.
 - Instead, utilize a bearer token.

Unique Twist: No X509!

- Bearer token: an opaque string which authorizes the bearer (holder) of the string to perform an action.
- Obvious implementations: SciTokens, Macaroons.
- Intent is to work toward a capability-based ecosystem.
 - **SciTokens**: token (based on JWT) can be independently verified to be a valid token from a VO.
 - **Macaroons**: capability-based scheme where token can be issued by site (symmetric encryption) and verified by site.
- Utilize these **tokens for authorization scheme instead of client X509.**

Token-based

- Once you have a token, it's easy to add it to the transfer — simply an HTTP header.
- How do you get a token? How does FTS get the token?
- In current pull request:
 - `fts-transfer-submit` works as normal.
 - File metadata
- Is this the right model? Discuss...

Positive Progress

- dCache and DPM, I think, contain all the implementations for this extension.
- Over the last two months, a new plugin was developed for Xrootd: <https://github.com/bbockelm/xrootd-tpc>
 - Shown to be compatible with dCache, DPM.
 - Does *not* support GSI delegation - only token-based.
 - Supports multi-streaming.
- With this plugin - and a few patches to FTS, GFAL2, DAVIX, and Xrootd - we can do TPC between dCache and Xrootd using tokens.

Next Levels?

- There are two token types supported:
 - VO-issued: Current intent is to have Rucio issue this.
 - Site-issued: fts-url-copy uses delegated proxy to get token from site. Rucio could generate tokens on behalf of the user
 - What about signed URLs?
- Thoughts on dCache support?
- What should we set as a goal for CHEP (July) and the WLCG workshop (March)? I see three dimensions:
 - Integration of capabilities with services (Rucio, FTS).
 - Number of sites in the testbed.
 - Number of TB moved.

Open PRs

- FTS3 support for bearer tokens: https://gitlab.cern.ch/fts/fts3/merge_requests/17
- Xrootd support for HTTP chunked transfer encoding: <https://github.com/xrootd/xrootd/pull/636>
- Xrootd bugfix for HTTP verbs: <https://github.com/xrootd/xrootd/pull/638>

References

- Xrootd/SciTokens plugin: <https://github.com/scitokens/xrootd-scitokens>
- Xrootd/TPC plugin: <https://github.com/bbockelm/xrootd-tpc>
- Bearer token acquisition: <https://github.com/bbockelm/x509-scitokens-issuer>