DC24 Technology Demos: Capability-Based Tokens

But First, Why?

Why evolve the system?

1. **Improve security & flexibility**: Aim for a system that allows for finer-grained privileges. Decreases risk of a stolen credential – and allows new workflows not possible before.
2. **Retire older software**: Much of the current authorization system software is seeing minimal improvements (and edging on abandoned). Little/no community, little/no investment.
3. **Sustainability**: Aligning with a larger community brings the capability to use other software – HEP doesn't carry the burden alone!

Capabilities and Authorization

- The core idea behind capabilities is the credential should say **what you can do** instead of **who you are**.
- Analogy:
  - In X.509 land, we basically pass around everyone's social security number. Effective & simple!
- With tokens, the authorization to do an action is based on a signed statement (the token!) from the experiment.
  - Other information *may* be used – e.g., the *sub* might map you to a Unix account for new files – but that's after the authorization decision.

HEADER: ALGORITHM & TOKEN TYPE


```
{
  "alg": "RS256",
  "kid": "key-rs256",
  "typ": "JWT"
}
```


PAYLOAD: DATA


```
{
  "scope": "read:/protected",
  "aud": "https://demo.scitokens.org",
  "ver": "scitoken:2.0",
  "iss": "https://demo.scitokens.org",
  "exp": 1694523232,
  "iat": 1694522632,
  "nbf": 1694522632,
  "jti": "ce007e2a-469d-47bd-988a-b10bc498395a"
}
```

The Three-Layer-Cake of LHC Transfers

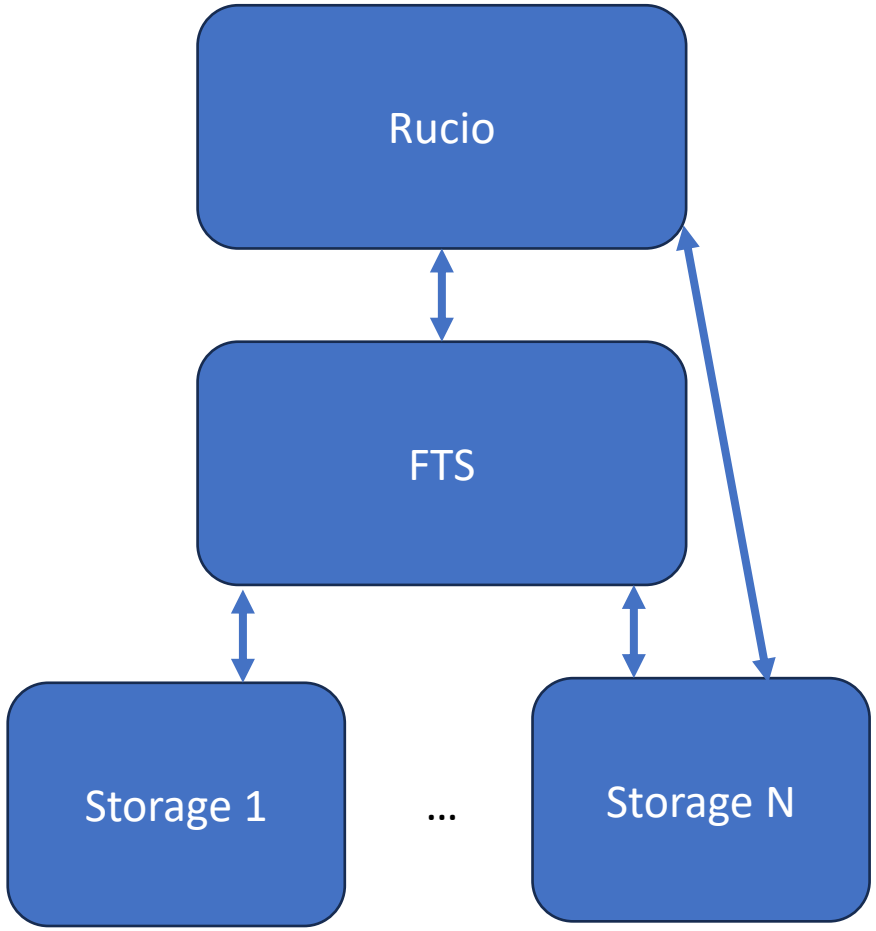

- Rucio manages a catalog of datasets, file locations, and placement policies on behalf of an experiment.
- Rucio decides which files need to be moved and issues a transfer request to FTS.
  - For some operations (deletes), Rucio interacts directly with the storage.
- FTS, working across multiple experiments, schedules and manages the transfers between the storage.

How FTS Used to Work (Authorization-wise)

- Client delegates its X.509 credential to FTS.
- Same X.509 credential used to initiate transfer.
  - X.509 delegated again to one endpoint (needed if encryption is desired).

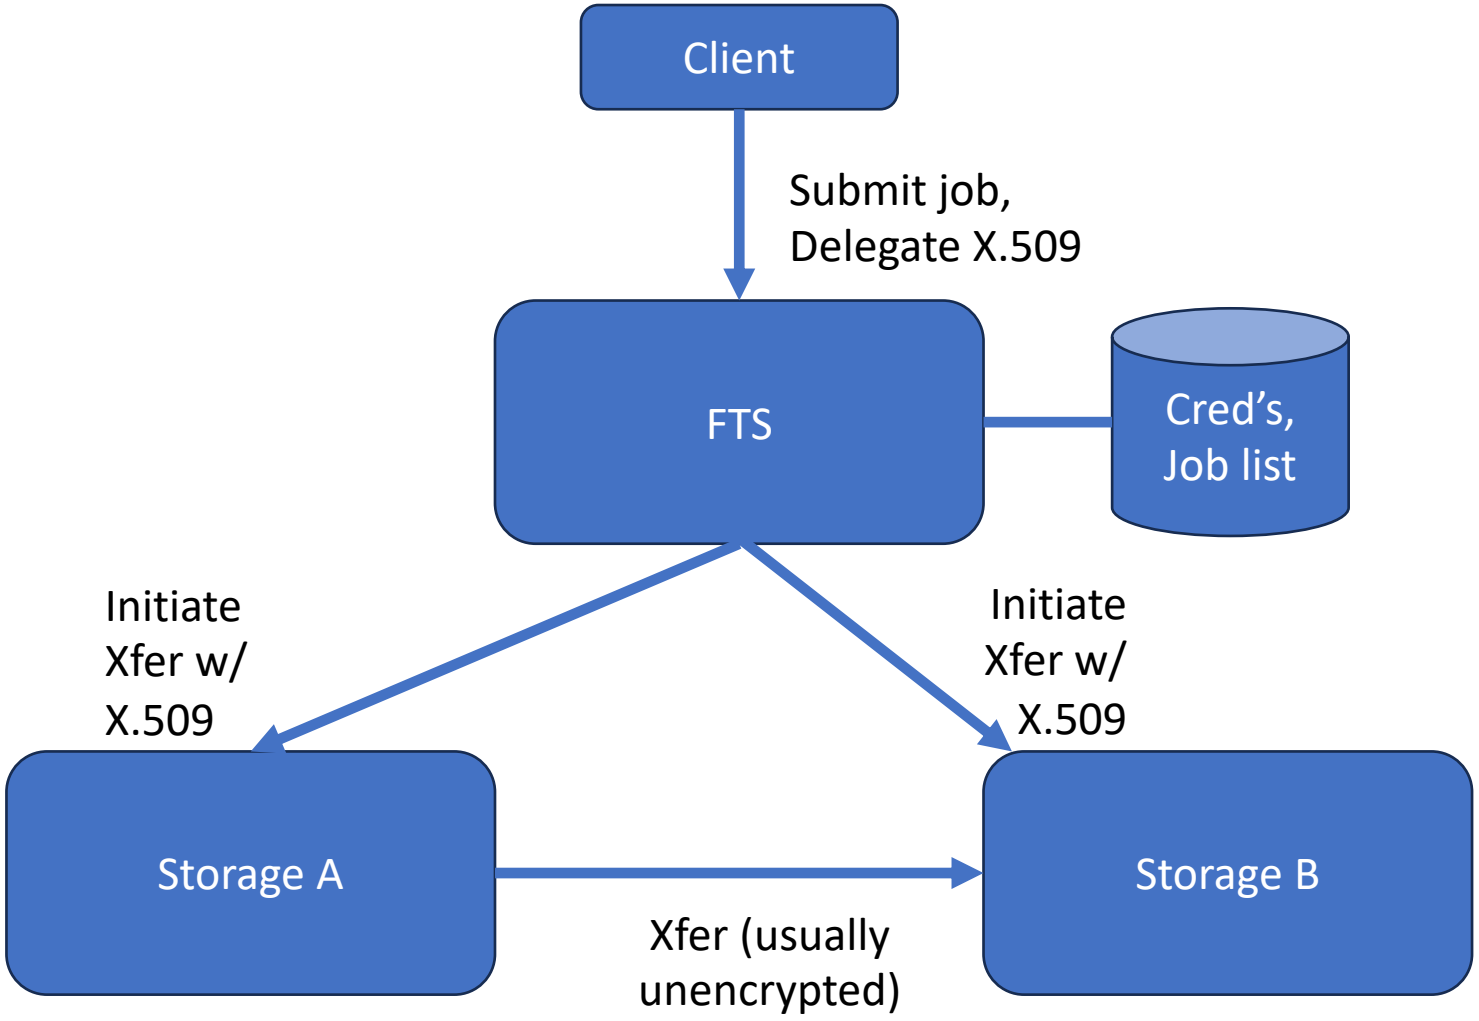

How FTS Currently Works (Auth'z)

- X.509 is still used to contact both storage endpoints.
- One endpoint (typically, the source) generates a local storage token for the transfer.
- The other endpoint is given the token; an HTTP download is initiated.

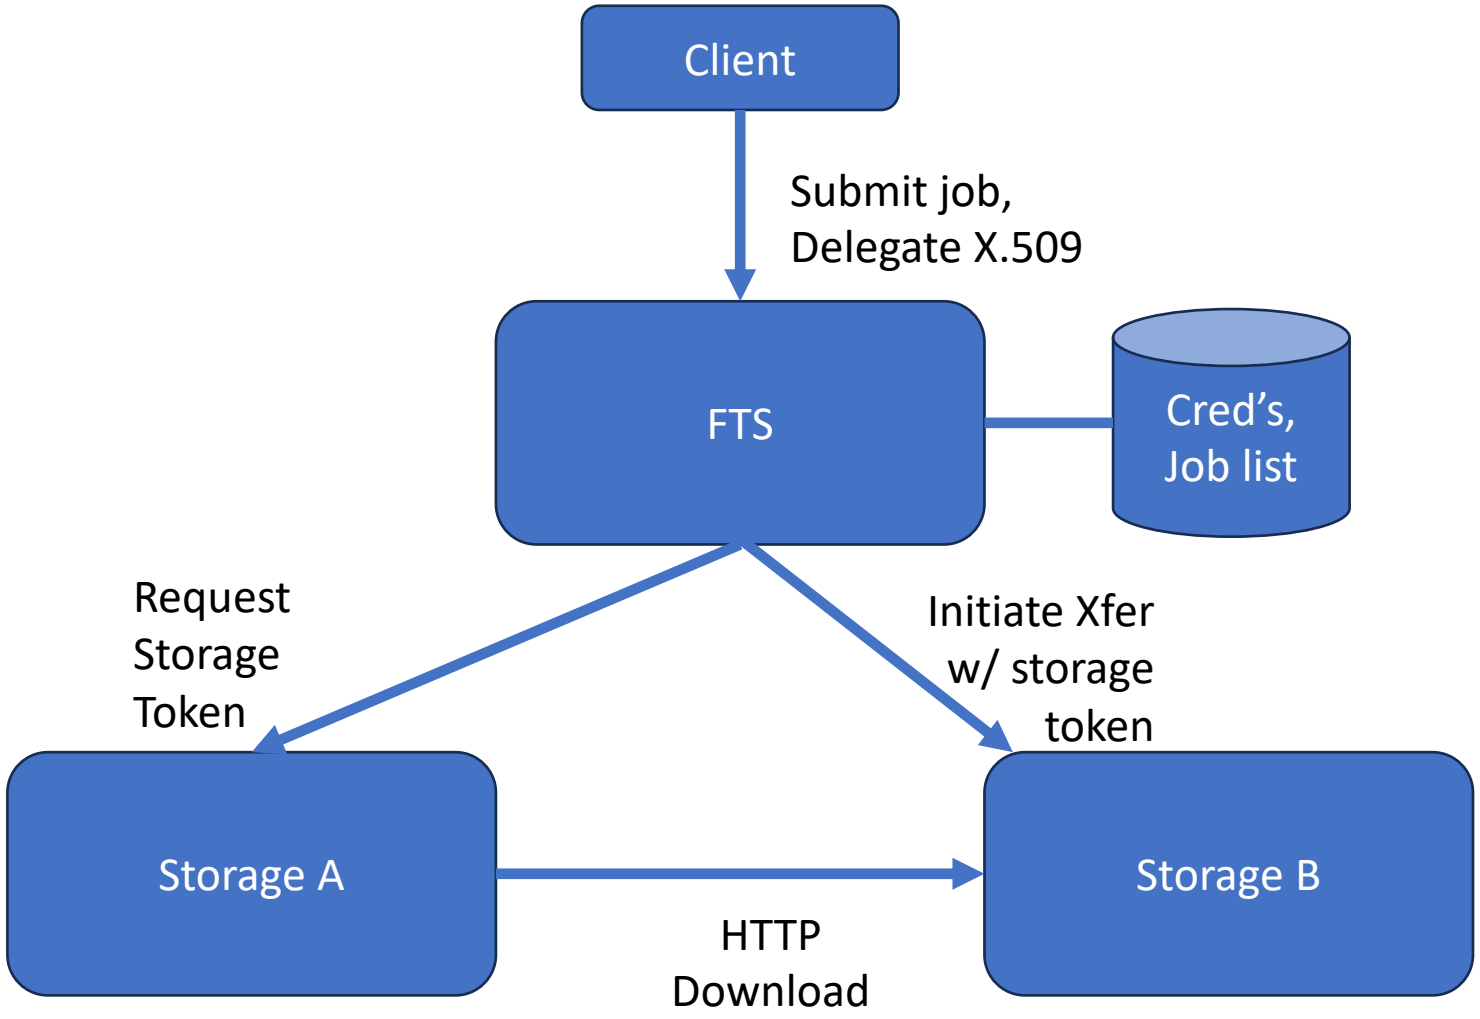

Design Considerations

- The goal is to have X.509-free transfers – however, there are multiple tokens to consider!
  - One token may authorize access to the FTS service.
  - Others will authorize access to the storage areas.
- Who is authorized to request the token? The client or FTS?
- How fine-grained should the token be?
  - One token for the whole experiment is equivalent to the setup today.
  - One token per transfer tightly couples the transfer system to the health of the token issuer.
  - Note: just because you can go “finer grained” doesn’t imply “finest grained” is required!

Current Approach

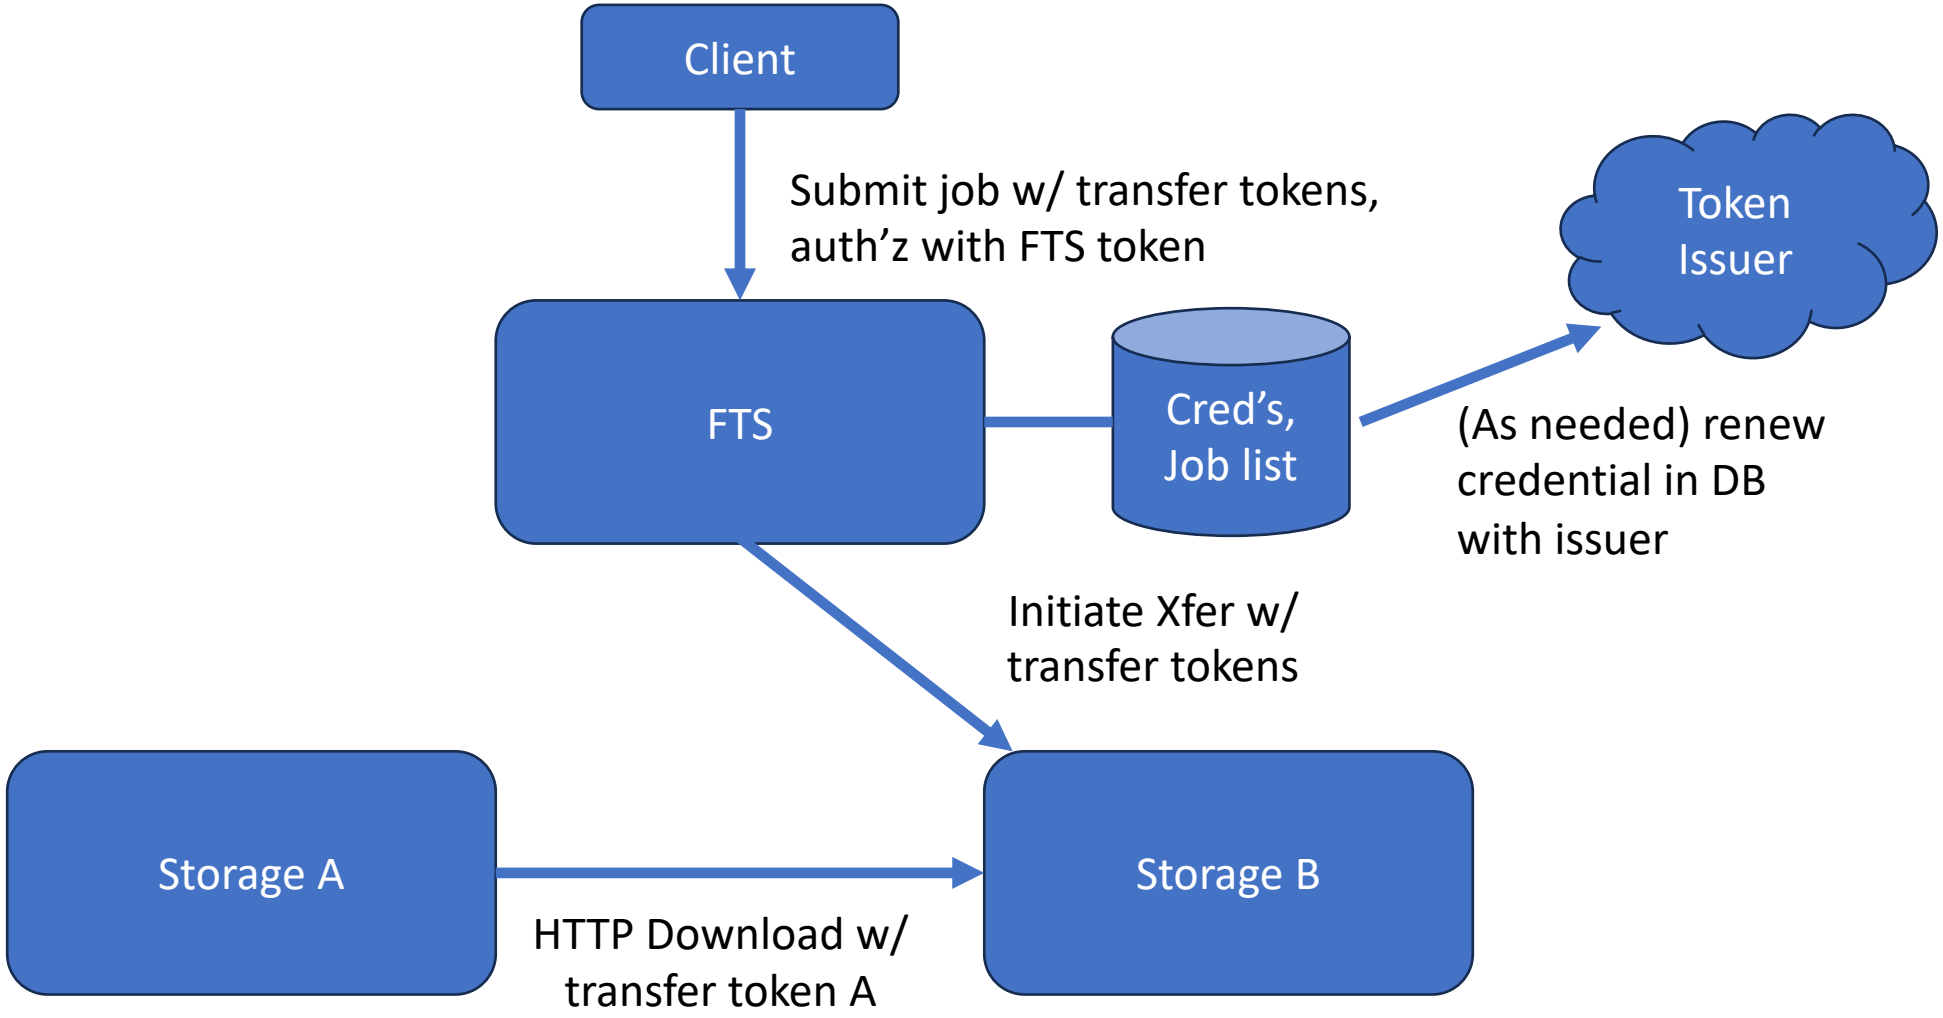

Approaches

- Separate out the “transfer job submit” API call from the “delegate token to FTS”.
- Tokens are specified and managed by the client. The service on top (Rucio) decides how many tokens it wants to manage.
- My goal:
  - Have Rucio manage O(# sites) tokens (e.g., a read & write token per site). Increases the granularity (improves security) without drastically increasing the reliance on the token issuer.

Timelines / Risks

- Goal is to do a MVP/demo at the November WLCG DC24 workshop.
  - Not all pieces would be working – particularly, the token renewal inside FTS is separate from the submission interface.
- The November milestone provides enough time prior to DC24 to change plans.
  - Risk: No one outside CERN is on the FTS team. Unclear how to provide effort to the project. Risk is shared with other experiments (DUNE) as well!
- Goal for DC24 is:
  - 5% of sites will participate. Likely already beyond

Questions?

This project is supported by the National Science Foundation under Cooperative Agreements OAC-1836650 and PHY-2323298. 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.