Toward a Post-Globus Toolkit Future

Brian Bockelman
OSG AHM 2018

Work supported by NSF OAC-1738962 and PHY-1148698

Globus Toolkit - Original Vision

- The original vision of the Globus Toolkit was to be *the* foundational layer for all grid computing.
    - A swiss-army-knife that did all the basics so your VO could focus on the high-level.
        - This ideal environment hit issues as the community's needs were specialized, developed better components, or simply adopted different models
- The LHC community investigated - and discarded - some components early on (can you remember these acronyms?):
    - **RLS server**. LFC was adopted instead, and even that is no longer used.
    - **MDS server**. BDII adopted instead; BDII has also largely been discarded.

Globus Toolkit - Gradual Progression

- In the RLS and MDS cases, the Globus component was exchanged for something similar non-Globus, then replaced with a different architecture.
    - The second step was done more recently with GRAM transitioning to HTCondor-CE. We are seeing projects at various phases in their lifecycle aiming to completely rework the concept of a CE.
- This is all a very natural part of the software lifecycle!
    - A big part of what OSG does is **help manage the software lifecycle** - reduce disruptions caused by transitions.
    - More important than “what technology are you using now?” is “**how do you get to the technology that comes next?**”

Globus Toolkit - Current Usage

- So, we no longer use GRAM, RLS, MDS, CAS, or any WSRF. *What is in use?*

- **GSI**: Reference implementation for our X509-proxy-based authentication / authorization ecosystem.

- **MyProxy**: X509 credential management. Used to store long-term credentials delegated to other services.

- **GSI-OpenSSH**: Set of SSH patches to enable GSI-based auth{z,n}.

- **GridFTP**: Reference implementation of the GridFTP protocol (plus extensions). Highly pluggable, so used in multiple storage systems.

Globus is going away

- Last June, Globus announced support for the Globus Toolkit was ending December 2017 (security-only support for another year).
    - Their organization's services planned to stop using GT components.
    - They didn't have a mechanism to provide sustainable support for the GT community.
- The GT support community didn't extend beyond the existing NSF project!

https://opensciencegrid.github.io/technology/policy/globus-toolkit/

Message #1:

Don’t Panic

Grid Community Forum

- There are several organizations that rely on similar functionality out of the Globus Toolkit — CERN, EGI, OSG, PRACE, XSEDE.

- Members of these organizations banded together to create the **Grid Community Forum** in order to maintain a fork of the Globus Toolkit, the **Grid Community Toolkit**.

- This mechanism will provide baseline support for the functionality we need.

    - Given the maturity level of the software, effort level is fairly manageable … until OpenSSL breaks its ABI.

    - This happens every 3-4 years: hence, we have a reasonable amount of time to plan for the future.

- Note that GridCF could potentially include other software stacks under its umbrella in the future.

Beyond GCF

- With a few years to work on new approach, how could the new ecosystem work?

  - **Transfers**. Migrate from GridFTP to HTTPS/WebDAV, particularly using the COPY verb for third-party-transfers.

  - **Authorization**. Move from an infrastructure that is identity based (“who are you?”) to capability based (“what are you allowed to do?”). From a bearer-token model, we can utilize existing standards such as OAuth2 for moving tokens between entities.

WebDAV TPC

- WebDAV TPC is done by FTS contacting one storage endpoint, asking it to COPY to/from a given URL.
    - The active endpoint performs the transfer, typically a HTTP GET or POST.
    - Important: **ANY URL** can be given, including GridFTP or XRootD.
    - "Storage B" needs to know *nothing* about WebDAV TPC; only needs GET/PUT semantics. Allows transfers with S3, for example.
- Already widely implemented, including plugin available for XRootD (`xrootd-tpc` in osg-upcoming).
- Tricky part: *authorization* with Storage B. For this, we are working on a concurrent transition away from X509 to bearer-token based.

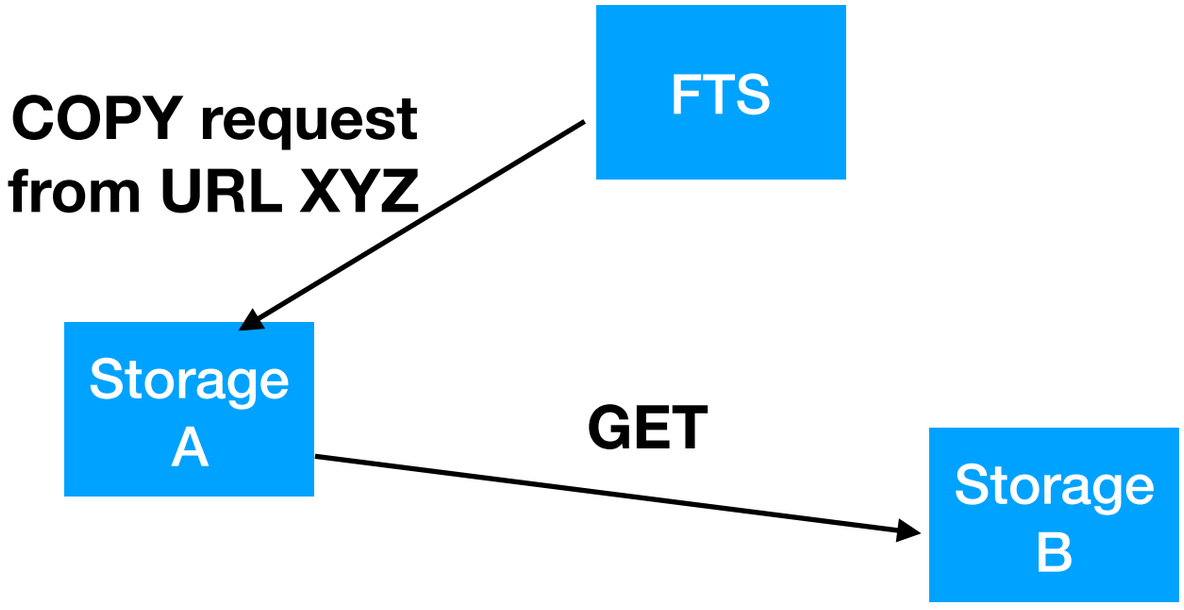

Authz revolution:

- **Identity-based:** authorization based on mapping who you are.
- **Capability/Token-based**: authorization based on something you are able to present.

Authorization Ecosystem

- In this ecosystem, the VO issuer would issue a signed token authorizing an action (“open file at Nebraska”) to whoever holds the token.

- Token is managed by a user agent (such as HTCondor) and distributed with the job.

- Job presents the token in the storage request (example: using HTTP header).

- Storage service checks signature against VO issuer’s public key, then authorizes the action relative to the VO’s storage area.

  - Storage service DOES NOT map users or manages access within filesystem.

  - User identifier still present for traceability purposes.

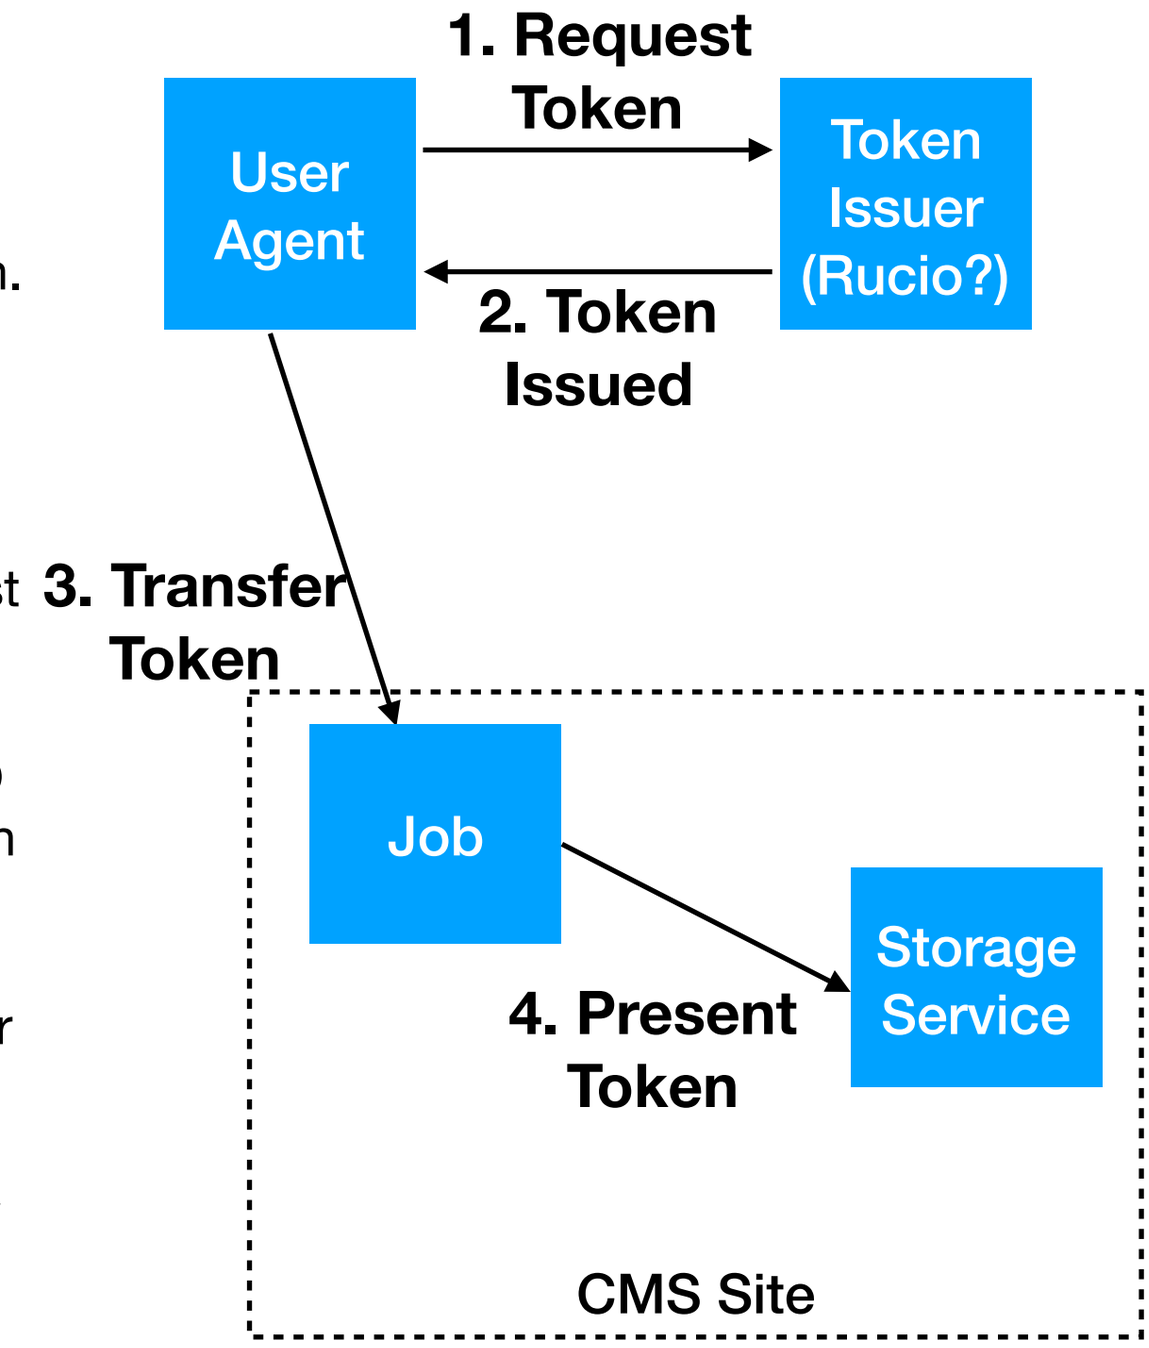

NOTE: User never explicitly manages their own tokens.

Put it together

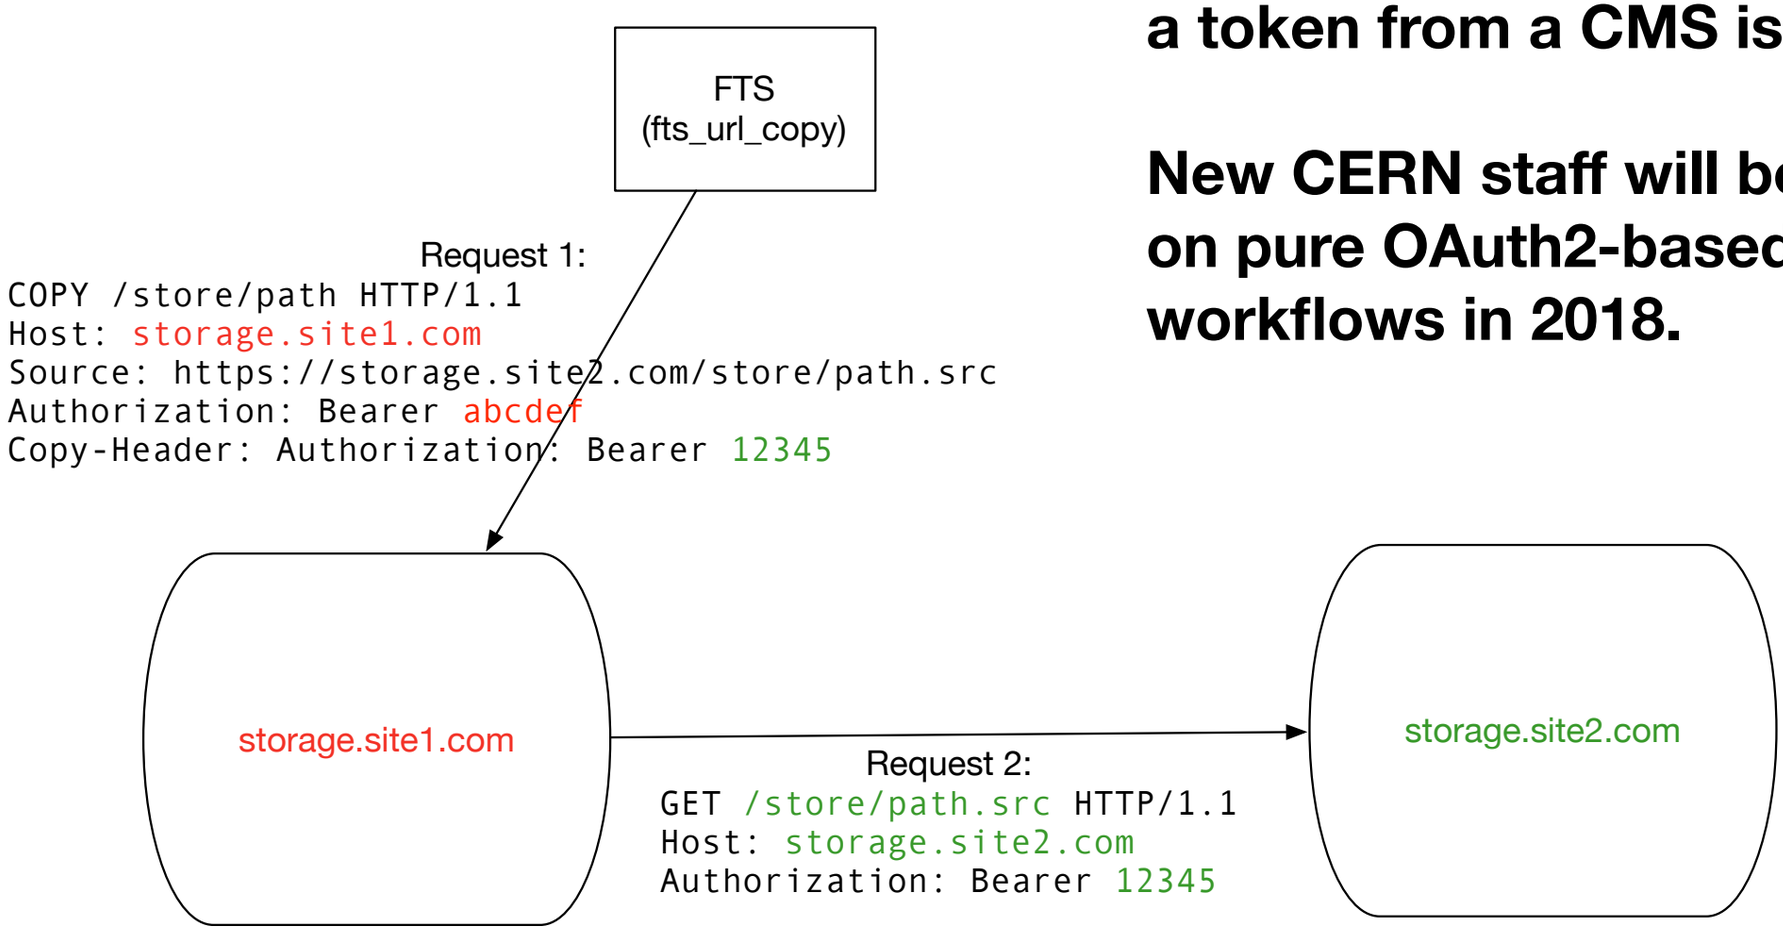

Currently, FTS acquires a token by exchanging a X509 proxy for a token from a CMS issuer.

New CERN staff will be working on pure OAuth2-based workflows in 2018.

No Globus, no user certificates, fine-grained security.

What do tokens look like?

eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjYzMjUxYWFmMjM2OWQ2Njc3ZjYwZmFhZDY0ZDJjYTcwMzFjZjNlYTVkMDYxM2M5YmU3OGZlYWYyZGVkNWM0NzcifQ.eyJqdGkiOiJiOGQ1NGE2Mi1jZDMzLTRiNGItYmI2NC0xMWI4MDQyNzJmMWQiLCJzdWIiOiJjbHVuZHN0IiwiZXhwIjoxNTIxNTYxMzgyLCJpc3MiOiJodHRwczovL3NjaXRva2Vucy5vcmcvY21zIiwiaWF0IjoxNTIxNTU3NzgyLCJzY3AiOlsid3JpdGU6L3N0b3JlL3VzZXIvY2x1bmRzdCIsInJlYWQ6L3N0b3JlIl0sInJiZiI6MTUyMTU1Nzc4Mn0.bmL-w3drGqTEyZCXt_06DTeg2cT6aWobDMQhKP00dEyk8FpO3YjqBz-ToNnmEDEeFuCRLgpwx_U14ttBulxkOw6hkOFbPGOilsLazK-6WcP2cpFkM9uuUJU_1_wLrCSkGBrFIjJwH4uzRWCIlttEAFnjy19QLCwgLsuGOBjNOdxTSh4Gp_SqjHYtMnXfI8T3q3krGCtSM39YmNwgc8oNYF7Bmq8WEP_Hqa-wr0B7qxMYb89BZQrZMjqve91klqI5Fvjd5RobA1oe5EMC9-FREInq8xCUu4TtpQAazC3UmL4vXipLUPGRXSioLGU6R74aP0AvIbc5sWvKsdeXCvea4Q

Decoded header:
{“alg":"RS256","typ":"JWT","kid":"63251aaf2369d6677f60faad64d2ca7031cf3ea5d0613c9be78feaf2ded5c477"}

Decoded payload:
'{"jti":"b8d54a62-cd33-4b4b-bb64-11b804272f1d","sub":"clundst","exp":1521561382,"iss":"https://scitokens.org/cms","iat":1521557782,"scp":**["write:/store/user/clundst","read:/store"]**,"nbf":1521557782}'

Full Stack

The scale is currently unimpressive, but we've been able to recently show the entire experiment vertical stack. For CMS, this is PhEDEx, FTS, Xrootd.

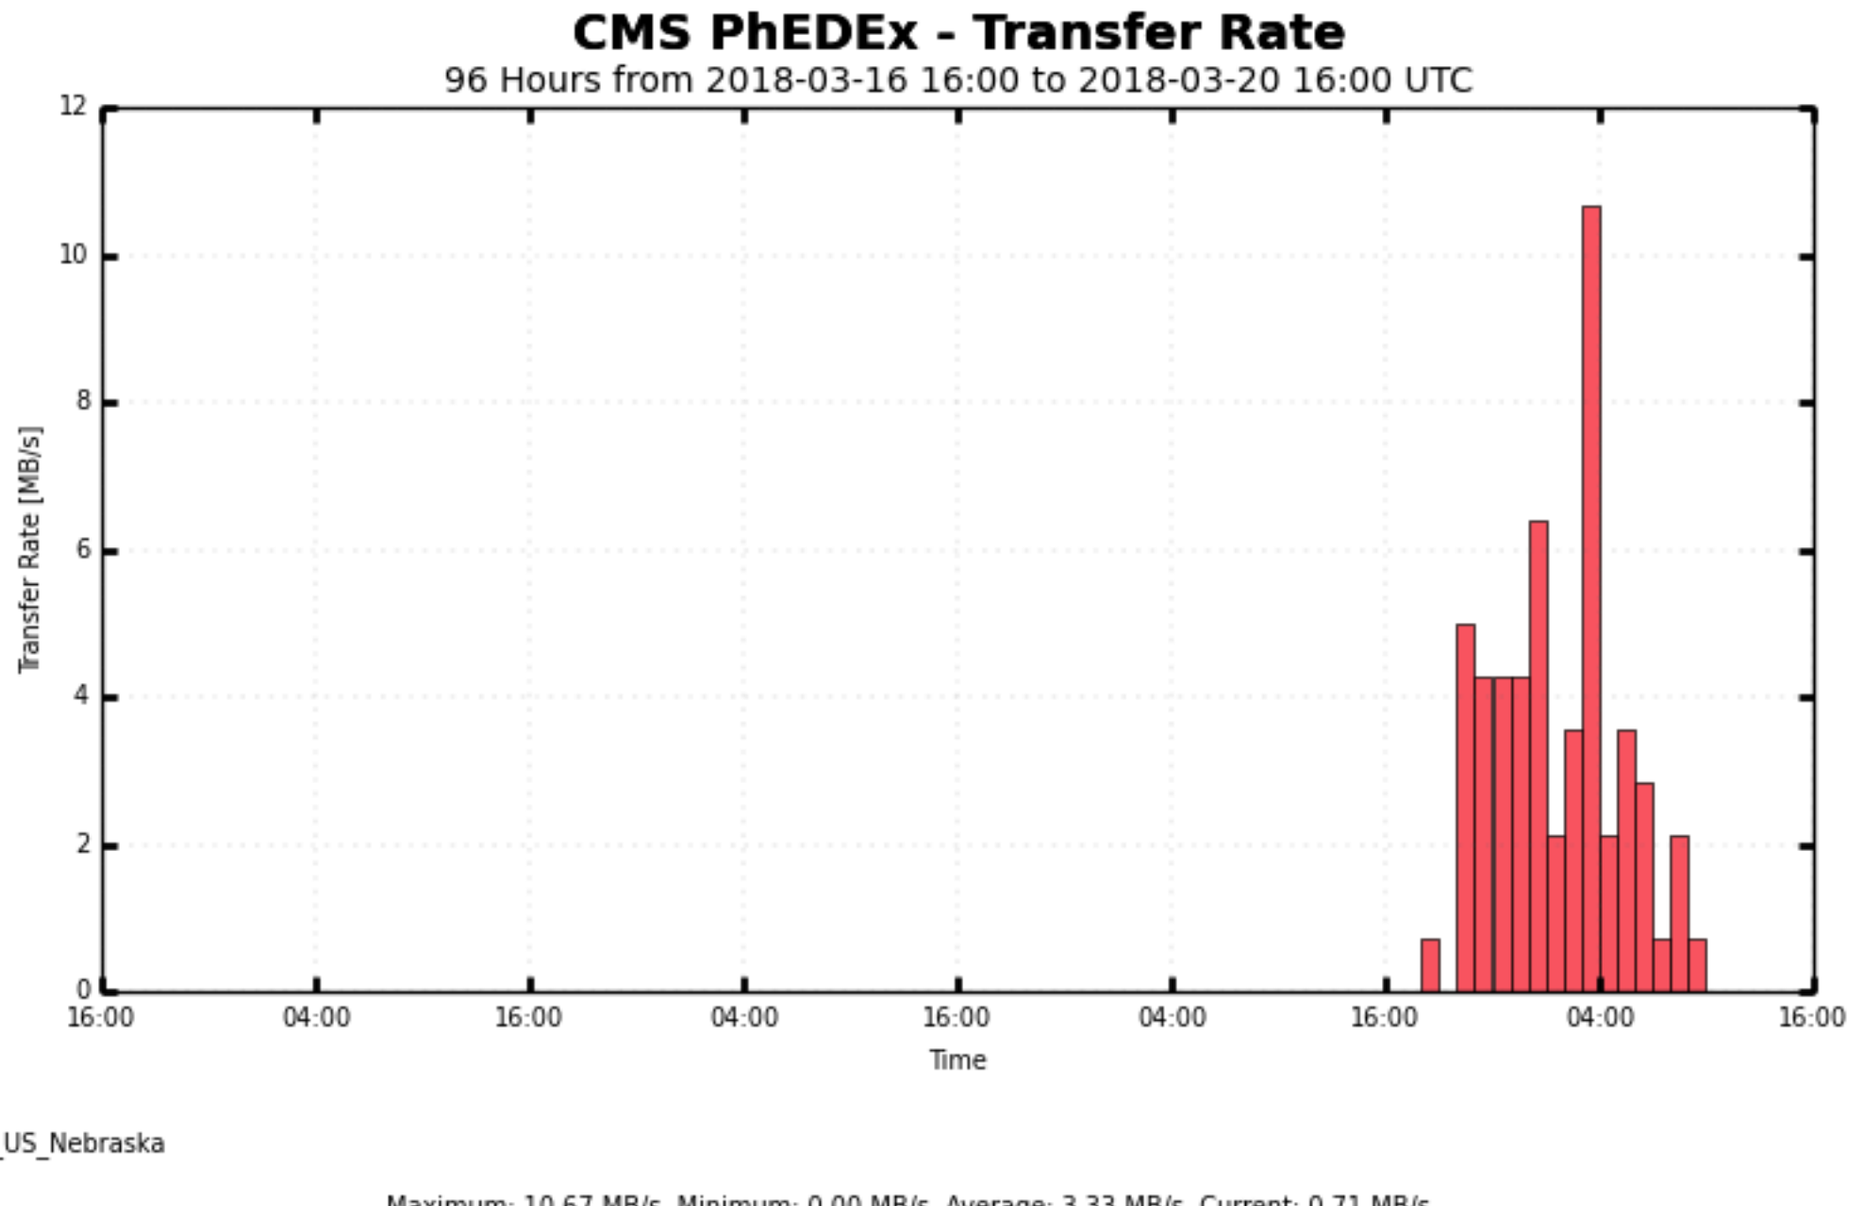

Looking Forward

- The Globus Toolkit itself is not an exceptionally large code base: OSG has managed larger transitions in terms of LoC before.

- However, it's a notable transition for two reasons:

- We are aiming to evolve past the identity-based authorization model. Model changes are harder than code changes!

- Transfer layer requires interoperating with the wider WLCG community: need careful collaboration to make sure we all go the same direction.