WLCG transition to tokens and Globus retirement

operations aspects

GDB meeting, March 10th, 2021

Maarten Litmaath
CERN

v1.0

Transition to tokens

- Since mid 2017, the WLCG AuthZ WG has looked into how we may migrate from cumbersome X509 user certificates to convenient modern mechanisms based on federated identities and JSON web tokens
  - As used in industry (Google, Facebook, GitHub, …)
  - And hence also in academia

- Users will ideally be able to “log into” the grid just using their home institute credentials and obtain VO-specific permissions from a new AuthZ service that hands out tokens to be used in grid workflows: INDIGO IAM

- The uptake of WLCG tokens is being pioneered in FTS third-party copies (TPC) using HTTPS / WebDAV, coordinated by the DOMA TPC WG
  - HTTPS / WebDAV can be used with X509 *and* tokens
  - GridFTP supports *only* X509

From VOMS-Admin to IAM

- IAM still has a VOMS service endpoint for legacy workflows requiring X509 + VOMS
- IAM does not have a VOMS-Admin service
  - User registration and group management go through a different SW stack that is suitable for the new landscape
- However, VOMS-Admin is also used by various grid services to obtain *lists of user DNs* per group or role
  - Certain use cases will become obsolete with tokens
  - Some may remain valid, though
- IAM supports similar functionality in principle
  - But currently only for entities that are *registered*
  - Registrations imply some complications
  - We may be able to let CRIC support some use cases

Relations to Globus

- A *range* of grid client and service implementations rely on Globus libraries to support X509 and GridFTP (details on the next pages)
- Moving away from those implies we stop depending on Globus
- A Globus retirement timeline discussion was prompted by OSG plans based on modernization incentives
  - X509 → JSON web tokens
  - GridFTP → HTTPS / WebDAV
  - Final OSG release dependent on Globus to be retired by Jan 2022
  - OSG will then stop contributing to Globus maintenance efforts
- The remaining relevant parts of Globus and a few related products are currently being maintained as the Grid Community Toolkit by the Grid Community Forum with contributions from several partners
- Technical details on Globus usage are being collected here

Supported components

- GSI
  - To deal with X509 in many products (see next pages)
  - A critical dependency of the other components
- GridFTP
  - Work ongoing in DOMA TPC WG to phase out the GridFTP *protocol*
  - Also used for job submissions to ARC
- MyProxy
  - Critical for most experiments
  - And for the new AAI while X509 / VOMS legacy services (if any) still need to be supported *behind the scenes* !
- GSI-OpenSSH
  - Critical for WLCG VOBOXes, particularly for ALICE
- UberFTP
  - In WLCG only used via CREAM, as far as we know

Development aspects

- Currently the code is mostly stable
- Some fixes may be needed for bugs e.g. encountered on CentOS 8
- The biggest concern probably is the missing support for TLS v1.3 in GSI
  - SSL/TLS transitions have typically been dictated by external circumstances not easily planned for
  - Significant development and debugging efforts might be required not only from experts in GSI code, but also from products it interacts with

Affected external MW (1)

- ARC
  - The GridFTP interface already has HTTPS alternatives: EMI Execution Service, REST (later this year)
- DPM
  - No problem expected to eject Globus dependencies eventually
- FTS
  - Globus dependencies mainly through GFAL2
- GFAL2
  - Only SRM and GridFTP plugins depend on Globus
- HTCondor
  - X509 *proxy* support depends on GSI
  - Job submission to ARC currently depends on GridFTP and GSI
  - HTCondor CE authorization often depends on Globus call-outs to LCMAPS or Argus

Affected external MW (2)

- LCMAPS (AuthZ plugin framework) – a (potential) dependency of various products
  - ARC, HTCondor CE
  - Globus GridFTP server (also used by StoRM), GSI-OpenSSH
    - For VOMS mappings
  - XRootD
    - OSG, (US)ATLAS, (US)CMS and GridPP appear to depend on it
  - Others?

- StoRM
  - SRM frontend (needed for tape at CNAF) currently depends on CGSI-gSOAP and Globus

- New AAI
  - For some use cases X509 may be needed *behind the scenes* during a transition period lasting up to a few years
  - RCauth online CA depends on MyProxy
  - CILogon ditto

Experiment dependencies (1)

- ALICE
  - MyProxy and GSI-OpenSSH are needed for WLCG VOBOX operations at most sites
- ATLAS
  - PanDA depends on MyProxy
  - Harvester: HTCondor-G job submission to ARC
  - Rucio:
    - No direct dependencies
    - GlobusOnline: optional transfer tool
    - dCache / StoRM tape access via SRM → GFAL

Experiment dependencies (2)

- CMS
  - X509 handling in various places may currently rely on GSI (to be further looked into)
  - CRAB depends on MyProxy
  - Several production services inside CMSWEB also depend on MyProxy
  - Rucio: see ATLAS details
- LHCb
  - DIRAC has no direct dependency
  - Indirect dependencies through GFAL2 etc.
- Other experiments (Belle-II, DUNE, …) are affected in very similar ways

Other communities

- WLCG sites need to support other communities that may not all be able to move to tokens on the same timescale as WLCG

- Sites may thus need to run GSI-based services for a while longer anyway
  - GridFTP (and SRM) could already be replaced with HTTPS / WebDAV and Xrootd
    - Both are supported by GFAL2 and FTS

- Such services then would still need to be supported
  - Possibly coordinated through EGI / EOSC

- The Grid Community Forum will play a crucial role there

Intermediate conclusions

- It looks possible for WLCG dependencies on GridFTP to have been removed by the end of 2021
  - But it may well be needed by other communities still
- For GSI the situation looks a lot less rosy
  - Many direct and indirect critical dependencies
    - Particularly through MyProxy
- It probably would be very expensive to try and remove such dependencies everywhere, while we have not fully moved to tokens
- This implies we need to ensure the remaining critical parts continue to be maintained
- At the same time we reduce our dependencies and hence risks where we can

Dependency plans

- Moving away from the GridFTP protocol allows our dependence on Globus GridFTP code to be stopped
  - Used by DPM, StoRM, FTS, GFAL2 and for ARC CE today
- That is mostly handled by the DOMA TPC WG and making good progress
  - ARC CE clients should switch to the REST interface this year
- We also need to have the remaining uses of SRM (for tape) replaced with alternatives that work with tokens
- An equivalent of MyProxy is desirable to obviate the spread of potentially powerful refresh tokens
  - Dave Dykstra's token client could be at the core of such a service
- There is no plan for a GSI-OpenSSH equivalent yet
  - As a fallback, SSH could be used with personal public keys

Timelines and milestones

- In the course of 2021 we need to make the Grid Community Toolkit independent of OSG
  - Transfer know-how etc. to partners in the Grid Community Forum
- By Feb 2022, OSG CEs have stopped supporting GSI and X509 proxies
  - HTCondor CEs elsewhere will support GSI and X509 longer
    - HTCondor said at OSG All Hands Meeting they will support GSI as long as GSI library remains supported, but official announcement is pending
- X509 will be supported longer by at least dCache, StoRM, XRootD, EOS
  - They have or will have their own GSI implementations
- Further progress depends on the deployment of services supporting WLCG tokens and the uptake of tokens by the experiments
- Drafts with tentative milestones: transition to tokens, Globus retirement

→ The WLCG Management Board requested an Ops Coordination WG to help coordinate the many parties involved and track the progress

WG for Transition to Tokens and Globus Retirement

- Representatives needed from all parties concerned
  - Grid Community Forum
  - MW product teams
  - Experiments
  - Sites
  - Infrastructures
- Activities may include:
  - Refinement of timelines and milestones
  - Coordination between stakeholders
  - Collaboration with various parties on functionality requirements
  - Deployment campaigns for new MW versions that support tokens and/or depend less on Globus
  - Configuration support for such MW versions
  - Testing campaigns for such MW versions
- Further details on the WG Twiki page