Authentication, now and then (???)

Christophe Haen

10th DIRAC Users Workshop
20/06/24

What are we talking about ?

2 main topics

- VOMS → IAM
- Token for StorageElements

VOMS → IAM

- **VOMS-Admin is going away**
  - Used to manage VO membership
- **VOMS is NOT going away**
  - Delivers proxy certificates with VOMS roles
  - IAM will deliver VOMS certificates
- **Deadline end of June**
  - For CERN managed VOs
  - CC7 EOL, no tested version for RHEL9
  - But there is... thanks to Robert Frank @ Manchester

Main changes

- **Users register in the VO via IAM web interface**
- **Users manage their certificates themselves**
  - But you still have to fix their mess…
- **Need configuration update**
  - Lsc files in your servers and all endpoints
  - VOMSServers in the CS

How do you transition ?

- **Get an IAM instance**
  - I recommend getting the same config as LHCb
  - Ask the CERN IAM team, as it evolves
- **Voms-importer**
  - Periodically sync VOMS content to IAM
- **Check that everything is okay**
  - Last DIRAC v8 version
  - dirac-admin-voms-sync -D -V <vo> -C
  - CompareWithIAM=True in VOMS2CSAgent

How do you transition ? (2)

- **Expect discrepancies !**
  - Different email addresses
  - Local IAM account
  - Etc
- **Fix them !**
- **Sync from IAM**
  - UseIAM = True option in VOMS2CSAgent
- **Turn off VOMS, use IAM web interface**

Good luck

- **Not too difficult but time consuming.**
  - Should be better now
  - Look at VOMS logs ! (some SE may still rely on it)
- **You'll need to have IAM correctly setup in DIRAC**
  - Your client should have "scim:read"
- **The script will help you, but probably a DIRAC migration doc would help… volunteers ?**
- **IAM itself is still not perfect, but work is ongoing**

And now....

HERE'S
WLCG TOKENS !
imgflip.com

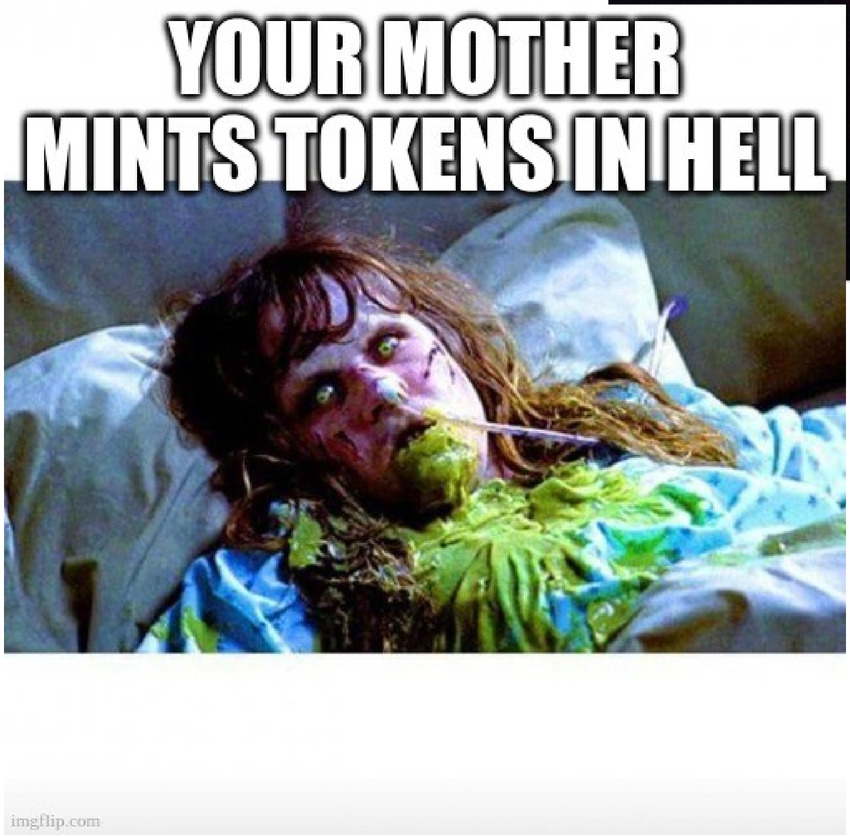
YOUR MOTHER
MINTS TOKENS IN HELL
imgflip.com

I'M YOUR
BOYFRIEND NOW, TOKEN !

Token for Storages: what ?

- Interacting with storage with IAM issued tokens, following the WLCG token profile, instead of VOMS roles.
- VOMS Roles: mapping on the server
  - Role User can write in /pnfs/gridka.de/<vo>/user, read /pnfs/gridka.de/<vo>/
  - Role Prod can read/write in /pnfs/gridka.de/<vo>

WLCG token profile

- 2 flavors: group based or capabilities

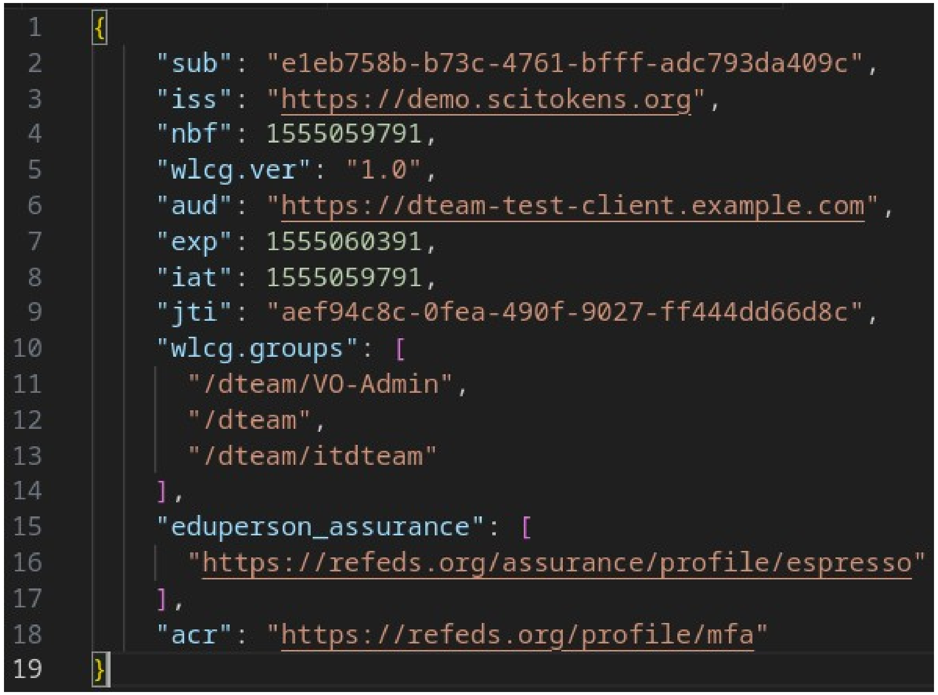

```
{
    "sub": "e1eb758b-b73c-4761-bfff-adc793da409c",
    "iss": "https://demo.scitokens.org",
    "nbf": 1555059791,
    "wlcg.ver": "1.0",
    "aud": "https://dteam-test-client.example.com",
    "exp": 1555060391,
    "iat": 1555059791,
    "jti": "aef94c8c-0fea-490f-9027-ff444dd66d8c",
    "wlcg.groups": [
      "/dteam/VO-Admin",
      "/dteam",
      "/dteam/itdteam"
    ],
    "eduperson_assurance": [
      "https://refeds.org/assurance/profile/espresso"
    ],
    "acr": "https://refeds.org/profile/mfa"
}
```

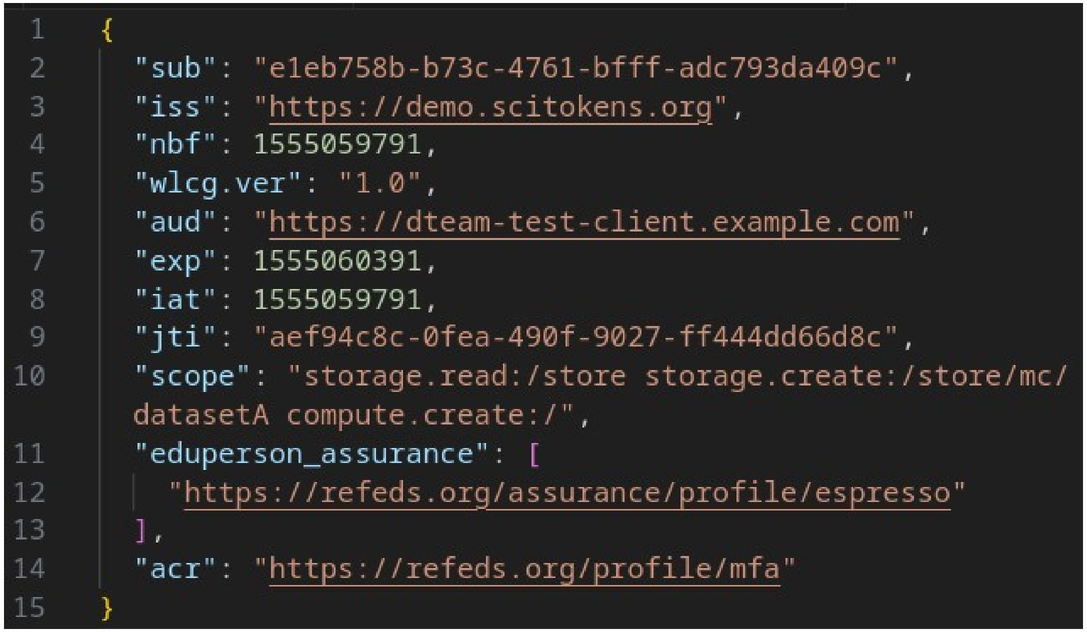

```
{
    "sub": "e1eb758b-b73c-4761-bfff-adc793da409c",
    "iss": "https://demo.scitokens.org",
    "nbf": 1555059791,
    "wlcg.ver": "1.0",
    "aud": "https://dteam-test-client.example.com",
    "exp": 1555060391,
    "iat": 1555059791,
    "jti": "aef94c8c-0fea-490f-9027-ff444dd66d8c",
    "scope": "storage.read:/store storage.create:/store/mc/datasetA compute.create:/",
    "eduperson_assurance": [
      "https://refeds.org/assurance/profile/espresso"
    ],
    "acr": "https://refeds.org/profile/mfa"
}
```

- VOs expressed preference for capabilities

Token for Storages: how ?

What you put in the token: easy

- **The capability**
  - Storage.read, storage.create, storage.modify, storage.stage
- **The target path**
  - storage.read:/home/joe
- **The audience**
  - For whom is the token intended
  - https://lhcbdcacheweb-kit.gridka.de
  - https://wlcg.cern.ch/jwt/v1/any

So let’s do it !

- **Aim: run the DataChallenge24 with tokens**
  - WLCG VO synchronized storage load test
- **Let the “direct storage interaction case” for later**
- **How it is supposed to be done: submit FTS transfers with 3 tokens:**
  - One to interact with FTS servers
  - One to read the source
  - One to write the destination
  -

LHCb Tokens during DC24

- Efficiency of token-based transfers are much lower, compared to certificate-based

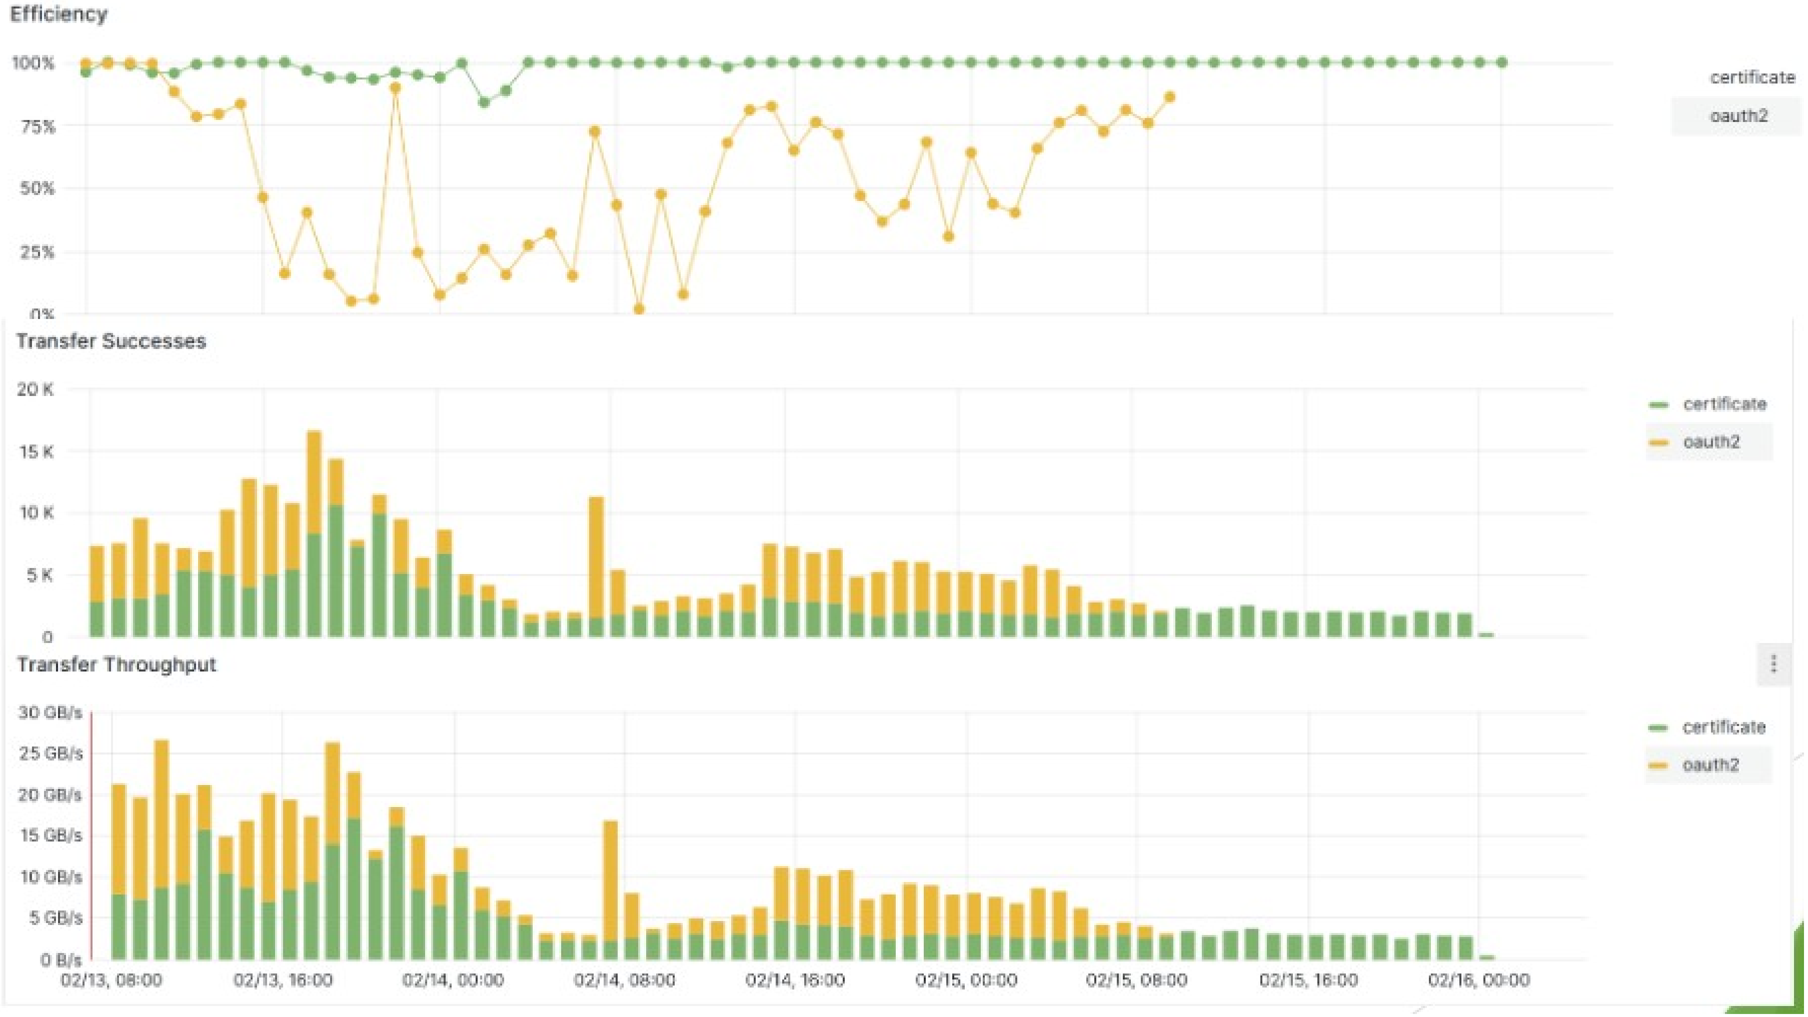

Credit: Alexander Rogovskiy

How we did it for the DC

- With this hotfix
- Ask a token to IAM via the DIRAC TokenManager
  - One for FTS
  - One with “storage.read:/path/to/source/file”
  - One with “storage.modify:/path/to/dest/file storage.read:/path/to/dest/file”

Where things went wrong

- **TokenManager did not scale: fixed**
- **IAM did not scale**
  - We request between 2 and 3 tokens per FTS job
  - FTS refreshes them
  - Up to 30mn to get a token
  - Work ongoing on their side
  - (Rucio caches 1 almighty token per SE, so no issue there)
- **Each storage implementation has a different understanding of the profile**

Profile issues: path

- “path” definition
  - Relative to a “VO path” that would be configured on storages (e.g. /eos/lhcb)
  - Different from SE path (e.g. /eos/lhcb/prod/mc for our MC-SE)
  - LFN /lhcb/MC/2024/hollidays.jpeg on MC-SE → path=/prod/mc/lhcb/MC/2024/hollidays.jpeg
- Assumes namespace uniformity throughout ALL grid SE → non sense
- Impossible to make e.g. a token valid for read everywhere
- Hacked something for DC24
- Proposed alternative solution, uphill battle
- Still convinced we should go for pre-signed URL

Outlook

- WLCG realized (after 6 years) that having 2 DOMA working groups focused on tokens would not cut it
- Proposal: create a task force
  - And I am not even kidding

Tokens after trying it

Questions ?
Comments ?
Sarcasms ?