ecosister

INFN CNAF

Support to experiments in the transition from X.509 authN/Z to SciTokens

International Symposium on Grids & Clouds (ISGC) 2024

Alessandro Pascolini
alessandro.pascolini@cnaf.infn.it

Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

Outline

- What is INFN – CNAF
- What are X509 and SciTokens
- Users AuthN/Z, state of the art
- User Support strategies
- Conclusions

INFN - CNAF

Worldwide LHC Computing Grid (WLCG) [1]

- **~170 computing centres** in more than 40 countries
- Providing **computing resources** to LHC and many other experiments worldwide
- According to their dimensions and resources, all the centers are divided in **Tier-0** (CERN), **Tier-1** and **Tier-2**

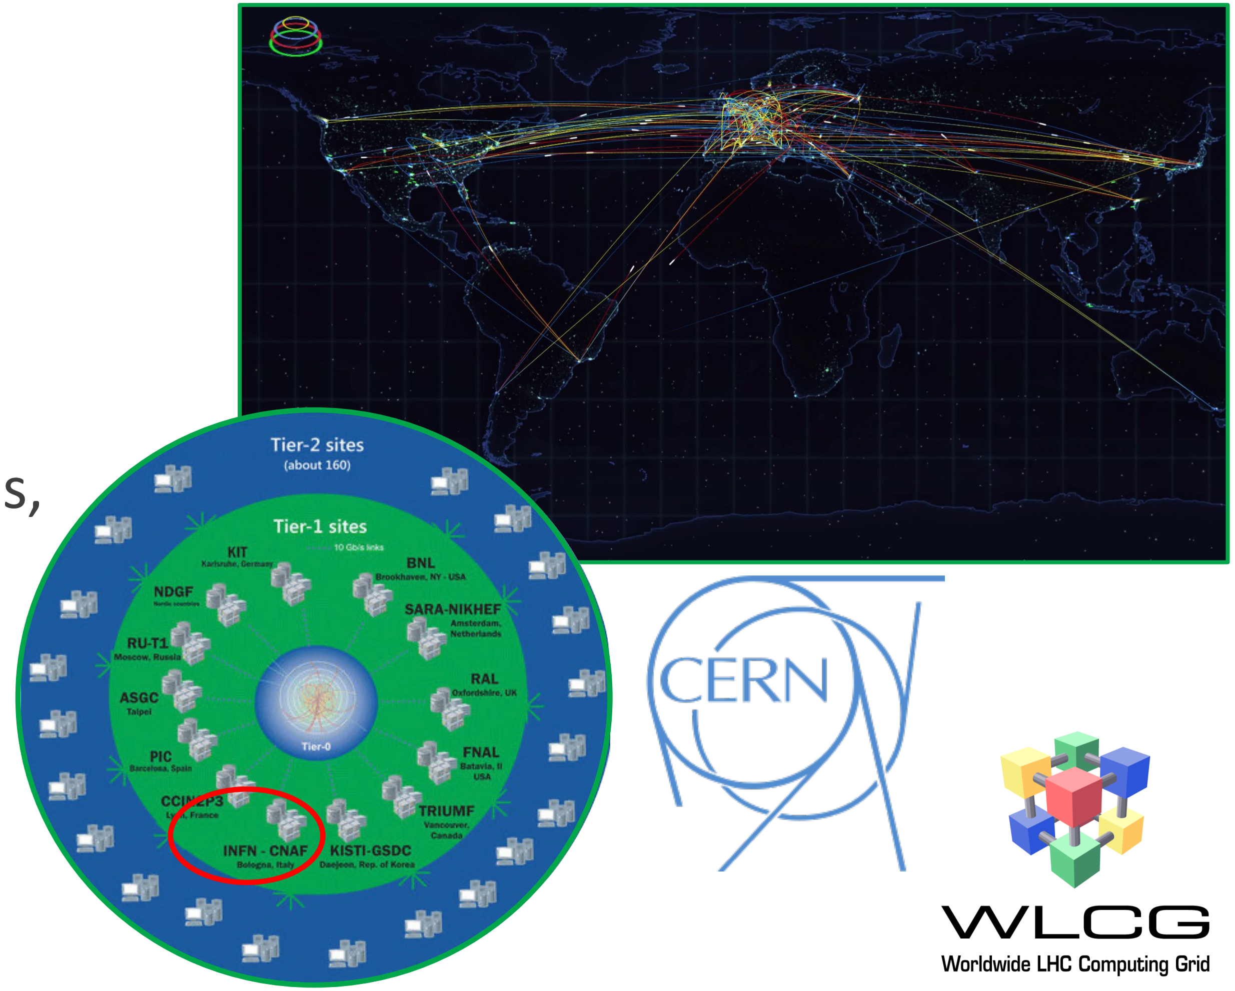

CERN

WLCG
Worldwide LHC Computing Grid

[1] https://wlcg.web.cern.ch/

INFN - CNAF

INFN – CNAF [2] hosts the **Italian Tier-1** since 2003

- Provides resources to more than **60 scientific communities**
  →(**~1500** local users)

- **~2.000 computing nodes**
  → **~60.000 cores** managed by an HTCondor **[3]** cluster
- **~70 PB disk** and **~130 PB tape**

[2] https://www.cnaf.infn.it/
[3] https://www.htcondor.org/

Not only WLCG

Supported scientific communities:

- High-Energy Physics: **8**
- Astroparticle Physics: **18**
- Gravitational Waves: **2**
- Nuclear Physics: **15**
- Dark Matter: **6**
- others: **10**

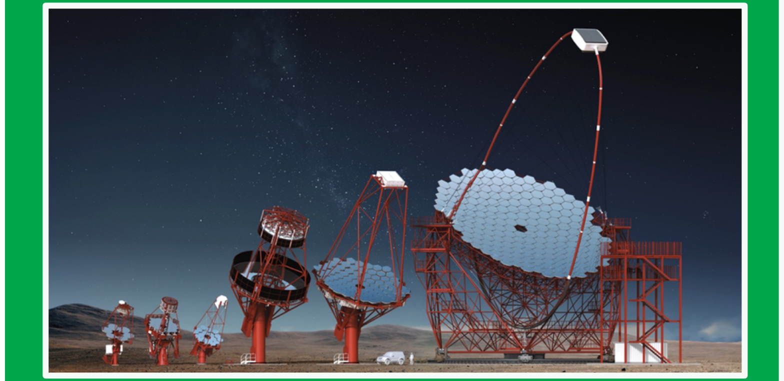

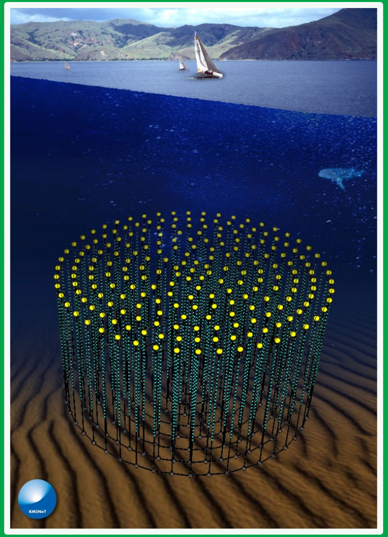

INFN – T1 Internal organization

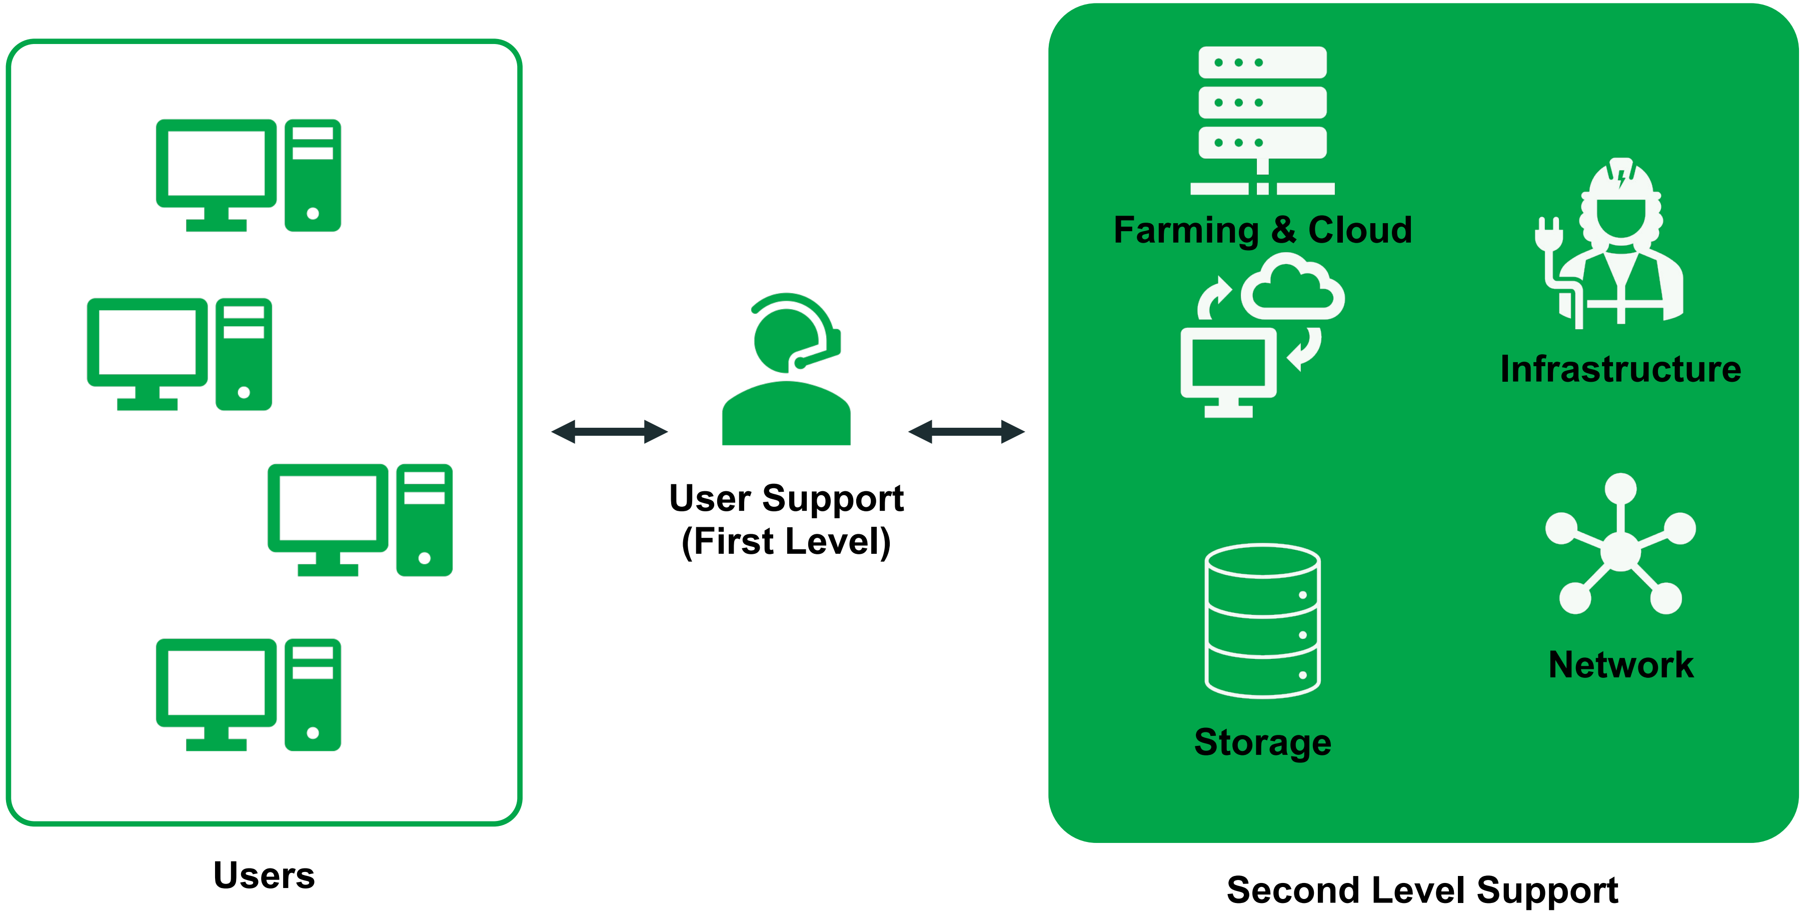

Users

User Support (First Level)

Farming & Cloud

Infrastructure

Storage

Network

Second Level Support

INFN – T1 Internal organization

User Support Team

New Entries !!

X509 GSI AuthN/Z

Pros:

- Widely used method of AuthN/Z
- E.g.: HTTPS protocol

Cons:

- Need for custom solutions to be integrated in other services
- No Fine-Grained AuthZ*
- Proxies last up to several days*

* **Security Issues!!**

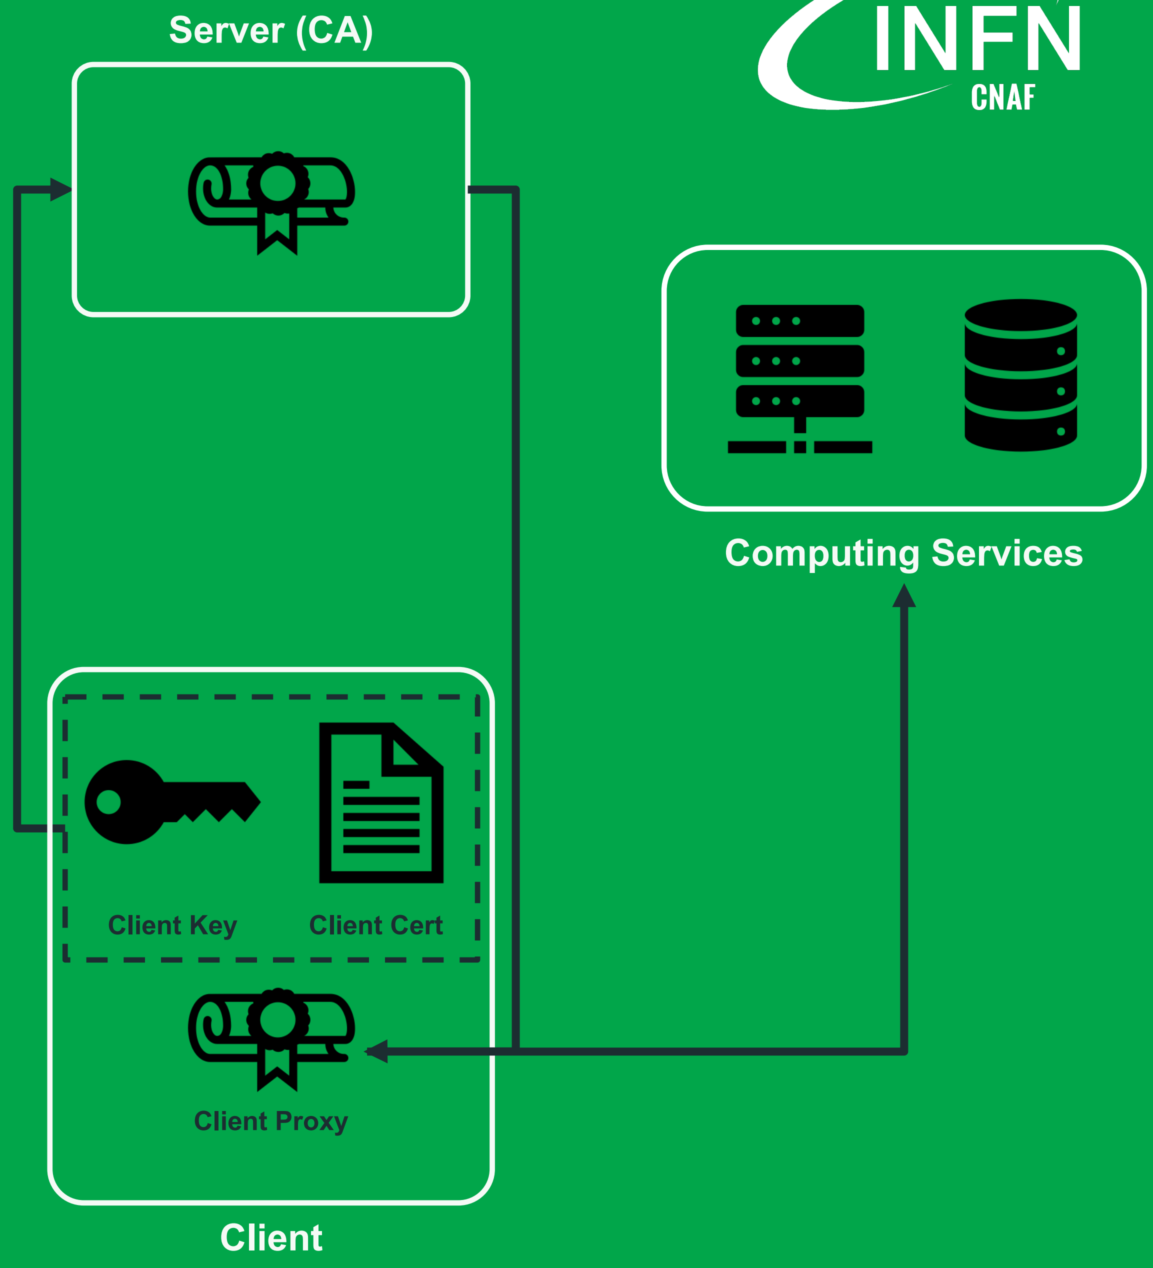

SciTokens

Pros:

- Based on **JWT technology**, widely integrated with other services' workflow
- Short-lived
- Provide **fine-grained** authZ based on **scopes** or **groups**

Cons:

- **Short-lived***
- oidc-agent is restricted to the user machine and **can't be forwarded**

*an issue as well, in some cases

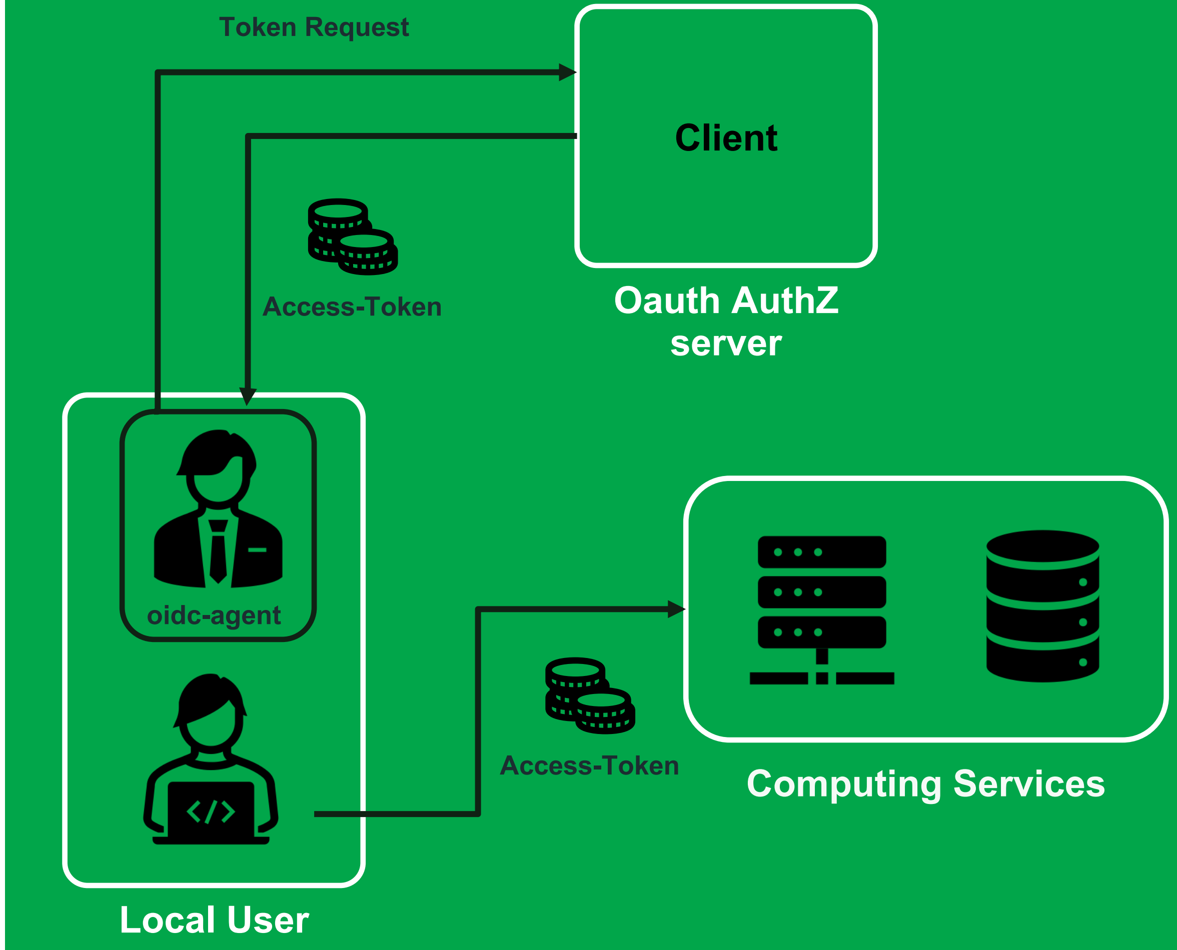

Users Toolkit – State of the art

Tokens Management

oidc-agent*

mytoken

htgettoken

Storage Access

StoRM*

StoRM WebDAV*

XrootD*

Job Submission

HTCondor-CE*

ARC-CE

* Currently supported at CNAF

Issues with tokens

- OAuth flow is quite complex
- → need to provide **simplified guides** to users
  → spot on assistance via **email** or **in person meeting**

- **Long lasting** job can't use a **short lived** token created during submission!
  → need a way to get fresh tokens during job execution!

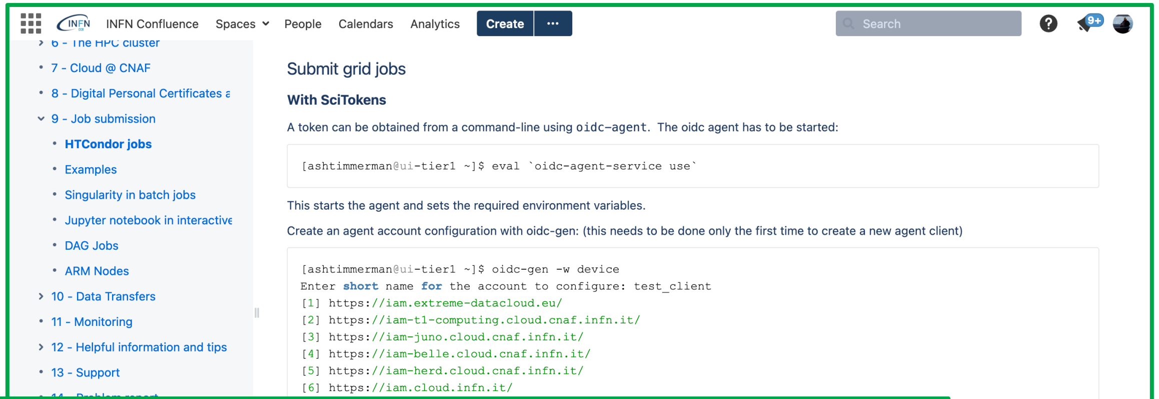

Submit grid jobs

With SciTokens

A token can be obtained from a command-line using `oidc-agent`. The oidc agent has to be started:

```
[ashtimmerman@ui-tier1 ~]$ eval `oidc-agent-service use`
```

This starts the agent and sets the required environment variables.

Create an agent account configuration with oidc-gen: (this needs to be done only the first time to create a new agent client)

```
[ashtimmerman@ui-tier1 ~]$ oidc-gen -w device
Enter short name for the account to configure: test_client
[1] https://iam.extreme-datacloud.eu/
[2] https://iam-t1-computing.cloud.cnaf.infn.it/
[3] https://iam-juno.cloud.cnaf.infn.it/
[4] https://iam-belle.cloud.cnaf.infn.it/
[5] https://iam-herd.cloud.cnaf.infn.it/
[6] https://iam.cloud.infn.it/
```

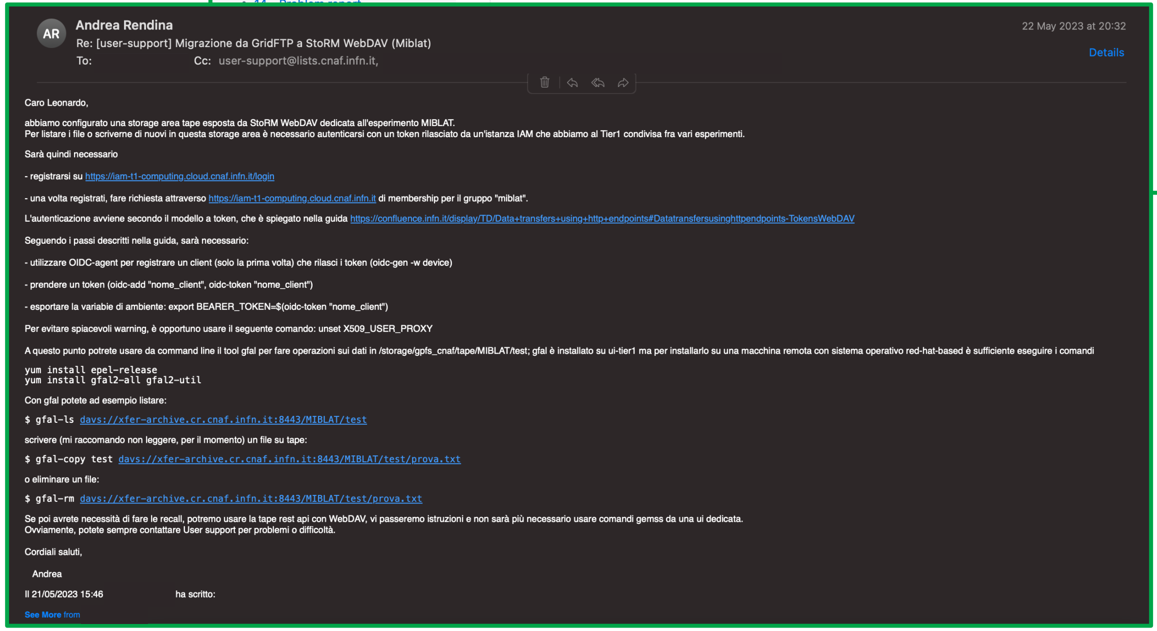

Andrea Rendina
Re: [user-support] Migrazione da GridFTP a StoRM WebDAV (Miblat)
To: [redacted] Cc: user-support@lists.cnaf.infn.it,

22 May 2023 at 20:32

Caro Leonardo,

abbiamo configurato una storage area tape esposta da StoRM WebDAV dedicata all'esperimento MIBLAT.
Per listare i file o scriverne di nuovi in questa storage area è necessario autenticarsi con un token rilasciato da un'istanza IAM che abbiamo al Tier1 condivisa fra vari esperimenti.

Sarà quindi necessario

- registrarsi su https://iam-t1-computing.cloud.cnaf.infn.it/login

- una volta registrati, fare richiesta attraverso https://iam-t1-computing.cloud.cnaf.infn.it di membership per il gruppo "miblat".

L'autenticazione avviene secondo il modello a token, che è spiegato nella guida https://confluence.infn.it/display/TD/Data+transfers+using+http+endpoints#Datatransfersusinghttpendpoints-TokensWebDAV

Seguendo i passi descritti nella guida, sarà necessario:

- utilizzare OIDC-agent per registrare un client (solo la prima volta) che rilasci i token (oidc-gen -w device)

- prendere un token (oidc-add "nome_client", oidc-token "nome_client")

- esportare la variabie di ambiente: export BEARER_TOKEN=$(oidc-token "nome_client")

Per evitare spiacevoli warning, è opportuno usare il seguente comando: unset X509_USER_PROXY

A questo punto potrete usare da command line il tool gfal per fare operazioni sui dati in /storage/gpfs_cnaf/tape/MIBLAT/test; gfal è installato su ui-tier1 ma per installarlo su una macchina remota con sistema operativo red-hat-based è sufficiente eseguire i comandi

```
yum install epel-release
yum install gfal2-all gfal2-util
```

Con gfal potete ad esempio listare:

```
$ gfal-ls davs://xfer-archive.cr.cnaf.infn.it:8443/MIBLAT/test
```

scrivere (mi raccomando non leggere, per il momento) un file su tape:

```
$ gfal-copy test davs://xfer-archive.cr.cnaf.infn.it:8443/MIBLAT/test/prova.txt
```

o eliminare un file:

```
$ gfal-rm davs://xfer-archive.cr.cnaf.infn.it:8443/MIBLAT/test/prova.txt
```

Se poi avrete necessità di fare le recall, potremo usare la tape rest api con WebDAV, vi passeremo istruzioni e non sarà più necessario usare comandi gemss da una ui dedicata.
Ovviamente, potete sempre contattare User support per problemi o difficoltà.

Cordiali saluti,

Andrea

Il 21/05/2023 15:46 [redacted] ha scritto:

See More from [redacted]

Issues with tokens

- **Long lasting** job can't use a **short-lived** token created during submission!
  → need a way to get fresh tokens during job execution!

Possible Solutions*

- DIY solution
- **mytoken [4]**
- htgettoken

*Marteen Litmaath's talk during GDB session on 27th March **[5]**

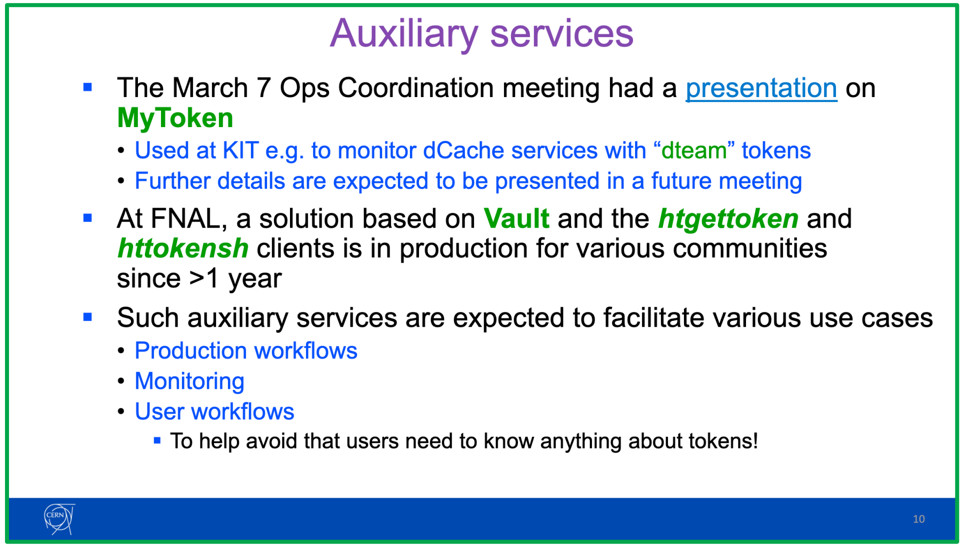

Auxiliary services

- The March 7 Ops Coordination meeting had a presentation on **MyToken**
  - Used at KIT e.g. to monitor dCache services with "dteam" tokens
  - Further details are expected to be presented in a future meeting
- At FNAL, a solution based on **Vault** and the ***htgettoken*** and ***httokensh*** clients is in production for various communities since >1 year
- Such auxiliary services are expected to facilitate various use cases
  - Production workflows
  - Monitoring
  - User workflows
    - To help avoid that users need to know anything about tokens!

[4] https://mytoken-docs.data.kit.edu/

[5] Token-Transition-update-240327

mytoken workflow

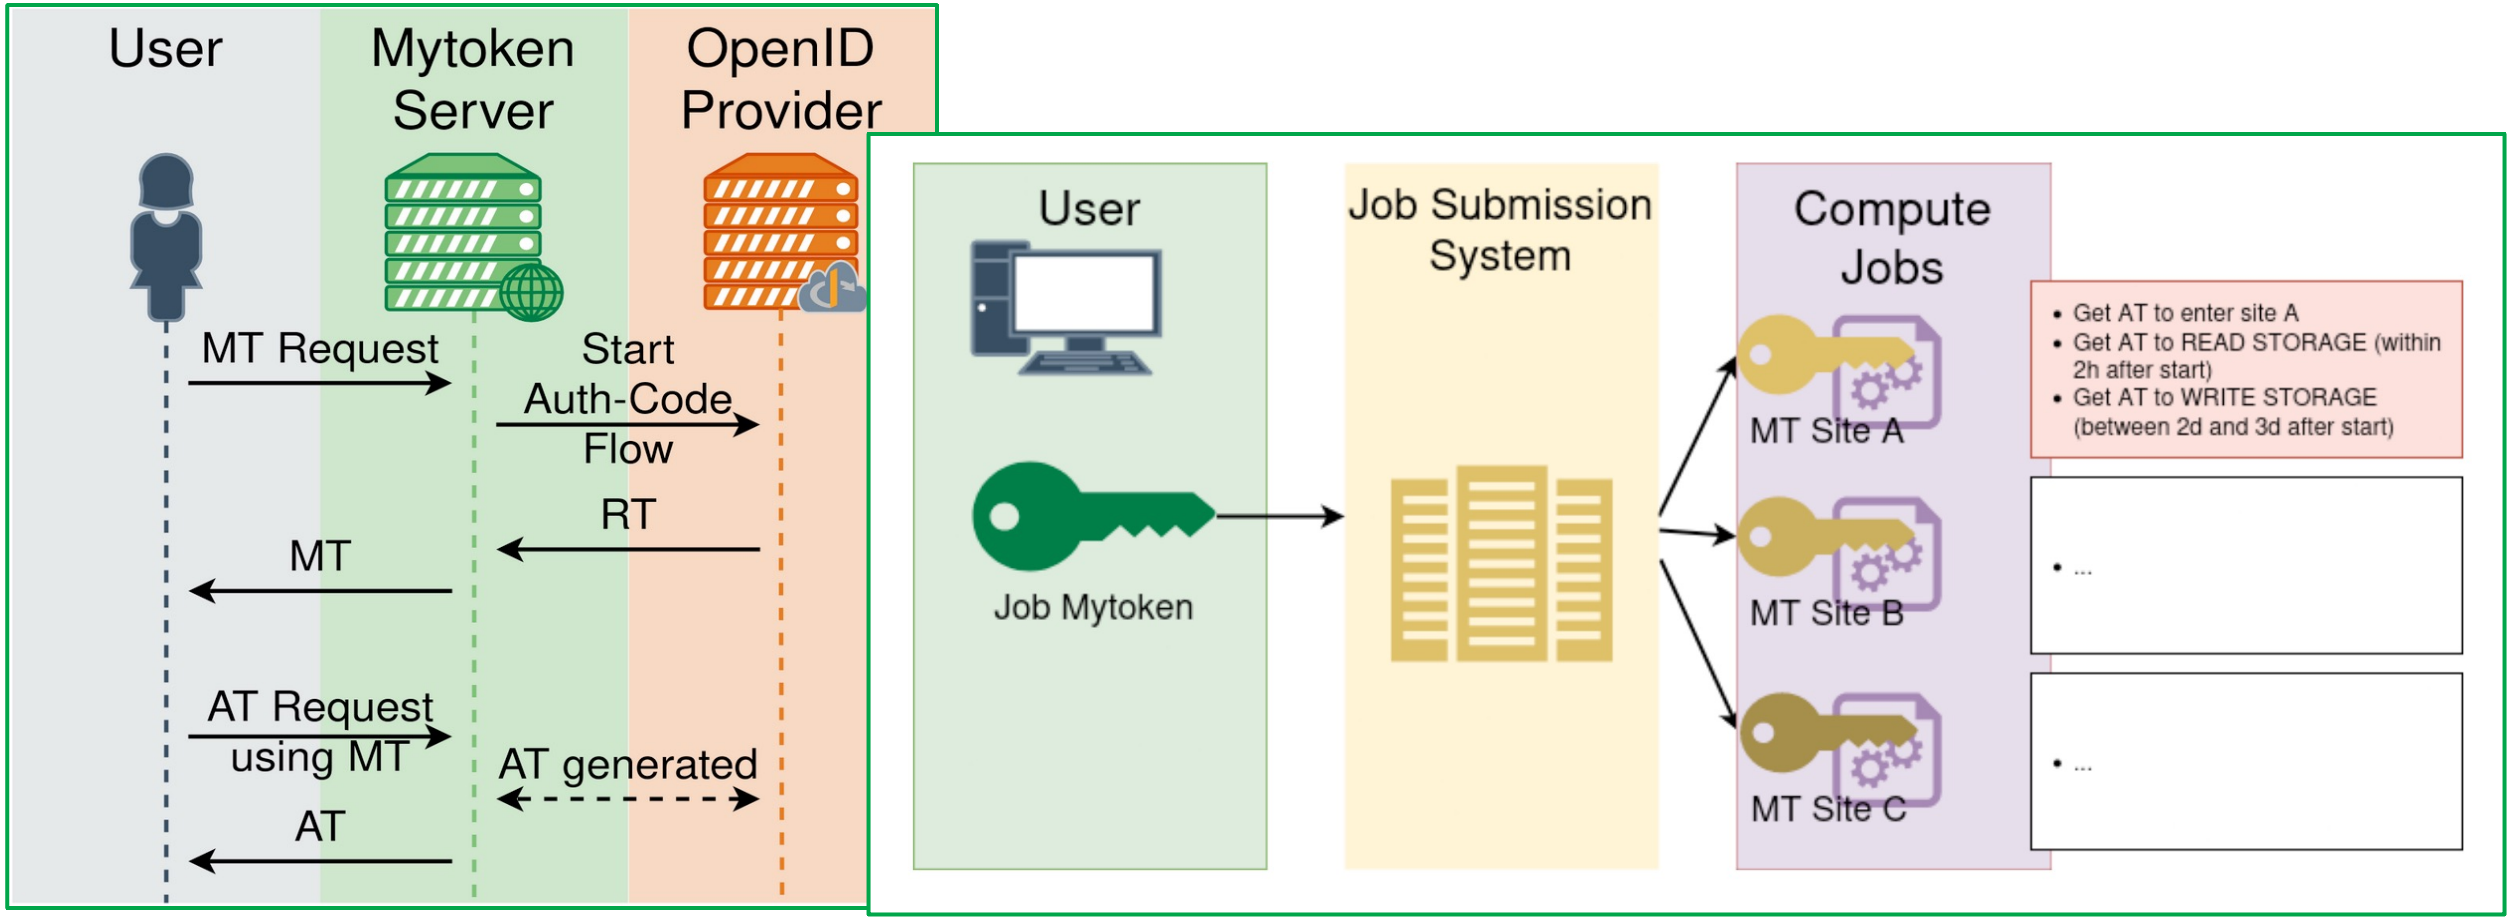

- Get AT to enter site A
- Get AT to READ STORAGE (within 2h after start)
- Get AT to WRITE STORAGE (between 2d and 3d after start)

- ...

- ...

https://cvs.data.kit.edu/talks/2403-mytoken-wlcg-ops/
https://cvs.data.kit.edu/talks/2306-mytoken-egi/

What is a mytoken?

- Extension on the concept of Refresh-Token
- JWT based
- Implements new features:

  →Rotation

  →Restrictions

  - **how much time** it lasts
  - from **which hosts**/country
  - how many times can be used

  →Capabilities

  - **AT** (get access-tokens)
  - tokeninfo (retrieve information about mytoken)

  → Profiles (includes the previous)
- It can be used like a refresh-token to connect to the oidc-agent on the **mytoken-server**

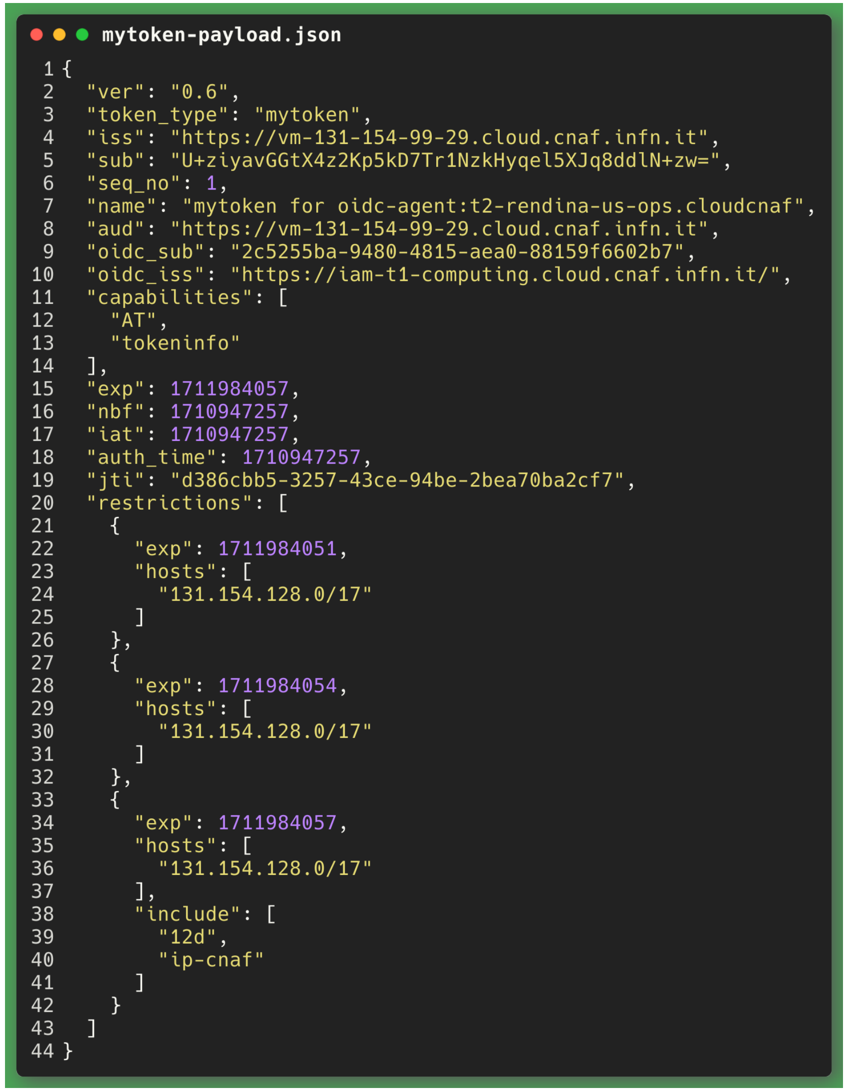

mytoken-payload.json

```json
{
  "ver": "0.6",
  "token_type": "mytoken",
  "iss": "https://vm-131-154-99-29.cloud.cnaf.infn.it",
  "sub": "U+ziyavGGtX4z2Kp5kD7Tr1NzkHyqel5XJq8ddlN+zw=",
  "seq_no": 1,
  "name": "mytoken for oidc-agent:t2-rendina-us-ops.cloudcnaf",
  "aud": "https://vm-131-154-99-29.cloud.cnaf.infn.it",
  "oidc_sub": "2c5255ba-9480-4815-aea0-88159f6602b7",
  "oidc_iss": "https://iam-t1-computing.cloud.cnaf.infn.it/",
  "capabilities": [
    "AT",
    "tokeninfo"
  ],
  "exp": 1711984057,
  "nbf": 1710947257,
  "iat": 1710947257,
  "auth_time": 1710947257,
  "jti": "d386cbb5-3257-43ce-94be-2bea70ba2cf7",
  "restrictions": [
    {
      "exp": 1711984051,
      "hosts": [
        "131.154.128.0/17"
      ]
    },
    {
      "exp": 1711984054,
      "hosts": [
        "131.154.128.0/17"
      ]
    },
    {
      "exp": 1711984057,
      "hosts": [
        "131.154.128.0/17"
      ],
      "include": [
        "12d",
        "ip-cnaf"
      ]
    }
  ]
}
```

Why mytokens?

Pros:

- Really close to our idea of resolving the issue
- Easy to configure
  → YAML config file
  → few steps to follow
  → utility scripts to configure server features
- Customizable
- Responsive developers to help troubleshooting
- Integrated with OIDC flow

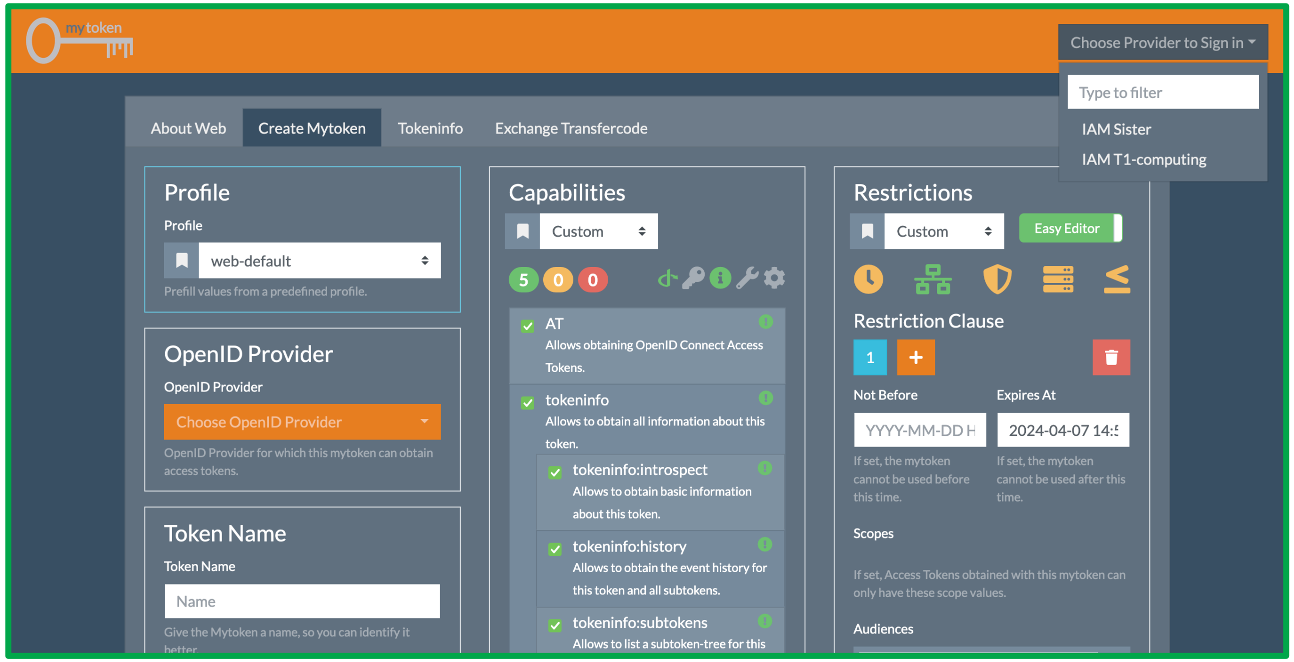

Profile

Profile: web-default

Prefill values from a predefined profile.

OpenID Provider

OpenID Provider: Choose OpenID Provider

OpenID Provider for which this mytoken can obtain access tokens.

Token Name

Token Name: Name

Give the Mytoken a name, so you can identify it better

Capabilities

Custom

☑ AT
Allows obtaining OpenID Connect Access Tokens.

☑ tokeninfo
Allows to obtain all information about this token.

☑ tokeninfo:introspect
Allows to obtain basic information about this token.

☑ tokeninfo:history
Allows to obtain the event history for this token and all subtokens.

☑ tokeninfo:subtokens
Allows to list a subtoken-tree for this

Restrictions

Custom

Restriction Clause

Not Before: YYYY-MM-DD H

If set, the mytoken cannot be used before this time.

Expires At: 2024-04-07 14:5

If set, the mytoken cannot be used after this time.

Scopes

If set, Access Tokens obtained with this mytoken can only have these scope values.

Audiences

Current activity

- Deployment of a self hosted my**token-server**
  → server configuration
  → setup **CNAF profile**
  (rotations, restrictions, capabilities, templates)
  → connect to Tier-1 **IAM instance**

- Test phase on how to get and use my**tokens**
  → client choice (my**token-client**, **oidc-agent**)
  → submission tests to manage files with AT requested via my**token** flow

Future actions

- Keep on testing the **mytoken** solution
  → scalability tests
  → security evaluation
  → understanding of users needs to implement new profiles

- Provide **mytoken** solution to INFN Tier-1 users in order to dismiss POSIX access to data from worker-nodes

Conclusions

User-support challenges in the future:

- support all collaboration and users transitioning from **X509** to **SciTokens**
- Keep updated guides on how to implement tokens into users' workflow
- Test new solutions to ease token usage

Acknowledgments

- Many thanks to **Gabriel Zackmann**
  → **main contributor** to mytoken project
  → **very helpful** on solving configuration issues we had

- How to deploy mytoken-server
  https://mytoken-docs.data.kit.edu/server
- How to configure mytoken-server
  https://mytoken-docs.data.kit.edu/dev
- mytoken-server git repo
  https://github.com/oidc-mytoken/server
- “mytoken - OpenID Connect Tokens for Long-term Authorization” G. Zackmann (Phd Thesis)

INFN
CNAF

Thank you!