Norfolk, Virginia, USA • May 8-12, 2023
CHEP 2023
Computing in High Energy & Nuclear Physics

# Adoption of a token-based authentication model for the CMS Submission Infrastructure

Marco Mascheroni[1], Jeff Dost[1], Saqib Haleem[2], A. Pérez-Calero Yzquierdo[3], Edita Kizinevic[4], Farrukh Aftab Khan[5], Hyunwoo Kim[5] and Nikos Tsipinakis[4] on behalf of the CMS Collaboration

1. Univ. of California San Diego (US), 2. National Center for Physics (PK),3. CIEMAT and PIC (ES), 4. CERN, 5. Fermi National Accelerator Lab. (US)

**CHEP 2023, May 11th 2023**

# Outline of the talk

- Why tokens?
- Token-based authentication in CMS Submission Infrastructure
- Current deployment status:
    - HTCondor CEs
    - ARC CEs
- Conclusions and Next steps

# The CMS Submission Infrastructure

- **The CMS Submission Infrastructure:** team in CMS Offline and Computing in **charge** of:
  - Organizing HTCondor Software Suite (**HTCSS**) and **GlideinWMS** operations in CMS,
  - Maintaining a **Global Pool, an infrastructure of distributed compute resources where reconstruction, simulation, and analysis of physics data takes place**
  - Communicate CMS **priorities to the development teams** of glideinWMS and HTCondor
- **In practice:**
  - We operate a set of federated HTCSS pools which aggregate **resources from 70 Grid sites, plus non-Grid resources**
  - We regularly hold **meetings with** HTCSS and glideinWMS **developers** where we discuss current **operational limitations, new feature requests and future scale requirements**

# The CMS Submission Infrastructure

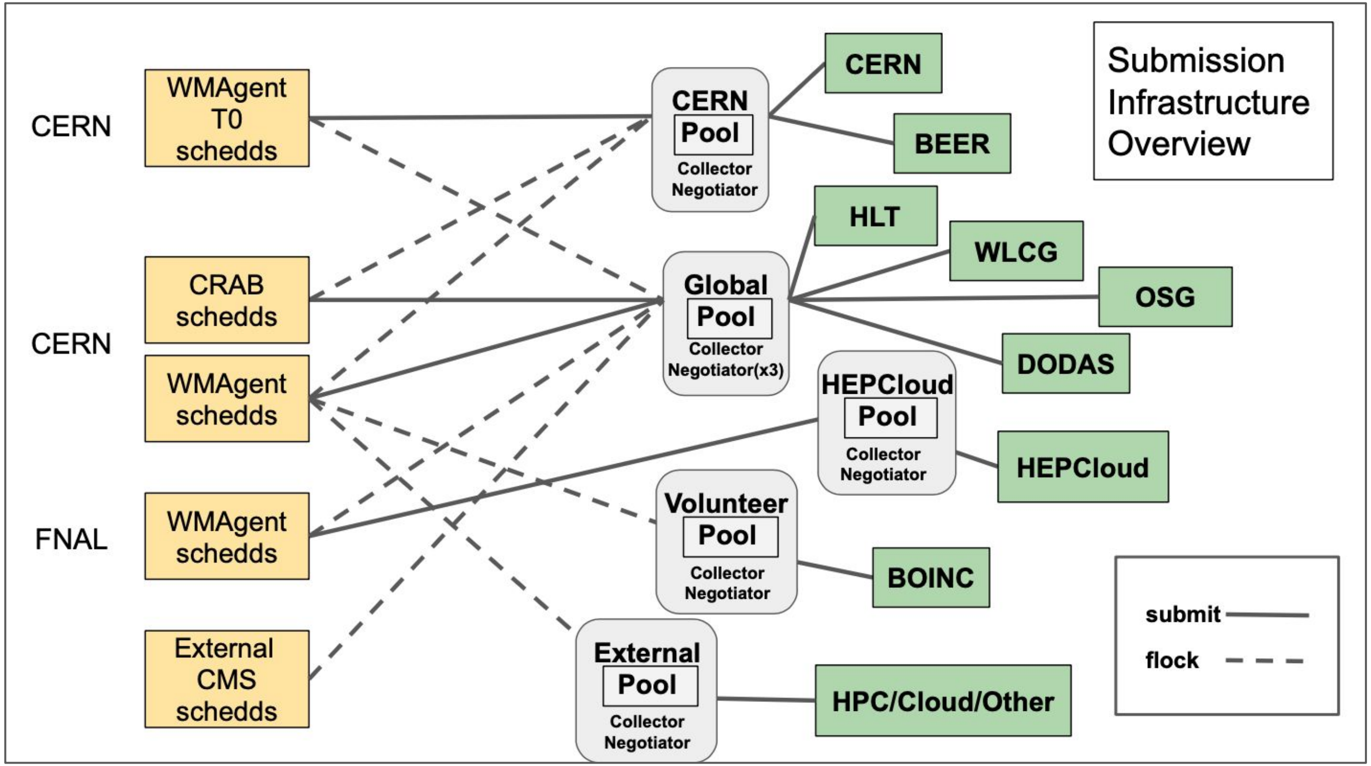

# Overall SI architecture

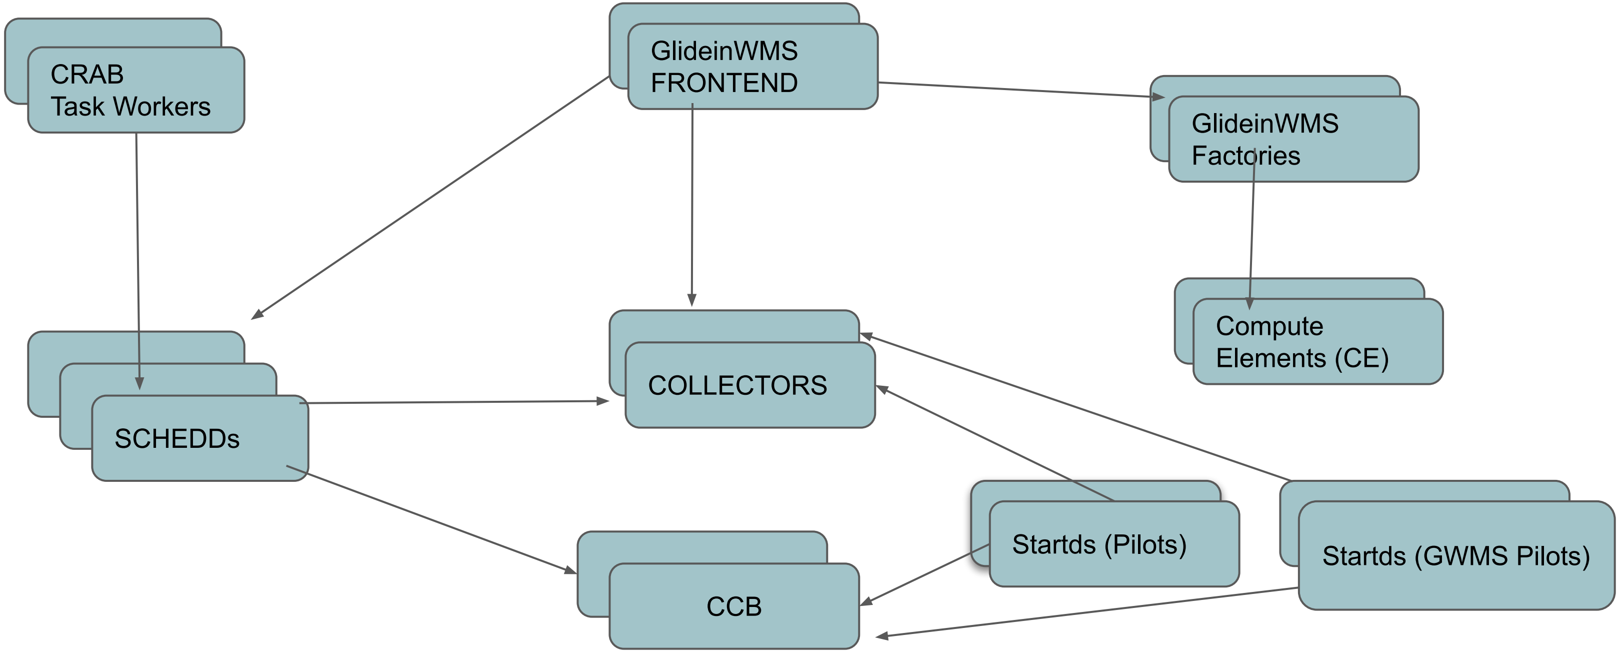

# Why tokens?

Motivations:

- Use de facto **standard** like OAuth **in place of Grid specific** ones (GSI)
    - Part of a general **WLCG** campaign

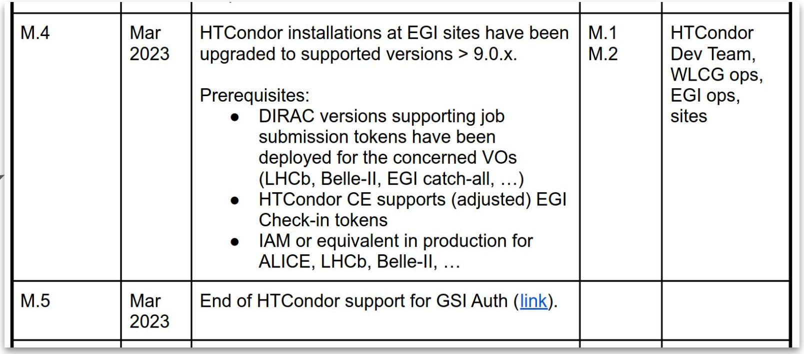

| | | | | |
|---|---|---|---|---|
| M.4 | Mar 2023 | HTCondor installations at EGI sites have been upgraded to supported versions > 9.0.x.<br>Prerequisites:<br>● DIRAC versions supporting job submission tokens have been deployed for the concerned VOs (LHCb, Belle-II, EGI catch-all, …)<br>● HTCondor CE supports (adjusted) EGI Check-in tokens<br>● IAM or equivalent in production for ALICE, LHCb, Belle-II, … | M.1<br>M.2 | HTCondor Dev Team, WLCG ops, EGI ops, sites |
| M.5 | Mar 2023 | End of HTCondor support for GSI Auth (link). | | |

Advantages:

- Authorization based on token **capability** (vs identity) allows for finer handling of communications between entities

- Easier to follow **principle of least privilege**
    - **Limited impact** in case of credential leak, for example
    - Only affect Factory<=>CE communication
        - Old pilot proxies were used in more places, e.g.: WN<=>Central Manager

# Two different technologies

- IDTOKENS for HTCondor Software Suite components
  - New **native** condor authentication mechanism
  - JWT tokens signed with a **simmeric key**
  - Tokens generated with condor **command line interface** (even remotely)
- SCITOKENS for Grid sites (HTCondor-CEs + ARC-CEs)
  - Uses WLCG IAM go generate tokens
    - A client is registered and can fetch tokens on the GWMS Frontend
  - JWT tokens signed using **private/public schema**
  - Compute Elements (CEs) verify authenticity of tokens and authorize pilot job submission

# Internal CMS Submission Infrastructure: token authentication with IDTOKENS

# WM to SI schedulers authentication

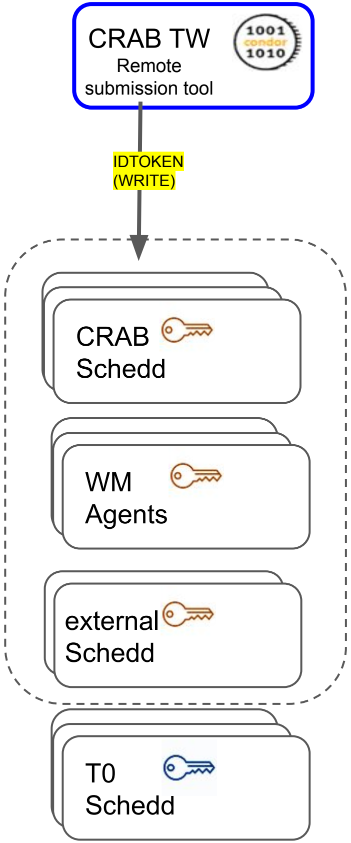

- Authentication of **Workload management tools** to schedulers switched almost **transparently**
  - **Local** authentication
- **User's proxy** used for **remote** authentication (CRAB)
  - Switched to single **service IDTOKEN**
  - Development was necessary in CRAB to **switch from multiple users** mapped with Argus to **single user**

# Condor components

- The whole Global pool is using IDTOKENS for authentication
  - **HTCSS** version 10 with **no GSI fallback** installed early April
- Multiple keys and tokens to **minimize impacts** in case of security incidents

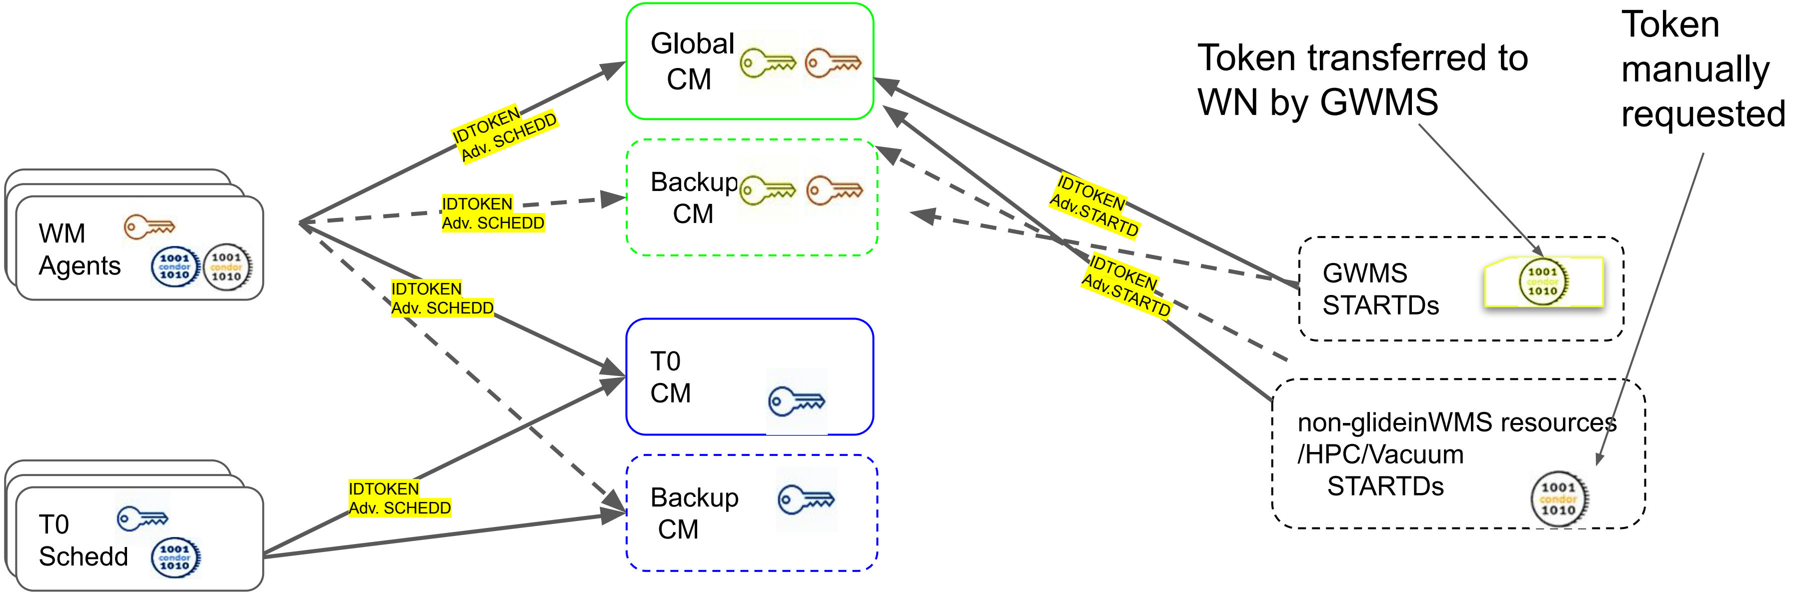

# GlideinWMS authentication

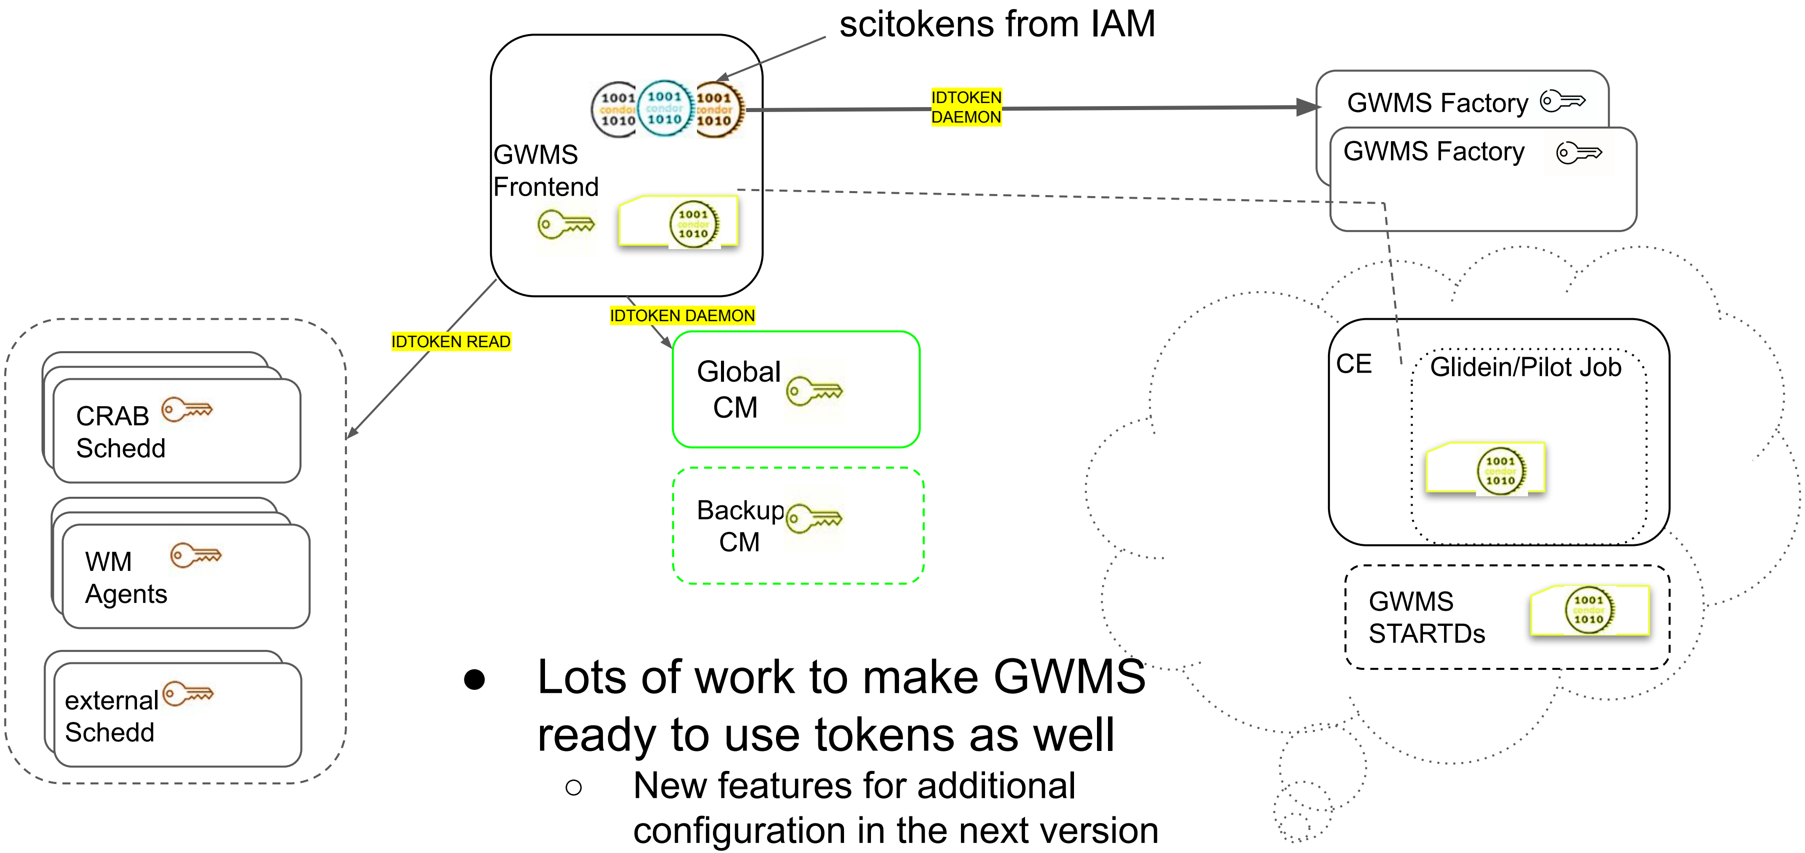

- Lots of work to make GWMS ready to use tokens as well
    - New features for additional configuration in the next version

# Overall authentication in SI

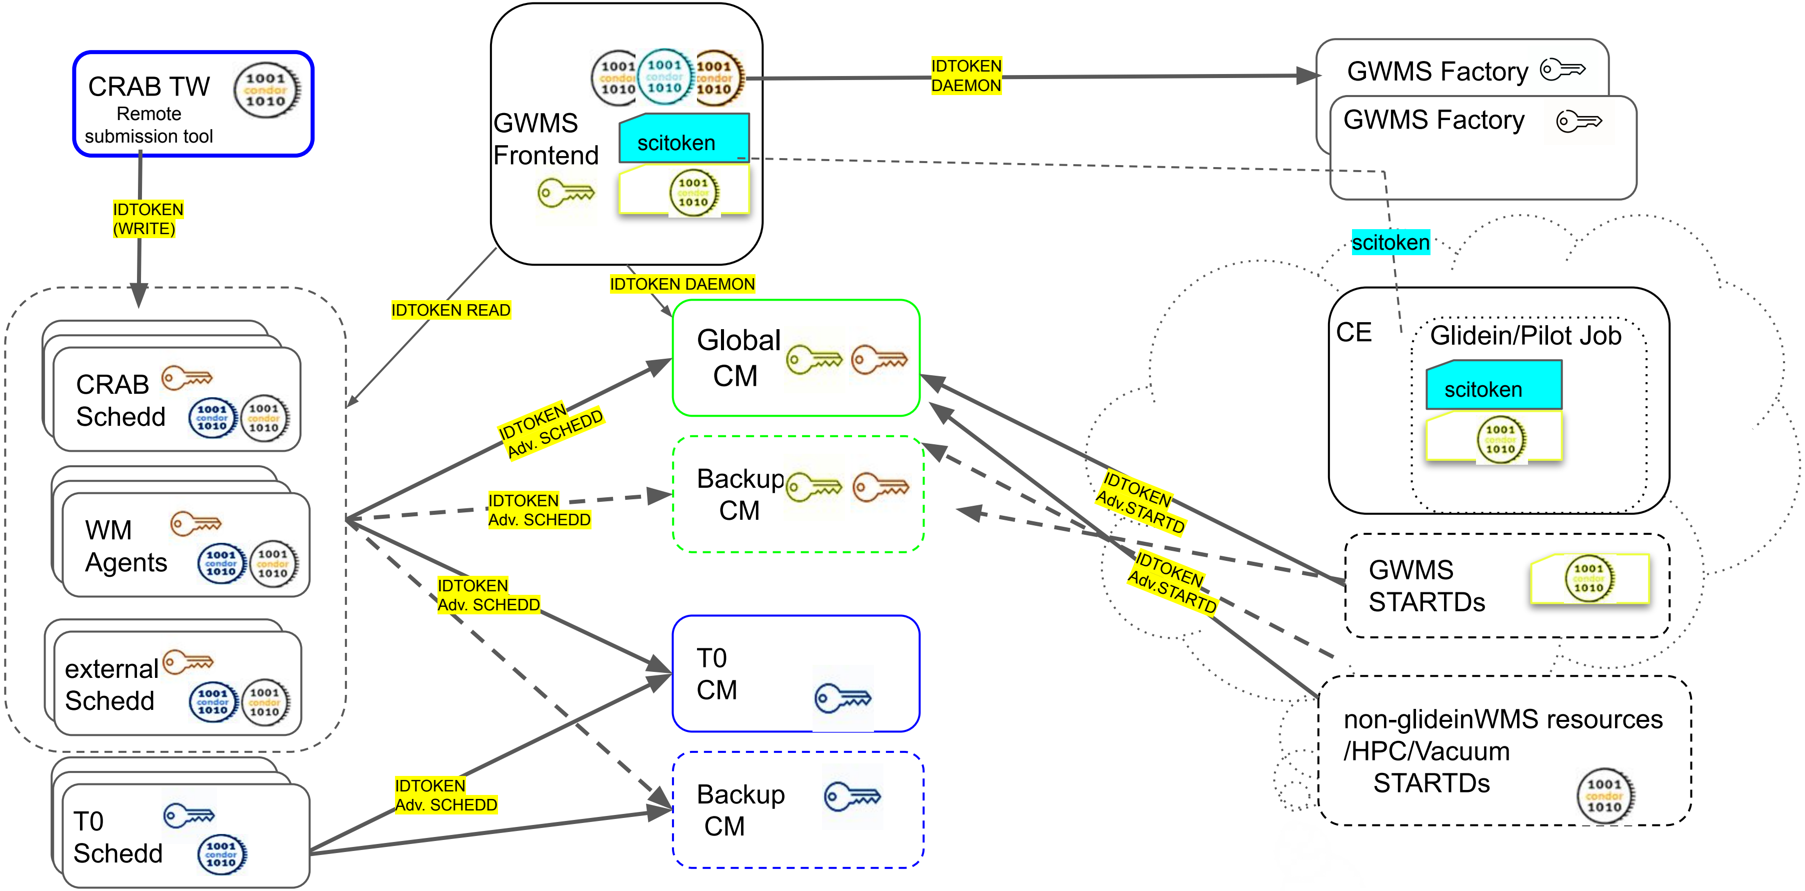

CRAB TW
Remote submission tool

IDTOKEN (WRITE)

GWMS Frontend

scitoken

IDTOKEN DAEMON

GWMS Factory

GWMS Factory

scitoken

IDTOKEN READ

IDTOKEN DAEMON

CRAB Schedd

WM Agents

external Schedd

T0 Schedd

Global CM

IDTOKEN Adv. SCHEDD

Backup CM

IDTOKEN Adv. SCHEDD

IDTOKEN Adv. SCHEDD

T0 CM

IDTOKEN Adv. SCHEDD

Backup CM

CE

Glidein/Pilot Job

scitoken

IDTOKEN Adv.STARTD

IDTOKEN Adv.STARTD

GWMS STARTDs

non-glideinWMS resources /HPC/Vacuum STARTDs

# Token authentication to sites with SCITOKENS

# Ways to submit jobs to the Grid

- Factory submits pilots using HTCondor
    - Sites use two CE technologies
- HTCondor submission
    - Completed 100% (both OSG and EGI sites, last site January)
- ARC-CEs submission:
    - Through LDAP interface and GSI proxies
        - Deprecated in HTCondor 10 series
    - New REST interface available
        - Support for **tokens**
        - Plain X509 proxies can be used (no GSI)

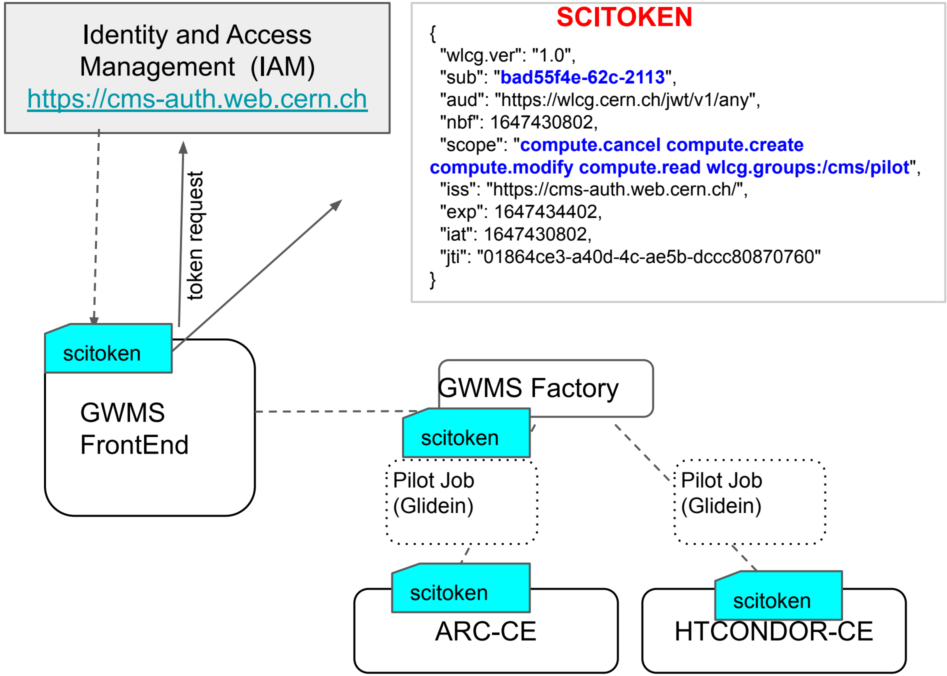

**SCITOKEN**

```
{
  "wlcg.ver": "1.0",
  "sub": "bad55f4e-62c-2113",
  "aud": "https://wlcg.cern.ch/jwt/v1/any",
  "nbf": 1647430802,
  "scope": "compute.cancel compute.create
compute.modify compute.read wlcg.groups:/cms/pilot",
  "iss": "https://cms-auth.web.cern.ch/",
  "exp": 1647434402,
  "iat": 1647430802,
  "jti": "01864ce3-a40d-4c-ae5b-dccc80870760"
}
```

# ARC CE REST interface

- **Current focus** in factory operations
- **Campaign with sites to enable ARC-CE REST** interface and configure them to **accept tokens**
    - T**ests in ITB** to not affect working sites
    - 12 CEs do not work out of 45 CEs (or 5 sites out of 25)
- Split configuration with **UCSD** factory **using REST** and **CERN using** old **LDAP** interface
    - To guarantee **seamless transition** (no fallback to proxy for pilot submission)
- **Few issues** found with **scale** usage
    - Memory leak of arc_gahp on rhel7
    - Job in run state forever under certain circumstances
    - Submission timeouts due to condor request serialization

# Conclusions

- CMS Submission Infrastructure internal components fully switched to IDTOKENS in Spring 2022
- Communication to HTCondor-CE sites with SCITOKEN to 100% of the sites by Autumn 2022
- Improved secret management
  - GSI secrets needed minimal maintenance as they used one certificate per-host
  - IDTokens need to be securely generated, stored, and distributed to their intended user
    - Introduced usage of Teigi for secret storage

Future Work:

- Complete token transition
  - Disable GSI fallback in the global pool
- Retire ARC-CE LDAP interface
- Security “drill exercises”

# Acknowledgements

Projects FPA2016-80994-C2-1-R, PID2019-110942RB-C21, BES-2017-082665 and PID2020-113807RA-I00 funded by:

MINISTERIO DE CIENCIA E INNOVACIÓN

UNIÓN EUROPEA

FONDO EUROPEO DE DESARROLLO REGIONAL

*"Una manera de hacer Europa"*

AGENCIA ESTATAL DE INVESTIGACIÓN

Ciemat Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas

cfp CIEMAT física de partículas

PIC port d'informació científica

US National Science Foundation
Grant No. 2121686

UC San Diego

# Backup Slides

# Abstract

The CMS Submission Infrastructure (SI) is the main computing resource provisioning system for CMS workloads. A number of HTCondor pools are employed to manage this infrastructure, which aggregates geographically distributed resources from the WLCG and other providers. Historically, the model of authentication among the diverse components of this infrastructure has relied on the Grid Security Infrastructure (GSI), based on identities and X509 certificates. In contrast, commonly used modern authentication standards are based on capabilities and tokens. The WLCG has identified this trend and aims at a transparent replacement of GSI for all its workload management, data transfer and storage access operations, to be completed during the current LHC Run 3. As part of this effort, and within the context of CMS computing, the Submission Infrastructure group is in the process of phasing out the GSI part of its authentication layers, in favor of IDTokens and Scitokens. The use of tokens is already well integrated into the HTCondor Software Suite, which has allowed us to fully migrate the authentication between internal components of SI. Additionally, recent versions of the HTCondor-CE support tokens as well, enabling CMS resource requests to Grid sites employing this CE technology to be granted by means of token exchange. After a rollout campaign to sites, successfully completed by the third quarter of 2022, the totality of HTCondor CEs in use by CMS are already receiving Scitoken-based pilot jobs. On the ARC CE side, a parallel campaign was launched to foster the adoption of the REST interface at CMS sites (required to enable token-based job submission via HTCondor-G), which is nearing completion as well. In this contribution, the newly adopted authentication model will be described. We will then report on the migration status and final steps towards complete GSI phase out in the CMS SI.