Transition to tokens: the CMS experience

Saqib Haleem, Marco Mascheroni,
Antonio Perez-Calero Yzquierdo,Edita Kizinevič
for the CMS Submission Infrastructure team

HTCondor Week 2022

Outline

- The CMS Submission Infrastructure (SI) overview.
- Transition status : from GSI to tokens
  - IDTOKENS
  - SciToken
- Conclusions and next steps

A complex infrastructure

- **The CMS SI model has evolved to running multiple federated pools**, with extensive use of **flocking**

- Multiple sets of specialized workflow managers (CRAB & WMAgent) attached to schedds

- **The main Global Pool:**
  - **Peaks at ~350k CPU cores**
  - **Up to 200k running jobs**
  - 50+ schedds
- Redundant infrastructure for **HA**

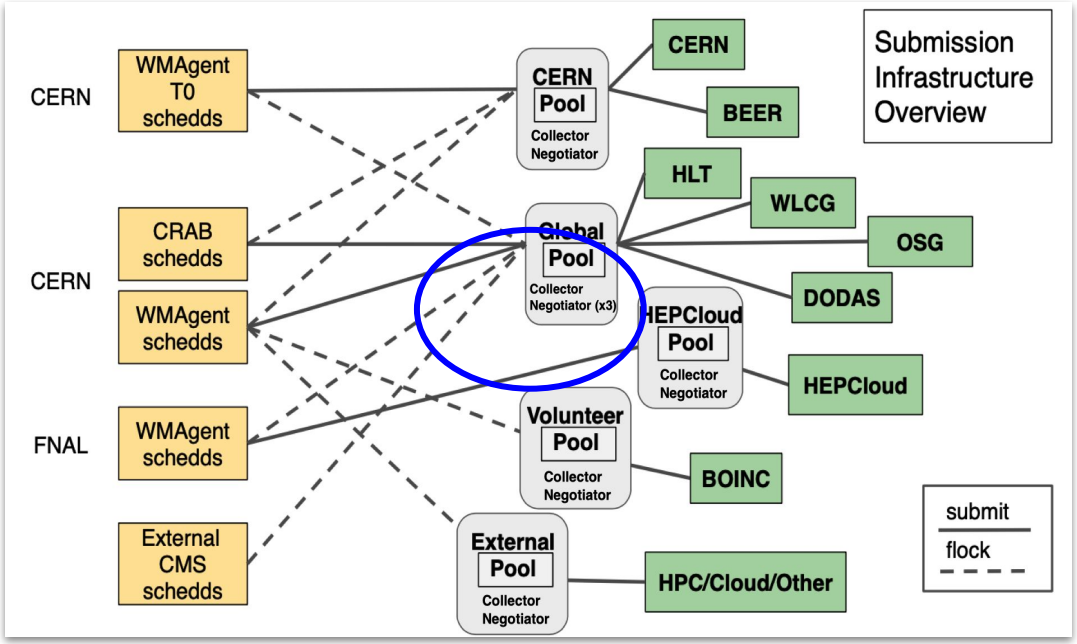

- Resource provisioning mainly with **GlidenWMS pilots** but also **vacuum-like** instantiated: DODAS, BOINC(CMS@Home), opportunistic (HLT), HPC...

Building dynamic HTCondor pools with GlideinWMS

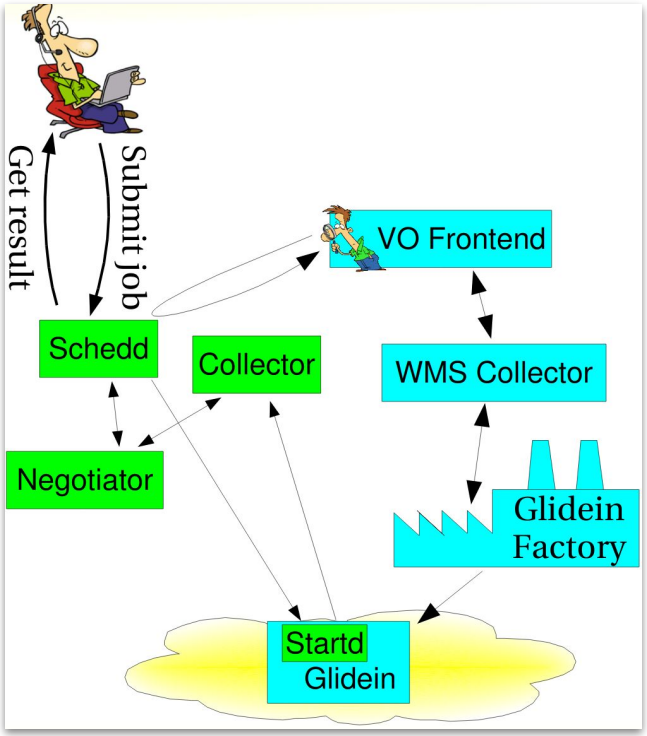

- CMS computing pool is build using two components:
  - **GlideinWMS** : Resource provisioning overlay batch system which grows and shrink based on Job pressure.
  - **HTCondor:** Batch system for Job scheduling.

Moving to tokens in the CMS Submission Infrastructure

- Motivation
  - Towards **industry** standards : Capabilities based authorization for distributed services (new)
  - Globus Toolkit **retirement**
  - Practical example: *multiple tokens with different capabilities instead of a single identity i.e. powerful pilot (GSI) proxy*
- Timeline
  - In coordination with WLCG/OSG timeline
  - OSG 3.6 release removed Globus toolkit dependency
  - **November 2022**: HTCondor GSI End Of Life

Components of the CMS Global Pool

- Authentication between SI Internal Components (IDTOKENS)
- Authentication between Factories <-> Sites (SciToken)

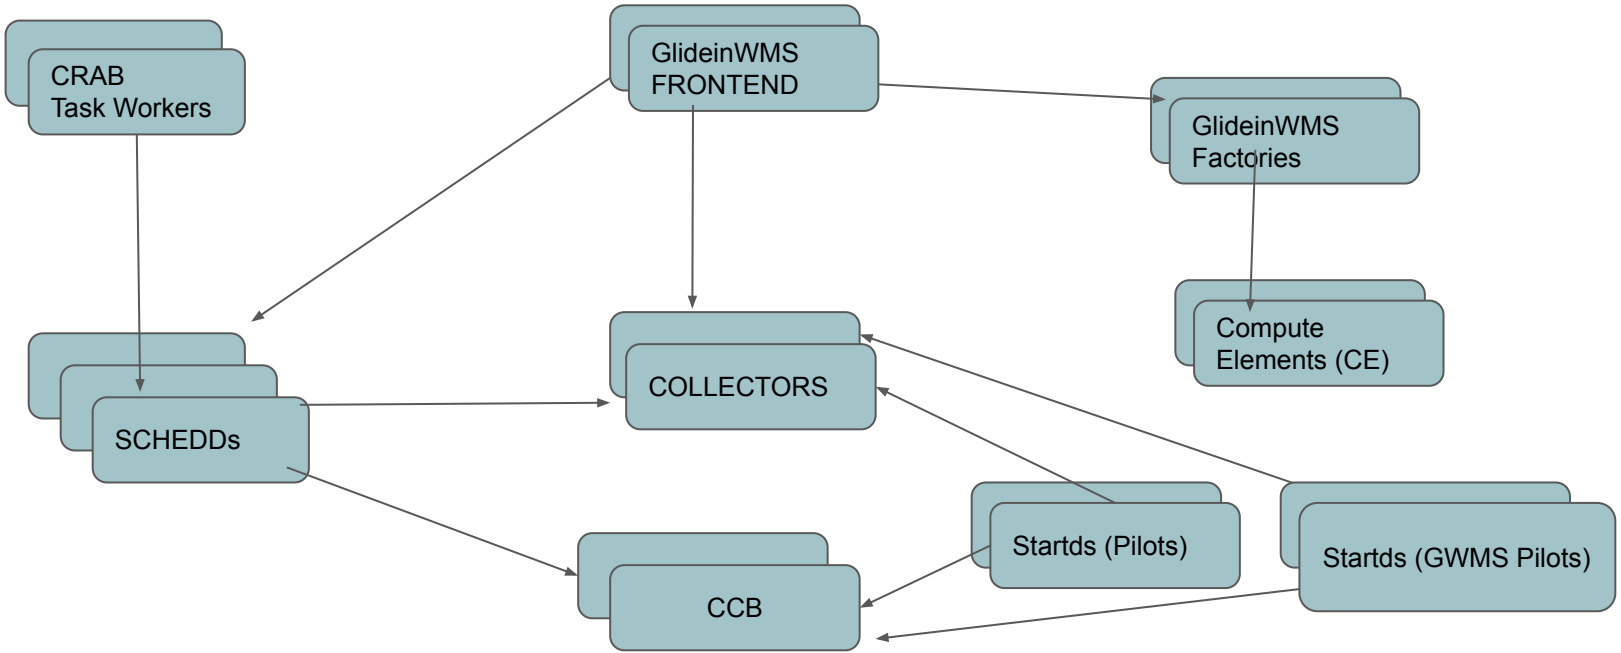

1.- IDTOKENS

IDTOKEN Implementation

- Whole Infrastructure is **already using** IDToken as preferred authentication method between components.
- 96% startds using IDTokens, 4% falling back to GSI due to token expiration.
- IDTOKENS created with need based authorization capabilities and limited lifetime.

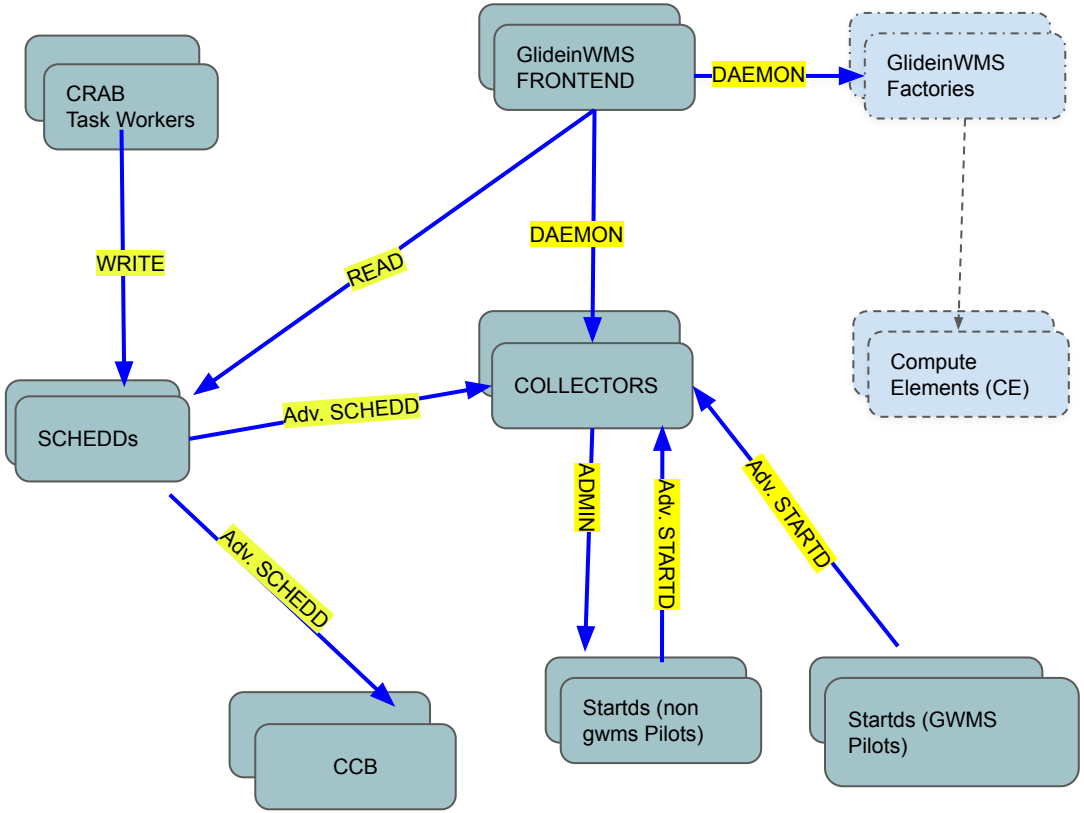

IDTOKEN authentication for GlideinWMS pilots

- **Same signing key** placed on both **Frontend** and **Collector**
- Frontend generates an IDTOKEN for each site.
- **IDTOKEN is transferred to** the factory, the CE, Batch System, **WN** (as pilot proxy before)
- **Startd is then authenticated and authorized** by the collector

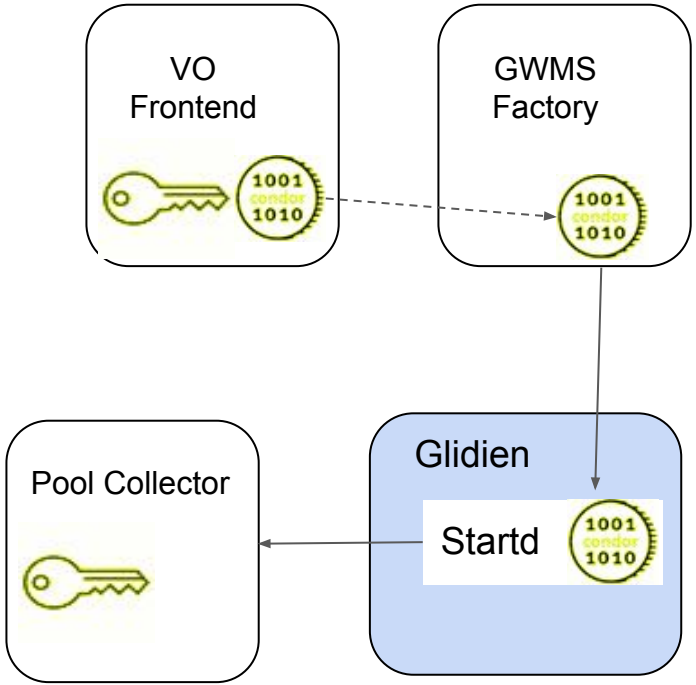

IDTOKEN Implementation

TRUST DOMAINS:

- Separate trust domains for:
  - Global and CERN pools
  - External schedds
  - Pilots and non-gwms startds.
- Each trust domain has its own signing key.
- External startds issues IDToken to CM which is required for admin operations: e.g **condor_drain**

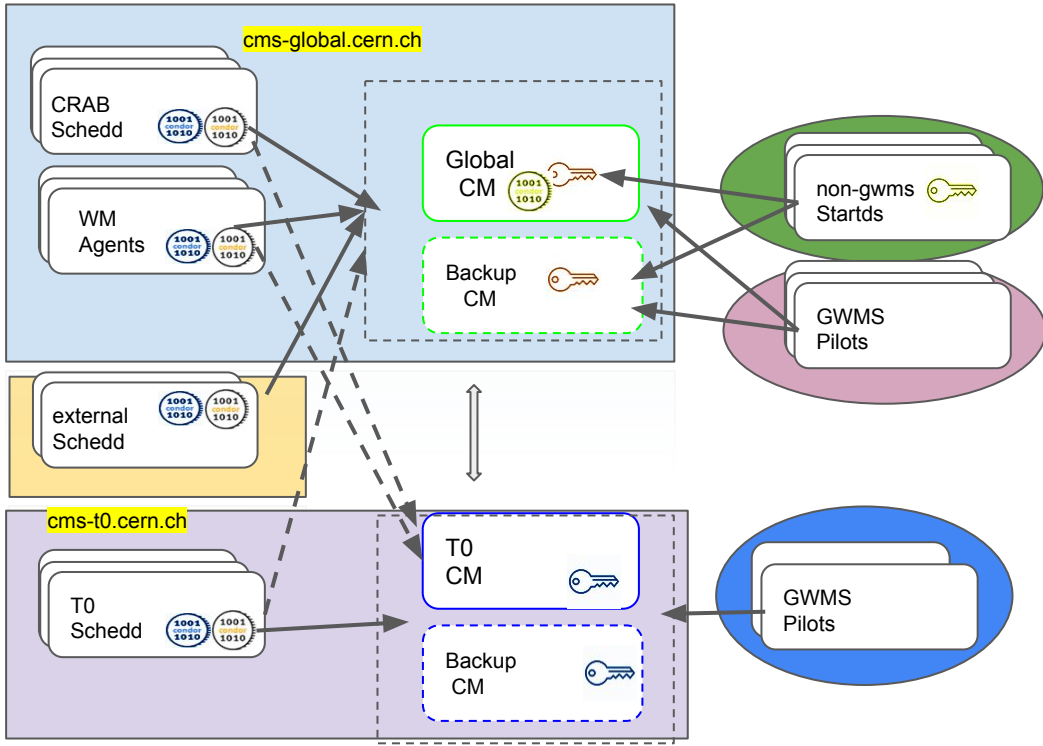

2.- SCITOKENS

SciToken Implementation

- SciToken authentication used for pilot submission between **Factory -> Compute Element (CE)**
  - HTCondor-CE (newer version)
  - ARC -CE (REST API)

- SciToken with different scopes/subjects are issued for different categories of pilots. e.g. local vs generic pilots.

- **CronJob:** Registered clients with CMS IAM fetches fresh token after every 10 minutes, and put it on FE, which is then used by factory for pilot submission.

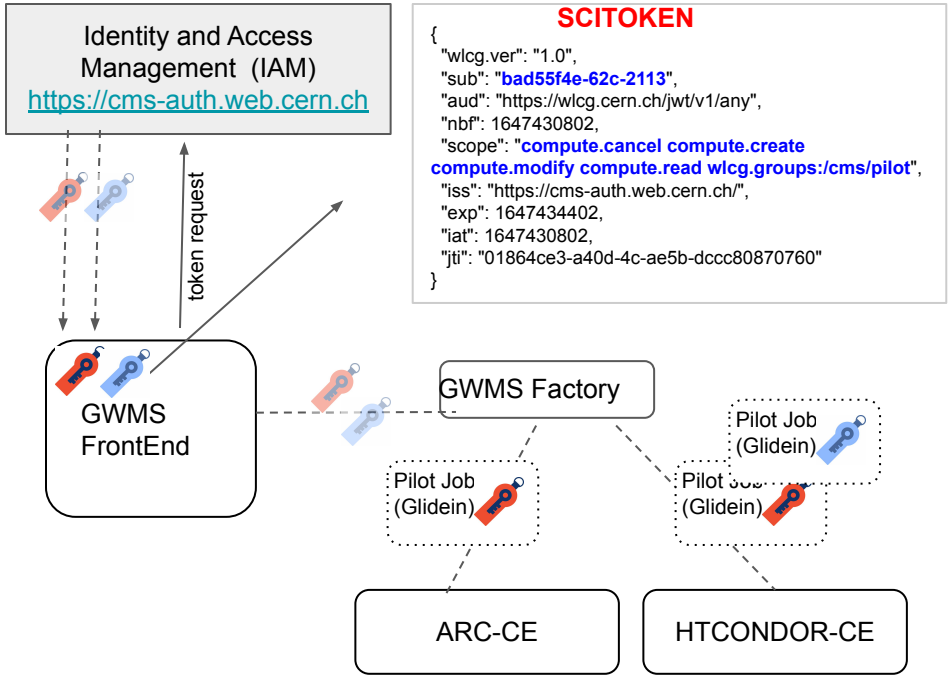

SciToken Implementation (HTCondor CEs)

- Nearly **70%** of the HTCondor CEs we interact with are already using recent enough HTCondor versions and supporting SciTokens authentication methods.
  - CERN last major site pending!
- CMS is working in a systematic way with each grid site to minimize disruption during transition (= **transparent from CMS Operations point of view)**
  - Separate glideinWMS FE group "main-token" created for submitting jobs with SciToken credentials.
  - Individual CEs are moved to token group after successful condor_ping test.
- Site admin perform mapping of different jobs based on scitoken's subject in htcondor-CEs e.g:

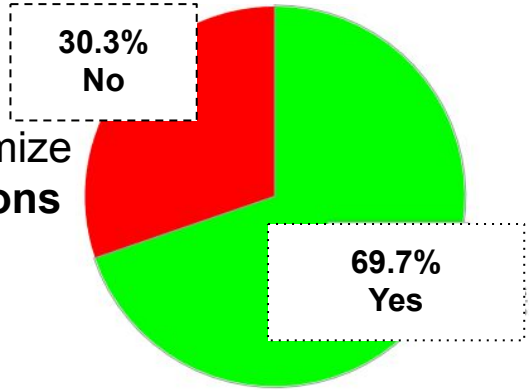

scitoken enabled CEs (HTCondor)

```
# CMS ITB generic pilots:
SCITOKENS /^https\:\/\/cms\-auth\.web\.cern\.ch\/,07f75a9a\-bb78\-4735-938b\-7e61b2b6d5c$/ cmspilot
# CMS ITB local pilots:
SCITOKENS /^https\:\/\/cms\-auth\.web\.cern\.ch\/,efbed8c1\-f9a7\-4063\-92f7\-f89c04c04a3$/ cmslocal
```

SciToken Implementation (ARC CEs)

- In the case of ARC CEs (about ⅓ of our total sites and CEs use this technology), our strategy so far has been to test that we can interact with them via x509 proxies but with the new REST interface
  - About **74%** of all ARC CEs we use already OK
- Secondly, we are already testing a HTCondor pre-release capable of submitting pilots with SciTokens on to an ARC CE
  - Tested with T2_IT_Rome, ok!

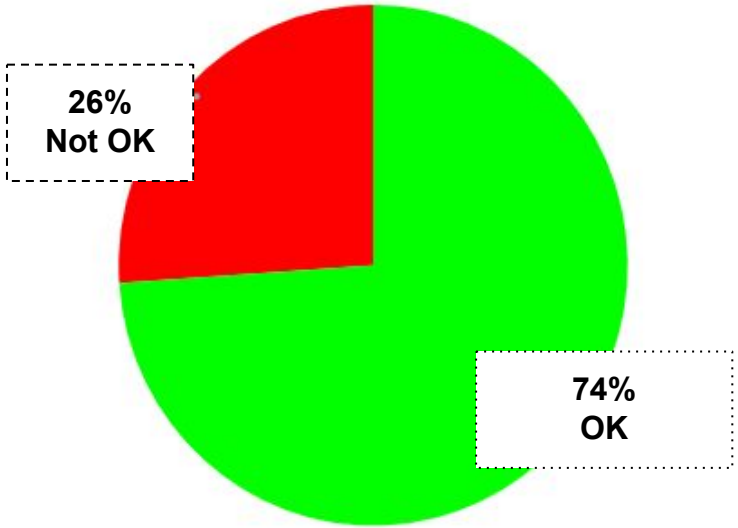

ARC-CEs REST interface status

Summary and next steps

Summary

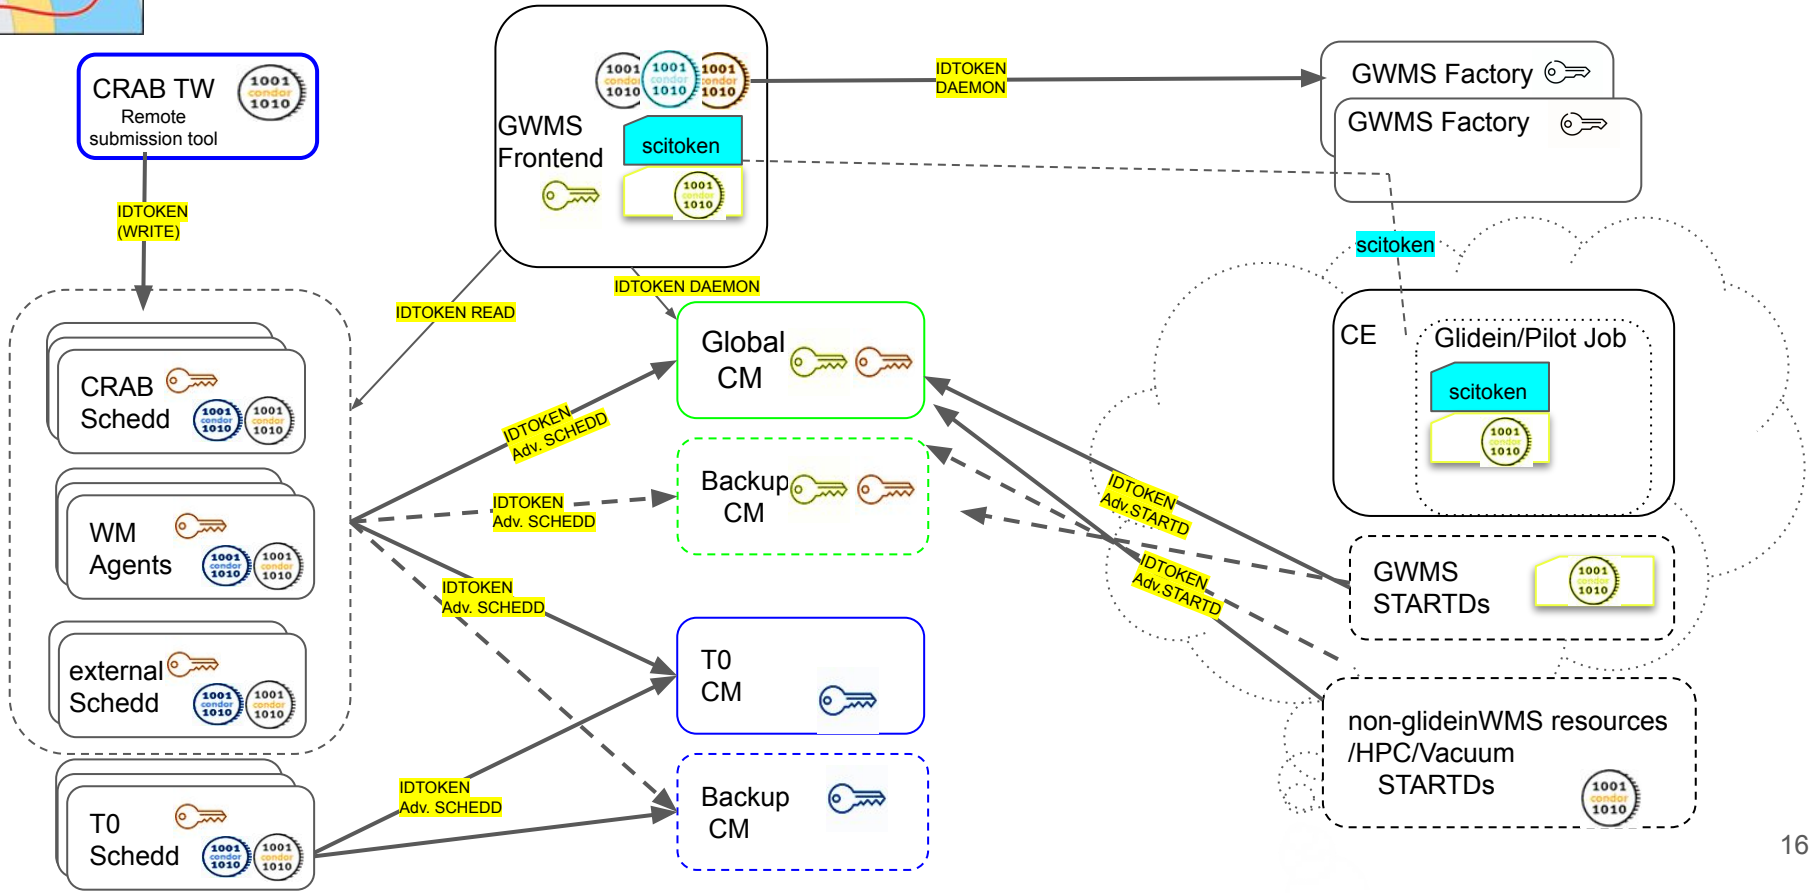
CRAB TW
Remote submission tool
IDTOKEN (WRITE)
GWMS Frontend
scitoken
IDTOKEN DAEMON
GWMS Factory
GWMS Factory
scitoken
IDTOKEN READ
IDTOKEN DAEMON
CRAB Schedd
WM Agents
external Schedd
T0 Schedd
Global CM
Backup CM
T0 CM
Backup CM
IDTOKEN Adv. SCHEDD
IDTOKEN Adv. SCHEDD
IDTOKEN Adv. SCHEDD
IDTOKEN Adv. SCHEDD
IDTOKEN Adv.STARTD
IDTOKEN Adv.STARTD
CE
Glidein/Pilot Job
scitoken
GWMS STARTDs
non-glideinWMS resources /HPC/Vacuum STARTDs

Conclusions and next steps

- CMS Submission Infrastructure internal components fully switched to IDTOKENS with fallback to GSI.
  - No dependency on external components (CRLs, Argus…)
- CMS ITB Pool (i.e. with same components) is running fine with condor feature release 9.8.0 (i.e. without GSI support).We can start dropping GSI fallback method soon in production pools.
- Still working with CMS WLCG sites to guarantee a seamless transition to Scitoken for Factory<->CE communication

We thank the HTCondor development team for the continued support to CMS Submission Infrastructure over the years, a model of excellent partnership!