Token transition state of affairs

WLCG Authorization WG

WLCG Workshop, Lancaster, Nov 8, 2022

v1.0

Introduction

- We aim for the transparent replacement of X509 + VOMS with tokens in a steadily increasing number and variety of workflows
  - Deployment started during LS2, to be mostly completed during Run 3
  - See the AuthZ WG pages and vCHEP 2021 talk for background information

- Agenda
  - INDIGO IAM recap and usage
  - Before and after diagrams
  - Smoothing the transition
  - Computing
  - CE token hackathon
  - AuthZ and IAM Workshop
  - Data management
  - Auxiliary services
  - Timelines & milestones
  - Phasing out VOMS-Admin
  - Globus retirement implications

INDIGO IAM recap

- INDIGO IAM will replace VOMS(-Admin)
  - Identity and Access Management
- Integrates federated identities through CERN SSO plugin
- Obtains LHC experiment membership details from the CERN HR DB
  - Just like VOMS-Admin
- Can issue fine-grained tokens to users and services
  - Details depend on the configuration per VO and per workflow
- Also has a VOMS endpoint for backward compatibility
  - Users will still have their X509 certificates linked for now
- It does **not** have a VOMS-Admin endpoint
  - Classic grid-mapfiles etc. have to be constructed using the IAM SCIM API
  - For any IAM instance, all SCIM API clients have to be **registered**

New AAI architecture for WLCG

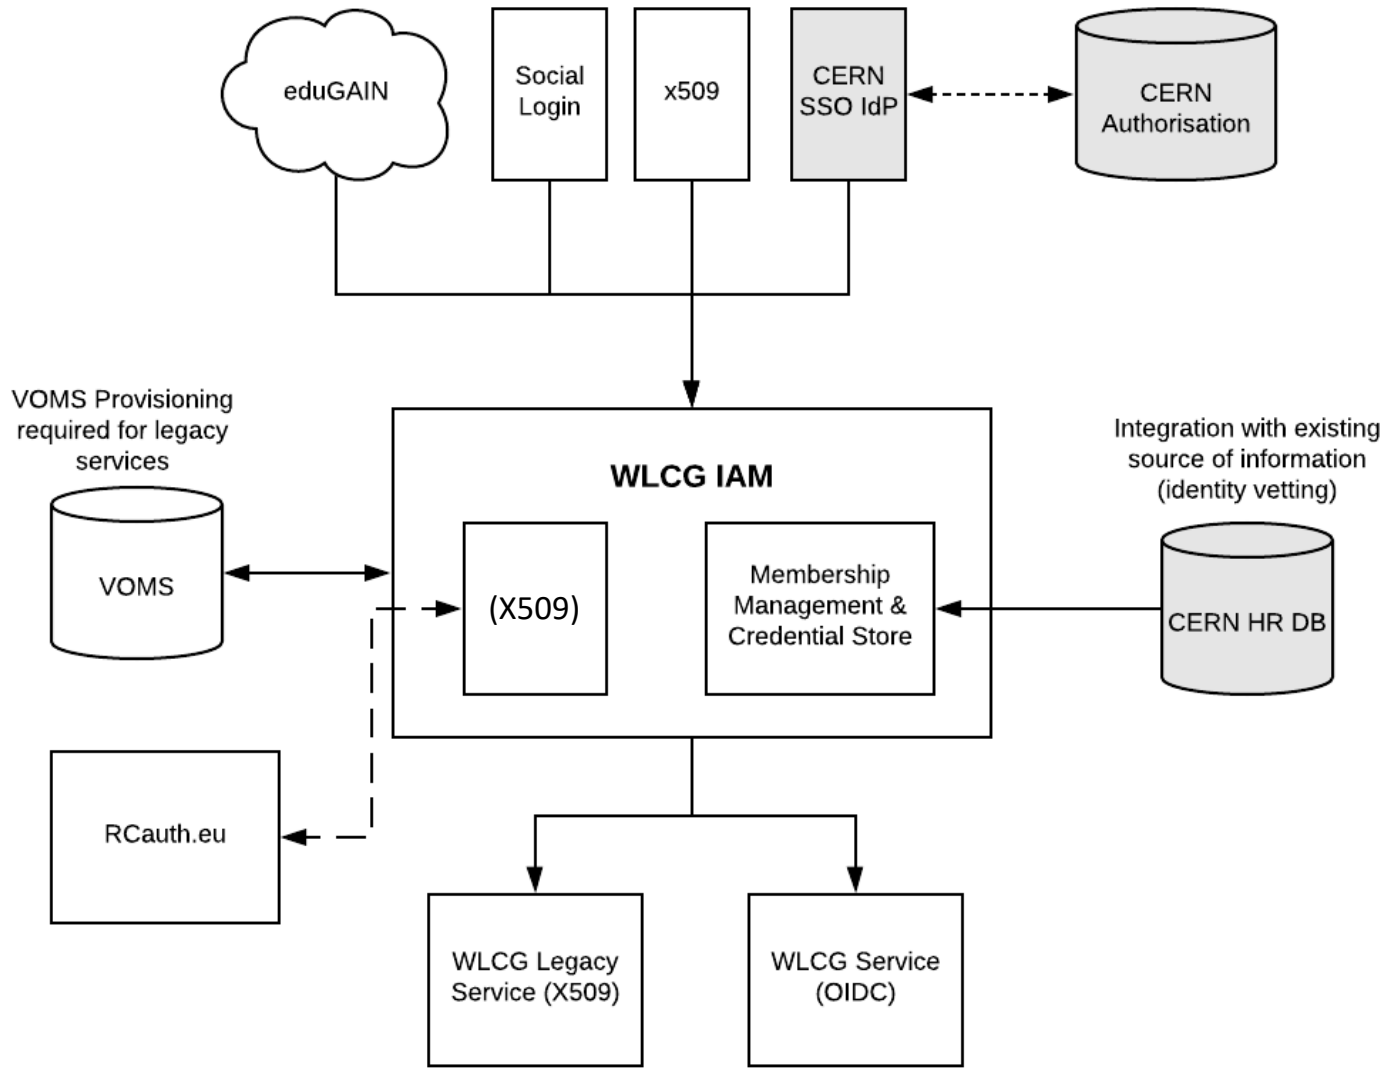

Workload management with X509 + VOMS

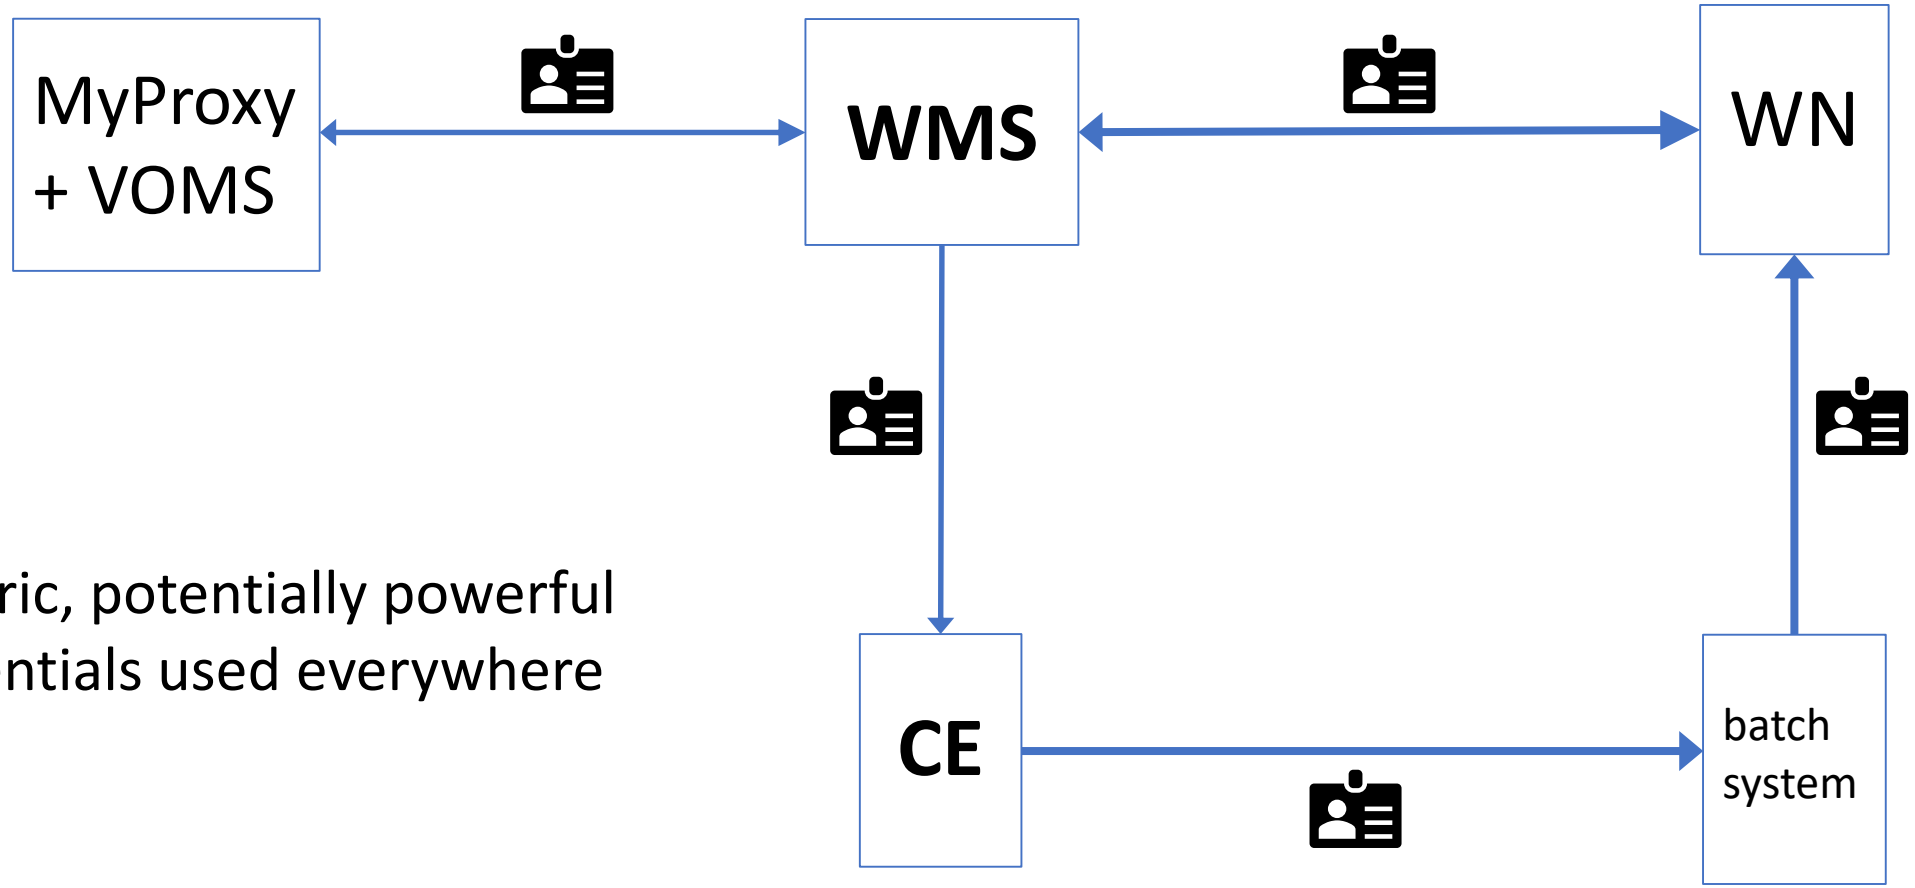

Generic, potentially powerful credentials used everywhere

Workload management with tokens

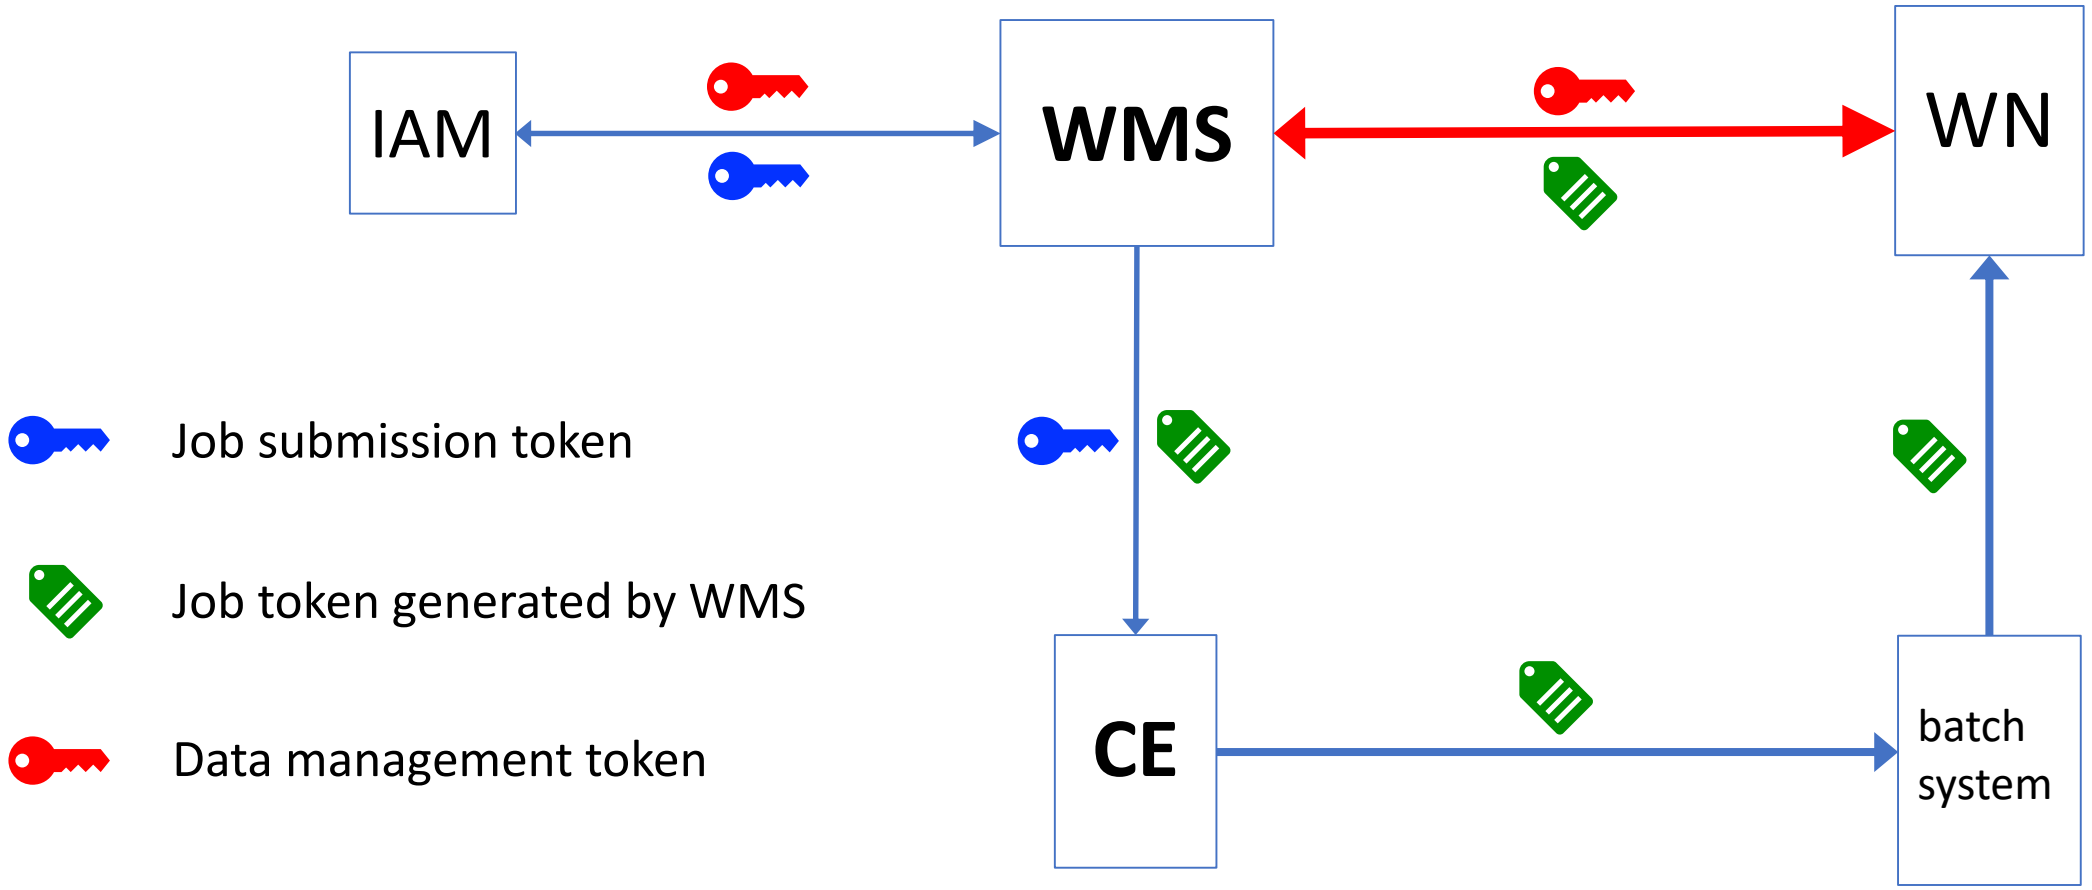

Data transfers with X509 + VOMS

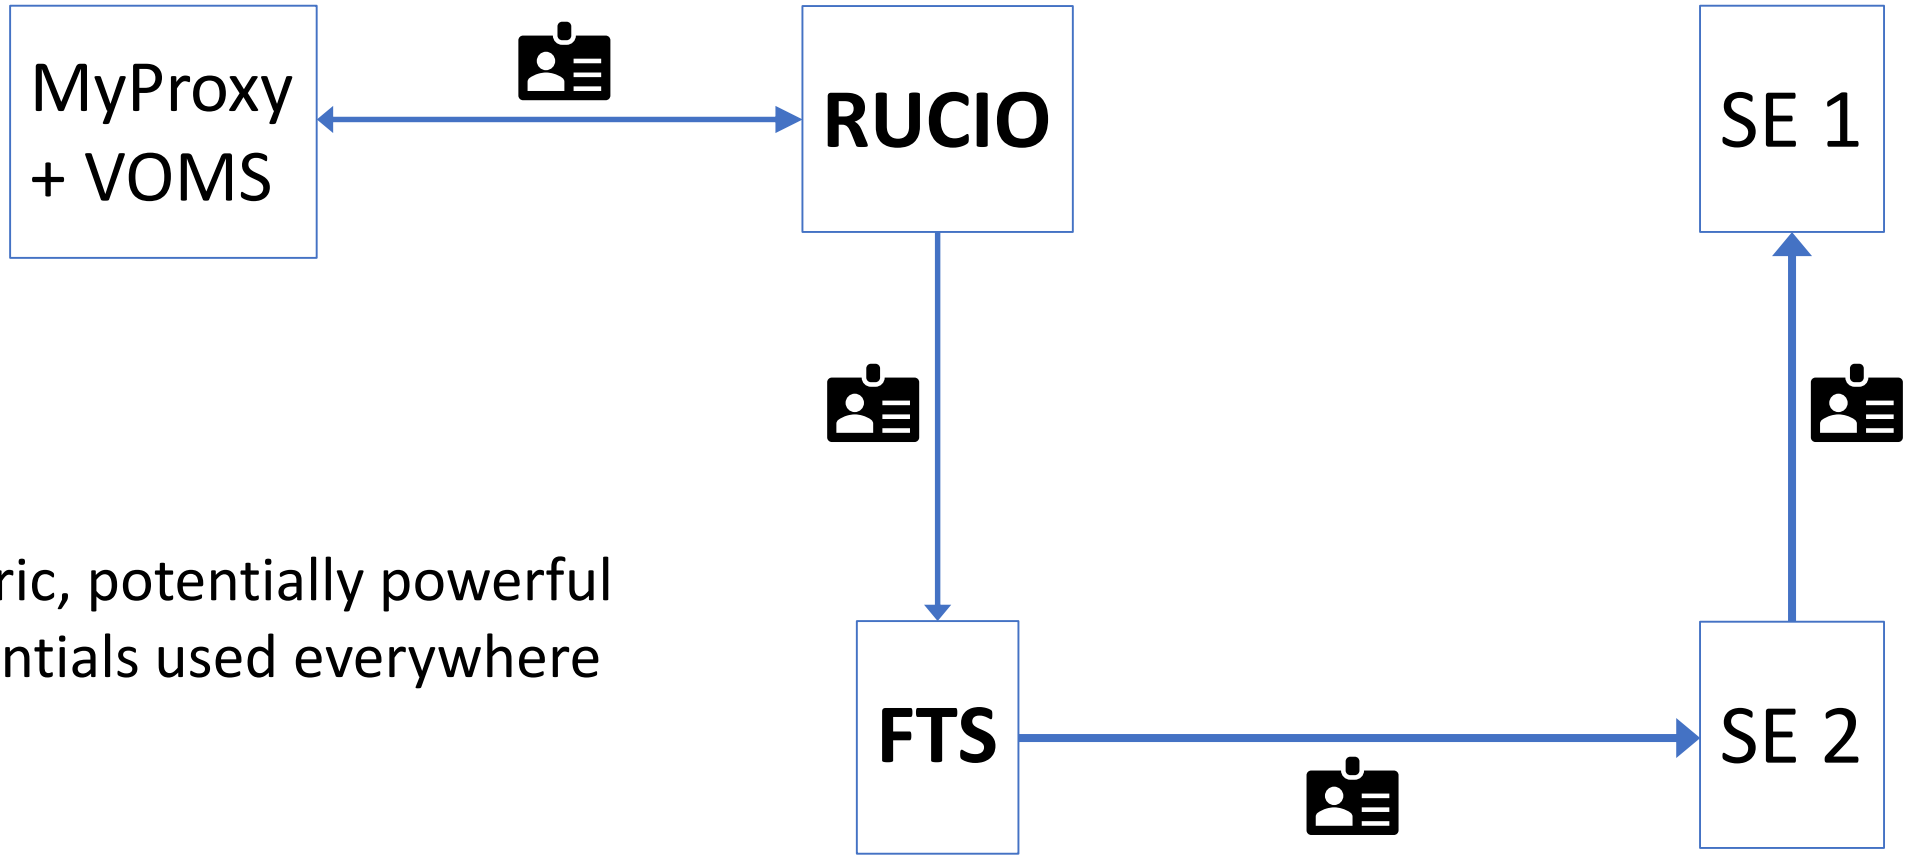

Generic, potentially powerful credentials used everywhere

Data transfers with macaroons

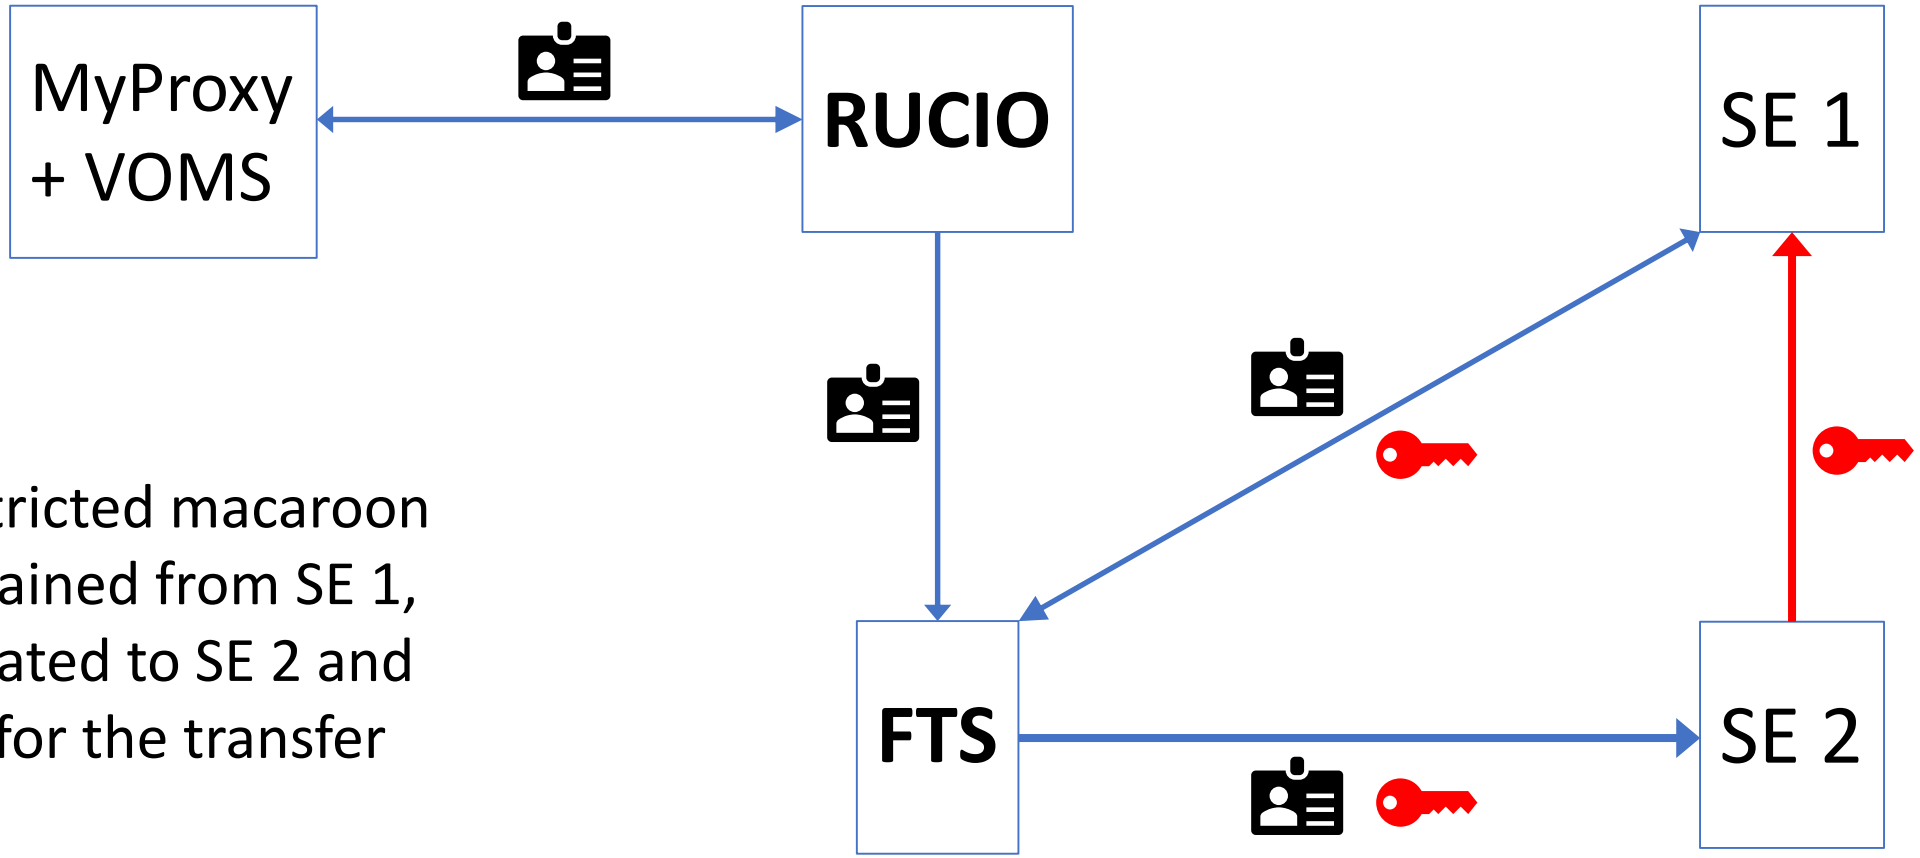

A restricted macaroon is obtained from SE 1, delegated to SE 2 and used for the transfer

Data transfers with WLCG tokens

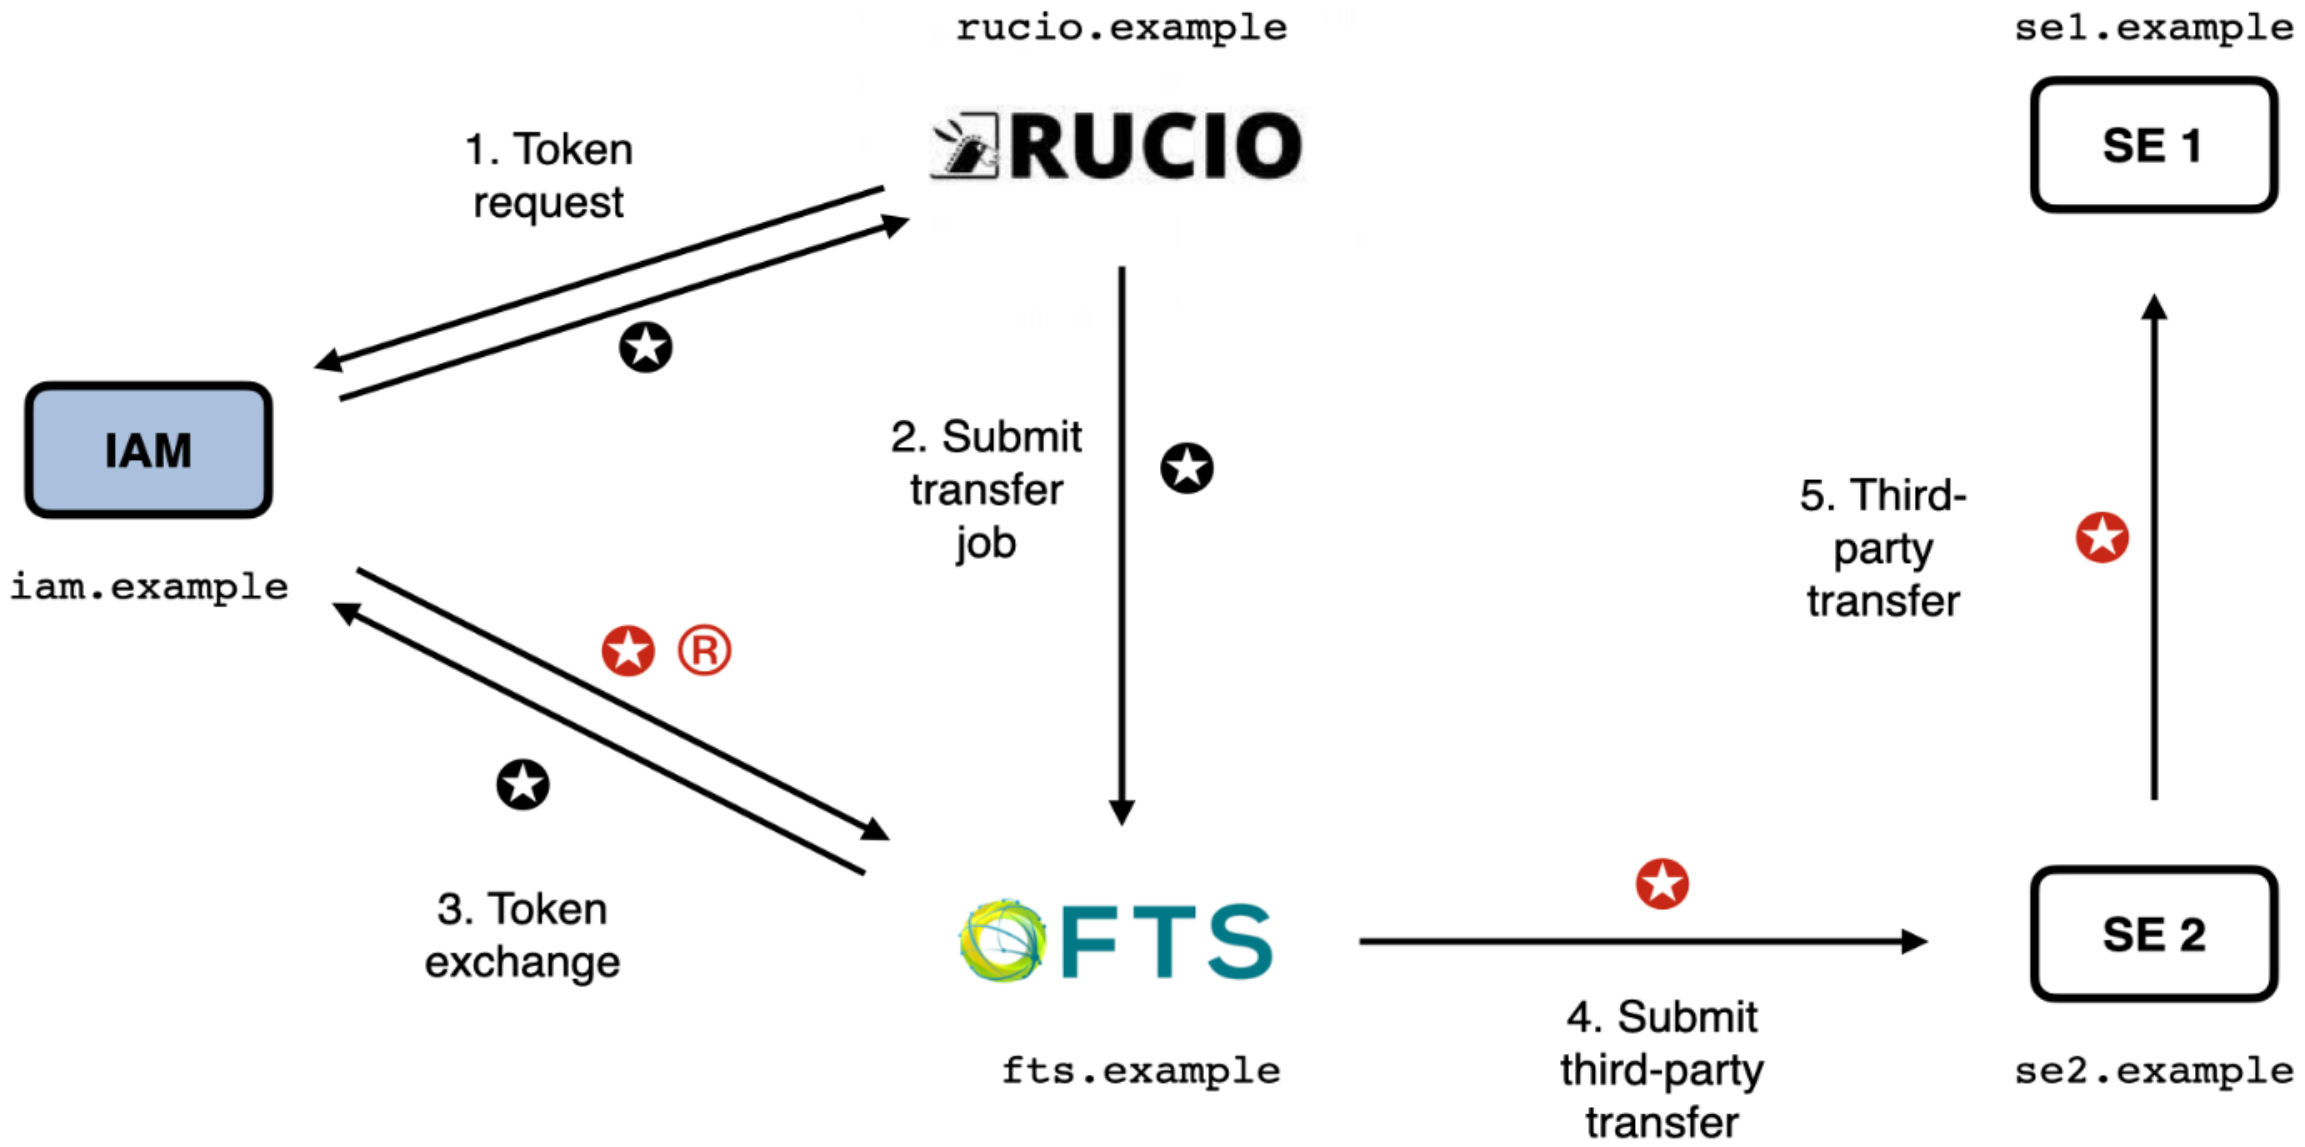

IAM usage

- IAM is in production for the four LHC experiments
  - Contents are automatically replicated *from* VOMS-Admin
    - Modulo known issues with workarounds
  - IAM VOMS endpoints can be used alongside the legacy VOMS services
    - In production for ATLAS and CMS
    - TBD for ALICE and LHCb
- ATLAS, CMS and SAM ETF can use tokens to *submit* jobs to HTCondor CEs
  - In particular the CEs on **OSG**, which only support tokens since May
- Those jobs *still* use X509 VOMS proxies for data management etc.
- For ALICE, LHCb, Belle-II, … these matters are WIP
  - See next pages

Smoothing the transition

- HTCondor GSI support EOL has been postponed to Feb 2023
- Pilot job submission tokens currently have lifetimes of a few days
  - Allow much more time to resolve IAM service incidents transparently
  - Those tokens are only used in security handshakes with CEs
- Long lifetimes will be reduced when sufficient positive experience has been obtained with the reliability and the support level of the IAM services
- In the future, *data management* tokens will require much shorter lifetimes for the desired level of security
- The service deployment will be made as HA as feasible
- Experiment experts can open tickets in SNow or directly in GitHub

Computing

- CE token support deployment campaign on EGI launched June 1: 70+%
  - HTCondor v9.0.x with tokens for ATLAS and CMS, others later
  - ARC CE REST interface, in particular to support job submissions via HTCondor-G
- Another campaign on EGI will be needed early 2023 to get all HTCondor CEs on supported versions > v9.0.x
  - Also EGI Check-in tokens should work by that time
- ALICE
  - HTCondor CE job submission can use tokens
  - IAM works for VOMS proxies, token setup TBD
- LHCb, Belle-II, …
  - DIRAC progress was presented in the June GDB
  - v8.0 provides token support for HTCondor CE
    - May be backported to v7.x series
  - Token-ready releases to be deployed
  - IAM services to be set up
- DUNE
  - Are using a CILogon token issuer at FNAL to submit jobs
  - Have tested accessing dCache with tokens
  - Need FTS and Rucio to work in a hybrid mode to be able to transition to tokens

CE token hackathon

- A CE token hackathon has been held at NIKHEF, Sep 15-16
  - A summary was presented in the Oct GDB
  - Important outcome: plugin API with callouts to support multiple profiles
    - SciTokens, WLCG, EGI Check-in
    - Still to be implemented !
    - Could also be the place to *reject blacklisted tokens* (cf. Argus)
  - Other discussion items included:
    - How to map tokens to accounts → LCMAPS successor would be desirable
    - How to equip jobs with tokens → depends on the VO job management framework
    - How to let the ARC CE do data management operations with tokens
    - User-friendly token client machinery
    - Token infrastructure security assessments

- Discussion has also started on how to extend the profile with further use cases
  - Access to various kinds of specific computing resources behind a CE

AuthZ and IAM workshop

- An AuthZ and IAM workshop was held at CERN on Oct 10-11
  - A summary was presented in the Oct GDB
  - The workshop featured a few presentations and lots of discussion
    - Admin tokens
    - Service accounts
    - User suspension flows
    - Refresh token flows
    - Automatic import and activation of accounts from CERN HR DB
    - Group management delegation
    - MFA integration
    - Flexibility and refinements of compute scopes
    - Request rate sustainability
    - Hackathon at the end of winter, likely at CERN
    - Engagement with other user communities
    - …

Data management

- Mostly covered in the DOMA session on Nov 7: token status & plans
- Workflow details involving Rucio/DIRAC and/or FTS vs. SEs have mostly been identified and implemented to various extents
  - May need to be re-discussed if major implementation or operational hurdles are encountered
- The token testbed covers *basic* functionality and interoperability
  - Most endpoints pass most tests
- Rucio and DIRAC should drive this → implications for the FTS
  - Will see further progress in the next months (cf. timeline on page 17)
  - FTS team will have 1 FTE joining
- SEs typically need to support concurrent use of X509 and tokens
  - Details documented here

Auxiliary services

- Services resembling MyProxy will help simplify things for users and robot workflows

- HashiCorp Vault + *htgettoken* already in production at FNAL
  - Presented in the April GDB, 2021, and at vCHEP, 2021
  - To be tried out at CERN
    - A Vault service already exists

- *MyToken* already in use for EOSC Synergy infrastructure monitoring
  - Presented in the AuthZ WG meeting of 2 Sep 2021
  - It allows tokens and workflows to be restricted in more ways

Timelines & milestones

- WLCG Token Transition Timeline v1.0 was published on August 22
  - Summarizes the progress of the last 2.5 years and presents a set of optimistic milestones to work towards during the course of Run 3
  - Subsequent versions will be added whenever there are significant updates of the timeline
  - Already put to use to get extra FTEs for the CERN IAM and FTS teams!

- Selected milestones (tentative dates)
  - M.2 (Dec 2022) – DIRAC deployments
    - To support at least HTCondor CE job submissions with tokens
  - M.3 (Feb 2023) – VOMS-Admin switched off for 1 or more experiments
    - See next page
  - M.4 (Mar 2023) – HTCondor installations on EGI are > 9.0.x
    - EGI Check-in tokens should be supported by then
  - M.8 (Mar 2024) – DC24 can be done with tokens
  - M.10 (Mar 2026) – users no longer need X509 certificates

Phasing out VOMS-Admin

- Prerequisites
  - Significant VO admin functionality issues in IAM need to be resolved
  - Experiments need to switch user, group and role management to IAM
  - The IAM services need to be made more robust (cf. *Halloween* outage)
    - Discussion ongoing about the token request rates to be sustained per VO
  - The CERN IAM team needs to have grown
  - Remaining VOMS-Admin use cases need to be moved or phased out

- Probably best done one experiment at a time
  - Its “vomses” files will need to have the legacy endpoints removed, ideally beforehand through a short campaign
  - The legacy hosts should remain up for a while longer in any case, with just the relevant VOMS services switched off
  - Clients will switch to the IAM VOMS endpoint, possibly after a warning
  - Legacy services could first be switched off temporarily as a test

Globus retirement implications

- OSG no longer depend on “Globus”, i.e. the Grid Community Toolkit (GCT)

- By March 2023, no supported HTCondor release depends on the GCT (link)
  - No more GSI support for job submissions to HTCondor CEs anywhere
  - HTCondor-G can still use X509 to submit jobs to ARC CEs via their REST interfaces

- The GridFTP protocol has mostly been phased out in WLCG
  - DOMA Bulk Data Transfer WG to follow up on remaining usage
  - Sites may need to keep supporting it for other VOs until those have also switched to HTTPS+WebDAV and/or Xrootd

- The GCT is maintained by the Grid Community Forum at best-effort level
  - Latest updates were released on June 8
  - RPMs are even published in EPEL 9
  - We will need various components until the transition to tokens has been completed