Traceability & Isolation WG

Vincent BRILLAULT, CERN/EGI-CSIRT

GDB March 2017, ISGC, Taipei

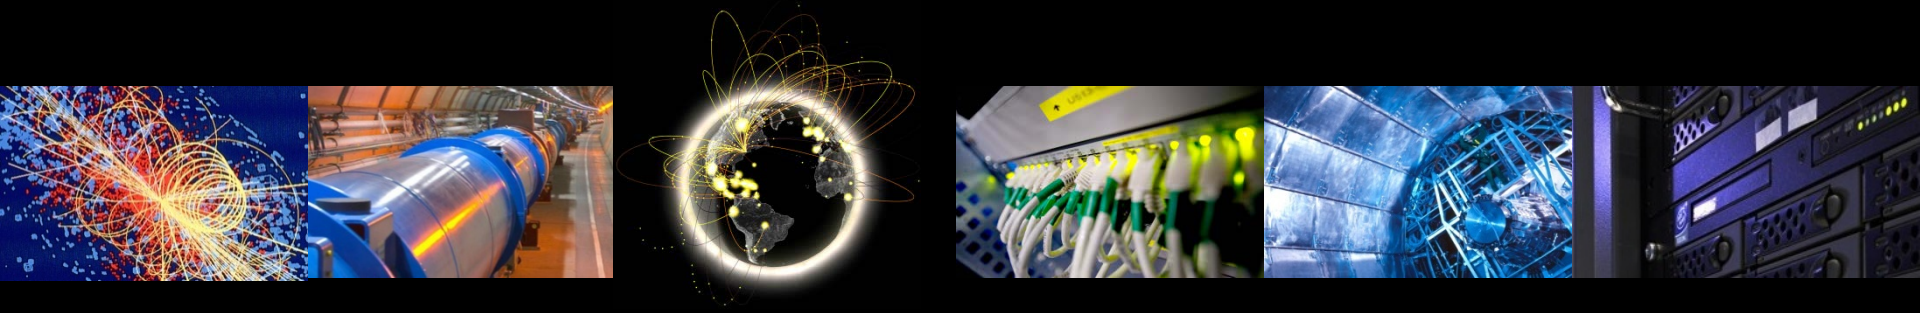

WLCG
Worldwide LHC Computing Grid

CERN

Traceability & Isolation Working Group

- Working Group created after GDB discussions in March and April 2016
- Kick-off meeting in May 2016, 5 meetings since
- All details and meeting link on web site: https://cern.ch/wlcg-traceability-isolation-wg
- Focus on evolution, not revolution

Mandate

Explore new traceability and isolation paradigms, propose a new model taking advantage of new technologies and VO frameworks while keeping full trustworthy traceability and isolation of users actions.

Split Taceability and Isolation

- glExec model:
  - glExec provides traceability (user certificate)
  - glExec provides isolation (uid change)
  - VOs only partially trusted?
    - Trust: push matched certificate and payload (same user)
    - Not trusted: traceability ?

Split Traceability and Isolation

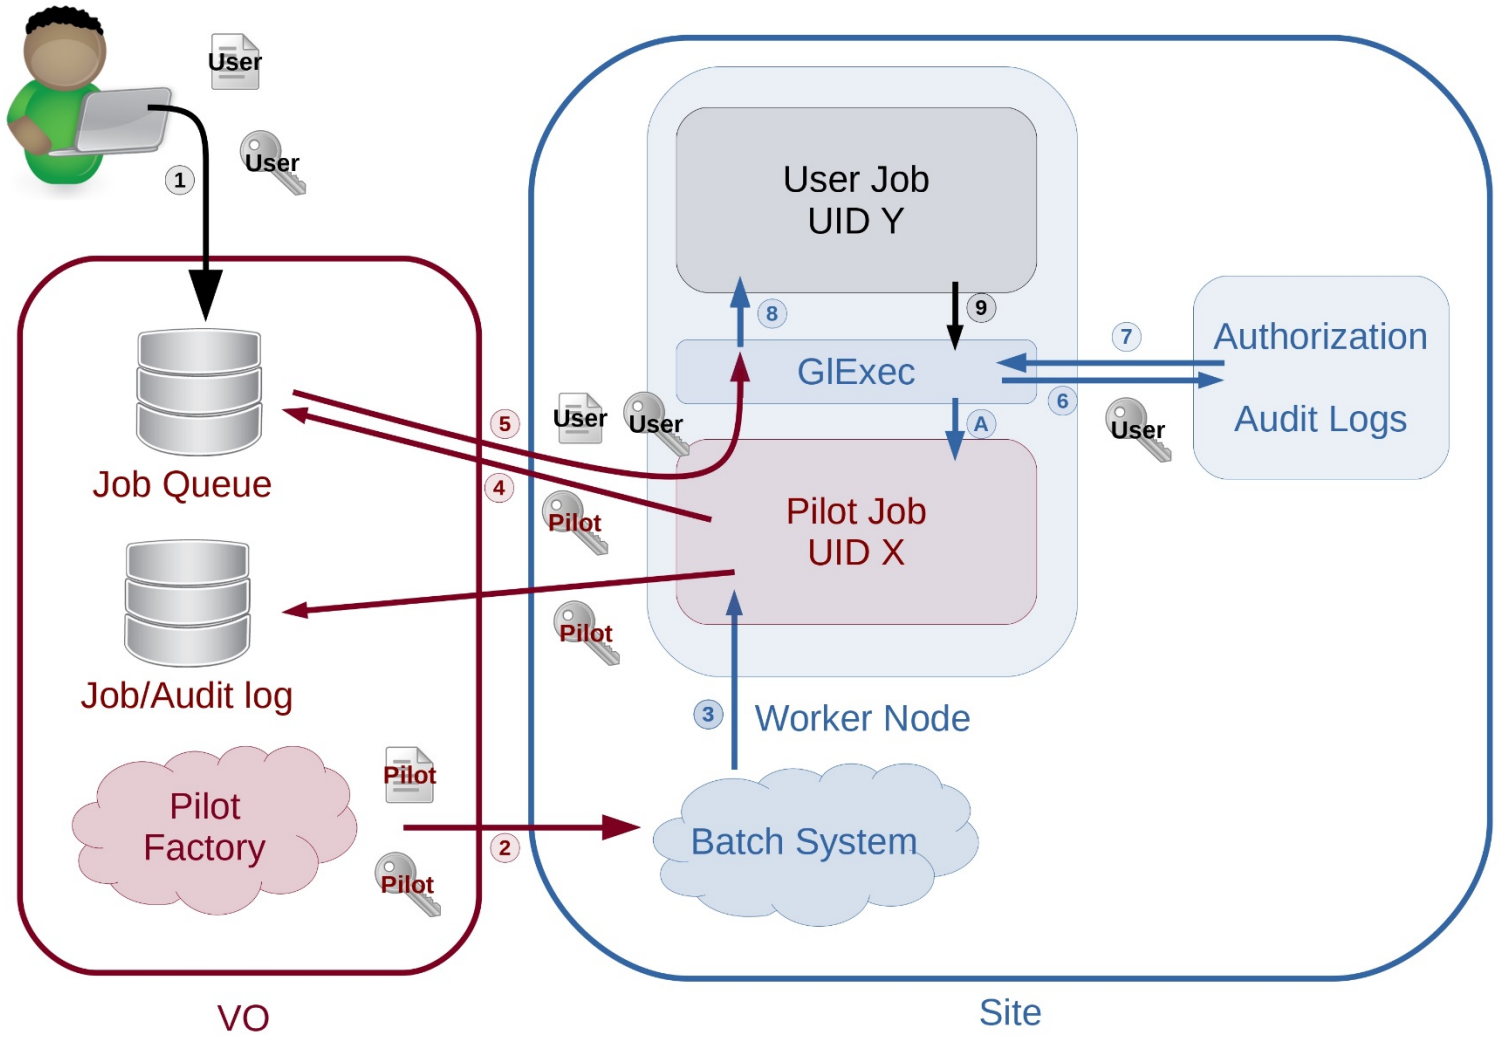

Split Traceability and Isolation

- glExec model:
  - glExec provides traceability (user certificate)
  - glExec provides isolation (uid change)
  - VOs only partially trusted?
    - Trust: push matched certificate and payload (same user)
    - Not trusted: traceability ?

- WG focus: split traceability and isolation
  - Isolation: use container technology (namespaces)
  - Traceability: use VO frameworks

Split Traceability and Isolation

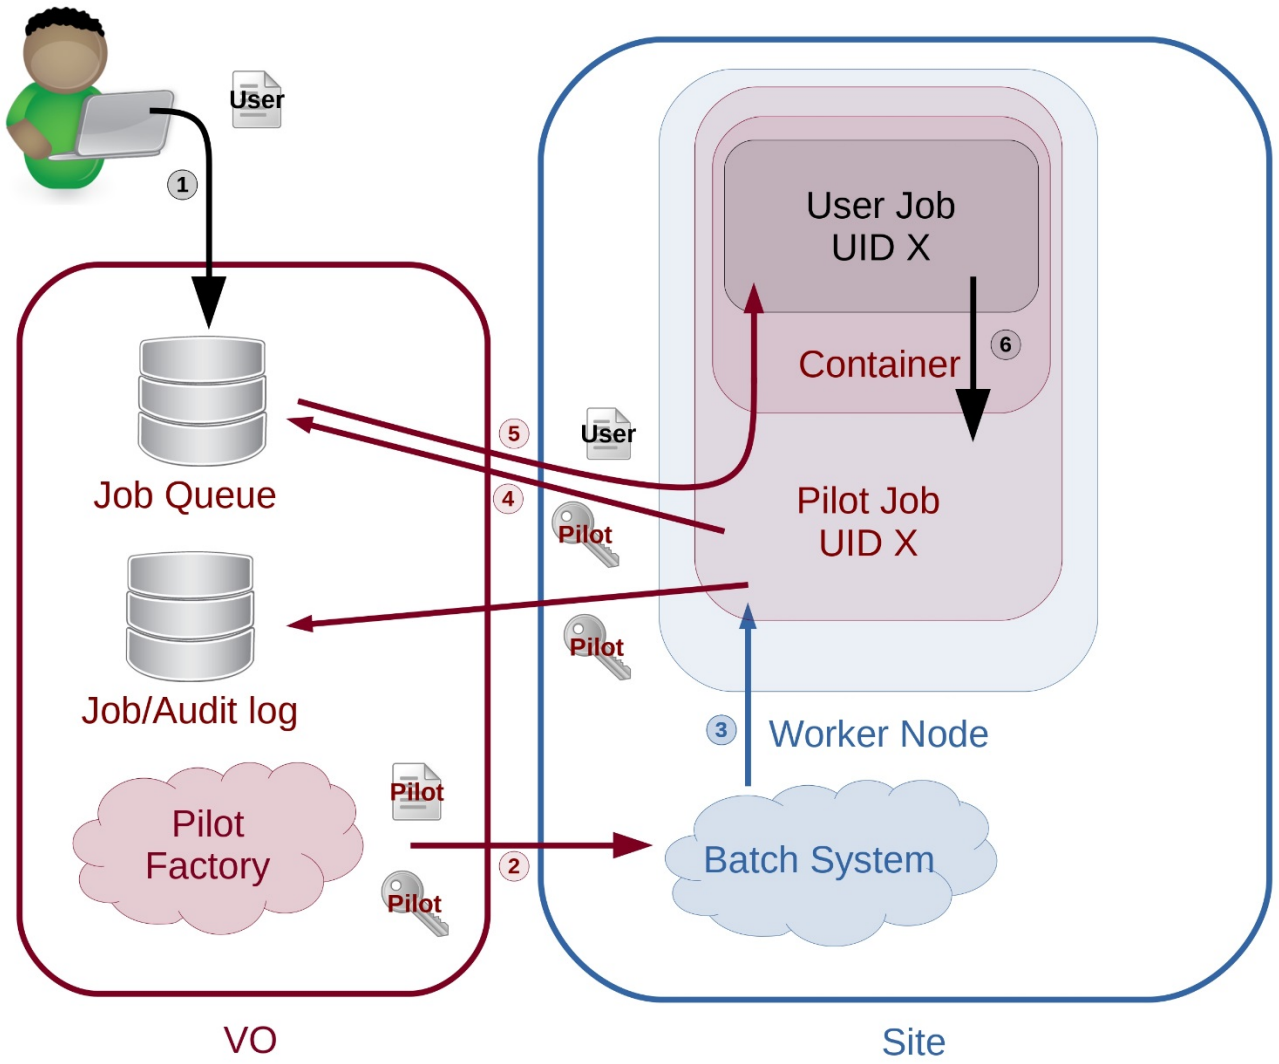

Split Traceability and Isolation

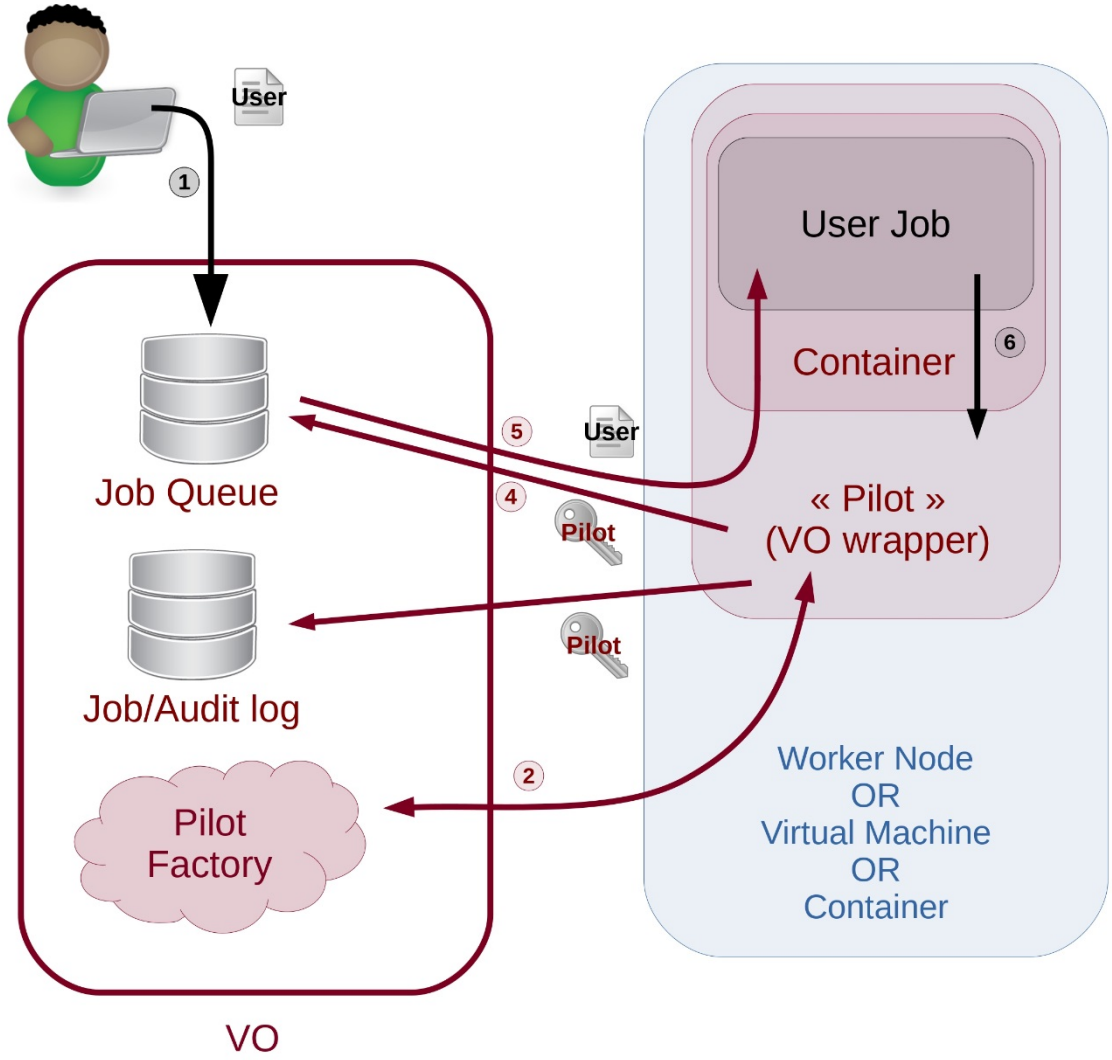

Working Group activities

- Evaluate new isolation solution
  - Using containment
  - Compatible with grid/cloud deployments
- Evaluate new traceability paradigm
  - Based on VO framework
  - Keep full traceability down to the user

New Isolation solution

- Existing tool identified by Brian Bockelman: Singularity (http://singularity.lbl.gov/)
  - Container solution initially coming from HPC
  - Not Docker: single binary, no root daemon
  - Not glExec: no UID mapping & switching
  - Isolation only: no external call-out (e.g. Argus)
  - Requires SUID on RHEL6 and RHEL7
    - No SUID required on recent upstream kernels
    - RedHat refused to backport it to RHEL 7.3*
- Now being tested by WG, esp. OSG-CMS

Singularity testing

- Good progress by OSG
  - Simplified installation: single RPM to install
  - Already deployed at ~15 sites
  - Already 1M singularity jobs run
  - ~200 lines of bash for setting up environment
  - CVMFS import from Docker for OSG users
  - Singularity in container possible (not default)
- CMS integration tests now ongoing
  - Could be used for RHEL7-only worker nodes
- Small HTCondor cluster with Singularity deployed for testing at CERN

Singularity: SUID…

- SUID will be required for some time (RHEL7.4?)
- No external security audit/review yet
  - Possibility of review identified in US
    But might not be possible before Autumn 2017
  - Looking for solutions in EU, none identified yet
- Convincing sites?
  - Large deployment exists (e.g. GSI Greencube)
  - Much simpler than GLExec?

Proposed traceability paradigm

- Sites still responsible for:
  - Which VO run on a Host/IP at a given time
  - Which VO was responsible for activity on a given slot on a worker node
- VOs now responsible for
  - Which user run at a given slot/host/IP & time
- Under discussion: data accesses?

Current Incident Response

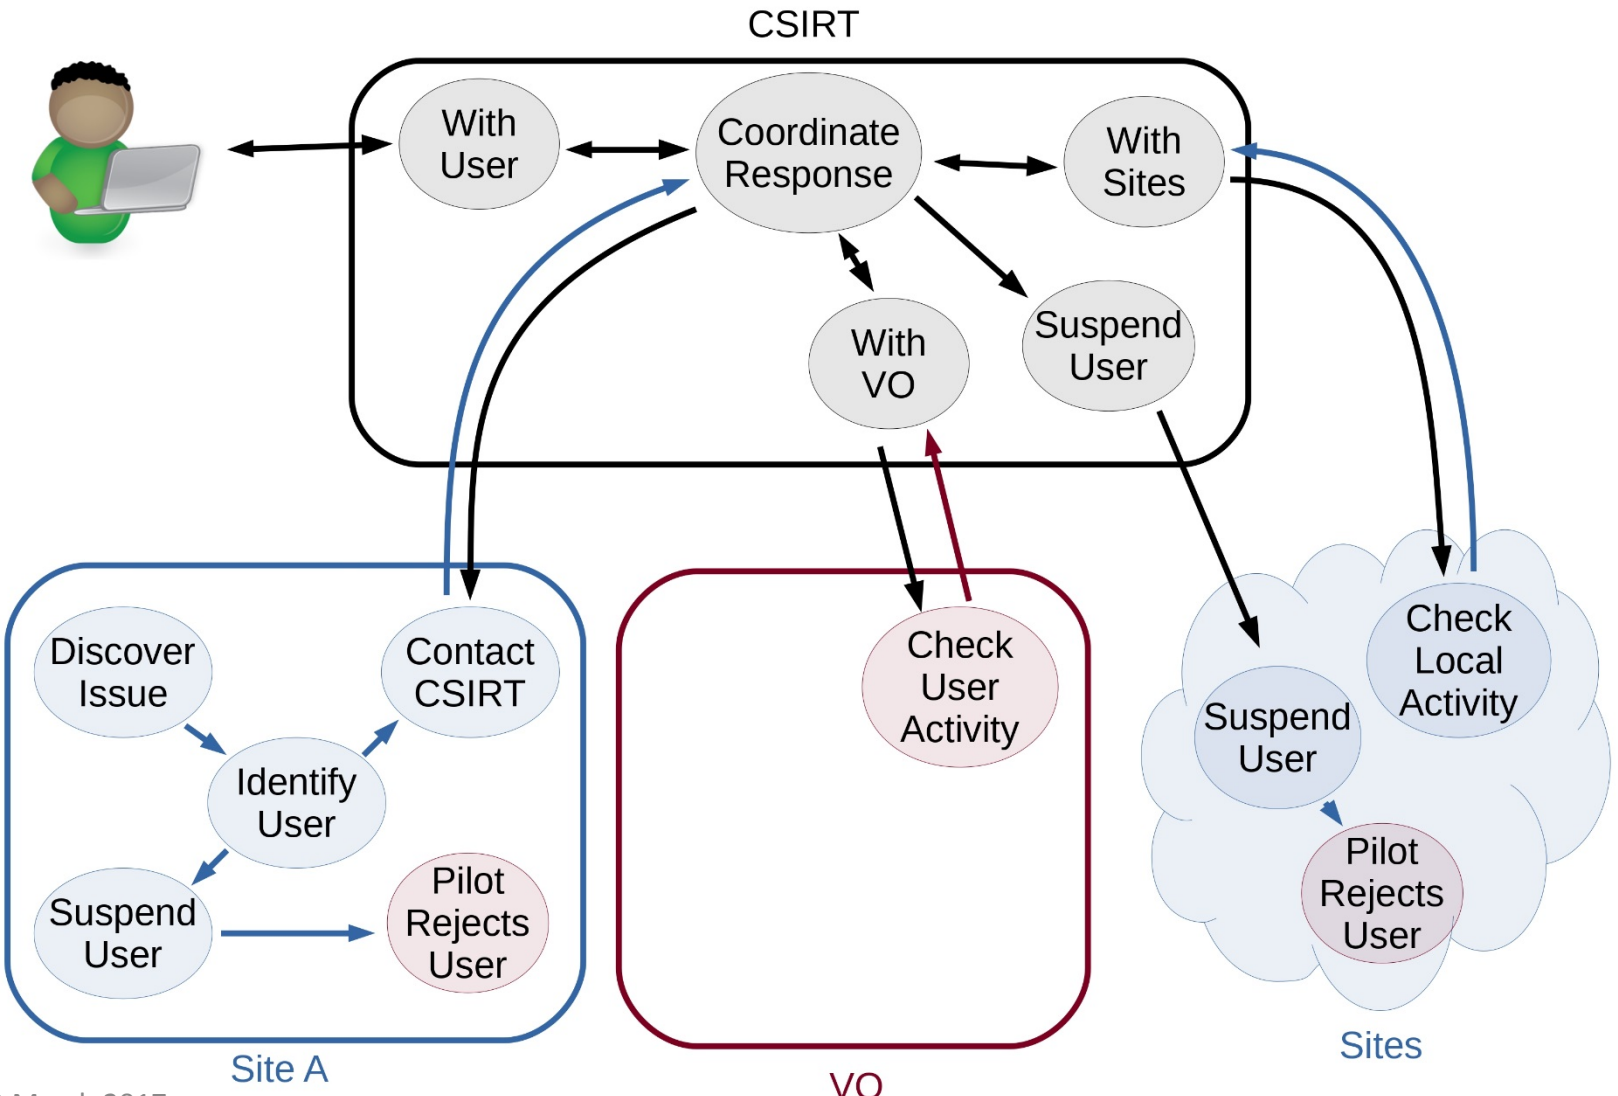

Proposed Incident Response

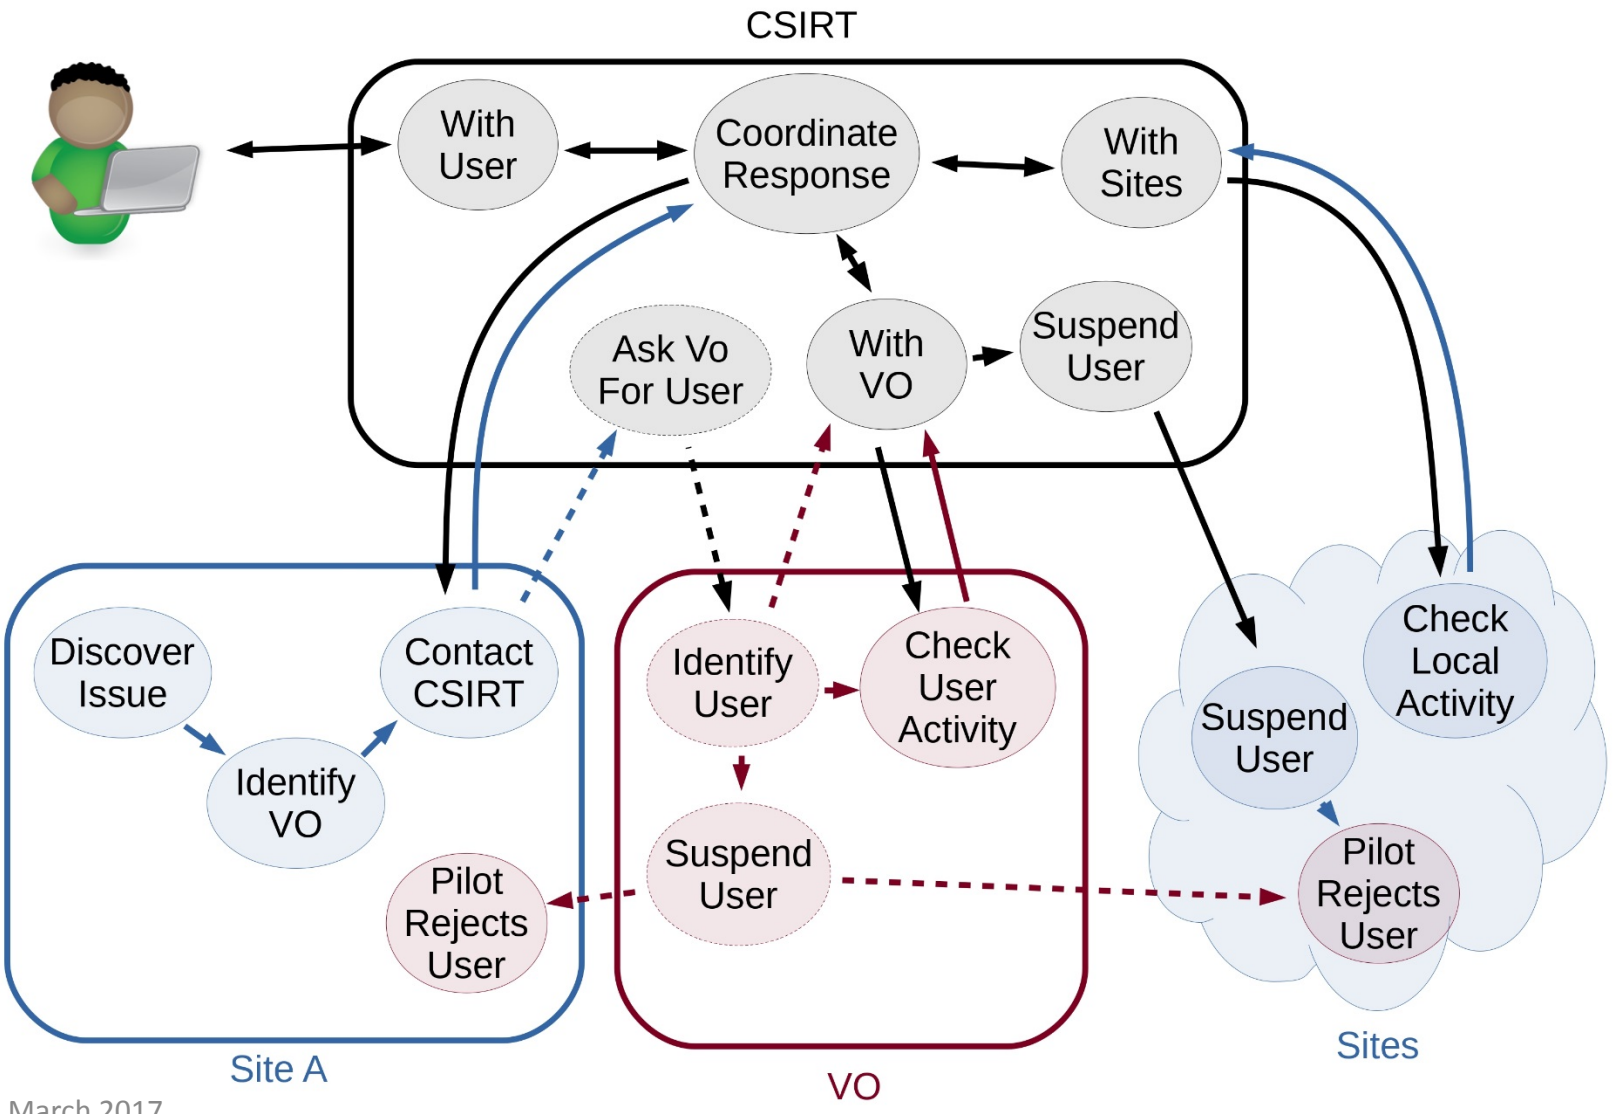

Proposed Incident Response

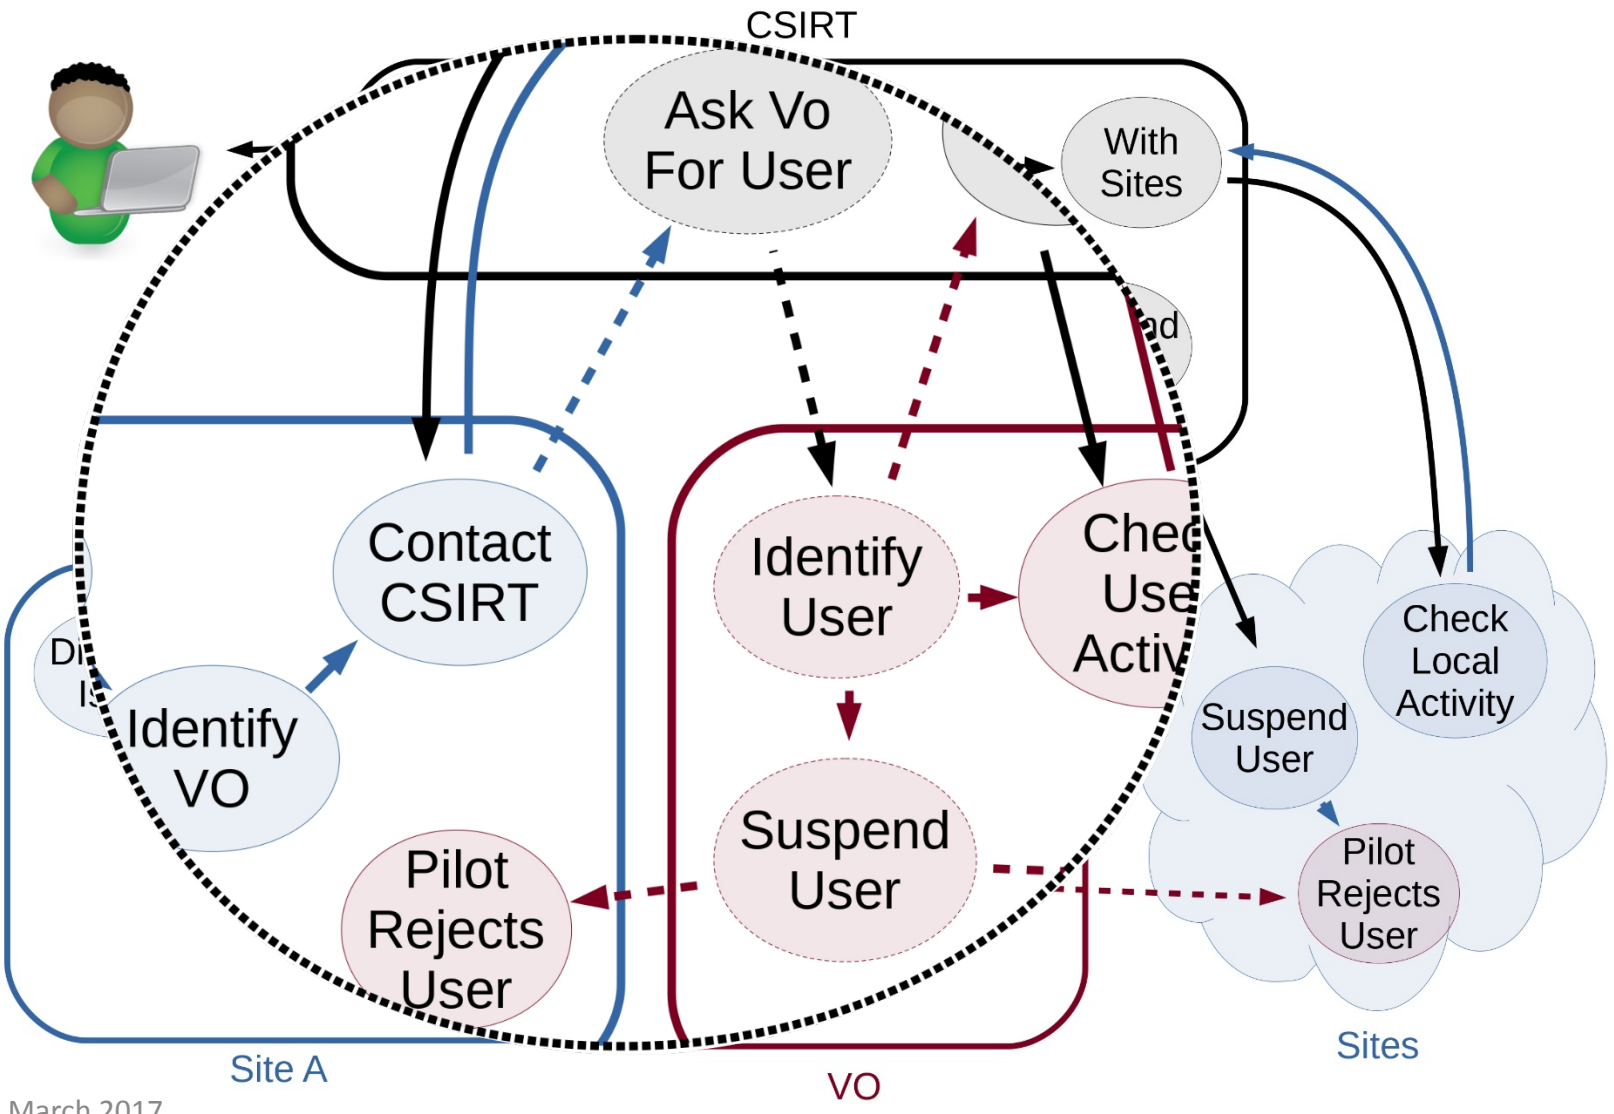

Traceability challenges

How to check if new model works?

- We can't risk to find issues during an incident
- VOs have performed a self-assessment
- Traceability challenges probably needed!

Data traceability?

- Proxies not required for compute traceability
- Is it still required for storage access?
- Discussion just started in WG:
  - Collecting VO data workflows for user jobs

Ongoing/future WG actions

- Testing Singularity
- Security review for Singularity
- Formalize data traceability model/requirement
- Traceability challenges & tests

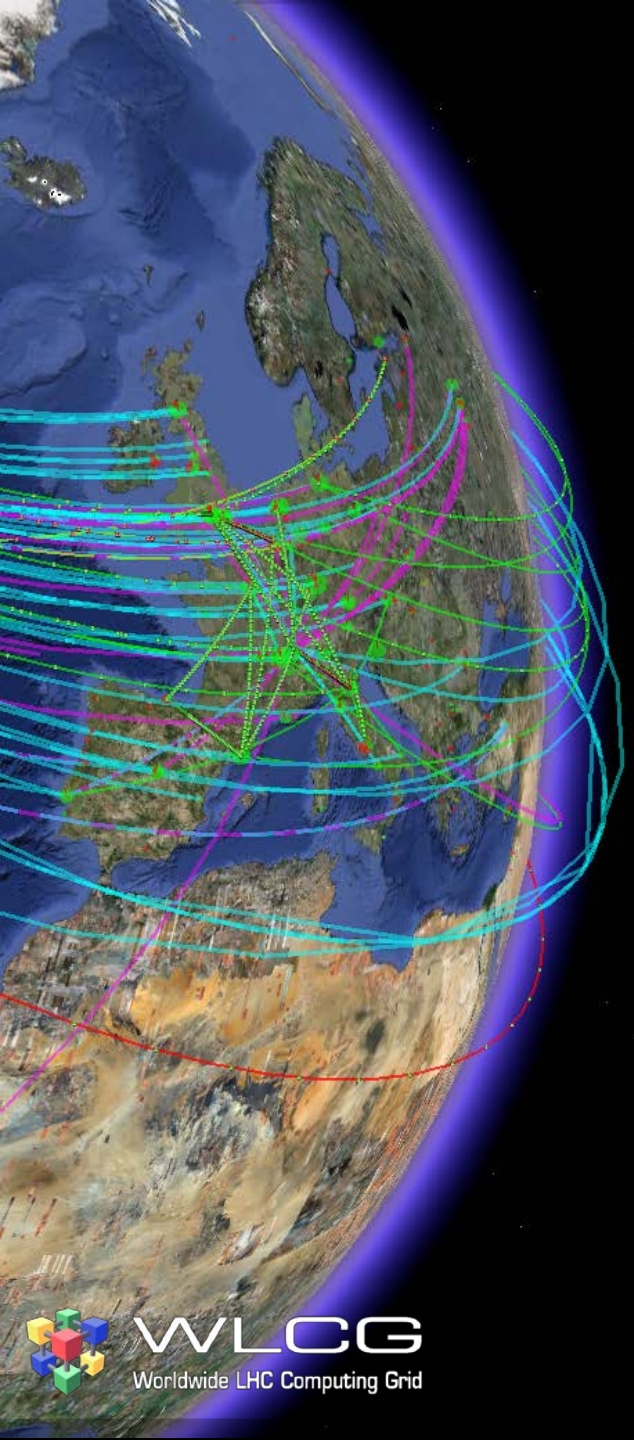

Thanks for your attention!

Any questions ?

WLCG
Worldwide LHC Computing Grid