Traceability & Isolation WG

Vincent BRILLAULT, CERN/EGI-CSIRT

Pre-GDB May 2016, CERN

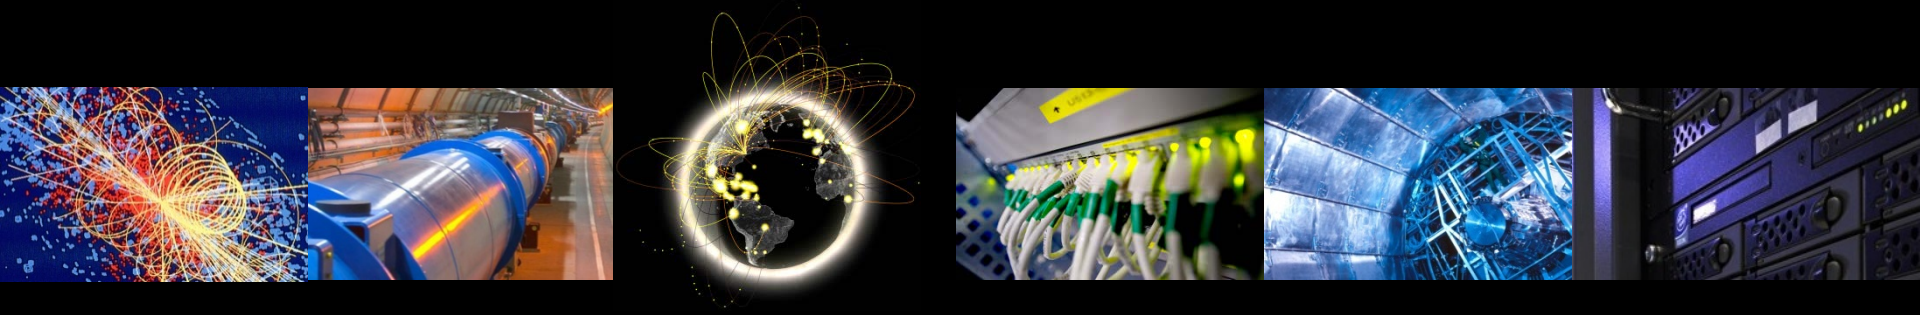

WLCG
Worldwide LHC Computing Grid

CERN

Welcome!

Kick-off meeting for the Working Group:

- Clarifying mandate, infrastructure & goal
- VOs current state of the art
- OSG recent evolution on traceability

➡ Agree on roadmap

WG Infrastructure

- wlcg-traceability-isolation-wg@cern.ch
- https://indico.cern.ch/category/4392/
- https://cern.ch/wlcg-traceability-isolation-wg
- “WLCGTraceabilityIsolationWG” share on https://cernbox.cern.ch

What is it?

Traceability

Knowing, per action:

- When
- What
- Where
- By Whom
- How

Isolation

Control interactions:

- System --- User
- User --- User
- System --- VO
- VO --- User
- User ---User

Why

Traceability

Resolve issues:

- Understand what really happened
- Prevent it from re-occurring

Isolation

- Un-expected "interactions":
  - Delete shared files
  - Write in the same file
  - Take over resources
- Impersonalization

Current issues

Traceability

- Blind areas:
  - VMs at sites
  - Commercial cloud
- Ignores completely VO framework

Isolation

- glExec frozen
- Ignore new tech:
  - Cgroups
  - Namespaces

Initial ideas

Traceability

- Trust VO framework
- Split traceability:
  - Site: Identify VO
  - VO:
    - User (who)
    - Payload (what)

Isolation

- Explore new tech:
  - Cgroups
  - Namespaces

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.

Example from CERN 2015: report

Dear Colleagues,

we are currently facing a DDOS attack (some physicist must have written a nice distributed script running on OpenStack to download in parallel data from {XXXXXXX}, which is also running at CERN).

Currently I have blacklisted:
{List of worker nodes}

but new machines are keep coming towards our host
{small except of new IPs}

Don't know if you can do something, or help us identify the responsible and educate him about how to gently write such type of bots.

Example from CERN 2015: our logs

18:13:16.877067+02:00 b635182c75 p:14022 s:7081 pp:13944 u:<X> g:1399 eu:<X> eg:1399 t:NULL /cvmfs/..../python2.7 TCP 188.185.194.218:60045 -> <tgt>:80

18:13:25.224751+02:00 b61e0aa520 p:7060 s:7410 pp:7038 u:<X> g:1399 eu:<X> eg:1399 t:NULL /cvmfs/..../python2.7 TCP 128.142.179.138:37998 -> <tgt>:80

18:13:36.701039+02:00 b62d21e061 p:22778 s:7037 pp:22761 u:<X> g:1399 eu:<X> eg:1399 t:NULL /cvmfs/..../python2.7 TCP 188.185.205.141:43095 -> <tgt>:80

18:13:43.671878+02:00 b62e9f240b p:12962 s:7267 pp:12945 u:<X> g:1399 eu:<X> eg:1399 t:NULL /cvmfs/..../python2.7 TCP 188.185.206.47:35350 -> <tgt>:80

18:13:53.306166+02:00 b6396015fb p:12458 s:7281 pp:12351 u:<X> g:1399 eu:<X> eg:1399 t:NULL /cvmfs/..../python2.7 TCP 188.185.202.224:41229 -> <tgt>:80

18:15:34.321037+02:00 b624cc96dd p:20855 s:7421 pp:20838 u:<X> g:1399 eu:<X> eg:1399 t:NULL /cvmfs/..../python2.7 TCP 128.142.46.90:37370 -> <tgt>:80

18:15:36.332852+02:00 b629fd8424 p:32127 s:7251 pp:32110 u:<X> g:1399 eu:<X> eg:1399 t:NULL /cvmfs/..../python2.7 TCP 188.185.201.160:46817 -> <tgt>:80

Example from CERN 2015: response

It seems that one of our batch users if DOSing (probably involuntarily) {XXXXXXX}.

I've attached the list of connection I've seen so far (still grepping).

The offending uid seems to be <X>.

Could you help us identify the offender and stop the "attack"?

Example from CERN 2015: now?

How will we solve such a case?

- Across multiple sites
- On VMs on multiple sites
- On commercial clouds

Milestone 0

Review the state of the art:

- What logs do we produce?
- How easy is it to investigate?

- How do we isolate users & pilot jobs?
- What models are pilot jobs using?

Questions to be discussed

- Granularity of user action tacking?
  - Storage/service access? Network actions?
  - Is the payload/job enough? Obfuscation?
- Users/pilots isolation requirements?
  - What do we need to protect?
- VO logs requirements?
  - Do we have enough with current logs?