



# ATLAS Rucio AuthN service incident of Aug 20

WLCG Operations Coordination meeting

2<sup>nd</sup> September 2021

Maarten Litmaath

v1.0

# Incident timeline

- SIR: <https://twiki.cern.ch/twiki/bin/view/LCG/RucioAuthSvcInc20210820>
- Fri Aug 20, 22:15 CEST – the Rucio AuthN service production hosts were dropped from their firewall set → no longer reachable from the internet → massive grid failures etc.
- 23:30 – SNow and GGUS alarm tickets opened
- DC operator mistakenly tried to contact the IT-DB group rota (see next page)
- Several experts in ATLAS and IT kept debugging the situation and found the right (Config Mgmt) team to be contacted for expert advice
- E-mails were sent to them, but an SMS was not (see next page)
- 03:55 – workaround implemented, main problem cured
- Sat morning – explanation from the Config Mgmt team why the problem occurred and how to get the dropped hosts included again
  - Puppet had been disabled for more than 30 days → the hosts dropped out of the Puppet DB from which the affected firewall set is built → fix: run Puppet

# Lessons learned and followup

- Review DC operator alarm procedure
  - Allow the operator to re-assess the correct response after new information has become available
    - Would have led to an SMS for the Config Mgmt team in this case
  - Remove ambiguity that led to the wrong initial response this time
  - Check if out-of-hours support is implemented consistently across services and groups
- Disseminate wider the behaviour of Puppet DB and that leaving Puppet disabled is rarely a good idea ([firewall doc](#))
- Investigate ATLAS suggestions for firewall set changes
  - Warnings about imminent removals → look possible
  - Changes (removals) only during working hours → probably not