



LHCONE security and use of Jumbo frames

WLCG Management Board

16th April 2024

edoardo.martelli@cern.ch

LHCONE meeting #52

Hosted by INFN Catania (IT) on the 9-10 of April 2024

Agenda: <https://indico.cern.ch/event/1349135/>

Co-located with

- GEANT Next-Generation-Networks workshop
- SKAO-LHCONE joint meeting

Next: meeting September-October 2024 in Beijing,
hosted by IHEP Tier1

LHCONE meeting #52

Two items need to be discussed with WLCG:

- IP network prefix tagging for LHCONE security (MultiONE)
- Use of Jumbo frames at sites

LHCONE security (MultiONE)

Trust of LHCONE

The major benefit of LHCONE is the trust in the connected sites: it allows the **LHCONE fat links to bypass slow and expensive security inspection**

Due to the inclusions of other collaborations (BelleII, DUNE...), the increasingly growing number of connected sites may reduce the trust

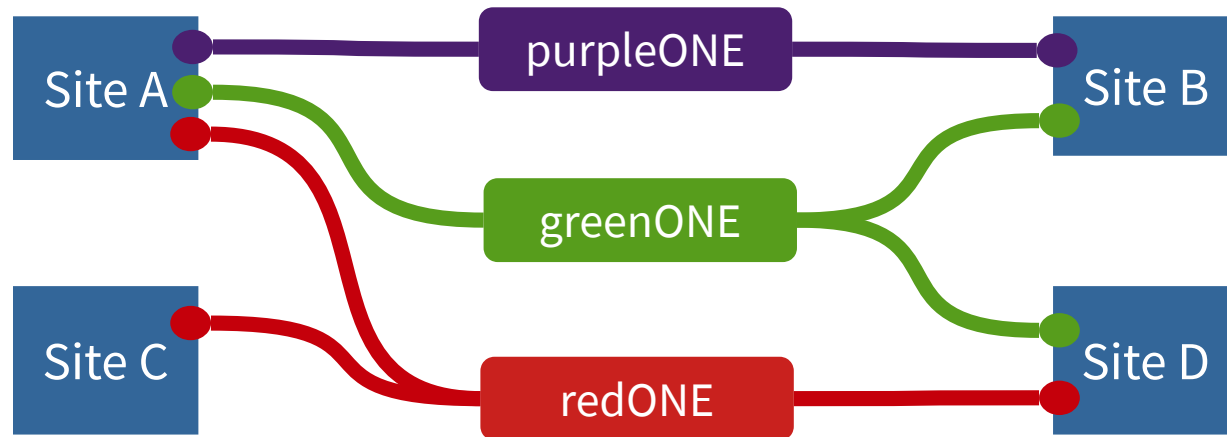
The MultiONE project aims to reduce the exposure of the sites, so to increase the trust in LHCONE

MultiONE as multiple LHCONE

Already proposed solutions based on the replication of LHCONE, one VPN per each collaboration

Major challenges:

- Difficult to route the traffic into the correct VPN
- Operationally complex for small sites, but also for NRENs



New proposal

Don't add any additional VPN (or maybe just another one for “Other-Big Sciences-beside-LHC”?)

Each prefix announced to LHCONE is tagged with BGP communities* that identify the collaborations served by the site

The tagging is done by the sites. Or by the connecting Network Provider, if a site is unable to do it

Later sites can decide to accept only the prefixes of the collaboration they are working with


** BGP is the routing protocol used in LHCONE. BGP communities are tags that can be added to the network prefixes announced by the sites to LHCONE*

Practical example

Tagging

Each LHCONE site:

- tags its prefixes announced to LHCONE with all the BGP communities that identify the collaborations the site is participating in




Practical example

Filtering

Each LHCONE site:

- accepts only the prefixes tagged with the BGP communities of its own collaborations



Benefits

- Reduced exposures of sites when filtering is implemented
- No additional VPNs to configure
- No changes at sites when a new site connects to LHCONE
- Filtering can be activated progressively
- Tags are in any case useful to document the use of the network and to double check what is declared in CRIC

Limitations

Not 100% secure:

- any sites will still be able to send packets to sites that tag and filter.
However TCP connections will not work.
- a malicious sites can tag its prefixes with all the existing tags and get the prefix accepted. This could be mitigated if Network Providers do the tagging and/or filtering

Implementation

Milestone: all prefixes tagged by LHCONE meeting #54 (Spring 2025)

Implement filtering at (some) WLCG sites during year 1 of LHC LS3 (2026)
and in time for DC26

Present the proposal to WLCG bodies (today, Hamburg)

Reach out all the LHCONE sites and request to implement the tagging,
while reviewing their prefix declarations in CRIC

Monitor the progress of the tagging in the LHCONE routing tables

Use of Jumbo frames

Jumbo frames

IPv4 and IPv6 frames have a standard maximum size of 1500 Bytes (IP header + payload)

Jumbo frames: whatever maximum size bigger than 1500B. 9000B is the most used value in the R&D community

Reducing the relative size of the IP header over the payload can reduce the load on the CPU of the sender of large data flows, thus **allowing greater throughput** for CPU intensive transfers

On the other hand, transfers between hosts **using different MTUs can lead to traffic blackholing if the networks in between are not properly configured**

Jumbo frames at LHCONE sites

Benefits are evident on long distance transfers, less on the short distance

Operational issues are also evident, but they can be mitigated by sharing deployment experiences

Agreed to put effort on testing at larger scale. CERN (historically reluctant) has agreed to push the testing on some production servers

Proposed actions

Present the activity to WLCG forums

Perform thorough testing focusing on performance gains and operational issues at large sites

Aims to large scale deployment before next Data Challenges:

- DC26: X% of sites with X close to 50
- DC28: 100% of sites

Questions? Comments?

edoardo.martelli@cern.ch

