



# Improving LHCONE security & Use of Jumbo frames

WLCG workshop – DESY, Hamburg

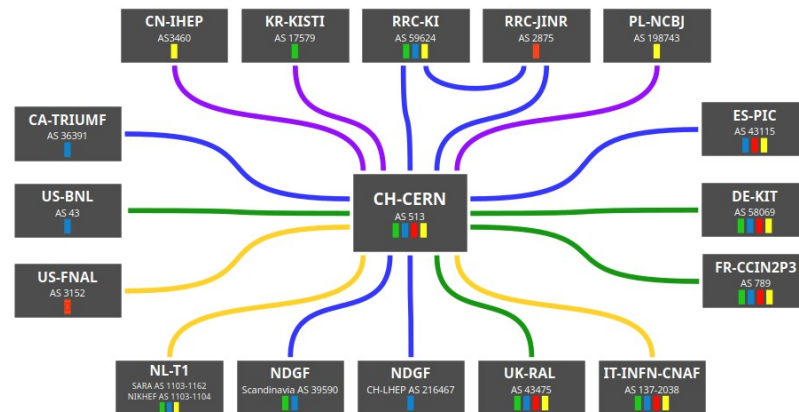
14<sup>th</sup> May 2024

[edoardo.martelli@cern.ch](mailto:edoardo.martelli@cern.ch)


Quick update

# LHCOPN - update

- 2.66Tbps of aggregated bandwidth to the Tier0
- Traffic stats: moved 619PB in the last 12 months. +27% compared to previous year
- CN-IHEP, new LHCb Tier1: connected at 20Gbps via CSTnet and GEANT
- NDGF-IHEP, new NDGF distributed Tier1 site at UniBern (CH): connected at 100Gbps
- TW-ASGC terminated Tier1 activities. All links decommissioned
- NLT1 tested 800Gbps connection
- IN2P3 upgraded to 2x100G, FNAL to 400Gbps



LHCOPN Total Traffic (CERN → T1s)



# LHCONE



## News

- DC24 impact well visible by all the LHCONE providers
- New LHCONE network providers:
  - SWITCH Switzerland (for University of Bern-LHEP)
  - FCCN Portugal (for NCG-INGRID-PT)

## ESnet Trans-Atlantic upgrades

- Now In Production:
  - 400G New York - London
  - 400G Boston - CERN
- Currently underway:
  - 400G Boston - London
- Trans-Atlantic capacity targets:
  - 3.2T in 2027, in advance of Run 4

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

# New proposal: prefix identification

Don't add any additional VPN

**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 on, sites can decide to drop prefixes of collaborations they are not working with

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

# Benefits

- Simple and commonly used technique, no additional VPNs to configure
- Tags are useful to document the use of the network and to double check what is declared in CRIC
- Reduced exposures of sites when filtering is implemented
- Tagging and Filtering can be implemented progressively
- No changes at sites when a new site connects to LHCONE



# 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 validate the tagging

# Implementation

- 1 - Reach out all the LHCONE sites and request to implement the tagging, while reviewing their prefix declarations in CRIC
- 2 - Monitor the progress of the tagging in the LHCONE routing tables

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

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

# Documentation

<https://twiki.cern.ch/twiki/bin/view/LHCONE/MultiOneBGPcommunities>

## Community format:

- Standard BGP community
- **Format: 61139:ExpID**
- The ExpIDs are defined by the SciTags initiative

## Defined Communities

ALICE	61339:5
ATLAS	61339:2
BelleII	61339:6
CMS	61339:3
DUNE	61339:8
JUNO	61339:12
LHCb	61339:4
NOvA	61339:13
Pierre Auger	61339:11
XENON	61339:14
perfSONAR servers	61339:60001
LHCONE backbone	61339:60002
Demo/Prototype	61339:60003

# Use of Jumbo frames

# Jumbo frames at LHCONE sites

Benefits of jumbo frames 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

Jumbo frames deployment:

- Perform thorough testing focusing on performance gains and operational issues at large sites (CERN IT will engage on this)
- Aiming to large scale deployment before next Data Challenges:
  - DC26: 50% of sites
  - DC28: 99% of sites(if proved to be useful)

*Questions? Comments?*

*edoardo.martelli@cern.ch*



Additional information




# 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



# Jumbo frames: what and why

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 survey

The LHCONE WAN networks have been supporting Jumbo frames at 9000B for many years

According to an informal survey among LHCONE sites (53 answers):

- 26 sites have Jumbo on the storage
- 19 sites of the 26 have Jumbo also on worker nodes
- 4 sites of the 26 didn't have any problem with implementing Jumbo
- 10 sites of the 20 without Jumbo don't want to implement Jumbo