

# multiONE

LHCOPN-LHCONE meeting at CERN  
14<sup>th</sup> of January 2020



**LHCONE limitations and the need for multiple “ONEs” (multiONE)**

# LHCONE

LHCONE is a worldwide Virtual Private Network (VPN) implemented by Research and Education Network providers (RENs)

Original AUP:

- **Only WLCG Tier1/2/3 sites can connect to LHCONE**
- At sites, only resources dedicated to WLCG can use LHCONE

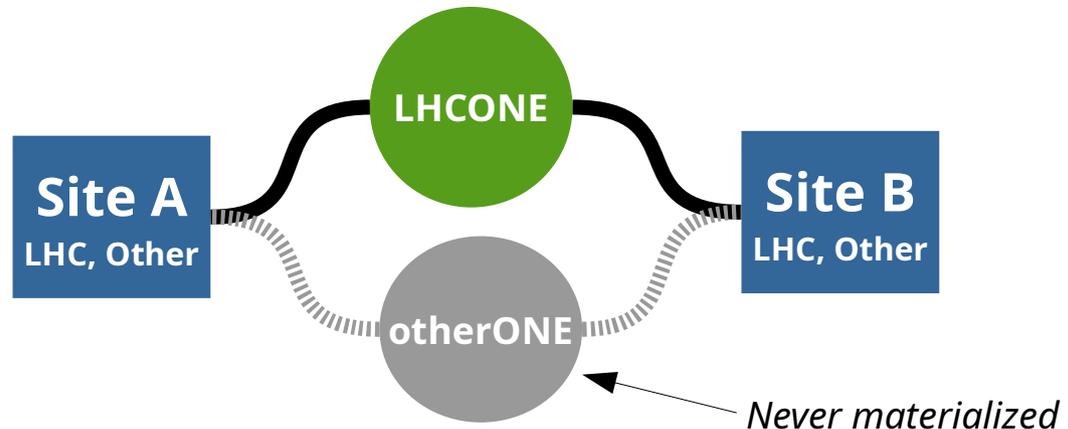
 Main advantage: **Sites can trust LHCONE to be safe and plug it directly into their datacentre**, bypassing bottlenecks (e.g. expensive security equipment)

# Adding more Collaborations

Other Collaborations would benefit of having their own “ONE”

- in fact, defining a new VPN is relatively simple for RENs;
- but **it's difficult for Sites participating to multiple collaborations to put the traffic in the corresponding VPN**

Thus, over the years, few HEP collaborations (Pierre Auger, XENON, BelleII, NOVA...) have simply joined LHCONE instead of building their own VPN



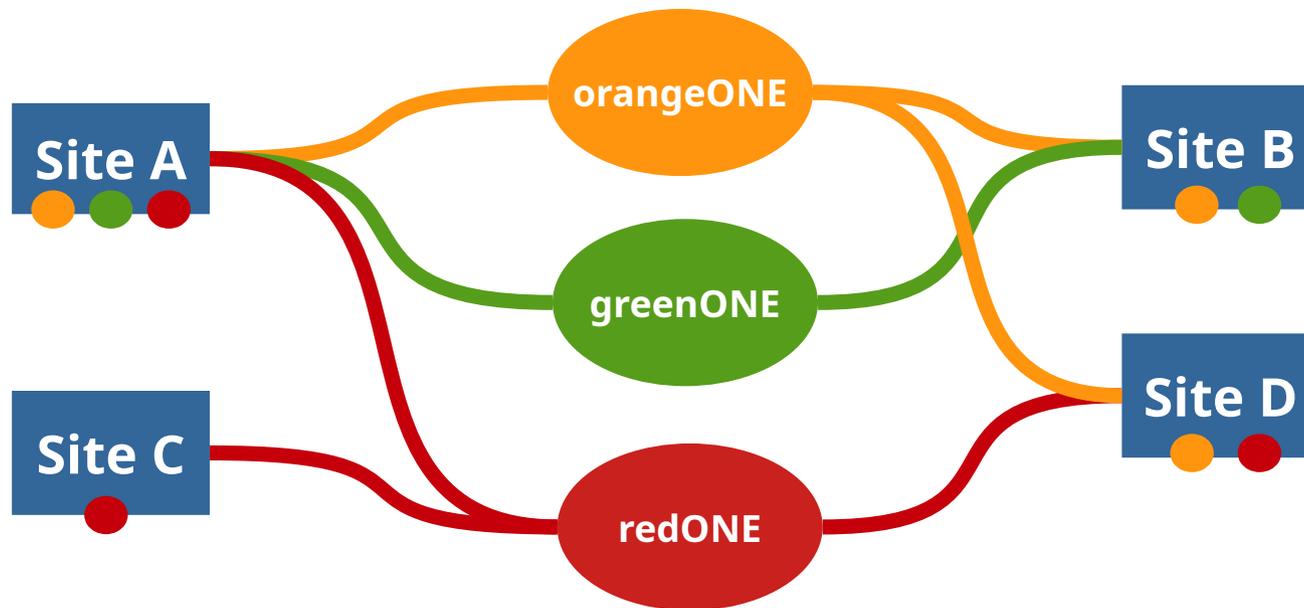
# Problems with just adding collaborations

- The more sites join LHCONE, the **less trustable** it becomes
- The more the traffic volume grows in a single domain, the **more difficult for RENs to shape the load** in their networks
- Funding agencies would prefer to have a clear distinction of **who is using the resources they fund** (in fact it was not always straightforward to accept new collaborations in LHCONE)

# multiple “ONEs”

A solution would be to implement a VPN for each Collaboration:

- Each site joins only the VPNs it is collaborating with, to reduce the exposure of their data-centre/Science-DMZ
- Each Collaboration funds its own VPN



# ...but there are issues with multiple VPNs

- Difficult to select what VPN to use for a Site that serves multiple Collaborations
- Even more difficult if the different Collaborations share the same servers and applications
- The simpler solution (static segregation of resources) is rather inefficient

# About multiONE

The growing LHCONE issue was discussed several times at LHCONE meetings

It has been agreed to start a project to verify if it is possible to use multiple VPNs for sites that participate to several science Collaborations

It's needed to identify a collaboration to prototype a working solution

# **DUNEONE proposal**

# DUNEONE proposal

It is proposed to build a VPN similar to LHCONE to connect protoDUNE and DUNE sites that are already connected to LHCONE, to allow those sites to prototype and test technical solution to correctly separate the traffic between the two VPNs. A two-phase project is proposed:

- Phase 1:** Migration of ongoing data transfers of the pre-processed data generated by the CERN-based protoDUNE detector(s) to FNAL (DUNE T0) archive facilities. This data movement is currently being carried over the LHCOPN.
- Phase 2:** Selective migration of Rucio-based data movement for DUNE's emerging distributed data storage facilities. This would be implemented on a site-by-site basis, as individual DUNE sites elect to participate in the project. Phase 2 would commence after satisfactory demonstration of proof-of-concept in the Phase 1 testing & evaluation, and consultation with the DUNE collaboration.

The tests will be structured in a manner to not disrupt production traffic. It also should be emphasized that **this project is targeted at proof-of-concept, not establishing a DUNE-wide service.**

# Involved parties

## Sites:

Phase 1 - CERN & FNAL

Phase 2 - DUNE sites involved in Rucio-based data movement

## Network providers:

Phase 1 - ESnet

Phase 2 - GEANT & NRENs that support DUNE data storage sites (as necessary)

# Examples of possible solutions

Here is a list of a few possible solutions that could allow the traffic separation. The list is not exhaustive and doesn't exclude other ideas. **Individual sites can use different solutions, as long as the objective of routing traffic into the correct VPN is achieved.**

- A) Dynamic, software defined allocation of computing resources: using network virtualization techniques (VXLAN, EVPN, Linux Name Space...), create virtual groups of computing resources that interface with the network to separate the traffic at the sources.
- B) Paired use of source-destinations address: making use of secondary IP addresses, computing resources should make sure to use appropriate source-destination address pairs, to allow simple destination routing.
- C) Packet tagging: using available fields in the IP header (DSCP, flow label..) set by the applications, make the network policy base routes those packets in the correct VPN.

# Requirements

In case of solutions based on mapping by IP addresses, it will be necessary for every site to allocate dedicated IP prefixes to the different collaborations. Since IPv4 is a scarce resource, **it is envisaged that the proposed solutions have to work with IPv6.**

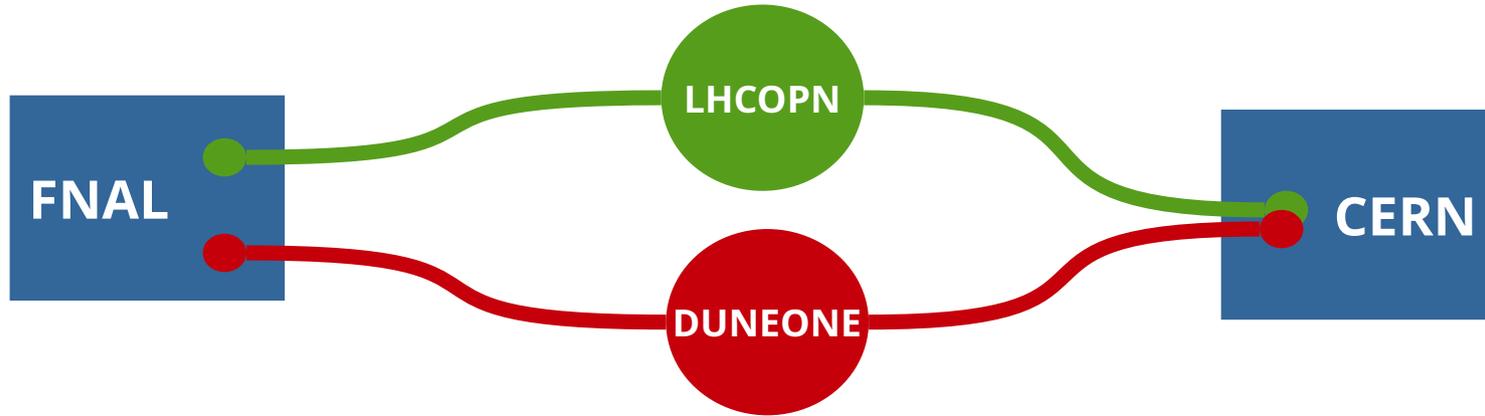
# Milestones

- 1 - Implementation of DUNEONE VPN by ESnet
- 2 - Successful migration of CERN  $\Leftrightarrow$  FNAL protoDUNE traffic over to DUNEONE
- 3 - Implementation of DUNEONE VPN by GEANT (and other NRENs, where necessary)
- 4 - Connection of participating DUNE storage sites to DUNEONE
- 5 - Evaluation at scale of DUNE data movement over DUNEONE

# Additional benefits

Clean up of the CERN-FNAL protoDUNE traffic, currently squatting the unused bandwidth of the FNAL LHCOPN connection to CERN

# DUNEONE prototype



# Conclusions

# Why do we want to do it?

- No urgent need nor specific request right now
- However size of LHCONE is already at its limits
- We need to be prepared when the next major collaboration (SKA?) will need its own ONE

# Next steps

- Explore more technical implementation
- Prototype different solutions
- Test solutions with a site that serves multiple experiments  
(e.g. with DUNEONE between CERN and FNAL)

*Questions?*

*edoardo.martelli@cern.ch*