



# **LHCOPN-LHCONE meeting #48**

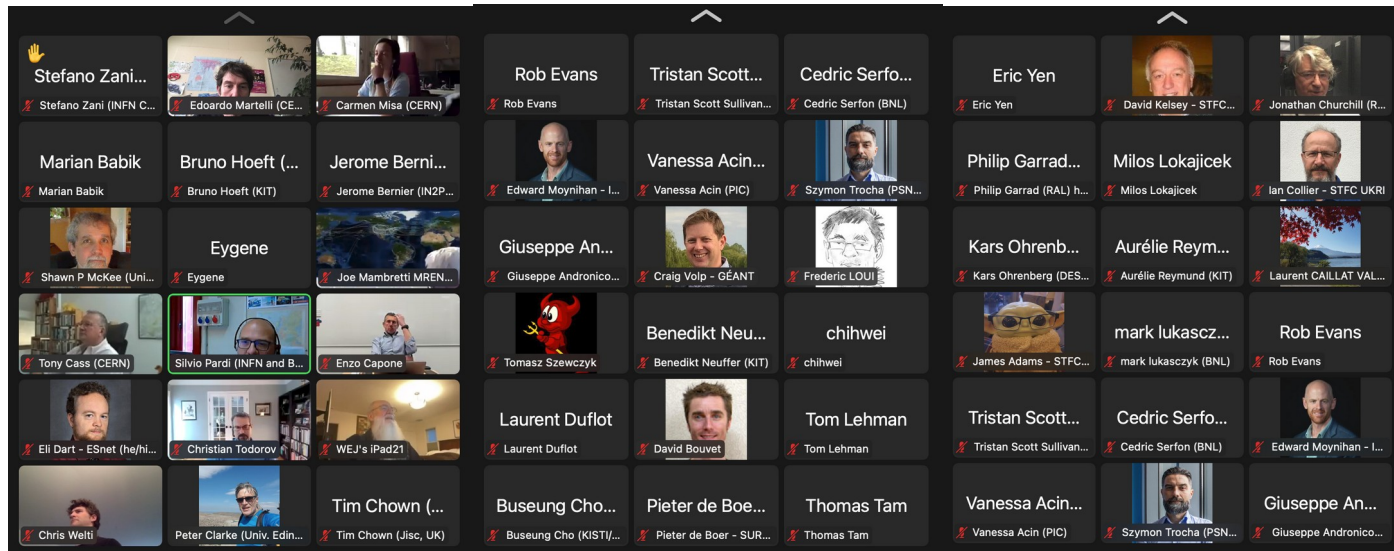
## **summary notes**

31 March 2022 – v1.1

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


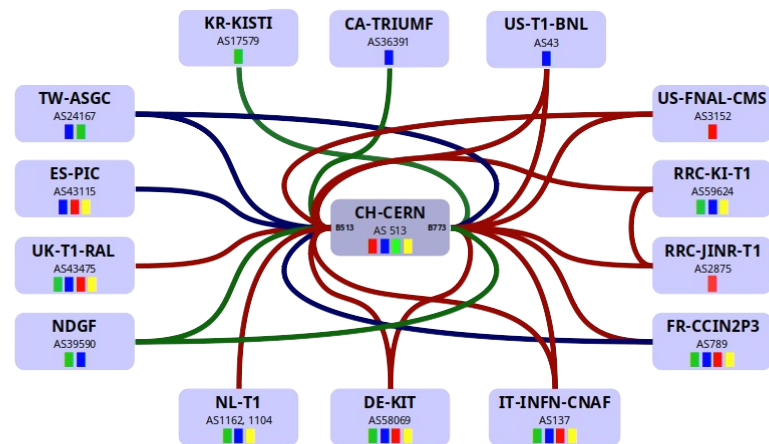
# Venue

- 29-30 of March 2022
- Two sessions of 3 hours in two days
- On video conference only, for the fifth time
- ~48 participants each day
- Agenda at <https://indico.cern.ch/e/LHCOPNE48>



# LHCOPN - update

- No major changes since last meeting. Total of 1.3Tbps from the Tier0 to the Tier1s
- Traffic stats: moved 433PB in one year. +60% compared to previous year.  
Reached again traffic level of end of Run2
- Upcoming links: PIC 100G (waiting for GEANT upgrade, using 50Gbs on LHCONe for now)
- FNAL access now 4x100G
- DE-KIT will announce two new IPv4 prefixes (157.180.228.0/22 and 157.180.232.0/22)
- Construction of new CERN data-centre (PCC) has started.  
Should be ready by mid 2023



WLCG data Challenge  
~10 Oct 2021



# LHCONE L3VPN - update



## Monitoring

- New LHCONE/OPN monitoring at CERN based on Grafana

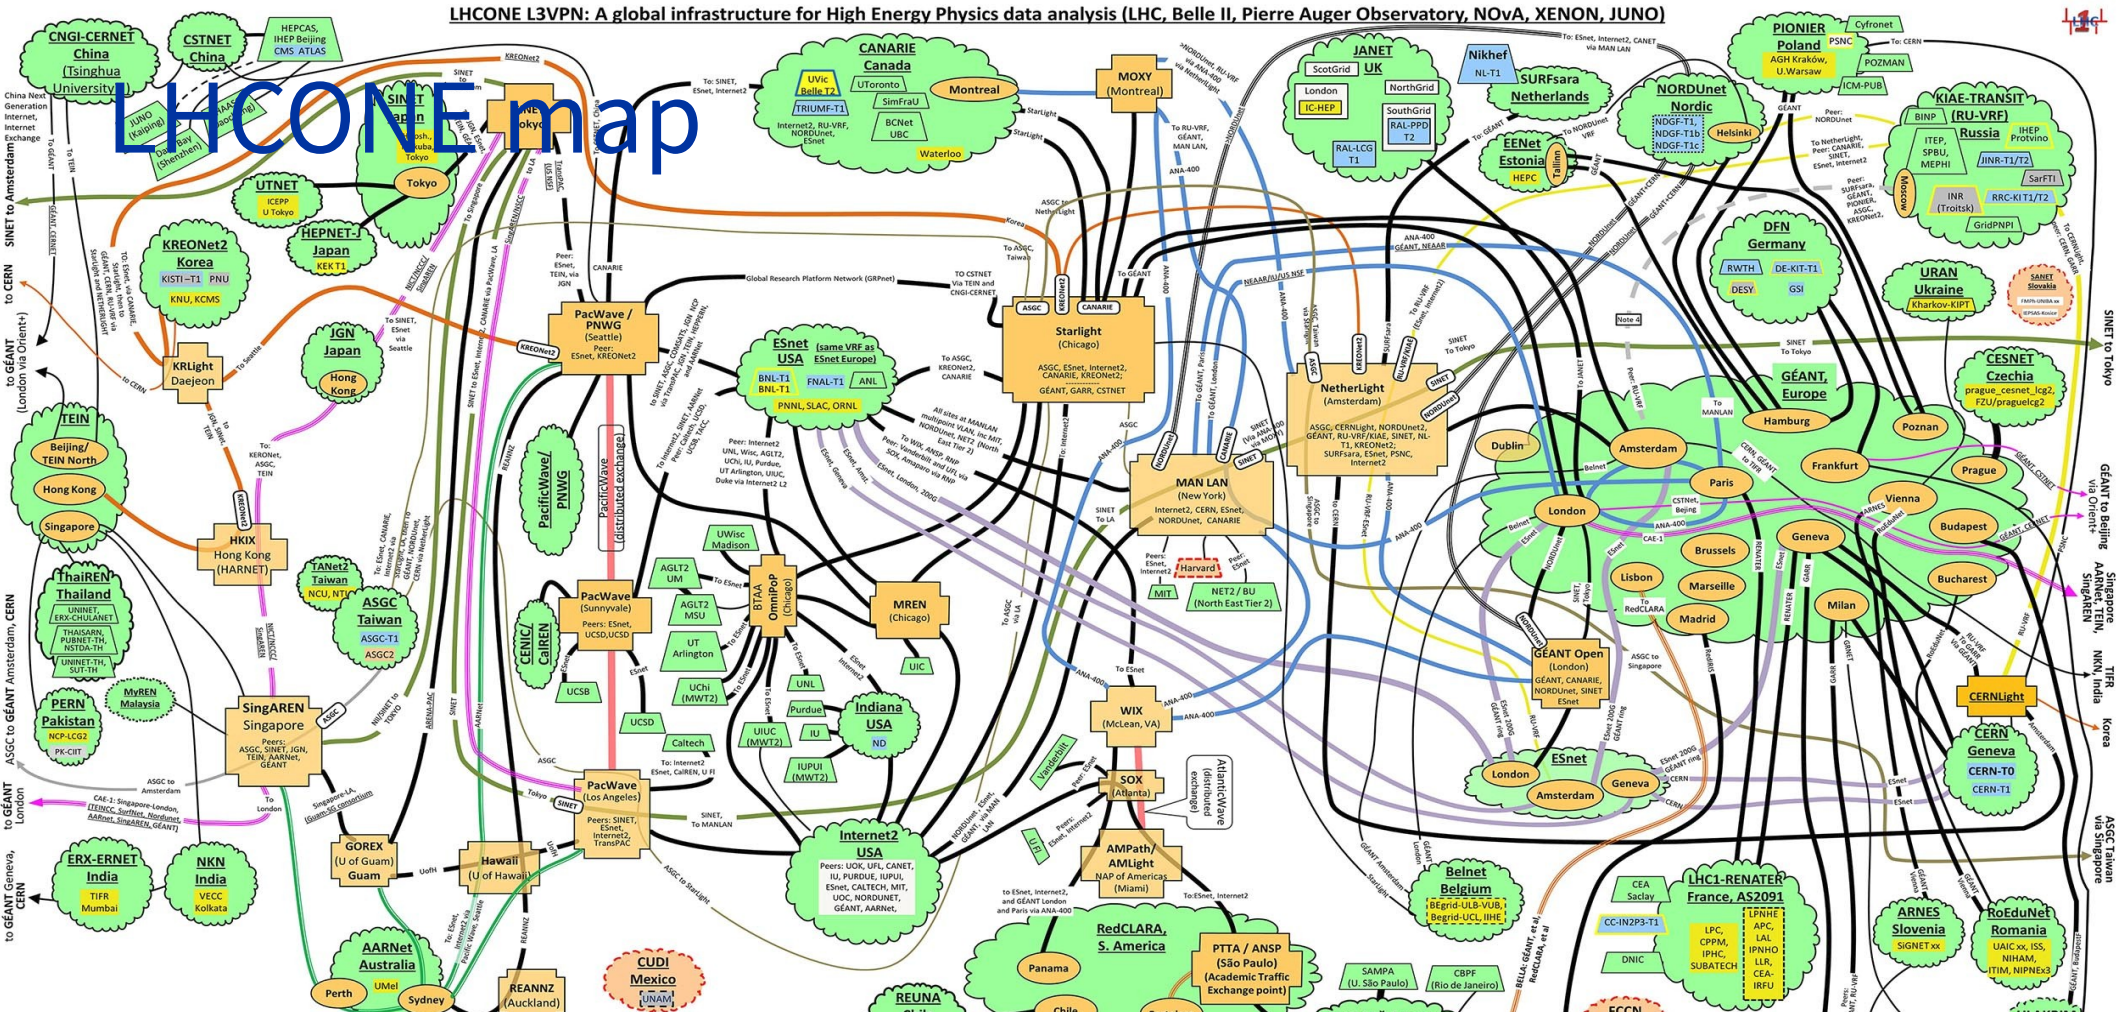
## Traffic statistics:

- ULAKBIM (Turkey) has started exchanging traffic
- large traffic increase in all regions
- High spikes during October 2021 and March 2022 WLCG challenges
- Ukraine traffic very low since start of the war (GEANT PoP in Kyiv disconnected for two weeks)





# LHCONE map



LHCONE Map Ver. 5.7 (draft 1), March 9, 2022 – WEJohnston, ESnet, wej@es.net

<ul style="list-style-type: none"> <li><b>GREEN</b> LHCONE VRF domain/aggregator - A provider network.</li> <li><b>ORANGE</b> Connector network - provides, e.g., an L2 path between VRFs.</li> <li><b>YELLOW</b> Provider network PoP router</li> <li><b>RED</b> Not currently connected to LHCONE</li> <li><b>ORANGE</b> Exchange point</li> </ul>	<ul style="list-style-type: none"> <li><b>SINET</b> AREN/site router at exchange point</li> <li><b>Communication links:</b> 1/10, 20/30/40, and 100Gb/s or N x 100G</li> <li><b>Underlined link information</b> indicates link provider, not use</li> <li><b>Dotted outline</b> indicates distributed site</li> <li><b>Blue dashed outline</b> indicating a WLCG federation site not currently on LHCONE</li> </ul>	<p><b>International infrastructure by provider/collaboration</b></p> <ul style="list-style-type: none"> <li><b>various</b></li> <li><b>GREEN</b> AARNet, Japan</li> <li><b>ORANGE</b> GEANT</li> <li><b>YELLOW</b> SINET, Japan, global ring</li> <li><b>PINK</b> ESnet transatlantic, USA</li> <li><b>PURPLE</b> NICT/NCCC/SingAREN</li> <li><b>GREEN</b> SINET</li> <li><b>ORANGE</b> NORDUnet</li> <li><b>YELLOW</b> KIAE, Russia</li> <li><b>PINK</b> KREONet2, Korea</li> <li><b>ORANGE</b> BELLA: GEANT, et al, RedCLARA, et al</li> </ul>	<ul style="list-style-type: none"> <li><b>GREEN</b> SINET</li> <li><b>ORANGE</b> NORDUnet</li> <li><b>YELLOW</b> KIAE, Russia</li> <li><b>PINK</b> KREONet2, Korea</li> <li><b>ORANGE</b> BELLA: GEANT, et al, RedCLARA, et al</li> <li><b>GREEN</b> Belle II Tier 1/2</li> </ul>	<ul style="list-style-type: none"> <li><b>PNU</b> LHC ALICE or LHCb site</li> <li><b>CNAF-T1</b> LHC Tier 1 ATLAS and CMS</li> <li><b>UCHI</b> LHC Tier 2/3 ATLAS and CMS</li> <li><b>KEK</b> Belle II Tier 1/2</li> <li><b>AUNL</b> Sites that are standalone VRFs</li> </ul>	<p><b>NOTES</b></p> <ol style="list-style-type: none"> <li>ONLY links involved in LHCONE are shown</li> <li>LHCOPN links are not shown on this diagram</li> <li>For map explanation see "Interpreting the LHCONE Map" at <a href="https://www.dragonbox.com/~/media/588/1/2/1/20220309/Interpreting%20the%20LHCONE%20Map.pdf">https://www.dragonbox.com/~/media/588/1/2/1/20220309/Interpreting%20the%20LHCONE%20Map.pdf</a></li> <li>From Enzo Capone, 4Mar2022: want to inform you that today at 17:00 GMT, GEANT has shutdown the peering between our VRF and KIAE, as per the decision of the GEANT Board of Directors. Mar. 2022, *</li> </ol>
--	---	--	---	--	---

# LHCONE monitoring - update



- perfSONAR 5: beta soon out. Now using Elasticsearch and Grafana
- 4.4.3 released in October, but still has problems hitting resource limits on busy nodes. Fixes will be in 4.4.4, soon out
- 100G mesh:
  - task force meeting every month to improve results
  - increase of TCP buffers to 1GB has improved performances
  - 25 100G servers available, more coming
  - new Grafana dashboard available
- New home page for monitoring tools
- AAAS: developing detection of bandwidth and path changes

# BelleII update



Currently in data taking. Collected 2.8PB of data since 2019

Good performance over LHCONE and General Internet

Long shutdown scheduled for late 2022

On-going:

- Migrating to DAVS protocols for data transfers
- Working also on Token authentication

Interested in participating in the packet marking activity



# JUNO data challenge



Data generated at IHEP (CN), then sent to INFN-CNAF (IT) and from there to CCIN2P3 (FR) and JINR (RU)

Expected volumes: 2PB/y raw data, 200TB/y reconstructed data

Data challenge organization on-going.

Goal: check transfer functionalities and performances, including network; set up proper monitoring

Previous data challenge has shown good results, with transfers around 8Gbps

Dashboard available ([link](#))





# Monitoring LHC data challenge with Stardust

Stardust: a new monitoring framework developed by ESnet



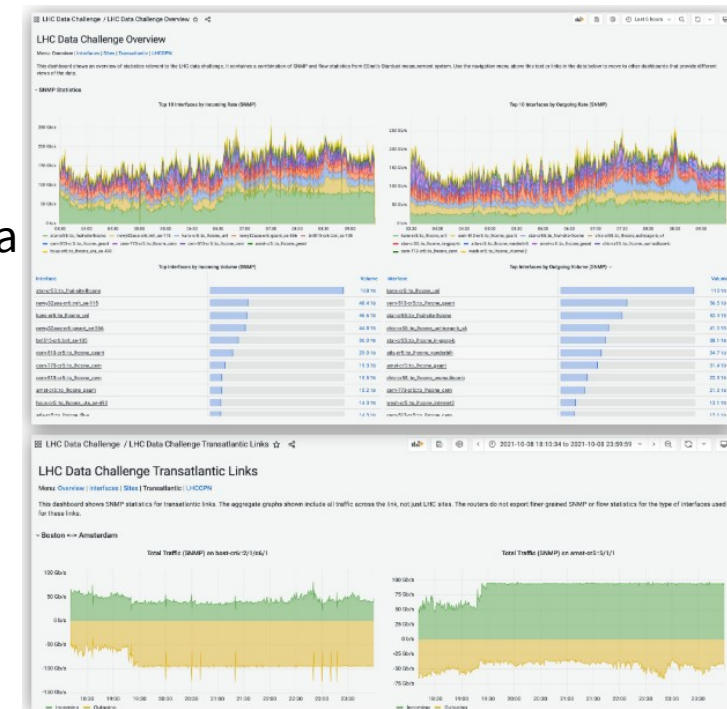
Designed to support multiple types of sources of measurements and metadata

Derived by NetSAGE and integrated with Open Source components

Allow access via web dashboards and APIs, but also to raw data

Designed to “tell stories” and visual explain data flows

Set up dashboard [\(link\)](#) for October’s WLCG data challenge using ESnet data



# Impact of war in Ukraine



## **Open discussion:**

- CERN hopes to not have to break the collaboration with Russia. Things have been made difficult by the statement made by Kurchatov supporting the invasion
- That statement doesn't seem to reflect what people working there think, though
- CERN is supporting the Russian physicists working in Geneva. There is lot of concern on the impact on them, if the collaboration would stop

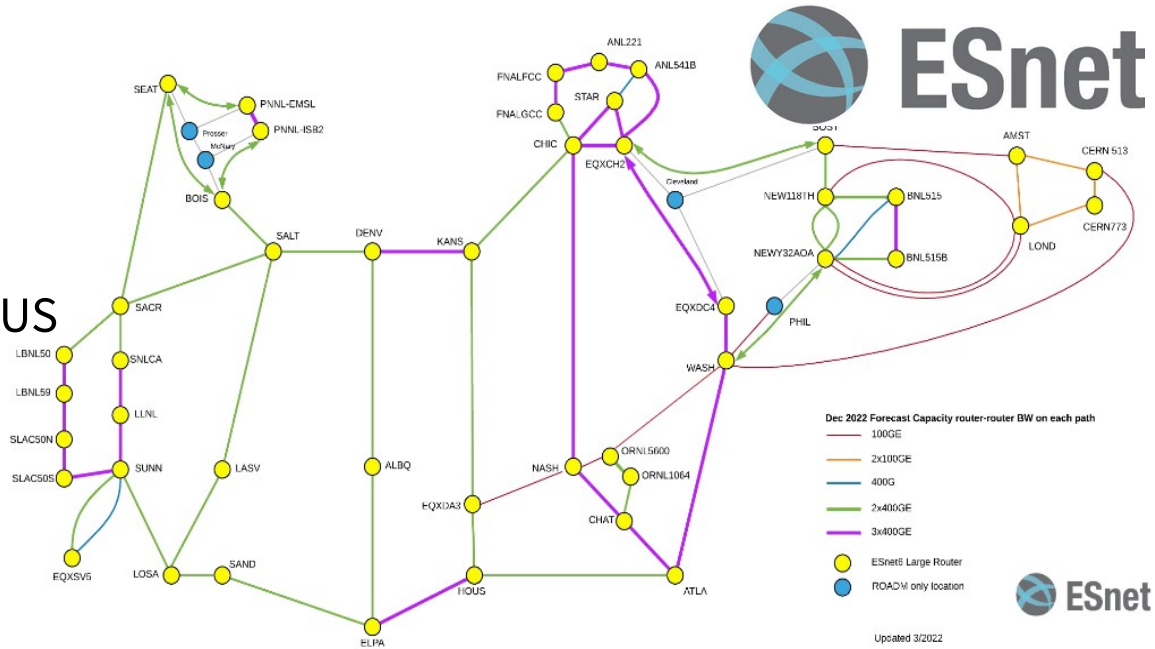
## **Observations on the network:**

- GEANT and CANARIE have stopped their LHCONE peerings with Russian LHCONE. General Internet peerings all up
- Large flow of data transfers for 18 days from Kurchatov; most likely data repatriation
- GEANT PoP in Kyiv disconnected several times because of long lasting fibre cuts
- Kharkiv Institute Tier2 is offline since 11th/3

# ESnet update

## ESnet6 network upgrade:

- Deployed 74x 400G optical links in the US
- Deployed 42x routers
- Most of provisioning done using new automation stack



## LHCONE

- All LHCONE connections in North America now on new ESnet6 routers

## Transatlantic

- Acquired 2x 400Gbps links on Amitié cable system
- available late 2022



# 400Gbps upgrades for data challenges in 2023

## **WLCG data challenge 2021**

- Achieved expectations (10% of HL-LHC)
- Network not saturated, but stressed especially at exchange points
- Current 100Gbps connections inadequate for future challenges

## **WLCG data challenge 2023:**

- 30% of HL-LHC
- Schedule not changed so far, but following challenges may be delayed
- 100Gbps links at exchange point will be upgraded to 400Gbps

Run3 is starting soon and traffic is already ramping up. LHCONE providers are monitoring the situation and planning necessary upgrades

# NetSage for LHC




NetSage is a tool to make network monitoring stats easily accessible and understandable

- It can ingest data from many sources: SNMP counters, net/s/flow, perfSONAR...
- It is committed to privacy, fully GDPR compliant

Setting up pilot for LHCONE and LHCOPN

- <https://lhc.netsage.global/>
- Already showing data from CERN  
LHCOPN/ONE border routers
- Looking for other LHC network operators  
willing to share their counters

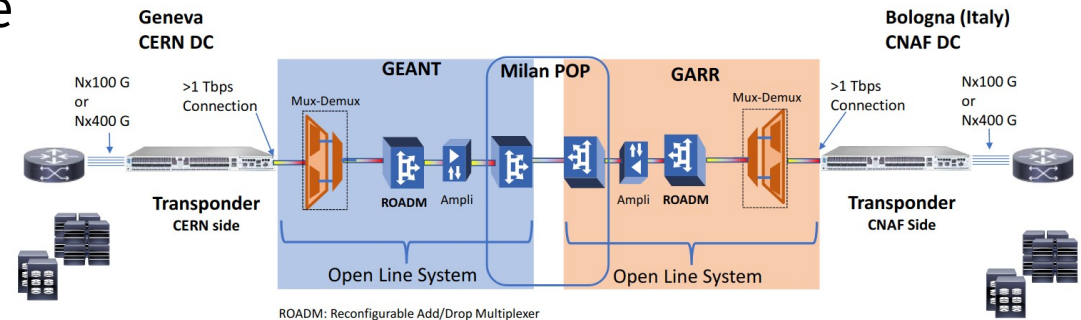




# CNAF-CERN DCI

## Pilot for a Tbps Data-Centre interconnections between INFN-CNAF and CERN

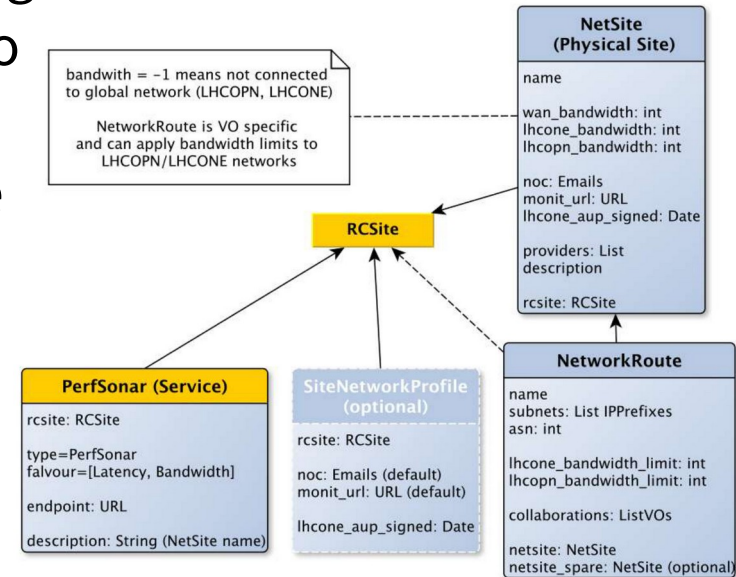
- to meet HL-LHC requirements at an affordable cost
- made using transmission devices at CNAF and CERN
  - no need to use expensive routers
- using optical channels on GARR and GEANT dark fibres
  - pilot also for their new Spectrum Sharing service
  - distance (~1000km) is also a challenge
- GARR will procure and manage the Infinera devices at both ends of the connection



# CRIC Database for LHCONE



- Network information related to LHCOPN, LHCONE and network monitoring has been added to CRIC
  - source: LHCONE/OPN Twiki and routers' routing tables
- WLCG operations has created tickets for all sites to check their records
- Route-set RS-LHCONE in the RIPE whois database built daily from CRIC. It can be used to build security filters



# NEA3R redundancy for LHCONE



- NEA3R has a second transatlantic link that could be used to add redundancy to LHCONE over the Atlantic
- The setup will discuss among the Network providers in the ANA group

NEA3R network



# NOTED demonstration at SC21



Ran a NOTED demo for Super Computing 2021

- NOTED used to detect FTS transfers between CERN-TRIUMF and KIT-TRIUMF.
- once detected, it requested a high bandwidth direct link to AutoGOLE/SENSE

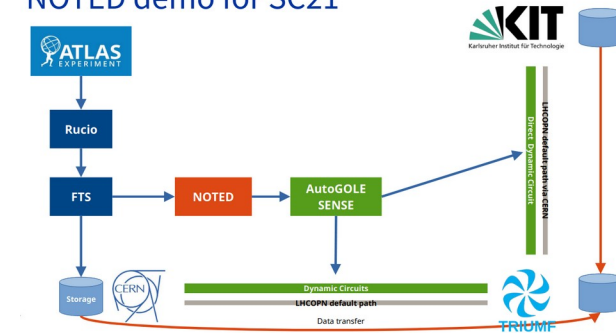
Results and observation:

- NOTED detection of transfers worked, but could have been more accurate
- Dynamic circuit services work, but a bit fragile

Next steps:

- NOTED 2.0 beta being tested
- Improve decision algorithm with Machine Learning

NOTED demo for SC21



# Research Network Technology WG - update

Main current activity: understand scientific traffic flows in detail

Website <https://www.scitags.org/>

Evaluating two approaches for flow tagging:

- Flow marking using UDP fireflies (works for both IPv4 and IPv6)
- Packet marking using IPv6 flowlabel

Fireflies tested during WLCG data challenge 2021.

Discussions on-going about the architecture of the fireflies collectors:  
in-line/distributed/central

scitags.org



# Flowlabel on programmable switches



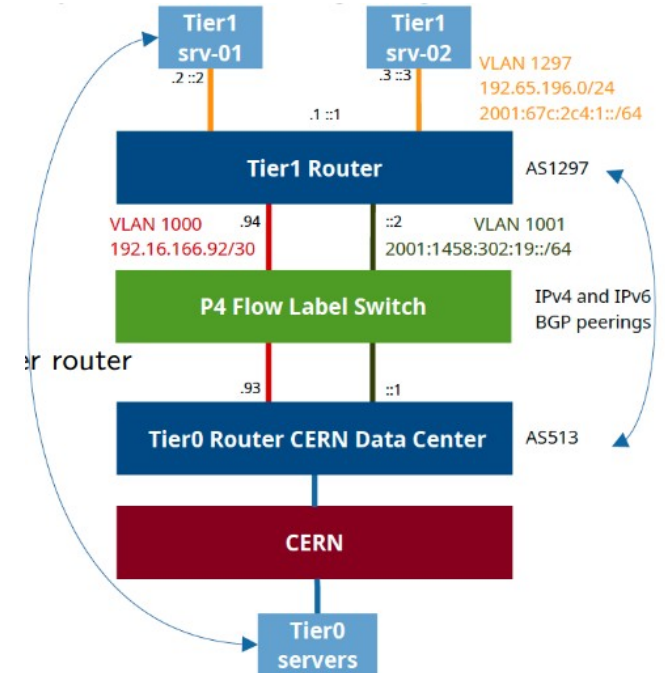
Project at CERN to program a switch to count and route packets with a flowlabel tag

Using EdgeCore Wedge switch

Implemented two solutions:

- L3 flowlabel-tag based routing
- L2 bridging with flowlabel accounting

L2 accounting implemented on a dummy Tier1 connected to LHCOPN. Further tests will be run with production traffic



# Conclusions

# Summary

- LHCOPN and LHCONE traffic has greatly increased in the last months, back to the level of the end of Run2
- ESnet deploying two 400G transatlantic links
- Several monitoring solutions being tested
- CRIC now authoritative for LHCONE data
- CNAF and CERN preparing a DCI pilot using spectrum sharing in GEANT and GARR
- Packet marking activity: progresses in tagging and accounting

# Actions for next meeting

- Implement perfSONAR info in CRIC
- Enlarge NetSAGE pilot for LHCONE monitoring
- Progress in packet marking

# Next meeting

24-25 of October 2022

In person meeting, at CERN

It can be the opportunity to meet LHC Experiments and WLCG

Agenda will be published here:

<https://indico.cern.ch/e/LHCOPNE49>



# References

Meeting agenda and presentations:  
<https://indico.cern.ch/e/lhcopne48>

*Questions?*

*edoardo.martelli@cern.ch*

