



# **LHCOPN-LHCONE meeting #47**

## **summary notes**

25 October 2021 – v1.0

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

# Venue

- On 11-12 of October 2021
- On video conference only, for the fourth time
- Two sessions of 3 hours in two days
- 76 subscribed
- 51 participants on the first day, 57 on the second
- welcome by new CERN IT department head, Enrica Porcari

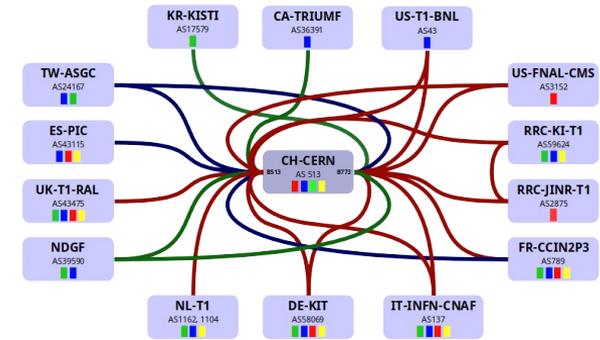
- Agenda at  
<https://indico.cern.ch/e/LHCOPNE47>



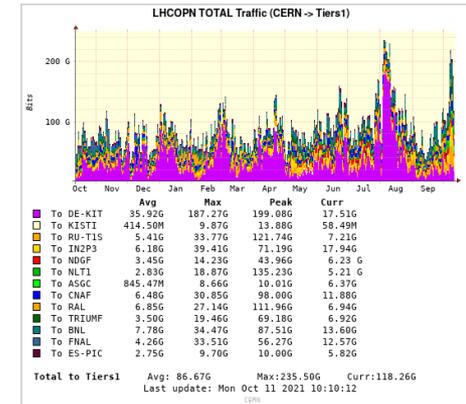
# LHCOPN - update

- No major changes since last meeting. Total of 1.3Tbps from the Tier0 to the Tier1s
- Traffic stats: moved 341PB in one year. +19% compared to previous year
- Upcoming links: PIC 100G (waiting for GEANT upgrade), IN2P3 second 100G, KISTI 100G to Amsterdam, NDGF 100G
- CNAF will activate load balancing over 2x100G (currently in primary/backup config)

LHCOPN



■ Alice ■ Atlas ■ CMS ■ LHCb ■ 100Gbps  
— 100Gbps — 100Gbps — 100Gbps — 100Gbps  
 edoard.martel@cern.ch 20210518



# LHCONE L3VPN - update

## **Topology:**

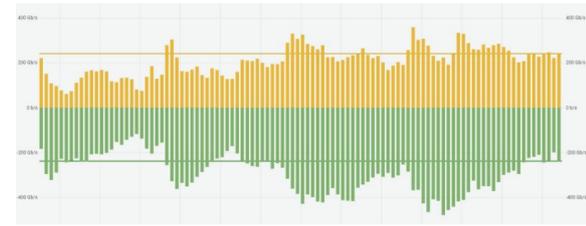
- new BELLA transatlantic link connects Europe to South-America
- new site in Canada: CA-UVic-Cloud
- ULAKBIM (Turkish NREN) connects three new sites in Turkey

## **Traffic volumes:**

- LHCONE traffic now count 40-50% of total traffic for most RENs
- large increase in Europe (GEANT, NORDunet, GARR)
- slight increase in North America and Asia
- High spikes during October WLCG challenges



# LHCONE monitoring - update



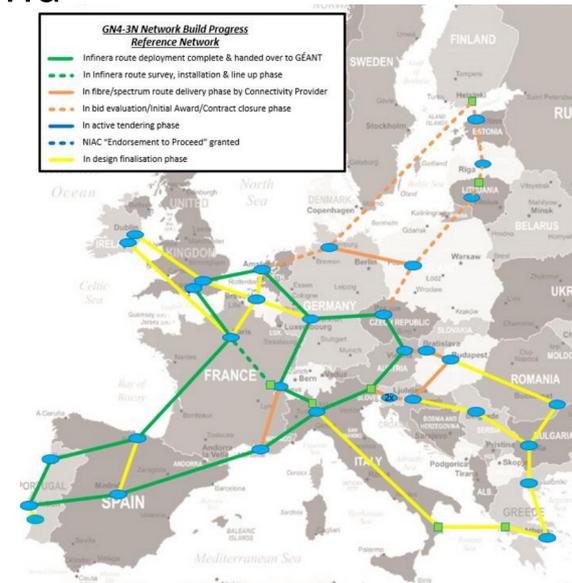
- SAND project ended, HEPiX NFV WG completed, new developer hired for perfSONAR
- perfSONAR 4.4.1 released in September, but introduces some issues on busy nodes. It has probably impacted negatively some of the meshes
- LHCOPN mesh improved, but still some nodes not sending data
- Created LHCONE mesh with few nodes in strategic locations
- New applications:
  - pS-Dash: to monitor a site and detect basic problems
  - AAAS: allows user to subscribe to alarms (not only for ATLAS)
- Monitoring of WLCG data challenges: exceeded 240Gbps goal. The network was not the bottle neck

**perfSONAR**



# GEANT update

- GEANT is upgrading and extending its network: new fibres, new contracts, new transmission devices
- On the optical level using Infinera FlexLS and Groove (just acquired by Infinera)
- Decided to implement a partially disaggregated setup and to allow spectrum sharing
- New routes: Paris-Bilbao, Bilbao-Porto, Porto-Lisbon, Lisbon-Madrid
- Project delayed by silicon shortage and pandemic restrictions
- Routers configuration manager now using Ansible
- Developed a new monitoring tool: BRIAN  
<https://public-brian.geant.org/>
- New network centric web site: <https://network.geant.org/>
- New flow analyses tool: Kentick
- Dropped support for 10G lambdas
- Routers will be replaced in 2022-23. Procurement not started yet



# ESnet update

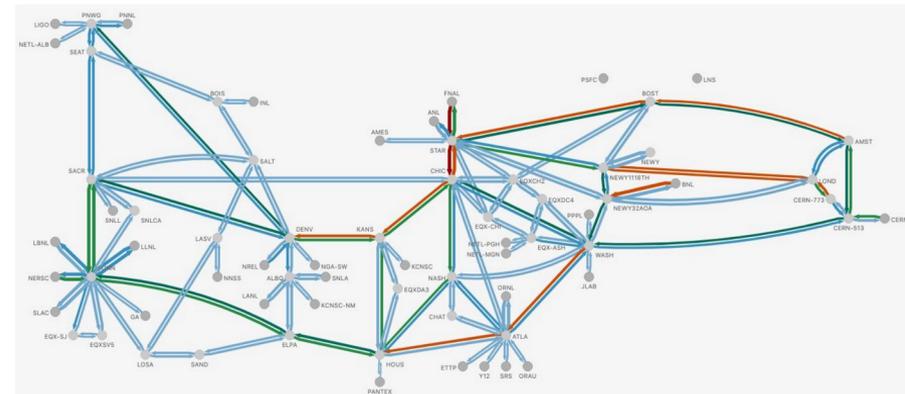
Published report on HEP network requirements study (link)

ESnet6 network upgrade:

- Transmission equipment from Infinera and Groove installation completed
- Phase 2 in progress: transponder, IP routers. Will be completed by H1 2022
- Automation: provisioning mostly ready; it was used to provide additional bandwidth during WLCG Data Challenge
- Statistics: new Startdust public website

R&D:

- Prototype of Firefly flow marking is working
- FABRIC infrastructure deployment is on going. It will be operated by ESnet



# Pierre Auger observatory - update

Astroparticle experiments located in Argentina

Data taking started in ~2000, it will continue at least till 2030

Primary data stored at CCIN2P3 in Lyon FR

Main storage sites with high traffic are in Europe. CNAF may become part of the collaborations

Total generated data volume will not exceed 1PB  
in 2 years



# LHCONE AUP review

LHCONE Acceptable Use Policy fully reviewed:

- full text review, more readable now
- security part shortened and pointer to relevant WLCG groups included
- more specific to WLCG

Published here:

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



# WLCG Network and Data challenges

WLCG has agreed to run Network and Data challenges every two years to demonstrate the ability to deliver what required by HL-LHC

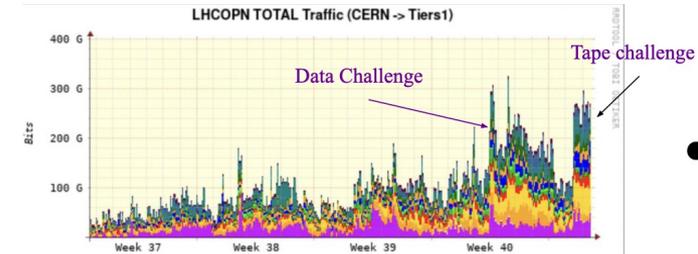
First challenge on going [completed at the end of the week]

- Objective: demonstrate 10% capacity needed for HL-LHC
- First week disks, second week tapes
- Two goals: demonstrate HTTPS-TPC performance; fill 10% network requirement
- Successful so far: the goal was 480Gbps, but we reached 1Tbps. Reached 10% of flexible model, which is twice the initial goal

Tape challenge is giving a more consistent signature on network stats because the data flows more constant and regular

Compared transfer monitoring and network monitoring. Nice match for Manchester Tier2. More difficult to compare for RAL Tier1 because of the more complex network and the large number of collaborations served, but the match was also good.

CERN LHCOPN/ONE traffic didn't match with FTS plots. It may be missing Xrootd



# NetSAGE to monitor LHC data

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

It can ingest data from many sources, like SNMP counters, net/s/flow, perfSONAR...

It is committed to privacy, fully GDPR compliant

NetSAGE instances can also be installed locally

NetSAGE instance at Indiana currently has access to few sources, like NEA3R, GEANT, TransPAC, SingAREN

The NetSAGE team would like to get feedback on existing features and new needed ones

The NetSAGE team offers to collect and show stats for LHCONE. It would need access to data, though. They would also need IP mapping to better understand the traffic. The CRIC database could help.

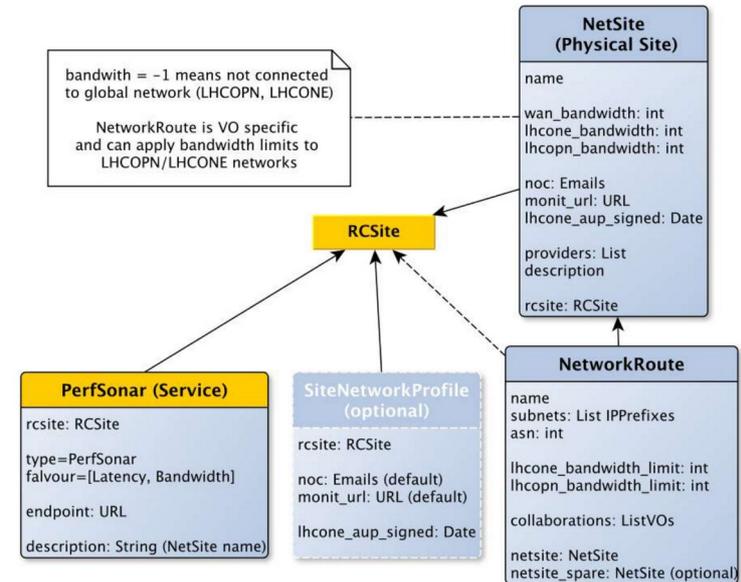
Action: discuss the possibility to use NetSAGE for LHCOPN and LHCONE monitoring



# Database for LHCONE prefixes



- Network information related to LHCOPN, LHCONE and network monitoring agreed to be added to CRIC
- CRIC developers have proposed a database schema and are now implementing it
- It is scheduled for availability in Q1 2022



# Research Network Technology WG - update

Main current activity: understand scientific traffic flows in detail

More participants joined the activity, now ~90

New website <https://www.scitags.org/>

Developed a reference implementation: flowd

Proposed two approaches:

- Flow marking using UDP fireflies (works for both IPv4 and IPv6)
- Packet marking using IPv6 flow label and/or header extensions

Very good progress on the firefly part. Implemented in few sites, being tested with the ongoing WLCG data challenge

# Programming switches

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

Acquired EdgeCore Wedge switch

Joined the GEANT P4 testbed

Installed the RARE/FreeRTR network operating system

Implemented prototype P4 program for flowlabel counting and routing



Figure: Intel Tofino P4-programmable Ethernet Switch ASIC



Figure: EdgeCore Wedge100BF-32QS

# NOTED demonstration at SC21

Planned a NOTED demo to be run at Super Computing 2021

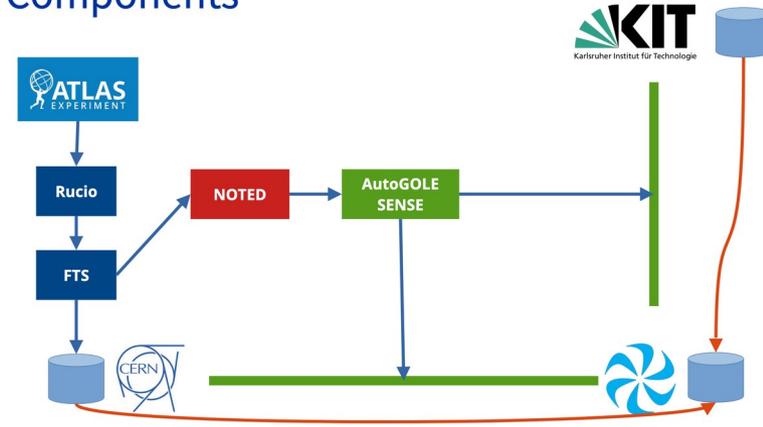
Supported by Starlight

NOTED will detect FTS transfers between CERN-TRIUMF and KIT-TRIUMF.

Once detected, it will request a high bandwidth direct link to AutoGOLE/SENSE



## Components



# FABRIC and FAB

**FABRIC** is an NSF R1-midscale project to build a US national scale programmable network with compute and storage at each node.

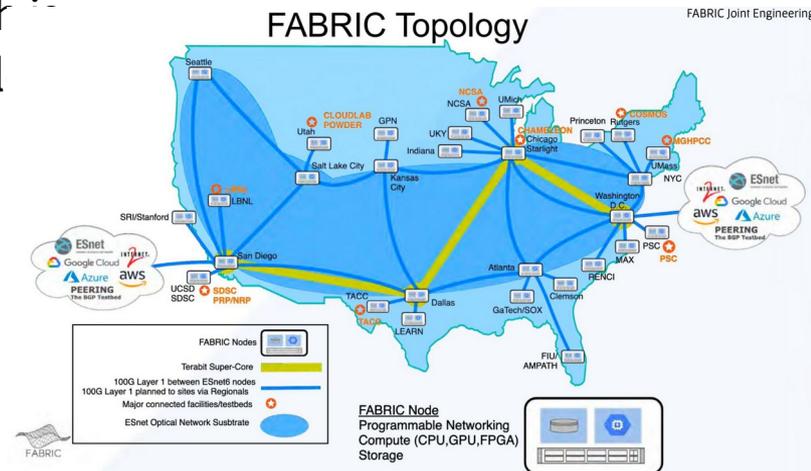
- Run computationally intensive programs & maintain information in the network
- Nodes have GPUs, FPGAs, and network processors (NICs) inside the network
- Interconnects national facilities: HPC, cloud & wireless testbeds, commercial clouds, Internet
- Design and test applications, protocols and services that run at any node in the network

**FAB** (FABRIC Across Borders) is a follow-on to FABRIC which is creating an international extension of this testbed to all at-scale testing for global science

Deployment status: main nodes in US being deployed.

CERN node planned for Q1 2022

Use case: accelerated data delivery



# Datacentre network architecture activity

Organized workshop in the WLCG pre-GDB context - [link to report](#)

Discussed whether keeping the activity in the LHCONE meetings

Agreed to have a dedicated slot when the meeting will be in person again

Discussion should be on planning and network design. Status reports fit better in HEPiX meetings

# Conclusions

# Summary

- LHCOPN: Traffic has started increasing again. Few more 100G upgrades coming
- LHCONE: few new sites (Turkey), new BELLA link to south America, noticeable traffic increase, AUP review completed
- GEANT and ESnet working on network upgrades
- First WLCG data challenge completed successfully. Networks ready to take more traffic
- R&D: NetSAGE for network monitoring, NOTED demo at SC21, Packet and flow marking

# Actions for next meeting

- Complete CRIC and route-set implementation
- Discuss use of NetSAGE for LHCONE monitoring

# Next meeting

End of March / beginning of April

In person meeting, if pandemic situation improves

Possible locations: CERN, Scandinavia

Agenda will be published here:

<https://indico.cern.ch/e/LHCOPNE48>

# References

Meeting agenda and presentations:  
<https://indico.cern.ch/e/lhcopne47>

*Questions?*

*edoardo.martelli@cern.ch*

