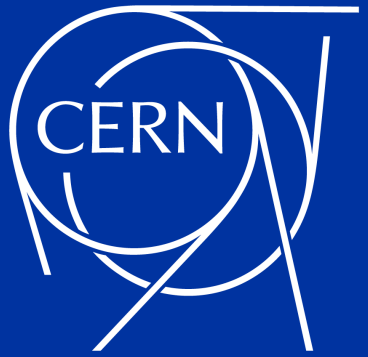

LHCOPN-LHCONE meeting #51 summary notes

27 October 2023 – v1.0

edoardo.martelli@cern.ch

Venue

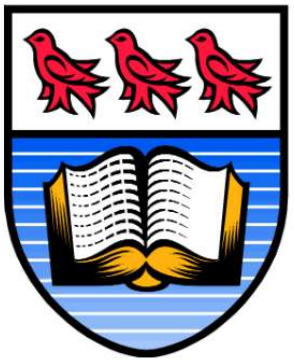

University of Victoria

- 18-19 of October 2023
- Hosted by UVic in Victoria (CA)
- co-located with HEPiX fall 2023
- ~20 people in presence and
  ~20 connected remotely

- Agenda and presentations at
  https://indico.cern.ch/e/LHCOPNE51

LHCOPN - update

- 2.1Tbps of aggregated bandwidth to the Tier0
- Traffic stats: moved 540PB in the last 12 months. +18% compared to previous year
- PL-NCBJ, new LHCb Tier1: connected at 20Gbps
- CN-IHEP, new LHCb Tier1: soon connecting at 20Gbps via CSTnet and GEANT
- NDGF-IHEP, new NDGF distributed Tier1 site at UniBern (CH): ready to activate 100Gbps link to CERN
- TW-ASGC terminated Tier1 activities. It will stay connected to LHCOPN until LHCONE is ready

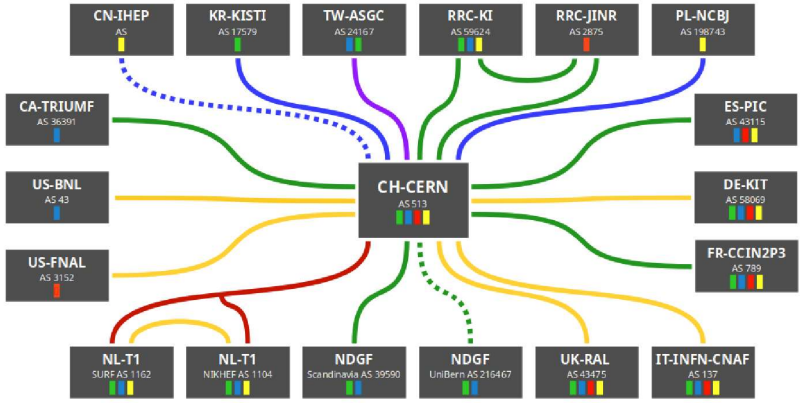

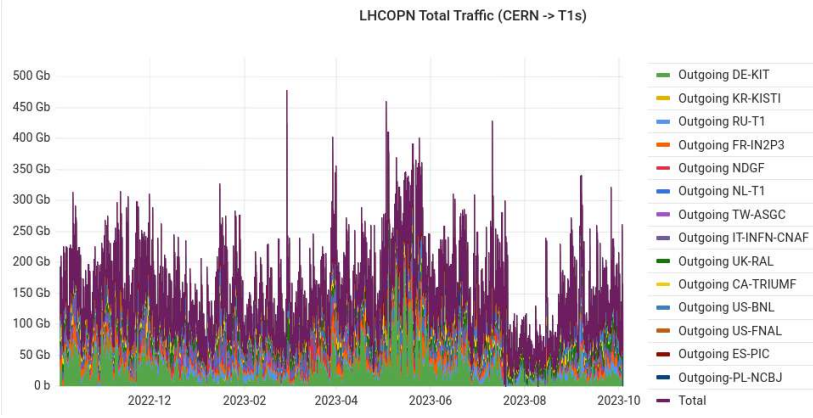

CNAF-CERN DCI project

CNAF-CERN data-centre interconnect (DCI) solution implemented using GEANt and GARR spectru sharing services

Four 400Gbps waves lightened

Tested with 2x100Gbps client interfaces

Next steps:

- test third party optics on client side
- test 400Gbps client (local) interfaces
- test with LHCOPN production traffic

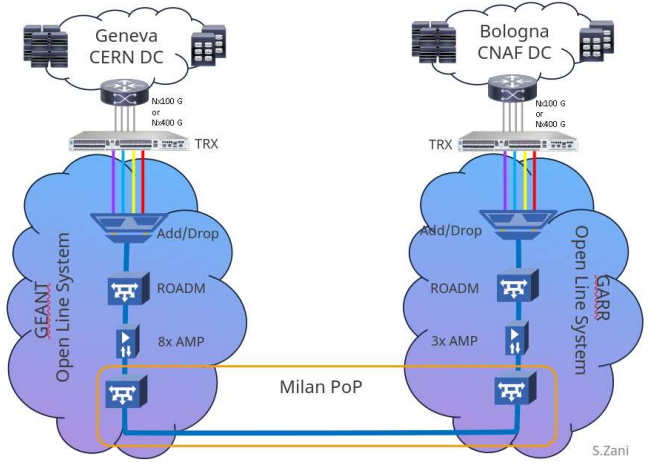

https://indico.cern.ch/event/1280363/contributions/5570060/attachments/2736873/4759609/DCI-CERN-CNAF-LHCOPN-2023.pdf

LHCONE L3VPN - update

News

New NRENs: SWITCH (just L2, no L3VPN), FCCN (PT)

New experiment DUNE

New link to China: collaboration with CSTNET. London-Singapore is provided by GEANT, Singapore-Bejing is provided by CSTNET

URAN: Ukranian site is back on line

GEANT: using CRIC content to analyse flows they collect

GEANT has seen a big traffic increase in the last three months. Most of it comes from the peering with ESnet

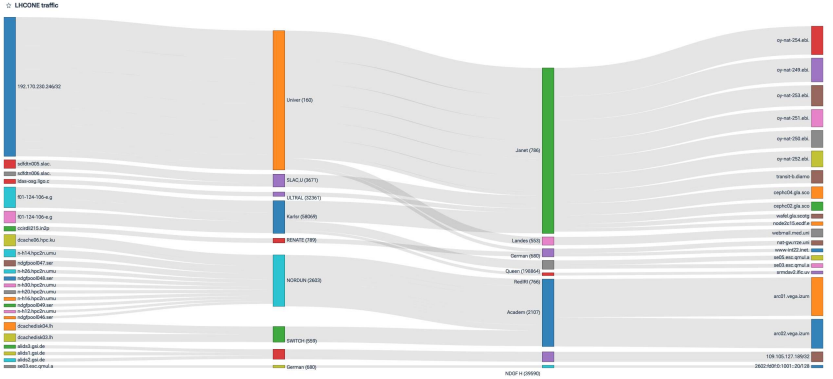

https://indico.cern.ch/event/1280363/contributions/5532174/attachments/2736874/4759612/2023-10-18_ECapone_LHCONE_L3VPN.pdf

ESnet update

Transatlantic, new In Production:
- 400G New York - London

Currently underway:
- 400G Boston - London (Estimated 12/23)
- 400G Boston - CERN (Estimated 12/23)

Trans-Atlantic capacity targets
- 1.5T in advance of DC24
- 3.2T in 2027, in advance of Run 4

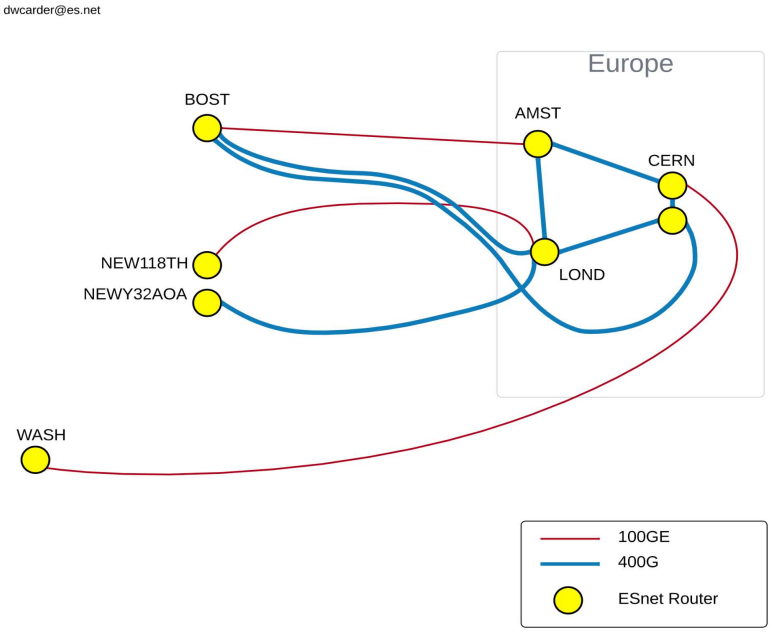

Europe, now In Production:
- 400G European Ring
- Amsterdam to CERN upgraded
- London to CERN upgraded
- Amsterdam to London upgraded

Cloud interconnections
- 5x100G to Google (one more pending)
- 3x100G to Oracle
- 6x100G to Microsoft
- 6x100G to Amazon

BelleII update

- Around 2.8PB of RAW Data Collected since 2019
- Currently in Long Shutdown for upgrade. Data taking will restart before the end of the year
- What should be exercised during DC24:
    - Technology that can be stressed: Network, DDM, FTS, Storages, Monitoring System, Protocols, IAM
    - Main goal: Emulate data transfer conditions in a Belle II high-lumi scenario
    - Current estimation for such scenario is 40 TB per day
- Plan for DC24 testing ready and execution already started

https://indico.cern.ch/event/1280363/contributions/5578045/attachments/2736776/4759435/BelleII-October-2023.pdf

GEANT update

Packet layer project renewal started

- 3 years project
- Will bring 400G connectivity in any PoP
- Selected Nokia 7750 routers

Automation platform will be enhanced with Workflow Orchestrator (tools developed by SURF)

https://indico.cern.ch/event/1280363/contributions/5622069/attachments/2736875/4759614/ECapone%20-%20GEANT%20Updates.pdf

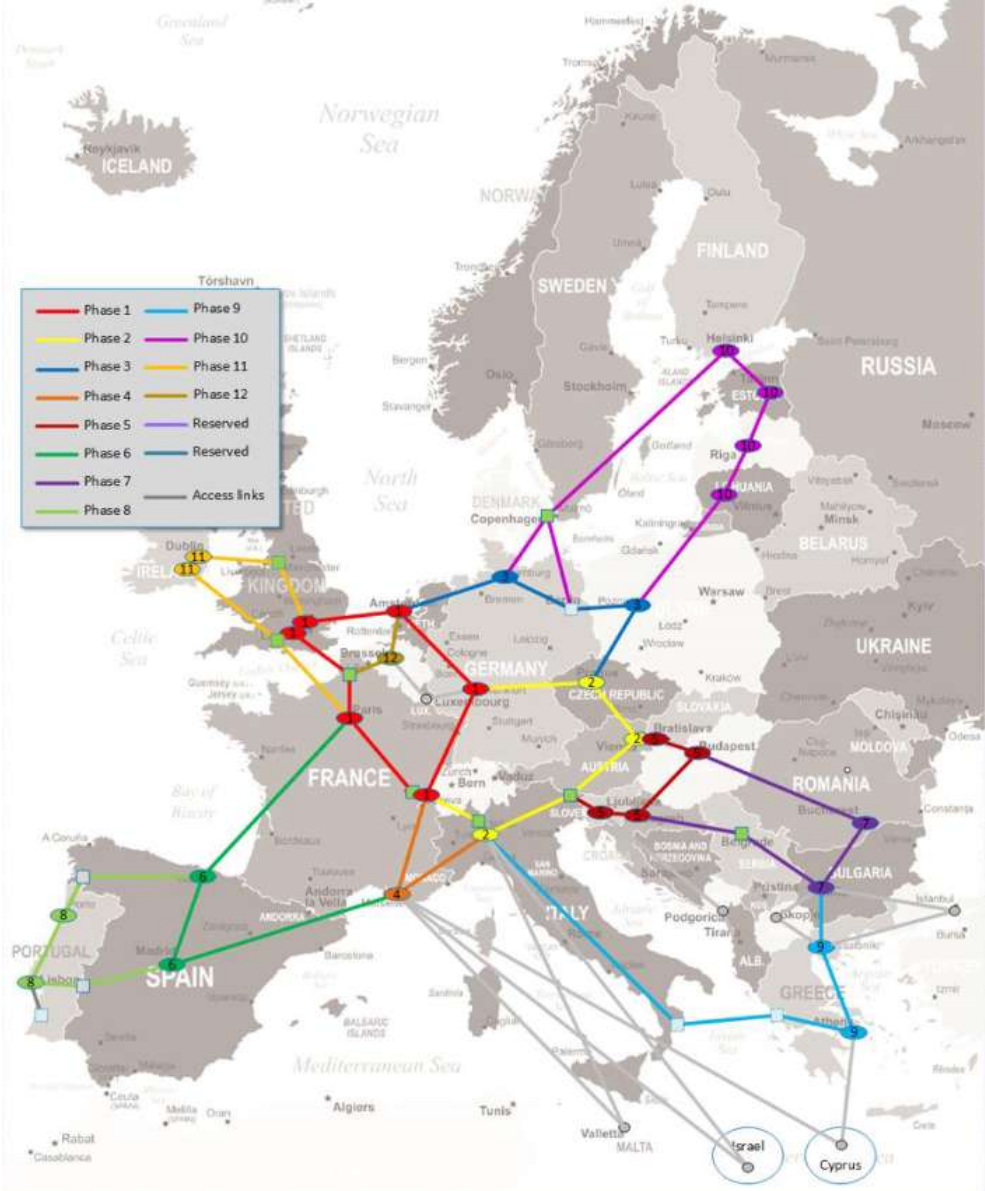

Misrouted packets in LHCONE

KIT has analysed the traffic coming from its LHCONE link and found a significant amount with source addresses that are not part of LHCONE (e.g. private addresses, cloud providers, others)

Agreed to start a working group to understand the origin of this traffic and try to remove it

https://indico.cern.ch/event/1280363/contributions/5584073/attachments/2736880/4759636/misrouted-packets.pdf

DUNE joins LHCONE

At the last LHCONE meeting in Prague (April 2023), DUNE has formally requested permission to join LHCONE

The LHCONE community expressed its approval to allow DUNE to join LHCONE

The request and the decision were presented to the WLCG Management Board of September 2023. The Board had no objections and endorsed the decision

https://indico.cern.ch/event/1280363/contributions/5532186/attachments/2736077/4758029/LHCOPNE-20231018-51-DUNE-joins-LHCONE.pdf

ATLAS Google project on cloud interconnect

- ATLAS Google project completed recently with very positive technical results
- TCO study highlighted the potential cost of egress
- Interest to leverage LHCONE to reduce (not eliminate) these costs and avoid hitting sites' commodity internet connection
- Tests with ESnet showed that the solution is not straightforward
    - Google Interconnect technology designed for bridging two data centres together through private IPs, e.g. Google resources with a University/Lab
    - Possibilities depending on each cloud provider
    - Adding cloud resources to the LHCONE requires more experience and work
- Further projects will require more detailed planning and possibly hiring additional support option to speed up support interactions
- NRENs willing to help with future tests

https://indico.cern.ch/event/1280363/contributions/5532184/attachments/2737715/4761330/ATLAS%20Google%20Project.%20Peering%20attempt%20between%20Google%20Cloud%20and%20LHCONE.pdf

Data Challenge 2024

List of the projects on networking:
- Packet marking
- Packet pacing, BBR performances
- perfSONAR for network alarms and debugging
- Site Network monitoring of in/out bandwidth
- Use of Jumbo frames to improve performances
- NOTED: FTS driven SDN
- Rucio SENSE
- ALTO FTS Rucio

https://indico.cern.ch/event/1280363/contributions/5532183/attachments/2736950/4759756/LHCOPNE-20231019-D24.pdf

Easy-to-use Network load generator and test results at US-ATLAS

BNL has developed some scripts to generate FTS transfers.

These scripts can be used to generate network load, for testing and DC24

Instructions about how to use it in the presentation

https://indico.cern.ch/event/1280363/contributions/5627741/attachments/2737626/4761143/Easy-to-use%20Network%20load%20generator%20and%20test%20results%20at%20USATLAS.pdf

The need for Packet pacing

Congestion is often seen on LHCOPN and LHCONE links.

Over-provisioning it is not the best solution, because average usage is ~25%

Packet pacing could help to reduce packet drops

BBRv3 (TCP congestion protocol developed by Google) will be tested for DC24

https://indico.cern.ch/event/1280363/contributions/5578062/attachments/2737701/4761305/20231019%20-%20LHCOPN%20Victoria%20-%20traffic%20pacing.pdf

Congestion driven NOTED

A new version of NOTED is being developed.

Its actions are triggered by network interfaces congestion alarms

It will be dry-run during DC24

https://indico.cern.ch/event/1280363/contributions/5551282/attachments/2736080/4758035/LHCOPNE-20231019-51-NOTED.pdf

NOTED Alarms

ID ↓	Alarm name	Version	NOTED status	NOTED action	SDN status	Max FTS Throughput [Gb/s]	Interface	NOTED alarm start
5	CH-CERN to DE-KIT	LHCOPN	Action	On-going SDN. FTS throughput [Gb/s]: 28.72	Provided	24.9	I513-e-rjup1-1_irb.3530	2023-10-11 08:27:56
4	CH-CERN to DE-KIT	LHCOPN	Running	Spectrum generated an alarm: NOTED is inspecting FTS.			I513-e-rjup1-1_irb.3512	2023-10-10 18:03:00
3	CH-CERN to DE-KIT	LHCOPN	Running	Spectrum generated an alarm: NOTED is inspecting FTS.			I513-e-rjup1-1_irb.3530	2023-10-10 18:02:43
2	CH-CERN to CA-TRIUMF	LHCOPN	Running	Spectrum generated an alarm: NOTED is inspecting FTS.			I513-e-rjup1-1_irb.2126	2023-10-10 18:02:28
1	CH-CERN to FR-IN2P3-LAPP	LHCONE	Running	Spectrum generated an alarm: NOTED is inspecting FTS.			I513-e-rjup1-1_irb.111	2023-10-10 18:02:13

LHC networking related SC23 NREs

Several Network Research Exhibition (NRE) submitted to Super Computing 2023 (Denver)
- NOTED: FTS transfers between CERN-TRIUMF, CERN-FNAL, KIT-TRIUMF
- SCITAGS: tagging and accounting at 400Gbps and beyond

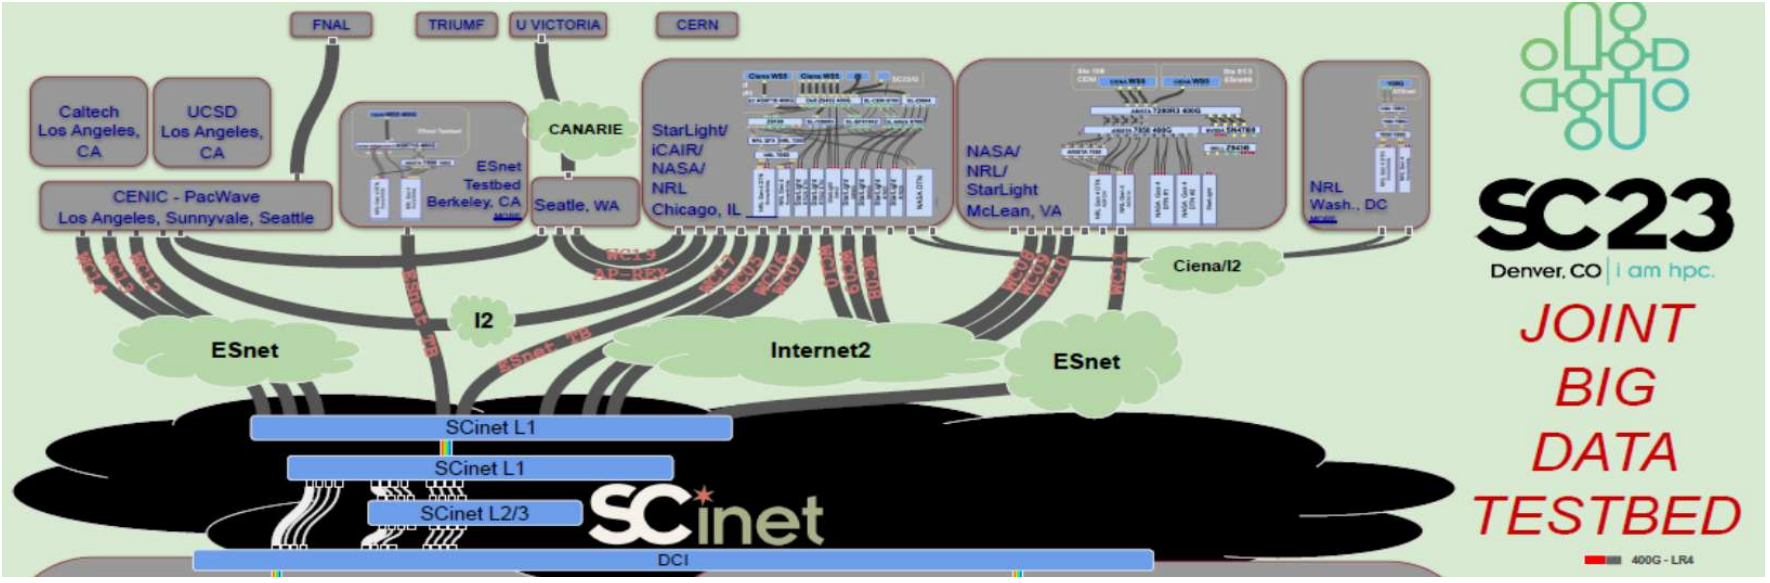

https://indico.cern.ch/event/1280363/contributions/5632150/attachments/2737720/4761335/LHC%20Networking%20SC23%20NRE%20Demonstrations.pdf

MultiONE: using BGP communities to identify collaborations and reduce exposure

New proposal for MultiONE implementation:
- Don’t add any additional VPN (or maybe just one for Other Big Sciences)
- 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 REN if they can’t do it
- Sites can/should then decide to accept only the prefixes of the collaboration they are working with

This proposal is less operationally complex then the previous one, since it use a common technique already used by RENs

Agreed to explore this option further. The proposal will be discussed on the Architecture mailing list and at the next meeting

https://indico.cern.ch/event/1280363/contributions/5615679/attachments/2736079/4758032/LHCOF

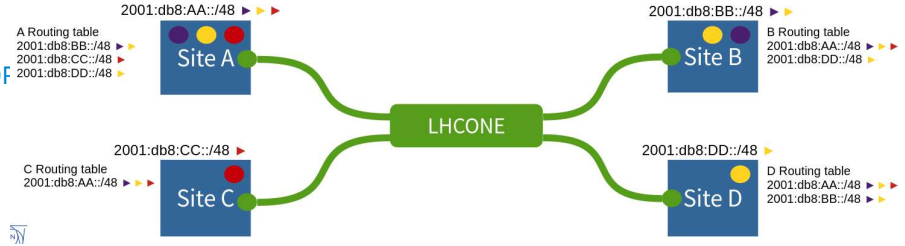

How to report the available bandwidth in CRIC

Doubts about how to express the LHCONE and the generic WAN bandwidth values in the CRIC NetSite table, whrn the two uplinks are provisioned as shared VLANs over a common physical connection

- No straightforward answer. An additional field may be necessary to signal that the bandwidth is shared

- A concrete proposal will be made and discussed again

- netsite: https://wlcg-cric.cern.ch/core/netsite/list/

- networkroute: https://wlcg-cric.cern.ch/core/networkroute/list/

Use of Jumbo frames

Revamp of Jumbo frame deployment initiative
Jumbo frames (large MTU, 9000 Bytes) can improve improve performance of data transfers
The NREN networks already support Jumbo frames, LHCOPN and LHCONE too
Transfer servers at LHCONE sites should support Jumbo frames

Discussion:
- Proposed to run survey to understand current use and interest in the community
- Results could be added in a new record In the NetSite table in CRIC

Conclusions

Summary

- LHCOPN: three new Tier1s
- LHCONE: three new NRENs and three new Tier2s
- ESnet: two more 400Gbps transatlantic links soon coming
- GEANT: Nokia routers will be deployed soon
- MultiONE: new solution with BGP communities.
- DC24: several network R&D projects will be tested
- Transfer performances: Jumbo frames and BBRv3 will be tested

Actions

- Discuss MultiONE with BGP communities proposal
- Survey use of Jumbo frames. Record results in CRIC
- Investigate unwanted packets
- Propose solution to correctly record shared bandwidth in CRIC

Next meeting

INFN Catania (IT)
9-10-11 April 2024

Agenda will be published here
https://indico.cern.ch/e/LHCOPNE52

References

Meeting agenda and presentations:
https://indico.cern.ch/e/lhcopne51

Questions?

edoardo.martelli@cern.ch