

# The case for LHCONE @ the EIC

Sep. 29, 2026

# A Brief History of LHCOPN/LHCONE

The LHCOPN (LHC Optical Private Network) was originally conceived to connect CERN's Tier-0 to the Tier-1's spread around the globe in 2005. The LHCONE (LHC Open Network Environment) followed in 2011 and connected the LHCOPN sites to the Tier-2/Tier-3 communities.

LHCONE a virtual overlay network built on top of existing Research and Education Networks (REN). It is not a replacement of LHCOPN, but a complimentary activity.

Provides a series of locations that are effectively entry points into a network that is private and in support of the LHC T1/2/3 sites.


Some of the following slides taken from presentations by Edoardo Martelli's (CERN) LHCOPN/LHCONE meetings and David Kelsey's (RAL) HEPIX IPv6 Working Group.

More details in CHEP 2015 presentation →

<https://iopscience.iop.org/article/10.1088/1742-6596/664/5/052025/pdf>

# LHCOPN Today

## LHCOPN



|                     |                     |
|---------------------|---------------------|
| <b>Line speeds:</b> | <b>Experiments:</b> |
| 20Gbps              | Alice               |
| 40Gbps              | Atlas               |
| 100Gbps             | CMS                 |
| 200Gbps             | LHCb                |
| 400Gbps             |                     |
| 800Gbps             |                     |
| IPv4                |                     |

**Last update:**  
- 20260413  
- edoardo.martelli@cern.ch

## Numbers


- 17 sites for 15 Tier1s + 1 Tier0
- 15 countries in 3 continents
- 3.28 Tbps to the Tier0
- Serbia has not started operations yet

<https://twiki.cern.ch/twiki/bin/view/LHCOPN/OverallNetworkMaps>



# Traffic Growth

## LHCOPN: traffic growth



Y-Axis: Gbps - Average bandwidth of previous 12 months

**Run1:** 2010-12

**LS1:** 2013-14

**Run2:** 2015-18

**LS2:** 2019-21

**Run3:** 2022-26

# LHCOPN/LHCONE and IPv6

Multi-year process to migrate from dual IPv4/IPv6 infrastructure to IPv6 only. This reduces operational complexity by removing dual-stack routing, simplify network traffic monitoring and minimize network security risks with the elimination of IPv4.

Task managed by WLCG IPv6 Task Force and the HEPIX IPv6 Working Group.


## Timeline:

- Tier-0 and Tier-1 migration to dual-stack (June 2019).
- Migrate applications and storage to prefer IPv6 (2024-2025).
- Begin turning off IPv4 (ie, BNL-FNAL IPv4 peering already turned off) from LHCOPN network (Sep. 2025).
- Aim to remove IPv4 from rest of LHCOPN in 2026, ahead of the WLCG Data Challenge (Feb. 22 – Mar. 5, 2027).
  - Data Challenge will include LHC and non-LHC (JUNO, DUNE, Belle-II, SKA, etc)
- Migrate any remaining IPv4 traffic to LHCONE but aim to become IPv6-only for the HL-LHC run4 period (2030 and beyond).
  - Evaluate removal of IPv4 traffic from U.S. Tier-2 sites in Feb. 2026 (AGLT2+NET2 for USATLAS and UNL+Purdue for USCMS) – no issues found, IPv4 moved from LHCONE to general REN.

# Current status of migration

LHCOPN % of traffic which is IPv6  
- last 2 years. (85% to 98%)

HEPiX




# IPv6 vs. IPv4 on LHCONE

From Eli Dart's presentation at  
ePIC collaboration meeting on  
May 18, 2026

## IPv4 vs IPv6 on LHCONE (30 days)

IPv4 vs IPv6 Traffic Rate Over Time



# EIC context

Should expect current and future EIC collaborating institutions in Asia, Europe and North America to support IPv6, since:

- IPv4 address space is effectively exhausted.
- Many of the overseas institutions also collaborate at the LHC and are part of LHCOPN/LHCONE.

Current hardware (servers, storage, network) is dual stack-enabled, and the past experiences of LHC experiments is informative about realistic timelines and outcome.

Benefits of joining include:

- Dedicated private, virtual network and existing operational model and policies.
- Leverage (lower costs) existing high-performance network infrastructure and existing governance, operational model and monitoring, instead of developing a new one from the ground up.
- Simplify data center management and operations – lower operational burden.

Several non-LHC experiments have joined the LHCONE network.

- Belle-II (KEK, Japan) in 2014 → BNL/SCDF is a Tier-1 site for Belle-II.
- XENON (Gran Sasso, Italy) in 2016.
- JUNO (IHEP, China) in 2021.
- DUNE (FNAL, USA) in 2023.

# Additional benefits of joining LHCONE

Interactions and collaborative activities with other LHCONE sites.

- Leverage large community for technical experience and information exchange.

Security policies that filter site ingress/egress but do not throttle LHCONE traffic.

Routing optimization

- Flexibility to route traffic between two LHCONE sites over different paths.

Enhanced traffic visibility and accounting (ie, packet marking).

# A potential model

Echelon 0 --1: dedicated, bandwidth-guaranteed point-to-point virtual circuits between BNL and JLAB, similar in concept to the connectivity provided by LHCOPN. Under discussion with ESNET.

Echelon 1--2: LHCONE provides the REN between Echelons 1 and any participating Echelons 2.

- Alternative is to go at it alone (see the slide titled “The RHIC approach”).
  - Increased capital and operational costs.
  - Additional networking and operational complexities for Echelons2 (specially if they are already on LHCONE).
  - Separate end-to-end monitoring and troubleshooting framework.

# The RHIC approach

Started operations in 2000 with 4 experiments (BRAHMS, PHENIX, PHOBOS and STAR), 5 years before LHCOPN and 11 years before LHCONE began operations.

- sPHENIX began operations in 2023.

Nearly all computing is/was done at BNL.

- PHENIX processed some data at KEK, and STAR made use of LBL's HPC computing clusters.
- sPHENIX and STAR made only light usage of OSG resources.
- Overall, distributed computing was/is not a priority.

RHIC network was self-funded and is (still) IPv4-based, and there were/are no plans to dual-stack or migrate to IPv6.

# Summary

LHCONE allows EIC to leverage an established network, operational and monitoring ecosystem rather than re-creating one specifically for EIC

Joining LHCONE requires a commitment to migrating to IPv6.

- Must be before 2030 (when HL-LHC starts).
- A similar 2-step process (IPv4 to dual-stack and then to IPv6-only) will require additional effort.
- Best for distributed computing models that rely on geographically remote sites.

Impact of not joining LHCONE:

- Increases capital costs by duplication of dedicated network infrastructure at BNL, JLAB and other collaborating institutions.
  - Option – migrate network traffic to the general REN and live with I/O throttling ( firewalls).
- Must create own governance, operational model, monitoring/trouble-shooting procedures, etc.
- Like RHIC, it would minimize role of distributed computing beyond Echelons 1.
- Recent DUNE experience is informative (data transfer rates 5-6x lower) --  
<https://twiki.cern.ch/twiki/bin/view/LHCONE/LhcOneWhyJoining>

# Additional information on LHCONE

## Regular meetings on the status of LHCOPN/LHCONE

- HEPIX ([www.hepix.org](http://www.hepix.org))
  - Update from April 2026 →  
<https://indico.cern.ch/event/1598655/contributions/6988152/attachments/3261920/5825358/IPv6-Kelsey-23Apr26-V3b.pdf>
- LHCOPN/LHCONE meetings
  - See <https://indico.cern.ch/event/1677691/> for information on the Oct. 1-2, 2026 meeting

## Additional documentation and information about LHCONE

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




# Back-up slides

# Traffic Snapshot

## LHCOPN Traffic – last 12 months

LHCOPN Total Traffic (CERN -> T1s)



### Numbers:

Moved ~812 PB in the last 12 months

+39% compared to previous year (583PB)



Ref: <https://monit-grafana-open.cern.ch/d/HreVOyc7z/all-lhcopn-traffic?orgId=16>

# LHC Long-Term Schedule

- Long-Shutdown (LS3) underway until mid-2030.
- Upgrades to improve LHCOPN/LHCONE connectivity.

