



# LHCOPN and LHCONE update

ATCF7 Jeju - 2<sup>nd</sup> November 2023  
edoardo.martelli@cern.ch

# Agenda

- CERN Tier-0
- LHCOPN
- LHCONE

# CERN Tier-0 update

# Highlights

## LHC

- Run3 in full swing. Cooling incident during the summer, now fixed
- Run4 planned for 2029

## Computing

- Storage now exceed 1 exabyte
- New Preveessin Data-Centre (PDC) building is ready.
  - Network and servers installation will start in November 2023

## Network

- Completed campus upgrade
- PDC network will use new overlay setup with BGP

*CERN Science Gateway is now open*



# PDC (Preveessin Data Centre) status

## 2023

- October: data-centre building, cooling, power, fibres: all ready
- November: installation of network equipment and first servers


## 2024

- Q2: ready for production



# CERN total traffic

LHCOPN+LHCONE+Internet



	min	max	avg
Internet IPv4 In	0 b	13.7 Gb	3.78 Gb
Internet IPv4 Out	0 b	107 Gb	20.8 Gb
Internet IPv6 In	0 b	24.2 Gb	7.57 Gb
Internet IPv6 Out	0 b	38.5 Gb	9.71 Gb
Internet Total In	0 b	28.1 Gb	11.3 Gb
Internet Total Out	0 b	124 Gb	30.6 Gb
WLCG IPv4+IPv6 In	0 b	141 Gb	42.1 Gb
WLCG IPv4+IPv6 Out	0 b	271 Gb	118 Gb
Total In	7.49 Gb	160 Gb	53.4 Gb
Total Out	13.1 Gb	319 Gb	148 Gb

## Numbers:

Sent out ~583 PB  
in the last 12  
months

+9% compared to  
previous year  
(536PB)

Cooling system  
incident  
reduced data  
production  
during the  
summer


Ref: [https://monit-grafana.cern.ch/d/cScW82Tnz/00-overview?orgId=14&var-source=long\\_term&var-bin=1d&from=now-1y&to=now](https://monit-grafana.cern.ch/d/cScW82Tnz/00-overview?orgId=14&var-source=long_term&var-bin=1d&from=now-1y&to=now)



# Experiments' DAQ lines to IT data-centre

Links for Data Acquisition from LHC experiments to CERN IT data-centre. Capacity in place for Run3:

- ALICE: 3.2 Tbps
- LHCb: 400 Gbps
- CMS: 400Gbps
- ATLAS: 200Gbps







# Networks at CERN: some numbers

- LHCONE capacity: 1.2 Tbps
- LHCOPN capacity: 2.1 Tbps
- Internet capacity: 1 Tbps
- Statefull Firewall capacity: 0.6 Tbps
- HTAR (firewall bypass) capacity: 0.6 Tbps (*waiting for new routers with more interfaces on Core routers*)



# LHC schedule

Here today



Last update: April 2023

<https://lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm>

# CERN Quantum Technology Initiative


Phase 2 (2024-2028) just approved

HYBRID QUANTUM  
COMPUTING AND  
ALGORITHMS

QUANTUM  
NETWORKS AND  
COMMUNICATIONS

CERN QUANTUM  
TECHNOLOGY  
PLATFORMS

COLLABORATION  
FOR IMPACT



QUANTUM  
TECHNOLOGY  
INITIATIVE

# QTI2 - Quantum Communications

Set up the CERN Quantum Networks Hub (physical layer)

## **Quantum Key (QKD), Quantum Communication**

- set up a test environment, to gain experience with required network capabilities and equipment
- use of White Rabbit for key synchronization
- interconnect with other NRENs Quantum networks

## **Optical Time and Frequency Distribution**

- identify experiments needing it, reach them with optical connections
- connect to metrology institutes in Europe. GEANT is planning an European network and is interested in a collaboration

**LHCOPN**

# LHCOPN


## Private network connecting Tier0 and Tier1s

### Secure:

- Dedicated to LHC data transfers
- Only declared IP prefixes can exchange traffic
- Can connect directly to Science-DMZ, bypass perimeter firewalls

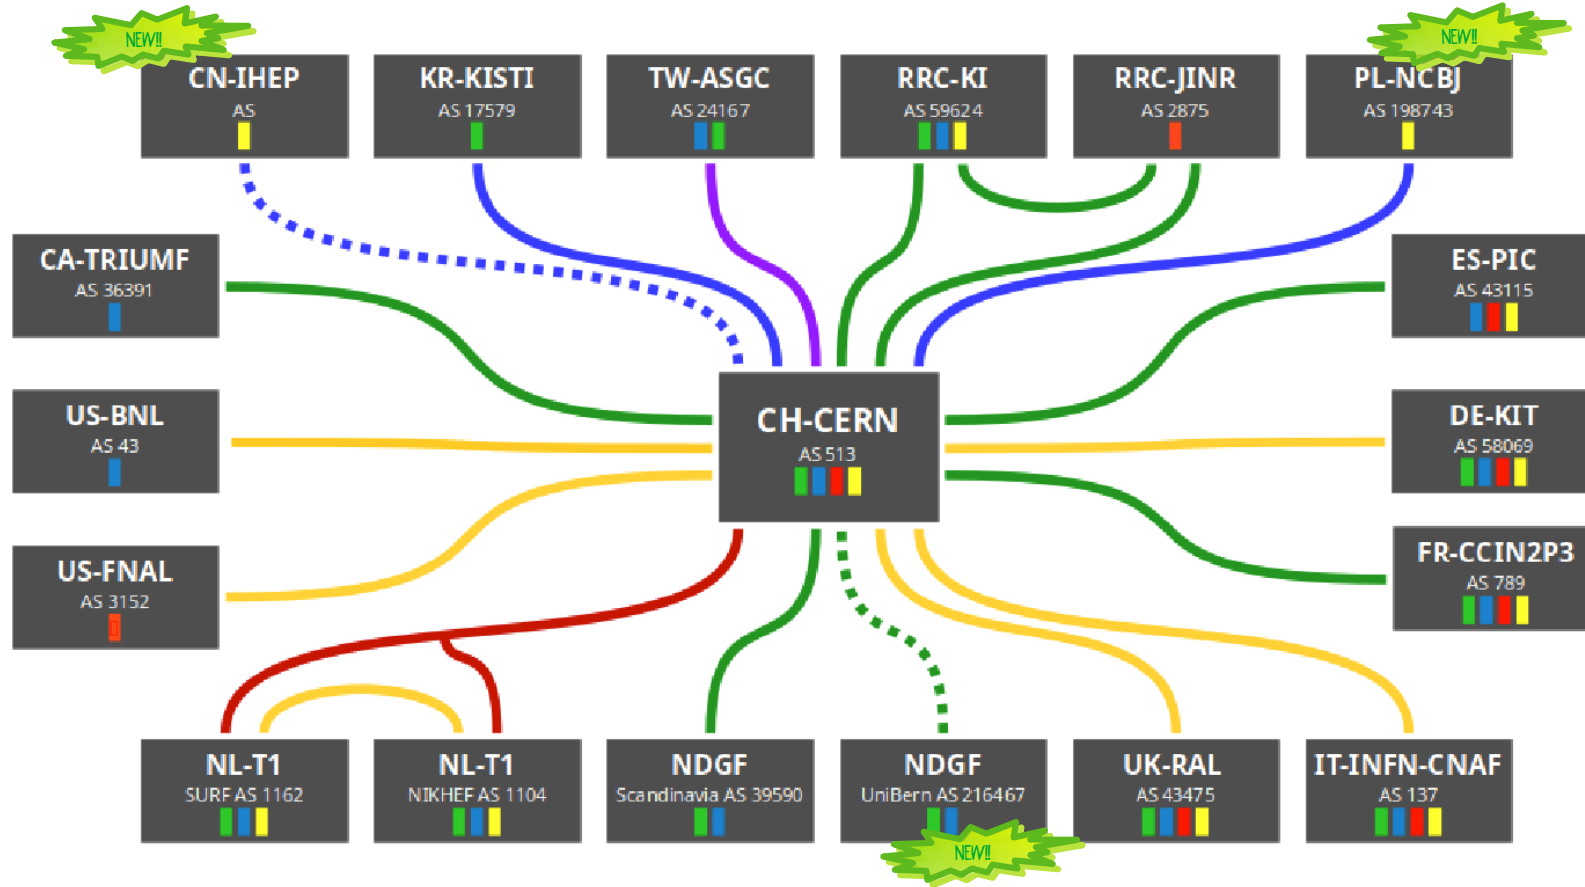
### Advanced routing:

- BGP communities for traffic engineering



LHCOPN

# LHCOPN



■ = Alice	■ = Atlas	■ = CMS	■ = LHCb	■ 10Gbps
				■ 20Gbps
				■ 100Gbps
				■ 200Gbps
				■ 400Gbps

edoardo.martelli@cern.ch 20231003

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

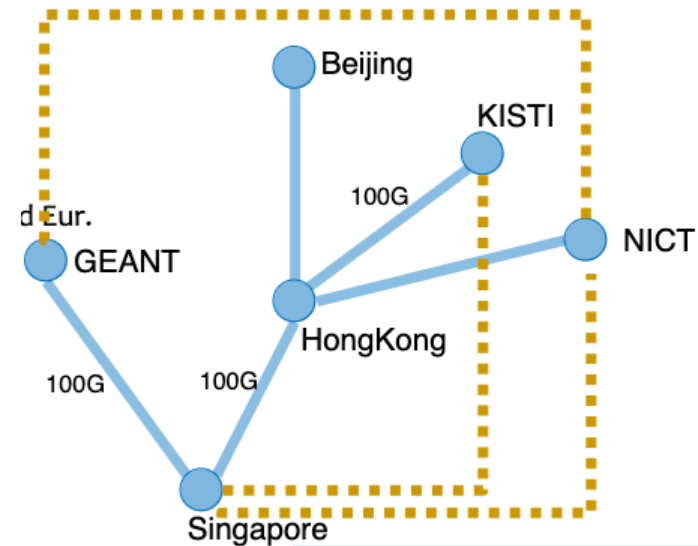
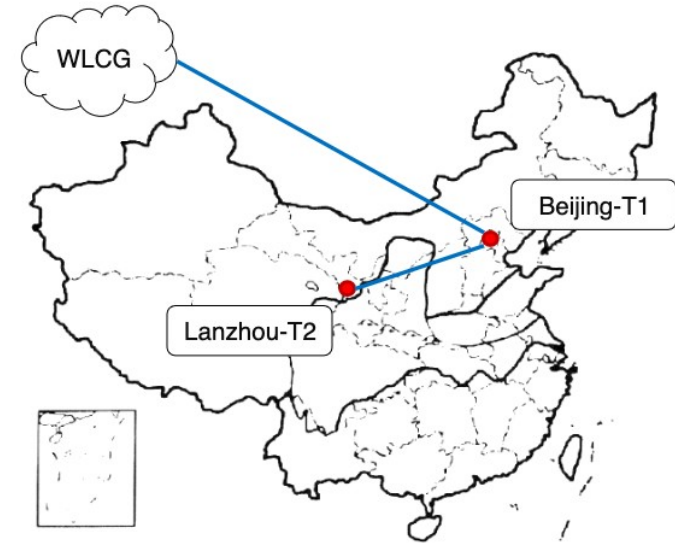
## Numbers

- 18 sites for 15 Tier1s + 1 Tier0
- PL-NCBJ just joined, CN-IHEP and NDFG-LHEP in the process to connect
- 15 countries in 3 continents
- 2.1 Tbps to the Tier0

# IHEP (CN): new LHCb Tier1

IHEP LHCb Tier-2 has started the procedure to become LHCb Tier-1

- CSTNet is the network service provider for IHEP
- International links
- All domestic connections will be upgraded from 10G to 100G
- New international connections will be deployed to improve the bandwidth between China and Europe
- LHCOPN: new link to CERN via CSTNet and GEANT



# NCBJ (PL) - new LHCb Tier1

NCBJ, National Centre for Nuclear Research in Warsaw has started the procedure to become a LHCb Tier1.

It hosts the Świerk Computing Centre (CIŚ)

- Computing: 1.4 PFLOPS, 36000 cores, 200 TB RAM
- Disk storage: 26 PB (Lustre, Isilon, Netapp, dCache)
- Tape storage: TSM4500, 16 PB (uncompressed)

Network resources:

- 100 Gbps link to PIONIER (academic internet, GEANT)
- 20 Gbps dedicated VLAN to LHCONE
- 2x10 Gbps dedicated link for LHCOPN connectivity





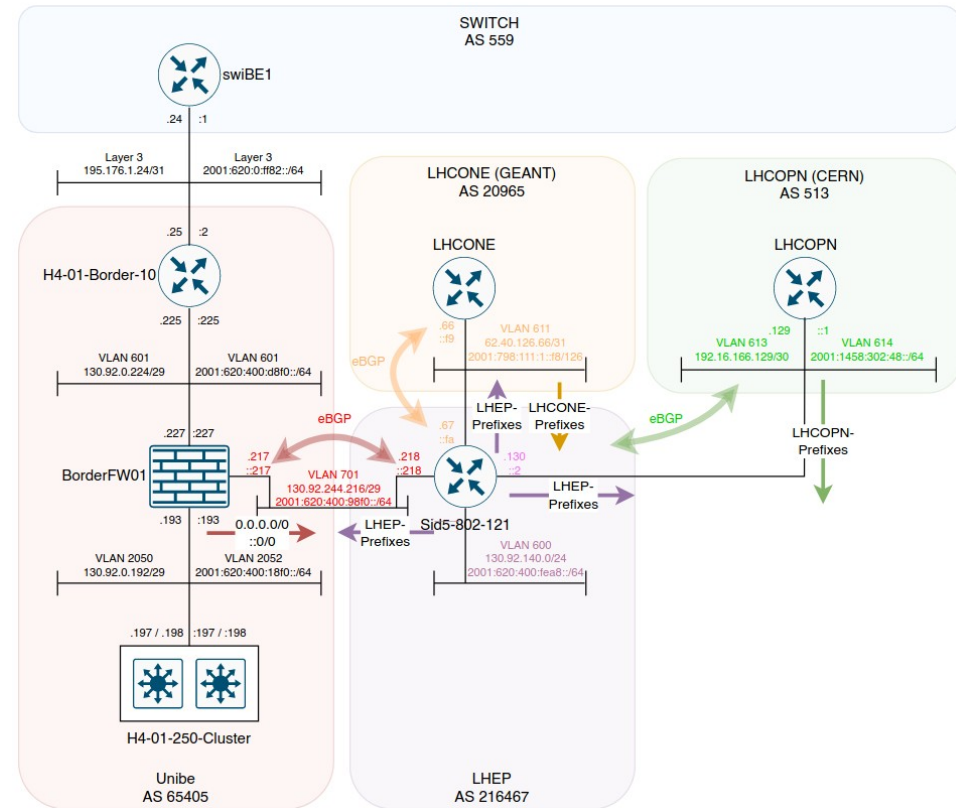
# LHEP (CH) new NDGF Tier1 site

LHEP at University of BERN has joined the NDGF distributed Tier1.


LHEP will be connected to CERN with a 100G LHCOPN link

The physical connection is provided by SWITCH (Swiss NREN) and it is already in place.

LHCOPN routing will be configured in November 2023



# LHCOPN latest news



## **NLT1:**

- SURF has completed the migration of SARA and NIKHEF behind 400G link and SURF ASN


## **UK-RAL:**

- Second 100G link used for new data-centre network

## **PL-NCBJ:**

- Primary link 2x10G in production. Backup link 2x10G just deployed.
- NCBJ has a 32bits AS number AS198743. Its LHCONE BGP community for Do-Not-Announce-To is an extended one (513:3:2138)

# LHCOPN latest news



## **NDGF-UniBern:**

- 100G link to CERN provided by SWITCH. Physical link ready, will be configured in November 2023
- Another 32bits ASN: 216467, LHCOPN community 513:3:19859
- 


## **CN-IHEP**

- Some delays with the procurement of the link to CERN. Situation now unblocked, the link should arrive very soon.

## **US-BNL:**

- BNL will increase its capacity to ESnet to 4x400Gbps. The LHCOPN link will have a capacity of 400Gbps (limited at CERN and over the Atlantic)

# LHCOPN latest news



## **KR-KISTI:**

- Link Daejon-Amsterdam upgraded to 100Gbps. Working on completing the full 100Gbps path


## **US-FNAL:**

- Working on the upgrade of the connection to ESnet

## **TW-ASGC:**

- TW-ASGC Tier1 ends in October 2023
- Working with ESnet to get LHCONE transit as ATLAS sponsored Tier2

# LHCOPN Traffic – last 12 months



	Mean	Max
Outgoing DE-KIT	53.0 Gb	251 Gb
Outgoing KR-KISTI	198 Mb	9.39 Gb
Outgoing RU-T1	6.60 Gb	42.6 Gb
Outgoing FR-IN2P3	11.5 Gb	89.4 Gb
Outgoing NDGF	6.87 Gb	83.3 Gb
Outgoing NL-T1	6.12 Gb	104 Gb
Outgoing TW-ASGC	1.14 Gb	9.69 Gb
Outgoing IT-INFN-CNAF	12.3 Gb	130 Gb
Outgoing UK-RAL	9.44 Gb	40.5 Gb
Outgoing CA-TRIUMF	5.96 Gb	75.7 Gb
Outgoing US-BNL	12.7 Gb	98.8 Gb
Outgoing US-FNAL	6.76 Gb	81.3 Gb
Outgoing ES-PIC	4.43 Gb	81.8 Gb
Outgoing PL-NCBJ	246 Mb	18.9 Gb
Total	137 Gb	479 Gb



## Numbers:

Moved ~540 PB in the last 12 months

+18% compared to previous year (457PB)

Peak at ~479Gbps

# CERN-CNAF Data-Centre Interconnect



**1.6 Tbps**

4 carriers

- DP-16QAM
- 27% SD-FEC
- 69 Gbaud

- Commissioning phase. Connected 2x 100Gbps client interfaces, 400G coming soon
- It will be used for IT-INFN-CNAF's LHCOPN link
- Plan to grow to 1.2 Tbps by HL-LHC

**LHCONE**

# LHCONE L3VPN service



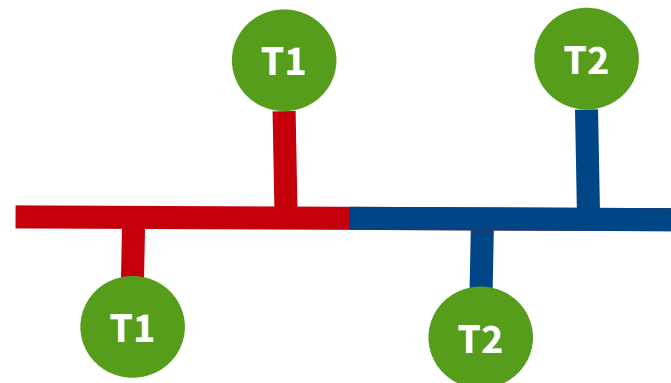
Private network connecting Tier1s and Tier2s

## Secure:

- Dedicated to LHC data transfers
- Only declared IP prefixes can exchange traffic
- Can connect directly to Science-DMZ, bypass perimeter firewalls

## Advanced routing:

- Multi domain L3 VPN
- BGP communities for traffic engineering





# Open to other HEP collaborations



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



# LHCONE L3VPN – latest news



## News

- DUNE now member of LHCONE. AUP updated
- CERN has upgraded its LHCONE connections to 400G with ESnet and 2x 400G with GEANT
- New sites:
  - Lawrence Berkeley National Laboratory (ESnet)
  - University of Massachusetts – Amherst (ESnet)
  - University of Bern-LHEP (SWITCH)
  - NCG-INGRID-PT (FCCN)
- New NRENs:
  - KIFU (HU)
  - SWITCH (CH)
  - FCCN (PT)

# LHCONE L3VPN – latest news



## **Traffic statistics:**

- continue increase
- first peak above 1Tbps seen in GEANT
- More IPv6 traffic than IPv4

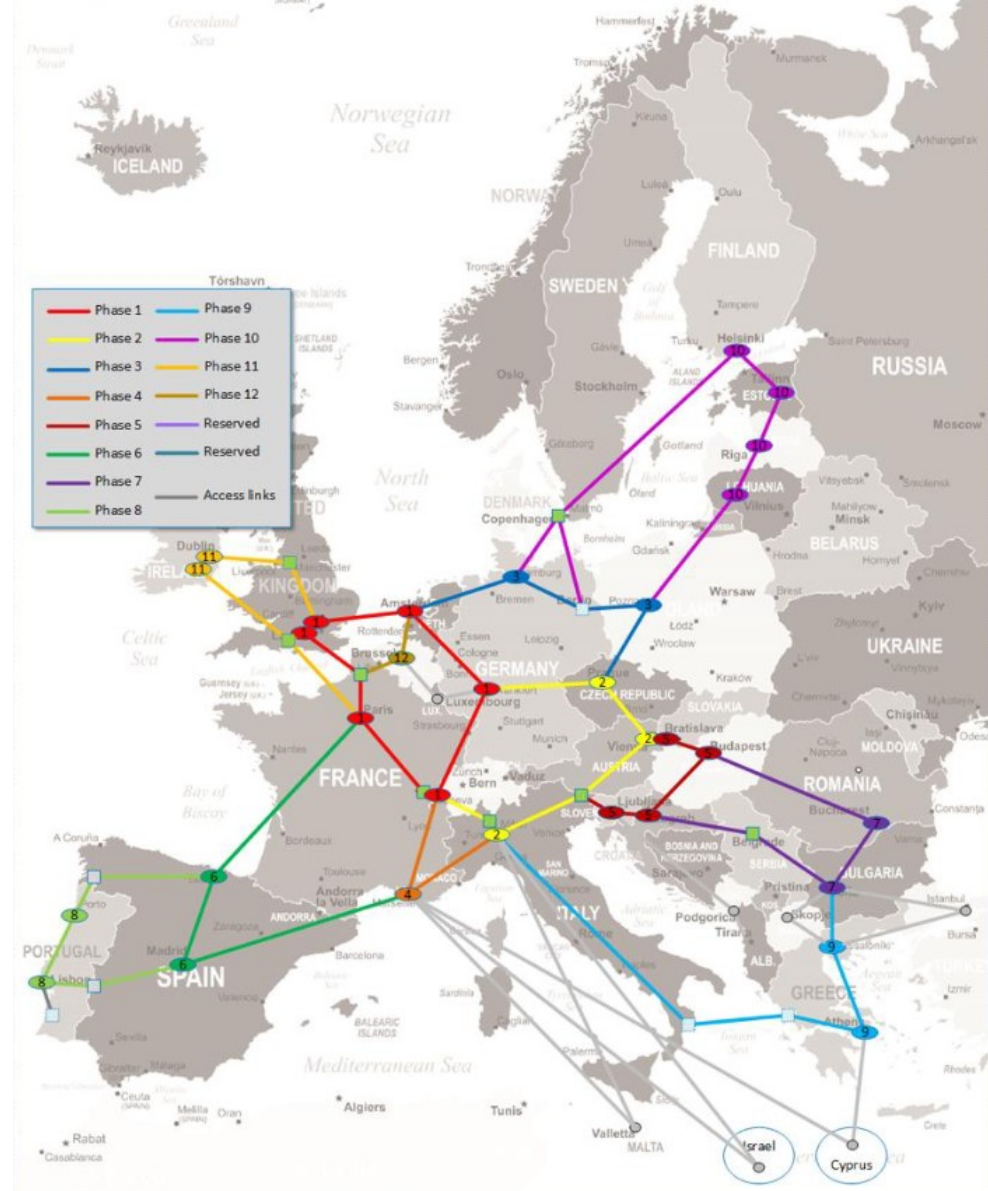
# GEANT update

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

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



# ESnet update

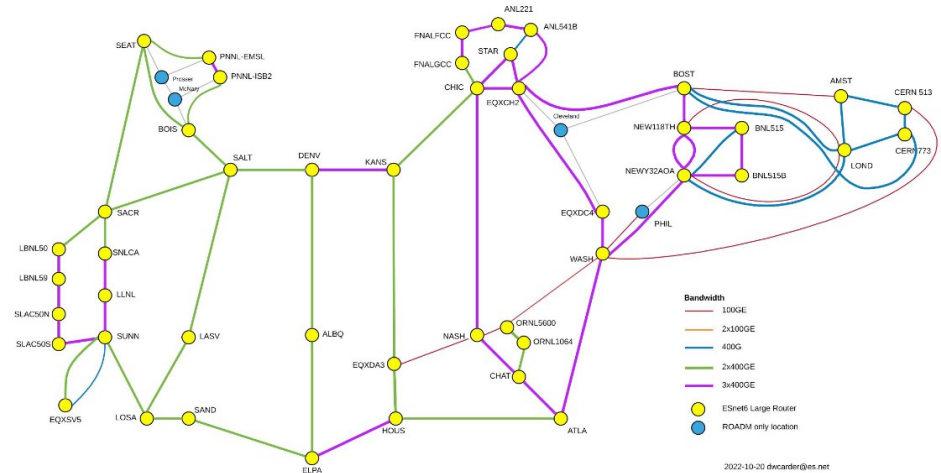
Upgraded all the links in Europe to 400Gbps

## Trans-Atlantic capacity targets:

- 500G now
- 1.5T in Q4 2023
- 3.2T in 2027, well in advance of Run 4

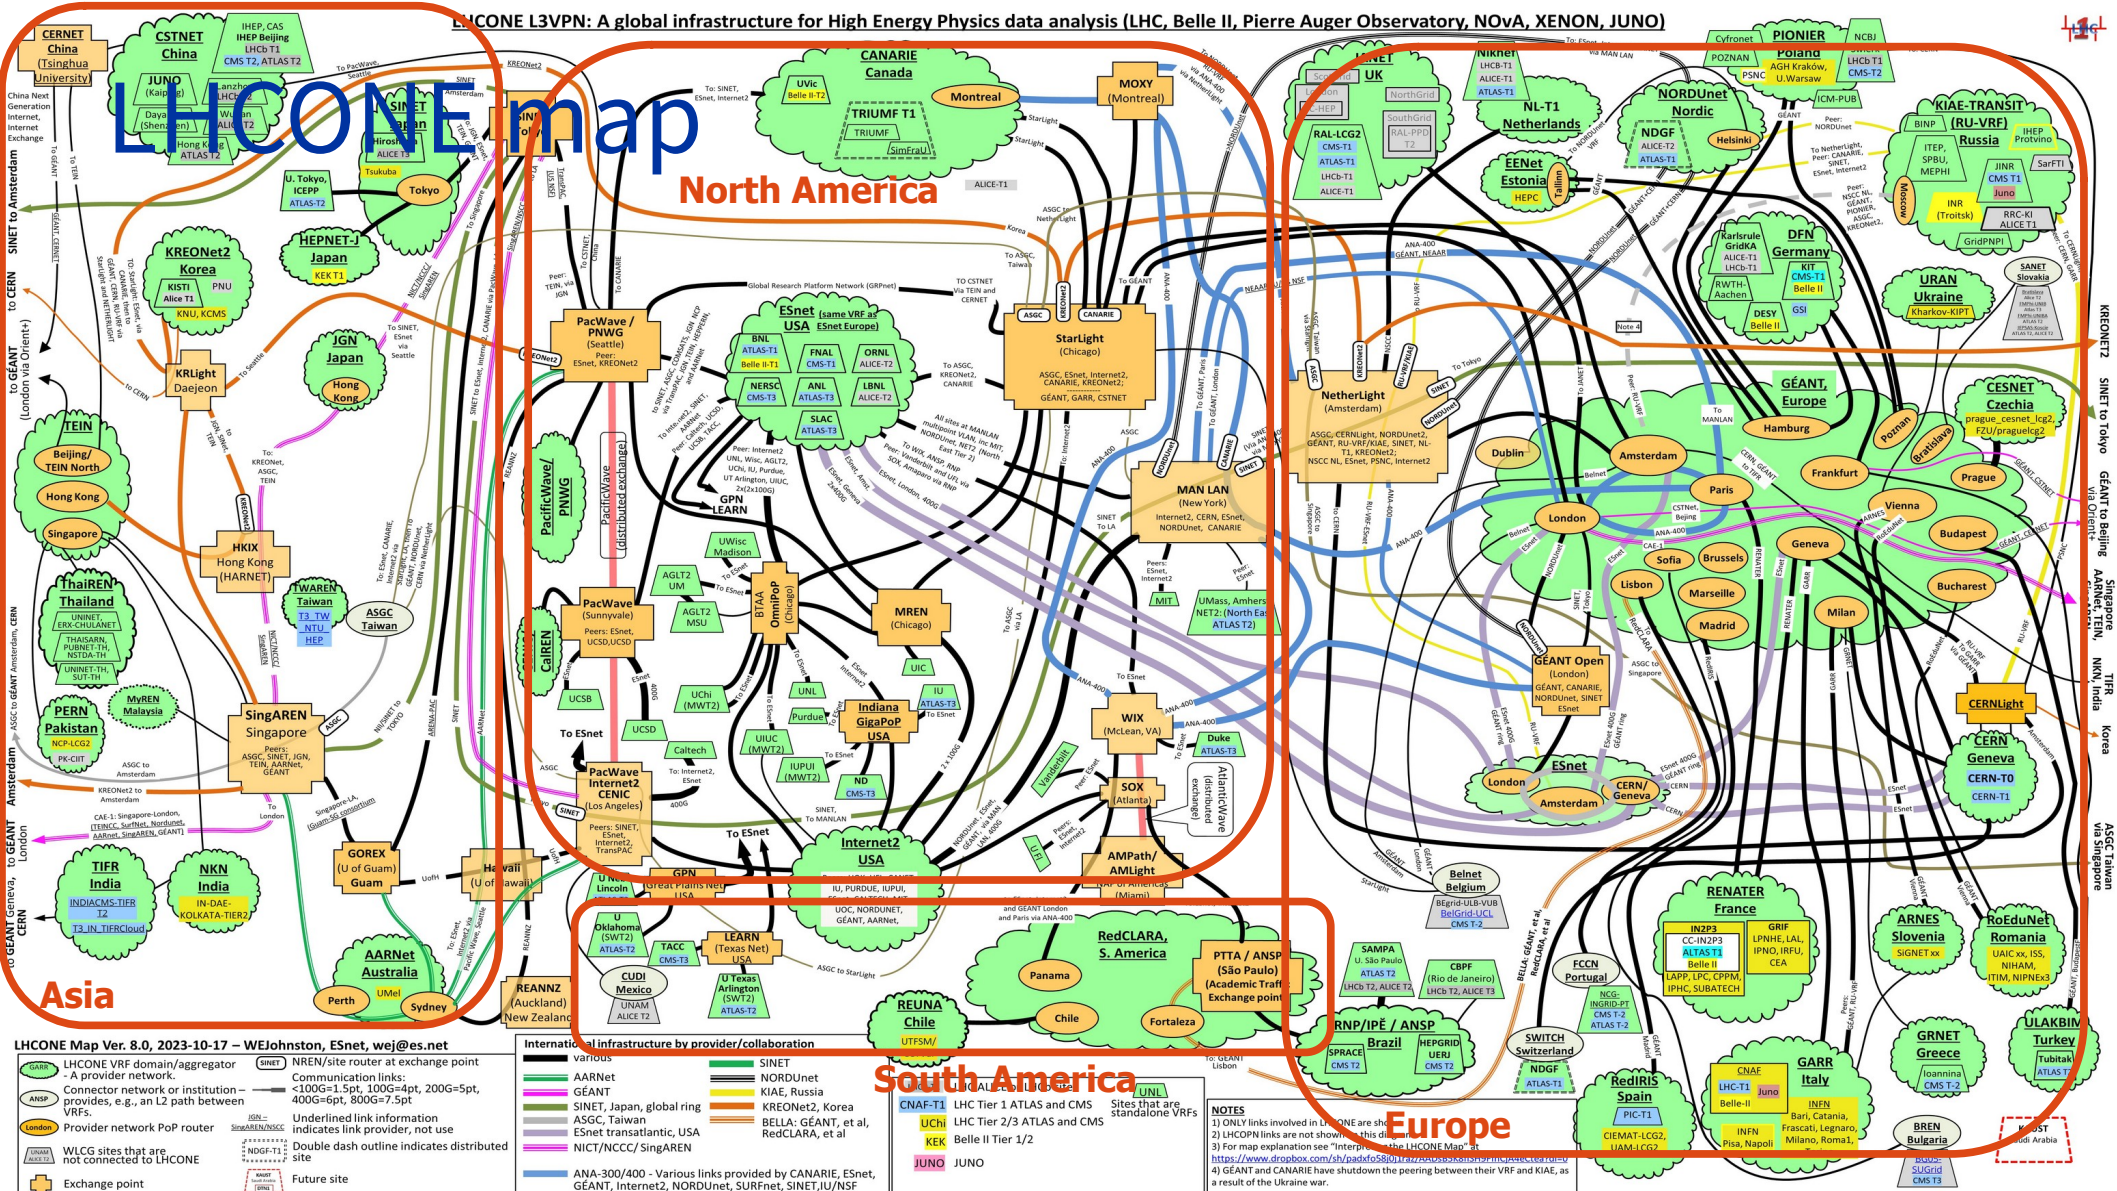
## US sites:

- BNL: connected at 2x 400Gbps
- FNAL: connected at 2x 400Gbps
- Tier2s will be upgraded to 400Gbps by 2027



# ESnet





# LHCONE map

## North America

## South America

## Europe

LHCONE Map Ver. 8.0, 2023-10-17 – WEJohnston, ESnet, wej@es.net

LHCONE VRF domain/aggregator – A provider network.

NREN/site router at exchange point – A provider network.

Connector network or institution – provides, e.g., an L2 path between VRFs.

Provider network PoP router

WLCG sites that are not connected to LHCONE

Exchange point

Future site

International infrastructure by provider/collaboration

- Various
- AARNet
- GEANT
- SINET, Japan, global ring
- ASCC, Taiwan
- ESnet transatlantic, USA
- NICT/NCC/C SingAREN
- SINET
- NORDUnet
- KIAE, Russia
- KREONET2, Korea
- BELLA: GEANT, et al.
- RedCLARA, et al

International infrastructure by provider/collaboration

- Various
- AARNet
- GEANT
- SINET, Japan, global ring
- ASCC, Taiwan
- ESnet transatlantic, USA
- NICT/NCC/C SingAREN
- SINET
- NORDUnet
- KIAE, Russia
- KREONET2, Korea
- BELLA: GEANT, et al.
- RedCLARA, et al

International infrastructure by provider/collaboration

- Various
- AARNet
- GEANT
- SINET, Japan, global ring
- ASCC, Taiwan
- ESnet transatlantic, USA
- NICT/NCC/C SingAREN
- SINET
- NORDUnet
- KIAE, Russia
- KREONET2, Korea
- BELLA: GEANT, et al.
- RedCLARA, et al

International infrastructure by provider/collaboration

- Various
- AARNet
- GEANT
- SINET, Japan, global ring
- ASCC, Taiwan
- ESnet transatlantic, USA
- NICT/NCC/C SingAREN
- SINET
- NORDUnet
- KIAE, Russia
- KREONET2, Korea
- BELLA: GEANT, et al.
- RedCLARA, et al

ONLY links involved in LHCONE are shown

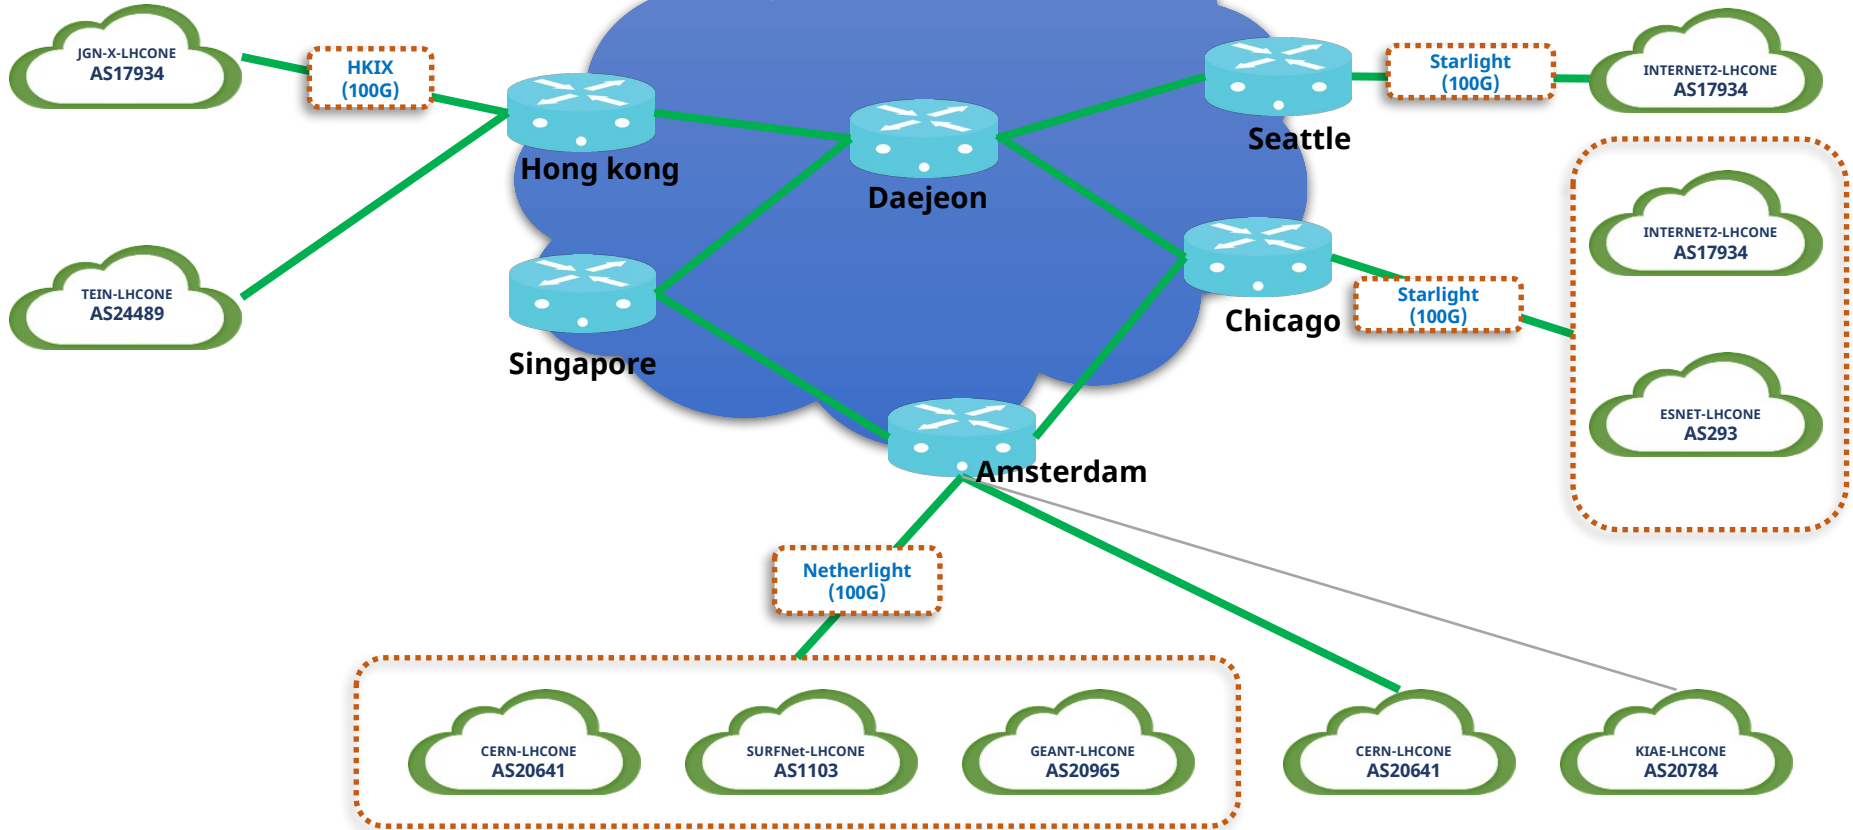
LHCOPN links are not shown

For map explanation see "International LHCONE Map"

GEANT and CANARIE have shutdown the peering between their VRF and KIAE, as a result of the Ukraine war.

# LHCONE - KREONet2

## KREONet2 LHCONE





# LHCONE - SINET

Amsterdam, NL:

GEANT LHCONE, SURFnet LHCONE, RRC-KI LHCONE, CERNLight LHCONE

Los Angeles, US:

ESnet LHCONE, Internet2 LHCONE, CANARIE LHCONE, AARNet LHCONE

New York, US:

GEANT LHCONE

Singapore, SG:

AARNet LHCONE, NKN LHCONE

Guam, US:

AARNet LHCONE



# LHCONE - TransPAC

- peerings with Internet2 and ESnet in Seattle.
- working on a peering with Taiwan in Japan

TransPAC also carries some LHC traffic to the US for AARnet, JGN, Kreonet, CSTnet, etc via Seattle.

# LHCONE monitoring

- perfSONAR 5 is out and being deployed Some bugs have pushed new release (latest 5.0.5)
- perfSONAR 5 uses Elasticsearch and Grafana
- 100G mesh: data is now shown correctly , but results are not great.  
Work in progress

**perfSONAR**

# LHCONE Looking Glass

Running looking-glass to analyse the routing tables of the VRFs

Implemented on a CERN router. Now peering with these VRFs:

- ASGC AS24167
- CANARIE AS6509
- CERNlight AS20641
- ESnet AS293
- KREOnet AS17579
- GEANT AS20965 (Geneva and Frankfurt routers)
- NORDUnet AS2603
- RU-VRF AS57484

The looking glass is accessible at <http://lhcone-lg.cern.ch/>

# Network information in CRIC

CRIC (Computing Resources Information Catalogue) is the database used by WLCG to document the available resources. It is used also to store network information related to LHCOPN and LHCONE

## **Easily accessible**

- Netsite: <https://wlcg-cric.cern.ch/core/netsite/list/> (login required)
- NetworkRoute: <https://wlcg-cric.cern.ch/core/networkroute/list/> (login required)
- Json view: <https://wlcg-cric.cern.ch/api/core/rcsite/query/?json> (no login)

# CRIC NetSite table

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

- NOC contact email
  - URLs to monitoring and other info pages
  - LHCONE AUP acknowledgment
  - LHCONE and LHCOPN participation
  - Network bandwidth: WAN, LHCOPN, LHCONE (-1 = not connected)
  - LHCONE providers
- ~140 entries

NetSite list Show 100 entries

RC Site	NetSite	NOC	monit URL	info URL	AUP	LHCONE active	LHCOPN active	WAN	LHCONE	LHCOPN	providers
AGLT2	<a href="#">US-AGLT2 Michigan State University</a>	aglt2-noc@umich.edu			✓	✓	✗	100	100	-1	ESnet
AGLT2	<a href="#">US-AGLT2 University of Michigan</a>	aglt2-noc@umich.edu			✓	✓	✗	80	80	-1	ESnet
ANLASC	<a href="#">US-ANL</a>	noc@anl.gov			✓	✓	✗	0	100	-1	ESnet
ARNES	<a href="#">SL-ARNES-NREN</a>				✓	✓	✗	0	9	-1	ARNES
ARNES	<a href="#">SL-IJS-Ljubljana</a>				✓	✓	✗	0	9	-1	ARNES
ARNES	<a href="#">SL-IZUM-Maribor</a>				✓	✓	✗	0	9	-1	ARNES
Australia-ATLAS	<a href="#">AU-Australia-ATLAS</a>				✓	✓	✗	20	10	-1	AARNET
BEgrid-ULB-VUB	<a href="#">BE-ULB-VUB</a>				✗	✗	✗	20	-1	-1	

# CRIC NetworkRoute table

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

One entry per set of prefixes sharing common routing policies

- set of v4 and v6 prefixes
  - “More specific” flag (in case of disaggregated prefixes)
  - Autonomous System number (ASN)
  - collaborations using these prefixes
  - LHCOPN and LHCONE bandwidth specific to the prefixes of the record
  - monitoring URL specifics to the prefixes of the record
- ~150 entries

RC Site	NetworkRoute	NetSite	ASN	monit URL	MS	Subnets	LHCONE limit	LHCOPN limit	collaborations
AGLT2	<a href="#">AGLT2_LHCONE_RT</a>	US-AGLT2 Michigan State University	229		✘	2001:48a8:68f7::/48, 2001:48a8:68f7::/50, 2001:48a8:68f7:4000::/50, 2001:48a8:68f7:8000::/50, 2001:48a8:68f7:c000::/50	100	-1	US-ATLAS, WLCG
AGLT2	<a href="#">AGLT2_MSU</a>	US-AGLT2 Michigan State University	229		✘	192.41.236.0/23, 192.41.238.0/28	100	-1	US-ATLAS, WLCG
AGLT2	<a href="#">AGLT2_MSU IPv6</a>	US-AGLT2 Michigan State University	237		✘	2001:48a8:68f7:8001::/64	100	-1	US-ATLAS, WLCG

# 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 are willing to help with future tests




# 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 than the previous one, since it uses 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



# WLCG Data Challenge 2024

## **Next data challenge (DC24)**

- Planned for 12-23 February 2024
- Target of 25% of HL-LHC requirements
- Network providers are increasing the network capacity
- Defined the list of features that will be tested during DC24
- Preparation Workshop at CERN 9-10 of November at CERN

# DC24 projects

List of the projects on Network:

- 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

# Conclusions

# Summary

- CERN networks: developing for Run4
- LHCOPN: Three new Tier1s connected
- LHCONE continue to grow. DUNE has joined it
- LHCONE R&D: several projects on going
- WLCG DC24: network projects already preparing

An aerial photograph of the CERN facility in Geneva, Switzerland, taken at sunrise. The sun is low on the horizon, creating a golden glow and long shadows. The image shows the large, circular, ribbed dome of the Large Hadron Collider (LHC) in the foreground on the left. A central road with multiple lanes runs through the center, flanked by various buildings and parking lots. In the background, there are rolling hills and mountains under a hazy sky. The overall scene is a mix of industrial architecture and natural landscape.

*Questions?*

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