



International networking for LHC

Communication Systems group 40th years anniversary

1st November 2023

Edoardo Martelli

LHC Computing Model

Back in the late 1990s, it was already clear that the expected amount of LHC data would far exceed the computing capacity at CERN alone. Distributed computing was the sensible choice.


The first model proposed was MONARC (Models Of Networked Analysis at Regional Centres for LHC Experiments), on which the experiments originally based their computing models.

In September 2001, CERN Council approved the first phase of the LHC Computing Grid project, led by Les Robertson of CERN's IT department. From 2002 to 2005, staff at CERN and collaborating institutes around the world developed prototype equipment and techniques. From 2006, the LHC Computing Grid became the Worldwide LHC Computing Grid (WLCG) as global computing centres became connected to CERN to help store data and provide computing power.

WLCG Tier Structure

WLCG uses a tier structure with the CERN data centre as Tier-0. CERN sends out data to each of the 11 major data centres around the world that form the first level, or Tier-1.

Each Tier-1 site is then linked to a number of Tier-2 sites, usually located in the same geographical region.



Datatag

DataTAG was one of the many projects that took care of prototyping the different components of the Computing Grid

DataTAG ran from 2002 to 2004 with the aim of testing high speed network devices and improve performance of data transfers over long distances.

DataTAG had a lab made of pair of routers provided by Cisco, Juniper and Alcatel. The routers were installed at CERN and at Starlight in Chicago. The two labs were connected by a transatlantic link, which started at 2.5Gbps and was later upgraded to 10Gbps

Some DataTAG routers in the External Network area of the B513 computer centre



*Elise Guyot, Olivier Martin and Edoardo Martelli,
IT-CS members of DataTAG*



First WLCG router at CERN

Juniper was initially selected as the platform to connect the Tier0-Tier1 links because of being one of the first routers supporting 10Gbs interfaces.

It was quite soon replaced by cheaper Force10 routers and sold to the second-hand market



Martin Fluckiger (IT-CS-EN) install a linecard on a brand-new Juniper MX router.

The birth of LHCOPN

From the early 2005, the network teams of the Tier0 and Tier1s sites and the Research and Education Network providers (REN) started meeting regularly to design the architecture of the network to connect the 11 Tier1s to the Tier0.

They designed a network with a star topology with CERN Tier0 in the centre.

The network was called **LHC Optical Private Network, LHCOPN**. The word **Optical** was chosen because SURFnet wanted to highlight the use of innovative transmission technologies of the time. **Private Network** because the use of the link was reserved to the 12 sites only to transfer WLCG data.

The initial 11 Tier1 sites were: CA-TRIUMF, DE-KIT, ES-PIC, FR-IN2P3, IT-INFN-CNAF, NDGF, NL-T1, TW-ASGC, UK-RAL, US-BNL, US-FNAL.

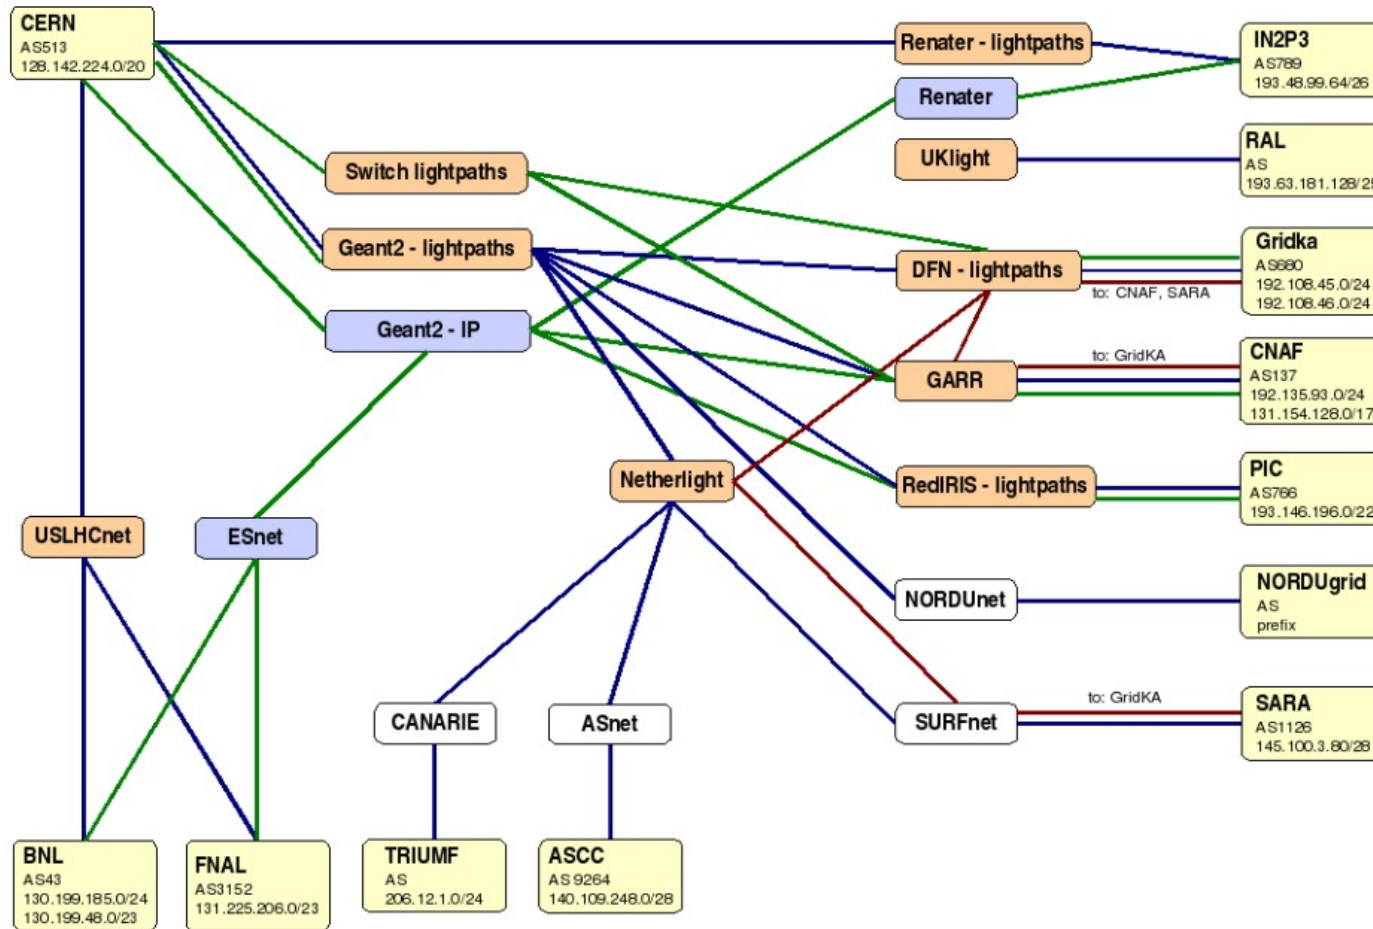
LHCOPN: LHC Optical Private Network

Will be used to move the data produced by the experiments to the Tier1 data centres all over the world.

Only the Tier0-Tier1s traffic will be carried by this network.

Every Tier1 will have a main dedicated 10Gbps link to CERN (lightpath) and a 10Gbps backup connection.

The lightpaths will be provided by Geant2 (Dante), Renater, USLHCnet (Caltech) and several NRENs [National Research and Education Network].



edoardo.martelli@cern.ch - last update: 20051005

LGC Network

The LGC network will interconnect the experiments and the server farms (CPUs, Disks, Tapes).


The backbone consists in eight Force10 E1200. They are being installed and configured in the Computer Center.

They will be interconnected with 10 times 10Gbps ethernet links. The core will be able to forward 1.2Tbps



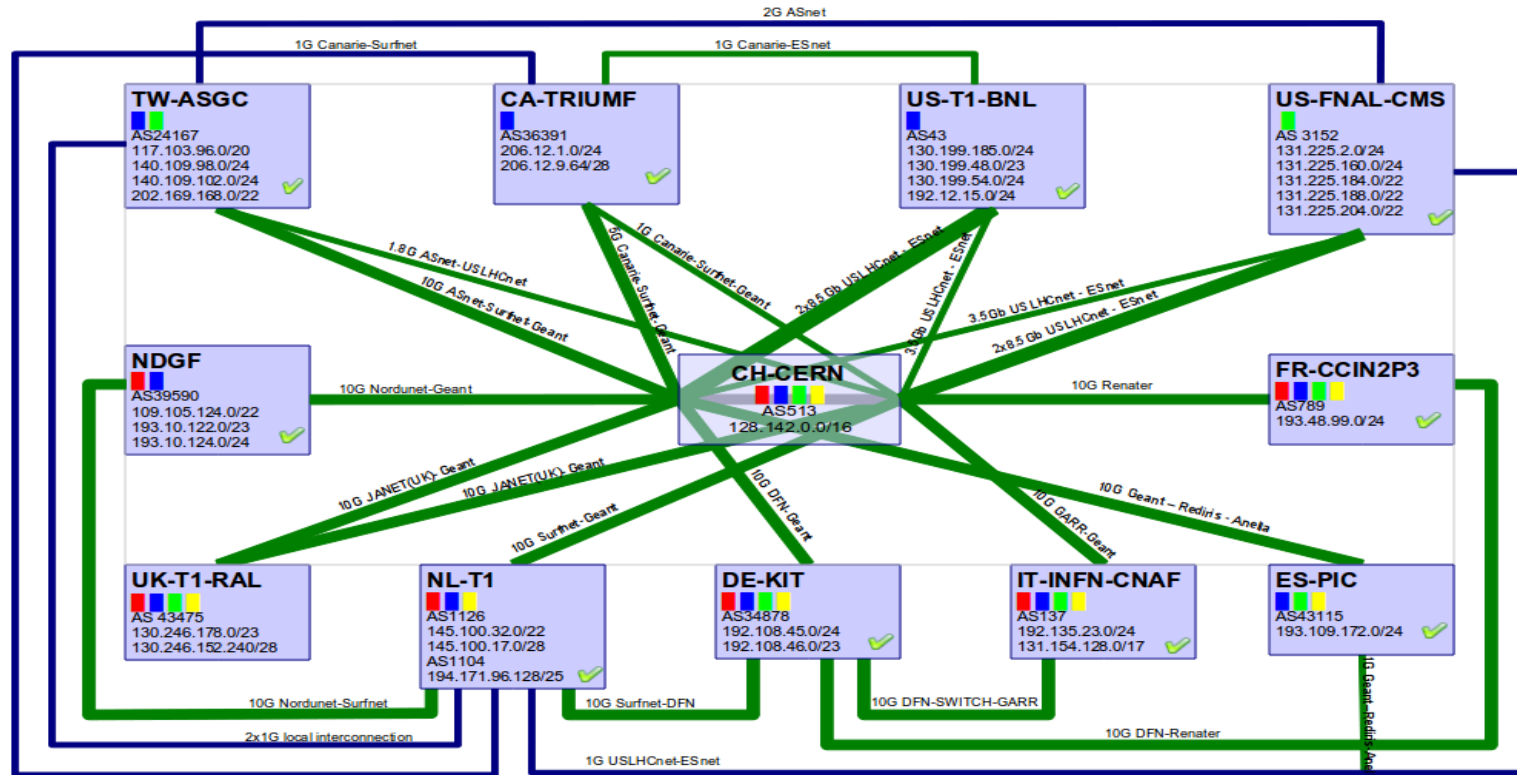
LHCOPN logo

The LHCOPN logo was designed in 2009 by Rosy Mondardini, WLCG communication person



LHCOPN 2010

LHCOPN



- T0-T1 and T1-T1 traffic
- T1-T1 traffic only
- Not deployed yet
- (thick) >= 10Gbps
- (thin) <10Gbps
- = Alice
- = Atlas
- = CMS
- = LHCb
- ✓ = internet backup available
- p2p prefix: 192.16.166.0/24
- eduardo.martelli@cern.ch 20100916



LHCOPN evolution

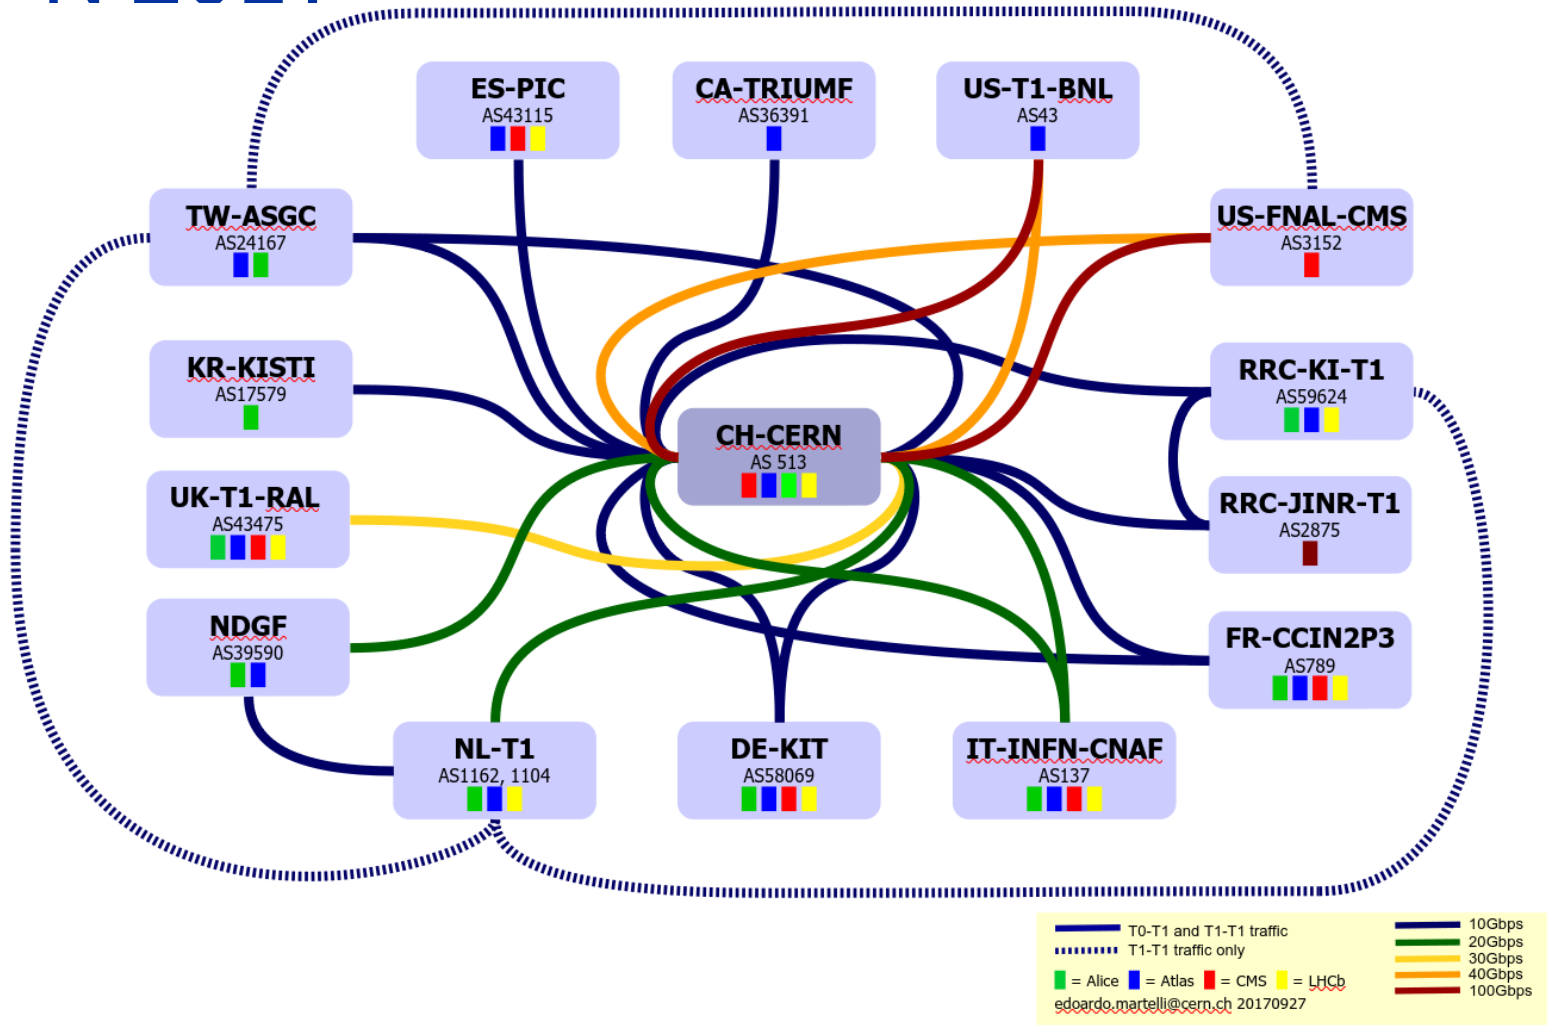
The topology later evolved in a star and a partial mesh, with direct links between some pairs of Tier1s.

In 2013 RU-KI and RU-JINR joined LHCOPN and in 2014 it was the time of KR-KISTI.

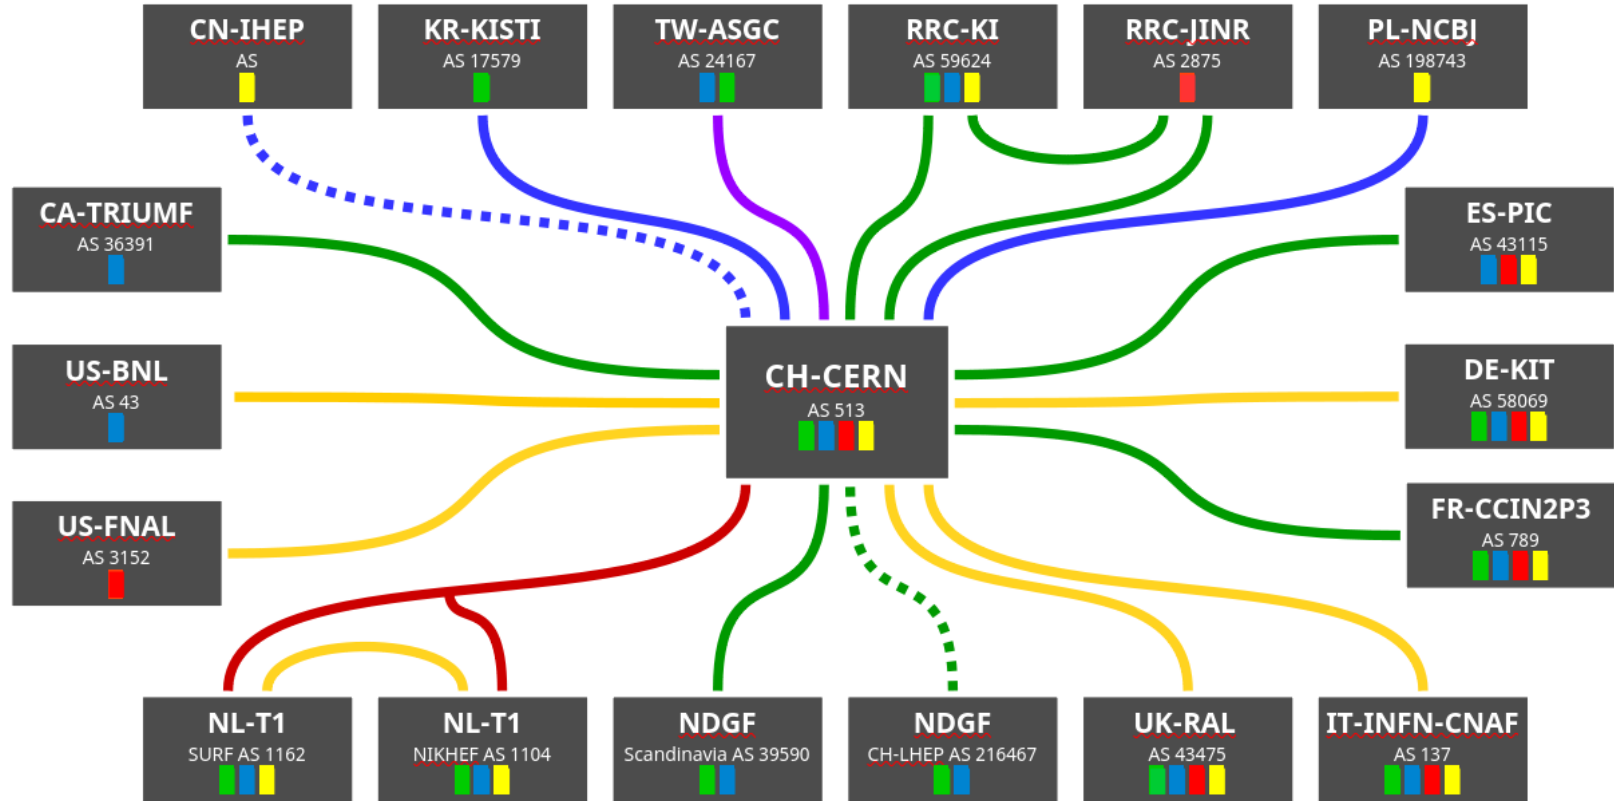
During the years the Tier1-Tier1 links were discontinued, replaced by LHCONE.

The membership of the network didn't change till 2023, when TW-ASGC announced the decision to discontinue the Tier1 role and disconnect from LHCOPN. At the same time three new sites have started the procedure to become Tier1 and are currently in the process to connect to LHCOPN: PL-NCBJ, CN-IHEP. CH-LHEP (part of NDGF)

LHCOPN 2017



LHCOPN 2023



■ = Alice
 ■ = Atlas
 ■ = CMS
 ■ = LHCb

— 10Gbps
— 20Gbps
— 100Gbps
— 200Gbps
— 400Gbps

edoardo.martelli@cern.ch 20231013

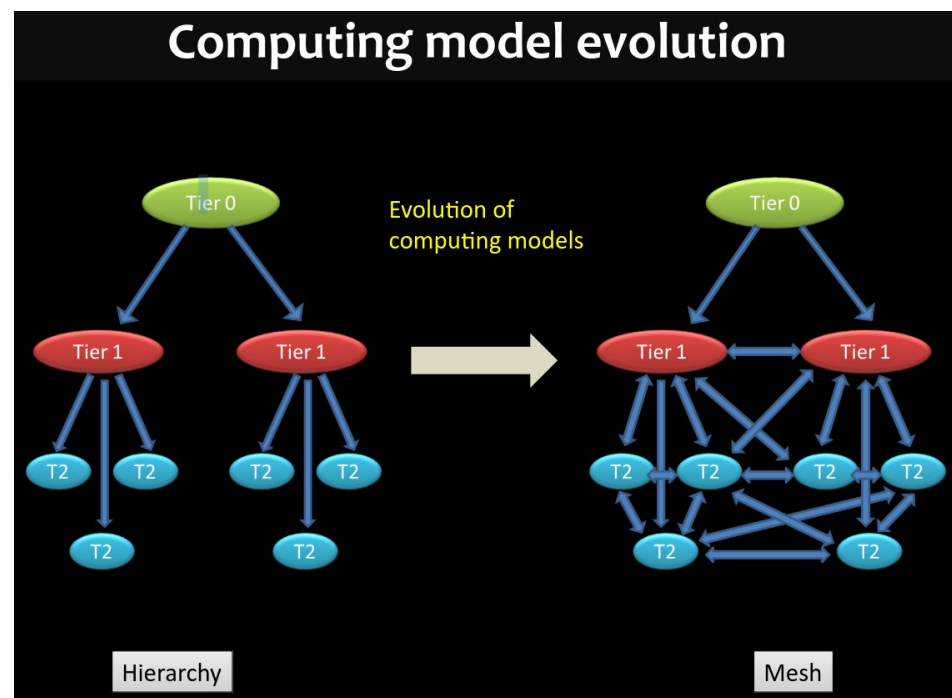


Evolving computing model

In 2010, during a successful Run1, the LHC experiment remarked that the network performed much better than expected, and better than other computing components.

They decided to evolve the computing model, such that the Tier2s were no longer bound to their closer Tier1

The LHCOPN community was involved to develop a network to better serve Tier1s and Tier2s



First community events

2010, June: Transatlantic connectivity Workshop at CERN

- Evolving computing models required better connectivity for Tier2s to reach any Tier1

2010, October 2010: LHCOPN meeting at CERN

- Tier2s connectivity WG mandate from LHCOPN community

2010. November-December: preparation of Tier2s connectivity proposals

- ESnet, GEANT, CERN+SURFnet work independently to prepare architecture proposals

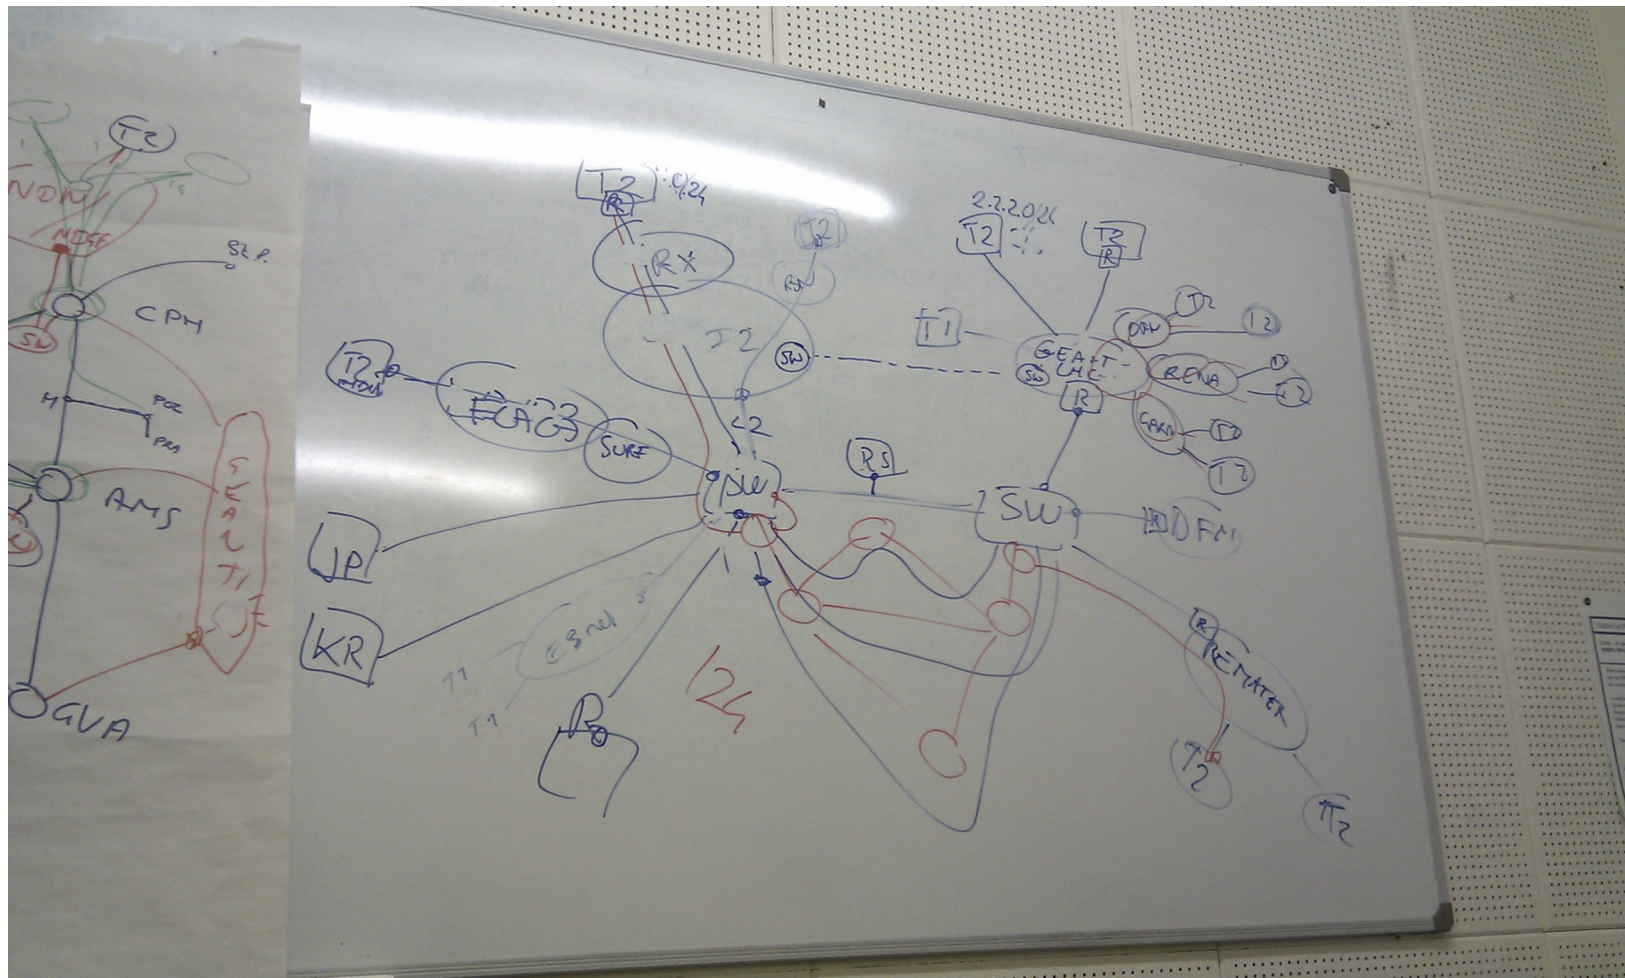


*Transatlantic connectivity workshop at CERN, 2010
David Foster (IT-CS group leader) on the left, Artur Barczyk (Caltech) standing*



Transatlantic connectivity workshop at CERN, 2010

LHCONE brain-storming, CERN October 2010



LHCONE architectures

2011 January: Workshop on connectivity for Tier2s at CERN

- First T2s connectivity Technical Meeting
- Two proposals: distributed IXP (L2VPN), Federated VRFs (L3VPN)

2011:

- GEANT, ESnet, CERN, Internet2, SURFnet, NORDUnet implement together a L2VPN prototype
- main issues: loops, unused transatlantic links

2011, December: LHCONE architecture workshop at SARA

- L2VPN turns out to be too complicated in a multi-domain network. Dropped.
- New design proposed: multidomain L3VPN
- It's a success. The main RENs implement their own LHCONE instance and peer together

First LHCONE meeting in Lyon, 2011

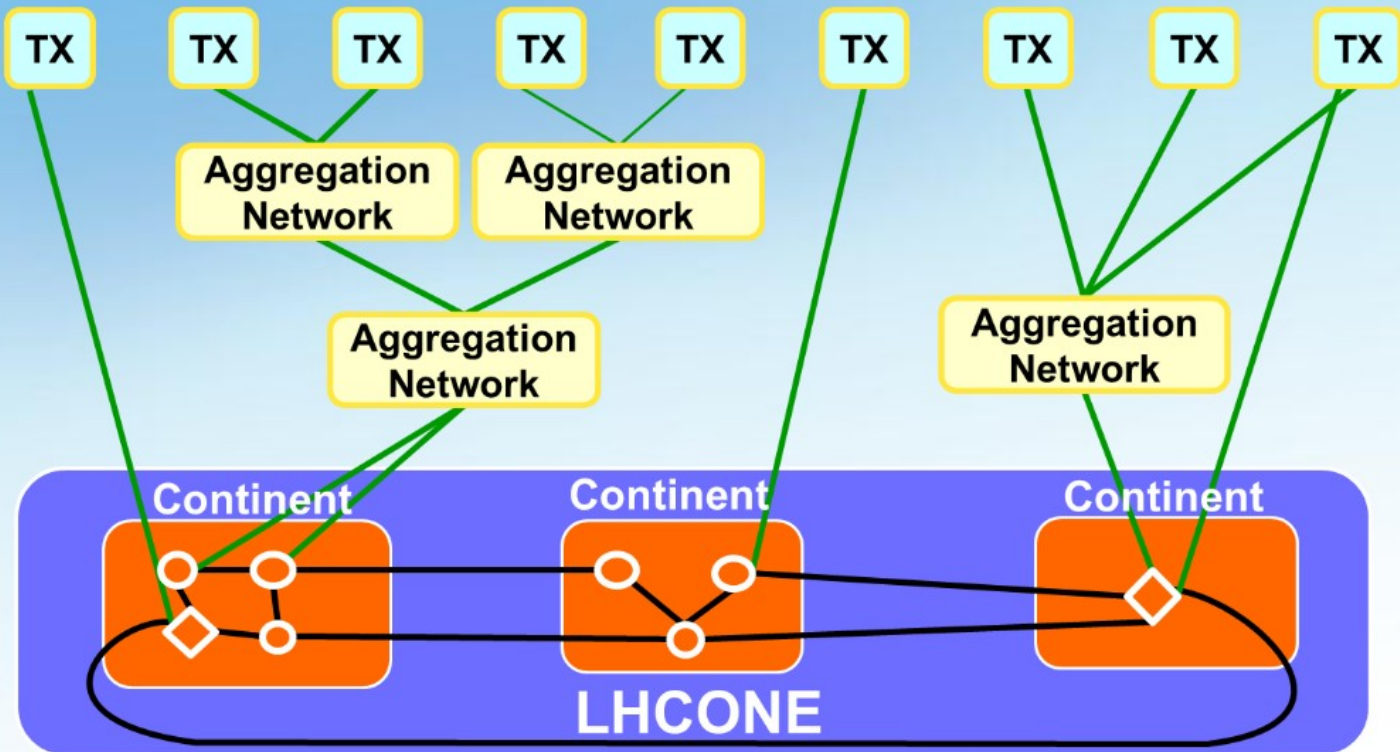
Three members of IT-CS:

- *David Foster (GL)*
- *John Shade*
- *Edoardo Martelli*



LHCONE L3VPN, January 2012

LHCONE High-level Architecture



Layer3 VPN proposal presented during the workshop during the LHCONE meeting at the Internet2 offices in Washington DC

LHCONE evolution

2013, January: LHCONE extends to Asia

- ASGC implement a VRF and reach the other VRFs via Netherlight (SURFnet)

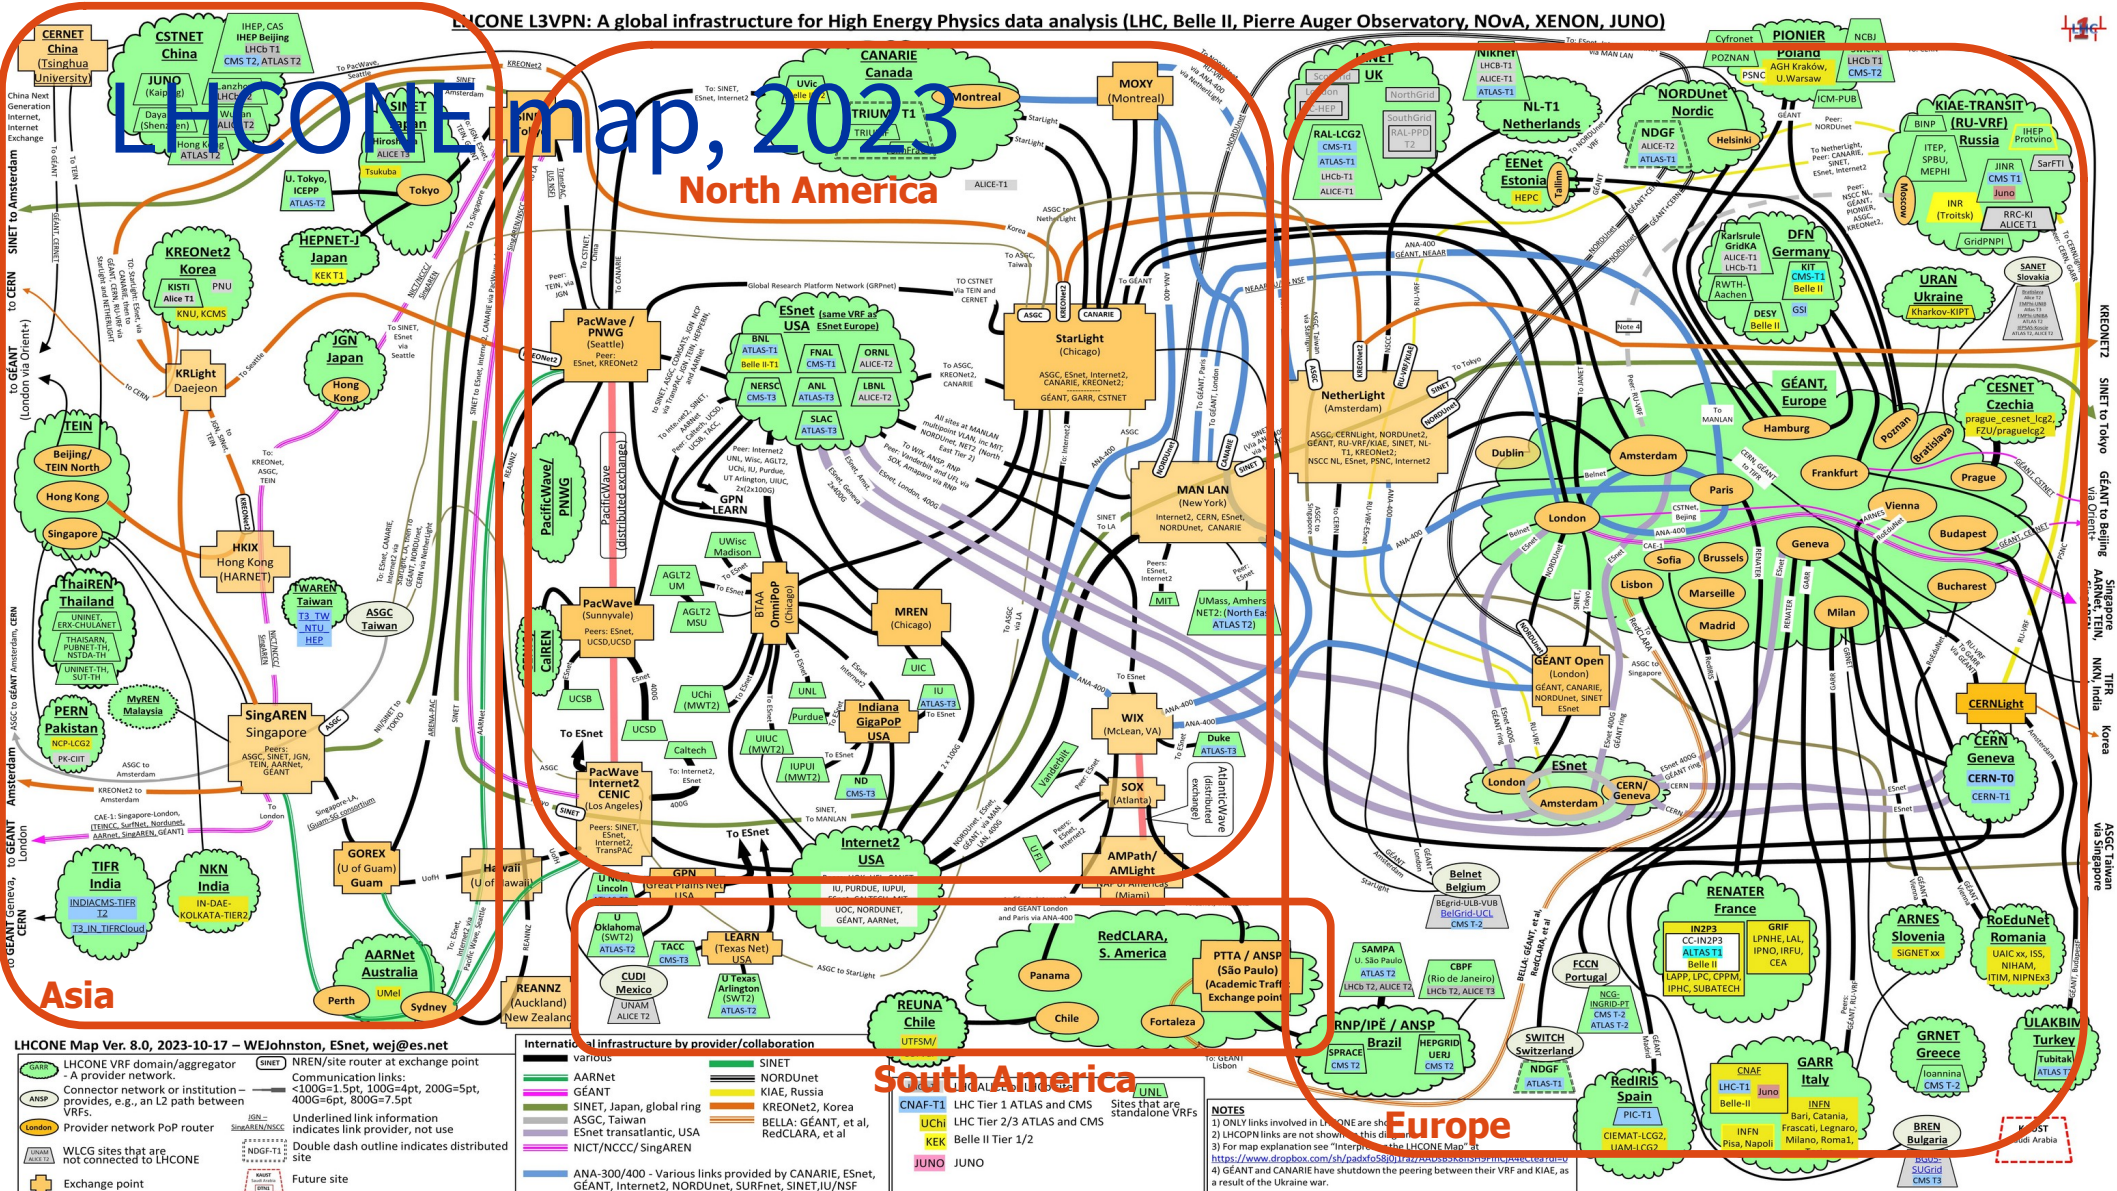
2018, LHCONE extends to South America

- RNP implements the LHCONE VRF for the Brazilian Tier2s

Over the years, other High Energy Physics collaborations and experiment has joined LHCONE, mostly because many of their sites were already connected to LHCONE. As of today, these collaborations are part of LHCONE: BelleII, DUNE, JUNO, NOvA, Pierre Auger Observatory, US-ATLAS, US-CMS, XENON and obviously WLCG

LHCONE map, 2023

North America



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

LHCONE VRF domain/router at exchange point
 - A provider network/aggregator
 - 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
- GÉANT
- SINET, Japan, global ring
- ASGC, Taiwan
- ESnet transatlantic, USA
- NICT/NCC/C SingAREN
- SINET
- NORdUnet
- KIAE, Russia
- KREONET2, Korea
- BELLA: GÉANT, et al.
- RedCLARA, et al

South America

UNL LHC Tier 1 ATLAS and CMS
 UChi LHC Tier 2/3 ATLAS and CMS
 KEK Belle II Tier 1/2
 JUNO JUNO

Sites that are standalone VRFs

Europe

ONLY links involved in LHCONE are shown
 LHCOPN links are not shown
 For map explanation see "Interpreting the LHCONE Map"
 GÉANT and CANARIE have shutdown the peering between their VRF and KIAE, as a result of the Ukraine war.

RENAT France
 GARR Italy
 RedIRIS Spain
 ARNES Slovenia
 RoEdUnet Romania
 GRNET Greece
 ULAKBIM Turkey
 BREN Bulgaria
 CERN Geneva
 CERN-TO
 CERN-T1

LHCONE logo

The LHCONE logo was designed around 2012 by Artur Barczyk, Caltech network architect



Questions?

edoardo.martelli@cern.ch

