LHCOPN

History of LHCOPN

Bruno Hoeft
Karlsruhe Institute of Technology

Agenda

- First initial steps
- Evaluate different topology
- IP-Design WG
- Mission Statement
- Operation WG
- Evolvement – bandwidth + additional sites
- Service challenges
- IPv6
- Responses of former colleagues

Disclaimer

- It is not possible to cover all LHCOPN history in these slides
- Slides are mainly about LHCOPN history
  - Tier-2/3… are part of the early discussions/efforts history and evolvement of LHCONE, they have no part in these slides
- LHC accelerator should start in 2007/8
  - This explains the pressure in planning how to achieve the required load

LHC Challenge

MONARC
Models Of Networked Analysis at Regional Centers

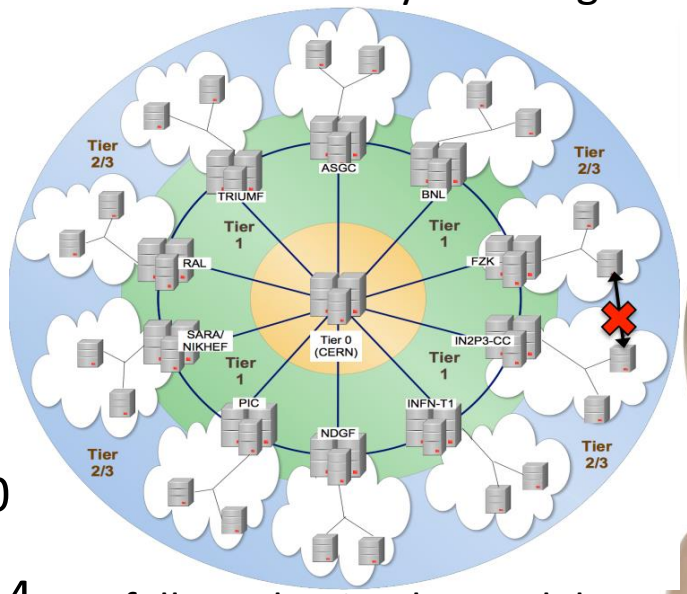

followed strict the model

- Towards the millennium a decision was made to a distributed computing engine for LHC (WLCG).
- Beginning of 2000 –
  the biggest challenge were the expected high data rates
  - detector data ⇔sent straight to the Tier-1s
    - for storing and first compute run (at CERN was no capacity)
  - extremely stable and “unbreakable” network between Tier-0 and Tier-1
  - In 2004 the connection to CERN was only 1 Gbps (most STM-4 and less)
  - **not** sufficient ⇔ expected high data load of LHC
- Political navigation
  - Dante (/Géant) as the aggregator and interconnect of the European NRENs
  - Besides that LHCOPN required “cross-border-fiber” between NRENs
  - Investments in transatlantic links were necessary (ESnet/USLHCnet)

First initial steps

First Meeting

- Davis Foster steering
- Kors Bos (chair of GDB) invited to „Data Challenge Meeting LHC“ at 12.-13.10.2004 in Amsterdam (SARA)
  - Small group of people (6 or 7) (CERN/NRENs/Tier-1s)
  - Only topic : how to build a network reaching the expected data rate
    - 100PetaByte/Year Detector Data
- Envisioned start of LHC : 2007 running up / 2008 first collision
- Optical Lightpath Network was just developed
    - NRENs establishing their first backbone with this technique

Next meeting (Jan. 20/21 2005)

- „T0/T1 Network Meeting“ – CERN much bigger (39 people):
  - of three Continents, nine Tier-1s and two Tier-2s, three experiment representatives, six NRENs
- Topics:
  - Who pays for what ?
  - Who ensures the service – operations – evolution ?
  - Who ensures the funding – longevity ?
  - How to manage security, firewalling, especially at high speed (10Gb/sec)?
  - Different working groups were built:
    - **Architecture** (Erik-Jan Bos)
    - **IP Design** (Edoardo Martelli)
    - **Operation** (Roberto Sabatino, Mathieu Goutelle)
    - **Security** (Robin Tasker)
    - **Monitoring** (Shawn McKee)
- The Group effort started
  ⇔ essential to keep all stake holder together
  - The Tier-1 Manager and Network Administrator
  - Géant/ESnet and all other NRENs
- During 2005 the abbreviation **LHCO**ptical**P**rivate**N**etwork was chosen

layer-2 approach

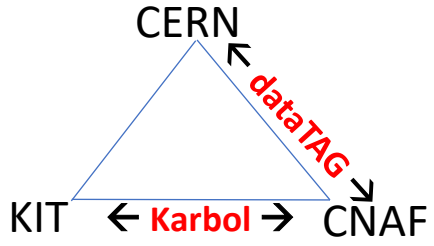

Next step of a triangle between CNAF (Bologna) KIT (Karlsruhe) and CERN

KarBol Project

Layer 2 Circuit between Karlsruhe and Bologna crossing DFN, GÉANT and GARR domains

Authors: Laura Leone, GARR *(laura.leone@garr.it)*
Marco Marletta, GARR *(marco.marletta@garr.it)*
Tiziana Ferrari, CNAF *(Tiziana.Ferrari@cnaf.infn.it)*

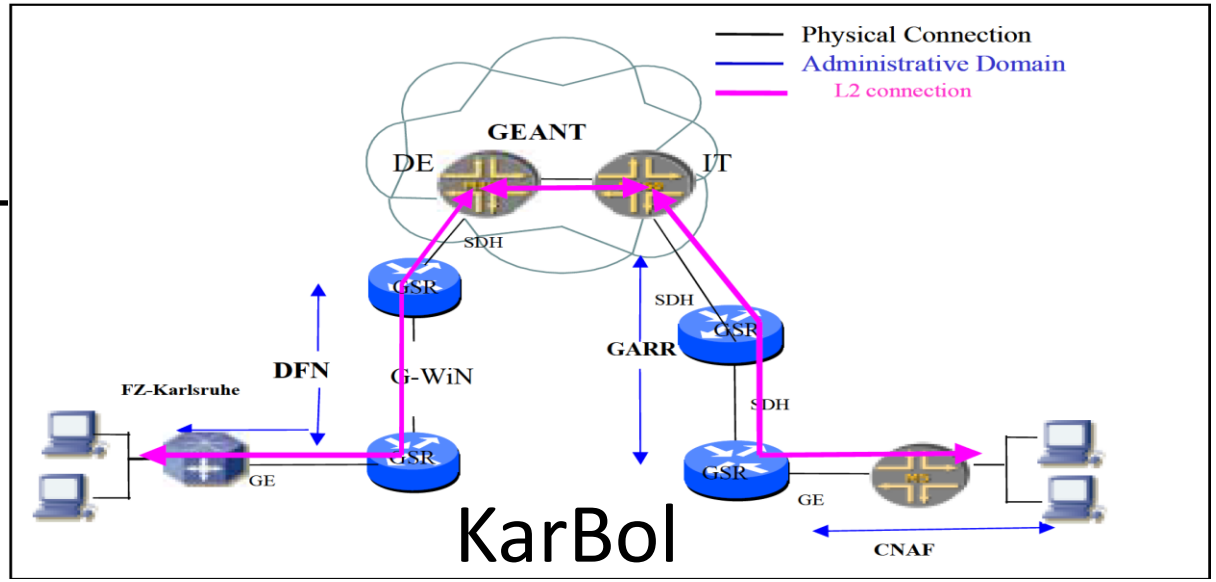

The MPLS-Layer-2-VPN approach was successfully deployed but dropped

- scalability issues
- the difficulty to deploy cross multidomain layer-2 circuits at that time

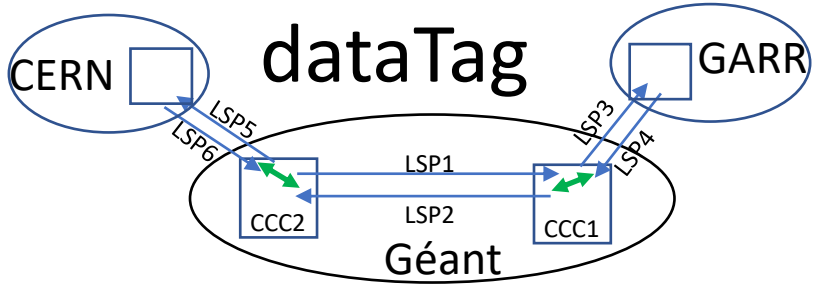

Evaluate different topology

different topologies

- ring topology did not support the required structure
- full mesh topology – suited best, but was too complex to deploy
- CERN centric Star topology was chosen

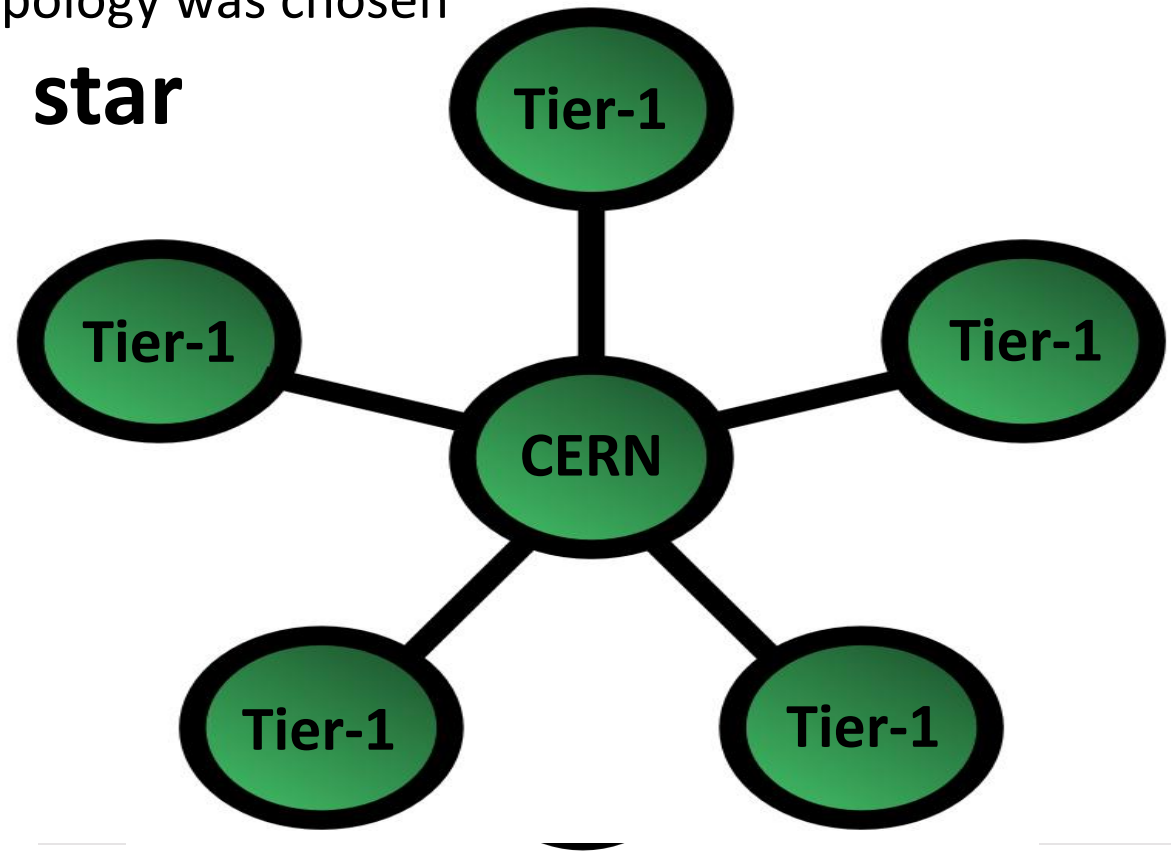

IP-Design WG

IP-Design

- IP-Network Group https://twiki.cern.ch/twiki/bin/view/LHCOPN/LHCopnRoutingDoc
  Layer-3 BGP deployed between CERN and Tier-1s
  - LHCOPN transitnet 192.16.166.0/24
    - CIDR /30 between Tier-0 and Tier-1
    - with BGP advertisement of the reachability
    - each Tier-0/1 needs its own AS-#

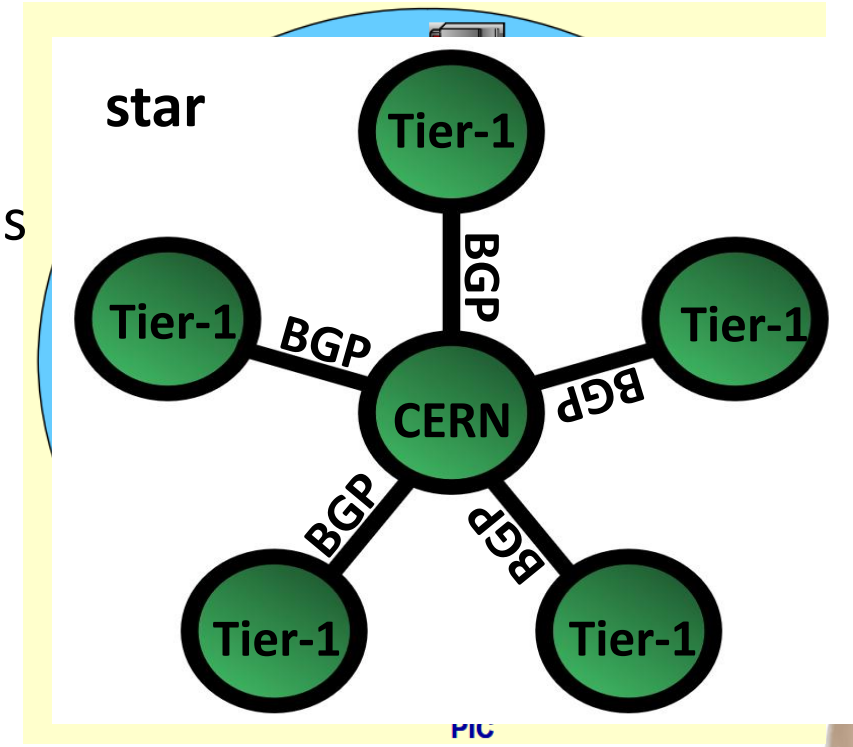

meeting Amsterdam, April 2005

- define LHCOPN model
- T1s to be connected
  →all 10 Tier-1s
  →Stated in the MOU
- already two examples of 2005

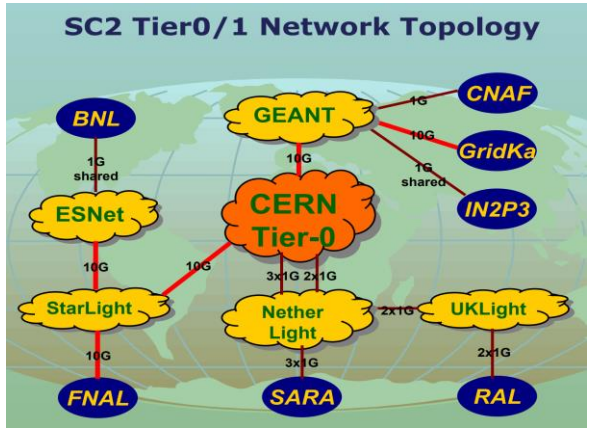

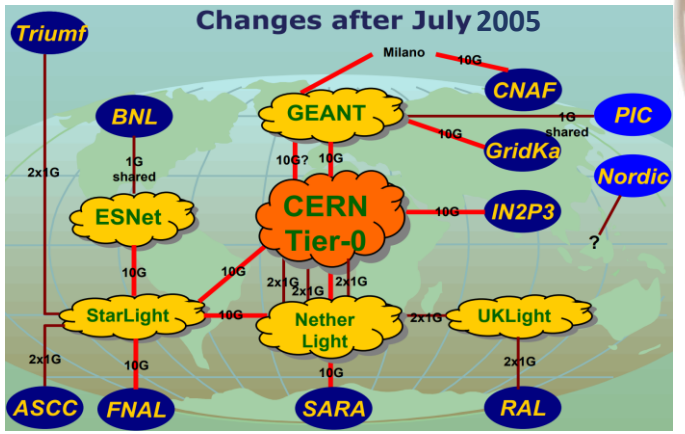

Mission Statement

LHC Mission statement

In the early state of building LHCOPN and after long and intensive discussions it was decided that LHCOPN shall be limited:

- Mission:
  - Tier-0 –Tier-1 transfers
    - transfer detector raw data from CERN to Tier-1s
  - Tier 1 – Tier 1
    - data exchange transit through CERN
    - a “best effort” base was agreed upon
  - Mission won’t be extended (no Tier-2/3 traffic through LHCOPN)

Operation WG

ENOC → EGEE NOC (Roberto Sabatino)

- Red Boxes (proposal by Géant → Géant administrated routing gear)
- Routing gear bought by Tier-1 sites
- managed by the ENOC

→ Tier-1s decided on a federated / collaborative management approach by Tier-0/1s sites

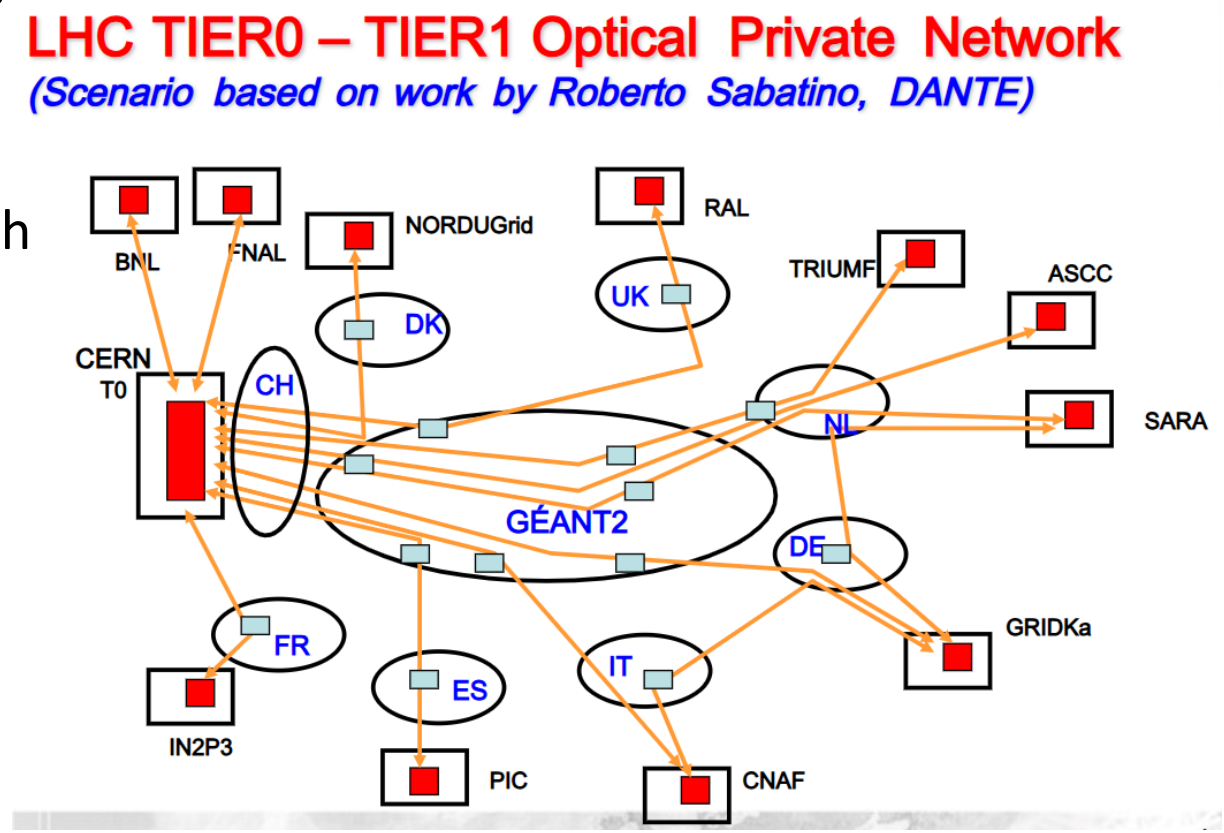

LHCOPN Operation

Manager: *Mathieu Goutelle* **2006** - **2008** - *Bruno Hoeft* **2012** - **2015** *Guillaume Cessieux*

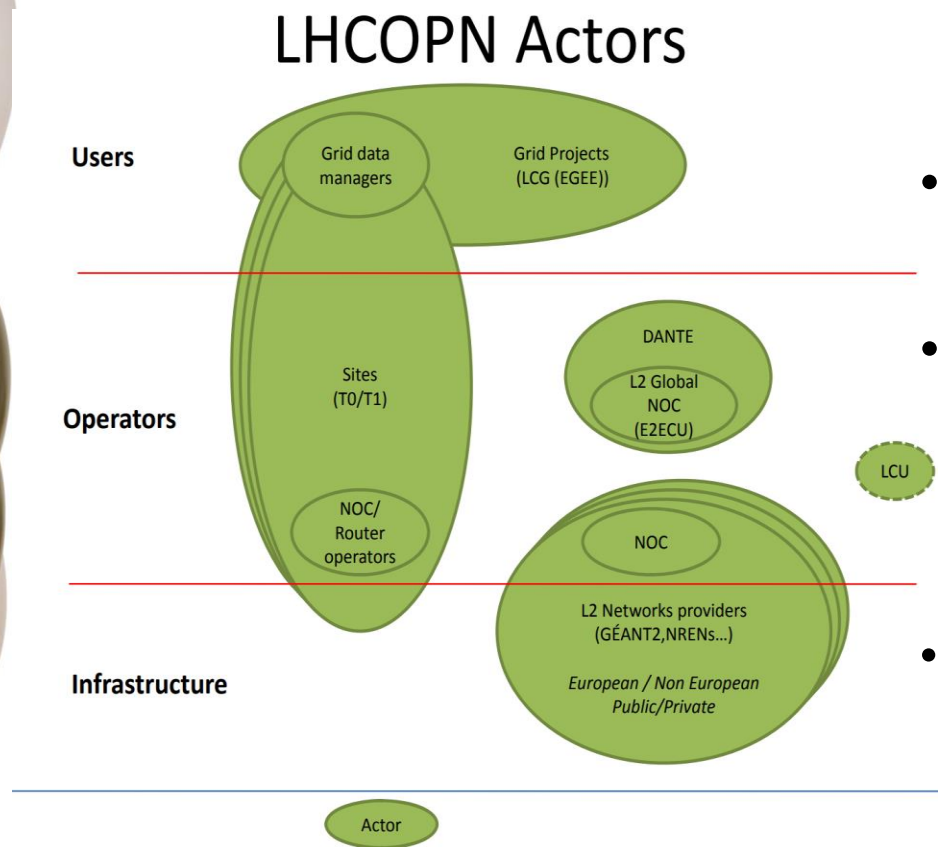

- 2005 Roberto Sabatino first sketch
  - Centralised model
  - LHC NOC resides in Géant → solves ASN and routing issues
- 2006 Mathieu Goutelle
  - LHC OPTICAL PRIVATE NETWORK Operational Handbook is written (part of EU funded EGEE-II-SA2 task)
- 2008 Guillaume Cessieux (EGEE-II,SA2)
  - Operational procedures of each site identifying weak spots (unsupported time slots)
  - Include backup links, incl. site connection between backup links
  - Operation of both layer-2 and layer-3 part of LHCOPN
- 2012 Bruno Hoeft
  - Tests of backup links
  - Mature ticketing system → GGUS

Evolvement – bandwidth + additional sites

T1 1Gb Uplink to CERN

- 10Gbps Uplinks ⇔ for the expected raw data traffic load from Tier-0 to Tier-1
- Upgrade process of Tier-1s to Tier-0 to **10Gbps** LHCOPN took till **2008**
- Crossborder connections accross different European NRENs between Tier-1s were deployed and considered as part of LHCOPN
  - CNAF Bologna ⇔ DE-KIT GARR/Switch/DFN (2005)
  - NLT1/Sara ⇔ DE-KIT Surf(Netherlight)/DFN (2007)
  - IN2P3 ⇔ DE-KIT Renater/DFN (10Gbps)(2006)
  - NLT1/Sara ⇔ NGDF (2008)
  - NLT1/Sara ⇔ TRIUMF (2008)
  - TRIUMF ⇔ BNL (2008)
- transit of Tier-1s through direct Tier-1 links:
  - e.g. SARA, CNAF, IN2P3 through DE-KIT

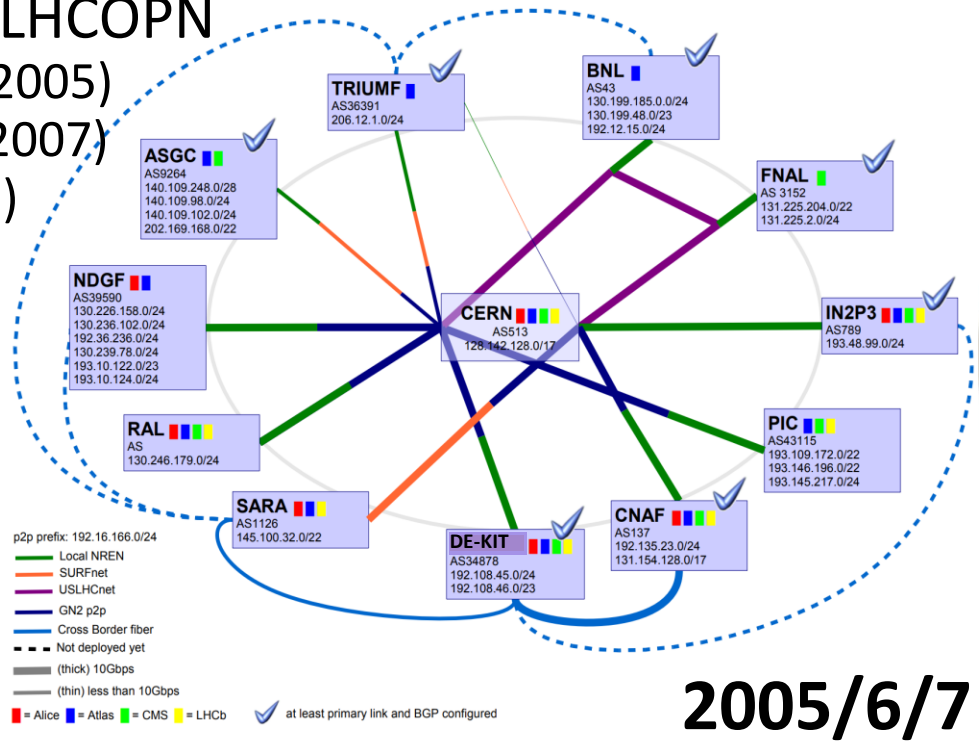

2005/6/7

LHCOPN status 2008/9/10/11/12

- Additional Tier-1 to Tier-1 links in production:
  - NL-T1 ⇔ TW-ASGC (1Gbps)
  - NL-T1 ⇔ CA-Triumf (1Gbps)
  - NL-T1 ⇔ US-FNAL-CMS (1Gbps)
  - TW-Triumf ⇔ US-FNAL-CMS (1Gbps)

  - DE-KIT to FR-CCIN2P3 (10Gbps)

- Upgrade 2012
  Tier-1 to Tier-1 link to 2Gpbs
  - NL-T1 ⇔ TW-ASGC (2Gbps)
  - NL-T1 ⇔ US-FNAL-CMS (2Gbps)
  - TW-Triumf ⇔ US-FNAL-CMS (2Gbps)

2009

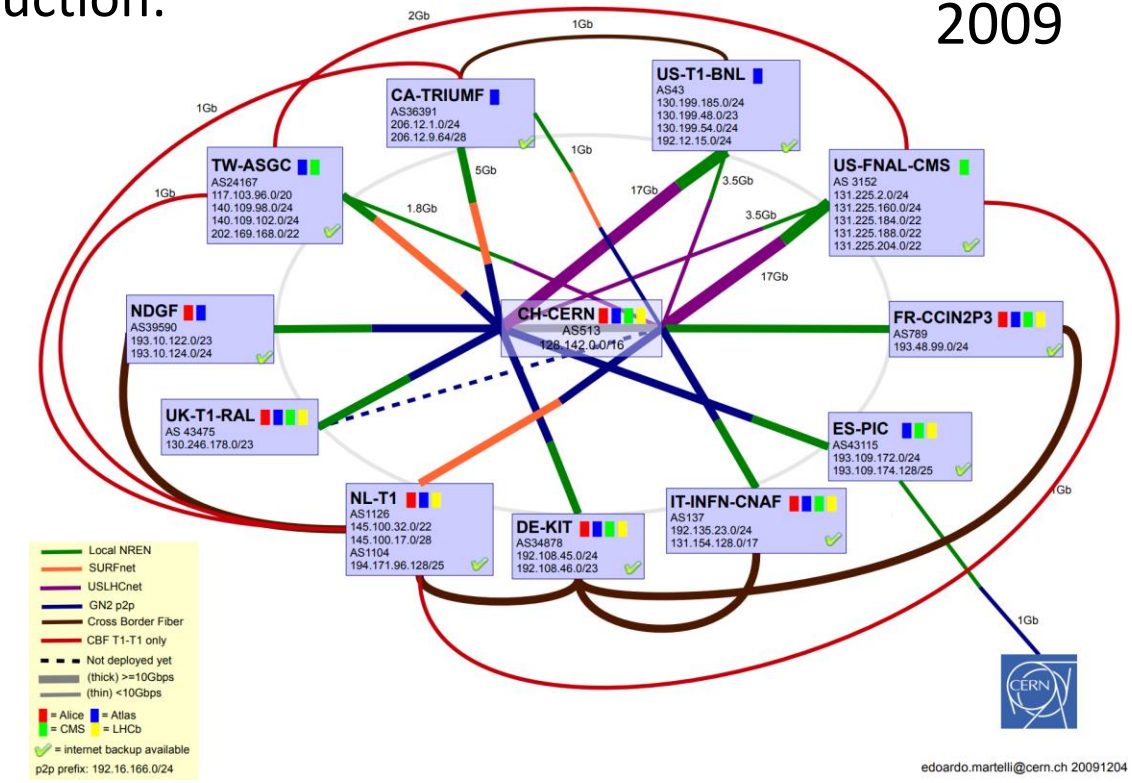

LHCOPN status 2013/4

- Tier-1 joined:
  - KR-KISTI
  - RRC-K1-T1

    +
  - RRC-JINR-T1

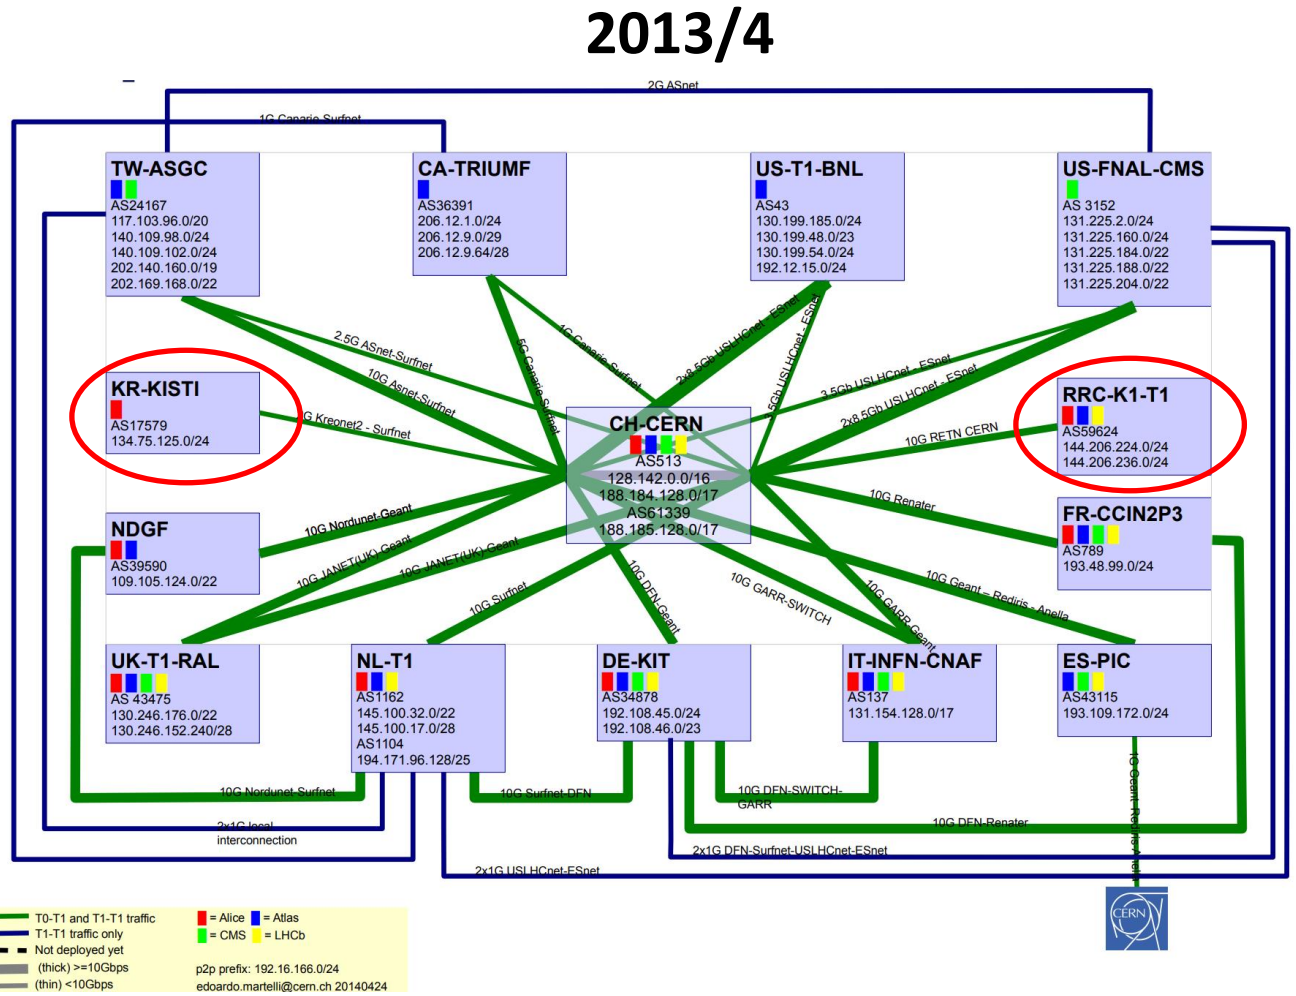

LHCOPN bandwidth upgrades

- ESnet connected with 100 Gbps (transatlantic) (2015)
  - FNAL and BNL connected each with 2*100 Gbps (2*40 Gbps backup) (Sept 2016)
- NDGF
  - 2* 20 Gbps (March 2018)
- UK-T1-RAL
  - 30 Gbps (Oct. 2017)
- DE-KIT
  - 100 Gbps (20 Gpbs backup) (Oct. 2018)
- IT-INFN-CNAF
  - 2* 100 from (June 2019)
- FR-CCIN2P3
  - 1*100 Gbps and 10 Gbps backup (Oct. 2018)
- NL-T1
  - 100 Gbps (Oct. 2019)
- RRC-KI-T1 and RRC-JINR-T1
  - 100 Gbps (June 2019)
- Triumf
  - 20 Gbps (Sept. 2020)

2020

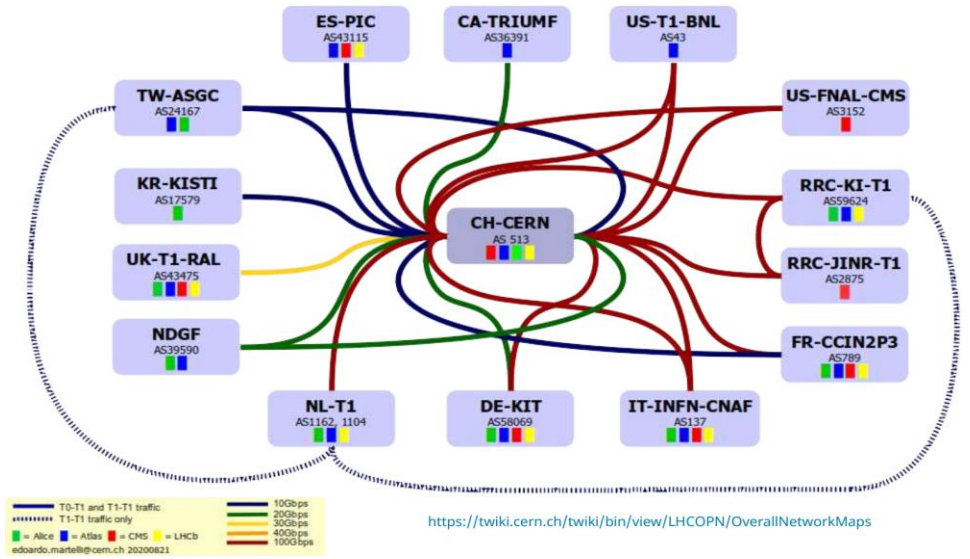

https://twiki.cern.ch/twiki/bin/view/LHCOPN/OverallNetworkMaps

LHCOPN bandwidth upgrades

- ESnet connected with 100 Gbps (transatlantic) (2015)
  - FNAL and BNL connected each with 2* 100 Gbps (June 2019)
- NDGF
  - 1* 100 Gbps and 20 Gbps backup Oct. 2022)
- UK-T1-RAL
  - 100 Gbps (March 2021)
- DE-KIT
  - 2* 100 Gbps (March 2021)
- IT-INFN-CNAF
  - 2* 100 from (June 2019)
- FR-CCIN2P3
  - 1*100 Gbps and 10 Gbps backup (Oct. 2018)
- NL-T1
  - 400 Gbps (Oct. 2022)
- RRC-KI-T1 and RRC-JINR-T1
  - After 2022 Russian Ukraine invasion decomissioned
- Triumf
  - 100 Gbps (Oct. 2022)

2022

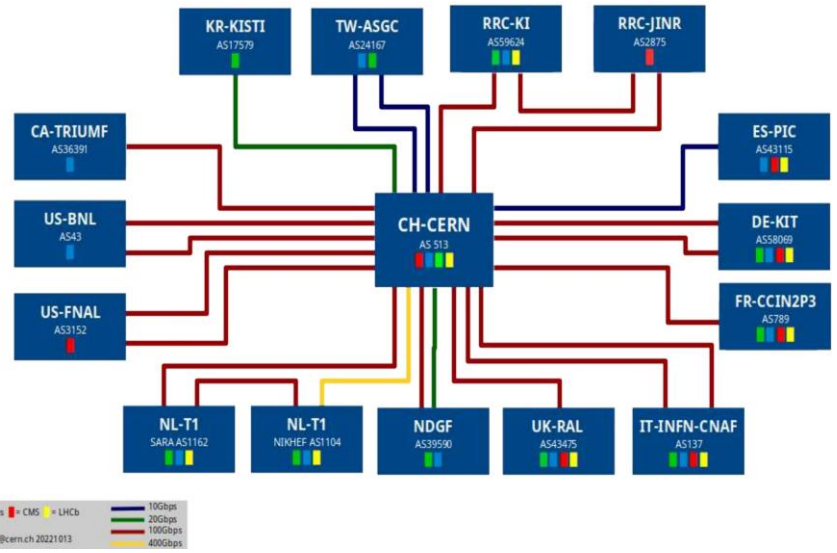

LHCOPN bandwidth upgrades

- ESnet connected with 100 Gbps (transatlantic) (2015)
  - FNAL 400 Gbps (April 2024) + BNL 400 Gbps (Oct. 2024)
- NDGF
  - 1* 100 Gbps and 20 Gbps backup Oct. 2022)
- UK-T1-RAL
  - 200 Gbps (Oct. 2023)
- DE-KIT connected with
  - 2* 100 Gbps (March 2021)
- IT-INFN-CNAF
  - 4x100 Gbps (April 2024) ⇔ DCI upgradable to 1.6Tb/s
- FR-CCIN2P3
  - 1*100 Gbps and 10 Gbps backup (Oct. 2018)
- NL-T1
  - 400 Gbps (Oct. 2022/2023/4)
- Triumf
  - 100 Gbps (Oct. 2022)
- PIC → 100 Gbps
  - 100 Gbps (April 2023)
- CN-IHEP 20 Gbps
  - 20 Gbps (April 2023)
- PL-NCBJ
  - 40 Gbps (April 2024)

2024

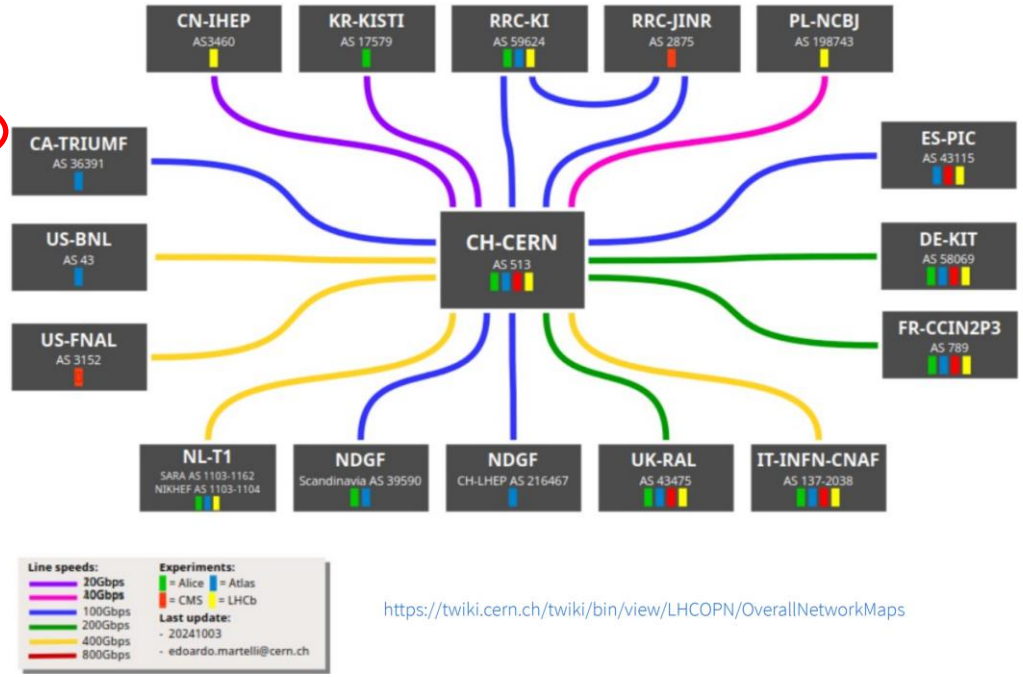

https://twiki.cern.ch/twiki/bin/view/LHCOPN/OverallNetworkMaps

LHCOPN traffic growth

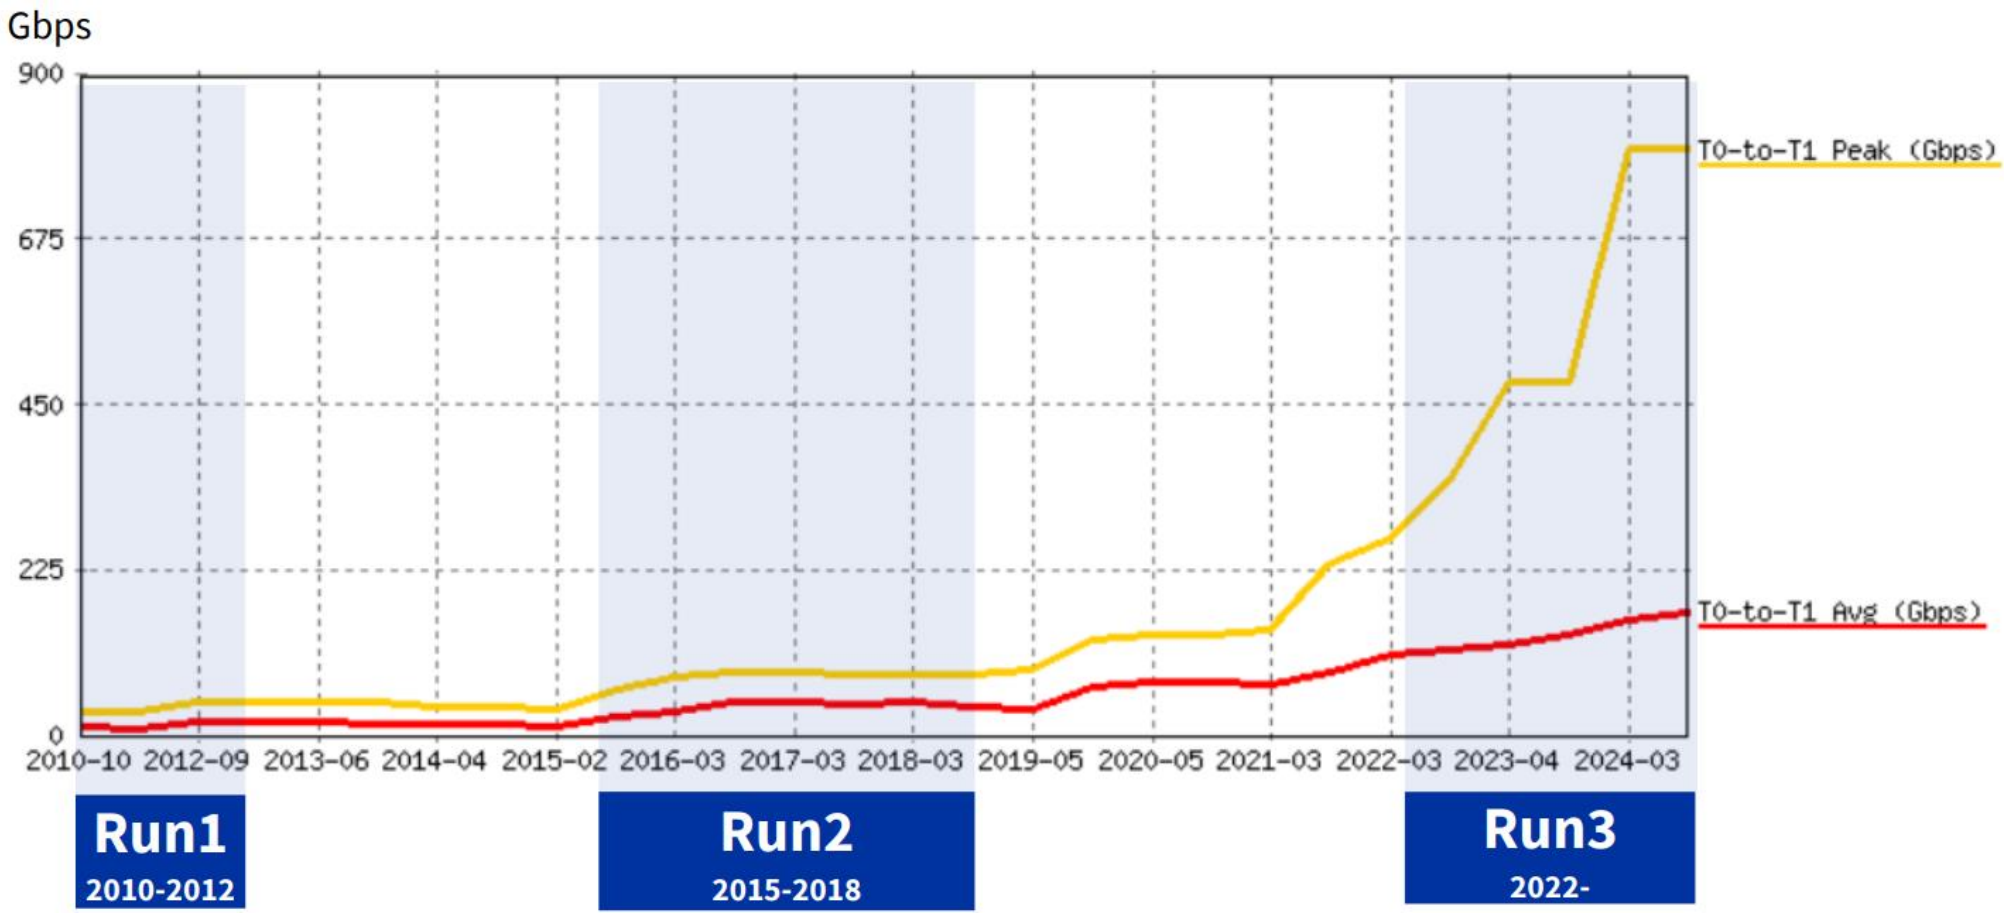

Backup status

- Because of limited resources at CERN detector data of the four LHC experiments should straight be transfered to the Tier-1 sites for the first computing run (reducing data and producing the so called DSTs)
  - For not loosing any data a detailed backup strategy was implemented
    - „normal operation“ LHCOPN
    - most of the Tier-1s were equipped with a second uplink which kept the connection to CERN
    - the third backup: Raw data transfered via general purpose internet (last resort)

- When LHC started in 2010 situation was quite different
  - The resources allowed that
    - all detector raw data were stored at CERN before sending to the Tier-1s
    - keeping pressure off the network

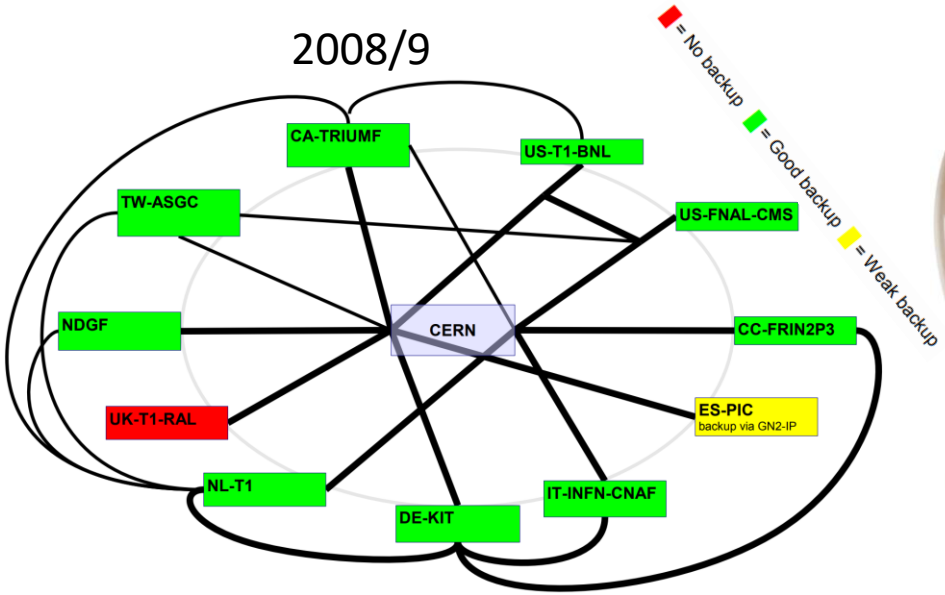

Service challenges

Service Challenge (SC) 1/2/3/4

To build up required disk/storage + tape, network and compute capacity

- SC1 ?
- SC2 March 23 – April 2 2005 → up to 600MB/s at CERN
  - Up to this point → for SC were special upgrades/deployments necessary for achieving the required load
  - From here SC shall be covered in "normal" Tier-1 operation

  From SC2 onwards SC has three phases preparation/execute/validate
- SC3 Q2 2005 → target data rate ~1 GB/sec at Tier-0
- SC4 Q1 2006 → target full data rate ~2 GB/sec at Tier-0
- SCs shall lead to required LHC (wlcg) performance and lead smoothly into standard service (Cosmics at Q4 2006)

Site	Average throughput (MB/s)
BNL	61
FNAL	61
GridKA	133
CCIN2P3	91
CNAF	81
RAL	72
SARA	106
TOTAL	600

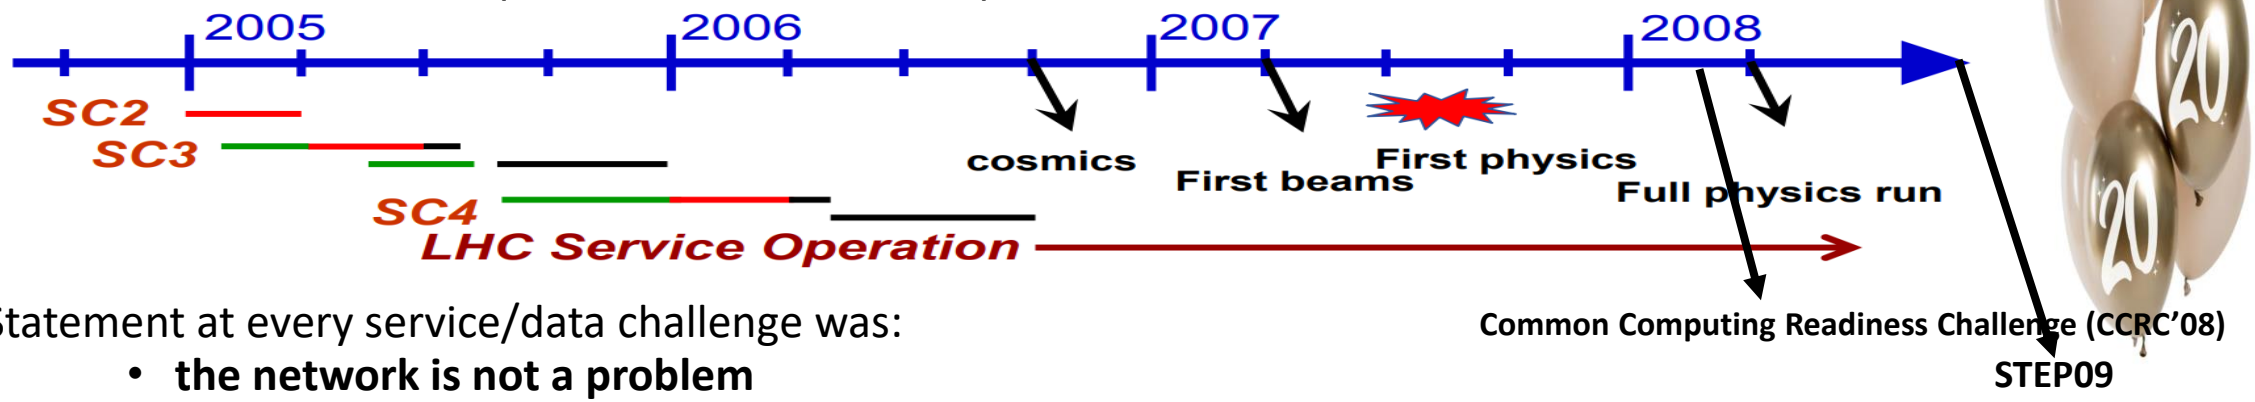

- Statement at every service/data challenge was:
  - **the network is not a problem**

IPv6

IPv6

- HEPiX-IPv6 Working Group started in 2011
  - First LHCOPN BGP peering between CERN and DE-KIT deployed
- LHCOPN IPv6 deployment T0/T1 started 2012/3
- IPv6 peering CERN and NDGF (Feb. 2013), PIC (June 2013)
- IPv6 peering FR-CCIN2P3 (Sept. 2014)
- IPv6 peering IT-INFN-CNAF and US-FNAL-CMS (Oct. 2015)
- IPv6 peering US-T1-BNL and CA-Triumf (Sept. 2016)
- IPv6 peering RAL (April 2017)
- Since 2018 - all Tier-1s peer with IPv6 prefixes

Today -Average: ~90% IPv6

Percentage of IPv6 traffic over the total - old accurate data

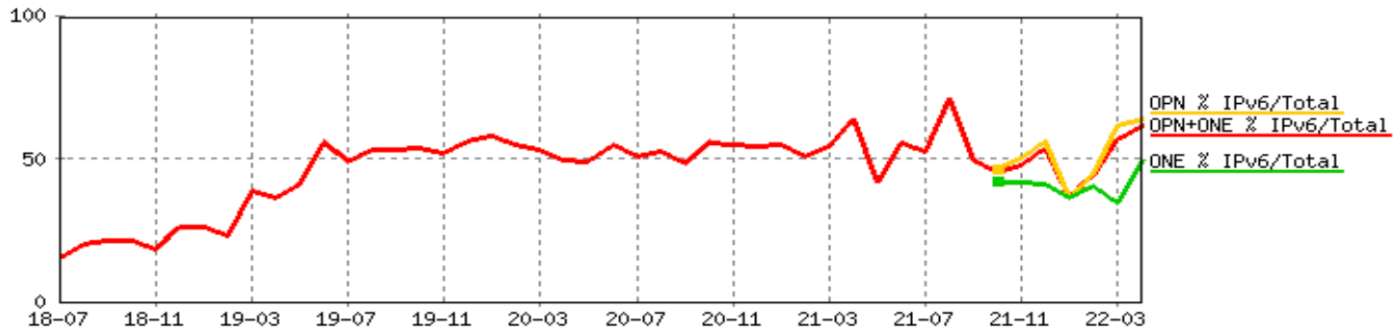

LHCOPN meetings during 20 years

2005 - 2006 - 2007 - 2008 - 2009 - 2010 - 2011 - 2012 - 2013 - 2014 - 2015 - 2016 - 2017 - 2018 - 2019 - 2020 - 2021 - 2022 - 2023 - 2024 - 2025

14. Nov. T0/1 Network Meeting – Seattle (SC05)
19. Juli T0/1 Network Meeting – CERN
07. Juli ESnet Network Meeting – CERN
08. Apr. T0/1 Network Meeting – Amsterdam
20.-21. Jan. T0/1 Network Meeting – CERN

21. Sept. T0/1 Network Meeting – Utrecht
16. Juni LHCOPN Meeting – Cambridge
04. Apr. LHCOPN Meeting – Rome
31. Jan. LHCOPN Meeting – Amsterdam

07. Dec. LHCOPN Operation Workshop – CERN
05.-06. Nov. LHCOPN Operation Workshop – CERN
16. Juni LHCOPN Meeting – Amsterdam
04. Apr. LHCOPN Meeting – Munich
12. Jan. LHCOPN Meeting – Cambridge

16.-17. Oct. LHCOPN Meeting – Copenhagen
16.-17. Jun. LHCOPN Meeting – CERN
10.-11. Mar. LHCOPN Meeting – Madrid

10.-11. Dec. LHCOPN Meeting – Bologna
31. Aug.-01. Sept. LHCOPN Meeting – Vancouver
21.-22. Apr. LHCOPN Meeting – Utrecht
15.-16. Jan. LHCOPN Meeting – Berlin

07.-08. Oct. LHCOPN Meeting – CERN
21.-22. Apr. LHCOPN Meeting – Barcelona
08.-09. Mar. LHCOPN Meeting – London

26.-27. Sept. LHCOPN Meeting – Amsterdam
13.-14. Jun. LHCOPN Meeting – Washington
10.-11. Feb. LHCOPN Meeting – Lyon

20.-21. Sept. LHCOPN Meeting – Oslo
03.-04. May LHCOPN Meeting – Stockholm
30.-31. Jan. LHCOPN Meeting – LBL Berkeley

03.-04. Dec. LHCOPN Meeting – Pasadena
17.-18. Jun. LHCOPN Meeting – Paris
31.Jan.-01.Feb. LHCOPN Meeting – Washington (DC)

15.-16. Sept. LHCOPN Meeting – Michigan
28.-29. Apr. LHCOPN Meeting – Rome

28.-29. Oct. LHCOPN Meeting – Amsterdam
01.-02. Jun. LHCOPN Meeting – LBL Berkeley
09.-10. Feb. LHCOPN Meeting – Cambridge

13.-14. Mar. LHCOPN Meeting – Taipei
19.-20. Sept. LHCOPN Meeting – Helsinki

16.-17. Oct. LHCOPN Meeting – KEK
04.-05. Apr. LHCOPN Meeting – BNL

16.-17. Oct. LHCOPN Meeting – FNAL
06.-07. Mar. LHCOPN Meeting – RAL

04.-05. Mar. LHCOPN Meeting – Umeå

16.-17. Sept. LHCOPN Meeting – VC
13. May LHCOPN Meeting – VC
13.-14. Jan. LHCOPN Meeting – CERN

11.-12. Oct. LHCOPN Meeting – VC
23.-24. Mar. LHCOPN Meeting – VC

24.-25. Oct. LHCOPN Meeting – CERN
29.-30. Mar. LHCOPN Meeting – VC

18.-19. Oct. LHCOPN Meeting – Victoria
18.-19. Apr. LHCOPN Meeting – Prag

09.-12. Oct. LHCOPN Meeting – Beijing
09.-10. Apr. LHCOPN Meeting – Catania

18.-20. Mar. LHCOPN Meeting – Manchester

2005 to 2011
24 LHCOPN Meetings

2011 to 2025
32 combined LHCOPN and LHCONE Meetings

Responses of former colleagues

Responses by some early contributor

- Roberto Sabatino (Géant/HEAnet)
  - 2004 was a very interesting time, with fibre optic networks being deployed in many cases by NRENs and GEANT for the first time, and the LHCOPN being the first large scale use case for services delivered over fibre networks. It posed many challenges, from having to work with new technology to delivering end to end services in a global, multidomain environment. It wasn't easy, but it was certainly exciting and through the collective persistence and dedication of all the stakeholders we got there. On time. To spec. We also worked together to deliver a novel concept: cross border fibre connections, where neighbouring NRENs (GARR, SWITH, DFN, RENATER) partnered in delivering connectivity to the LHCOPN over their own fibre networks, thus providing additional resilience and diversity from the pan-European network, GEANT.
  - LHCOPN has been, and still is, instrumental in pushing the boundaries of what can be done with new technology , novel approaches and remains a perfect example of global collaboration in support of science.
- Enzo Valente
  - Chep 2000 in Padua → agreed with Les Robertson on distributed WLCG
  - LHCOPN did a greate Job

Responses by some early contributor - 2

- Kors Bos
  - I got to see the LHCOPN doing its good work
  - saw the Higgs in my last weeks at CERN → a good time to retire
- Guillaume Cessieux
  - Sends his greetings
- David Foster
  - “father” of LHCOPN
  - the lhcopn was very political as it involved Géant and the NRENs directly, but it gave a groundbreaking goal to bring the community together in a common objective.
  - It was not an easy task to navigate the technical and political issues, of which there were many. But I think that 20 years later we can declare it a major success. It is one of the things I am most proud of doing in my time at CERN. I created something that was fundamental to the success of the LHC.
  - Setting the ambitious goal of a 10 Gb/s CERN centric star network

Success story of LHCOPN

- Operation still stable
  - Thanks to all collaborators
    - NRENs
    - Sites
    - Different working groups (operation/monitoring/…)
    - …
  - LHCOPN is still dedicated light path (or assured dedicated bandwidth)
    - approaches to run LHCOPN as overlay did not succeed
  - Bandwidth upgrade from 1 Gb/s (2004/5) to 400 Gb/s today
  - Extreme stable LHC production support
  - Aggregated LHCOPN throughput:

2005 – SC2	3.2 Gb/s
2024 – DC24	in: 656 Gb/s out: 945 Gb/s

  - All Service and Data challenges stated that:

problems where never network caused

Celebrating 20 LHCOPN

Thank you for your attantion