Evolution of networks at CERN and in Europe

Tbilisi, 21st October 2014

edoardo.martelli@cern.ch

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Summary

- CERN networks
- CERN data centres
- LHCOPN
- LHCONE
- What's upcoming

CERN Networks

CERN Topology

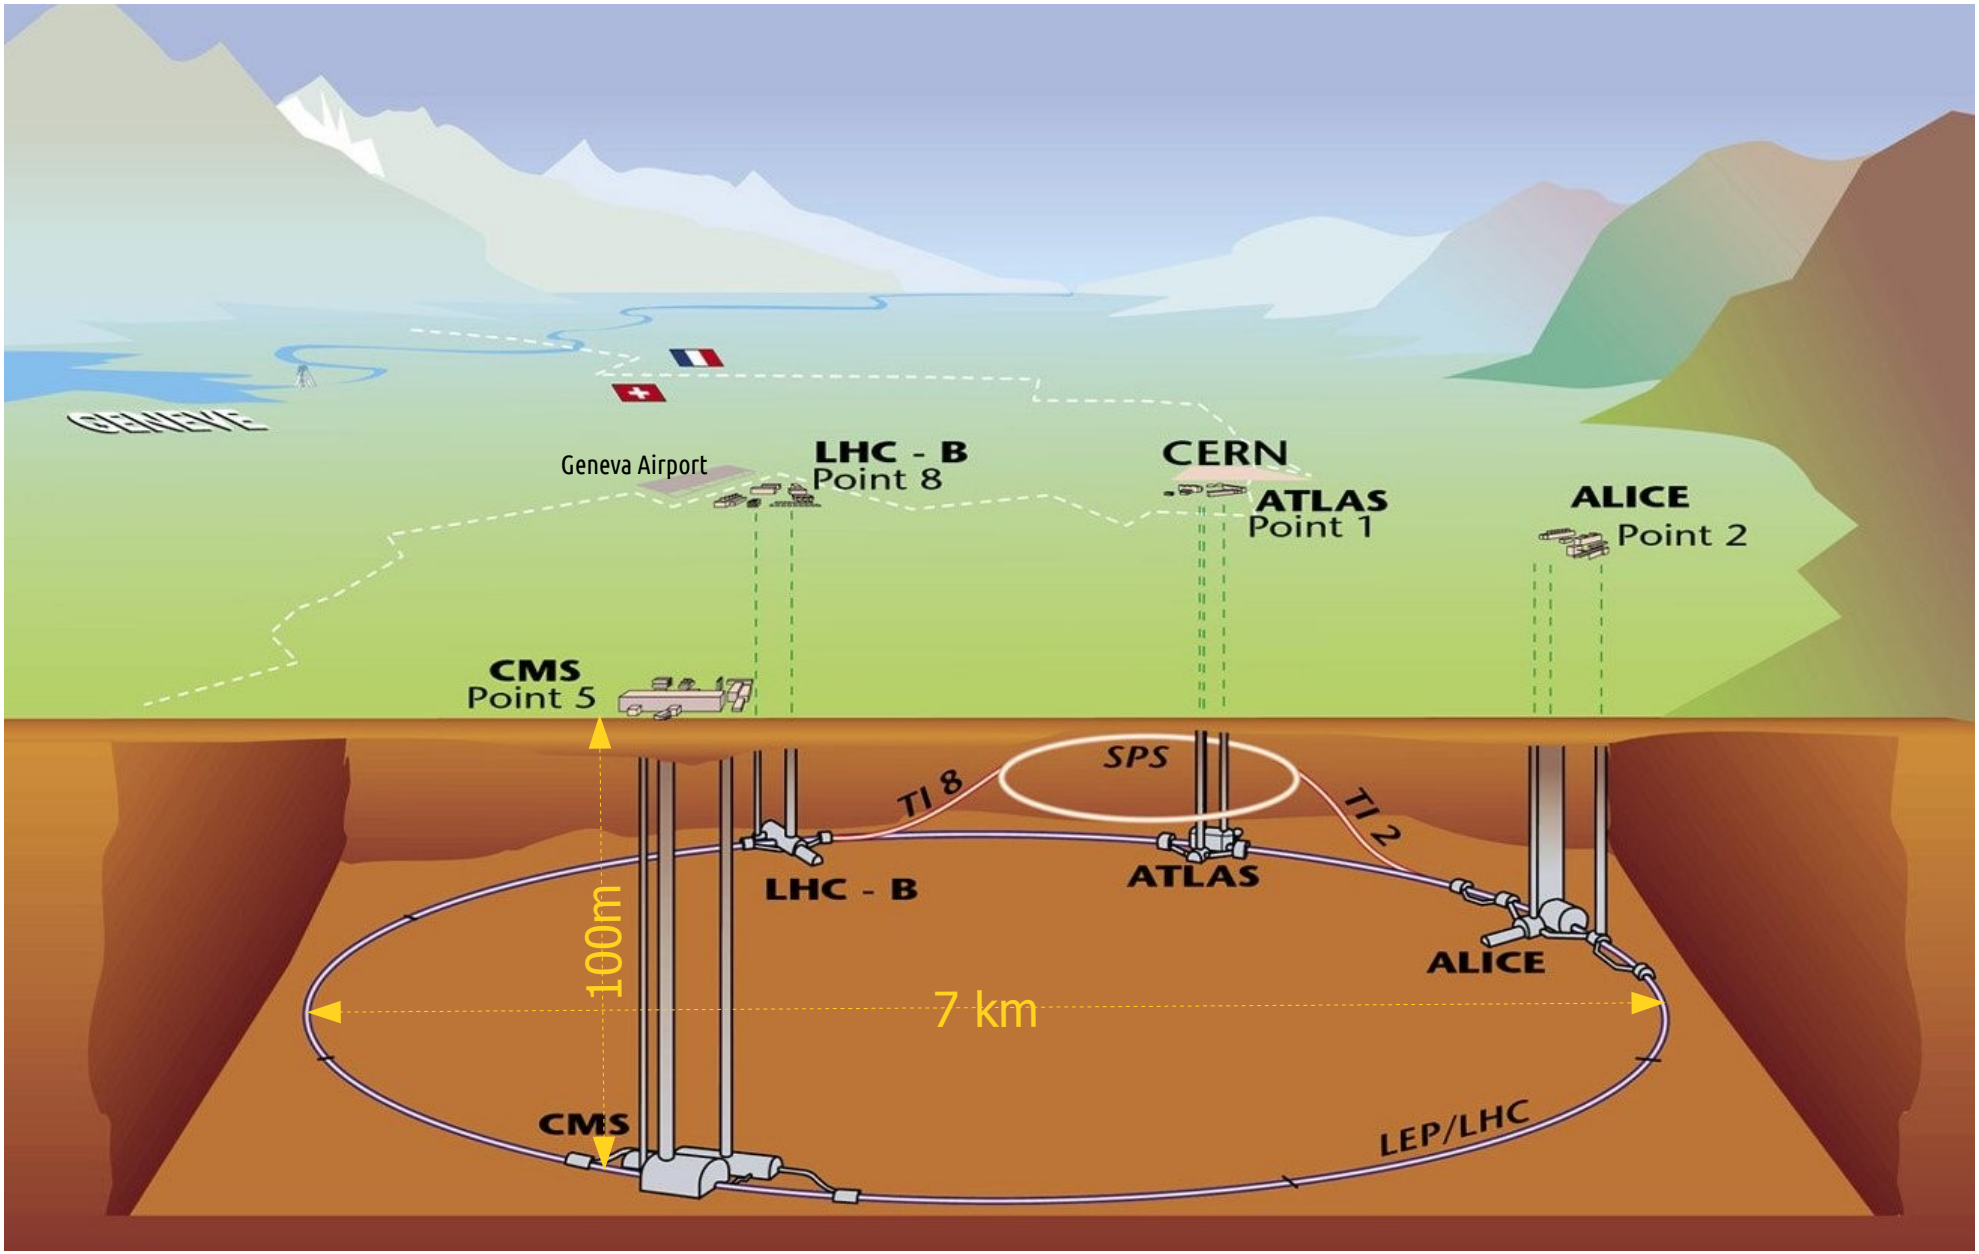

Network connections for LHC DATA

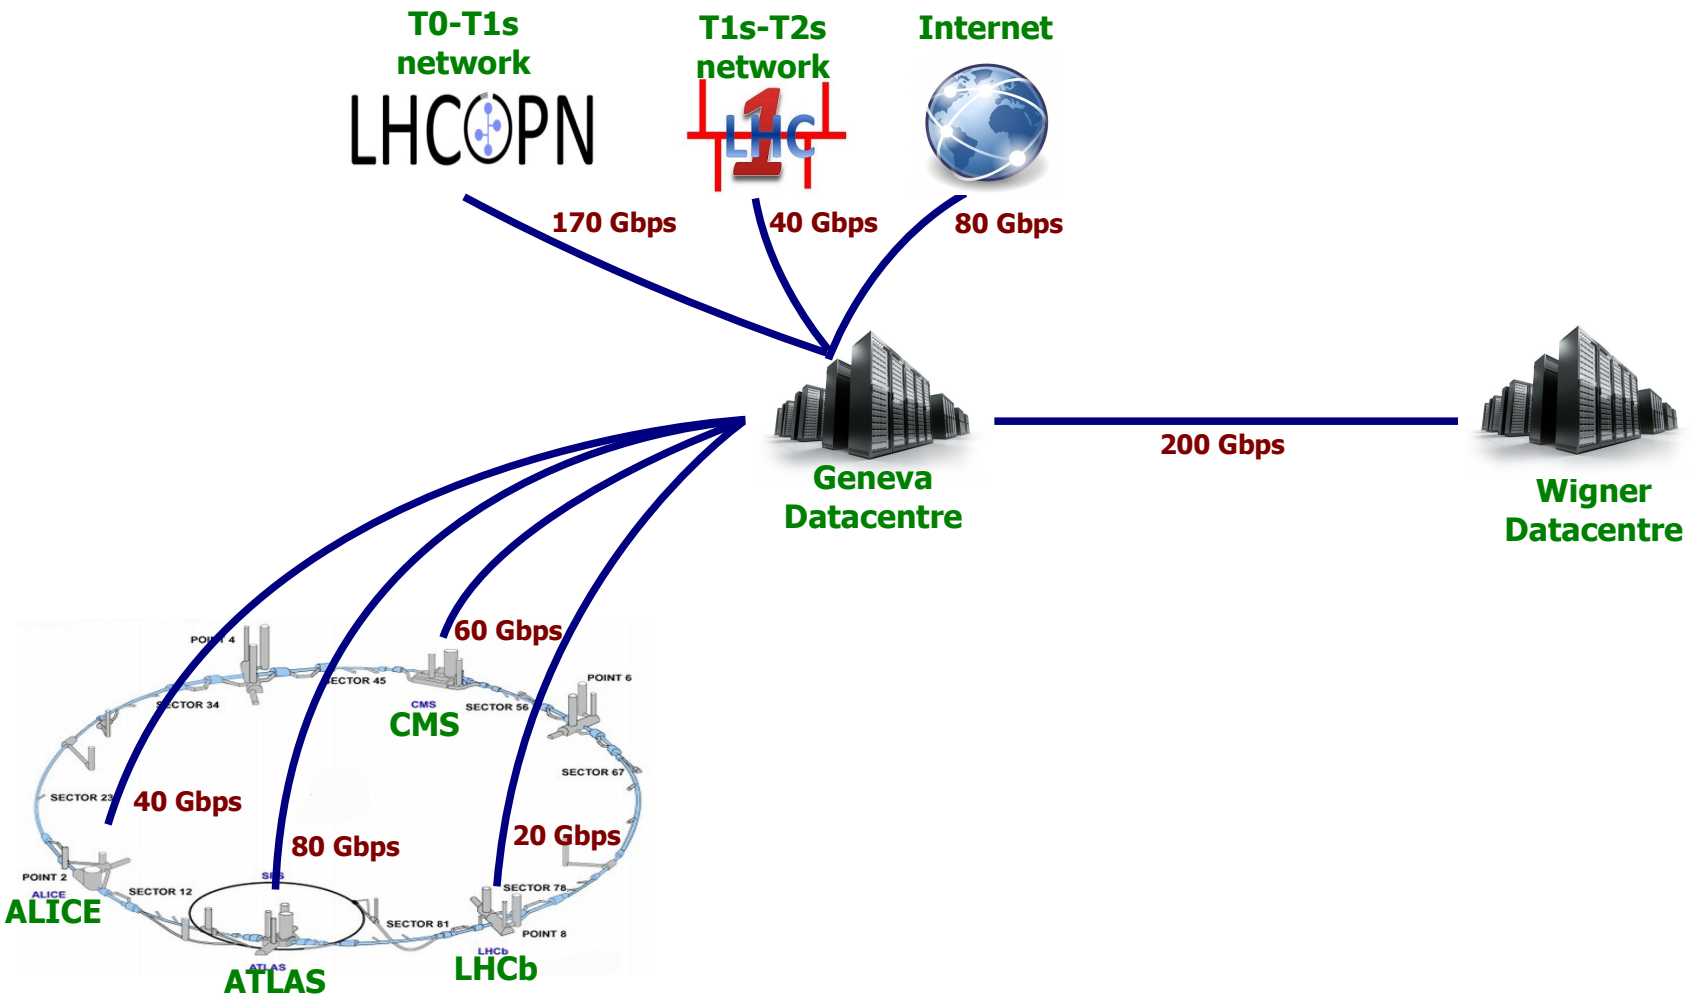

Networks in details

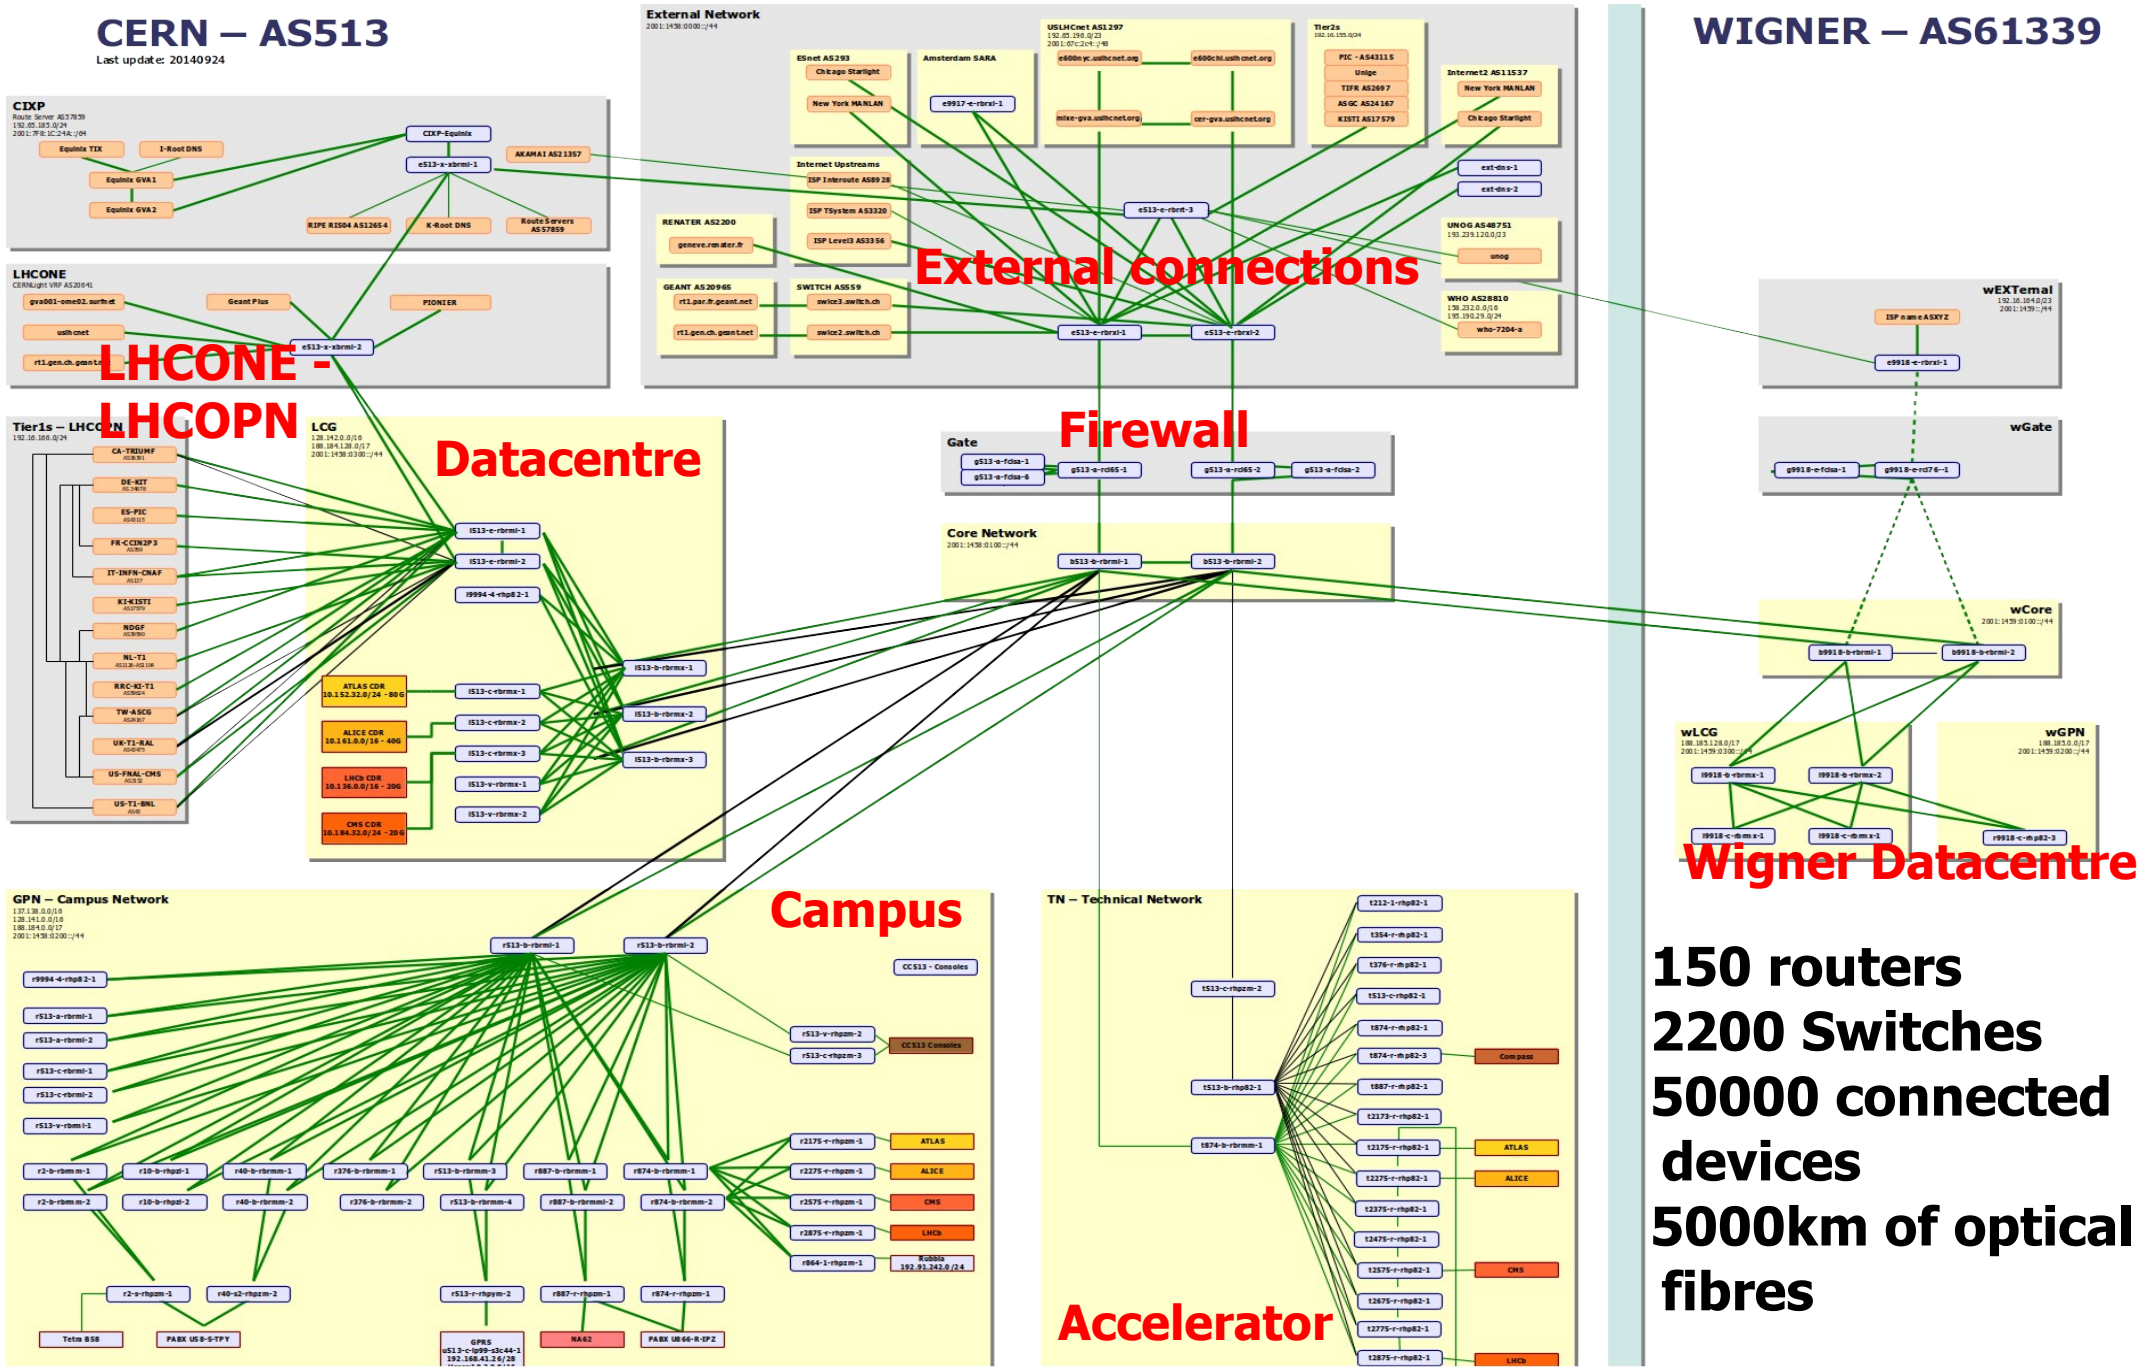

150 routers
2200 Switches
50000 connected devices
5000km of optical fibres

Main Firewall

Firewall rules stored in network database

Routers and Firewall configuration updated very 15 minutes

Statefull firewall bypass for large, well known flows

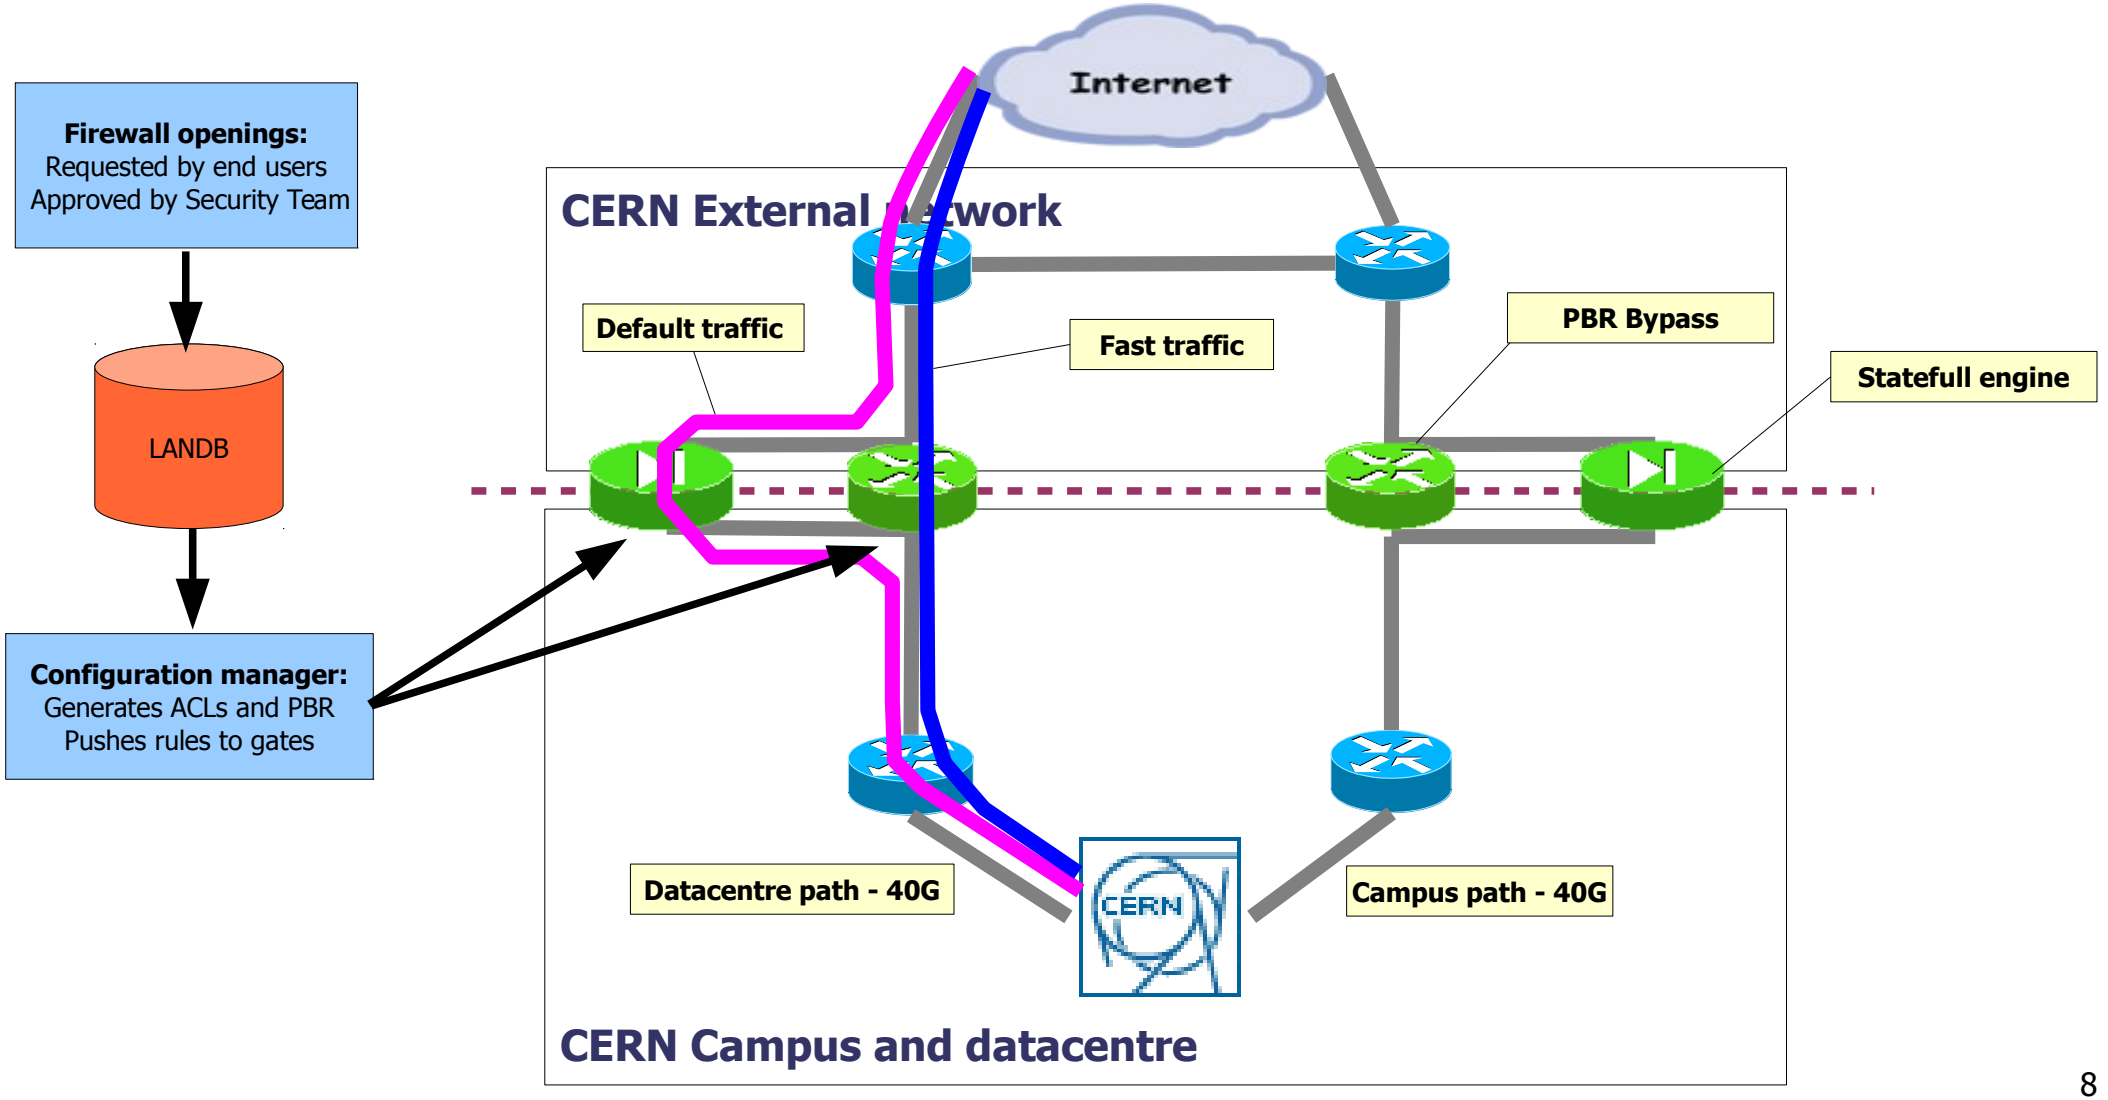

IPv6

IPv6 dual stack network deployment completed in March 2014!

Main driver: IPv4 address exhaustion due to VMs utilization

More information:
http://cern.ch/ipv6

Network Provisioning and Management System

- >250 Database tables
- ~200,000 Registered devices
- >1,000,000 lines of codes
- >15 years of development

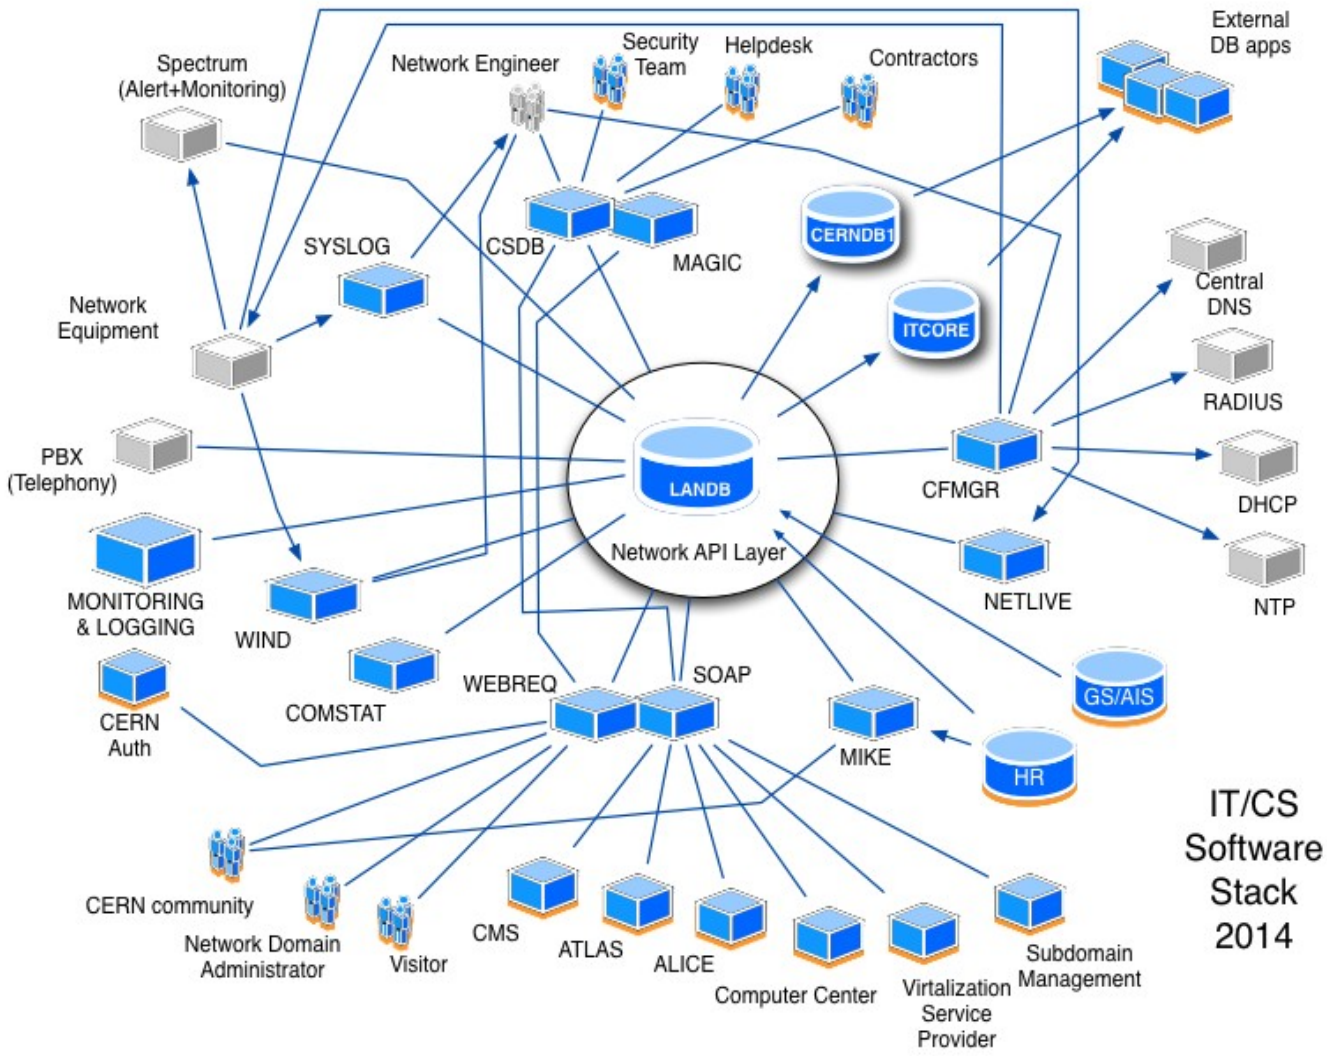

Monitoring and Operations

The whole network is monitored and operated by the CERN NOC:

- 7500 devices

- 150 alarms treated per day

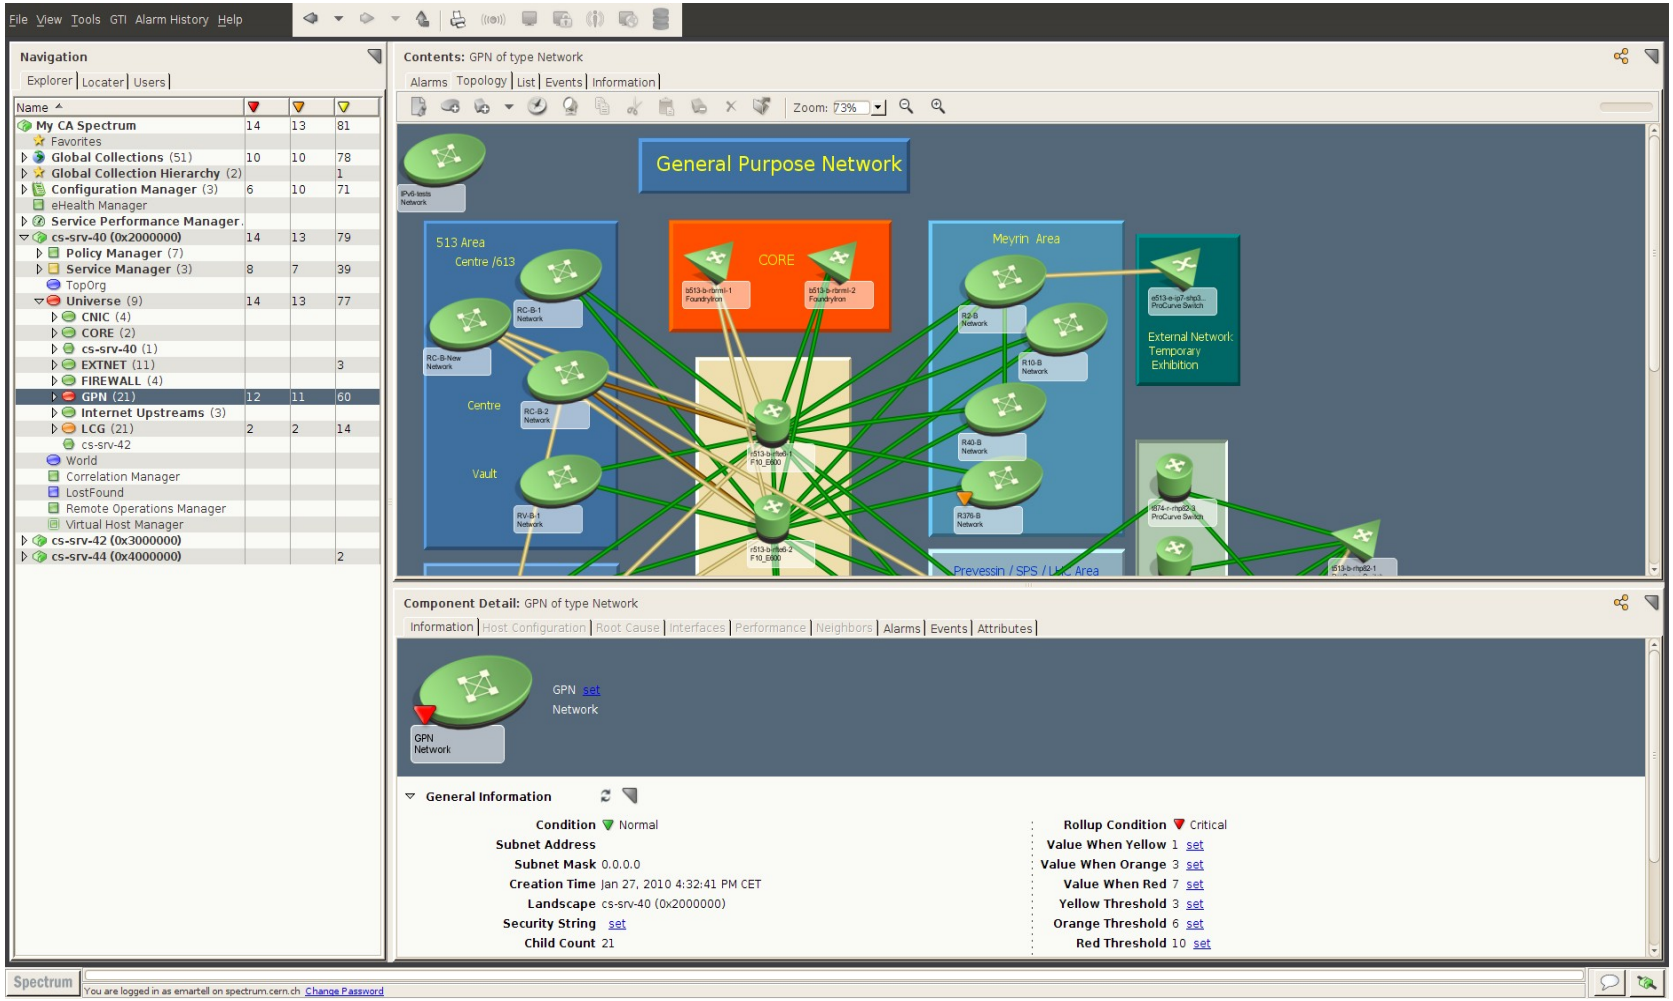

CERN Data Centres

CERN Data Centres

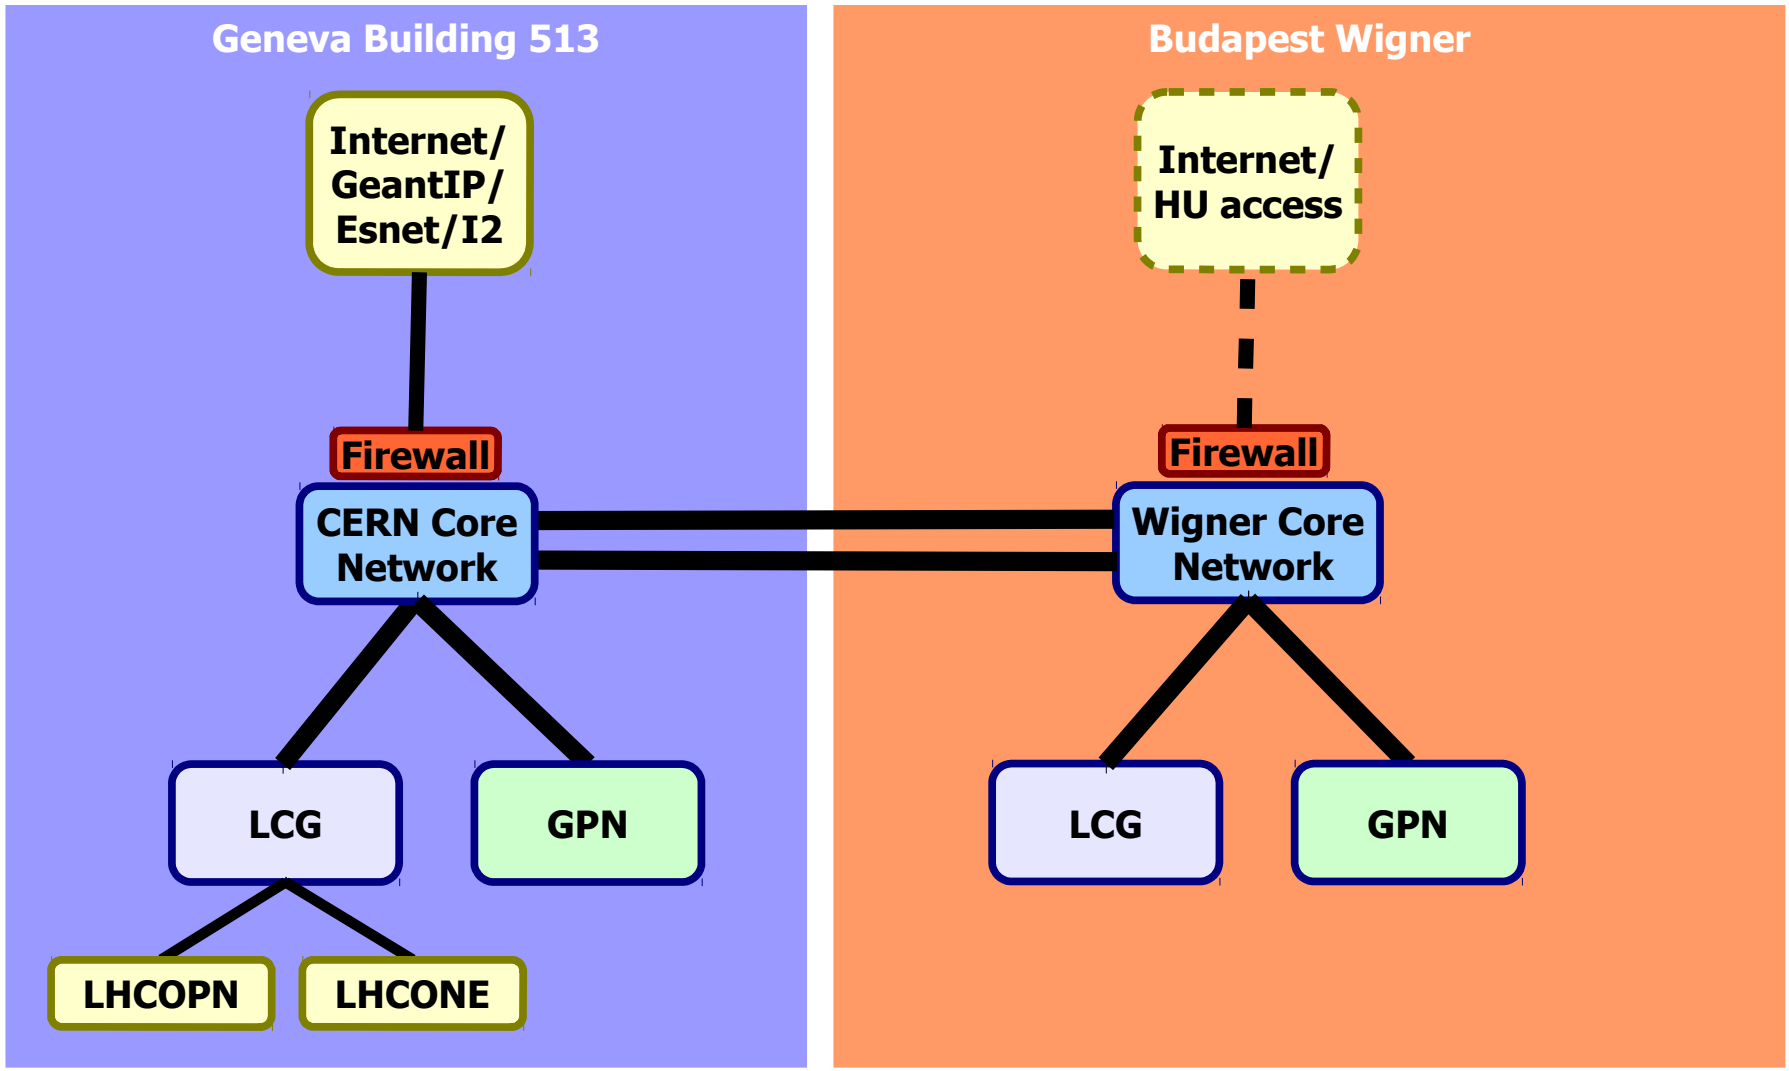

2x100Gbps circuits CERN-Wigner

Two independent paths across Europe

Data Centre architecture

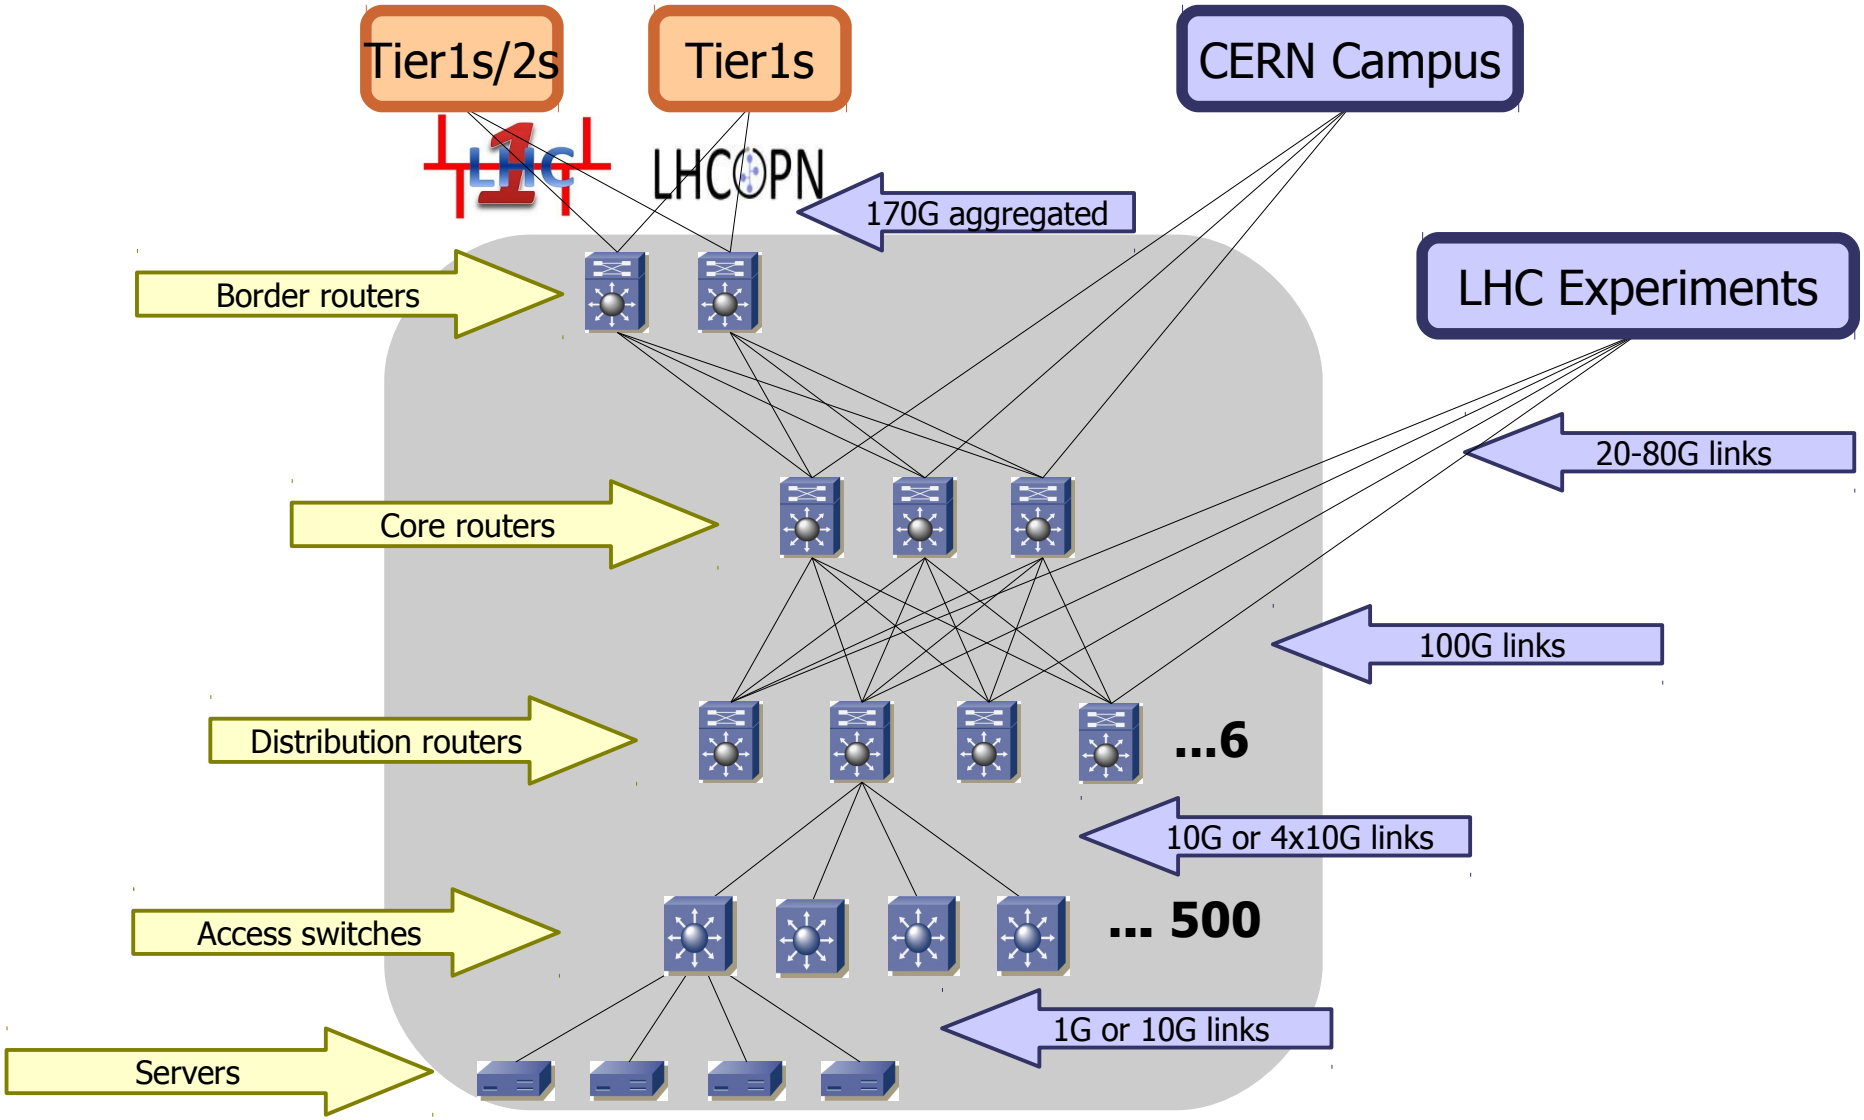

CERN Tier0 resources

MEYRIN DATA CENTRE	last_value
Number of Cores in Meyrin	93,937
Number of Drives in Meyrin	65,716
Number of Memory Modules in Meyrin	66,167
Number of 10G NIC in Meyrin	3,708
Number of 1G NIC in Meyrin	18,776
Number of Processors in Meyrin	18,018
Number of Servers in Meyrin	9,808
Total Disk Space in Meyrin (TB)	99,329
Total Memory Capacity in Meyrin (TB)	344

WIGNER DATA CENTRE	last_value
Number of Cores in Wigner	20,544
Number of Drives in Wigner	10,921
Number of Memory Modules in Wigner	10,247
Number of 10G NIC in Wigner	1,211
Numer of 1G NIC in Wigner	2,292
Number of Processors in Wigner	2,570
Number of Servers in Wigner	1,288
Total Disk Space in Wigner (TB)	32,584
Total Memory Capacity in Wigner (TB)	83

NETWORK AND STORAGE	last_value
Tape Drives	141
Tape Cartridges	50,622
Data Volume on Tape (TB)	98,604
Free Space on Tape (TB)	13,549
Routers (GPN)	135
Routers (TN)	24
Routers (Others)	101
Star Points	631
Switches	3,229

October 2014

WLCG Networks

WLCG resources

WLCG sites:

- 1 Tier0 (CERN)
- 13 Tier1s
- ~140 Tier2s
- >300 Tier3s worldwide
- ~250,000 CPUs
- ~ 150PB of disk space

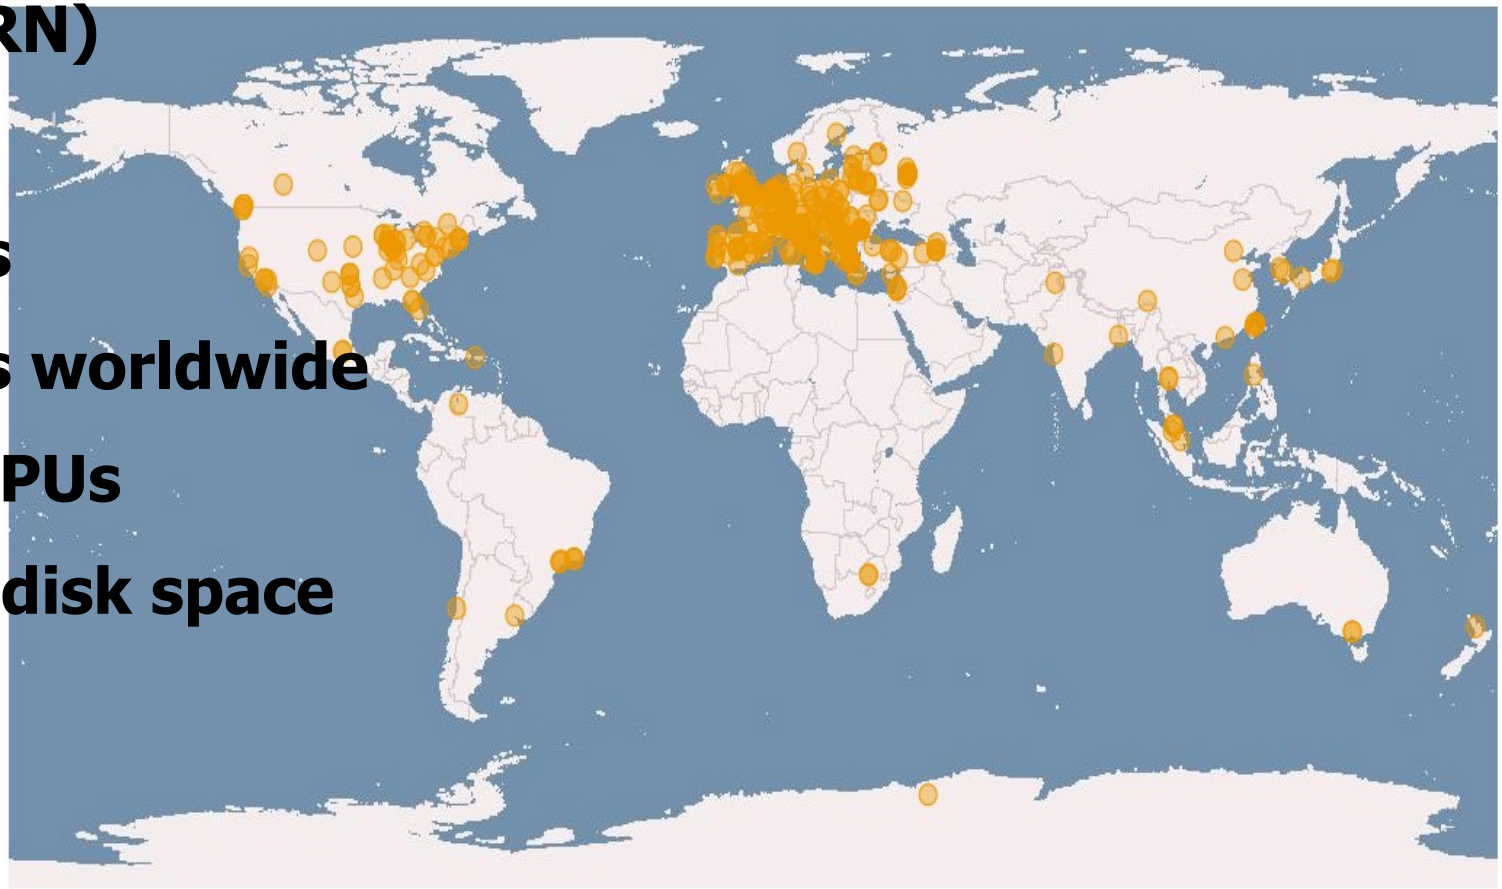

LHCOPN

LHC Optical Private Network

LHCOPN:

Private network connecting **Tier0 and Tier1s**

Reserved to LHC data transfers and analysis

Dedicated large bandwidth links

Resilient architecture

Topology

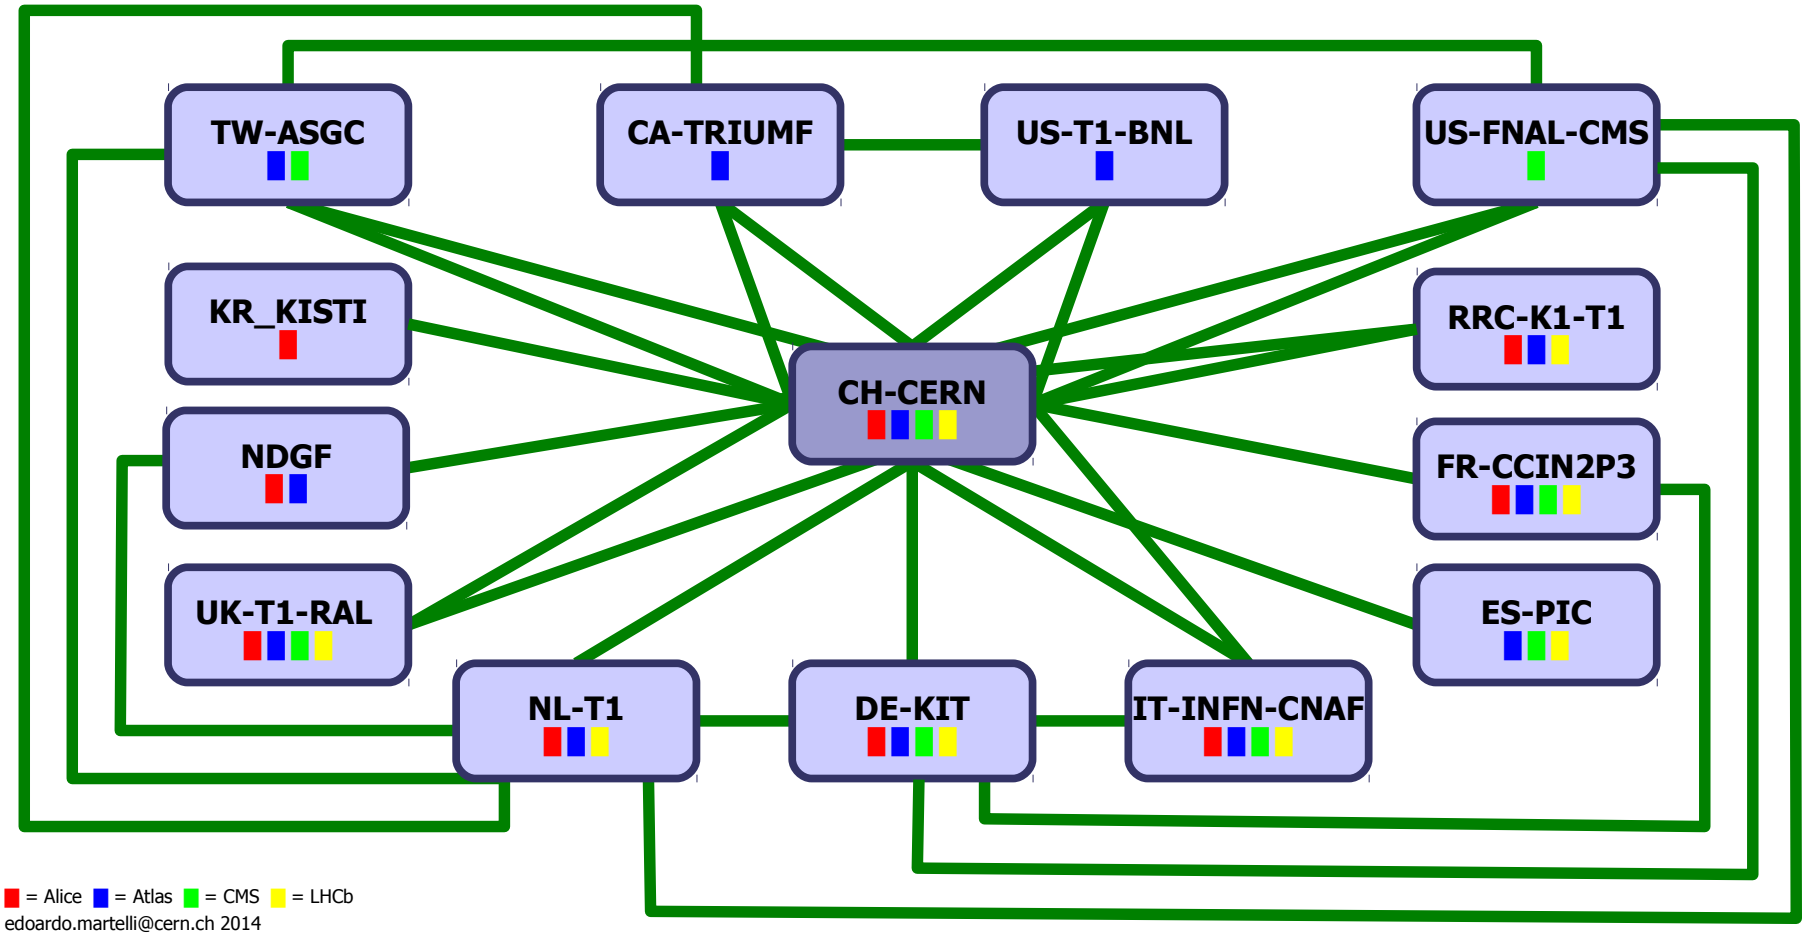
TW-ASGC
CA-TRIUMF
US-T1-BNL
US-FNAL-CMS
KR_KISTI
RRC-K1-T1
CH-CERN
NDGF
FR-CCIN2P3
UK-T1-RAL
ES-PIC
NL-T1
DE-KIT
IT-INFN-CNAF
= Alice = Atlas = CMS = LHCb
edoardo.martelli@cern.ch 2014

Technology

- Single and bundled long distance 10G Ethernet links

- Multiple redundant paths. Star and Partial-Mesh topology

- BGP routing: communities for traffic engineering, load balancing

- Security: only declared IP prefixes can exchange traffic

LHCOPN future

- The LHCOPN will be kept as the main network to exchange data among the Tier0 and Tier1s

- Links to the Tier0 may be soon upgraded to multiple 10Gbps or 100Gbps (waiting for Run2 to see the real needs)

- T1-T1 traffic and T0-T1 backups moving to LHCONE

LHCONE

LHC Open Network Environment

Computing model evolution

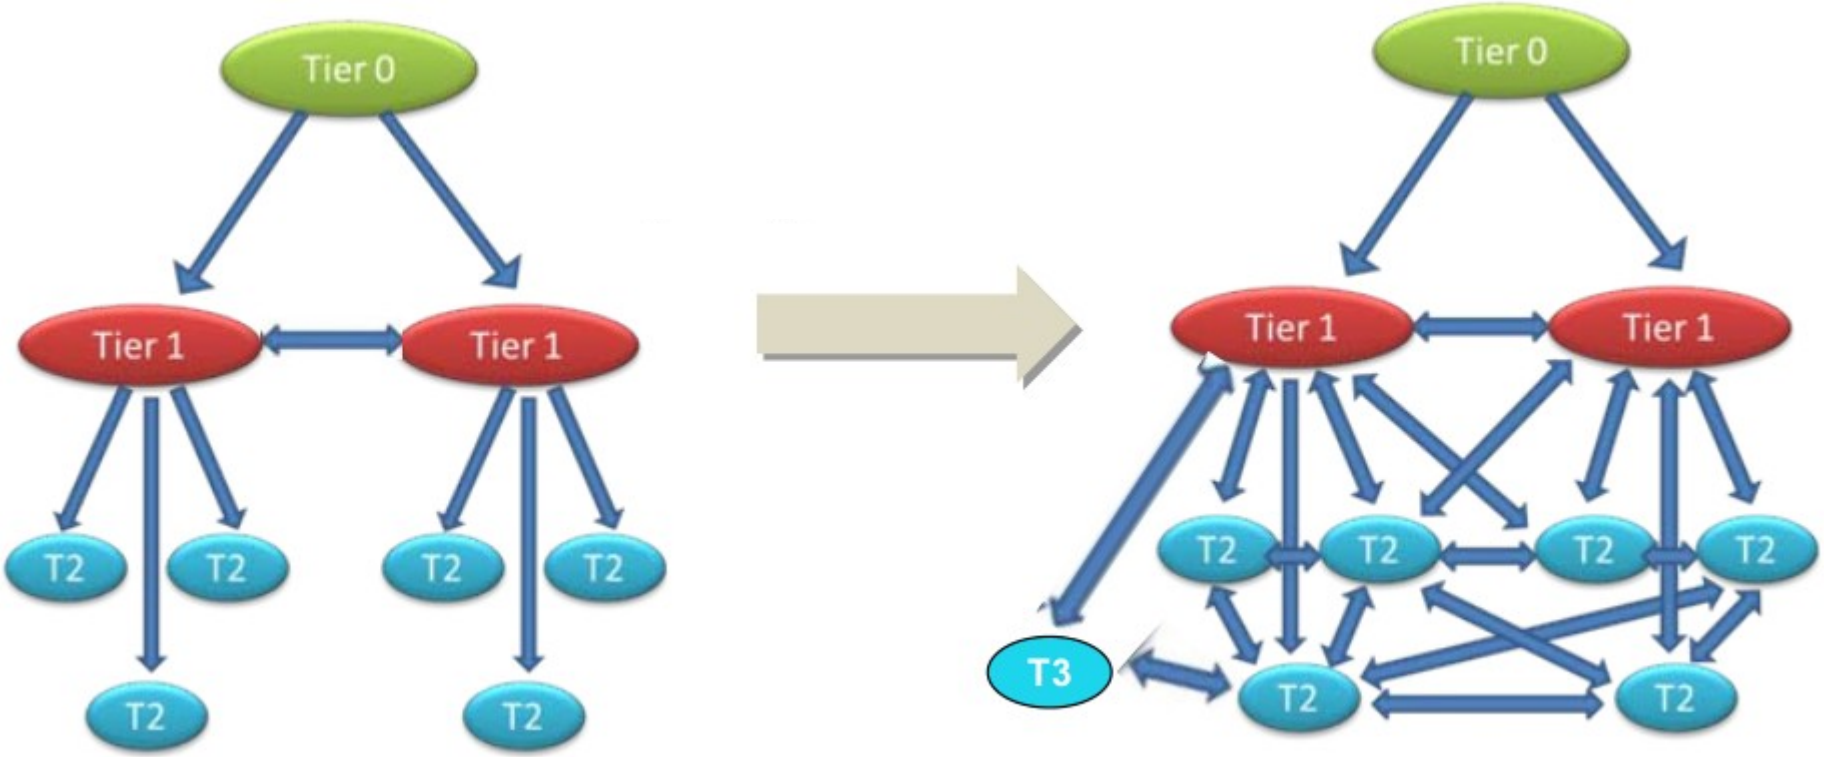

Original MONARCH model

Model evolution

New computing model

- Better and more dynamic use of storage
- Reduced load on the Tier1s for data serving
- Increased speed to populate analysis facilities

Needs for a faster, predictable, pervasive network connecting Tier1s and Tier2s

Requirements from the Experiments

- Connecting any pair of sites, regardless of the continent they reside
- Site's bandwidth ranging from 1Gbps (Minimal), 10Gbps (Nominal), to 100G (Leadership)
- Scalability: sites are expected to grow
- Flexibility: sites may join and leave at any time
- Predictable cost: well defined cost, and not too high

LHCONE:

Conceived to provide better network connectivity to all WLCG sites

Worldwide effort driven by WLCG and the Research and Education Network Providers of all continents

LHCONE concepts

- Serving any LHC sites according to their needs and allowing them to grow

- Sharing the cost and use of expensive resources (like transoceanic links)

- Traffic separation: no clash with other data transfer, resource allocated for and funded by the HEP community

LHCONE services

L3VPN (VRF): routed Virtual Private Network - *operational*

P2P: dedicated, bandwidth guaranteed, point-to-point links - *development*

perfSONAR: monitoring infrastructure - *operational*

LHCONE L3VPN

Layer3 Virtual Private Network

LHCONE L3VPN:

Layer3 (routed) Virtual Private Network

Dedicated worldwide backbone connecting **Tier1s, Tier2s and Tier3s** at high bandwidth

Reserved to HEP data transfers and analysis

Advantages

Bandwidth dedicated to HEP data analysis, no contention with other research projects

Well defined cost tag for WLCG networking

Trusted traffic that can bypass firewalls

LHCONE L3VPN architecture

- TierX sites connected to National-VRFs or Continental-VRFs
- National-VRFs interconnected via Continental-VRFs
- Continental-VRFs interconnected by trans-continental/trans-oceanic links

Acronyms: **VRF** = Virtual Routing Forwarding (virtual routing instance)

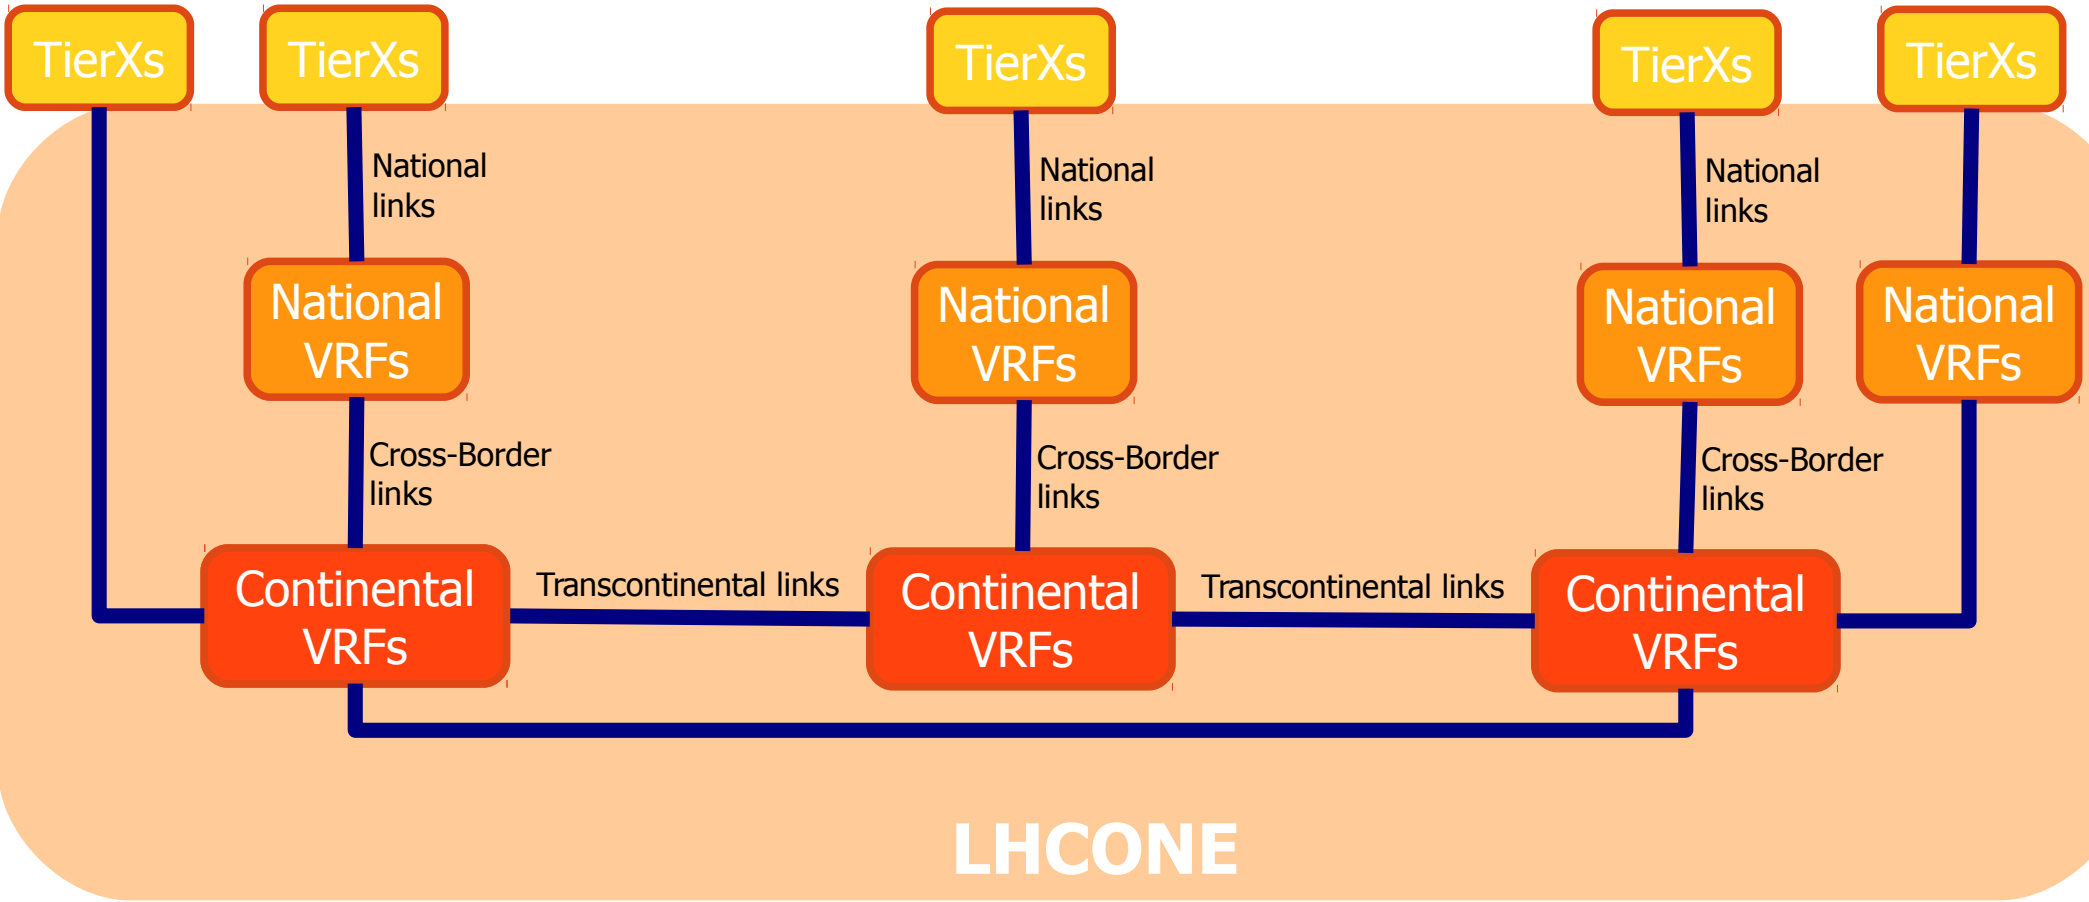

Current L3VPN topology

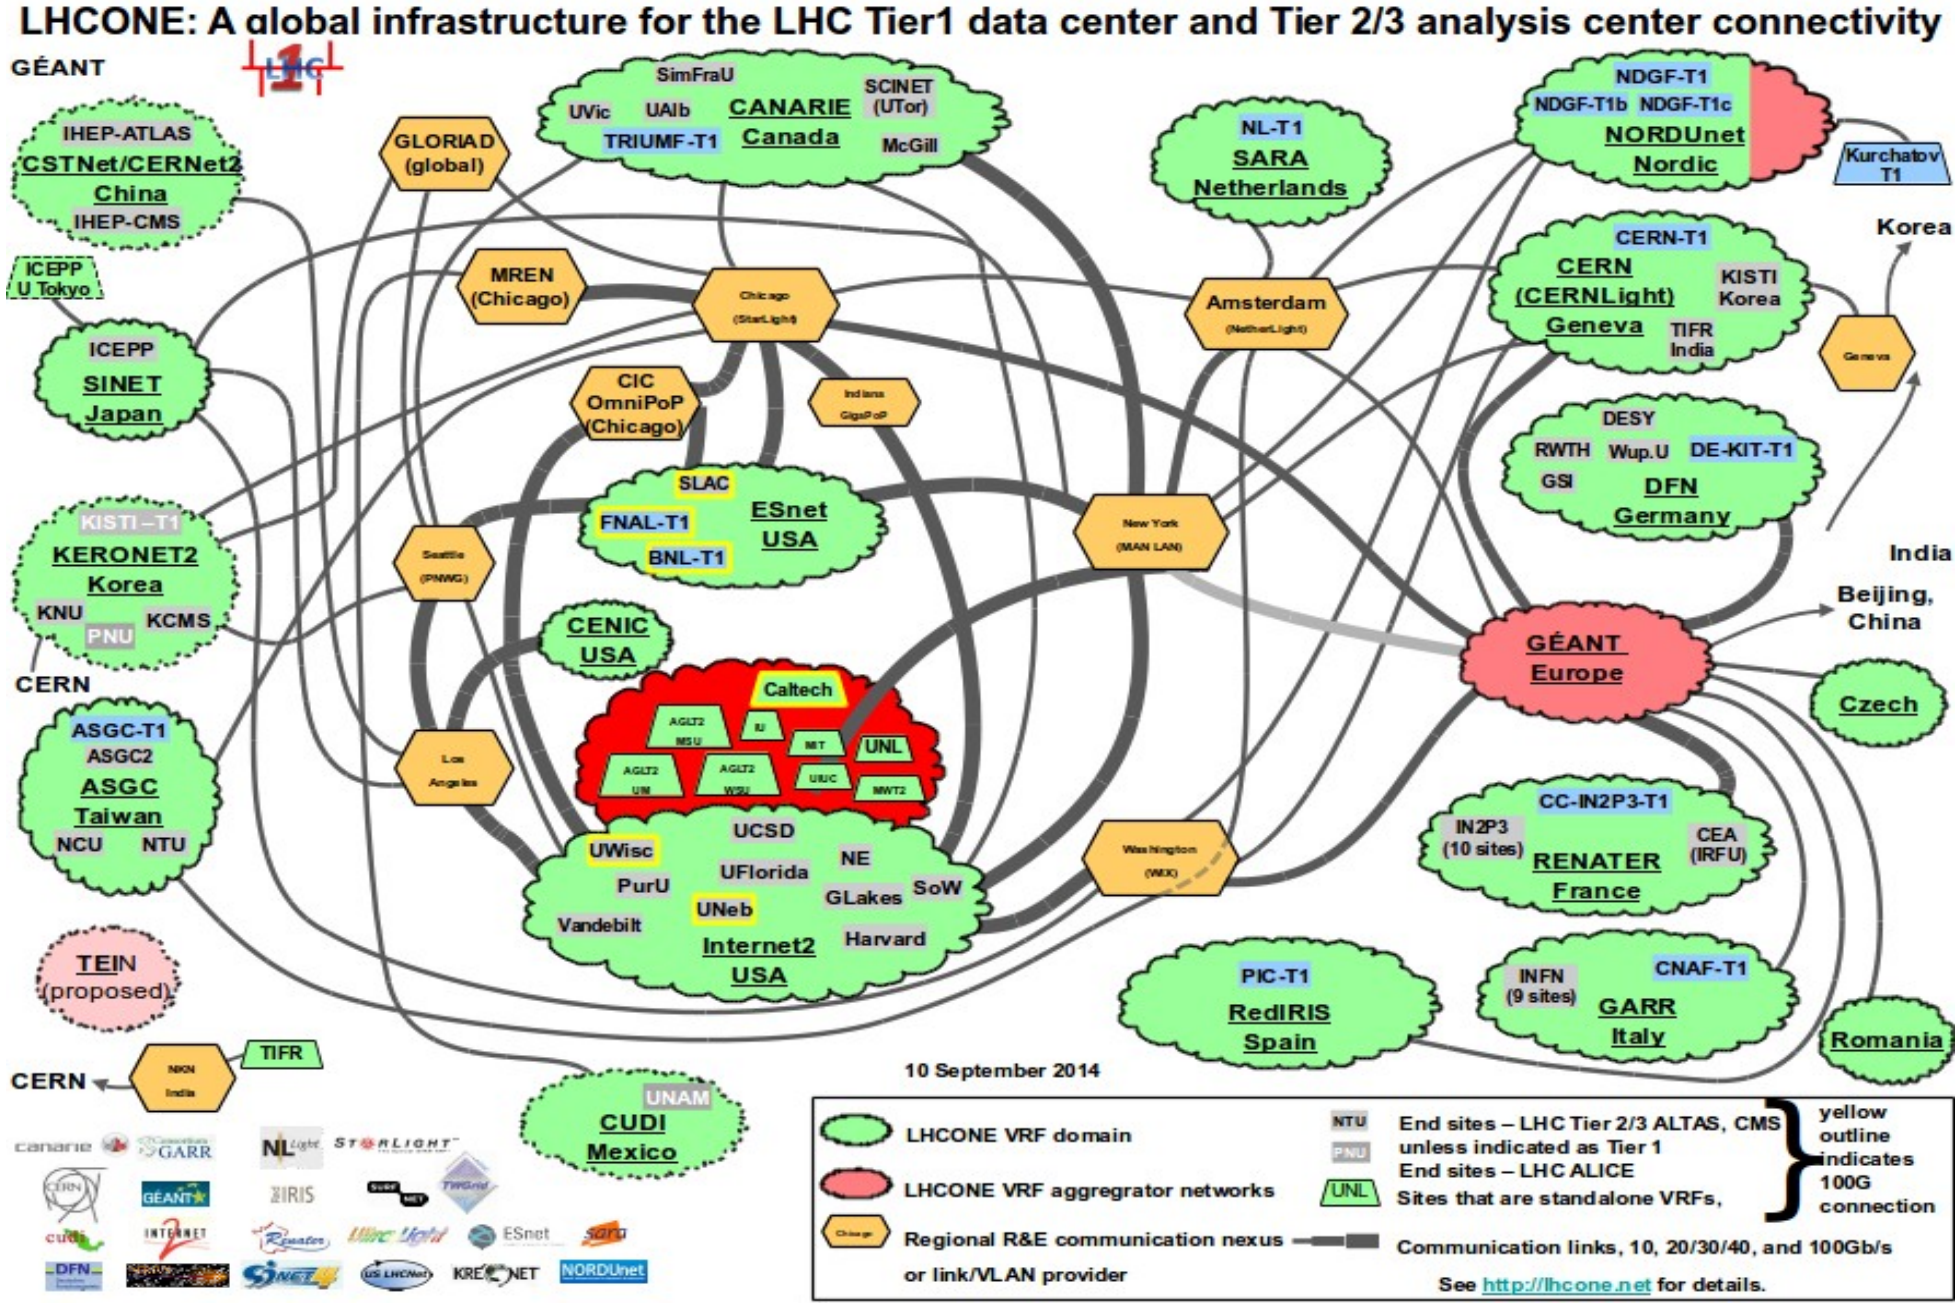

credits: Bill Johnston, ESnet

Status

Over 15 national and international Research Networks

Several Open Exchange Points including NetherLight, StarLight, MANLAN, CERNlight and others

Trans-Atlantic connectivity provided by ACE, GEANT, NORDUNET and USLHCNET

55 end sites connected to LHCONE:
- 10 Tier1s
- 45 Tier2s

Credits: Mian Usman, Enzo Capone - Dante
More Information: https://indico.cern.ch/event/318811/contribution/5/material/slides/0.pdf

L3VPN in Europe

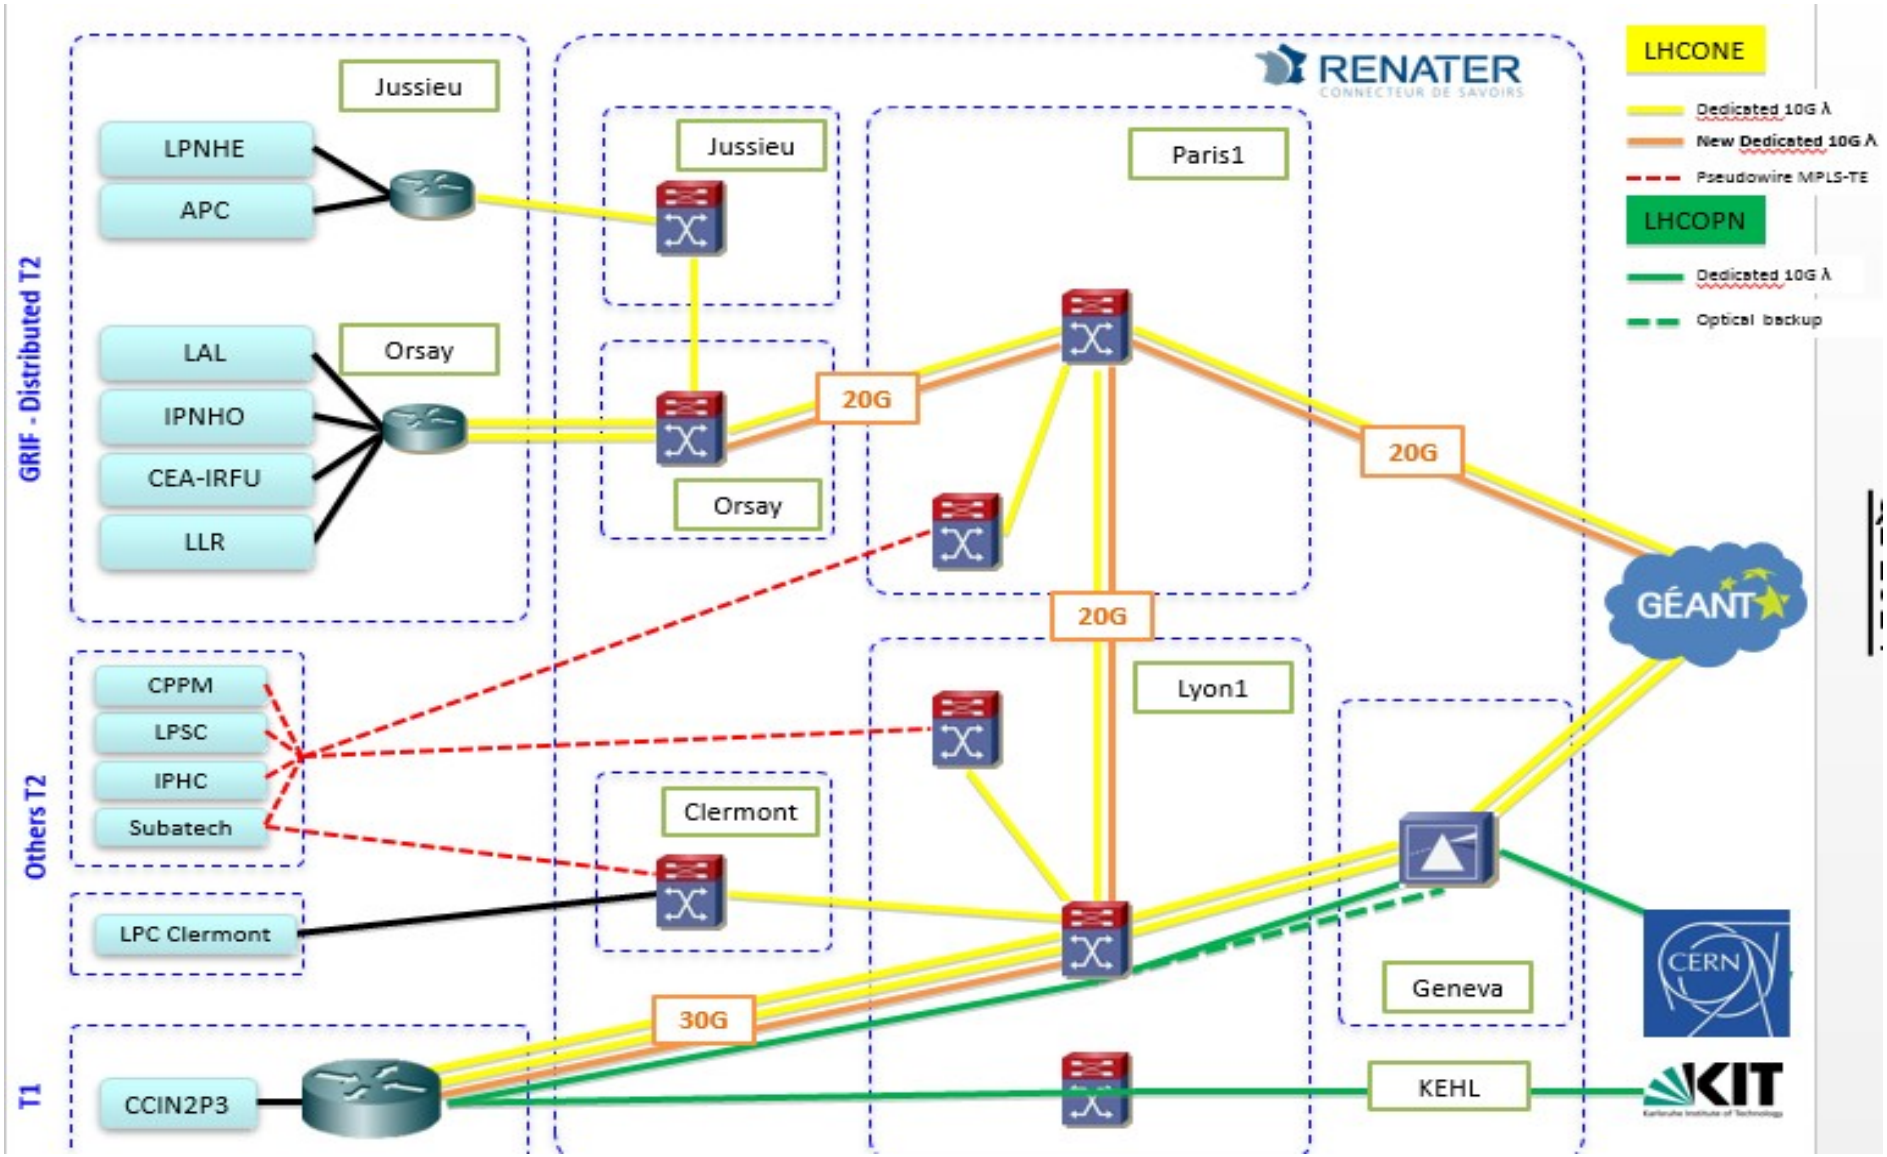

Credits: Mian Usman, Dante
Source: https://indico.cern.ch/event/289680/contribution/6/material/slides/0.ppt

L3VPN in Europe

LHCONE L3VPN Service in Europe is deployed in:

- ARNES (Slovenia)
- CERNLight(Switzerland)
- CESnet(Czech Republic)
- DFN (Germany)
- GARR (Italy)
- GEANT
- NORDUnet
- RENATER (France)
- RoEduNet(Romania)
- SURFnet (Netherlands)

RENs in these countries connect:

- 6 Tier 1s in Europe
- 30 Tier 2s in Europe
- Indian T2
- Russian T1

LHCONE P2P
Guaranteed bandwidth point-to-point links

LHCONE P2P description

On demand point-to-point (P2P) link system over a multi-domain network

Provides P2P links between any pair of TierX

Provides dedicated P2P links with guaranteed bandwidth (protected from any other traffic)

Accessible and configurable via software API

Status

Work in progress: still in design phase

Challenges:

- multi-domain provisioning system
- intra-Tier connectivity (extend links inside data centre)
- TierX-TierY Layer3 routing over dynamic links
- interfaces (APIs) for WLCG software

LHCONE perfSONAR

LHCONE perfSONAR:

LHCONE Network monitoring infrastructure

Probes installed at:
- VRFs interconnecting points
- TierXs

Accessible to any TierX for network healthiness checks

Status

Endorsed by WLCG to be a standard WLCG service

Probes already deployed in many TierXs.

Being deployed in the VRF networks

More information:
https://twiki.cern.ch/twiki/bin/view/LCG/PerfsonarDeployment

LHCONE evolution

LHCONE evolution

- VRFs' interconnection links will be upgraded to 100Gbps. 100Gbps Transatlantic link being tested.

- Few Tier sites connecting at 100Gbps

- Operations need to be improved, especially how to support a TierX in case of performance issue

LHCONE evolution - 2

- L3VPN opened to Belle-II collaboration. Other HEP projects might join as well

- LHCONE AUP (Acceptable Use Policy) being defined (link here)

- Asian connectivity will be improved

LHCONE Asian Workshop

- full report here: https://indico.cern.ch/event/318813/

- created lhcone-asia-pacific@cern.ch mailing list to reach all the concerned parts

- Work with TEIN to create the Asia-Pacific VRF hub which will keep the traffic local

- NRENs will create their LHCONE VRFs and connect them to TEIN and to other existing VRFs

Possible Asian-Pacific VRF setup

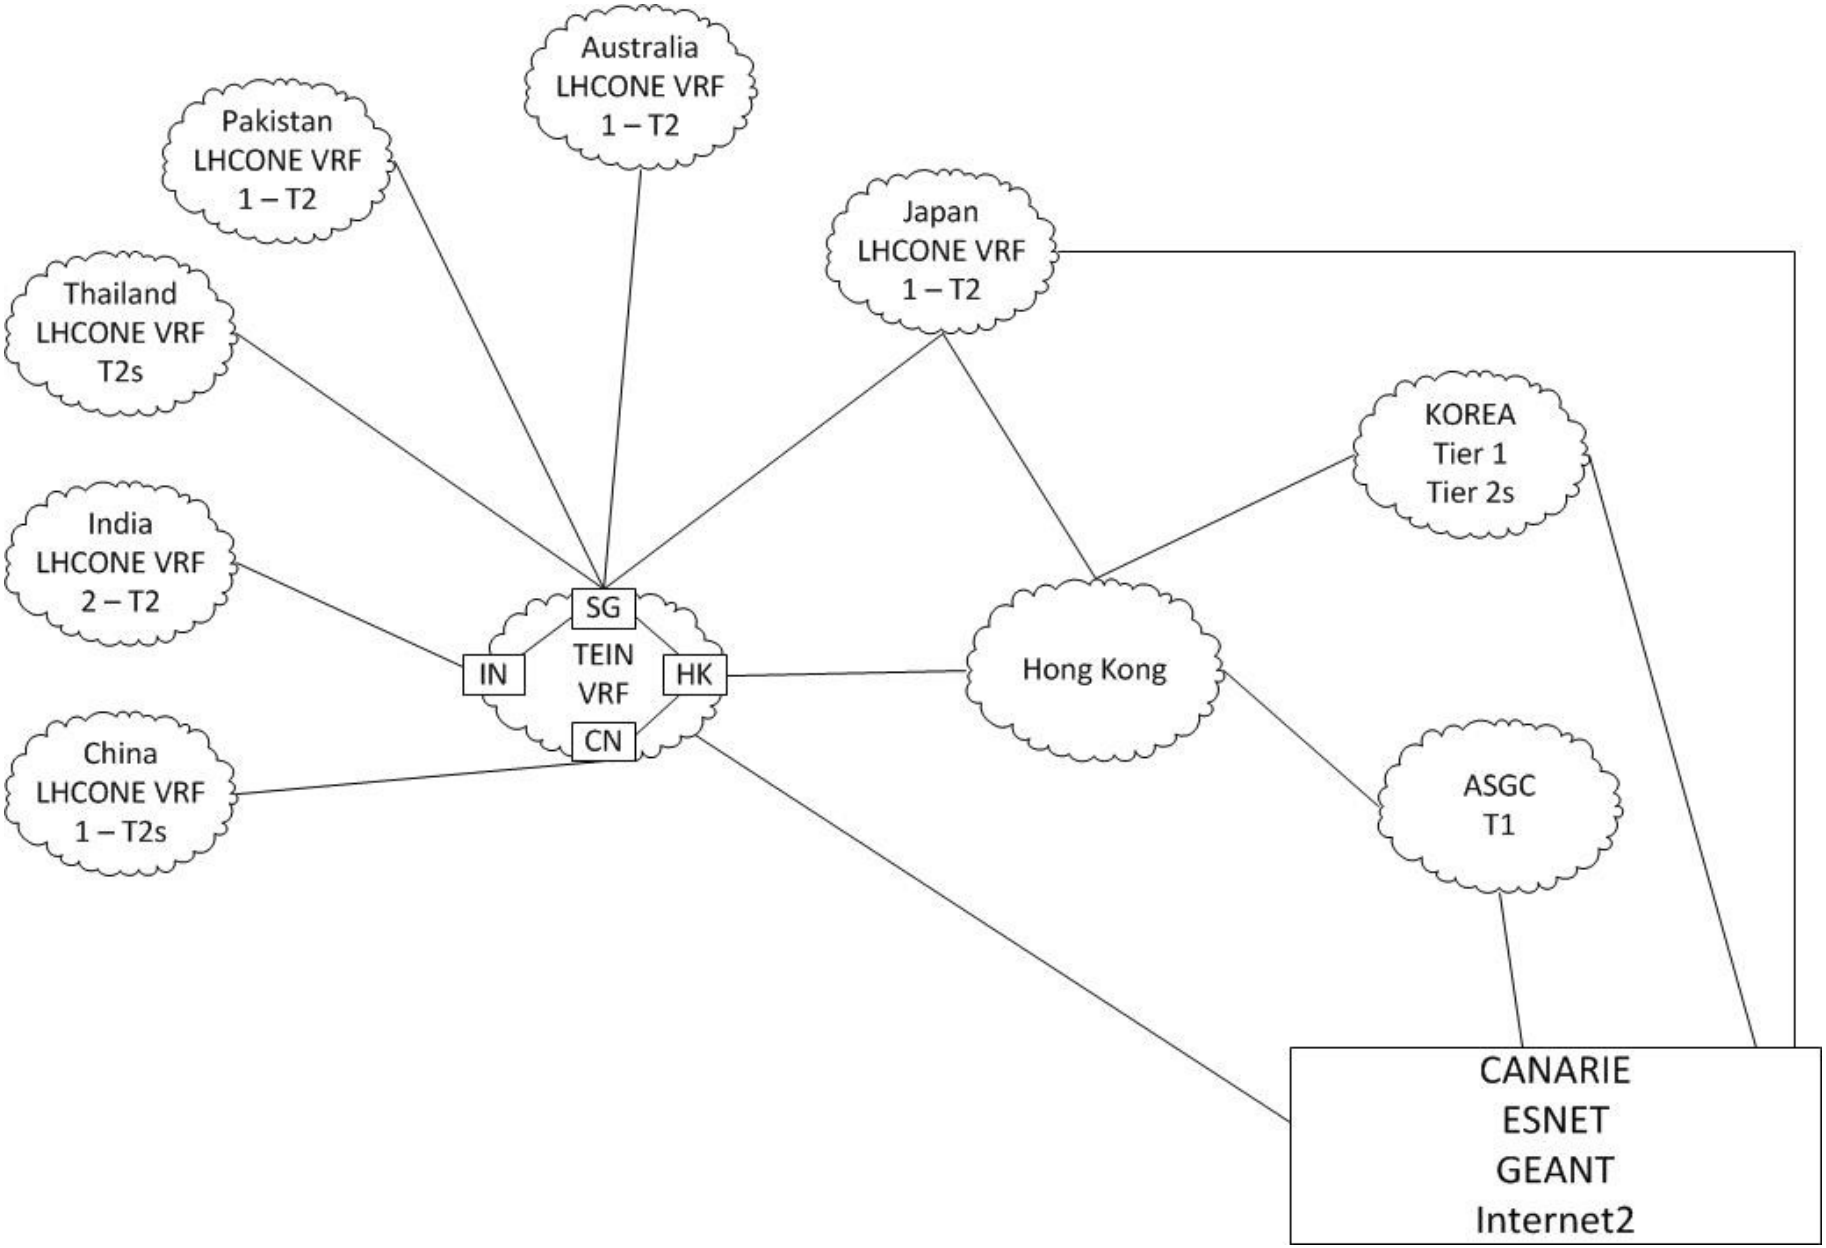

Chances for the Caucasus region

- In 2013, the NRENs in the South Caucasus region, AzRENA (Azerbaijan), GRENA (Georgia) and ASNET-AM (Armenia) became Partners of the GÉANT Project

- A Caucasus LHCONE VRF connected to GEANT would be desirable

E@Pconnect project

- E@Pconnect project (formerly known as EPIC) aims to enable the entire research and education community in the South Caucasus region to participate in collaborative activities on an equal footing with their peers within Europe and throughout the world

- The E@Pconnect project's implementation, initially foreseen for early 2015, has been slightly delayed. The decision is expected for Q1 2015 and the project to be launched in mid-2015

- The following countries have been identified as beneficiaries in the project: Armenia, Azerbaijan, Belarus, Georgia, Moldova, Ukraine

Geant connections to Asia

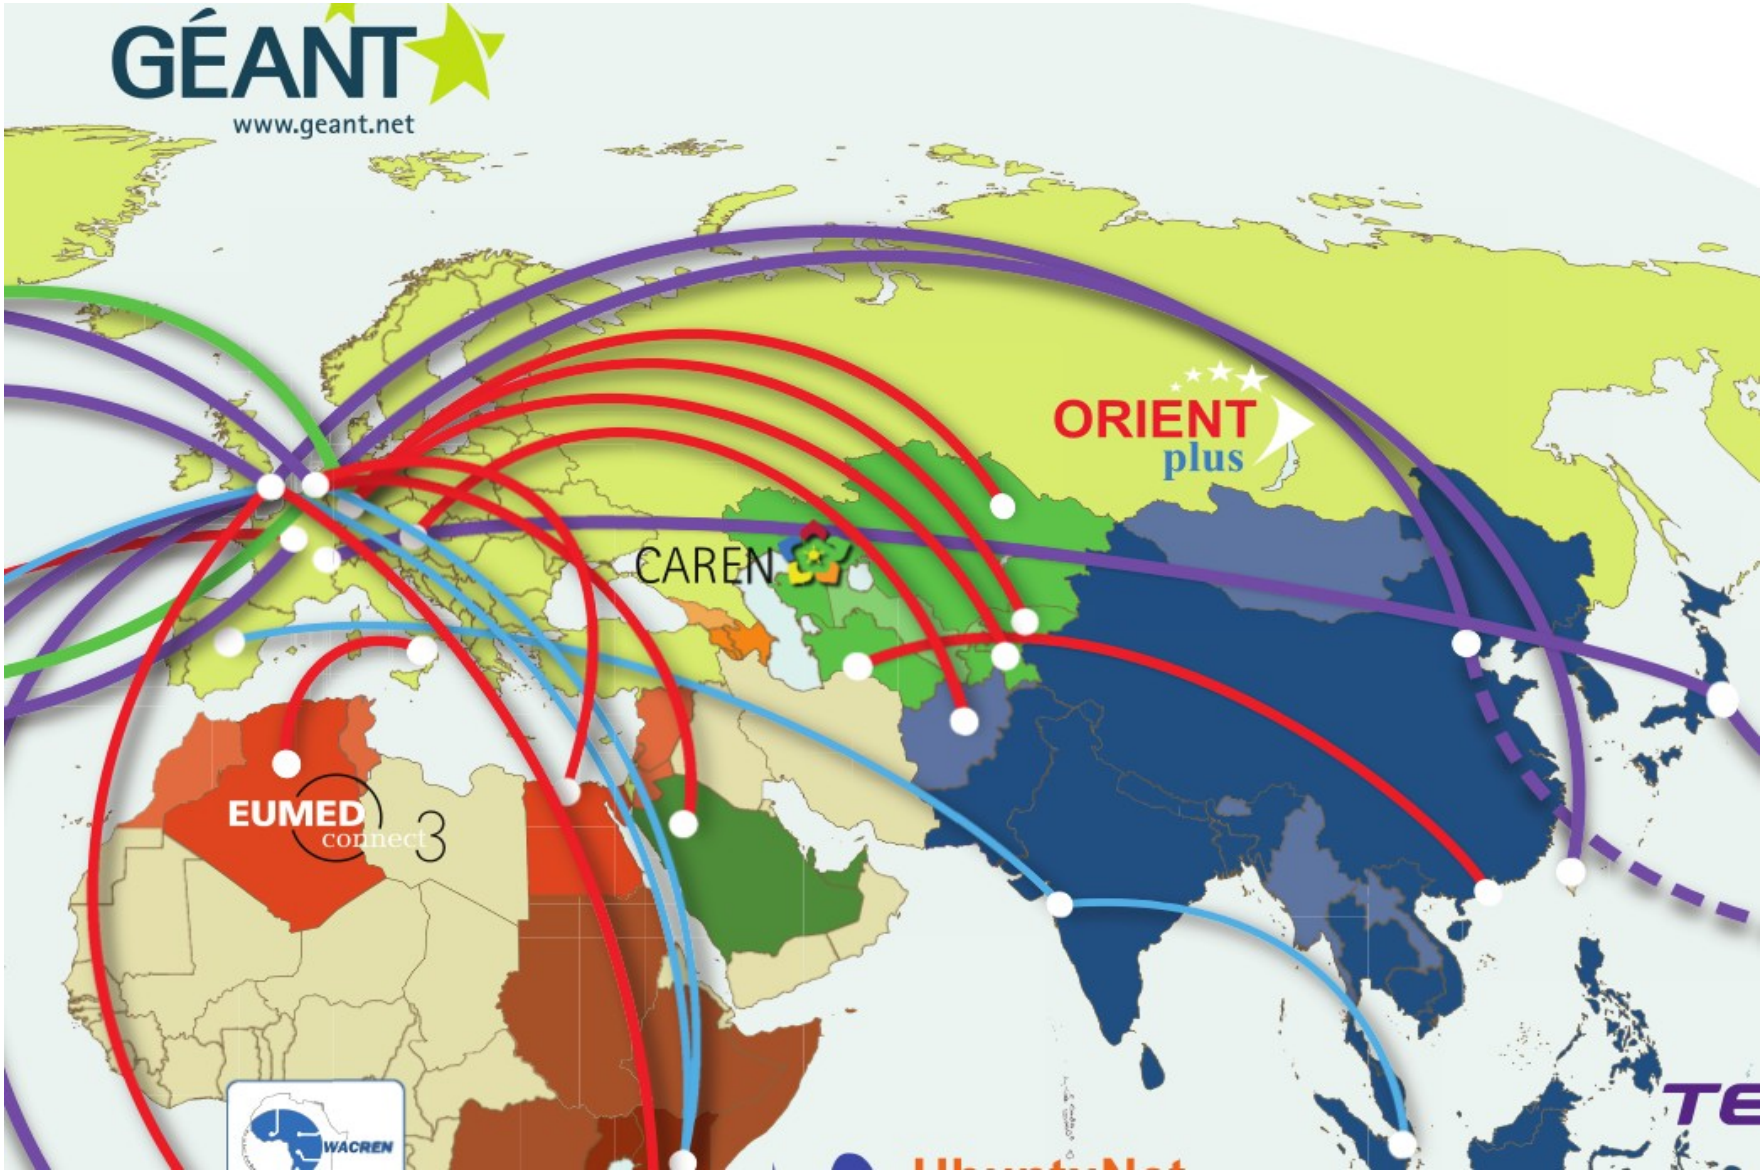

Conclusions

Conclusions

- CERN moving to 10G-100G technologies everywhere to accommodate soon increasing LHC data production

- Tier1s and Tier2s are improving their connectivity to CERN Tier0 and among themselves, to better exchange data

- Connectivity outside Europe and North America needs to be improved

Questions?