

WLCG Network Update

Alice workshop
Torino, 24th February 2015
Edoardo.Martelli@cern.ch

Summary



- Networking for WLCG
- LHCOPN
- LHCONE
 - services
 - how to join
- Recent networking activities at CERN



Networking for WLCG

Worldwide LHC Computing Grid




WLCG sites:

- 1 Tier0 (CERN)
- 13 Tier1s
- ~170 Tier2s
- >300 Tier3s worldwide




WLCG
Worldwide LHC Computing Grid

Computing model evolution



Original MONARCH model



Model evolution

**Computer Networks even more essential
component of WLCG**

**Data analysis in Run 2 will need more network
bandwidth between any pair of sites**

LHCOPN

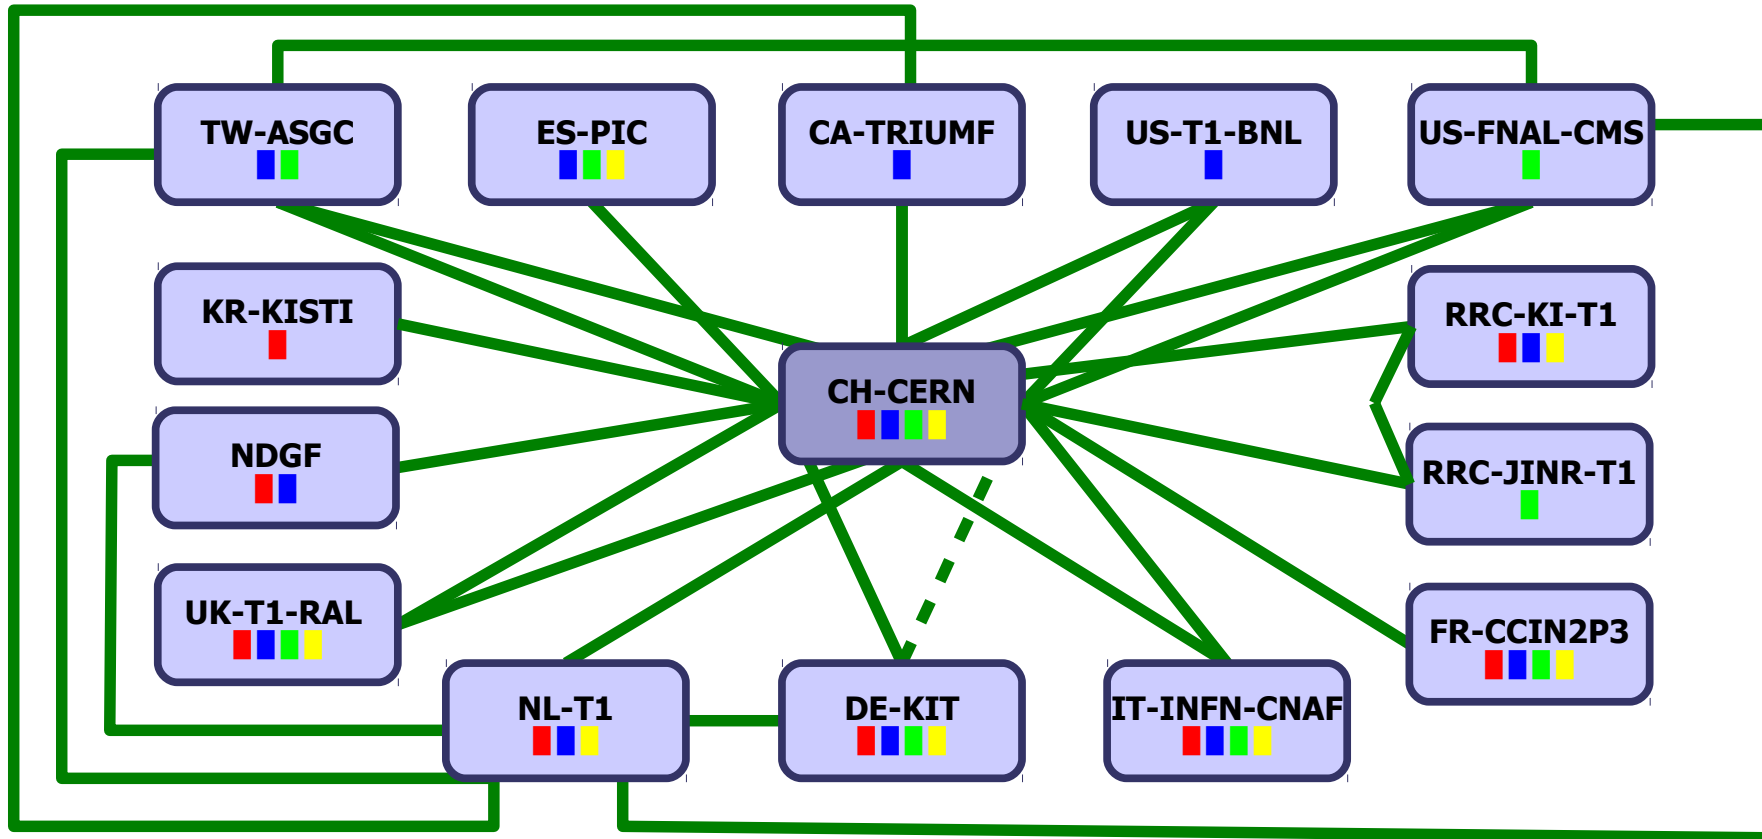
T0-T1 private network

LHCOPN is:



- a private network connecting Tier0 and Tier1s
- reserved to LHC data transfers and analysis
- dedicated large bandwidth links
- highly resilient with redundant architecture

Topology



■ = Alice ■ = Atlas ■ = CMS ■ = LHCb
edoardo.martelli@cern.ch 20150216

Technology



- **High bandwidth:** single and bundled long distance 10G and 100G Ethernet links
- **L3 Routing:** BGP peerings managed by the connected sites
- **Secure:** only declared IP prefixes can exchange traffic

On going development



- The LHCOPN will be kept as the main network to exchange data among the Tier0 and Tier1s
- Backup functionalities and Tier1-Tier1 traffic progressively moving to LHCONE (the Tier1-Tier2 network)
- Links to the Tier0 being upgraded to multiple 10Gbps and 100G

LHCONE

T1-T2-T3 private network

LHCONE principles



LHCONE is a network:

- **connecting any pair of sites**, regardless of the continent they reside
- **scalable**: sites are expected to grow
- **flexible**: sites may join and leave at any time
- **with a predictable cost tag**, sharing expensive resources
- **dedicated to HEP**: no clash with other data transfer, resource allocated for and funded by the HEP community

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

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

perfSONAR: monitoring infrastructure

LHCONE L3VPN

What LHCONE L3VPN is:



Layer3 (i.e. 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 HEP networking

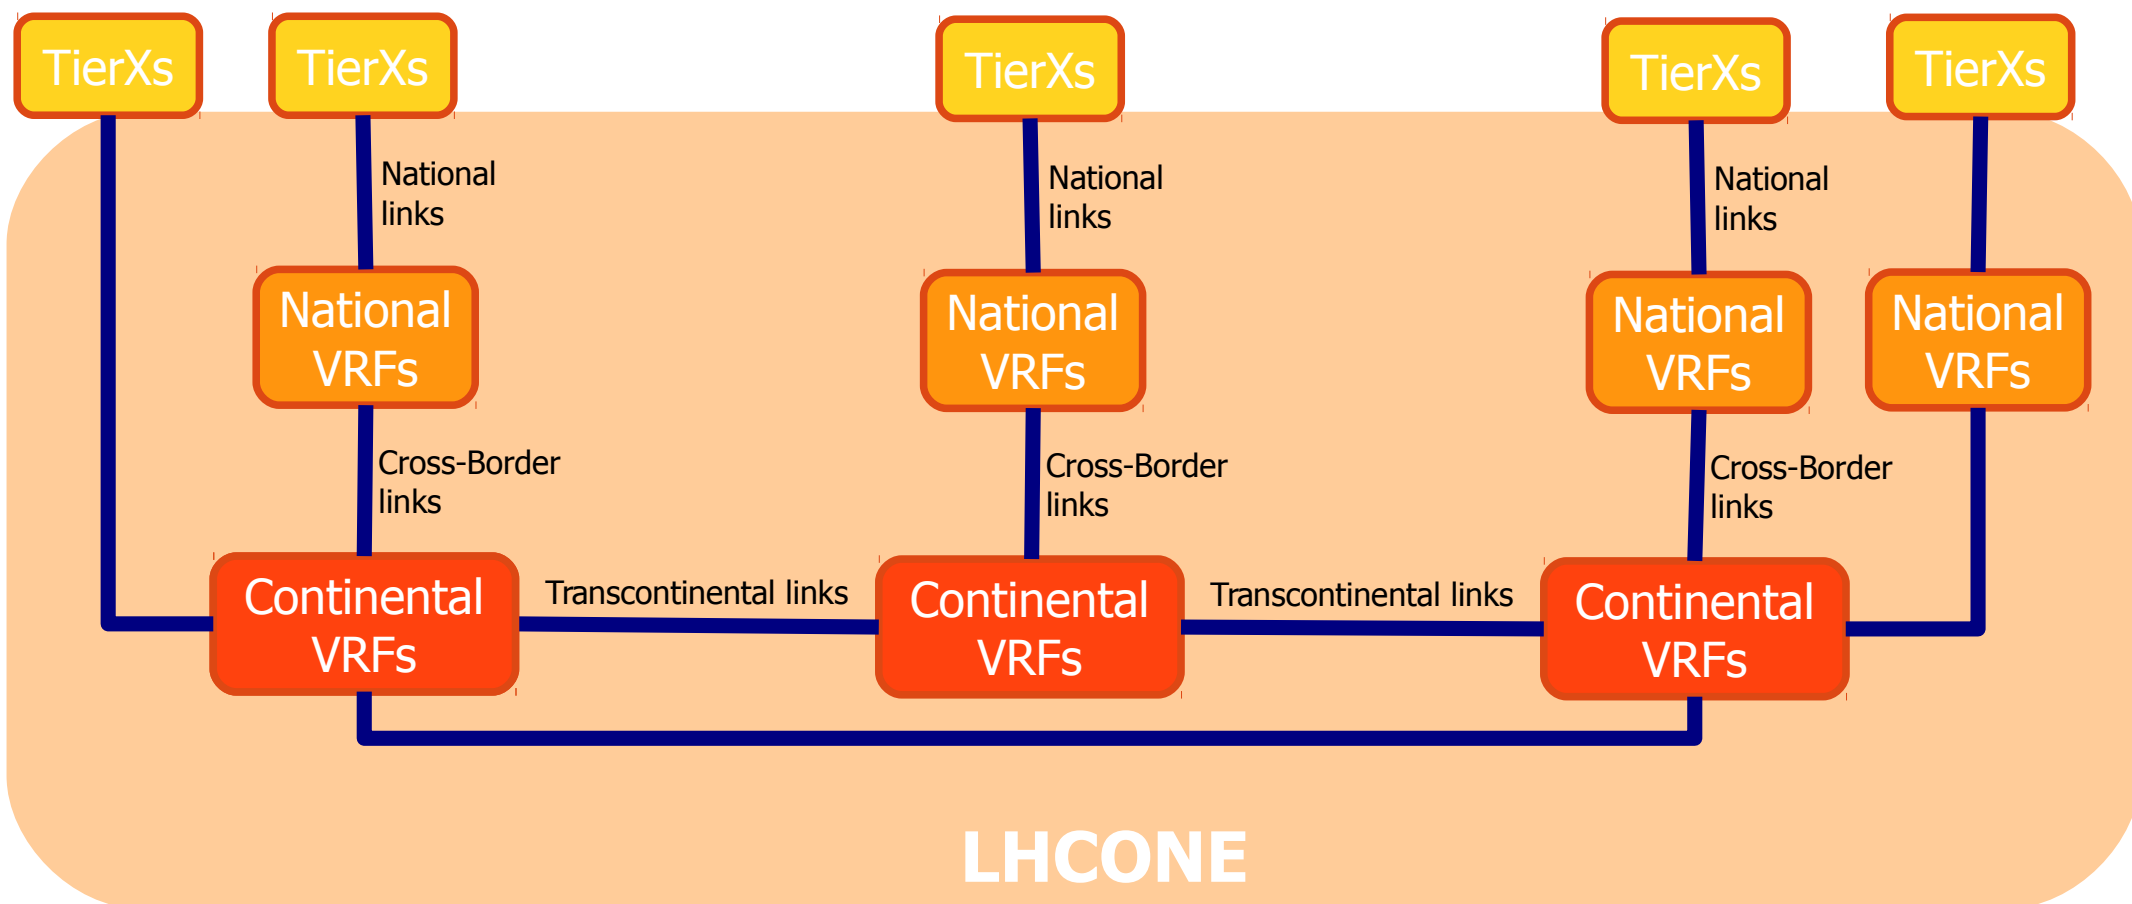
Trusted traffic that can bypass firewalls

LHCONE L3VPN architecture

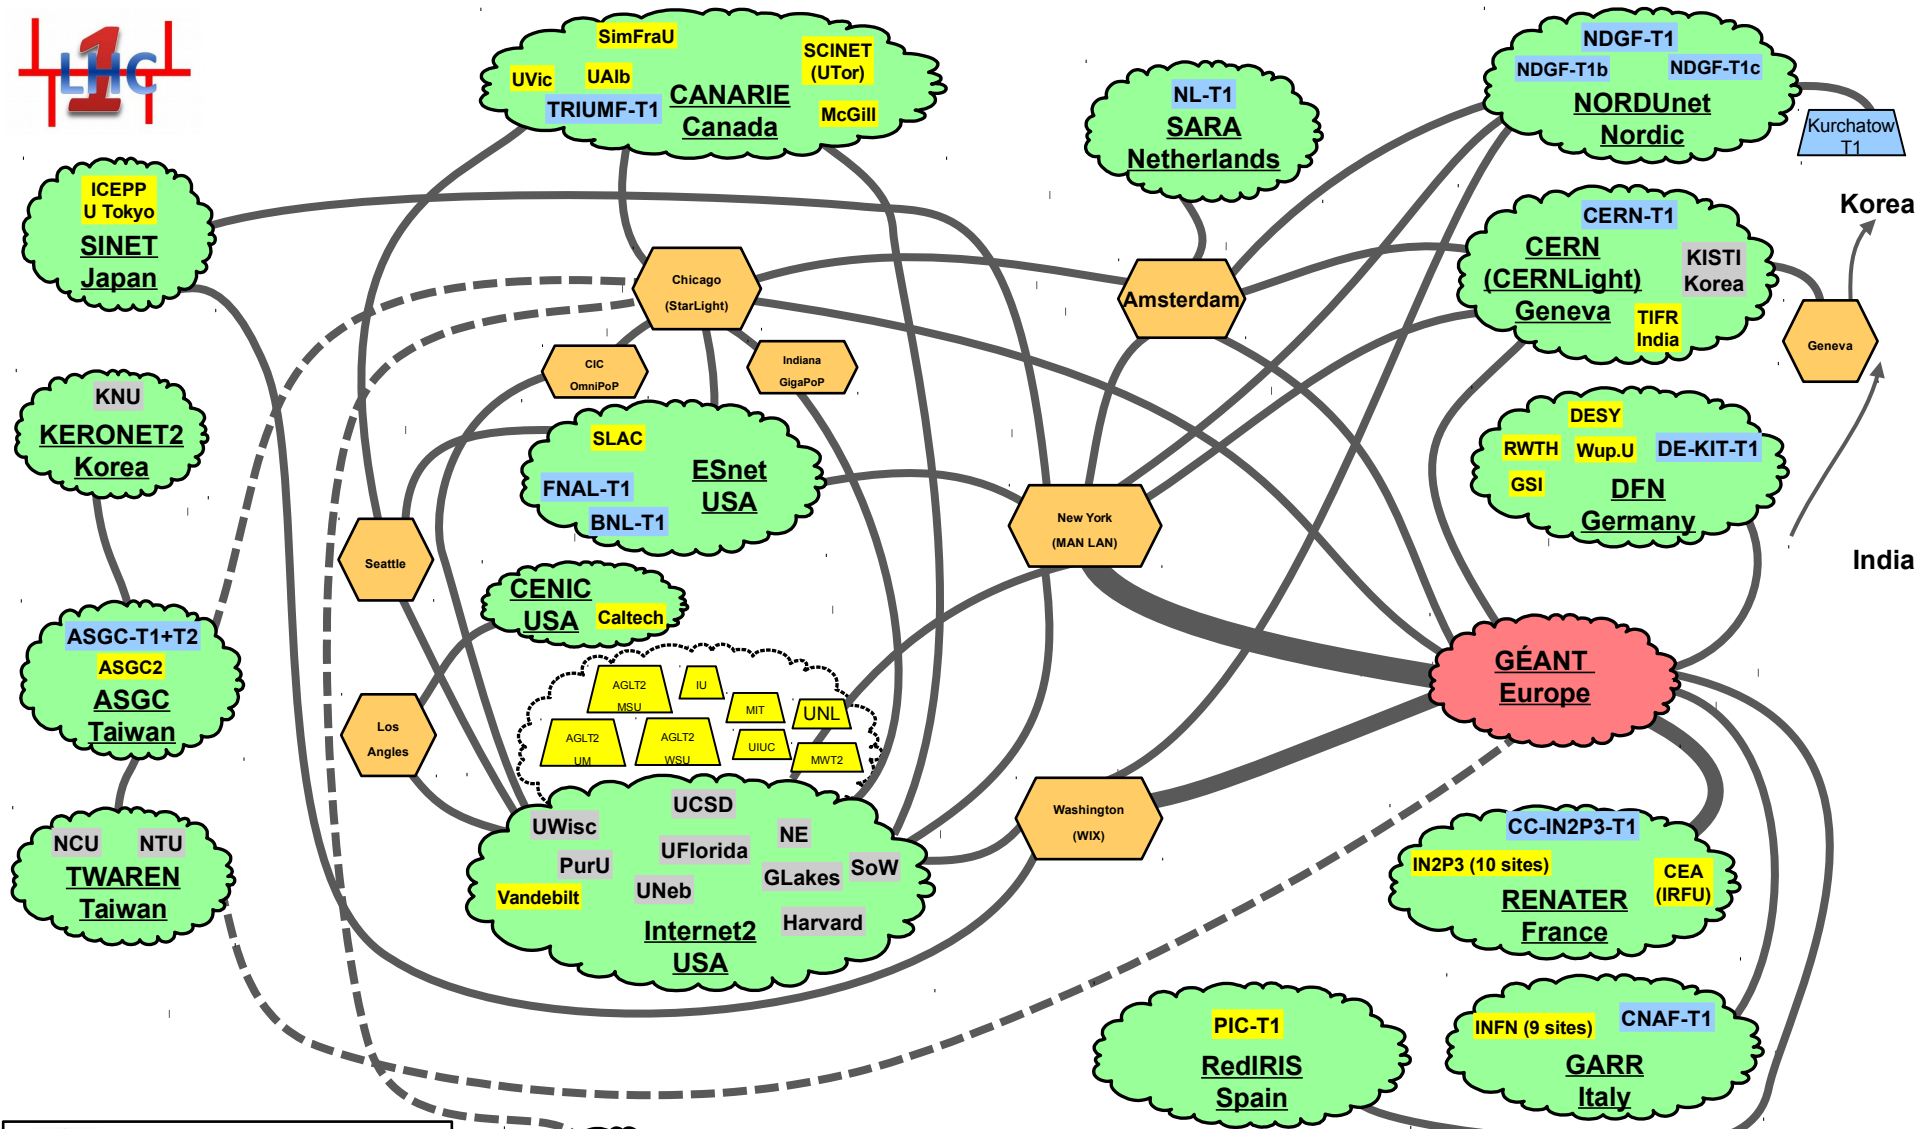


- TierX sites connected to National-VRFs or Continental-VRFs
- National-VRFs interconnected via Continental-VRFs
- Continental-VRFs interconnected by trans-continental/trans-oceanic links

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



Current L3VPN topology



LHCONE VRF domain
 LHCONE VRF aggregator networks
 End sites – LHC Tier 2/3 unless indicated as Tier 1
 Sites that are standalone VRFs
 Regional R&E communication nexuses

Communication links, 10, 20, 30, and 100Gb/s See <http://lhcone.net> for details.

August 7, 2014

credits: Bill Johnston, ESnet

Over 15 national and international Research Networks

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

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

~55 end sites connected to LHCONE:

- 10 Tier1s
- 45 Tier2s

Credits: Mian Usman, Dante

More Information: <https://indico.cern.ch/event/318811/contribution/5/material/slides/0.pdf>

Open to other HEP collaborations



L3VPN has been recently opened to the BelleII experiment

The LHCOPN AUP (Acceptable Use Policy) has been recently defined to regulate the utilization of the L3VPN service
(<https://twiki.cern.ch/twiki/bin/view/LHCONE/LhcOneAup>)

How to join the L3VPN

Pre-requisites



The TierX site needs to have:

- Public IP addresses
- A public Autonomous System (AS) number
- A BGP capable router

How to connect



The TierX has to:

- Contact the Network Provider that runs the closest LHCONE VRF
- Agree on the cost of the access
- Lease a link from the TierX site to the closest LHCONE VRF PoP (Point of Presence)
- Configure the BGP peering with the Network Provider

TierX routing setup

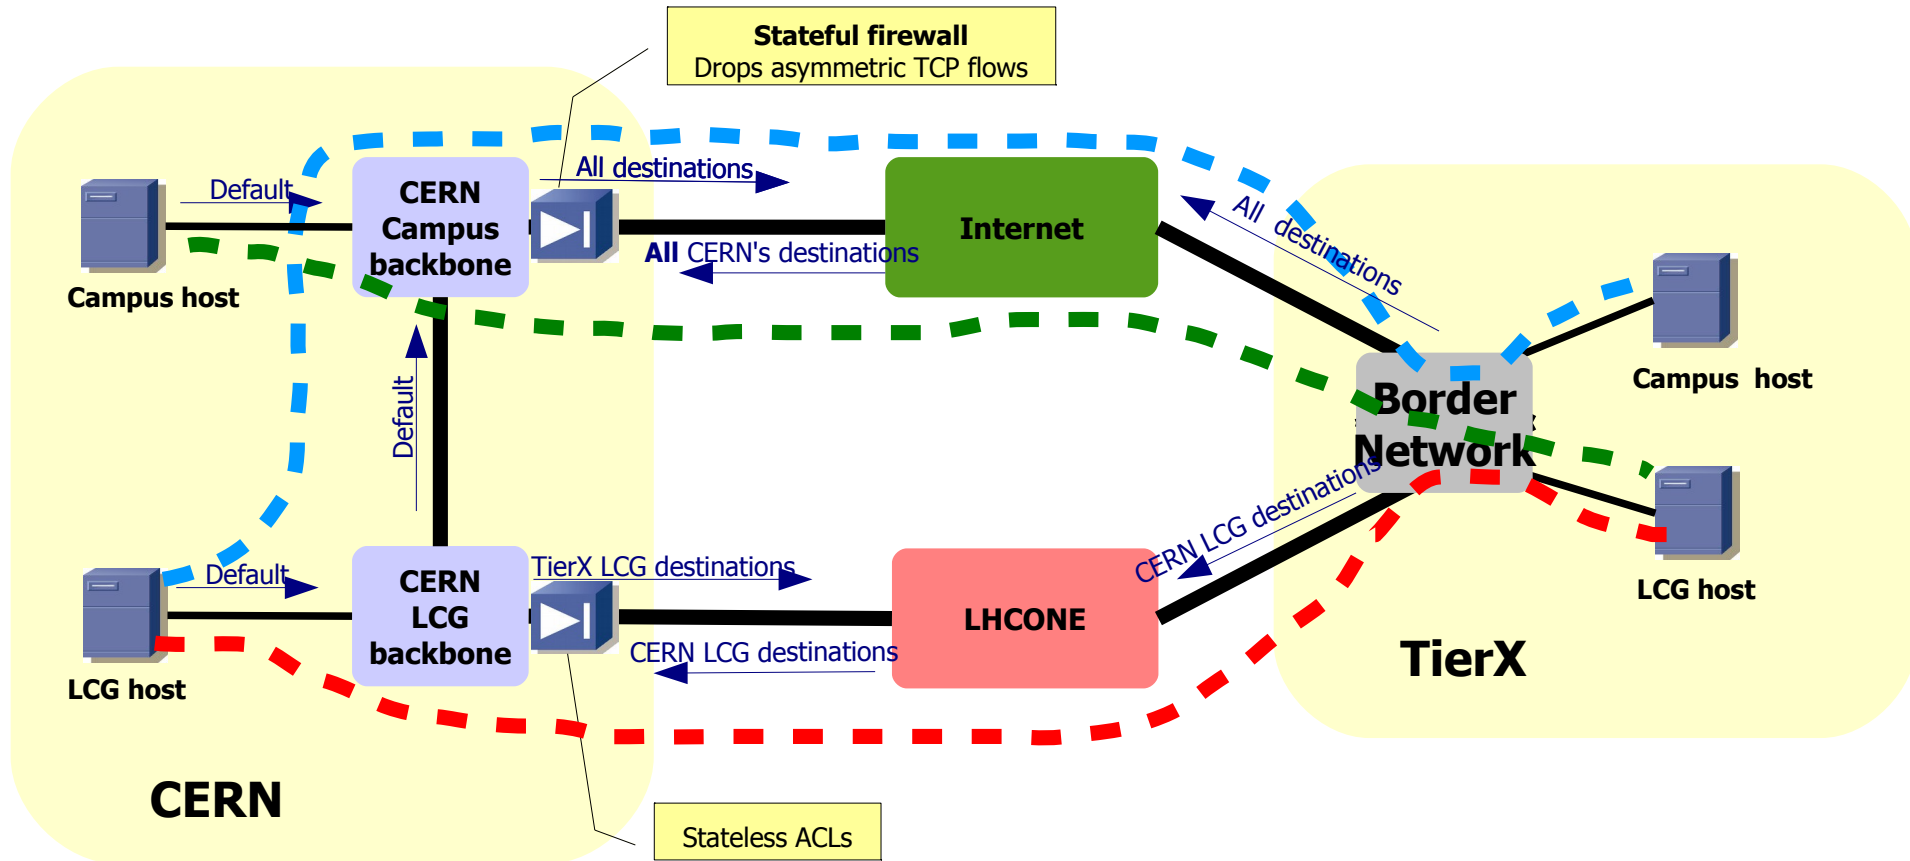


- The TierX announce only the IP subnets used for WLCG servers
- The TierX accepts all the prefixes announced by the LHCONE VRF
- The TierX **must** assure traffic symmetry: injects only packets sourced by the announced subnets
- LHCONE traffic may be allowed to bypass the central firewall (up to the TierX to decide)

Symmetric traffic is essential



Beware: statefull firewalls discard unidirectional TCP connections



- LHCONE host to LHCONE host
- CERN's LHCONE host to TierX not LHCONE host
- CERN's not LHCONE host to TierX's LHCONE host

LHCONE P2P

Guaranteed bandwidth point-to-point links

What LHCONE P2P is (will be):



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

Provides P2P links between any pair of TierX. The P2P links have guaranteed bandwidth (protected from any other traffic)

Accessible and configurable via software API

Work in progress: still in design phase

Challenges:

- multi-domain provisioning system
- intra-TierX connectivity
- TierX-TierY routing
- Integration with WLCG software

LHCONE perfSONAR

- framework for active and passive network probing
- developed by Internet2, Esnet, Geant and others

What is LHCONE perfSONAR



LHCONE Network monitoring infrastructure

Probes installed at:

- VRFs interconnecting points
- Sites

Accessible to any Site for network healthiness checks

Endorsed by WLCG to be a standard WLCG service

Probes already deployed in many sites

Being deployed in the VRF networks

More information:

<https://twiki.opensciencegrid.org/bin/view/Documentation/DeployperfSONAR>

LHCONE evolution

LHCONE evolution



- VRFs have upgraded internal links to 100Gbps
- Bigger Sites now connected with 100Gbps links
- VRFs interconnection links being upgraded to 100Gbps. 100Gbps Transatlantic becoming the norm
- Operational procedures being improved. Strong commitment of all the network operators
- perfSONAR deployment progressing well

More information



LHCONE meetings:

<https://indico.cern.ch/category/5793/>

LHCONE websites:

<http://lhcone.net>

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

Mailing lists:

lhcone-operations@cern.ch

lhcone-architecture@cern.ch

lhcone-asia-pacific@cern.ch

Recent Networking activities at CERN

Long Shutdown 1 activities



Replaced all the switches of the Technical Network (LHC control network)

Re-cabled major LHC starpoints

Upgraded Technical Network core routers to improve stability and security

Upgraded major Campus routers to improve stability

Experiments' Data Acquisition links upgraded

External connectivity: Geant



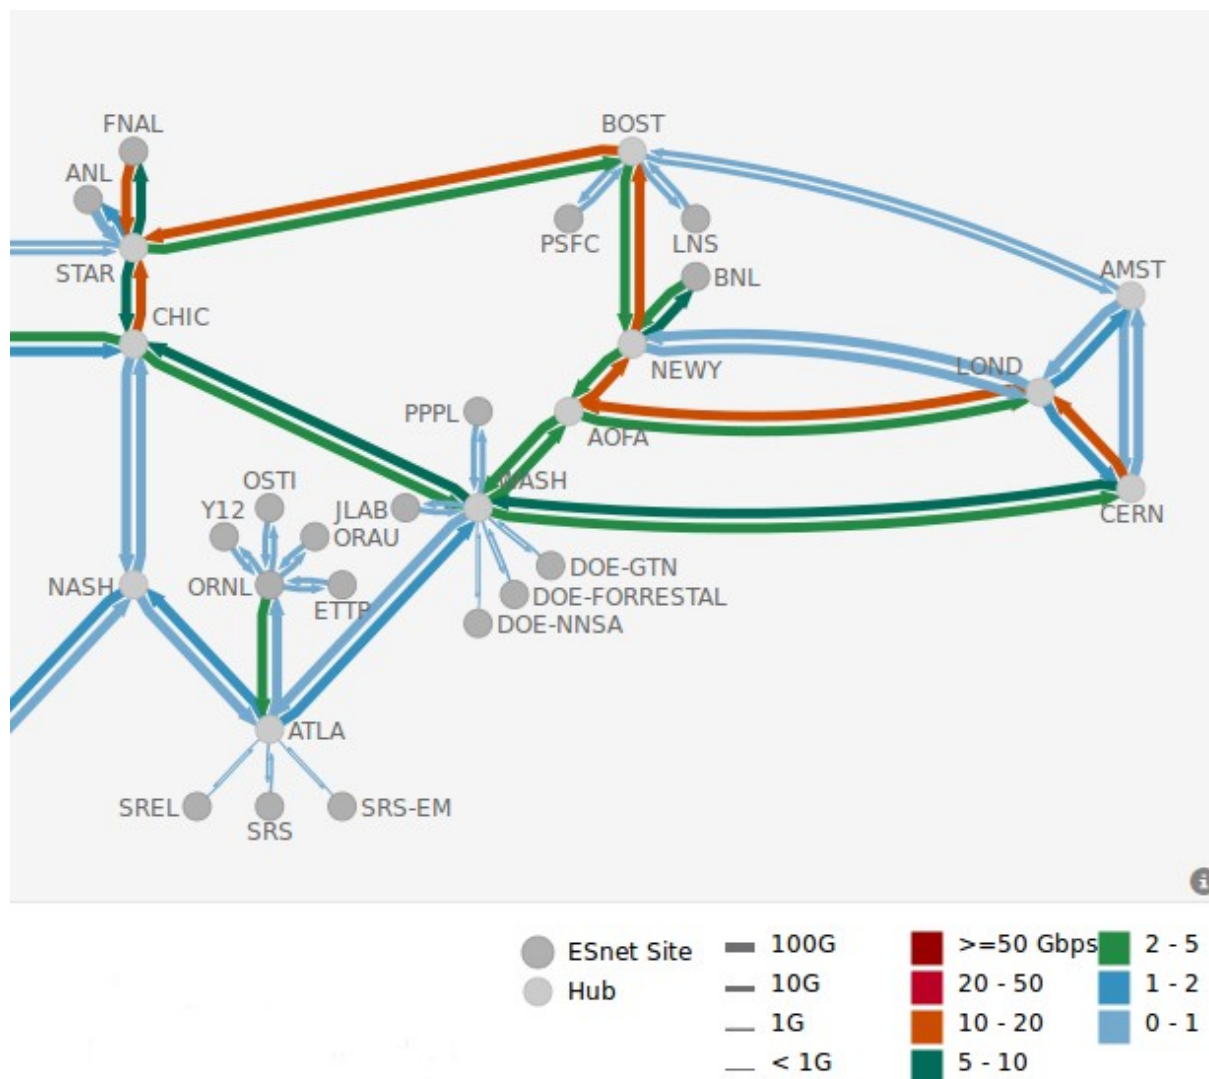
Connection with Geant LHCONE has been upgraded to 100G (was 30G)

Connection to Geant IP (general purpose) being upgraded to 100G (currently 50G)

External connectivity: ESnet



ESnet (U.S. Energy Science Network) has extended its network to Europe with three 100G transatlantic link. They have established three new PoPs in Europe, one of them at CERN



CERN Tier0 datacentre



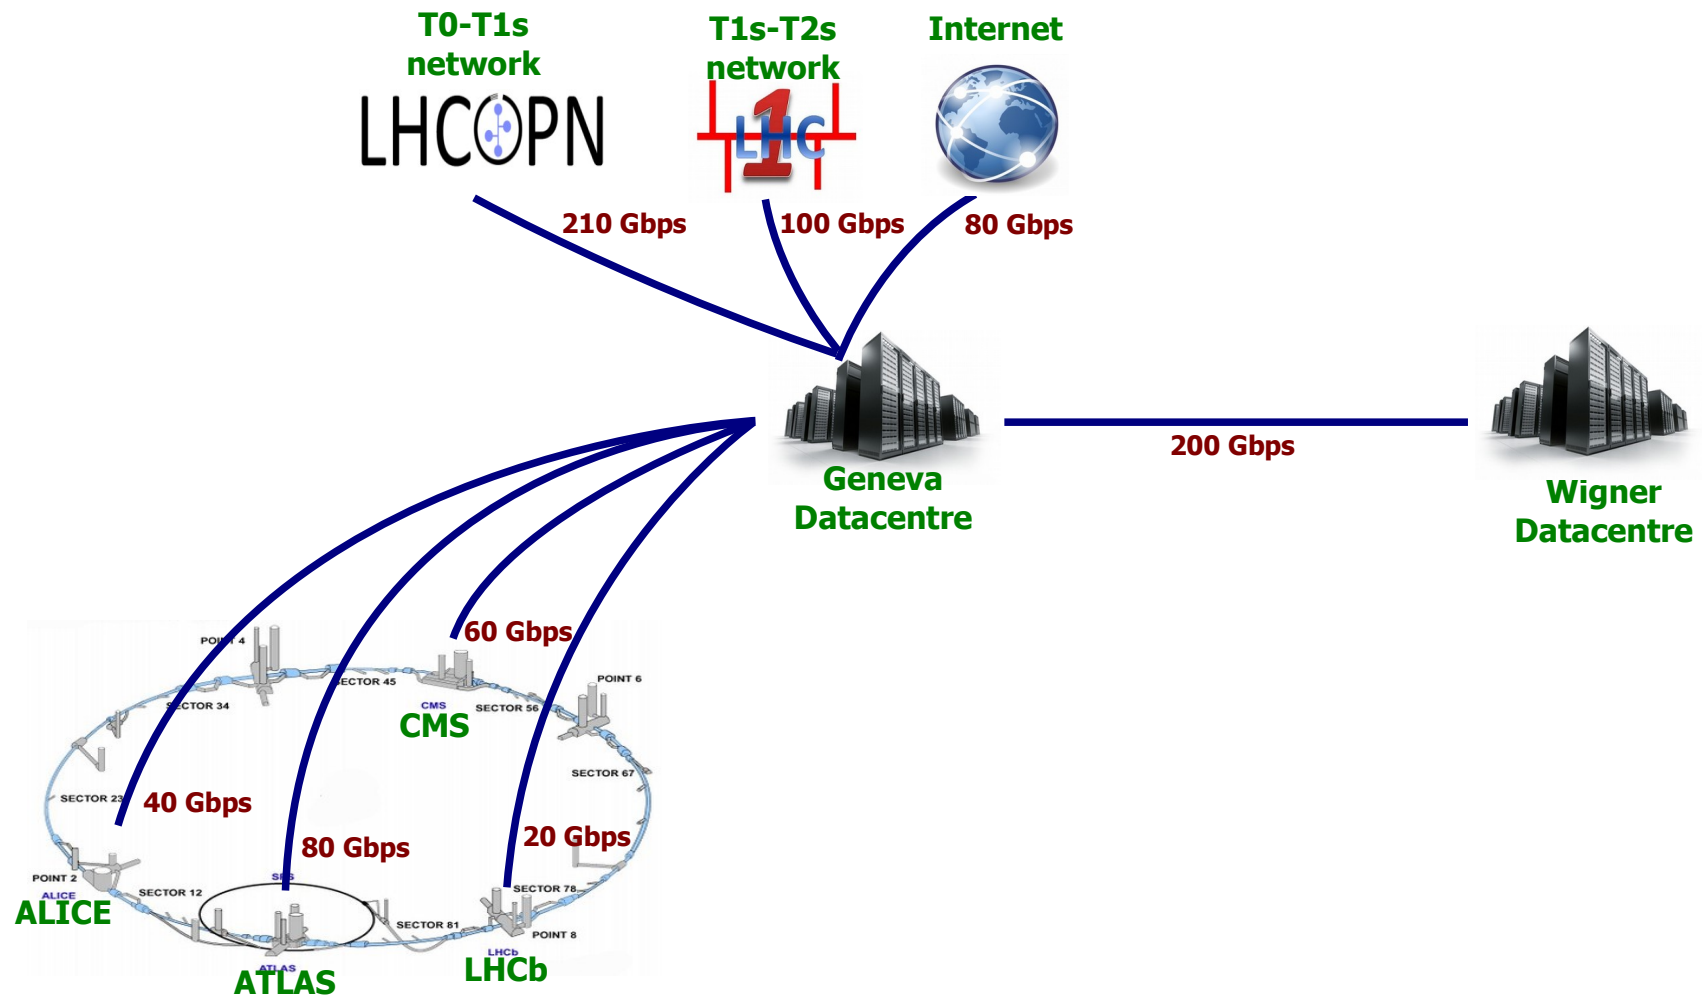
Wigner Budapest: Installation in the 3rd room has started

Current situation:

MEYRIN DATA CENTRE		WIGNER DATA CENTRE		NETWORK AND STORAGE	
	last_value		last_value		last_value
● Number of Cores in Meyrin	119,093	● Number of Cores in Wigner	20,576	● Tape Drives	115
● Number of Drives in Meyrin	78,124	● Number of Drives in Wigner	15,075	● Tape Cartridges	23,213
● Number of Memory Modules in Meyrin	79,124	● Number of Memory Modules in Wigner	10,263	● Data Volume on Tape (TB)	105,248
● Number of 10G NIC in Meyrin	4,280	● Number of 10G NIC in Wigner	1,211	● Free Space on Tape (TB)	30,546
● Number of 1G NIC in Meyrin	22,203	● Number of 1G NIC in Wigner	2,224	● Routers (GPN)	134
● Number of Processors in Meyrin	21,311	● Number of Processors in Wigner	2,574	● Routers (TN)	26
● Number of Servers in Meyrin	11,557	● Number of Servers in Wigner	1,290	● Routers (Others)	99
● Total Disk Space in Meyrin (TB)	143,948	● Total Disk Space in Wigner (TB)	47,694	● Star Points	633
● Total Memory Capacity in Meyrin (TB)	443	● Total Memory Capacity in Wigner (TB)	83	● Switches	3,490

Source: https://meter.cern.ch/public/_plugin/kibana/#/dashboard/elasticsearch/Overview:%20Data%20Centre
February 2015

Main datacentre connections




Services: IPv6



IPv6 deployment completed at CERN.

Testing of IPv6 compliance of WLCG applications is progressing
(HEPiX activity)



Questions?