
CERN

LHCONE services and status

Alice Forum
KISTI Daejeon, 23rd September 2015
Edoardo.Martelli@cern.ch

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

Summary

- Networking for WLCG
- LHCONE
  - Services
  - Status

Networking for WLCG

Worldwide LHC Computing Grid

WLCG sites:

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

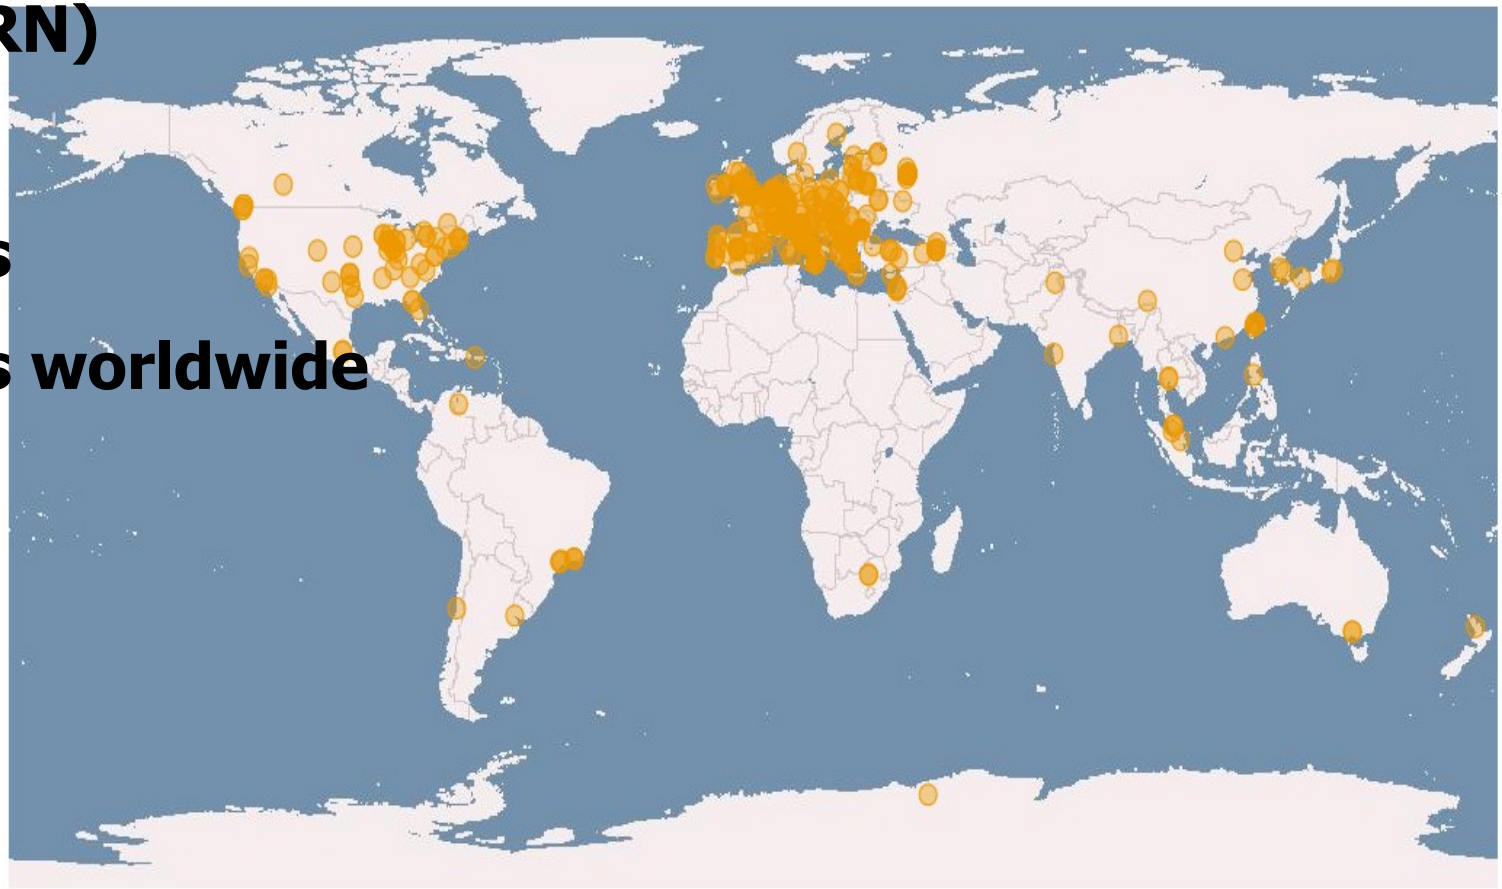

WLCG
Worldwide LHC Computing Grid

Original computing model

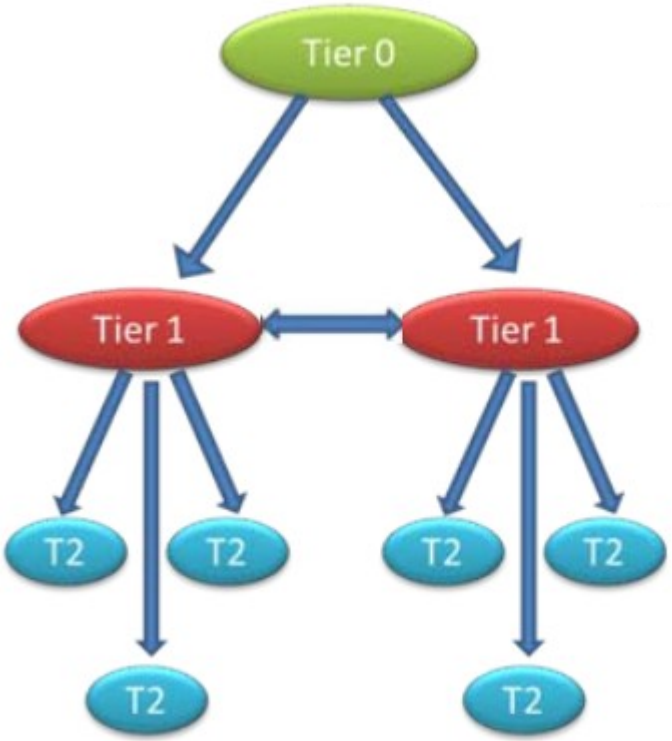

Original MONARCH model

LHCOPN: Tier0-Tier1 network

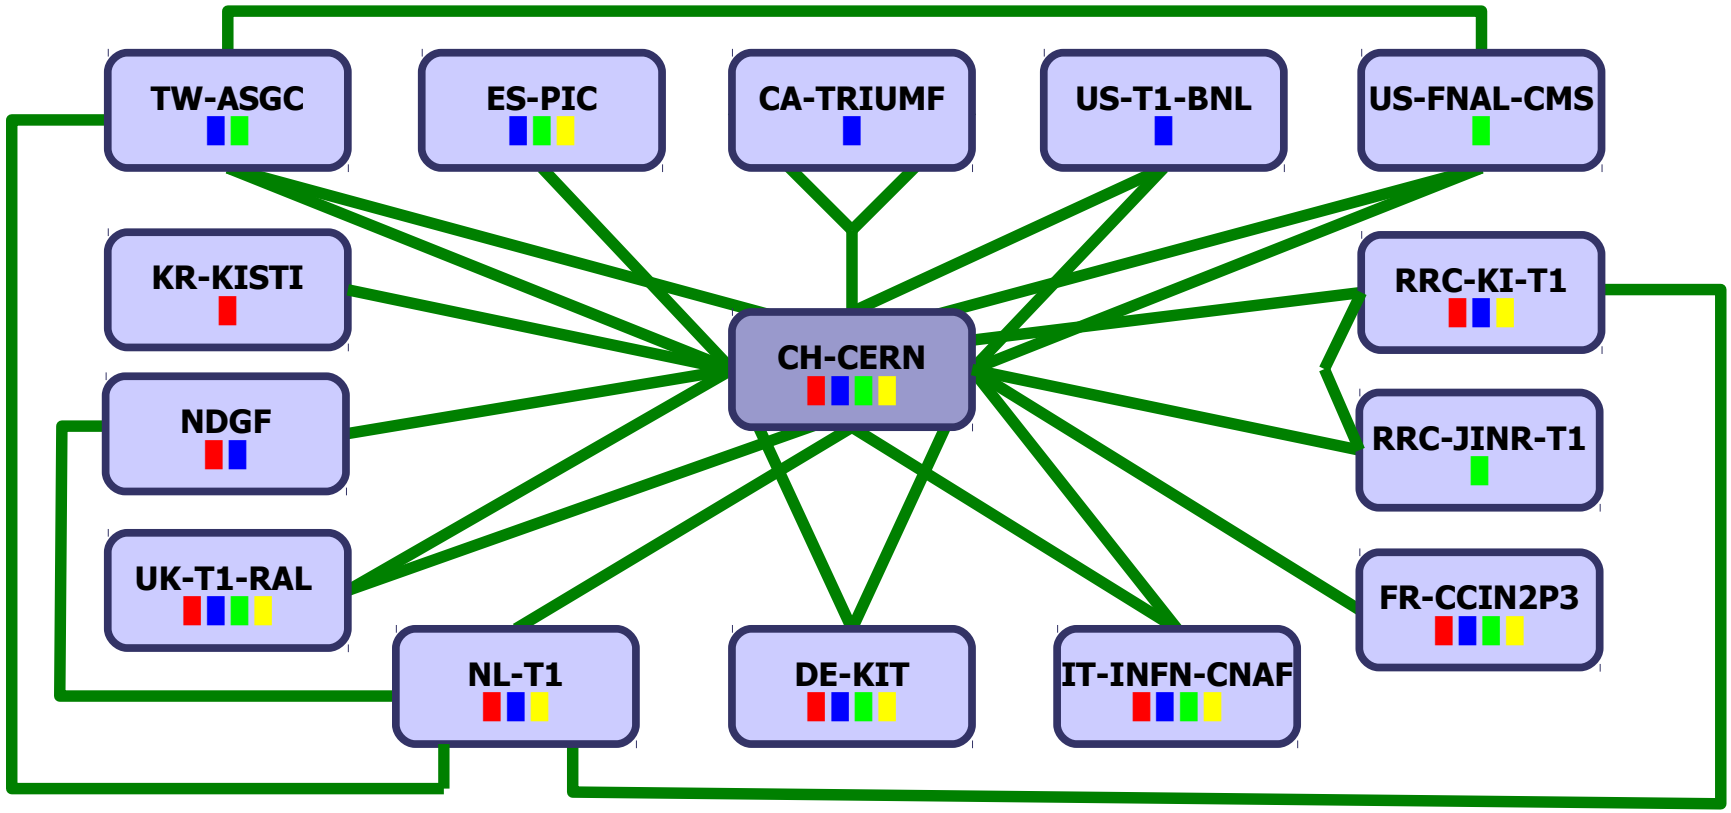

■ = Alice ■ = Atlas ■ = CMS ■ = LHCb

edoardo.martelli@cern.ch 20150821

Computing model evolution

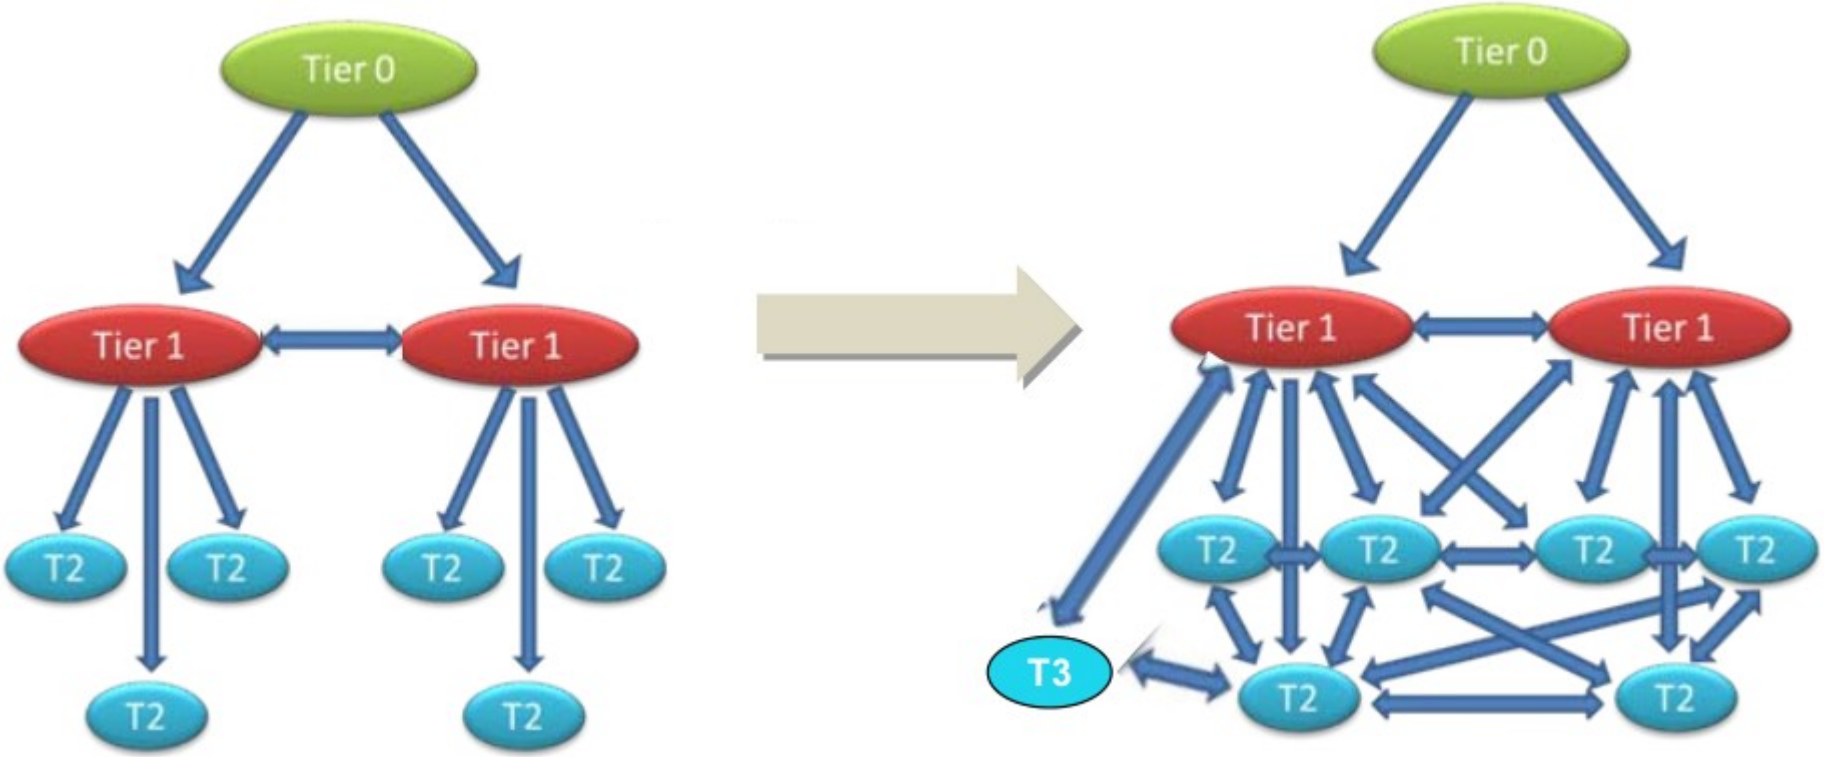

Original MONARCH model

Model evolution

Role of Networks in WLCG

Computer Networks even more essential component of WLCG

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

LHCONE
T1-T2-T3 private network

LHCONE principles

LHCONE is a multi-domain network:

- **connecting any pair of sites**, regardless of the continent or network they reside in
- **scalable**: sites are allowed to grow
- **flexible**: sites may join and leave at any time
- **with a predictable cost tag:** by sharing expensive resources
- **dedicated to HEP**: 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 – *under development*

perfSONAR: monitoring infrastructure- *operational*

LHCONE L3VPN service

LHCONE L3VPN service

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

Benefits

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 (i.e. virtual routing instance)

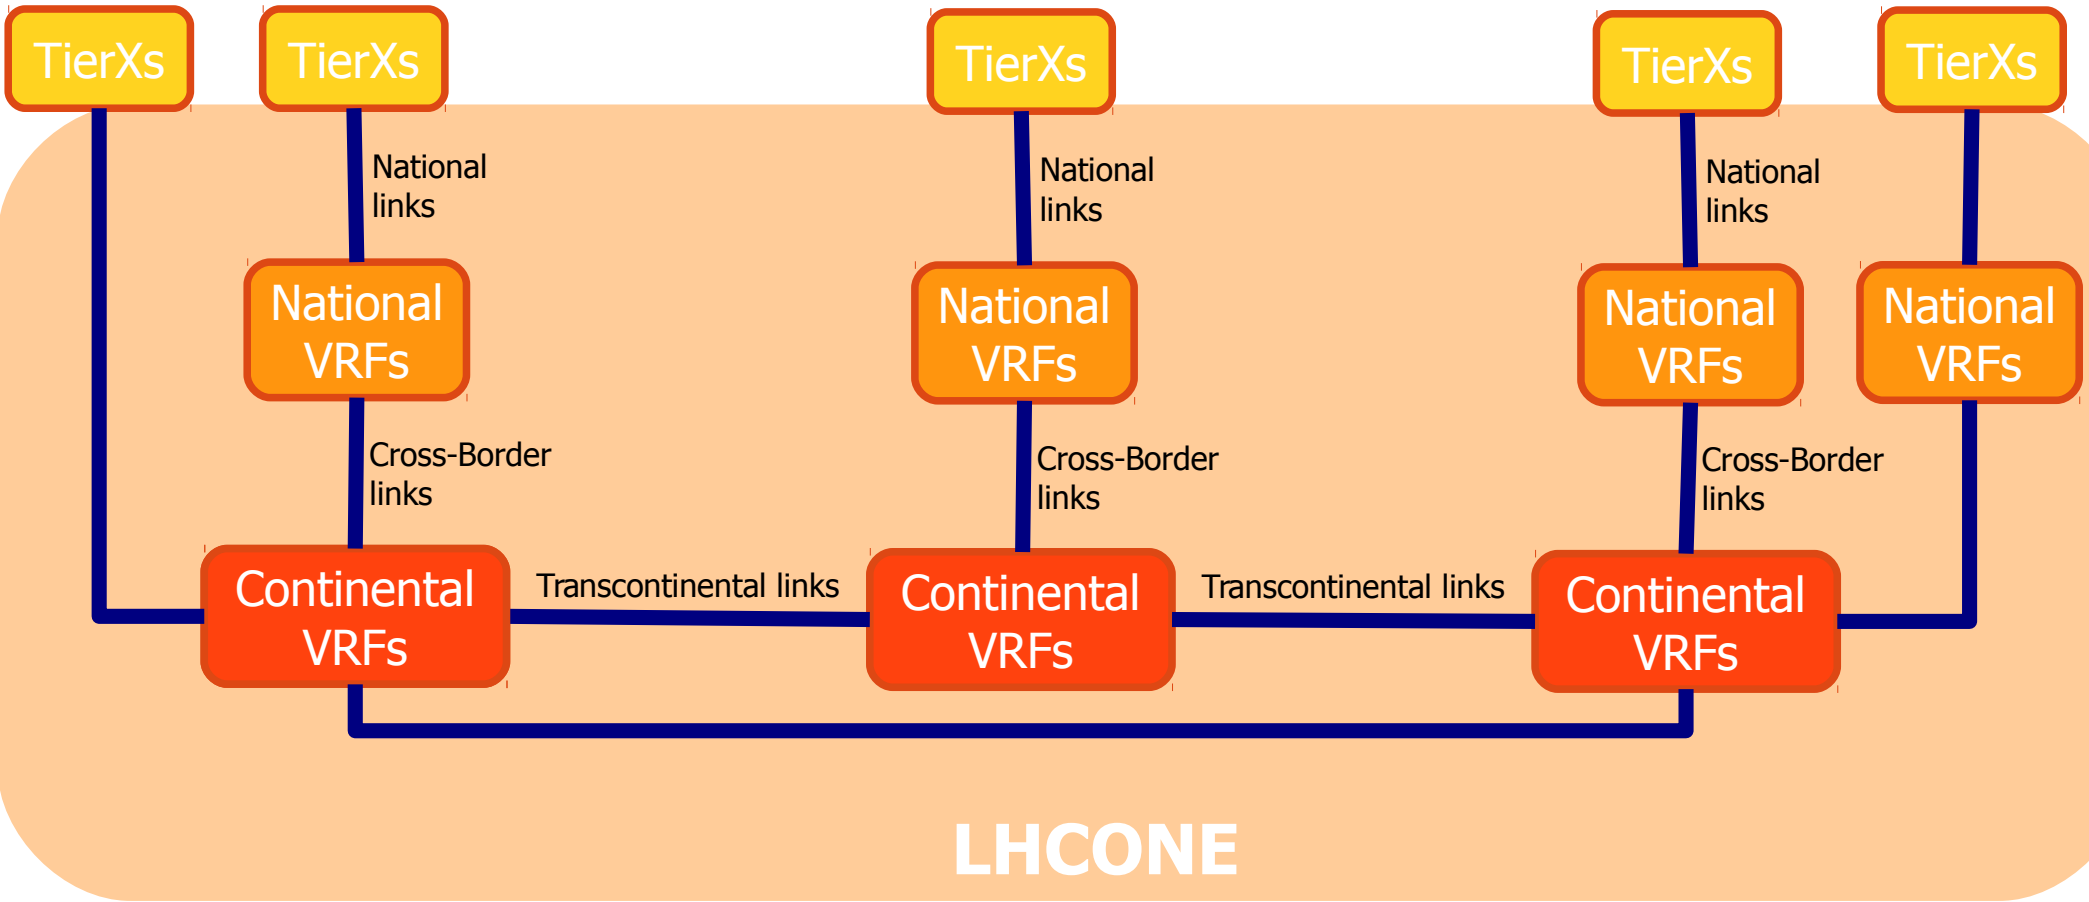

LHCONE P2P service

LHCONE P2P service

On demand point-to-point (P2P) links 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

Status

Work in progress: still in the prototyping phase

Challenges:

- multi-domain provisioning system
- intra-TierX connectivity
- TierX-TierY routing
- integration with WLCG software
- bandwidth allocation and protection

LHCONE perfSONAR service

perfSONAR

- framework for active and passive network probing

- developed by Internet2, Esnet, Geant and others

perfSONAR

http://www.perfsonar.net/

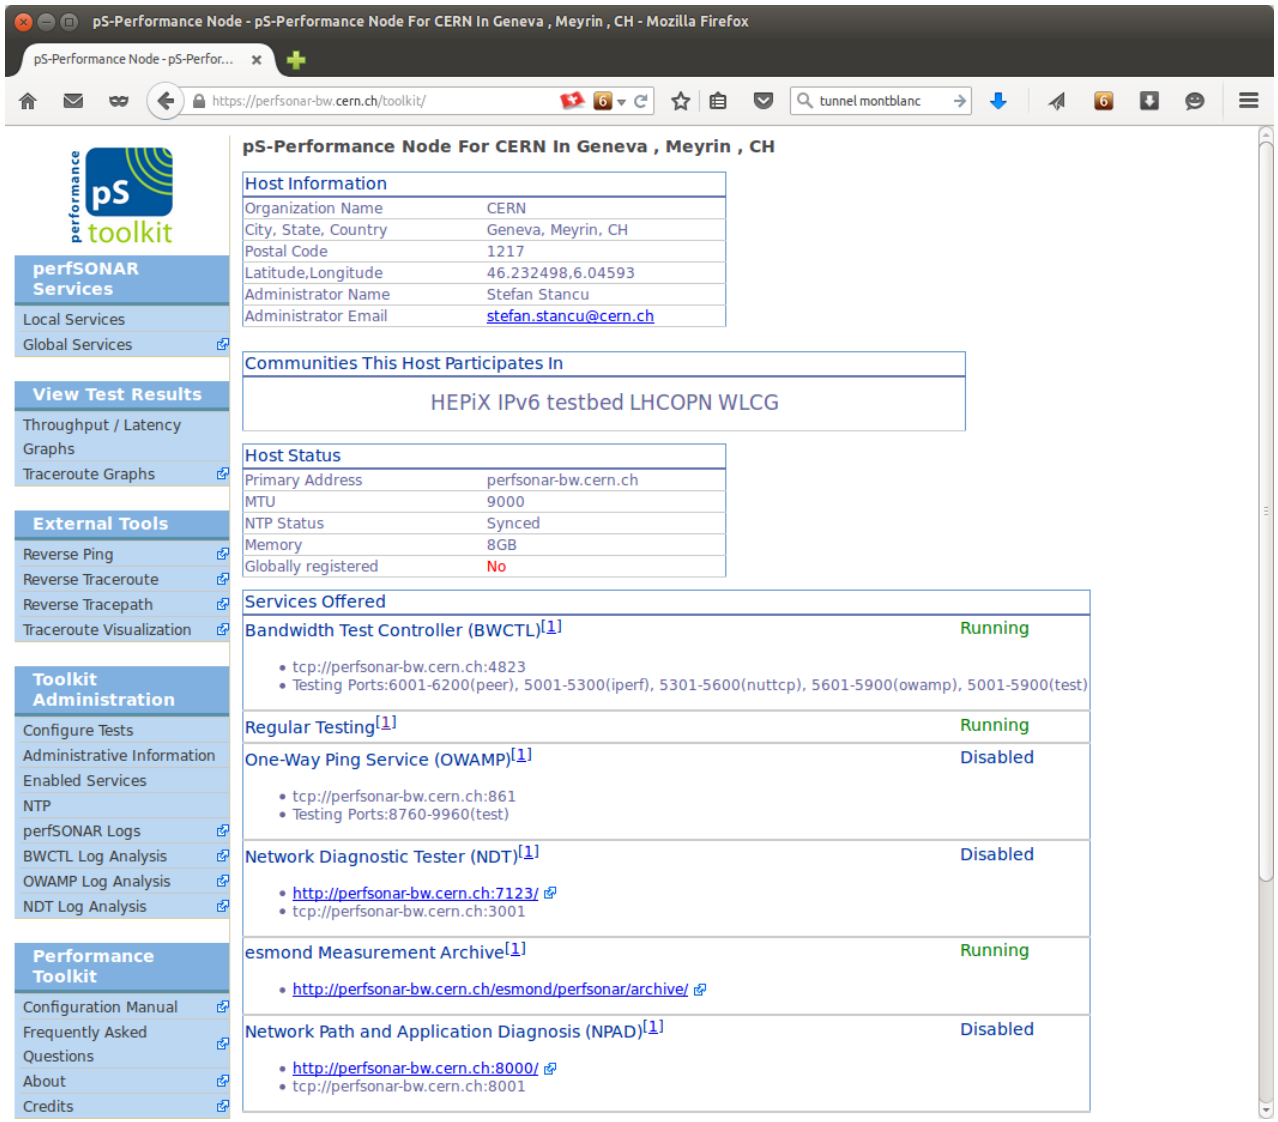

LHCONE perfSONAR service

LHCONE Network monitoring infrastructure

Probes installed at:
- VRFs interconnecting points
- Sites

Accessible to any Site for network healthiness checks

LHCONE perfSONAR MaDDash

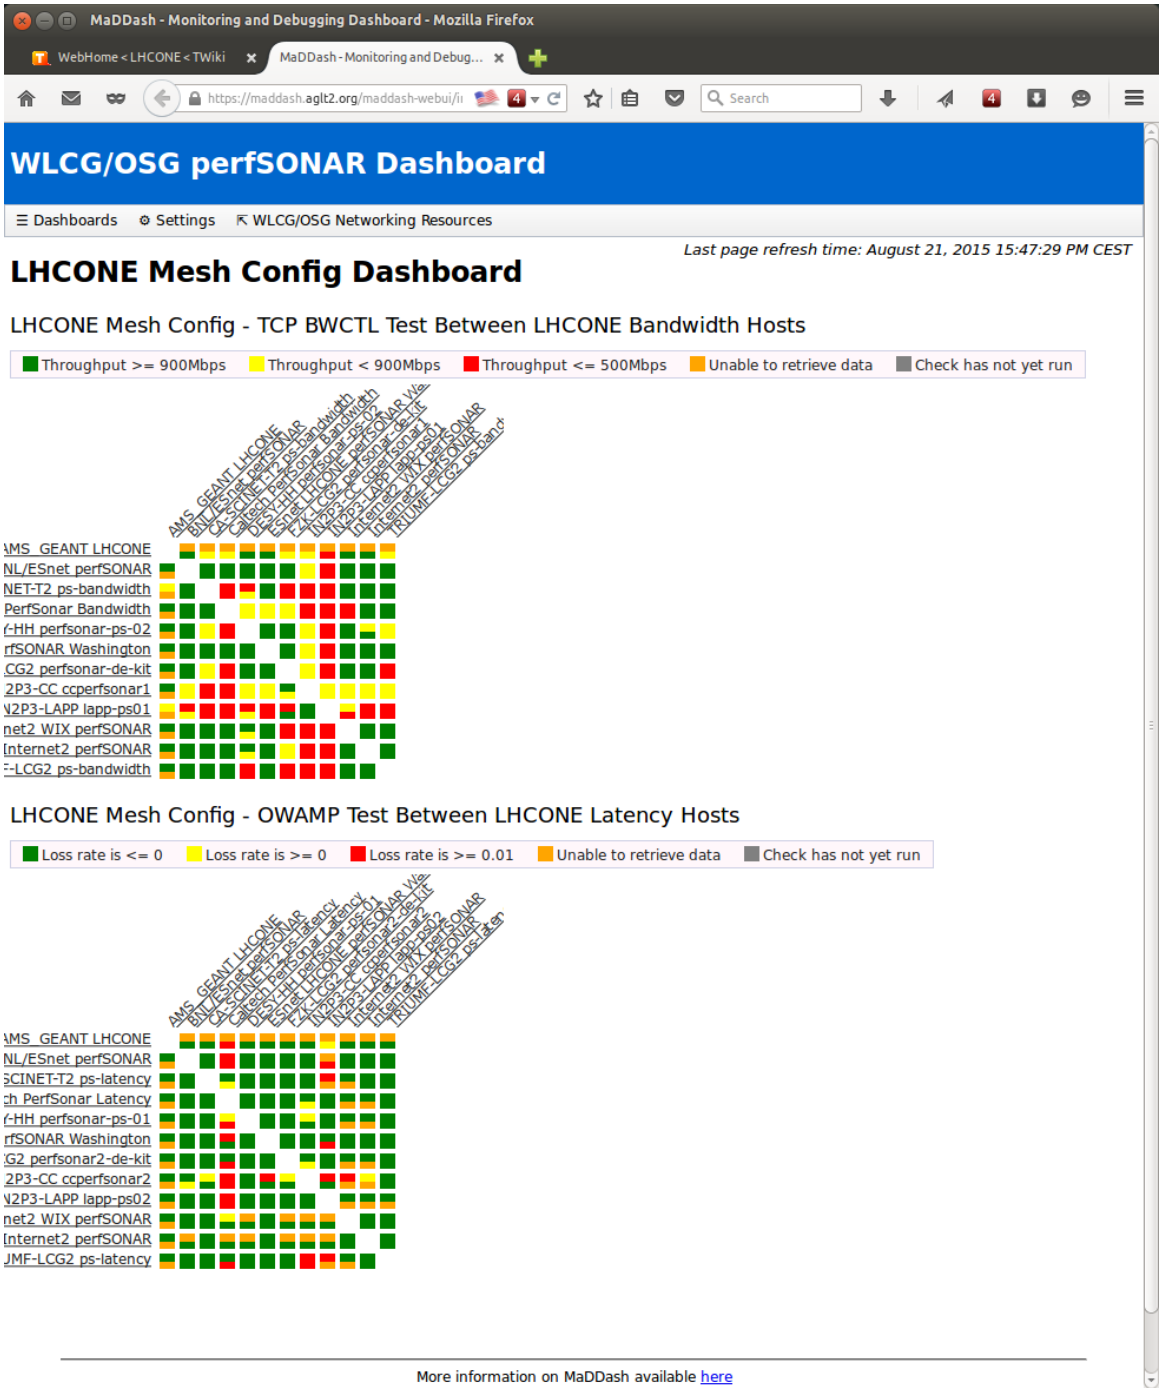

More information

Deployment:
https://twiki.opensciencegrid.org/bin/view/Documentation/DeployperfSONAR

MaDDash
https://maddash.aglt2.org/maddash-webui/index.cgi?dashboard=LHCONE%20Mesh%20Config

perfSONAR
http://www.perfsonar.net/

LHCONE Status

Status

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, Esnet, GEANT, NORDUNET, SURFnet

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

More Information:
- https://indico.cern.ch/event/318811/contribution/5/material/slides/0.pdf
- https://twiki.cern.ch/twiki/bin/view/LHCONE/LhcOneVRF

Current L3VPN topology

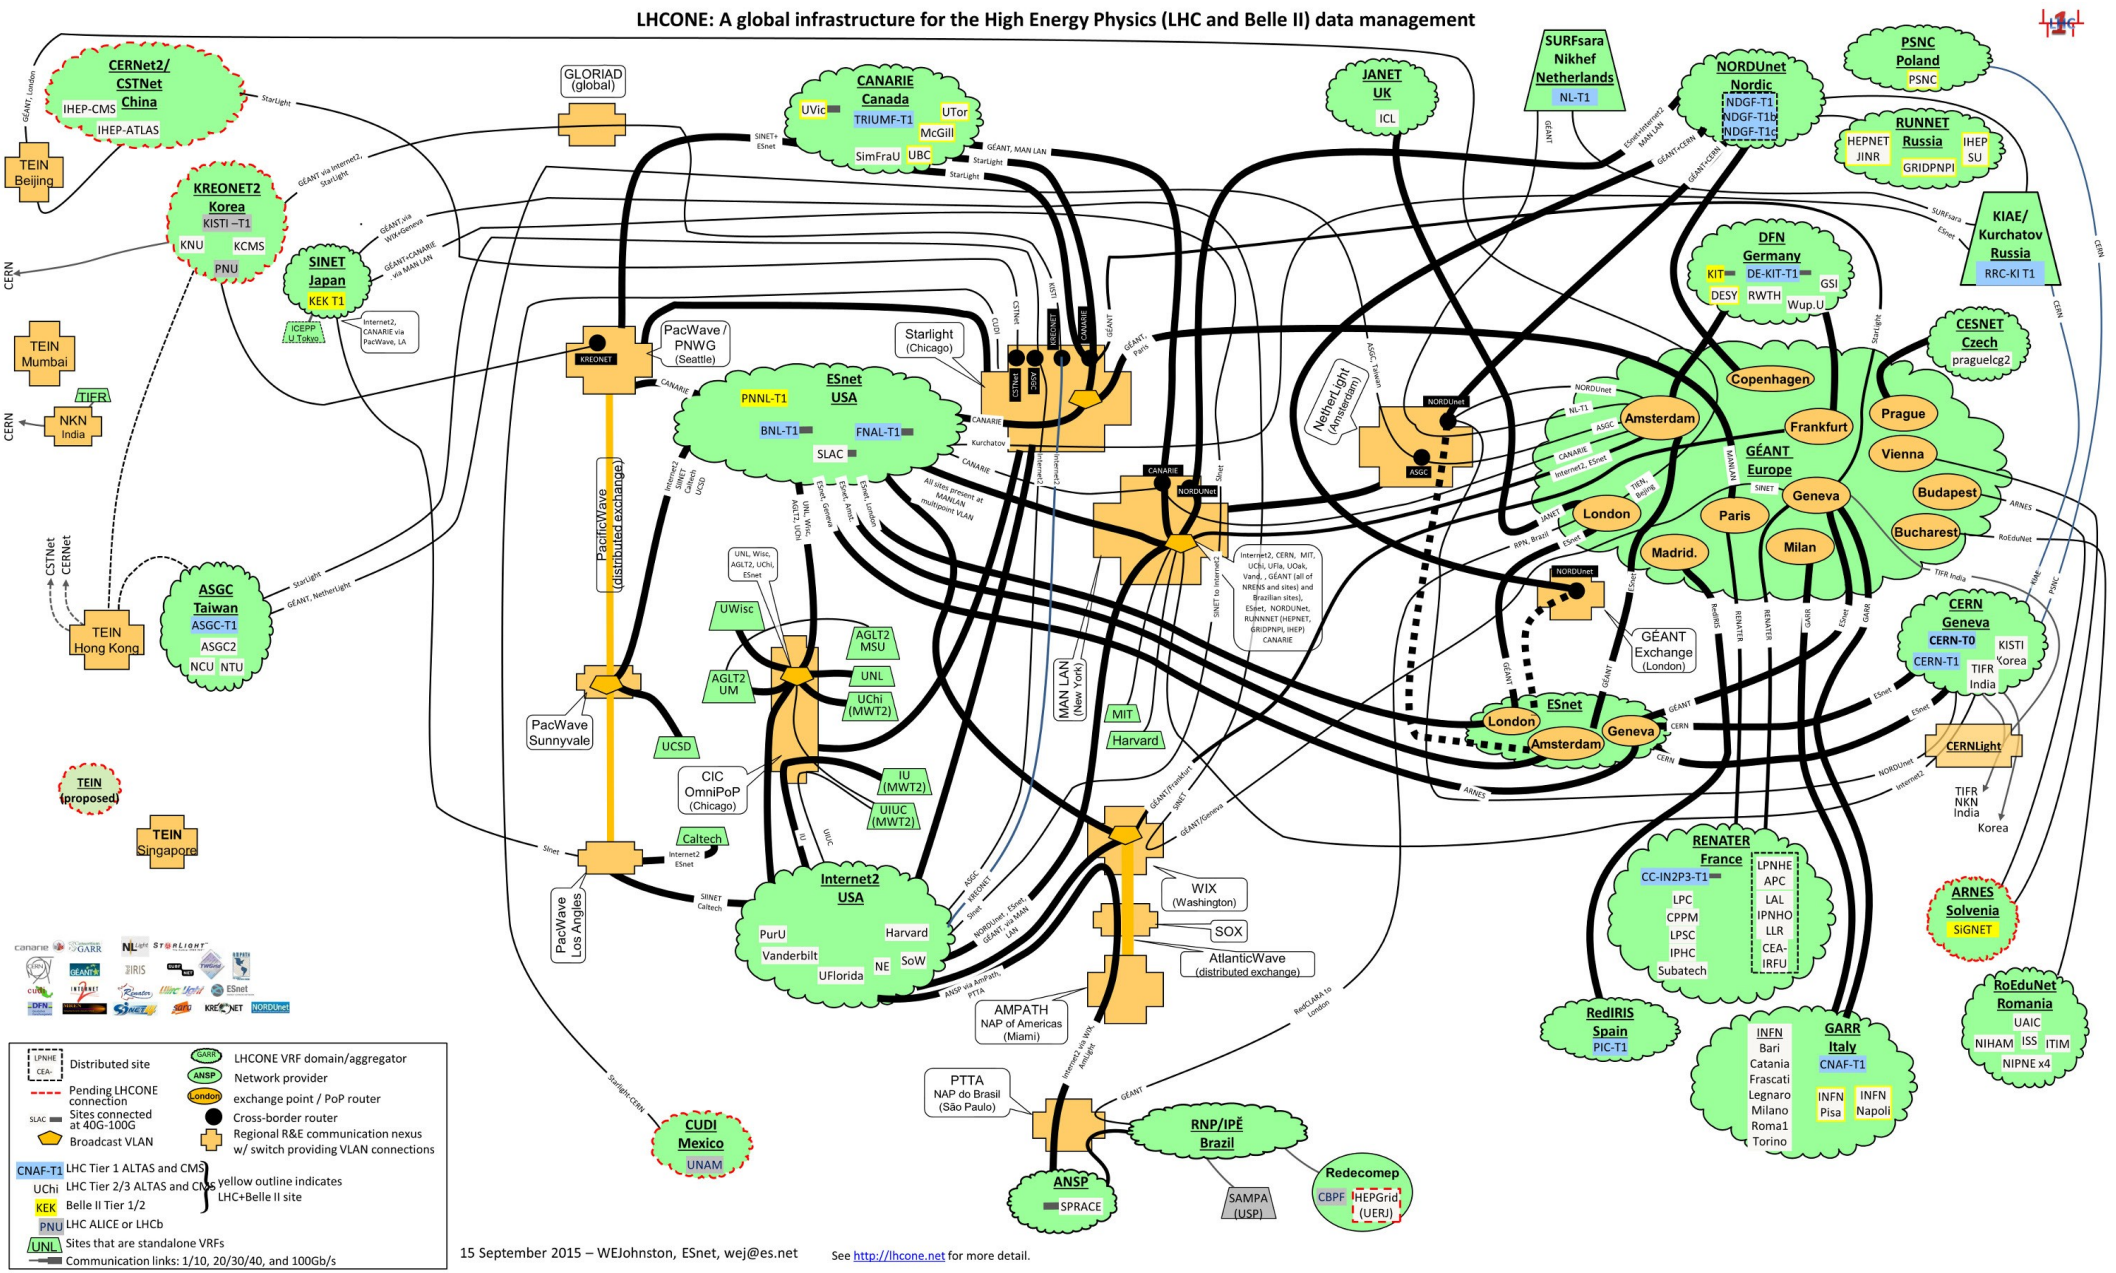

credits: Bill Johnston, ESnet

Geant VRF traffic

Weekly

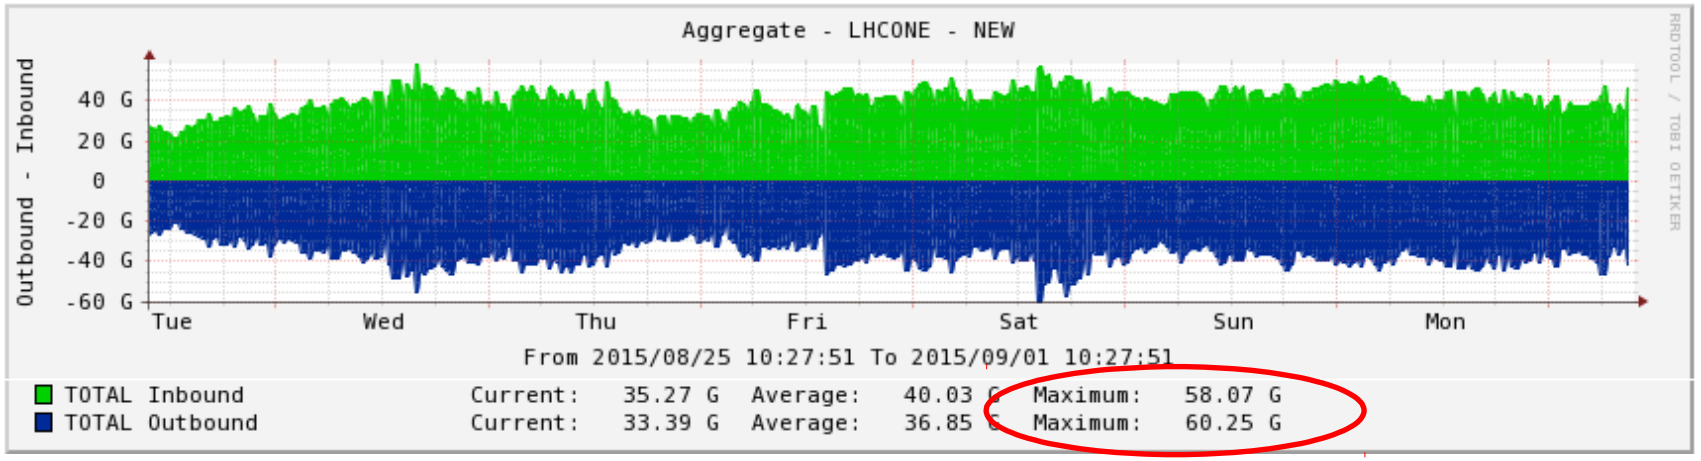

Yearly

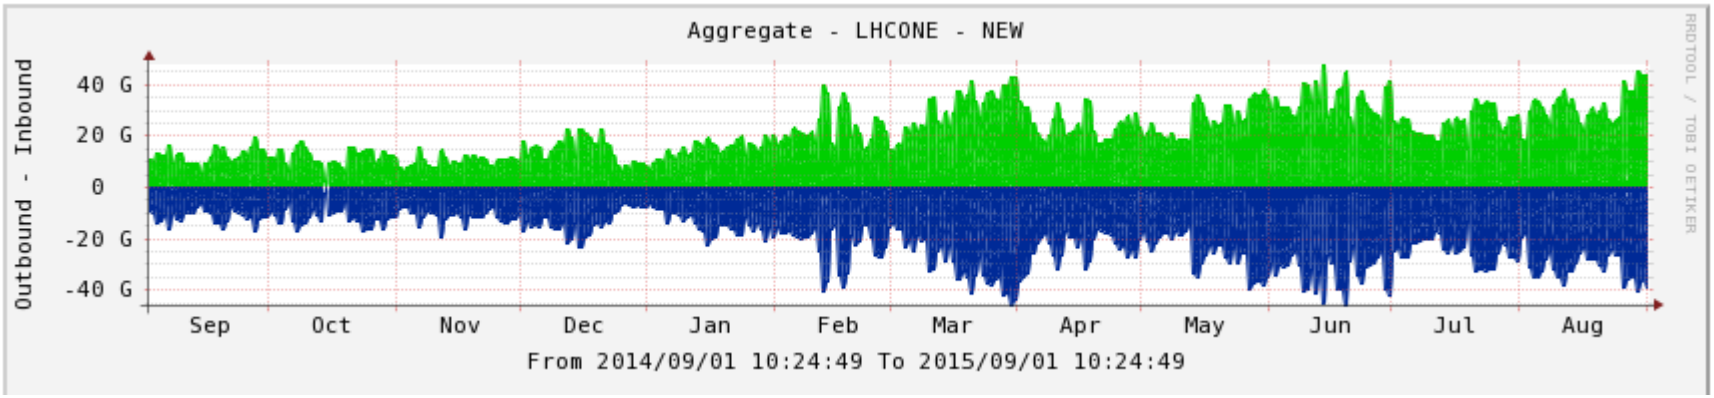

https://tools.geant.net/portal/links/cacti/graph_view.php?action=tree&tree_id=30&leaf_id=6688 (login needed)

Internet2 VRF traffic

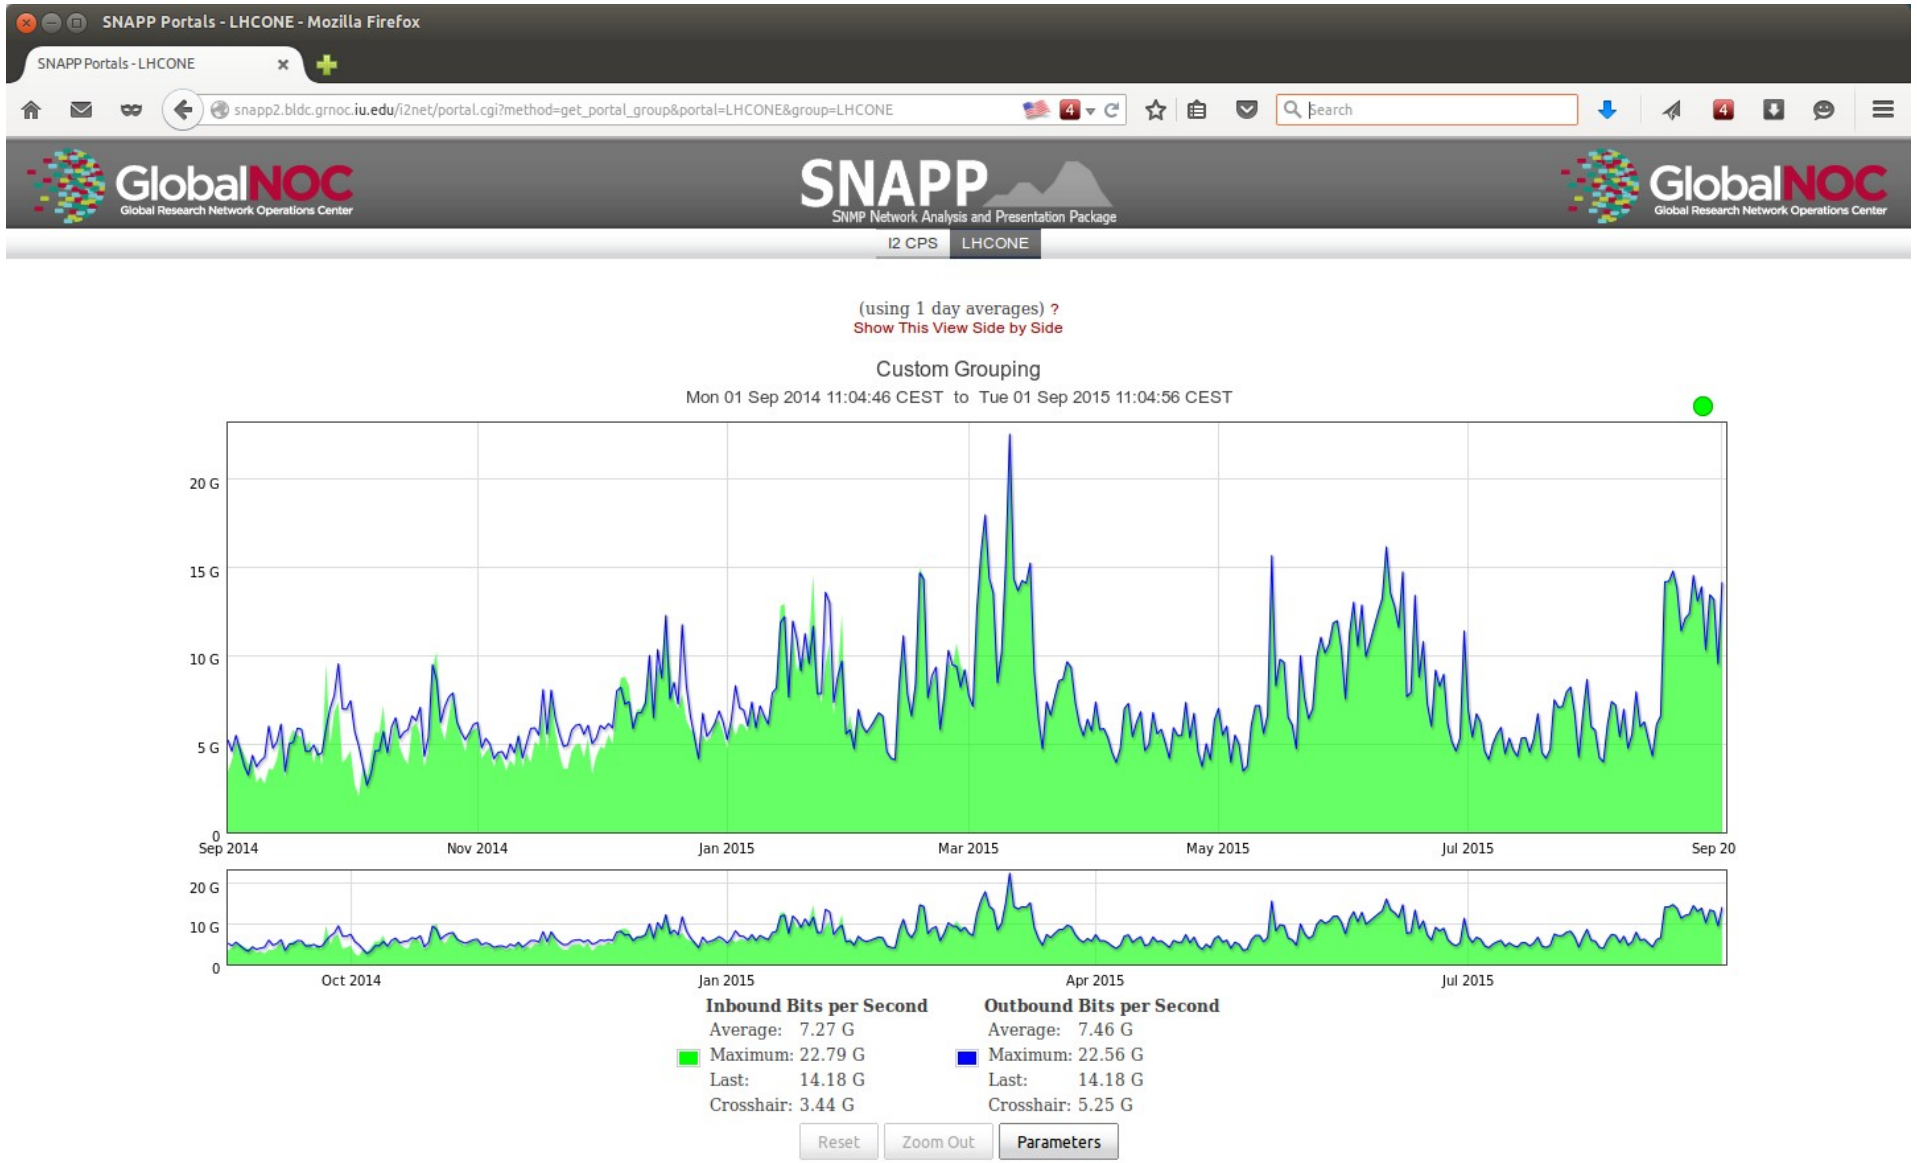

http://snapp2.bldc.grnoc.iu.edu/i2net/portal.cgi?method=get_portal_group&portal=LHCONE&group=LHCONE

Tier0 traffic

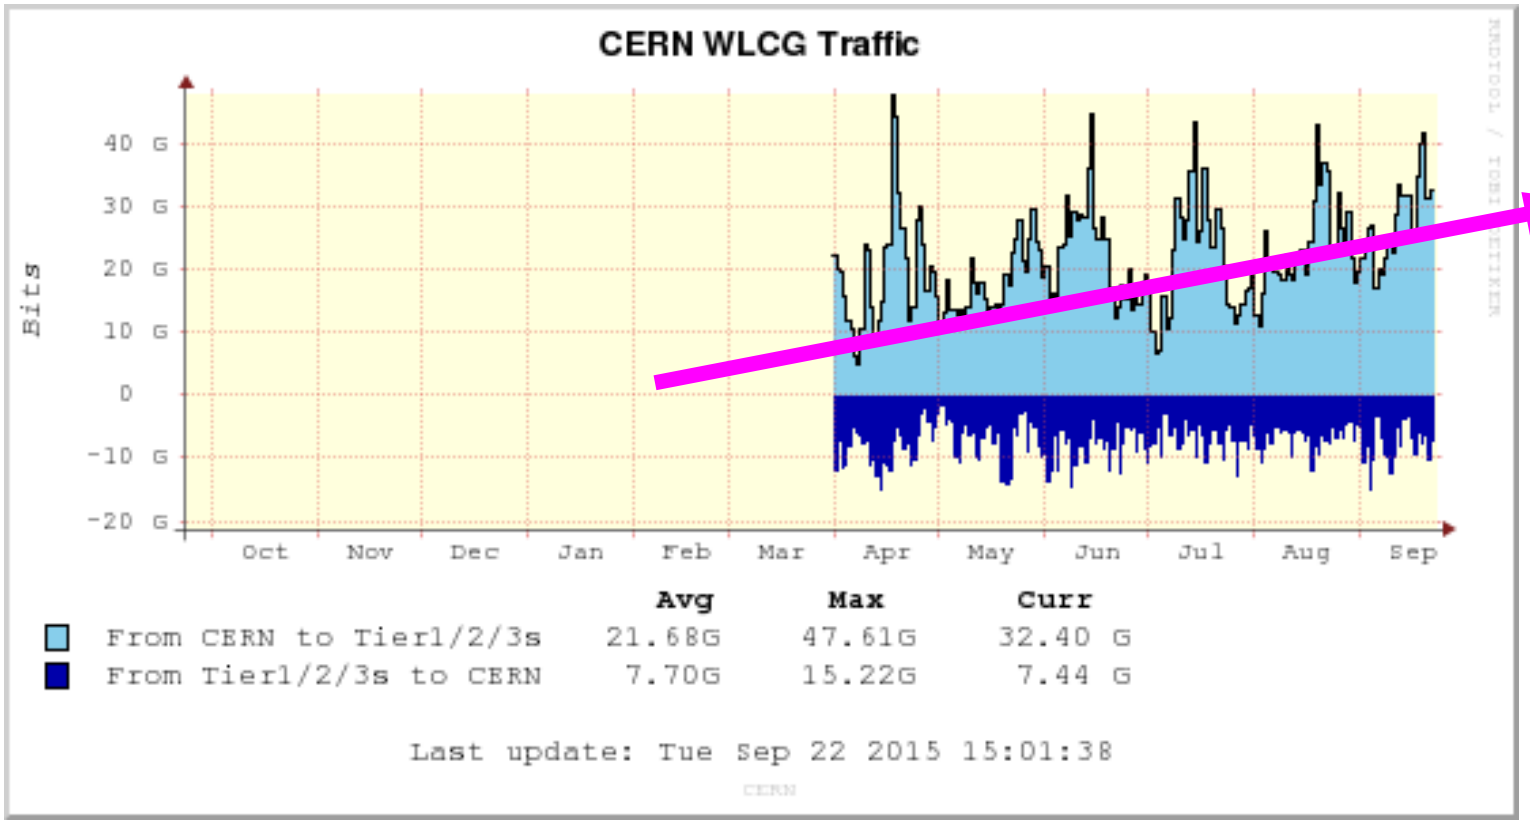

Open to other HEP collaborations

L3VPN has been recently opened to:

Belle II experiment

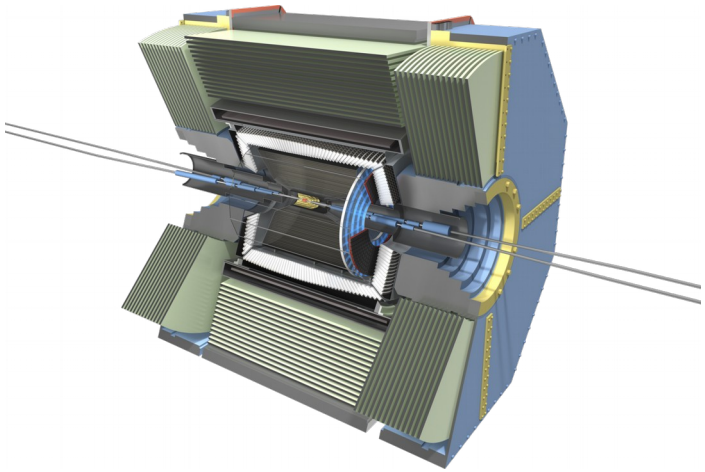

Pierre Auger Observatory

Acceptable Use Policy

The LHCONE 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)

Being currently audited

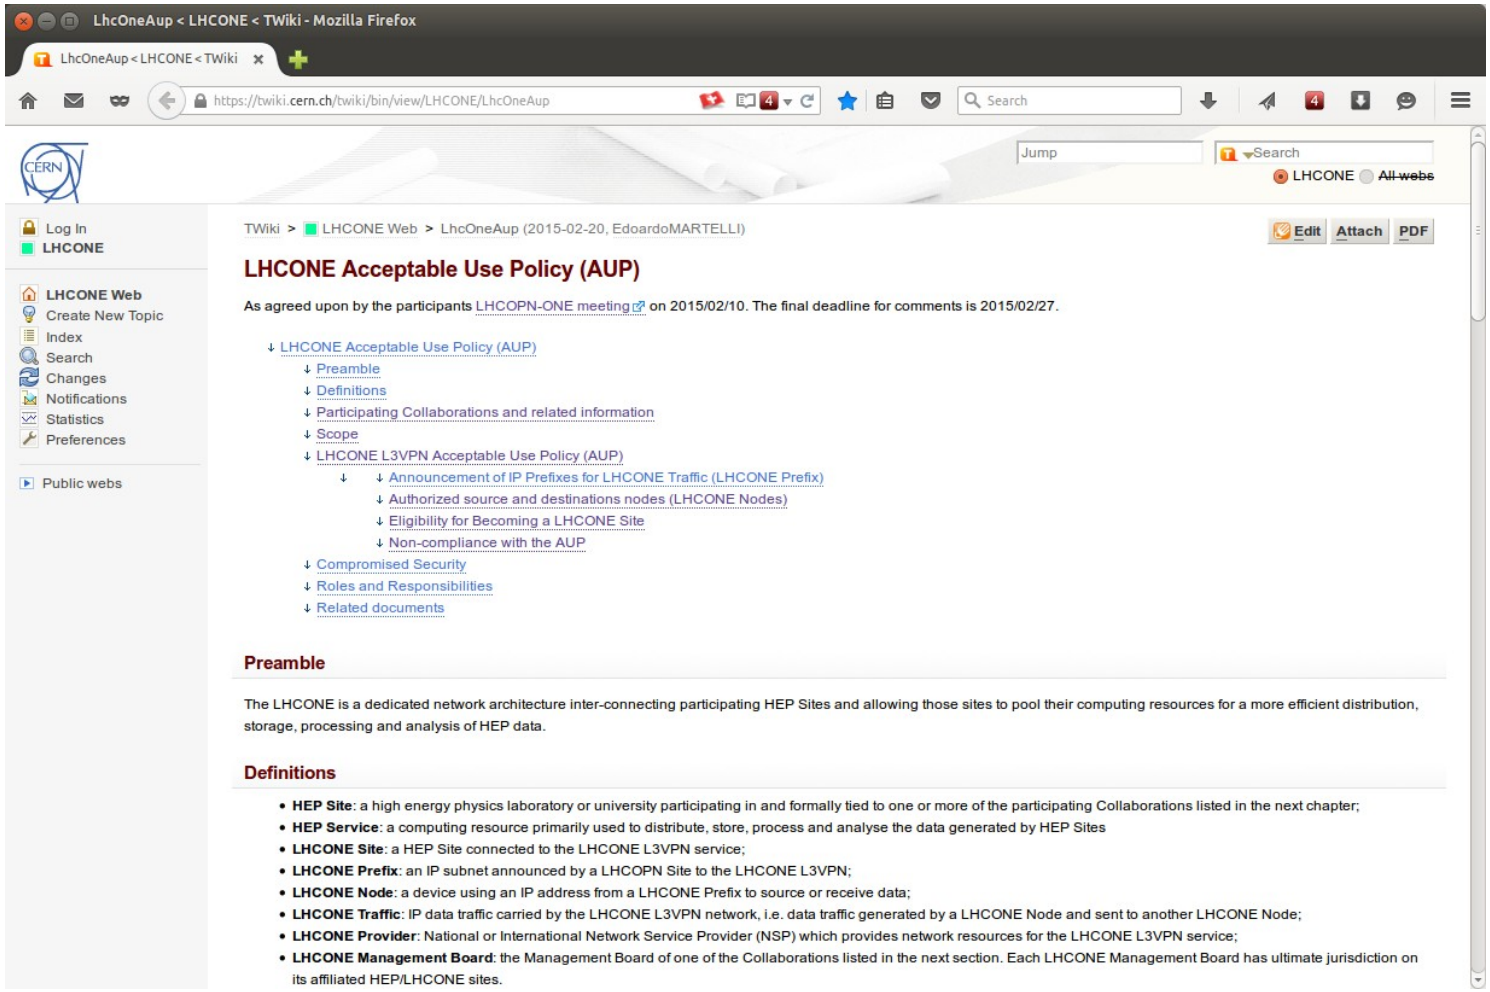

More information

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

Join the lhcone-asia-pacific@cern.ch

If you have a CERN account:
https://e-groups.cern.ch/e-groups/Egroup.do?egroupId=10143913

Otherwise ask e.m@cern.ch

Questions?


CERN