LHC networking update

3rd Asia Tier Centre Forum
Daejeon, South Korea

12th October 2017

edoardo.martelli@cern.ch

Content

- LHCOPN
- LHCONE
- Future developments

LHCOPN update

LHCOPN

Private network connecting Tier0 and Tier1s

- Dedicated to LHC data transfers and analysis
- Secured: only declared IP prefixes can exchange traffic
- Advanced routing: communities for traffic engineering, load balancing.

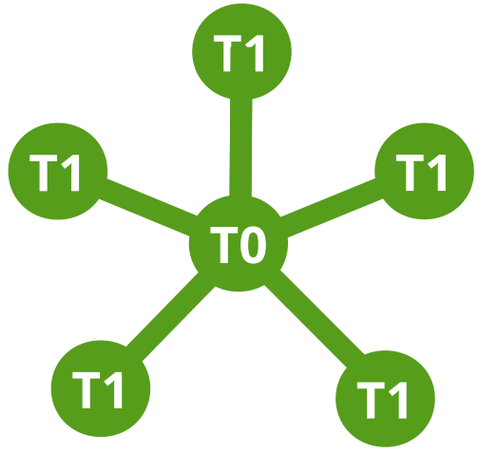

LHCOPN

LHCOPN

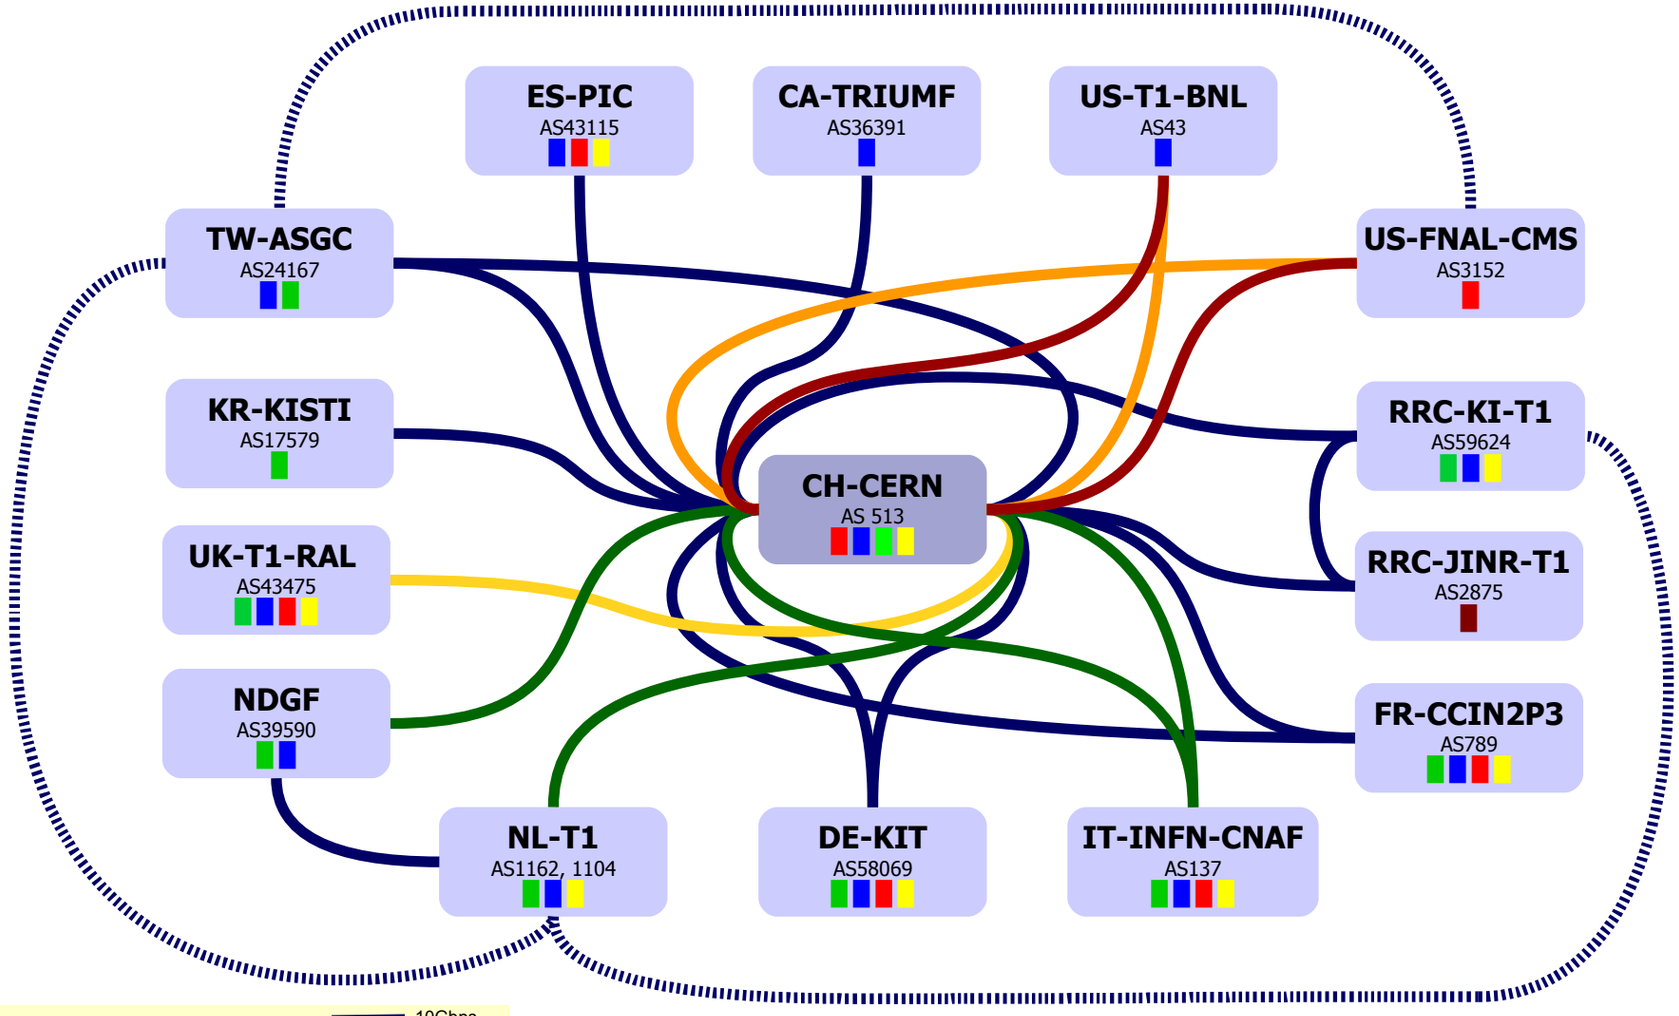

https://twiki.cern.ch/twiki/bin/view/LHCOPN/OverallNetworkMaps

Numbers

- 14 Tier1s + 1 Tier0
- 12 countries in 3 continents
- Dual stack IPv4-IPv6
- 410 Gbps to the Tier0
- Moved ~160 PB in the last year

Tier1s' latest news

CA-TRIUMF: is moving to SFU; TRIUMF dismantled by 2020

UK-T1-RAL: added 3rd 10G link; IPv6 configured

NL-T1: will provision more bandwidth on backup path via NORDUnet

IT-INFN-CNAF: new core network devices (Cisco)

FR-CCIN2P3: new core network devices (Cisco)

US-T1-BNL: implemented IPv6

Tier0's latest news

CH-CERN

Tender for datacentre routers and switches: selected Juniper QFX

2nd network hub in Prevessin (FR) site:

- construction completed
- first two routers installed
- preparing for REN installations (GEANT, ESnet..)

Withdrawn plan of shared datacentre for Experiments' Data Acquisitions in Prevessin (FR)

Still considering construction of datcentre extension in Prevessin (FR)

EOS storage dual stack IPv4/IPv6 (LHCb, ALICE)

LHCOPN

Latest developments

Traffic volume constantly growing
- already increased of 200% since the beginning of Run2

Tier1s are considering to upgrade links to 100Gbs
- already done for US sites
- waiting deployment of new network router at CERN for affordable 100G interfaces

Almost fully dual-stack IPv4-IPv6
- 13 Tier1s and the Tier0 peering over IPv6
- dual-stack perfSONAR installed in all of them

LHCOPN

LHC performances in 2017

2017 operations affected by beam instabilities. Suspected some air got trapped in vacuum chambers during last Long Shutdown

Target integrated luminosity for 2017 is 41 fb^{-1}, despite problems and shorter running time in 2017 due to EYETS (last year it achieved 40 fb^{-1})

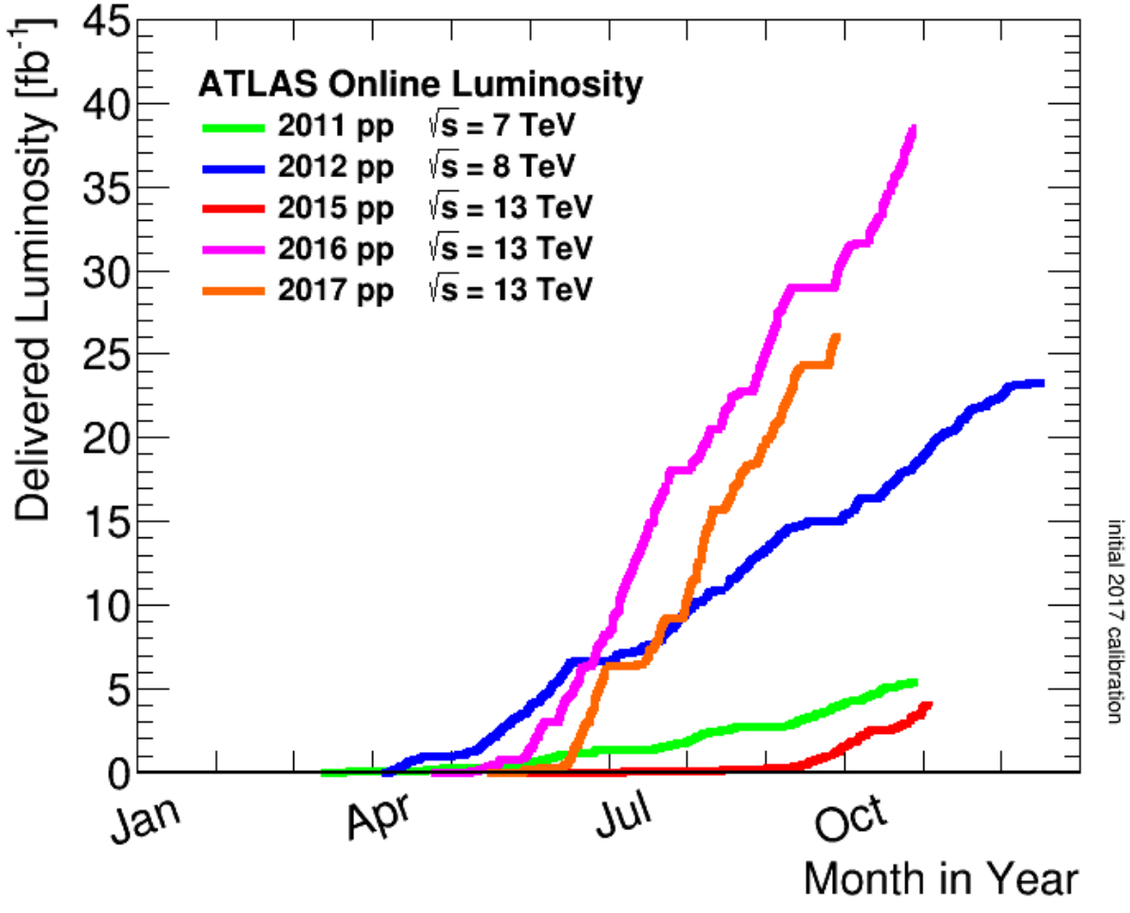

LHCOPN

Last 12 months traffic statistics

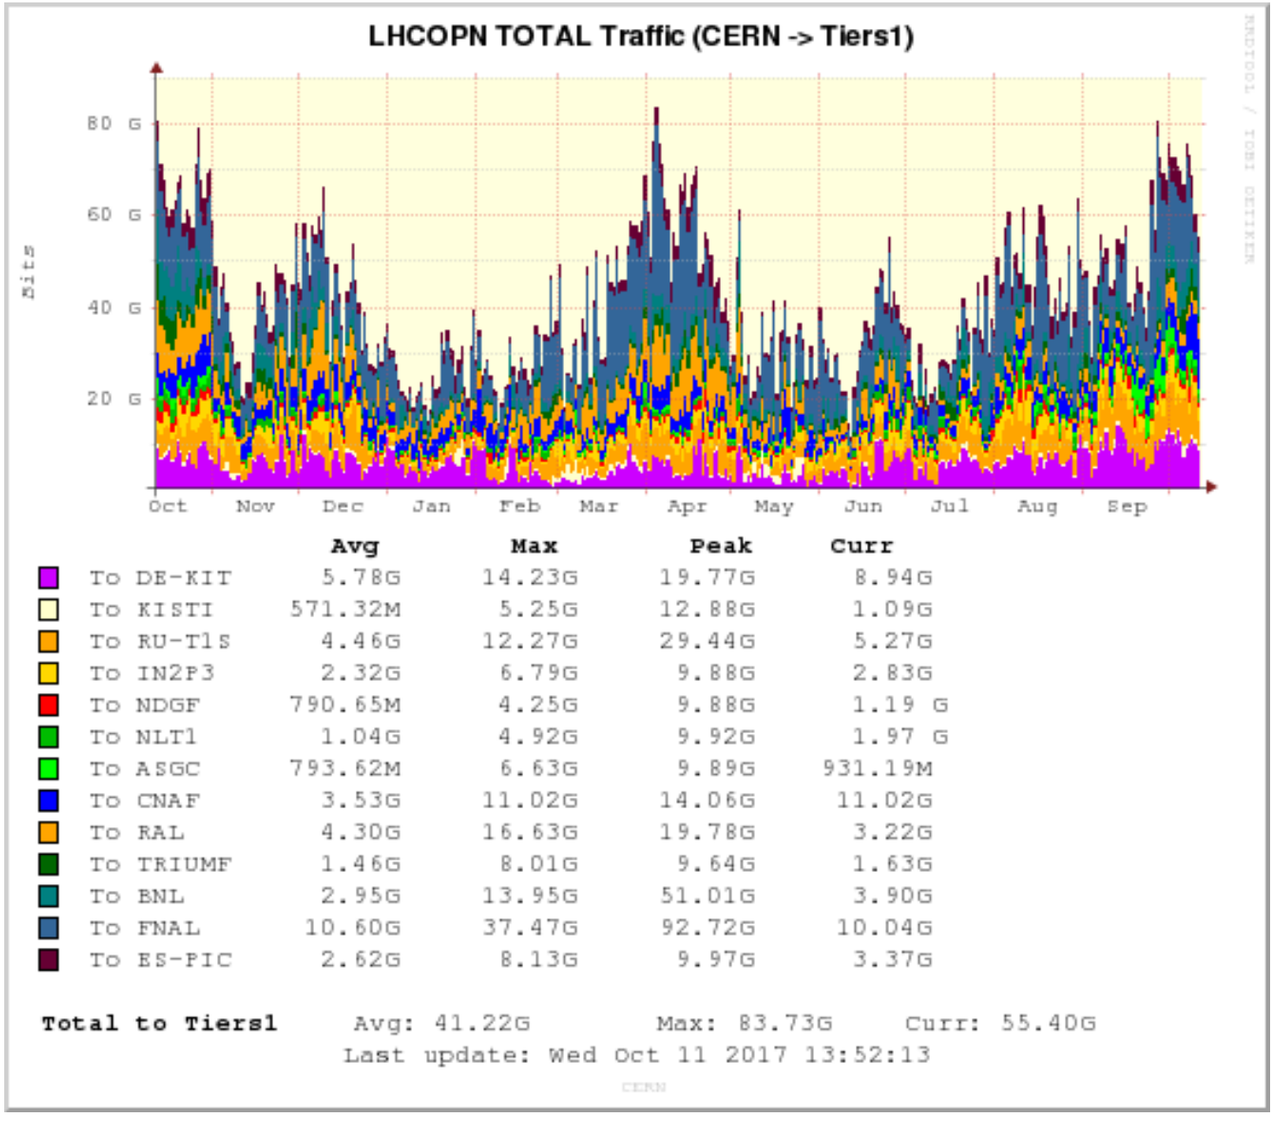

No major increase compared to previous year

LHCOPN

Plans

Deploy more 100G links
- now more cost effective than multiple 10G (at least in Europe and North America)

Complete IPv6 deployment:
- connect all Tier1s
- use IPv6 for production data transfers

Considering allowing BelleII traffic
- will be discussed next week at LHCOPN/ONE meeting

LHCOPN

LHCONE update

LHCONE

- Network serving HEP sites according to their needs and allowing them to grow
- Sharing the cost of expensive resources for common benefit
- Traffic separation: no clash with other data transfers, resource allocated for and funded by the HEP community
- Trusted peers: common security policies

Open to other HEP collaborations

LHCONE L3VPN service

Layer3 (routed) Virtual Private Network

Dedicated worldwide backbone connecting **Tier1s, T2s and T3s** at high bandwidth

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

Trusted traffic that can bypass slow perimeter firewalls

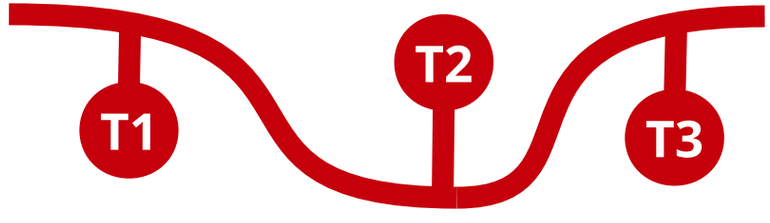

L3VPN Current topology

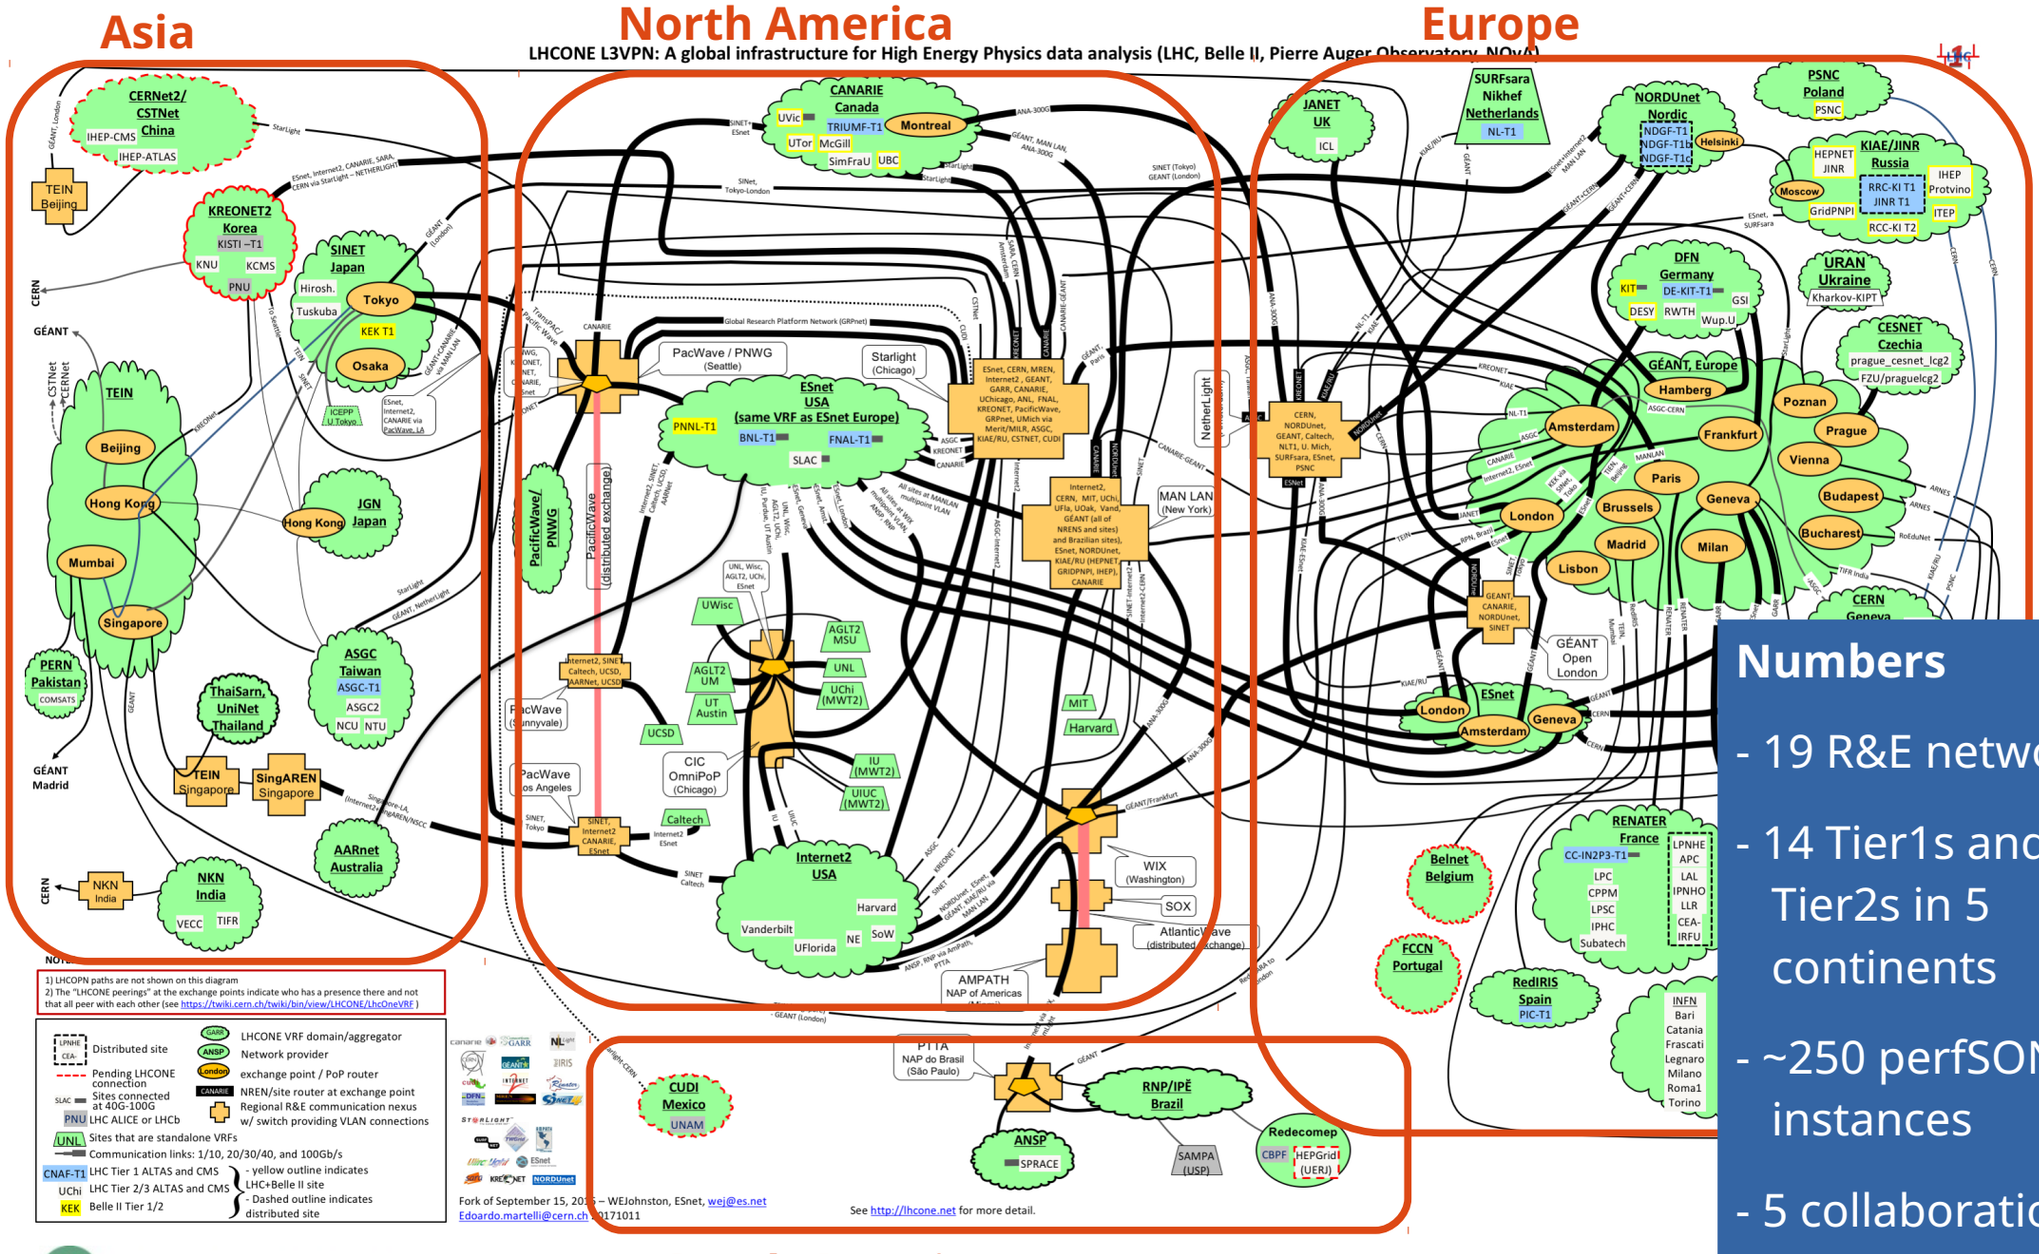

Latest changes

Recently connected:

- Russian VRF
- Belgium VRF
- Greek VRF
- Ukraine VRF
- Japan JGN-X VRF
- Korean VRF

Asia-Pacific VRFs

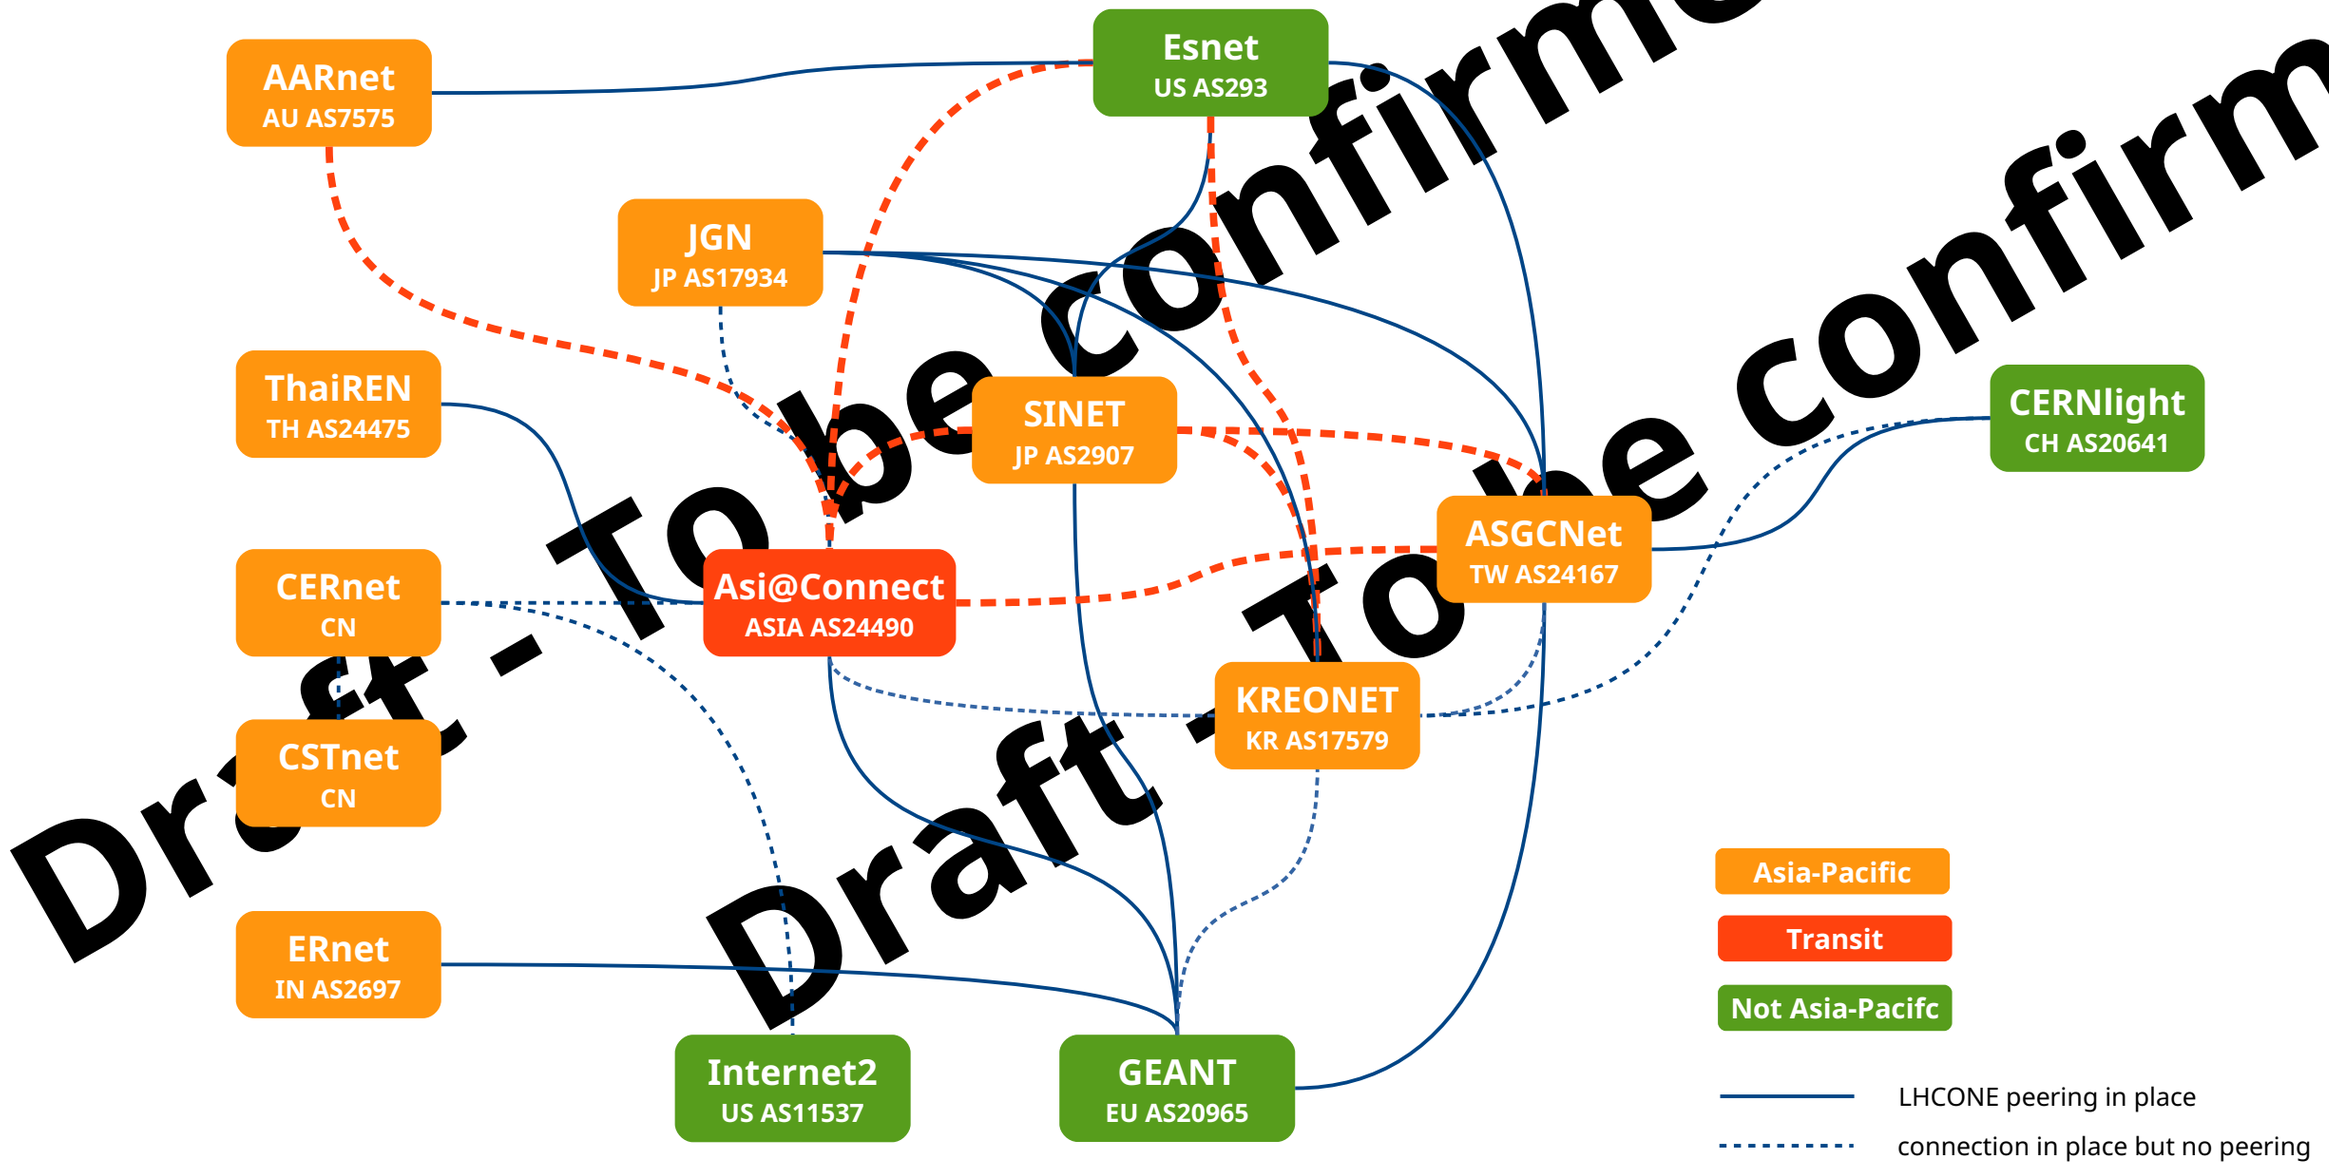

LHCONE perfSONAR service

- LHCONE Network monitoring infrastructure based on perfSONAR monitoring suite

- perfSONAR probes installed at the VRFs interconnecting points and at the TierX sites

- Accessible to any TierX to check network healthiness and debug issues

LHCONE perfSONAR: status

~280 perfSONAR instances registered in GOCDB/OIM

~250 Active perfSONAR instances

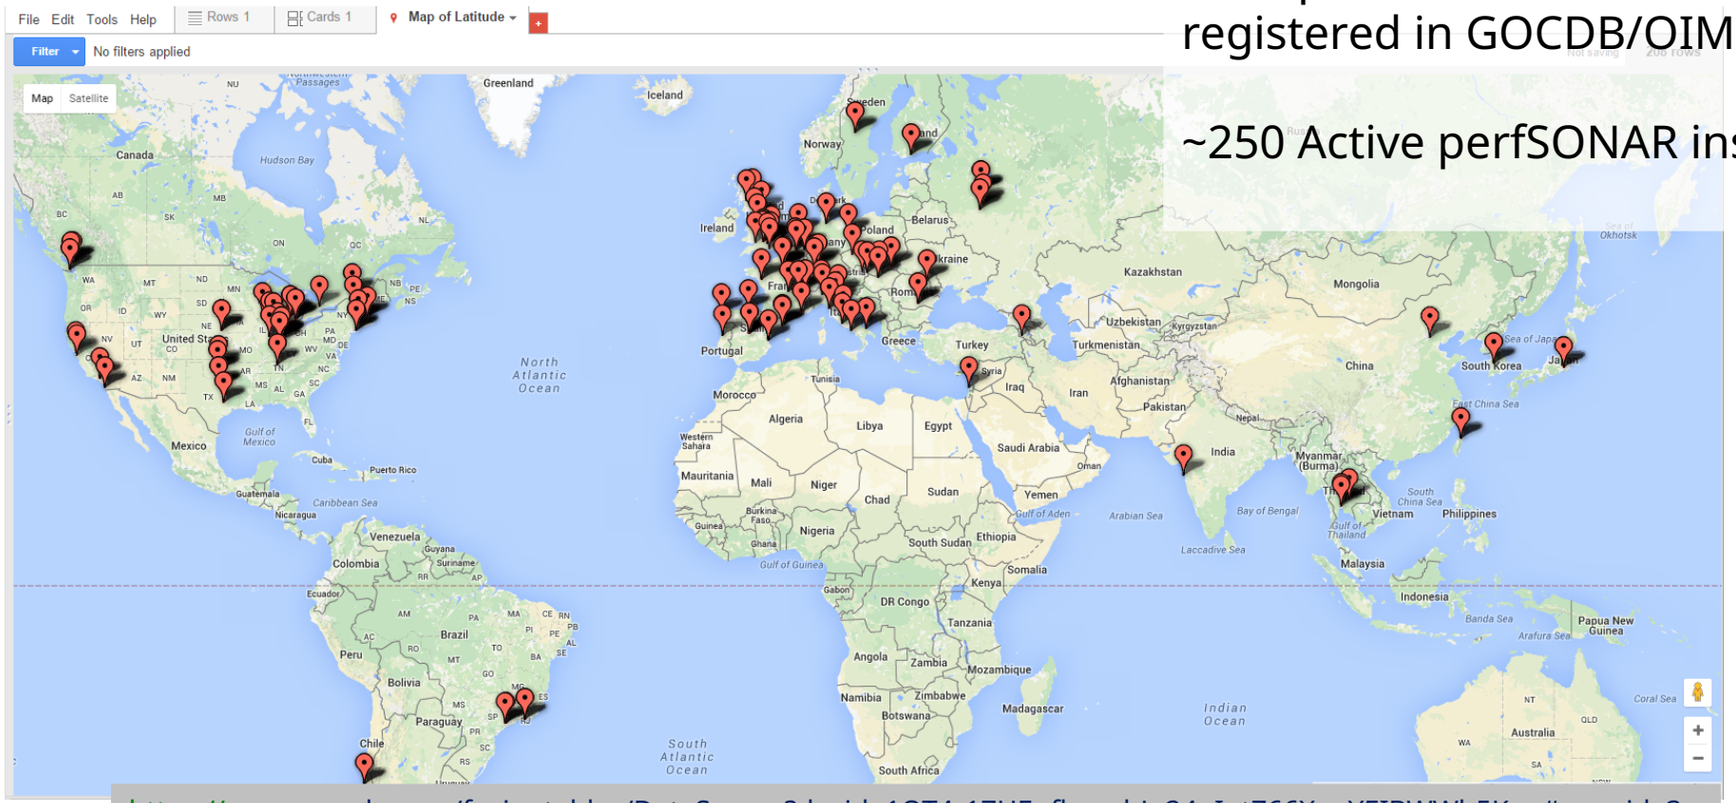

https://www.google.com/fusiontables/DataSource?docid=1QT4r17HEufkvnqhJu24nIptZ66XauYEIBWWh5Kpa#map:id=3

- Initial deployment coordinated by WLCG perfSONAR TF
- Commissioning of the network followed by WLCG Network and Transfer Metrics WG

Slide credit: Shawn McKee, University of Michigan

perfSONAR: gathering and storing metrics

- **OSG is providing network metric data for its members and WLCG via the Network Datastore**
  - The data is gathered from all WLCG/OSG perfSONAR instances
  - Stored indefinitely on OSG hardware
  - Data available via Esmond API
  - In production since September 14th 2015
- **The primary use-cases**
  - Network problem identification and localization
  - Network-related decision support
  - Network baseline: set expectations and identify weak points for upgrading

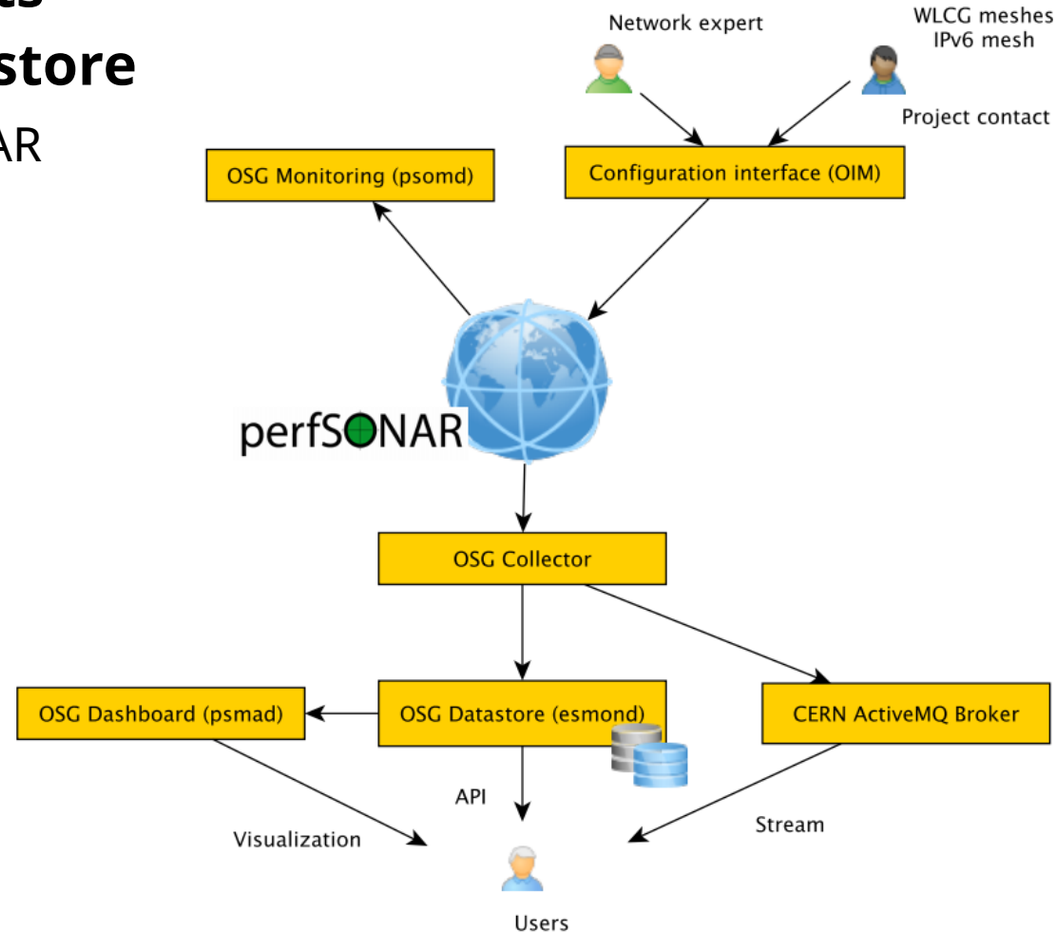

Slide credit: Shawn McKee, University of Michigan

perfSONAR update

Completed MCA (Mesh Configuration Admin) tool

WLCG is working on ETF (Experiment Test Framework) to monitor perfSONAR services

ATLAS is working on getting network metrics into an analytics platform:

- packet loss from perfSONAR
- network stats and flows from ESnet routers
- stats from CERN routers will be added

Looking for more router sources

Future developments

Networks have to follow LHC data growth

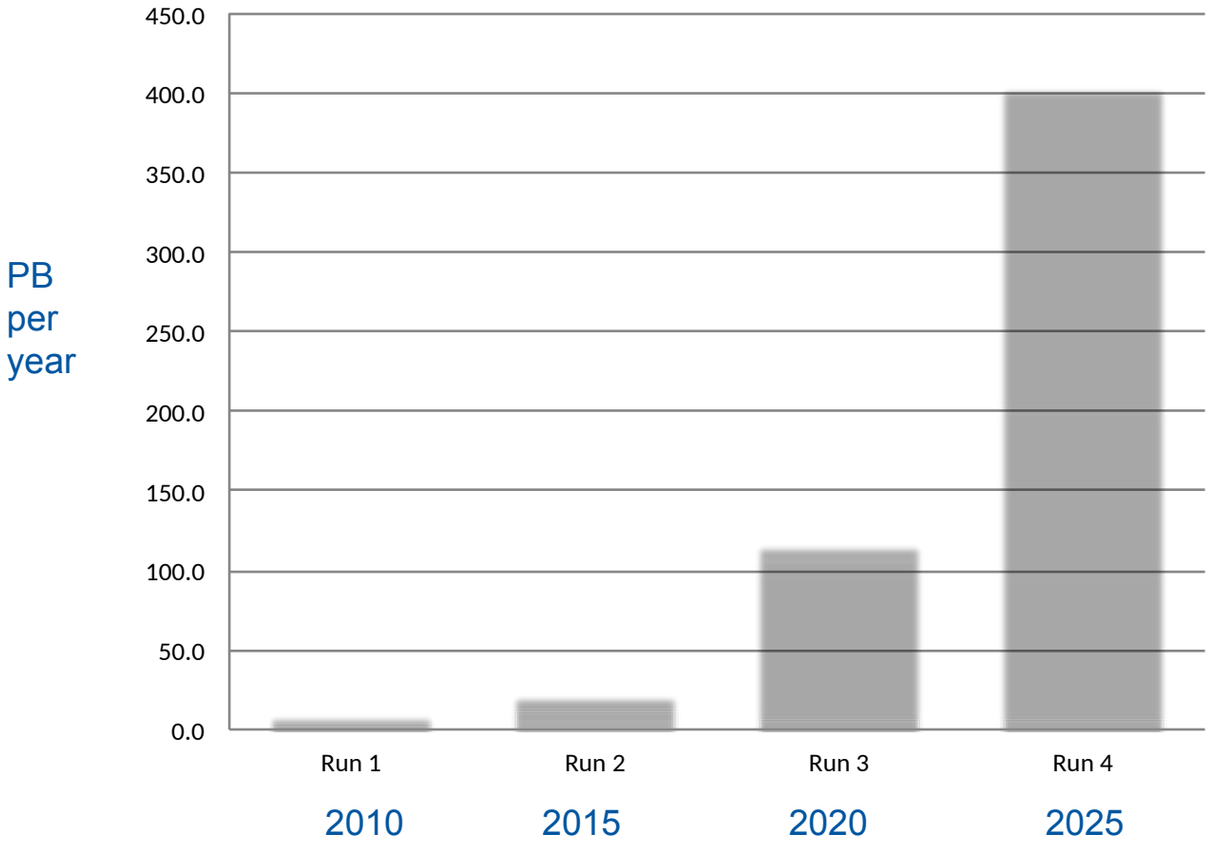

CMS
ATLAS
ALICE
LHCb

LHC data expected to record 400PB/year in Run4

Computing needs expected to be around 50x current levels, if budget available

Networks must grow accordingly

Use of Commercial Cloud Services

Evaluation of use of Commercial Cloud Services for HEP computing have been on-going

Research and Education Networks have designed and deployed solutions to better connect Cloud Service Providers to their customers

Main issues:
- deliver traffic from cloud datacentres to users in different continents
- avoid or not cloud-to-cloud traffic
- not all the research networks allow commercial traffic

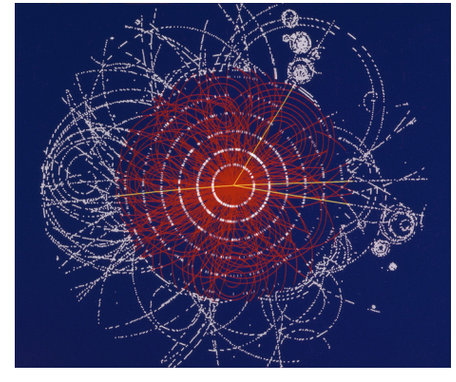

Helix Nebula Science Cloud

HNSciCloud

Objective: procure innovative IaaS level cloud services

Update:

- Prototype phase on going: testing of capabilities
- participants to the Prototype Phase:
    - T-Systems, Huawei, Cyfronet, Divia
    - IBM
    - RHEA Group, T-Systems, exoscale, SixSq
- Two will be selected for the Pilot phase: assessment on performance, scalability and security

LHC experiments' requirements for Run3

pre-GDB meeting on networks held in January 2017. Gathered LHC experiments, WLCG sites and REN operators

Collected requirements for Run3:
- increase network capacity of 5-10 times
- improve monitoring and share statistics of the networks

Acknowledged the usefulness of gathering Experiments, Sites and Network operators all together. Should be repeated every 1-2 years

https://indico.cern.ch/event/571501/

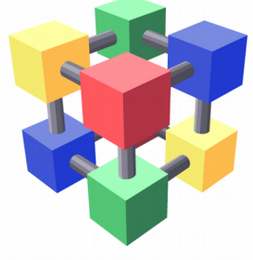

WLCG
Worldwide LHC Computing Grid

Community White Paper

The HEP community is aiming to produce a Community White Paper (CWP) which will describe the community strategy and a roadmap for software and computing R&D in HEP for the 2020s.

The CWP will identify and prioritize the software research and development investments required:

- to achieve improvements in software efficiency, scalability and performance and to make use of the advances in CPU, storage and network technologies
- to enable new approaches to computing and software that could radically extend the physics reach of the detectors
- to ensure the long term sustainability of the software through the lifetime of the HL-LHC

http://hepsoftwarefoundation.org/activities/cwp.html

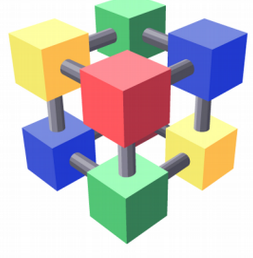

WLCG
Worldwide LHC Computing Grid

Possible change of computing model

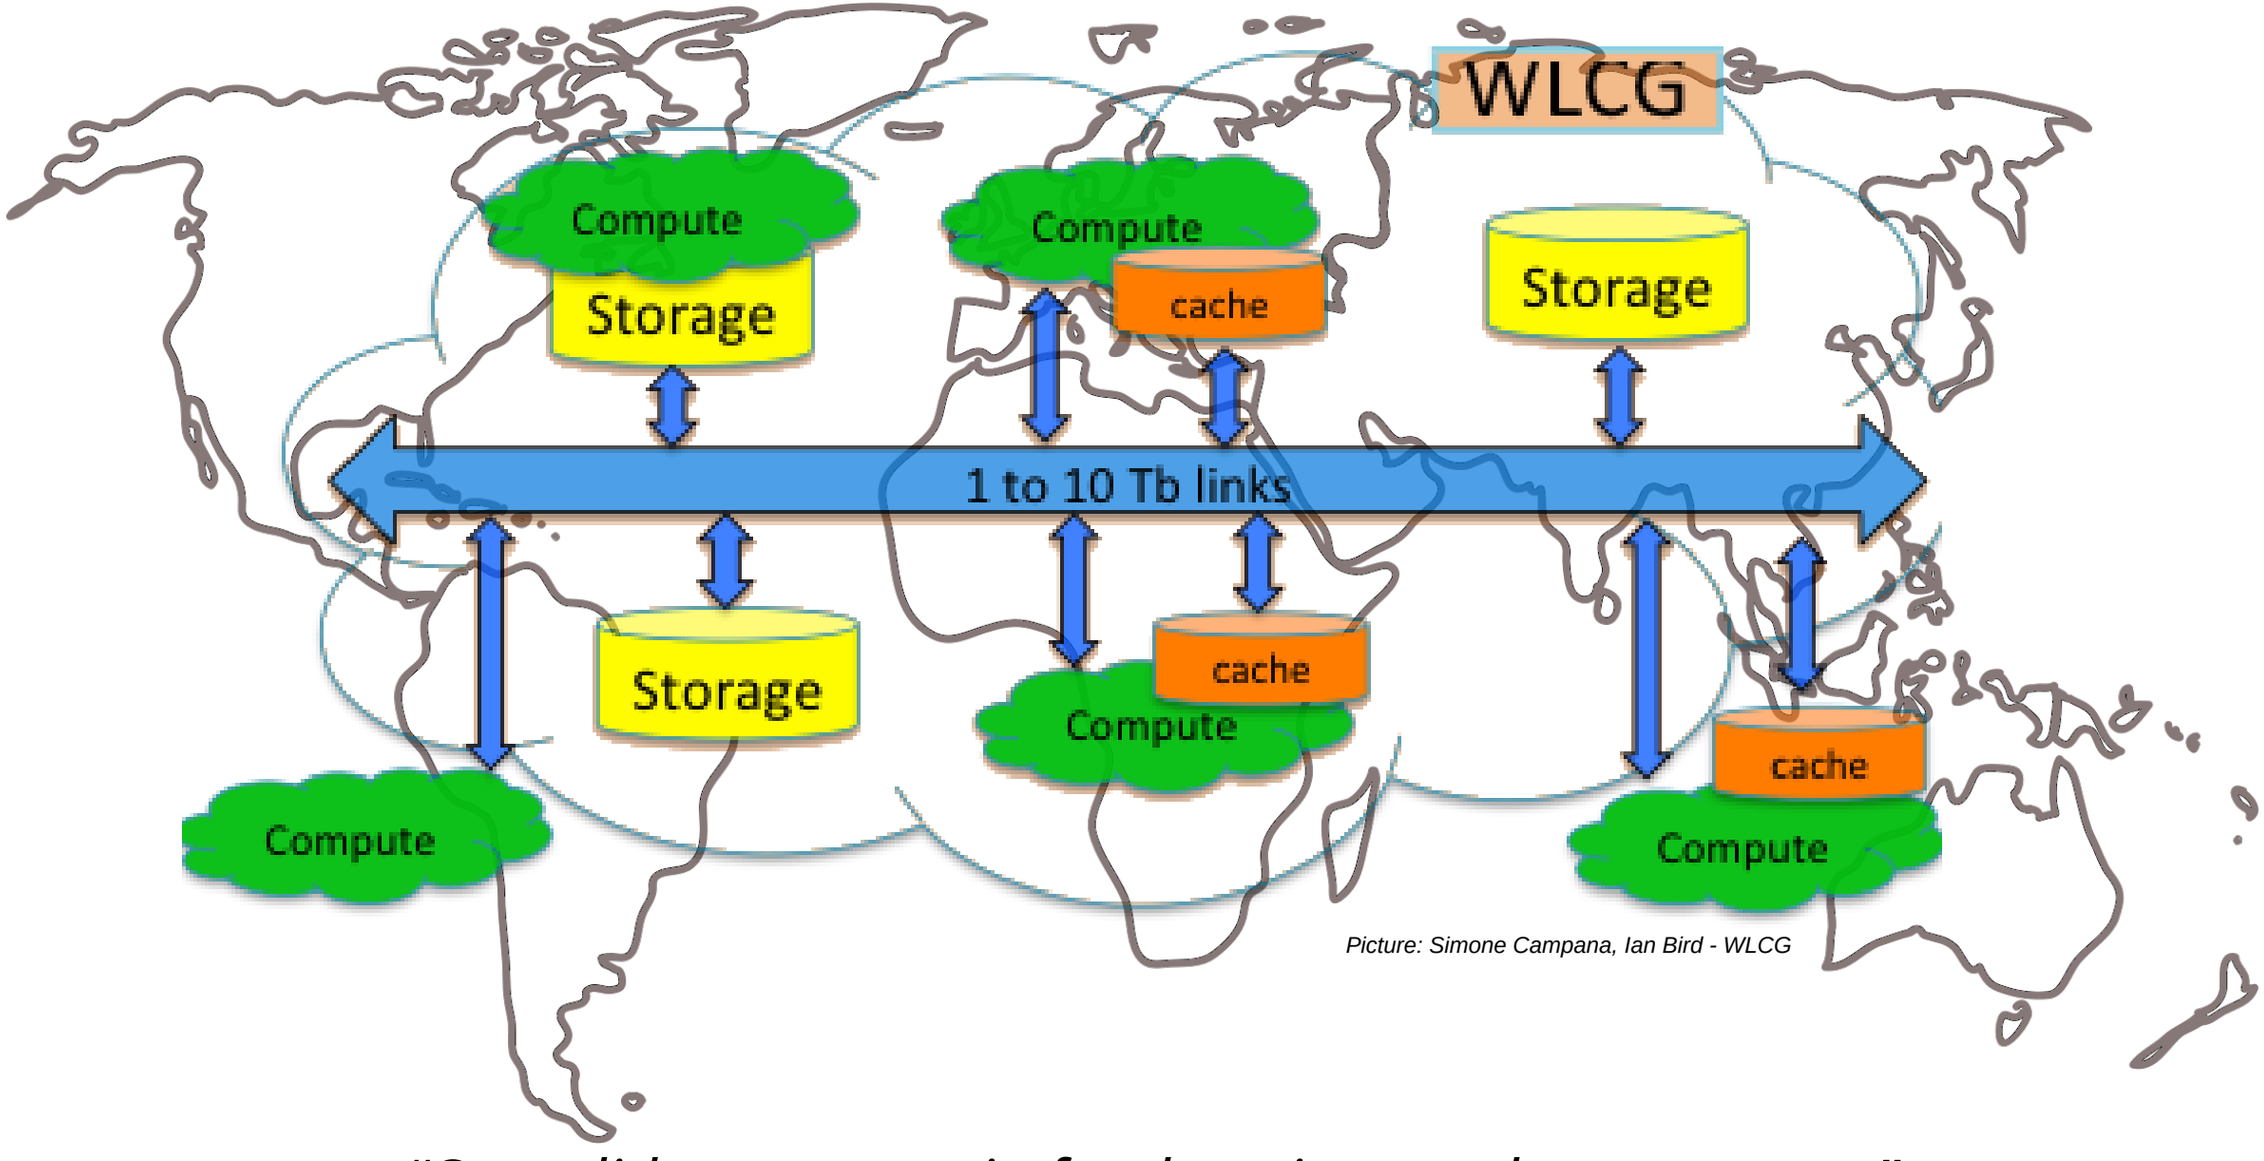

Picture: Simone Campana, Ian Bird - WLCG

"Consolidate storage in few locations to decrease costs"

GNA - Global Network Architecture

Sharing of R&E network resources for common benefit

Extend benefits gained with ANA-300G experience (sharing of three 100G transatlantic links)

Pillars:
- resource sharing
- aligning operational standards
- maximizing funding and investments
- knowledge sharing and outreach
- increasing global collaboration

https://gna-re.net/

Conclusion

Summary

LHCOPN:

- increasing capacity
- IPv6 deployed and started being used
- started 100G adoption

LHCONE:

- expanding in capacity
- more connections in Asia-Pacific
- consolidating perfSONAR infrastructure

Future developments:

- use of commercial cloud resources
- planning for data deluge in Run4

More information

Next LHCOPN/ONE meeting:
Date: 16-17 October 2017
Location: KEK, Tsukuba – Japan
http://indico.cern.ch/event/646629/

Previous LHCOPN/ONE meetings:
Helsinki, September 2016: https://indico.cern.ch/event/527372/
BNL New York, March 2017: http://indico.cern.ch/event/581520

Websites:
LHCOPN: https://twiki.cern.ch/twiki/bin/view/LHCOPN/WebHome
LHCONE: https://twiki.cern.ch/twiki/bin/view/LHCONE/WebHome

Questions?

edoardo.martelli@cern.ch