

LHCOPN and LHCONE an introduction

APAN workshop
Nantou, 13th August 2014
Edoardo.Martelli@cern.ch

Summary



- WLCG
- LHCOPN
- LHCONE
 - L3VPN
 - P2P
 - perfSONAR


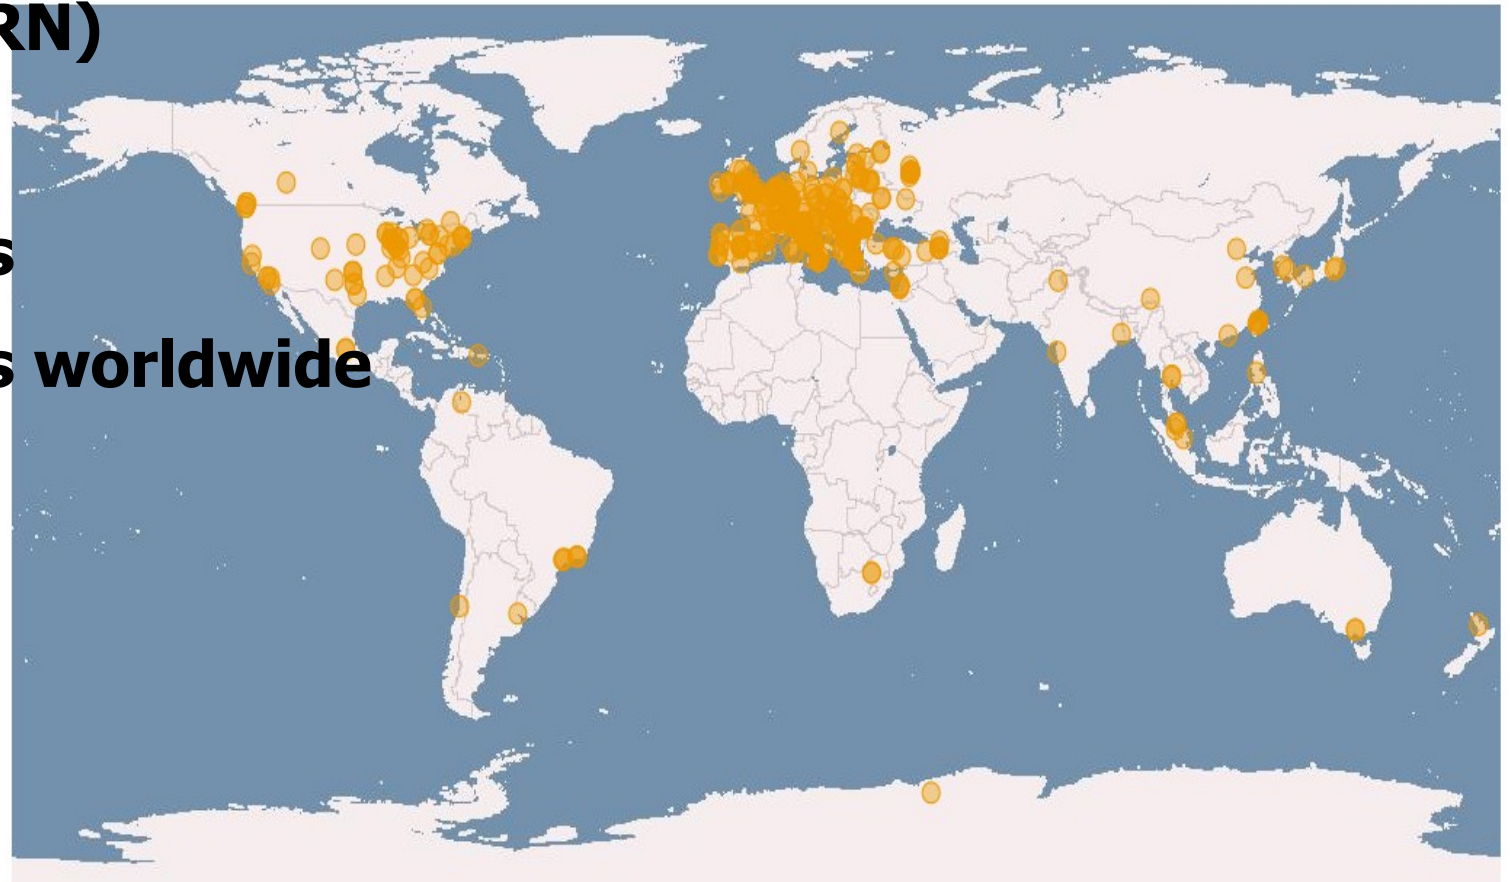
WLCG

Worldwide LHC Computing Grid



WLCG sites:

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



WLCG
Worldwide LHC Computing Grid

Computer Networks are an essential component of the WLCG

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

Networks for WLCG



Two dedicated, private data network have been built for WLCG:

- LHCOPN (Tier0-Tier1s)
- LHCONE (Tier1s-Tier2s)

LHCOPN

LHC Optical Private Network

What is the LHCOPN:



Private network connecting **Tier0 and Tier1s**

Reserved to LHC data transfers and analysis

Dedicated large bandwidth links

Highly resilient architecture

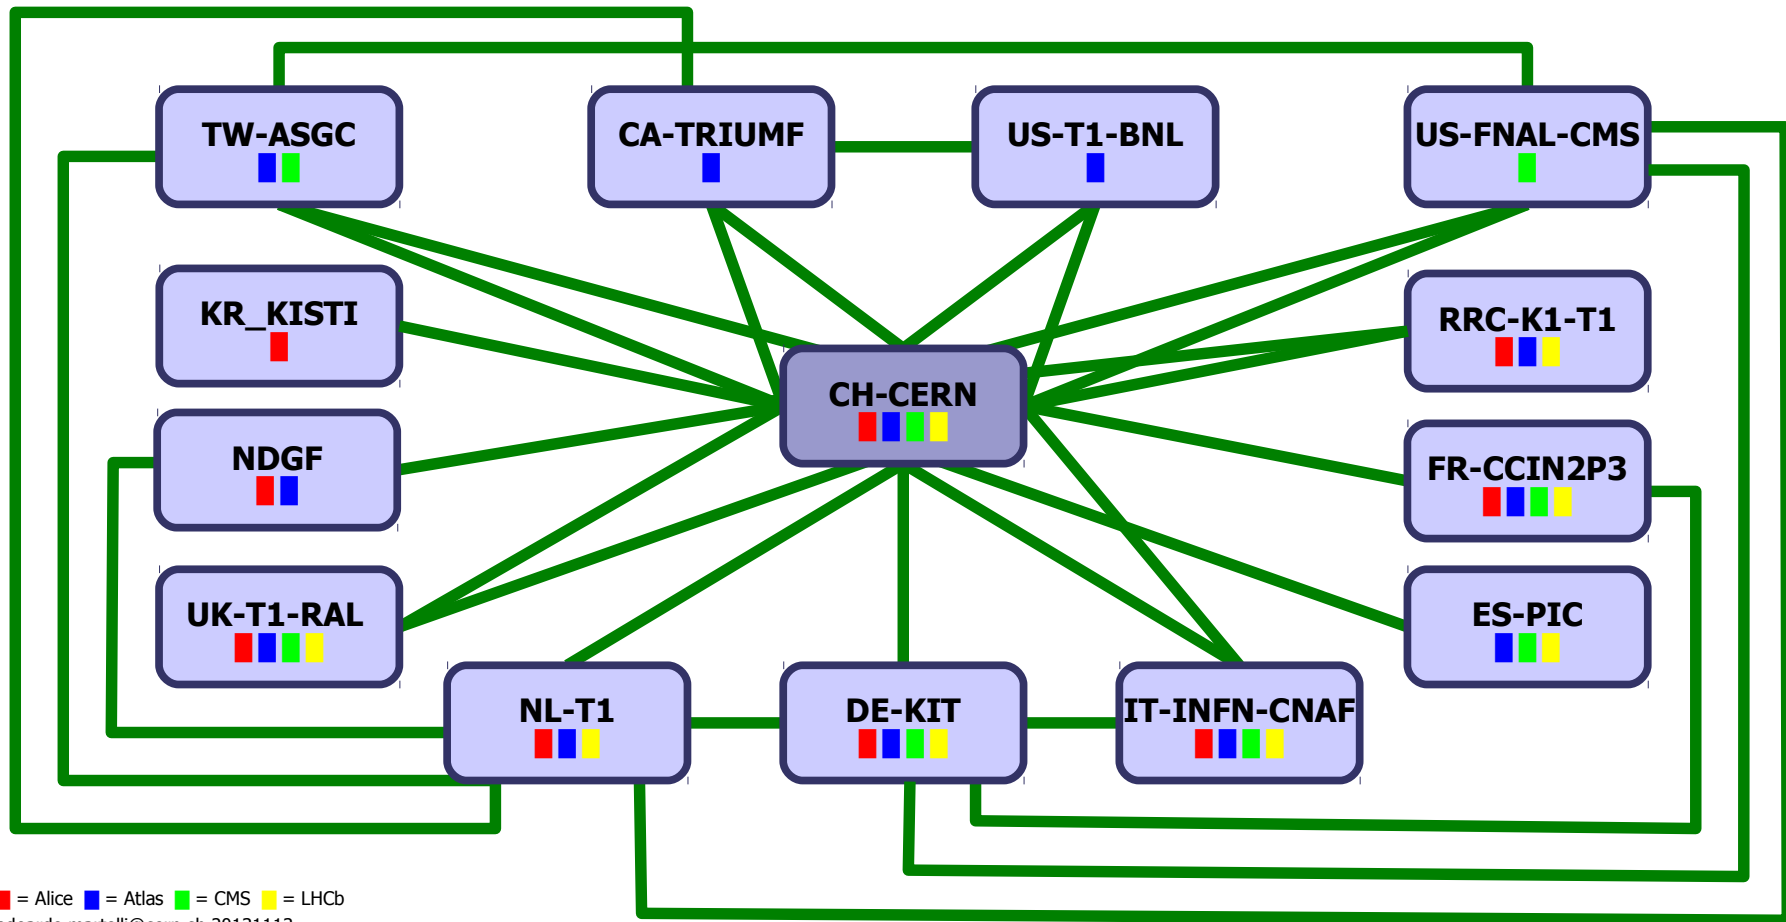
A collaborative effort



Layer3 (IP): designed, built and operated by the Tier0-Tier1s community

Layer1-2 (transmission): links provided by Research and Education network providers:
Asnet, ASGCnet, Canarie, DFN, Esnet, GARR, Geant, JANET, Kreonet, Nordunet, Rediris, Renater, Surfnets, SWITCH, TWAREN, USLHCnet

Topology



■ = Alice ■ = Atlas ■ = CMS ■ = LHCb
edoardo.martelli@cern.ch 20131113

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)

LHCONE

LHC Open Network Environment


New computing model



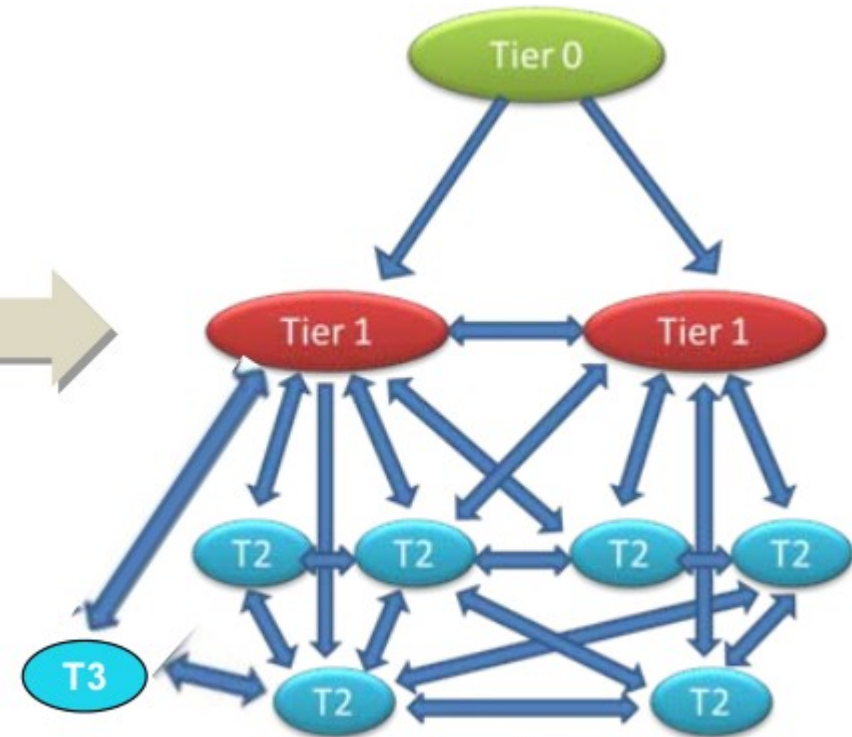
- Better and more dynamic use of storage
- Reduced load on the Tier1s for data serving
- Increased speed to populate analysis facilities

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

Computing model evolution



Original MONARCH model



Model evolution

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 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)
- A collaborative effort among Research & Education Network Providers
- Traffic separation: no clash with other data transfer, resource allocated for and funded by the HEP community

LHCONE is a community effort.

All stakeholders involved: TierXs, Network Operators, LHC Experiments, CERN.

L3VPN (VRF): routed Virtual Private Network -
operational

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

perfSONAR: monitoring infrastructure -
operational

LHCONE L3VPN

What LHCONE L3VPN is:



Layer3 (routed) Virtual Private Network

Dedicated worldwide backbone connecting
Tier1s and Tier2s (and Tier3s) at high
bandwidth

Reserved to LHC data transfers and analysis

Advantages



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

Well defined cost tag for WLCG networking

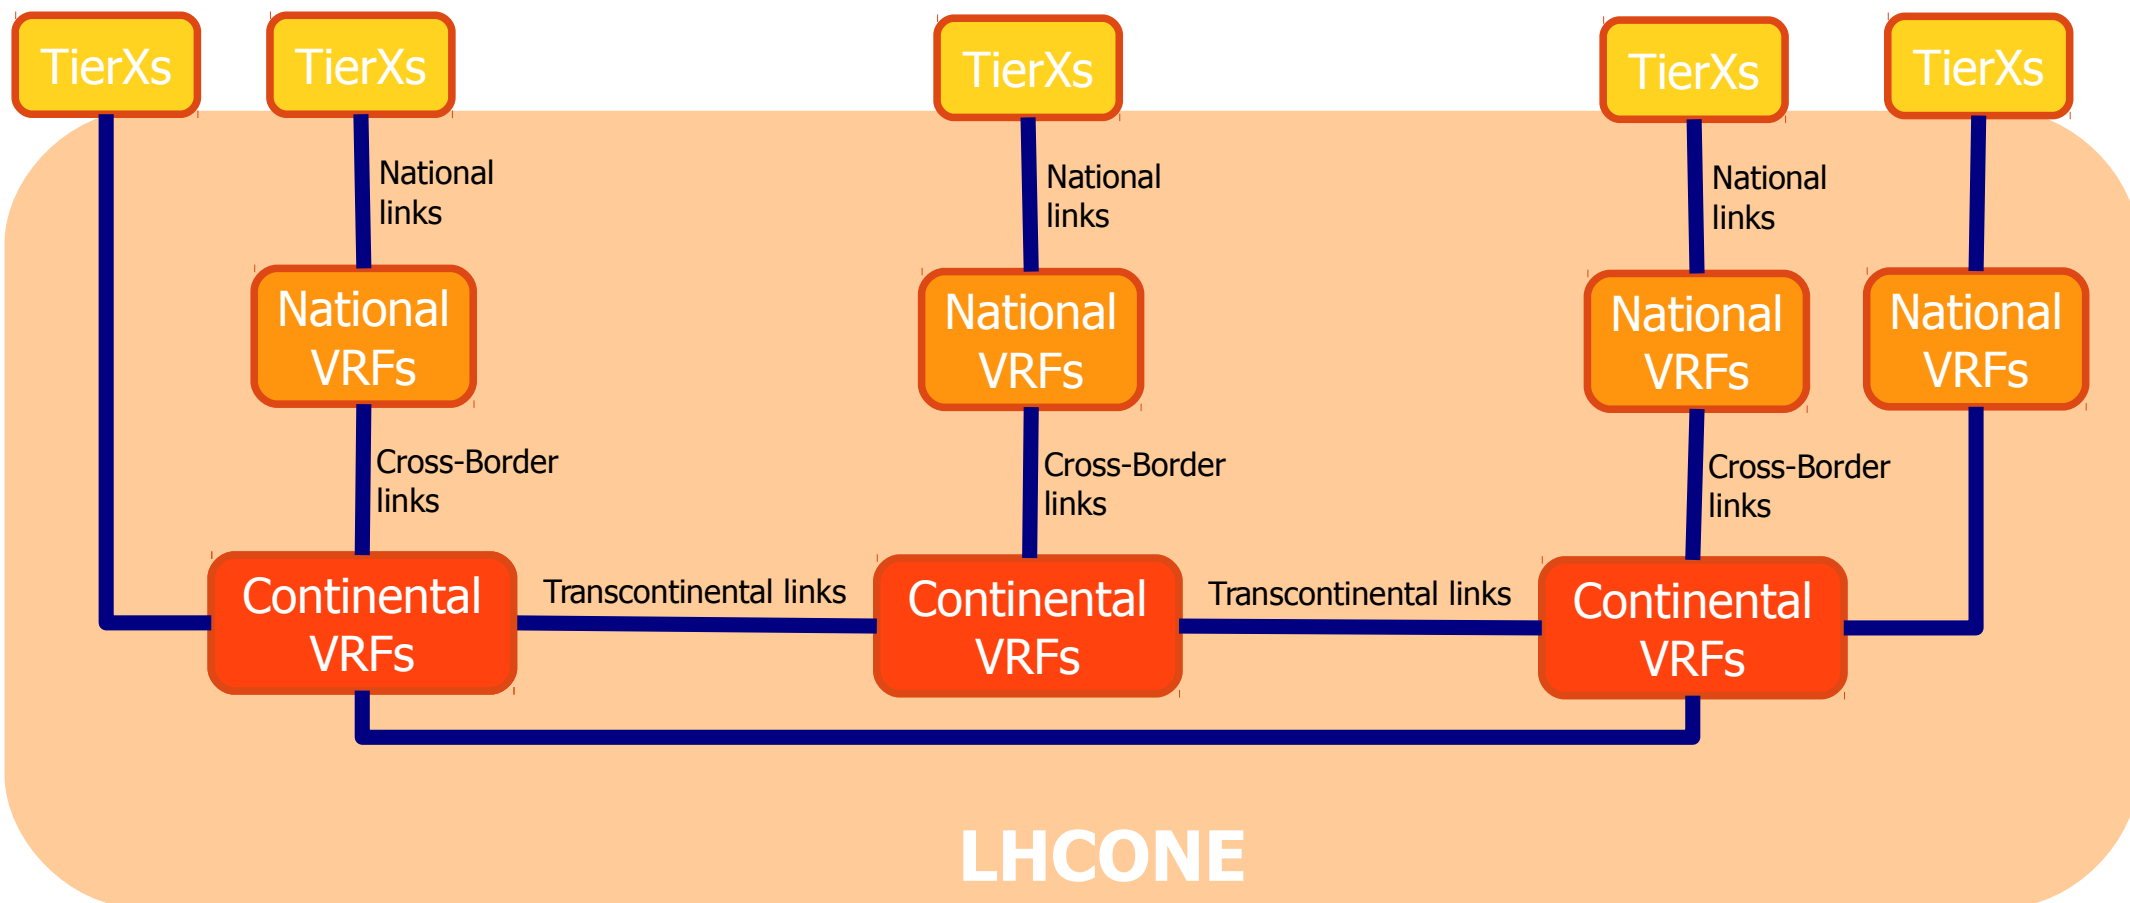
Trusted traffic that can bypass firewalls

LHCONE L3VPN architecture

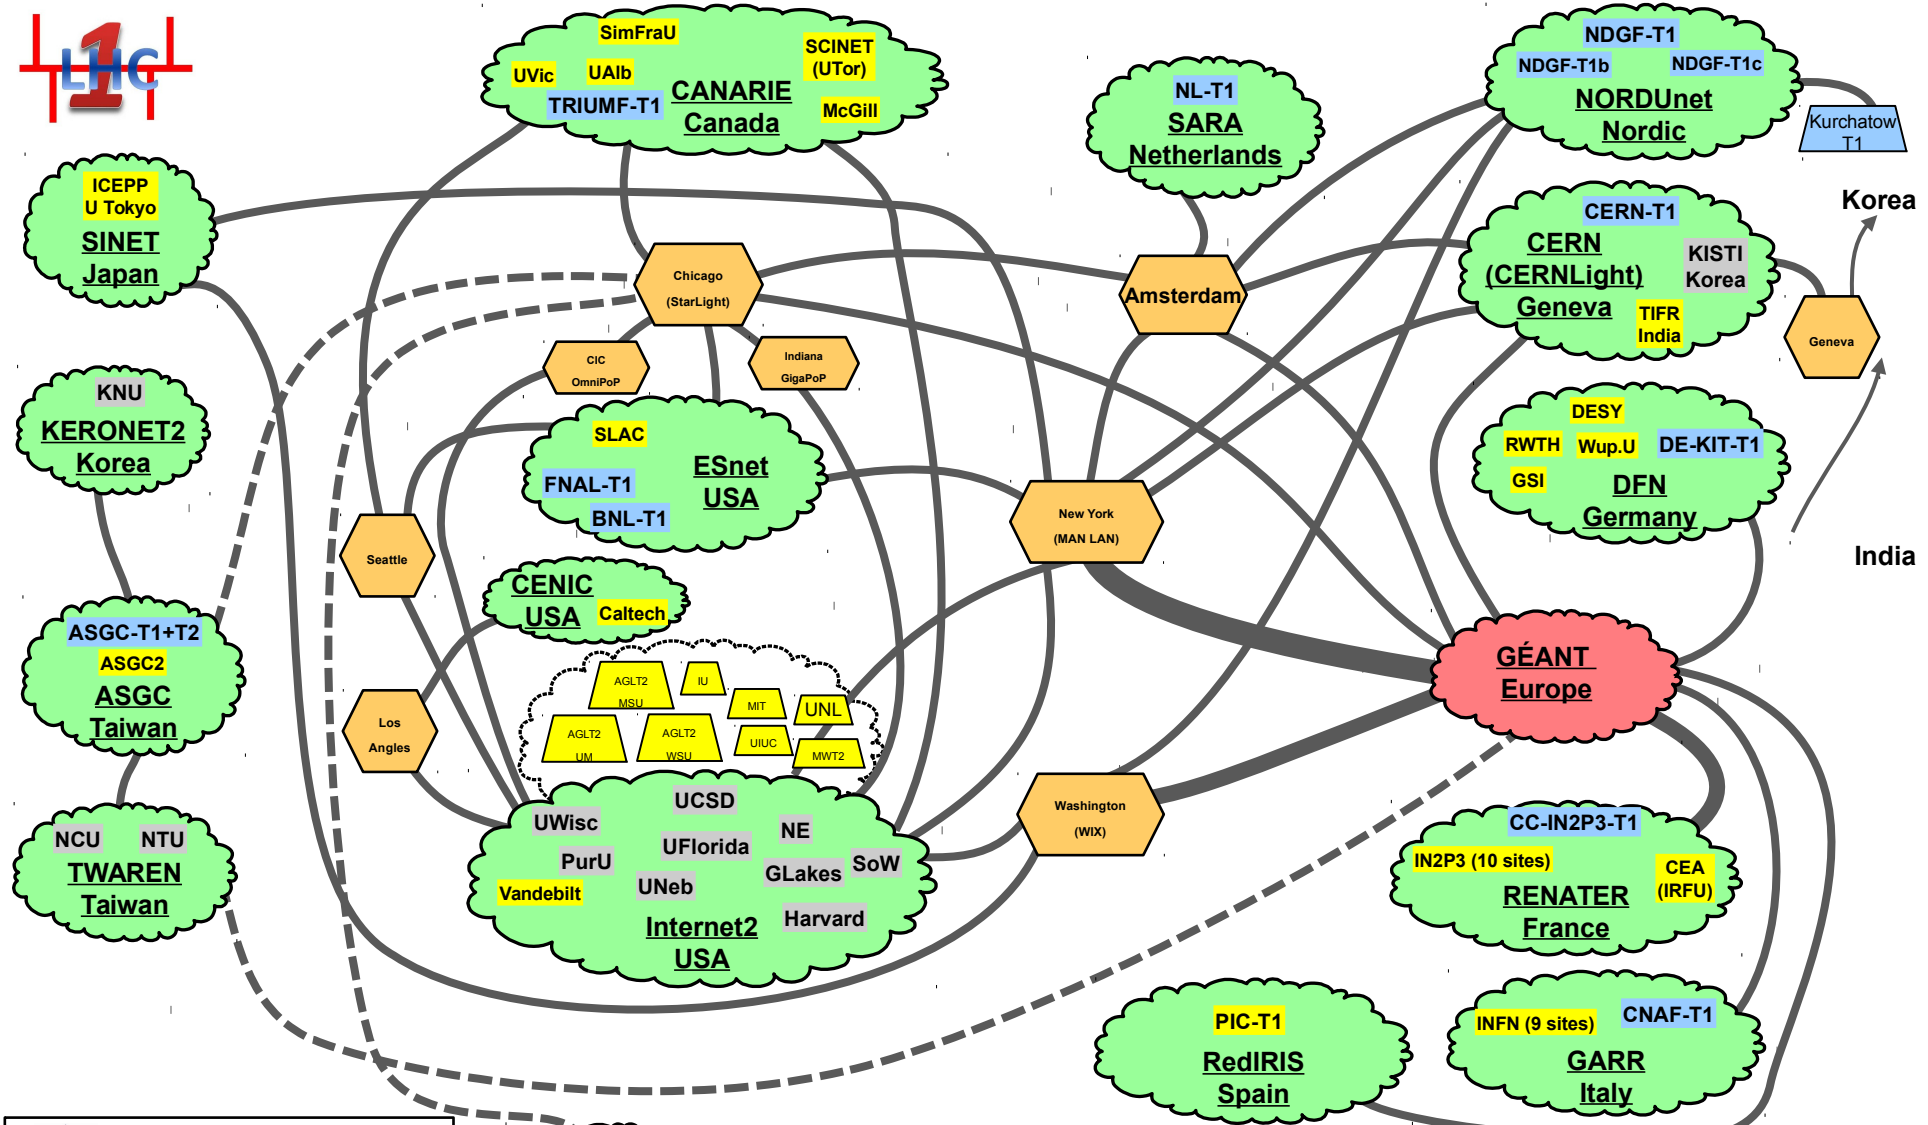


- TierX sites connected to National-VRFs or Continental-VRFs
- National-VRFs interconnected via Continental-VRFs
- Continental-VRFs interconnected by trans-continental/trans-oceanic links

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



Current L3VPN topology



UNAM
CUDI
Mexico

	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

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, Dante

More Information: <https://indico.cern.ch/event/289680/contribution/6/material/slides/0.ppt>

Operations



Usual Service Provider operational model: a TierX must refer to the VRF providing the local connectivity

Bi-weekly call among all the VRF operators and concerned TierXs

LHCONE P2P

Guaranteed bandwidth point-to-point links

What is LHCONE P2P:



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

Work in progress: still in design phase

Challenges:

- multi-domain provisioning system
- intra-TierX connectivity
- TierX-TierY IP routing
- APIs for WLCG software

LHCONE perfSONAR

What is LHCONE perfSONAR



LHCONE Network monitoring infrastructure

Probes installed at:

- VRFs interconnecting points
- TierXs

Accessible to any TierX for network healthiness checks

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



- VRFs have started upgrading internal links and links to TierXs to 100Gbps
- VRFs interconnecting links will be upgraded to 100Gbps. 100Gbps Transatlantic link being tested.
- Operations need to be improved, especially how to support a TierX in case of performance issue
- perfSONAR deployment will be boosted

Conclusions

Conclusions



- New Computing Models will rely even more on good and abundant network connectivity
- Tier2s need to improve their network connectivity
- LHCONE-L3VPN is a viable solution already adopted by many Tier1/2s

More information



Last LHCONE workshop:

<https://indico.cern.ch/event/289679/>

LHCONE websites:

<http://lhcone.net>

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

Weekly audio conference:

Monday 14:30 GMT, alternating every second week
architecture and operations

Mailing lists:

lhcone-operations@cern.ch

lhcone-architecture@cern.ch

Questions?