Corso reti avanzato per neo assunti (PNRR)

Netgroup

Introduction

This is the second part of the course, and it addresses some advanced and more specific network related concepts.

It will:

- Security
- Overview of the INFN network infrastructure
- Describe the network context on the international peerings
- Give a glance to new technologies

SECURITY ACCESS CONTROL LISTS

Definition of ACL (Access Control List)

In computer security, an access-control list (ACL) is a list of permissions associated with a system resource.

In networking, ACLs are lists of rules to be applied to specific traffic flows.

ACLs are generally managed by specific hardware devices like **Routers, Switches, and Firewalls.**

An ACL contains a sequential list of **permit** or **deny** statements and are evaluated from the first to the last. (e.g. cisco “exits” on first match).

In general Security ACLs are applied on the **WAN access interface** to protect the entire network and or on the internal devices to protect specific parts of the network

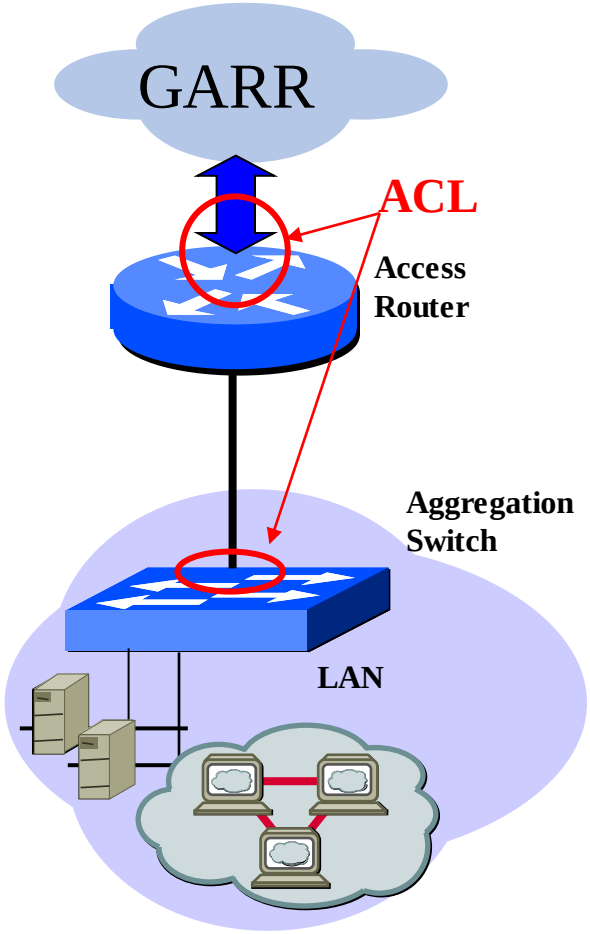

ACL Creation (Example)

In an ACL it is possible to discriminate a data flow by: IP Source, IP Destination, Protocol used (IP, TCP , UDP, ICMP…) and eventually the port number.

Ip access-list extended <ACL Name>

permit/deny <Protocol> <Source Address> <port> <Destination Address> <port> <log>

The addresses definition can be a single host or a network.

To identify a network, **it is necessary to use the wildcard address.**

Example of ACL (Cisco CLI)

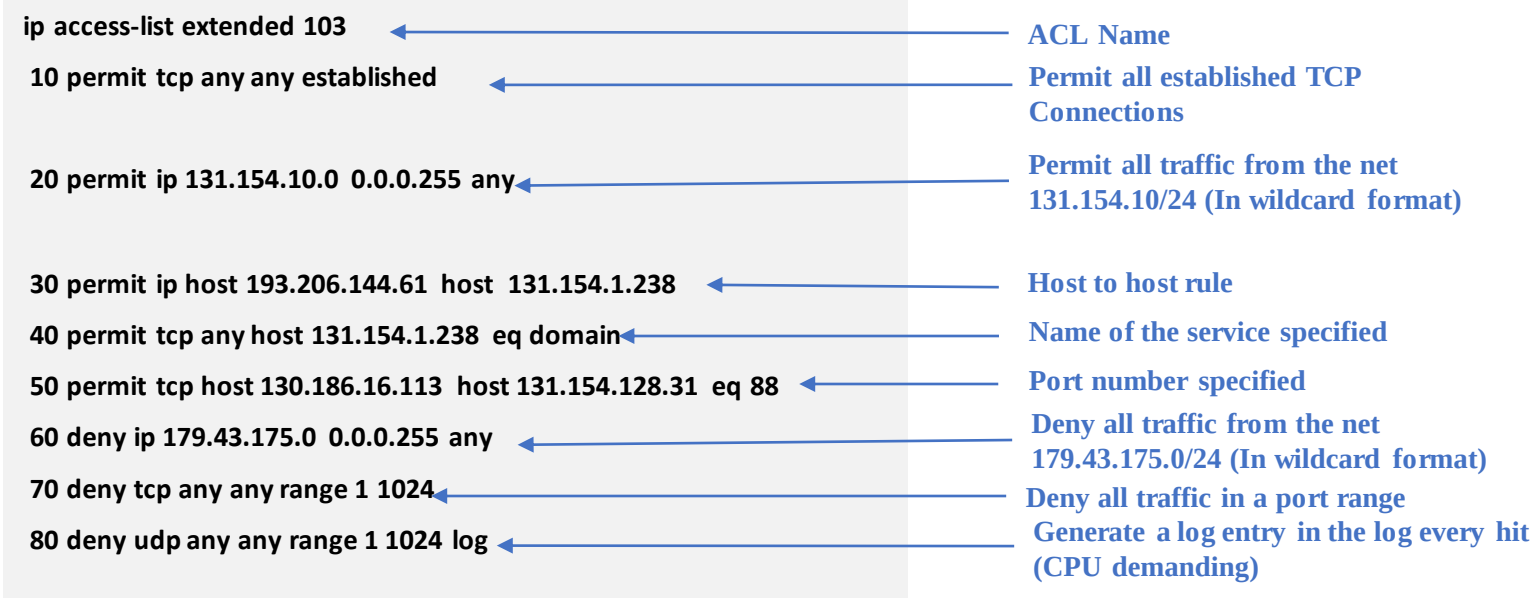

Netmask / Wildcard correspondence

Slash	Netmask	Wildcard mask
/32	255.255.255.255	0.0.0.0
/31	255.255.255.254	0.0.0.1
/30	255.255.255.252	0.0.0.3
/29	255.255.255.248	0.0.0.7
/28	255.255.255.240	0.0.0.15
/27	255.255.255.224	0.0.0.31
/26	255.255.255.192	0.0.0.63
/25	255.255.255.128	0.0.0.127
/24	255.255.255.0	0.0.0.255
/23	255.255.254.0	0.0.1.255
…	…	…

Application of ACL (Access Control List)

ACLs are generally applied to: **Physical Interfaces (ports),** VLANs, IP Interfaces associated to VLANs

ACLs are applied for a specific traffic direction (**Input** or **output**)

ACLs control traffic in **only one direction,** so if you want to control inbound and outbound traffic you must create two distinct ACLs per interface.

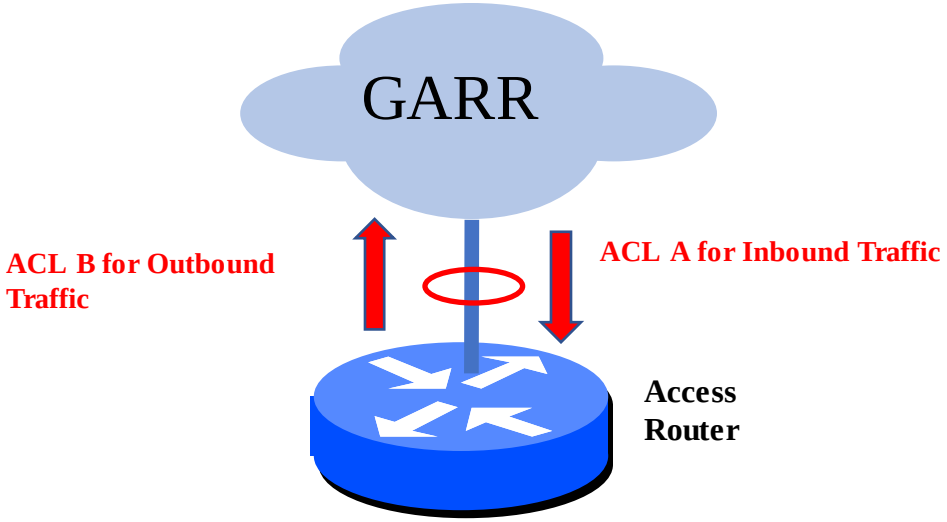

Example of ACL application (Cisco CLI)

interface Vlan1000 ← Interface to which apply the ACL

description general-internet

ip address 193.206.128.XX 255.255.255.252

no ip redirects

no ip unreachables

ip mtu 9178

ip nat outside

ip access-group 103 in → ACL (103) acting on IPv4 Input traffic

ip access-group 104 out → ACL (104) acting on IPv4 Output traffic

ipv6 address 2001:760:FFFF:110::XX/127

ipv6 traffic-filter IPv6-103 in → ACL (IPv6-103) acting on Input IPv6 traffic

ipv6 traffic-filter IPv6-104 out → ACL (IPv6-104) acting on Input IPv6 traffic

FIREWALL

Definition of Firewall

A firewall is a network security appliance that restricts data communication traffic to protect your network from unwanted accesses.

Firewalls can be:

Stateful firewalls: aging on the connections in a similar way as simple ACLs do but using dedicated hardware resources (at line rate). These firewalls can evaluate also the status of the TCP connection.

Next Generation Firewalls:

- Deep Packet Inspection (ASIC/FPGA)
- Application recognition (not only on the base of the ports used for the connections)
- Can recognize **in real time** malicious traffic coming from Internet or directed to Internet (for example by malware) using continuously updated **signature** database
- Can perform URL Filtering and Online Antivirus control
- SSL inspection (Need to use the firewall as an SSL proxy)
- Signature DB Continuously updated **(with relative recurring costs for subscriptions)**

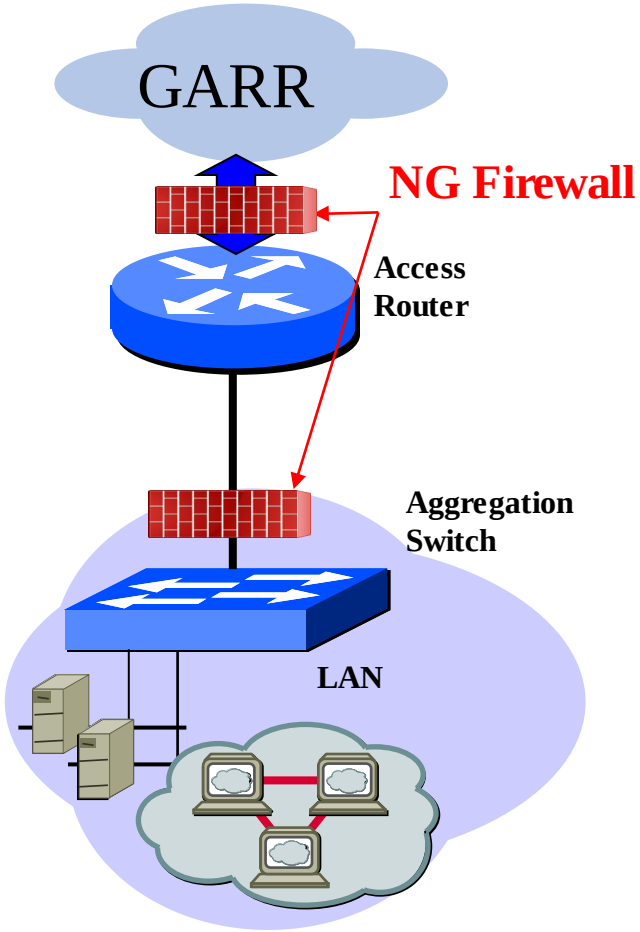

NG Firewall features

NGFW provides **dashboards** giving a status at a glance of the traffic and potential threats detected.

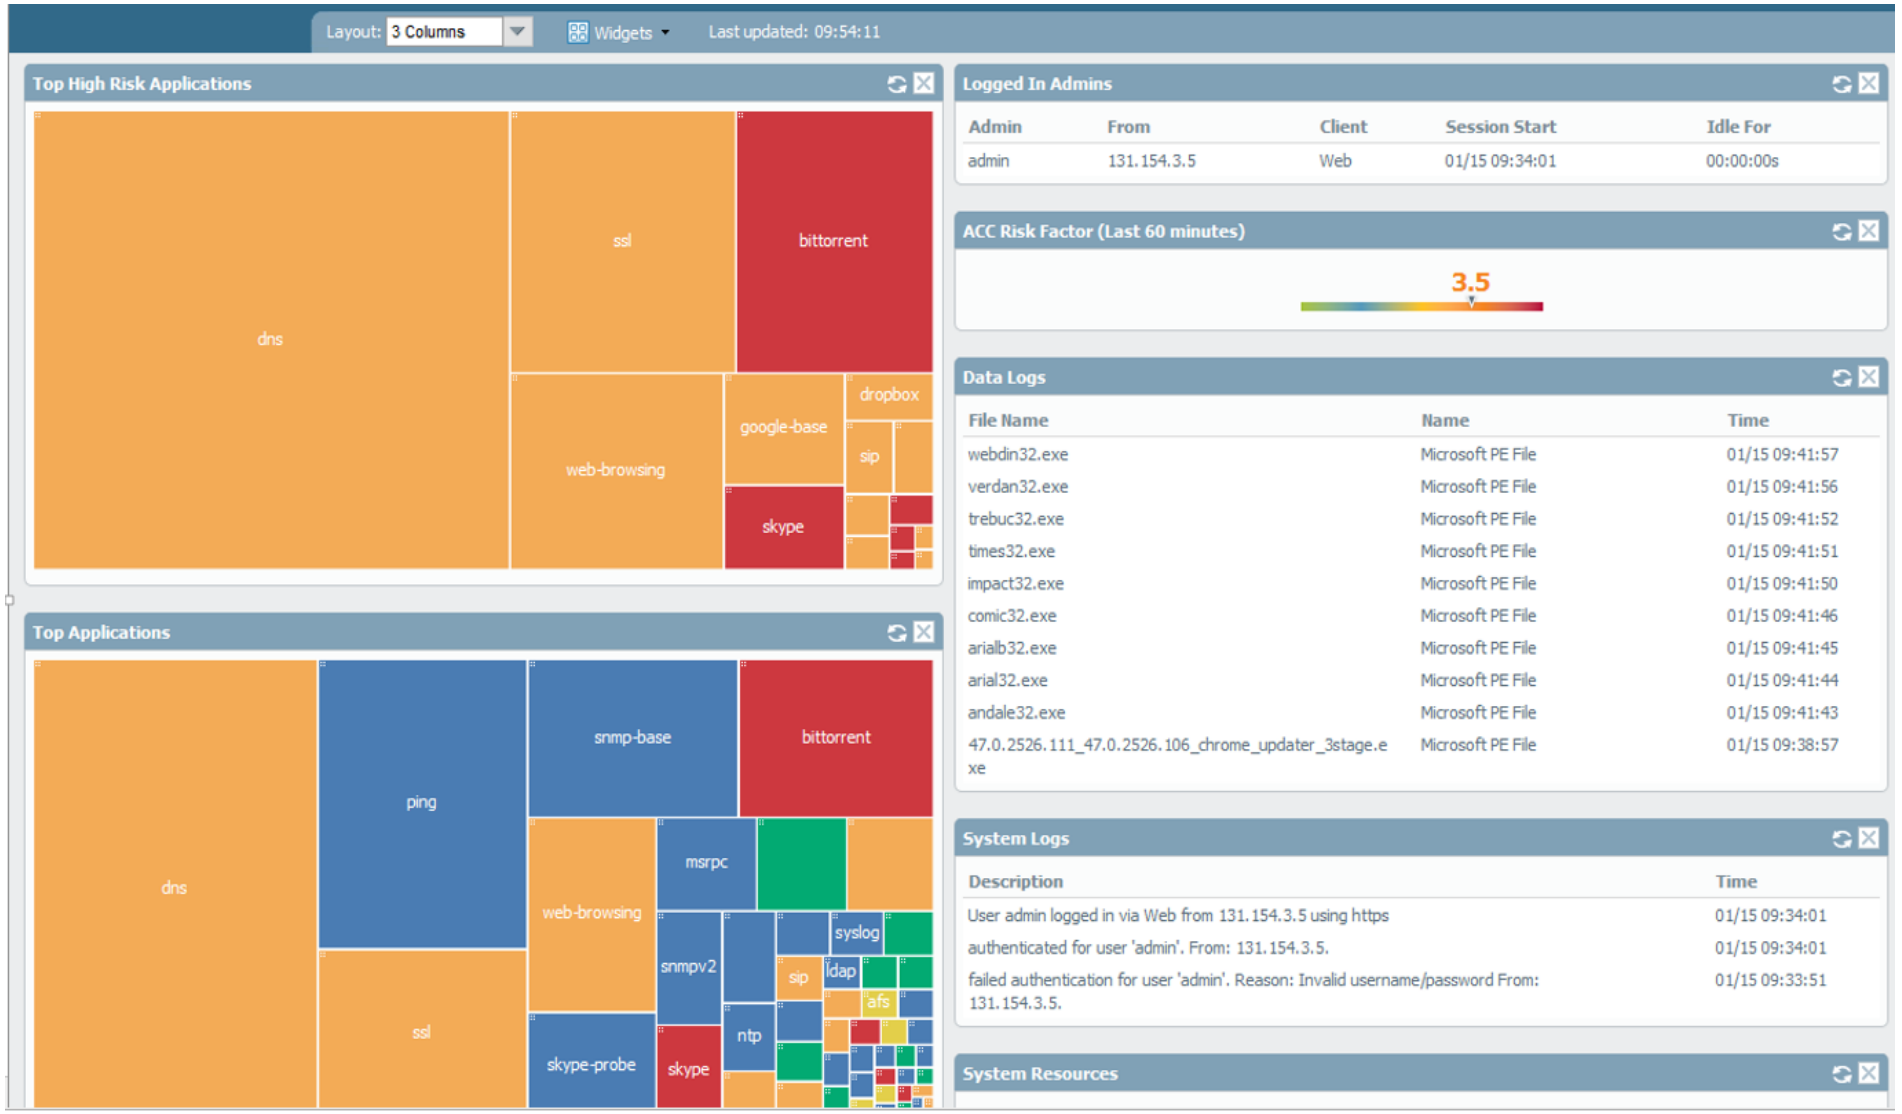

NG Firewall features

NGFW provides an automatic classification of the threats

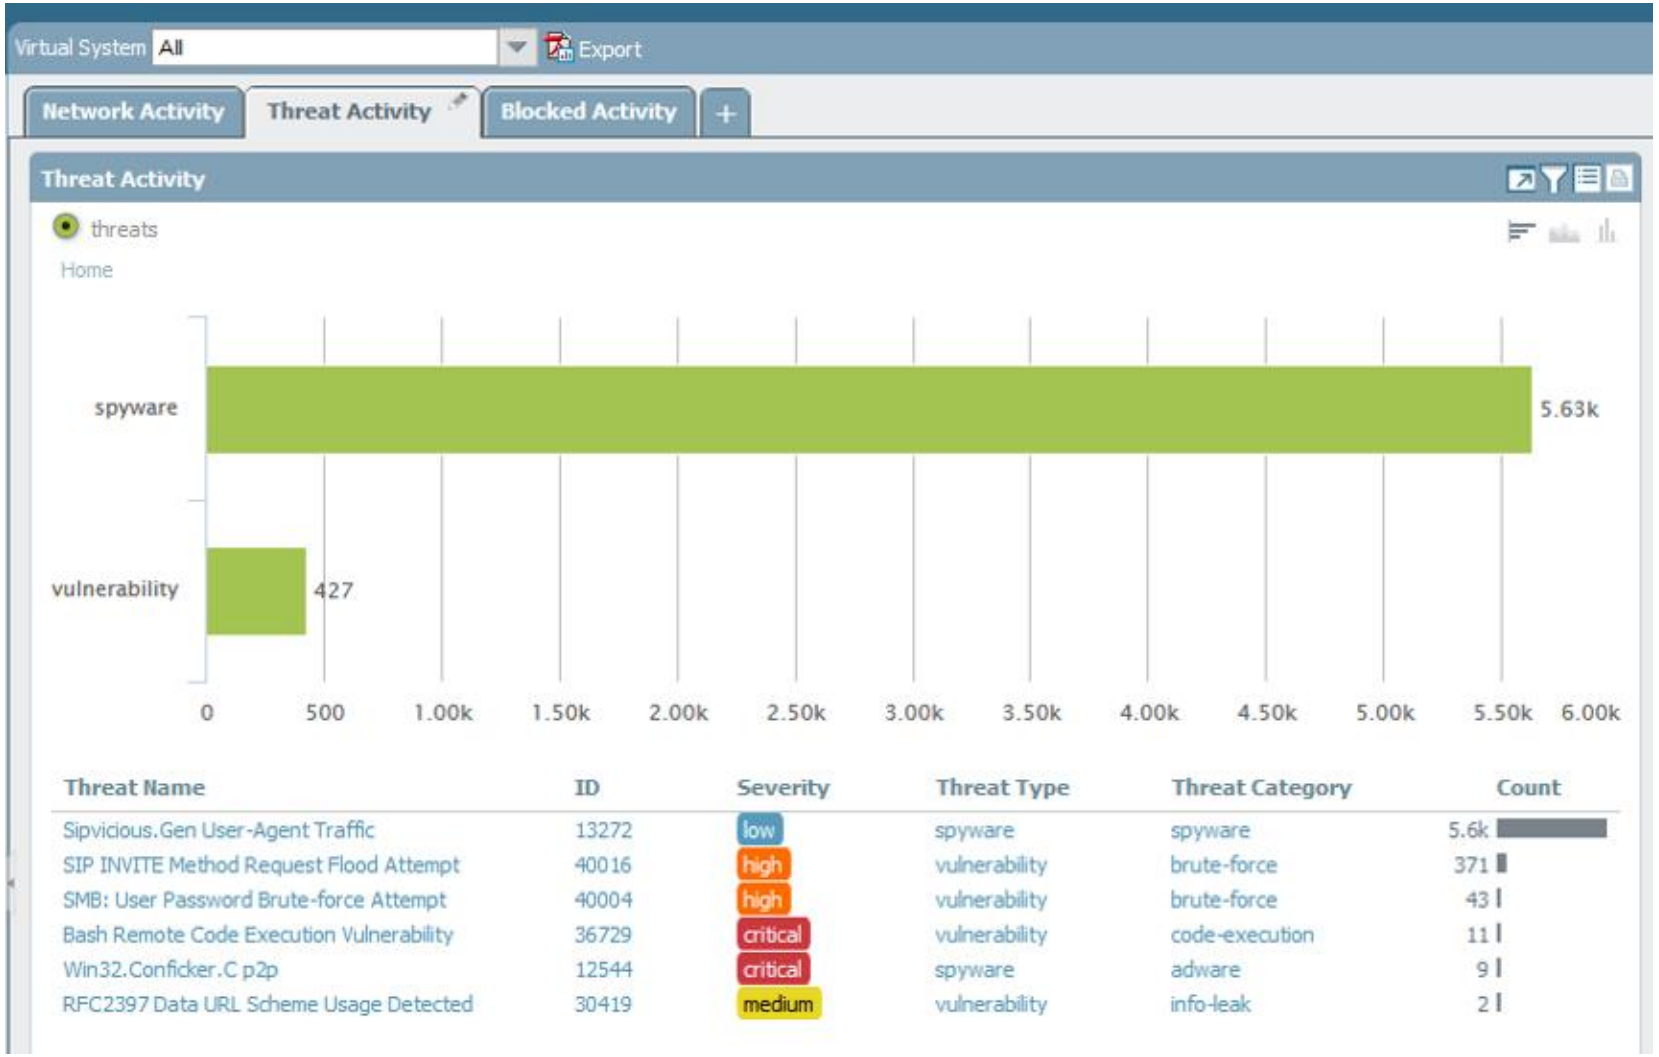

NG Firewall features

For each threat it is possible to view all the involved hosts (source and destination)

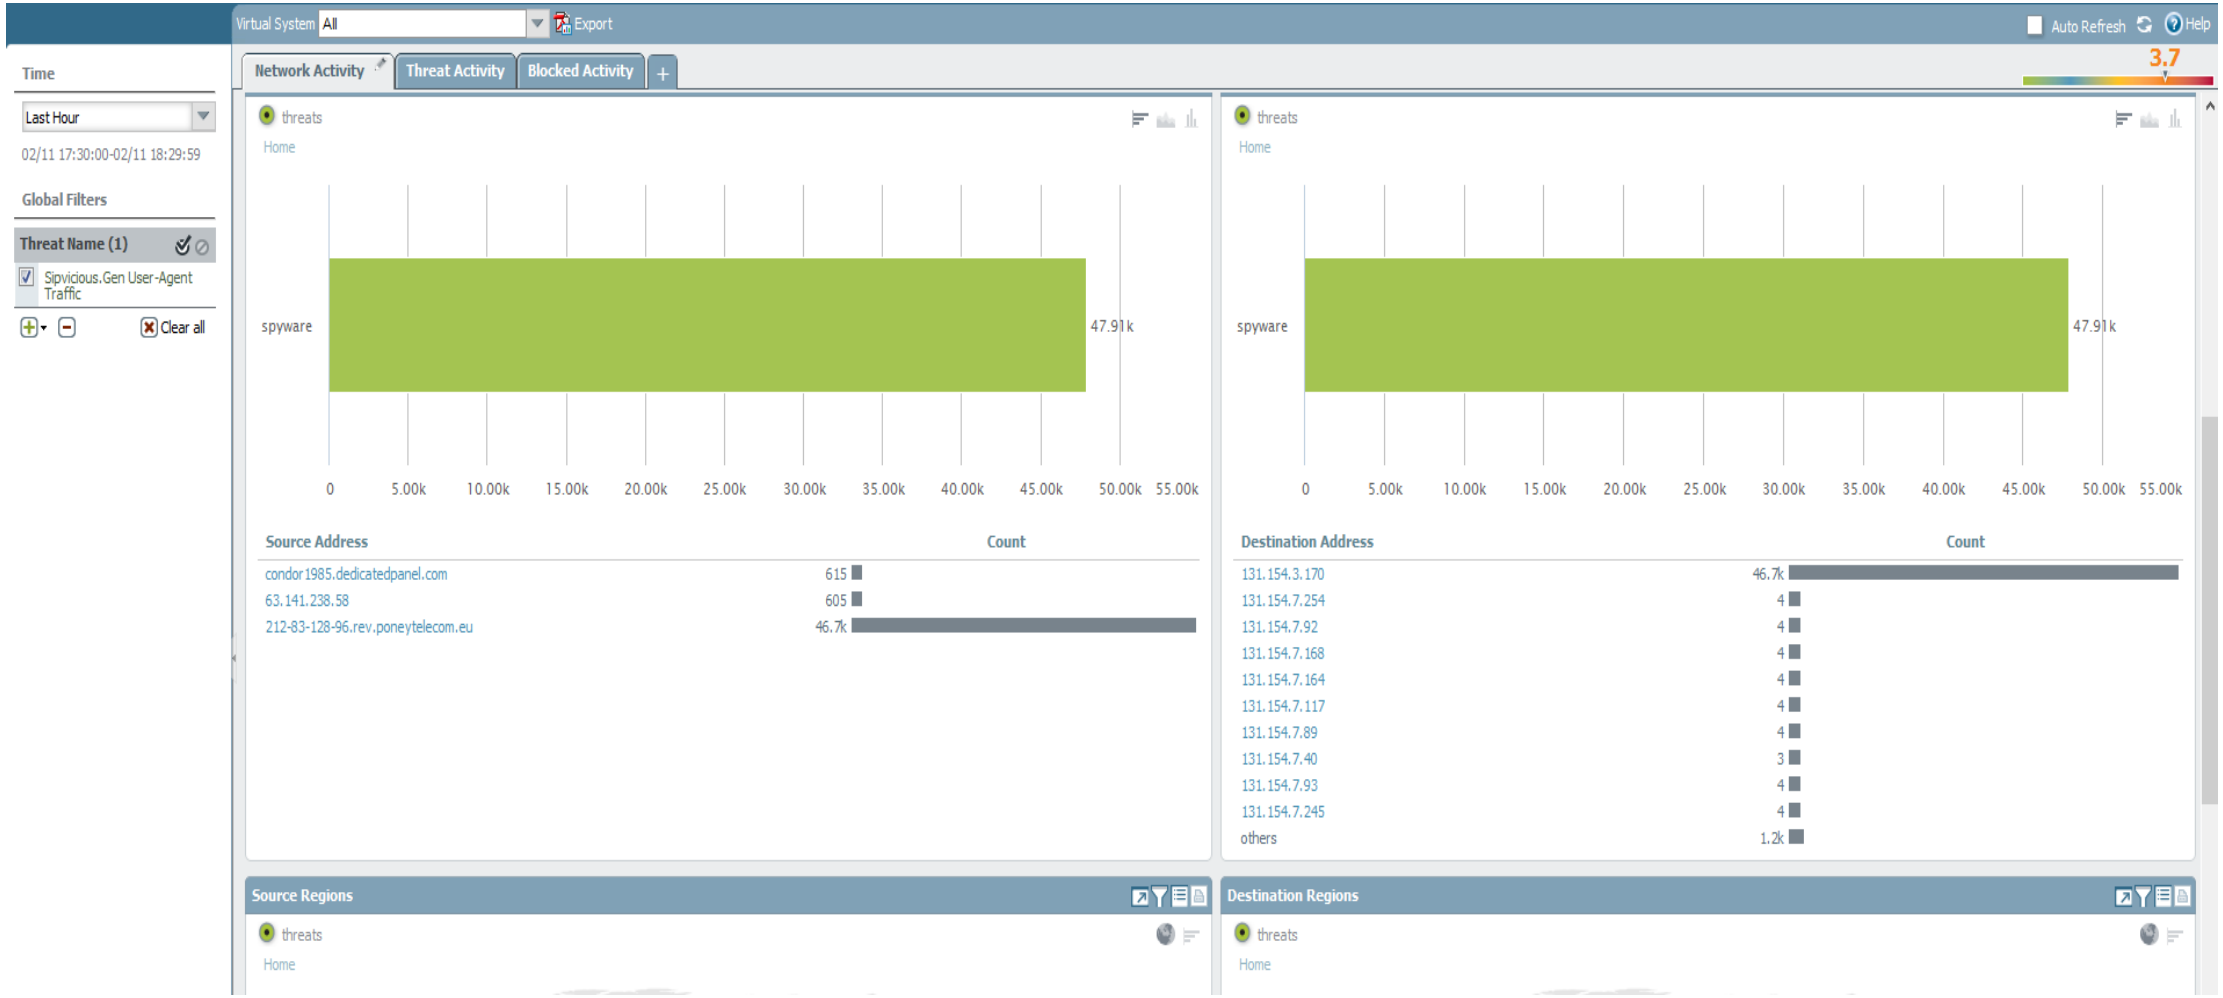

NG Firewall features

For each recognized threat it is possible to setup the specific rule to eventually block the malicious traffic in many different modes (drop, block-IP, alert, reset both sides…).

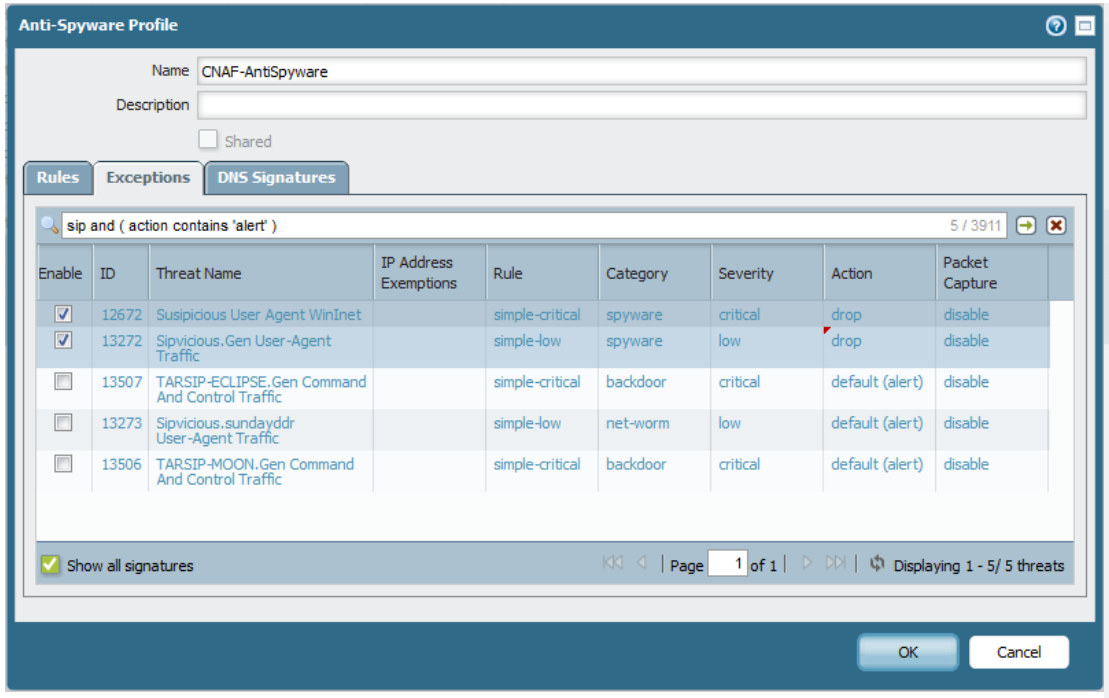

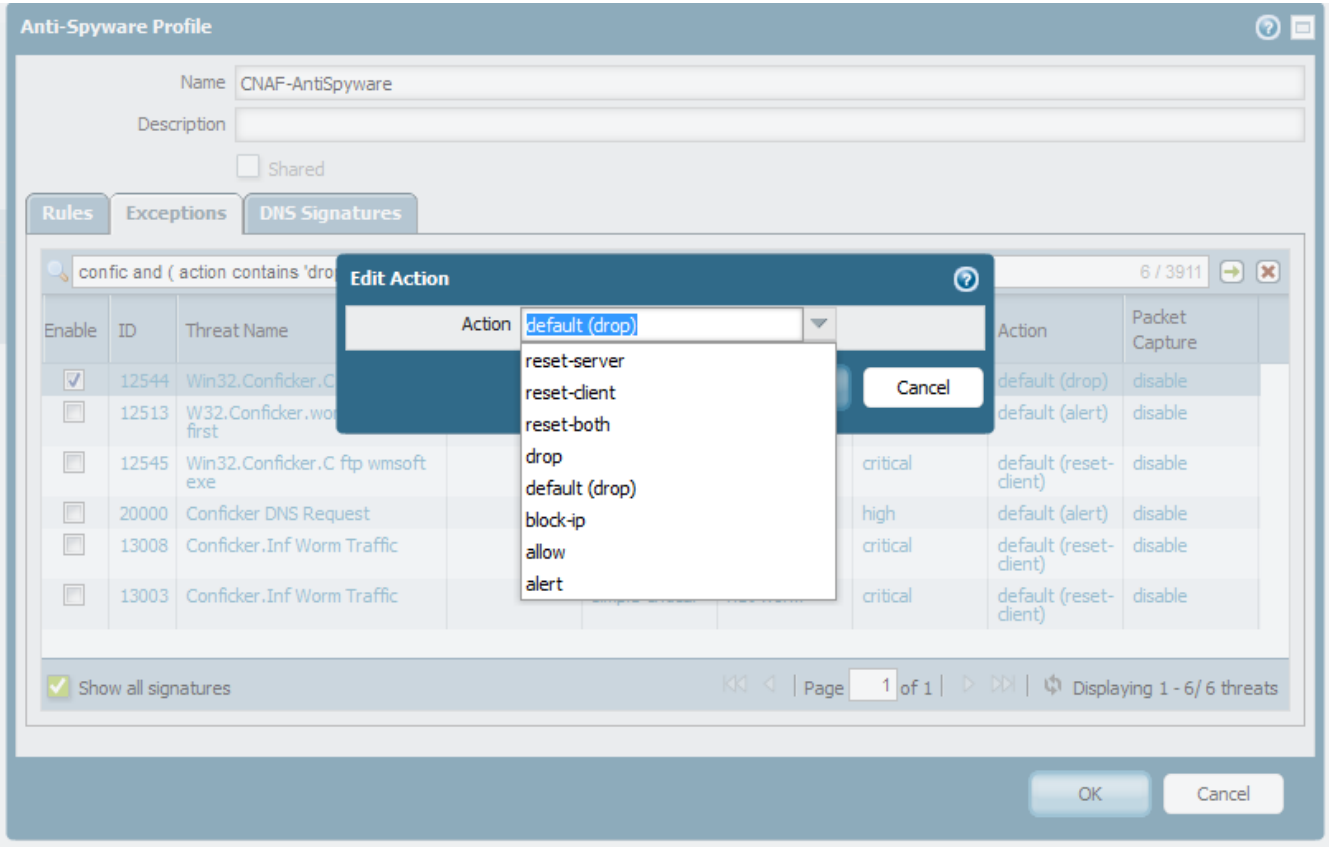

Firewall features

Automatically block suspicious activities

NGFWs can block traffic according to specific policies, manually defined traffic profiles or based on the automatic classification of detected threats.

URL Filtering

URL filtering can be fine-tuned defining the specific forbidden URLs or simply using the continuously updated lists of URL considered related to malware or in general malicious.

INFN does not enforce strict URL control to preserve users' privacy.

Only malicious URL filtering related to virus or malware is enabled.

DDOS protection

It is not easy to mitigate a well-organized Distributed Denial of Service attack but some NGFWs provide DoS protection mechanisms allowing a maximum number of sessions for each client and for specified applications. Setting maximum “hit rate” per client could help with DoS mitigation.

Firewall (Transparent or Routed)

Transparent mode:

Firewalls can be configured to be completely transparent to the network. In this configuration firewalls don’t have any IP address and are traversed by all traffic configuring the ports to simulate a virtual “wire” in certain cases not participating neither to LACP or other L2 protocols.

Routed Mode:

Firewalls are participating in the routing process (behave as a router) and can operate for example as NAT. In this case firewall have an IP address and are a visible hop in the network path.

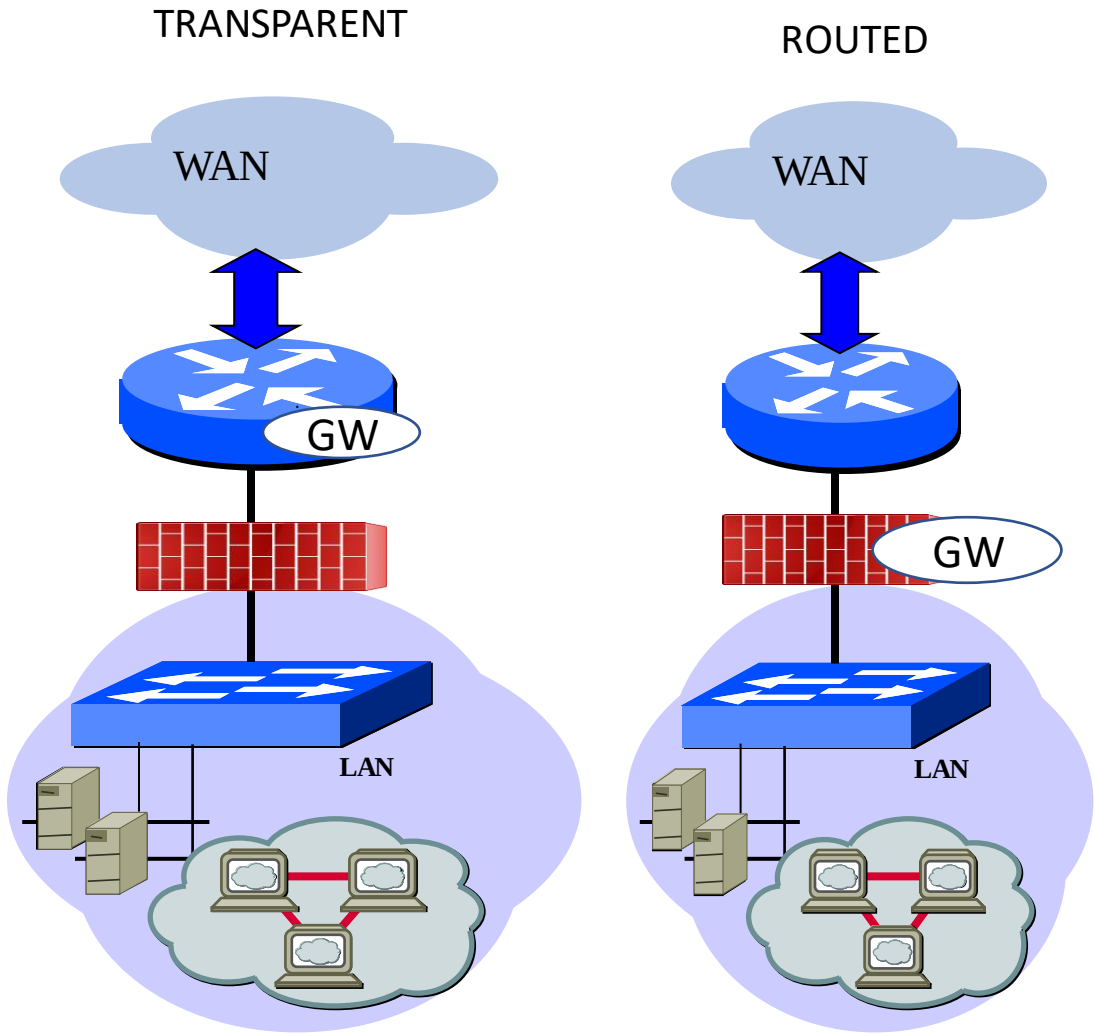

NG Firewall

Next Generation Firewalls can operate at **line rate** as IDS/IPS.

The limits in term of throughput depends on the HW resources (ASIC or FPGA adopted).

In general NG Firewalls are expensive devices and in order to reduce costs, it's recommended, where possible to use **bypass rules (whitelists)** in order to avoid deep inspection on trusted traffic.

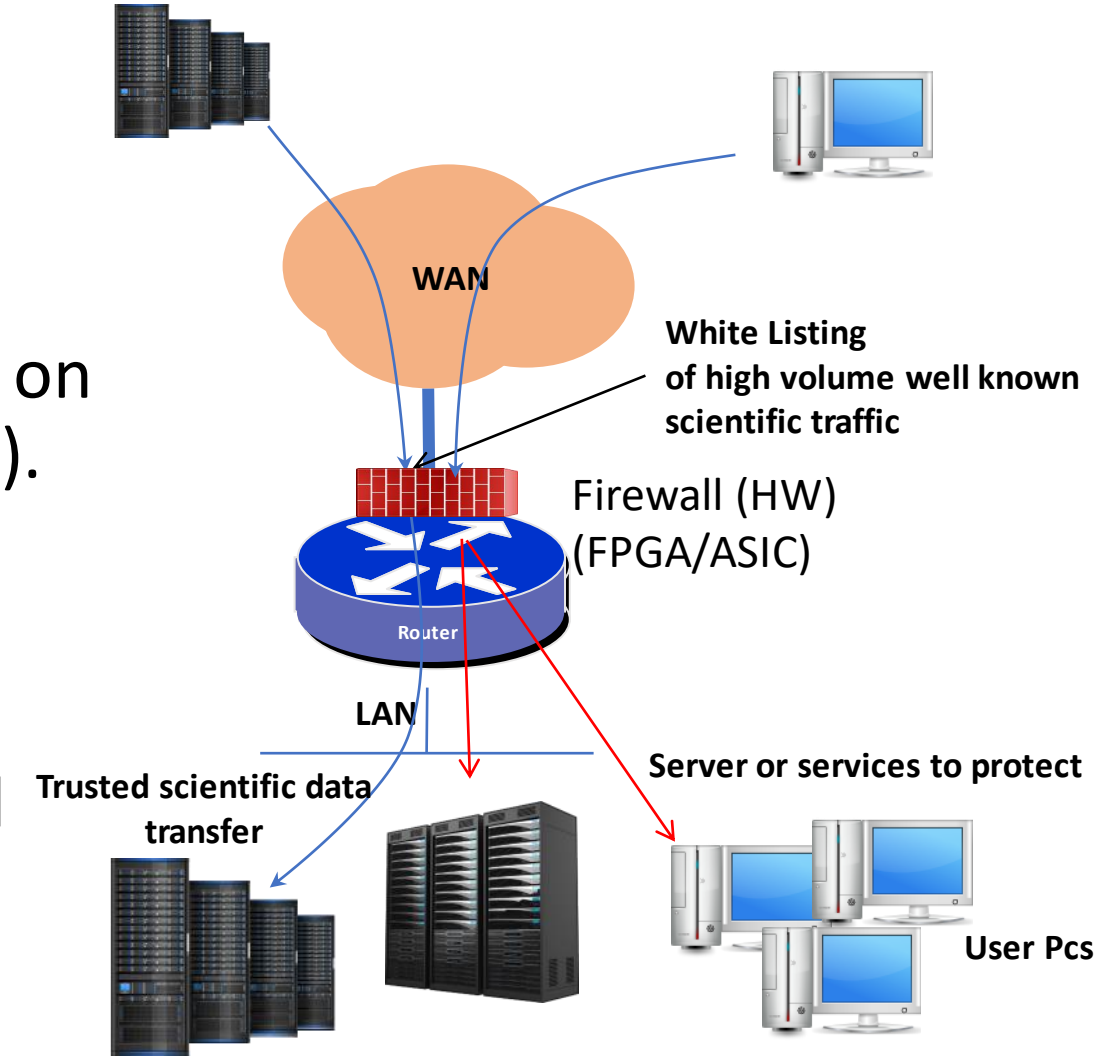

INFN Wi-Fi: radius servers and authentication

- INFN wi-fi infrastructure is based on a distributed authentication model (radius servers), using **INFN-dot1x** as SSID. Only CNAF site uses a different SSID for INFN users: CNAF-dot1x.
- An INFN user located at any INFN site associates with the access point using INFN-dot1x as the SSID and EAP-TTLS as the authentication protocol. This allows tunneling of credentials within a TLS encrypted channel.
- The username syntax is : **uid@*site*.infn.it**
- Local users are authenticated by the local radius.
- Roaming users are authenticated by their own radius server via proxy mechanism and the response will reach the access point which will allow or deny access to the network.

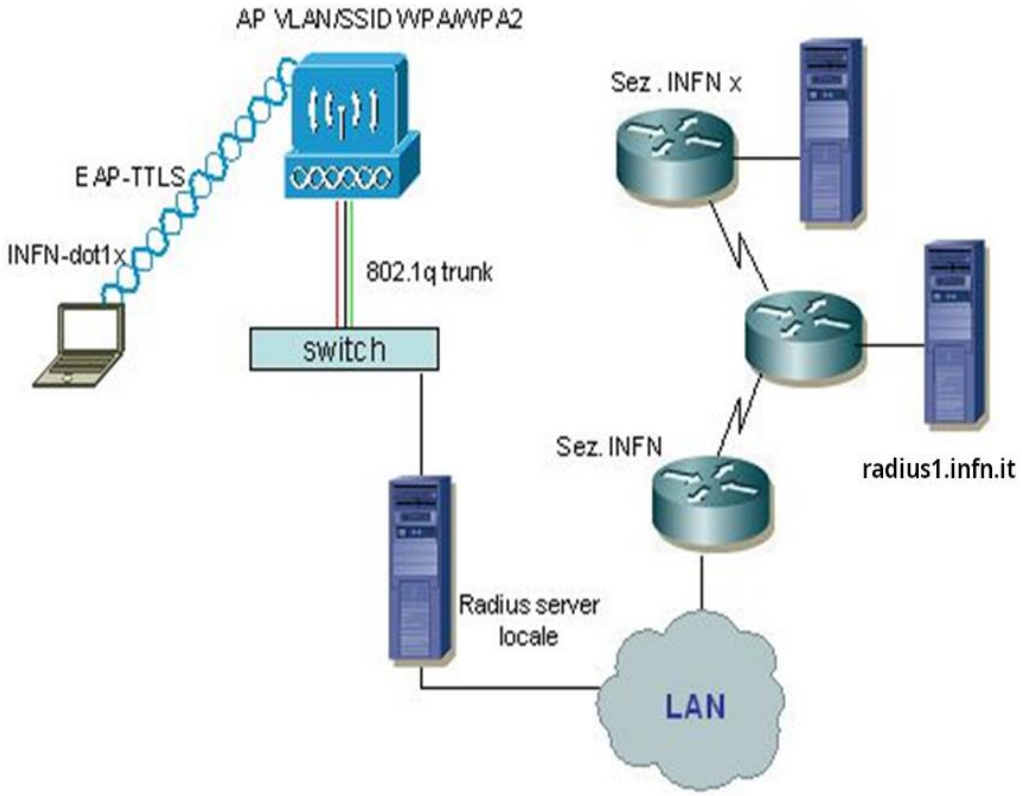

Eduroam (**Edu**cation**Roam**ing) uses the same technology but gains access to all education and research organizations included in the worldwide federation.

INFN-WEB Captive Portal

- Many INFN sites provide a captive portal authenticated SSID for temporary guests: **INFN-WEB**
- Once the guest device associates to the WIFI SSID, the user must open a web browser (some devices recognize automatically the authentication system)
- It will be redirected to a web login page connected to a local authentication system

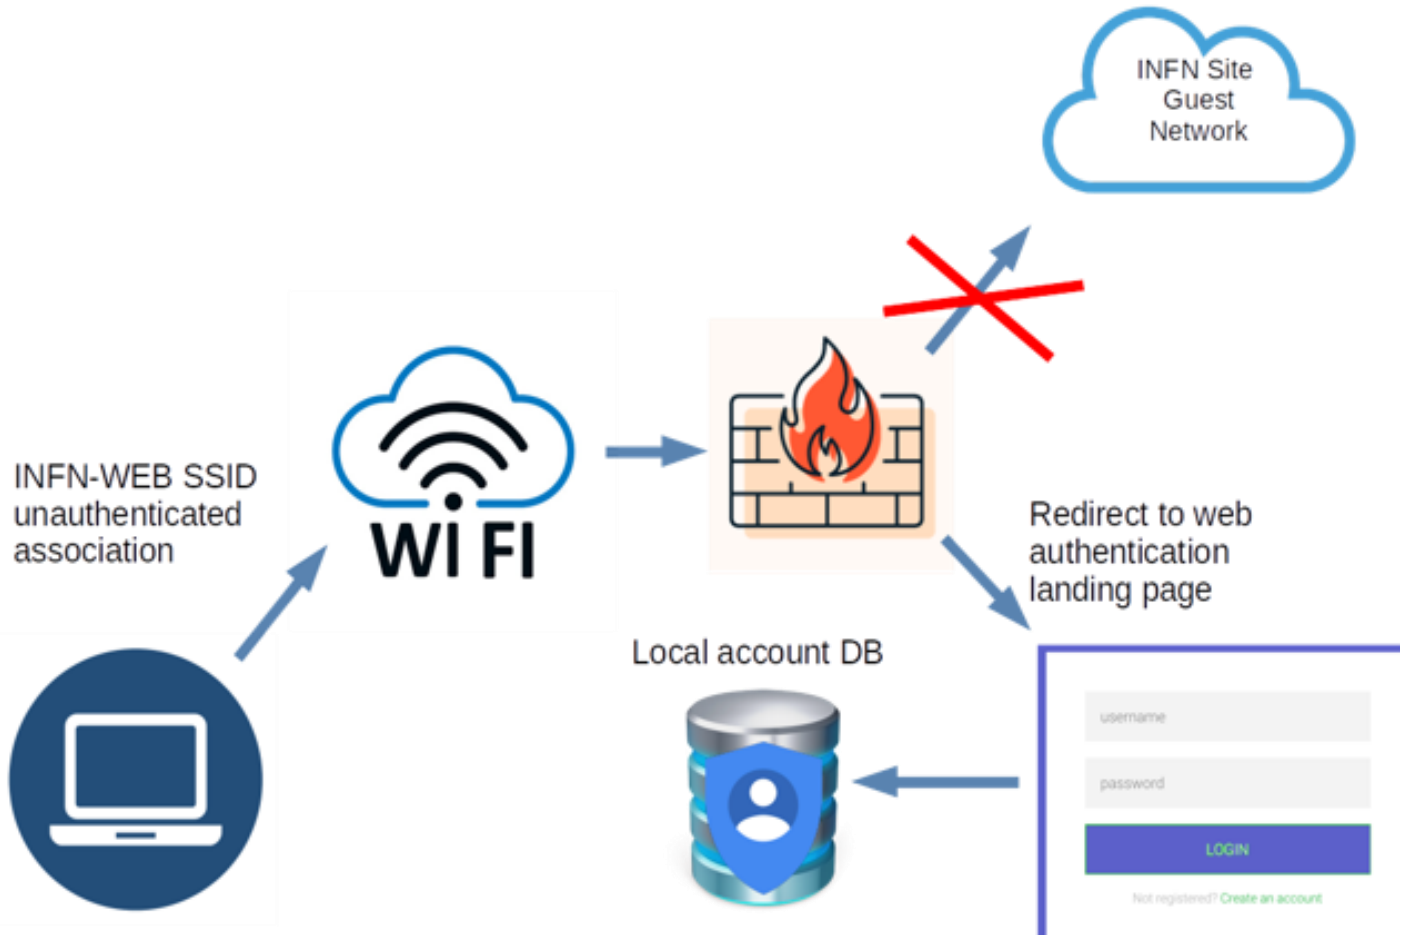

INFN-WEB Captive Portal

- If web authentication ends successfully, the client mac address is registered in the system and the firewall allows access to the site guest network

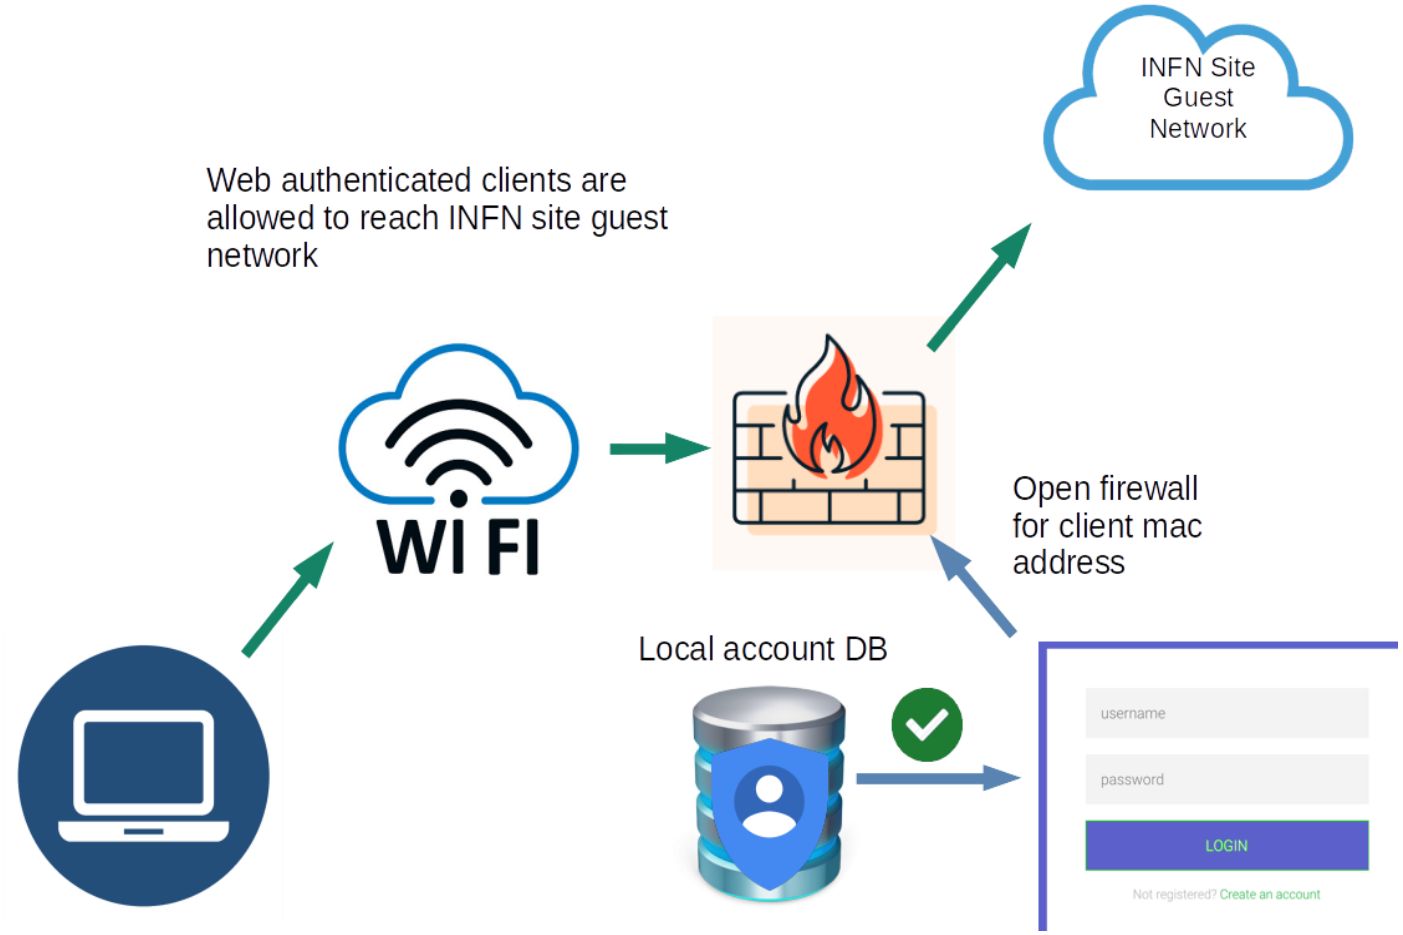

INFN Wi-Fi Summary

- **INFN-dot1x (CNAF-dot1x)**
  - INFN people only (preferred on INFN sites)
- **EDUROAM**
  - Worldwide education and research people (outside INFN sites)
- **INFN-WEB**
  - Local Guests
  - Not INFN or education people

Every INFN sites can implement local different security policies

Virtual Private Network (VPN)

- VPN is a mechanism for creating a **secure connection** between a computing device and a computer network, or between two networks, using an insecure communication medium such as the public Internet.
- A VPN can **extend a private network** and enables users send and receive data across public networks **as if the public networks' devices were directly connected to the private network.**
- The benefits of a VPN include security, reduced costs for dedicated communication lines, and greater flexibility for remote workers. Encryption is common, although not an inherent part of a VPN connection.

VPN Types

- Remote access VPN
  - to connect users from internet (Client/Server)
- Site to site
  - to stretch the LAN over multiple sites using internet (Peer to Peer)

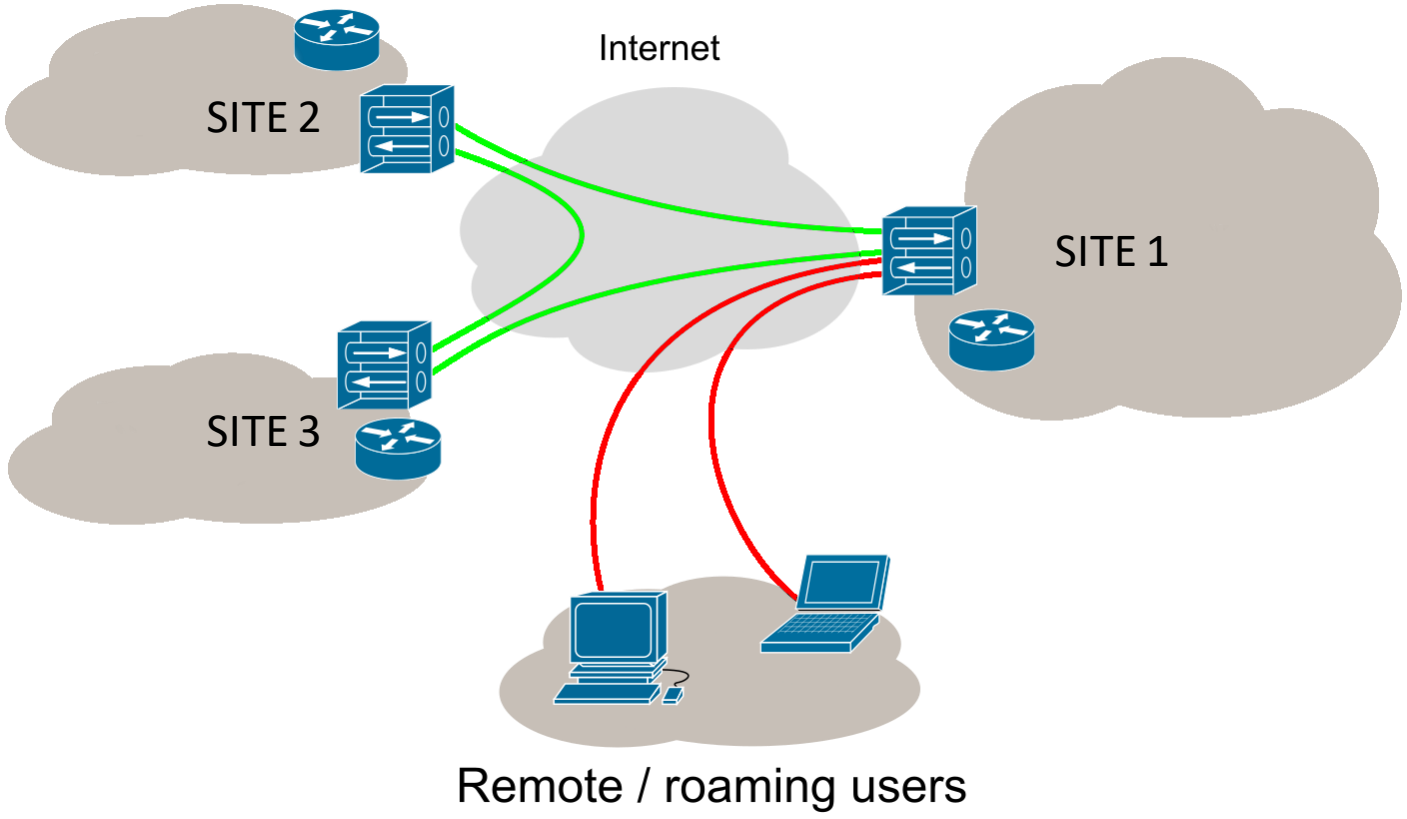

How does it work

- Authentication
  - User-created remote-access VPNs may use passwords, biometrics, two-factor authentication, or other cryptographic methods (INFN AAI or Certificate)
- Secure tunnel creation and routing rules
  - Using a secure protocol (e.g. IPsec or SSL) a “tunnel” from client to server is created and a virtual network is added to the client.
  - A routing rule is added to the OS routing table to permit:
    - All the traffic flows thru the tunnel **TUNNEL ALL**
    - Only the traffic to the LAN is routed to the tunnel and the internet traffic is routed normally aka **SPLIT TUNNEL**

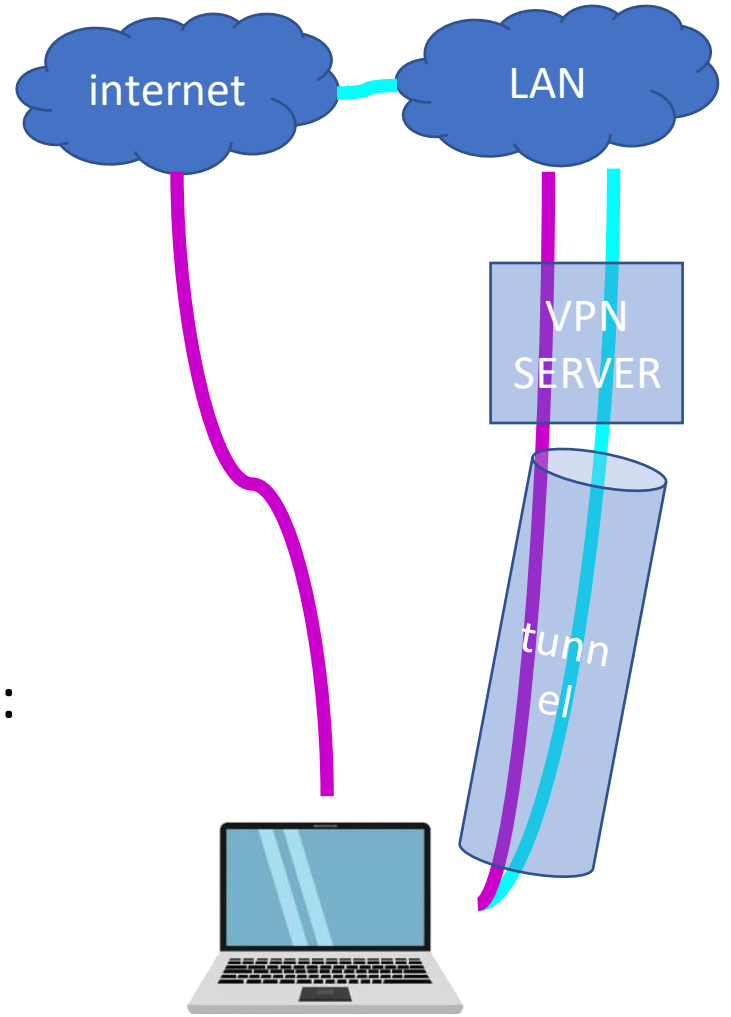

Science DMZ

WLCG (Worldwide LHC Computing Grid)

Particle Physics scientific collaborations are international and have a worldwide scope.

WLCG project is the global collaboration that includes all the computing centers (about 170) distributed in more than 40 countries providing HEP community the computing resources.

“The mission of the WLCG project is to provide global computing resources to store, distribute and analyse the ~200 Petabytes of data expected every year of operations from the Large Hadron Collider (LHC) at CERN”

The distributed computing model of the experiments and the necessity to “Move” large amount of data between different data centers have a huge impact on the network evolution.

https://wlcg.web.cern.ch/

WLCG
Worldwide LHC Computing Grid

Network challenges inside the datacenters (LAN)

- LHC Computing Centers are the typical example of **High Throughput Computing** (HTC).
- These computing centers are made of a huge number of CPUs (Standard CPUs) installed in hundreds of servers and storage systems of the order of tens of PB Petabyte of data**.**
  - **INFN CNAF TIER1 datacenter consists of about 60000 CPU cores, 50PB disk storage**
  - **INFN Bari Italian T2 consists of about: 20000 CPU cores, 20PB disk storage**
- Every Job must have access to data at high throughput to perform the necessary analysis.
- In general, the LAN of an LHC datacenter has to provide an aggregate throughput between CPU and storage of several **Tbps** (Terabit per second).

Network challenge between datacenters (WAN)

- WLCG Datacenters TIER classification
  - **TIER-0: CERN** Datacenter where data are generated
  - **TIER-1: 15 sites in 12 Countries** (Main Datacenters for data custodial and analysis)
  - **TIER-2: 140 sites in 40 Countries** for analysis and simulation
- In order to guarantee the data transfer from CERN to the TIER1s it has been created the **LHC Optical Private Network** (LHC OPN https://lhcopn.web.cern.ch/).
- Connectivity between all T2s and between T1s and T2s is provided by **LHC Open Network Environment** (LHC ONE https://lhcone.web.cern.ch/) network as an L3VPN on the Internet (But using links dedicated to scientific applications).
- Other scientific communities asked to be part of LHCONE in last 5 years and now is supporting the WAN traffic of: LHC, BELLEII, Pierre Auger Observatory, Juno, and Xenon.

LHCOPN

Every TIER1 has a dedicated high speed link to CERN.

INFN CCR

Optical Private Network

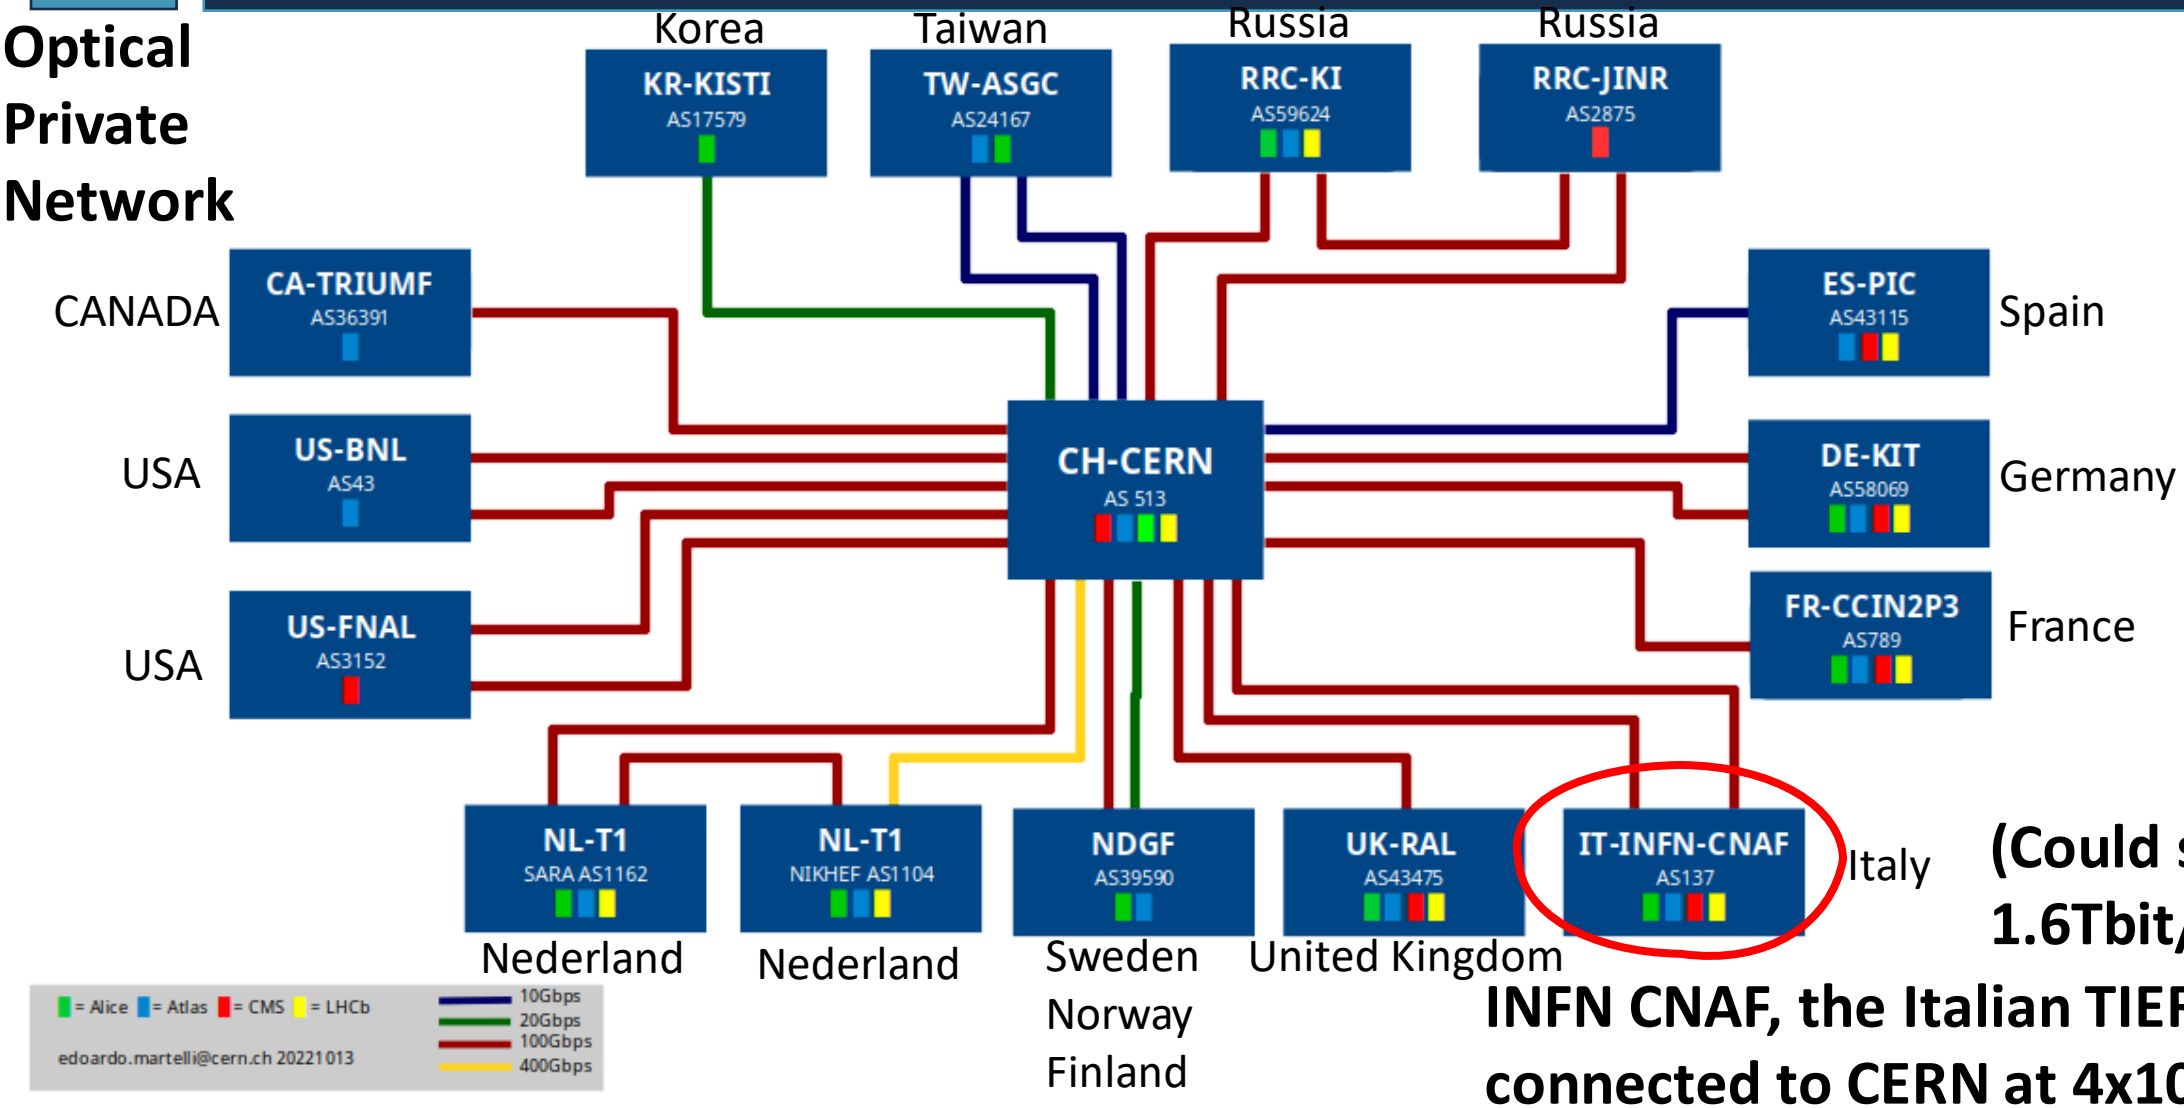

(Could scale to 1.6Tbit/s)

INFN CNAF, the Italian TIER1 is connected to CERN at 4x100Gbps

LHC ONE (Open Network Environment)

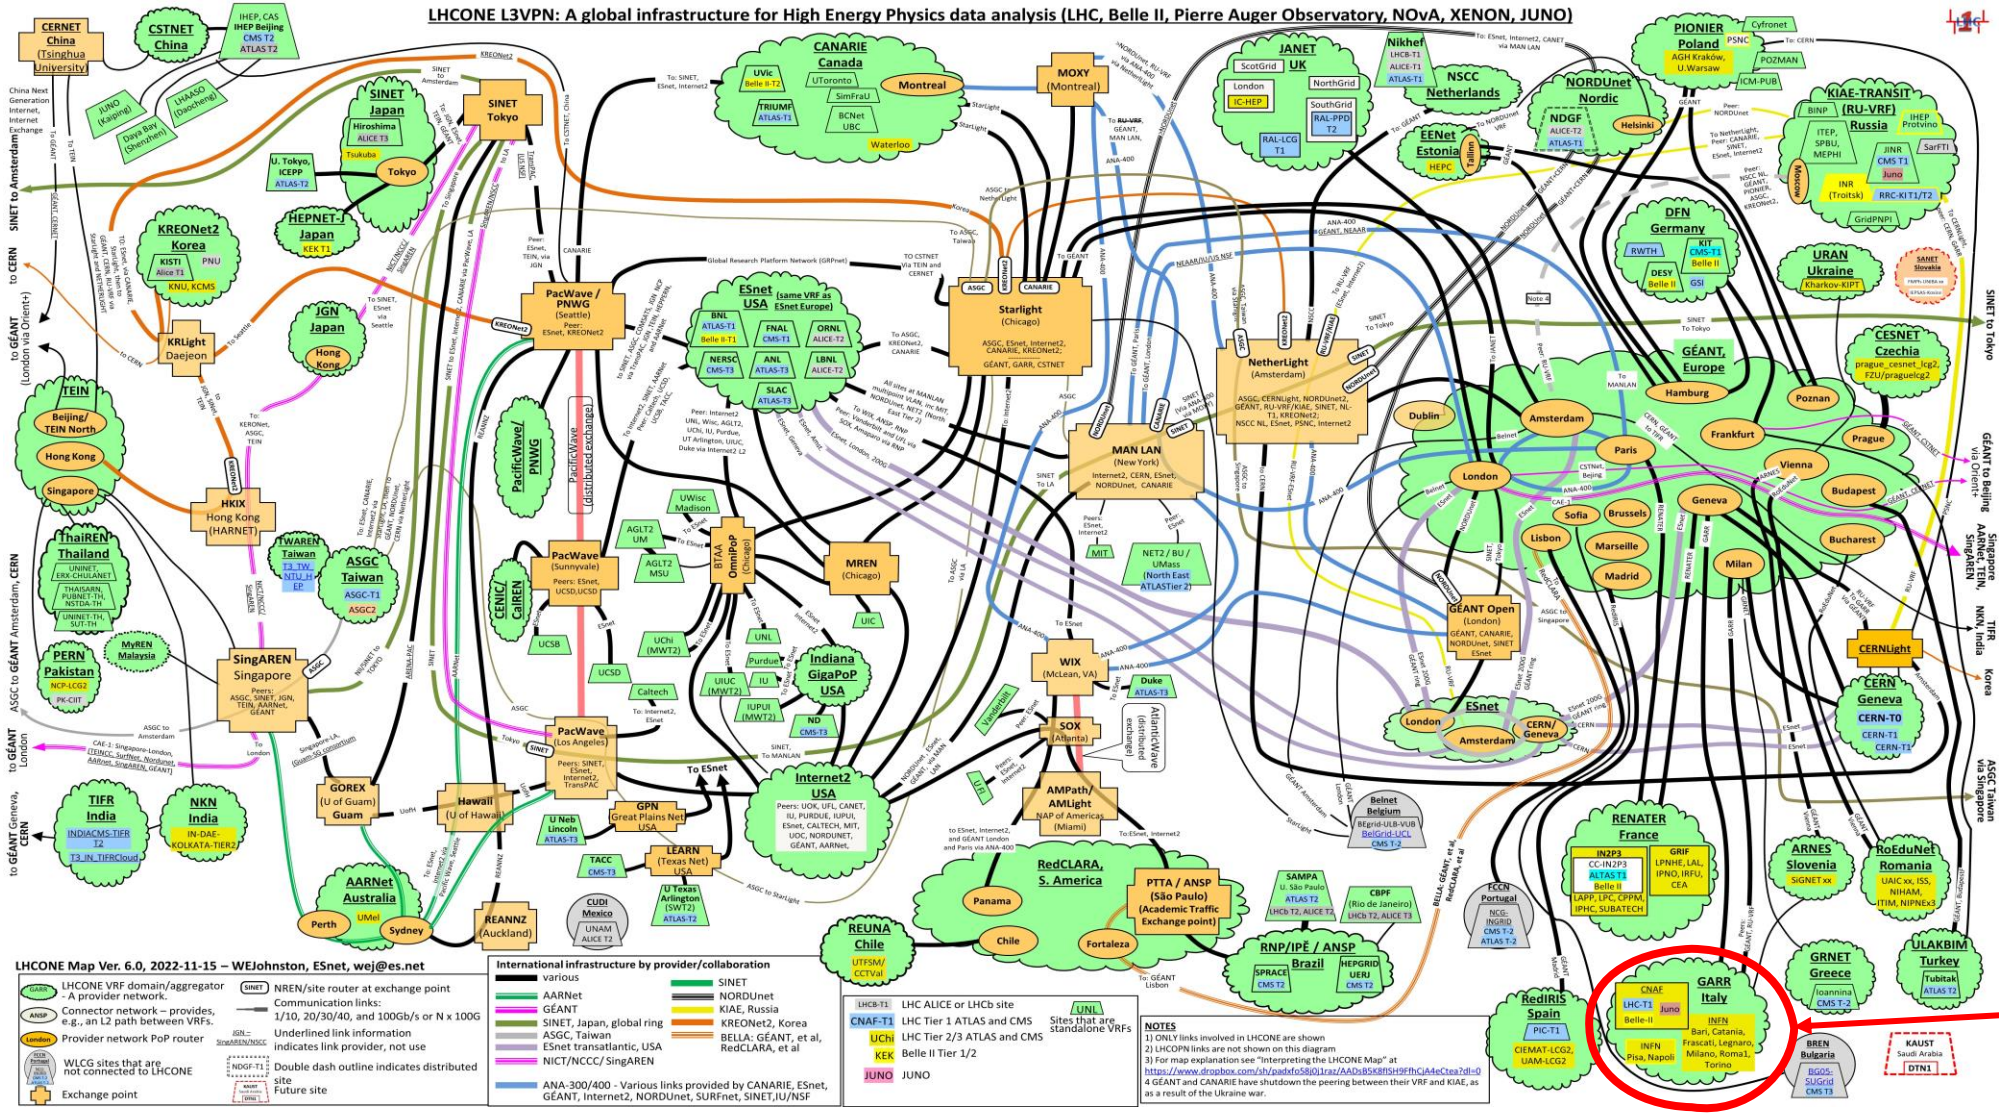

This worldwide Overlay Network is quite complex and requires the collaboration between:

1)HEP Sites (Datacenters)

2)NREN (National Research ad Education Networks)

The Italian NREN is **GARR** (www.garr.it)

You are Here!

Thanks to William Johnston (ES NET)

LHCOPN and LHCONE in Italy

Italian TIER1 and TIER2s datacenter connectivity

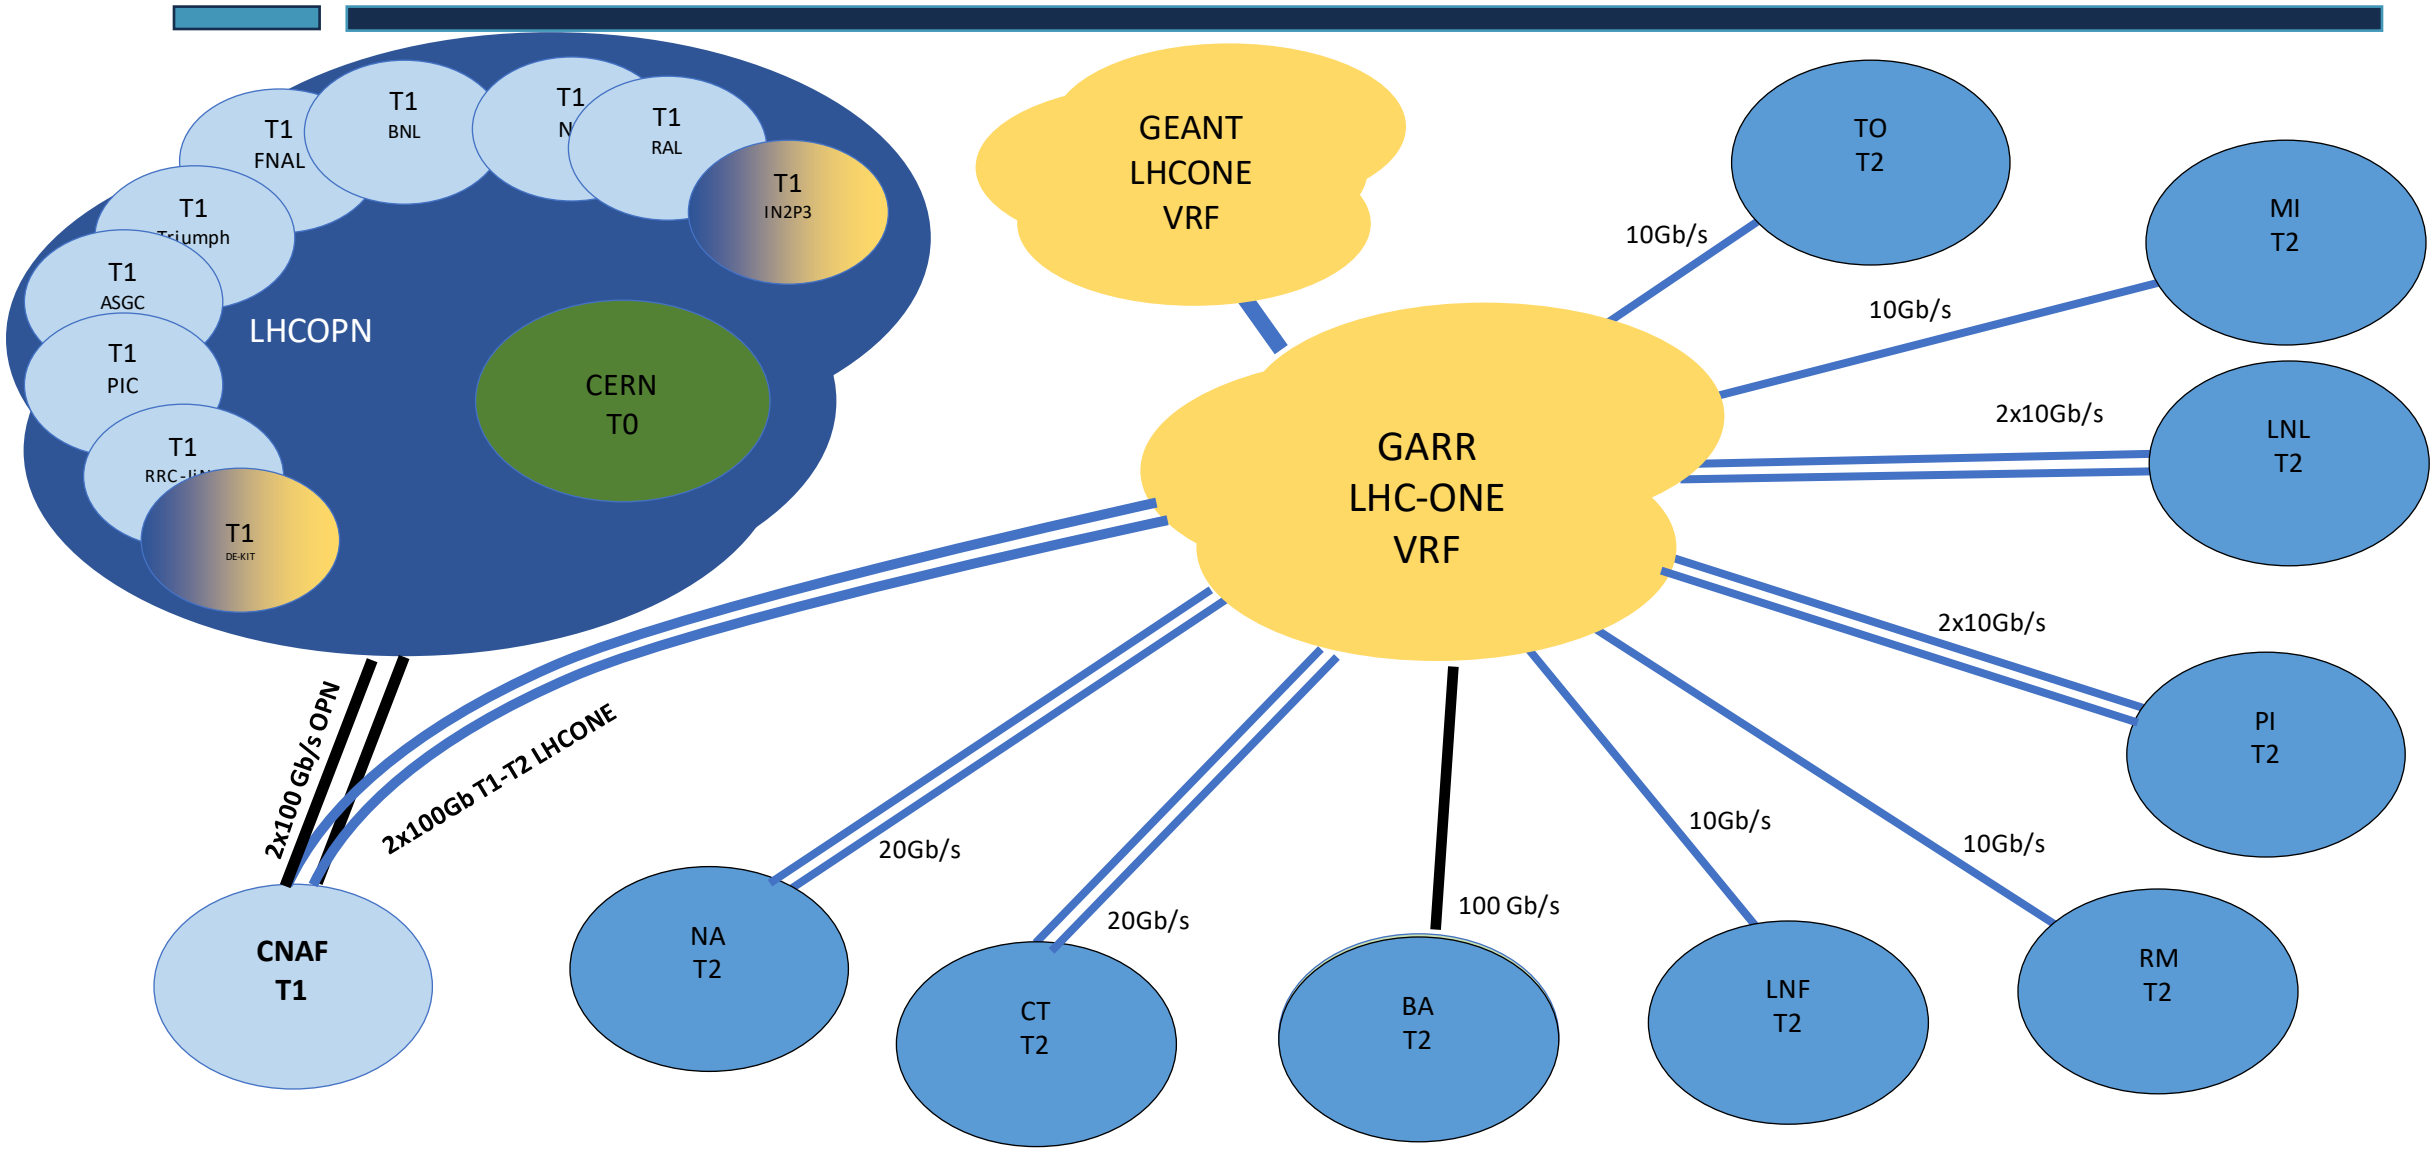

Global Collaboration for global connectivity

In order to guarantee the global connectivity many **NRENs** are involved on the management and evolution of the complex network of international links

At the time of writing this document **transatlantic links consists of about 800Gbps**

Trans-Atlantic connectivity provided by ESnet, GEANT, Internet2, NORDUnet and SURFnet

- Washinghton-London (I2+Canarie)
- London-NY (Internet2+Canarie)
- London-NY (I2+EsNet)
- Washinghton-Geneva (ES-Net)
- Paris-NY (GEANT)
- NY-Amsterdam- (NII-Sinet)
- Amsterdam-Montreal (Nordunet,SURFnet)
- Boston-Amsterdam (ES-Net)

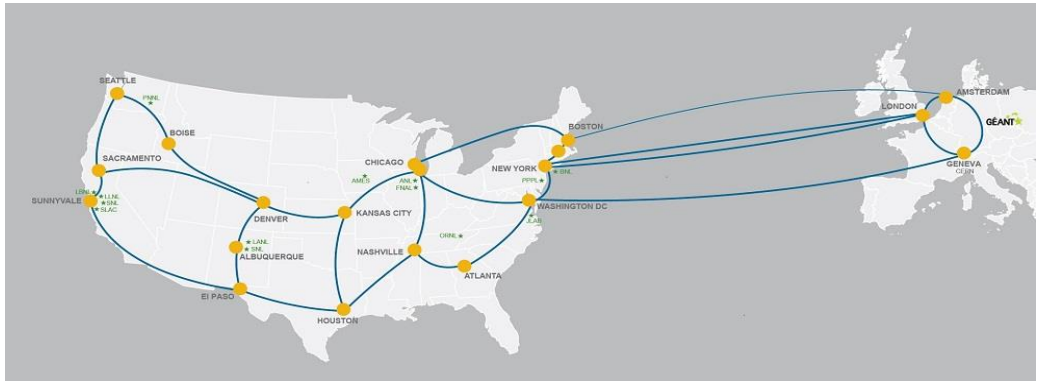

Trans-Pacific connectivity provided by ASGCnet, KREOnet, SINET, TransPAC

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

GARR (The Italian Research Network)

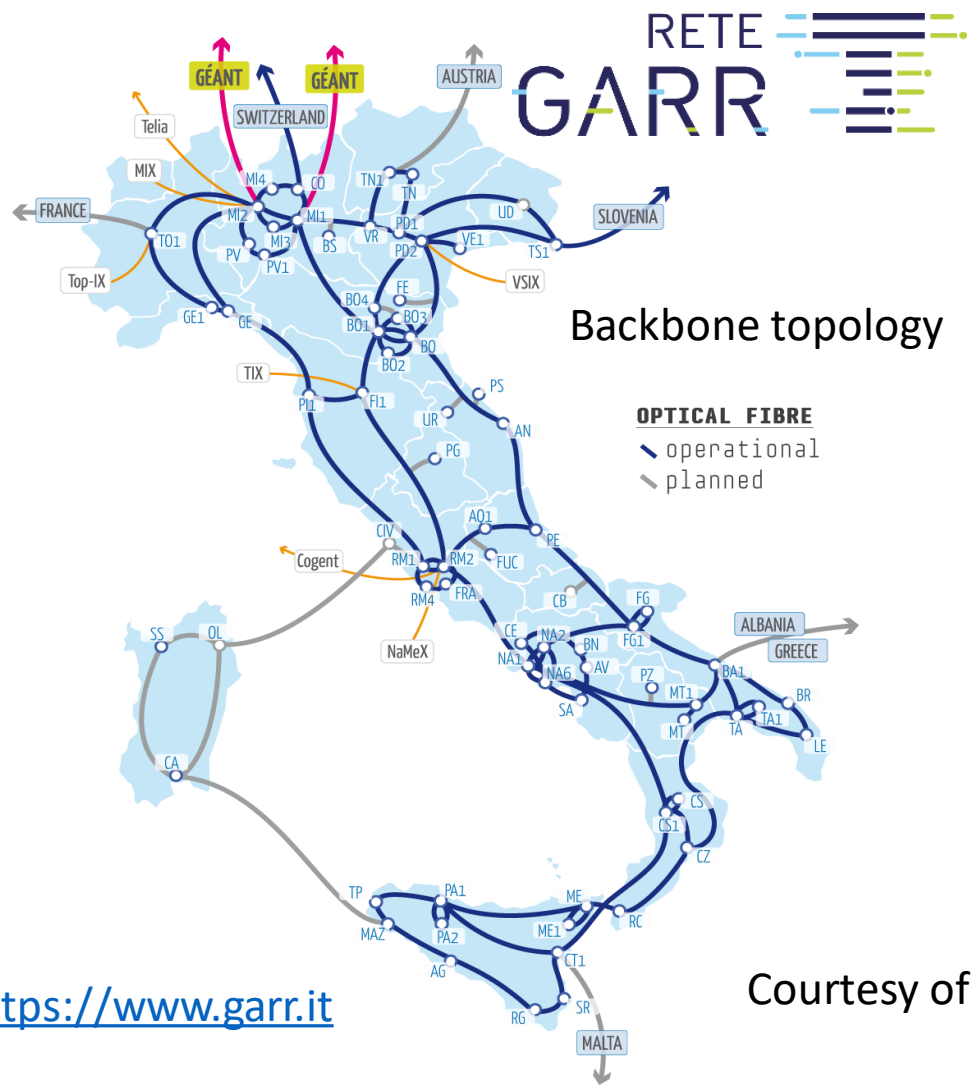

Backbone topology

Courtesy of GARR

https://www.garr.it

GARR (Gruppo Armonizzazione Rete della Ricerca) provides network access and services to about 1000 organisations

- 100 Universities
- 350 Research Institutes and Laboratories
- 60 Biomedical Research Institutes
- 65 Music Conservatories, Art Academies, Libraries, Museums & other Cultural Institutions
- 500 Schools

The Backbone is multi Terabit capable and It is made of more than 100 PoPs (Point of Presence)

GARR is part of the **European Research Network** GÉANT

GARR-T (Next generation network)

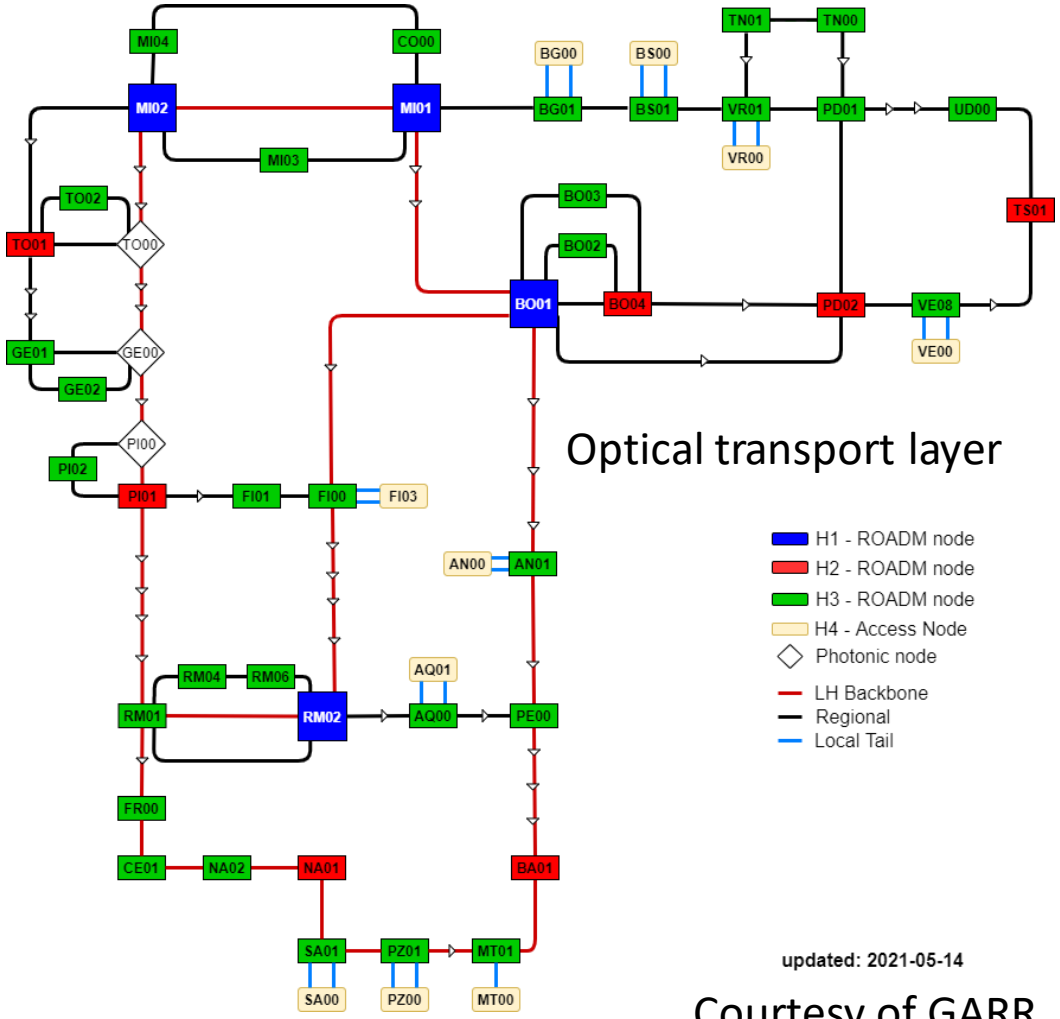

Optical transport layer

- H1 - ROADM node
- H2 - ROADM node
- H3 - ROADM node
- H4 - Access Node
- Photonic node
- LH Backbone
- Regional
- Local Tail

updated: 2021-05-14

Courtesy of GARR

Packet Layer

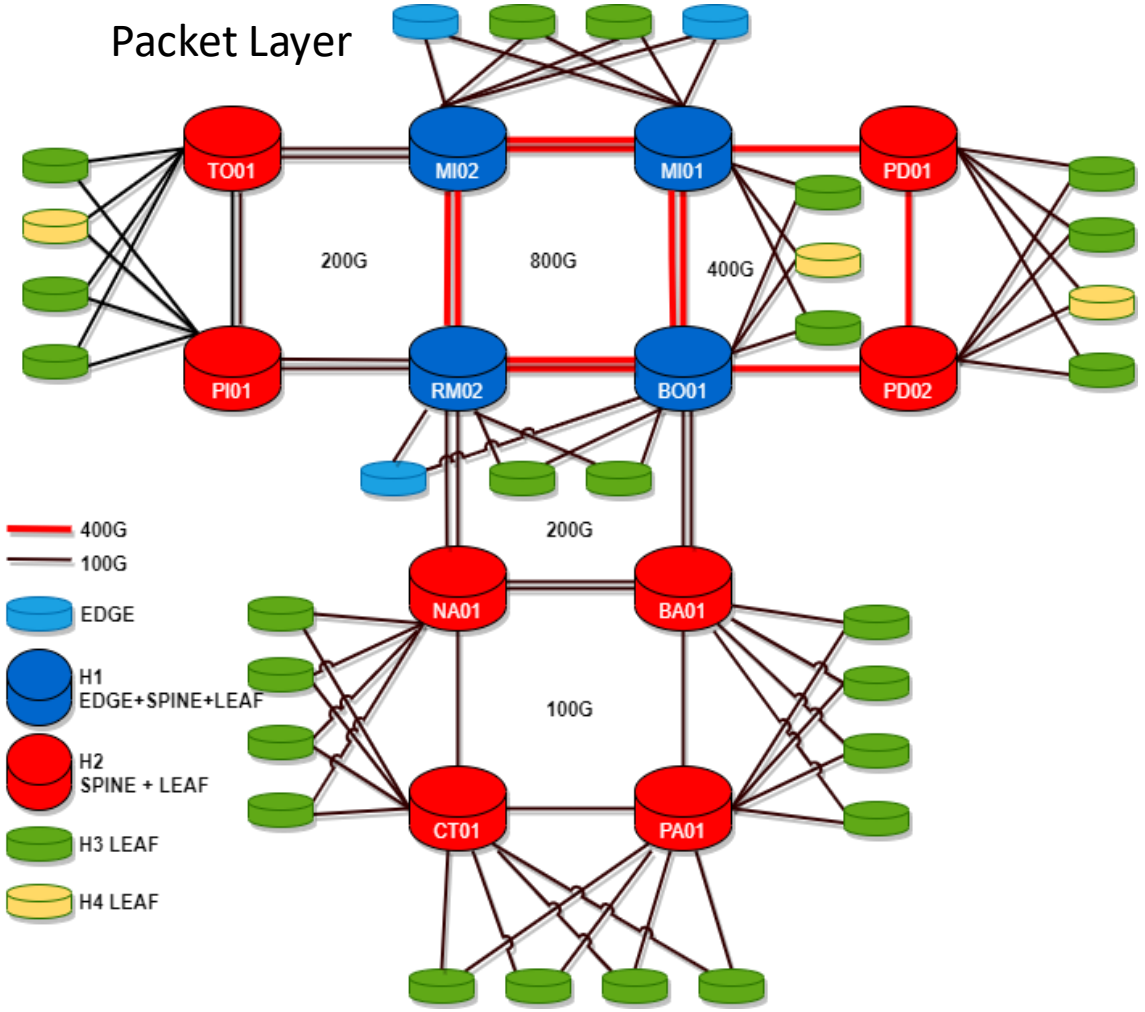

GARR-T some numbers

Courtesy of GARR

OPTICAL Network

Fiber network	**6155 km**
New optical fiber infrastructure	**740 km**
line system directions	**128**
PoP ROADM	**42**
In Line Amplifier (ILA)	**35**
New Metro POP	**9**
Doubling POP in some City	**6**
100GEth Services	**130**
400GEth Services	**11**
Day1 backbone capacity	**17.4 Tbps**

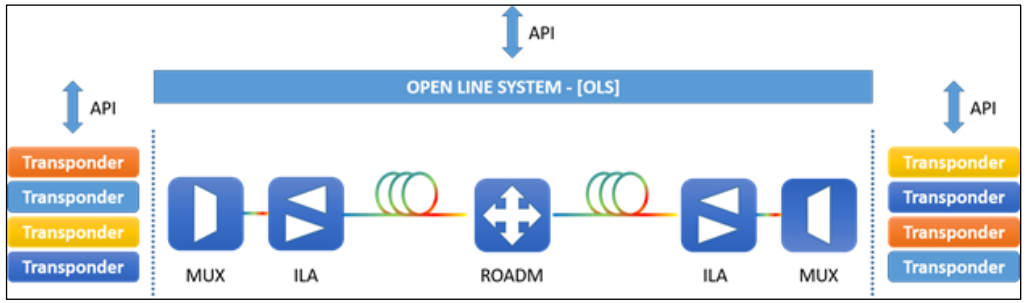

PACKET Network

Device Capacity	Capacity MAX	Day 1 Capacity
Total LEAF/EDGE/CSD	144.6 Tbps	61.28 Tbps
Total SPINE	307.2 Tbps	108.8 Tbps

Backbone					End user Acces		
400G	100G	40G	10G	1G	100G	10G	1G
34	**622**	**32**	**265**	**59**	**67**	**290**	**876**

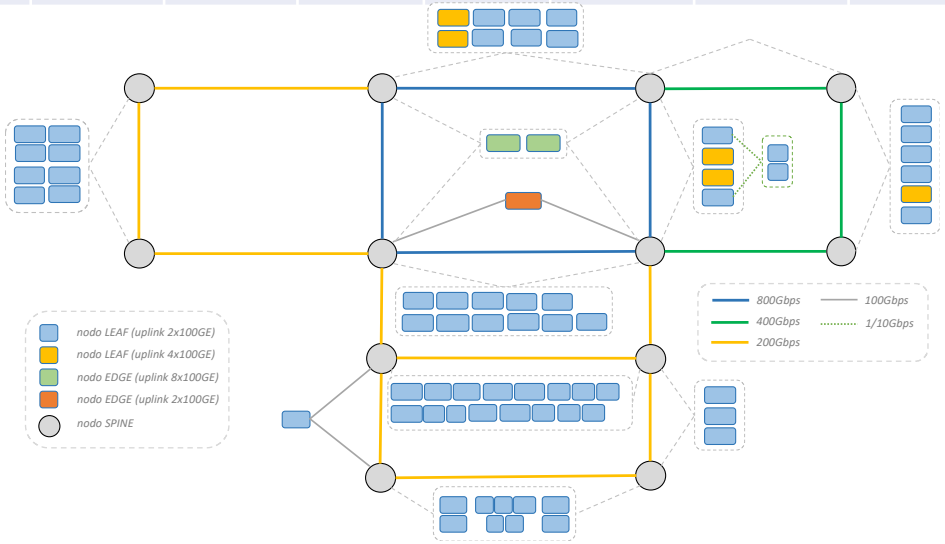

GARR Accepable Use Policy (AUP)

An acceptable use policy (AUP) is a document stipulating constraints and practices that a user must agree to for access to a network, the internet or other resources.

e.g., *"All users to whom access to the GARR Network and its services shall be provided must be recognized and identifiable."*

"Responsibility for the content of materials being produced and disseminated using the network and its services is attributable to the persons producing and disseminating said materials."

"damaging, destroying, or seeking unauthorized access to data, or violating other users' confidentiality, including interception or dissemination of passwords, confidential cryptographic codes, and any other personal data, as it is defined by legislation pertaining to privacy protection"

https://www.garr.it/it/regole-di-utilizzo-della-rete-aup

https://www.garr.it/en/acceptable-use-policies

Science DMZ- Accepable Use Policy

As it is considered a scientific DMZ with very high bandwidth uplinks there are some rules to respect in order to take part to LHCONE

- The OPN/ONE access must be distinct from the General Internet access
- An agreement on LHCONE Acceptable Use Policy (AUP)
  - Traffic injected into the LHCONE must originate only from addresses that belong to a LHCONE prefix;
  - Traffic injected into the LHCONE must be directed only to addresses that belong to a LHCONE prefix.
  - Compliance with WLCG Security policies
  - Not all links could be protected by a NG Firewall

  https://twiki.cern.ch/twiki/bin/view/LHCONE/LhcOneAup

New Technologies

DataCenter Interconnection
and
Spectrum sharing

Optical DCI (General Concept)

A possible solution to implement a Tbps DCI between two remote datacenters is the use of packet/optical **transponders** in the sites connected through an Open Line System transport infrastructure provided by the Network Research Networks.

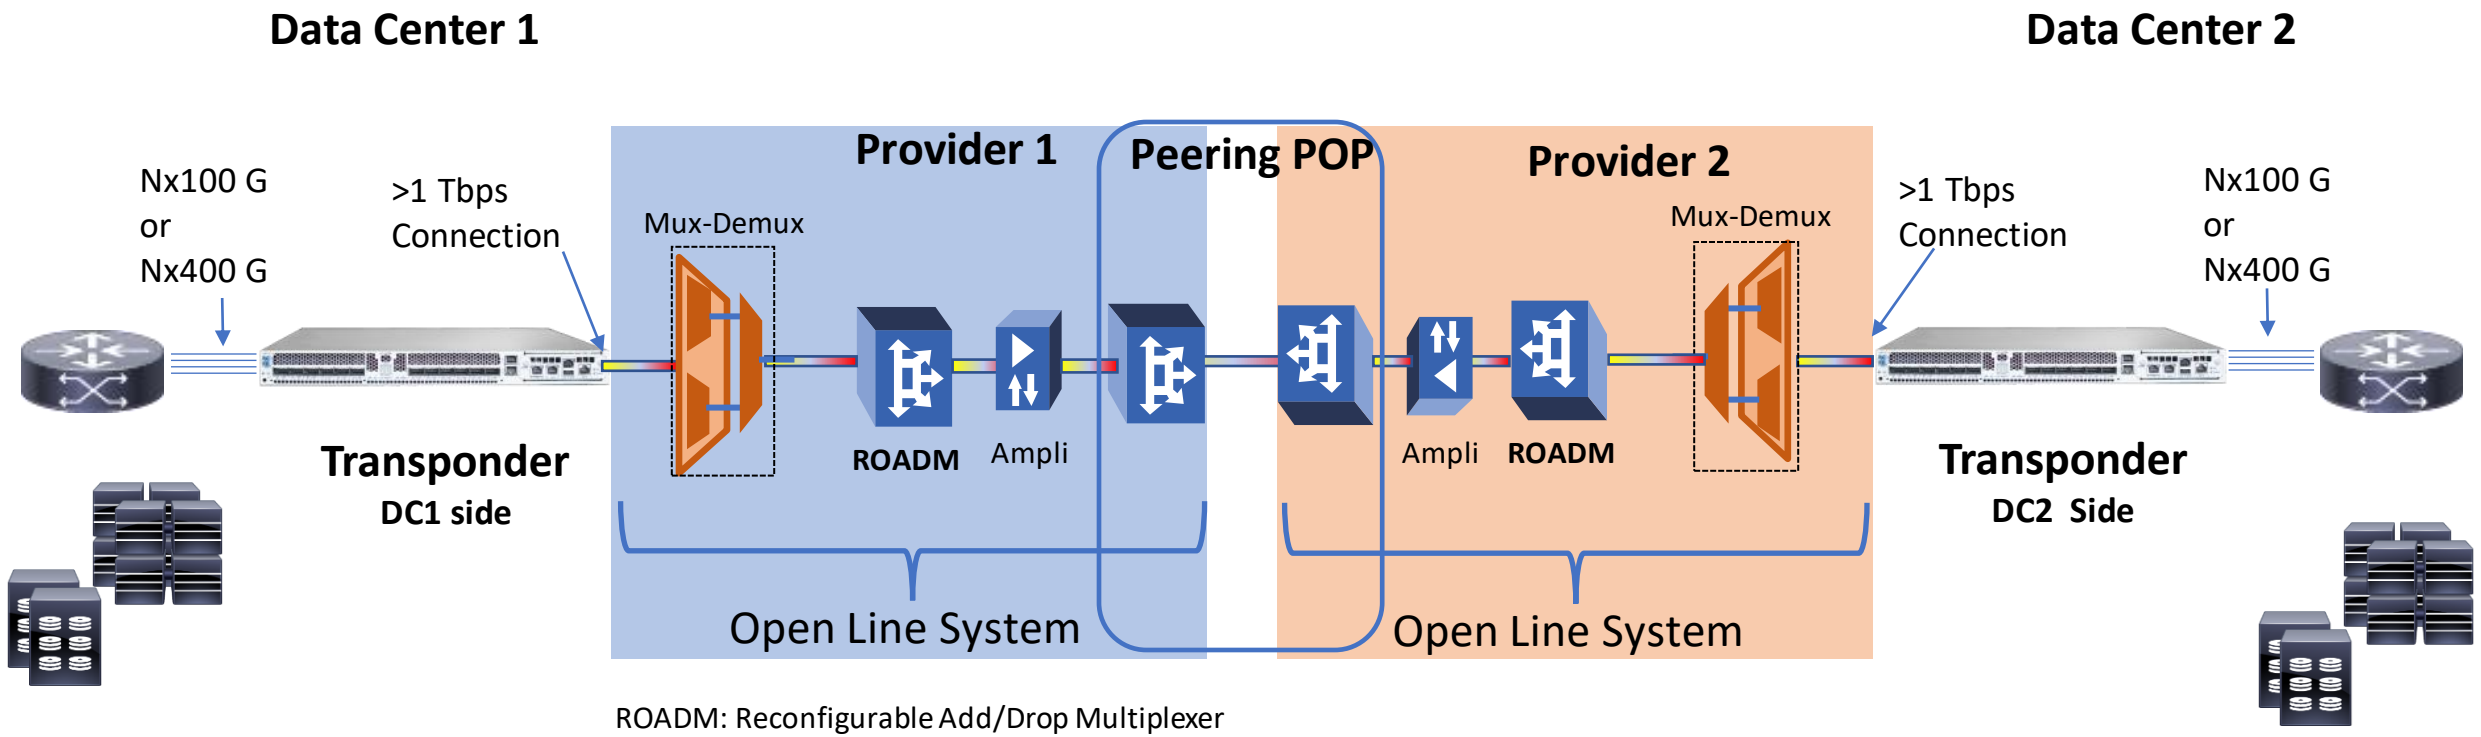

ROADM: Reconfigurable Add/Drop Multiplexer

Data Center Interconnection (DCI)

Physical Elements of an optical DCI

Enabling Technologies

Optical Network (disaggregated model):

- **Open Optical Line System**
  - ROADM *(Reconfigurable Optical Add/Drop Multiplexer)*
  - Amplifier
  - Multiplexer/Demultiplexer

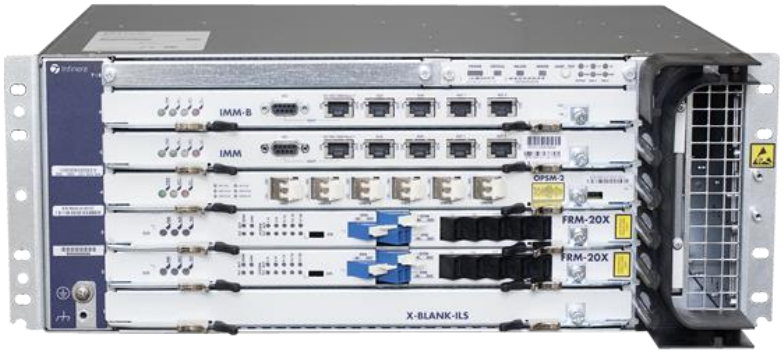

- **Transponder boxes** (manageable with SMNP,gRPC) with:
  - Optical interfaces on line side
  - Packet interfaces (100/400G Ethernet) on LAN side

- **Pluggable coherent transceivers** (For example ZR e ZR+) able to directly connect to DWDM line systems

It has been proposed in WP7-T2 di GN4-3 (GEANT) a multi domain **Spectrum Connection Service (SCS)** that allows the creation of high speed "Optical Links" between sites belonging to different National Research Networks.
CERN-CNAF has been identified as a good use case (~1000 Km fiber distance)

DCI CERN-CNAF

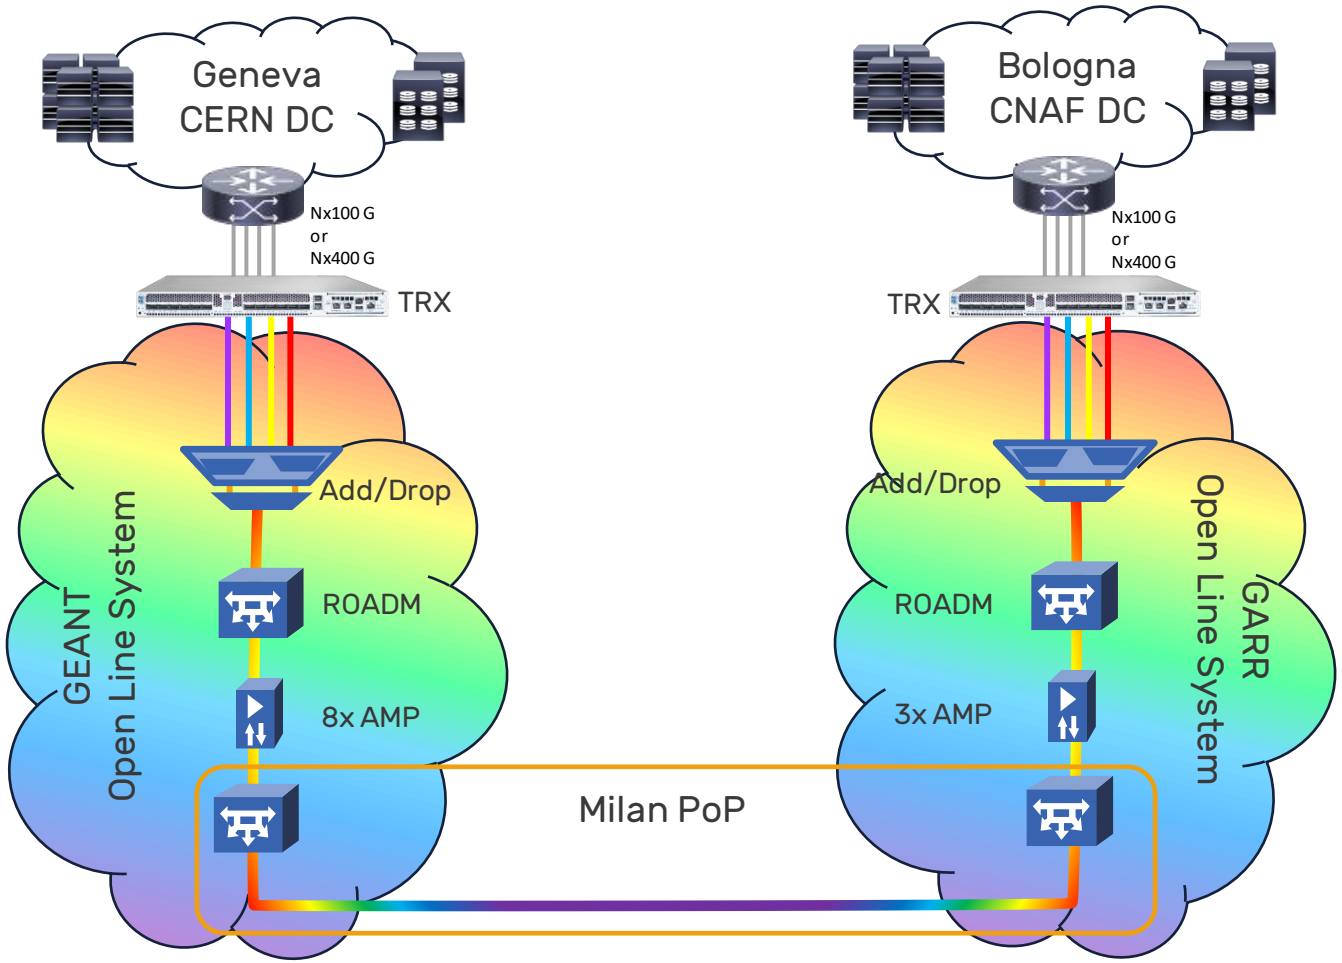

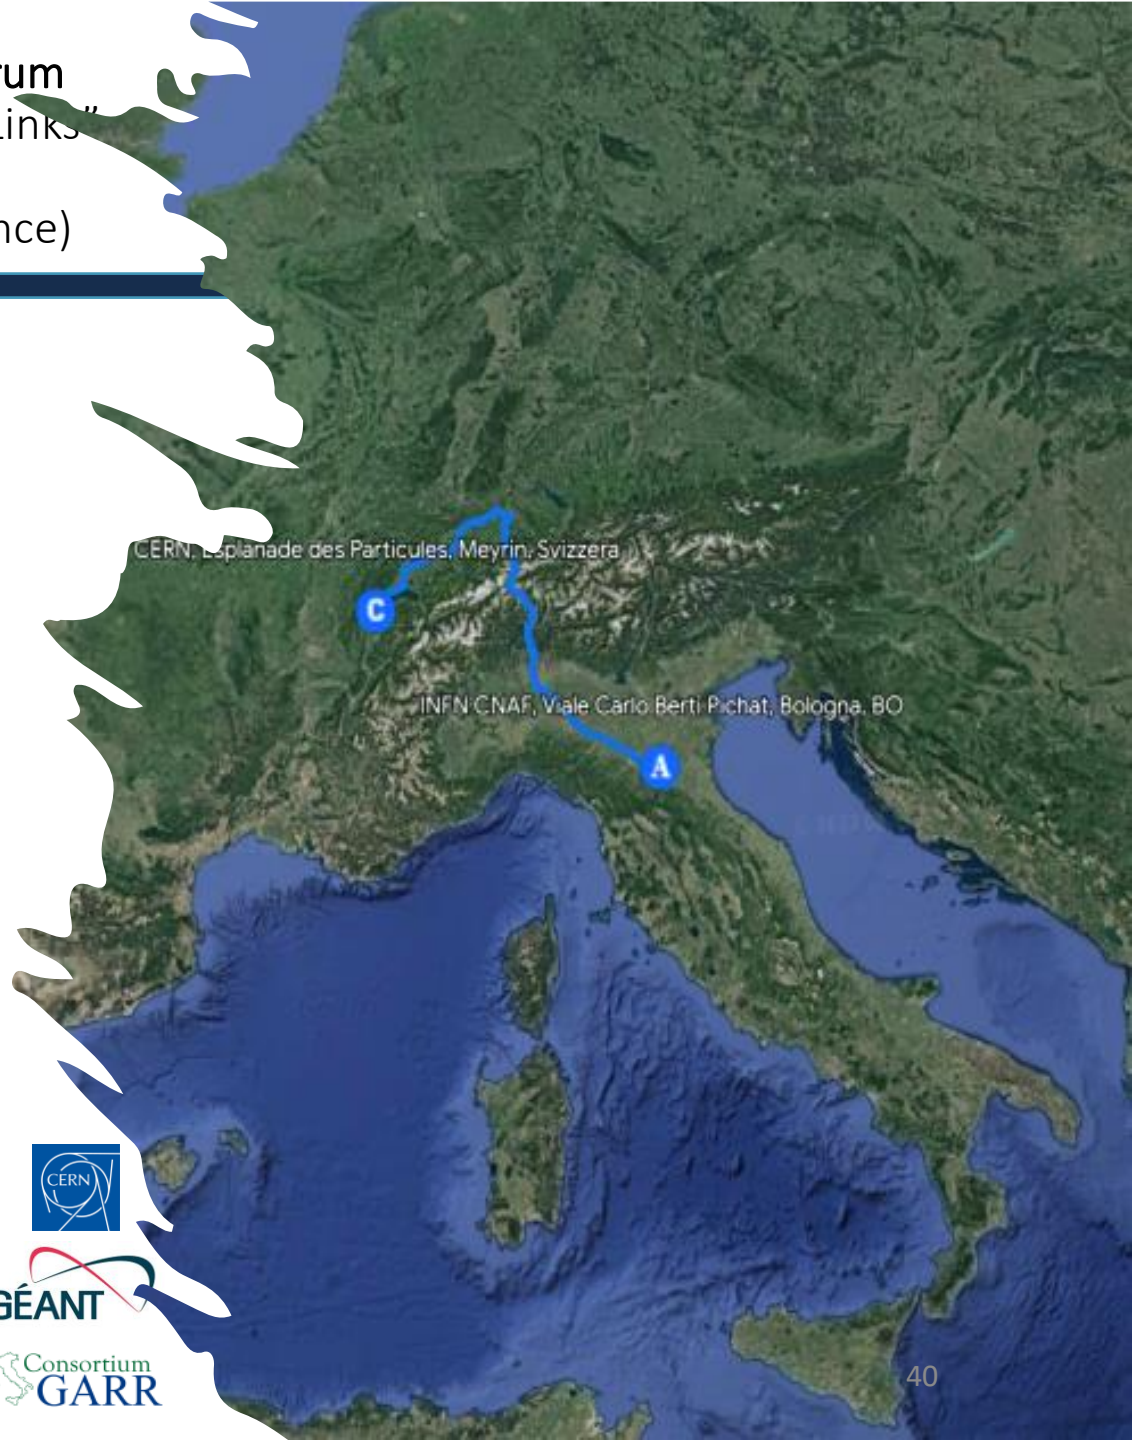

CERN
GÉANT
Consortium GARR

Current state of "pilot" DCI (CERN-CNAF)
1,6 Tbps potentially available today

Istituto Nazionale di Fisica Nucleare
Commissione Calcolo e Reti

Now:
Setup completed in both sites (CERN e CNAF).
We used a 16 QAM* /69 Gbaud with 100GHz spectrum for each channel on the distance of ~1000 km.
With this setup we can reach 400Gbps on every line port and having 2 CHM2T boards we already have 1.6Tbps using 400GHz that is the 10% of the whole C-Band extended frequency (4.8THz)
On client side we can have connected CERN Routers at 4x100Gbit to route all the OPN traffic and we could scale up to 1.6Tbit using 16 100Gps or 4 400Gpbs ethernet ports.

On February WLCG data challenge has reached peaks over 300Gbit/s

The old connection via GARR routers with 2x100Gbit is now the backup connection

*Quadrature Amplitude Modulation

Considerations

Direct connections with optical circuits have several advantages in terms of latency and costs but there is no resiliency because it's an end-to-end circuit and there's no automatic re-routing.

Backup connection are highly recommended considering that can occur up to 5 fibre-cut per year on 1000km links

	Packet switching	Optical Circuit
Latency	Higher	Lower
Jitter	Higher	Lower
Capacity	Statistically muxed	Dedicated
Security	IP-address based	Almost for free
Scalability	Lower	Higher
Consumption	Higher	Lower
Provisioning	Days (if available)	Months

It's not true if we are operating on an existing line system based on overprovisioned fiber circuits.

Next generation datacenter networks

(Special Thanks to Giancarlo Viola and Nino Ciurleo @GARR)

Datacenter Network Tier 1 @CNAF

Is based on CISCO Virtual Port Channel (aka MLAG VLT VLAG …) enables two switches (cores) to be one single L2 switch but not at L3 (they still be two different routers).

- HSRP active-active guarantee the distributed routing

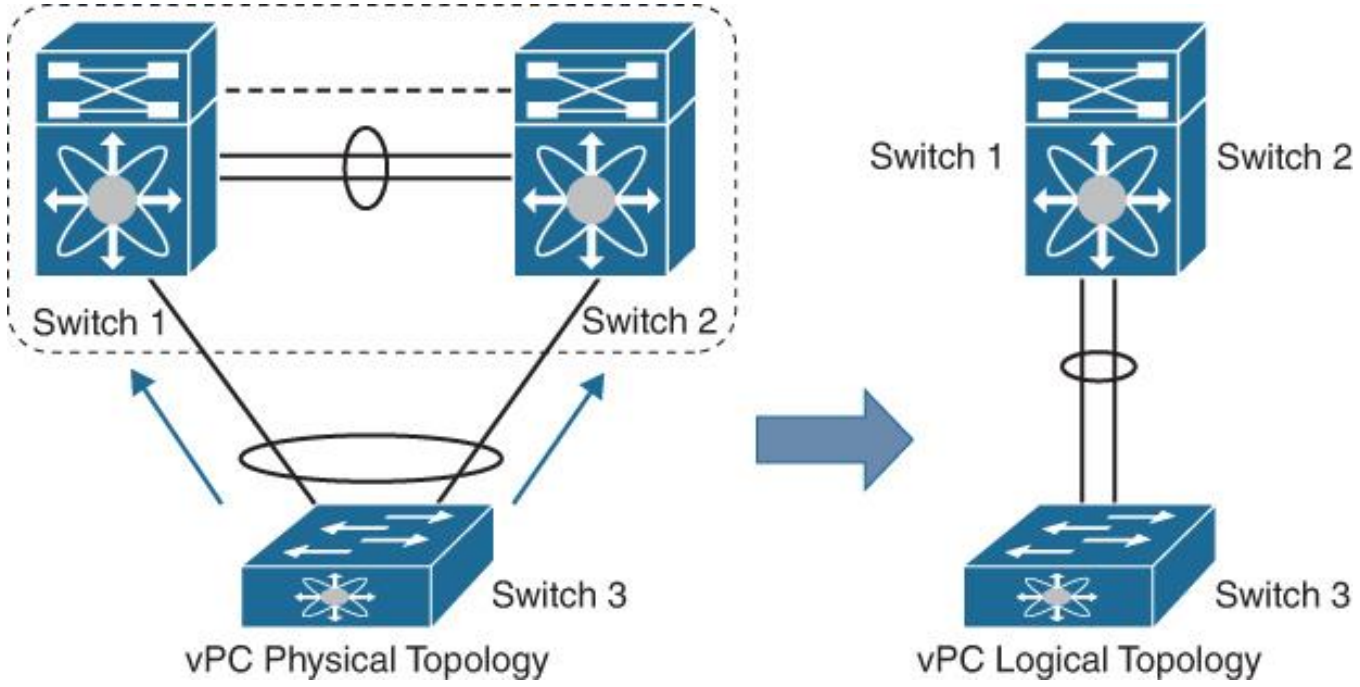

Tier1 @CNAF

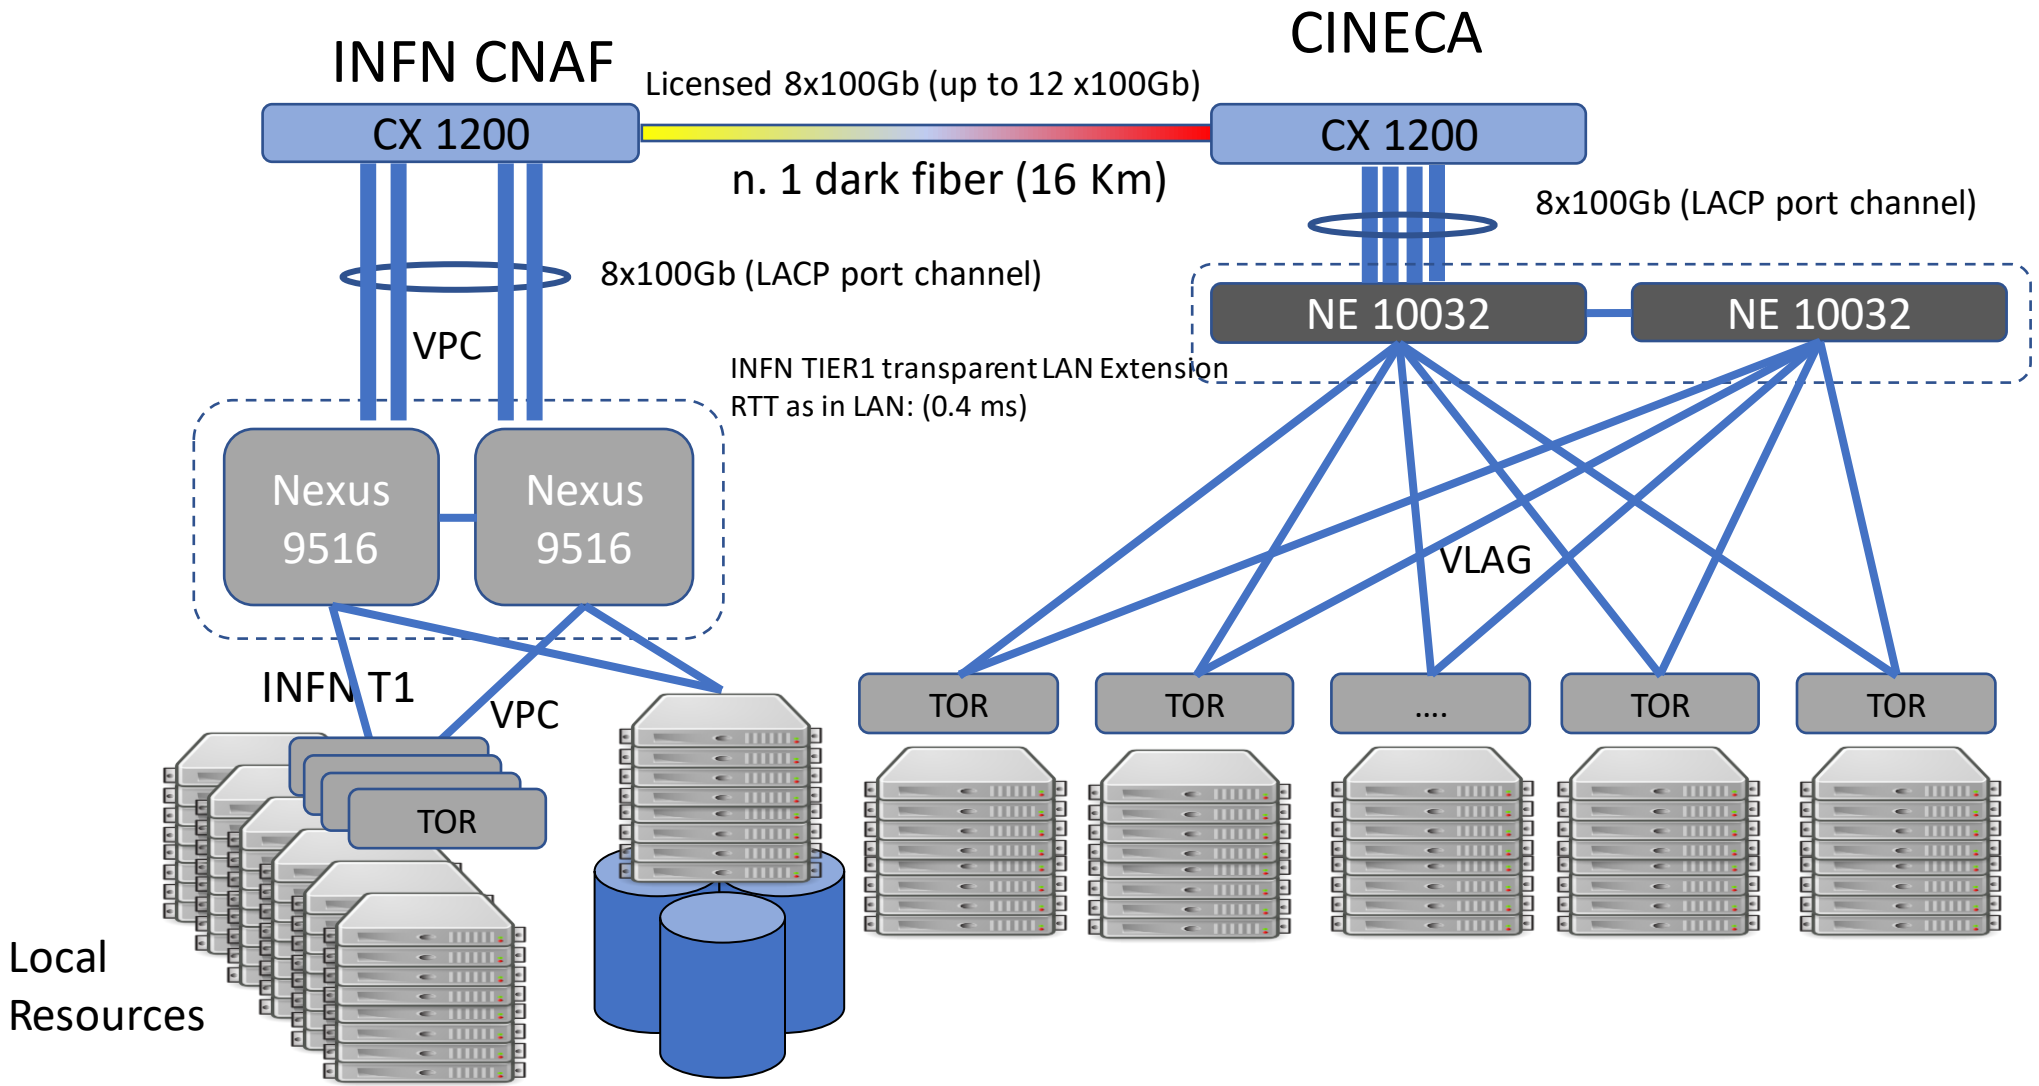
INFN CNAF
CINECA
Licensed 8x100Gb (up to 12 x100Gb)
CX 1200
CX 1200
n. 1 dark fiber (16 Km)
8x100Gb (LACP port channel)
8x100Gb (LACP port channel)
VPC
NE 10032
NE 10032
INFN TIER1 transparent LAN Extension
RTT as in LAN: (0.4 ms)
Nexus 9516
Nexus 9516
VLAG
INFN T1
VPC
TOR
TOR
TOR
....
TOR
TOR
TOR
Local Resources

Spine – Leaf topology + IP Fabric

To scale out the Ethernet fabric limits (MLAG) and VLAN number IP Fabric can be implemented (is it worth it?)

- Loop “Free”
- Scalability
- Redundancy
- Flexibility
- Distributed Routing
- Expensive
- High Complexity

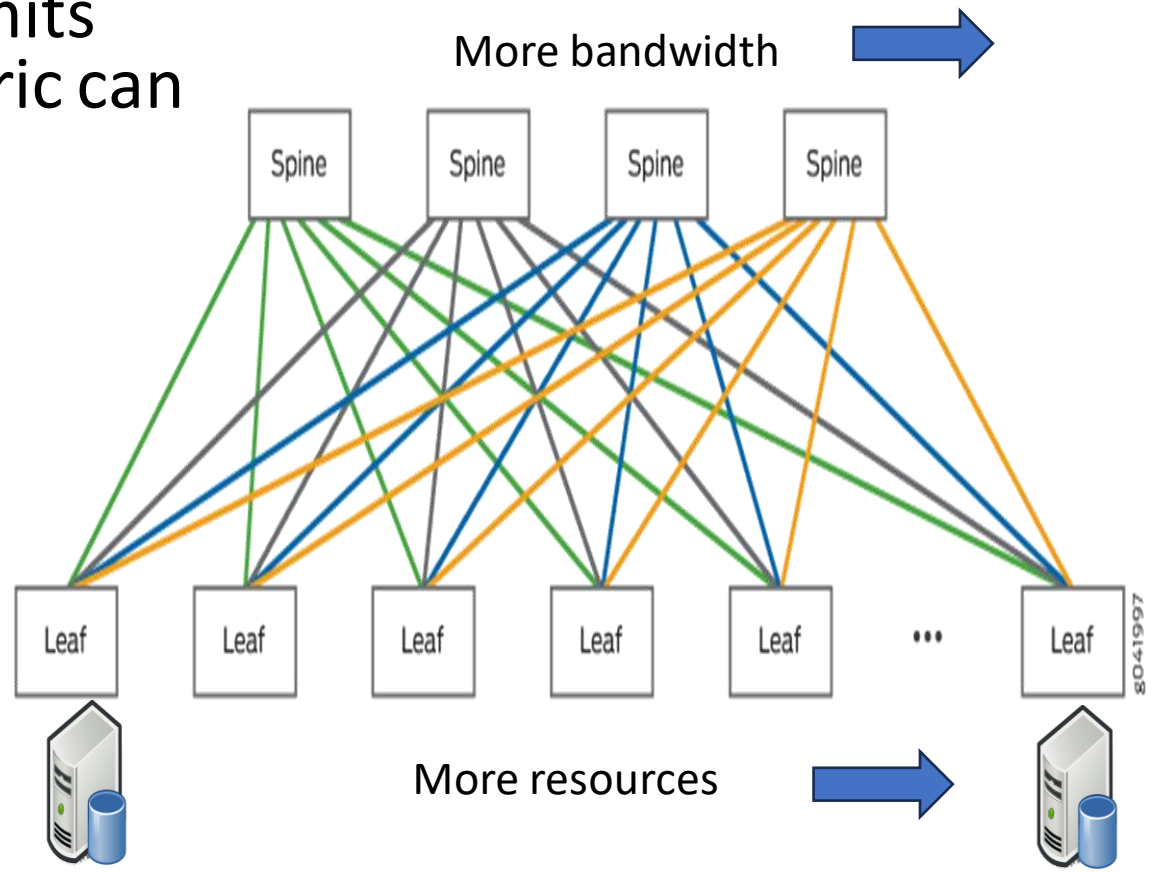

IP Fabric: WAN on LAN

IP fabric is based on WAN protocols applied to LAN and consists of two main building blocks

- **Underlay network**
  - The primary network that enables the IP reachability of network elements and where the services can be implemented
    - IS-IS OSPF BGP
- **Overlay network**
  - The network that “simulates” the L2 layer on top the L3 underlay network
    - **Control Plane**
      - BGP + EVPN
      - Mac learning -> Mac address routing
    - **Data Plane**
      - VXLAN

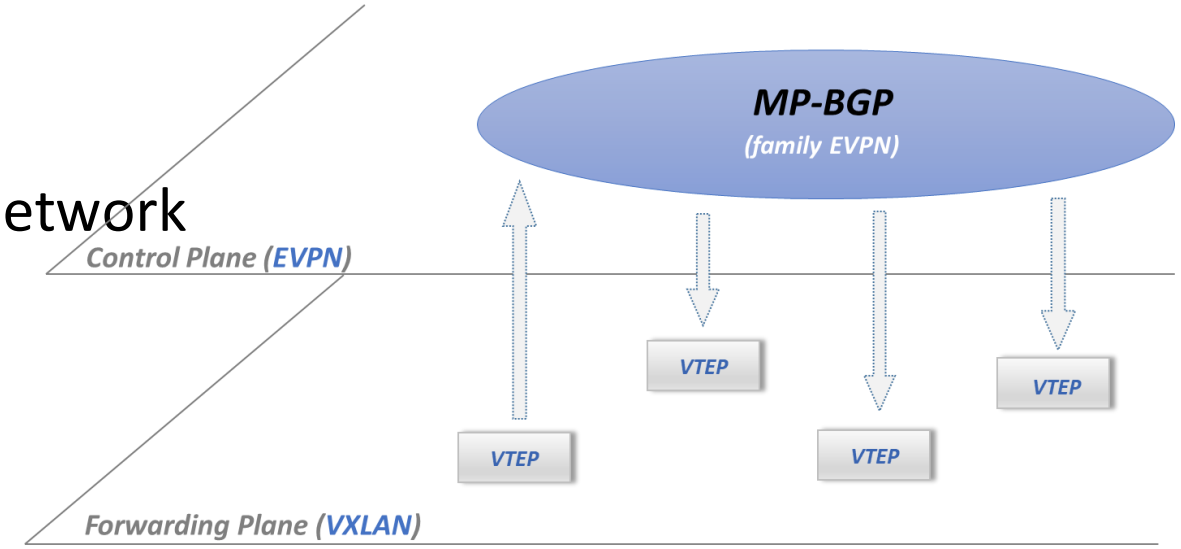

IP Fabric: Underlay

- The best practice is to define a loopback interface on every network element involved in the IP fabric this interface identifies the network element.
- Choose an IGP: Underlay network is can be based on IS-IS OSPF or BGP.
- Build a full mesh using loopback interfaces
  - for example, if you choose BGP you will have all to all (Point to Point) BGP sessions and every network device announces his own loopback.

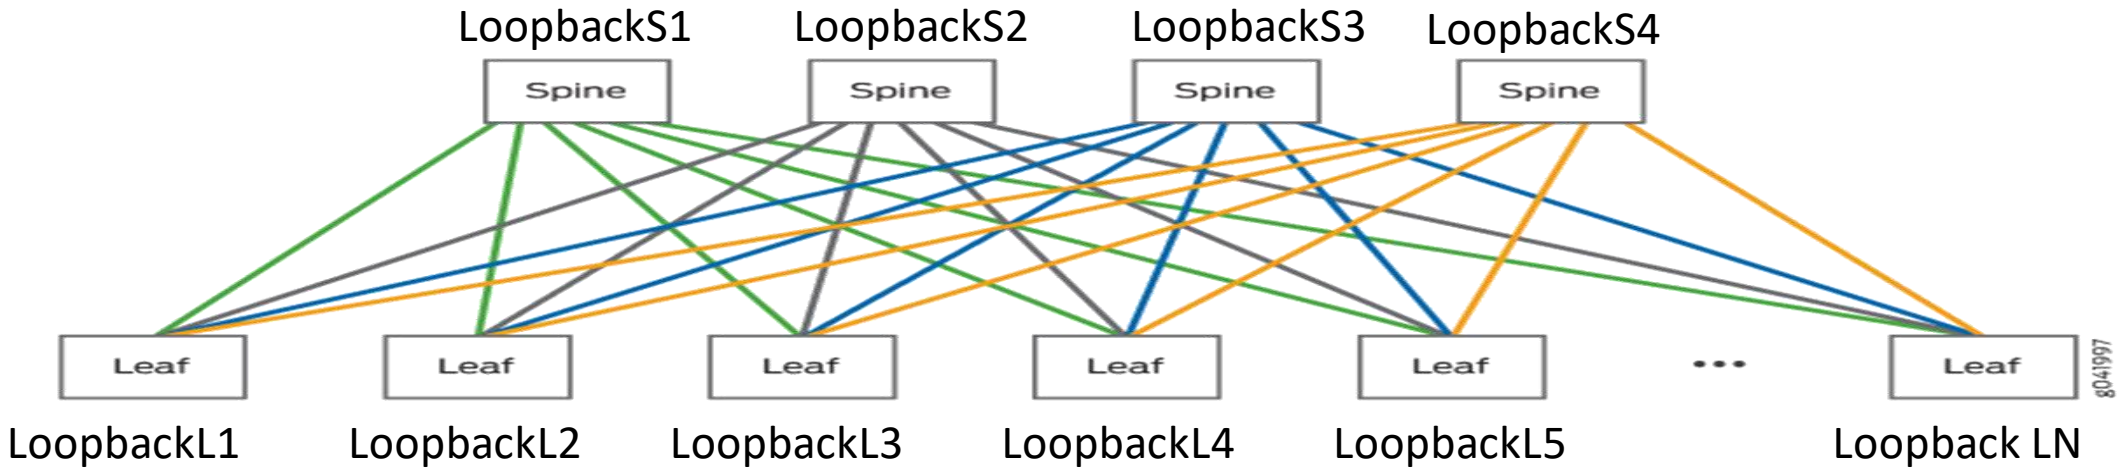

IP Fabric: Underlay

```
show ip route vrf default

 B I      10.1.20.1/32 [200/0] via 192.168.108.131, Ethernet45
                               via 192.168.108.133, Ethernet46
 B I      10.1.20.2/32 [200/0] via 192.168.108.131, Ethernet45
                               via 192.168.108.133, Ethernet46
 B I      10.1.20.3/32 [200/0] via 192.168.108.131, Ethernet45
                               via 192.168.108.133, Ethernet46
 B I      10.1.20.4/32 [200/0] via 192.168.108.131, Ethernet45
                               via 192.168.108.133, Ethernet46
 B I      10.1.20.5/32 [200/0] via 192.168.108.131, Ethernet45
                               via 192.168.108.133, Ethernet46
 B I      10.1.20.6/32 [200/0] via 192.168.108.131, Ethernet45
                               via 192.168.108.133, Ethernet46
```

IP Fabric: Underlay

```
BGP summary information for VRF default
Router identifier 10.8.20.3, local AS number 64743
Neighbor            AS Session State AFI/SAFI     AFI/SAFI State   NLRI Rcd   NLRI Acc
---------- ----------- ------------- ------------- ----------- ---------- ----------
10.8.20.64   64743 Established   IPv4 Unicast   Negotiated   98          98
10.8.20.64   64743 Established   IPv6 Unicast   Advertised   0           0
10.8.20.64   64743 Established   L2VPN EVPN     Negotiated   6775        6775
10.8.20.65   64743 Established   IPv4 Unicast   Negotiated   98          98
10.8.20.65   64743 Established   IPv6 Unicast   Advertised   0           0
10.8.20.65   64743 Established   L2VPN EVPN     Negotiated   6775        6775
```

Overlay Network the Control Plane

The control plane is implemented by BGP adding EVPN family that is a BGP extension that enables the distribution MAC addresses, Ips and more…

Via specific BGP announcements hosts belonging to the same LAN can communicate each others in a “simulated” layer2

- ARP suppression (BUM traffic reduction)
- Distributed default gateway
- Datacenter Stretching (can be implemented geographically)

Overlay Network the Data Plane

The **data plane** is implemented by **Virtual eXtensible LAN (VXLAN)** that is is a mechanism where Ethernet Frames are encapsulated inside UDP packets and sent through the overlay network a tunnel established on the underlay network. This is also called MAC-In-UDP and is implemented by a VXLAN Tunnel Endpoint (VTEP).

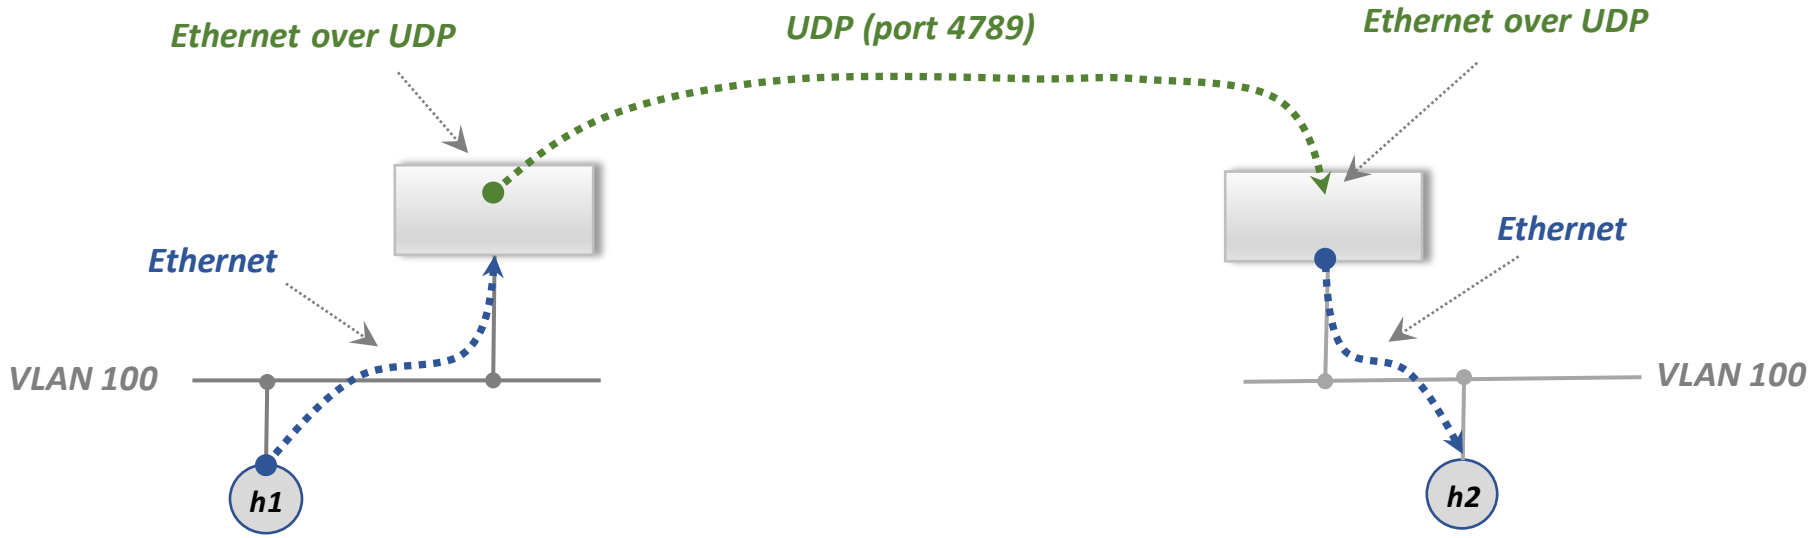

VXLAN Tunnel Endpoint (VTEP)

VXLAN Tunnel Endpoint (VTEP)

- A VTEP is a function implemented on a network device (or hypervisor) via software or dedicated hardware (ASIC)
- Encapsulates/de-encapslutates Ethernet frames in UDP (MAC-in-UDP) packets
- Manages the FIB MAC - VNI - destination VTEP or port

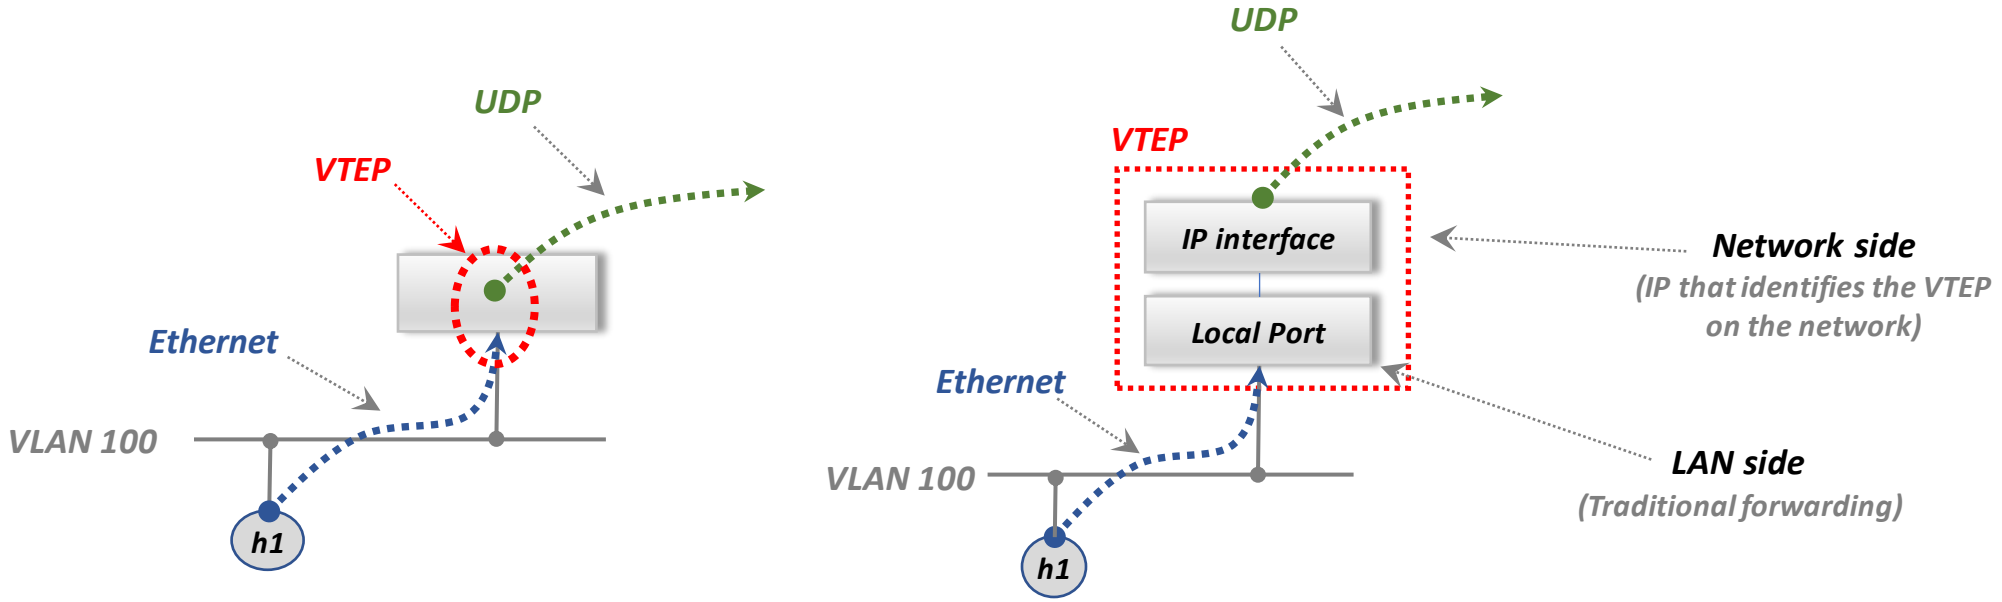

Mac table example

```
show mac address-table
          Mac Address Table
------------------------------------------------------------------
Vlan    Mac Address       Type        Ports      Moves   Last Move
----    -----------       ----        -----      -----   ---------

1199    807f.f854.e21c    DYNAMIC     Vx1        0       6 days, 15:59:43 ago
1199    968e.d343.0f64    DYNAMIC     Vx1        0       6 days, 15:59:43 ago
1199    968e.d343.2b94    DYNAMIC     Vx1        0       6 days, 15:59:45 ago
1199    968e.d343.3e7a    DYNAMIC     Vx1        0       6 days, 15:59:43 ago
1199    968e.d343.9bb0    DYNAMIC     Vx1        0       6 days, 15:59:44 ago
1199    968e.d343.a82e    DYNAMIC     Vx1        0       6 days, 15:59:44 ago
1199    968e.d343.b646    DYNAMIC     Vx1        0       6 days, 15:59:44 ago
1199    968e.d343.beec    DYNAMIC     Vx1        0       2 days, 22:48:29 ago
1699    0000.0000.0001    STATIC      Router
1699    0050.5682.3b65    DYNAMIC     Et16       2       9 days, 13:31:26 ago
1699    0050.5682.5607    DYNAMIC     Po47       2       9 days, 15:07:59 ago
```

VXLAN table example

```
show vxlan address-table
          Vxlan Mac Address Table
----------------------------------------------------------------------

VLAN   Mac Address      Type      Prt   VTEP             Moves   Last Move
----   -----------      ----      ---   ----             -----  ---------
1182   4c73.4f20.7820   EVPN      Vx1   10.8.20.64               6 days, 16:05:03 ago
1182   4c73.4f23.7020   EVPN      Vx1   10.8.20.65               6 days, 16:04:59 ago
1182   968e.d343.a82e   EVPN      Vx1   10.4.30.5                6 days, 15:59:59 ago
1182   968e.d343.b646   EVPN      Vx1   10.4.30.7                6 days, 15:59:58 ago
1182   968e.d343.beec   EVPN      Vx1   10.4.30.3                6 days, 15:59:59 ago
1187   4c73.4f21.e820   EVPN      Vx1   10.4.20.65               6 days, 15:59:57 ago
1187   807f.f854.e21c   EVPN      Vx1   10.4.20.64               6 days, 15:59:57 ago
1187   968e.d343.a82e   EVPN      Vx1   10.4.30.5                6 days, 15:59:59 ago
1187   968e.d343.b646   EVPN      Vx1   10.4.30.7                6 days, 15:59:59 ago
1187   968e.d343.beec   EVPN      Vx1   10.4.30.3                6 days, 15:59:59 ago
```

IP Fabric: Overlay

```
show ip route vrf l3vpn-infra-site2site

 B I      10.1.56.5/32 [200/0] via VTEP 10.1.30.5 VNI 123 router-mac 96:8e:d3:43:9b:b0
 B I      10.1.56.6/32 [200/0] via VTEP 10.2.30.3 VNI 123 router-mac 96:8e:d3:43:2b:94
 B I      10.1.56.7/32 [200/0] via VTEP 10.3.30.3 VNI 123 router-mac 96:8e:d3:43:0f:64
 B I      10.1.56.8/32 [200/0] via VTEP 10.1.30.5 VNI 123 router-mac 96:8e:d3:43:9b:b0
 B I      10.1.56.9/32 [200/0] via VTEP 10.1.30.5 VNI 123 router-mac 96:8e:d3:43:9b:b0
 B I      10.1.56.0/24 [200/0] via VTEP 10.2.30.3 VNI 123 router-mac 96:8e:d3:43:2b:94
                               via VTEP 10.1.30.5 VNI 123 router-mac 96:8e:d3:43:9b:b0
                               via VTEP 10.3.30.3 VNI 123 router-mac 96:8e:d3:43:0f:64
                               via VTEP 10.1.30.3 VNI 123 router-mac 96:8e:d3:43:3e:7a
 B I      10.1.100.0/24 [200/0] via VTEP 10.8.20.64 VNI 123 router-mac 4c:73:4f:20:78:20
                                via VTEP 10.8.20.65 VNI 123 router-mac 4c:73:4f:23:70:20
 B I      10.2.5.0/24 [200/0] via VTEP 10.8.20.64 VNI 123 router-mac 4c:73:4f:20:78:20
                              via VTEP 10.8.20.65 VNI 123 router-mac 4c:73:4f:23:70:20
```

IP Fabric: ESI Load Balancing

For Dual Homed hosts with normal LAG (LACP) EVPN uses the Ethernet Segment Identifier on routers

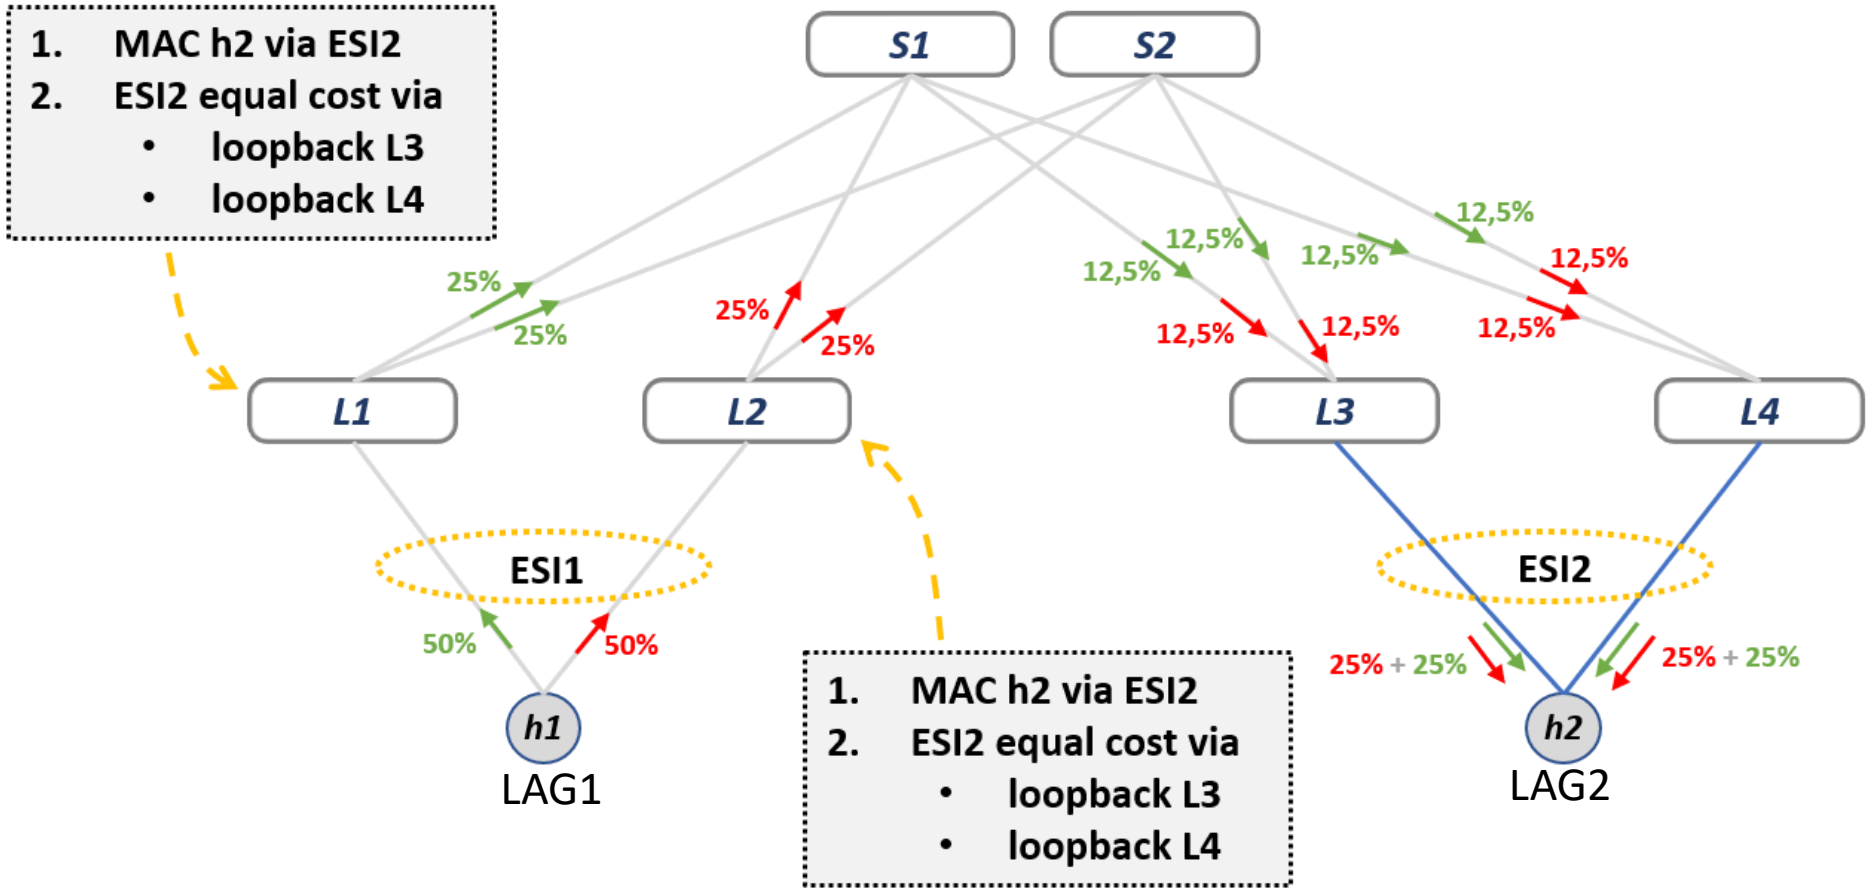

Orchestration and Automation

- Network devices can be configured via API
- To configure and manage a big and complex network, some vendor has implemented his own “orchestrator”. (e.g., Arista Cloud Vision) that “masquerade” the complexity.
- It is possible to automate deployment and configuration with Zero Touch Provisioning (ZTP) and for example Ansible.
- This interaction can be with the orchestrator (e.g., cloud vision) or directly with network devices

Arista switch API example

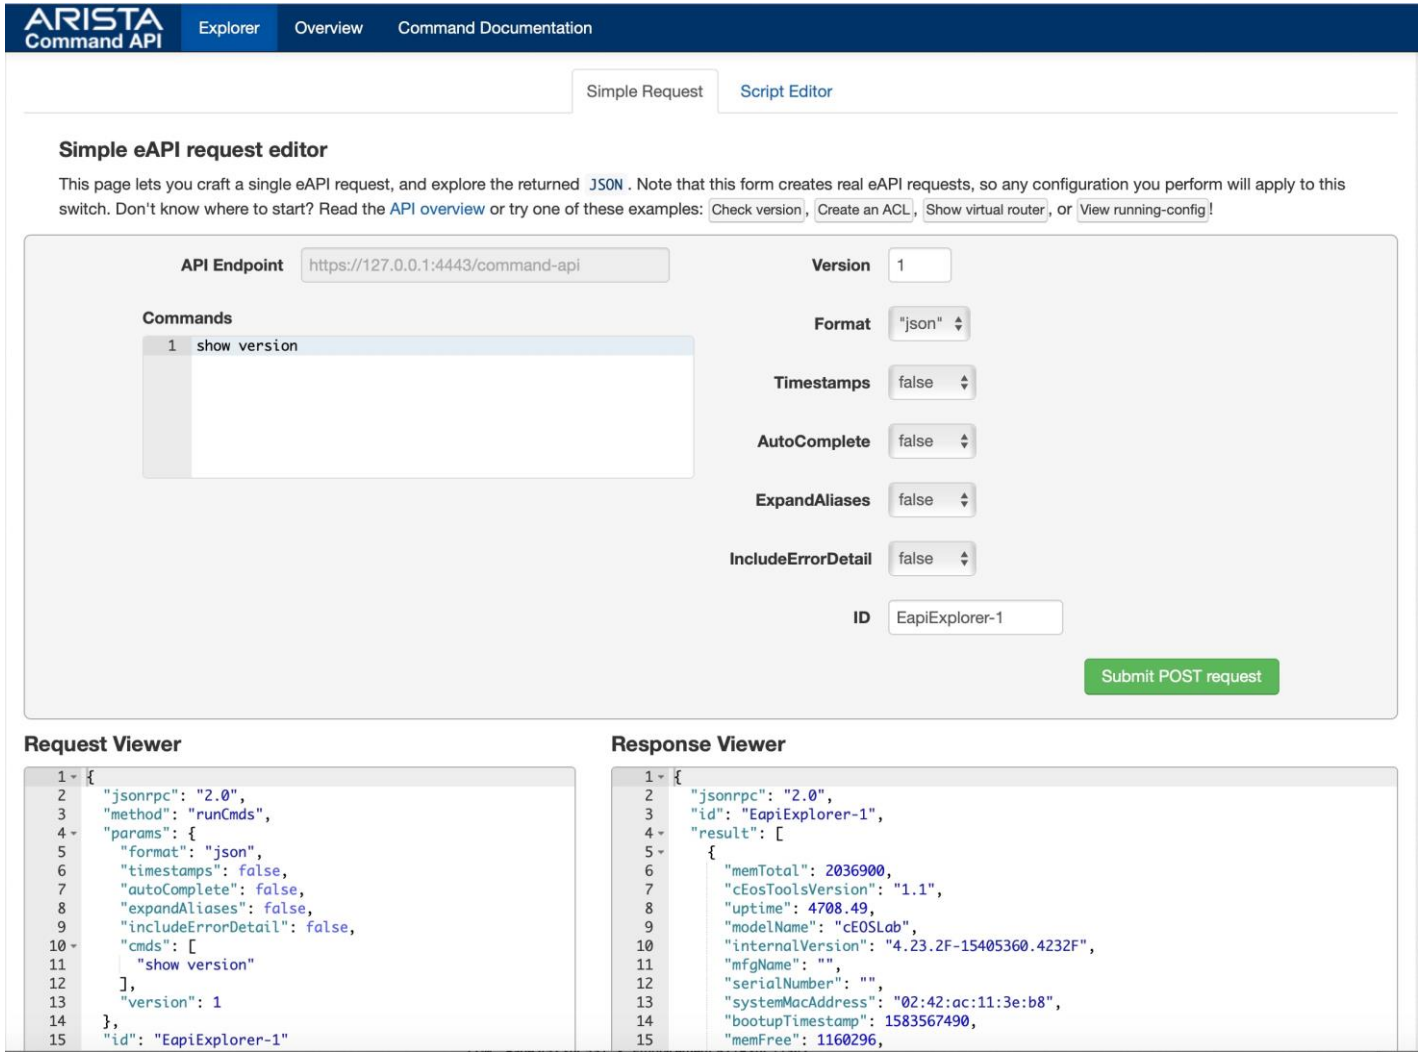

Arista Cloud Vision example

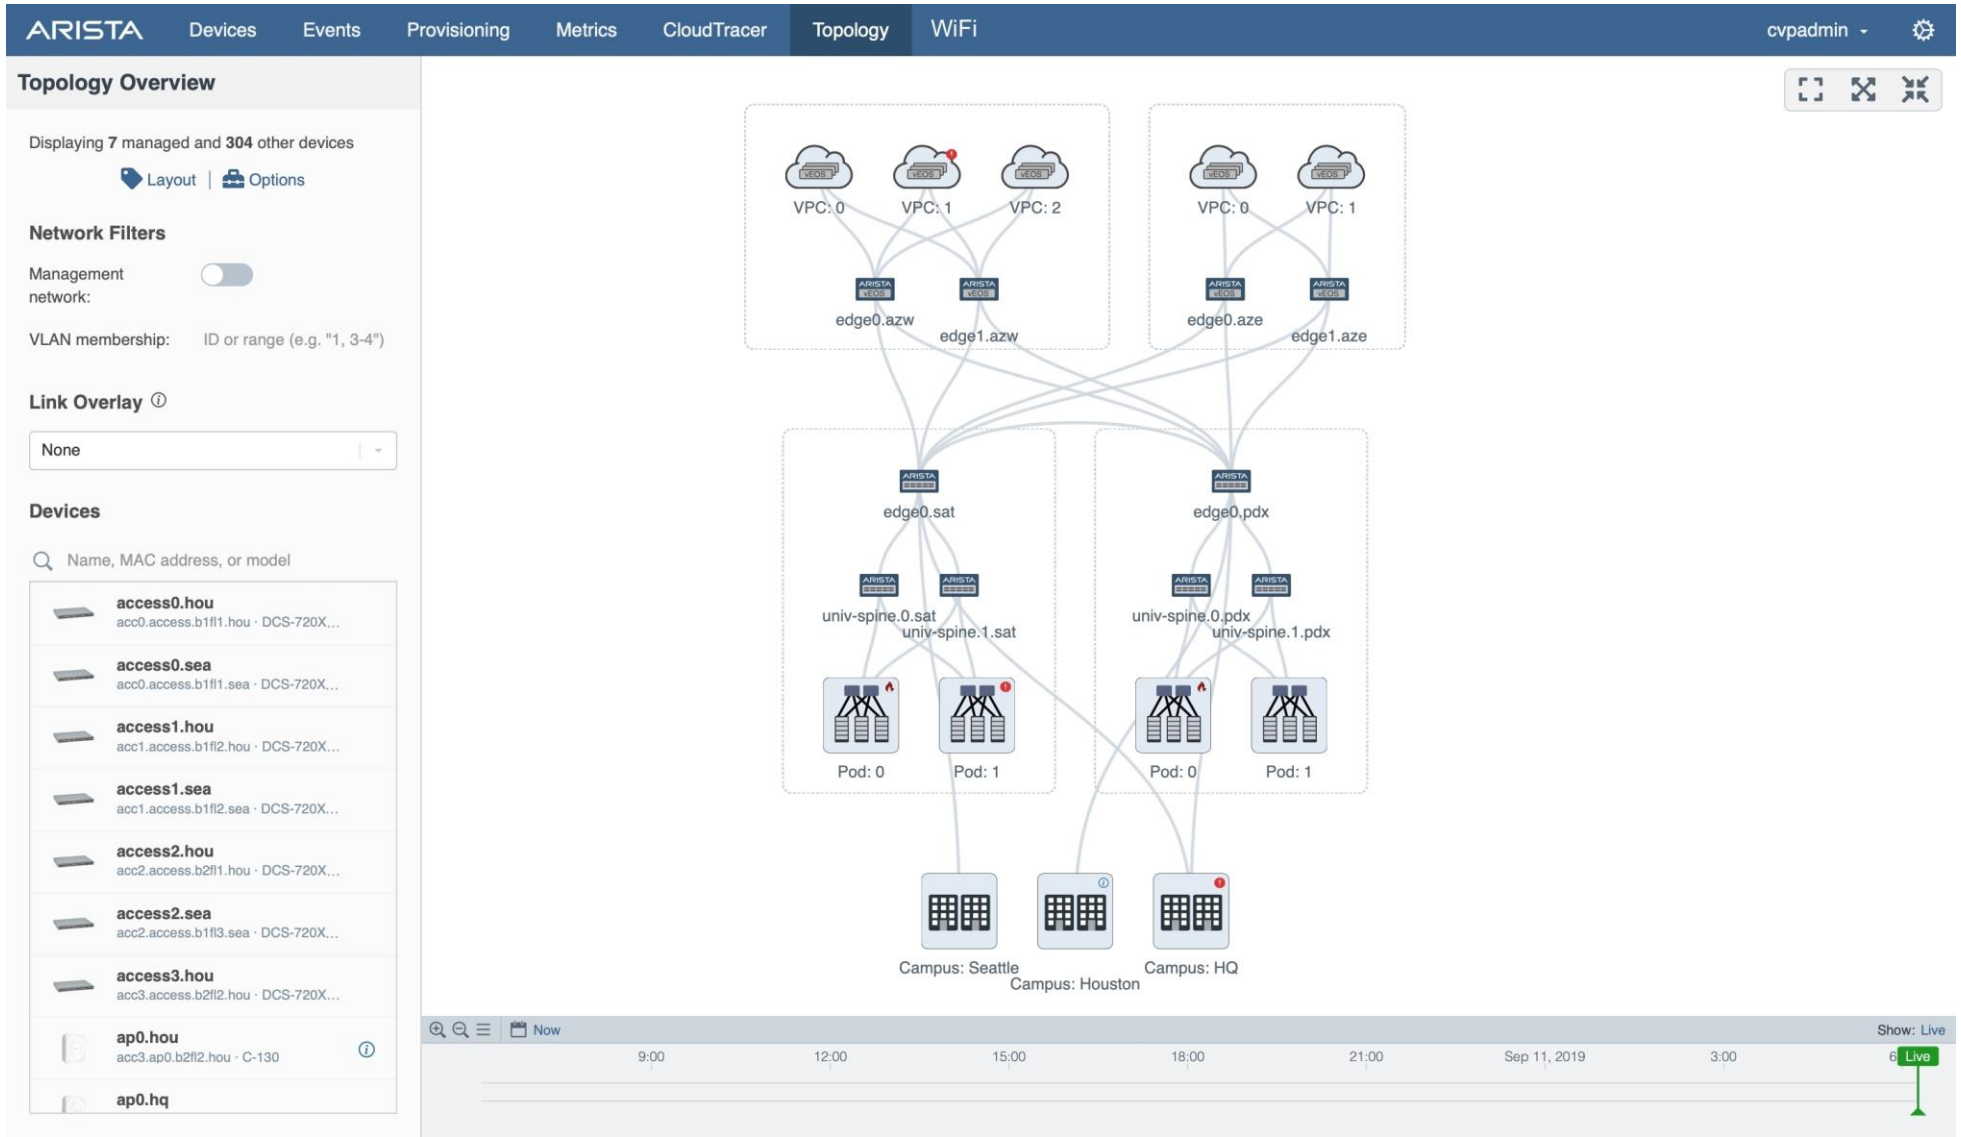

NETGROUP

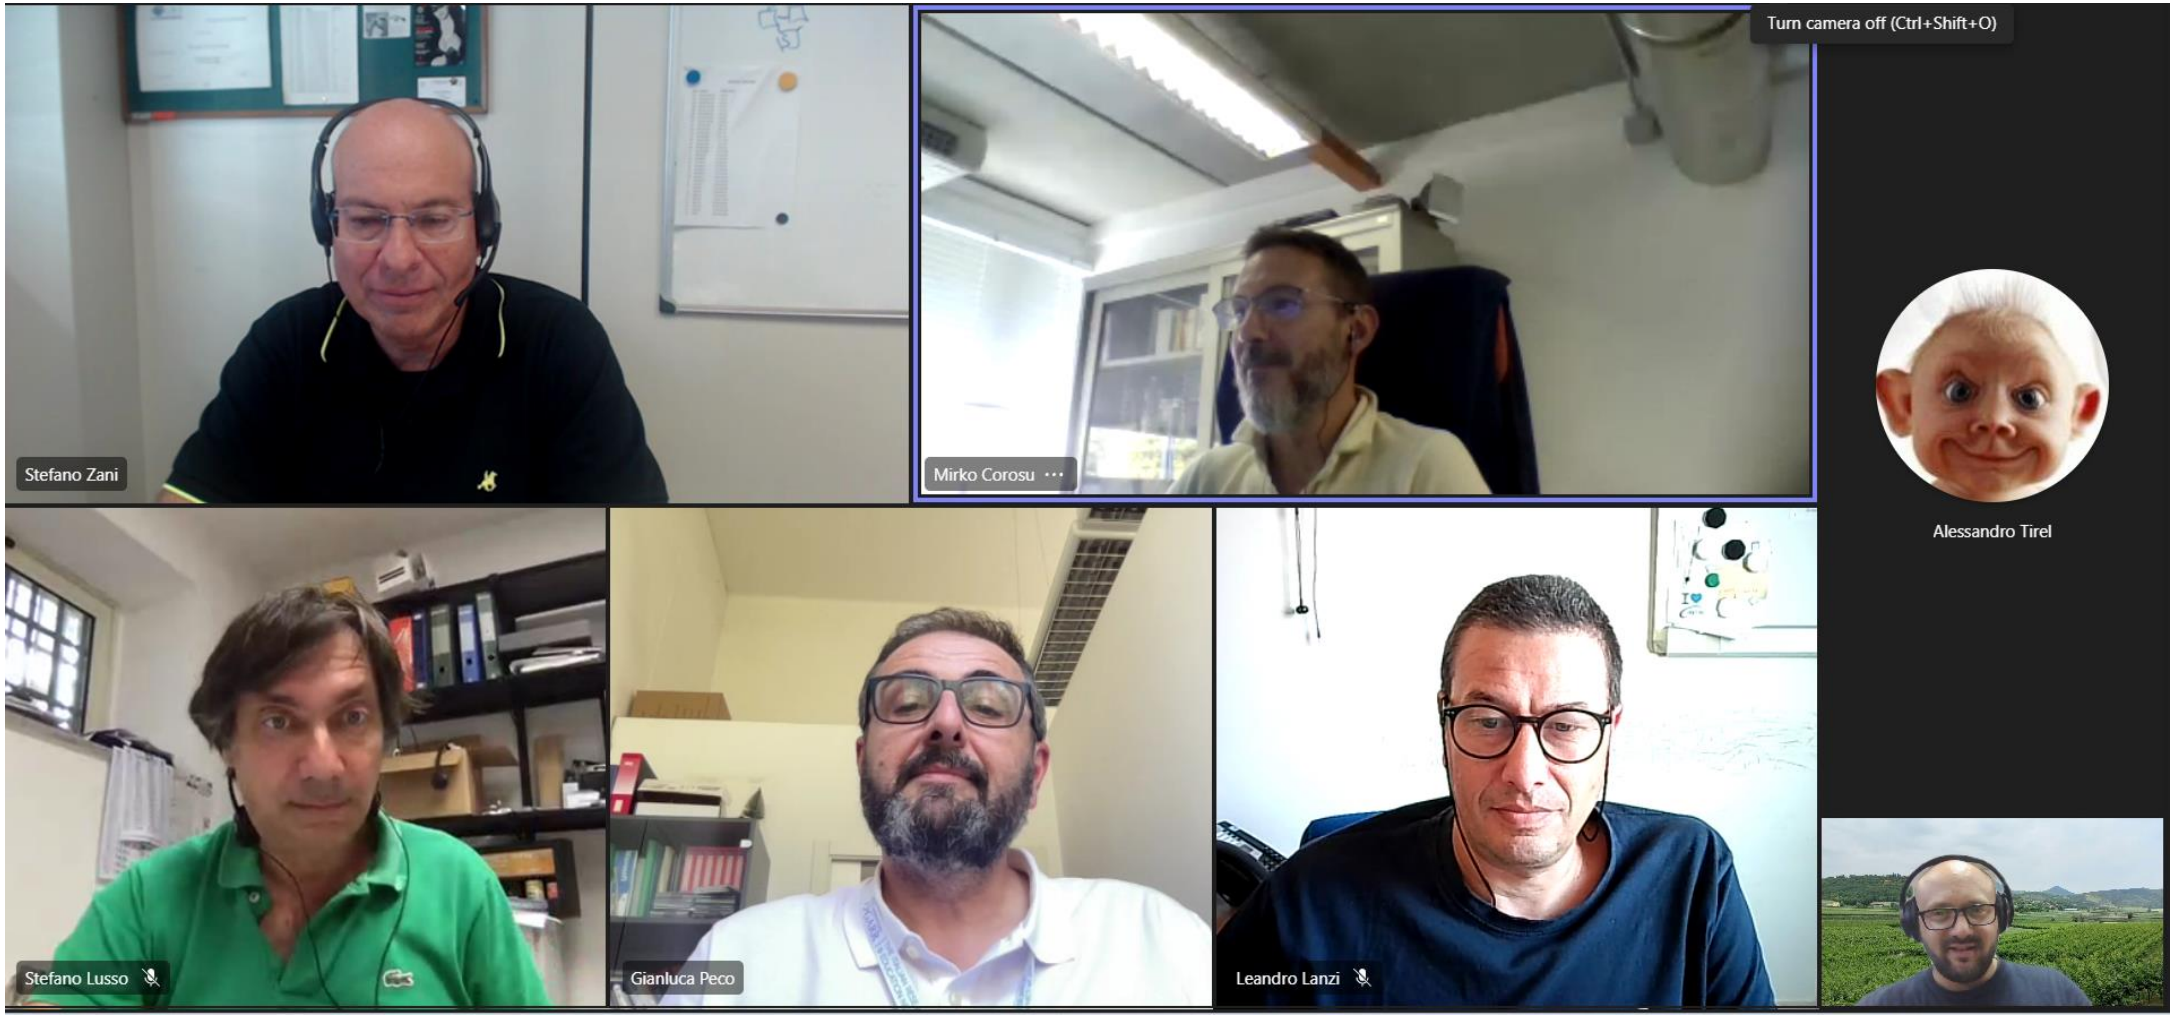
Stefano Zani
Mirko Corosu
Turn camera off (Ctrl+Shift+O)
Alessandro Tirel
Stefano Lusso
Gianluca Peco
Leandro Lanzi