

CERN IPv6 deployment

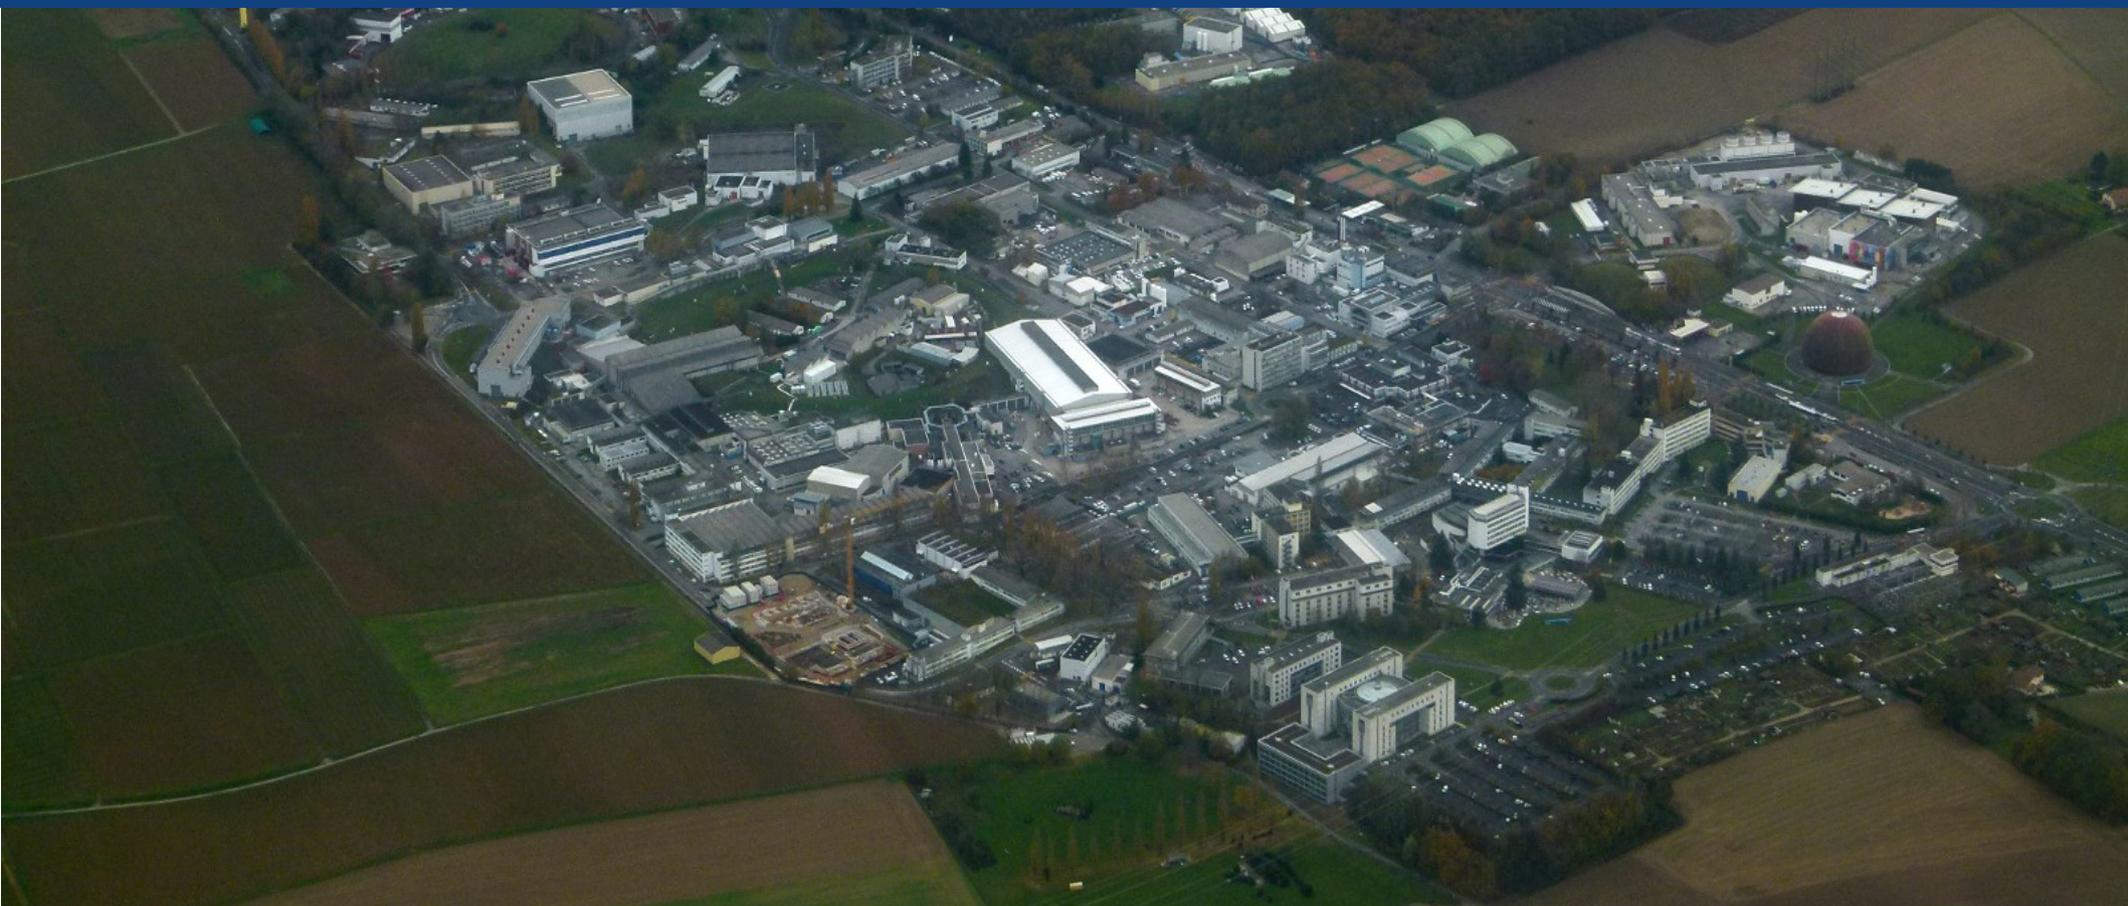
HEPiX WG meeting - CERN

16th of January 2020

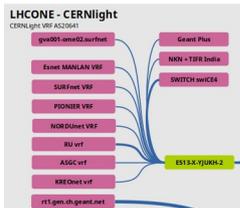
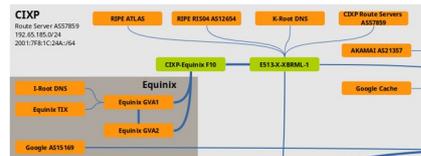
edoardo.martelli@cern.ch



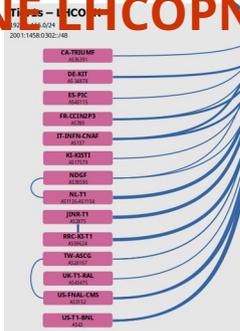
CERN networks



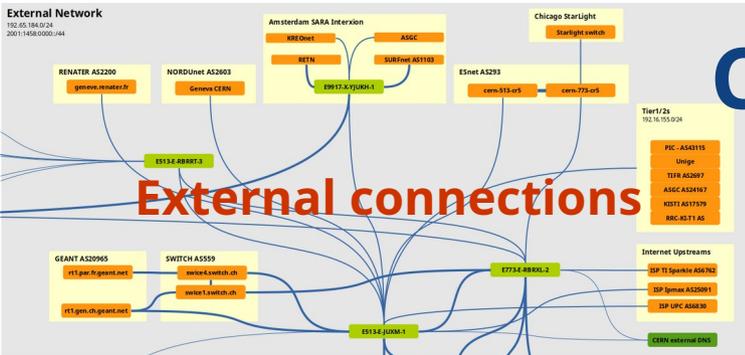
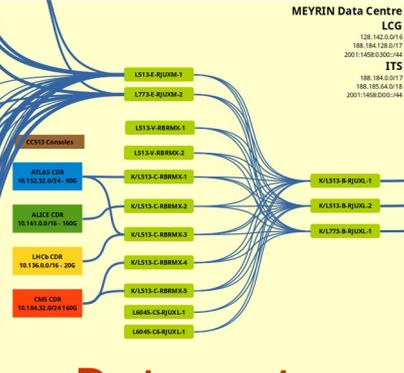
CERN networks



LHCONE LHCOPN



Datacentre

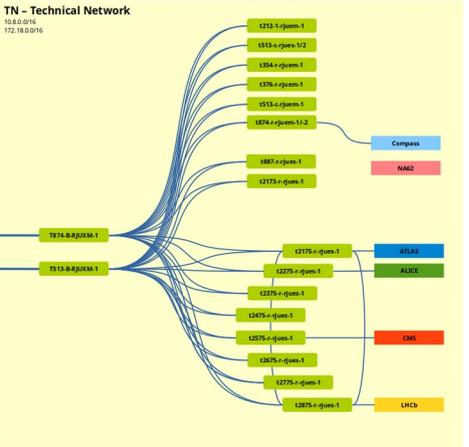


External connections

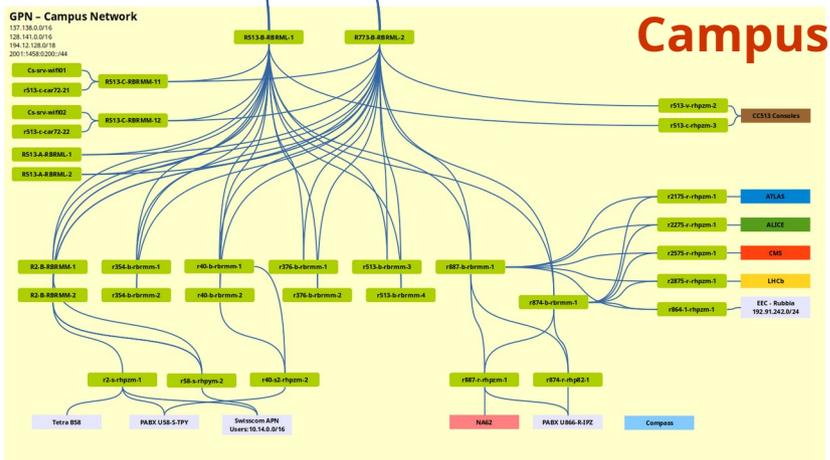
Firewall



TN: LHC control



Campus



Figures:

- 230 routers
- 3800 Switches
- 50000 connected devices

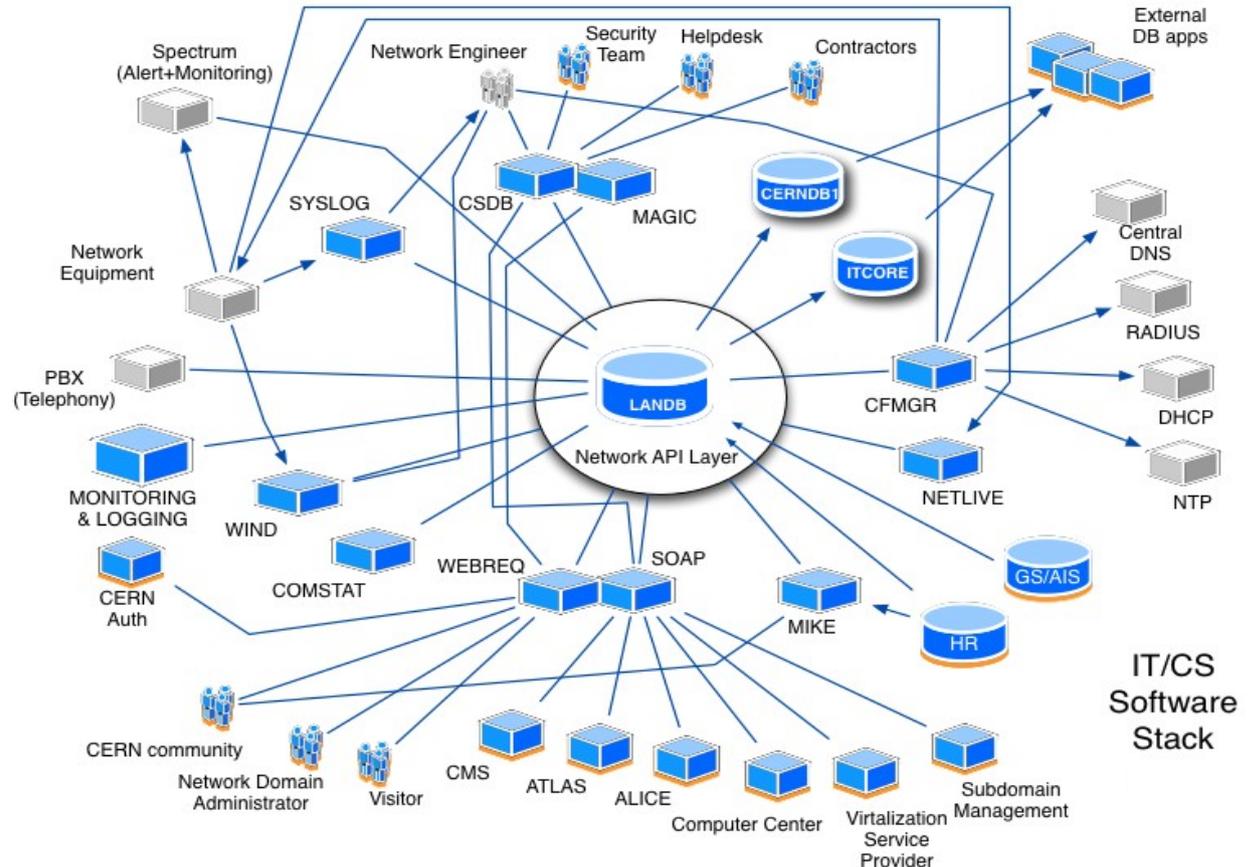
Network Provisioning and Management

LANDB

- 250+ Database tables
- 200,000+ Registered devices
- 1,000,000+ lines of codes
- 15+ years of development

Source of configurations of:

- all network devices
- network services (DNS, DHCP..)



IPv6 project



IPv6 deployment driver

CERN started using IPv6 in 2001, but for many years there was no reason for it

Large Virtual Machines deployment started in 2010. It was soon **planned to have 130,000 VMs with public IP addresses** to crunch the data from LHC after its upgrade in 2014

CERN IPv4 address pools (5x /16) already mostly fully used at that time

Approval and resources

IPv6 deployment project approved in Q1 2011

Allocated resources:

- *Network design/testing/deployment:*
1x Network Engineer FTE for 2 years.
- *Network database and NMS applications:*
2x Software Developers FTE for 2 years

To be ready for production by 2013

IPv6 service definition

- Dual Stack
- At least one IPv6 address for every IPv4 address in use
- Identical performance as IPv4, no degradation
- Common provisioning tools for IPv4 and IPv6
- Same network service portfolio as IPv4
- Common security policies for IPv4 and IPv6

IPv6 project - work plan

Software tools:

- New IPv6-compatible Network Database schema
- IPv6 Address assignments in Network Database
- Development of NMS tools (Network device configurations, DHCPv6, DNS...)
- Network-Database interface for end-users

Network:

- Testing on installed base
- Configuration of Network devices
- Deployment of Network services (DNS, DHCPv6, Radius, NTP)
- Training of Support-Lines and Advanced Users

IPv6 deployment



Network Database (LANDB)

Main tasks:

- Addition of IPv6 network tables and records
- Address plan
- Assignment of IPv6 subnet/addresses to all IPv4 entries
- Web interfaces for engineers and users

Challenges:

- New schema compatible with all legacy queries
- Consistently populate IPv6 records

Limitations:

- IPv4 still needed

CSDBweb (engineer interface)

CSDB WEB

- ManUTP++
- Admin
- ManSPIP
- GTI
- Inventory
 - Search
 - Equipment
 - Batch insert
 - Model change
 - Statistics
- Firewall
- Data Export
- Fiber
 - Trunk
 - Trunks List
 - Channel
 - Channels List
- MTP++
- Multicast
- NetLive
- Blocking
- DNS domains
- Syslog
- Syslog Configuration
- Vm Cluster
 - Vm clusters list

FIREWALL FILTER

[Insert](#) [Update](#) [Delete](#)

Filter information

Filter name: [Show gates](#)

Type: Status:

IPv4 / IPv6:

Responsible: [Myself](#)

Description:

Rules

Traffic rules									
<input type="checkbox"/>	Seq	Action	Protocol	Bidirect	IPv4/IPv6	Left Address	Ports	Right Address	Ports
<input type="checkbox"/>		<input type="text" value="ALL"/>	<input type="text" value="ALL"/>	<input type="text" value="ALL"/>	<input type="text" value="ALL"/>				
<input type="checkbox"/>	35	Deny	IP	→	Both	[N->DHCP] [2001:1458:202::] [128.141.0.0/0.0.255.255]		[Any] [::] [0.0.0.0/255.255.255.255]	
<input type="checkbox"/>	45	Deny	IP	→	Both	[N->LCG] [2001:1458:301::] [128.142.0.0/0.0.255.255]		[Any] [::] [0.0.0.0/255.255.255.255]	
<input type="checkbox"/>	55	Deny	IP	→	Both	[N->RLAN] [2001:1458:201::] [137.138.0.0/0.0.255.255]		[Any] [::] [0.0.0.0/255.255.255.255]	
<input type="checkbox"/>	60	Deny	IP	→	Both	[UNKNOWN]		[Any]	

Webreq (end-users interface)

Device Information

- **Device Name:** RIPE-ATLAS-PROBE [Last Operation]
- **Location:** 0031 S-0012
- **Manufacturer:** UNKNOWN
- **Model/Type:** UNKNOWN
- **Generic Type:** UNKNOWN
- **Description:** RIPE MEASUREMENT PROBE
- **Tag:**
- **Serial Number:**
- **Operating System:** UNKNOWN **Version:** UNKNOWN
- **CERN Inventory number:**
- **Network Interface Card(s):** 00-20-4A-C8-24-98/ETHER-AUTO-10/100
- **Responsible for the device:** MARTELLI EDOARDO IT CS
EDOARDO.MARTELLI@CERN.CH / Tif: 72613
- **Main User of the device:** MARTELLI EDOARDO IT CS
EDOARDO.MARTELLI@CERN.CH / Tif: 72613
- **HCP Response:** This system **CAN** obtain an IP address automatically [[more info](#)]
- **IPv6 Ready:** This system **IS NOT** IPv6 ready
- **Last changed:** 21-02-2014 (15:51)

Interface(s) Information

[>>Network Service HELP<<](#) [>>Network Interface Card\(s\) HELP<<](#)

Interface Name	IP Address	Service Name	Internet Connectivity
RIPE-ATLAS-PROBE.CERN.CH	137.138.32.177 2001:1458:201:b459::100:3f	S31-S-IP3	Y
Subnet IPv4 Mask: 255.255.255.192 Default IPv4 Gateway: 137.138.32.129		Name IPv4 Servers: 137.138.16.5, 137.138.17.5 Time IPv4 Servers: 137.138.16.69, 137.138.17.69	
Subnet IPv6 Netmask: 64 Default IPv6 Gateway: 2001:1458:201:b459::1		Name IPv6 Servers: 2001:1458:201:1000::5, 2001:1458:201:1100::5 Time IPv6 Servers: 2001:1458:201:1040::69, 2001:1458:201:1140::69	
IP Aliases: NONE			
Bound Interface Card(s): NONE			

Network device configuration

Main tasks:

- Test IPv6 functionalities and performance on all running devices
- Development of NMS tools for automatically generated configurations and ACLs
- Finally add commands for IPv6 addressing and routing on all the devices

Challenges:

- Translate firewall ACLs
- DHCPv6 configuration and functionalities

Limitations:

- Management and Monitoring still over IPv4
- IPv6 PBR not line rate (at the time, now fixed)

DHCPv6

Main tasks:

- Configuration of routers to replicate IPv4 behavior (only registered MAC addresses can get a lease)
- Configuration of servers from LANDB

Challenges:

- DHCPv6 is not DHCPv4
- DUID: Clients may request lease without telling their MAC address (saved by [RFC6939](#))

Limitations:

- Router Advertisements always needed (risk of SLAAC)
- No DHCPv6 clients on some O.S. (Android)
- *Main source of trouble tickets*

DNS

Main tasks:

- Make DNS servers answer over IPv6
- Generate zones with IPv6 records: AAAA by default on ipv6.cern.ch, but user configurable

Challenges:

- For once mostly straightforward

Users control

End users can control the configuration of the DNS and firewall openings with two flags:

IPv6 DNS and firewall == Yes

- publish the IPv6 address (AAAA record) in the zone cern.ch
- activate IPv6 openings in the central firewall

IPv6 DNS and firewall == No

- publish the IPv6 address (AAAA record) in the zone ipv6.cern.ch
- deactivate IPv6 openings in the central firewall

IPv4 equivalent flag: **IPv4 DNS and firewall**

DHCPv4 and DHCPv6 leases always provided

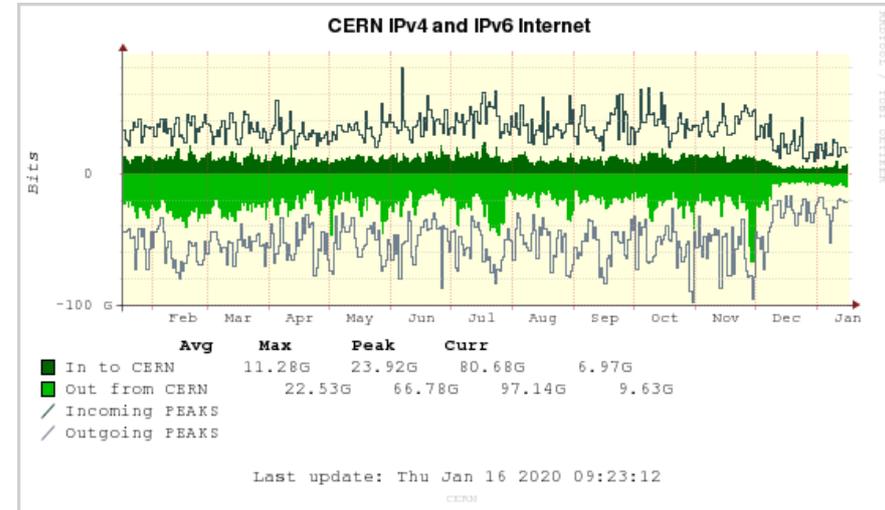
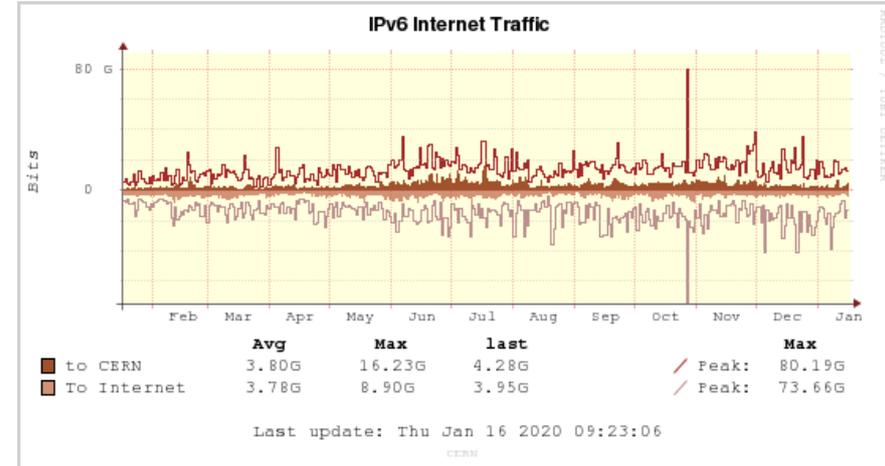
IPv6 status



CERN Internet Access

More incoming IPv6 than outgoing

- Input: ~25% of total
- Output: ~12% of total



Sources:

<https://netstat.cern.ch/monitoring/network-statistics/ext?q=CERN&p=EXT&mn=IPv6-Internet&t=Yearly>

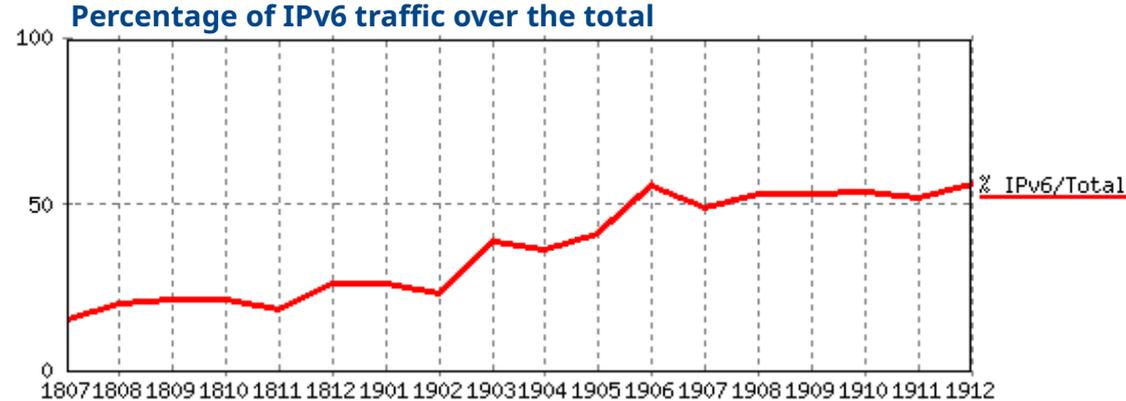
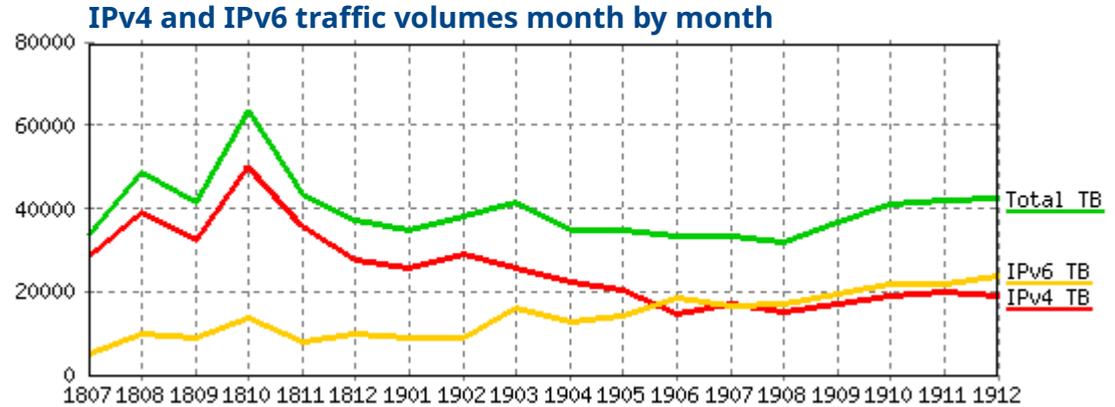
<https://netstat.cern.ch/monitoring/network-statistics/ext?q=CERN&p=EXT&mn=01-Total-Internet-traffic&t=Yearly>

LHCOPN and LHCONE

LHCOPN and LHCONE traffic measured on the CERN routers:

IPv6 surpassed IPv4 in June 2019

Mostly thanks to dual-stack EOS clusters of LHC experiments



Sources:

<https://twiki.cern.ch/twiki/bin/view/LHCOPN/LHCOPNEv4v6Traffic>

Lessons learnt

- Catching up with 20+ years of IPv4 experience and development takes a lot of time
- The network is the easy part
- DHCPv6 is definitely not DHCPv4
- Have a staged deployment with a large variety of early adopters
- It works!



Questions?