IPv6 deployment at CERN

RIPE IPv6 working group
Warsaw, 15th May 2014
edoardo.martelli@cern.ch

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

Agenda

- CERN Network
- IPv6 deployment project
- IPv6 deployment status
- Challenges and lessons learnt

CERN Network

CERN Accelerators complex

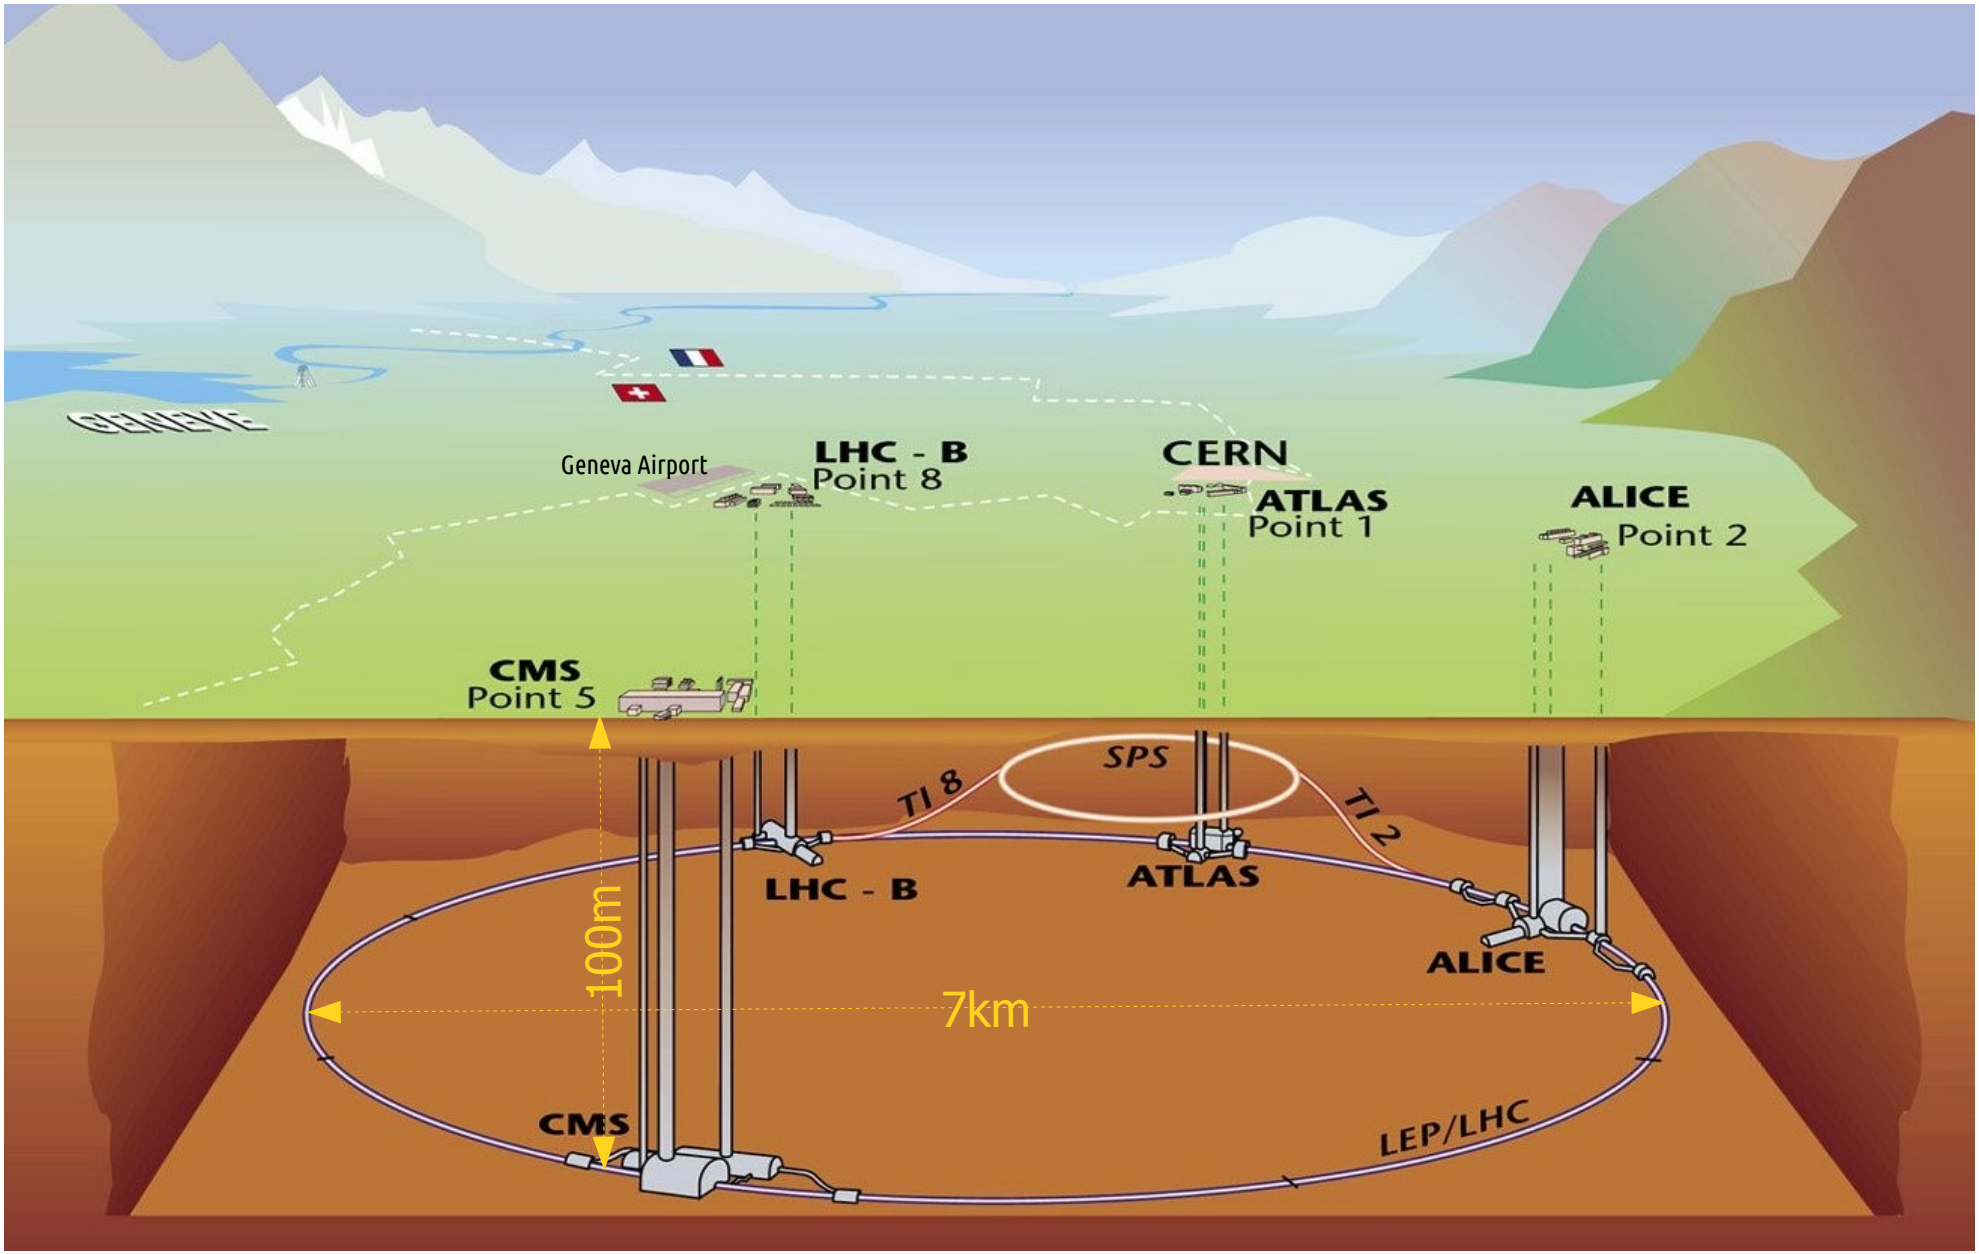

CERN data network

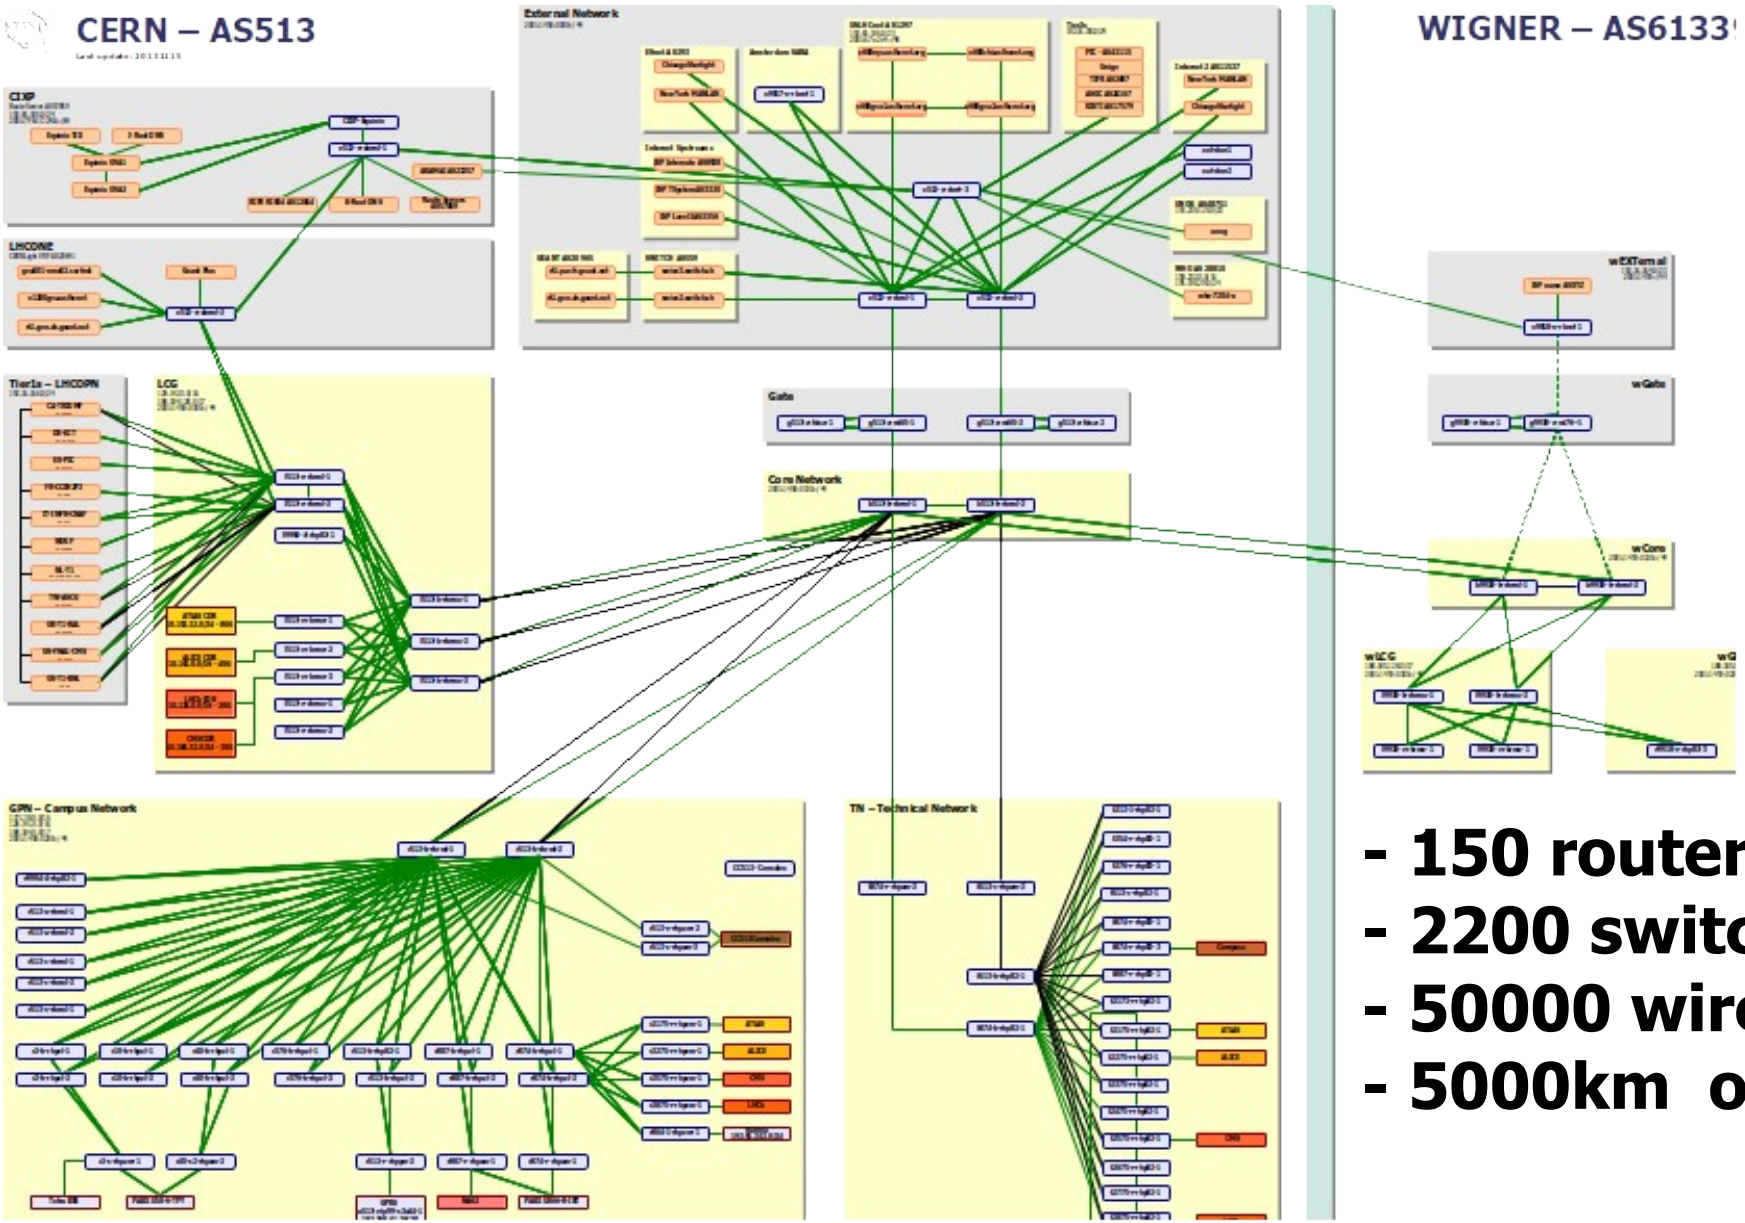

- 150 routers
- 2200 switches
- 50000 wired devices
- 5000km optical fibres

Network Provisioning and Management System

- 250 Database tables
- 100,000 Registered devices
- 50,000 hits/day on web user interface
- 1,000,000 lines of codes

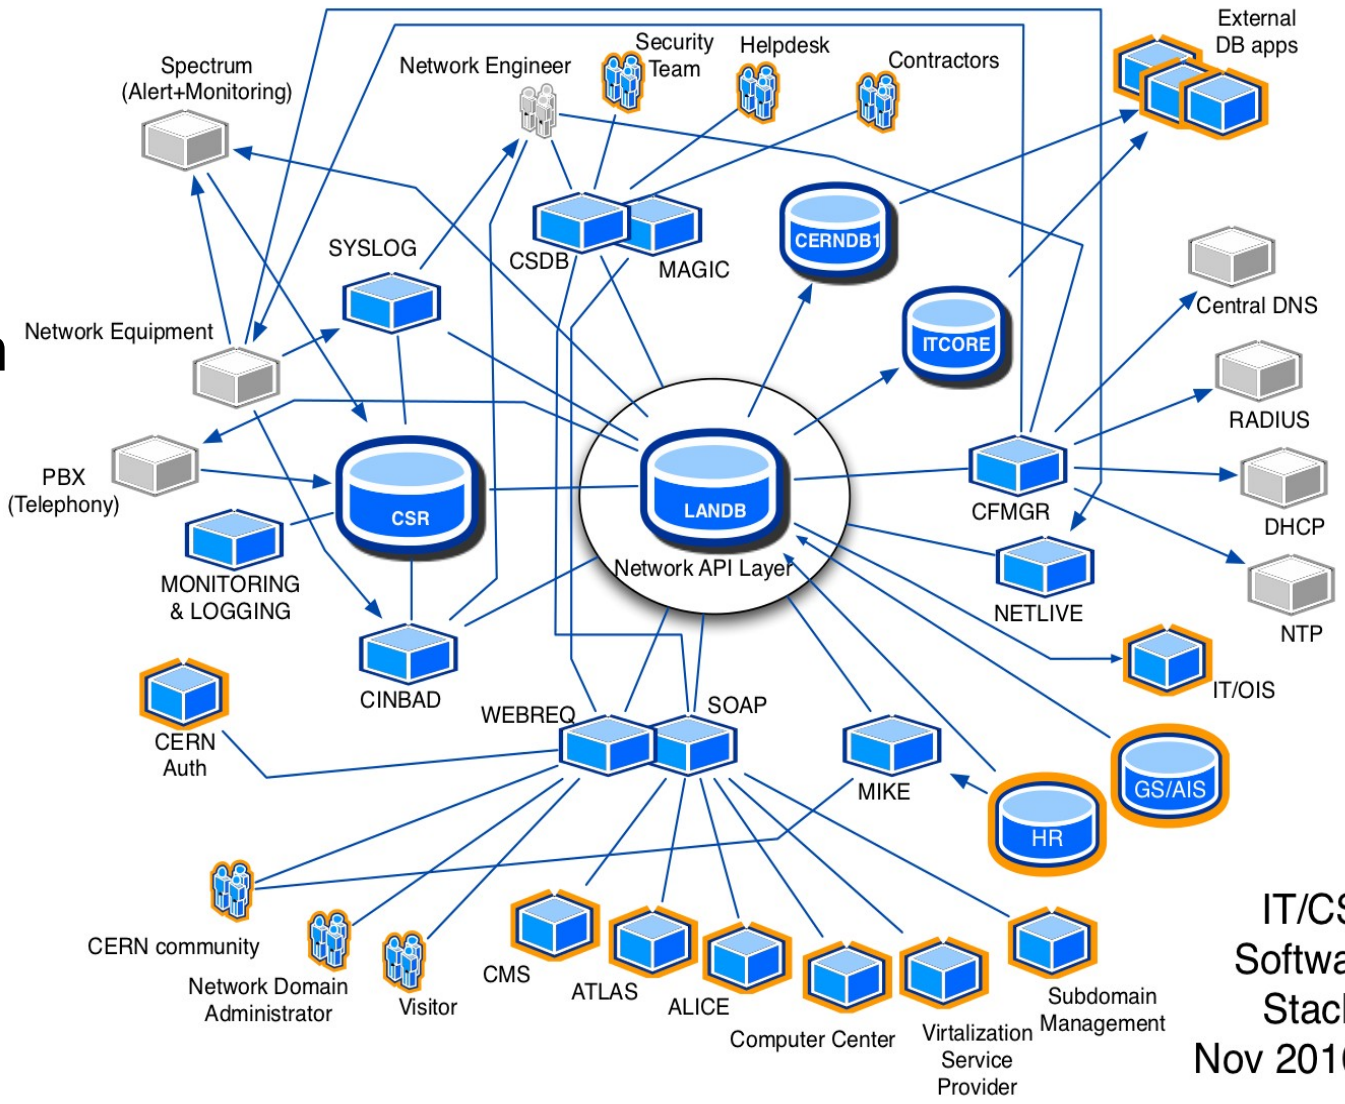

IPv6 deployment project

Driver

CERN started playing with IPv6 in 2001, but for many years there was no reason for it.

Large **Virtual Machines** deployment ramped up in 2010.
Planned to have 130,000 VMs with public IP addresses to crunch the data from LHC after its upgrade in 2014.

Approval and resources

IPv6 deployment 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

IPv6 service definition

- Dual Stack
- At least one IPv6 address for every assigned IPv4 address
- Identical performance as IPv4, no degradation
- Common provisioning tools (NMS) for IPv4 and IPv6
- Same network services portfolio as IPv4
- Common security policies for IPv4 and IPv6

Workplan

- Testing of available network devices
- New compatible Network-DB schema
- Address assignments in Network-DB
- NMS tools development
- Network devices configuration
- Network services (DNS, DHCPv6, Radius, NTP)
- Network-DB Web interface for end-users
- Training for Support-Lines and Powe-Users

To be ready for production in 2013

How did it go?

Well :-)

“Dual stack”

Network database:

- IPv6 now main navigation data
- New schema compatible with all legacy queries
- IPv6 address tables fully populated

Network:

- All campus, data centres, firewalls, external interfaces are dual stack (except: LHC accelerator control network, LHC detectors data acquisition networks).

Same routing architecture (BGP and OSPF).

"An IPv6 address to every IPv4 ones"

- Every device with an IPv4 address has an IPv6 address assigned in the Network DB

- All assigned IPv6 addresses have a name in **ipv6.cern.ch**

```
# host ping.ipv6.cern.ch
ping.ipv6.cern.ch has IPv6 address 2001:1458:201:1c80::100:175

# host TELEPHONE-62470.ipv6.cern.ch
TELEPHONE-62470.ipv6.cern.ch has IPv6 address fd01:1458:204:27a::100:2e
```

- Dynamic (portable) devices get a name in **dyndns6.cern.ch**

```
# host myiphone.dyndns6.cern.ch
myiphone.dyndns6.cern.ch has IPv6 address 2001:1458:202:180::101:8a26
```

"Identical performance"

All production network devices can forward IPv6 packets at wire speed

Only exception: policy base routing for statefull firewall bypass. Not a show stopper, because of low IPv6 traffic volume.

"Common provisioning tools"

NMS:
- routers configuration generator for all the vendors
- DHCPv6, DNS configurations from Network-DB
- ACL generator for firewalls from Network-DB

CSDBweb (Network-DB interface for engineers): IPv6 everywhere there is IPv4

WebReq (Network-DB interface for end-users): All IPv6 info visible together with IPv4, IPv6-ready flag settable

CSDBweb (engineering)

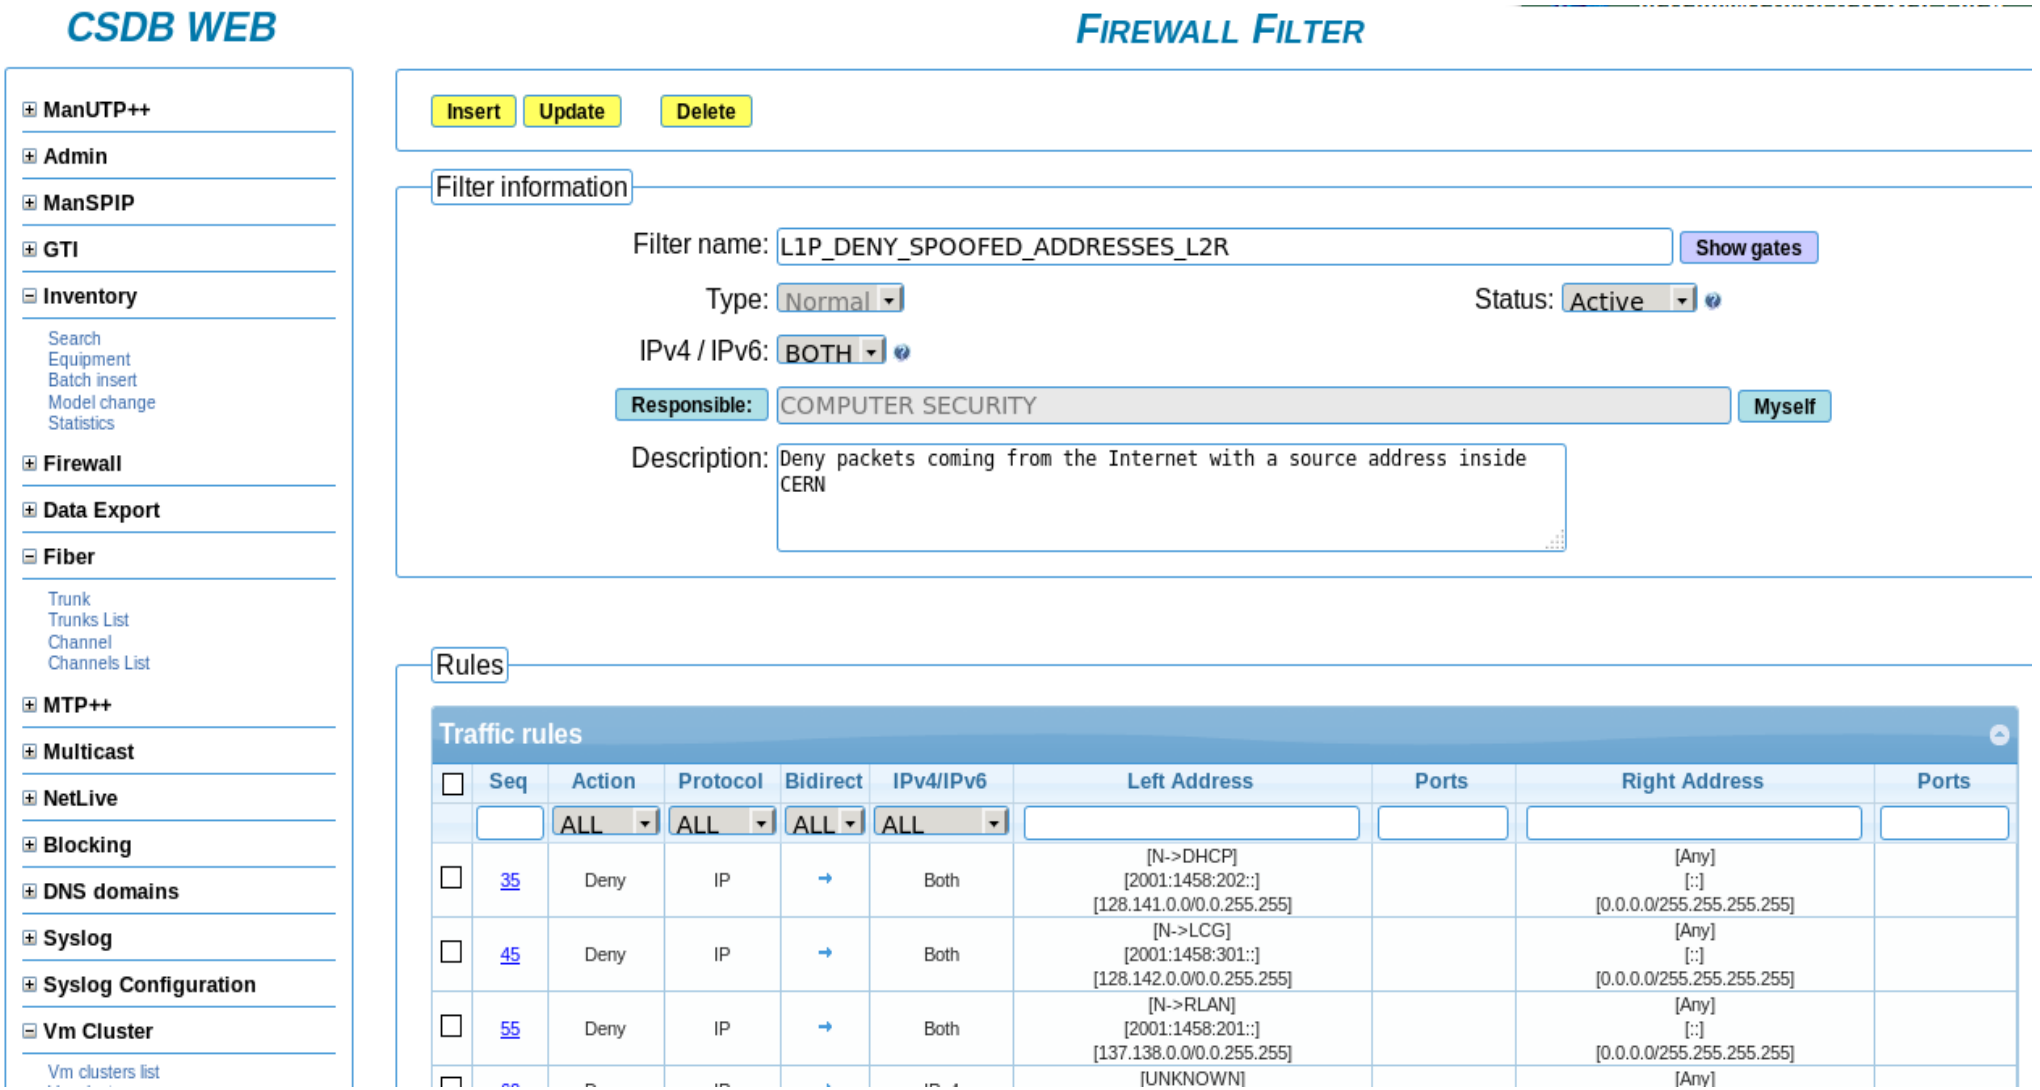

Webreq (end-users)

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 / Tlf: 72613
- Main User of the device: MARTELLI EDOARDO IT CS EDOARDO.MARTELLI@CERN.CH / Tlf: 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	

IPv6-ready flag

Users can declare their own devices as “IPv6-ready”

IPv6-ready means:

- IPv6 connectivity is OK
- all running server applications are listening on both v4 and v6 sockets

Consequences:

- Firewall: IPv6 equivalent of IPv4 security openings applied to the central firewall
- DNS: DEVICENAME.cern.ch returns A and AAAA records, reverse returns DEVICENAME.cern.ch (and host certificates work)

"Same network services as IPv4"

DNS:
- queryable over IPv6
- announced in the DHCPv6 leases

NTP:
- reachable over IPv6 (ip-time-1.ipv6.cern.ch and ip-time-2.ipv6.cern.ch)

DHCPv6:
- Static devices: same servers and daemon of SHCP for IPv4
- Dynamic devices: different servers because running the very latest version (classes only works in 4.3.0)

"Common security policies"

Firewall rules database

- schema and management tools developed
- Most IPv4 rules automatically translated into IPv6
- IPv6 only rules manually created

ACL compiler generator

- Antispoofing ACLs applied to all router interfaces
- All firewalls managed by the NMS

IPv6 on a day of May

DHCPv6 active leases: 5000 avg, 10000 peak (55% of DHCPv4)

DNS queries over IPv6: 210,000/hour (4% of queries over IPv4)

Internet traffic: 5% of ISP traffic (but 0.2% of external traffic)

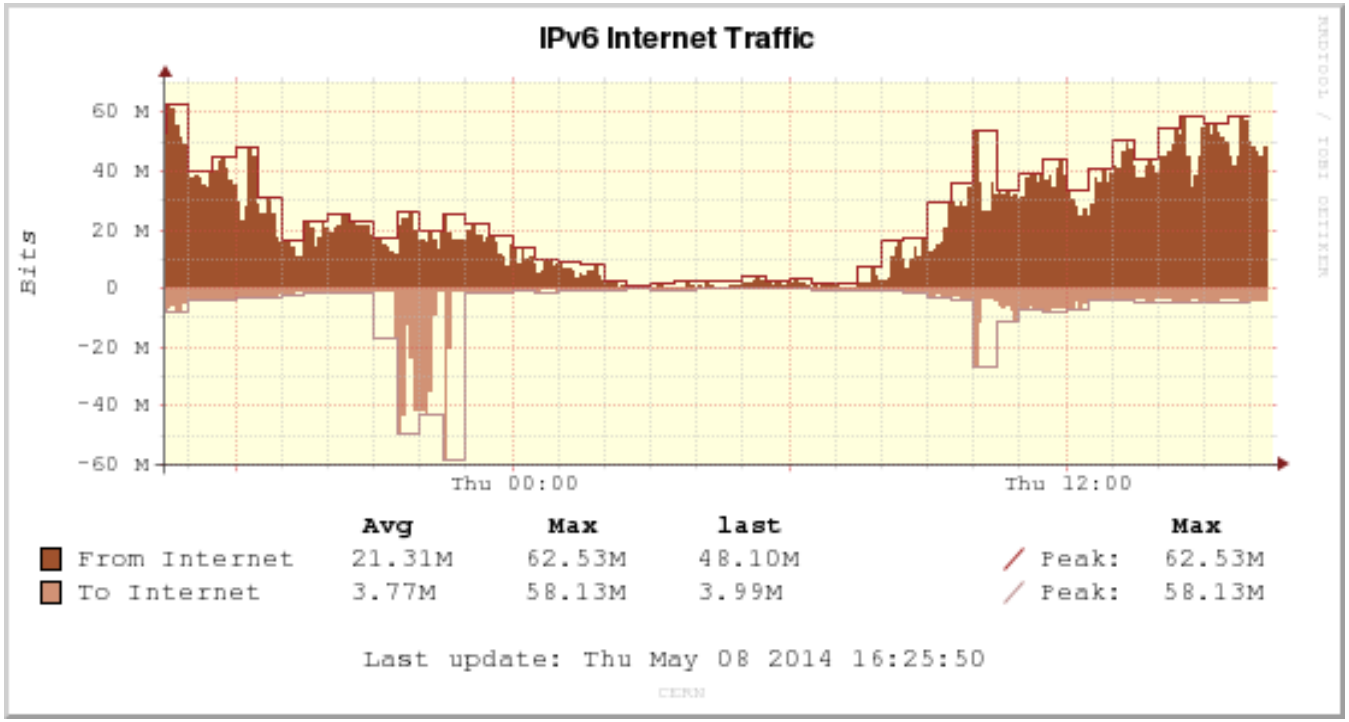

Timeline

2001: CERN IPv6 testing started
2003, June: public IPv6 prefix assigned to CERN
2003, September: IPv6 deployed in the CERN External Network: CERN prefix announce to NRENs. Direct and Reverse DNS over IPv6.
2003, November: IPv6 Land Speed record in collaboration with Caltech
2009, November: CERN IPv6 prefix visible in the whole IPv6 Internet.
2011, January: IPv6 deployment project approved
2011, February: IPv6 address plan issued
2011, March: Development LANDB schema includes IPv6 information.
2011, July: IPv6 connectivity in part of LCG, CORE and GPN backbones (Brocade routers)
2011, July: Prototype of DNS servers
2011, August: Pilot IPv6 services for LCG and GPN users
2012, March: LANDB (Network database) with IPv6 tables in production
2012, March: CSDWEB support of IPv6 information
2012, March: training of Netcom and CD about new CSDB
2012, July: CSDB supports IPv6 for MANUTP and MTP and Blocking
2012, October: cfmgr Brocade and HP compilers can generate IPv6 configurations
2013, March: all routers in the Computer Centre of Building 513 support IPv6 for end-users
2013, March: WEBREQ support of IPv6 information (not for end-users yet)
2013, April: DHCPv6 for static devices (IP services)
2013, April: All LCG routers have dual-stack services
2013, June: NTP service ready: ip-time-1.ipv6.cern.ch and ip-time-2.ipv6.cern.ch
2013, September: DHCPv6 for portable devices
2013, September: DNS replies over IPv6 from ip-dns-1.ipv6.cern.ch and ip-dns-2.ipv6.cern.ch
2013, October: Gates software (landb schema and translation of existing IPv4 rules, csdweb, webreq, cfmgr gate update).
2013, October: DNS automaticaly configured from LANDB information
2013, November: All GPN routers have dual-stack services
2013, November: LANDB IPv6 information available from the SOAP interface
2013, November: WEBREQ shows IPv6 information to any user
2014, January: Automatic IPv6 configuration in the central firewall
2014, January: Leased dynamic addresses published in dyndns6.cern.ch
2014, February: IPv6-ready flag fully functional (DNS and Firewall)
2014, February: Netcom and IT Service desk trained
2014, February 18: DHCPv6 leases to any device in the IT buildings 31,28,600
2014, April 1st: DHCPv6 leases to any device in the IT datacentre in building 513
2014, May 6th: DHCPv6 leases to any registered device connected to a portable socket or WIFI

Latest information

Check the current status at

http://cern.ch/ipv6/content/implementation-plan-and-status

Challenges and lessons learnt

Benefits

Simplified management of addresses
- one subnet size fits all (/64)
- no-brain address planning for new deployments
- reduced risk of future renumbering

[Hopefully] Future proof

Challenges

- Size of routing tables and ACLs have doubled
- New issues to be solved by Support lines
- DHCPv6 still in an early stage
- New security threats to take into account
- Legacy applications don't understand IPv6, and some will never do

Challenges: DHCPv6

DHCPv6 rationale: Network-DB driven address assignment for automatic configuration of DNS and firewall, user traceability, light access control

Drawbacks

- RAs necessary for default-gateway and mask-length: two protocols to maintain and control, no predictive load-balancing for multi-router subnets, all available prefixes exposed

- MAC address authentication not always works: DHCPv6 clients don't have to use the MAC address of the interface they send the request via. Waiting for implementation of RFC6939 to fix it. DUID management is not an option

- Not all devices use DHCPv6 by default (iOS up to v6, Android up to 4.4, old MacOS/Linux/Windows versions, industrial devices...)

Compliance of physics applications

Physics Community actively reviewing IPv6 compliance of its applications

Pushing developers to code IPv6 support and correct bugs

IPv6 compliance of WLCG applications

Software Component	Type	Used by Experiment	Version	IPv6 Compliance
AliEN	LHC Experiment Application	ALICE		
ARC CE	Middleware	ATLAS, CMS		YES
ARGUS	Middleware	ALICE, ATLAS, CMS, LHCb		Unknown
BDII	Middleware	ATLAS, CMS, LHCb	EMI 2	YES
BestMAN	Middleware	ATLAS, CMS		Unknown
CASTOR	Middleware	ALICE, ATLAS, CMS, LHCb		NO
cfengine	Monitoring			Unknown
CMS Tag Collector	LHC Experiment Application	CMS		Unknown
CMSSW	LHC Experiment Application	CMS		Unknown
cmsweb	LHC Experiment Application	CMS		Unknown
CRAB 2	LHC Experiment Application	CMS		Unknown
Cream CE	Middleware	ALICE, ATLAS, CMS, LHCb	1.16.2	YES
CVMFS	Other Application	ALICE, ATLAS, CMS, LHCb	2.1.15	YES
Dashboard Google Earth	Monitoring	ALICE, ATLAS, CMS, LHCb		Claimed
dCache	Middleware	ALICE, ATLAS, CMS, LHCb	2.6.19	Claimed
dCache	Middleware	ALICE, ATLAS, CMS, LHCb	1.9.12	NO

http://hepix-ipv6.web.cern.ch/wlcg-applications

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.

Don't rush. Have a staged deployment with a large variety of early adopters. And keep in touch with them: they may not report all the problems.

Only the deployment on the live network will prove it can cope with the two protocols.

Don't rely on support from applications developers: there are already enough bugs to fix without adding IPv6.

Give IPv6 a chance

IPv4 connectivity was deployed at CERN in ~1986

The World Wide Web was conceived in 1989 and made its debut in 1992

Decnet was phased out in 1999

Tim Berners-Lee was *not* a network engineer

More information:

http://cern.ch/ipv6