IPv6 Network Service

IT GLM, 23 May 2011

edoardo.martelli@cern.ch

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

Summary

- Why IPv6?
- What is IPv6?
- CERN IPv6 Network Service
- v4/v6 coexistence risks
- IT-CS work plan
- Implications for IT
- Progressing together
- Costs

Why IPv6?

IPv4 ends

IPv4 address pool soon depleted

IPv4 exhaustion predictions

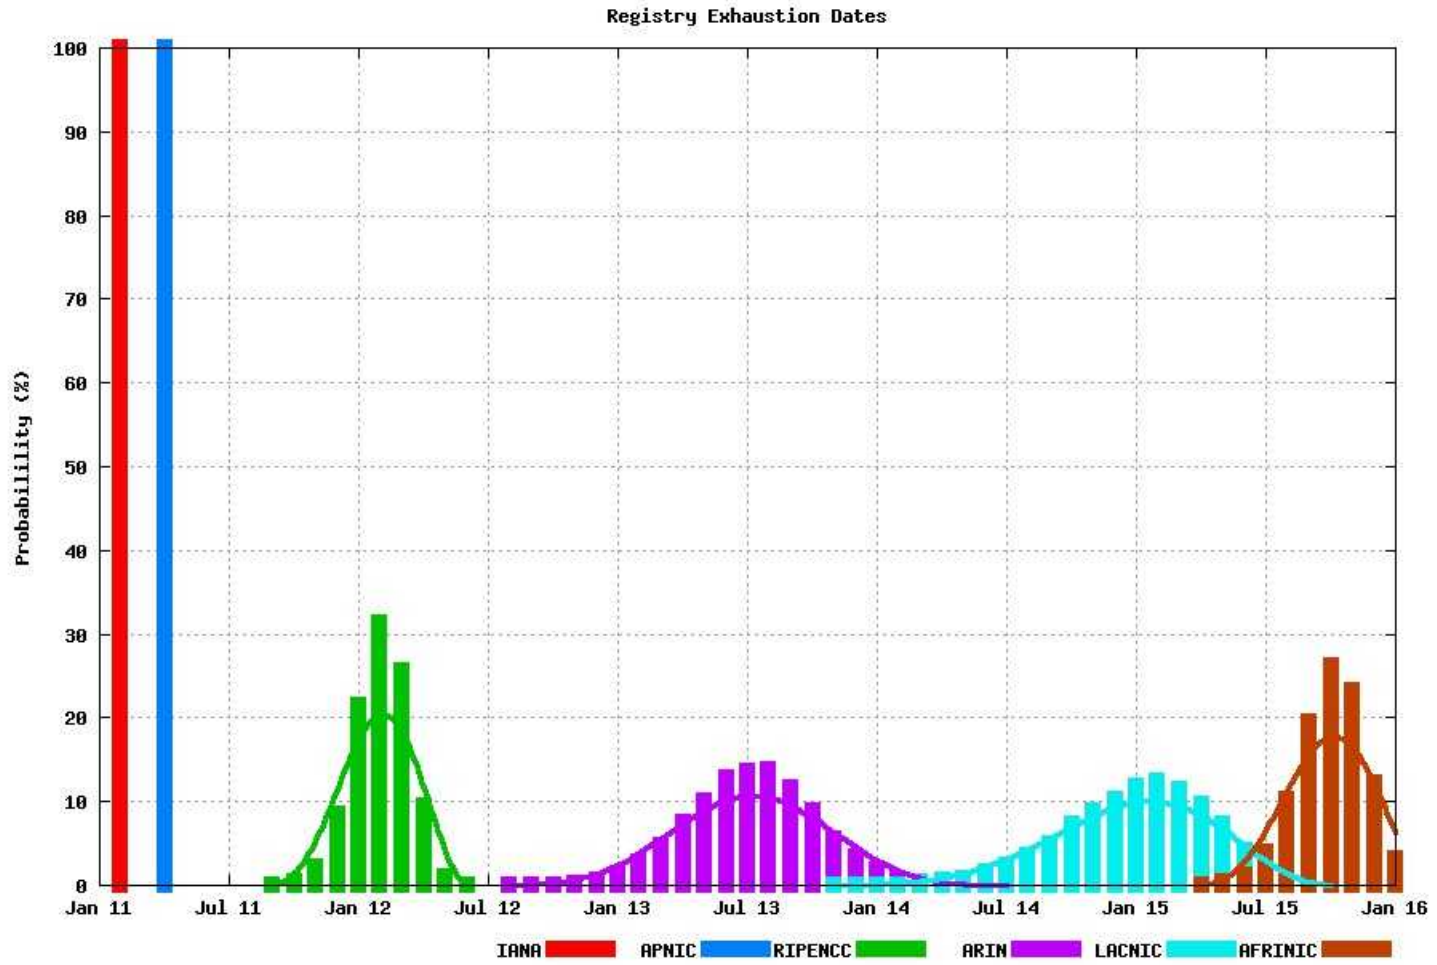

http://www.potaroo.net/tools/ipv4/rir.jpg

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

IPv4 exhaustion consequences

CERN IT Department

- Problematic for new players to join the IPv4 Internet
- Difficult to deploy new large services based on IPv4 (virtualization, clouds...)

As a consequence: part of the Internet will become IPv6 only

CERN unable to reach all its users

IPv6 necessary to keep reaching the whole Internet

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

What is IPv6?

IPv6 in a nutshell

```
2001:1458:a137:b138:c000:d000:e000:f001/64
  Site      Subnet           Host        Length
```

- **128 bits**, written in 8 groups of 4 hexadecimal digits

- 64 bits for network address, 64 bits for host address (recommendation)

- typical major site allocation: /32. It gives 2^32 subnets available (the whole IPv4 address space) . Every subnet has 2^64 host addresses available.

- NAT not available

Huge address space

Using the recommended sizes,
4 billions organizations can get:
- 4 billions subnets
- each subnet with
18,446,744,073,709,552,000 hosts

~25,000 addresses for every square meter on Earth for each one of the /64 subnets.

Transition strategies

Many NAT/Tunneling “solutions”:

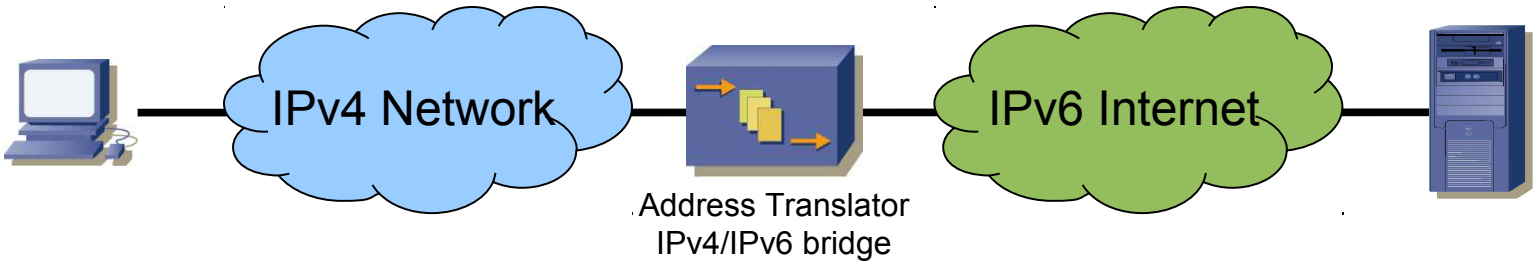

DUAL-STACK:

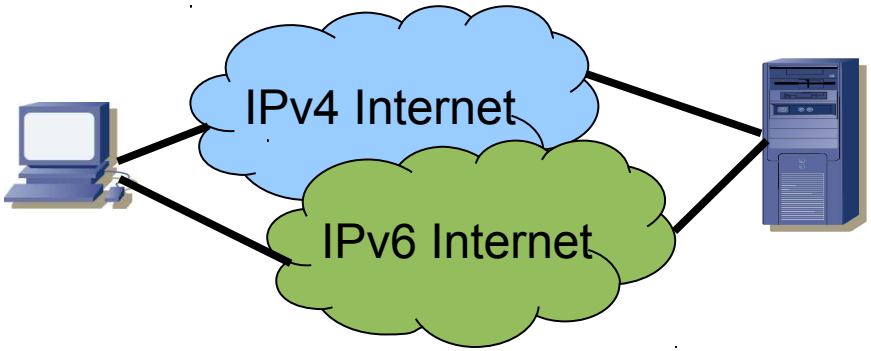

Transition strategies

Many NAT/Tunneling “solutions”:

DEPRECATED

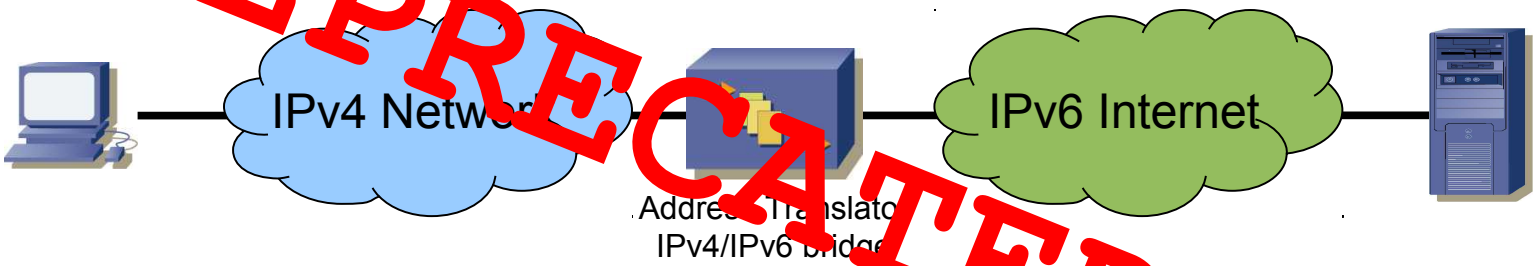

DUAL-STACK:

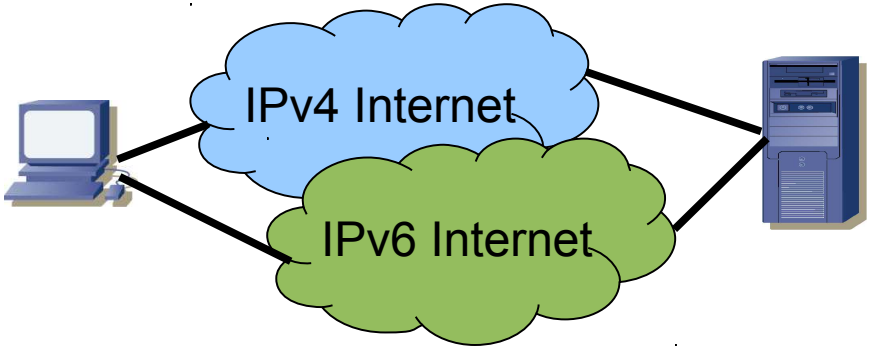

Dual Stack: only viable solution

CERN IPv6 Service

Strategy

IPv6 ≥ IPv4

The CERN IPv6 service must be at the same level of the IPv4 service.
Plus the advantages peculiar to IPv6.

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

Service Description

- Dual Stack
- One IPv6 address assigned to every IPv4 one
- Identical performance as IPv4
- Common provisioning tools for IPv4 and IPv6
- Same network services portfolio as IPv4
- Common security policies for IPv6 and IPv6
- Connectivity to IPv6 only systems!

CERN IPv6 prefixes

Public prefix **2001:1458::/32**
(globally routed, full Internet connectivity)

Local prefix **FD01:1458::/32**
(private addresses like 10.0.0.0,
no Internet connectivity)

IPv6 User services

At least one IPv6 sub-prefix per physical subnet, public and/or local.

Subnet size: /64 (i.e. 64 bits for the network address, 64 bits for the host address)
Hosts per subnets: 2^{64} (recommended size).

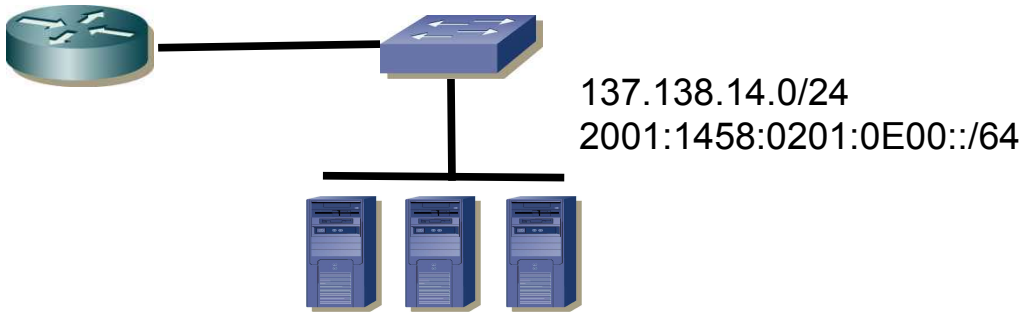

Keep control to ensure stability and security

Addresses assigned by LANDB:

- IPv6 addresses assigned by DHCPv6 servers. Static or Dynamic assignments based on the MAC address (same principles as IPv4).

Avoid Risks:

- IPv6 autoconfiguration disabled and rogue Router Advertisement messages filtered.

Network Services

DNS, DHCPv6, Radius and NTP will be available over the IPv6 network.

The existing IPv4 DNS, Radius and NTP servers will provide the IPv6 services.

DHCPv6 and DHCP(v4): two services running on the same physical server.

LANDB

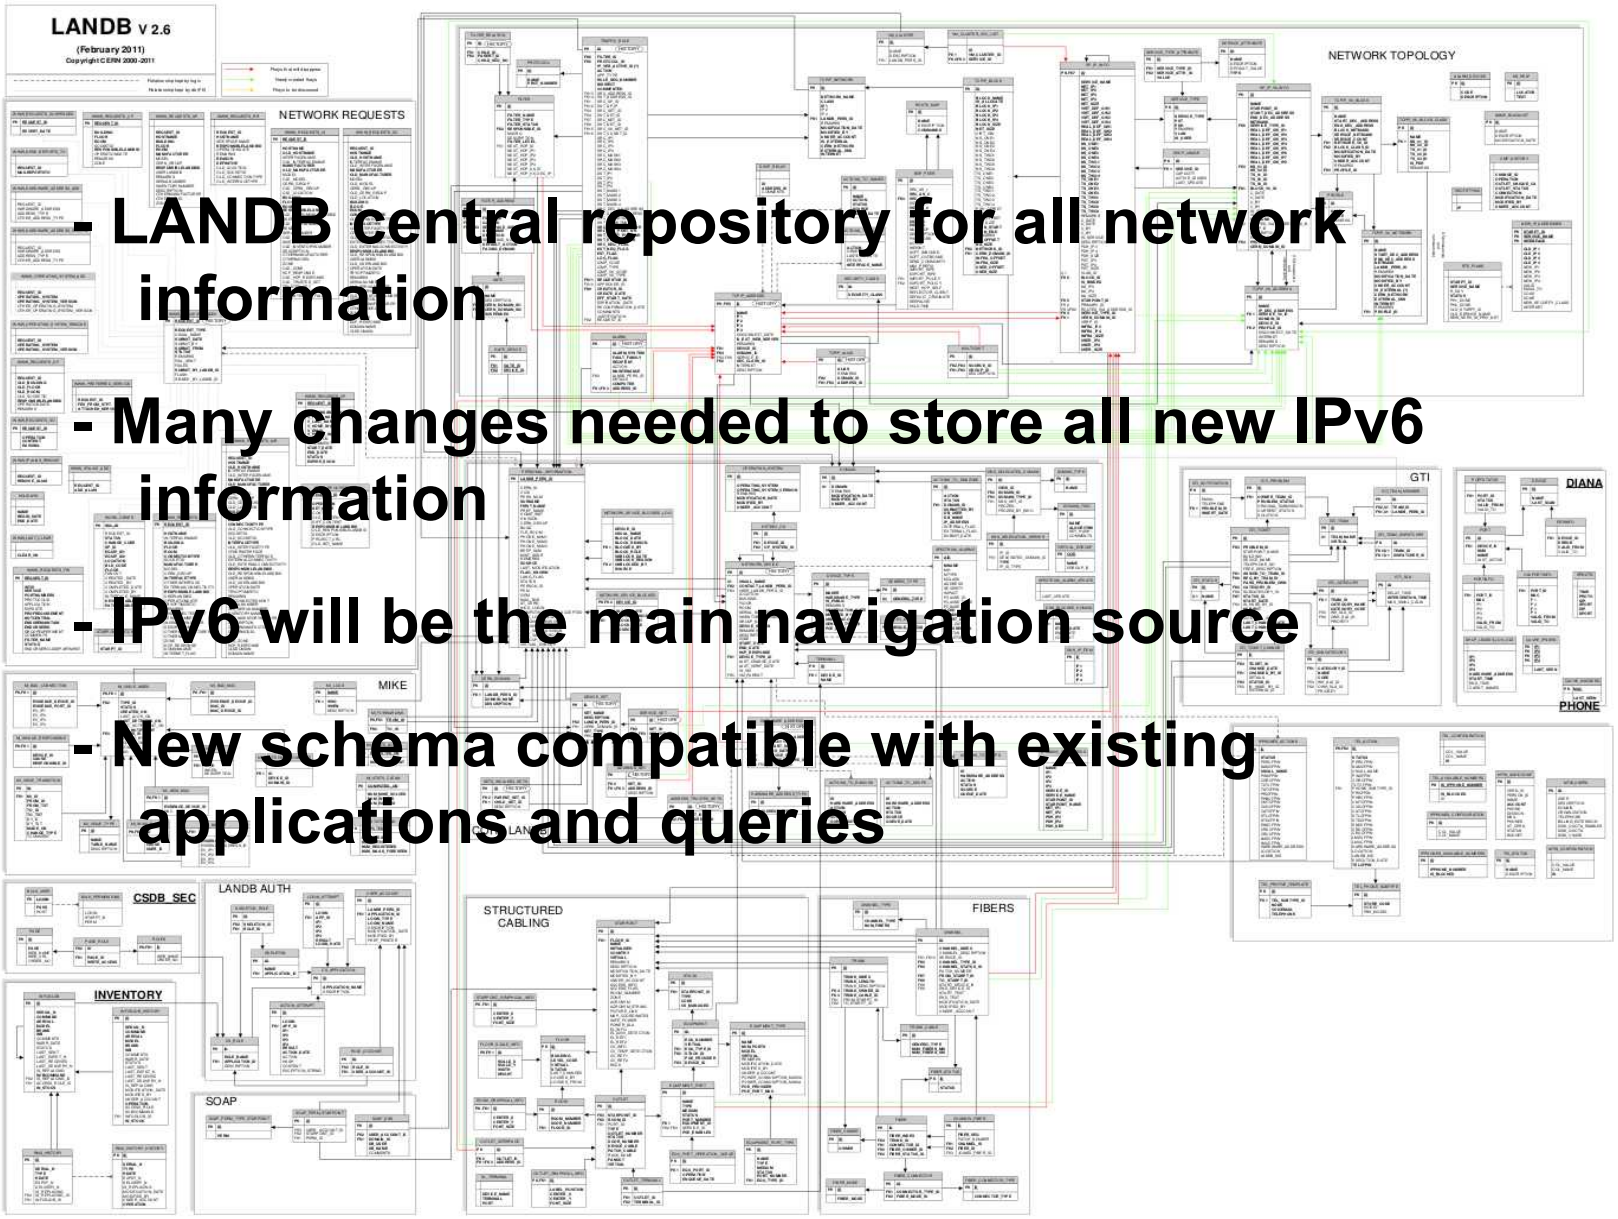

- LANDB central repository for all network information
- Many changes needed to store all new IPv6 information
- IPv6 will be the main navigation source
- New schema compatible with existing applications and queries

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

Monitoring

IPv6 will be monitored as the equivalent IPv4 counterpart

But initial monitoring not at the same level as IPv4 (upcoming missing features).

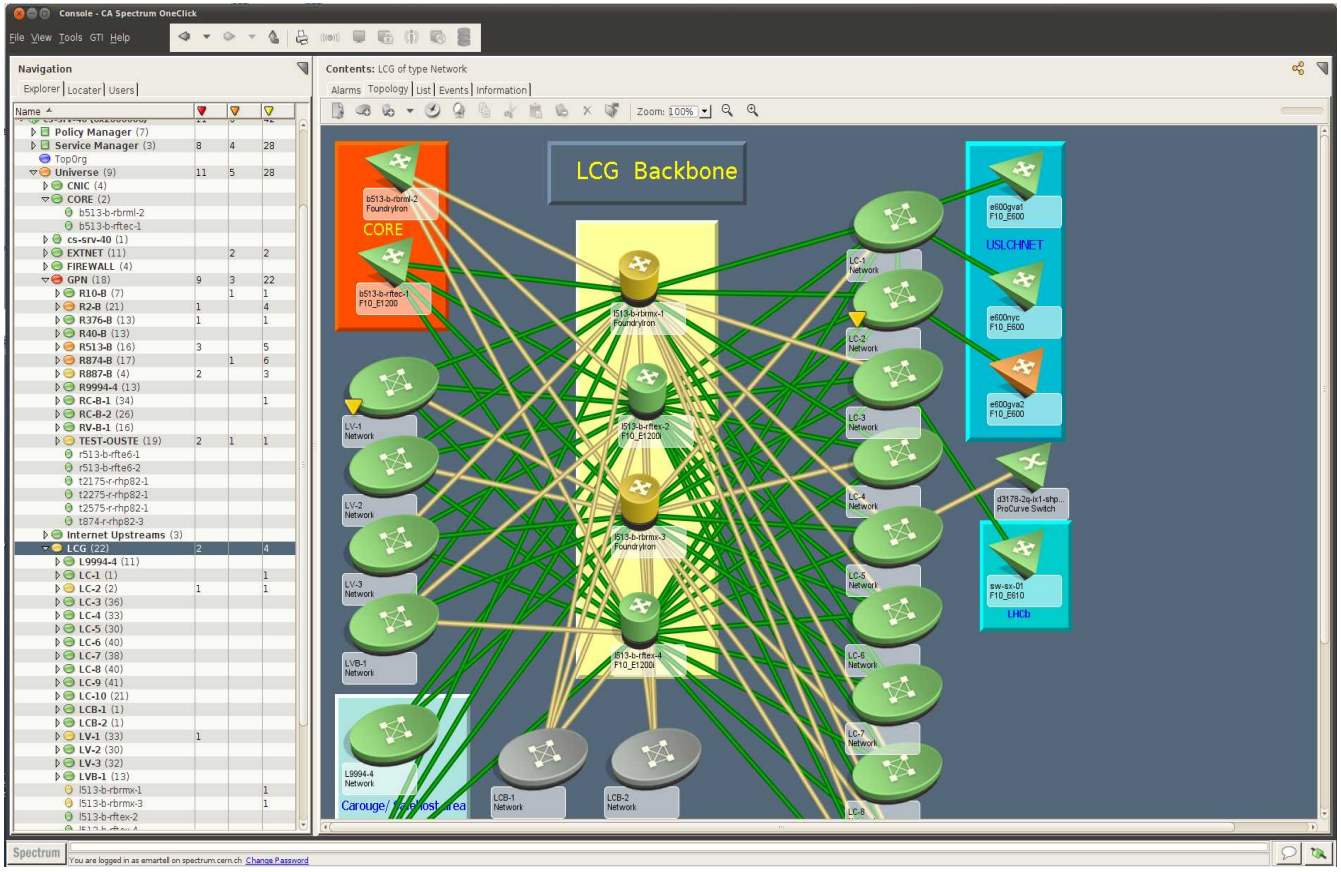

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Security

The same IPv4 security policies will be applied to the IPv6 service.

Every existing IPv4 firewall and CNIC rules will be extended with IPv6 information.

Firewall rules concerning host addresses: the IPv6 opening counterpart will be **activated only when the host administrator will declare the server IPv6 ready**.

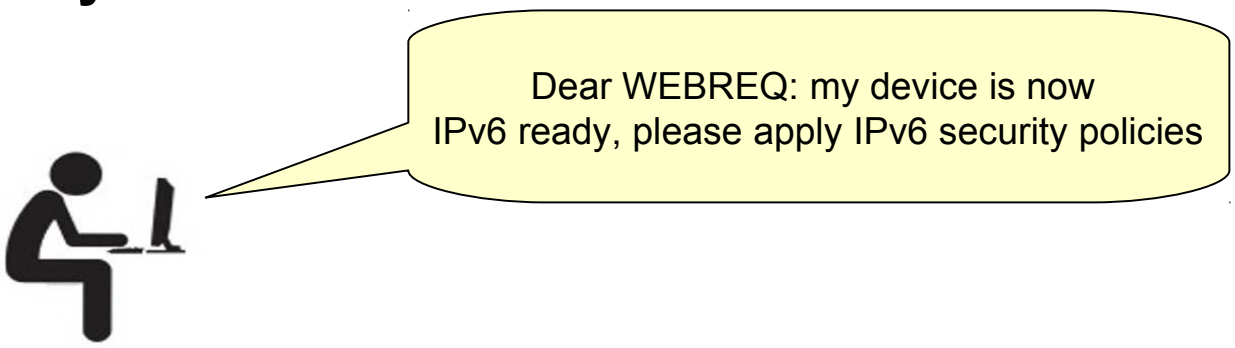

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Coexistence risks

Typical client's behavior

The choice of the IP protocol to be used is up to the client application, based on the DNS reply and its own settings.

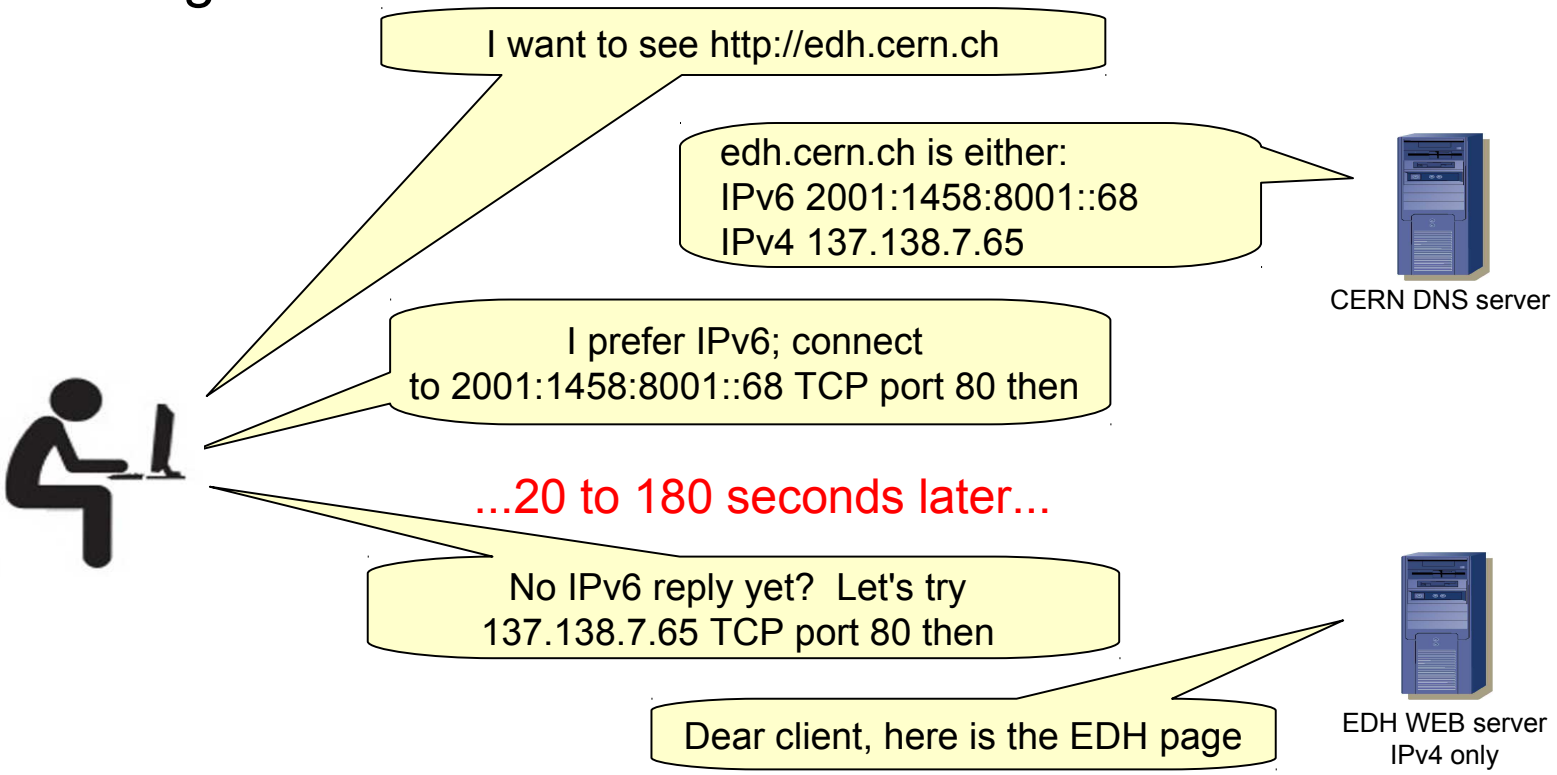

Problem: DNS is not application aware!

Control by DNS

Servers cannot decide which IP protocol the client will use.

IPv6 can be avoided by the DNS not returning the IPV6 address

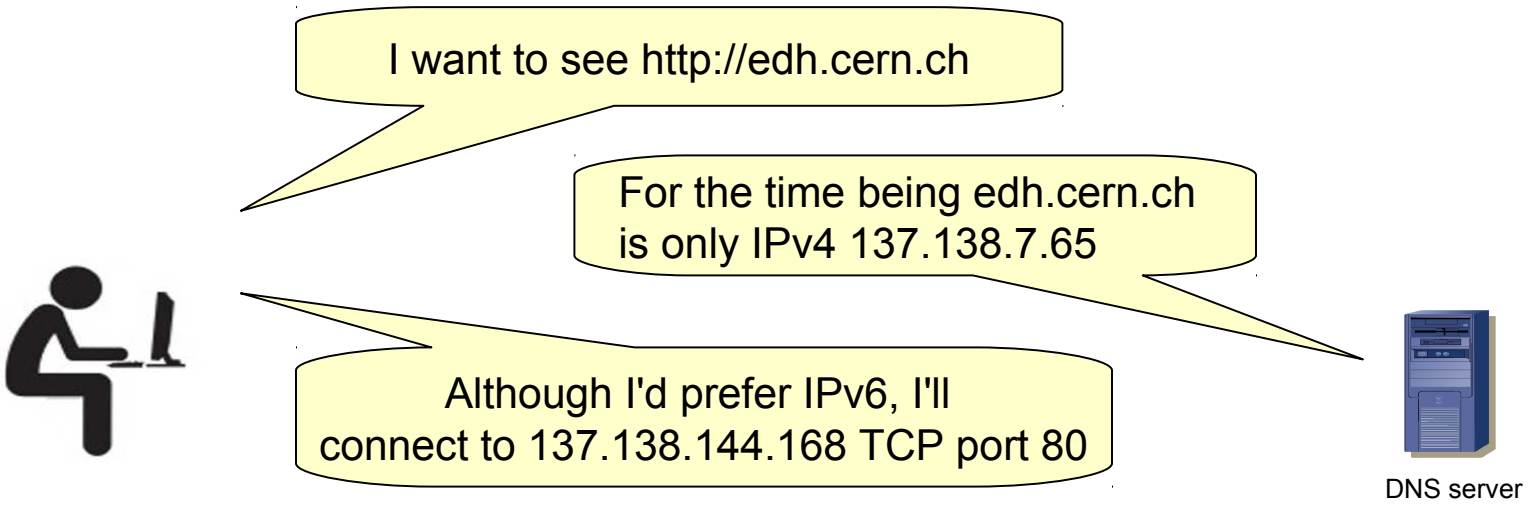

LANDB flag: IPv6 ready

The DNS device name **.cern.ch** will be resolved only with the IPv4 address **until the user declares to LANDB to be IPv6 ready via WEBREQ**.

IPv6 ready means:
- **IPv6 connectivity is OK**
- **all the server's applications are listening on both IPv4 and IPv6 protocols**

Consequences:
- **IPv6 security openings activated**
- **name.cern.ch returns IPv4 and IPv6 addresses**

LANDB flag: Not IPv6 ready

Not IPv6 ready means:
- Still testing IPv6 or Only Client machine

Consequences:
- No IPv6 security openings
- name.cern.ch for IPv4 address
- name.ipv6.cern.ch for IPv6 address

Issue with remote sites

If broken IPv6 connectivity, clients will wait up to 180secs before falling back to IPv4

If degraded IPv6 connectivity, fall back will never occur

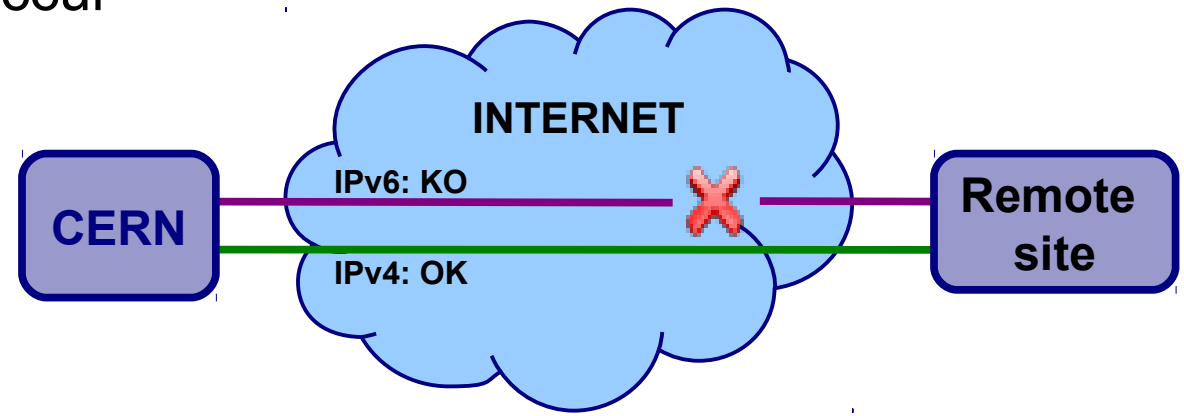

Client's perception: there's a server issue

IT-CS Work Plan

IT-CS workplan

- Testing of network devices ✔
- New LANDB schema ✔
- Addressing plan in LANDB
- Provisioning tools (cfmgr and csdbweb)
- Network configuration
- Network services (DNS, DHCPv6, Radius, NTP)
- User interface (webreq)
- IPv6 Service ready for production in 2013

Implications for IT

Everybody is concerned

Network managers

System managers

Application managers

Developers

Operation managers

System managers

Most recent versions of Windows, Linux and MacOS support IPv6.

Installation of a DHCPv6 client may be necessary.

Upgrade/replace old OSes with no/broken IPv6 support.

Local firewall configuration

Application managers

In house and open source applications
(i.e. CDB, QUATTOR, LEMON, CASTOR, GridFTP, EDH...):
- understand IPv6 addresses
- connect/listen over IPv6 and IPv4
See some recommendations in RFC4038

Commercial applications
(i.e. Oracle, LSF, printers, PLCs...):
- Ask vendors to implement IPv6 support
- Upgrade the applications

Progressing together

Share experiences

CERN IPv6 Forum organized by IT/CS

Representatives from:
- each IT group
- each department
- each experiment

Mandate: to be agreed.

Mailing list: ipv6-forum@cern.ch

Share experiences

IT/CS involved in HEPiX IPv6 WG

Other groups encouraged to join.

More information:
http://indico.cern.ch/contributionDisplay.py?contribId=45&confId=118192

mailing list: ipv6@hepix.org

Testbed setup proposal

Starting in August 2011:

- One dual stack IPv4/IPv6 switch in the Computer Centre of building 513
- Address configuration via DHCPv6
- DNS service over IPv6
- Global IPv6 connectivity via a statically configured statefull firewall
- Servers running Virtual Machines with IPv6 capabilities

Collaboration of all groups needed!

Testbed setup

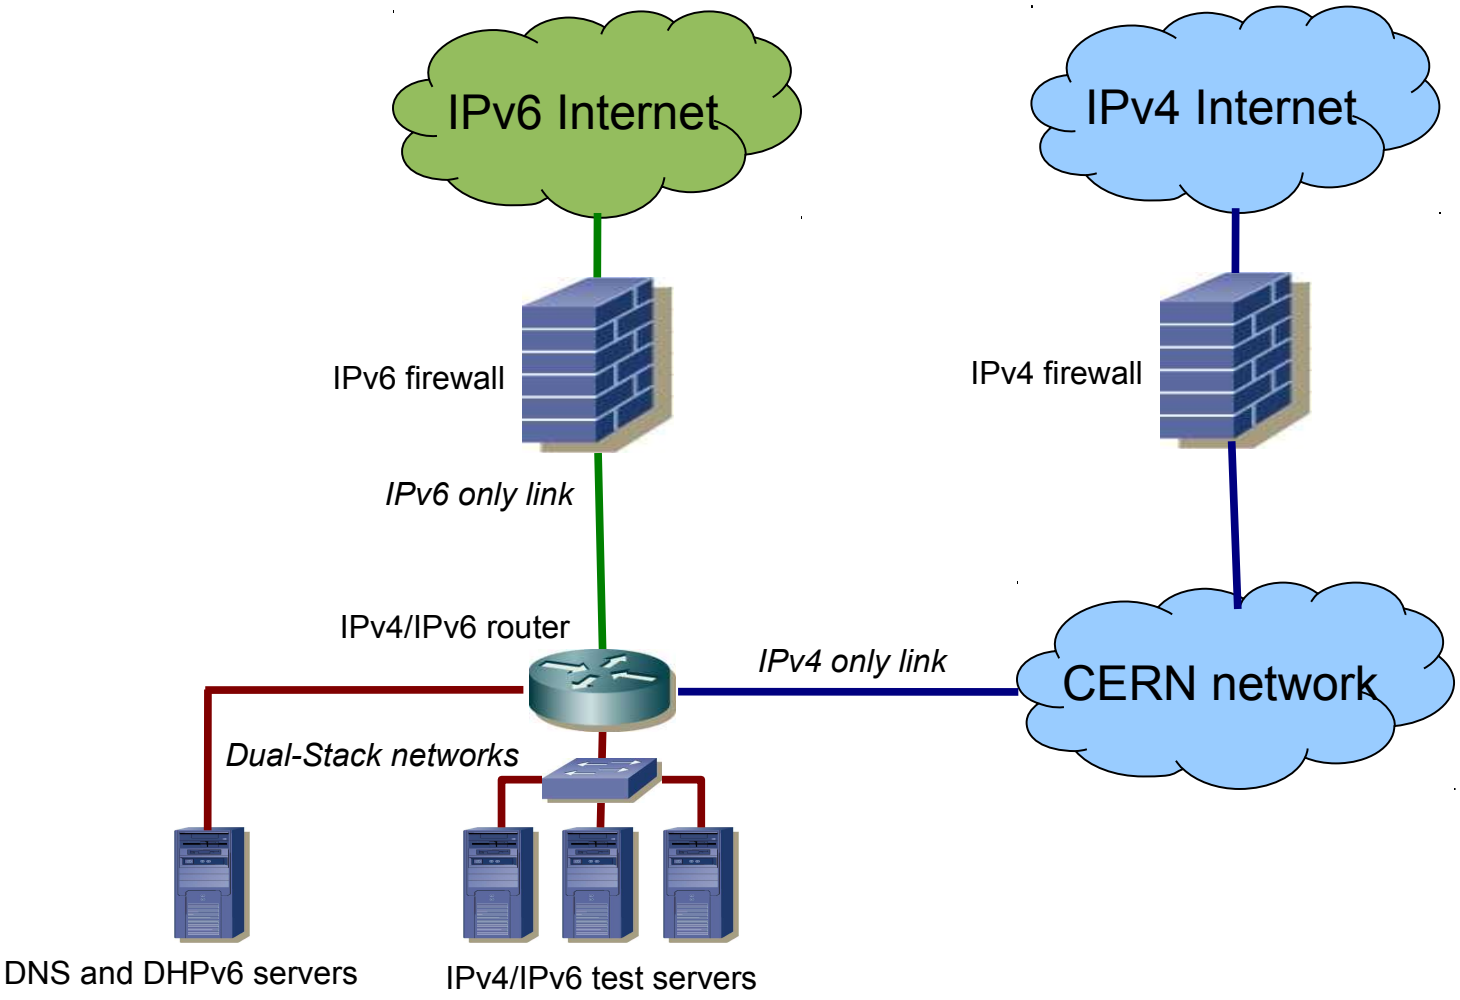

Costs... unknown yet

Costs for IT/CS

Network design/testing/deployment:
at least one Expert Network Engineer FTE for 3 years.

IT/CS database and applications:
at least three Expert Software Developers FTE for 3 years

IPv6 Internet upstreams.

Training for the group members and Support Lines

Tools for Operations (Spectrum, analyzers..)

Costs for IT and CERN

Developers and testers

Software Licenses (Oracle, LSF...)

Time for software upgrades

Money and time for hardware upgrades

Training

Operational issues

Conclusions

Conclusions

- IPv6 is necessary
- Implementation already started
- It will take time
- It will be expensive
- New operational problems will arise
- Everybody is concerned

Questions?

More information:

http://cern.ch/ipv6