CERN IPv6 Network

HEPiX 2011 - Vancouver

edoardo.martelli@cern.ch

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

Summary

- Why IPv6?
- What will change?
- How to go there?
- CERN IPv6 Network Service
- IT Work plan
- Progressing together

Why IPv6?

IPv4 address pools soon depleted

Registry exhaustion date predictions:

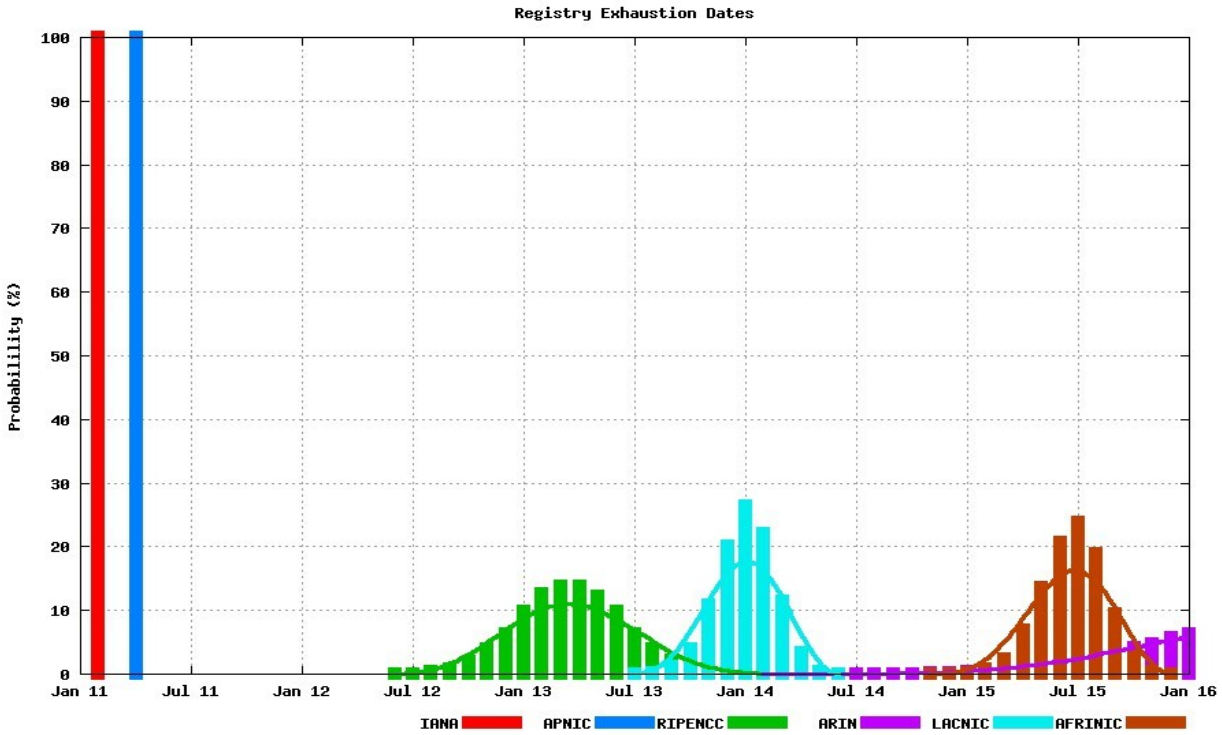

http://www.potaroo.net/tools/ipv4/rir.jpg

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, mobile devices...)
- Part of the Internet will become IPv6 only

IPv6 necessary to let the Internet growing and to deploy new services.

What will change with IPv6?

More addresses

Typical address:

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

128 bits (were 32 in IPv4); a typical major site allocation gives:

- 2^32 subnets available (the whole IPv4 space)

- 2^64 host addresses per subnet (25000 hosts per square meter on earth, per subnet)

Few new acronyms

NDP (Network Discovery Protocol):
replaces ARP

SLAAC
(StateLess Address AutoConfiguration):
allows interface autoconfiguration

EUI-64 (Autogenerated addresses)**, RA** (Router Advertisements)**, DHCPv6, OSPFv3**…
not much else ;-)

Change your mindset

- Public addresses for everything (End-to-End connectivity everywhere)
- No NAT (not even designed)
- No fear to waste
- No more Broadcast, but Multicast
- Multiple addresses per interface, even in the same IPv6 subnet

How to get there?

Transition strategies

Bridging:

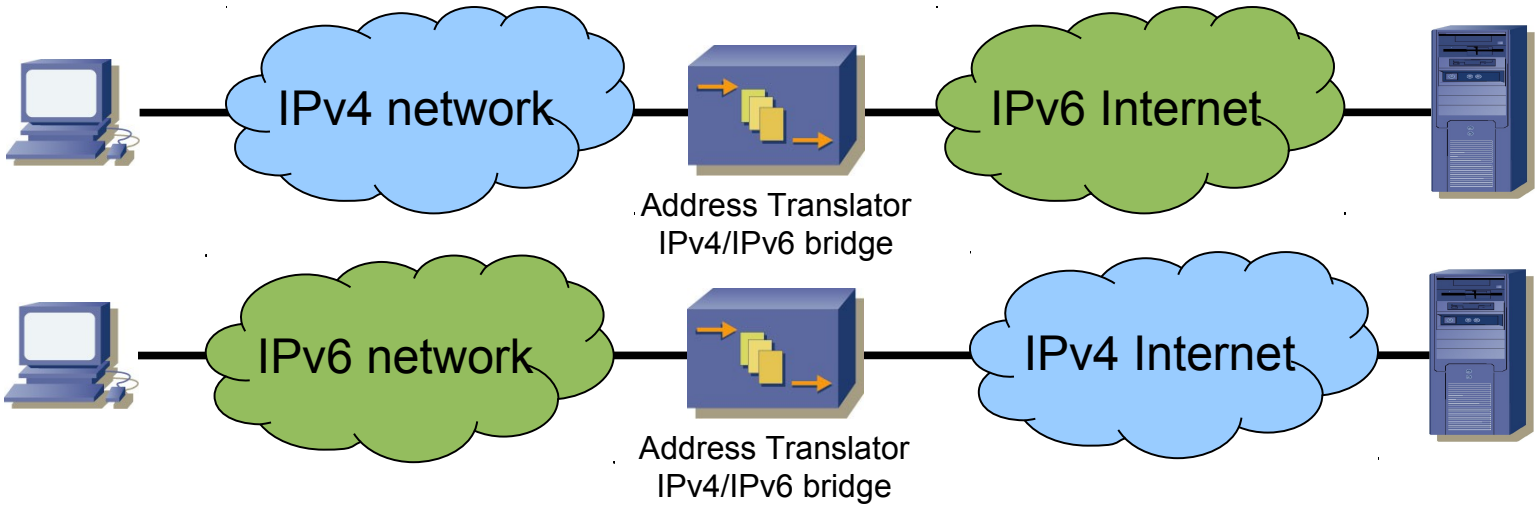

Dual Stack:

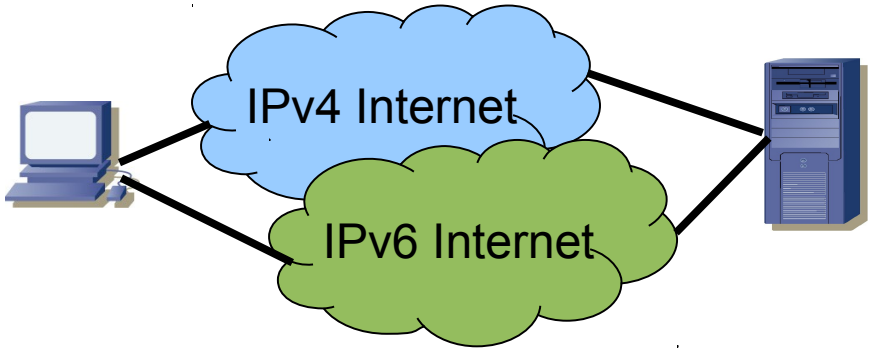

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Transition strategies

Bridging

- doesn't scale
- no end-to-end connectivity
- all typical issues of NAT
- may be good for an easy start

Dual-Stack

- The way to go!

CERN IPv6 Service

IPv6 ≥ IPv4

The IPv6 service must be at the same level of the IPv4 service (performance, reliability, security...)

Service Description

- Dual Stack
- One IPv6 address 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 IPv4 and IPv6

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

One (or more) IPv6 subnet per physical subnet, together with the IPv4 subnet.

Subnet size: /64 everywhere

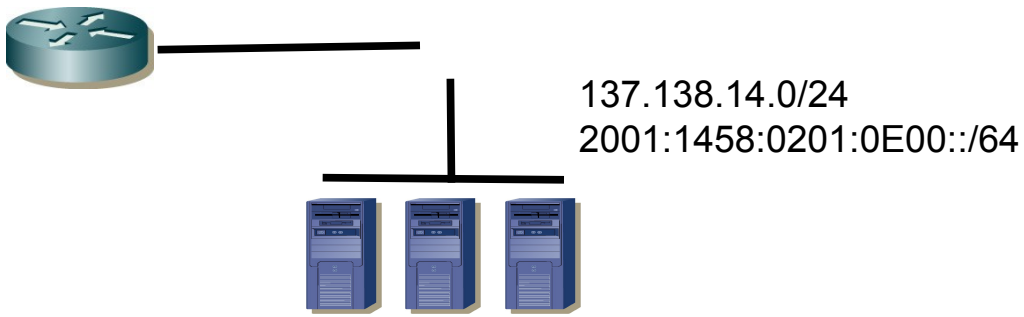

Performance

IPv6 deployed only in network devices capable of IPv6 line rate performance (no software processing)

- to avoid bad IPv6 perception
- to not impact performance of IPv4 network

Addresses assigned by DHCPv6

- IPv6 host addresses assigned by DHCPv6 servers, based on the MAC address declared in the Network database (LANDB)

No Auto Configuration

- SLAAC disabled and rogue Router Advertisement messages filtered
- EUI-64 addresses dropped at the central firewall

DNS, DHCPv6, Radius and NTP:
available over IPv6

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

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

Network Database

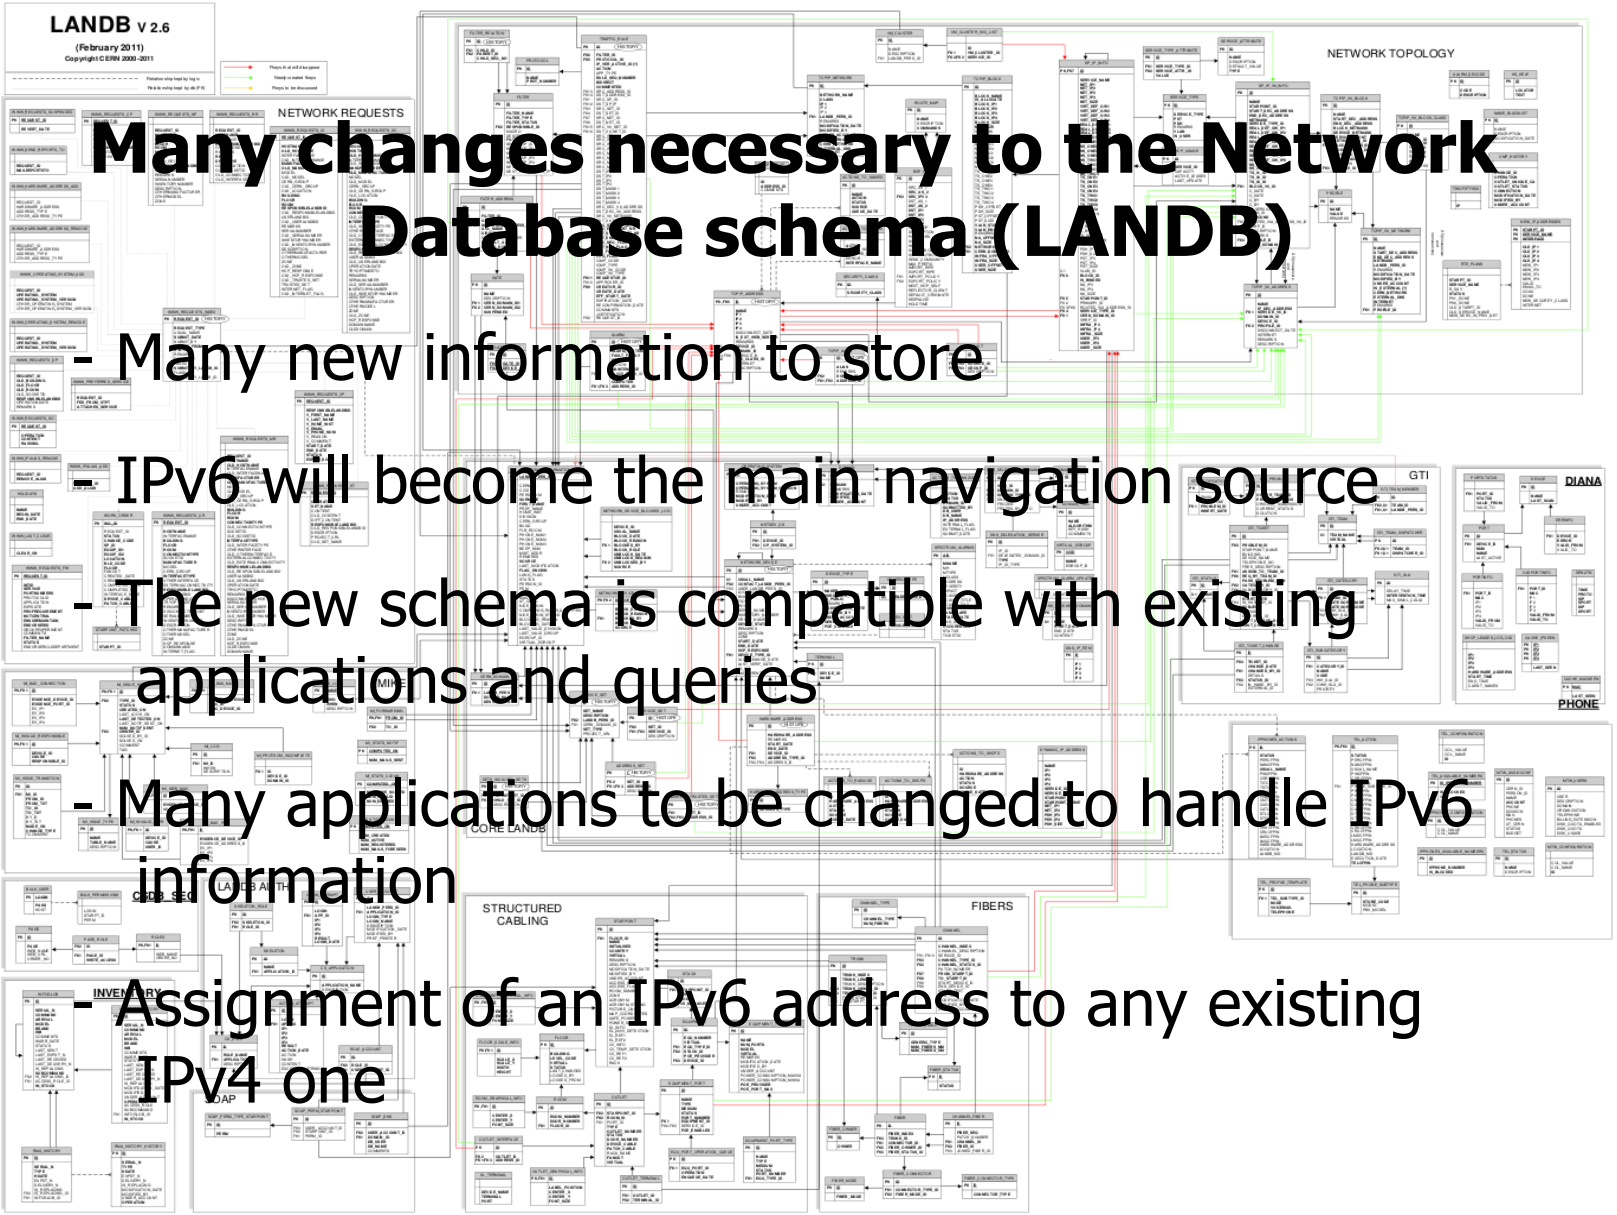

Many changes necessary to the Network Database schema (LANDB)

- Many new information to store
- IPv6 will become the main navigation source
- The new schema is compatible with existing applications and queries
- Many applications to be changed to handle IPv6 information
- Assignment of an IPv6 address to any existing IPv4 one

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

Monitoring

IPv6 will be monitored as the equivalent IPv4 counterpart

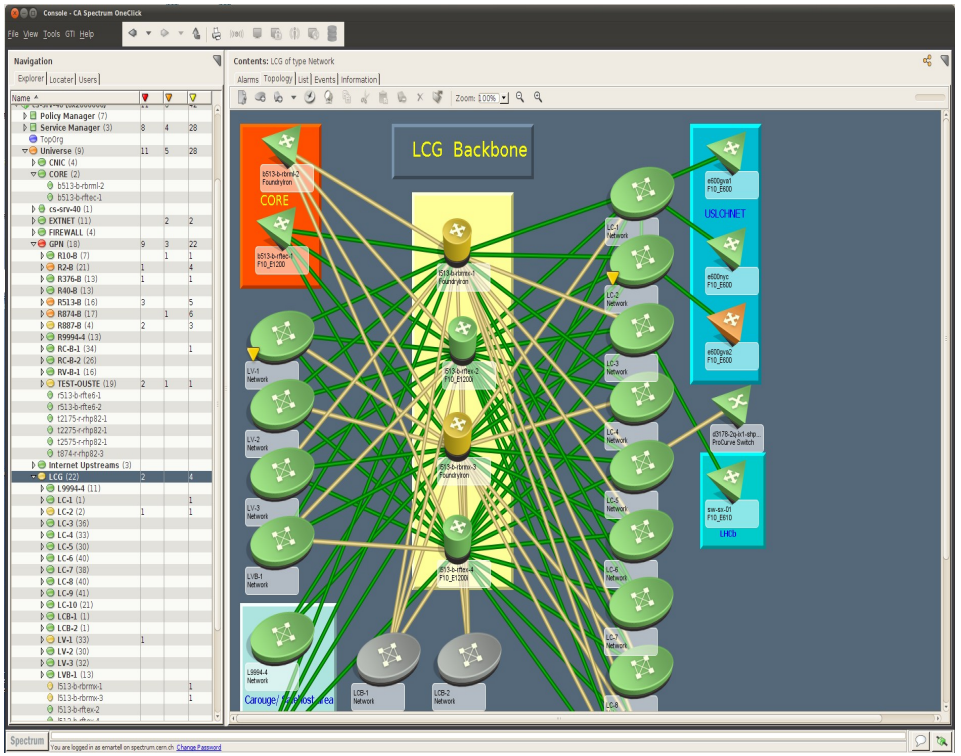

Security

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

Every existing IPv4 firewall rule will be extended with IPv6 information

(central firewall rules are stored in the Network Database)

IPv6-Ready flag

The DNS name **DEVICE-XYZ.cern.ch** will be resolved only with the IPv4 address until the responsible of the device declares it is "IPv6 Ready"

"IPv6 Ready" means:

- IPv6 connectivity is OK
- All the server's applications are listening on both IPv4 and IPv6 protocols

Consequences:

- IPv6 openings activated in the central firewall
- DEVICE-XYZ.cern.ch returns both IPv4 and IPv6 addresses

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

Not IPv6 ready

Not "IPv6 ready" means:

- Still testing IPv6 or Client-Only machine

Consequences:

- No IPv6 security openings in the central firewall
- Name DEVICE-XYZ.**cern.ch** for IPv4 address
- Name DEVICE-XYZ.**IPv6.cern.ch** for IPv6 address

Work Plan

Deployment Plan

- Testing of network devices and configurations
- Addressing Plan
- New Network Database schema for IPv6 info
- Mapping of IPv6 to existing IPv4 info
- Changes to Network Management software:
  - handling of IPv6 information (network DB)
  - use of IPv6 information (provisioning tools)
- Security
- Native IPv6 Internet upstreams (no tunnels)
- Training of Operation Teams
- Site wide deployment

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

Everybody is concerned

Network managers

System managers

Application managers

Developers

Operation managers

System managers

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

Installation of a DHCPv6 client may be necessary.

Local firewalls configuration

Application managers, developers

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 and maybe the hardware

Operation managers

New operational procedures

A whole set of unknown issues to tackle

It takes time and money to...

- make the good address plan
- find and remove network devices firmware's bugs
- get implemented missing features
- design and deploy the Network Database schema
- develop necessary new software
- check and correct the existing software
- train all the people and update the procedures

Progressing together

Share experiences

CERN IPv6 Forum

Representatives from:
- each IT group
- each Department
- each Experiment

Mailing list: ipv6-forum@cern.ch

Web site: http://cern.ch/ipv6

Share experiences

Involvement in the HEPiX IPv6 WG

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

mailing list: ipv6@hepix.org

Testbed for Users

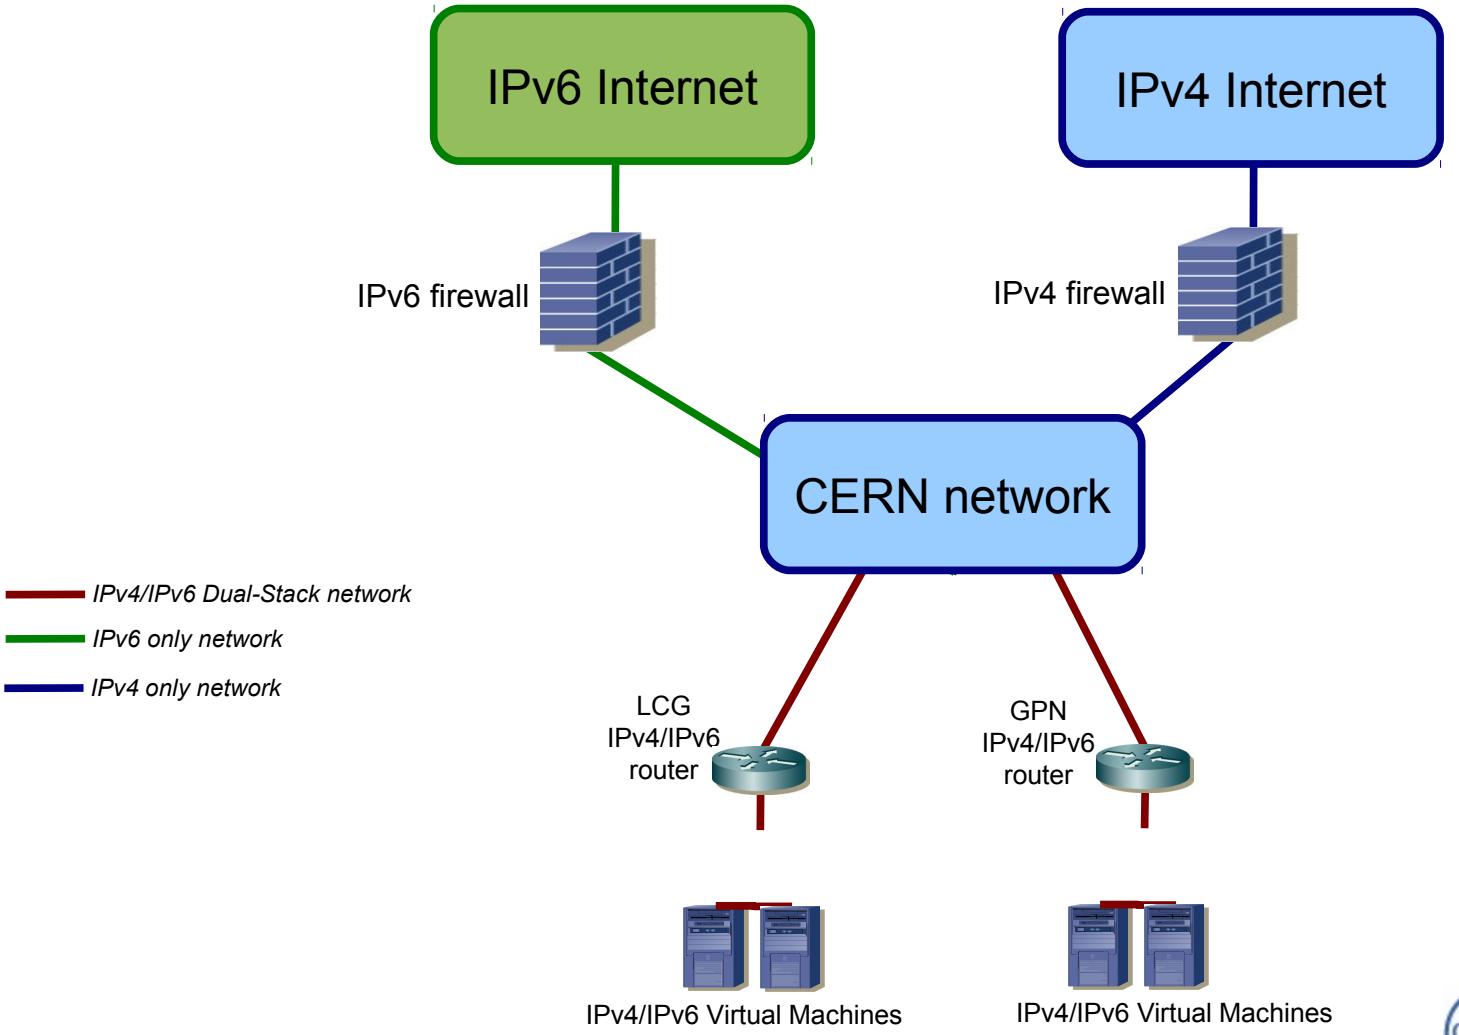

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Conclusions

Conclusions

- **IPv6 will soon be necessary**

- **The deployment will take time**

- **It will be expensive**

- **New operational problems will arise**

- **Everybody is concerned**

… it's time to start

More information:

http://cern.ch/ipv6

mailto:ipv6@cern.ch