

My thoughts on the transition to IPv6

How it is happening and how to join it

Marco d'Itri

`<md@linux.it>`

`@rfc1036`

Italian Hacker Camp 2018 - 4 August 2018

Marco, an IPv6 early adopter

- IPv6 user since 2000.
- Enabled IPv6 for 130000 web sites in 2007.
- Manages the network of a cloud infrastructure provider.

A few assumptions

Asymmetry

Using v4 NAT or IPv6 transition technologies is easy enough for access networks, but almost impossible for servers.

There is plenty of content

Google (Youtube), Facebook, Netflix, Apple and others are IPv6-enabled.

Over 50% of the traffic of typical access networks can use IPv6.

Inevitability

General adoption of IPv6 will happen, so you may as well join us before it will be too late for your business.

Transition steps

- Before the transition (when we still had new IPv4 addresses).
- The transition (now).
- After the transition (when IPv4 connectivity will be optional).

The transition is dual stack: it does **not** mean providing no connectivity to the IPv4-only internet to customers.
But often end users will not get a v4 IP anymore!

There will be no flag days, and the transition is in progress!

When we still had new IPv4 addresses:

- Anybody could get enough IPv4 addresses, just by asking nicely.
- Some had IPv6 connectivity, but it was rarely actually needed.
- Everybody had to have IPv4 connectivity.

In the RIPE region this stage ended on 14 September 2012.

The transition: you are here

Now:

- Each LIR can have a single /22 from RIPE: if they need more IPv4 addresses then they must buy them on the market.
- Adoption of IPv6 is growing.
- Everybody still needs to reach the IPv4 Internet, but some can do this only by using IPv6 transition technologies.

The large content providers use IPv6, the small ones often do not and still do not have strong incentives to deploy it.

Cloud infrastructure providers are suffering.

The transition: how long?

How long will it last?

- Years? Decades?
- This depends on how much expensive it will be for ISPs to keep buying new IPv4 addresses and manage gateways.
- This depends on how much trouble gateways will cause to customers.

Now adoption must be driven by the access networks.

IPv6 adoption in the leading countries

How much is adoption growing, actually?

IPv6 enabled users by country:

- Belgium: 54.64%
- USA: 38.24%
- Germany: 38.07%
- Greece: 34.90%
- India: 33.29%
- Uruguay: 31.42%
- Switzerland: 28.38%
- Malaysia: 27.70%
- Luxembourg: 26.82%
- **Italy: 1.87%**

(Source: Google.)

The future:

- There will still be IPv4-only islands, but it will be up to them to make themselves accessible from the IPv6 Internet.
- v4-v4 traffic will be encapsulated in VPNs.
- IPv6 connectivity will be enough: the transition technologies will gradually be retired.

At some point there will be no incentives to keep maintaining the IPv6 to IPv4 gateways.

How to start your transition

From the network border, one thing at a time:

- Start with (parts of) your network core.
- At least one IPv6 transit, then add peerings.
- DNS servers.
- Mail servers (but beware of SPF/DKIM requirements).
- Web servers.
- If you are a service provider then check your provisioning systems.

If you are an access network then this may be slightly more complex.

A simple IPv6 addressing plan

This will work well as-is for small to medium networks.

From your /29:

- Allocate a /32 (of $2^3 = 8$) to each geographic location where you will have independent transit.
- Assign a /48 (of $8 \cdot 2^{16} = 524288$) to colocation or business access customer.
- Assign a /56 (of $8 \cdot 2^{24} = 128\text{M}$) to residential access customer.
- Assign a /64 when subnetting is technically impossible (e.g. dedicated servers, mobile phones...).

Bigger networks can get more addresses and should probably aggregate internally, but there is no excuse to assign small networks to customers.

A simple IPv6 addressing plan (2)

Try to encode useful information in your network and host addresses.

E.g. use a /124 for PtP networks

- The interface on each side can be named ::a and ::b.
- Nibble-aligned for easy management of rDNS.
- Leaves many bits available for aggregation and encoding information.

You may encode information like the POP, router, interface, etc... This makes addresses self-documenting and easy to assign automatically.

Any questions?



<https://www.linux.it/~md/text/ipv6-ihc2018.pdf>
(Google ... Marco d'Itri ... I feel lucky)

