

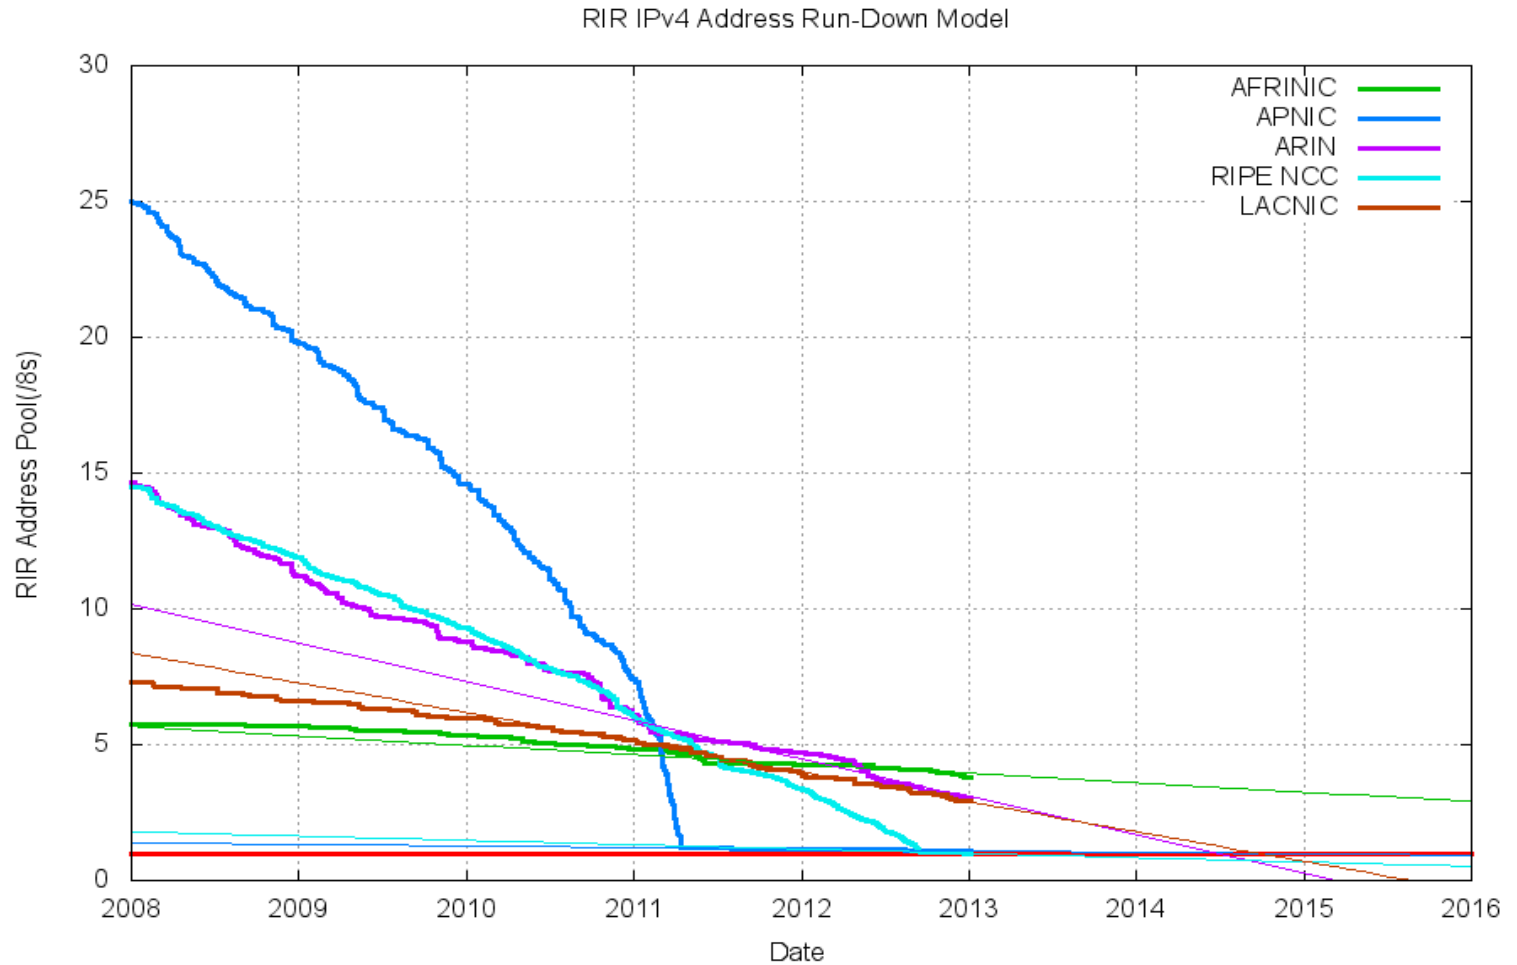
IPv4 shortage and CERN

24 January 2013

edoardo.martelli@cern.ch

- IPv4 shortage
- IPv4 and IPv6 coexistence
- Tunnels and Translations
- CERN strategy
- Conclusions

IPv4 shortage



IANA Unallocated Address Pool Exhaustion: 03-Feb-2011

Projected RIR Address Pool Exhaustion Dates and remaining /8s (16M blocks):

APNIC:	19-Apr-2011 (actual)	0.8938
RIPE NCC:	14-Sep-2012 (actual)	0.9462
ARIN:	07-Jun-2014	3.0049
LACNIC:	23-Sep-2014	2.8778
AFRINIC:	27-Feb-2021	3.8043

[as of 7th of January 2013]

- 128.141.0.0/16 (64K) - GPN dynamic addresses (~65% used)
- 128.142.0.0/16 (64K) - LCG servers in the CC (~40% used)
- 137.138.0.0/16 (64K) - GPN static addresses (~92% used)
- 188.184.0.0/16 (64K) - GPN static addresses (~5% used)
- 188.185.0.0/16 (64K) - Wigner datacentre
- 194.12.128.0/18 (16K) - Network infrastructure (~35% used)

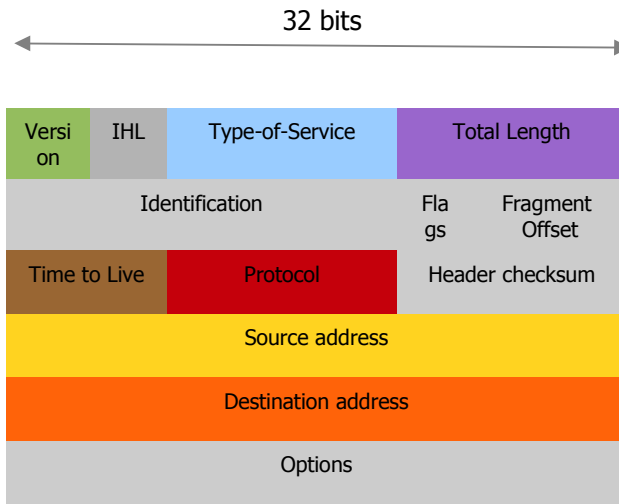
[as of 7th of January 2013]

Allocation of 188.184.0.0/16 started in October 2012: **5% allocated in only 2 months**

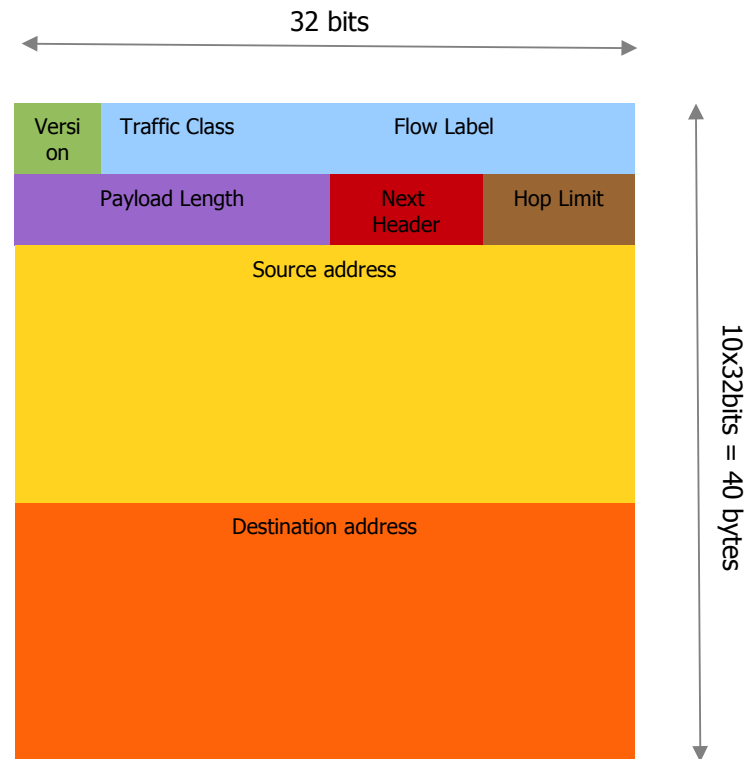
CERN can ask only for one additional /22 (1K)

IPv4 and IPv6 coexistence


IPv4 header



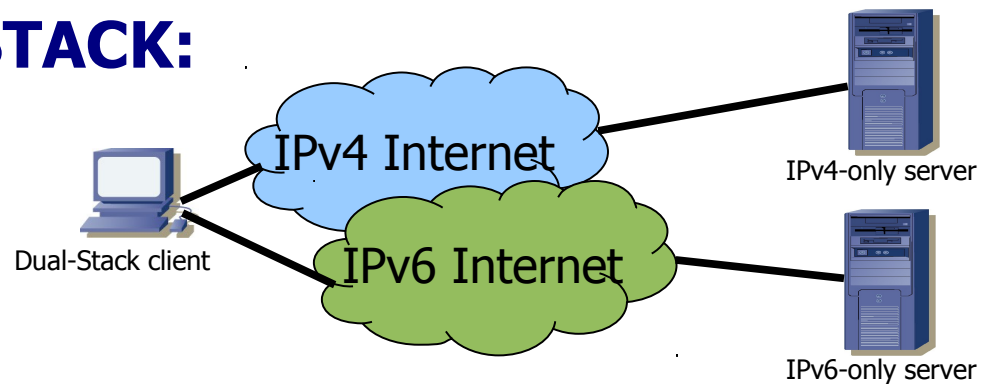
IPv6 header



Several NAT/Tunneling options:



DUAL-STACK:



Tunnelings:

- + rapid deployment (few changes)
- + cheap
- limited performance/doesn't scale well
- missing some protocol features

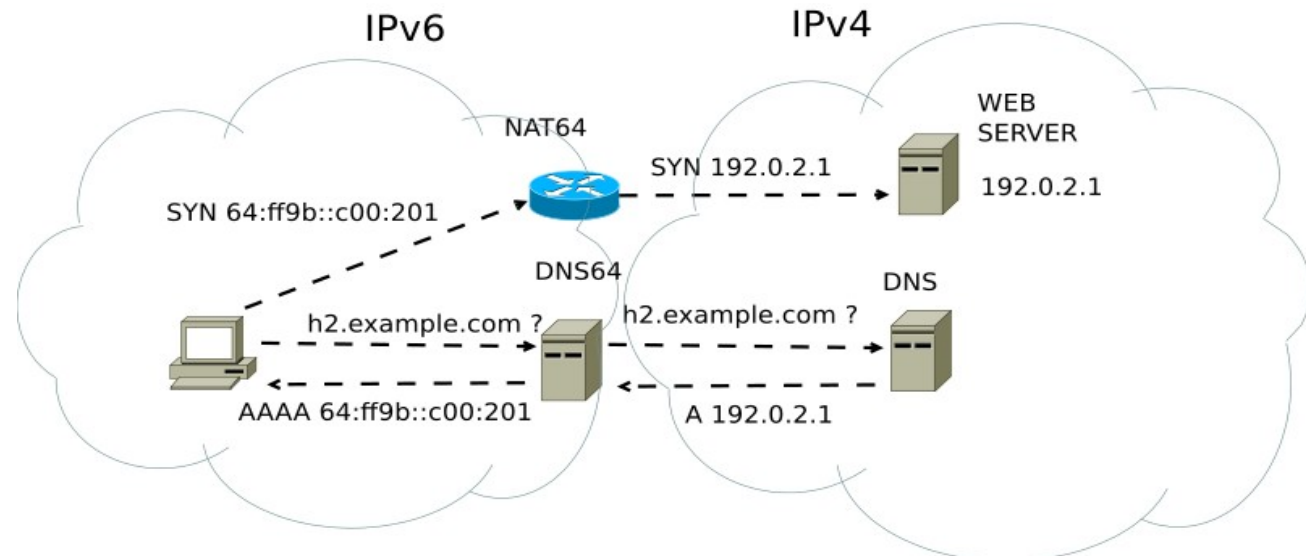
Dual-stack:

- + best performance
- + full features
- + scale well / long term solution
- re-configuration of all devices
- expensive

Tunnel and Translation protocols

NAT64 allows IPv6-only clients to reach IPv4-only servers. In general, NAT64 is designed to be used when the communications are initiated by IPv6 hosts. Static address mapping exists to allow the reverse.

The v4-v6 bridge/NAT device works in conjunction with a special DNS server that converts v4 addresses in local v6 ones.



Stateless IP/ICMP Translation (SIIT) allows communications between an IPv4 host and an IPv6 host by translating the packet headers.


Good for bidirectional reachability

It maps one v4-address to one v6-address

DS-Lite allows communications between IPv4 hosts in IPv4 islands. IPv4 clients use private IPv4 addresses.

IPv4 client packets are encapsulated into IPv6 packets when crossing the IPv6-only ISP backbone.


IPv4 packets are decapsulated and NATed by special DS-Lite CGN devices (Carrier Grade NAT), then routed to the IPv4 Internet.



6to4 allows communications between IPv6 hosts in IPv6 islands.


IPv6 packets are encapsulated into IPv4 packets when crossing the IPv4 Internet.

IPv6 encapsulated packets are exchanged between well-known 6to4 routers and relay.



Derived from 6to4 but designed to operate entirely within the end-user's ISP's network, to avoid problems due to misconfigured 6to4 routers.

Developed and currently used by Free.fr for their ADSL customers.



4rd is a mechanism to facilitate IPv4 residual deployment across IPv6 networks.

It is the reverse of 6rd.

Teredo allows IPv4-only clients to reach IPv6 only servers by establishing IPv4 tunnels to well-known Teredo relays


Similar to 6to4 but with more limitations.

MAP allows IPv4 communication between IPv4 islands. Similar to DS-lite + CGN but with the NAT functions delegated to the CPE device (customer router)

Still an IETF draft.

CERN strategy

- **Dual Stack**
- One IPv6 address assigned to every IPv4 one
- Identical performance as IPv4, no degradation
- Common provisioning tools for IPv4 and IPv6
- Same network services portfolio as IPv4
- Common security policies for IPv4 and IPv6

- Testing of network devices: *completed*
 - IPv6 Testbed for CERN users: *available*
 - New LANDB schema: *in production*
 - Addressing plan in LANDB: *in production*
 - Provisioning tools (cfmgr and csdbweb): *almost done*
 - Network configuration: *on going*
 - Network services (DNS, DHCPv6...): *on going*
 - User interfaces (webreq): *on going*
 - User training
 - IPv6 Service ready for production in 2013
- 

Current VMs adoption plan may cause IPv4 depletion during 2014

Then:

A) IPv6-only VMs

or

B) VMs with private IPv4 addresses

- + Unlimited number of VMs
- Several applications don't run over IPv6 (PXE, AFS, ...)
- Very few remote sites have IPv6
- + Will push IPv6 adoption in the WLCG community

NAT64 or SIIT may be used:

<http://tools.ietf.org/html/draft-anderson-siit-dc-00>

- + Works flawlessly inside CERN domain
- Needs NAT to reach not-CERN IPv4-only hosts:
 - may not work fairly with some application
 - still need public IPv4 addresses for external services
 - reduced performance

Conclusions

- IPv4 shortage will soon hit CERN
- Applications will have to live either with private ipv4 addresses or ipv6-only stacks
- Use of IPv6 in the WLCG have to start as soon as possible

More information:

<http://cern.ch/ipv6>