NAT64 and DNS64 in 30 ~~seconds~~ minutes

Ivan Pepelnjak (ip@nil.com)
NIL Data Communications

NIL Assist

NIL Non-Stop

CISCO PARTNER Learning Solutions

CISCO PARTNER Gold Certified

NIL Podatkovne komunikacije
Data Communications

IPv6 adoption theory: the “famous” S-curve

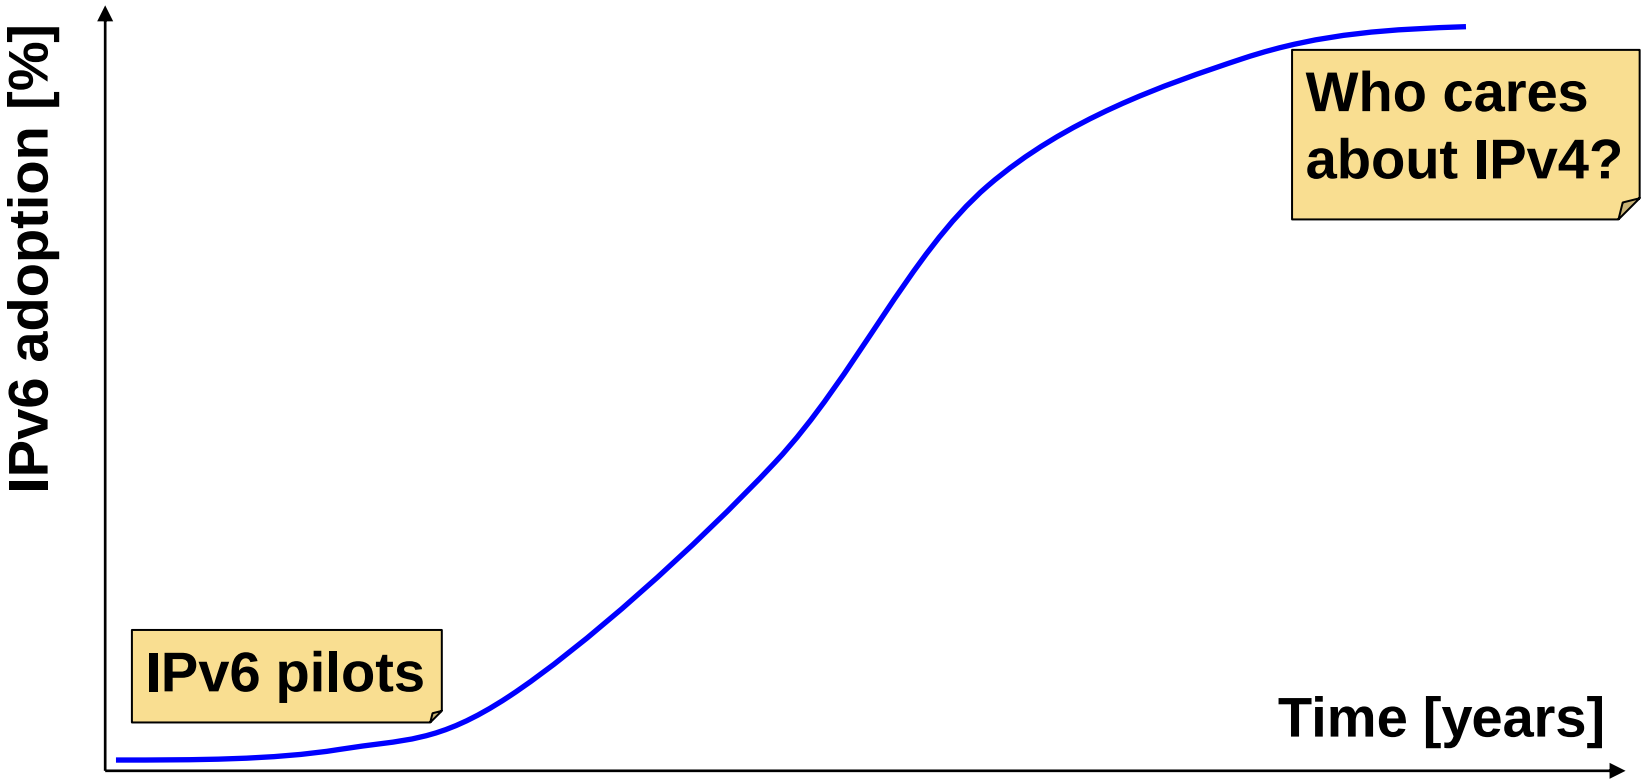

IPv6 adoption: the “ivory-tower” beliefs

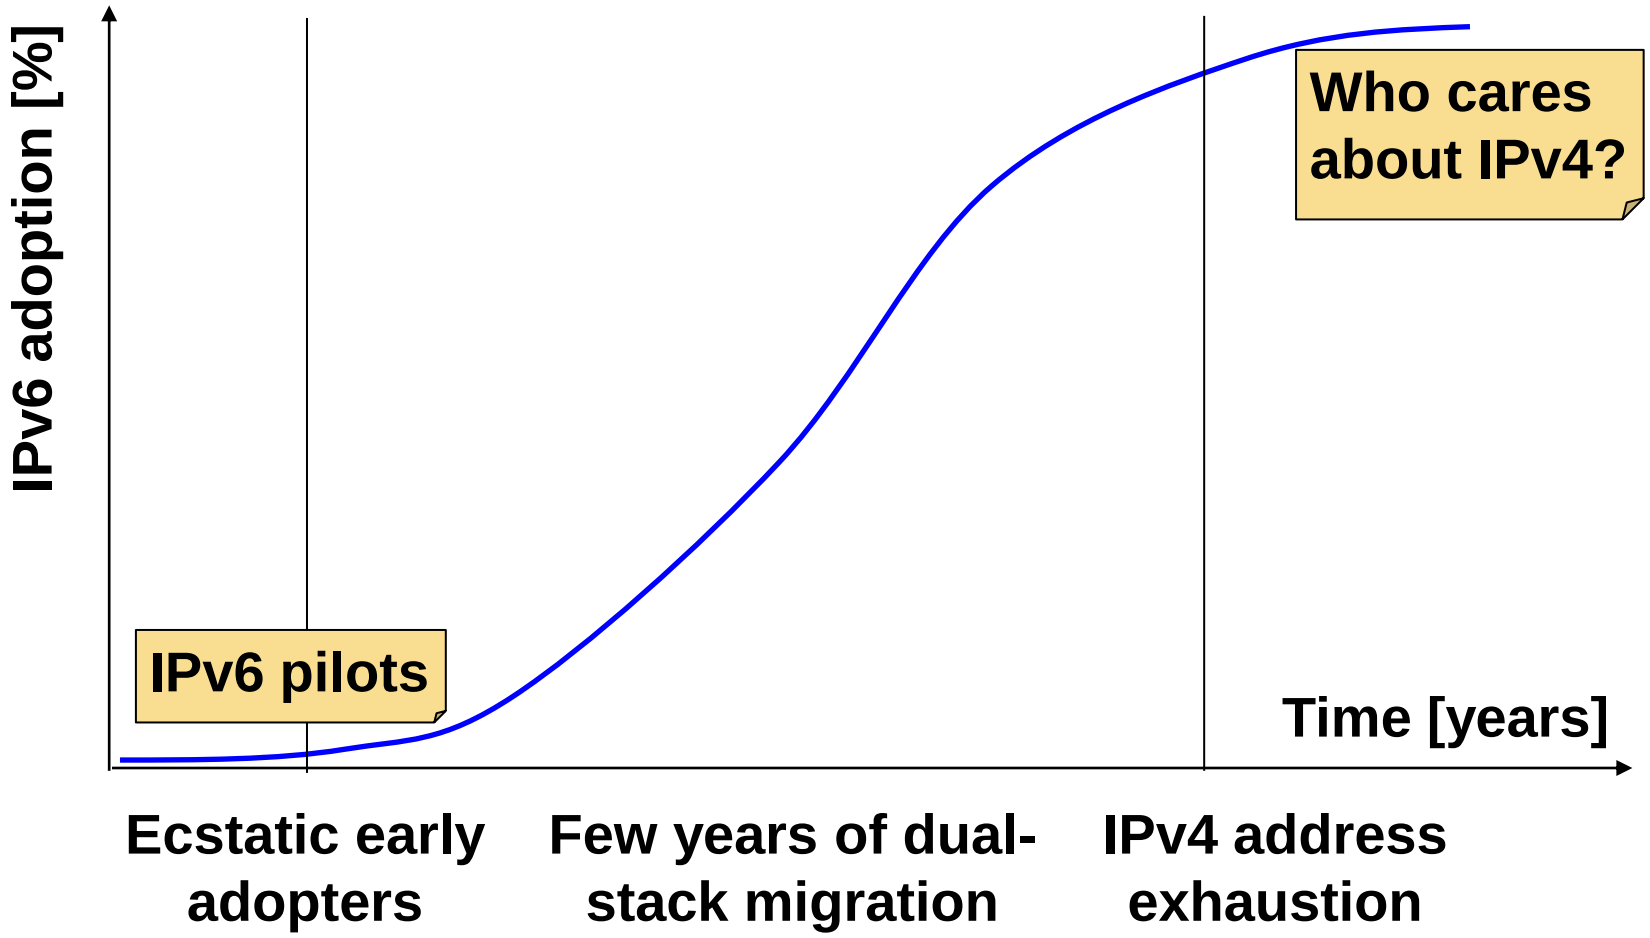

IPv6 adoption: the unpleasant reality

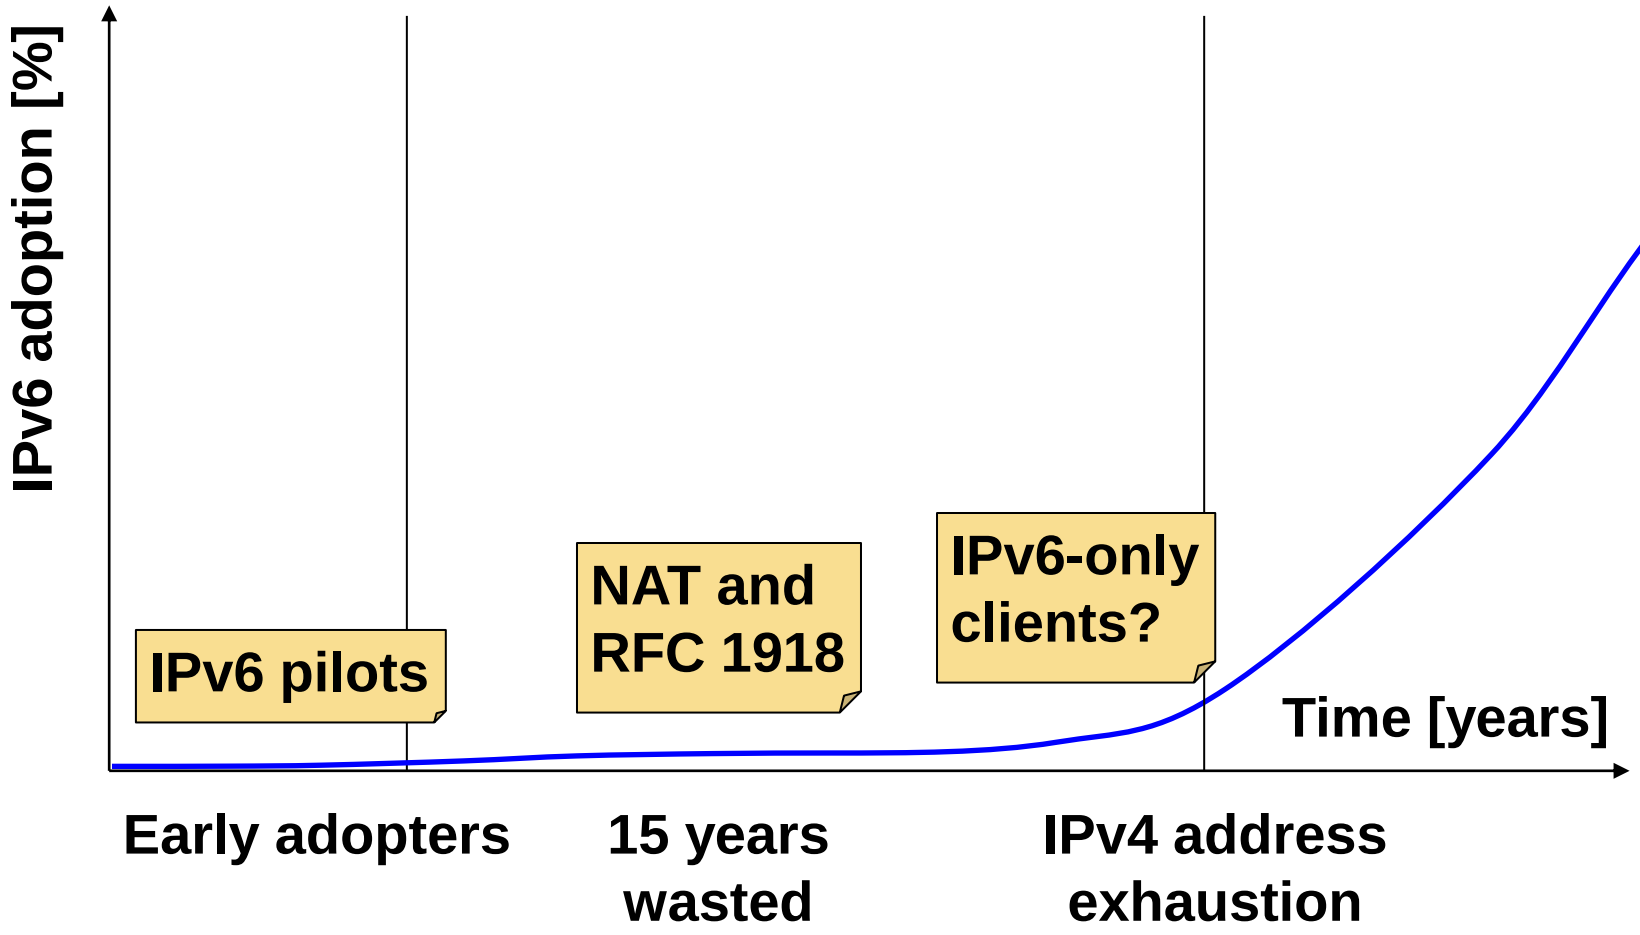

Options

Facts:

- In 2 years some clients will not get public IPv4 addresses
- These clients will have to reach IPv4 content

Options:

- CGN (large-scale NAT44)
- NAT444 (CGN + CPE NAT44)
- DS-Lite (NAT44 + 4-over-6 tunnel)
- A+P (DS-Lite with preconfigured port ranges)
- NAT64

NAT options: IPv4 only

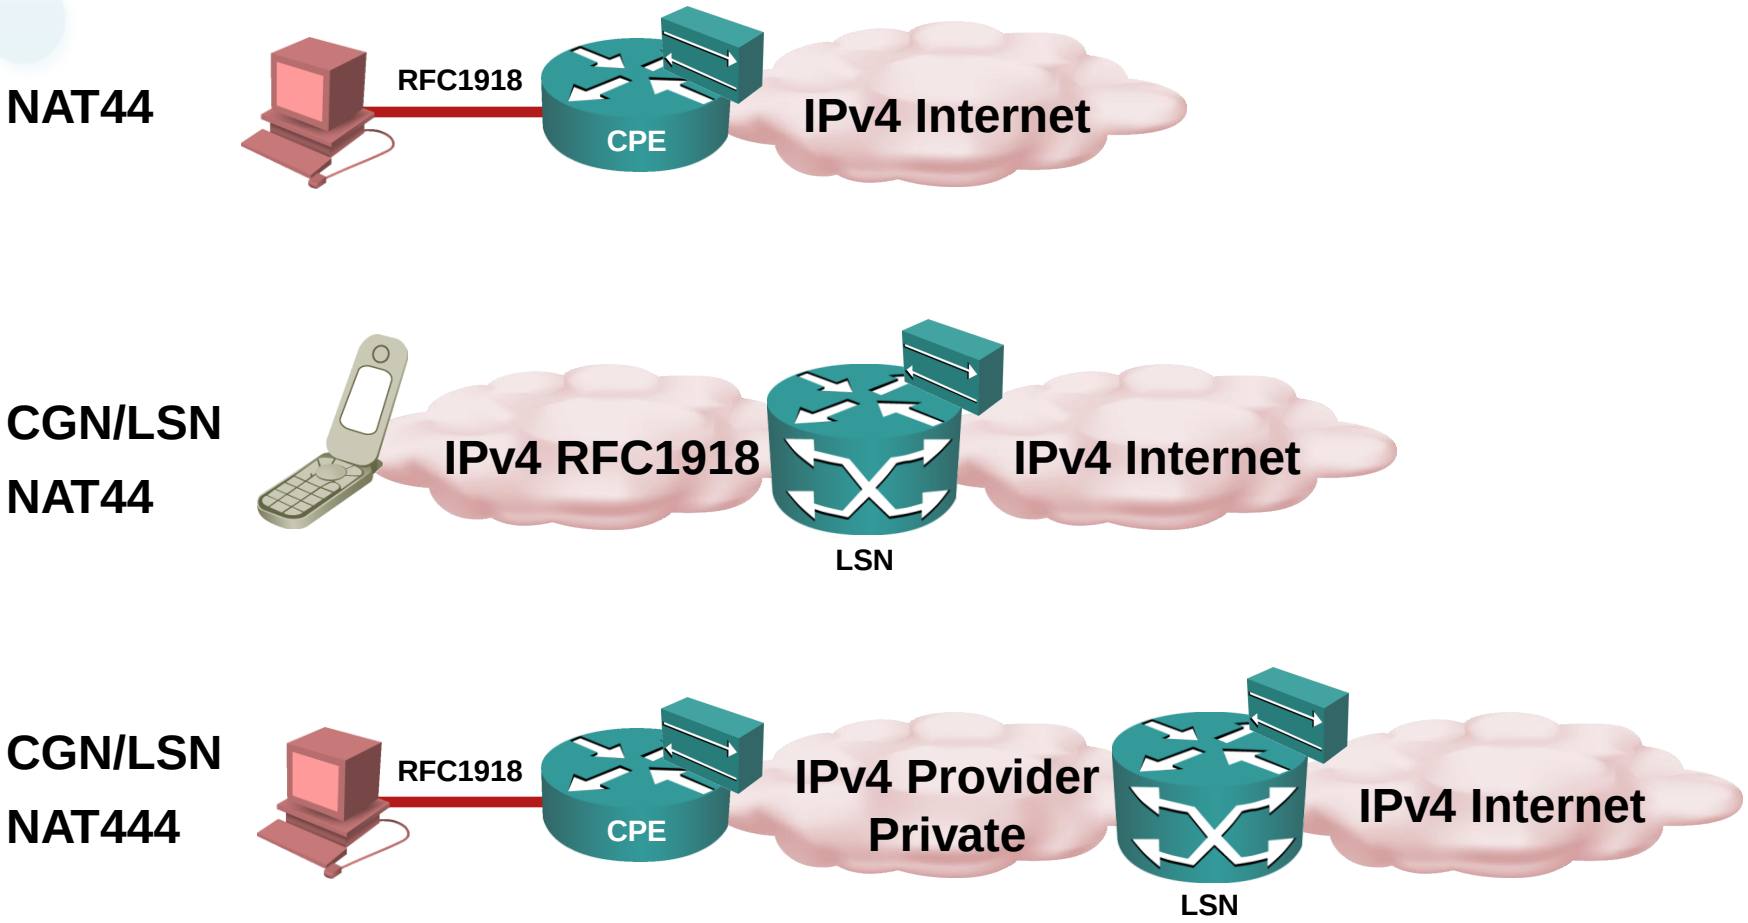

NAT options: IPv6 + IPv4

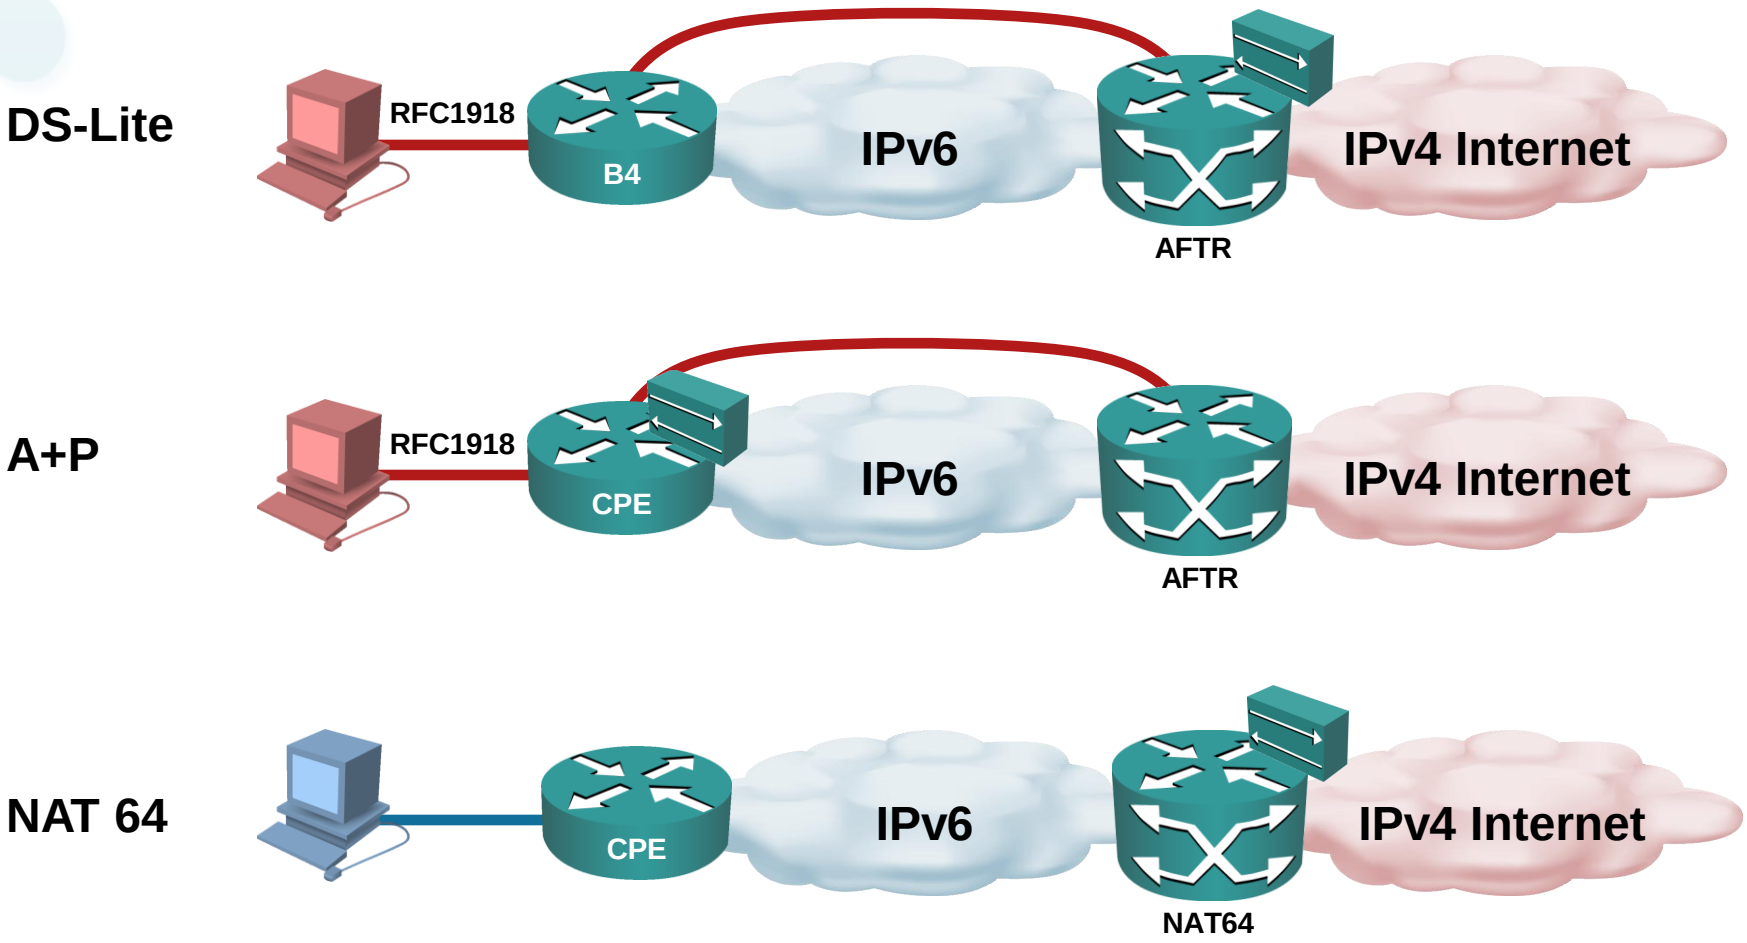

NAT is bad ... Is it really?

Facts:

- Any NAT is worse than end-to-end Internet
- Dual NAT is worse than NAT (scrap NAT444)
- NAT with ALG is really bad (scrap NAT-PT, see RFC 4966)
- NAT is OK for outbound client-server sessions
- NAT + STUN/TURN works for peer-to-peer sessions
- We need some NAT to survive past IPv4 address exhaustion

Personal opinion:

- NAT64 or DS-Lite/A+P are reasonable options

What went wrong with NAT-PT

NAT-PT (RFC 2766) = NAT64 + NAT46 + DNS ALG

- Academic “we will bring world peace” approach

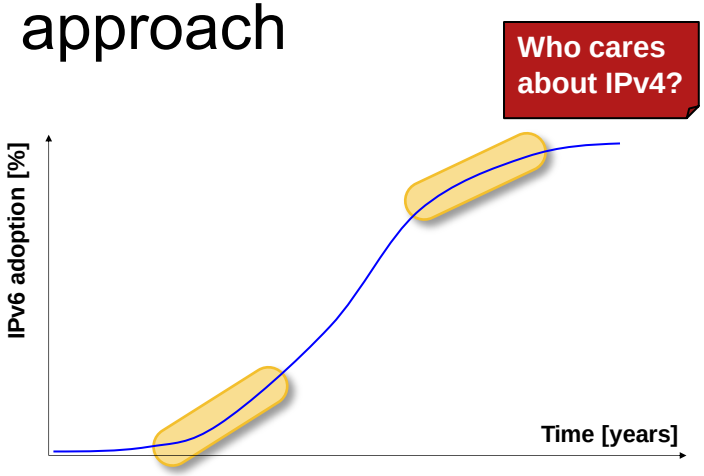

DS-Lite = NAT44 over IPv6

- Well-known solution (and problems)
- Large-scale

NAT64 = limited scope

- IPv6 client to IPv4 server
- NAT46 is useless

NAT64 topology

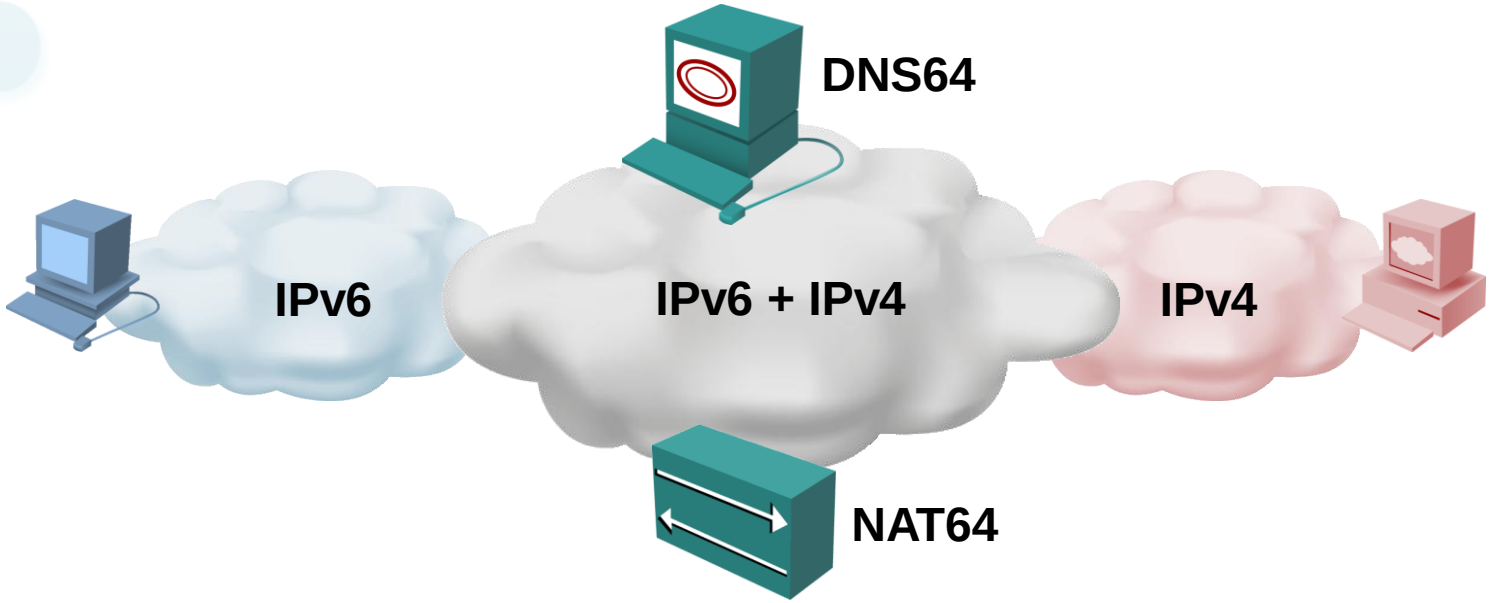

- An IPv6 prefix (well-known or network-specific) is dedicated to mapped IPv4 addresses
- DNS64 converts A records into AAAA records using NAT64 prefix, serves A and AAAA records to the client
- NAT64 router advertises NAT64 prefix into IPv6 network to attract traffic toward IPv4 servers

DNS64 in action


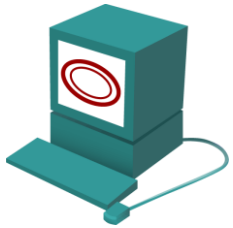


Q: AAAA for example.com

Q: AAAA for example.com

R: name error

Q: A for example.com

R: example.com (A) = 192.0.2.33

DNS64 translation for WKP

R: example.com (AAAA) = 64:FF9B::192.0.2.33
example.com (A) = 192.0.2.33

DNS64 in action (end-to-end IPv6)

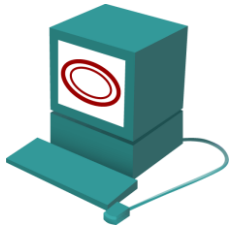

Q: AAAA for example.com

Q: AAAA for example.com

R: example.com (AAAA) = 64:FF9B::192.0.2.33

R: example.com (AAAA) = 64:FF9B::192.0.2.33

Native IPv6 communication w/o NAT64

NAT64 in action


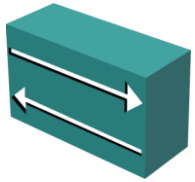


TCP SYN S=C-v6 D=WKP-v6

Translate WKP-v6 into IPv4
Pick free IPv4 addr/port from pool
Build NAT session entry

TCP SYN S=NP-v4 D=S-v4

TCP ACK S=S-v4 D=NP-v4

Translate NP-v4 + port into C-v6

TCP ACK S=WKP-v6 D=C-v6

NAT64: dirty details

NAT64 prefix

- Any /32, /40, /48, /56, /64 or /96 prefix
- WKP = 64:FF9B::/96
- Recommendation: use /64 for NSP

Stateful NAT64

- Very similar to PAT (stateful NAT44)
- Individual TCP and UDP sessions + ICMP replies are translated
- Source IPv6 address + port number used in lookup

Stateless NAT64

- Each IPv6 address is translated into one IPv4 address
- Only ICMP packets and IP headers are translated
- Limited use: IPv6 only servers

NAT64 versus DS-Lite

NAT64

- IPv6 to IPv4 NAT
- Native transport
- DNS 64 = DNS ALG
- No CPE or network modifications
- IPv6-only hosts
- NAT64 largely unknown

DS-Lite

- IPv4 to IPv4 NAT
- 4over6 Tunnel
- No DNS(SEC) interaction
- Requires CPE support
- Does not need host IPv6 (not even dual-stack)
- NAT44 well tested

NAT64 in enterprise networks

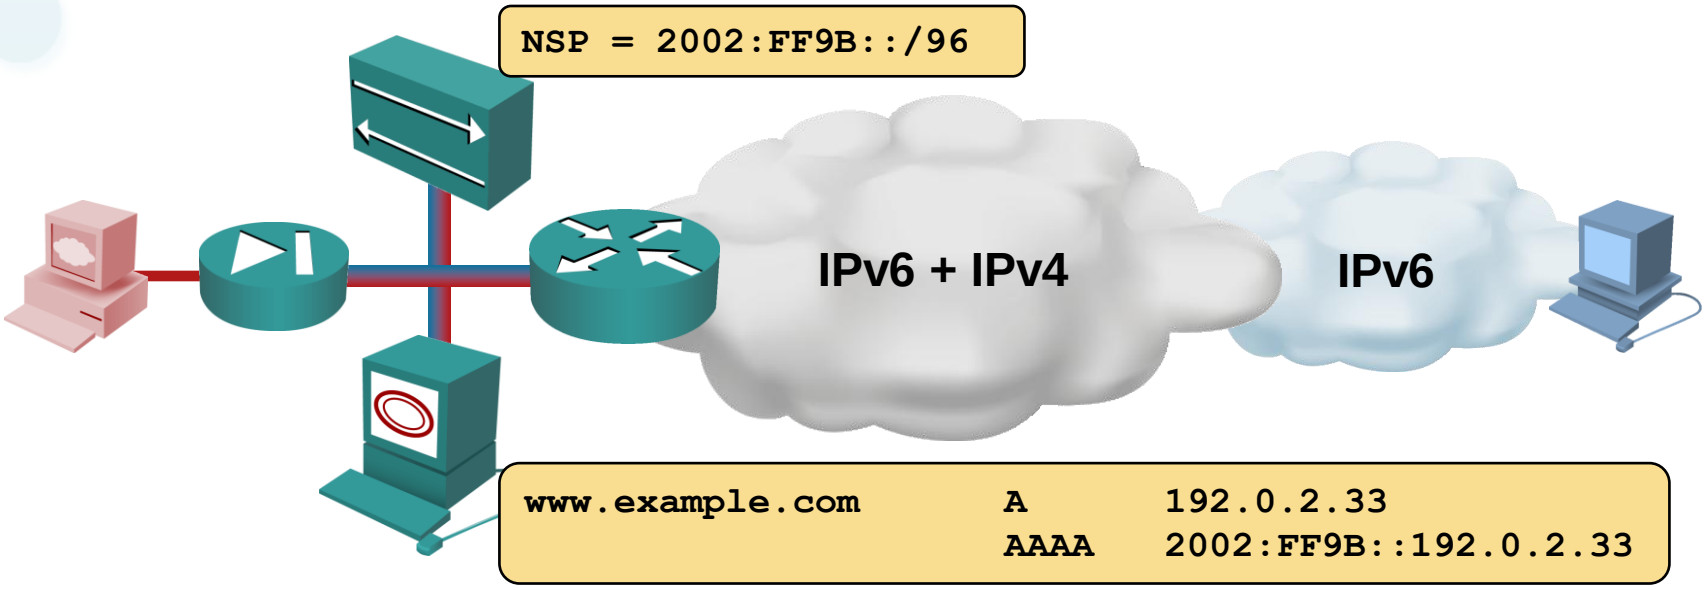

- Use NAT64 to make IPv4-only servers available to IPv6 clients
- Static entries in DNZ zone; DNS64 is not needed

Implementations

- **Open-source:** Ecdysis
- **Microsoft:** Forefront UAG DirectAccess
- **Cisco:** CGv6
- **Ericsson:** field trials

NAT64 is also (sort-of) part of NAT-PT

Conclusions

- We are not prepared for IPv4 address exhaustion
- We will not survive without NAT
- Best options: NAT64 or DS-Lite/A+P
  - Push NAT64 – it promotes IPv6 clients
- NAT64 is not NAT-PT
  - 6-to-4 only
  - DNS ALG not in the forwarding path
- NAT64 also solves legacy server problems