IPv6 at the RIPE NCC

James Aldridge

Overview

- Network
- Servers
- Services
- RIPE Meeting experiments

The Network - Background

- The RIPE NCC network connects three locations in Amsterdam using Gigabit Ethernet over dark fibre:
  - The Office (Singel 258)
  - NIKHEF (Kruislaan 409)
  - Telecity (Kuiperbergweg 13)
- Connections to the AMS-IX at NIKHEF and Telecity
- Switched layer 2 network carrying multiple VLANs

The Network - Background

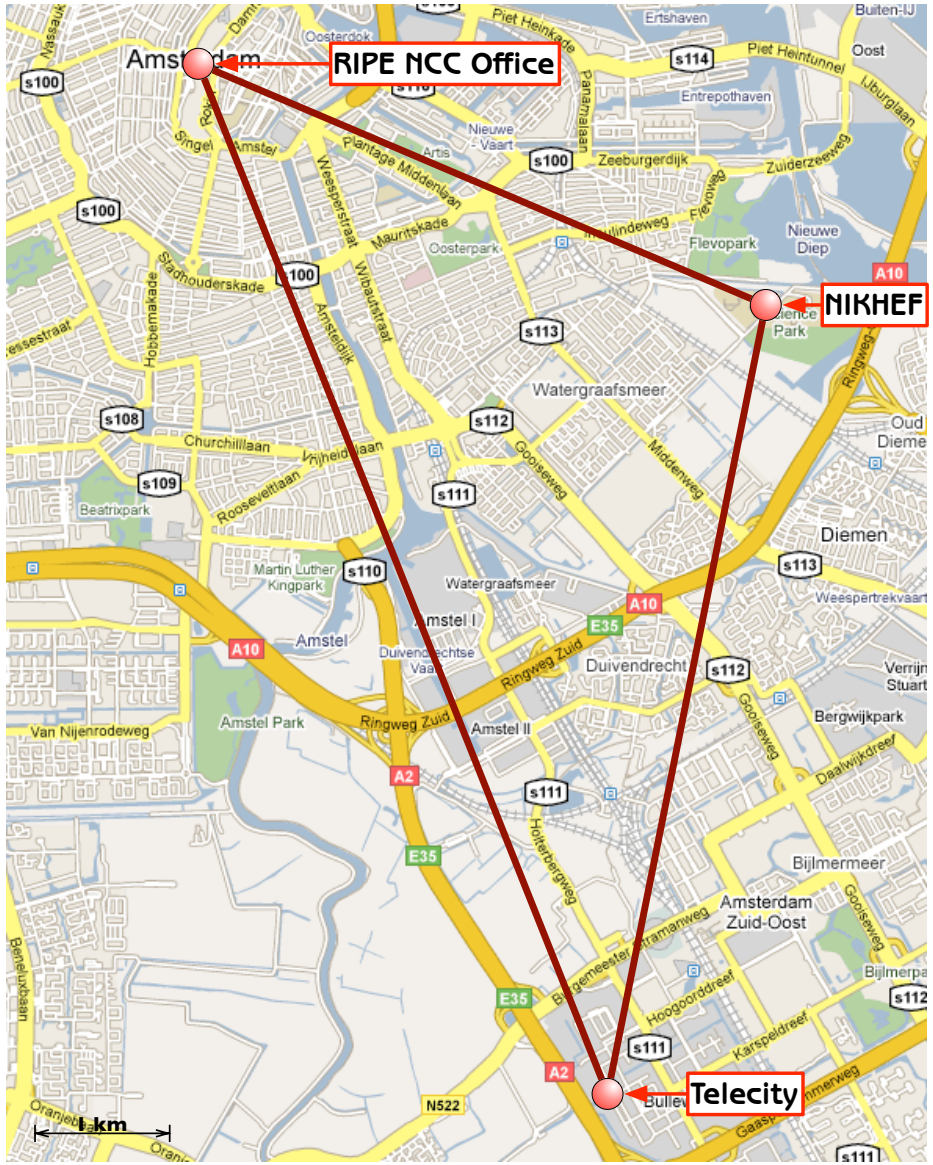

The Network - Background

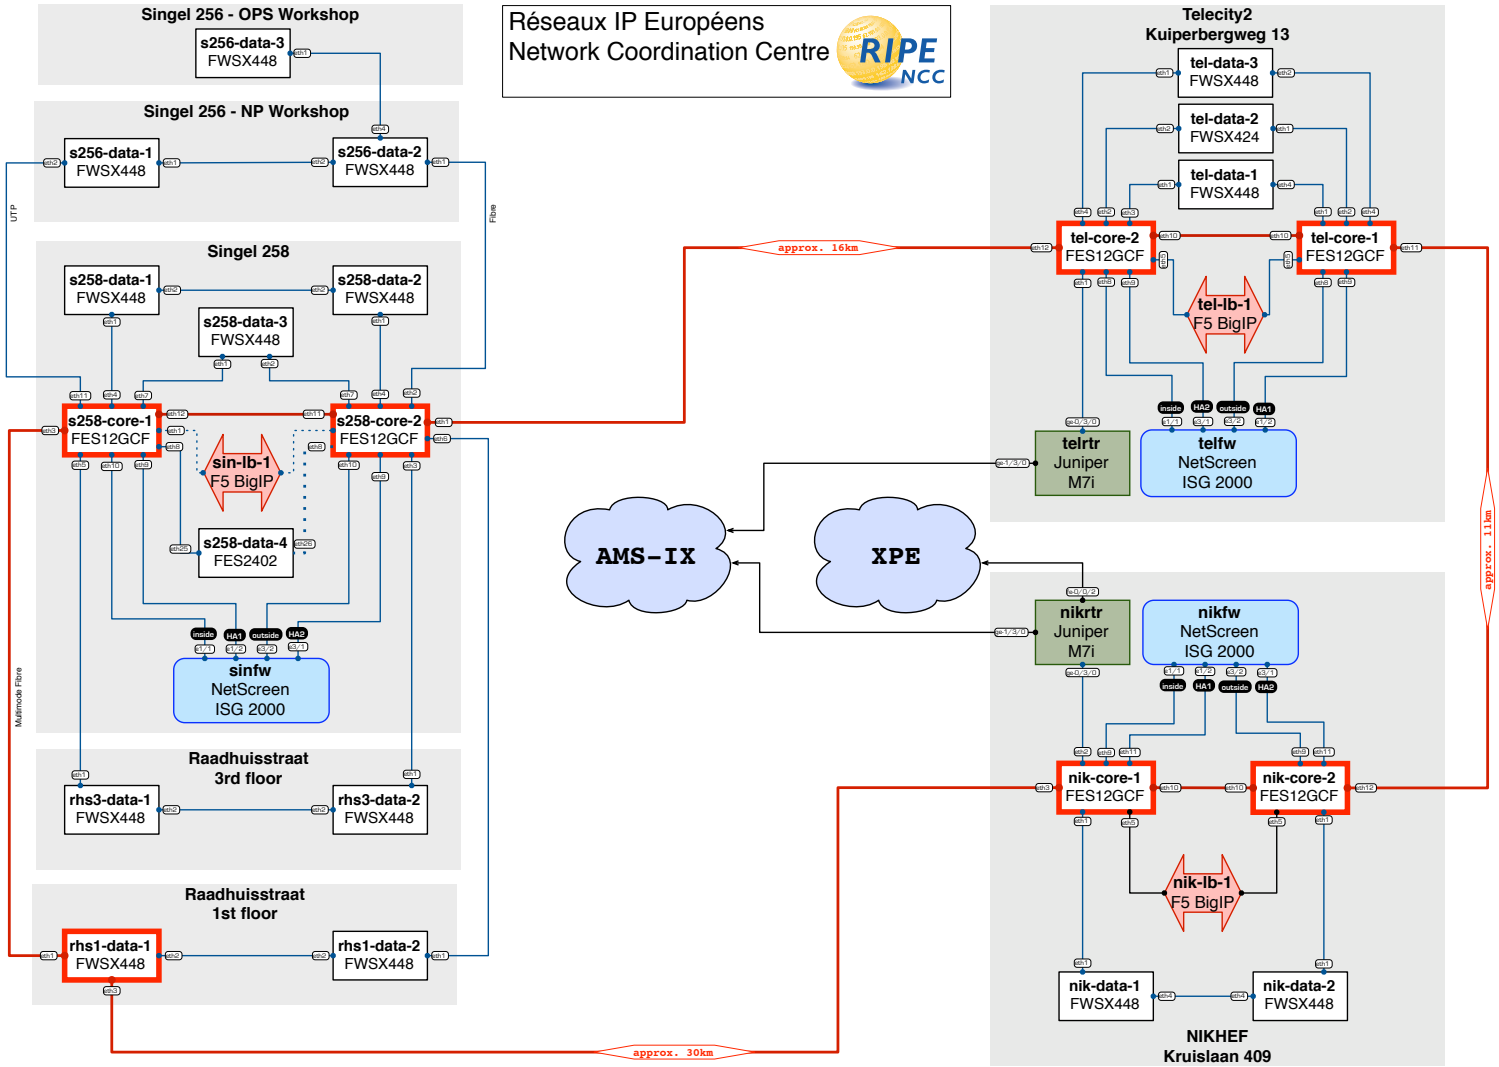

The Network - Prior to 2006

- /42 IPv6 Assignment from SURFNET since 2002
- Layer 2 switches (of different performance) from multiple vendors
- Two overlaid networks:
  - IPv4 using Cisco 7206vxr routers at each corner of the triangle and 100 Mbps connections to the AMS-IX at NIKHEF and Telecity
  - IPv6 using a separate Cisco 3206 at NIKHEF with a dedicated 10 Mbps connection to the AMS-IX
- We needed to upgrade the network to replace old hardware

The Network - Since 2006

- Replaced the layer 2 switching fabric with Foundry switches
- Replaced the Cisco 7206vxr routers at NIKHEF and Telecity by Juniper M7i's
- Introduced a cluster of Juniper Netscreen ISG2000 firewalls
- Moved IPv6 to the M7i's and to use the same (now Gigabit) AMS-IX connections as for IPv4
- We have an open peering policy at the AMS-IX and have about one third the number of IPv6 peers compared to IPv4

The Network - What Happened?

- IPv4 support was robust and reliable
- IPv6 support was good on the Juniper routers but lacked some features (e.g. VRRP) which we were used to having with IPv4
  - Subsequent JunOS releases have fixed these issues
- Initial IPv6 support on the Netscreen firewalls could have been better ...

The Firewalls (1)

- "Full dual-stack support" in ScreenOS 5.4.0
  - This didn't mean reliability but did mean that we could open cases with Juniper for any issues we saw.
- One main problem initially:
  - The firewalls would stop passing any IPv6 traffic and required a reboot to recover.
  - After a month of debugging Juniper came up with a patched version of ScreenOS

The Firewalls (2)

- Everything then went well until there was a firewall failover:
  - IPv6 stopped working until master recovered
  - No NSRP for IPv6 until ScreenOS 6.2… so we installed that…
- Session counters would grow until IPv6 stopped working
  - Could recover by performing a manual failover
  - Caused some sleepless nights for our on-call engineers

The Firewalls (3)

- Reported bug to Juniper and got a patch
  - ScreenOS 6.2.0r1cu3.0
- Now firewall would crash and cause a transparent failover before IPv6 stopped working.
- Better… but still not entirely satisfactory

The Firewalls (4)

- More debugging with Juniper..
- We finally got ScreenOS 6.2.0r1cu4.0 installed last week
- All fine since then…

Load Balancers

- More recently we have deployed hardware load balancers for a numbers of services.
- We looked at products from Foundry, Cisco and F5 before finally settling on a cluster of BigIP 3400s.
- The original software would properly load-balance IPv4 sessions but would only act as a v6-to-v4 proxy for IPv6.
- Recent software updates have allowed full IPv4 and IPv6 load balancing.
- We still do some v6-to-v4 proxying.

Servers

- Most servers at the RIPE NCC run one or other distribution of Linux:
  - Slackware
  - Debian
  - CentOS
- Some other operating systems for particular roles
- Behaviour of these systems with IPv6 varies

Router Advertisements, etc.

- We have experienced various issues with the handling of router advertisements (or the lack thereof) by different operating systems.
- Most systems will accept the link-local address of the router as a default gateway while others need to have the global IPv6 address of the gateway statically configured
- Very much a case of trial and error and depends on OS, kernel, etc.

Router Advertisements, etc.

- For servers we have currently settled on enabling RA's from the routers and firewalls but with the "managed address configuration" bit set
- Workstations get configured using stateless autoconfiguration

Services (1)

- At this point all RIPE NCC services are supported over IPv6
  - Web
    - Straight forward Apache2 installation
    - About 2% of connections come over IPv6
  - Email
    - Initially delayed by use of unsupported home-written software
    - Now using “off the shelf” packages
  - FTP
    - Firewall issues with Extended Passive Mode and IPv6
    - Resolved in the latest ScreenOS releases

Services (2)

- LIR Portal
  - IPv6 proxy on load balancers
- See Erik's presentation for details of:
  - RIPE Database
  - DNS
  - Information Services

IPv6 at RIPE Meetings

IPv6 at RIPE Meetings

- Two Juniper J2320 routers
  - Provide resilient dual-stack network
  - Also a couple of older Cisco routers for other purposes
- IPv6 connectivity depends on the location of each RIPE Meeting
  - Most host organisations can now offer native IPv6 connectivity.
  - Occasionally we still resort to a tunnel back to Amsterdam.

IPv6 at RIPE Meetings

- IPv6 Experiments at RIPE 56 in Berlin
  - In May 2008 we built a couple of IPv6-only networks as an experiment and demonstration of one possible transition mechanism in the event of IPv4 exhaustion: NAT-PT and DNS-ALG

Building an IPv6-Only Network

- Three Options:
  - IPv6-only with no transition mechanisms
    - Only those parts of the Internet which have transitioned to IPv6 are accessible
    - No access to IPv4-only sites
    - Not particularly interesting
  - IPv6-only with NAT-PT and DNS ALG
    - For everything except Windows XP
  - IPv6, local IPv4-based resolver and NAT-PT and DNS ALG
    - Just for Windows XP
- For RIPE 55 we built the last two

Transition Methods

- NAT-PT
  - Network Address Translation - Protocol Translation
  - RFC2766
  - We used IOS 12.4(15)T5 Advanced IP Services (the release of the week) but “IOS 12.4(15)T3 or later should also work”.
- DNS ALG
  - DNS Application Layer Gateway
  - DNS Proxy synthesizes **AAAA** records for those DNS entries which have only **A** records
  - We used “`totd`” under FreeBSD

DNS ALG

- A DNS **A** record is of no use on a pure IPv6 network, so what do we do if we receive only an **A** record in response to a query?
- Local DNS proxy (`totd`) has a hack: takes the IPv4 address returned in the **A** record, embeds it within a particular IPv6 prefix and returns a synthesized **AAAA** record.
- NAT-PT knows the prefix and strips it back to IPv4 when a packet leaves the pure IPv6 network destined for this IPv6 address

DNS ALG

- Doesn't work if an application forces an IPv4 or IPv6 transport
  - ping/ping6
  - traceroute/traceroute6
- The usual NAT problems...
  - protocols which embed IP addresses
  - need some additional form of proxy for these

DNS ALG / NAT-PT

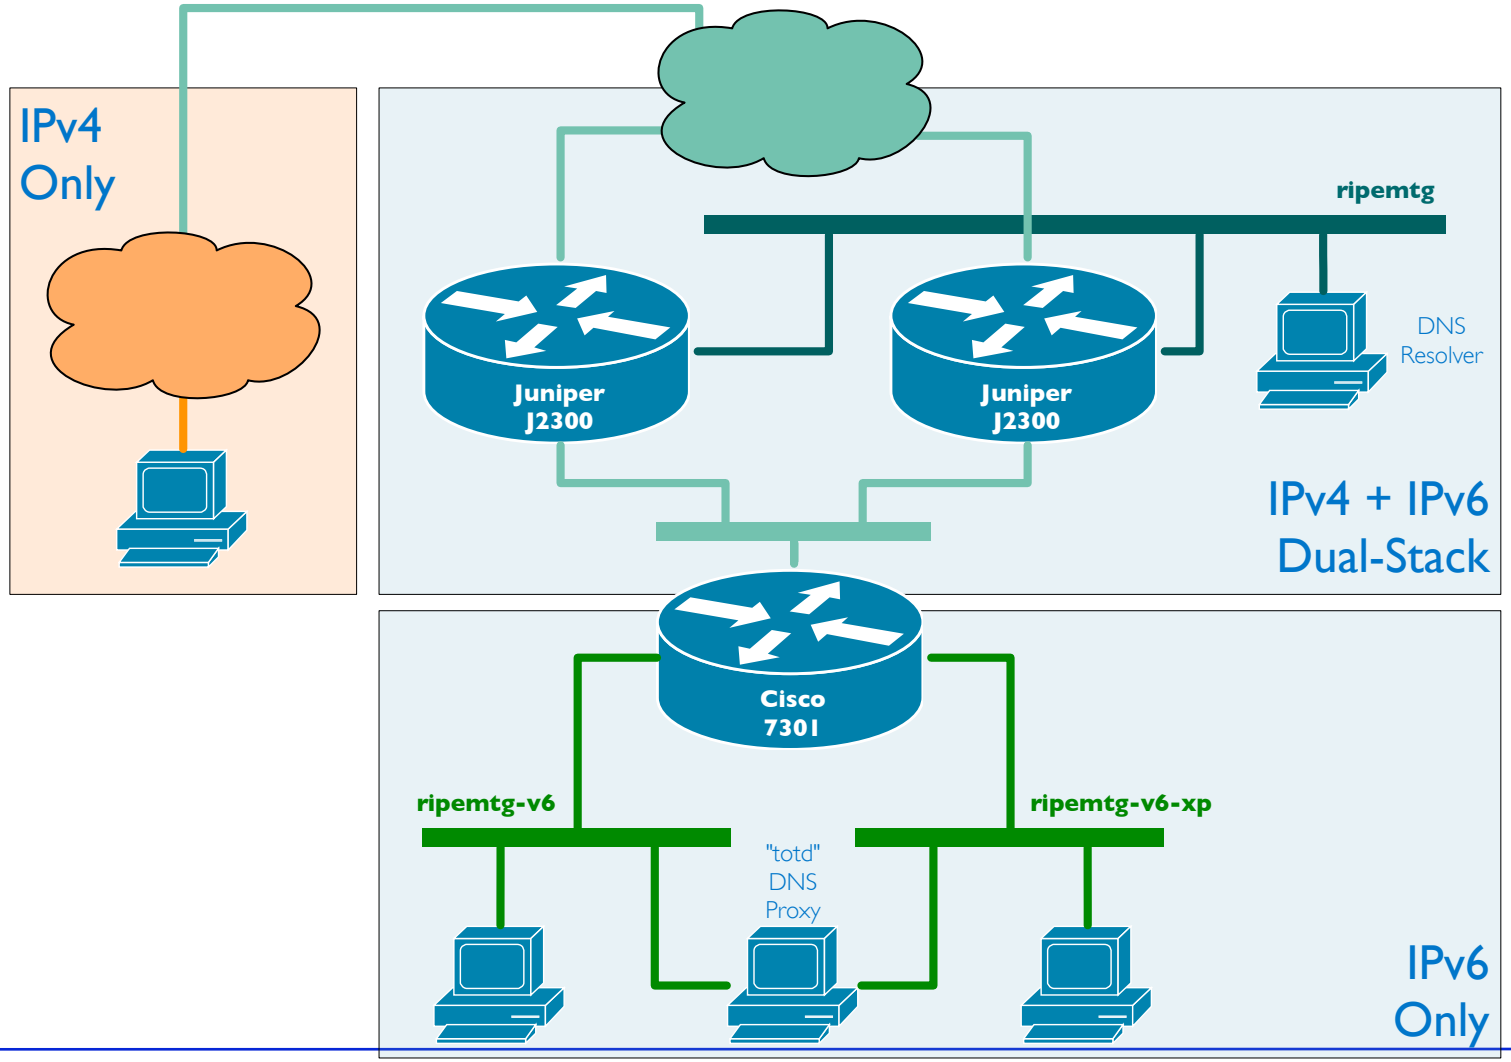

DNS ALG / NAT-PT

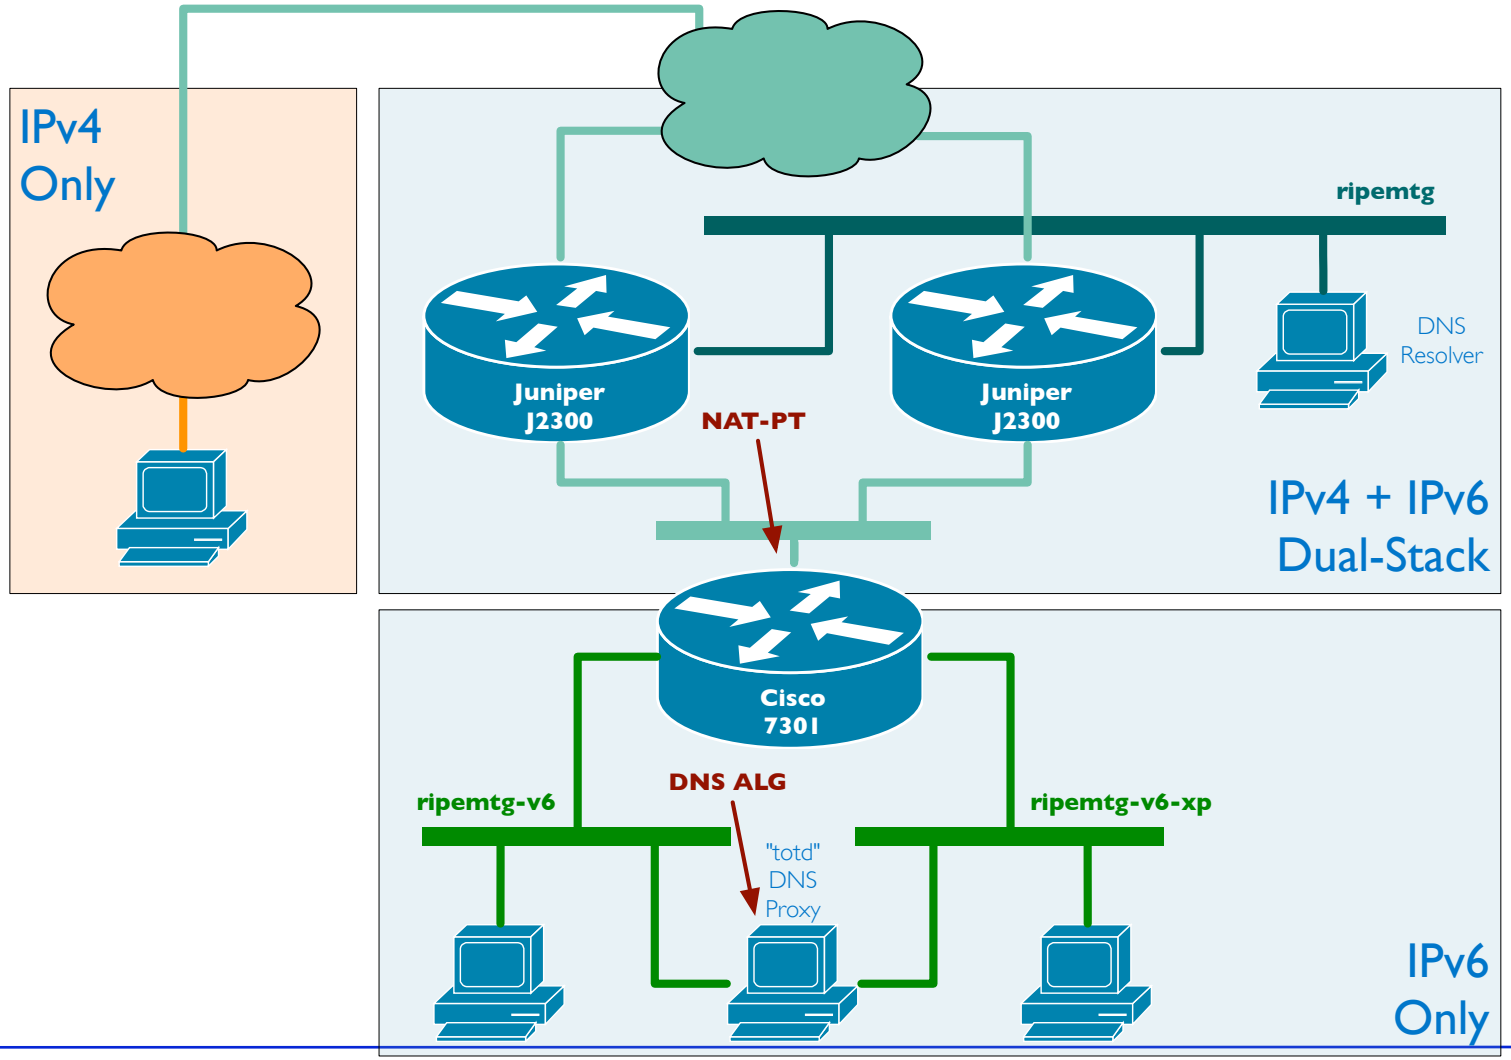

DNS ALG / NAT-PT

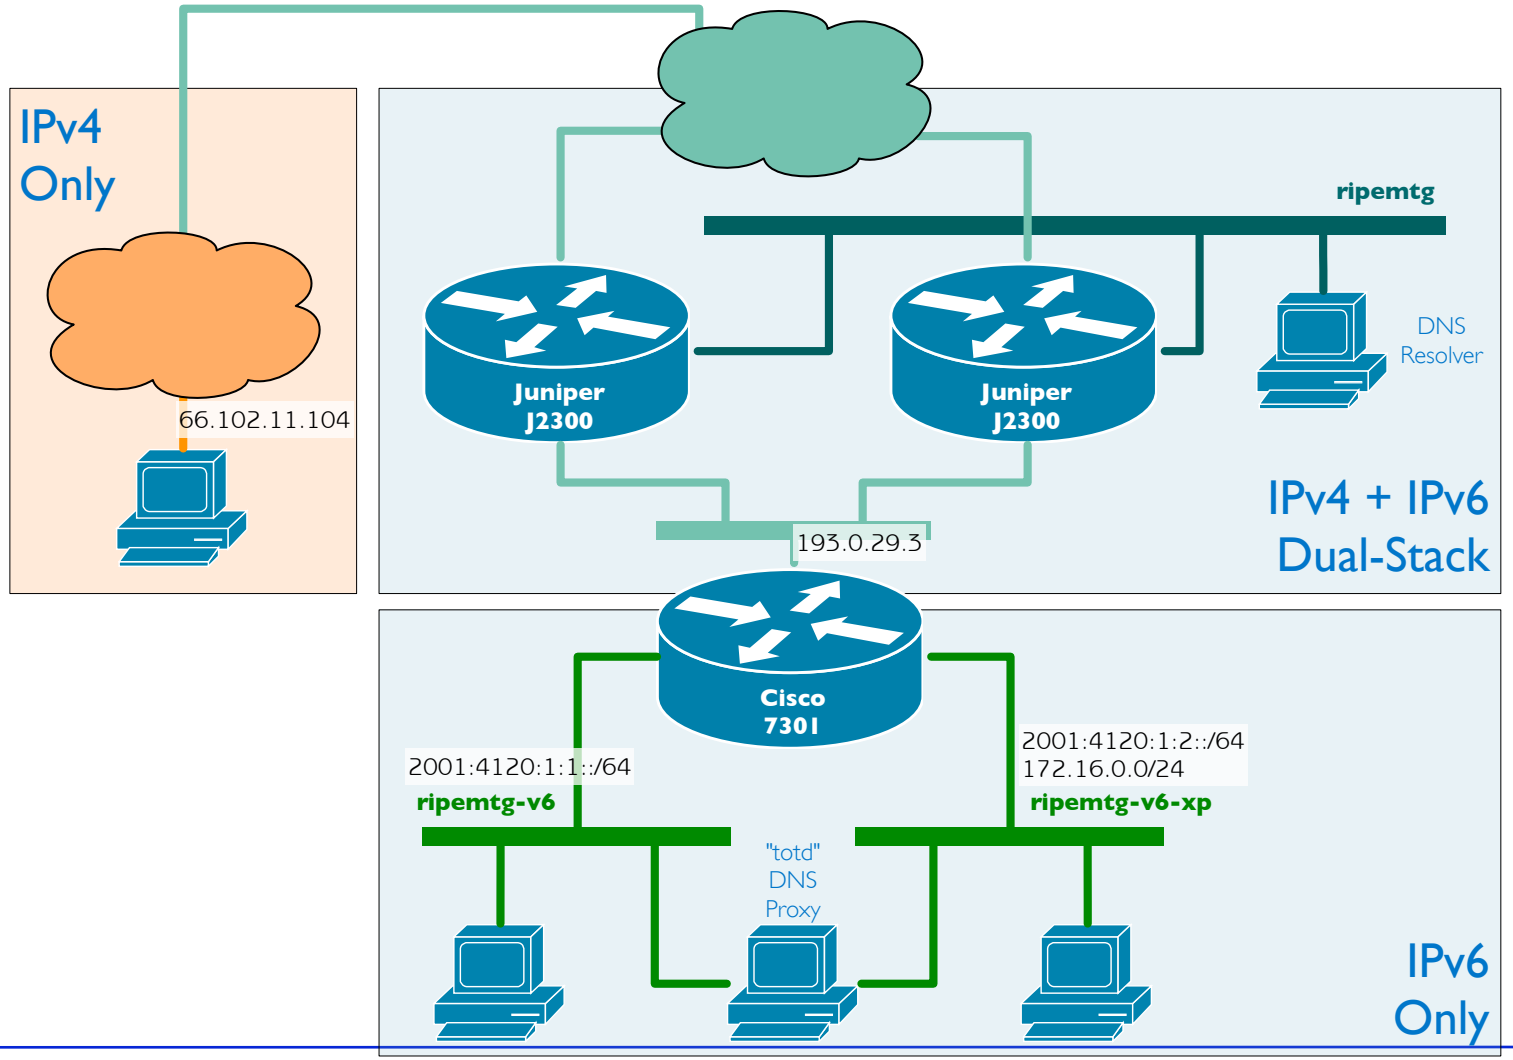

DNS ALG / NAT-PT

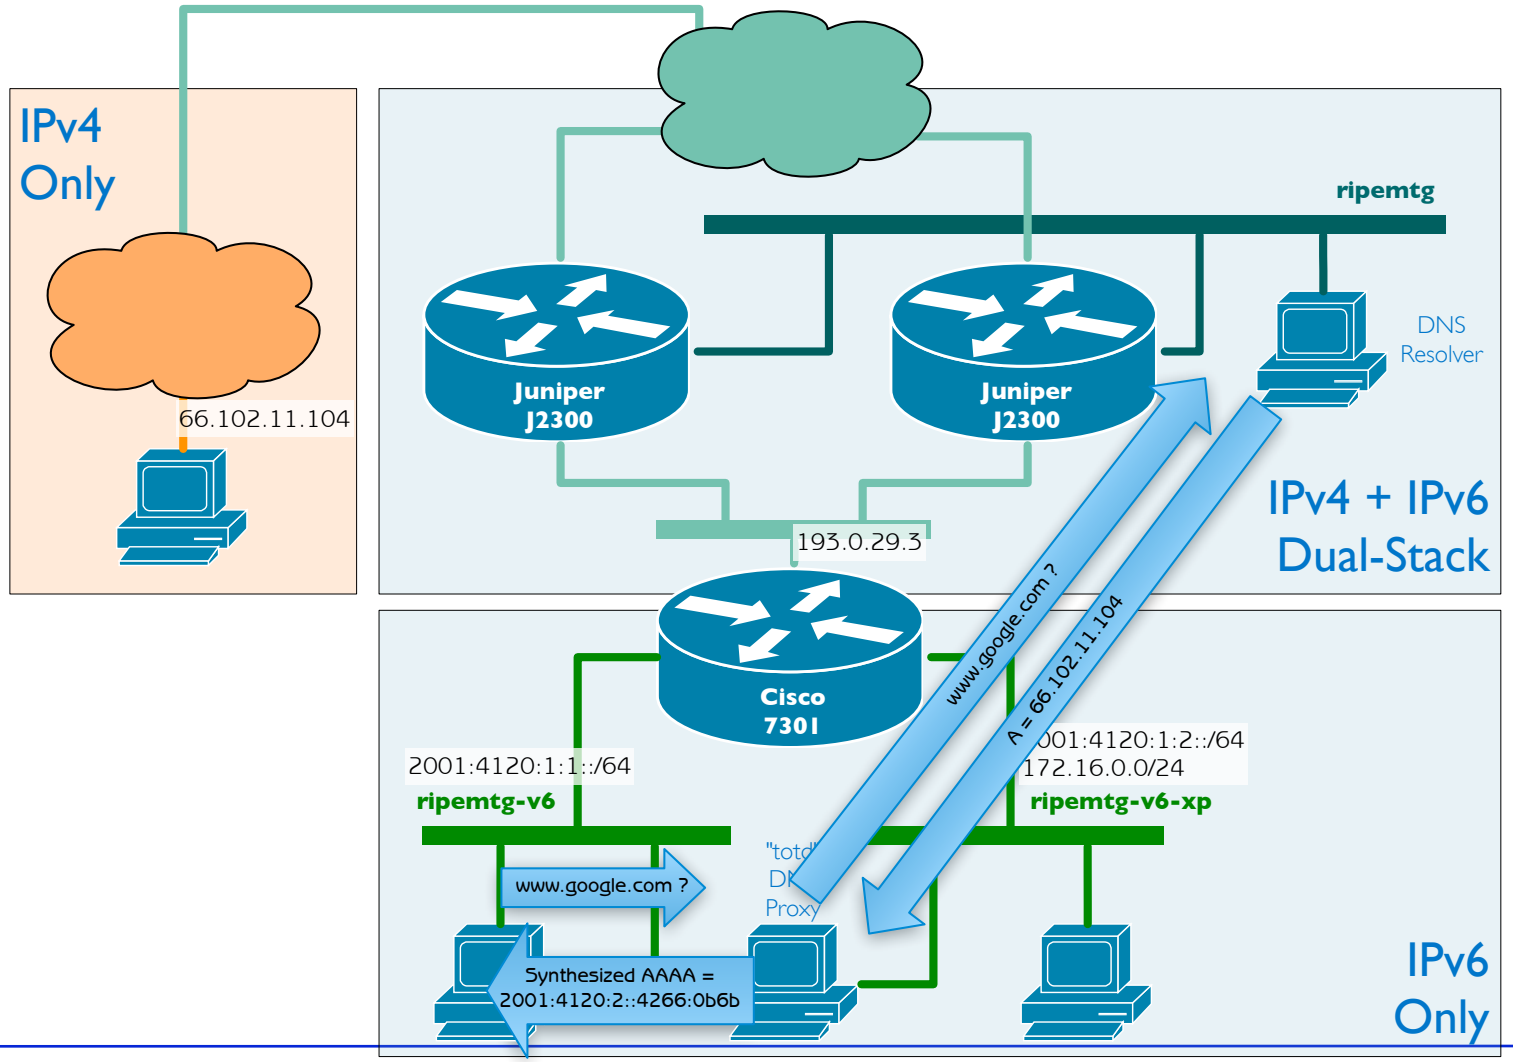

DNS ALG / NAT-PT

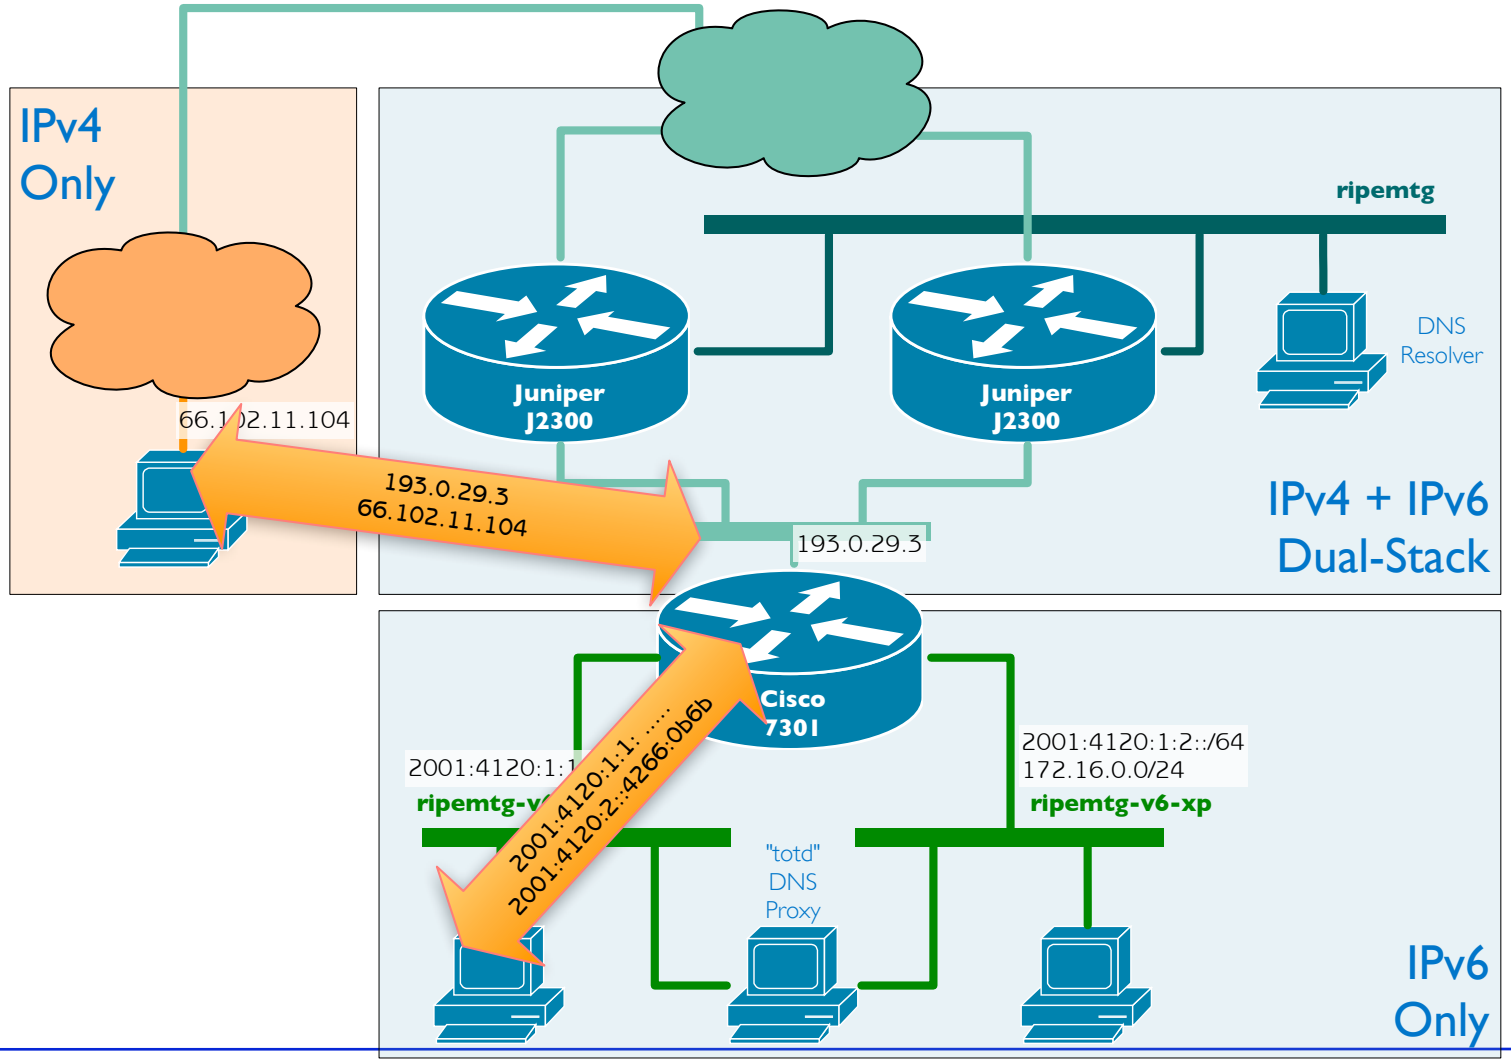

Cisco NAT-PT Configuration

- On each interface:

```
ipv6 nat
```

- NAT-PT configuration:

```
ipv6 nat v6v4 source list NATPT interface Loopback0 overload
ipv6 nat prefix 2001:4120:2::/96 v4-mapped NATPT
ipv6 access-list NATPT
  permit ipv6 2001:4120:1:1::/64 2001:4120:2::/96
  permit ipv6 2001:4120:1:2::/64 2001:4120:2::/96
```

Cisco NAT-PT Configuration

- This configuration maps all traffic to a single IPv4 address.
  - Problems with growth in the size of the mapping table
- An alternative is to map to a (small) range of addresses
  - The Cisco documentation for this wasn't too clear
  - Time constraints during setup
- With approximately 100 users (4Mbps) the CPU load on the 7301 rose to about 10%
- NAT-PT RFC2766 has since been marked as "historic" but with no replacement yet.

Questions?

Over to Erik ...