



IPv6 at RIPE NCC


Erik Romijn
<eromijn@ripe.net>



I like statistics



People at this meeting who have seen this presentation before



Has seen before




Has not seen before



**At the RIPE NCC
IPv6
is very important**







RIPE NCC services without support for IPv6



RIPE NCC services without IPv6

THIS SLIDE IS INTENTIONALLY LEFT BLANK




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)
 - (Krasnapolsky hotel)
- AMS-IX connection at NIKHEF and Telecity
- Switched layer 2 network carrying multiple VLANs



The Network - Background





RIPE NCC network

- /42 IPv6 Assignment from SURFnet since 2002
- L2 switching with Foundry switches
- Routing on Juniper M7i's
- Cluster of Juniper Netscreen ISG2000 firewalls
- IPv6 to the M7i's and using the same (Gigabit) AMS-IX connections as for IPv4
- 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 (2)

- 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 (3)

- More debugging with Juniper..
- We finally got ScreenOS 6.2.0r1cu4.0 installed a few months ago
- All fine since then...



Load Balancers

- Load balancers for some services
- Running on a cluster of F5 BigIP 3400s
- Could originally only do 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 run Linux:
 - Slackware
 - Debian
 - CentOS
- Some other operating systems for particular roles
- Your mileage with these systems and IPv6 will vary



Router Advertisements, etc.

- Various issues with handling of RA's (or the lack thereof) by different OSes
- Most systems will accept the link-local address of the router as a default gateway
- 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:
 - RAs enabled from routers/firewalls
 - “Managed address configuration” bit set
- Workstations get configured using stateless autoconfiguration



IPv6 at RIPE meetings



IPv6 at RIPE Meetings

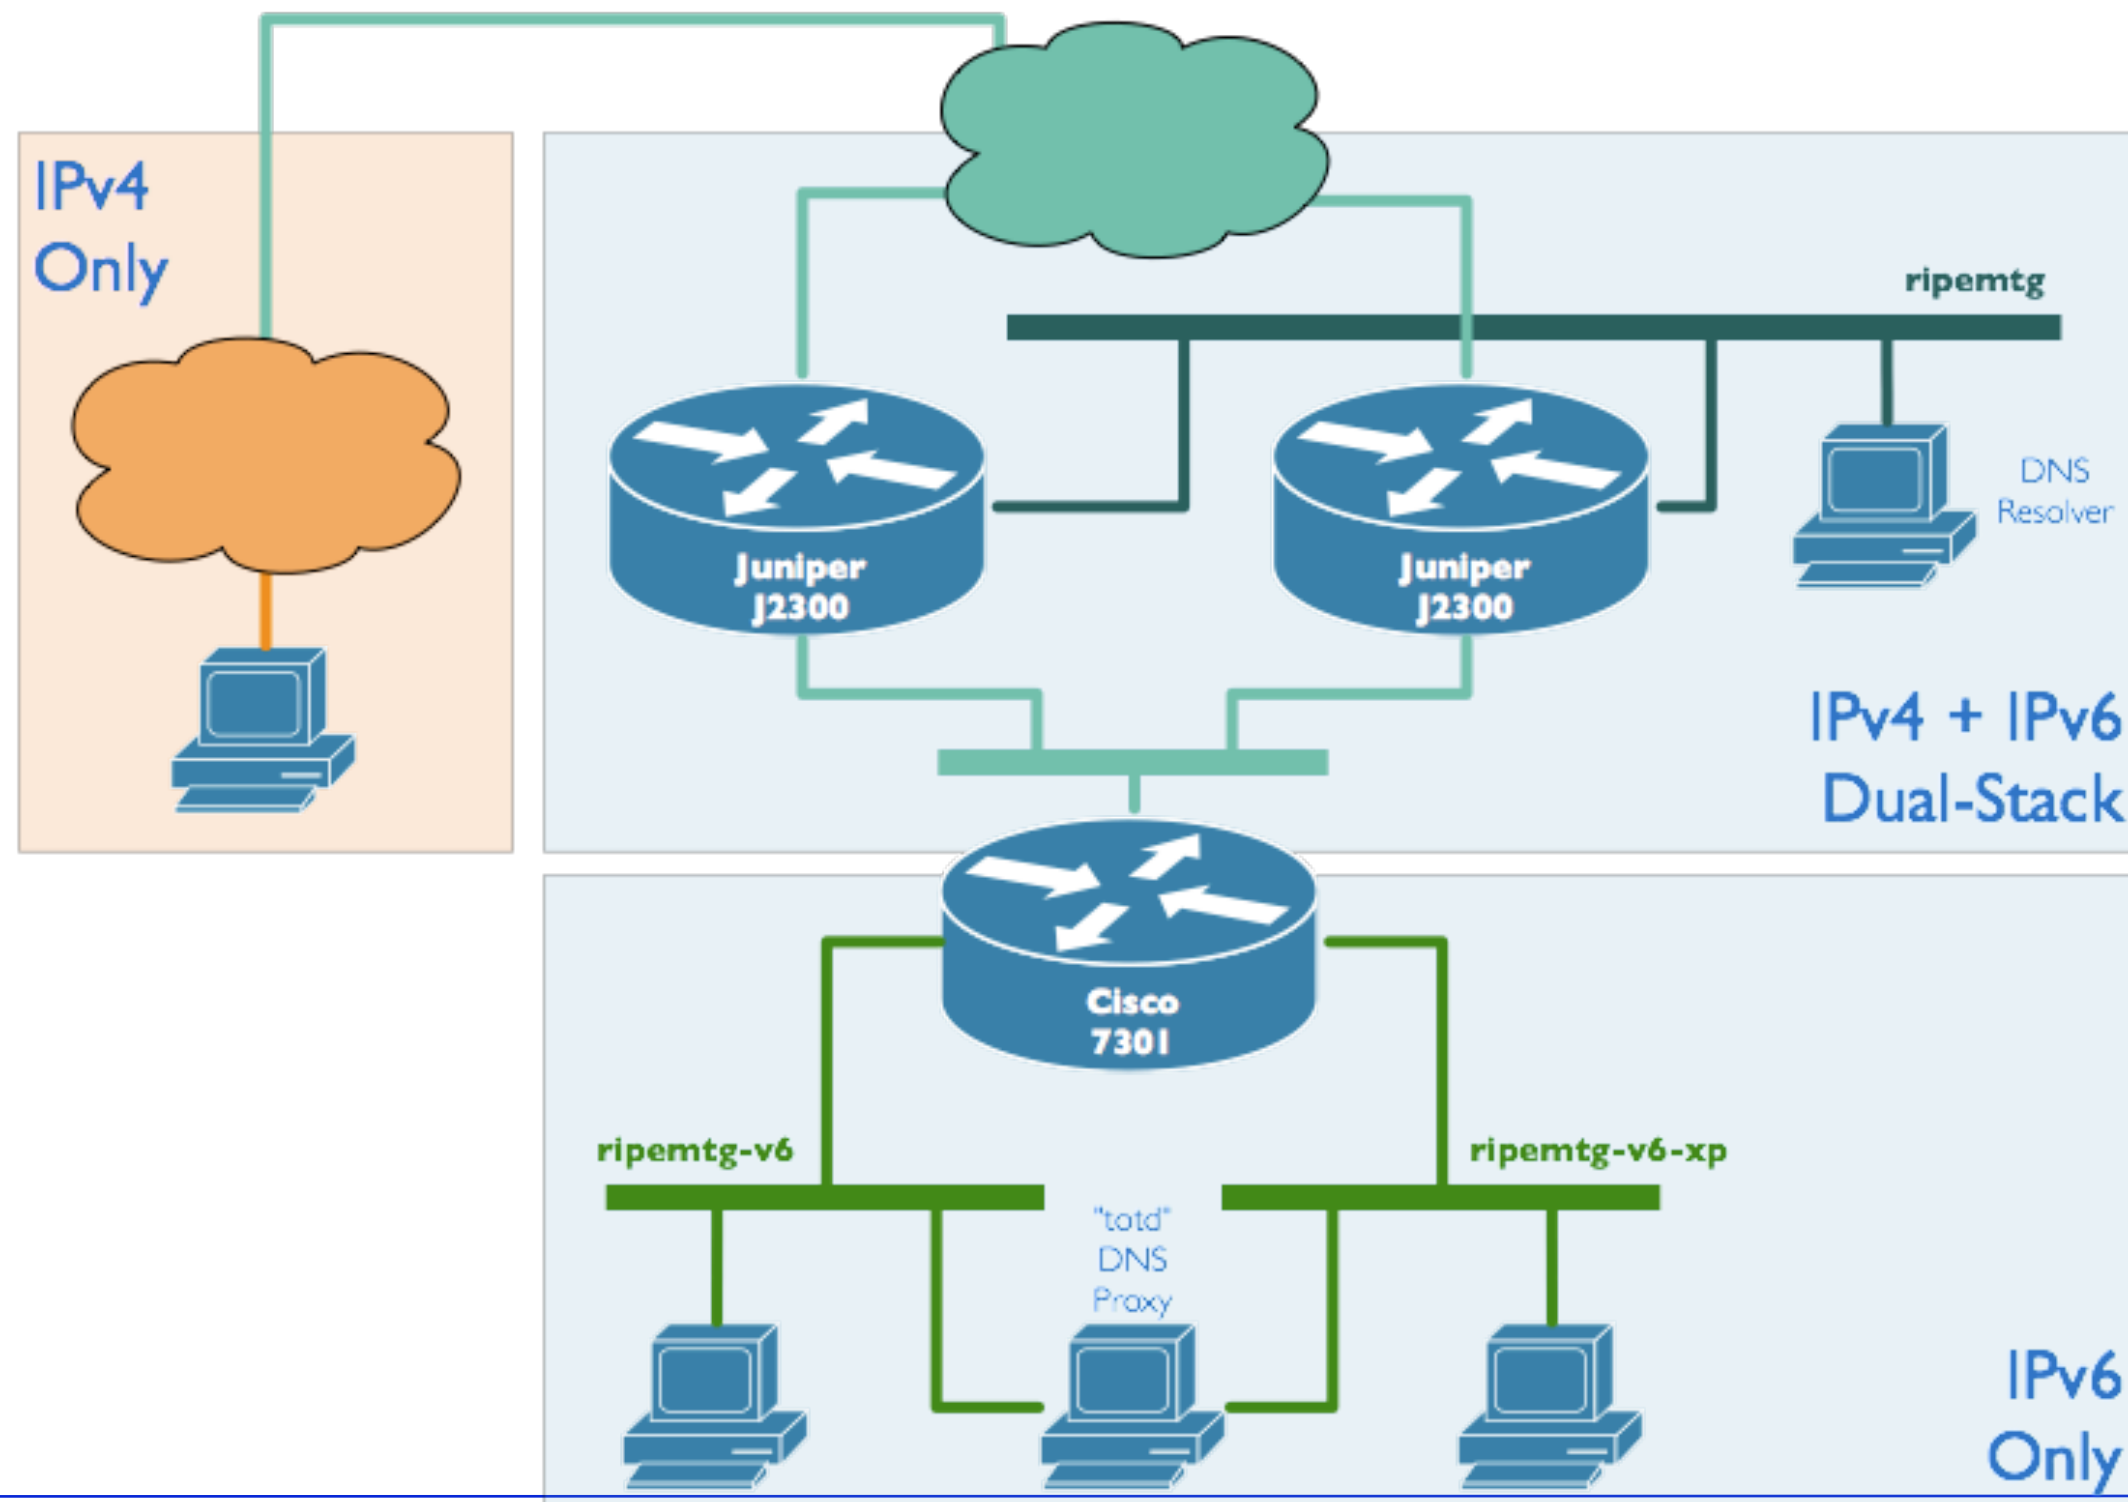
- Two Juniper J2320 routers
 - Provide resilient dual-stack network
 - Also a couple of older Cisco routers for other purposes
- IPv6 connectivity depends on the location of each RIPE Meeting
 - Most host organisations can now offer native IPv6 connectivity.
 - Occasionally we still resort to a tunnel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




IPv6 at RIPE Meetings

- IPv6 Experiments at RIPE 56 in Berlin
 - In May 2008
 - Two IPv6-only networks
 - Experiment and demonstration of NAT-PT and DNS-ALG
- One hour where wireless had no IPv4

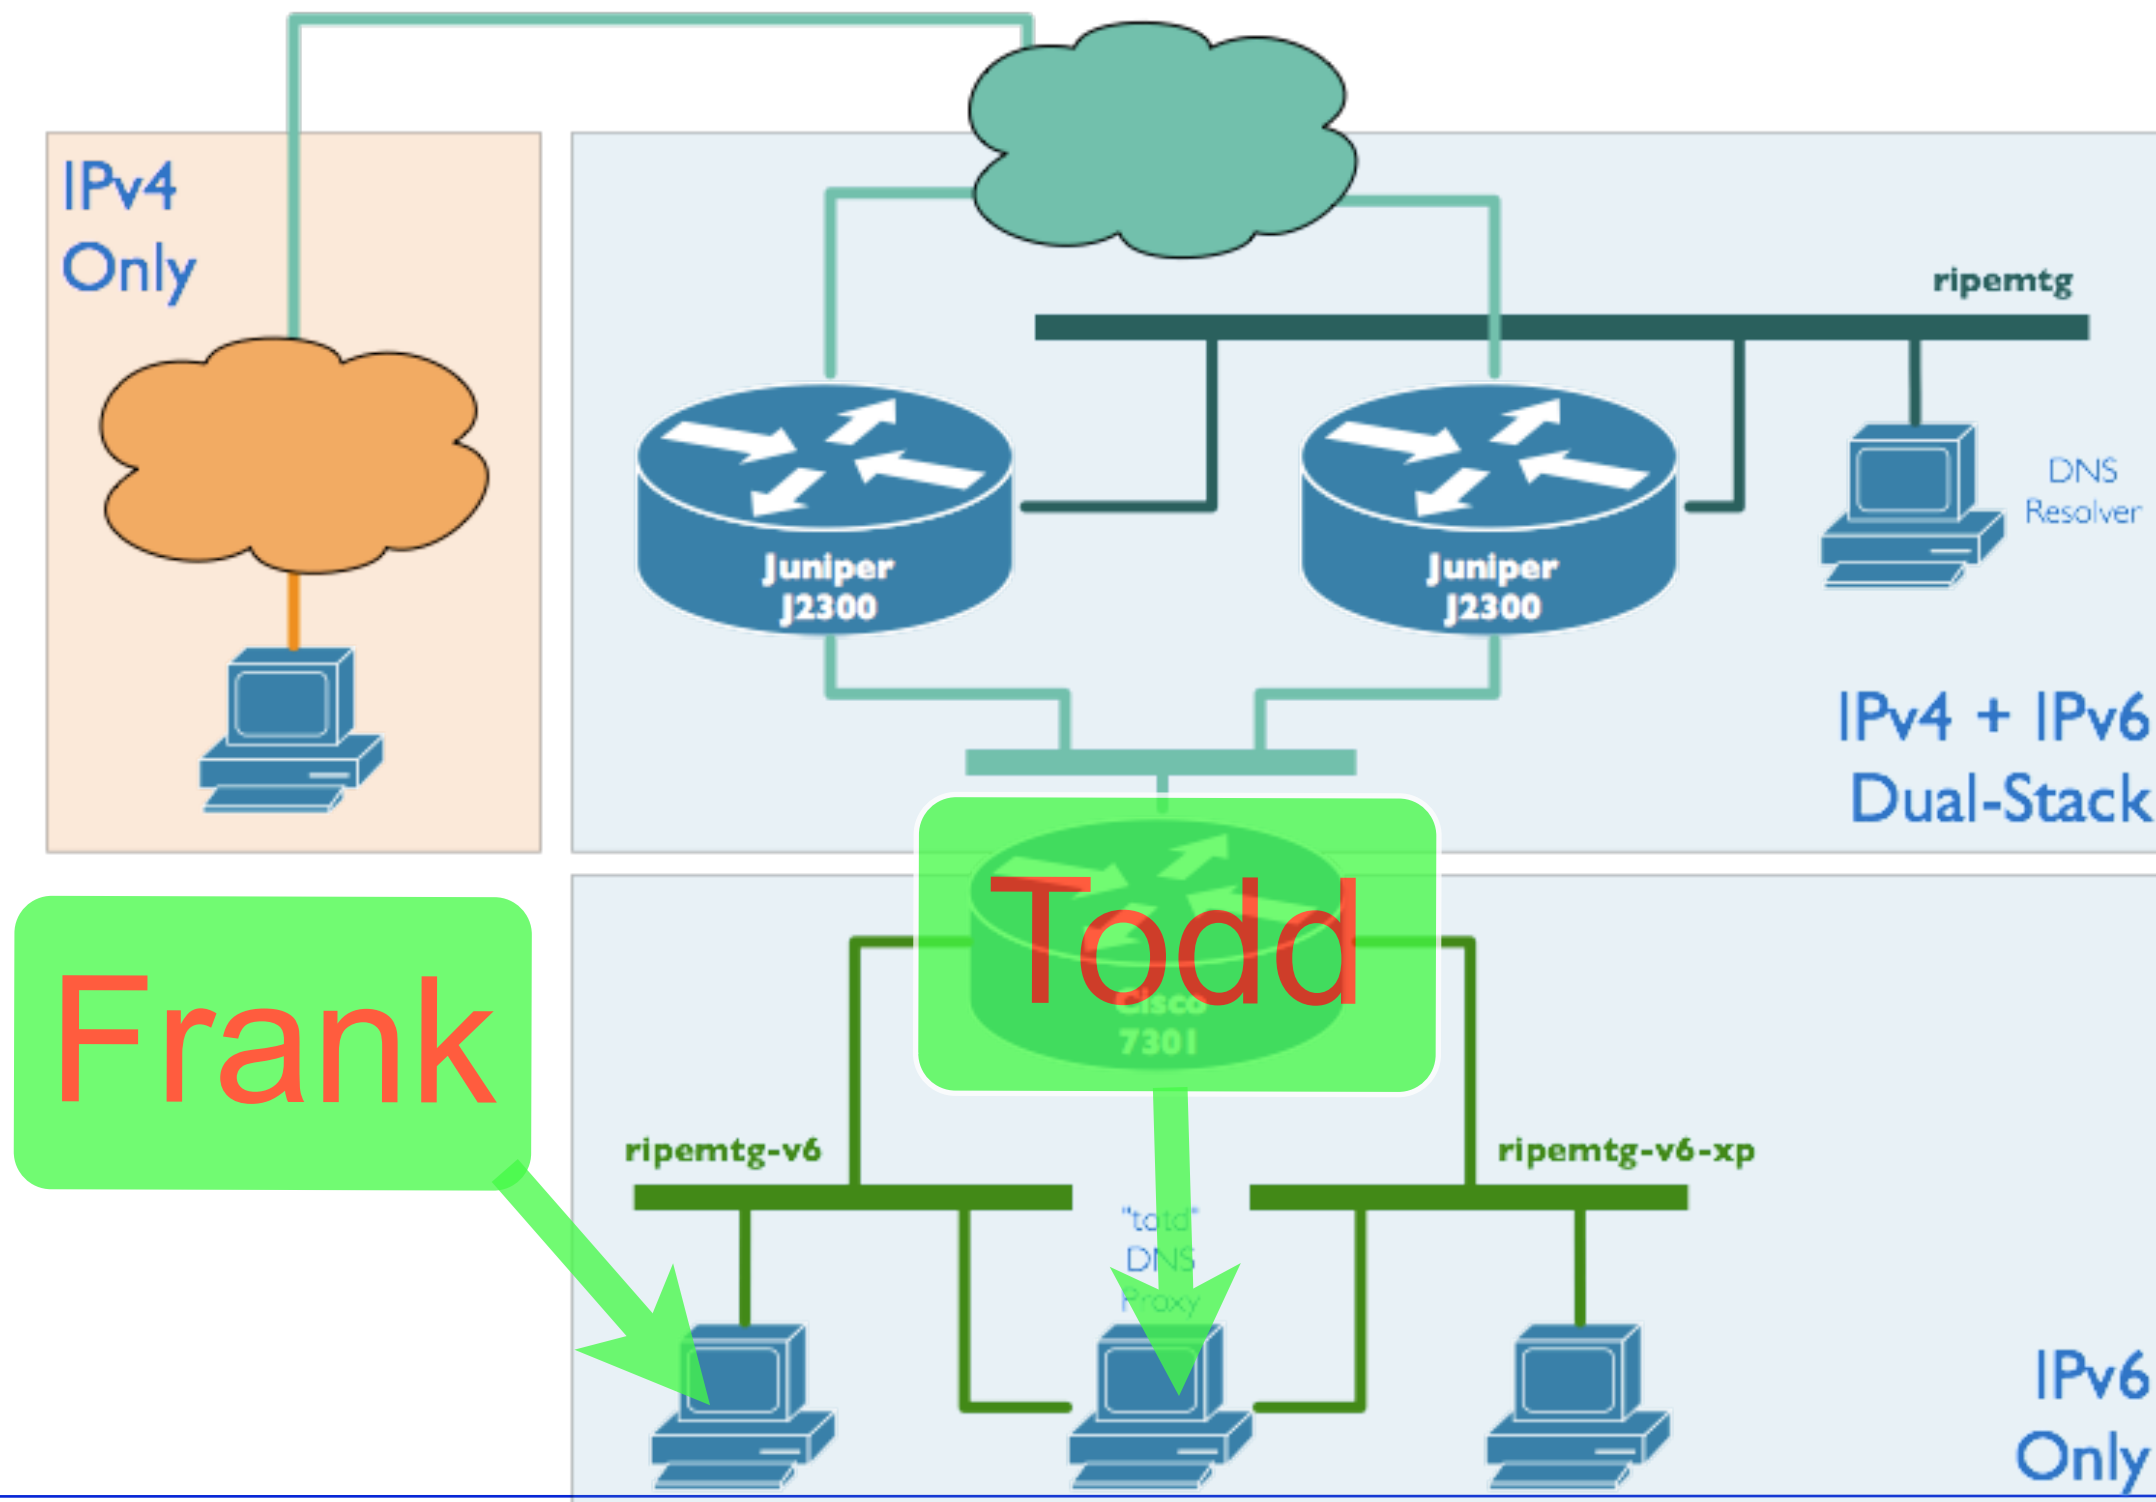
DNS ALG / NAT-PT




DNS ALG / NAT-PT




DNS ALG / NAT-PT



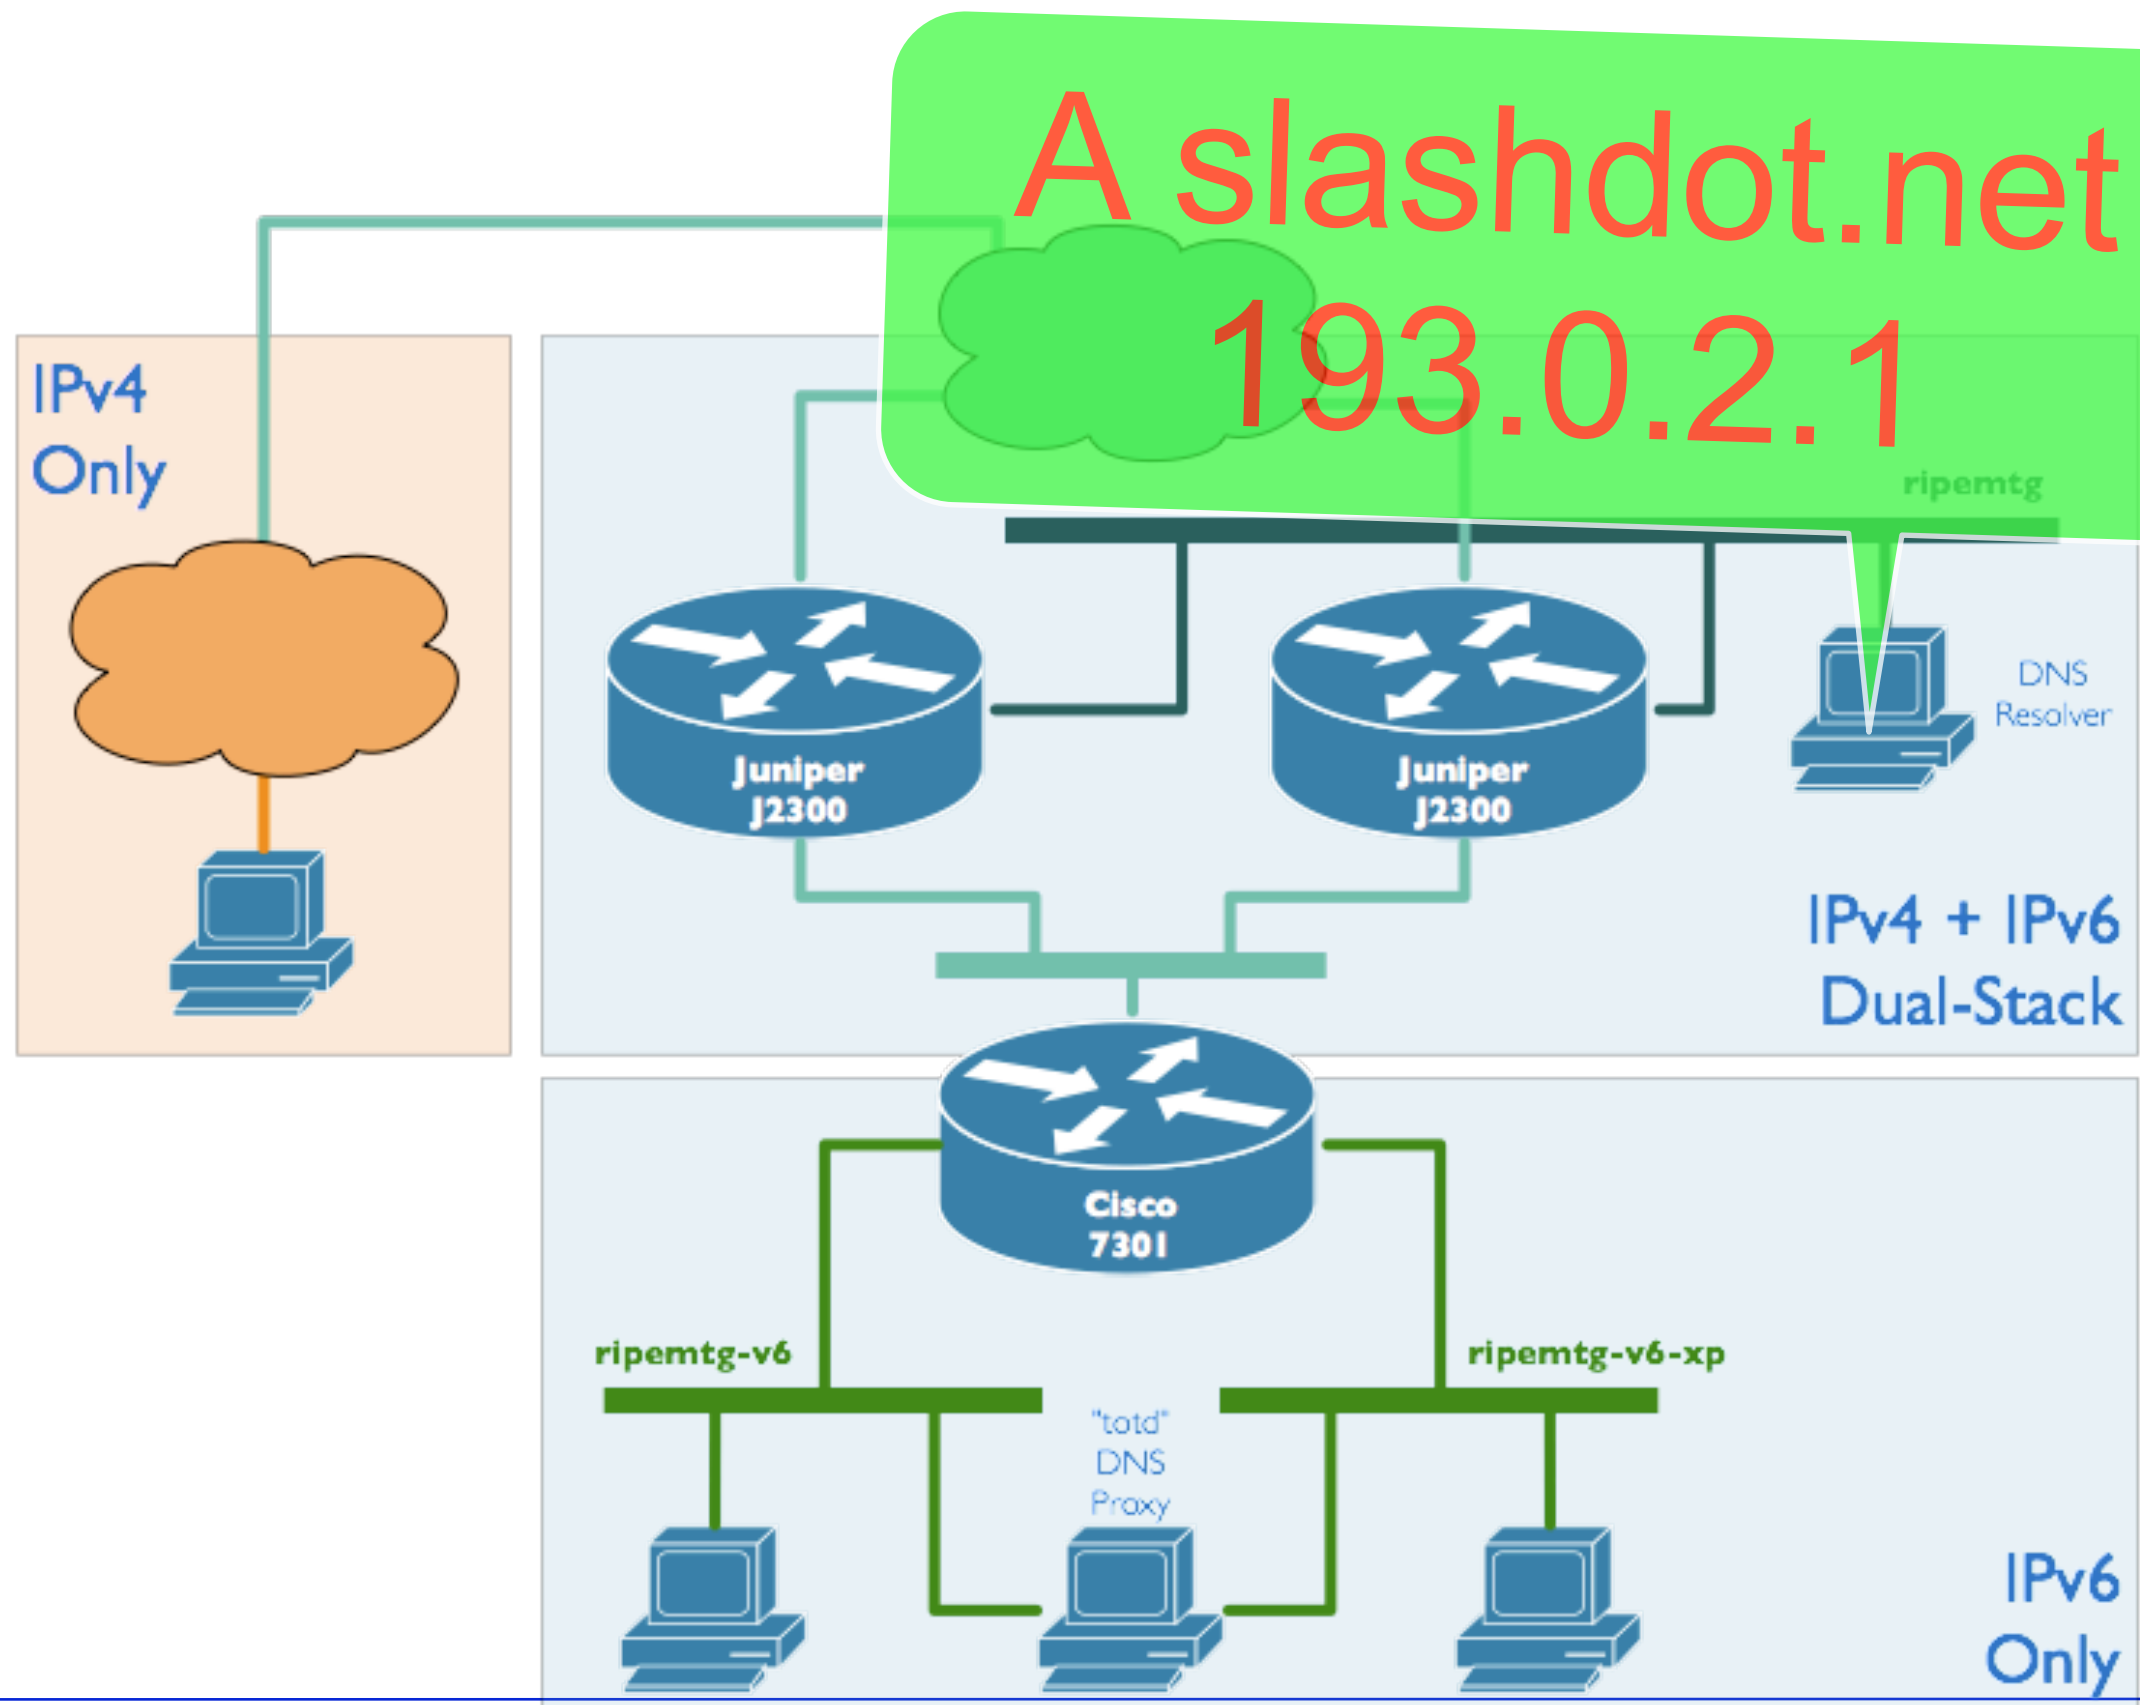
DNS ALG / NAT-PT



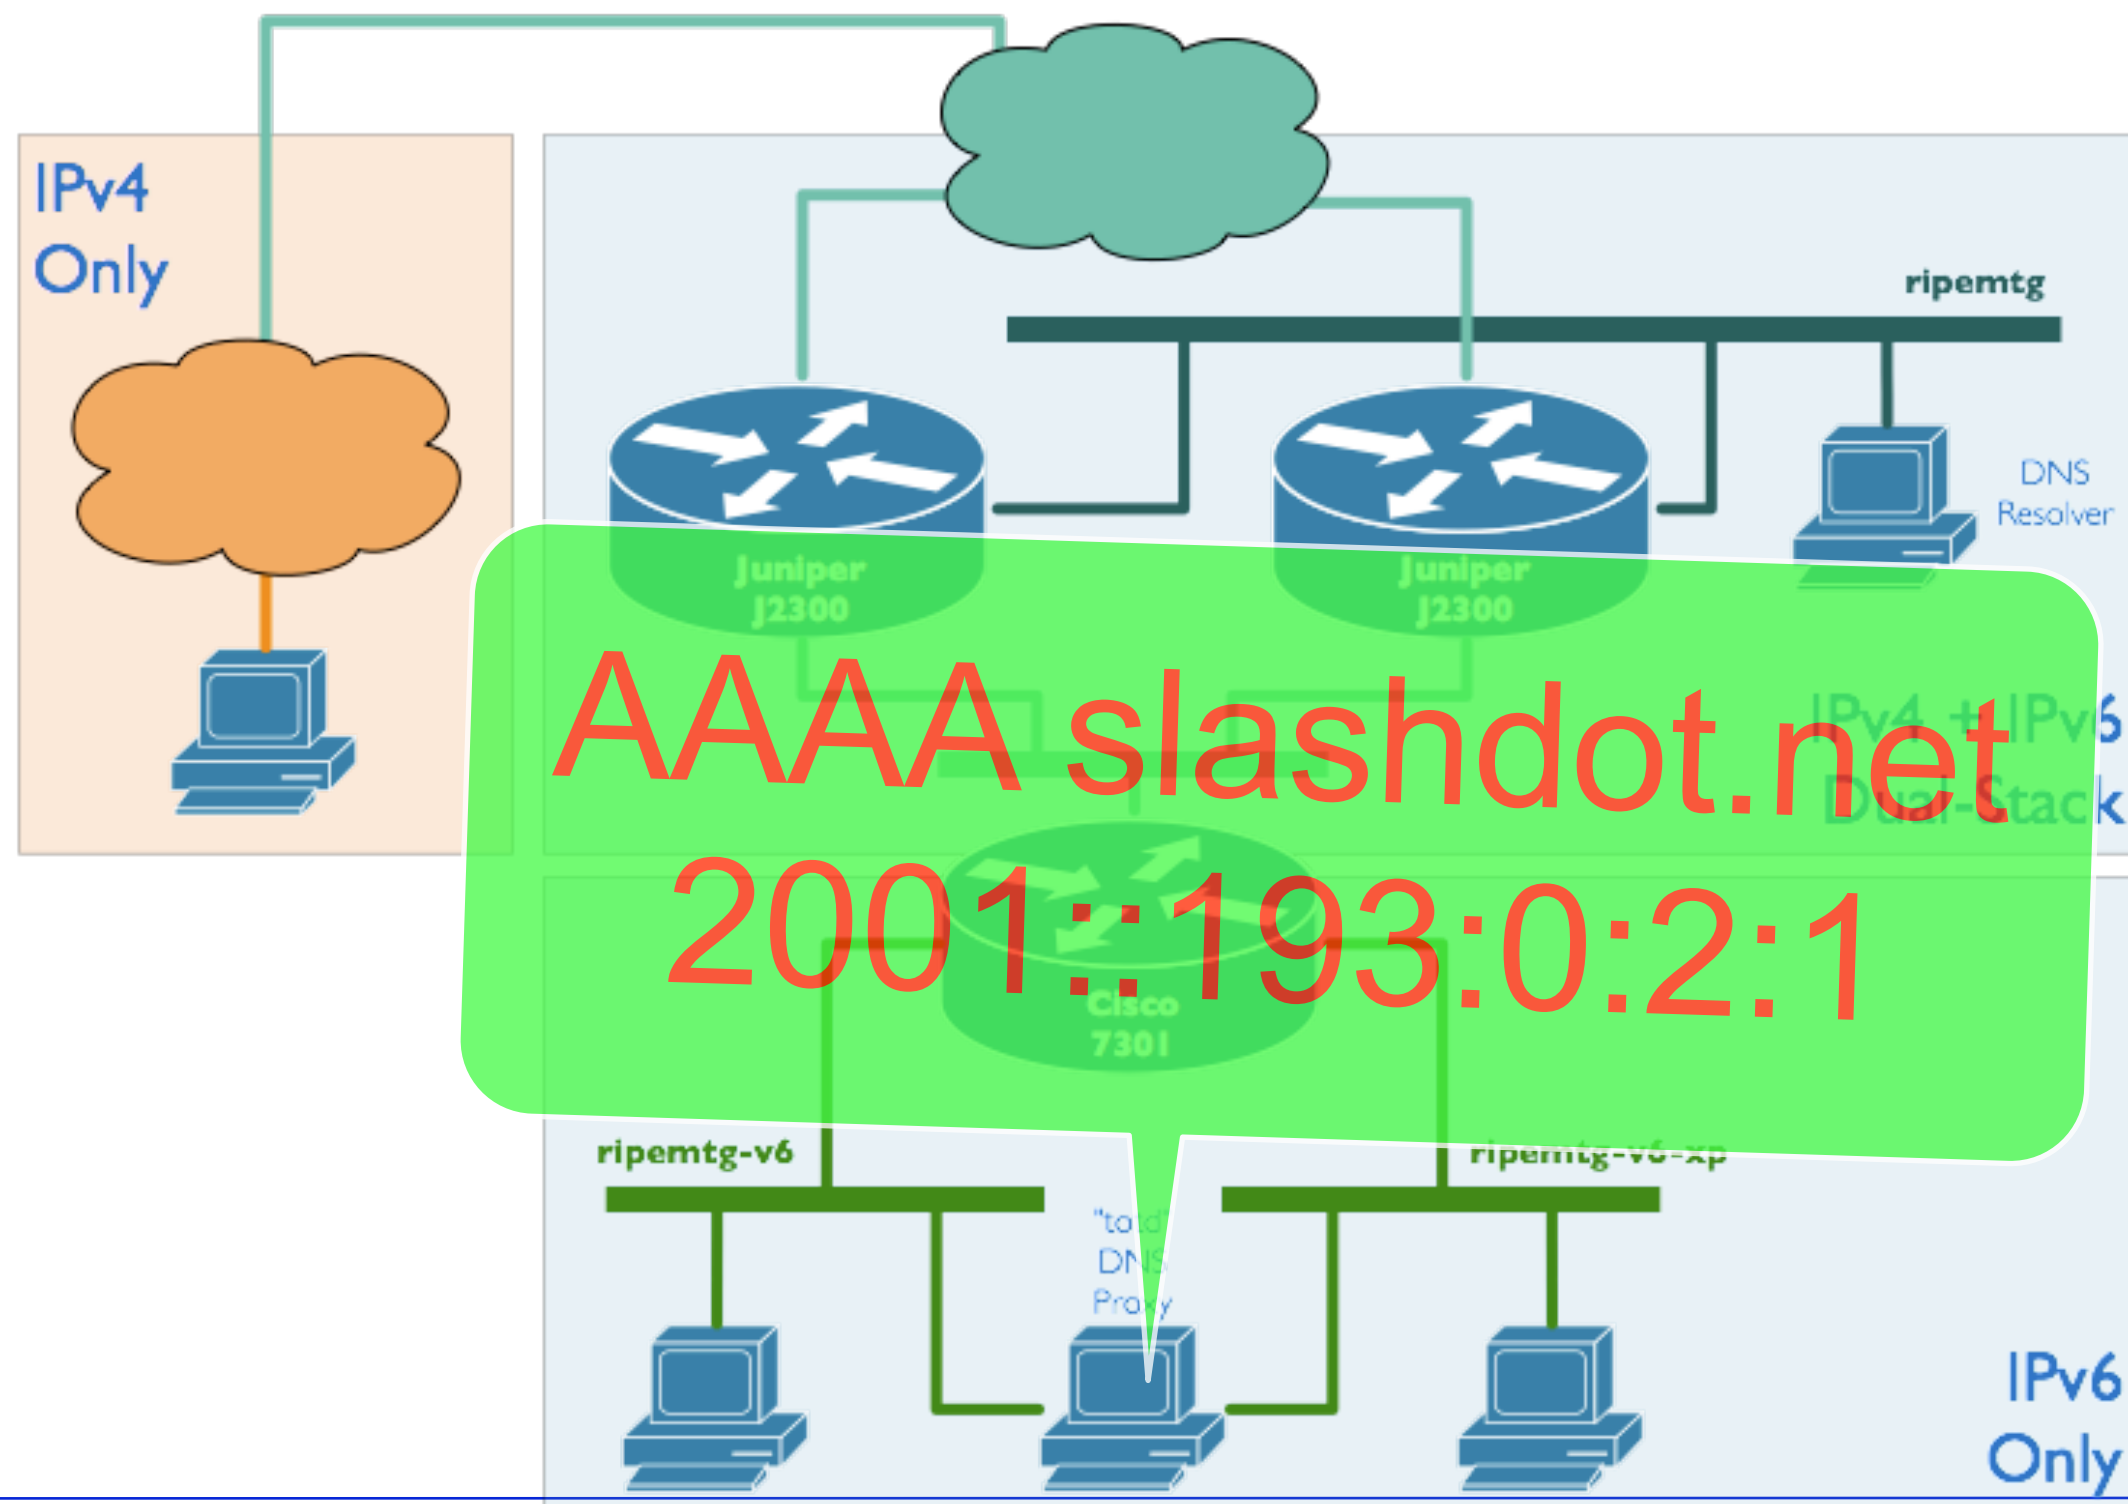
DNS ALG / NAT-PT




DNS ALG / NAT-PT



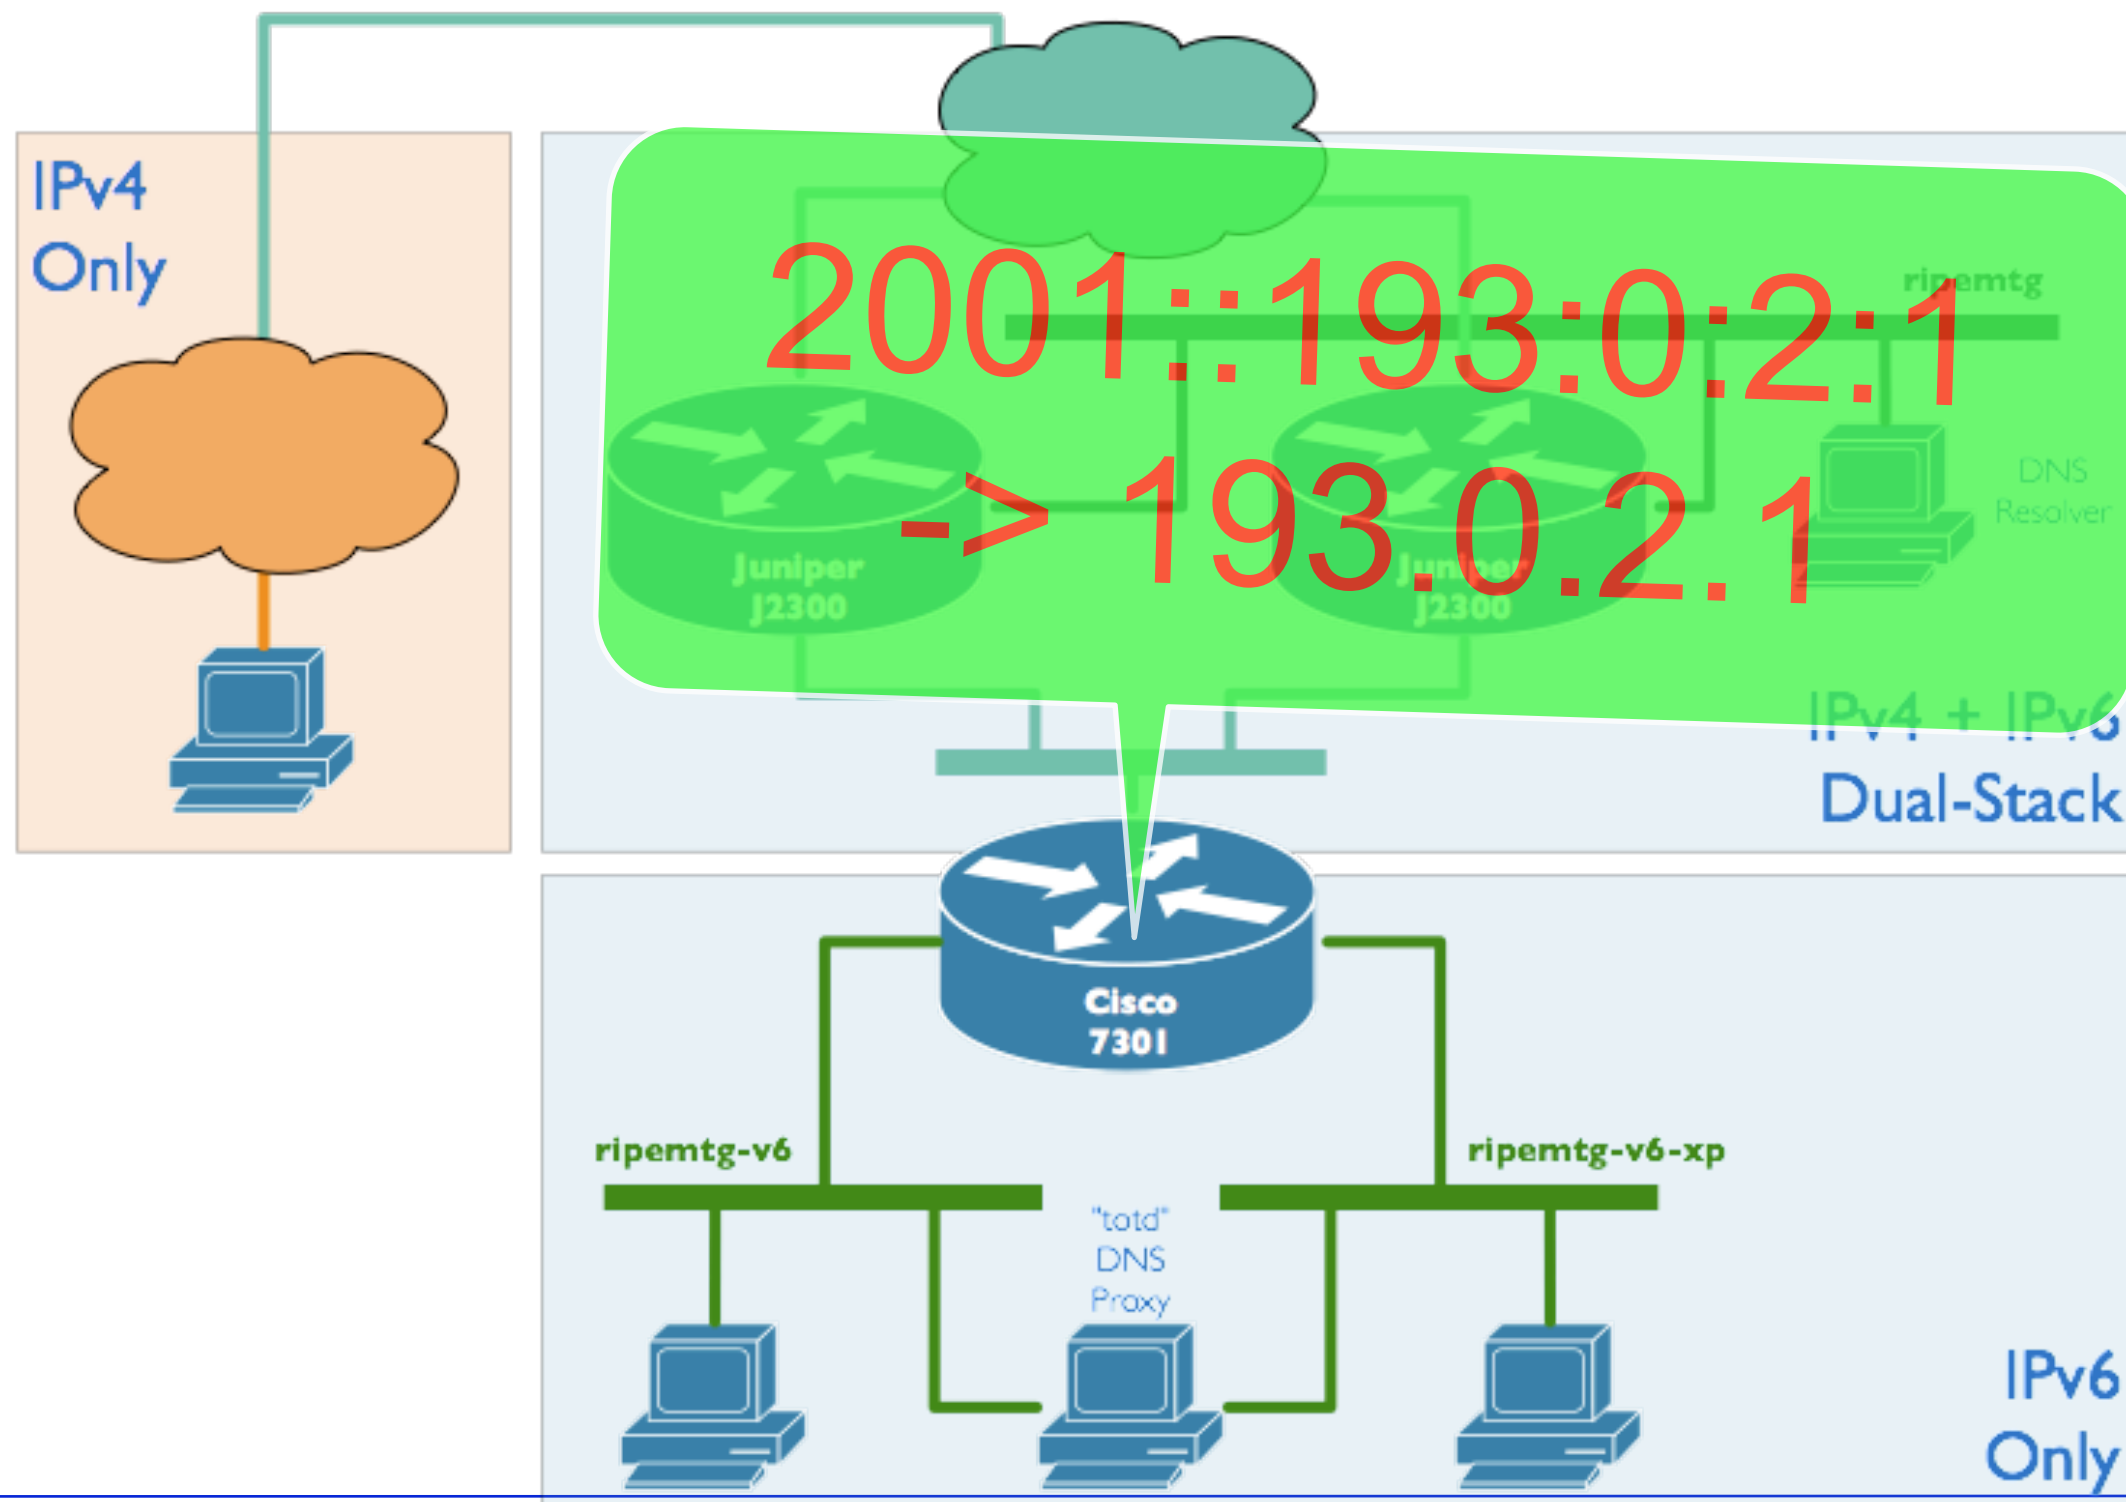
DNS ALG / NAT-PT




DNS ALG / NAT-PT



DNS ALG / NAT-PT



DNS ALG / NAT-PT





NAT-PT and DNS-ALG experiences

- Setup took:
 - 20 hours of hard work
 - Two experienced network engineers
 - One Cisco engineer with access to the developers of the implementation
 - Careful match of settings and software version
- But it worked!
- With about 4Mbps, cpu load on C7301 went to 10%
- More in the tech team presentation from RIPE56



Various services



Services (1)

- Basic services on 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 (1)

- LIR Portal

- IPv6 proxy on load balancers

- RIPE database

- Native IPv6
- Average of ~7500 queries per minute
- ~21 queries per minute over IPv6: about 0.27 %



DNS



DNS services


- K-root server:
 - AAAA for k-root in root zone since Feb 2008
 - 1% of queries is done over IPv6
 - 25% of queries is for AAAA
 - 8 anycast nodes with IPv6
- Reverse delegation:
 - Authoritative for /8s administered by RIPE NCC
 - Delegates down to holders of IP space
 - Also 1% of queries over IPv6



DNS services – K-root IPv6 queries


IP Version Carrying DNS Queries

From Feb 04, 2008, 11:00:26 To Feb 05, 2008, 11:00:26 CET





DNS services – K-root





Experiences from IPv6 DNS

- “It just works.”
- No problems observed, no strange bugs seen
- Biggest challenge: get IPv6 transit for nodes
- Continuously monitored by DNSMON



Hostcount++




Hostcount++

- Counts hosts in the RIPE region
- Walks through (the relevant part of) the DNS tree
- Sources:
 - Zone transfers for IPv4 and IPv6
 - Reverse tree for IPv4
(Reverse IPv6 tree is too large.)
- No zone transfer means no IPv6 count!
- Help us count IPv6:
 - Allow 193.0.0.0/22 for zone transfer




Hostcount++ – RIPE region data






Hostcount++ – RIPE region data

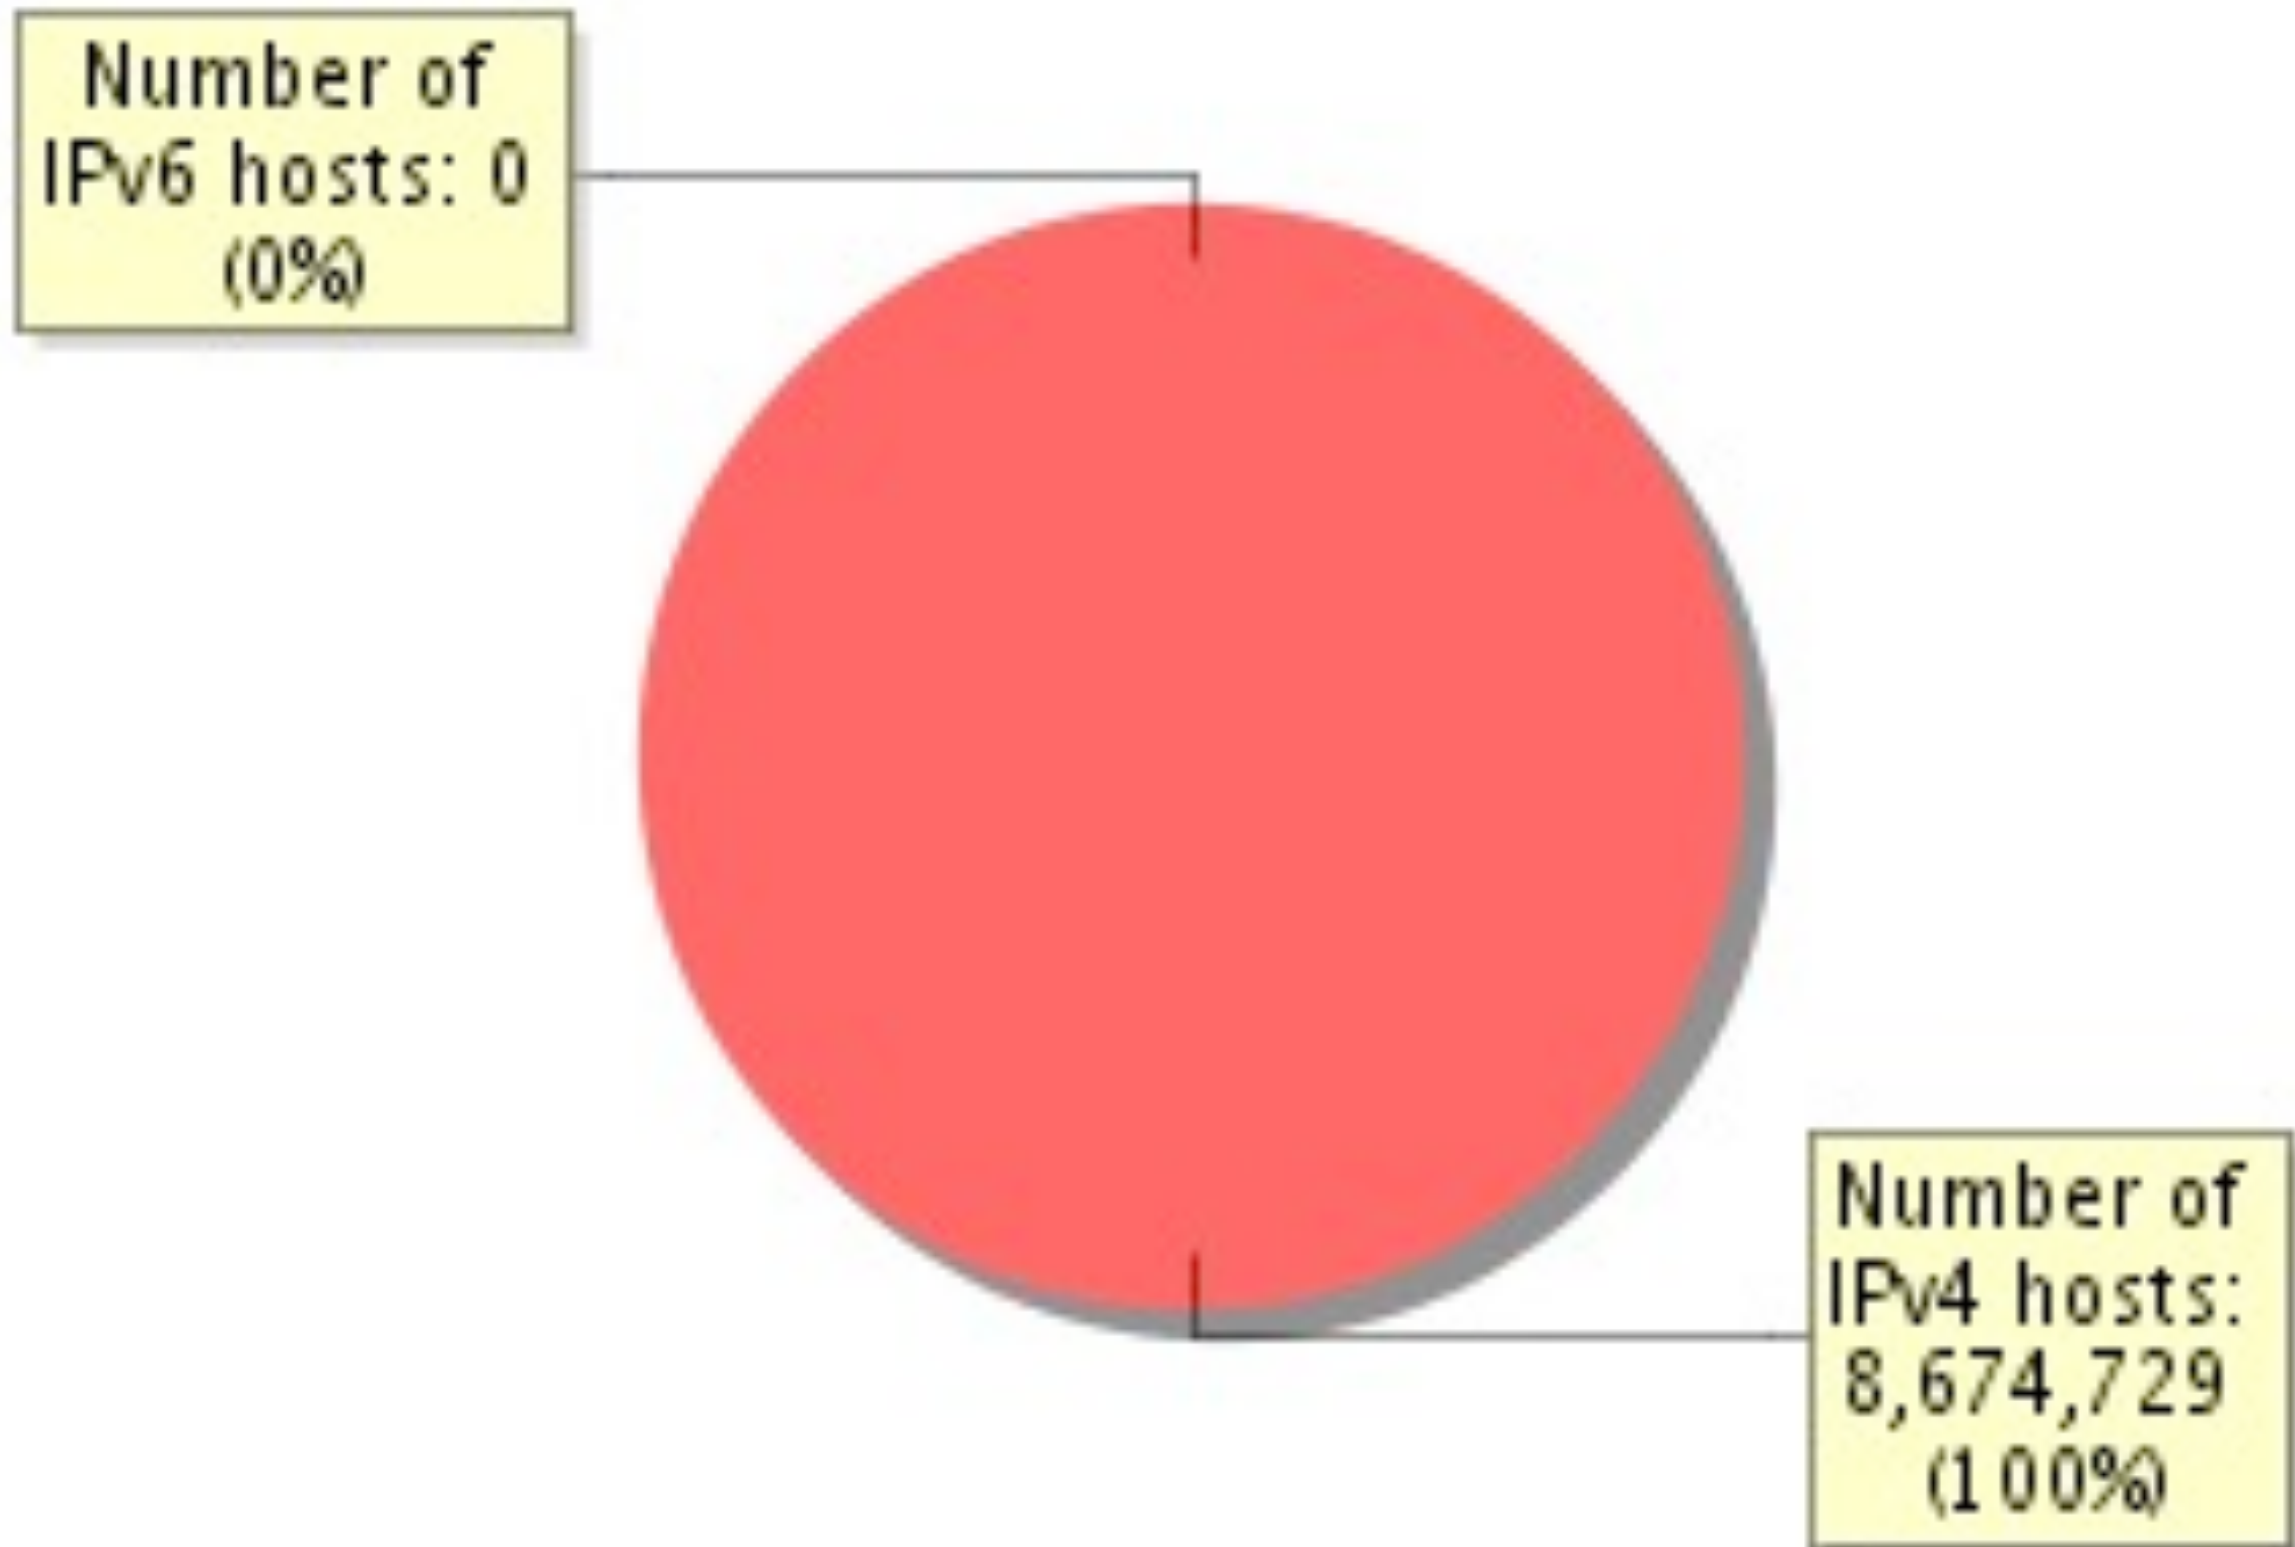





Hostcount++ – data for .uk



Hostcount++ – data for .nl



Hostcount++ – data for .nl





Routing Information Service (RIS)



Routing Information Service (RIS)

- Collects BGP routing data using collectors all over the world
- IPv4 since Dec 1999, IPv6 since Nov 2002
 - 12 out of 15 collectors have IPv6 peers
- Historical overview kept forever
- Data presented here based on a subset!
 - Taken from LINX collector



RIS data

- 305K IPv4 prefixes
- 1800 IPv6 prefixes


(this includes some internal prefixes, which are not seen by others)

- No surprises here



RIS data

■ Avg. updates per day per prefix



IPv6 prefixes almost 1.5 times as unstable



Test traffic measurements (TTM)



TTM

- One-way delay measurements between custom boxes hosted in remote networks
- 77 boxes online, 36 measuring IPv6


DNSMON

- Measures reachability and latency for root and TLD name servers using TTM grid



RIS data

Average median latency in ms between 18 boxes,
600 measurements, 5 years ago




In 2004, IPv6 latency was 38% higher than IPv4



RIS data

Average median latency in ms between 36 boxes,
2000 measurements, now



IPv6 latency is 17% higher



TTM data – IPv6 tunneling in 2004

Destination Testbox																				
	tt01	tt100	tt103	tt105	tt107	tt13	tt25	tt35	tt42	tt52	tt55	tt56	tt72	tt73	tt77	tt85	tt86	tt94	tt97	tt98
Source Testbox	tt01	noVal	1280	1480	1500	1500	1480	1500	1500	1500	1500	1280	1500	1500	1500	1476	1476	1500	1500	1500
	tt100	noVal	noVal	noVal	noVal	1500	noVal	1500	1280	1480	1480	1480	1480	1500	1462	1476	1476	1480	1480	1480
	tt103	1280	noVal		1280	1280	1280	1280	1280	1280	1280	1280	1280	1280	1280	1280	1280	1280	1280	1280
	tt105	1480	noVal	1500		1476	1480	1480	1476	1480	1480	1476	1480	1476	1476	1480	1480	1480	1480	1480
	tt107	1500	noVal	1280	1500		1500	1280	1500	1500	1500	1500	1280	1500	1500	1462	1476	1476	1500	1500
	tt13	1500	noVal	1280	1480	1500		1480	1500	1500	1500	1500	1280	1500	1500	1500	1476	1476	1500	1500
	tt25	1476	noVal	1476	1500	1476		1476	1428	1476	1476	1280	1476	1476	1476	1476	1476	1476	1476	1476
	tt35	1480	noVal	1280	1480	1500	1500	1476		1500	1500	1500	1280	1500	1500	1462	1476	1476	1480	1480
	tt42	1500	noVal	1280	1500	1500	1500	1480	1500		1500	1500	1500	1500	1500	1500	1500	1500	1500	1500
	tt52	1500	noVal	1280	1500	1500	1500	1480	1500	1280		1500	1280	1500	1500	1462	1476	1476	1500	1500
	tt55	1500	noVal	1280	1500	1500	1500	1480	1500	1480	1500		1280	1500	1500	1462	1476	1476	1500	1500
	tt56	1476	noVal	1280	1476	1476	1476	1280	1476	1280	1280	1476		1476	1476	1462	1476	1476	1476	1476
	tt72	1500	noVal	1280	1500	1500	1500	1480	1500	1280	1500	1500	1280		1500	1462	1476	1476	1500	1500
	tt73	1500	noVal	1280	1500	1500	1500	1480	1500	1500	1500	1500	1280	1500		1500	1500	1500	1500	1500
	tt77	1500	noVal	1280	1500	1476	1500	1476	1476	1500	1476	1476	1280	1476	1500		1476	1476	1500	1500
	tt85	1476	noVal	1280	1500	1476	1476	1280	1476	1500	1280	1476	1280	1476	1476	1462		1500	1476	1476
	tt86	1476	noVal	1280	1500	1476	1476	1280	1476	1500	1476	1476	1280	1476	1476	1462	1500		1476	1476
	tt94	1500	noVal	1280	1480	1500	1500	1480	1500	1500	1500	1500	1280	1500	1500	1500	1476	1476		1500
	tt97	1500	noVal	1280	1480	1500	1500	1480	1500	1500	1500	1500	1280	1500	1500	1500	1476	1476	1500	
	tt98	1500	noVal	1280	1480	1500	1500	1480	1500	1500	1500	1500	1280	1500	1500	1500	1476	1476	1500	1500



Questions?

Erik Romijn <eromijn@ripe.net>
For TTM & RIS: IS Demo stand

