FROM AN IPV4 GLOBAL INTERNET TO A MIX OF IPV4 NATED AND IPV6 WORLD….

Alain Durand- Dir of Software Engineering PSG/CTO,
adurand@juniper.net

2011

JUNIPER PERSPECTIVE ON IPV4 EXHAUSTION AND IPV6 DEPLOYMENT

IPV4 REALITY CHECK:
IANA FREE POOL HAS EXHAUSTED

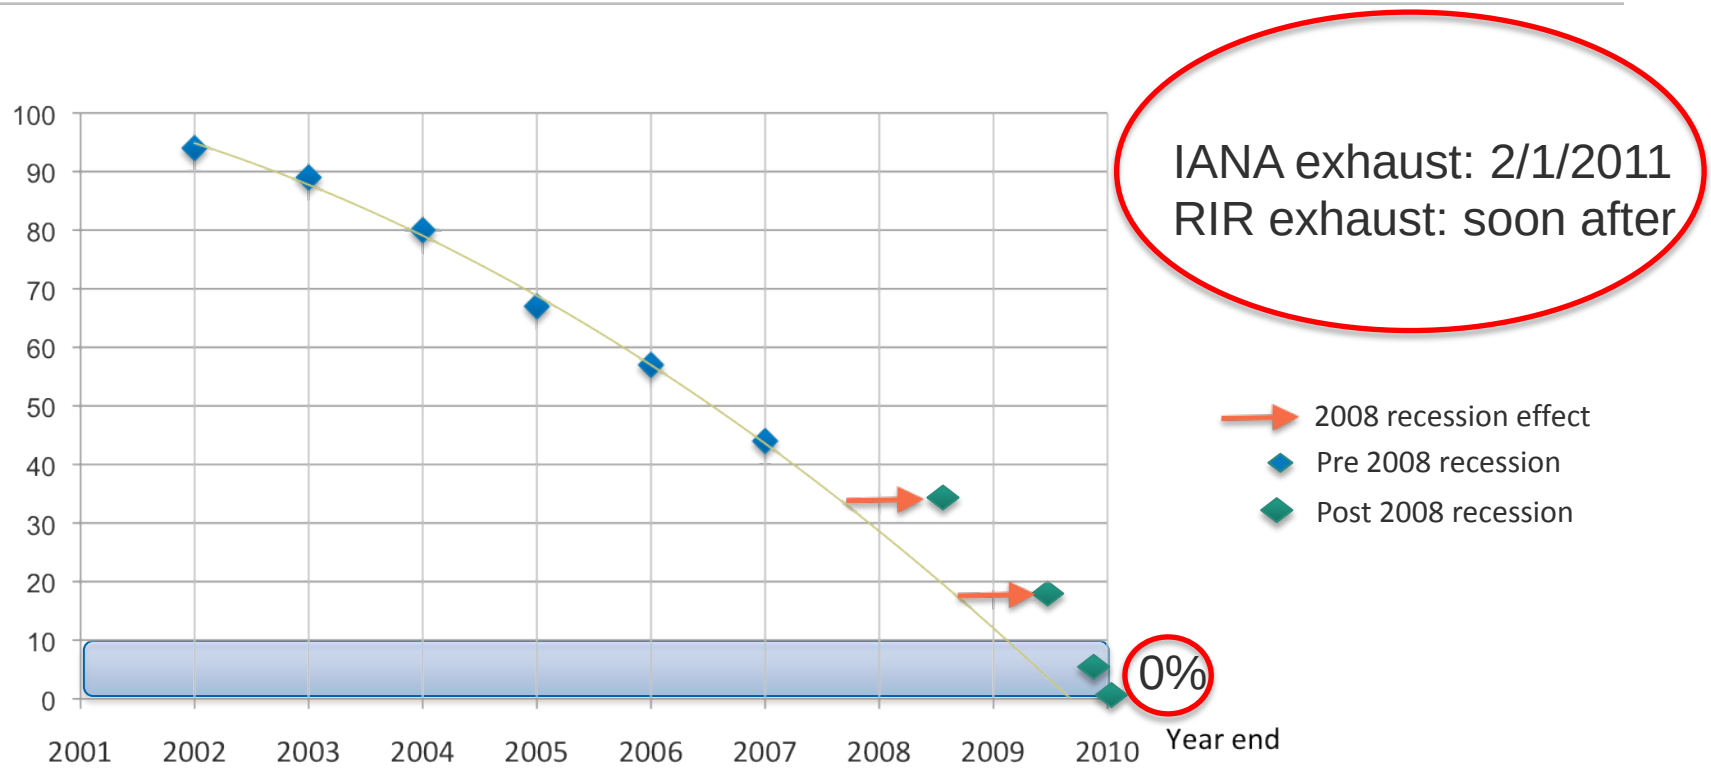

After completion:
Existing IPv4 addresses will not stop working.
Current networks will still operate.

Copyright © 2010 Juniper Networks, Inc. www.juniper.net

IPV6 REALITY CHECK: THE IPV4 LONG TAIL

Post IPv4 allocation completion:

- Many hosts & applications in customer residential networks (eg Win 95/98/2000/XP, Playstations, consumer electronic devices) are IPv4-only.
- Most software & servers in enterprise network are IPv4-only
  - They will not function in an IPv6-only environment.
  - Few of those can or will upgrade to IPv6.
- Content servers (web, email,…) are hosted on the Internet by many different parties. It will take time to upgrade those to IPv6.

Current measurement:

0.15% of Alexa top 1-million web sites are available via IPv6

(This number has not changed in the last 12 months)

Source: http://ipv6monitor.comcast.net

Copyright © 2010 Juniper Networks, Inc.

IS IPV6 TAKING OFF?

A number of very large ISPs and very large content providers are deploying IPv6 and various transition technologies **now**.

- Still early in the adoption curve.
- However, momentum is building.
- Can’t be ignored.

IPv6 does not solve the immediate problem of IPv4 address exhaust.

- Maintaining IPv4 service after IPv4 exhaustion is #1 priority for most players.
- This implies some form or another of IPv4 address sharing: NAT
- Many transition technologies to choose from
  - Impact on routing and network architecture

Copyright © 2010 Juniper Networks, Inc.

INDUSTRY IPV6 SCORE CARD

Function	Element	Status
Network	Core Router: T	
	Edge Routers: MX, 6PE	
Servers	Linux 2.6+	
	Datacenter equipments, CDN	
End-user clients	Windows 7 (Many XP boxes out there)	
	MacOS 10.x	
	Game consoles Wii, PS3, Xbox	
Software	Web Browser: Firefox, IE, Safari	
	Skype	
	On-line PC games	
	SSL VPN	
Content	**Web content available over IPv6**	
CE	**CPEs**	

Copyright © 2010 Juniper Networks, Inc. www.juniper.net

IS IPV6 TAKING OFF? MEASUREMENTS

COMCAST IPV6 MONITOR /1

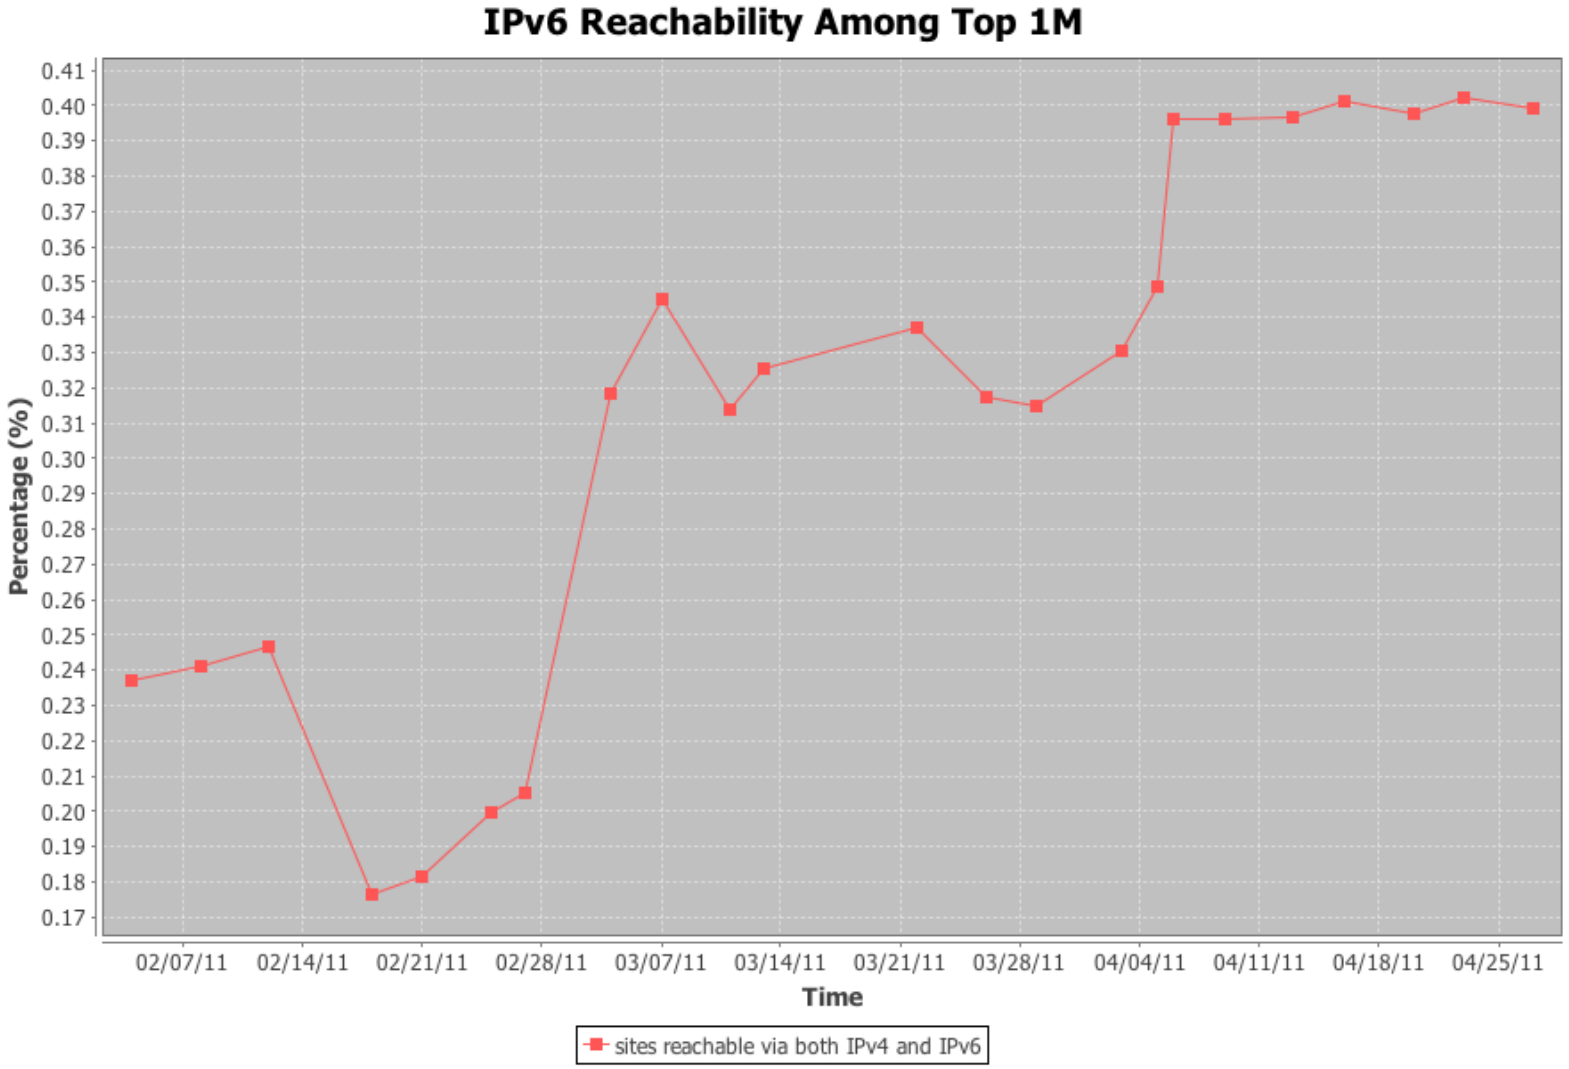

Copyright © 2010 Juniper Networks, Inc. www.juniper.net

COMCAST IPV6 MONITOR /2

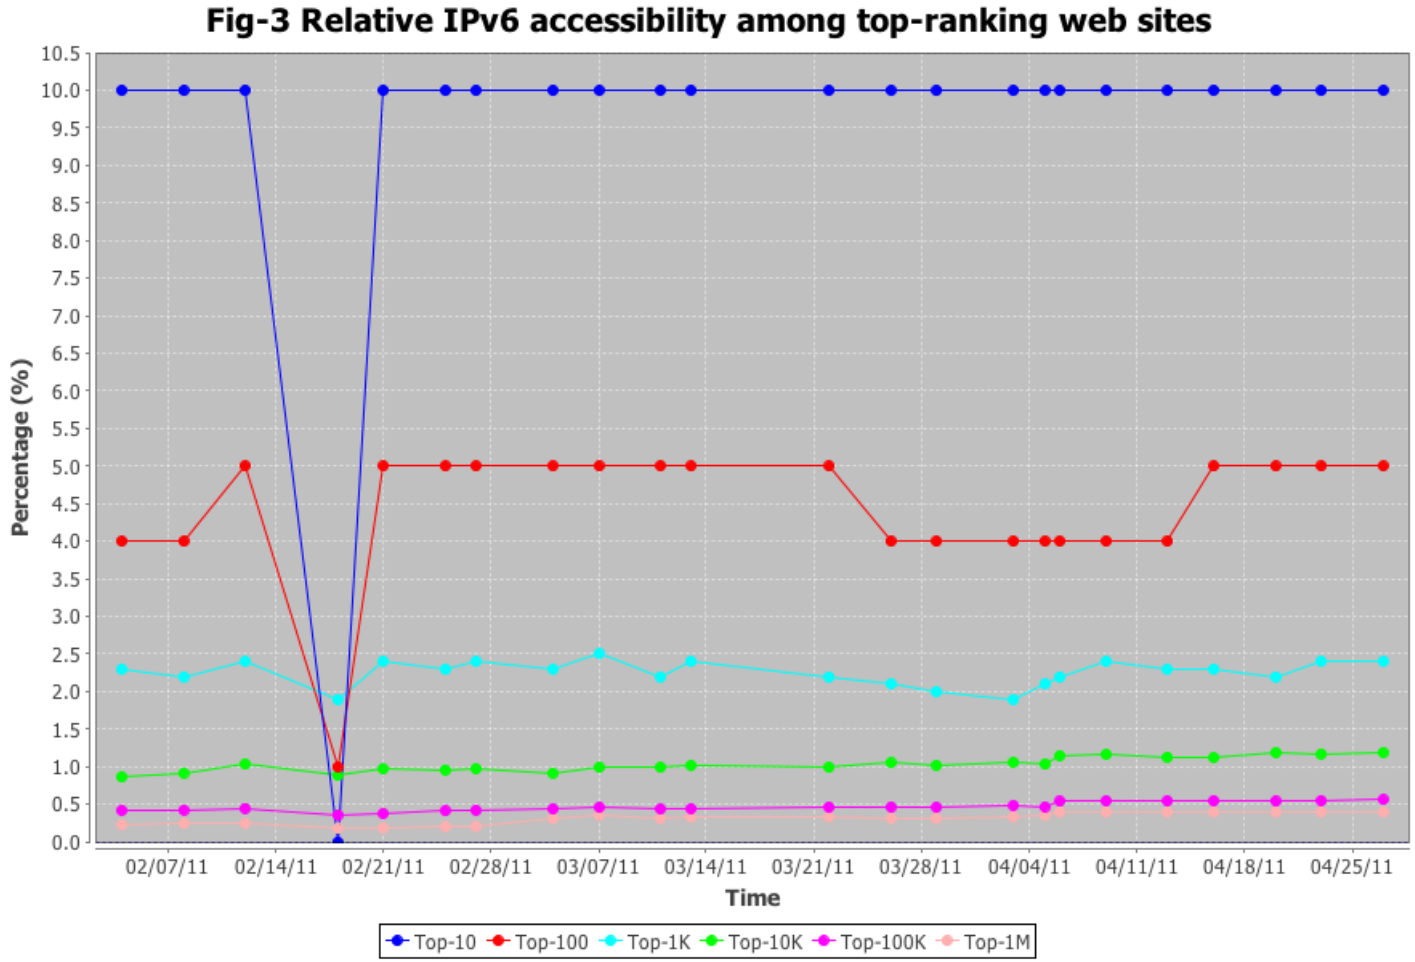

There is a direct correlation between content popularity and IPv6 presence.

Source: http://ipv6monitor.comcast.net

Copyright © 2010 Juniper Networks, Inc. www.juniper.net

COMCAST IPV6 MONITOR /3

Site (Top 1000)	Alexa rank	IPv4 download Time (ms)	IPv6 download Time (ms)	Delta IPv6/IPv4
google.com	1	196	344	76%
google.com.hk	18	205	331	61%
google.co.jp	28	200	369	85%
google.cn	87	519	1333	157%
netflix.com	121	287	322	12%
free.fr	167	1192	1165	-2%
comcast.net	186	838	528	-37%
scribd.com	252	419	667	59%
seznam.cz	286	1083	1015	-6%
comcast.com	390	338	478	42%
nu.nl	524	964	2627	173%
softlayer.com	578	188	306	63%
sapo.pt	647	1814	6722	271%
opera.com	753	758	1054	39%
telegraaf.nl	802	4106	6903	68%
novinky.cz	901	1216	3634	199%
doctissimo.fr	906	1352	3199	137%
01net.com	939	1212	2397	98%

User experience with IPv6, measured by download time, is **not** on par with IPv4.

Source: http://ipv6monitor.comcast.net

Copyright © 2010 Juniper Networks, Inc.

ROUTES

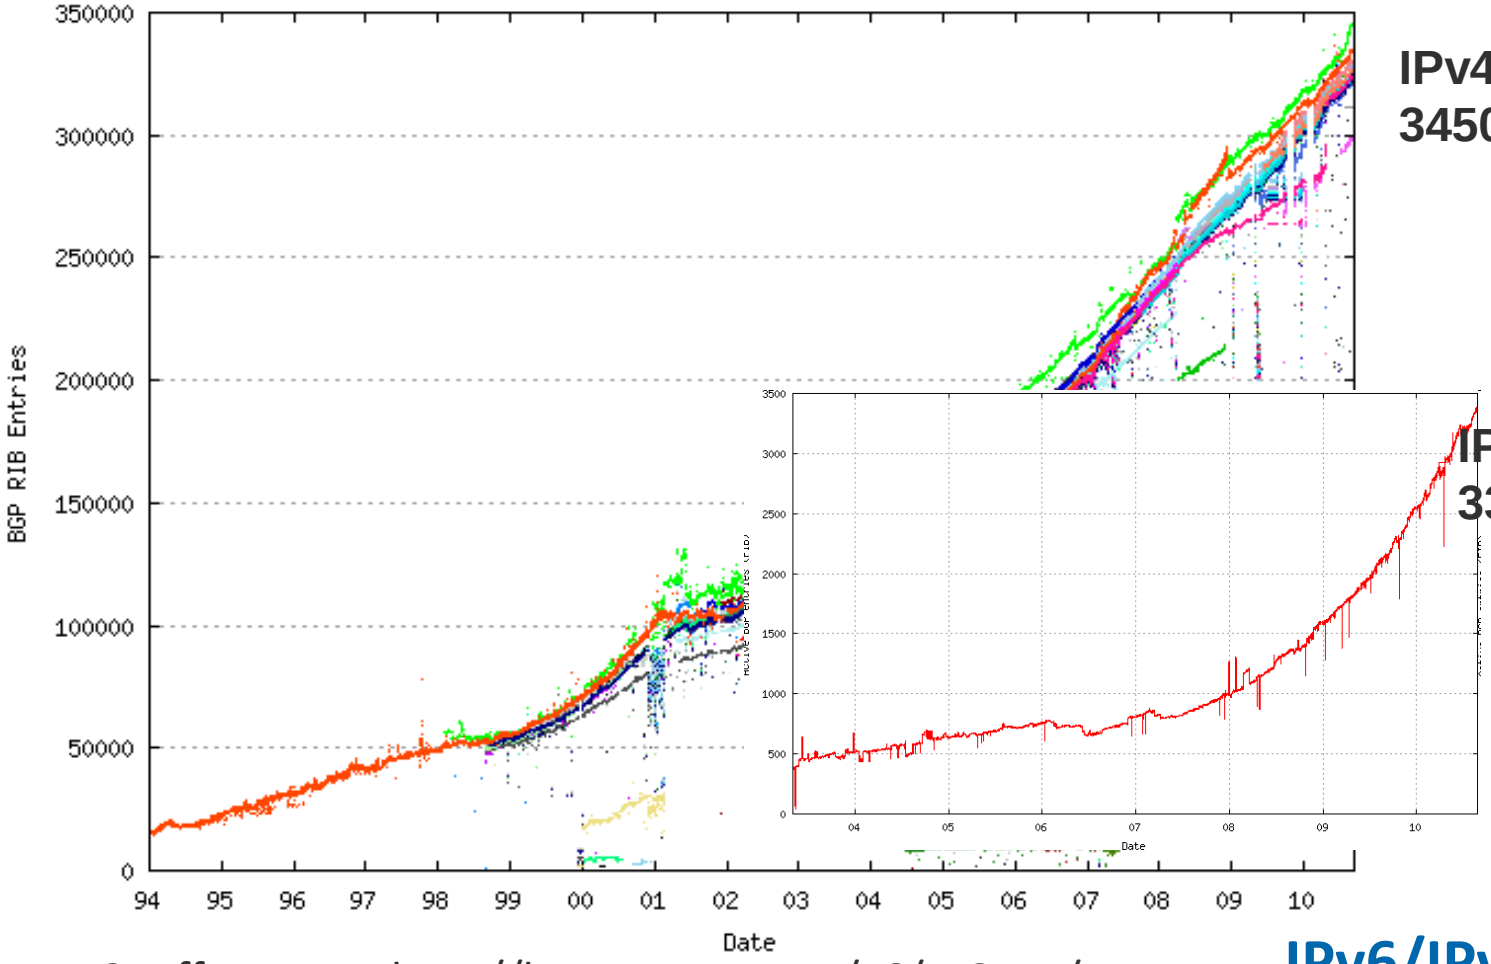

IPv4 routes: 345075

IPv6 routes: 3389

Source : Geoff Houston http://bgp.potaroo.net/v6/as6447/

IPv6/IPv4 = 1%

Copyright © 2010 Juniper Networks, Inc. www.juniper.net

DNS QUERY

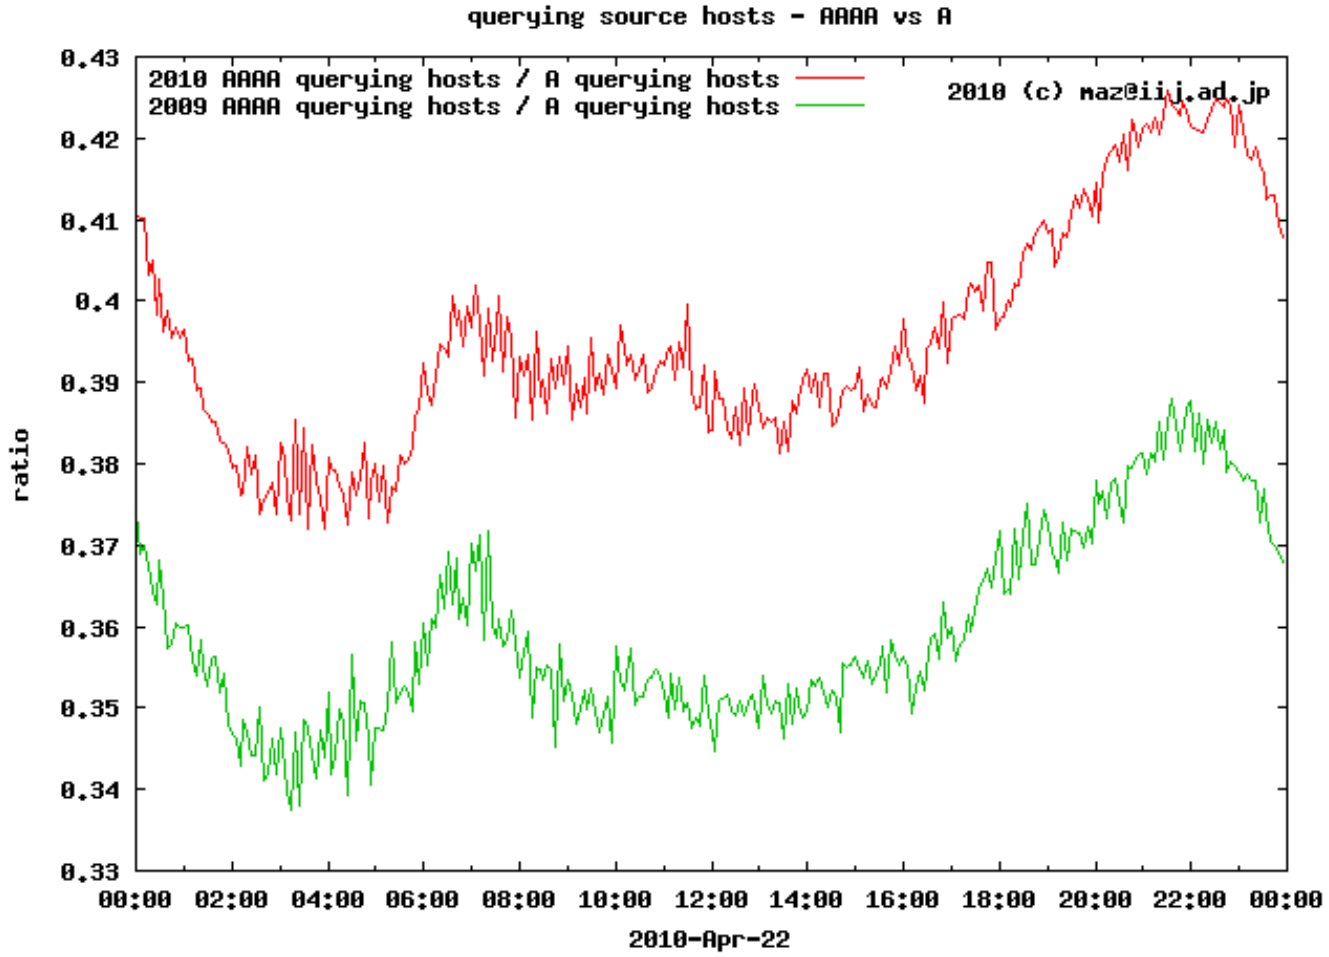

Source : Yoshinobu Matsuzaki @ IIJ

IPv6/IPv4 = 0.41%

Copyright © 2010 Juniper Networks, Inc. www.juniper.net

TRAFFIC DATA FROM ARBOR

Source: http://www.arbornetworks.com/en/ipv6-report.html*

IPv6 traffic << 1%

* Study made in 2008

Copyright © 2010 Juniper Networks, Inc.

CASE STUDIES OF ISP NETWORKS

CASE STUDY 1: INCUMBENT

Incumbent ISP in a mature market

- Business has been growing a lot in the last couple years, but growth has slowed down.
- Saturated market

ISP can reclaim address internally

Redesigning networks to get more address efficiency
More aggressively NATing wireless subscribers

As a consequence:

- ISP does not see the urge to move to IPv6 right now.
  - Wait until technology mature.
  - Synchronize IPv6 deployment with roll-out of next gen service

Copyright © 2010 Juniper Networks, Inc.

CASE STUDY 2: OLD/NEW ACCESS TECHNOLOGY

ISP offer two access technologies, a legacy one and a new one

- Growth & ARPU is happening in the new technology, not the older
- Deploying IPv6 in legacy environment may be costly

Migrate IPv4 addresses from legacy environment to new environment:

- Carrier Grade NAT (CGN) & 6rd in legacy world

- Public IPv4 & native IPv6 in new world

Issue: cost of replacing CPEs to support IPv6

- With 6rd offered as an optional service, a service provider can offload the cost of replacing CPEs in the old technology to the end-users who want to be early adopters of IPv6.

Copyright © 2010 Juniper Networks, Inc.

CASE STUDY 3: NEW CUSTOMERS, NEW NETWORKS

ISP makes a clear distinction between current, existing customers and post-exhaustion customers.

Build new IPv6-based networks for new customers.

IPv4 is a service overlayed on top of IPv6 with DS-Lite (with or without a Carrier-Grade NAT)

Enabling customers to run their applications expecting incoming connections (Eg: Set-Top box control, P2P):

- PCP (Port Control Protocol) to open-up pin-holes on CGN

ISP offers **new** IPv6 CPEs to **new** customers.

Copyright © 2010 Juniper Networks, Inc.

CASE STUDY 4: MOBILE

The key issue is license cost:

	Dual-Stack (NAT44)	IPv6-only (NAT64)
License cost 2G & 3G/3GPPr8 (using separate PDP contexts for IPv4 & IPv6)	Two licenses: 1 for IPv4 PDP + 1 for IPv6 PDP	1 for IPv6 PDP
License cost LTE and 3G/3GPPr9 (using a combined PDP context for IPv4&IPv6)	1 for IPv4/IPv6 PDP/bearer	1 for IPv6 PDP/bearer

Preferred

Going IPv6-only + NAT64 works **ONLY if** all applications are converted to IPv6 and there is no connectivity to external devices such as PCs.

Dual-Stack remains the preferred/simplest general solution.

Copyright © 2010 Juniper Networks, Inc.

CASE STUDY 5: BUSINESS ISP

ISP has a corporate mandate to prepare for IPv6

Issue: ISP will have to support legacy IPv4 devices/apps operated by their customers as well.

Reduce drastically (to just a few?) the number of IPv4 addresses allocated to business customers. NAT is performed by the business CPEs.

Copyright © 2010 Juniper Networks, Inc.

CASE STUDY 6: INTERNATIONAL ISP

ISP is incumbent is a region/country and want to expand internationally. Need to offer IPv6 quickly.

6PE is a good way to jumpstart IPv6 global presence.

IPS will have to migrate to native IPv6 at some point in the future.

Copyright © 2010 Juniper Networks, Inc.

OBSERVATIONS ABOUT TRANSITION TECHNIQUES

All transition techniques (NAT444+6RD, NAT64, DS-Lite) revolve around the notion of sharing IPv4 addresses via some form of NAT.

They all require the exact same amount of IPv4 addresses to be shared in a NAT pool.

- The difference is how packets are transported to the NAT

Sharing addresses among customers introduces issues:

- LEA/Abuse/Logging/Geo-location/Access control

Copyright © 2010 Juniper Networks, Inc.

2011:
IPV6 ALONE IS NOT THE ANSWER TO IPV4 ADDRESS DEPLETION

IPV6 AS A L2.5

IPV6 IS A DRAMA IN FOUR ACTS

Prelude: IPv4 exhaustion happens in 2011.

Act I: NAT solves IPv4 exhaust.

Act II: IPv6 to simplify IPv4 service delivery.
IPv6 networks with IPv4 overlays enable the management of a large number of customers while maintaining an IPv4 service.

Act III: Emergence of IPv6 content.
The decoupling of deploying IPv6 networks from the deployment of IPv6 applications & content solves the chicken and egg problem. IPv6 traffic is a cap& grow strategy around NAT scaling issues.

Act IV: IPv4 dies (very slowly) .
IPv4 & IPv6 co-exist until IPv6 become pervasive.

Copyright © 2010 Juniper Networks, Inc.

EXAMPLE #1 OF IPV6 UNDER-LAYER (L2.5): DS-LITE

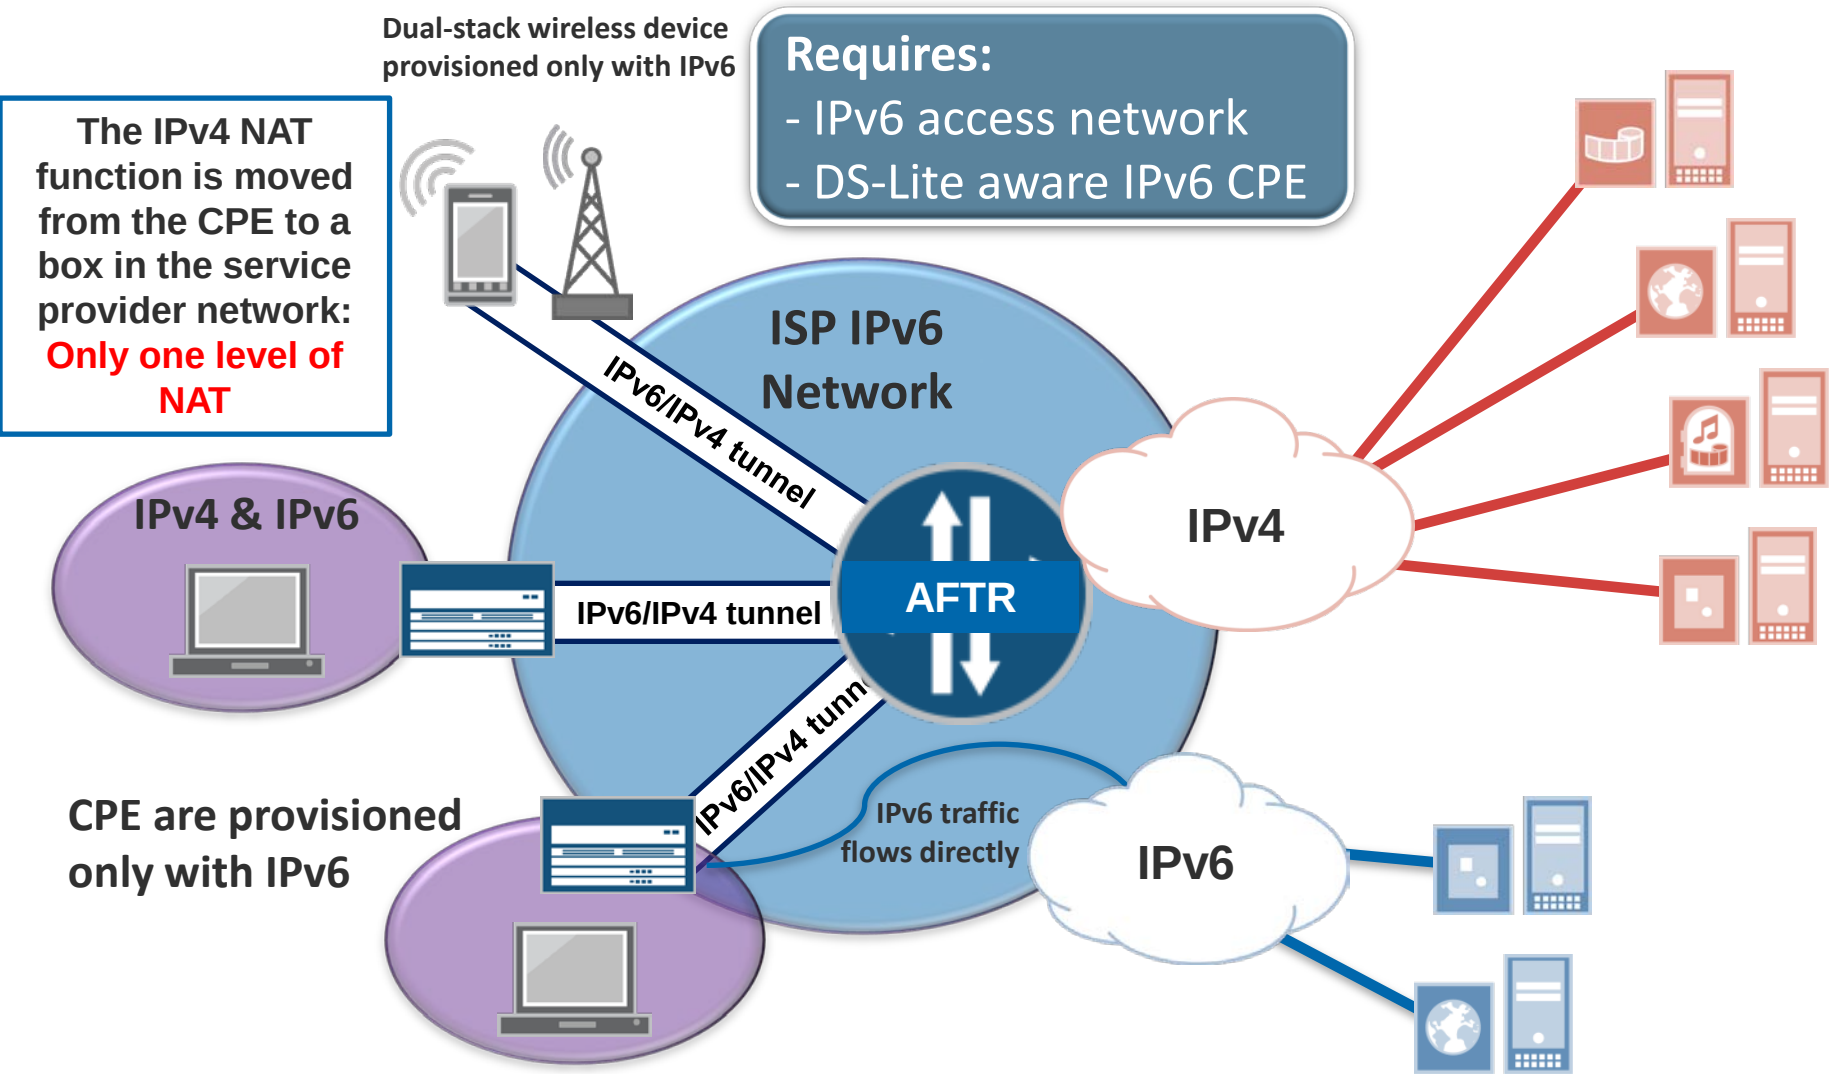

Copyright © 2010 Juniper Networks, Inc.

EXAMPLE #2 OF IPV6 UNDER-LAYER (L2.5): VIRTUAL IPV4 AGGREGATION

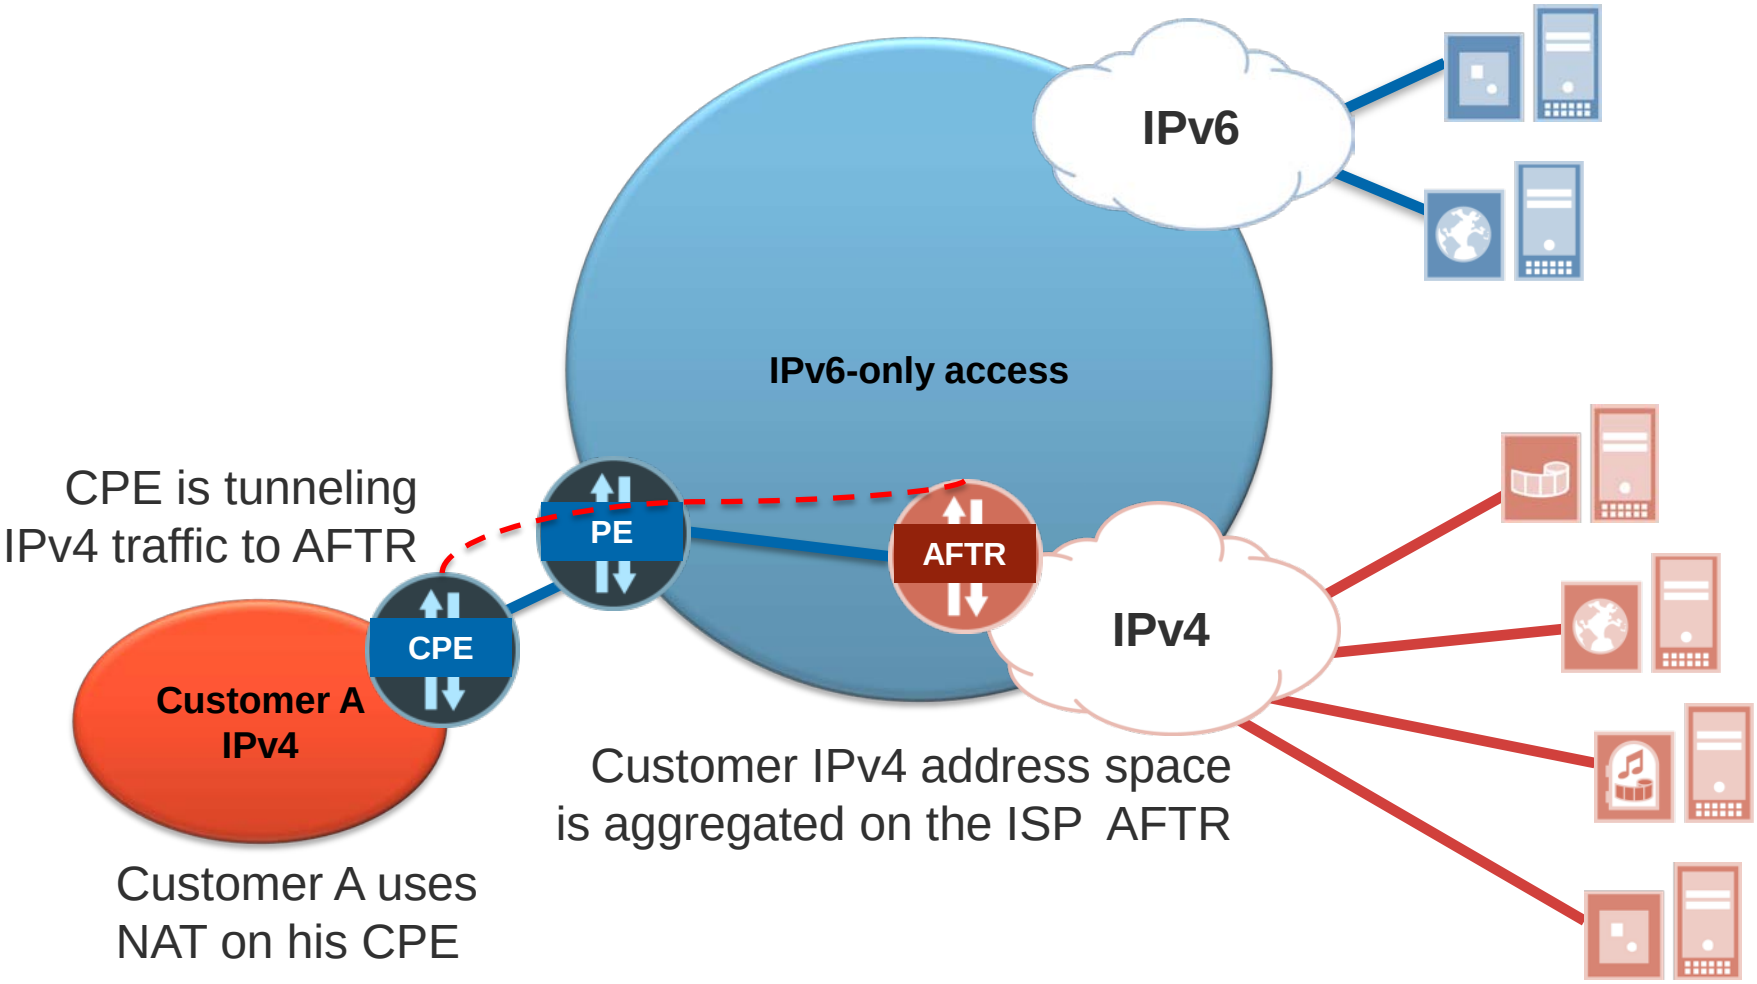

Copyright © 2010 Juniper Networks, Inc.

DELIVERING SERVICES OVER IPV6

DELIVERING IPV6 SERVICES WITH IPV6 SERVERS

Cannot offer IPv6 services until everything is IPv6 ready.

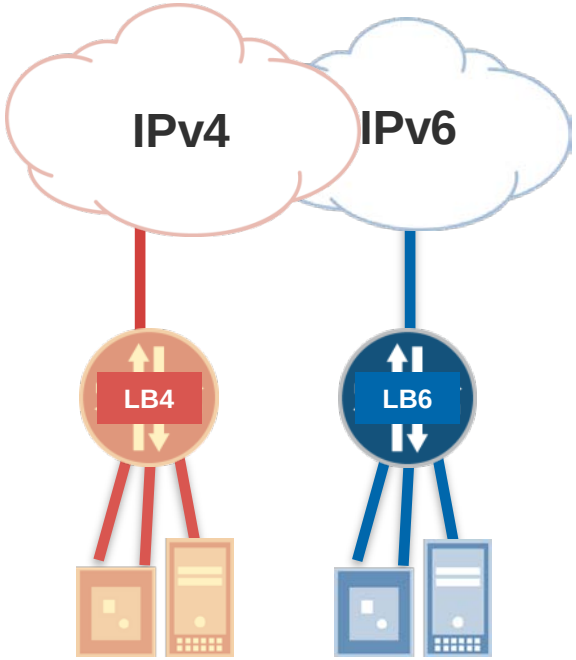

IPv6 servers
Dedicated IPv6 load balancers

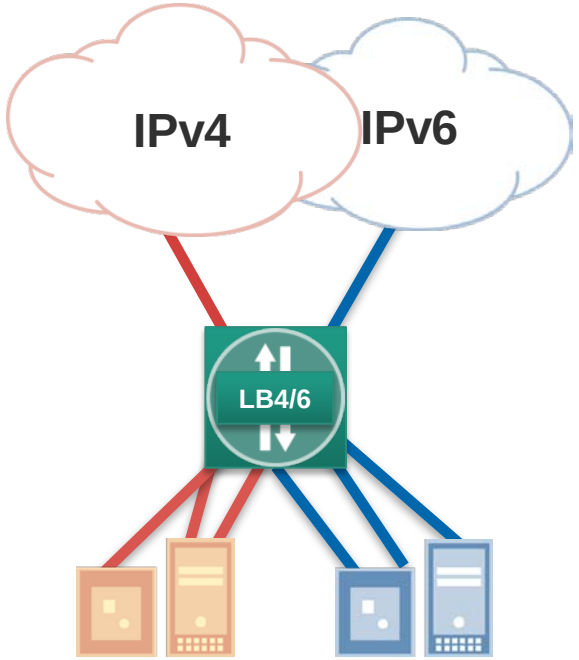

IPv6 servers
Shared IPv4/IPv6 load balancers

Copyright © 2010 Juniper Networks, Inc.

ENABLING IPV6 SERVICE DELIVERY BASED ON IPV4 SERVERS

Decouple IPv6 deployment in the network from IPv6 deployment on the servers

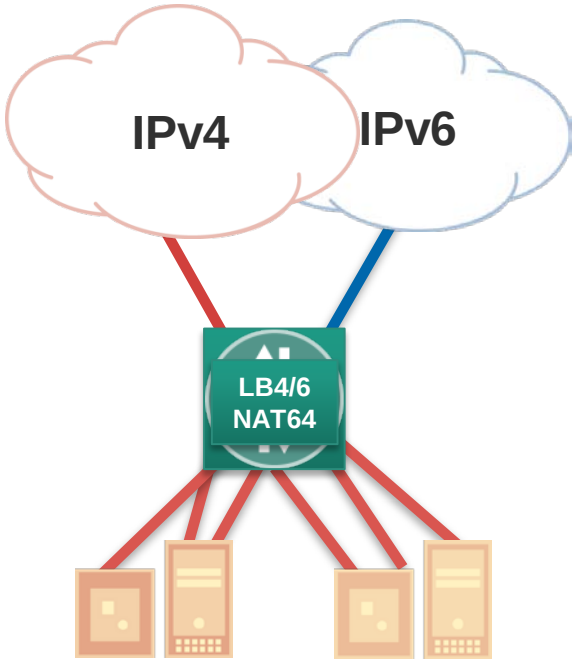

IPv4-only servers
IPv4/IPv6 road balancers with NAT64

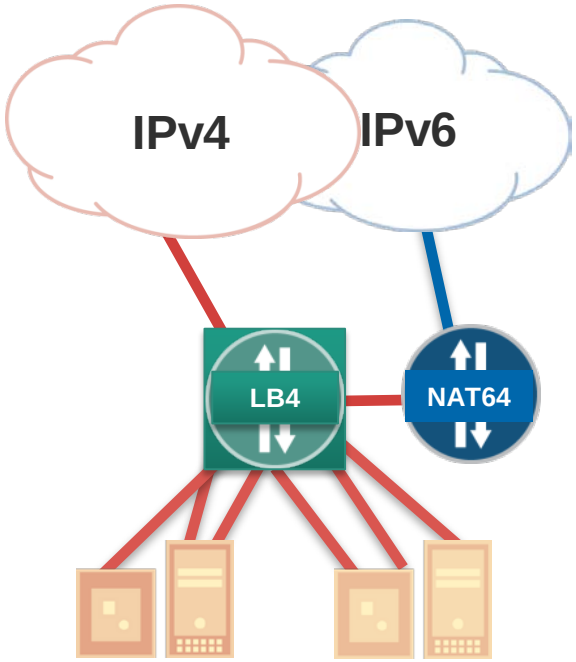

IPv4-only servers
IPv4-only load balancer
NAT64 as a front-end to load balancer (enable gradual introduction of IPv6, Good to quickly establish IPv6 presence while traffic is low)

Copyright © 2010 Juniper Networks, Inc.

MAKING WWW.JUNIPER.NET IPV6 REACHABLE

http://ipv6.juniper.net reachable over IPv6 since Jan. 8th

- Uses ‘translator in the cloud’

Commitment to participate to the IPv6 world day on June 8th with http://www.juniper.net

Copyright © 2010 Juniper Networks, Inc.

PROBLEM STATEMENT: GETTING CONTENT AVAILABLE OVER IPV6 QUICKLY

How to get example.com web site available over IPv6 quickly and at the lowest possible cost?

- ☐ Get everything dual-stack (Network, Load-balancer, Servers…)
- ☐ Get the network dual-stack and leave the servers IPv4

  (Easier, as the engineering teams dealing with servers are often not the same as the ones dealing with the network)

- ☑ Don't touch anything and let some else handle the problem…

An IPv6->IPv4 translator in the cloud can do this translation for you.

Copyright © 2010 Juniper Networks, Inc.

PRODUCT TO BUILD: “TRANSLATOR IN THE CLOUD” TO QUICKLY DELIVER IPV6 SERVICE

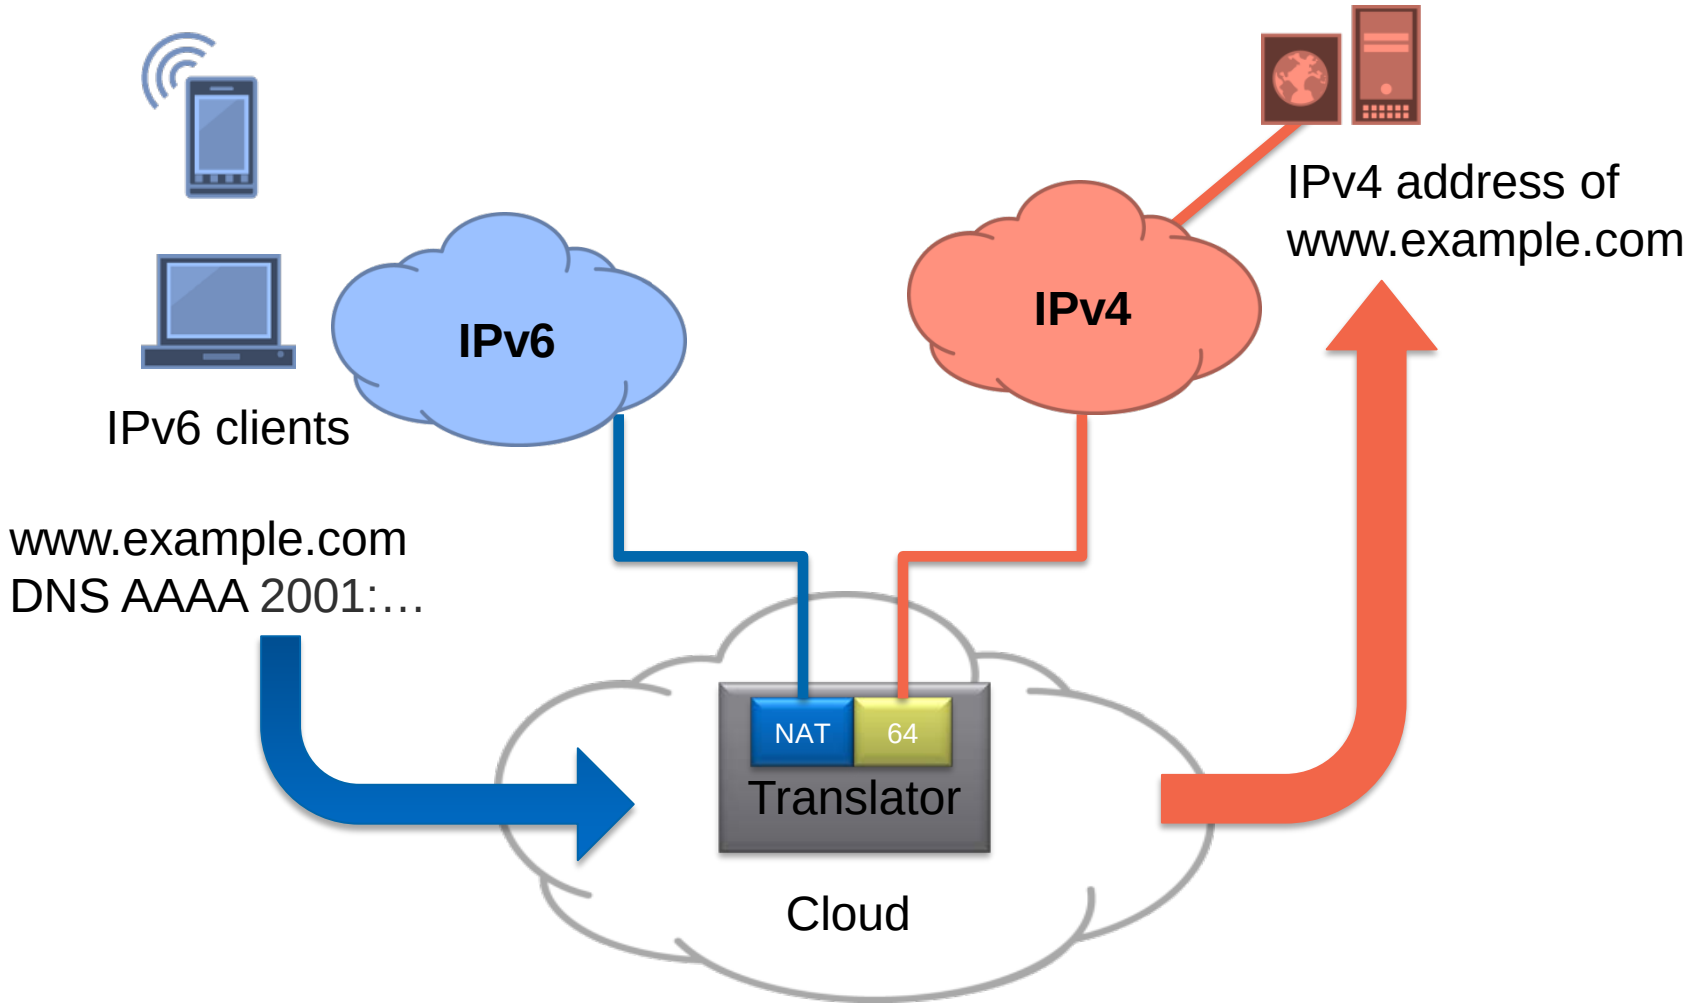

Copyright © 2010 Juniper Networks, Inc.

TECHNOLOGY: PCP

TECHNOLOGY: PCP (NEW DEVELOPMENT)

PCP: Port Control Protocol

PCP objectives are to enable applications to receive incoming connections in the presence of an ISP NAT/Firewall.

Instead of ‘working around’ NATs like other NAT traversal techniques like STUN/TURN/ICE, PCP enables an explicit dialog between applications and the NAT.

PCP can be seen as a ‘carrier-grade’ evolution of UPnP-IGD and NAT-PMP.

The work on PCP is done at IETF in a new working group co-chaired by Alain Durand (Juniper) & Dave Thaler (Microsoft).

Copyright © 2010 Juniper Networks, Inc.

PCP IN A NUTSHELL

Applications negotiate ports with the ISP NAT to establish external presence.
Application asks: “I’d like to get port 5000 for 48 hours”, NAT PCP server responds: “I give you port 6003 for 12 hours”.

No more keep-alive!

Better radio efficiency
Better battery life

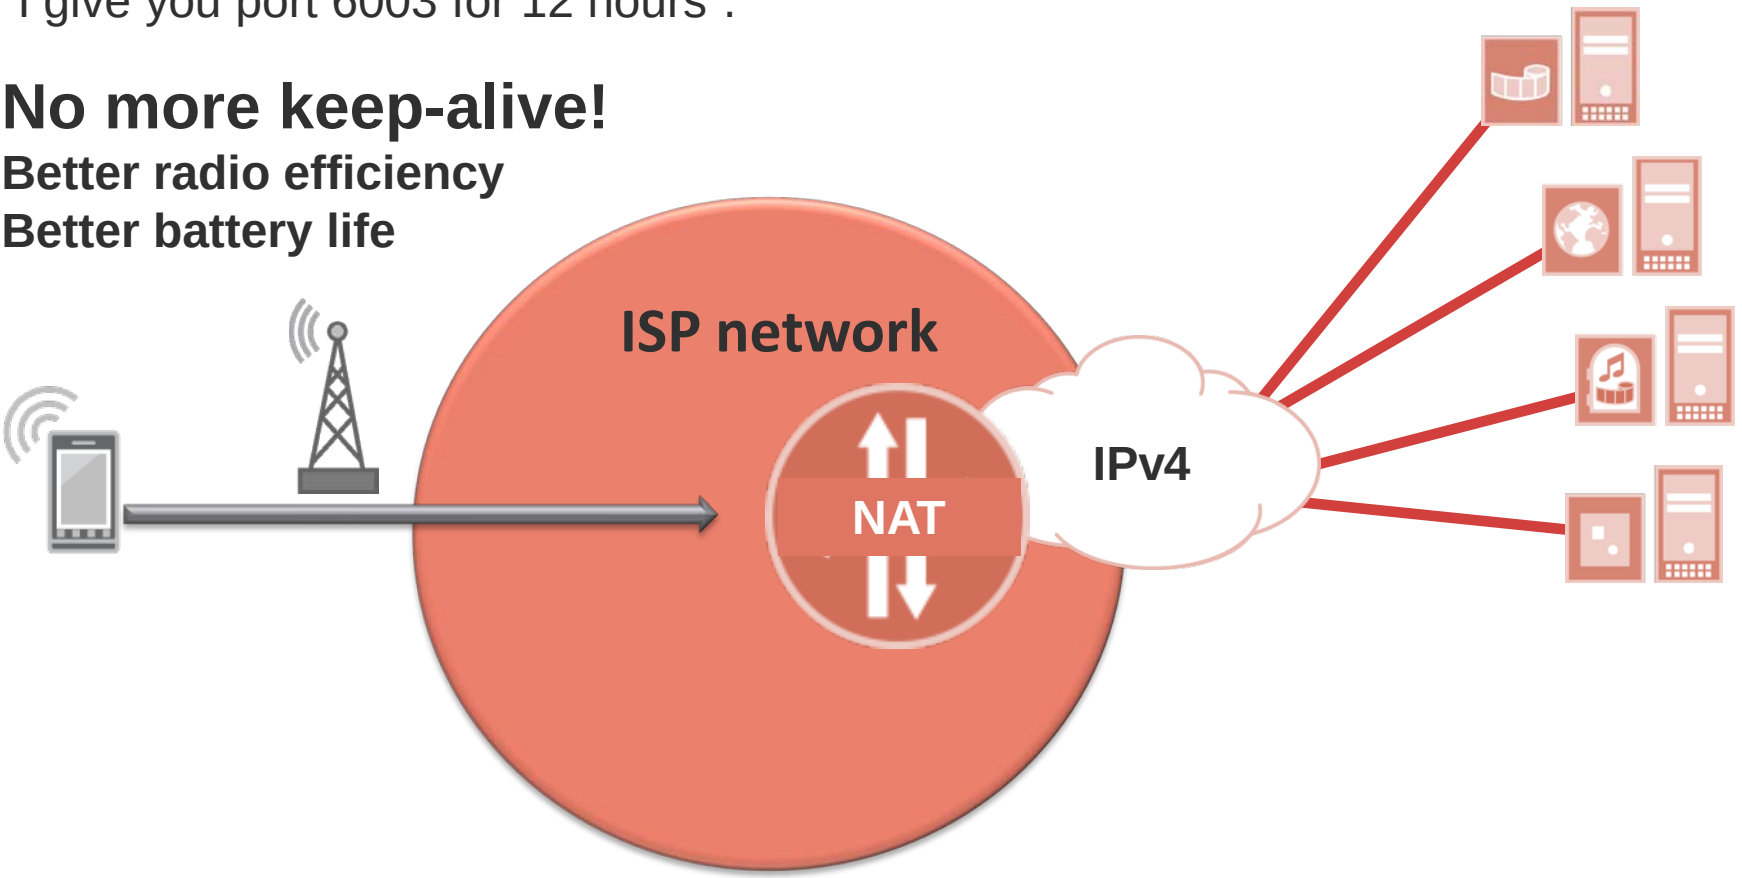

Copyright © 2010 Juniper Networks, Inc.

PCP IN WIRELINE

Applications negotiate ports with the ISP NAT to establish external presence.
CPE can relay legacy or PCP requests.
PCP can be used to signal long live session that would benefit from HA

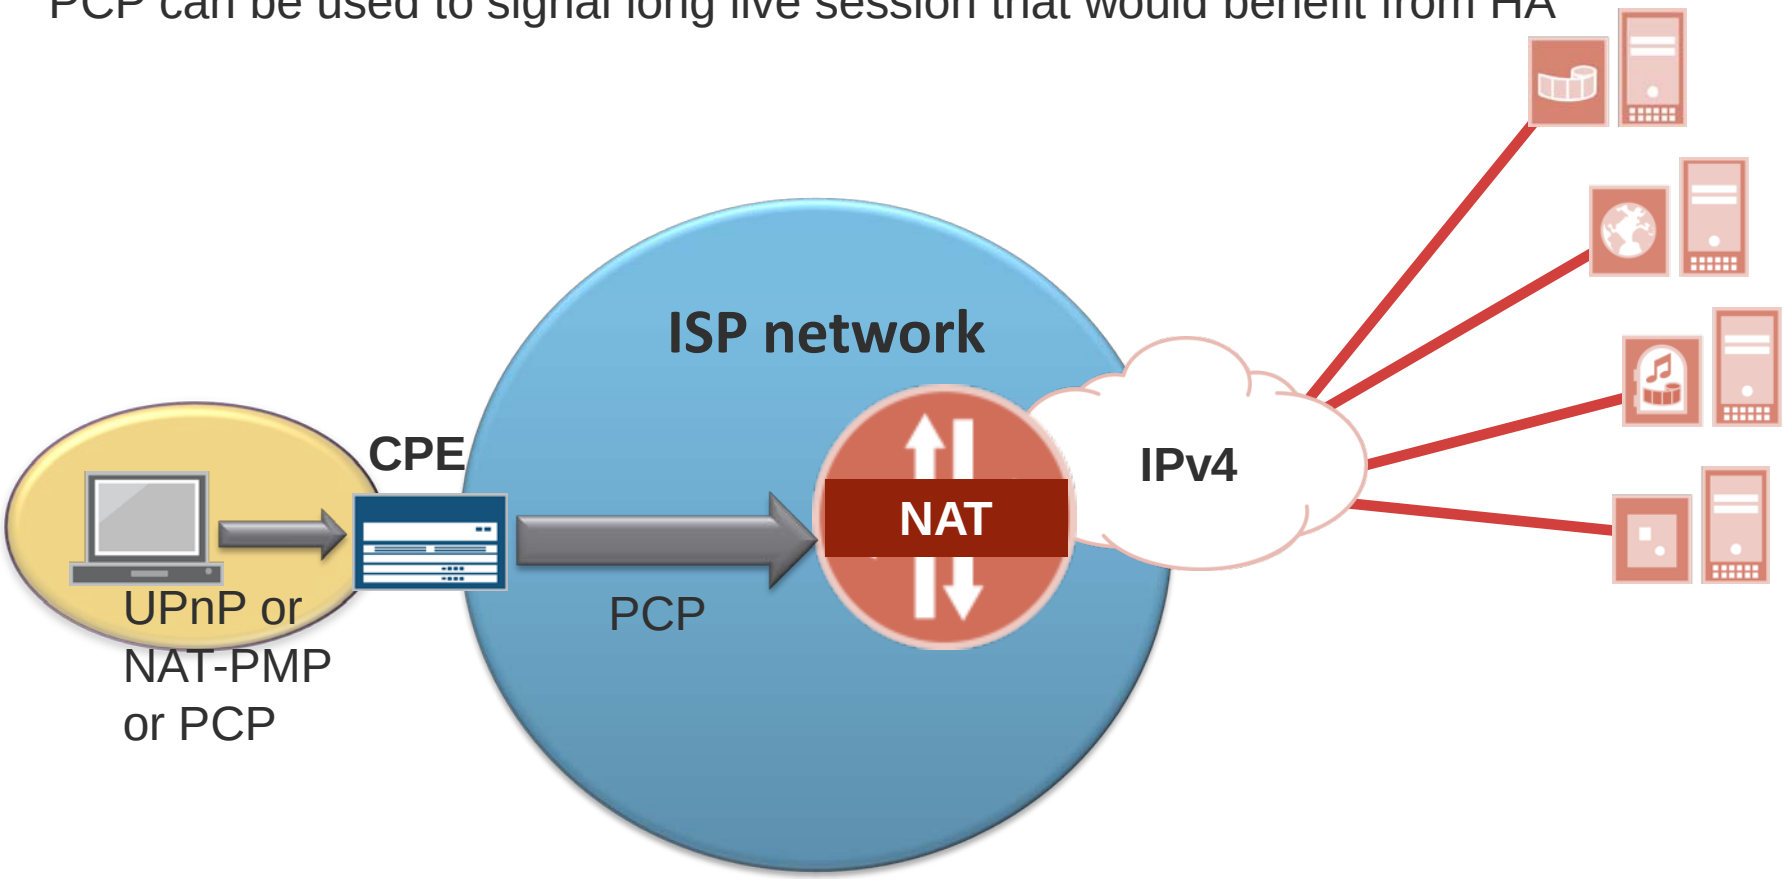

Copyright © 2010 Juniper Networks, Inc.

CONCLUSION

Now is the time to get serious about IPv6.
In doing so, it is critical to **preserve IPv4 service**.

Key hot topics are:

- Replacing every CPE to enable IPv6
- Making the operation of IPv4 NAT technologies scale
  - Control plane: PCP
- Getting content on IPv6

Copyright © 2010 Juniper Networks, Inc.

everywhere