

No more IPv4 Impact on applications and measuring IPv6 deployment

Distinguished Engineer, Cisco
evyncke@cisco.com

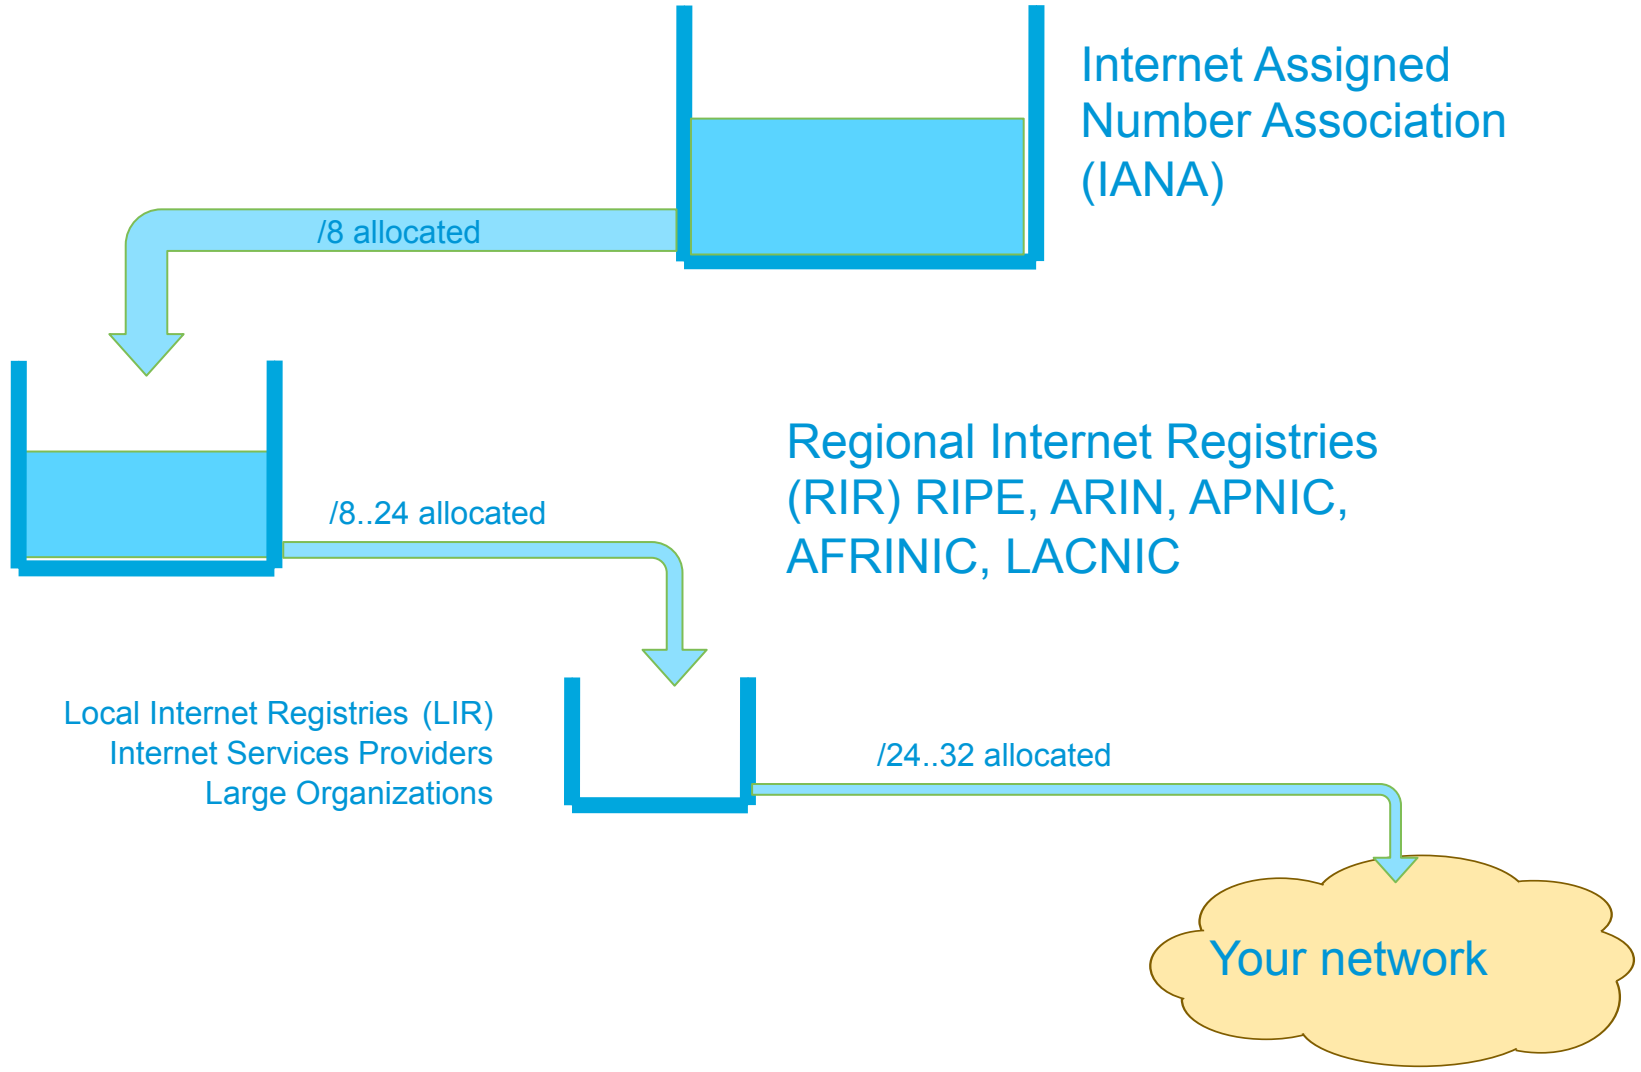
Eric Vyncke, IPv6 Council Belgium Co-Chair
Eric.Vyncke@ipv6council.be
www.ipv6council.be

[@evyncke](#)

PGP fingerprint = F2CB F935 6F18 1A1E A274 6359 77FA 1A2D 30F9 E979

February 2014

A Word About IPv4 Address Allocation



No Doubt Anymore: I



14 September 2012 Last updated at 15:08 GMT

Europe hits old internet address limits

By Mark Ward
Technology correspondent, BBC News

Europe has almost exhausted its stock of old-style internet addresses.

Strict rationing of these addresses - called IPv4 - has been started by the body that hands them out in Europe.

From now on, companies can only make one more application for IPv4 addresses and, if successful, will only get 1,024 of them.

In addition, any application for more old addresses must demonstrate how an organisation is using the new, replacement, addressing scheme.



Europe's stock of old-style net addresses has effectively run dry.

APNIC IPv4 Address Pool Reaches Final /8

Published on: Friday, 15 April 2011

We wish to inform you that as of Friday, 15 April 2011, the APNIC pool reached the Final /8 IPv4 address block, bringing us to [Stage Three of IPv4 exhaustion in the Asia Pacific](#).

/8 address policy

Requests will now be assessed under [section 9.10](#) in "Policies for IPv4 address space management in the Asia Pacific region".

APNIC's objective during [Stage Three](#) is to provide IPv4 address space for new entrants to the region and for those deploying IPv6.

From now on, [all new and existing APNIC account holders](#) will be entitled to receive a maximum allocation of a /22 from the Final /8 address space.

NOW on IPv6

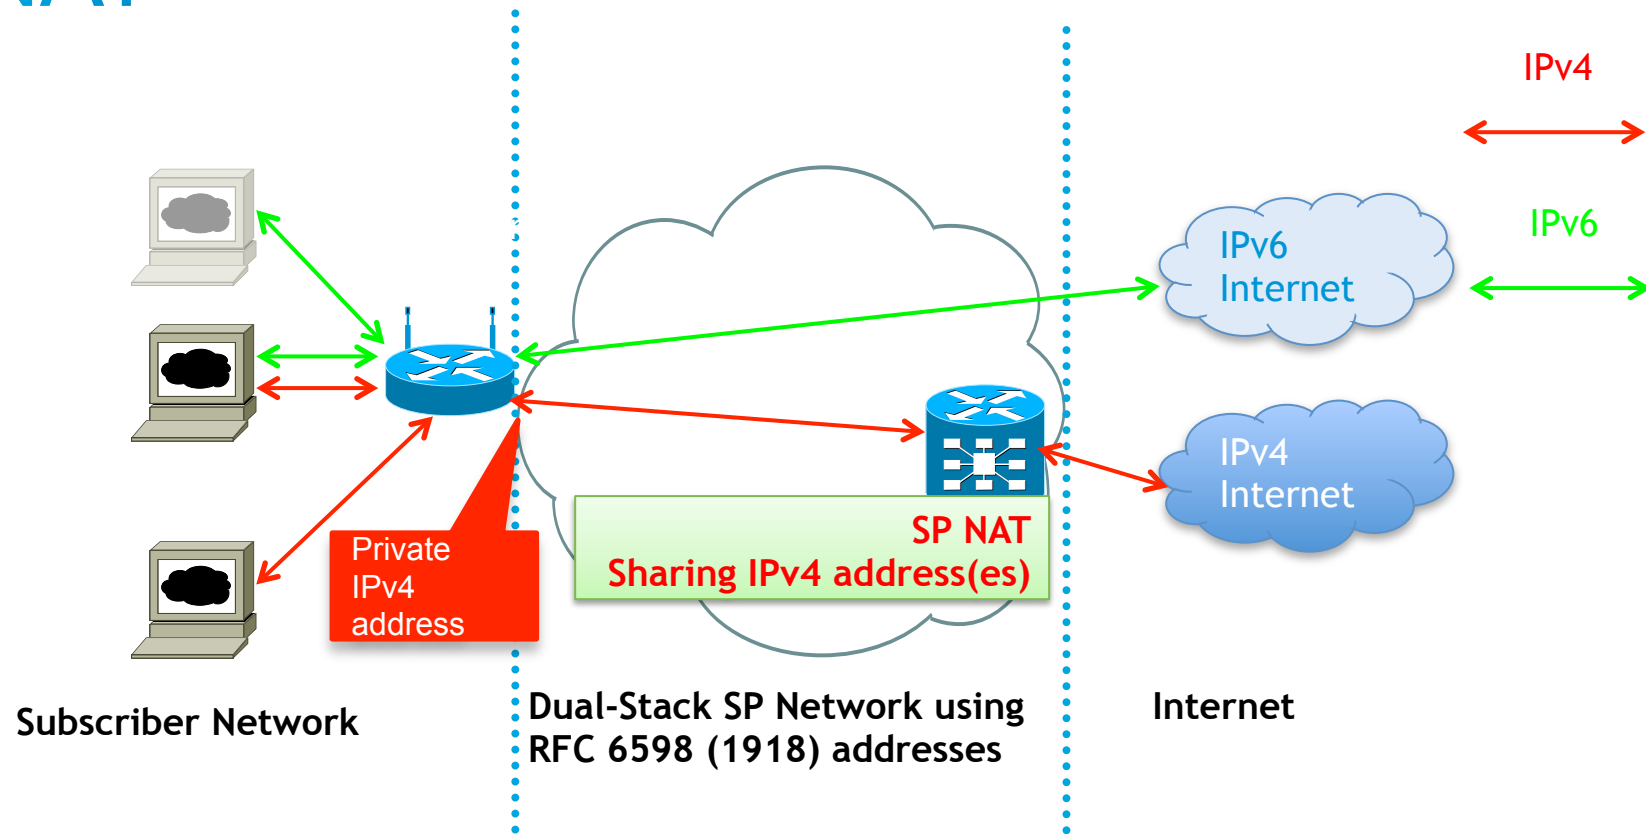
APNIC encourages Asia Pacific Internet community members to deploy IPv6 within their organizations. For more information, please refer to APNIC for [information regarding IPv6 deployment](#), statistics, training, and related operational policies.

"We have been anticipating for quite some time," states Raúl Echeberria, representative of the five RIRs. "The future of the Internet is in IPv6."

Impact on Internet Connectivity

- 3 approaches currently worked on by SP
 - Dual-stack:** SP allocates both IPv4 and IPv6 global addresses to his subscribers
 - Shared IPv4 address:** SP relies on NAT in his network to share a couple of global IPv4 addresses and allocates private IPv4 addresses to his subscriber
 - IPv6 only**
- It is expected that SP will deploy a combined and phased approach
 - Sharing IPv4 global address is part of the solution
 - IPv6-only as ultimate the goal, there is no alternative plan

Service Providers Dual-Stack (IPv6 + IPv4) with SP IPv4 NAT



- More likely scenario:

IPv6 being available all the way to the consumer

SP core and customer has to use IPv4 NAT due to v4 depletion

Impact on Internet Applications

- Enterprises should expect their customers, partners, and remote employees to have a mix of connectivity


Public IPv4-only

Public IPv4 and IPv6

Shared IPv4-only

Shared IPv4 and IPv6

IPv6 only



Impact of Shared IPv4 Address on Applications

- Applications which could **work poorly or even not at all** when one side uses a shared IPv4 address

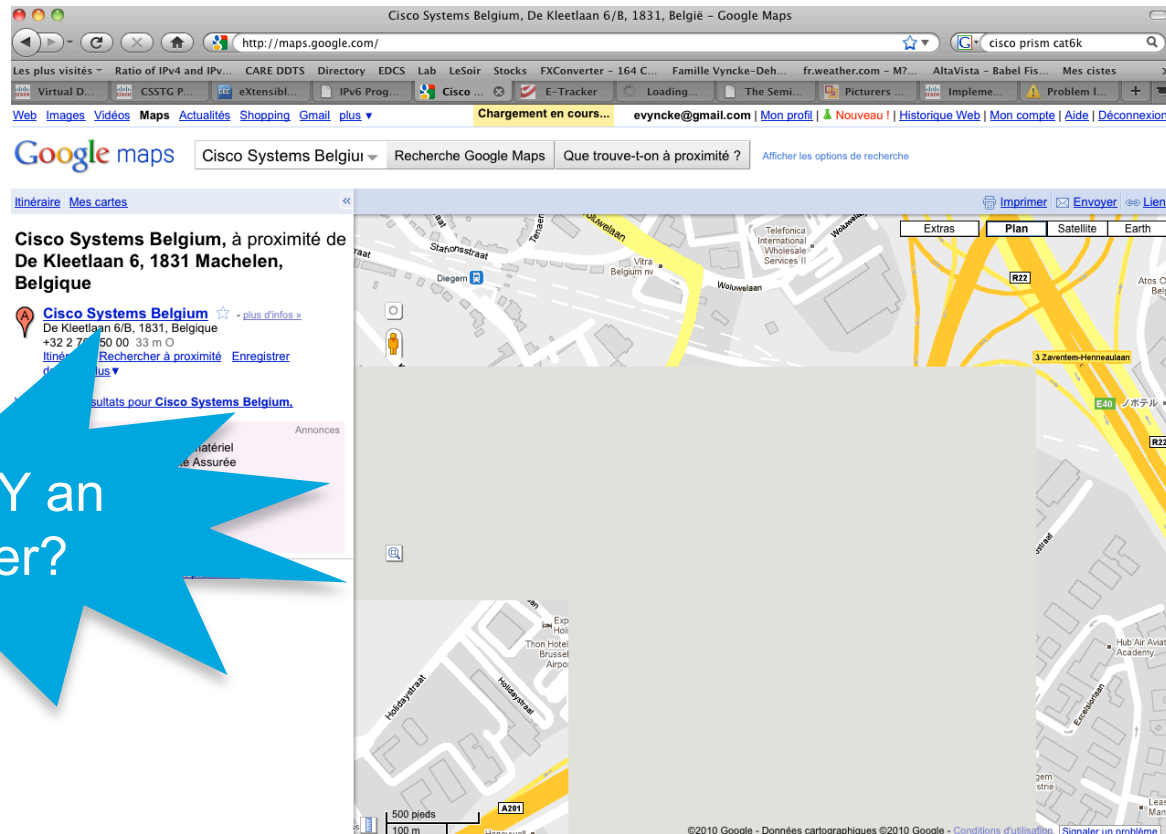
Multiple TCP connections (like AJAX) in parallel

Assuming that one IPv4 address = one user (for logging, for load balancing, for access control...)

Expecting inbound connections (like active FTP)

Using an application not yet supported by the NAT devices

Examples of Broken IPv4 Applications



Is SPDY an answer?

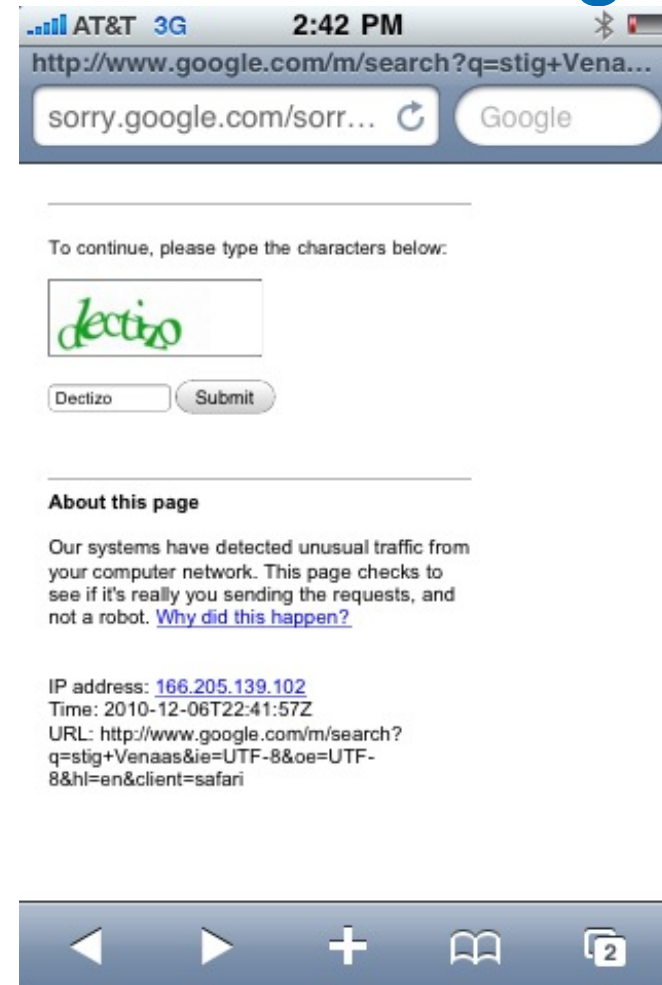
Shared IPv4 Address and Reputation

- Every IPv4 address has a reputation
 - Either blacklist or more sophisticated (senderbase.org)
 - Used to detect spam, botnet members, ...
- It is fine as long as:
 - One IPv4 == One legal entity (subscriber)
- What if
 - One IPv4 == 10.000 entities/subscribers through SP NAT?

Shared IPv4 Address and Rate Limiting

Applications throttle use per IPv4 address

- When address is sharing by 1000's of people the usage threshold is crossed
- And rate limiters are triggered even for legit traffic
- Example with AT&T using NAT for mobile phones



Shared IPv4 Address Forensic

- Servers currently keep the remote IPv4 address in their log
- Law Enforcement Agencies (LEA) can request any ISP to get the subscriber ID of this IPv4 address on a specific time
- With SP NAT, there will be 10,000 subscribers using this IPv4 address

Shared IPv4 Address and Forensic (Cont.)

- SP will have to keep all the translation log (data retention)
 - <time, subscriber internal IP, subscriber internal TCP/UDP port, **subscriber external TCP/UDP port**, Internet IP, Internet TCP/UDP port>
 - <10:23:02 UTC, 10.1.2.3, 6543, **23944**, 91.121.200.122, 80>
- **AND, the server will have to extend the log to include the TCP/UDP port**
 - See also RFC 6302 “Internet-Facing Server Logging”**
- *“At 10:23:02 who was using the shared port 23944?”*



For Your
Reference

Logging Remote TCP Port in Apache


```
LogFormat "[%h]:%{remote}p %l %u %t \"%r\" %>s %b" common
```

```
[10.55.185.71]:53958 - - [09/Sep/2011:10:10:26 +0200] "GET /nav/ HTTP/  
1.1" 200 7112
```


```
[10.55.185.71]:53958 - - [09/Sep/2011:10:10:27 +0200] "GET /nav/nav.js  
HTTP/1.1" 200 33519
```

```
[2001:6f8:1468:1000::4]:3268 - - [09/Sep/2011:10:10:49 +0200] "GET /  
ping_ws.php HTTP/1.0" 200 53
```

Ca



Courtesy of Jason Fesler, Yahoo (V6 World Congress 2012)



Let's Stop Talking about the legacy IPv4
Let's introduce IPv6

Jason Fesler

Courtesy of Jason Fesler, Yahoo (V6 World Congress 2012)

Short Introduction to IPv6

IPv6 in One Slide

- IPv6 is IPv4 with larger addresses
 - 128 bits vs. 32 bits
 - NAT no more needed => easier for applications
 - Simpler hence more security
- Data-link layer unchanged: Ethernet, WiFi, xDSL, ...
- Transport layer unchanged: UDP, TCP, ...
- Applications “unchanged”: HTTP, SSL, SMTP, ...
- IPv6 is not really BETTER than IPv4 because it is ‘new’
 - IPv6 has been specified in 1995...
 - IPsec is identical in IPv4 & IPv6
 - Only benefit is a much larger address space

Address Representation

- Format:

x:x:x:x:x:x:x:x where x is 16 bits hexadecimal field

2001:0000:130F:0000:0000:09C0:876A:130B

Case insensitive

Leading zeros in a field are optional:

2001:0:130F:0:0:9C0:876A:130B

Successive fields of 0 are represented as ::, but only once in an address:

2001:0:130F::9C0:876A:130B

2001:~~130F~~::9C0:876A:130B

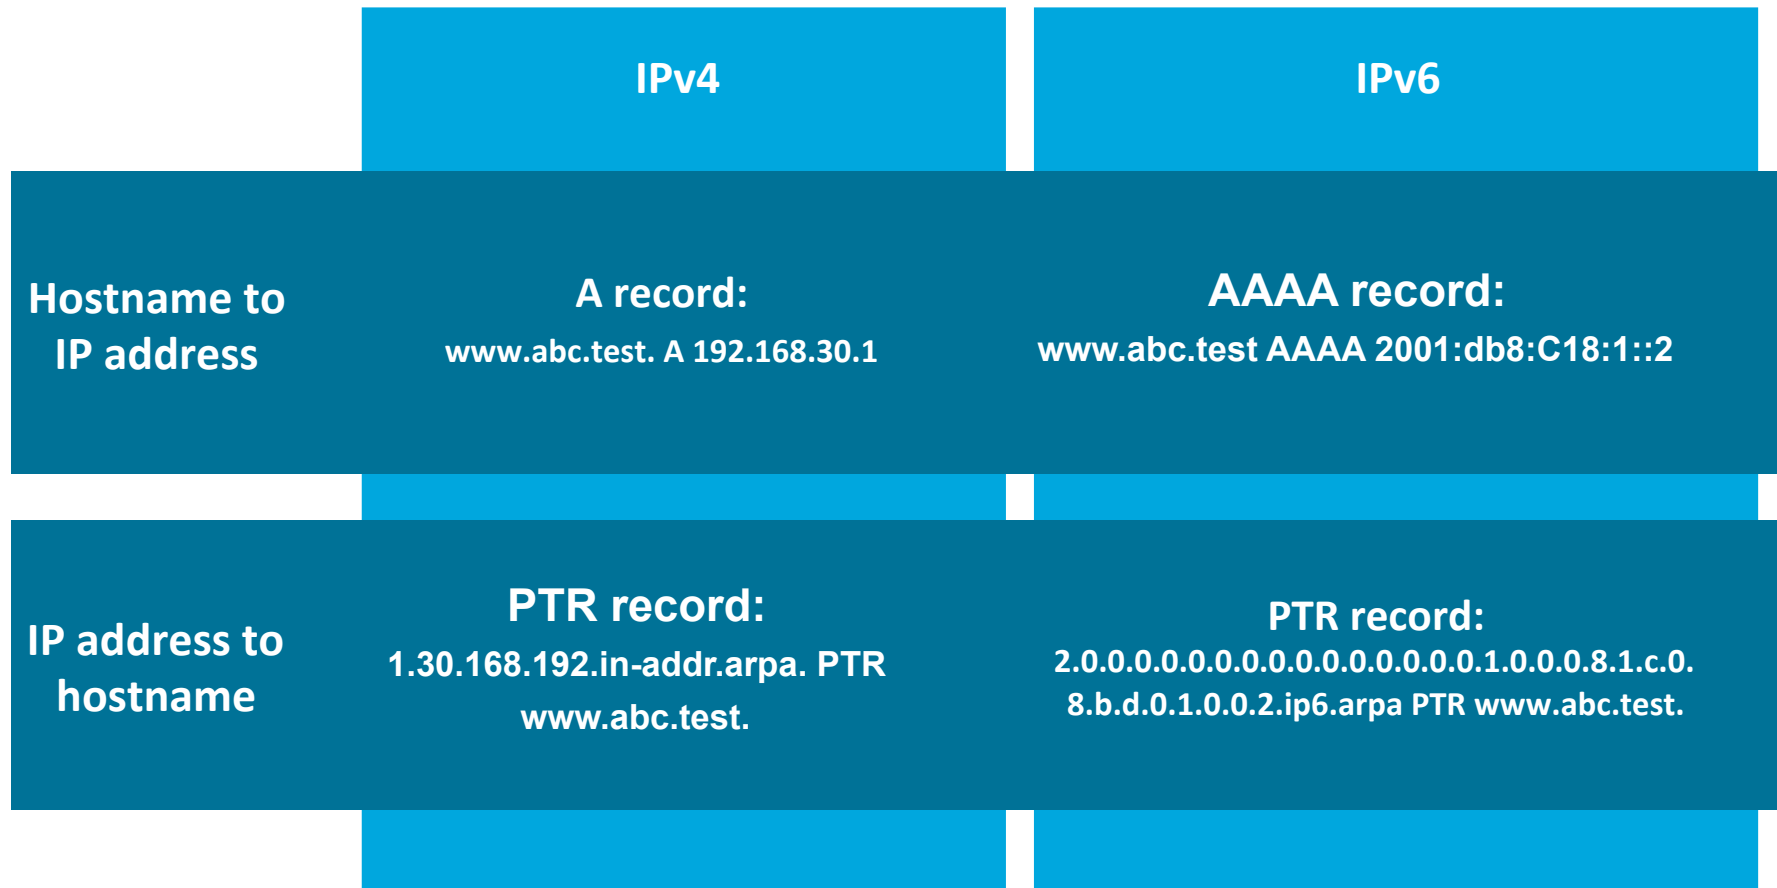
FF01:0:0:0:0:0:0:1 => FF01::1

0:0:0:0:0:0:0:1 => ::1

0:0:0:0:0:0:0:0 => ::

IPv6 and DNS

Actual DNS requests are transported over IPv4 or IPv6 and is not related to the request content.



IPv6 Readiness: Devices



- Hosts are ready
 - Since Windows Vista: enabled by default, disabling it = no more support from Microsoft
 - Mac OS X, iOS, Android, Linux, */BSD: enabled by default
- Network devices (routers, switches, phones, ...) are also ready
 - Some minor performance drop
 - IPv6 routing protocols are identical to IPv4: OSPF, BGP, ...
 - Usually IPv6 is for free
 - Just beware of FIB/RIB size which can double with IPv6
 - Low cost residential CPE has appeared in 2011
- → add IPv6-support as a requirement for all your RFP
 - Use IPv6 Forum 'IPv6-ready logo'
 - If public sector, copy the text from RIPE 554



IPv6 Readiness: Applications

- Microsoft: Active Directory, Exchange, Filesharing, SQL server, SharePoint...
- Public domain: apache, PHP, bind named, ...
- Commercial applications: Infoblox, ...
- Your own application????
 - Not too complex to port application to dual-stack, some changes in socket library
 - But it will require extensive/expensive testing...

Innocent W2K3 -to- W2K8 Upgrade

Windows 2003

```
C:\>ping svr-01

Pinging svr-01.example.com [10.121.12.25] with 32 bytes of
data:
Reply from 10.121.12.25: bytes=32 time<1ms TTL=128
Reply from 10.121.12.25: bytes=32 time<1ms TTL=128
Reply from 10.121.12.25: bytes=32 time<1ms TTL=128
Reply from 10.121.12.25: bytes=32 time<1ms TTL=128
```

Upgraded Host to Windows 2008

```
C:\>ping svr-01

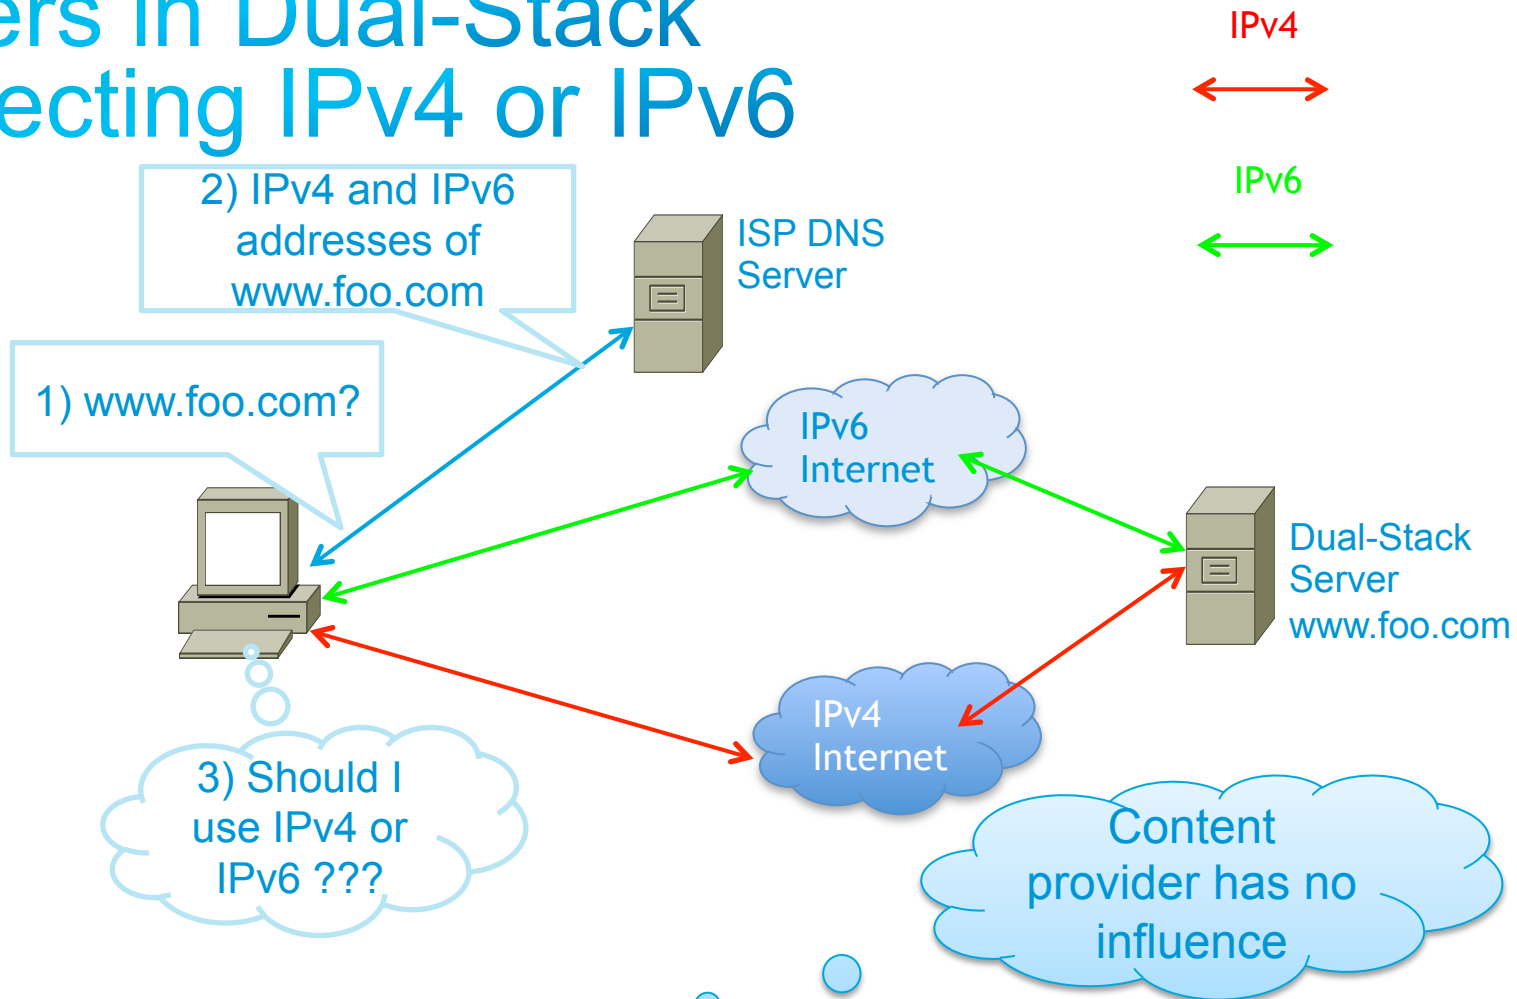
Pinging svr-01 [fe80::c4e2:f21d:d2b3:8463%15] with 32 bytes
of data:
Reply from fe80::c4e2:f21d:d2b3:8463%15: time<1ms
Reply from fe80::c4e2:f21d:d2b3:8463%15: time<1ms
Reply from fe80::c4e2:f21d:d2b3:8463%15: time<1ms
Reply from fe80::c4e2:f21d:d2b3:8463%15: time<1ms
```

ALL recent OS have IPv6 enabled by default and prefer it...

=> Enable IPv6 host security and IPv6 IPS

IPv4-IPv6 Coexistence

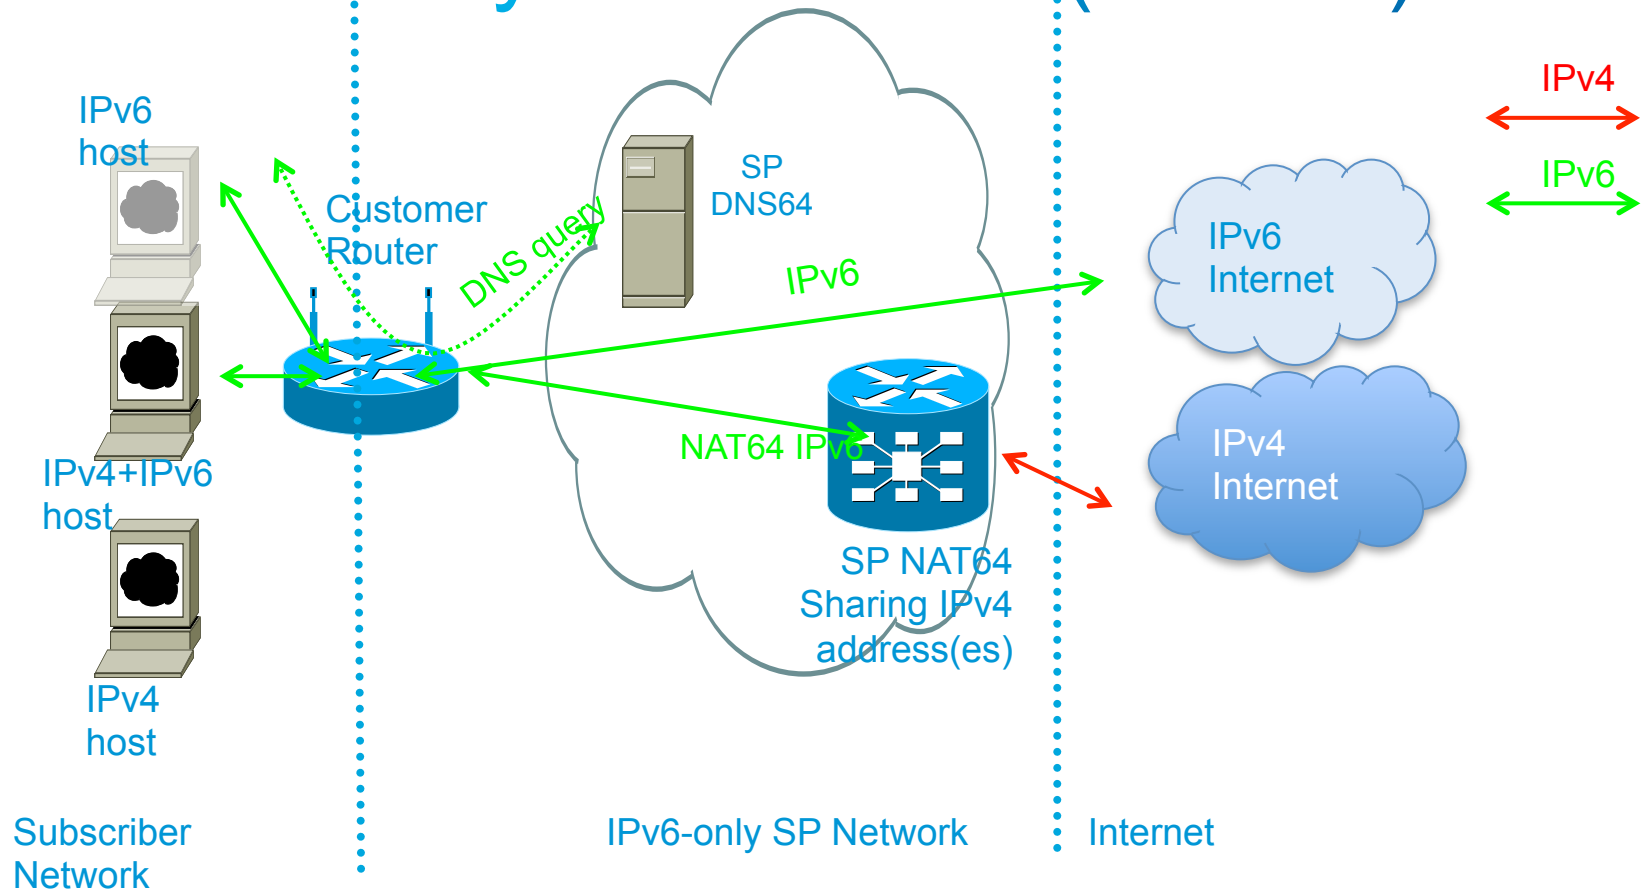
Users in Dual-Stack Selecting IPv4 or IPv6



Decision by the USER/INITIATOR:

- RFC 6555: Happy Eyeball, try both and keep the fastest
- RFC 6724: local policy, usually IPv6 is preferred

Address Family Translation (NAT64)

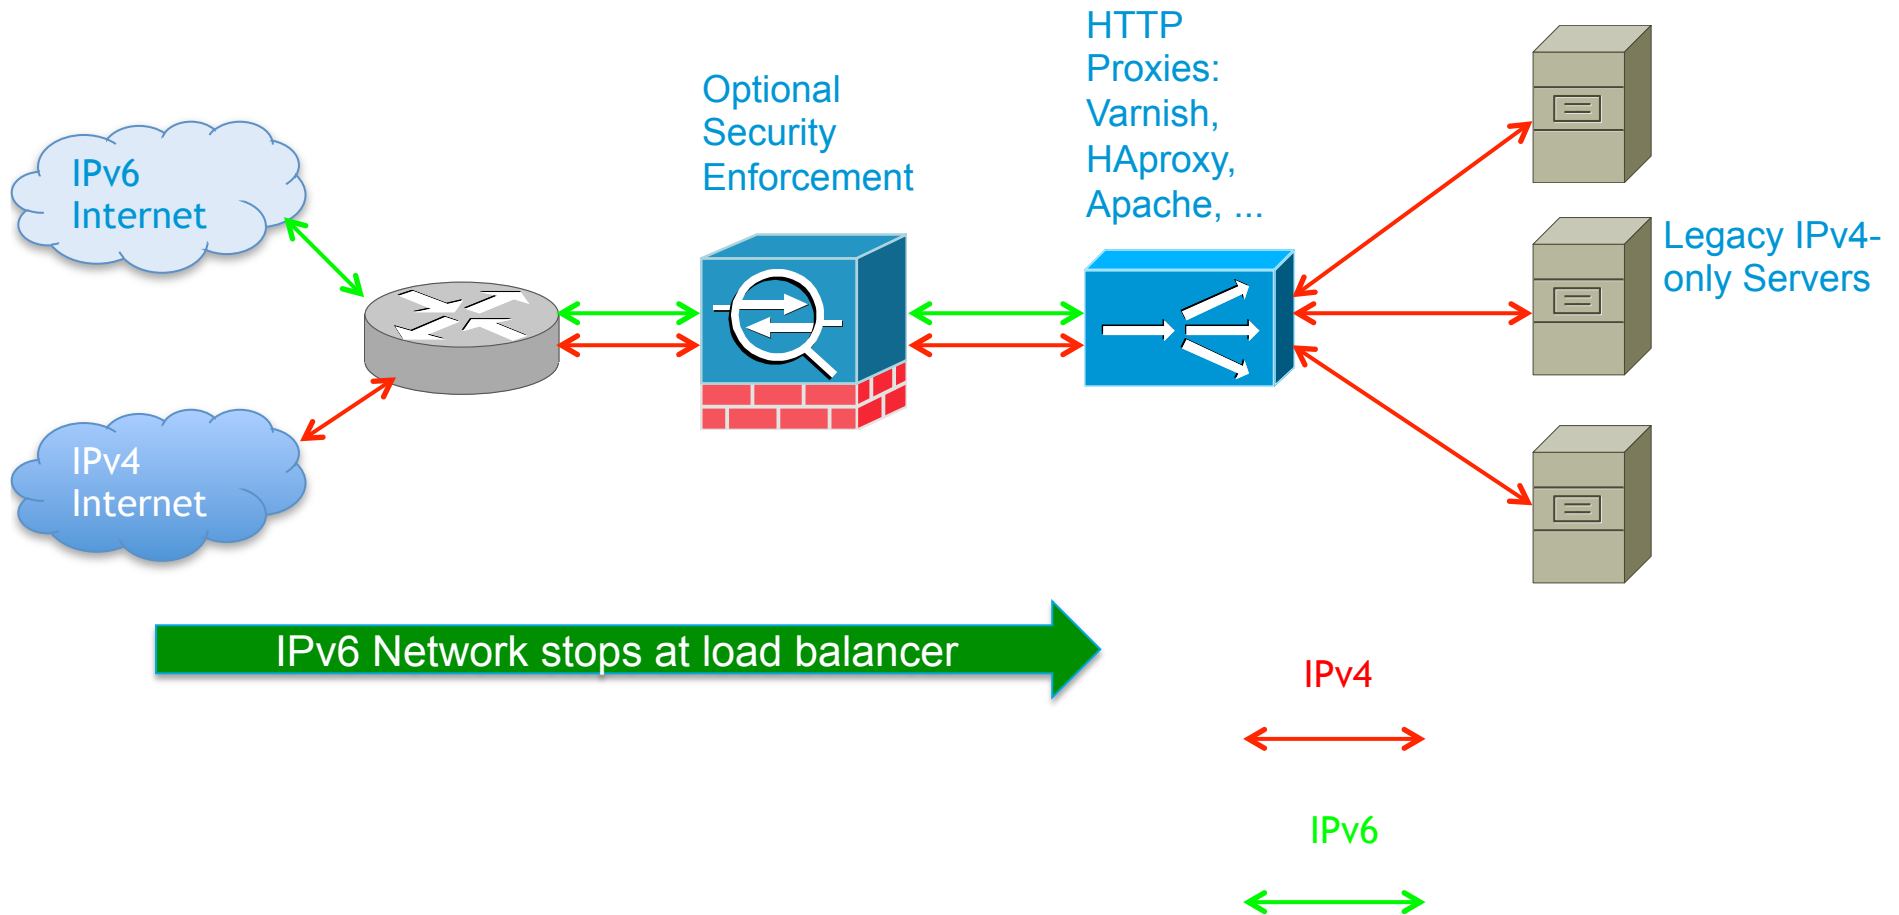


- Service Provider deploys IPv6-only infrastructure:

Only IPv6 is available to the consumer

IPv4 Internet available via Address Family Translation on SP NAT device

Dual-Stack Internet Edge with HTTP Proxy



Impact on Applications

Impact of IPv6 on Applications...

- Is the software package available?
- What about the network part?
- No more NAT!
- What about the logs and DB?
- What about the API?

Socket library

RFC 6724

IPv6-Enabled Applications

- Most open-source packages have been IPv6 enabled for several years

BIND `listen-on-v6 { any; };`

SSH (including WinSCP, Putty)

On by default, else `ListenAddress ::`

Apache

Sendmail, Qmail,

Postfix `inet_protocols=all`

VLC, Mozilla

- Specific IPv6 open-source

Dibbler: DHCP server

Wide-DHCP

Linux-DHCP

Ip6tables

radvd

IPv6-only Networks?

- Mostly experimental, such as FOSDEM or CERNET or ...
- T-Mobile (US) 4G mobile network is IPv6 only
- IETF has a SUNSET4 working group for 2 years now
- 6th June 2014, a day without IPv4

<https://www.facebook.com/groups/ipv6.only.day/> 6th June 2014

<https://secure.avaaz.org/en/petition/>

[IT_professionals_worldwide_Turn_off_the_Legacy_IP_Protocol_IPv4_for_the_day_of_06th_June_2014](#)

- And BSD 10.0 without IPv4 ☺

<https://www.zabbadoz.net/users/bz/blog/20140131-e9e7b6ead61ee2fd3e7337576663b128.html>

The Network Part of IPv6 Applications

Path MTU Discovery

- No fragmentation by routers
 - Path MTU Discovery **MUST** work
 - Free flow of ICMPv6 must be guaranteed
 - All routers must be able to send ICMP
 - Even if using link-local address on infrastructure
- Let's get real: it does not in 2011
 - Set the MSS to 1220 (minimum MTU of 1280 – 40 – 20)**
 - Or *setsockopt()* and IPV6_MTU

No more NAT

- With IPv6, NAT is no more required
- Real peer-to-peer applications are now trivial
 - No need for a relay 'in the cloud'
- Additional benefits, no need for periodic keep-alive packets (used to keep the translation state in stateful NAT devices)
 - Saving mobile battery 😊

IPv6 and Back End Databases

IPv6 Addresses are Larger...

- User Interface

Needs to accept wider data

Needs to accept : in addresses

Multiple ways to represent the same address

1. Use lower case (uppercase is reserved for MAC addresses)
2. Canonicalization with `inet_pton()` & `inet_ntop()`

Needs to redesign GUI to fit field in the screen ☹

- Logging the remote address

Needs to support 2 address families

Needs to store 39 characters instead of 15

- Audit

Need to adapt to IPv6

Multiple IPv6 addresses per end host

Multiple ways to write an IPv6 address

Privacy extension ...
Changing IPv6 address

The socket() API in C or C++

- Mostly unchanged, but should use some new functions
 - `inet_ntop, inet_pton`
 - `getaddrinfo, getnameinfo`
 - `struct sockaddr_storage`

- Either one socket

AF_INET6 in all calls

IPv4 address in compatibility mode (::ffff:192.0.0.2)

S = socket(AF_INET6, ...)

- Or two sockets (one IPv4 & one IPv6) for specific options

Happy eye-ball: clients open two connections and use the 'better' one

The IPv6 Support in Other Language

- Most of the scripting languages Python, Perl, PHP, ...

IPv6 support built-in

Scripts with FQDN simply work 😊

Beware of logging remote IP* addresses

Some OS (Windows, Mac OS/X) also have Happy Eyeballs

RFC 6724: Source Address Selection (SAS)

- Policy
 - Several local host addresses
 - Several remote server addresses
 - Which one to use?
- Scope is important in RFC 6724
 - Link-local vs. global

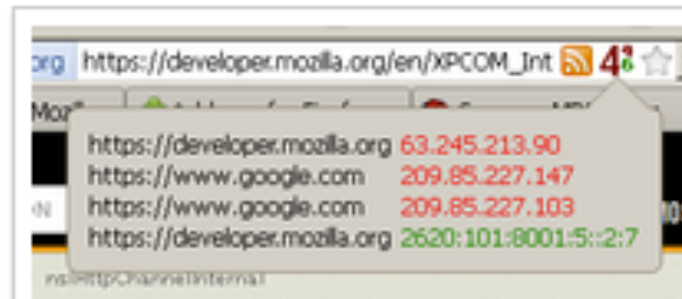
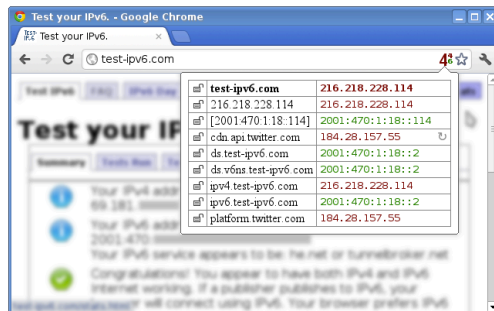
Prefix	Precedence	Label
::1/128	50	0
::/0	40	1
::ffff:0:0/96	35	4
2002::/16	30	2
2001::/32	5	5
fc00::/7	3	13
::/96	1	3
fec0::/10	1	11
3ffe::/16	1	12

Is IPv6 Used? How to Measure ?



How do you know you are using IPv6 ?

- test-ipv6.com
- Browser (Chrome & Firefox)
IPvFoo & IPvFox



- iPhone/iPad

Ipv6config
IPv6toolkit



or IPView



- Android

IPv6 & more



or IPv6 config



1000 EUR Question

How to monitor deployment?

- Looking at future?

Interviews

Monitoring 'precursors' (IPv6 prefixes, devices, ...)

- Looking at current state of the Internet

Open metrics => open results

Worldwide view

IPv6-enabled ISP Subscribers

- Measure the web traffic

Using a small web bug (1x1 pixel transparent GIF)

Used by Google, and others

Can precisely measure IPv4, IPv6 and dual-stack

E.g. <http://www.vyncke.org/countv6/>

Web servers log files

Easier but cannot check IPv4, IPv6

Google Ads

Used by APnic

All ads are in flash and flash can try to download IPv4, IPv6, dual-stack and report

- Participate in dual-stack peer-to-peer networks such as BitTorrent

More on Web bug

- By inserting a 1x1 pixel transparent image (or IFRAME)

Or even better 3 images:

IPv4-only

IPv6-only

Dual-stack then check whether IPv4 or IPv6 was preferred


- IPv6 brokenness (old 6to4 CPE mainly): can reach IPv4-only but not dual-stack (because IPv6 was preferred but broken), cfr Tore Anderson's experiment
- IPv6 non-managed tunnels (Sixxs, Gogonet, Hurricane Electric) because the AS is different between IPv4 and IPv6
- Can generate a lot of data and CPU load => sampling on large sites

Worldwide IPv6 Users

The “mother” of deployment measures



IPv6 Google Users Evolution



<http://www.vyncke.org/ipv6status/compare.php?metric=p&countries=de,nl,ch,lu,us,be>

Monitoring IPv6 Web Content

- Simple and easy
- Try to connect to www.example.org over IPv6
 - Make a AAAA DNS request
 - Then try to connect to it by TCP on port 80
- But, also try www6.example.org IPv6.example.org, ...
 - This is an indicator of 'IPv6 under test' or 'IPv6 pilot'
- Can be extended to measure email and DNS servers of a domain

www.alexa.com is your friend

- List of most visited web sites
- Freely available top-1-million
 - == Which are popular sites visited by the whole Internet

Then class by country top level domain: .be, .ch, .lu

Issues:

www.ice-watch.com is actually in Belgium

www.youtu.be is actually outside of Belgium


Estimation of IPv6-readiness of a country

- Paying list of visited web sites by country
 - == Which are popular sites actually visited by users of this country

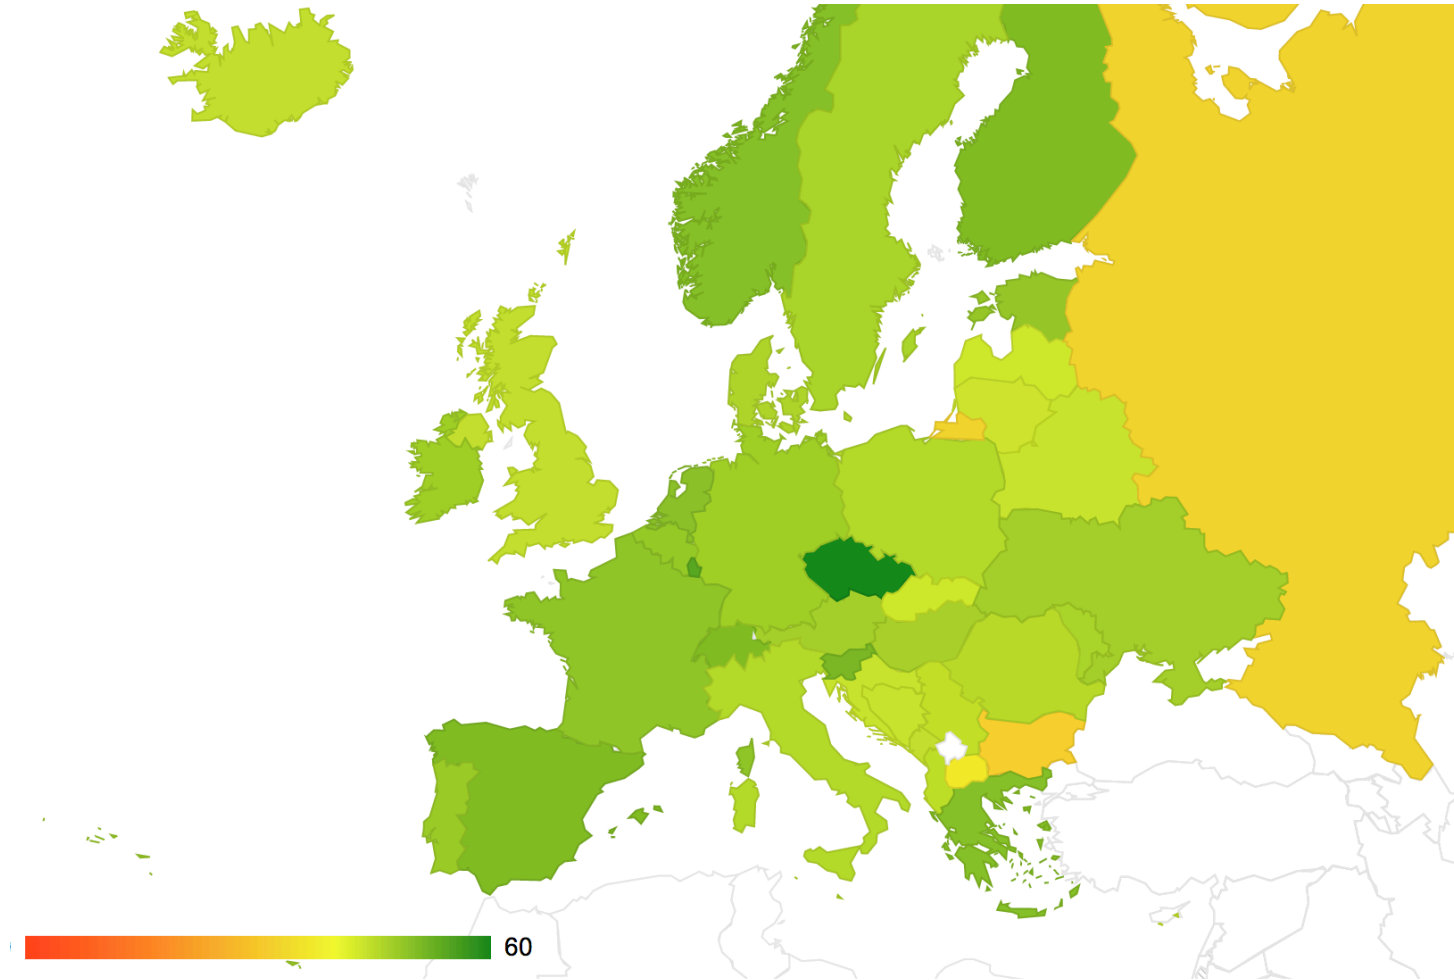
E.g. google.ch, facebook.com, google.com, live.com, ...

Assuming a long-tail distribution, can estimate the amount of IPv6 traffic IF all subscribers were dual-stack

Worldwide IPv6 Web Servers



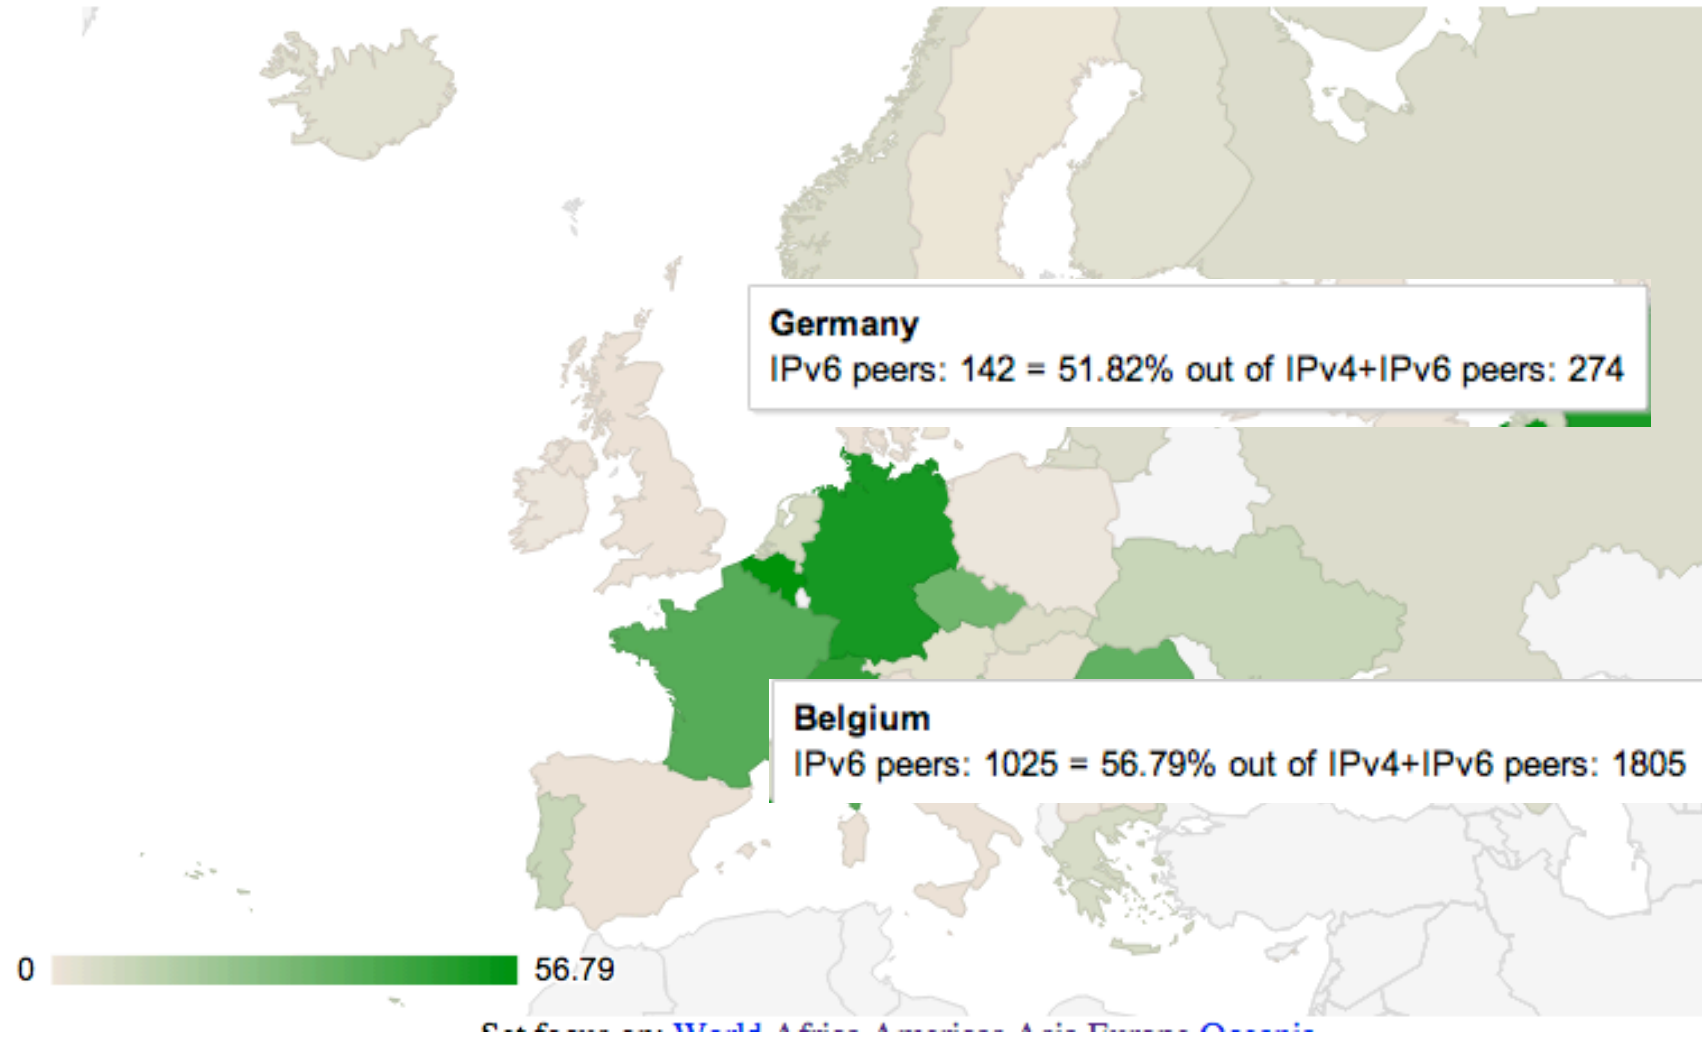
Web Servers Location



More on BitTorrent

- BitTorrent is a dual-stack peer-to-peer file exchange
 - Some trackers have dual-stack peer lists
 - Private Exchange is also dual-stack
 - Distributed Hash Table does not appear to be dual-stack
- draft-vyncke-ipv6-traffic-in-p2p-networks
- Of course, only pretend to have the file, do not upload/download it

European BitTorrent Jan-2013



Questions?

Thank you.

