The Buggers' Dilemma:

Eavesdroping and Traceback on the Internet

Geoff Huston

Chief Scientist

APNIC

Eavesdropping in the Telephony World

- Telephony is a network-centric architecture
- The network is aware of the address and location of attached endpoints
- Traffic is in the clear
- Interception and eavesdropping can be performed as a network operation

Eavesdropping in the Telephony World

internet

The internet

- Telephony is a network-centric architecture
- The network is aware of the address and location of attached endpoints
- Traffic is in the cl
- In
- Everything else is as a network operation

The internet is just a telephone network for computers. Everything else remains the same! Right?

Internet Eavesdropping - 80's-90's

- Modem tap to tape recorder to modem to transcript
- Switches with eavesdrop port
- Routers with eavesdrop port

- Data was in the clear, IP addresses were static, and eavesdropping was a case of performing a binary decode of the data stream

Internet Eavesdropping - 80's-90's

- Modem tap to tape recorder to modem to transcript
- Switches with eavesdrop port
- Routers with eavesdrop port

The Internet is just a telephone network for computers. Everything else remains the same! Right?

- [illegible], IP addresses were [illegible], and eavesdropping was a case of performing a binary decode of the data stream

Encryption becomes a service

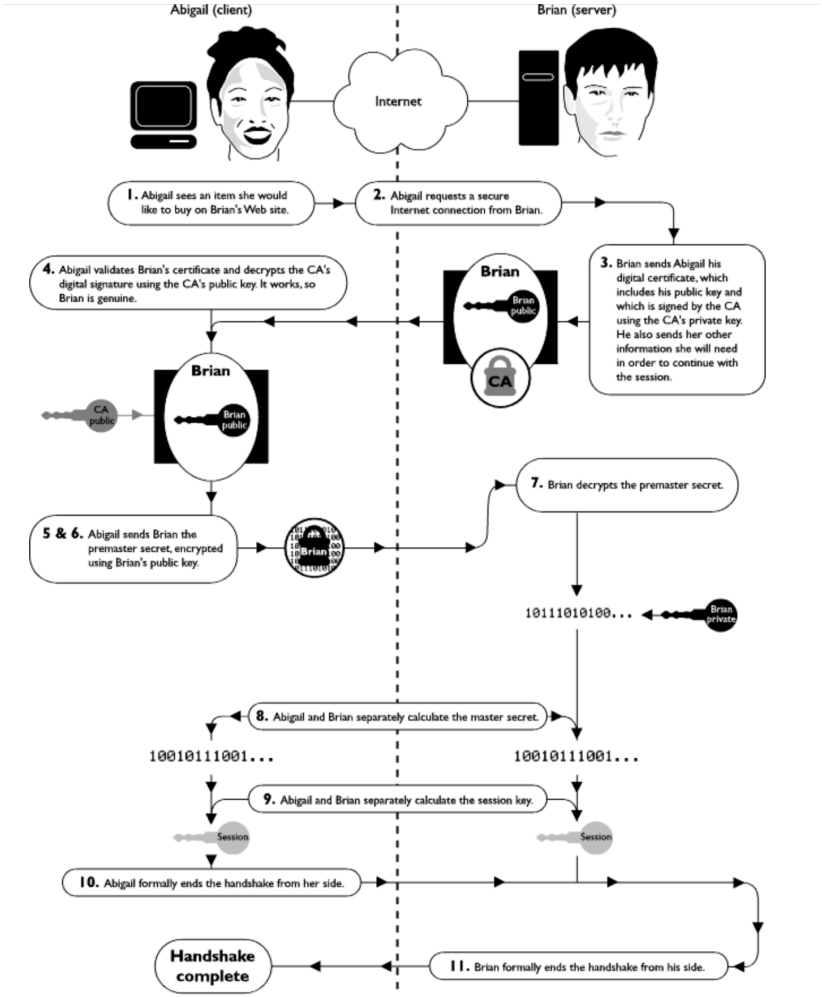

With the introduction of "Secure Sockets" in the mid-1990s it was feasible for services to encrypt their sessions

But this was not for everyone – it required money and tech knowledge

The Internet is a Belief System

Last Wednesday we had a CloudFlare Board meeting. We went over our plans for launching Universal SSL and how doing so may hurt our revenue given that SSL is one of the reasons people upgrade to a paid plan. But everyone on CloudFlare's Board was unanimous: even if it does hurt revenue in the short term, it's the right thing to do.

Brad Burnham, who is the partner at Union Square Ventures who led our last round of financing, reminded me during the meeting of the Joi Ito essay about how the Internet is a belief system. Inherent to Joi's point is that small groups of people, working together, can create great things. That, fundamentally, is the Internet.

The team behind Netscape first introduced SSL back in February 1995, originally intended to facilitate ecommerce online. As the Internet grew in importance, governments, ISPs, and hackers began to intercept, throttle, and censor traffic as it flowed across the network to serve their ends. In response, SSL's importance expanded beyond ecommerce to help ensure a free and open web. As Google and the IETF work on the next generation Internet protocols like SPDY and HTTP/2, it's no wonder encryption is at their heart. And so, in order for CloudFlare to fulfill its mission of helping build a better Internet, we knew one of the most important things we could do was enable Universal SSL for all our customers — even if they don't pay us.

Having cutting-edge encryption may not seem important to a small blog, but it is critical to advancing the encrypted-by-default future of the Internet. Every byte, however seemingly mundane, that flows encrypted across the Internet makes it more difficult for those who wish to intercept, throttle, or censor the web. In other words, ensuring your personal blog is available over HTTPS makes it more likely that a human rights organization or social media service or independent journalist will be accessible around the world. Together we can do great things.

The Internet is a belief system. At CloudFlare, we're proud today that we're playing a part in helping advance that belief system. And, having proven that Universal SSL is possible at our scale, we hope many other organizations will follow in turning SSL on for all their customers and at no additional cost.

2014 - https://blog.cloudflare.com/introducing-universal-ssl/

The Internet is a Belief System

Last Wednesday we had a CloudFlare Board meeting. We went over our plans for launching Universal SSL and how doing so may hurt our revenue given that SSL is one of the reasons people upgrade to a paid plan. But everyone on CloudFlare's Board was unanimous: even if it does hurt revenue in the short term, it's the right thing to do.

Brad Burnham, who is the partner at Union Square Ventures who led our last round of financing, reminded me during the meeting of the Joi Ito essay about how the Internet is a belief system. Inherent to Joi's point is that small groups of people, working together, can create great things. That, fundamentally, is the Internet.

The team behind Netscape first introduced SSL back in February 1995, originally intended to facilitate ecommerce online. As the Internet grew in importance, governments, ISPs, and hackers began to intercept, throttle, and censor traffic as it flowed across the network to serve their ends. In response, SSL's importance expanded beyond ecommerce to help ensure a free and open web. As Google and the IETF work on the next generation Internet protocols like SPDY and HTTP/2 it's no wonder encryption is at their heart. And so, in order for CloudFlare to fulfill its mission of helping build a better Internet, we knew one of the most important things we could do was enable Universal SSL for all our customers — even if they don't pay us.

But over time what's expensive becomes cheap and Universally available

Having cutting-edge encryption may not seem important to a small blog, but it is critical to advancing the encrypted-by-default future of the Internet. Every byte, however seemingly mundane, that flows encrypted across the Internet makes it more difficult for those who wish to intercept, throttle, or censor the web. In other words, ensuring your personal blog is available over HTTPS makes it more likely that a human rights organization or social media service or independent journalist will be accessible around the world. Together we can do great things.

The Internet is a belief system. At CloudFlare, we're proud today that we're playing a part in helping advance that belief system. And, having proven that Universal SSL is possible at our scale, we hope many other organizations will follow in turning SSL on for all their customers and at no additional cost.

2014 - https://blog.cloudflare.com/introducing-universal-ssl/

Lets ALL Encrypt!

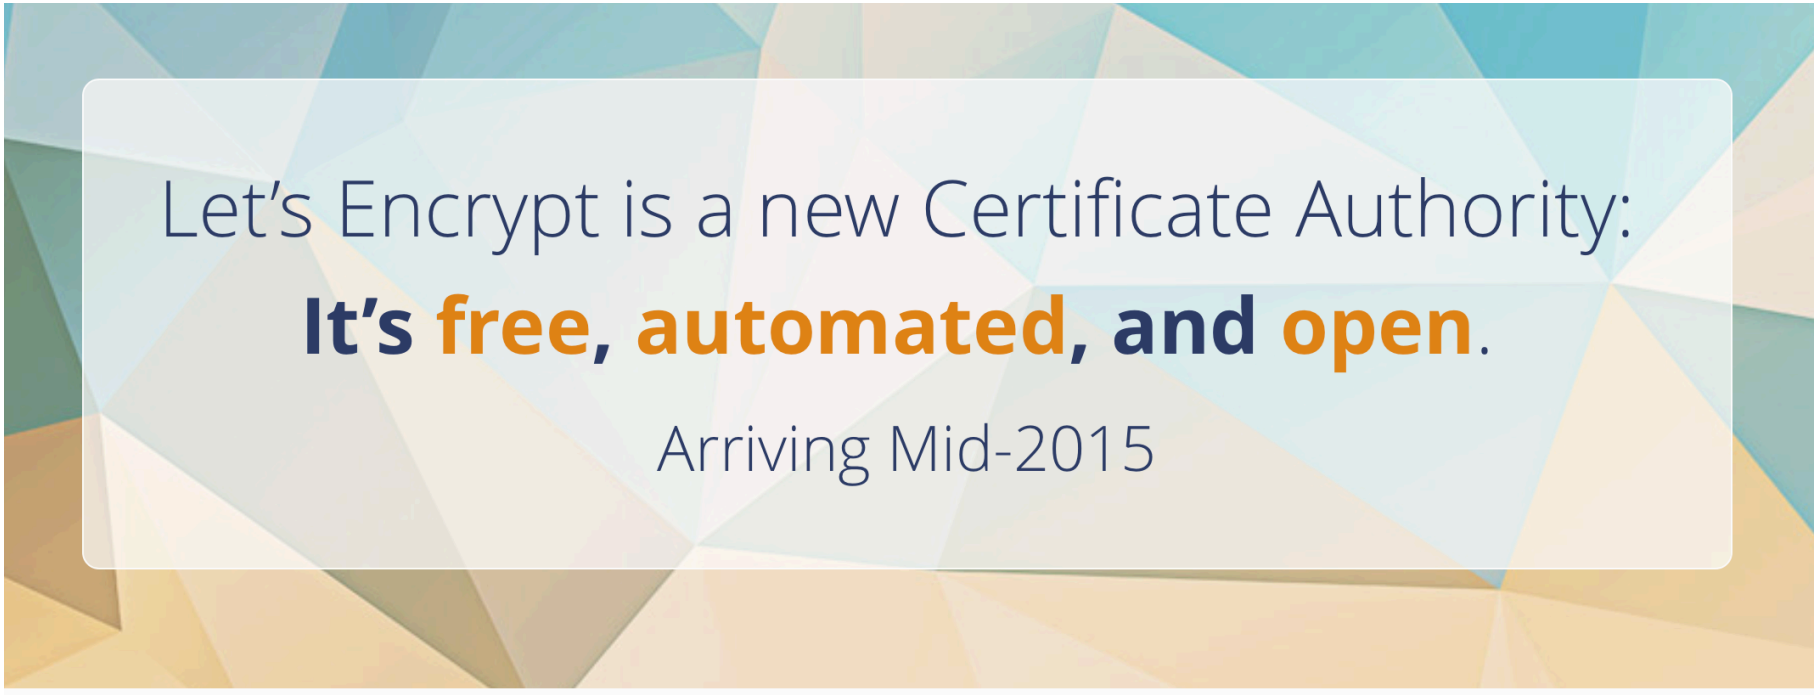

FROM OUR BLOG

Apr 23, 2015

Updated Draft ISRG CP and CPS

Today we're publishing an updated draft of our Certificate Policy (CP) and the first public draft of our Certification Practice Statement (CPS).
Read more

MAJOR SPONSORS

mozilla Akamai CISCO

IdenTrust part of HID Global AUTOMATTIC

Good Security is Relative

For traffic encryption for you and I the aim is to make it expensive for the eavesdropper

- So the compromise between efficiency and protective strength tends towards the adequate as distinct from the ideal

The aim of universal encryption is to increase the cost to the eavesdropper to the point where general surveillance is not affordable

Defense is expensive

The defender has to defend everything, the attacker only needs to exploit just one vulnerability…

Heartbleed

The Heartbleed Bug

The Heartbleed Bug is a serious vulnerability in the popular OpenSSL cryptographic software library. This weakness allows stealing the information protected, under normal conditions, by the SSL/TLS encryption used to secure the Internet. SSL/TLS provides communication security and privacy over the Internet for applications such as web, email, instant messaging (IM) and some virtual private networks (VPNs).

The Heartbleed bug allows anyone on the Internet to read the memory of the systems protected by the vulnerable versions of the OpenSSL software. This compromises the secret keys used to identify the service providers and to encrypt the traffic, the names and passwords of the users and the actual content. This allows attackers to eavesdrop on communications, steal data directly from the services and users and to impersonate services and users.

The bug that keeps on giving

CRYPTOME

Donate for the Cryptome archive of files from June 1996 to the present

---

26 April 2015

GoodCrypto Attacked

---

https://goodcrypto.com/news/2015/03/26/surveillance-system-used-for-censorship-in-europe/

Published here to resist censorship.

Surveillance system used for censorship in Europe

Censorship attack combines packet injection and Heartbleed

We all know there is censorship online. It happens in China. It happens to "terrorists". But we don't believe it will happen to us.

As Eben Moglen[1] and Kaspersky[2] have pointed out, companies developing crypto are prime targets no matter where they are. So you don't have to be a bad guy for the NSA to attack you. You just have to protect people from the NSA. Even protecting yourself is often enough. NSA prefers their victims to be defenseless.

Detection in the wild

In early 2015 people were still downloading our ISO file for GoodCrypto. But suddenly installations stopped.

After a lot of checking we noticed that the downloads got HTTP 200 result codes, but the lengths were all too short. This isn't supposed to happen. A 200 result means success. These weren't successful downloads, but the web logs said they were. Ordinary log checks didn't show the bug.

The bug that keeps on giving

CRYPTOME

Donate for the Cryptome archive of files from June 1996 to the present

MiTM attack in the UK using key compromise by exploiting Heartbleed vulnerabilities on the client's side and presumably applying the attack through an interception appoach such as the UK's "cleanfeed"

We a … terrorists". But we don't believe it will happen to us.

As E… Kaspersky[2] have pointed out, companies developing crypto are prime targets no matter where they are. So you don't have to be a bad guy for the NSA to attack you. You just have to protect people from the NSA. Even protecting yourself is often enough. NSA prefers their victims to be defenseless.

Detection in the wild

In early 2015 people were still downloading our ISO file for GoodCrypto. But suddenly installations stopped.

After a lot of checking we noticed that the downloads got HTTP 200 result codes, but the lengths were all too short. This isn't supposed to happen. A 200 result means success. These weren't successful downloads, but the web logs said they were. Ordinary log checks didn't show the bug.

Who's winning?

Pervasive security is a theme across much of the IETF's current technology work:

- DNS: Secure DNS, qname minimization, client-resolver opportunistic encryption, DANE
- Addresses: Address PKI, Secure routing
- Transport: Opportunistic session encryption

The true capabilities and budgets of the security agencies are not clearly known:

- But the greater the take up of encryption and secure infrastructure the greater the cost and effort of surveillance

Who's winning?

Pervasive security is a theme across much of the IETF current technology work:

- DNS: Secure DNS, qname minimisation, opportunistic encryption
- Addresses

The security agencies are not

- The greater the take up of encryption and secure infrastructure the greater the cost and effort of surveillance

We think we are winning – we're just not sure what "winning" means! "who "we" are, and

After the fact

Traceback and forensics in today's Internet

Traceback- Version 1

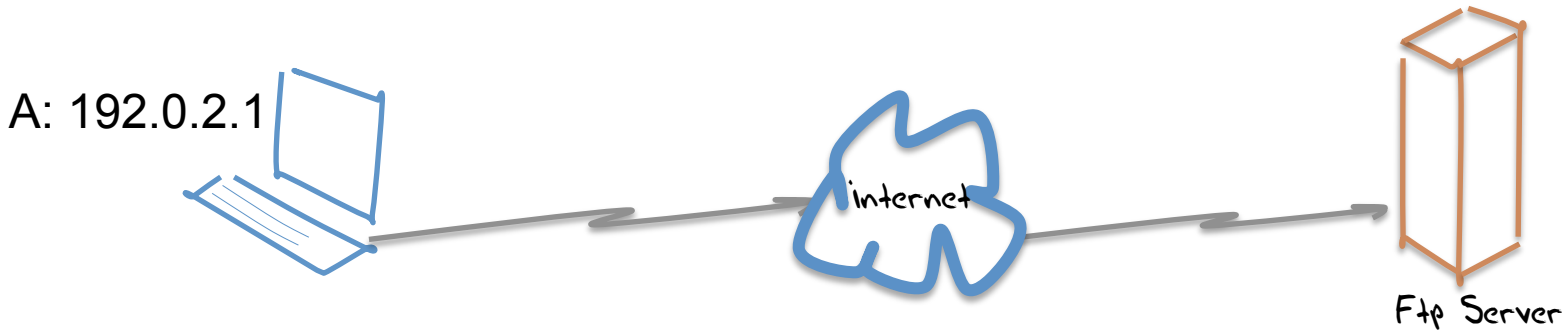

Lets start by looking waaaay back to the internet of the 1980's

Assumptions:

- Each end site used a stable IP address range
- Each address range was recorded in a registry, together with the end user data
- **Each end device was manually configured with a stable IP address**
- The networks uniformly route IP addresses

- Traceback is keyed from the IP address

Traceback - Version 1

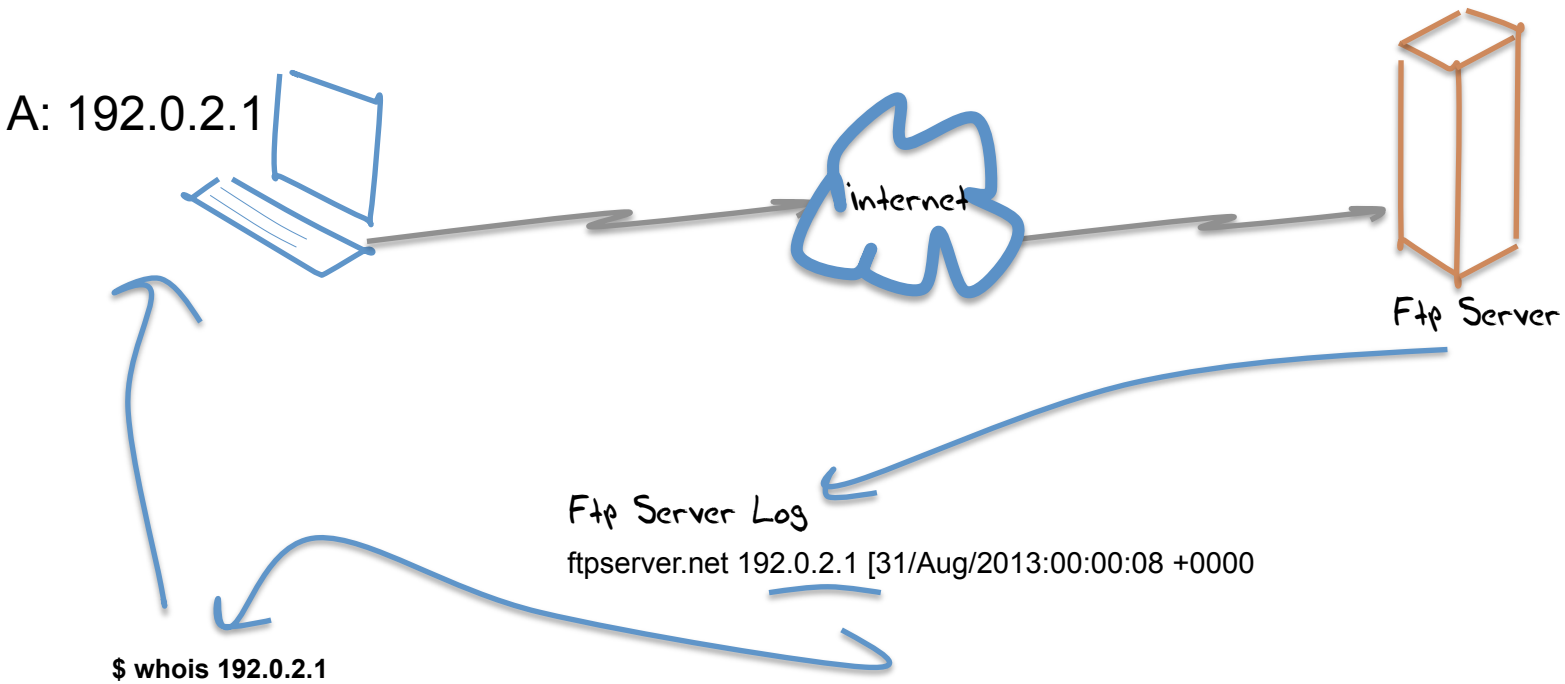

There was a rudimentary whois service and it listed all end users!

Assumptions:

- Each end site used a stable IP addr
- Each address range was recor
  together with the end use
- **Each end device w .gured with a stable IP ad**
- The netw .e IP addresses
- ed from the IP address

This model largely fell into disuse by the late 1990's it was replaced by a combination of provider-address blocks, dynamic addressing to end users (AAA tools) and CPE NATs

+ NATs

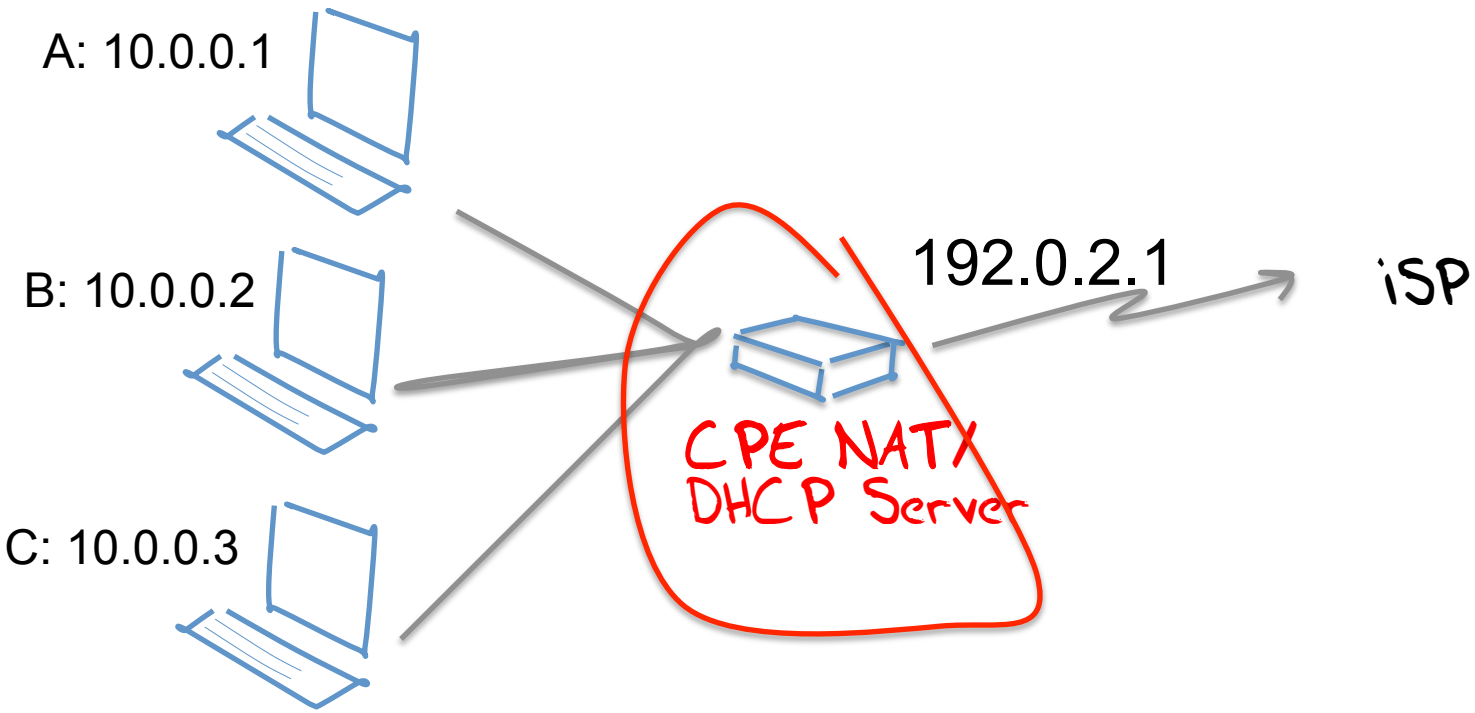

Traceback - Version 2

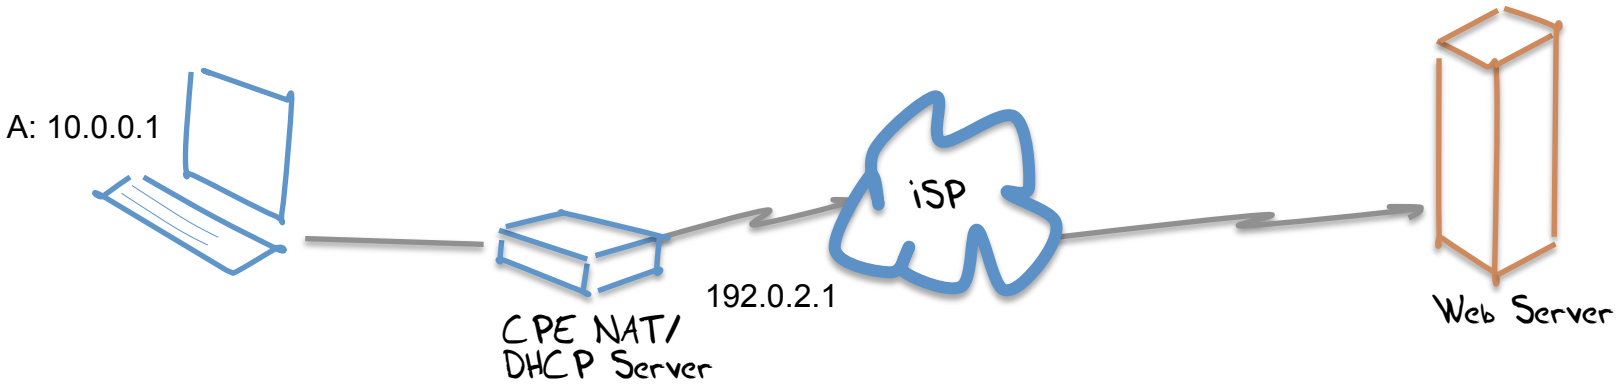

Traceback - Version 2

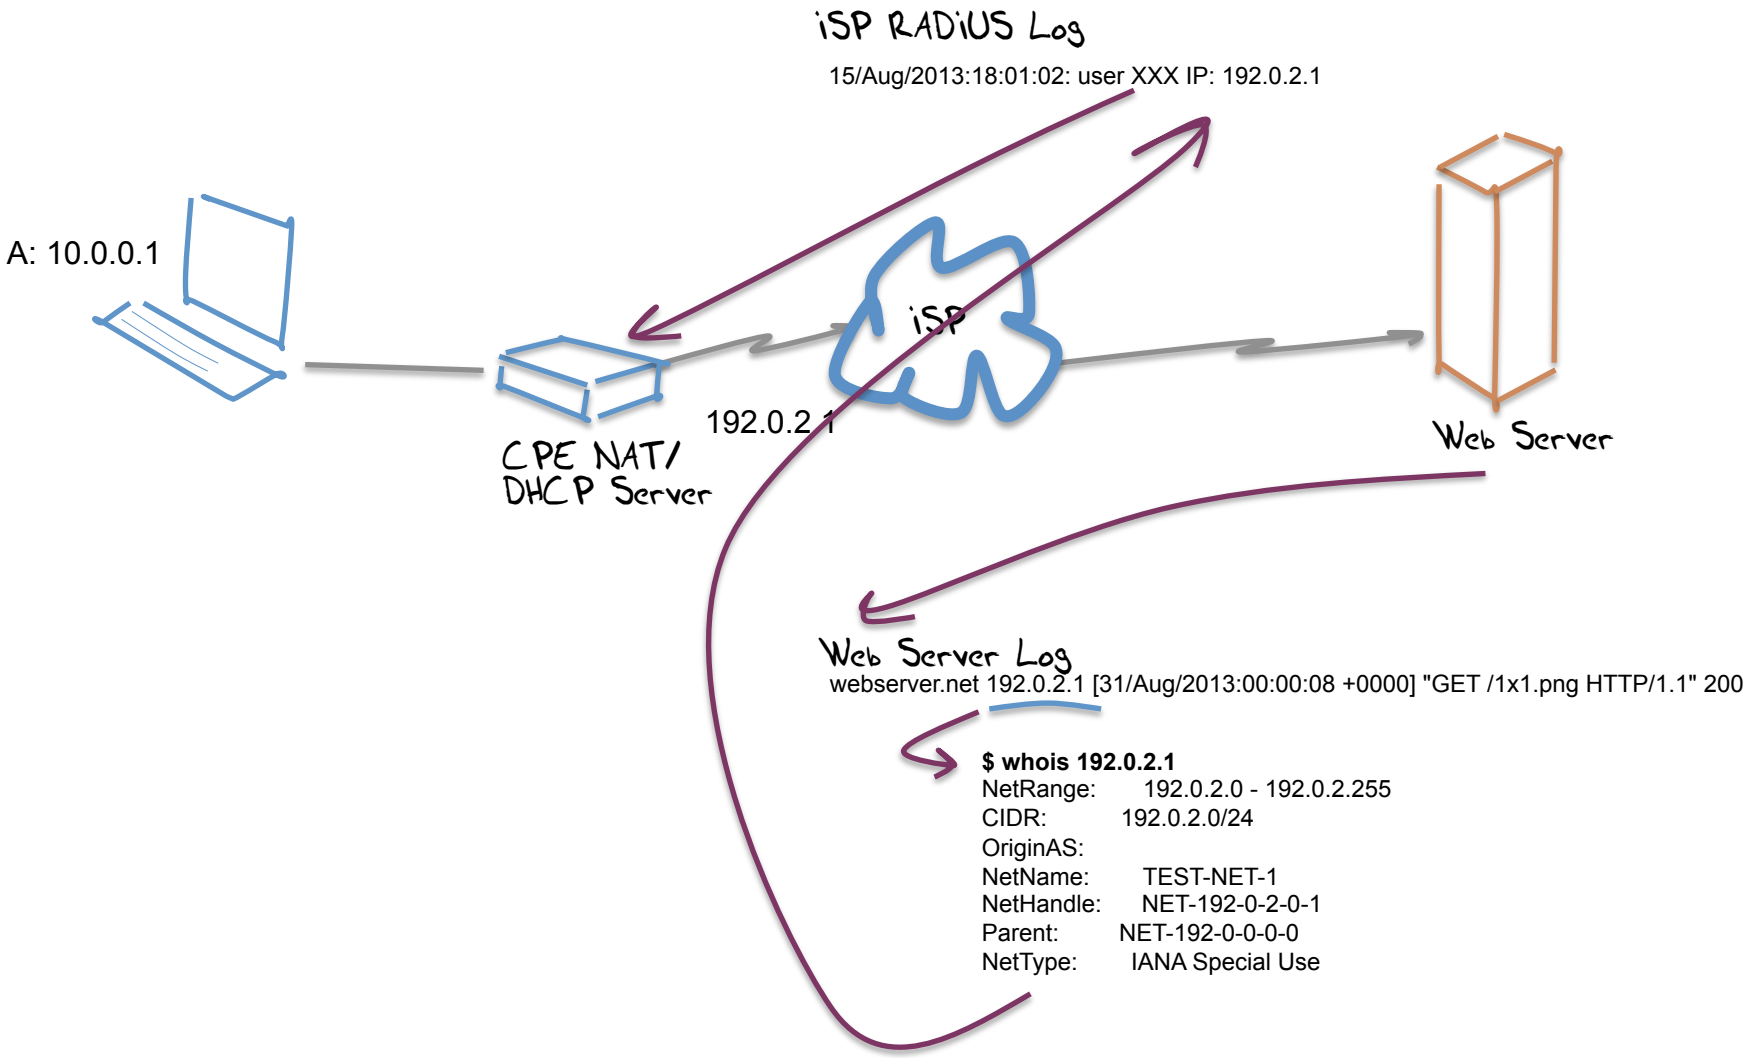

Assumptions

- The ISP operates an address pool
- Each end site is dynamically assigned a single IP address upon login (AAA)
- The single public address is shared by the private devices through a CPE NAT

- Traceback to an end site is keyed by an IP address and a date/time
- Network data gets you to the CPE NAT, but no further

Assumptions

- The ISP operates an address pool
- Each end site is dynamically assigned an IP address upon login
- The single address is shared by the private devices behind the CPE NAT
- Traceback to an end site is keyed by an IP address and a date/time
- Network logs get you to the CPE NAT, but no further

Individual devices are anonymous to the network.
All that is visible to the network is the single shared address

Why?

- Why are we sharing IP addresses between devices?
- Surely there was nothing wrong with allowing each connected device to use its own dedicated address

IETF Meeting - August 1990

Internet Growth (Continued):
Continued Internet Growth

Frank Solensky
Racal InterLan
solensky@interlan.com

- A preliminary analysis of data presented earlier in the conference projects the "size" of the Internet on several metrics, assuming continued exponential growth.
  - NIC Assigned Network Numbers
  - NIC "Connected" Status Nets
  - BBN's snapshots
  - NSFnet Policy Routing DataBases
- As was mentioned during the discussion period, a logistic curve would likely be a more realistic model: This will be the subject of further analysis. Note, however, that the limit that this approaches may turn out to be beyond the capacity of the class A-B-C numbering scheme

NIC

"Connected" IP Network Numbers

- Assigned Numbers RFC defines connected networks as connected to research and operational internet.
- Does not reflect whether the net is, in fact, entered in any routing table.

$y = \beta e^{\alpha t}$ where y = predicted number of nets, t = time (in months) since Jan. 1983

	Class A	Class B	Class C	Class A-B
β	12.069	24.412	887.879	3032.211
α	.012163	.040721	.011690	.013467
growth rate per yr.	15.613%	61.440%	14.497%	17.413%
y	125	16,382	2,097,150	49,147
$\hat{x}$	192.193	159.839	624.438	206.346
	Jan 6, 1999	Apr 26, 1996	May 14, 2038	Mar 27, 2000
r	.9293	.9870	.7942	.9548

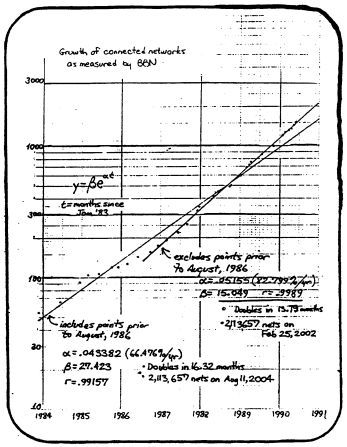

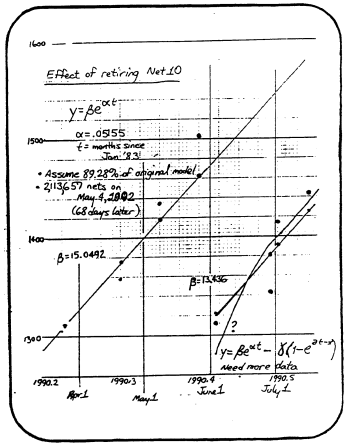

Assignment of IP Network Numbers

- Reflects organizations' desire for IP Address assignment; that is, to be listed in RFC-1166.
- Does not reflect "connectivity"

$y = \beta e^{\alpha t}$ where y = predicted number of nets, t = time (in months) since Jan '83

	Class A	Class B	Class C	Class A-B
β	11.883	21.446	1531.793	2899.462
α	.013175	.049411	.027187	.015587
growth rate per yr.	17.008%	78.389%	37.973%	20.394%
y	125	16,382	2,097,150	49,147
$\hat{x}$	178.605	134.35	265.64	181.58
	(Nov 19, 1997)	(Mar. 11, 1994)	(Feb. 18, 2005)	(Feb 17, 1998)
r	.9491	.9842	.9800	.9749

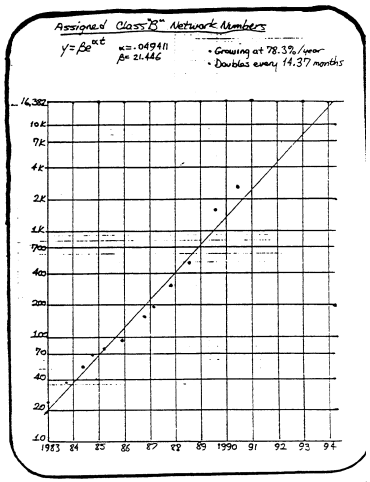

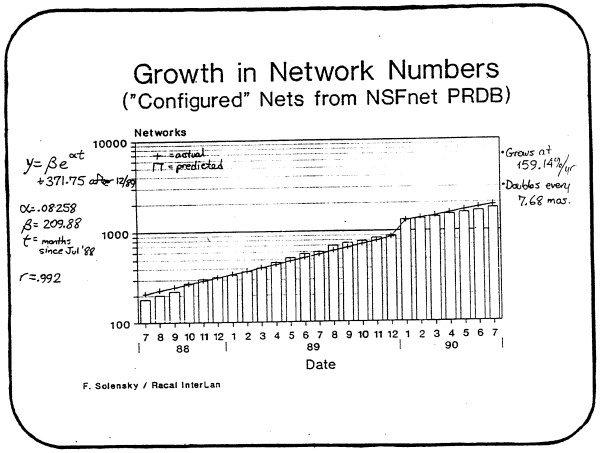

IETF Meeting - August 1990

Depletion Dates

• Assigned Class "B" network numbers	Mar. 11, 1994
• NIC "connected" Class B network numbers	Apr. 26, 1996
• NSFnet address space*	Oct. 19, 1997
• Assigned Class "A-B" network numbers	Feb 17, 1998
• NIC "connected" Class A-B network numbers	Mar. 27, 2000
• BBN snapshots*	May 4, 2002

* all types: may be earlier if network class address consumption is not equal.

IETF Meeting - August 1990

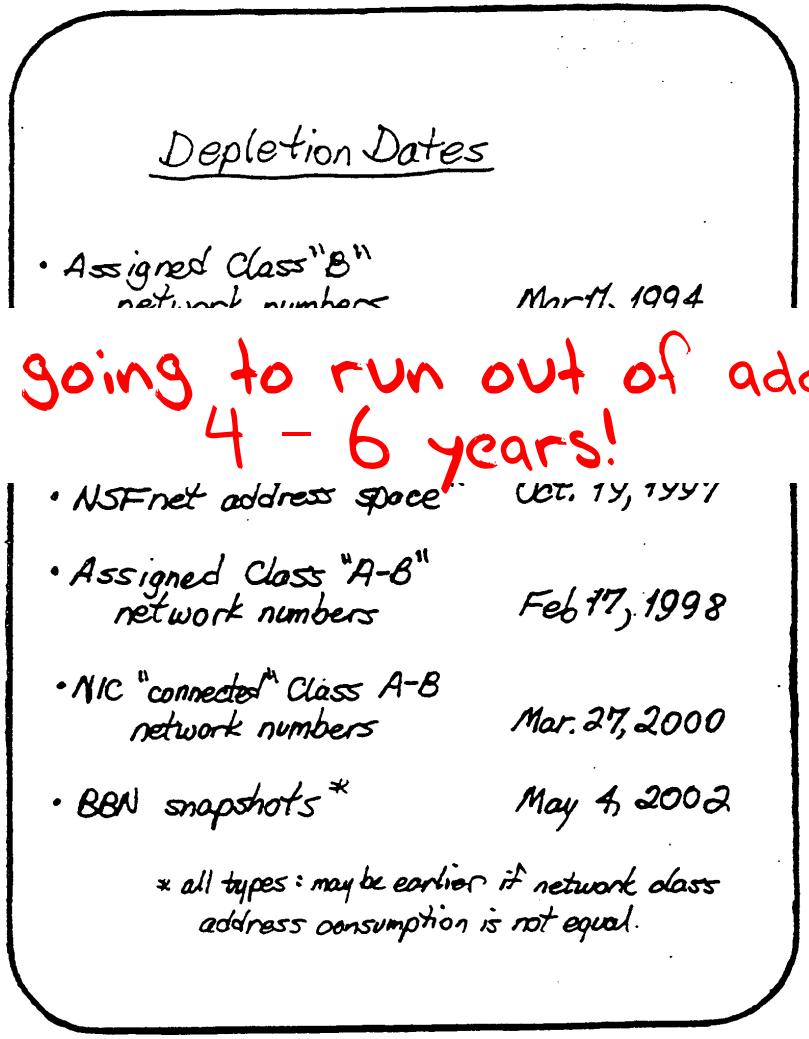

We were going to run out of addresses in 4 - 6 years!

The Response!

- The short term
  - Stop “wasting” addresses
- The long term
  - We need a new protocol

The Response!

- The short term
  - Stop "wasting" addresses

Change the routing protocols to support variable host/net boundaries in addressing

Share IP addresses behind Network Address Translators

- The long term
  - We need a new protocol

IPv6!

The Response!

- The short term
  - Stop "wasting" addresses
- The long term
  - We need a new protocol

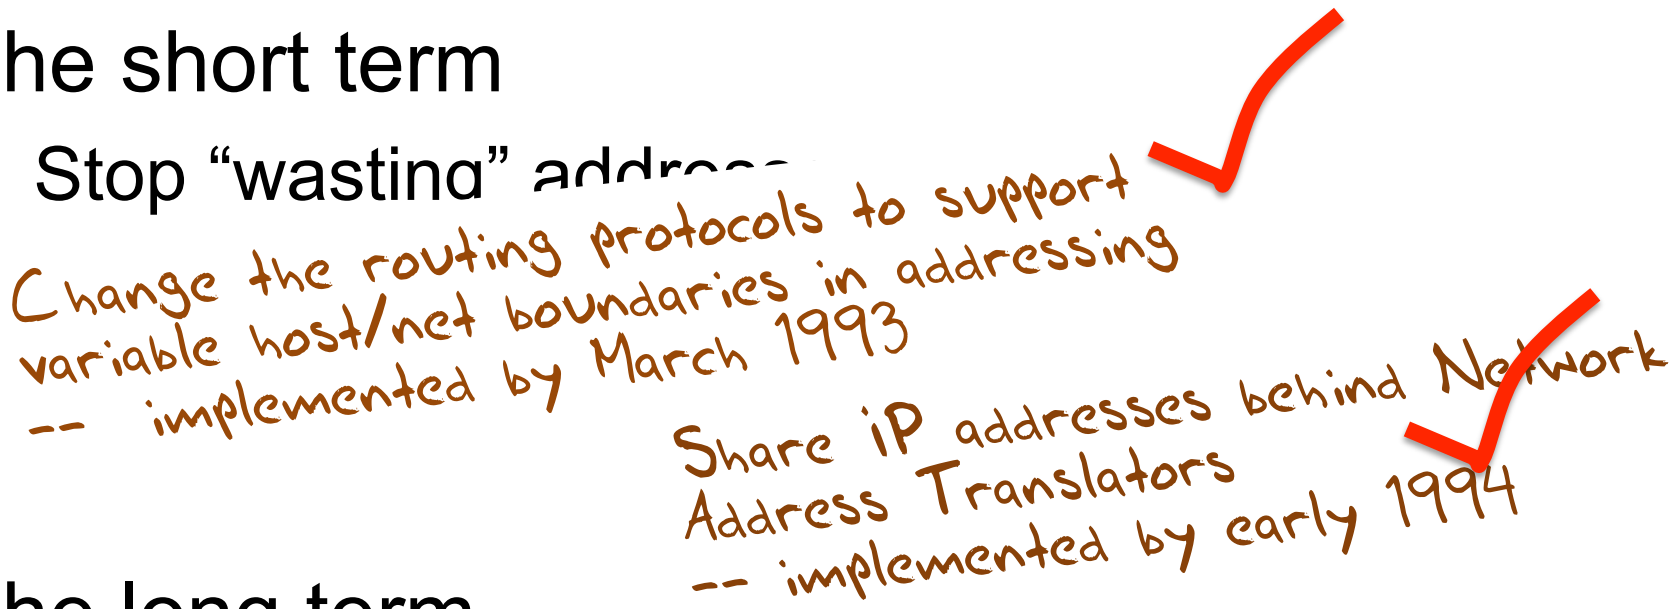

The IPv6 Transition Plan - as planned

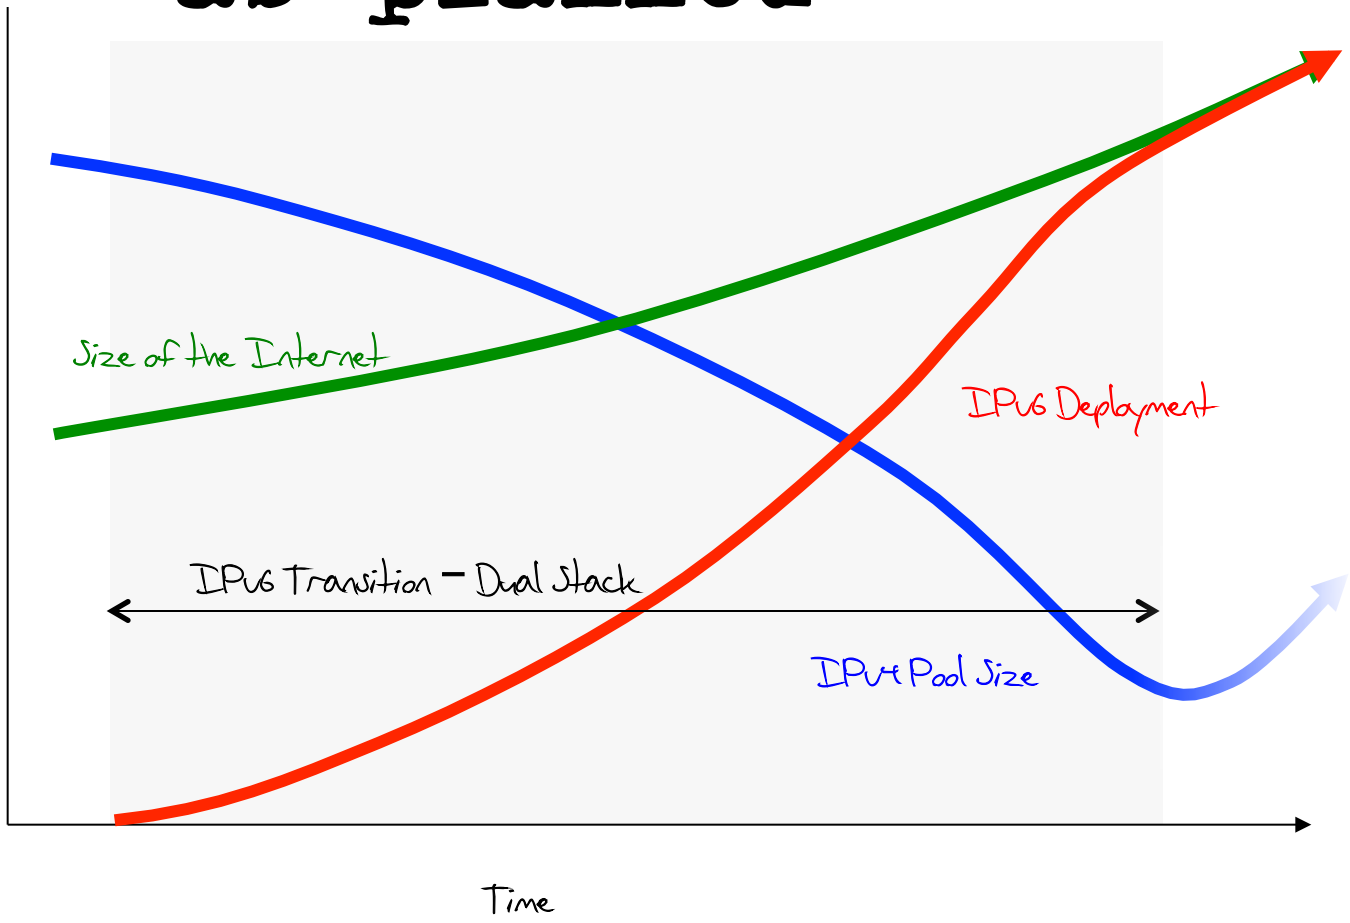

For this to work we have to start early and finish BEFORE IPv4 address pool exhaustion

The IPv6 Transition Plan - as implemented

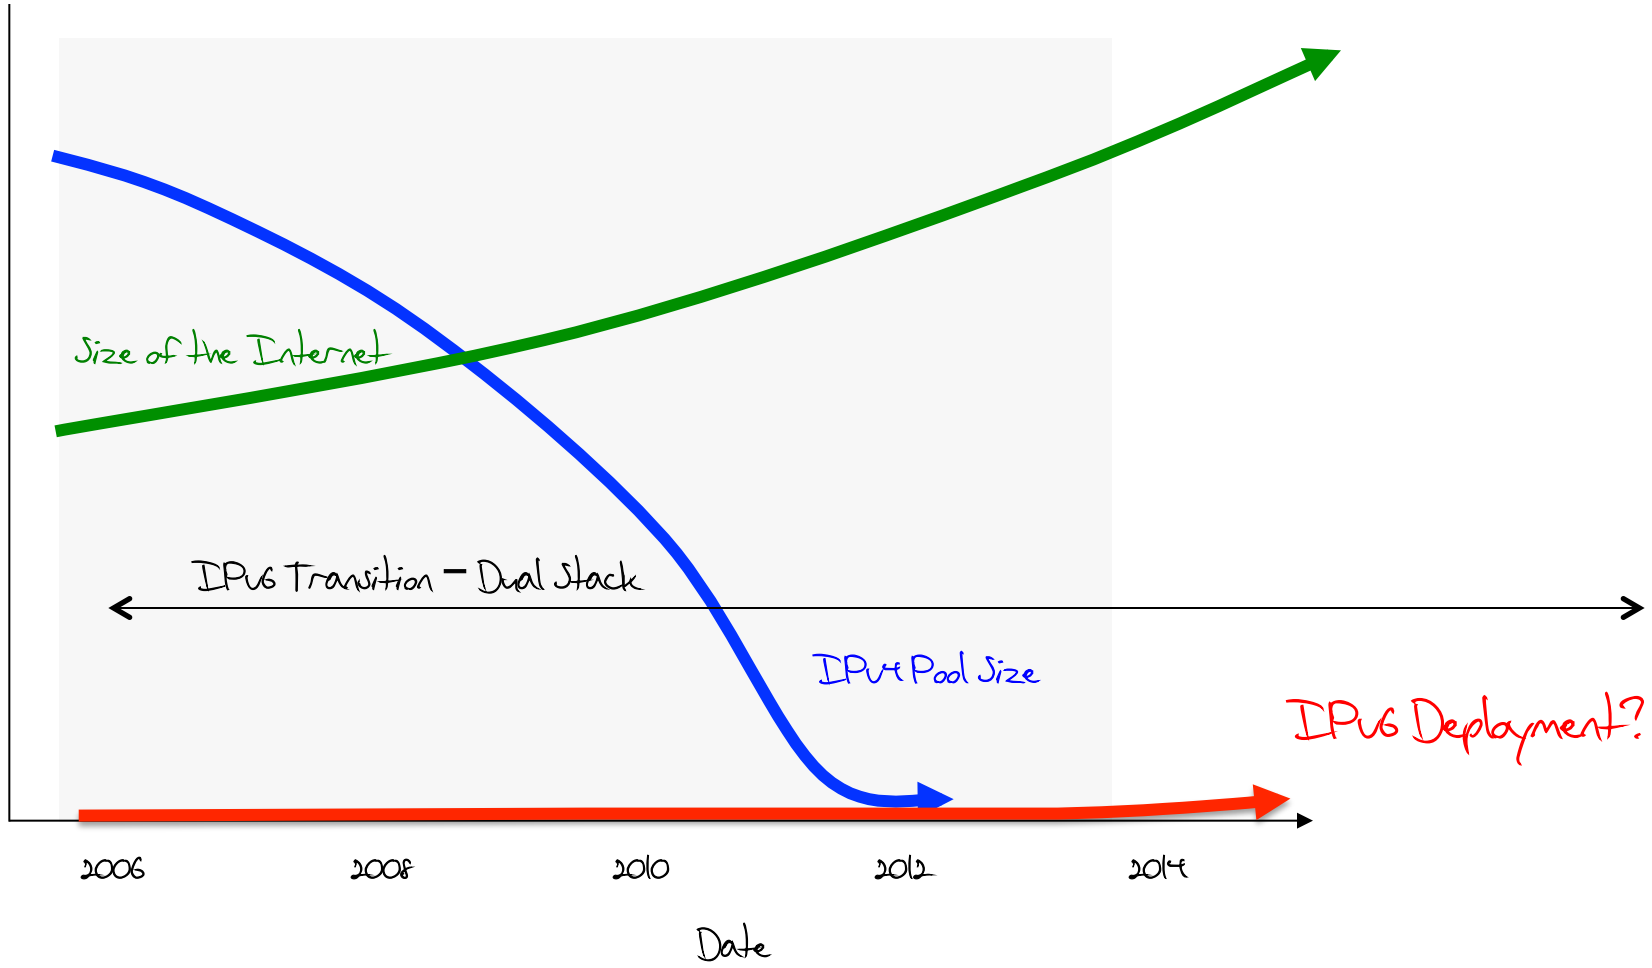

Where's IPv6 Today?

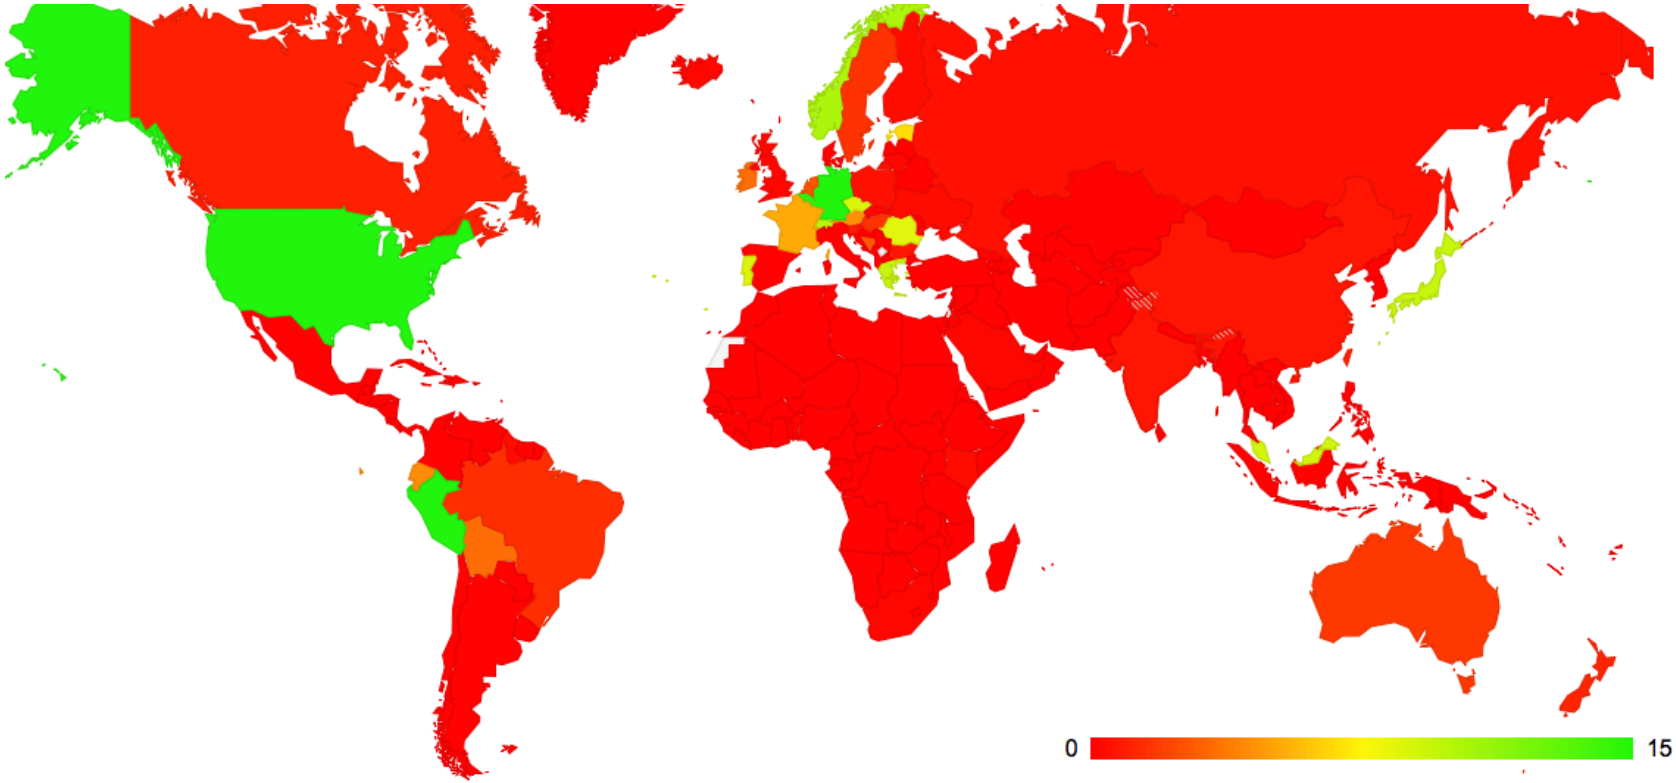

How much is IPv6 Today?

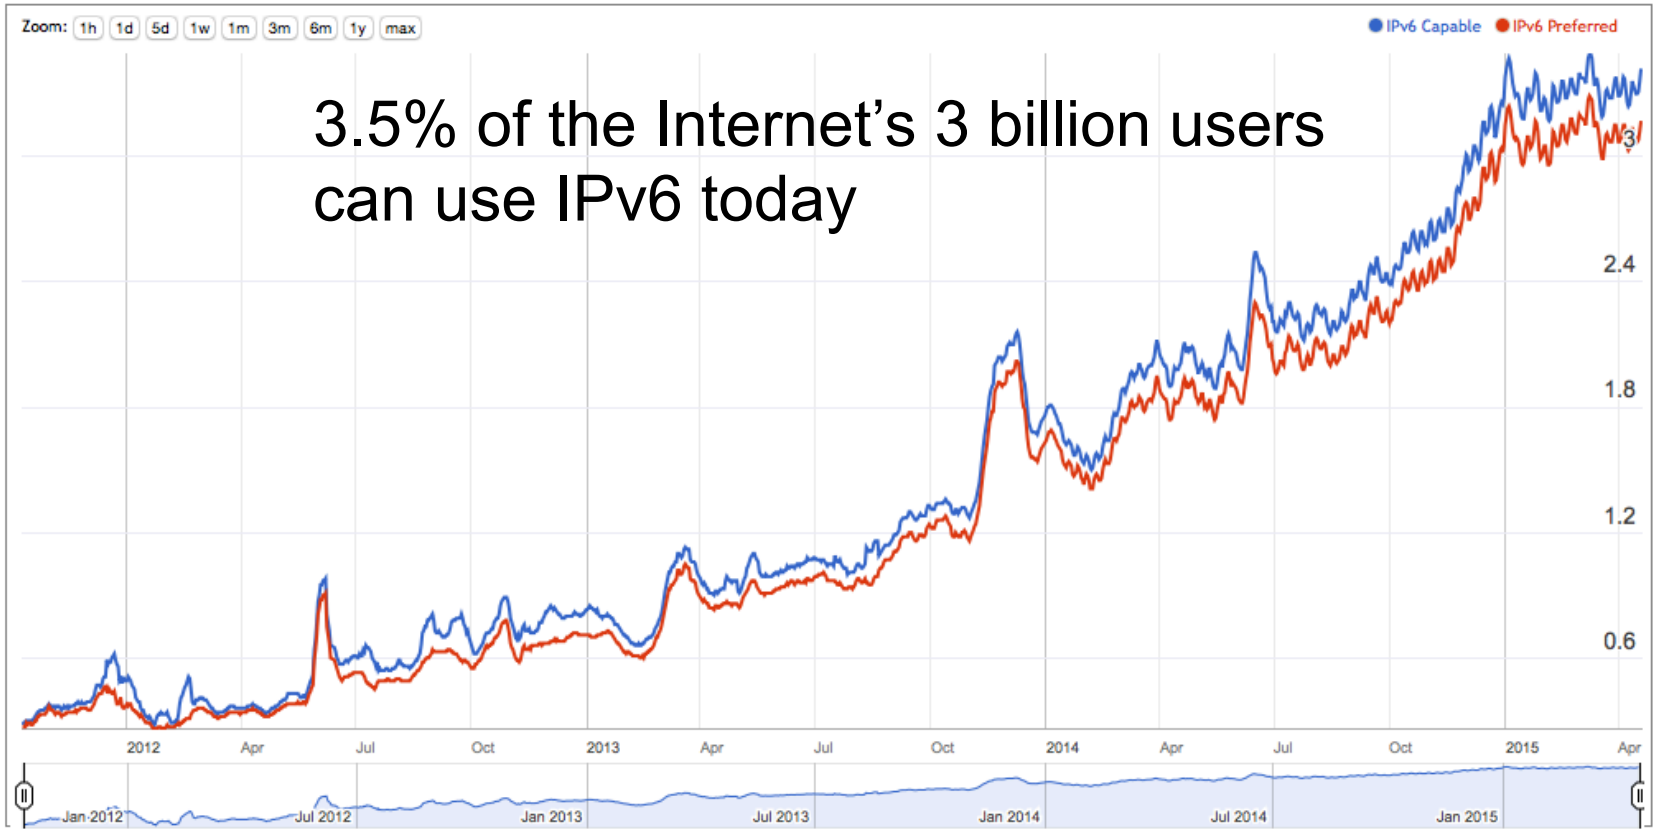

Running on Empty

To get from "here" to "there" requires an excursion through an environment of CGNs, CDNs, ALGs and similar middleware 'solutions' to IPv4 address exhaustion

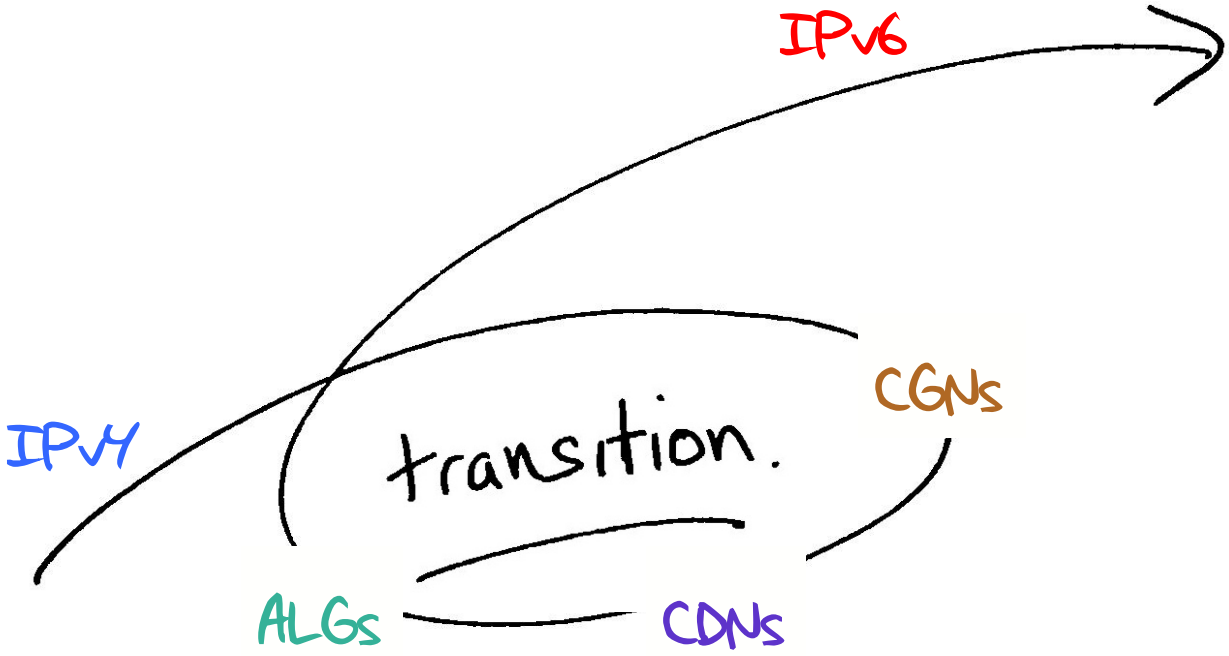

IPv4 Address Exhaustion

What are ISP's doing in response?

- It's not viable to switch over to IPv6 yet
- But the supply of further IPv4 addresses to fuel service platform growth has dried up
- How will ISPs continue to offer services to customers in the interim?

CGNs…

- What we are seeing is the increasing use of address sharing using Carrier Grade NATs as a means of extending the useable life of the IPv4 Internet while we are still waiting for IPv6 to be viable in its own right
- This has some significant implications for LEA functions, principally in traceback and ISP meta-data record keeping practices

Carrier Grade NATs

By sharing public IPv4 addresses across multiple customers!

Yes, that's my phone using net 10!

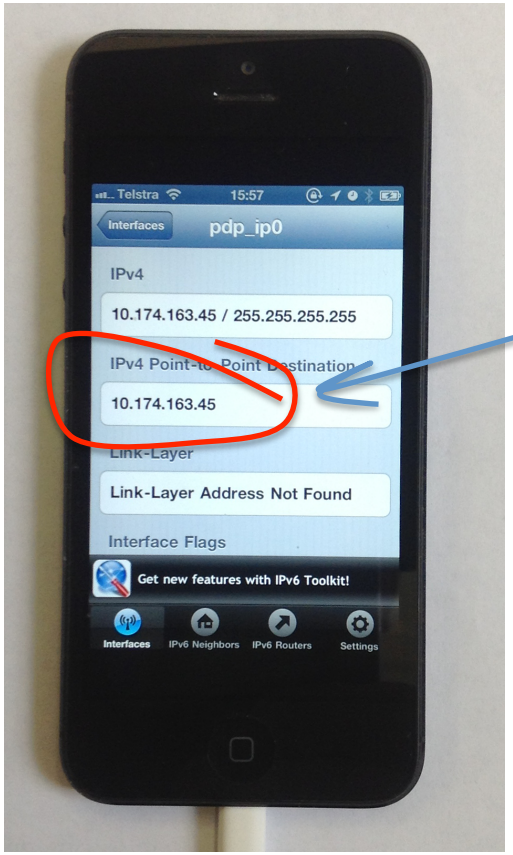

Carrier Grade NATs

By sharing public IPv4 addresses across multiple customers!

BT Begins Customer Tests of Carrier Grade NAT

Posted by **timothy** on Tuesday May 07, 2013 @09:27AM
from the party-line-but-with-less-yelling dept.

judgecorp writes

> "BT Retail has started testing Carrier Grade NAT (CGNAT) with its customer. CGNAT is a controversial practice, in which IP addresses are shared between customers, limiting what customers can do on the open Internet. Although CGNAT goes against the Internet's original end-to-end principles, ISPs say they are forced to use it because IPv4 addresses are running out, and IPv6 is not widely implemented. BT's subsidiary PlusNet has already carried out CGNAT trials, and now BT is trying it on "Option 1" customers who pay for low Internet usage."

http://tech.slashdot.org/story/13/05/07/1232234/bt-begins-customer-tests-of-carrier-grade-nat

NATs + CGNs

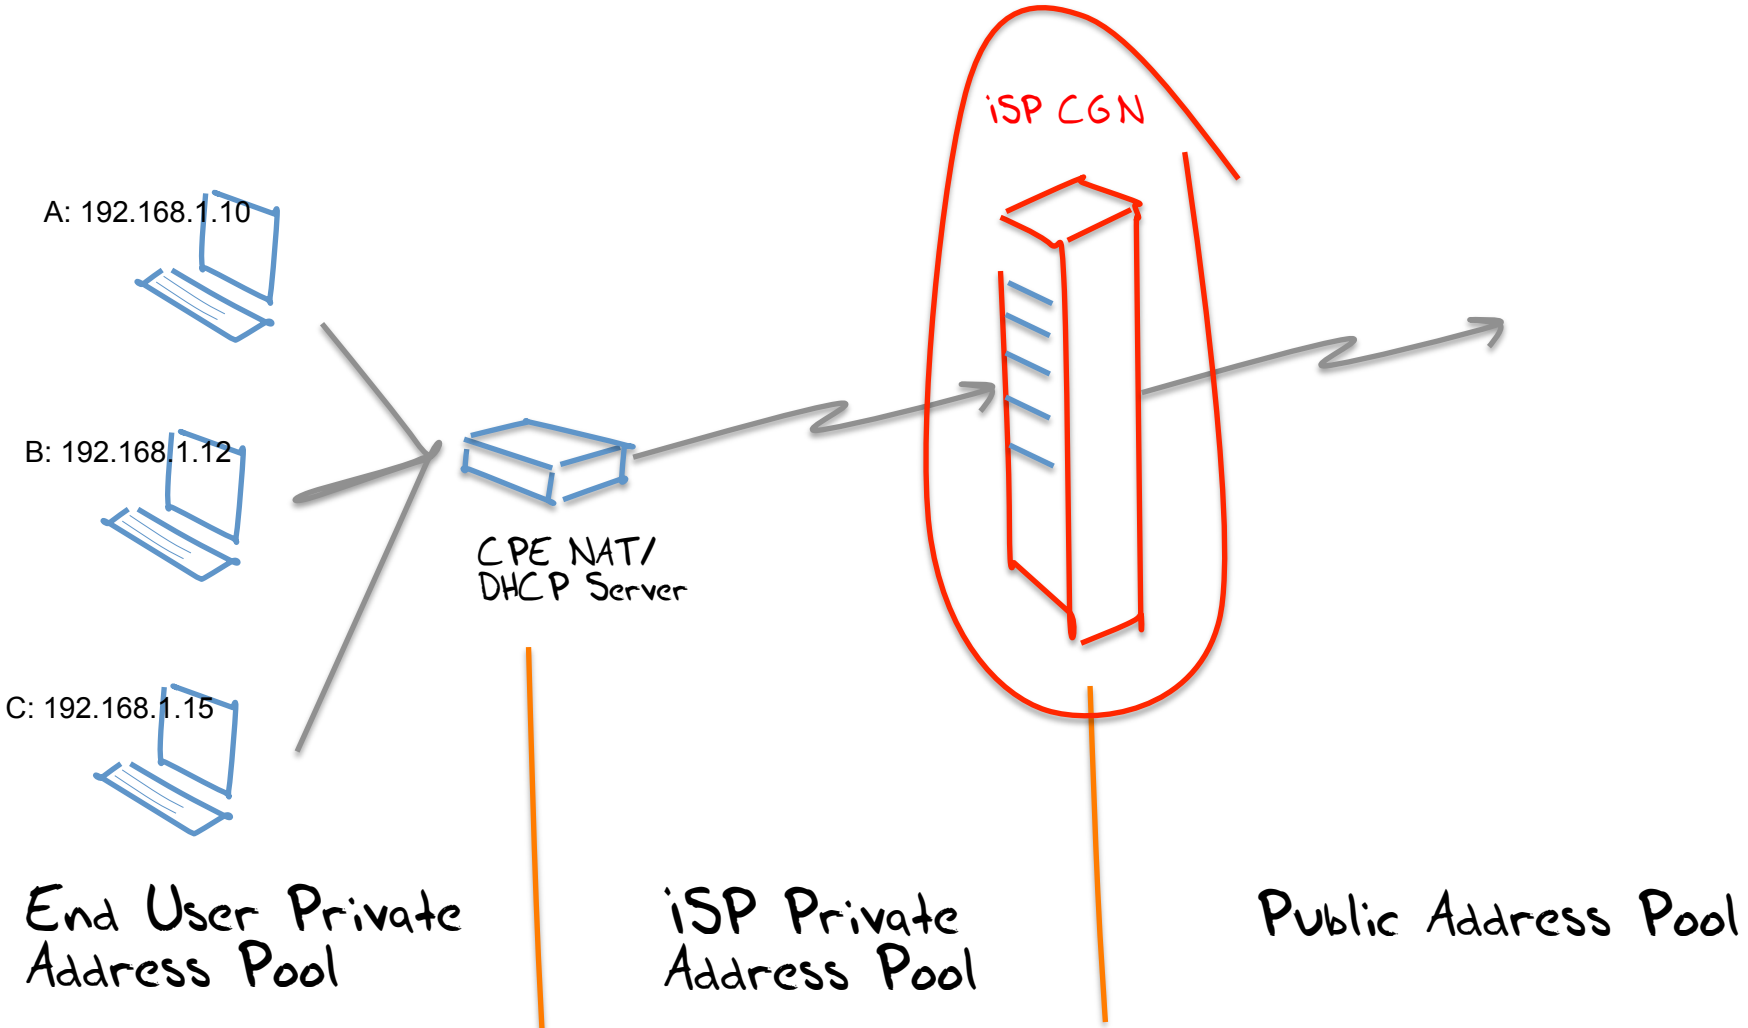
ISP CGN
A: 192.168.1.10
B: 192.168.1.12
C: 192.168.1.15
CPE NAT/
DHCP Server
End User Private
Address Pool
ISP Private
Address Pool
Public Address Pool

NATs + CGNs + Connections

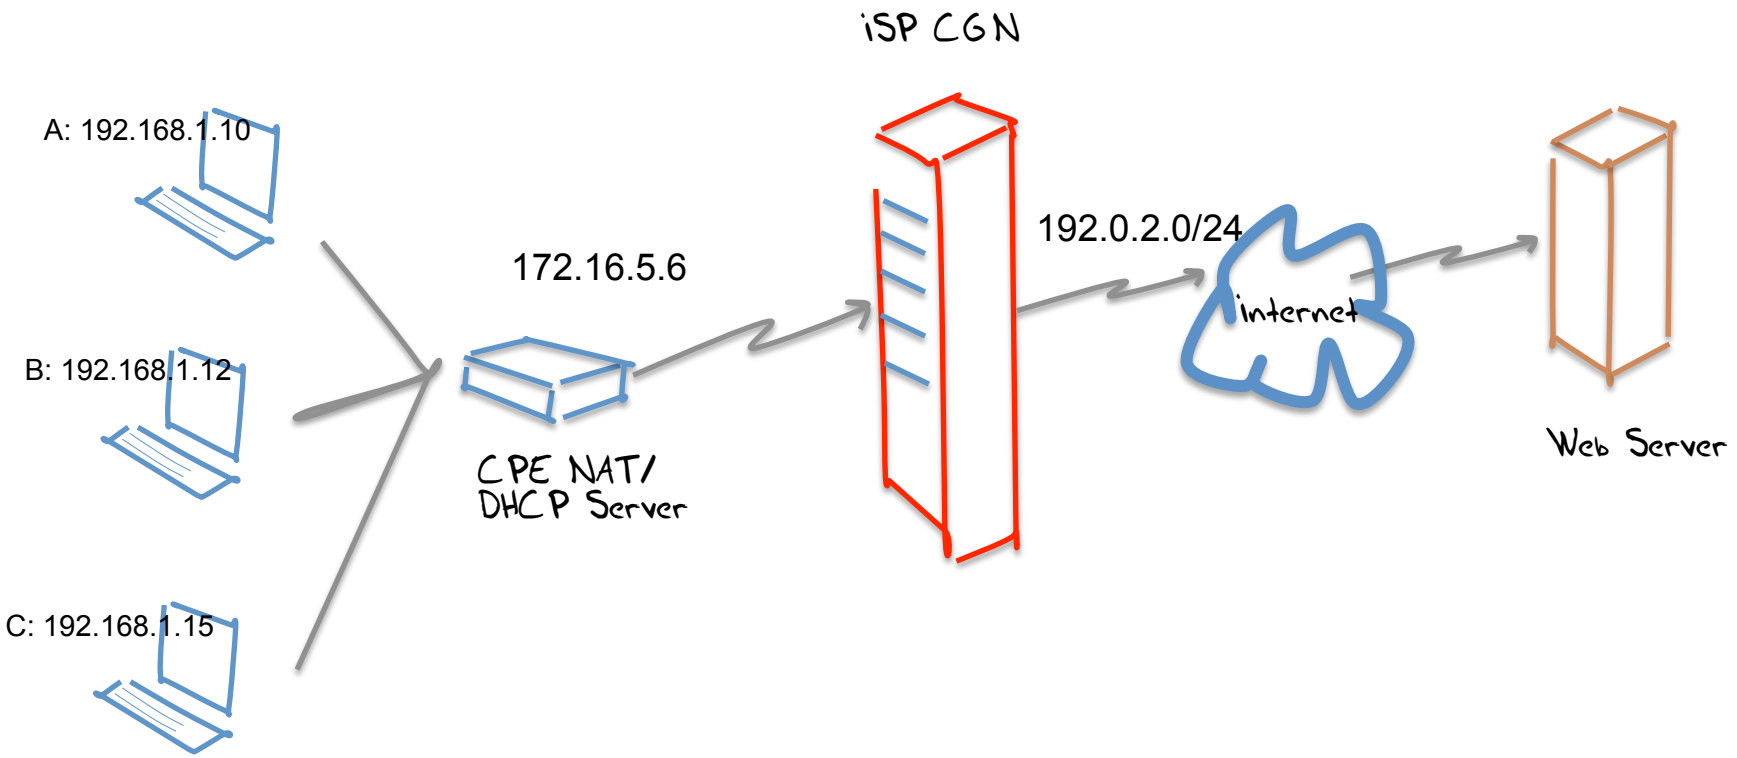

Assumptions

- The ISP operates a public address pool and a private address pool
- The access into the public address pool is via an ISP-operated NAT (CGN)
- Each end site is dynamically assigned a single private IP address upon login (AAA)
- The site is dynamically addressed using a private address range and a DHCP server
- The single public address is shared by the private devices through a CPE NAT

Traceback - Version 3

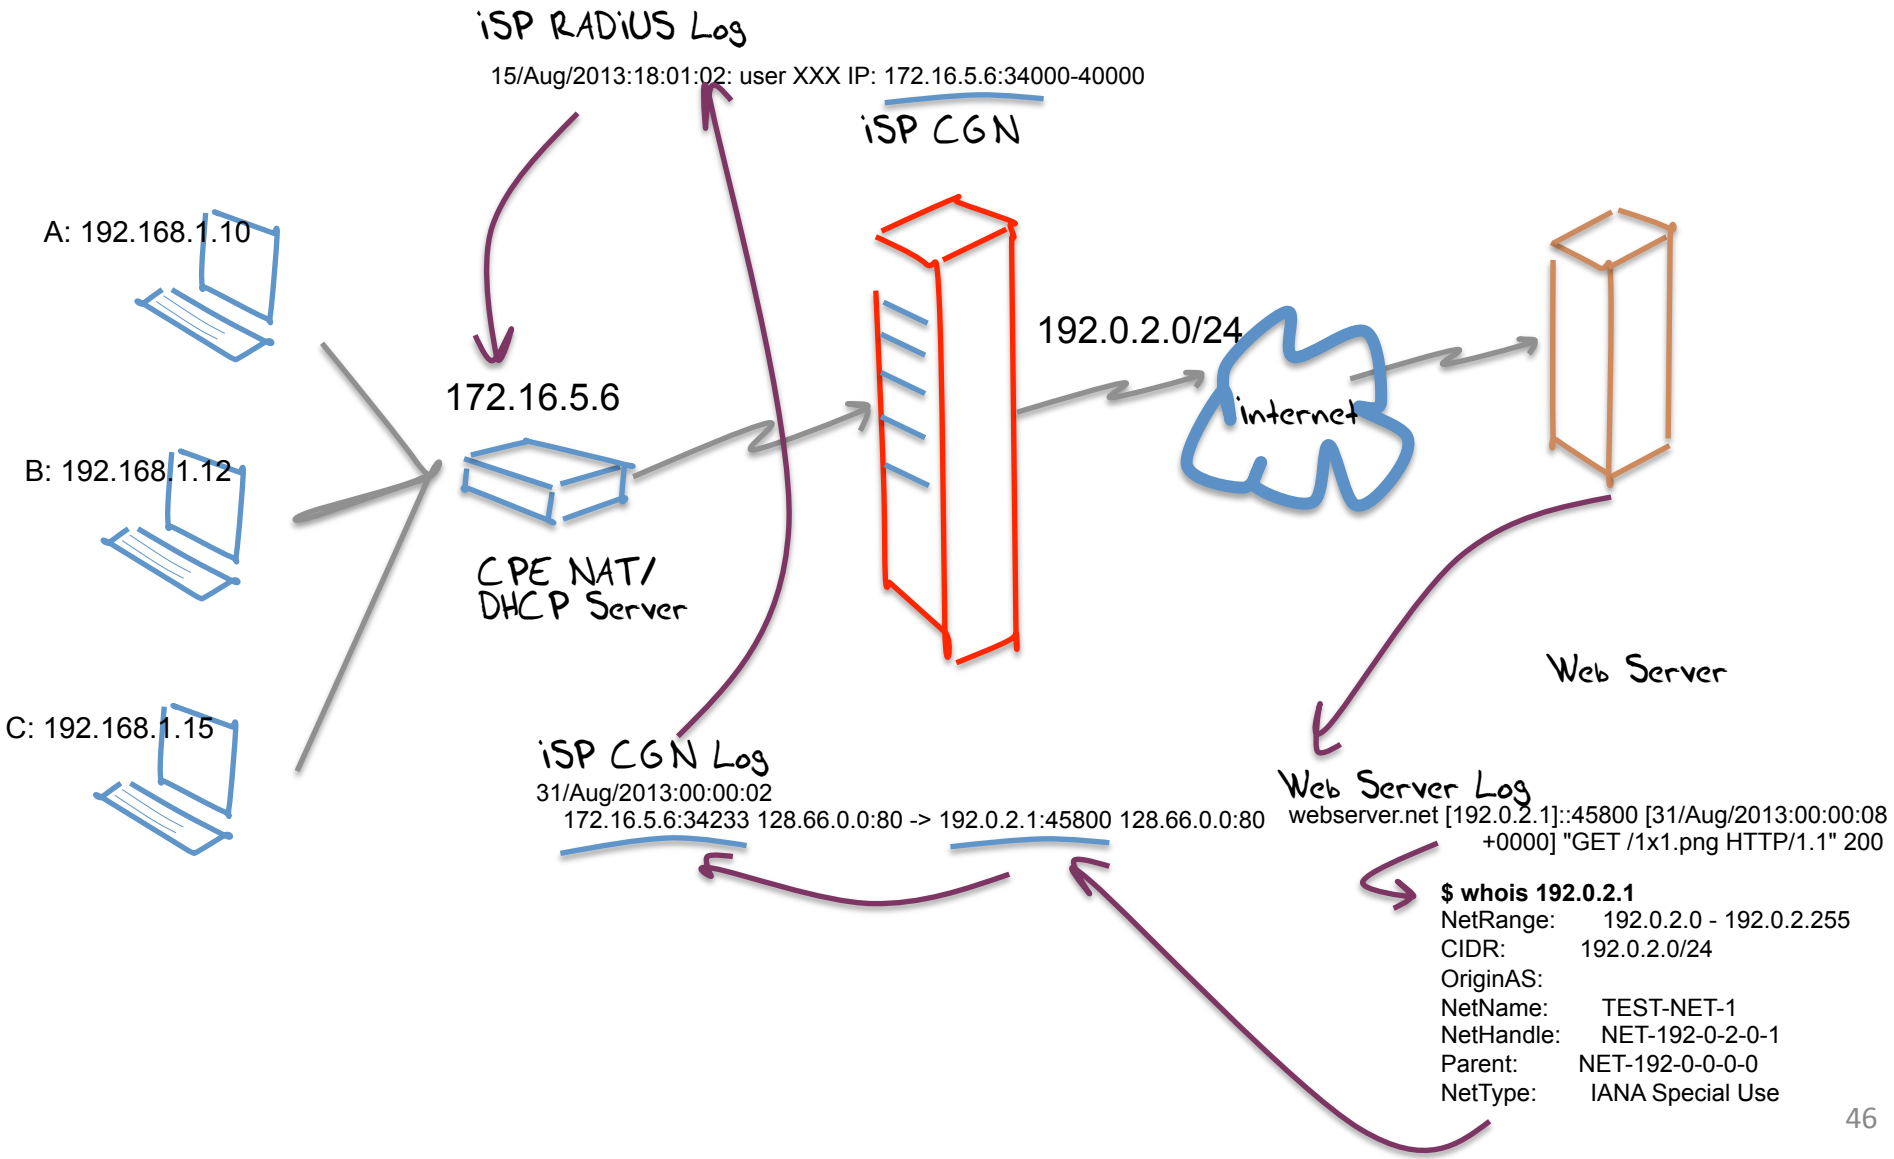

Assumptions

- Traceback to an end site is keyed by an **IP address** AND **a port address**, AND a **date/time** (uSec!)
  - Requires access to:
    - WHOIS records to identify the ISP,
    - the ISP's CGN logs to identify the ISP's private address, and
    - the ISP's AAA logs to identify the end site

Assumptions

- Traceback to an end site is keyed by an **IP address** AND **a port address**. AND a **date/time** (uSec!)

Nobody logs this!

  - Requires access to:
    - WHOIS records to identify the ISP,
    - the ISP's CGN logs to identify the ISP's private address, and
    - the ISP's AAA logs to identify the end site

ISP CGN Logging

CGN bindings are formed for EVERY unique TCP and UDP session

That can be a LOT of data to retain…

The Horror (log volumes)

150 - 450 bytes/connection

\+ 33k - 216k connections per sub per day

---

5 - 96 MB / user / day

That's potentially over 1 PB per 1M subs per month

It's also over 20Mbps for just the log stream…

http://www.nanog.org/meetings/nanog54/presentations/Tuesday/GrundemannLT.pdf

It could be better than this…

- Use Port Blocks per customer

or

- Use a mix of Port Blocks and Shared Port Pool overflow

and

- Compress the log data (which will reduce storage but may increase search overhead)

Or it could be worse…

Challenges in Address Exhaustion:

1. This is a deregulated and highly competitive environment
   There is no plan, just the interplay of various market pressures

2. Varying IPv4 Address Exhaustion Timelines
   Differing time lines create differing pressures in the market

3. Regional Diversity
   One network architecture is not an assured outcome!

What does this mean for the Internet?

What does this mean for the Internet?

We are going to see a LOT of transition middleware being deployed!

What does this mean for the Internet?

We are going to see a LOT of transition middleware being deployed!

And we are going to see a significant diversity in what that middleware does

What does this mean for LEAs?

LEAs have traditionally focused on the NETWORK as the point of interception and tracing

They are used to a consistent model to trace activity:

- get an IP address and a time range
- traceback based on these two values to uncover a set of network transactions

What does this mean for LEAs?

In a world of densely deployed CGNs and ALGS then the IP address loses any coherent meaning in terms of end party identification.

What does this mean for LEAs?

In a world of densely deployed CGN, then the IP address loses any[illegible]ning in terms of end party ident[illegible]

Today's traceback approaches won't work any more!

What does this mean for LEAs?

And instead of shifting to a single “new” model of IP address use, we are going to see widespread diversity in the use of transition mechanisms and NATs in carrier networks

Which implies that there will no longer be a useful single model of how to perform traceback on the network

Or even a single coherent model of “what is an IP address” in the network

Variants of NAT CGN Technologies

Variant:	Address Compression Ratio
CGN with per-user port blocks	10:1
CGN with per-user port blocks + pooled overflow	100:1
CGN with pooled ports	1,000:1
CGN with 5-tuple binding maps	>>10,000:1

The same public address and port is used simultaneously by multiple different internal users

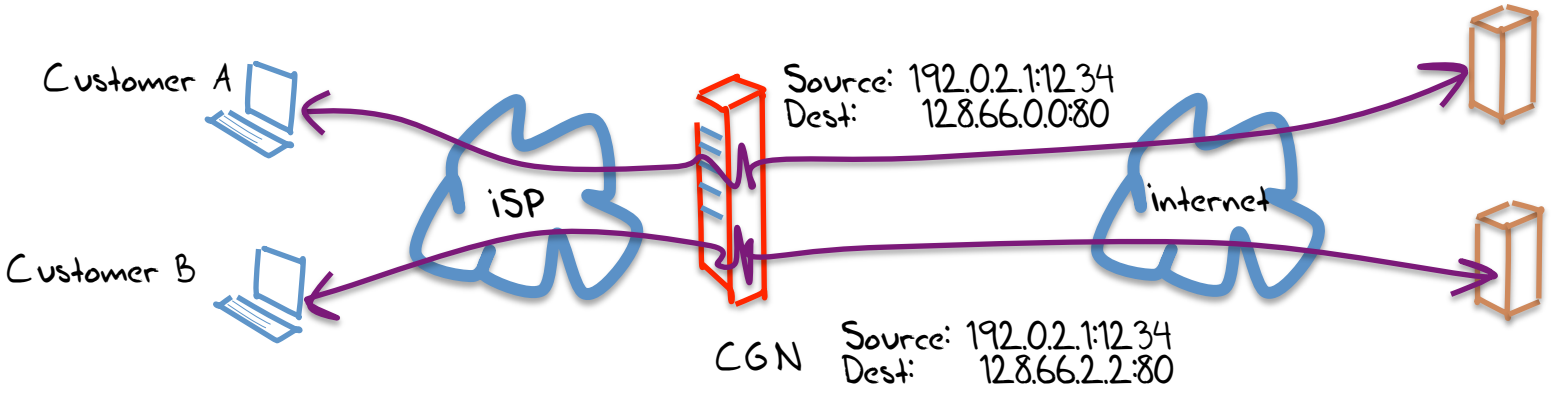

It gets worse …

Adding IPv6 to the CGN Mix

- The space is not exclusively an IPv4 space.
- While CGNs using all-IPv4 technologies are common today, we are also looking at how to use CGN variants a mix of IPv6 and IPv4

  For example: Dual-Stack Light connects IPv4 end users to the IPv4 Internet across an IPv6 ISP infrastructure.

- We can expect to see many more variants of ISP's address transform middleware when you are allowed to add IPv6 into the mix

++IPv6: Transition Technologies

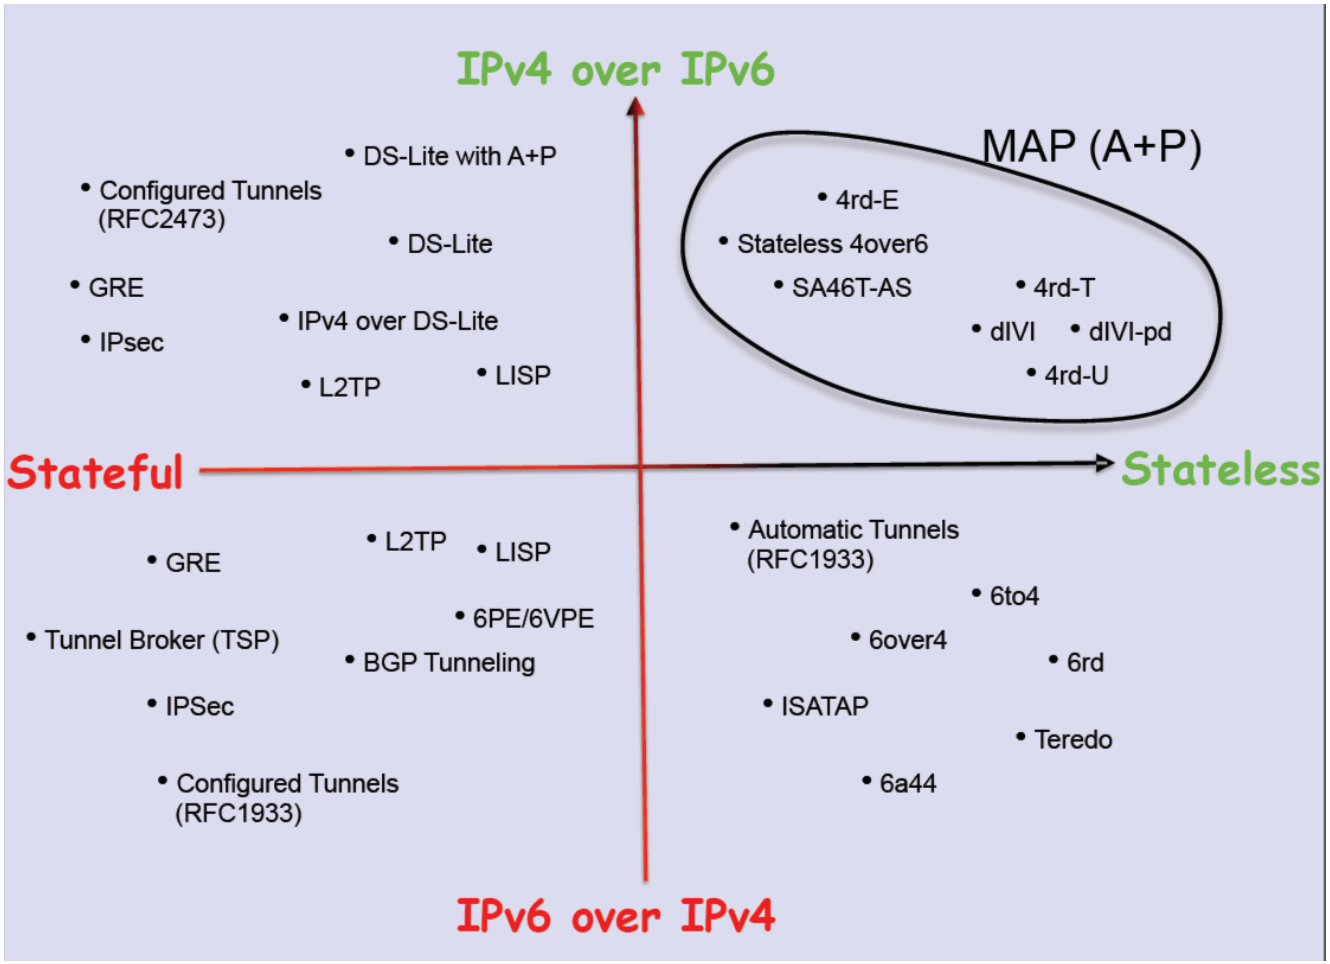

Randy Bush, APPRICOT 2012: http://meetings.apnic.net/__data/assets/pdf_file/0016/45241/120229.apops-v4-life-extension.pdf

Transition Technologies Example: 464XLAT

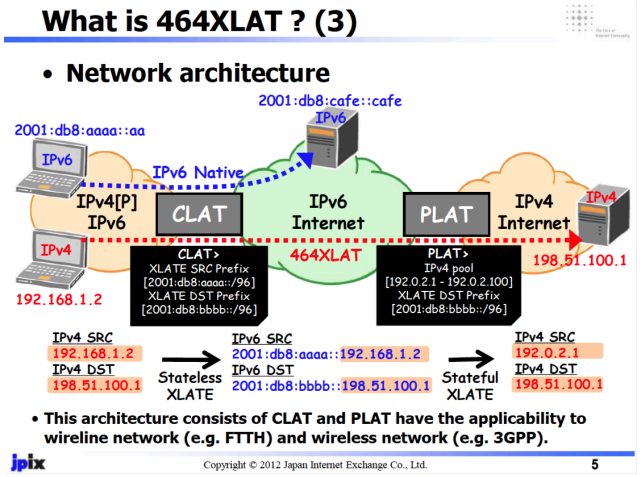

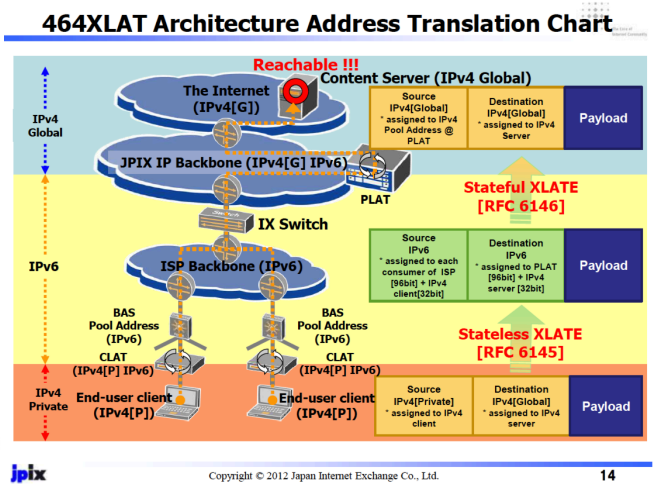

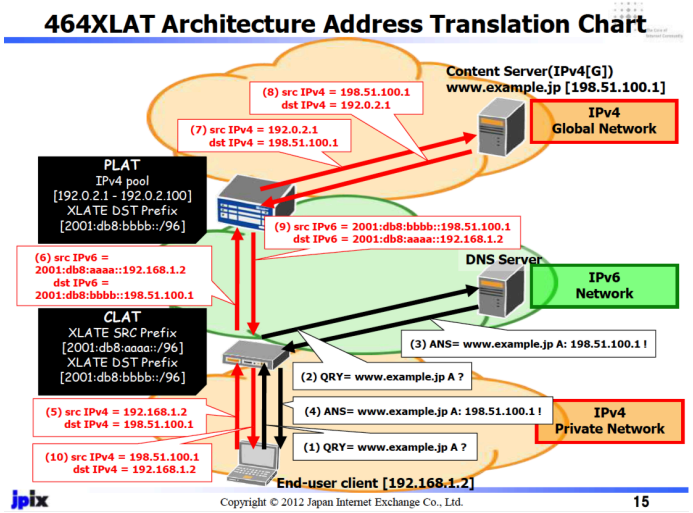

Masataka Mawatari, Apricot 2012, http://meetings.apnic.net/__data/assets/pdf_file/0020/45542/jpix_464xlat_apricot2012_for_web.pdf

What does this mean for LEAs?

The risk we are running at the moment is that there is no longer be a single consistent model of how an IP network manages IPv4 and IPv6 addresses

What does this mean for LEAs?

What's the likely response from LEAs and regulators?

One likely response is to augment the record keeping rules for ISPs:

> record *absolutely everything*, and keep the records for 2 years
>
> [Australian Data Retention, 2015]

What does this mean for ISPs and LEAs?

But what are the new record keeping rules?

In order to map a "external" IP address and time to a subscriber as part of a traceback exercise then:

- for **every** active middleware element you now need to hold the **precise** time and the **precise** transforms that were applied to a packet flow
- and you need to be able to **cross-match** these records accurately

What does this mean for ISPs and LEAs?

But what are the new record keeping rules?

In order to map a "external" [illegible]ress and time to a subscriber as part of [illegible]back exercise then:

- for **every** [illegible]ddleware element you now need to hold th[illegible]ime and the **precise** transforms that were applied [illegible] packet flow
- and you need to be able to **cross-match** these records accurately

Degree of difficulty -- approaching 10/10 !

What does this mean for ISPs and LEAs?

How many different sets of record keeping rules are required for each CGN / dual stack transition model being used?

And are these record keeping practices affordable?

(granularity of the records is shifting from “session” records to “transition” and even individual packet records in this diverse model)

Are they even practical within today’s technology capability?

Is this scaleable?

Is it even useful any more?

Traceback in tommorrow's Internet?

The traceback toolkit:

- precise time, source and dest IP addrs, protocol and port information
- Access to all ISP middleware logs
- CDN SP logs
- Network and Middleware deployment maps
- V6 Transition technology map used by the ISP
- A thorough understanding of vendor's equipment behaviour for various applications
- A thorough understanding of application behaviours

Making it hard...

The V6 transition was challenging enough

The combination of V4 exhaustion and V6 transition is far harder

The combination of varying exhaustion times, widespread confusion, diverse agendas, diverse pressures, V4 exhaustion and V6 transition is now amazingly challenging

Making it very hard...

The problem we are facing is that we are heading away from a single service architecture in our IP networks

Different providers are seeing different pressures and opportunities, and are using different technology solutions in their networks

And the longer we sit in this “exhaustion + transitioning” world, the greater the diversity and internal complexity of service networks that will be deployed

“Toto, I've a feeling we're not in Whois-land ~~Kansas~~ any more!”

All this will makes the entire record and trace problem for ISPs and LEAs harder

At some point along this path of escalating network complexity and diversity its likely that our networks will be simply be unable to traceback individual use in any coherent manner

If this is where the Internet is heading, then from an LEA perspective the tracking and tracing story is looking pretty bad

Does it ever get easier?

is there light at the end of this tunnel?

The Transition to IPv6

- Once we get to complete this transition we no longer need to use IPv4
- Which means that we can throw aware these CGNs and their associated records
- And the entire exercise of record keeping and traceback gets a whole lot easier

Traceback - IP Version 6

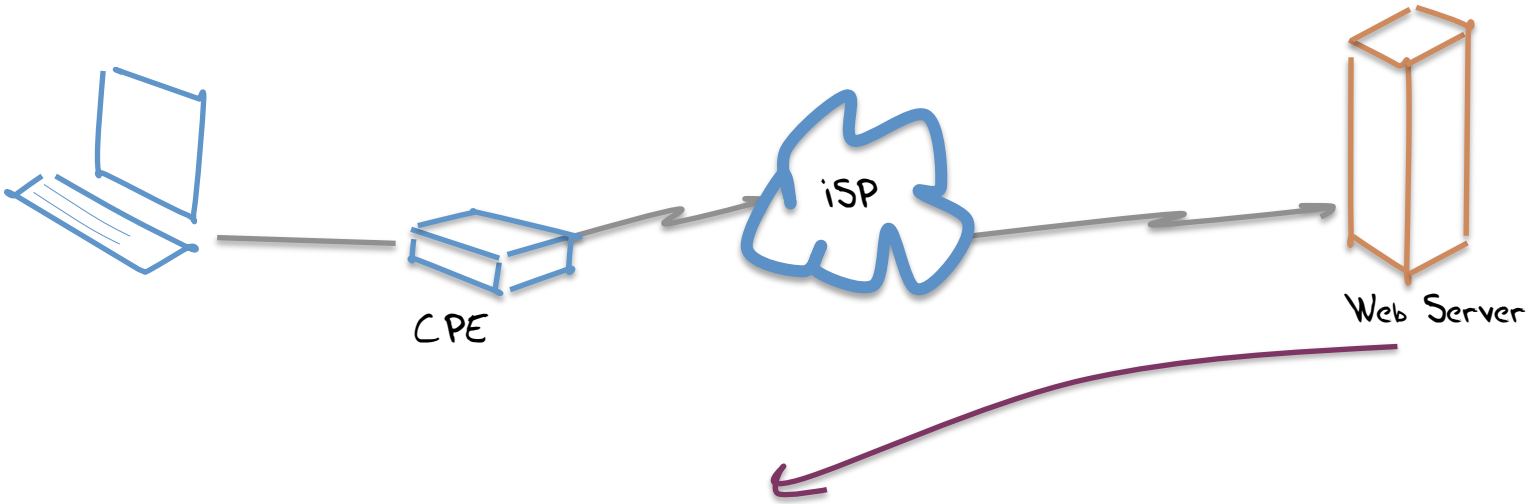

Traceback - IP Version 6

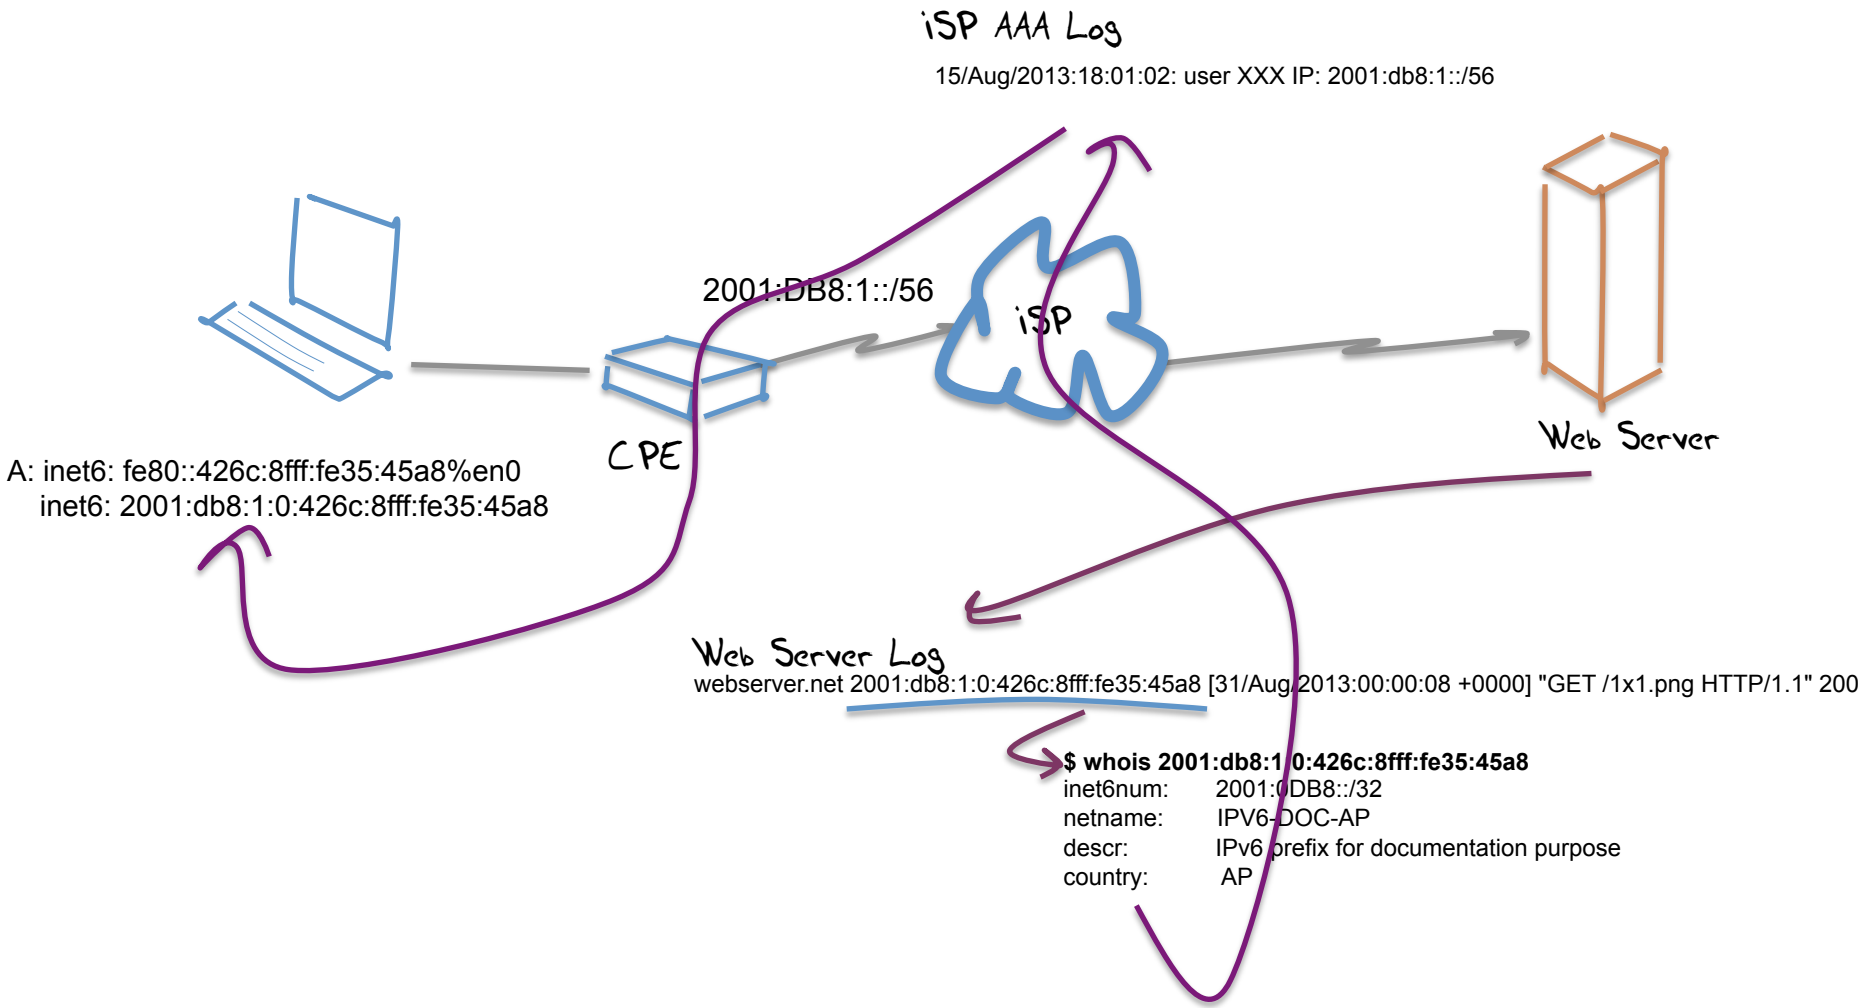

IPv6 makes it easy again. Right?

Yes.

The semantics an IPv6 address in an IPv6 network are much the same as the original model of IPv4 addresses in a non-NATTed IPv4 Internet

Which is good.

But it's not completely the same as the original IPv4 model…

IPv6 makes it mostly easy again. Right?

IPv6 Privacy Addresses introduce ephemeral public IPv6 addresses into the mix

There are no logs of the privacy address, as it's self assigned

> IPv6 Privacy addresses are used in Windows, Max OSx, some variants of Linux. We will see this in mobile networks as well in the coming months.

So IPv6 may not be able to track back to the device every time. Sometimes the best you can get is the home site and no closer!

As long as the /64 network address can trace to the end customer / mobile device then this will not be a critical problem – but the network's address architecture is now a critical piece of knowledge

The Bottom Line

Compared to the byzantine complexities of the emerging CGN world of the IPv4 Internet, it certainly appears that an IPv6 Internet makes the conventional activities of record keeping and logging far easier once more

Typically, these IPv6 addresses will map all the way back to the MAC address of the device that is attached to the network

With IPv6 Privacy Addresses these address records do not necessarily resolve back to individual devices all the time, but they should give consistent visibility to the granularity of the home/end site network based on IPv6 address without massive record generation

Thank You!