

Forensic Tracing in the  
Internet:  
An Update

Geoff Huston

Chief Scientist

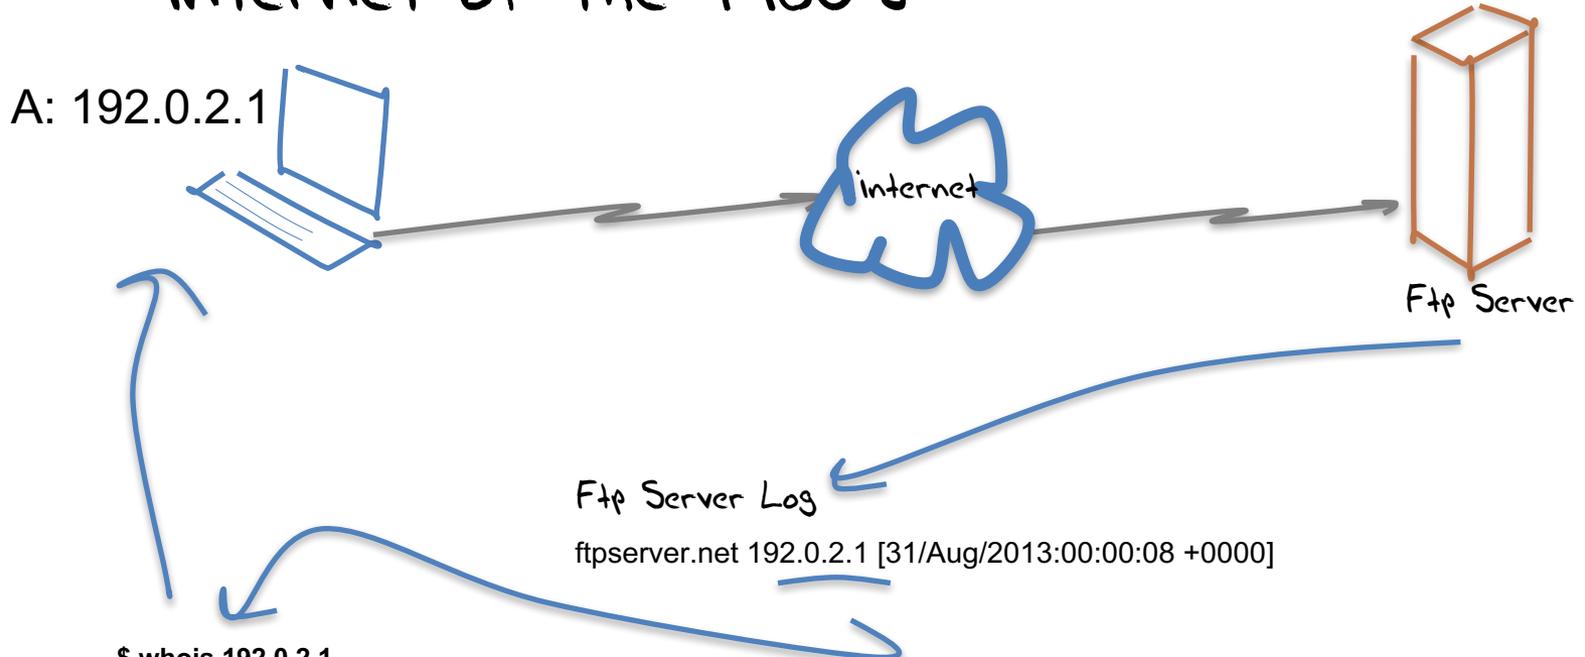
APNIC

# The story so far...

- The status of the transition to IPv6 is not going according to the original plan:
  - We have exhausted the remaining pools of IPv4 addresses in all regions except Africa - this was never meant to have happened
  - We we meant to have IPv6 fully deployed by now
- What we are seeing is the pervasive use of Carrier Grade NATs as a means of extending the useable life of the IPv4 Internet
- Around 10% of users use both IPv6 and IPv4 – the other 90% are IPv4 only
- It appears that most IPv4 use today uses NATs in the path
- This has some major implications for LEA functions, principally in traceback and metadata record keeping

# Traceback - Version 1

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



**\$ whois 192.0.2.1**

NetRange: 192.0.2.0 - 192.0.2.255  
NetName: TEST-NET-1  
Contact: User Contact Details

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

# Assumptions:

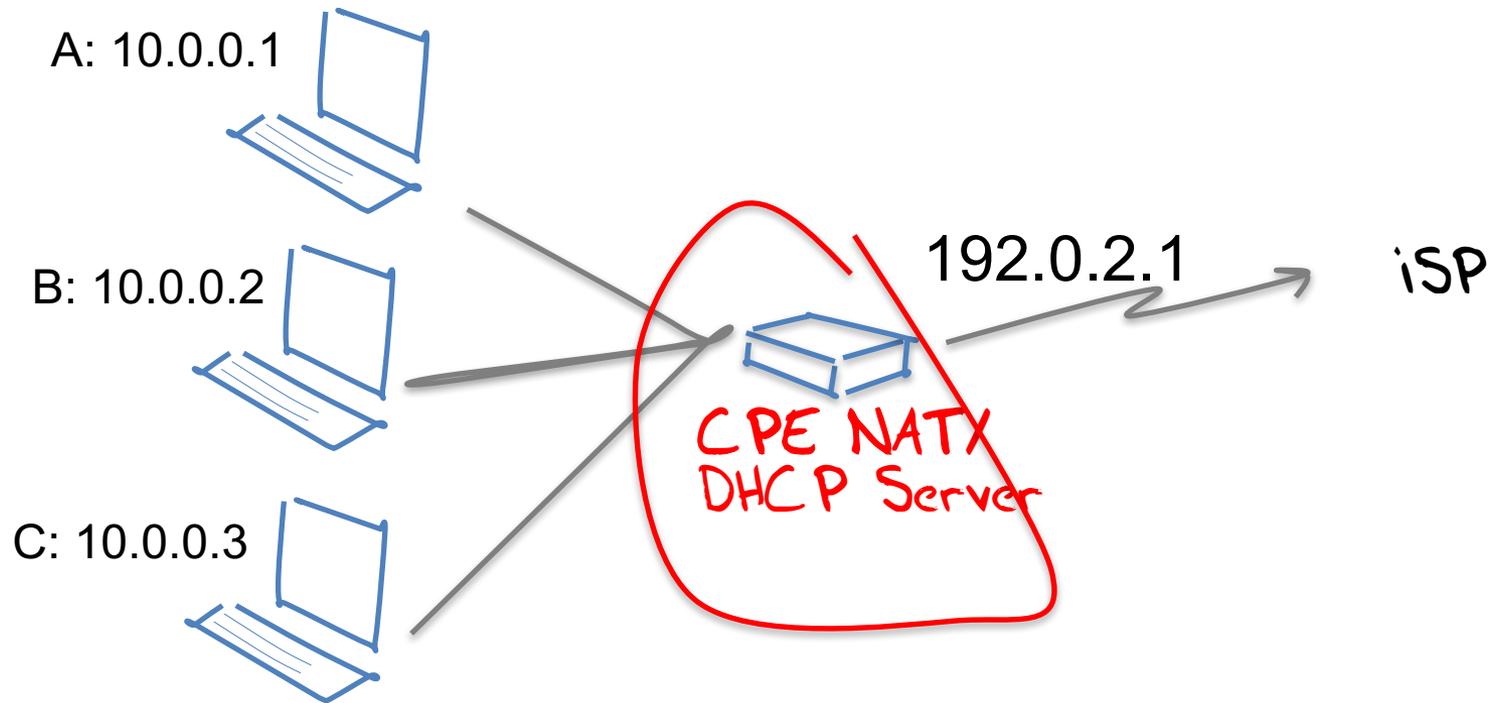
- Each end site used a stable IP address range
- Each address range was recorded in a registry, together with the end user data
- Each end device was manually configured with a stable IP address
- Traceback was keyed from the IP address

# Assumptions:

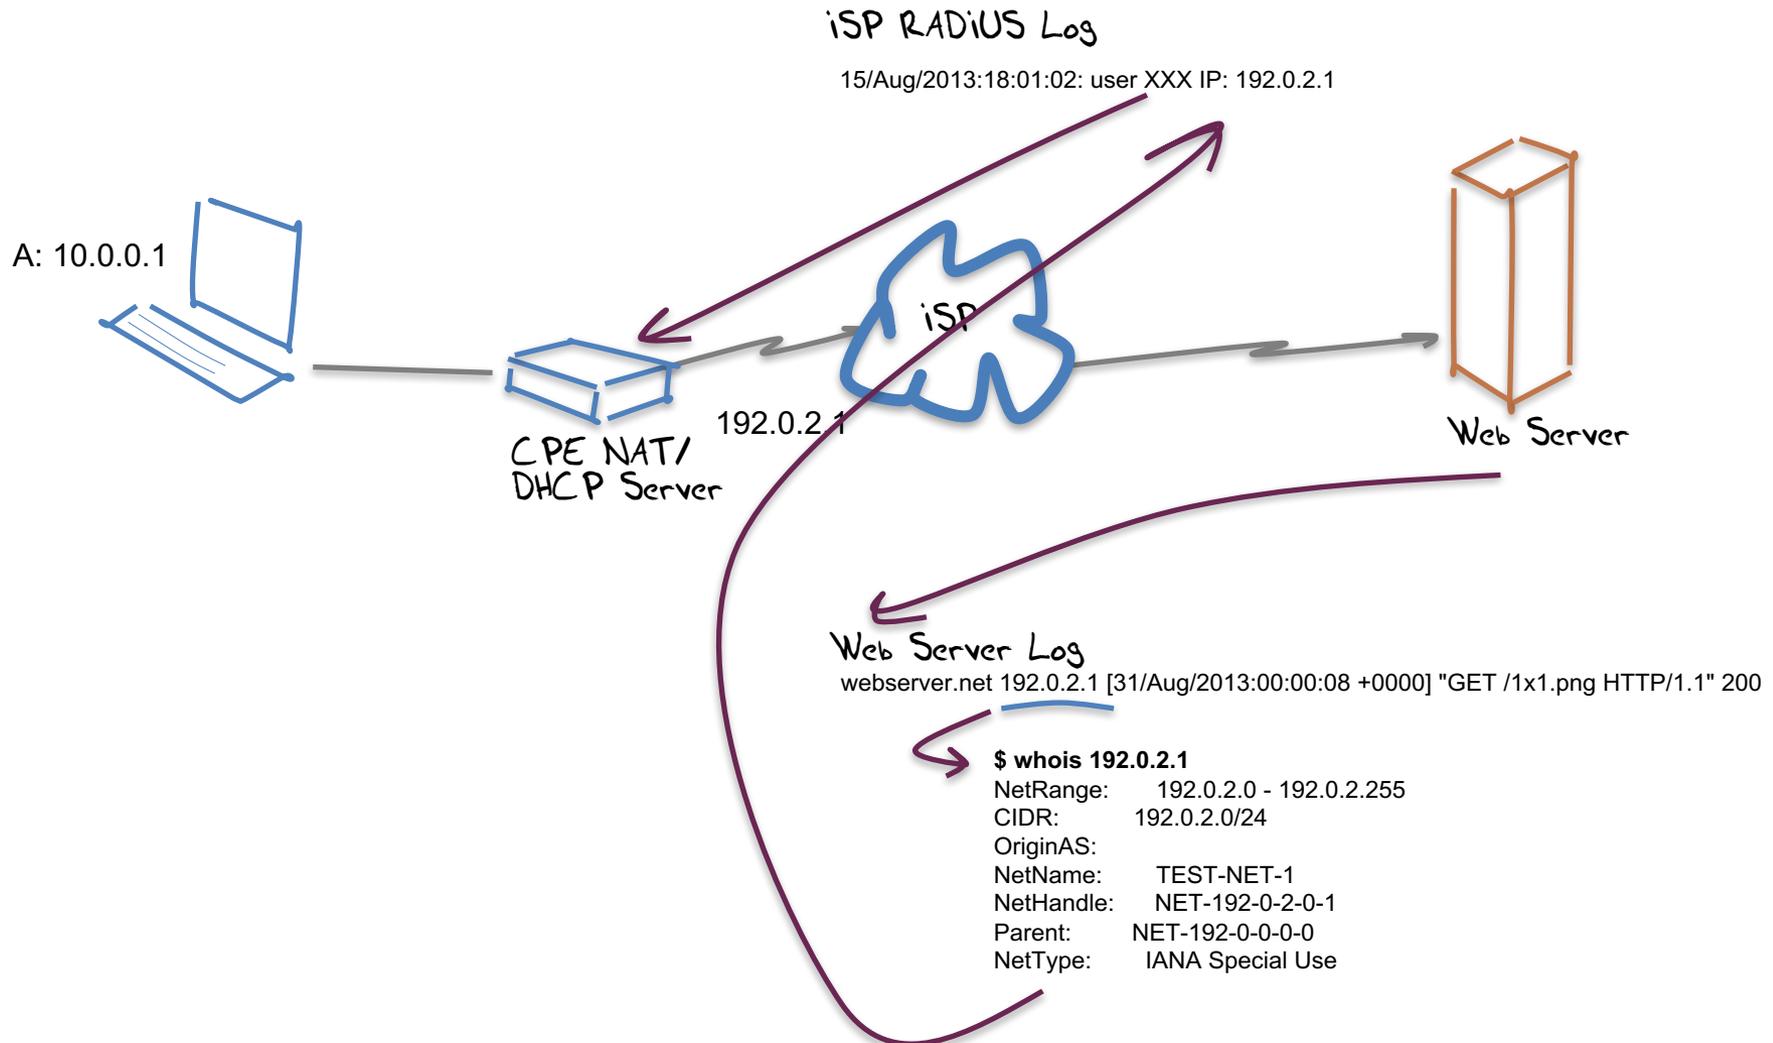
- Each end site used a stable IP address
- Each address range was recognized together with the end user
- Each end device was associated with a stable IP address
- Traceback was possible for any IP address

*This model largely fell into disuse in the 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



# Assumptions

- The ISP operates an address pool
- Each end site is dynamically assigned a single 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

# Changes

- Traceback to an end site is keyed by an IP address and a date/time
  - Requires access to WHOIS records to identify the ISP and the ISP's AAA logs to identify the end site
- No traceback to an individual device – the trace stops at the edge NAT

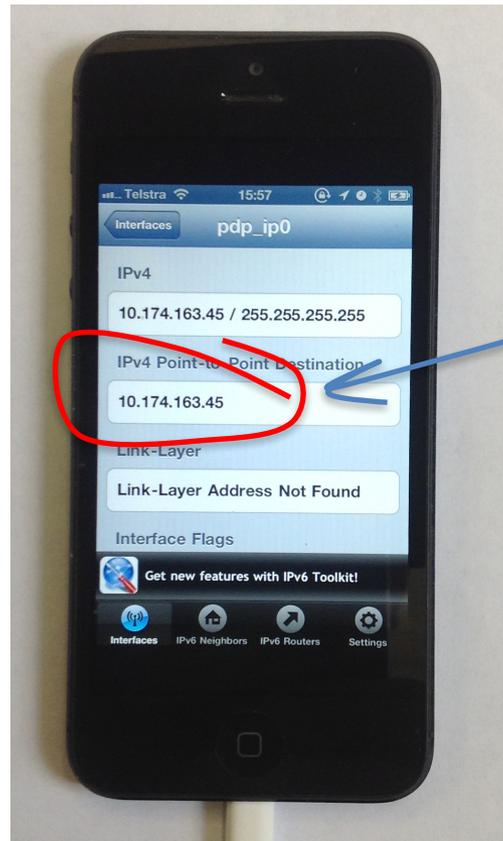
# IPv4 Address Exhaustion

What are ISP's doing in response?

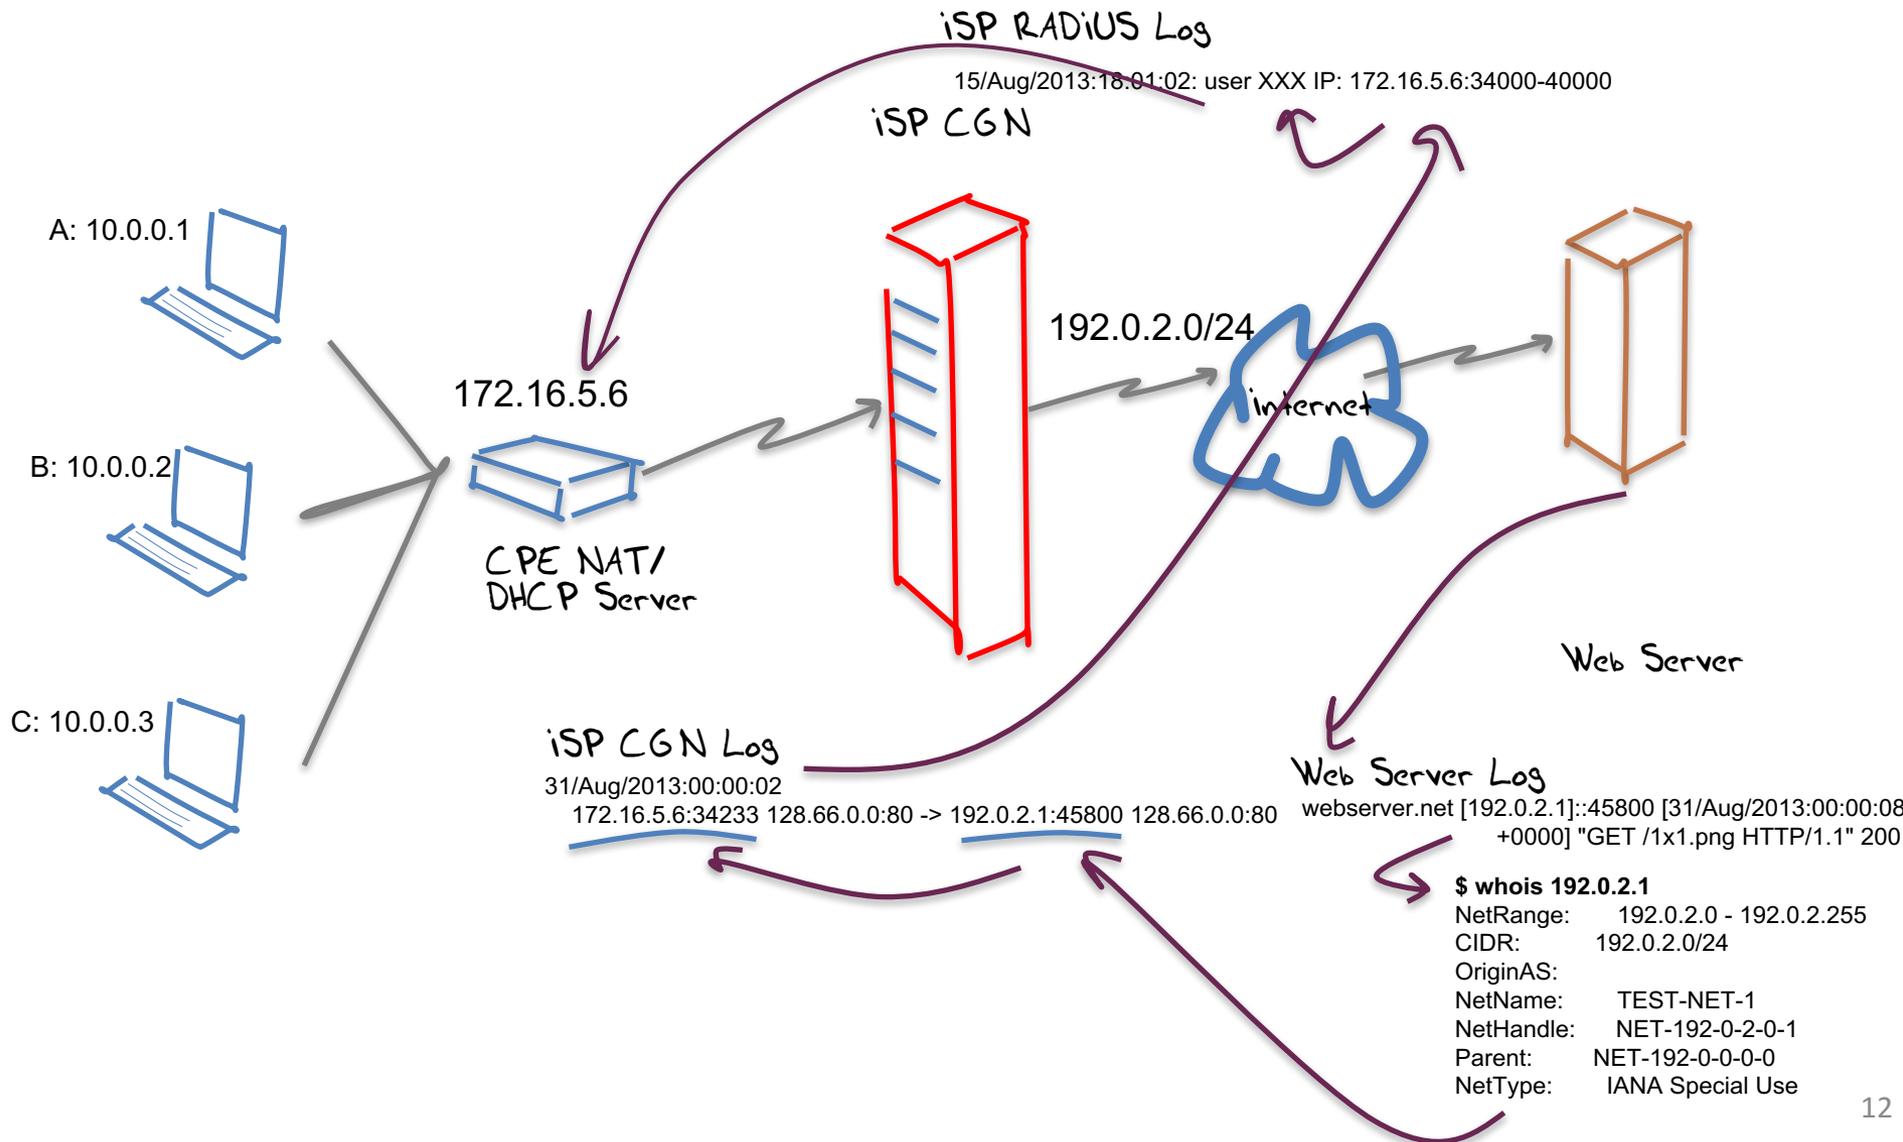
- It's not viable to switch over to all-IPv6 yet
- But the supply of further IPv4 addresses to fuel service platform growth has dried up
- How do ISPs continue to offer IPv4 services to customers in the interim?

# Carrier Grade NATs

By sharing public IPv4 addresses across multiple customers!



# Traceback - Version 3



# 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

# Assumptions

- Traceback to an end site is keyed by a source IP address and a source port address, and a date/time
- 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...*

# 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...

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

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

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

# What does this mean for Forensic tracing?

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
- trace back based on these two values to uncover a set of network transactions

# What does this mean for Forensic tracing?

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

# What does this mean for Forensic tracing?

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

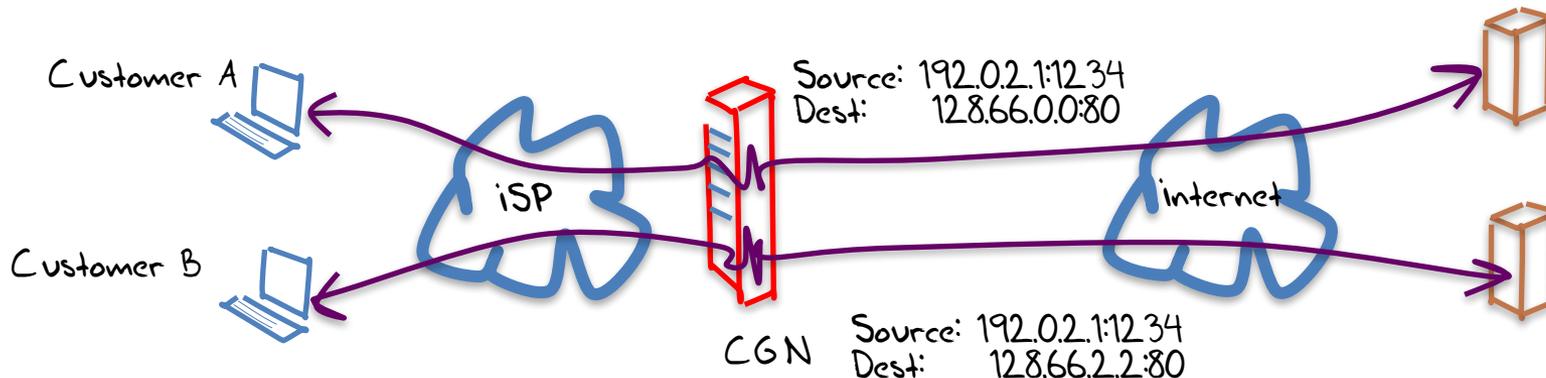
Which implies that there will no longer be a useful single model of how to perform traceback on the network

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 <u>5-tuple binding maps</u>	>>10,000:1

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



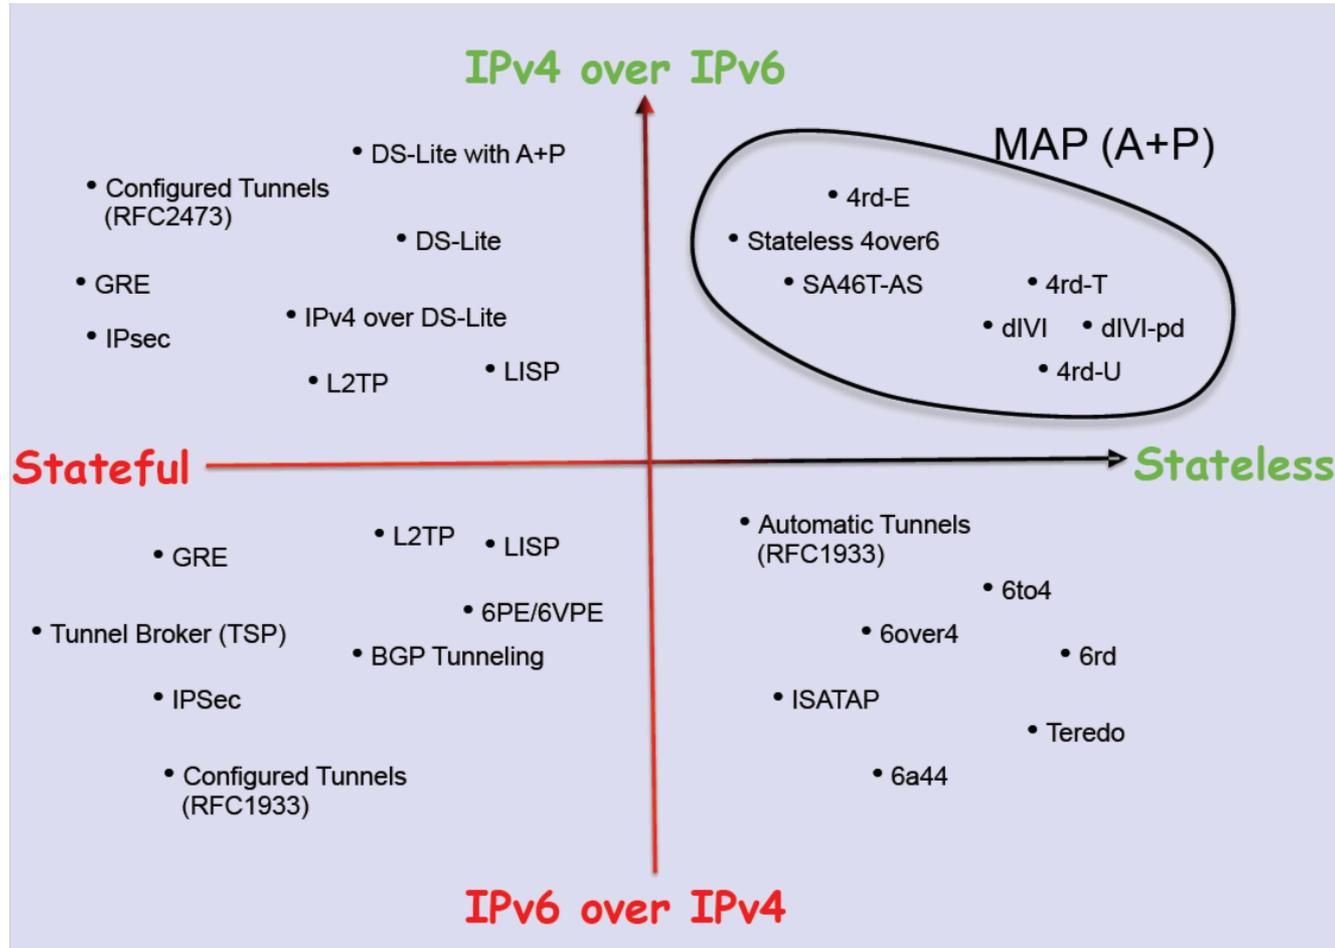
# Adding IPv6 to the CGN Mix

- 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 with 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 see many more variants of ISP's address transforming middleware when they IPv6 into the mix

# ++IPv6: Transition Technologies

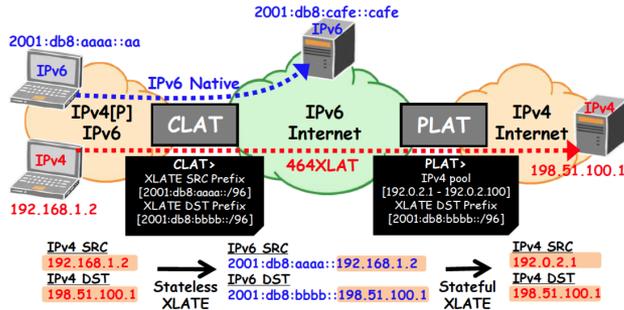


# Transition Technologies

## Example: 464XLAT

### What is 464XLAT ? (3)

#### • Network architecture



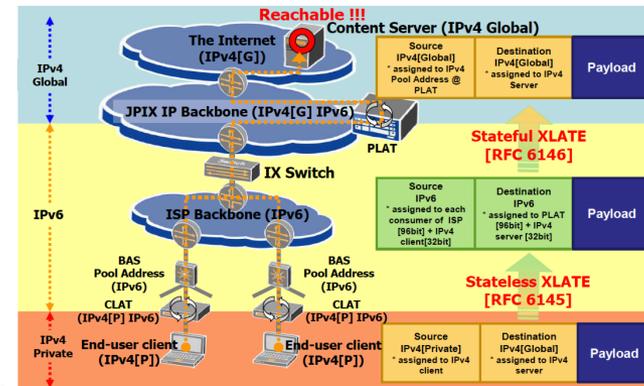
- This architecture consists of CLAT and PLAT have the applicability to wireline network (e.g. FTTH) and wireless network (e.g. 3GPP).

jpix

Copyright © 2012 Japan Internet Exchange Co., Ltd.

5

### 464XLAT Architecture Address Translation Chart

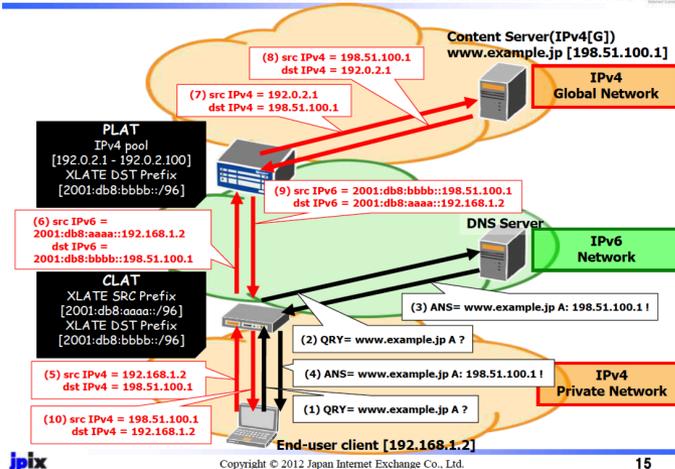


jpix

Copyright © 2012 Japan Internet Exchange Co., Ltd.

14

### 464XLAT Architecture Address Translation Chart



jpix

Copyright © 2012 Japan Internet Exchange Co., Ltd.

15

# What does this mean for Forensic tracing?

There is no single consistent model of how an IP network manages IPv4 and IPv6 addresses

There is no fixed relationship between IPv4 and IPv6 addresses

What you see in terms of network trace information is strongly dependent on **where** the trace data is collected

# 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

# 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” process and time to a subscriber as part of a back exercise then:

*Degree of difficulty: approaching 10/10!*  
for every software element you now need to hold the time and the **precise** transforms that were applied to packet flow

and you need to be able to **cross-match** these records accurately

# 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 years”*

# 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 scalable?

Is it even useful any more?

# 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

Does it ever get easier?

is there light at the end of this tunnel?

# That was then

The material so far refers to the Internet of late 2013

Three years later, has it got any easier?

Or has it just got harder?

# Sessions are the Key

We assumed that there is a “session” that maps between a service and a client, and this session is visible in some manner to the network

The forensic task was to take a partial record of a “session” and identify the other party to the session by using ancillary information (whois registries, web logs, metadata data sets, etc)

But maybe the entire concept of a “session” no longer exists! Do we still use “sessions” in applications?

What is changing?

Spying in America

# How Edward Snowden changed history

A damning account of a devastating intelligence breach

Jan 14th 2017



257



**How America Lost Its Secrets: Edward Snowden, the Man and the Theft.** By Edward Jay Epstein. *Knopf*; 350 pages; \$27.95.

THE effects of Edward Snowden's heist of secrets from America's National Security Agency (NSA) in 2013 can be divided into the good, the bad and the ugly, writes Edward Jay Epstein in a meticulous and devastating account of the worst intelligence disaster in the country's history, "How America Lost Its Secrets".

# The new Paranoid Internet Service Architecture

The entire concept of open network transactions is now over

We are shifting into an environment where user information is deliberately withheld from the network, withheld from the platform and even withheld from other applications

We circulate large self-contained applications that attempt to insulate themselves completely from the host platform

Application Service Providers see the platform provider as representing a competitive interest in the user, and they want to prevent information leakage from their application to the platform

Application Service Providers see other applications as as representing a competitive interest in the user, and they want to prevent information leakage from their application to other applications in the same platform



# Use Multipath TCP to create backup connections for iOS

If you're a network administrator, you can use Multipath TCP with iOS to strengthen connectivity to your destination host.

iOS supports [Multipath TCP \(MPTCP\)](#) and allows an iPhone or iPad to establish a backup TCP connection to a destination host over a cellular data connection.

Welcome to Project Fi,  
a wireless service from Google

By designing around how people live, we've created a service that feels like it was built for you.

WATCH THE INTRO ▶

## Project Fi intelligently shifts between multiple networks.

Project Fi automatically connects you to the best available signal, whether that's Wi-Fi or one of our three 4G LTE partner networks.

SEE COVERAGE DETAILS →



**The Chromium Projects**

Home  
Chromium  
Chromium OS

**Quick links**  
Report bugs  
Discuss  
Sitemap

**Other sites**  
Chromium Blog  
Google Chrome  
Extensions  
Google Chrome Frame

Except as otherwise noted, the content of this page is licensed under a [Creative Commons Attribution 2.0 license](#), and examples are licensed under the BSD License.

**QUIC, a multiplexed stream transport over UDP**

QUIC is a new transport which reduces latency compared to that of TCP. On the surface, QUIC is very similar to TCP+TLS+HTTP/2 implemented on UDP. Because TCP is implemented in operating system kernels, and middlebox firmware, making significant changes to TCP is next to impossible. However, since QUIC is built on top of UDP, it suffers from no such limitations.

**Key features of QUIC over existing TCP+TLS+HTTP2 include**

- Dramatically reduced connection establishment time
- Improved congestion control
- Multiplexing without head of line blocking
- Forward error correction
- Connection migration

**Documentation**

- [QUIC overview](#)
- [QUIC FAQ](#)
- [QUIC wire specification](#)
- [QUIC crypto design doc](#)
- [Getting started with the QUIC toy client and server](#)
- [QUIC tech talk](#)
- [QUIC Discovery](#)
- [QUIC FEC v1](#)



# Use Multipath TCP to create backup connections for iOS

Use Multipath TCP with iOS to create backup connections to a host.

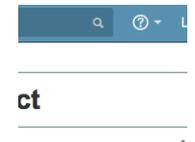
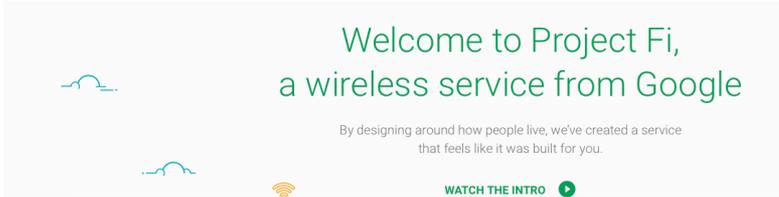
Use Multipath TCP connection to a host.

These technologies are already deployed, and enjoy significant use in today's network

They break down the concept of a "session" and display the encrypted traffic across multiple networks, and even multiple protocols

They use opportunistic encryption to limit third party access to information about users' actions

The result is that only the endpoints see the entirety of a session, while individual networks see disparate fragments of pseudo-sessions

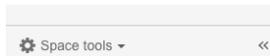


### Table of Contents

- QUICK START GUIDE
- LATEST NEWS: Watch the IETF 97 EDU DNS Privacy Tutorial!
- Overview of DNS Privacy Status
- DPRIVE Working Group
- Reference Material
- Thanks
- Contact

Home  
Quickstart  
Reference  
Contact

- [Getting started with the QUIC toy client and server](#)
- [QUIC tech talk](#)
- [QUIC Discovery](#)
- [QUIC FEC v1](#)



Current contributors to this project include Simon St. Laurent, NLnet Labs and No Mountain Software.

sinodun Spaces

**S** DNS Privacy

Pages 0

DNS Privacy

Links

- DNS Privacy Project homepage
- DPRIVE
- getdns
- NLnet Labs
- Sinodun

PAGE TREE

- DNS Privacy - The Problem
- DNS Privacy - Ongoing Work
- IETF DNS Privacy Tutorial
- DNS Privacy daemon - Stubby
- DNS-over-TLS test servers
- Using a TLS proxy
- DNS-over-TLS implementations
- DNS-over-TLS clients
- DNS Privacy reference material
- Project Notes

Space tools <<

## DNS Privacy Project

### DNS Privacy Project Homepage

Created by Sara Dickinson, last modified on Dec 13, 2016

#### Welcome to the DNS Privacy project home page

This site is the home of a collaborative open project to promote, implement and deploy DNS Privacy. The goals of this project include:

- Raising awareness of the issue of DNS Privacy
- Empowering users to take advantage of DNS Privacy tools and resources (client applications, DNS Privacy resolvers)
- Evolving the DNS to support DNS Privacy in particular developing new DNS Protocol standards
- Working towards full support for DNS Privacy in a range of Open Source DNS implementations including: getdns, Unbound, NSD, BIND and Knot (Auth and Resolver)
- Co-ordinating deployment of DNS Privacy services and documenting operational practices

Current contributors to this project include Sinodun IT, NLnet Labs and No Mountain Software.

Table of Contents

Google Public DNS

Secure <https://dns.google.com/query?name=www.europol.europa.eu&typ...>

Google Public DNS

RR Type

DNSSEC  ON  EDNS Client Subnet

Result for www.europol.europa.eu/A with DNSSEC validation:

```
{
  "Status": 0,
  "TC": false,
  "RD": true,
  "RA": true,
  "AD": false,
  "CD": false,
  "Question": [
    {
      "name": "www.europol.europa.eu.",
      "type": 1
    }
  ],
  "Answer": [
    {
      "name": "www.europol.europa.eu.",
      "type": 1,
      "TTL": 1799,
      "data": "136.243.175.123"
    }
  ]
}
```

You may also resolve directly at: <https://dns.google.com/resolve?name=www.europol.europa.eu>.

Even going the DNS is "dark!"

# The Bottom Line

It's no longer just an issue with IPv4 and NATs and a visible reluctance to shift to IPv6

Networks, platforms and applications now regard each other with mutual suspicion

Platforms seek to hide users' activities from the network

Applications seek to hide their information from the platform and from other applications

The DNS is sealing itself into private tunnels that resist external examination, intervention and intervention

"Sessions" are being deconstructed into opaque fragments

Opportunistic encryption is being applied ubiquitously

Its not just "the IPv6  
transition" any more

These are not just temporary steps to make IPv4 last longer for the transition to IPv6

Even if we complete the transition to an all-IPv6 Internet, this paranoia, complexity and deliberate obfuscation will not go away

This is now the Internet we have to live with

We are never coming back from here – this is the new “ground state” for the Internet!

Does it ever get easier?

is there light at the end of this tunnel?

No!

Thank You!

This pack: <http://bit.ly/2iCcQGK>

Me: [gih@apnic.net](mailto:gih@apnic.net)

