

# IPv6 Transition and CGNs: A Progress Report

Geoff Huston

Chief Scientist

APNIC


# The story so far..

- The status of the transition to IPv6 is not going according to the original plan:
  - Over the past two years APNIC (Asia Pacific) and the RIPE NCC (Europe and the Middle East) have exhausted their supplies of general use of IPv4 Addresses
  - ARIN have some 14 months to go, but this assumes a very constrained availability over this period
  - We we meant to have IPv6 deployed by now. This is not happening.

# The story so far..

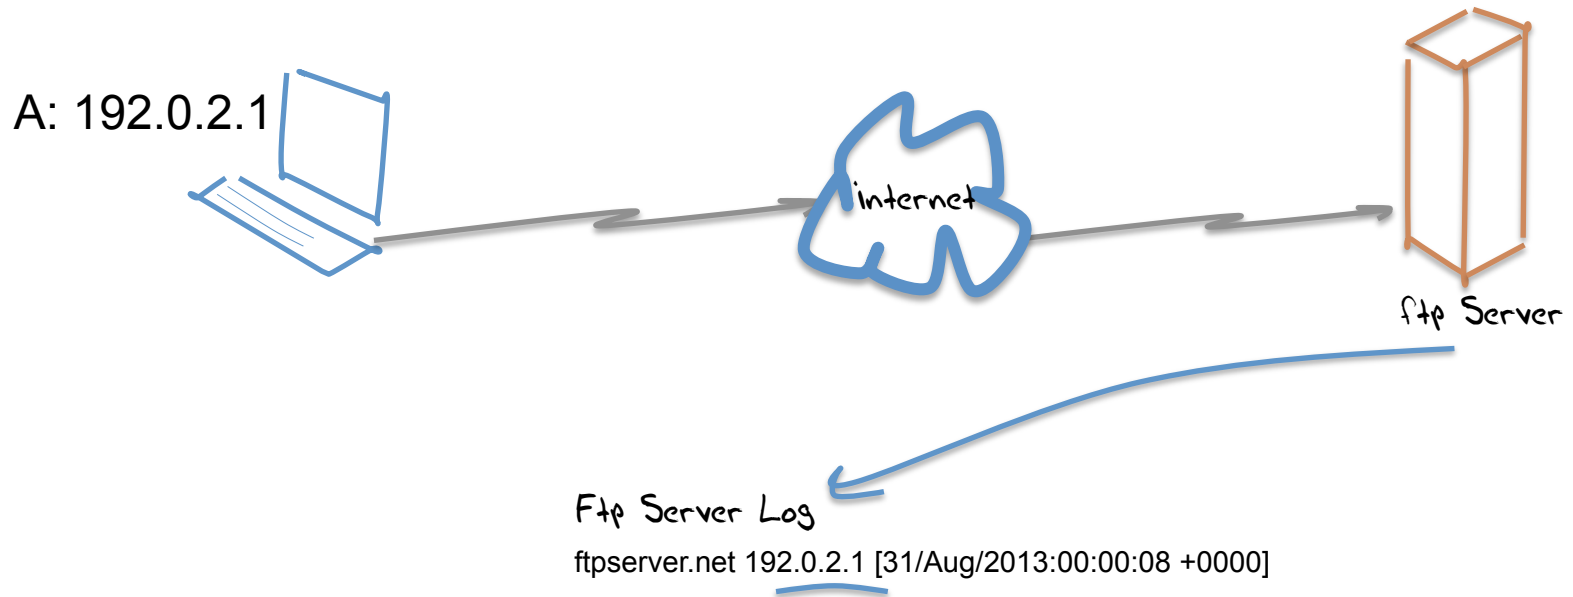
- What we are seeing is the increasing use of Carrier Grade NATs as a means of extending the useable life of the IPv4 Internet while we are still waiting for IPv6 to be viable in its own right
- This has some significant implications for LEA functions, principally in traceback and Internet meta-data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

# Traceback - Version 1




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

# Traceback - Version 1



No web - just ftp  
And server logs were pretty basic

# Traceback - Version 1



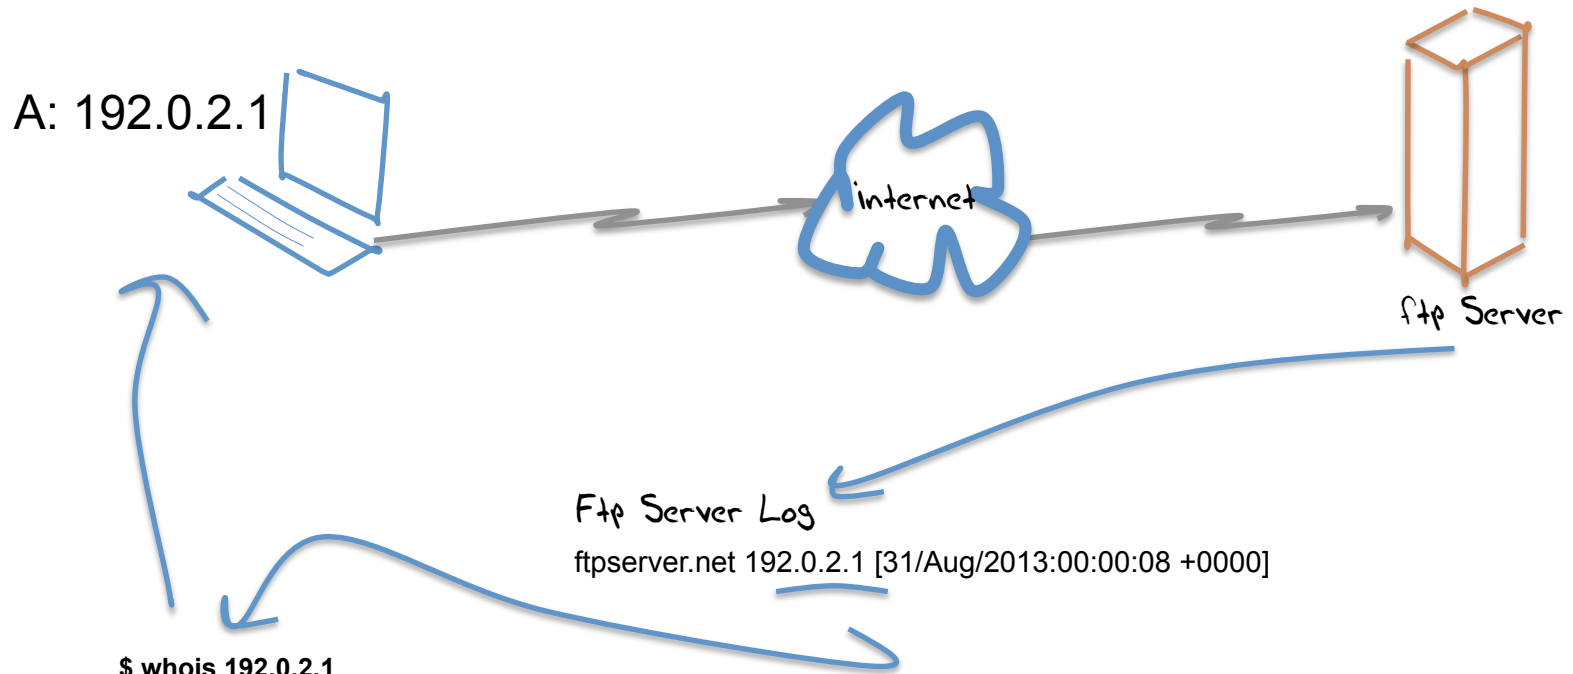
Ftp Server Log  
ftpsrvr.net 192.0.2.1 [31/Aug/2013:00:00:08 +0000]

**\$ 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!

# Traceback - Version 1



**\$ whois 192.0.2.1**

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

# 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

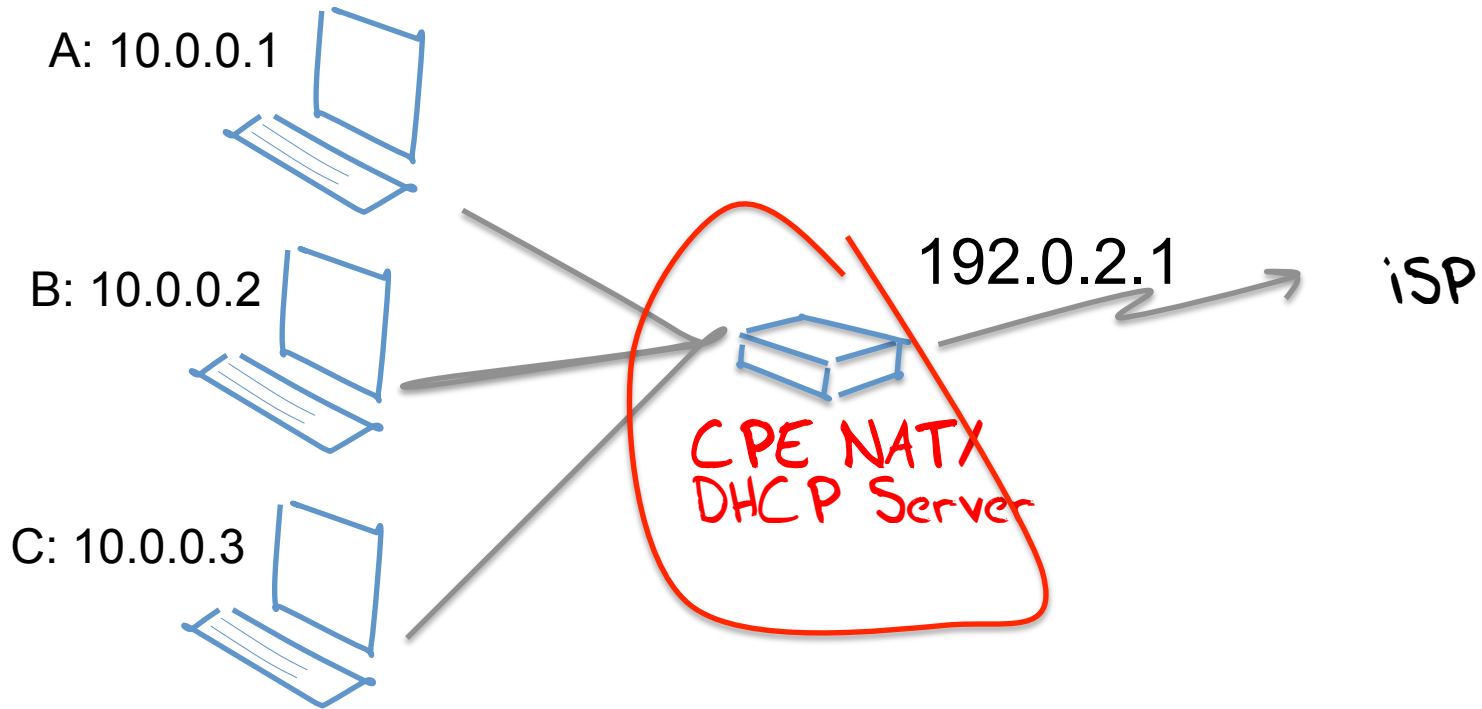


# Assumptions:

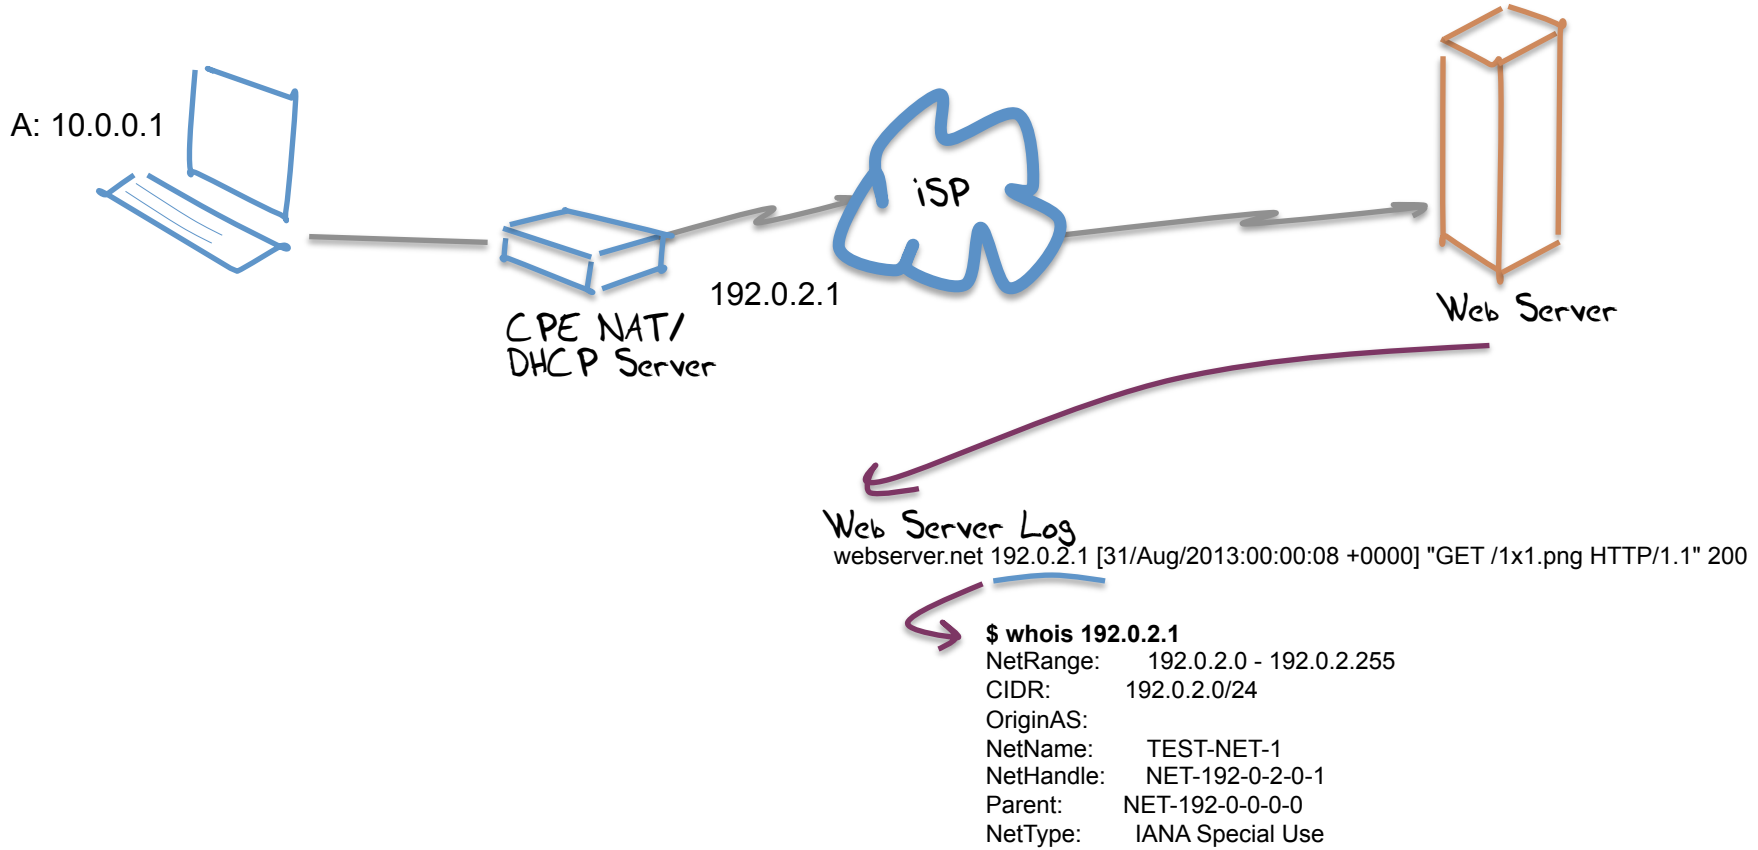
- Each end site used a stable IP address
- Each address range was recorded together with the end user
- Each end device was associated with a stable IP address
- Traceback was possible to the IP address

*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



# Traceback - Version 2

iSP RADIUS Log

15/Aug/2013:18:01:02: user XXX IP: 192.0.2.1



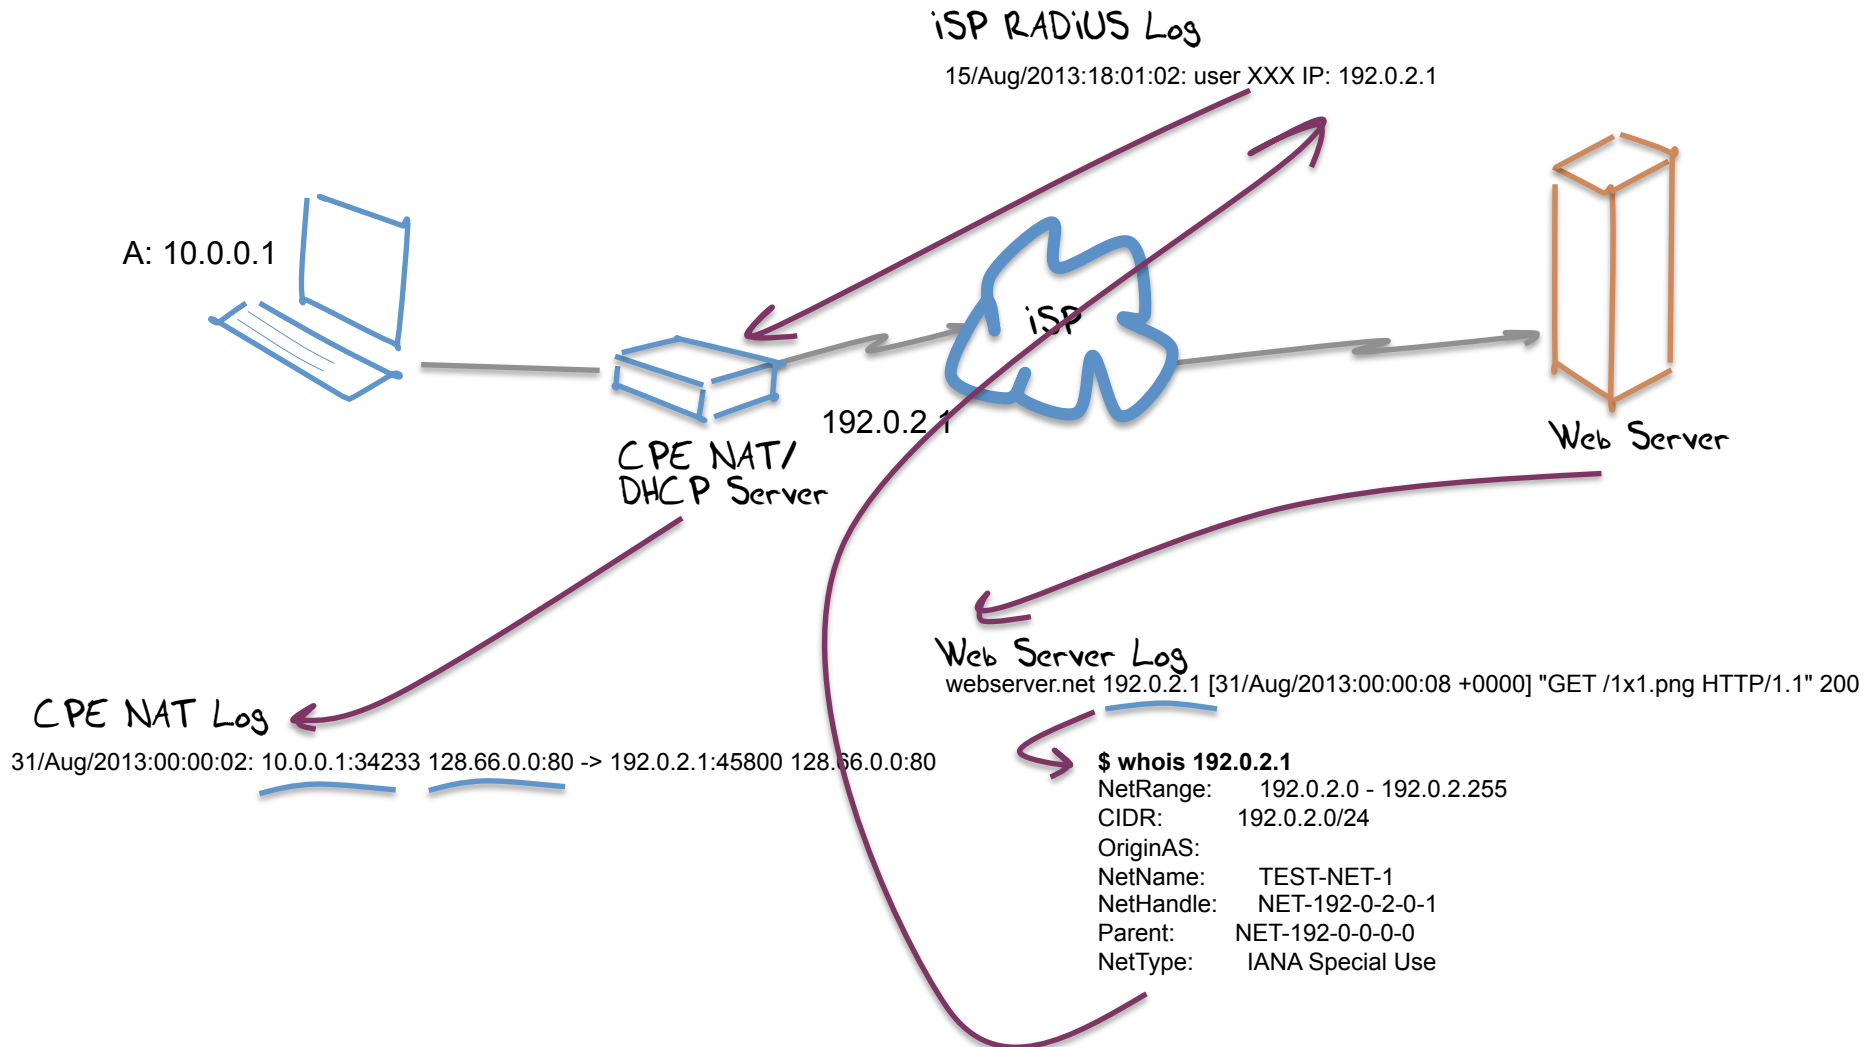
Web Server Log

webserver.net 192.0.2.1 [31/Aug/2013:00:00:08 +0000] "GET /1x1.png HTTP/1.1" 200

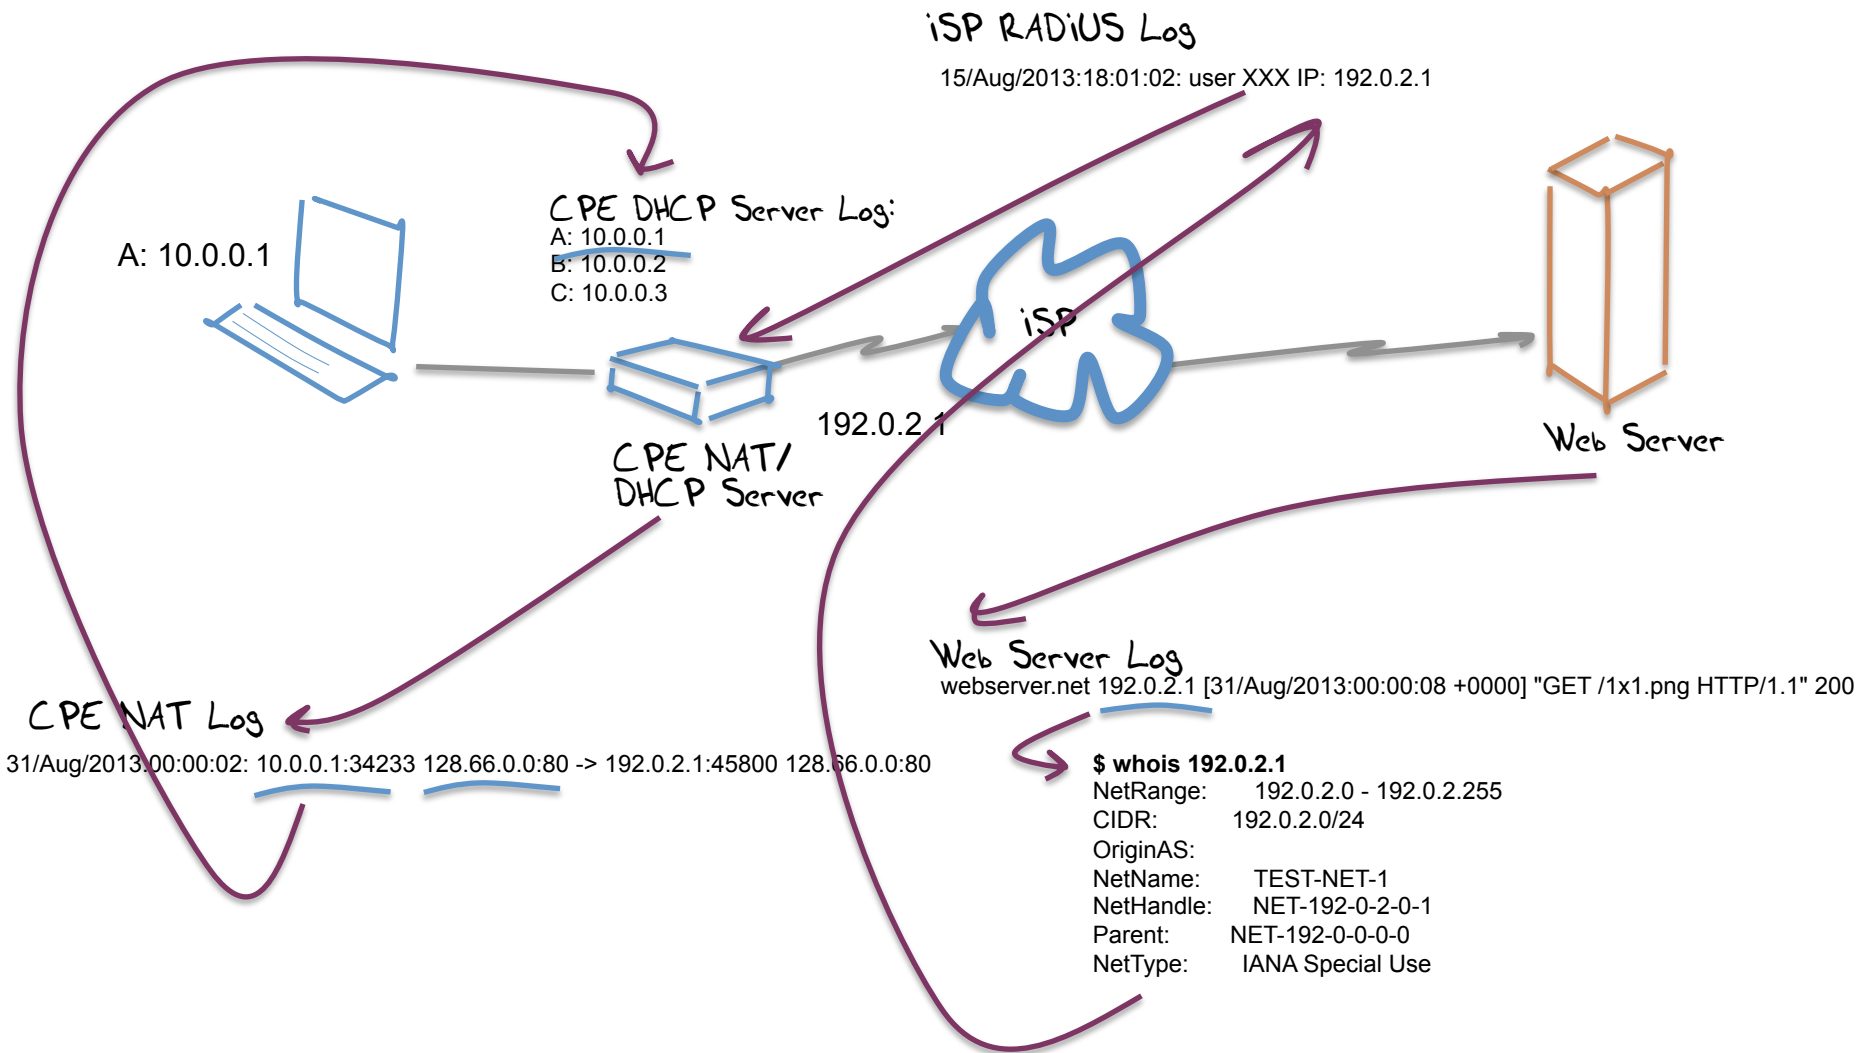
\$ whois 192.0.2.1

NetRange: 192.0.2.0 - 192.0.2.255  
CIDR: 192.0.2.0/24  
OriginAS:  
NetName: TEST-NET-1  
NetHandle: NET-192-0-2-0-1  
Parent: NET-192-0-0-0-0  
NetType: IANA Special Use

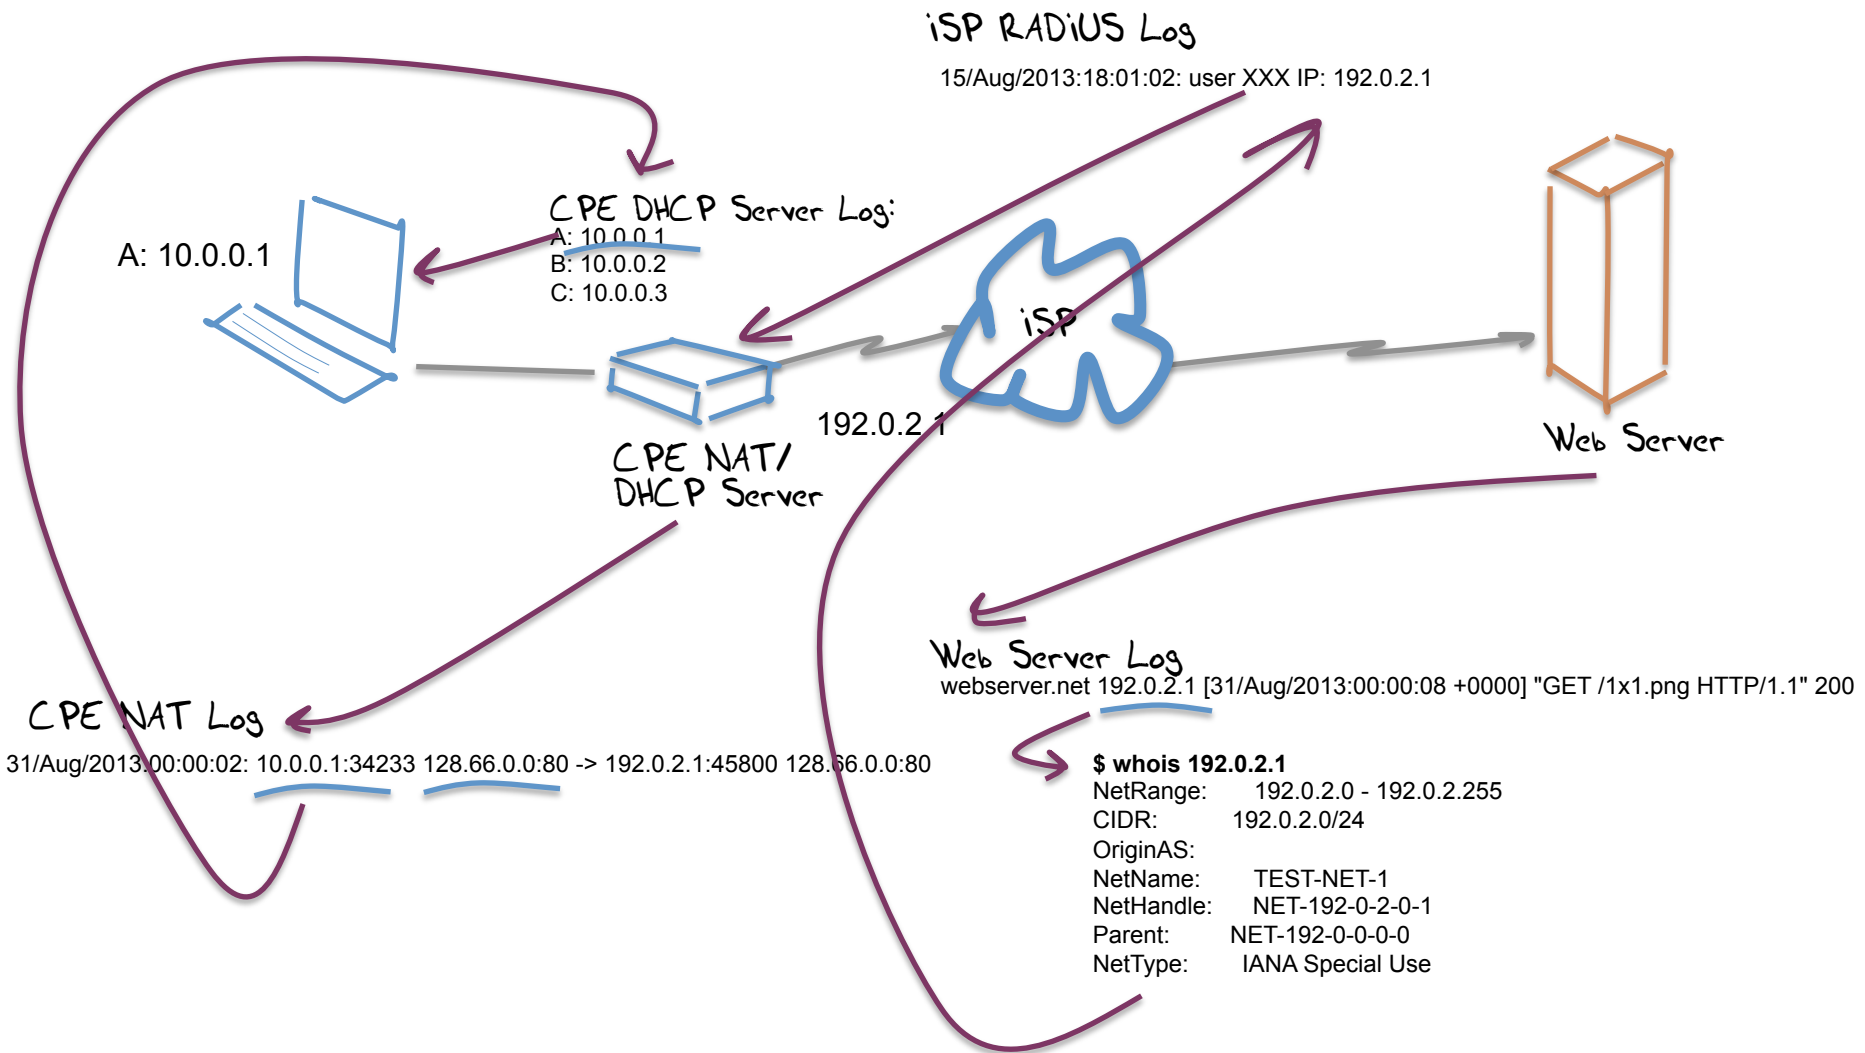
# Traceback - Version 2



# Traceback - Version 2



# Traceback - Version 2



ISP RADIUS Log

15/Aug/2013:18:01:02: user XXX IP: 192.0.2.1

CPE DHCP Server Log:


A: 10.0.0.1  
B: 10.0.0.2  
C: 10.0.0.3

A: 10.0.0.1

CPE NAT/  
DHCP Server

192.0.2.1

ISP



Web Server

Web Server Log

webserver.net 192.0.2.1 [31/Aug/2013:00:00:08 +0000] "GET /1x1.png HTTP/1.1" 200

CPE NAT Log

31/Aug/2013:00:00:02: 10.0.0.1:34233 128.66.0.0:80 -> 192.0.2.1:45800 128.66.0.0:80

\$ whois 192.0.2.1

NetRange: 192.0.2.0 - 192.0.2.255  
CIDR: 192.0.2.0/24  
OriginAS:  
NetName: TEST-NET-1  
NetHandle: NET-192-0-2-0-1  
Parent: NET-192-0-0-0-0  
NetType: IANA Special Use

# 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



# Assumptions


- 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
- Further traceback within the end site relies on access to the site's DHCP and NAT logs (assuming that there are such logs!)

# Assumptions

- Traceback to an end site is keyed by an IP address and a date/time
    - Requires access to the ISP's AAA logs to identify the end site
  - Further traceback within the end site relies on access to the site's DHCP and NAT logs (assuming that there are such logs!)
- This model is commonly used today*

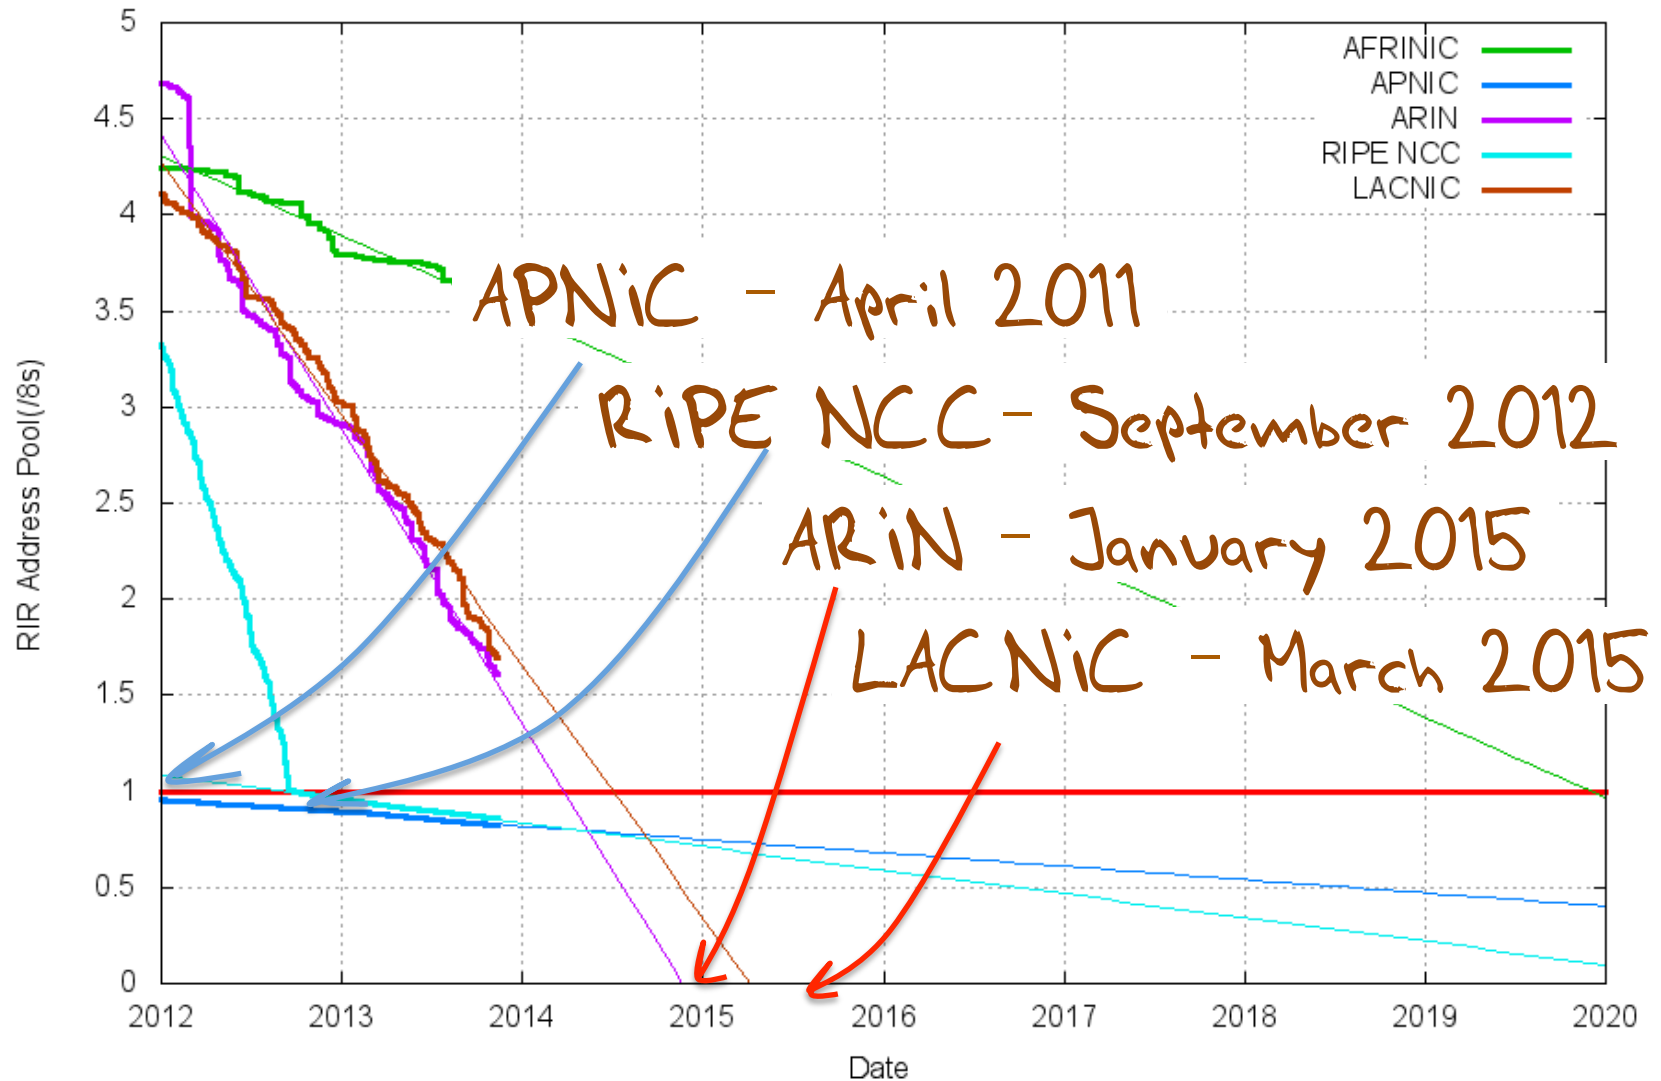
# IPv4 Address Exhaustion

RIR IPv4 Address Run-Down Model



# IPv4 Address Exhaustion

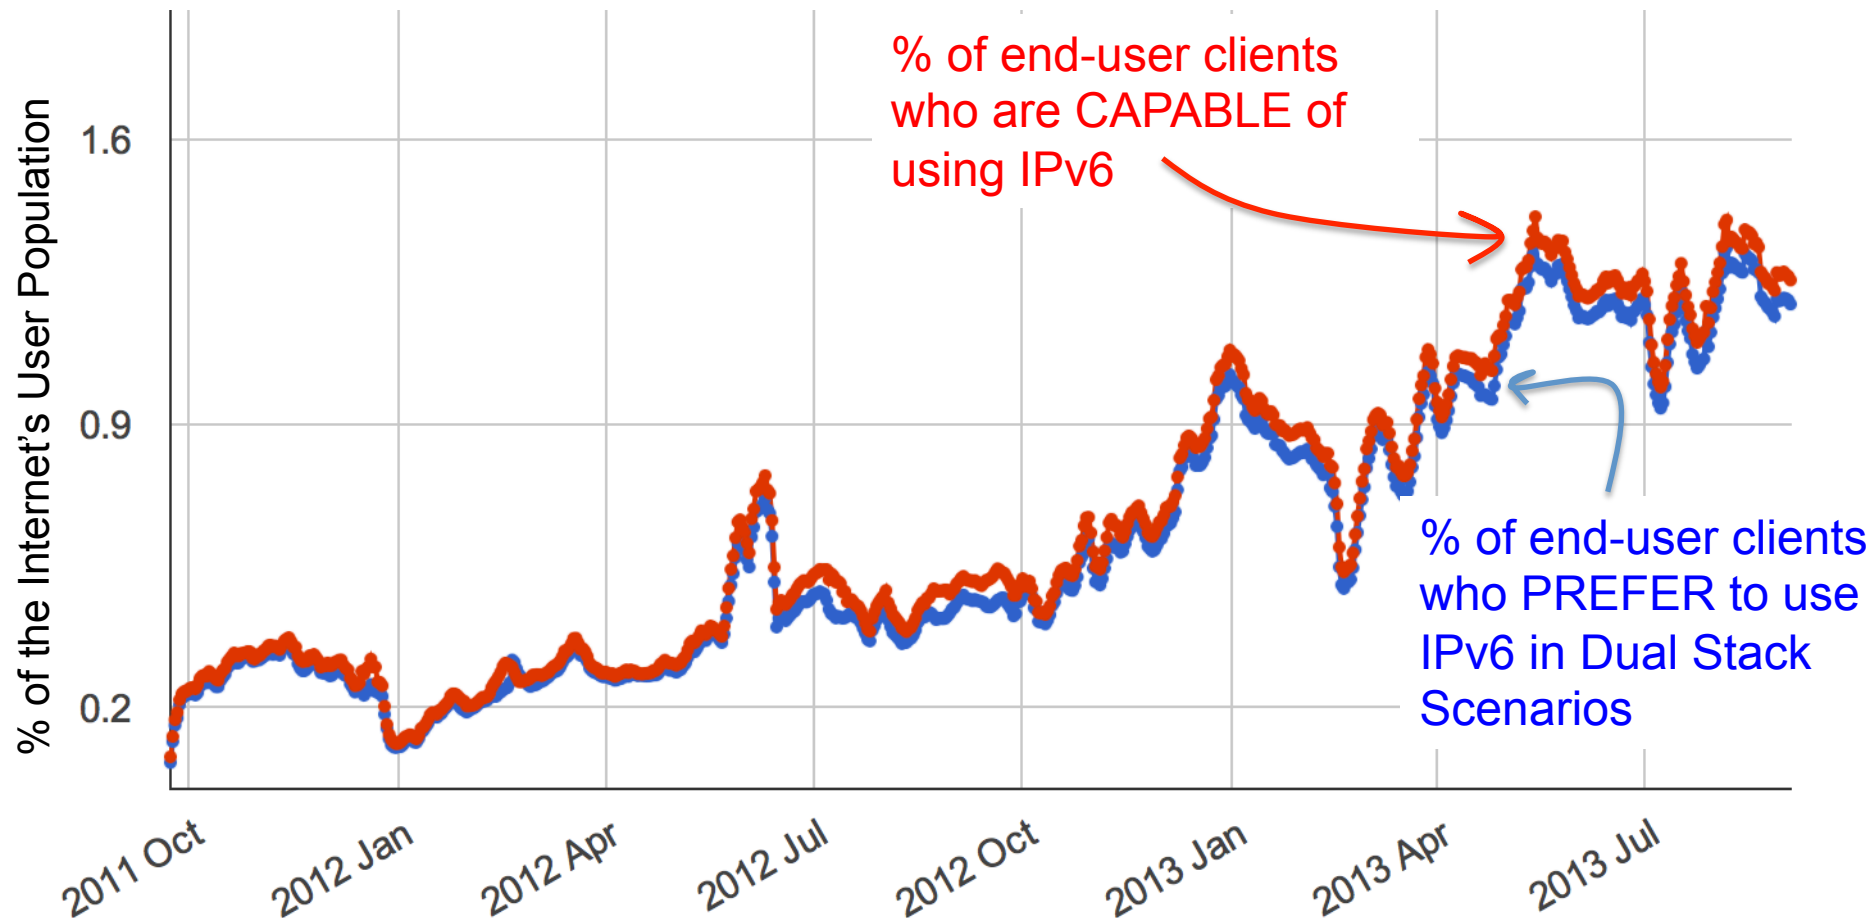
RIR IPv4 Address Run-Down Model



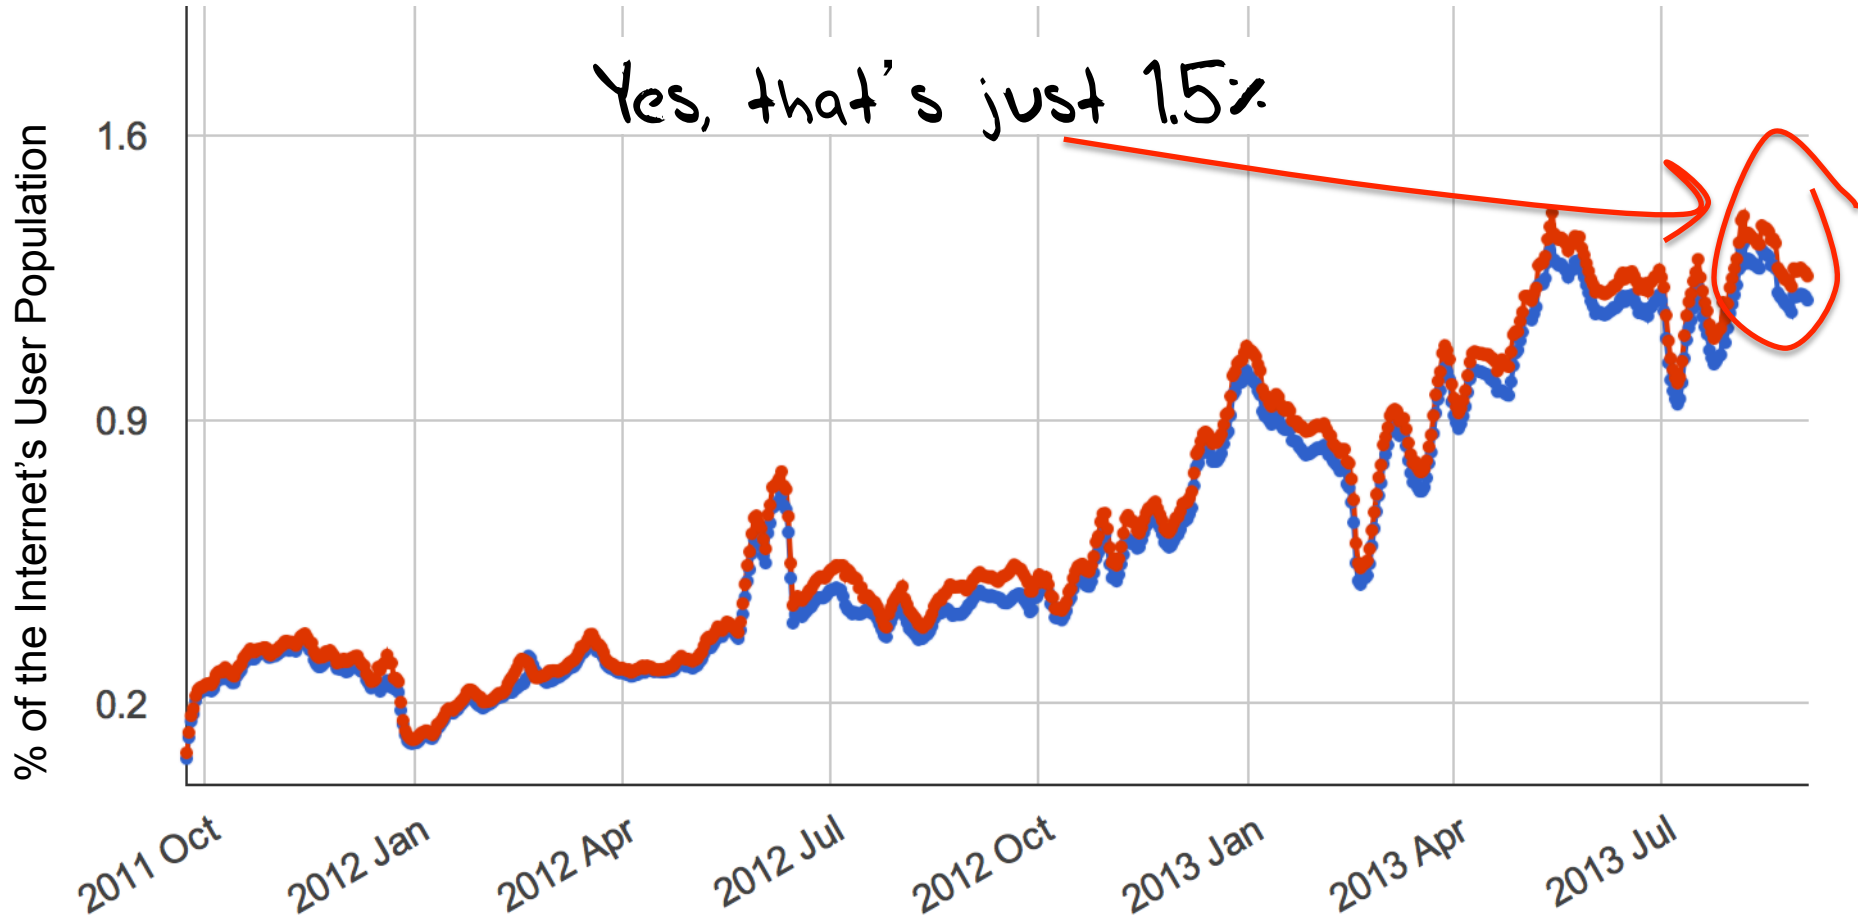
# IPv4 Address Exhaustion

- We are seeing ISPs in Asia Pac and Europe run out of IPv4 addresses
- And while the plan was to have IPv6 fully deployed by now, that has not happened

# IPv6 Penetration in the Internet



# IPv6 Penetration in the Internet



# IPv4 Address Exhaustion

What are ISP's doing in response?

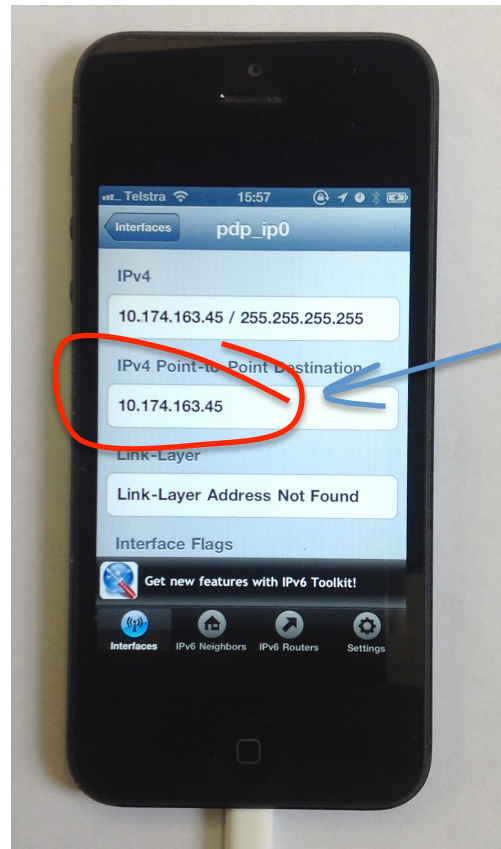
- It's not viable to switch over to IPv6 yet
- But the supply of further IPv4 addresses to fuel service platform growth has dried up / will soon dry up
- How will ISPs continue to offer IPv4 services to customers in the interim?



# 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.

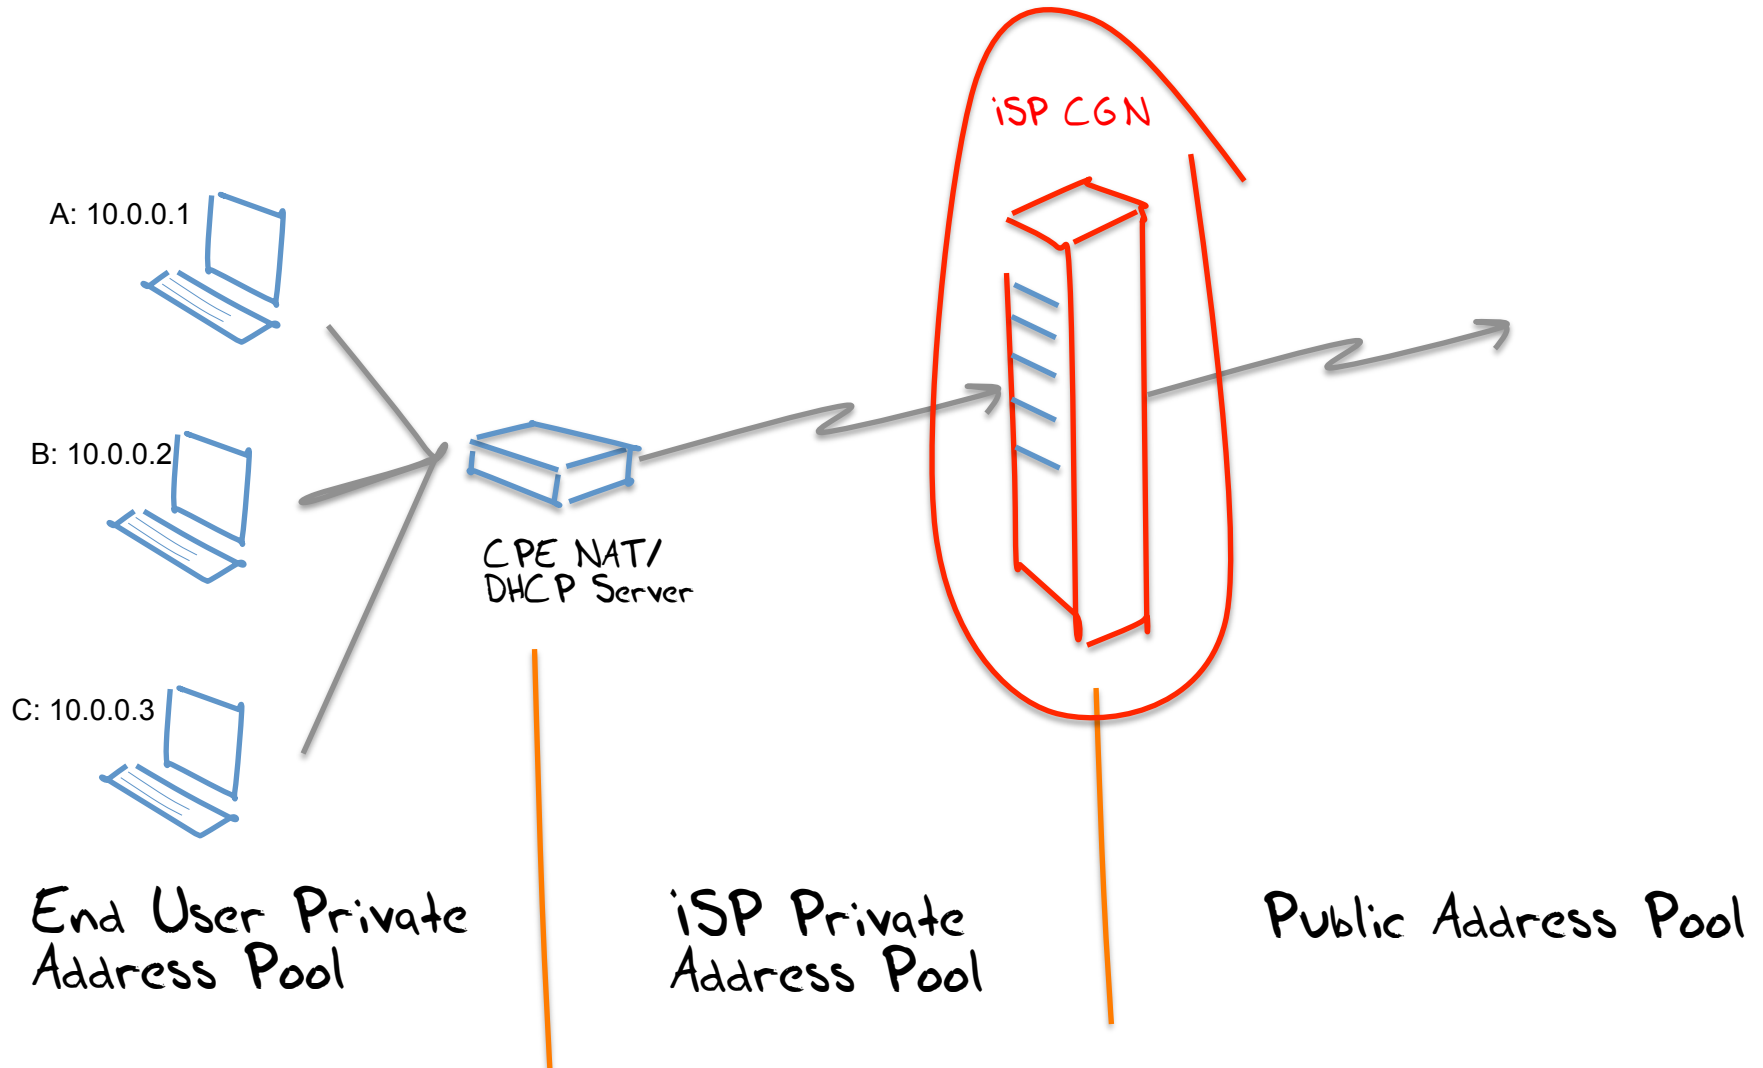


judgecorp writes


"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."



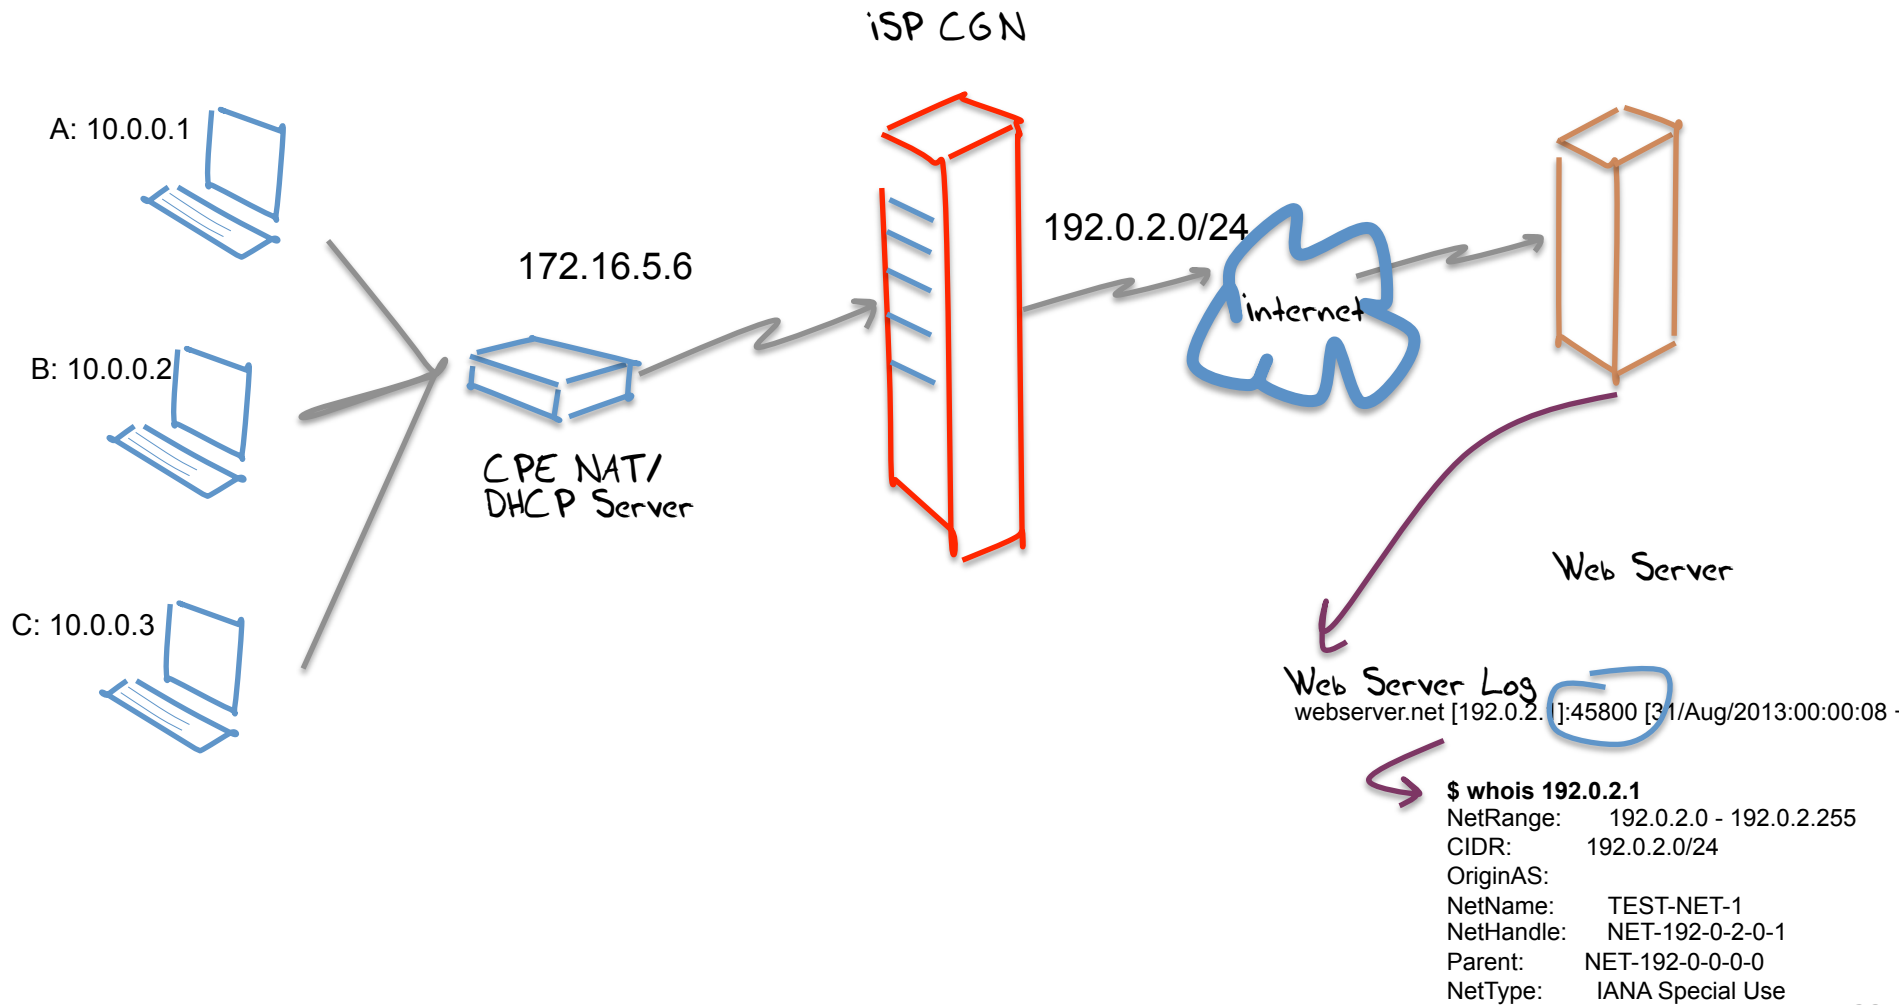
# NATs + CGNs



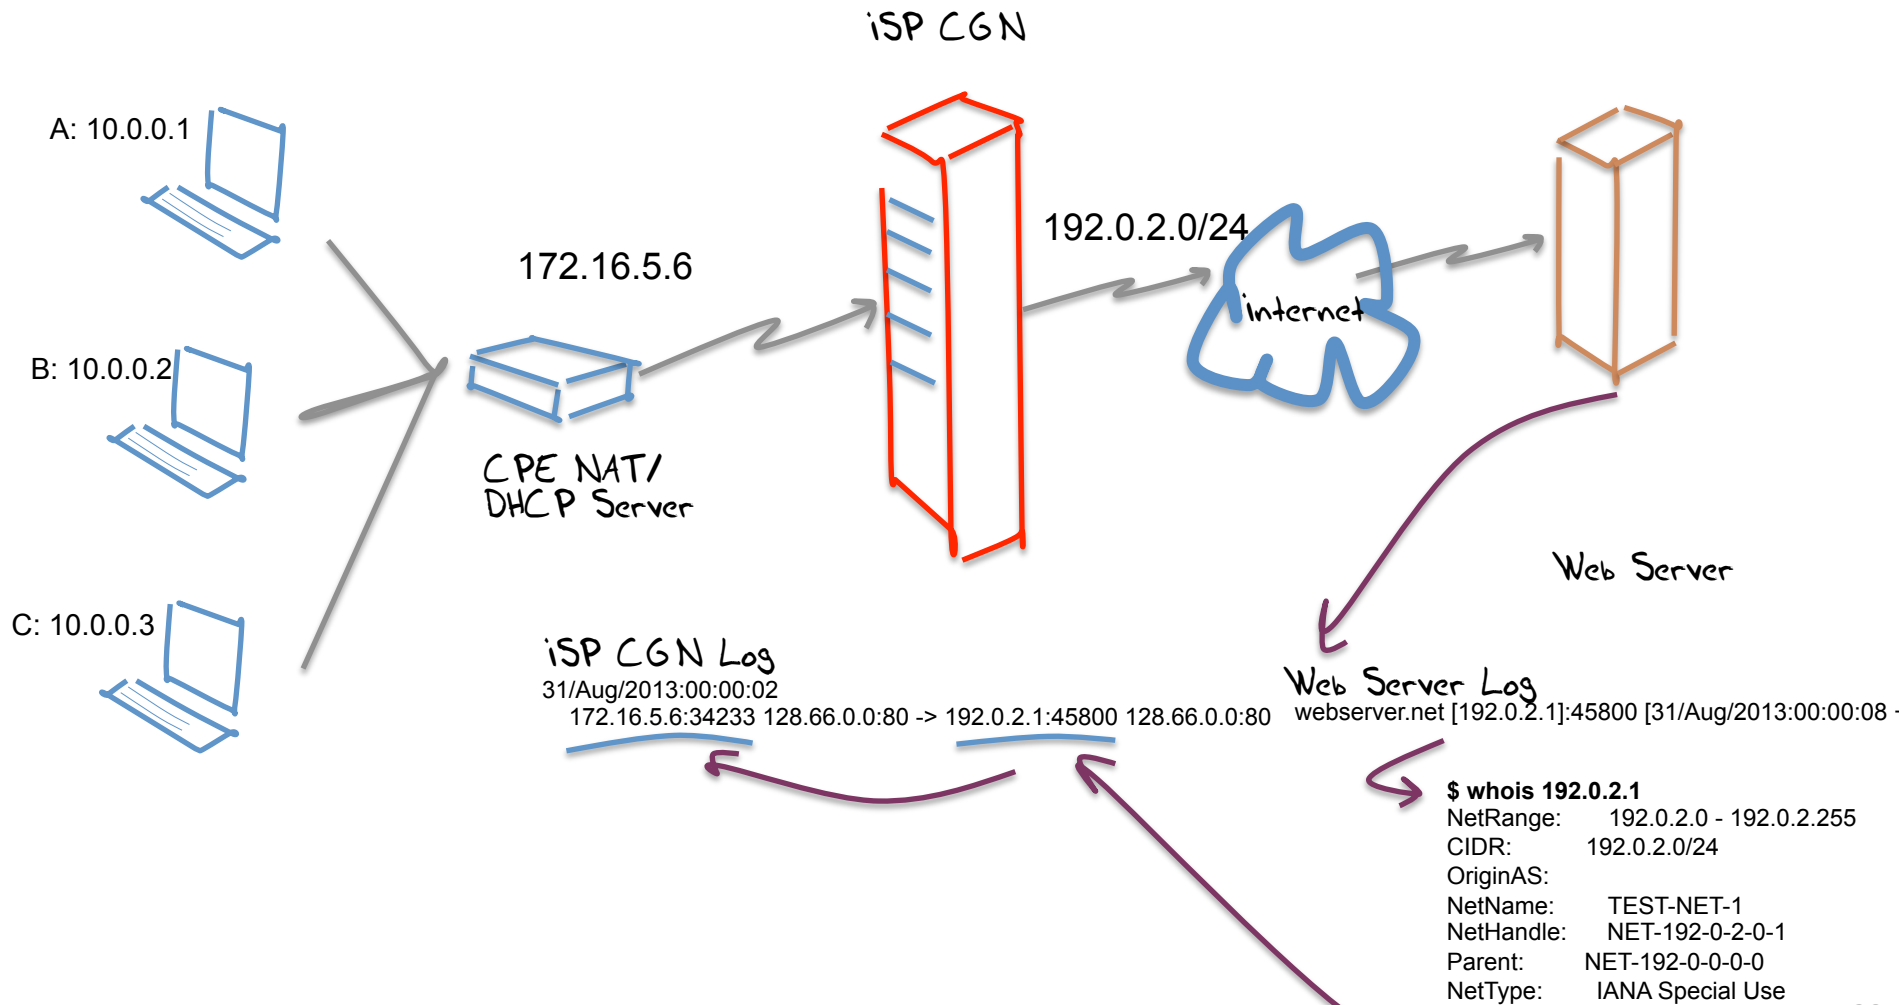
# NATs + CGNs + Connections



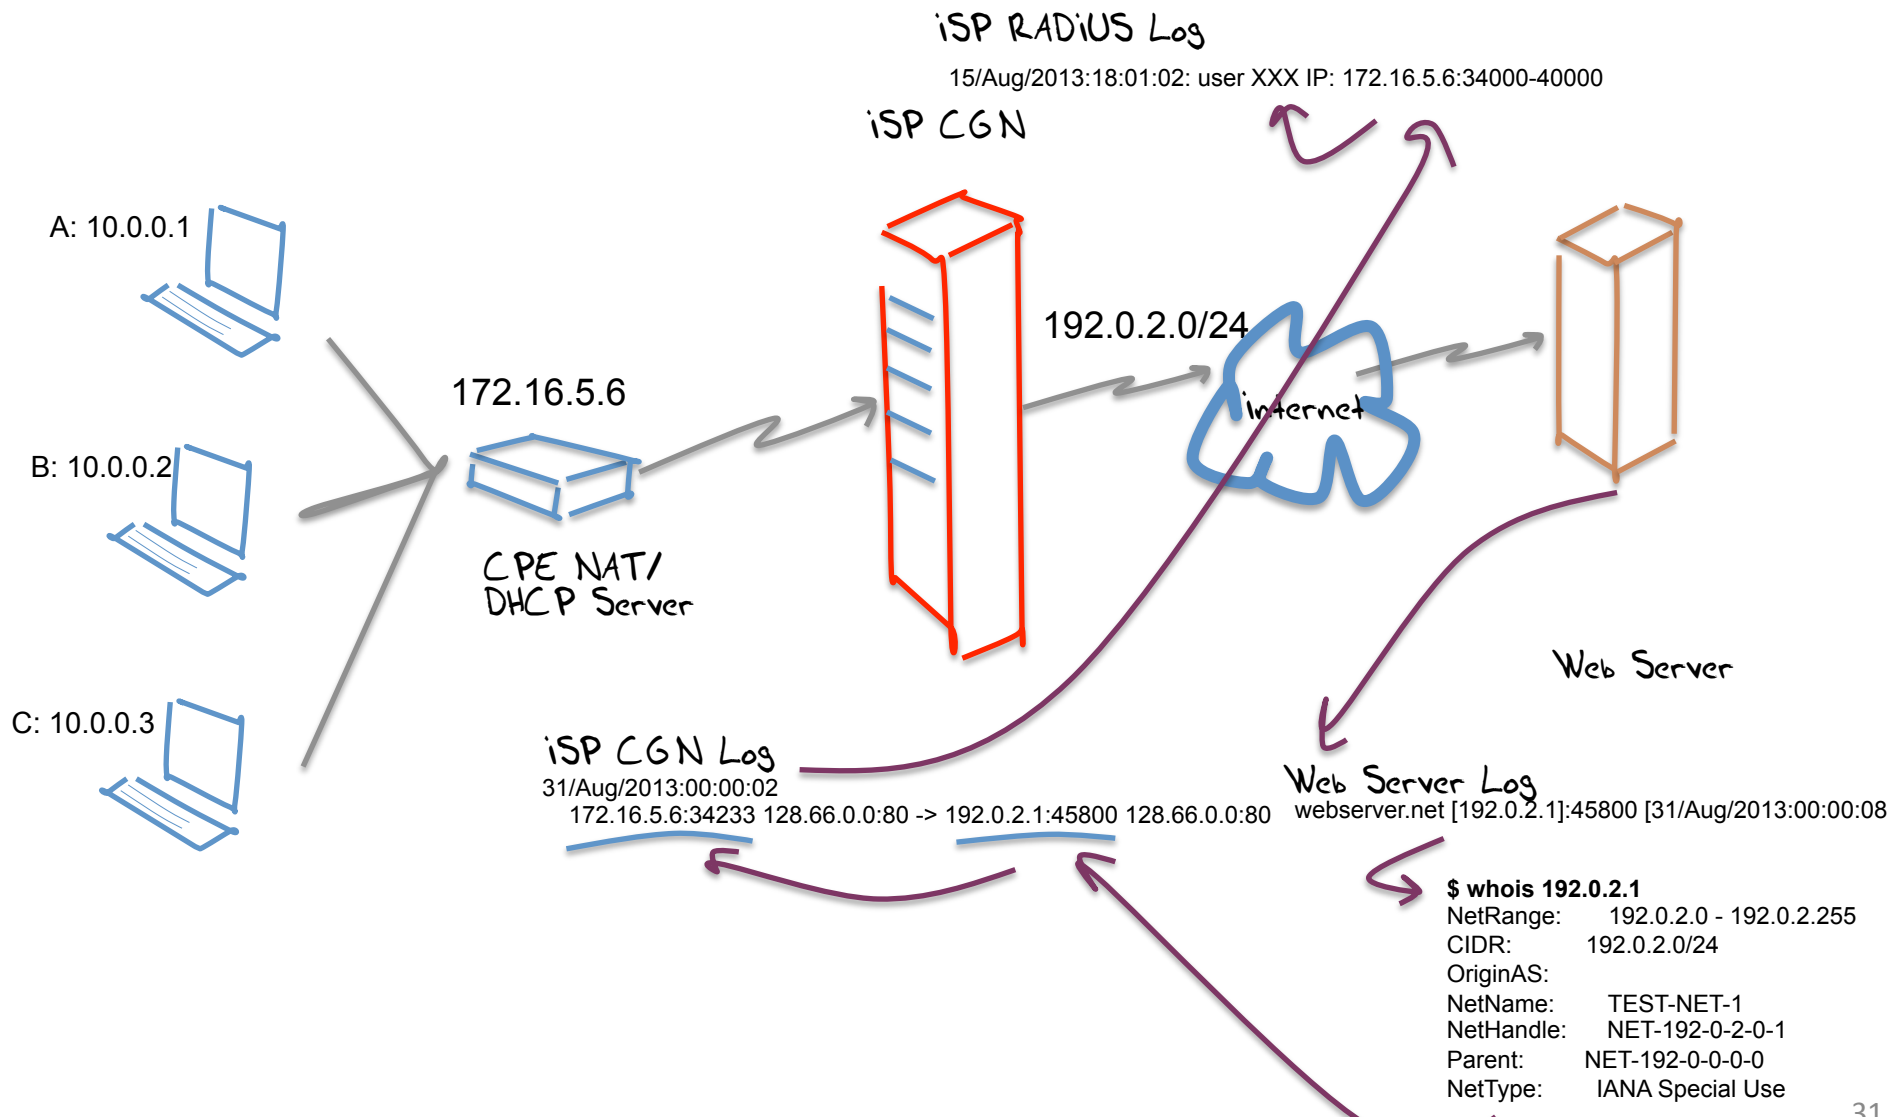
# Traceback - Version 3



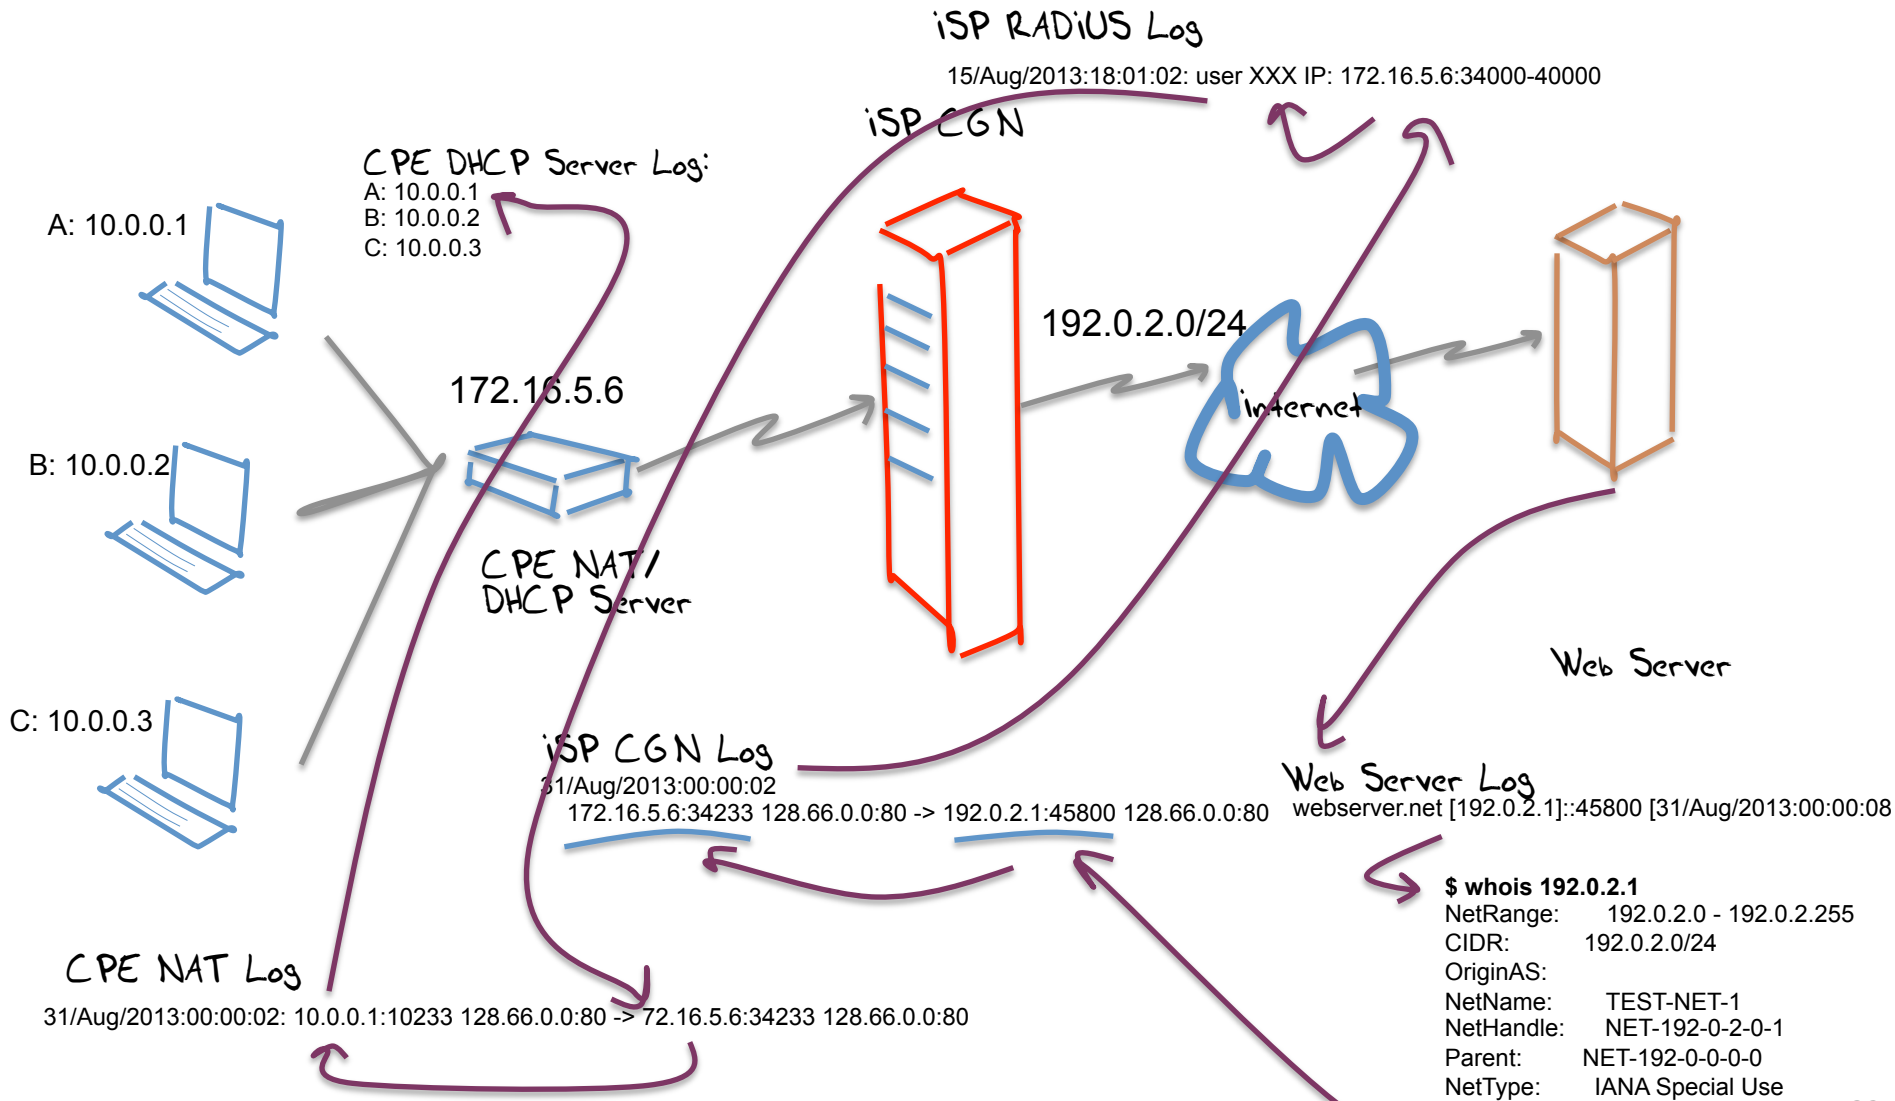
# Traceback - Version 3



# Traceback - Version 3

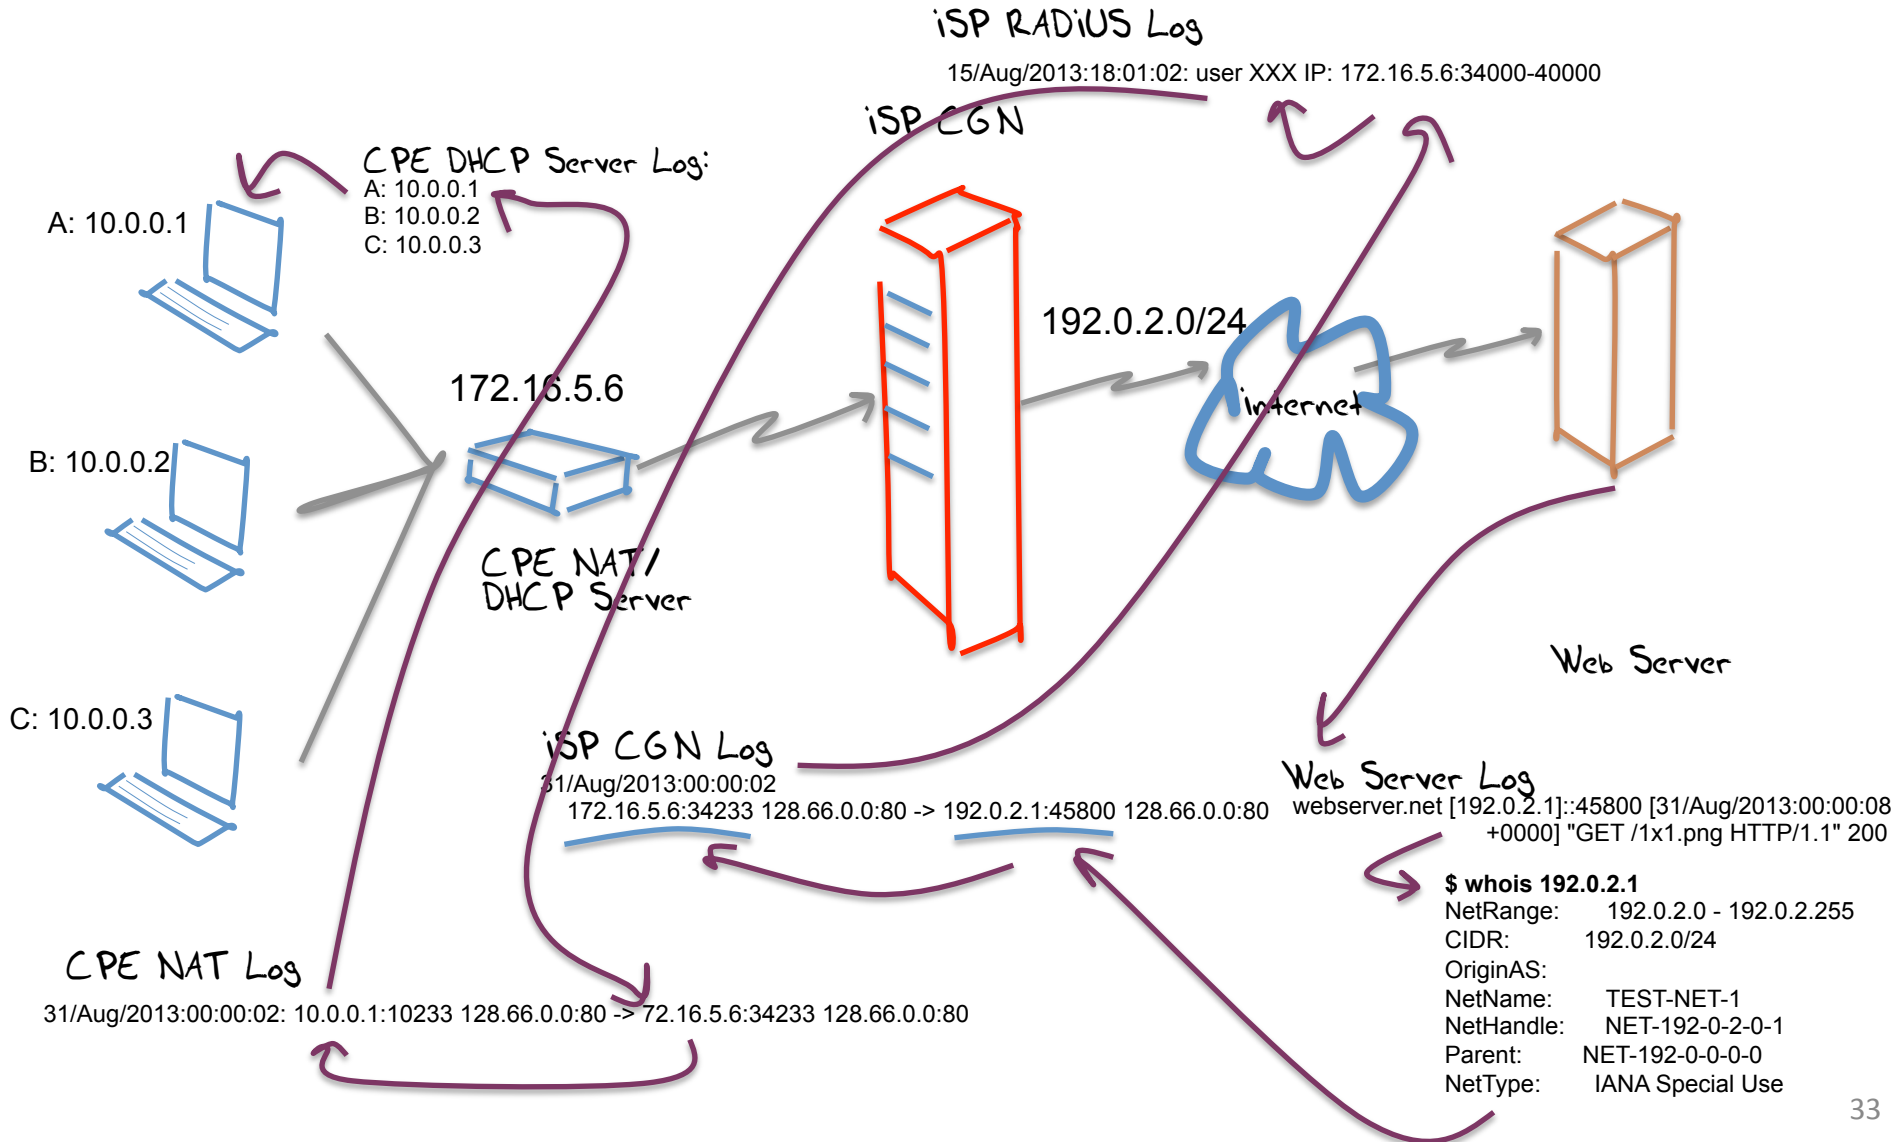


# Traceback - Version 3





# 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 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
- Further traceback within the end site relies on access to the site's DHCP and NAT logs  
(assuming that there are such logs!)

# 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..

# Challenges:

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 LEAs?

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
- trace back 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 coherent meaning in terms of end party identification.

# What does this mean for LEAs?

In a world of densely deployed CGN  
the IP address loses coherence  
party identification

*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 NAT444 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 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




# 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

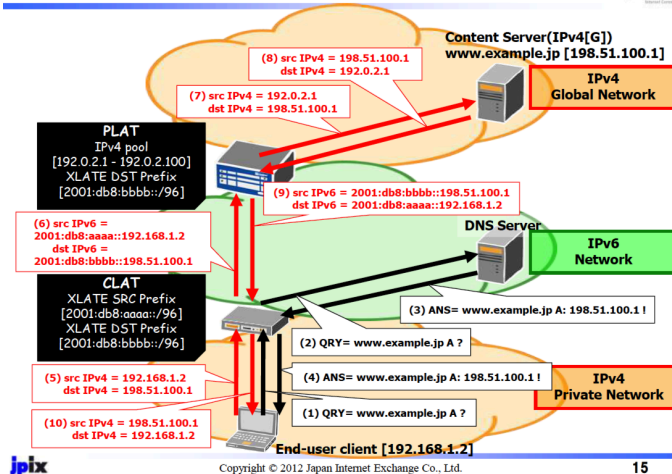


jpix

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 LEAs?

The risk we are running at the moment is that in the near future there will 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 years”*

# 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:

for every hardware element you now need to hold the time and the **precise** transforms that were applied to a 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 scalable?

Is it even useful any more?

# Traceback in tomorrow'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

# Even harder?

All this will make the entire record and trace problem for ISPs and LEAs far harder than it is at present!

At some point along this path of escalating network complexity and diversity its likely that our networks will be simply be unable to track 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 away these CGNs and their associated records
- And the entire exercise of record keeping and traceback gets a whole lot easier

# Traceback - IP Version 6



## Web Server Log


webserver.net 2001:db8:1:0:426c:8fff:fe35:45a8 [31/Aug/2013:00:00:08 +0000] "GET /1x1.png HTTP/1.1" 200

**\$ whois 2001:db8:1:0:426c:8fff:fe35:45a8**  
inet6num: 2001:0DB8::/32  
netname: IPV6-DOC-AP  
descr: IPv6 prefix for documentation purpose  
country: AP

# Traceback - IP Version 6

iSP AAA Log

15/Aug/2013:18:01:02: user XXX IP: 2001:db8:1::/56

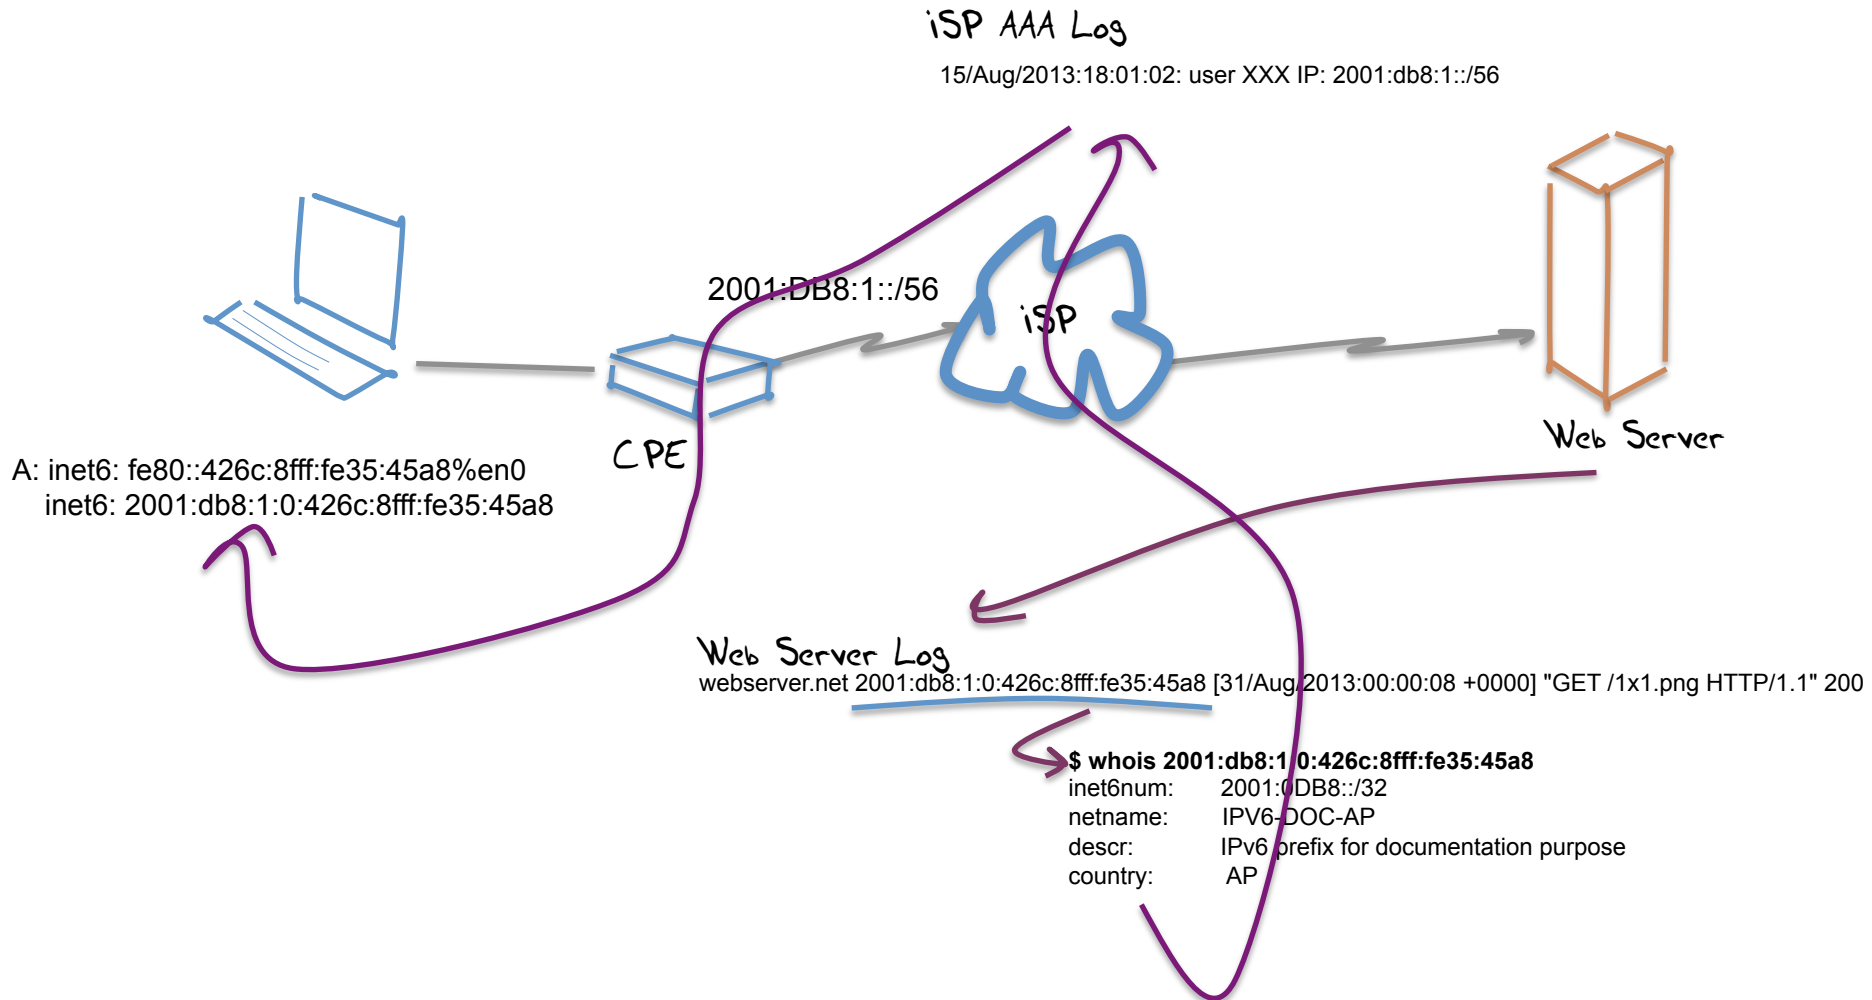


Web Server Log

webserver.net 2001:db8:1:0:426c:8fff:fe35:45a8 [31/Aug/2013:00:00:08 +0000] "GET /1x1.png HTTP/1.1" 200

```
$ whois 2001:db8:1:0:426c:8fff:fe35:45a8
inet6num: 2001:0DB8::/32
netname: IPV6-DOC-AP
descr: IPv6 prefix for documentation purpose
country: AP
```

# 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 <sup>mostly</sup> easy again. Right?

Firstly, there are private addresses creeping into IPv6

The use of ULAs is taking hold in IPv6 networks

Apple use ULA addresses in “back to my Mac”, and many enterprise networks use ULAs internally because network renumbering in IPv6 is hard.

So we can expect more ULAs + V6NATs, not less.

Which, if they leak out to the public Internet, would not be so good

IPv6 makes it <sup>mostly</sup> easy again. Right?

Secondly, IPv6 Privacy Addresses introduce ephemeral public IPv6 addresses into the mix

IPv6 Privacy addresses are used in Windows, Mac OSx, some variants of Linux

The privacy address cannot trace back to the original device. At best it can track back to the original site network, but no further

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

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!

# 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

And that's a big win over where IPv4 + CGNs is heading!

Thank You!