
cisco
You make possible

IPv4 Exhaustion: IPv6 Transition and NAT Architecture

For Service Providers

Rajiv Asati, Distinguished Engineer

BRKSPG-2602

cisco Live!
Barcelona | January 27-31, 2020

Cisco Webex Teams

Questions?

Use Cisco Webex Teams to chat with the speaker after the session

How

1. Find this session in the Cisco Events Mobile App
2. Click “Join the Discussion”
3. Install Webex Teams or go directly to the team space
4. Enter messages/questions in the team space

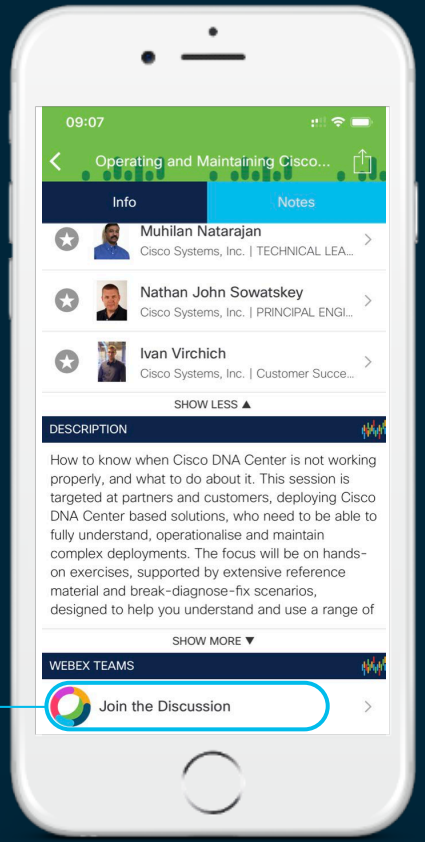

Nanog@nanog.org

Tuesday, February 27, 2018 at 11:32 AM

Show Details

Couple questions please. When you put thousands of customers behind a cgnat boundary, how do you all handle customer complaints about the following.

1 - for external connectivity to the customers premise devices, not being able to access web servers, web cameras, etc, in their premises?

2 - from the premise natted device, when customers go to a university or bank web site, how do you handle randomly changing ip addresses/ports that may occur due to idle time and session tear-down in nat table such that the bank website has issues with seeing your session ip change?

Hmm....CGNAT issue or something else ?

PlayStationNetwork blocking of CGNAT public addresses

etwork blocking of CGNAT public addresses

Simon Lockhart sent by NANOG

Friday, September 16, 2016 at 9:12 AM

To: nanog@nanog.org

All,

We operate an access network with several hundred thousand users. Increasingly we're putting the users behind CGNAT in order to continue to give them an IPv4 service (we're all dual-stack, so they all get public IPv6 too). Due to the demographic of our users, many of them are gamers.

We're hitting a problem with PlayStationNetwork 'randomly' blocking some of our CGNAT outside addresses, because they claim to have received anomalous, or 'attack' traffic from that IP. This obviously causes problems for the other legitimate users who end up behind the same public IPv4 address.

Despite numerous attempts to engage with PSN, they are unwilling to give us any additional information which would allow us to identify the 'rogue' users on our network, or to identify the 'unwanted' traffic so that we could either block it, or use it to identify the rogue users ourselves.

Has anyone else come up against the problem, and/or have any suggestions on how best to resolve it?

Finally the fourth ape!
He is the sum of the first three: He sees nobody, hears nobody and speaks to nobody.

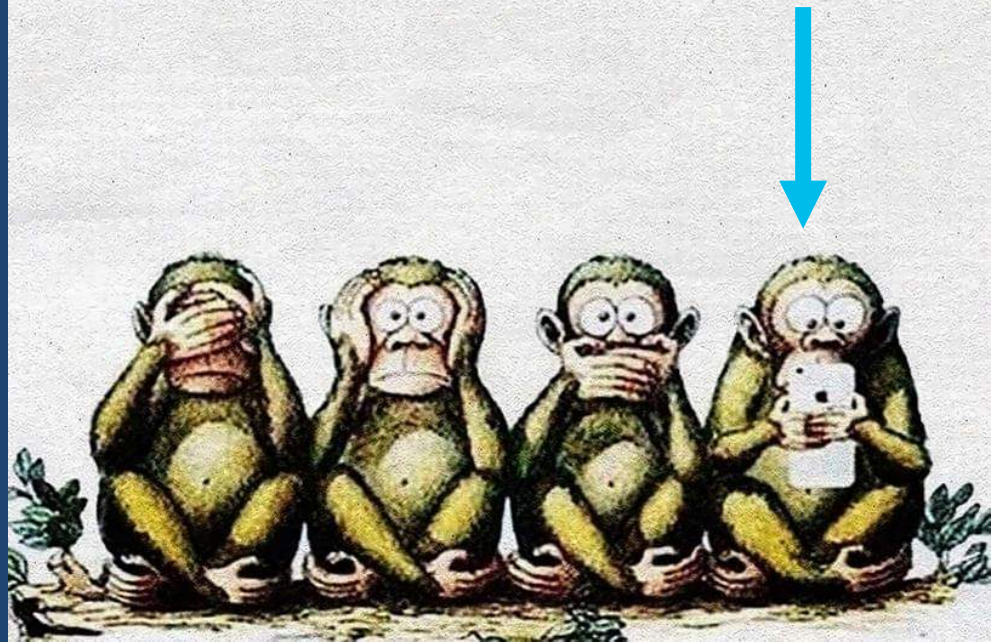

IPv4 – Classic
But spare parts have run out

IPv6 – Next Gen
Getting to full parity and end-end use takes time

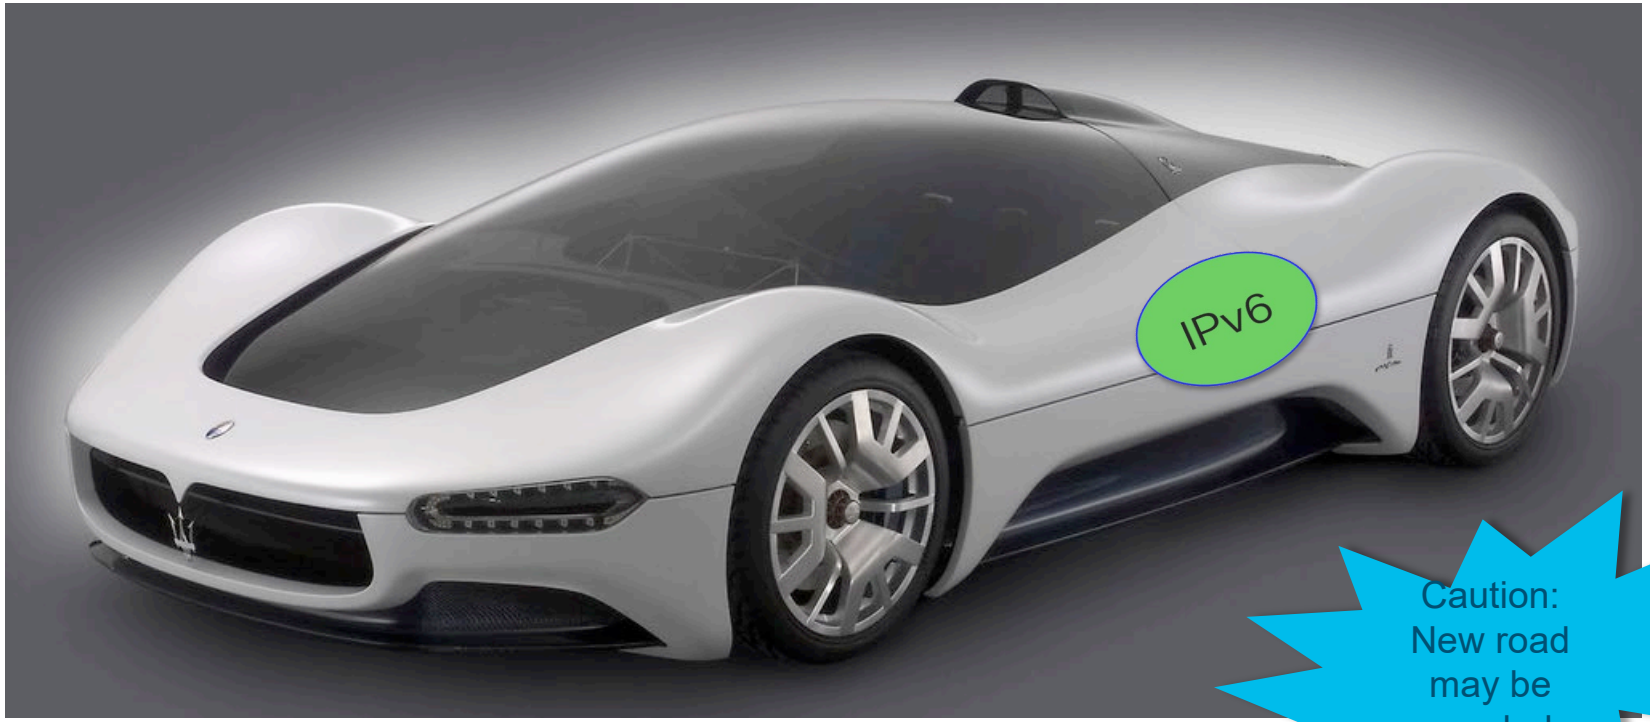

Transition Technologies help to continue Driving classic IPv4 around

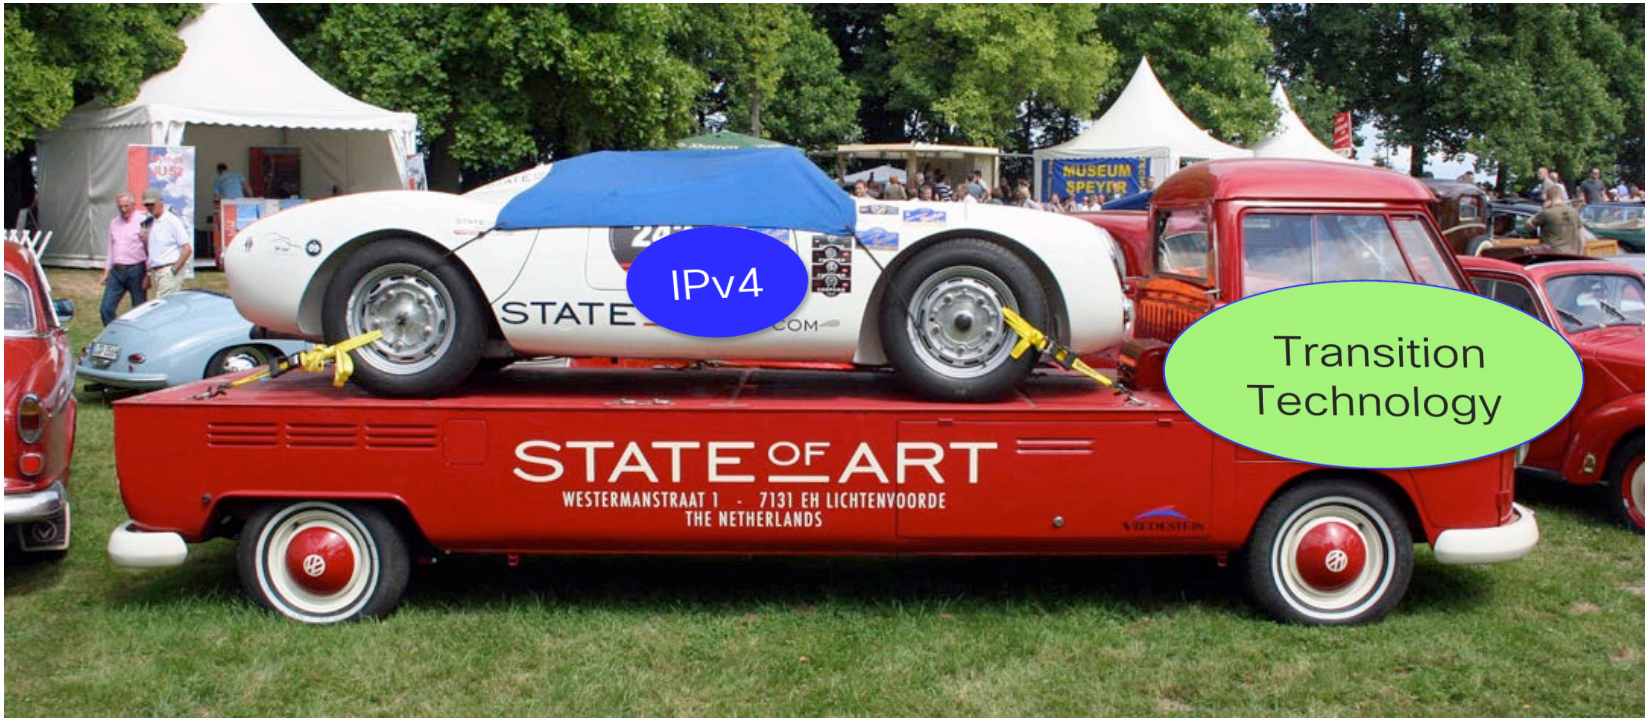

Abstract

- Any **Service Provider** that has exhausted its IPv4 address pool, should not only have to deploy/offer IPv6, but also employ IPv4 sharing.
  - This is because some content may be reachable only via IPv4 internet, even though majority is available via IPv6 internet.

- This session discusses few technologies such as MAP-T/E, 464XLAT, DS-Lite and CGN 64/44 etc. that facilitate IPv4 sharing with and without IPv6.
  - It contrasts stateful and stateless translation techniques as well.
  - 6rd is included as a reference as well.

- This session is intended for Service Providers.

Agenda

- Introduction
- Overview of Transition Technologies
  - Single-Stack IPv4 – Obtain more IPv4
  - Single-Stack IPv4 – CGN 44, 6rd
  - Dual Stack – Impact (& Happy Eyeballs)
  - Single-Stack IPv6 – DS-Lite, MAP-T/E
  - Single-Stack IPv6 – CGN 64
- IPv4 Address Sharing - Impact
- Conclusion

Recommended Approach (2005 – 2016)

RFC 4213: Deploy Dual-Stack and then move to IPv6-only

- **Dual-Stack** at the Clients
  - Windows, OSX, iOS, Android, Linux etc.
- **Dual-Stack** at the DC/Servers
  - Windows, Linux etc.
- **Dual-Stack** at the Network
  - Routers: IOS, XR, NXOS etc.
  - Switches: NXOS, CatOS, IOS etc.

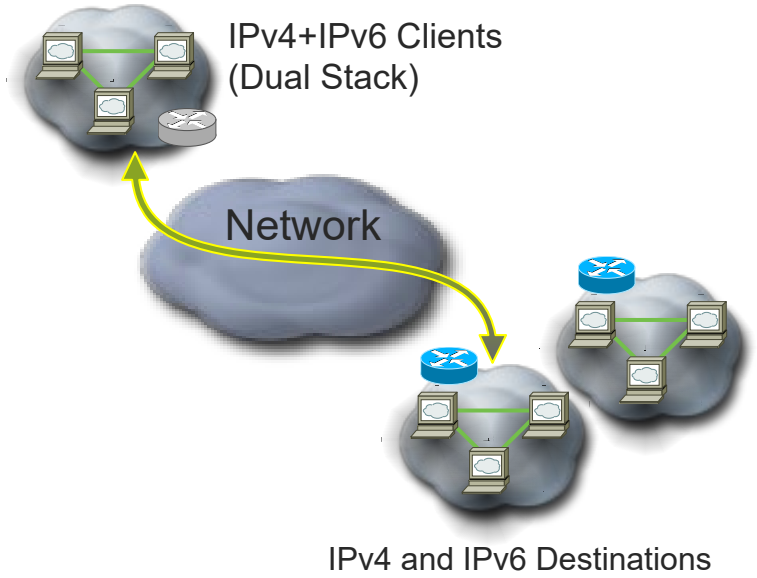

Recommended Approach (2005 – 2016)

RFC 4213: Deploy Dual-Stack and then move to IPv6-only

- Dual-Stack at the Clients
  - Wind
- Dual-S
  - Wind
- Dual-Stack at the Network
  - Route XR, NXOS
  - Switc

But IPv4 exhaustion means **every client** can NOT be assigned a public IPv4 address

But IPv4 exhaustion means **every network device** may NOT be assigned a public IPv4 address

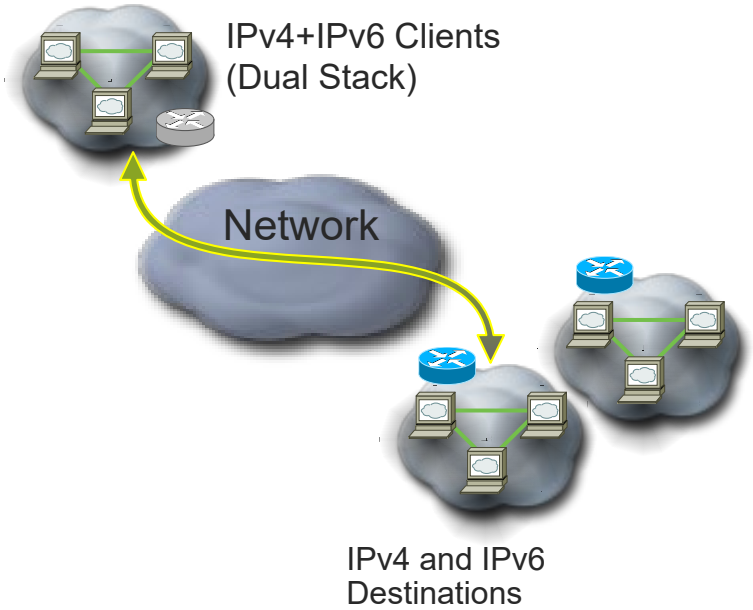

Recommended Approach (2005 – 2016)

IPv4 Address Exhaustion : Different Impact

Impact on ISPs

- Lack of IPv4 addresses for users
- Harder to grow the business
- IPv4 address sharing requires NAT
- Stateful NAT requires NAT logging

Impact on Users

- IPv4 Address sharing results in shared reputation (more on this later)
  - Breaks applications
  - Complicates operating servers
  - Limits UDP/TCP ports per user

Recommended Approach (2005 – 2016)

RFC 4213: Deploy Dual-Stack and then move to IPv6-only

~90% of Desktop hosts and ~99% of Mobile hosts support IPv6

- **Dual-Stack** at the Clients – IPv6 support
  - Desktop (Windows, OSX, Linux, Chrome OS etc.) & Mobile (iOS, Android…)

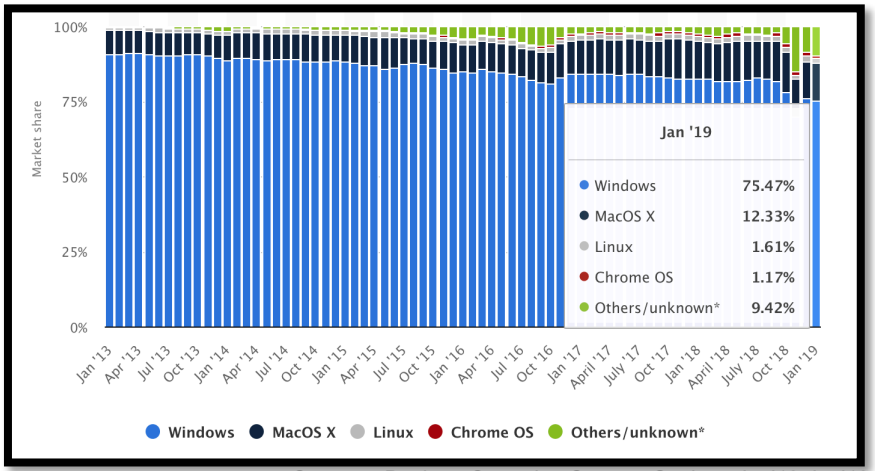

Source – Desktop Operating System, Statista, Jan'13-Jan'19

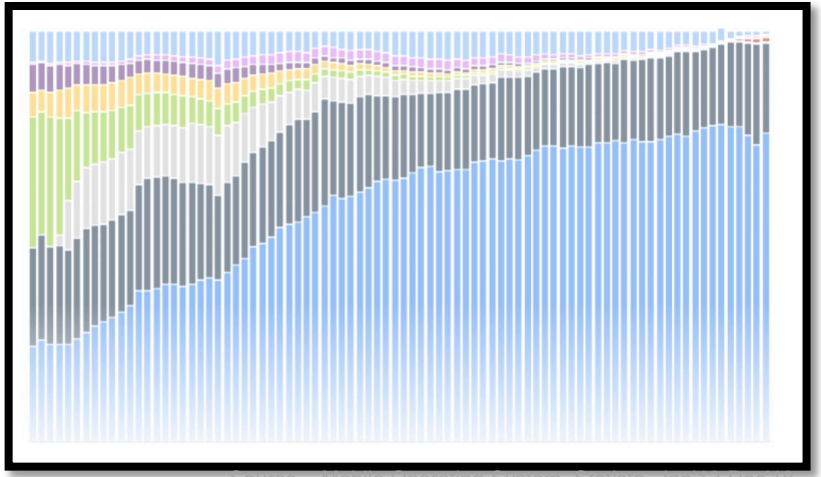

Source – Mobile Operating System, Statista, Jan'12-Dec'18

Recommended Approach (2017 onwards)

RFC 4213: Deploy Dual-Stack and then move to IPv6-only

- Single-Stack at the Clients
  - Windows, OSX, iOS, Android, Linux etc.
- Single-Stack at the DC/Servers*
  - Windows, Linux etc.
- Single-Stack at the Network
  - Routers: IOS, XR, NXOS etc.
  - Switches: NXOS, CatOS, IOS etc.
- Assume Dual-stack Destinations

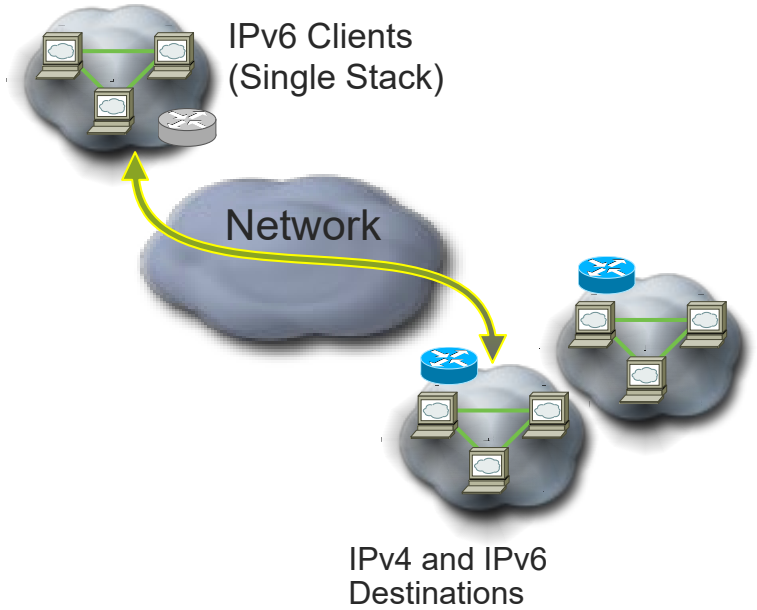

* RFC7755 prescribes Single-stack IPv6 for DC

Recommended Approach (2017 onwards) Which path suits you?

- Path towards IPv6 for Networks and CPEs –
  - ~~can't enable IPv6 in Network or CPE~~
  - enable IPv6 (dual-stack) in Networks, but not on CPEs
  - enable IPv6 (dual-stack) in Networks and on CPEs
  - Remove IPv4 or build IPv6-only Networks, CPEs stay on Dual-stack
  - Remove IPv4 or build IPv6-only Networks, and/or IPv6-only CPEs
- Your path may mean –
  - IPv6 co-existing with IPv4 !
  - IPv6 interoperating with IPv4 !
  - IPv4 address sharing by CPEs!

IPv6 Adoption Continues to increase...30% globally

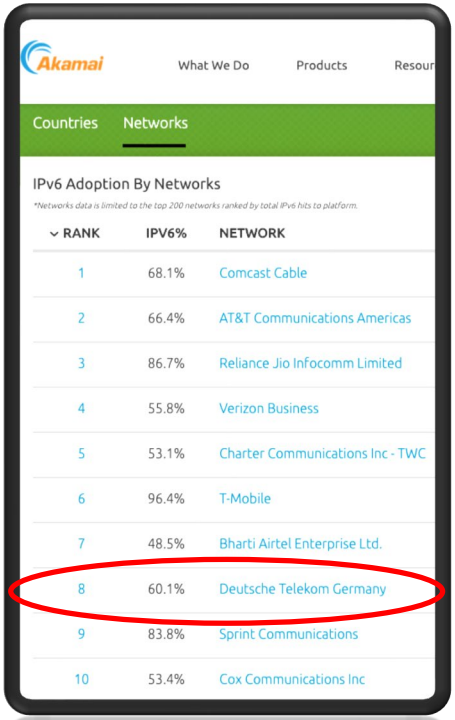

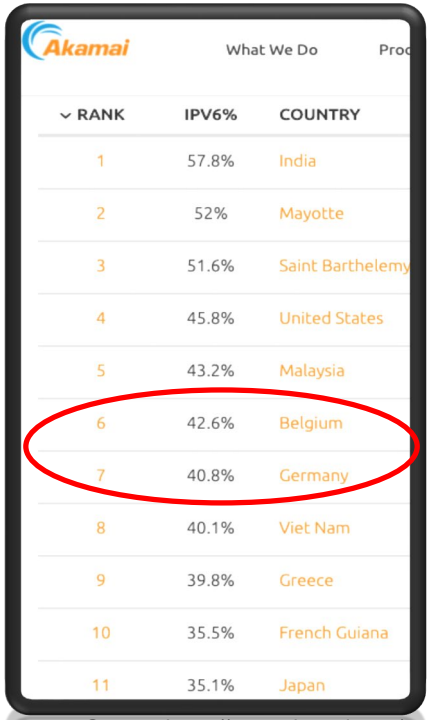

Source: https://www.akamai.com/

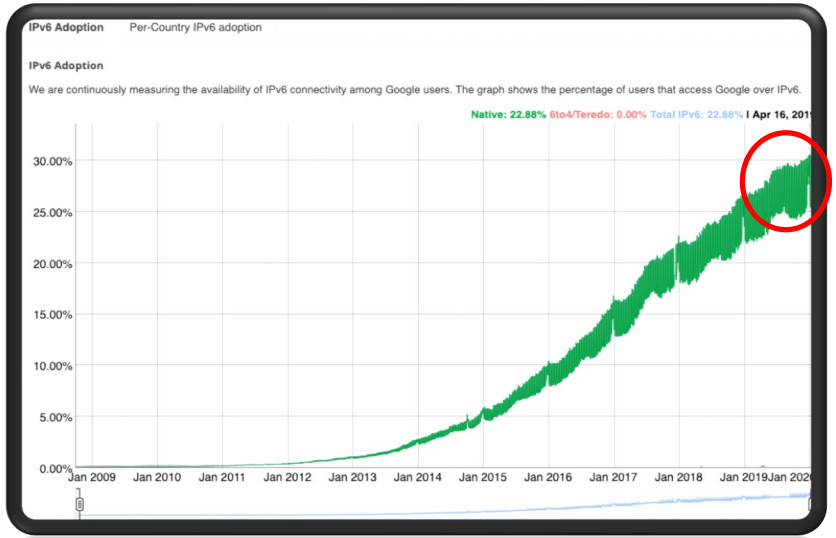

Source: https://www.google.com/

Towards IPv6 ...with or without IPv4

Transition Technologies in One Slide

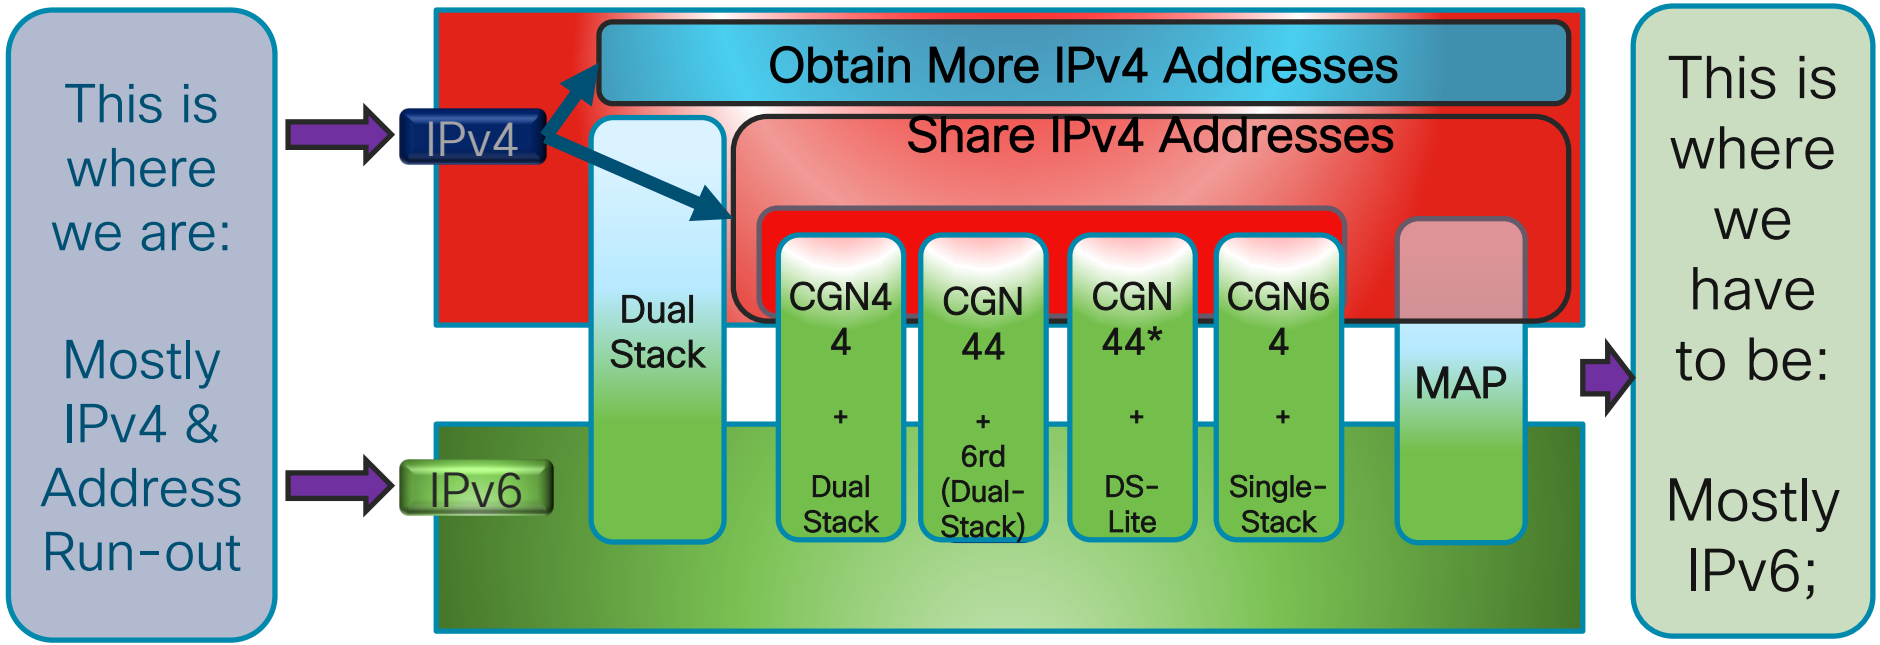

1. CGN = Carrier Grade NAT - Stateful
2. Modified to support DS-Lite

Agenda

- Introduction
- Transition Technologies - Summary
- Transition Technologies - Overview
  - Single-Stack IPv4 – Obtain more IPv4
  - Single-Stack IPv4 – CGN 44, 6rd
  - Dual Stack – Impact (& Happy Eyeballs)
  - Single-Stack IPv6 – DS-Lite, MAP-T/E
  - Single-Stack IPv6 – CGN 64
- IPv4 Address Sharing - Impact
- Conclusion

Towards IPv6 ...with or without IPv4

Transition Technologies in One Slide

* Allows both arbitrary and algorithmic mapping
** Changes needed if IPv6 is not supported by existing CPE

	Options	CPE LAN IPv4 or IPv6	CPE WAN IPv4 or IPv6	Tunnel or Translate?	In-network "State"?	Arbitrary IP addressing of CPE?	Extra CPE features?
0	Single-Stack	IPv4	IPv4	-NA-	-NA-	Yes	No
1	Single-Stack	IPv4	IPv4	Translate	Yes (CGN44)	Yes	No
2	Dual-Stack	IPv4 + IPv6	IPv4+IPv6	-NA-	-NA-	Yes	No**
3	Dual-Stack	IPv4 + IPv6	IPv4+IPv6	Translate	Yes (CGN44)	Yes	No**
4	DS-Lite	IPv4 + IPv6	IPv6	Both	Yes (CGN44)	Yes	Yes
5	6rd	IPv4 + IPv6	IPv4	Tunnel	No	No	Yes
6	6rd + CGN	IPv4 + IPv6	IPv4	Both	Yes (CGN44)	No	Yes
7	MAP	IPv4 + IPv6	IPv6	Either	No	Yes*	Yes
8	Single-Stack	IPv6	IPv6	Translate	Yes (CGN64)	Yes	Yes|No

Agenda

- Introduction
- Transition Technologies - Summary
- Transition Technologies - Overview
  - Single-Stack IPv4 – Obtain more IPv4
  - Single-Stack IPv4 – CGN 44, 6rd
  - Dual Stack – Impact (& Happy Eyeballs)
  - Single-Stack IPv6 – DS-Lite, MAP-T/E
  - Single-Stack IPv6 – CGN 64
- IPv4 Address Sharing - Impact
- Conclusion

Single-Stack IPv4: Obtain IPv4 Addresses

Host/CPE gets Public IPv4 prefix(es)

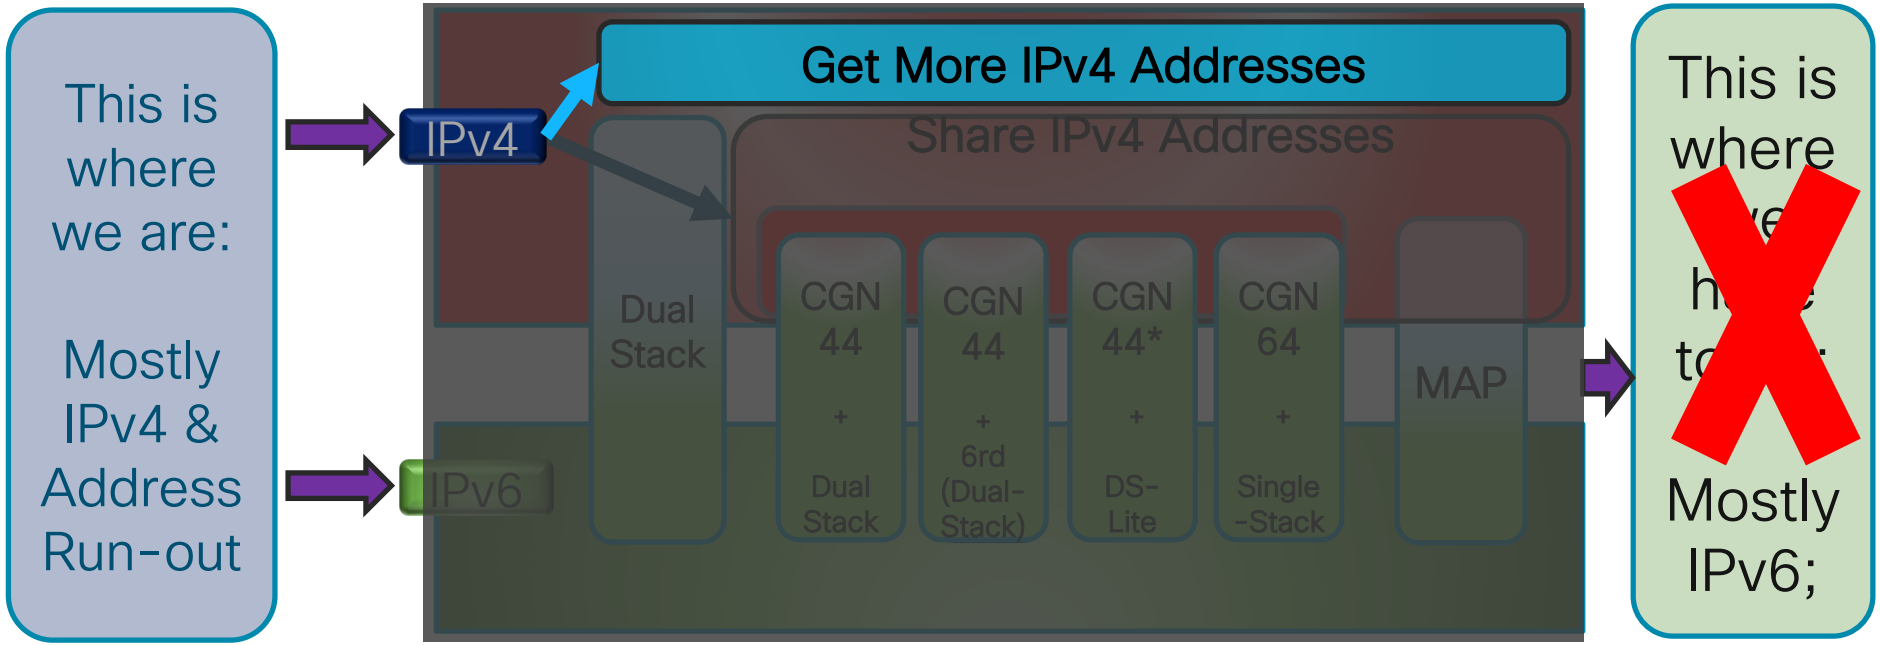

Single-Stack IPv4: Obtain IPv4 Addresses

- Obtain IPv4 addresses
  - RIR: May Not have any left. ☹
  - Open market: USD ~~$10-$15~~ $20-$25 per IPv4 address

Source – https://auctions.ipv4.global

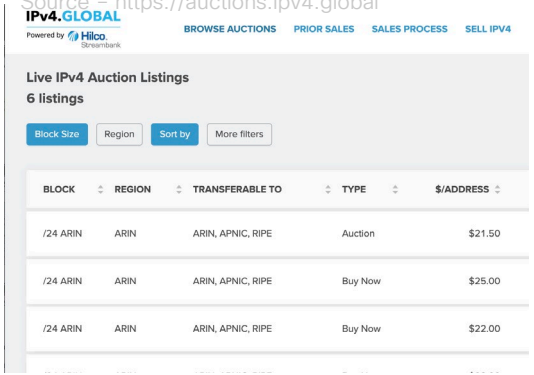

ADVANTAGES:

- No CGN, no address sharing, no operational changes
- No need to press for IPv6 deployment

DISADVANTAGES :

- If business growing, delaying the inevitable
- Geo-location needs to be updated (mileage varies)
- No IPv6 deployed
- Reputation might be bad

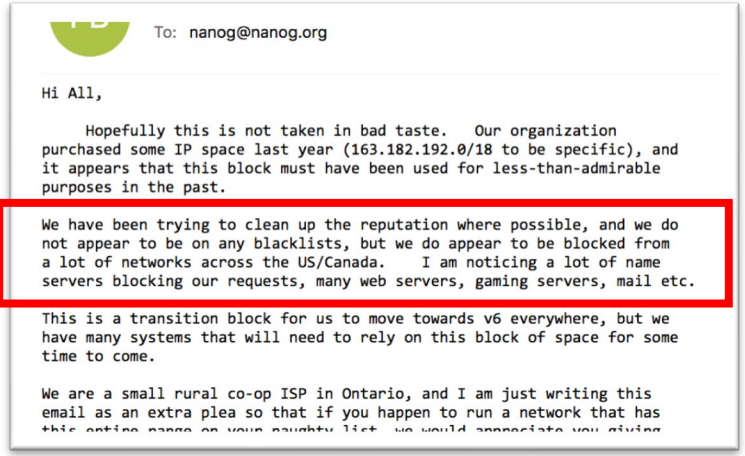
To: nanog@nanog.org

Hi All,

Hopefully this is not taken in bad taste. Our organization purchased some IP space last year (163.182.192.0/18 to be specific), and it appears that this block must have been used for less-than-admirable purposes in the past.

We have been trying to clean up the reputation where possible, and we do not appear to be on any blacklists, but we do appear to be blocked from a lot of networks across the US/Canada. I am noticing a lot of name servers blocking our requests, many web servers, gaming servers, mail etc.

This is a transition block for us to move towards v6 everywhere, but we have many systems that will need to rely on this block of space for some time to come.

We are a small rural co-op ISP in Ontario, and I am just writing this email as an extra plea so that if you happen to run a network that has

Single-Stack IPv4: Obtain IPv4 Addresses

- Check the following before purchasing:

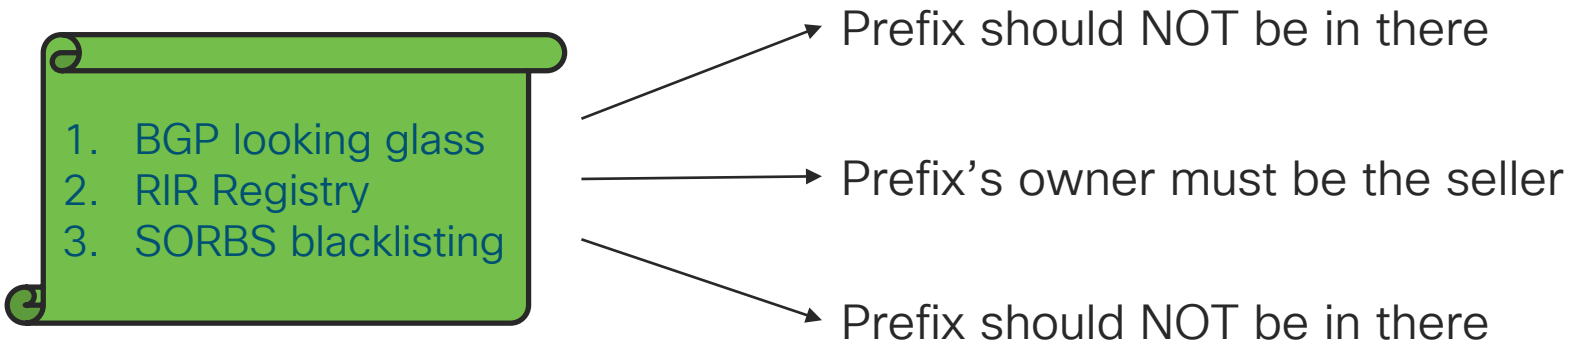

Agenda

- Introduction
- Transition Technologies - Summary
- Transition Technologies - Overview
  - Single-Stack IPv4 – Obtain more IPv4
  - Single-Stack IPv4 – CGN 44, 6rd
  - Dual Stack – Impact (& Happy Eyeballs)
  - Single-Stack IPv6 – DS-Lite, MAP-T/E
  - Single-Stack IPv6 – CGN 64
- IPv4 Address Sharing - Impact
- Conclusion

Carrier Grade NAT 44 (CGN44)

Transition Technologies in One Slide

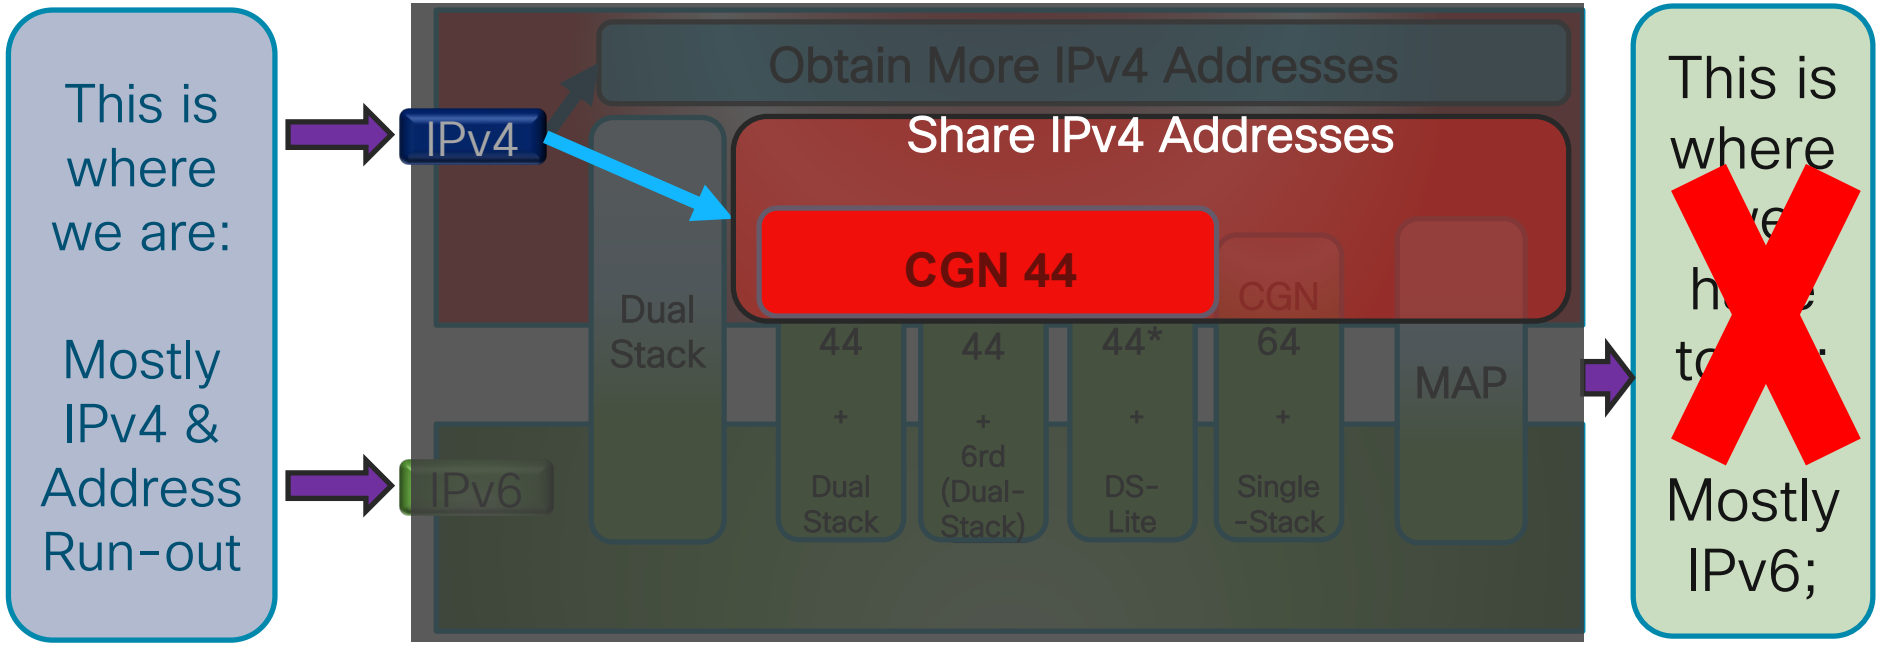

1. CGN = Carrier Grade NAT - Stateful
2. Modified to support DS-Lite

CGN 44

- Single-Stack IPv4 requires CGN 44, if IPv4 address sharing
  - Single-Stack IPv6 would require CGN 64, covered later
- Carrier Grade Network Address Translation
  - Address and Port Translator (NAPT), really
  - RFC5389 : Endpoint independent Mapping/Filtering (EIM and EIF)
  - Similar to residential NAT (Linksys, etc.), but large scale
  - Port Logging (e.g. syslog, netflow v9)
  - Per-user port limit
- In case of IPv4-only Clients with CGN44: Using 100.64.0.0/10 instead of private IPv4 space is an option
  - In case of IPv6-only Clients with CGN64: Using GUA should be the only option, covered later

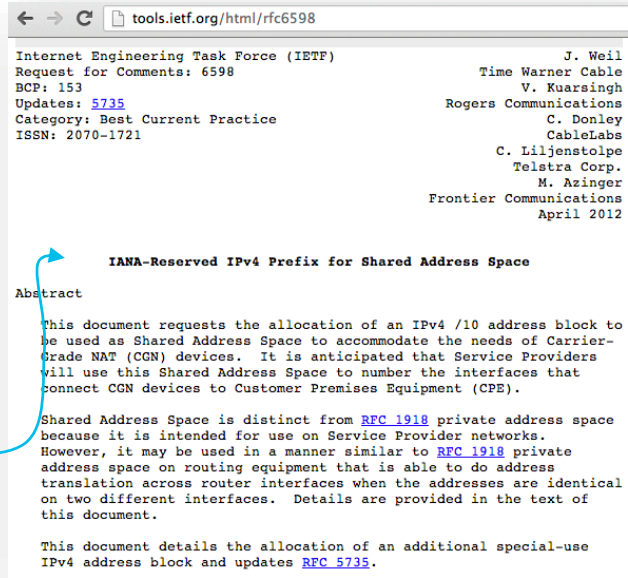

CGN 44

Supported on ASR9K, ASR1K, FirePower, CRS

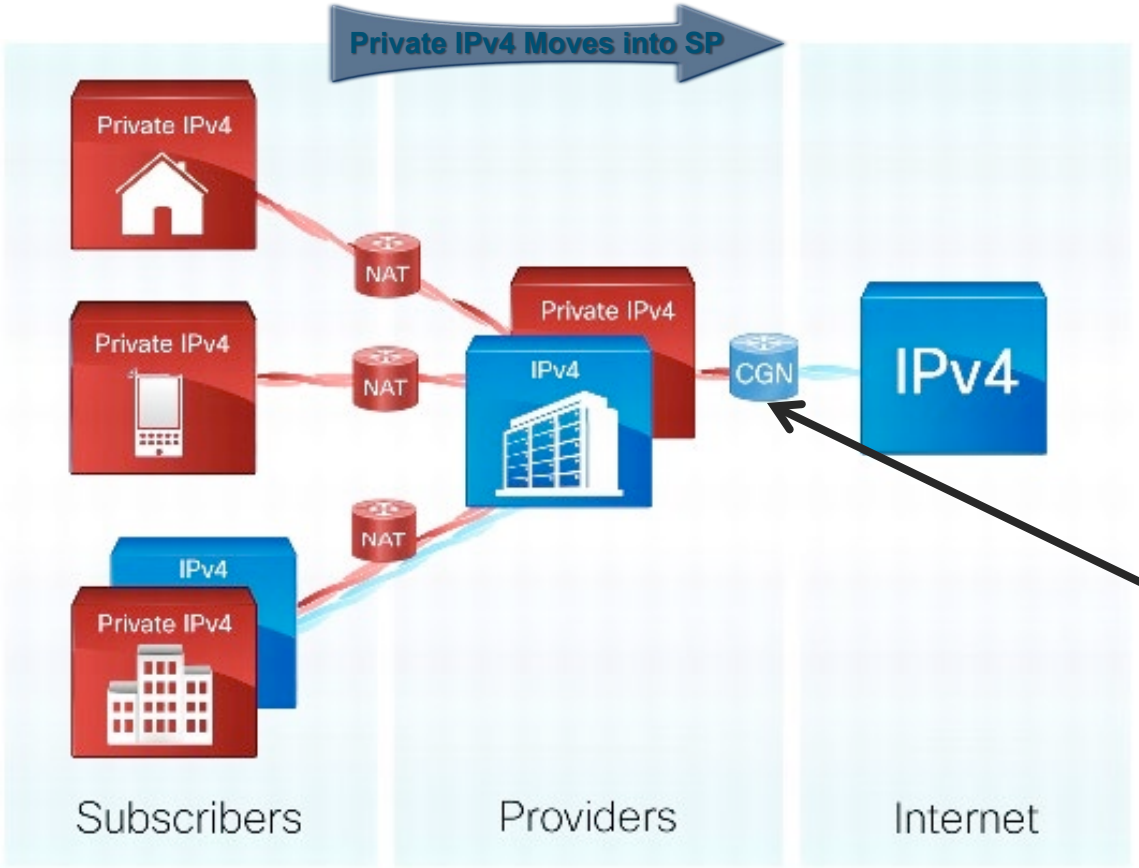

Stateful NAT function inside SP network

CGN 44

Supported on ASR9K, ASR1K, FirePower,CRS

- Nicknamed NAT444 = NAT44 in home, NAT44 in ISP
- Advantages:
  1. Very well known technology
  2. No dependency on CPE router
- Disadvantages:
  1. Logging = huge storage
  2. Port Forwarding breaks
  3. Certain Applications may NOT sufficiently work
  4. Network/Routing Design Headache
  5. IPv4 address sharing efficiency
  6. DoS possibility, if unwanted incoming traffic is dropped in slow path, or is fragmented
  7. Any application hardcoding a specific port# may not work without UPnPv2+PCP

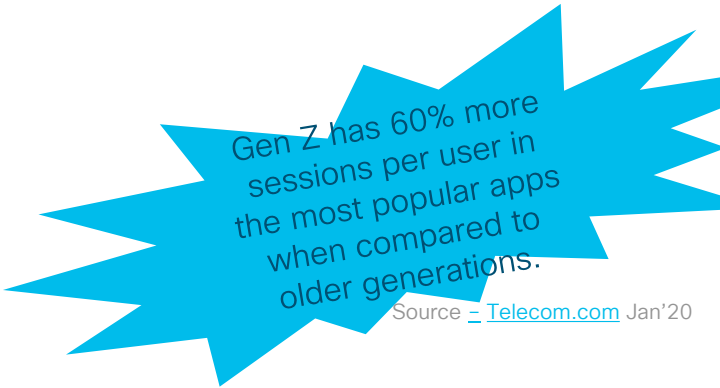

Source – Telecom.com Jan'20

CGN
ALG, Logging

ALG, NAT Logging etc. issues applicable to all these solutions relying on stateful NAT

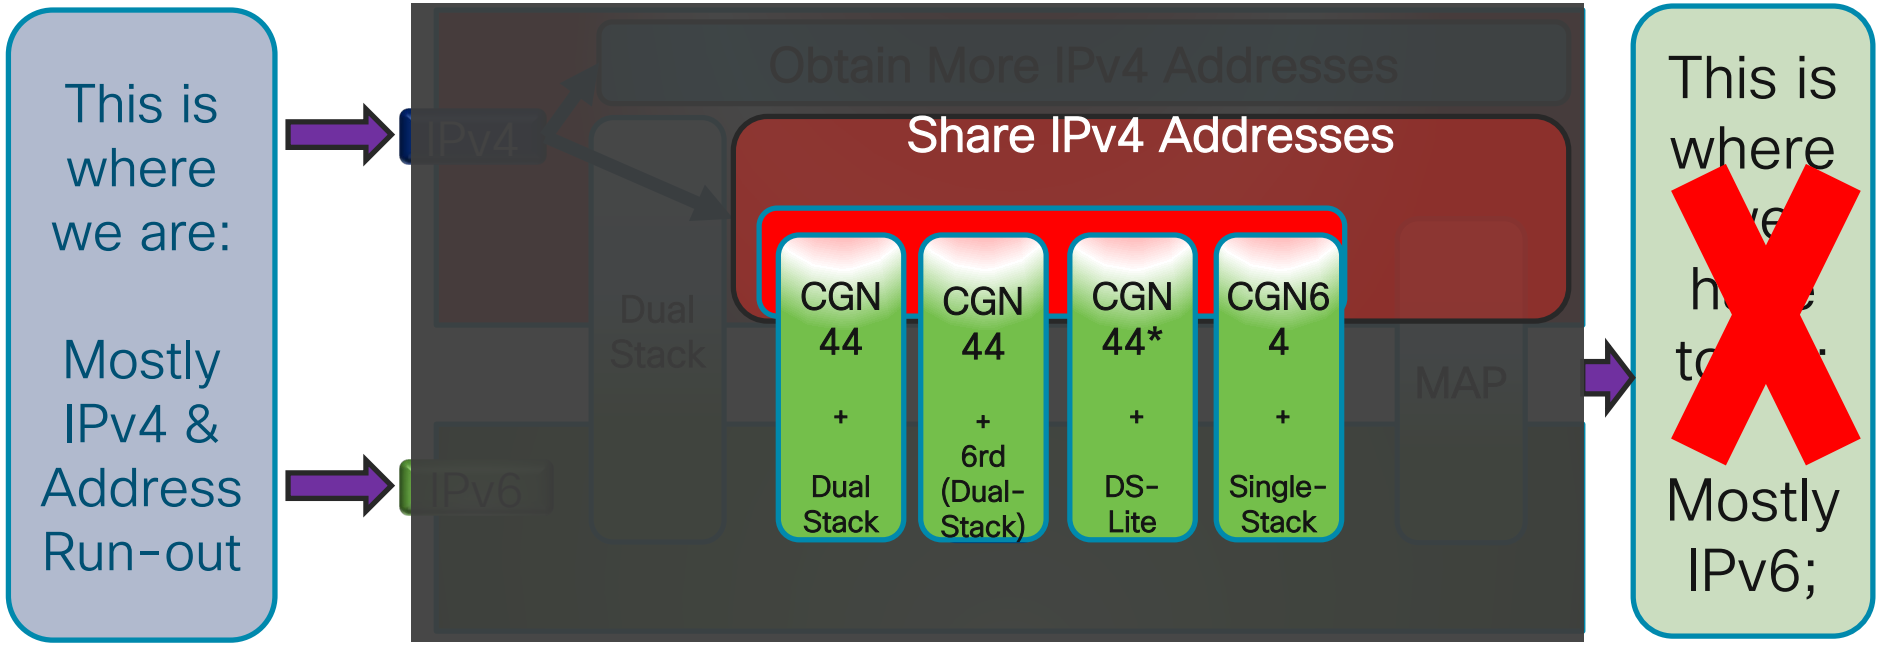

CGN
Logging Source Port Ranges

Supported on ASR9K, ASR1K, CRS

Reference

- Stateful NAT requires logging (NAT44, NAT64, DS-Lite…)
  - NAT mappings are temporary (similar to DHCP addresses)
- Logging each NAT mapping creates large logs!
- Bulk port allocation (BPA) reduces logging, at the expense of reduced efficiency of IPv4 address sharing ☺ ☹
  - Bulk size of **N** ports, logs reduced by 1/**N**
  - Acceptable compromise !!!
- Recommended

42.5TB over 60 days for 200K subscribers, 72K flows/second

(each syslog comprised private source IP:port, public source IP:port, protocol, and timestamp, resulting in ~100B in ASCII). See note below.

See BRKSPG-3334 from CiscoLive2014 for more details

CGN
Logging Destination

Supported on ASR9K, ASR1K, CRS

- Server Log combined with CGN log identifies subscribers
  - **Timestamp (new)**
  - Source IP address, **source port (new)**, destination IP address, destination port
  - RFC6302
- Some servers don't enable source port logging, or don't have good timestamp
  - Note that majority support logging source port, but don't do so by default, see RFC7768 and draft-daveor-cgn-logging
- Tempting to log destination IP (and port) at CGN
  - Consider privacy and legal issues
  - Incompatible with bulk port allocation, increases logging costs
- Not recommended in general

See BRKSPG-3334 from CiscoLive2014 for more details

CGN – Common Sane Practices

- Use Bulk Port Allocation, if logging
- Limit number of users sharing an IPv4 address *
- Monitor KPIs with threshold
  - outbound SSH connections
  - Incoming fragmented traffic
  - Incoming dropped traffic
  - …
- Test, Test, Test as many apps as possible

See RFC6888 for CGN requirements

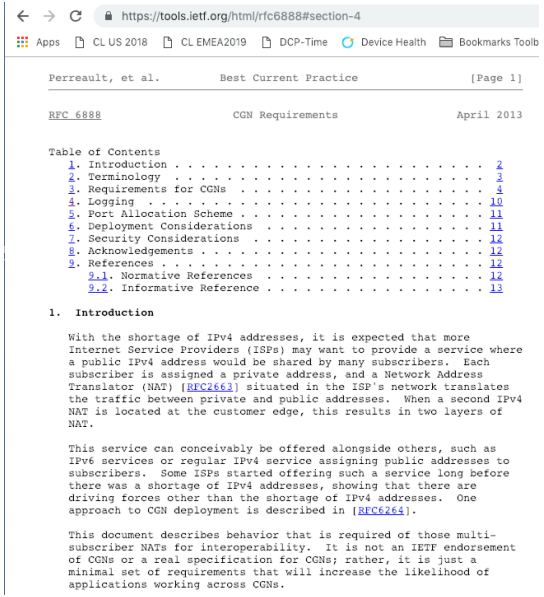
https://tools.ietf.org/html/rfc6888#section-4

Apps | CL US 2018 | CL EMEA2019 | DCP-Time | Device Health | Bookmarks Toolb

Perreault, et al. Best Current Practice [Page 1]

RFC 6888 CGN Requirements April 2013

Table of Contents

1. Introduction 2
2. Terminology 3
3. Requirements for CGNs 4
4. Logging 10
5. Port Allocation Scheme 11
6. Deployment Considerations 11
7. Security Considerations 12
8. Acknowledgements 12
9. References 12
   9.1. Normative References 12
   9.2. Informative Reference 13

1. Introduction

With the shortage of IPv4 addresses, it is expected that more Internet Service Providers (ISPs) may want to provide a service where a public IPv4 address would be shared by many subscribers. Each subscriber is assigned a private address, and a Network Address Translator (NAT) [RFC2663] situated in the ISP's network translates the traffic between private and public addresses. When a second IPv4 NAT is located at the customer edge, this results in two layers of NAT.

This service can conceivably be offered alongside others, such as IPv6 services or regular IPv4 service assigning public addresses to subscribers. Some ISPs started offering such a service long before there was a shortage of IPv4 addresses, showing that there are driving forces other than the shortage of IPv4 addresses. One approach to CGN deployment is described in [RFC6264].

This document describes behavior that is required of those multi-subscriber NATs for interoperability. It is not an IETF endorsement of CGNs or a real specification for CGNs; rather, it is just a minimal set of requirements that will increase the likelihood of applications working across CGNs.

* Tricky because you would want higher sharing ratio, given IPv4 shortage

6rd and 6rd with CGN 44

Transition Technologies in One Slide

Reference

Reference

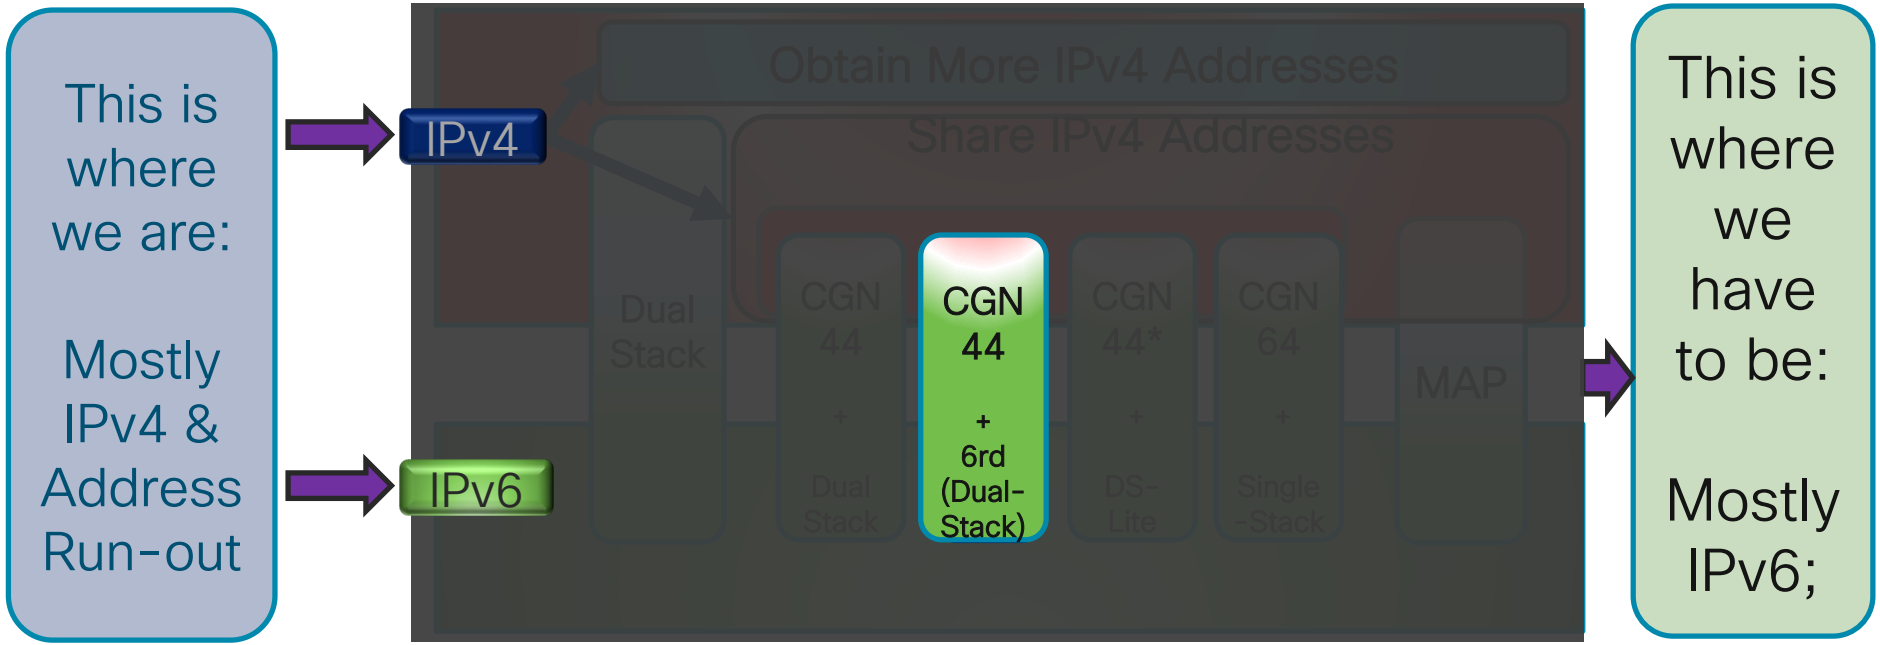

6rd – IPv6 over (Public) IPv4

Supported on ASR9K, ASR1K, CRS

Reference

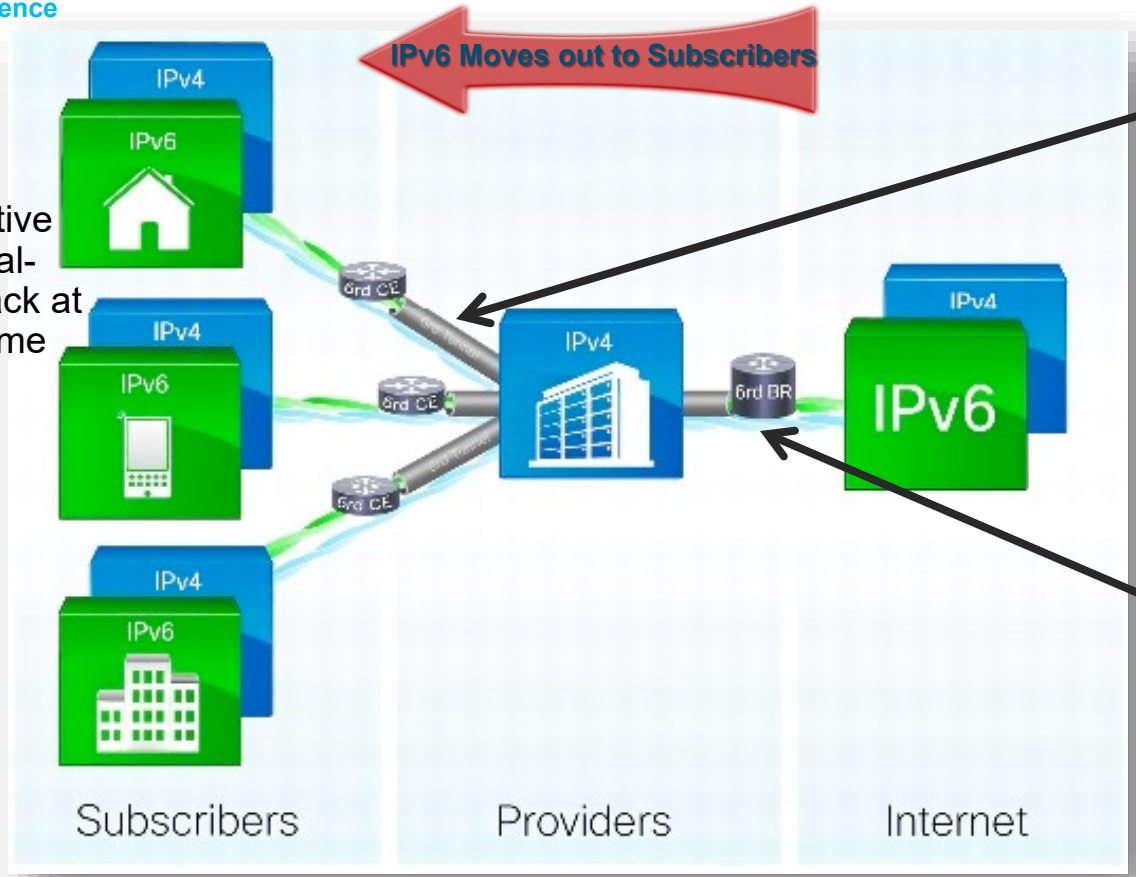

IPv6-over-IPv4 tunnels

Stateless Tunneling function (on routers) inside SP network

6rd + CGN = IPv6 over (Private) IPv4

Reference

Supported on ASR9K, ASR1K, CRS

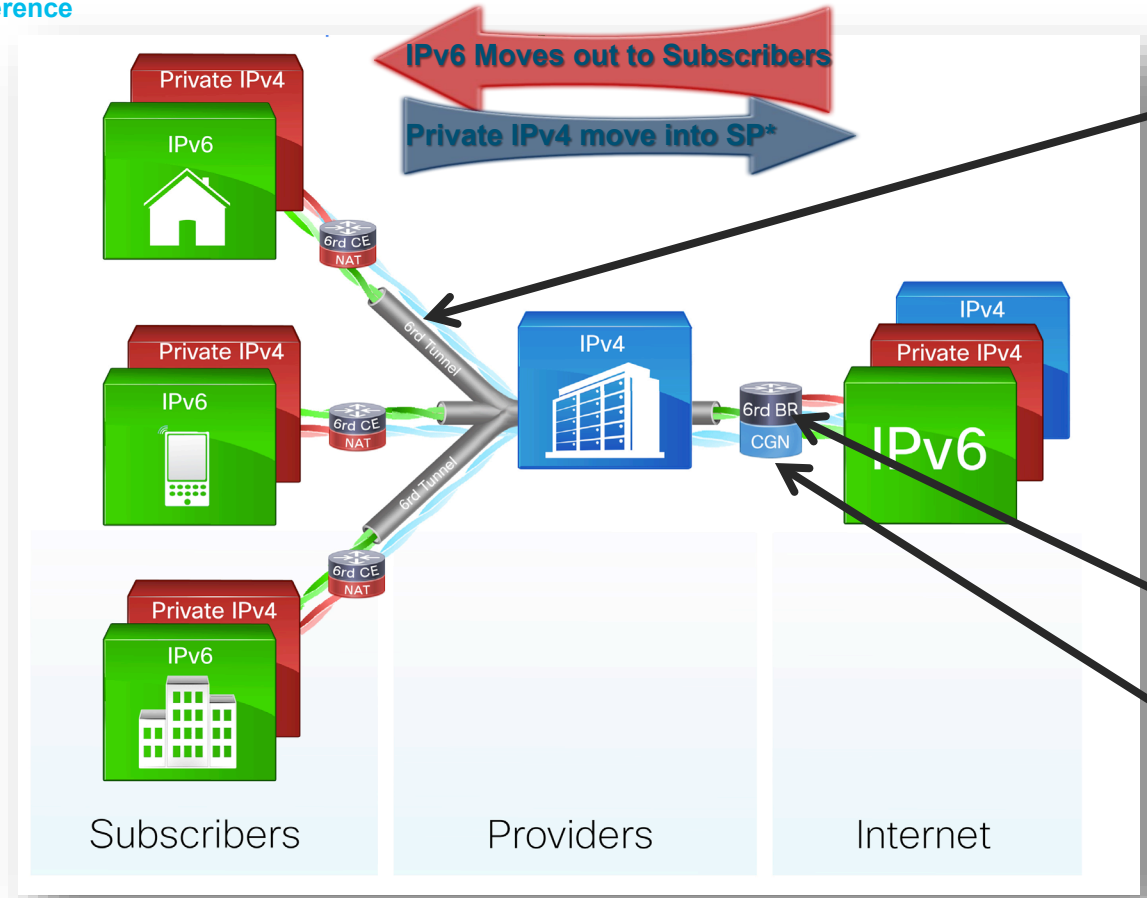

IPv6-over-IPv4 tunnels

Stateless Tunneling function (on routers)

Stateful NAT function (on routers) inside SP network*

* Assuming RFC1918 usage

Agenda

- Introduction
- Transition Technologies - Summary
- Transition Technologies - Overview
  - Single-Stack IPv4 – Obtain more IPv4
  - Single-Stack IPv4 – CGN 44, 6rd
  - Dual Stack – Impact (& Happy Eyeballs)
  - Single-Stack IPv6 – DS-Lite, MAP-T/E
  - Single-Stack IPv6 – CGN 64
- IPv4 Address Sharing - Impact
- Conclusion

Dual-Stack

Clients/CPE gets both IPv4 and IPv6 prefixes

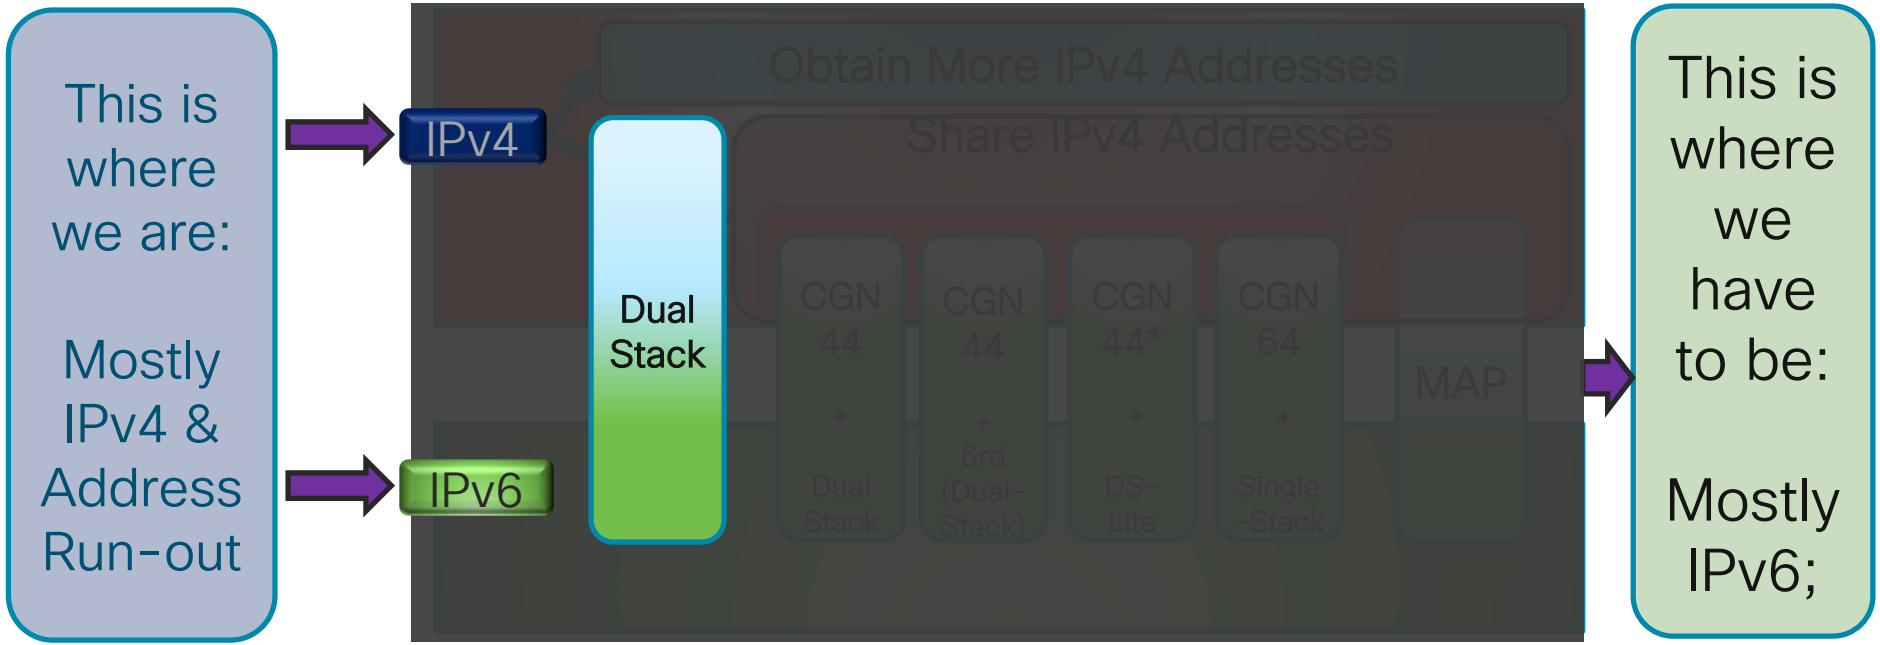

Dual-Stack
Do I use IPv6 or IPv4 ?

- Dual-stack client connecting to dual-stack server
- IPv6 is preferred by default (RFC6724)
- If IPv6 is slower, then users blame IPv6 and may **disable IPv6**! ☹
- **IPv6 better not be slower than IPv4**
  - Who can guarantee that ! ☹
- What if IPv6 is broken altogether?
- What if IPv6 is broken to few websites?

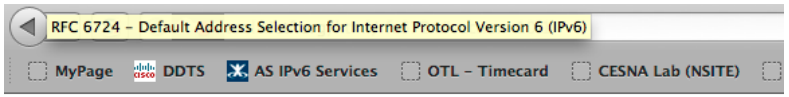

10.3.1. Handling Broken IPv6

One problem in practice that has been recently observed occurs when a host has IPv4 connectivity to the Internet but has "broken" IPv6 connectivity to the Internet in that it has a global IPv6 address but is disconnected from the IPv6 Internet. Since the default policy table prefers IPv6, this can result in unwanted timeouts.

This can be solved by configuring the table to prefer IPv4 as shown above. An implementation that has some means to detect that it is not connected to the IPv6 Internet MAY do this automatically. An implementation could instead treat it as part of its implementation of Rule 1 (avoid unusable destinations).

Dual-Stack : What if IPv6 is Broken or slower to a certain website ?

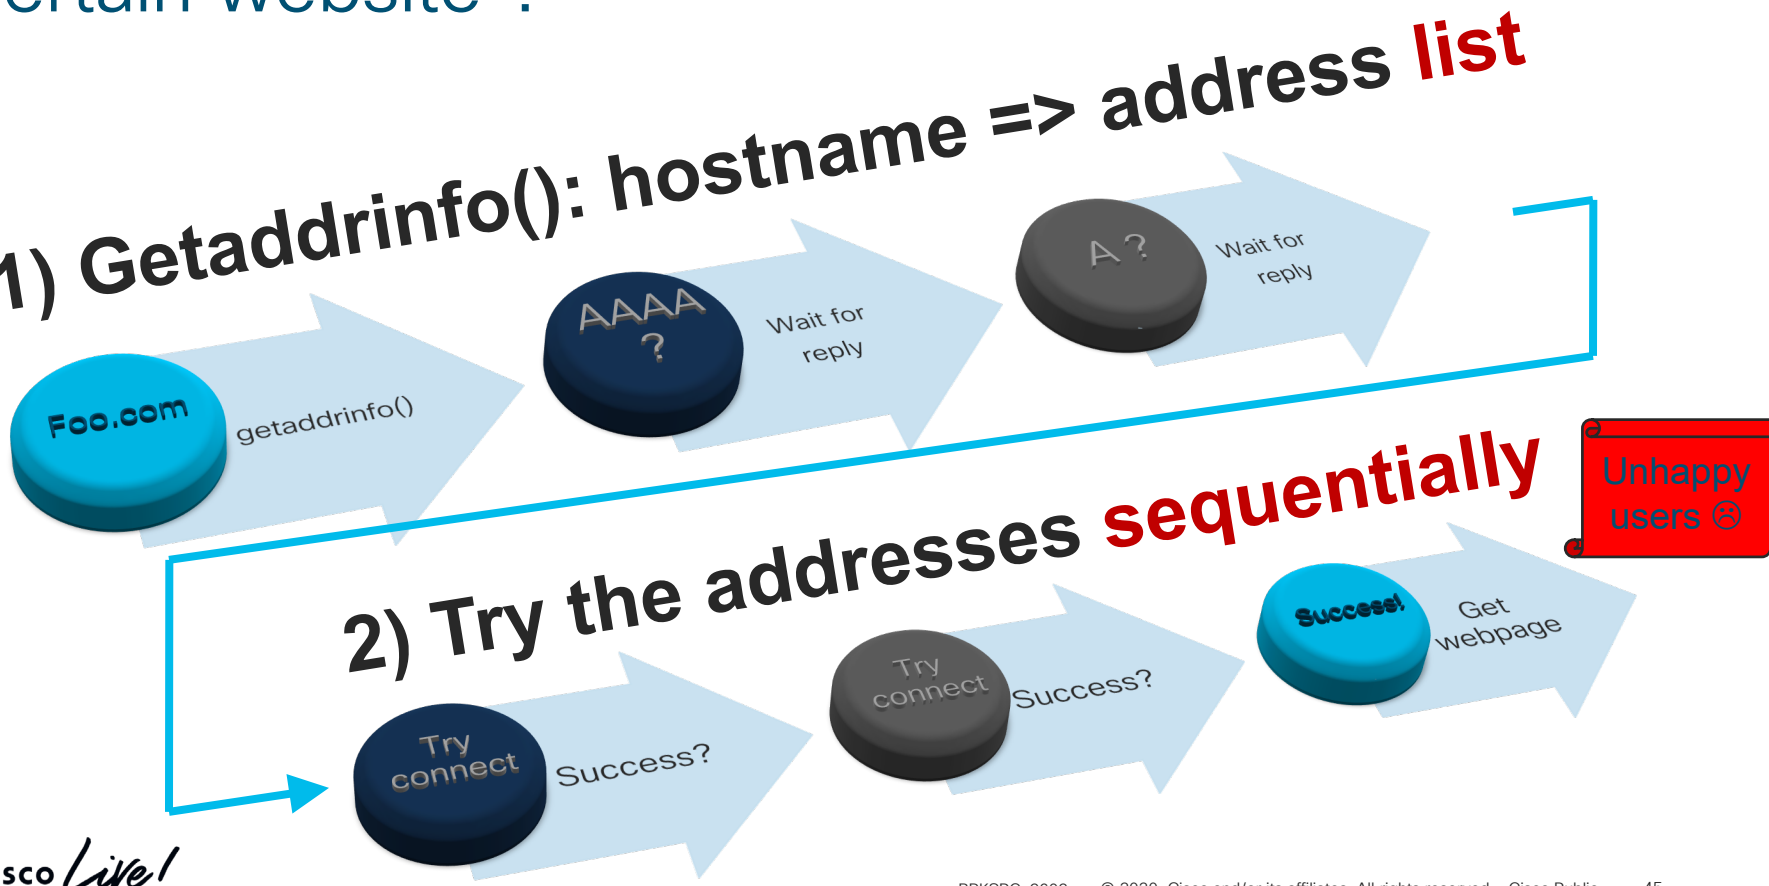

Dual-Stack Solution – Happy Eyeballs (RFC6555)

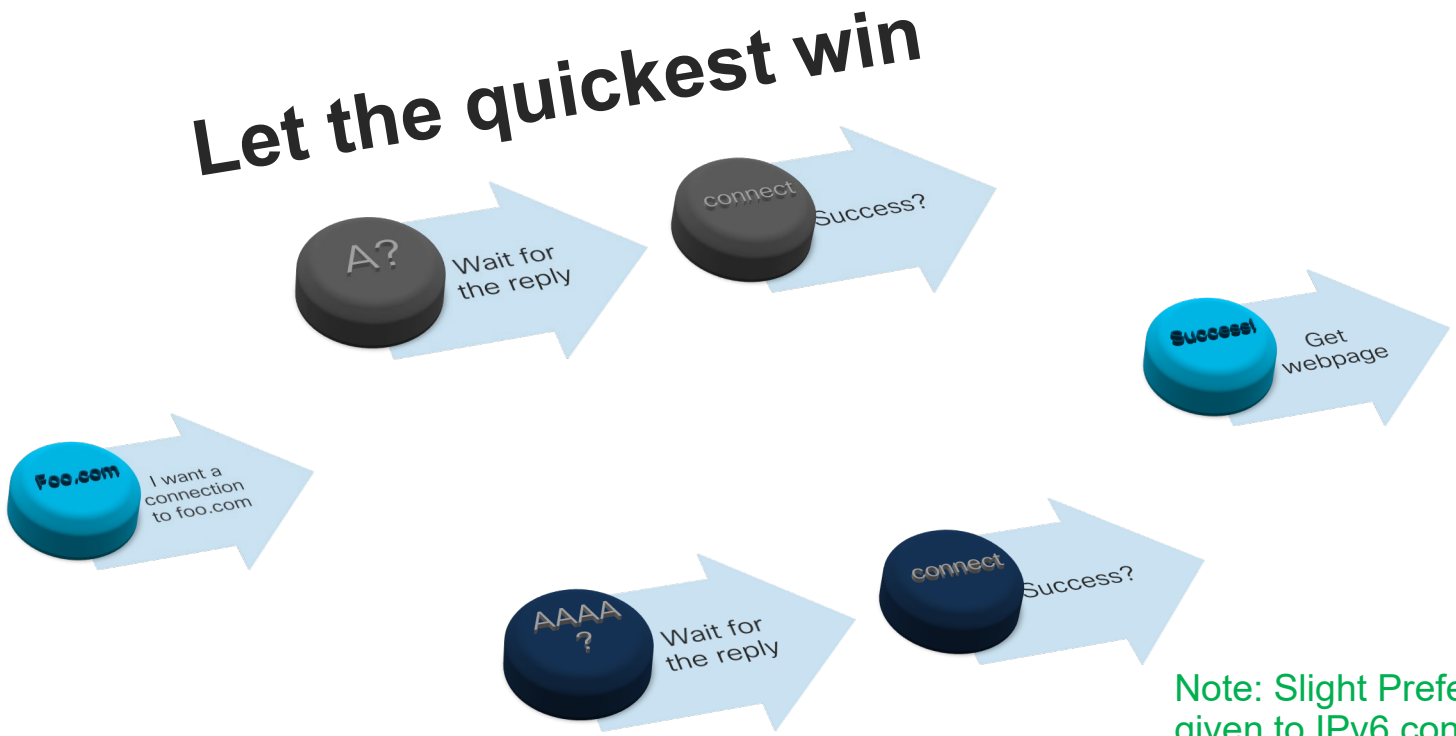

Note: Slight Preference is given to IPv6 connection

Dual-Stack
Happy Eyeballs Optimization (RFC8305)

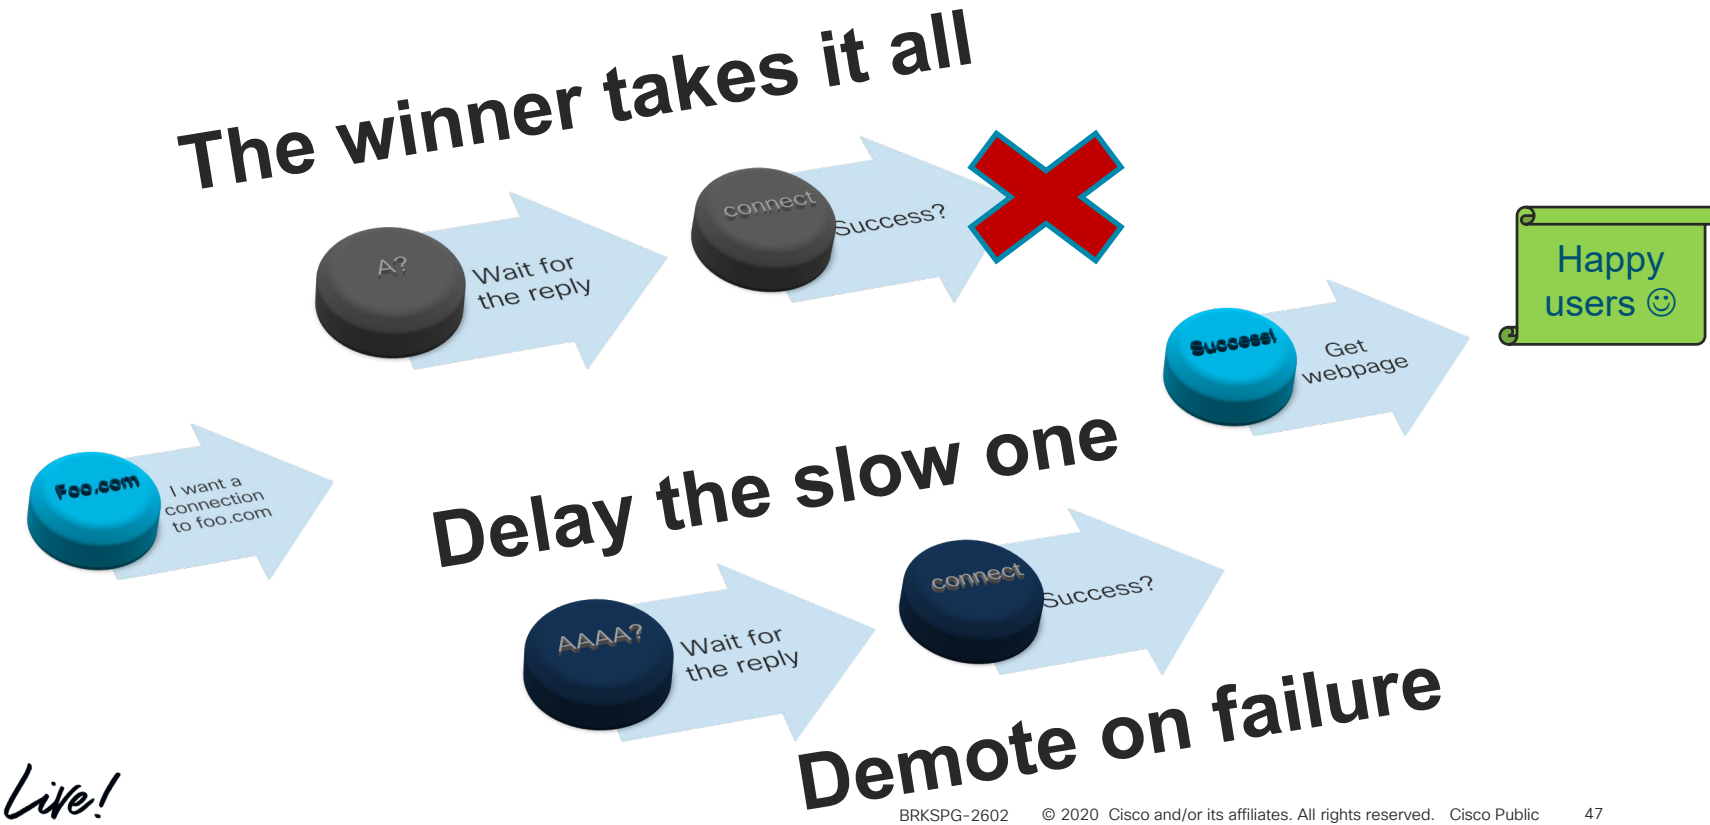

Dual-Stack
Happy Eyeballs (RFC6555 and RFC8305)

- Users are happy
  - Aimed initially at web browsing
  - Web browsing is *the* most common application
  - Fast response even if IPv6 (or IPv4) path is down
- Network administrators are happy
  - Users no longer trying to disable IPv6
  - Reduces IPv4 usage (reduces load on CGN)
- Content providers are happy
  - Better geolocation and DoS visibility with IPv6

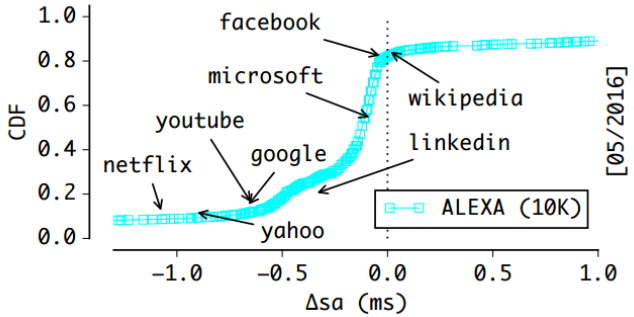

Figure 7: CDF of absolute difference of TCP connect times between IPv4 and IPv6 as of May 2016. 18% of the top 10K ALEXA websites are faster over IPv6 today, although 91% of the rest are at most 1 ms slower.

Source: http://seclists.org/nanog/2016/Jun/809

Dual-Stack
Happy Eyeballs Implementations

- Google Chrome and Mozilla Firefox: Yes ☺
  - Utilizes long-established 250-300ms 'backup' thread
  - Follows getaddrinfo() address preference
- Apple Safari, iOS*, OSX* : Yes ☺
  - DNS AAAA sent before A query on the wire
  - If AAAA reply comes first, then v6 SYN sent immediately
  - If A reply comes before 25ms of AAA reply, then v4 SYN sent
  - Else, Heuristics based Address selection algorithm is applied
- Microsoft Windows OS and Internet Explorer : NO ☹
  - Not even something like happy eyeballs
- Cisco WebEx : Yes ☺
- Cisco AnyConnect: No ☹

RFC6555 Compliant

* http://lists.apple.com/archives/Ipv6-dev/2011/Jul/msg00009.html

* https://www.ietf.org/mail-archive/web/v6ops/current/msg22455.html

Agenda

- Introduction
- Transition Technologies - Summary
- Transition Technologies - Overview
  - Single-Stack IPv4 – Obtain more IPv4
  - Single-Stack IPv4 – CGN 44, 6rd
  - Dual Stack – Impact (& Happy Eyeballs)
  - Single-Stack IPv6 – DS-Lite, MAP-T/E
  - Single-Stack IPv6 – CGN 64
- IPv4 Address Sharing - Impact
- Conclusion

Single-Stack IPv6 …DS-Lite

Transition Technologies in One Slide

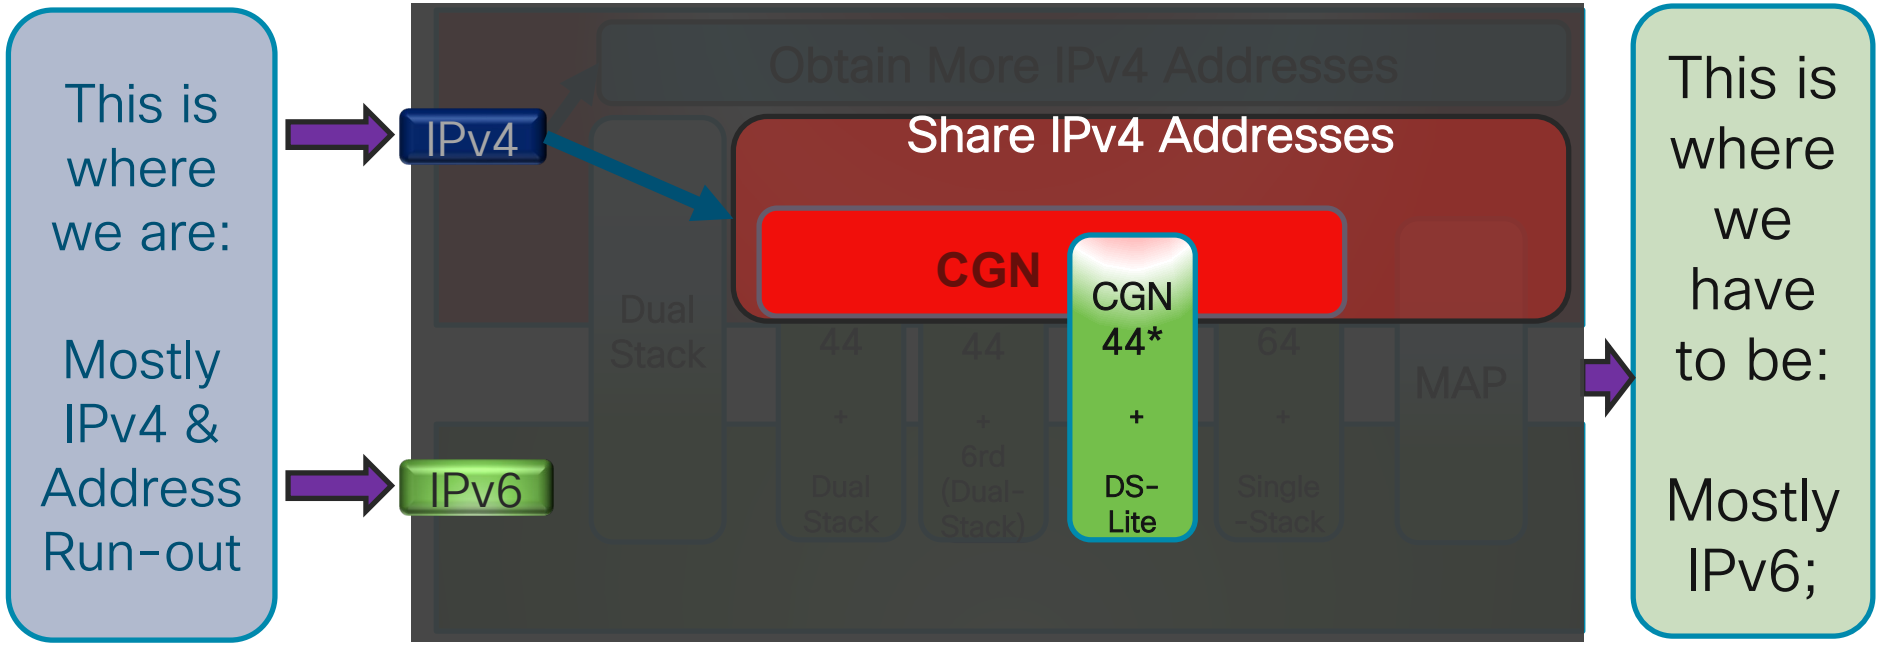

Note: DS-Lite requires CGN

Single-Stack IPv6: DS-Lite (RFC6333) IPv4 over IPv6 Access

Supported on ASR9K, ASR1K, CRS

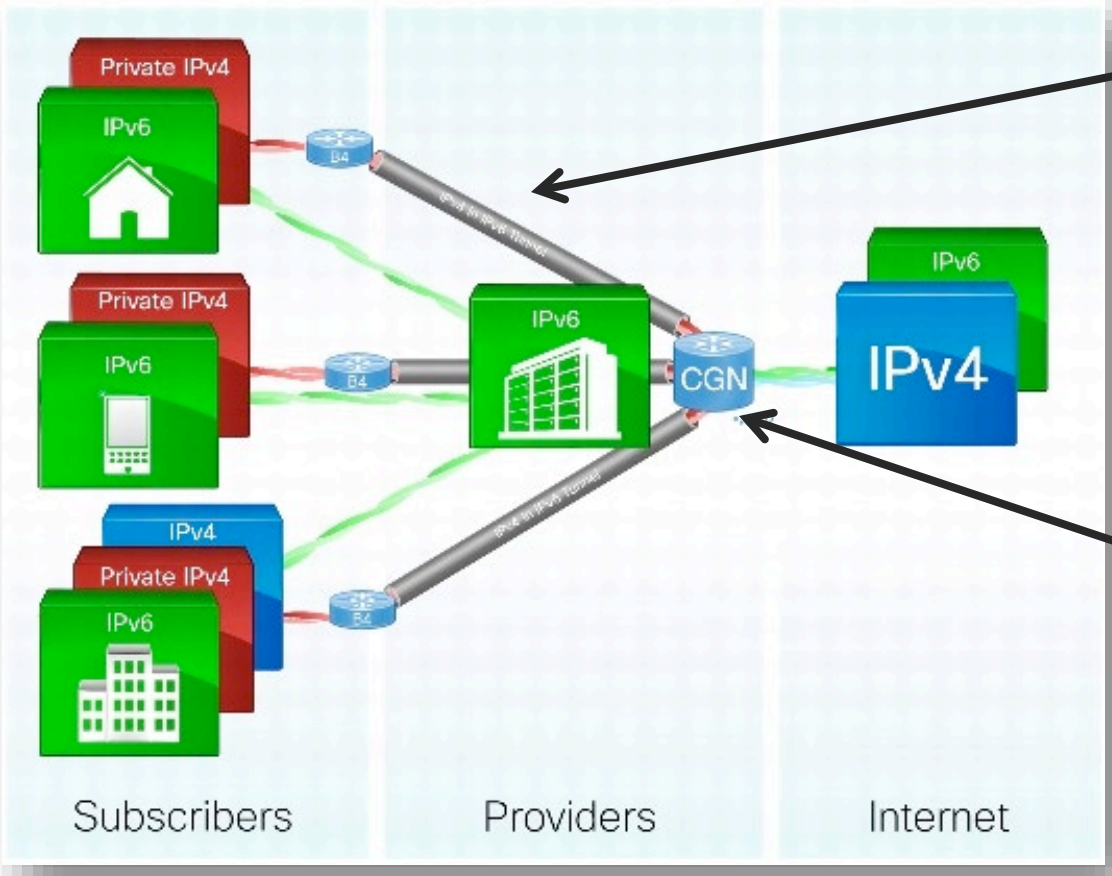

IPv4-over-IPv6 tunnels

Stateful NAT 44 function (on routers etc.) inside SP network

DS-Lite

- **Advantages:**
  - Leverages IPv6 in the network; Helps with IPv6-only Network

- **Disadvantages:**
  - Dependency on CPE router
  - NAT disabled on CPE router
  - Content Caching function may break
  - DPI function may break
  - QoS function may break
  - All disadvantages of stateful CG NAT also apply

Single-Stack IPv6: MAP

Mapping of Address and Port (RFC7599)

This is where we are:

Mostly IPv4 & Address Run-out

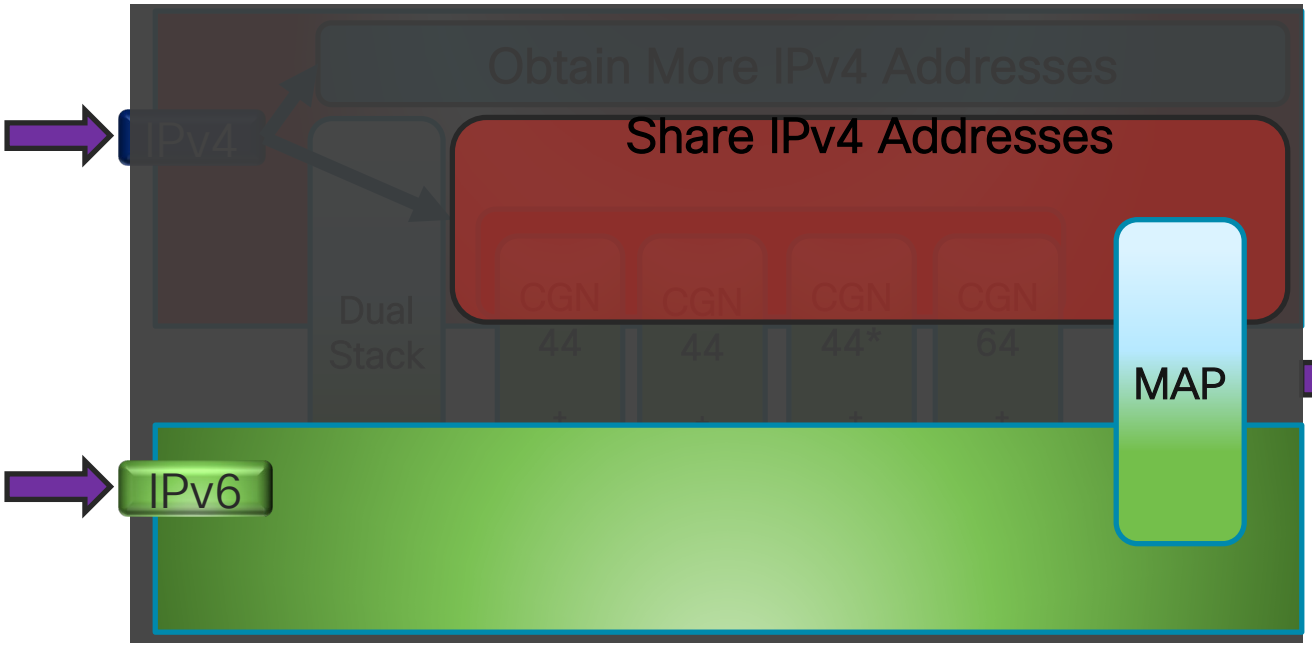

This is where we have to be:

Mostly IPv6;

MAP (Mapping of Address and Port)

BR function Supported on ASR9K, ASR1K

- MAP has CPE and Border Relay (BR) function;
- CPE learns of MAP info (e.g. DHCPv6 option (RFC7598))
  - Each CPE could formulate a shared IPv4 address with unique TCP/UDP port-range(s) via MAP info (=rules)
  - All or part of IPv4 address can be derived from the assigned IPv6 prefix (allows for IPv6 route summarization)
- Stateless BR in SP network
  - Can use anycast, can have asymmetric routing
  - Uses Algorithmic encoding to map IPv6 and IPv4 headers
  - No single point of failure, no need for high availability hardware

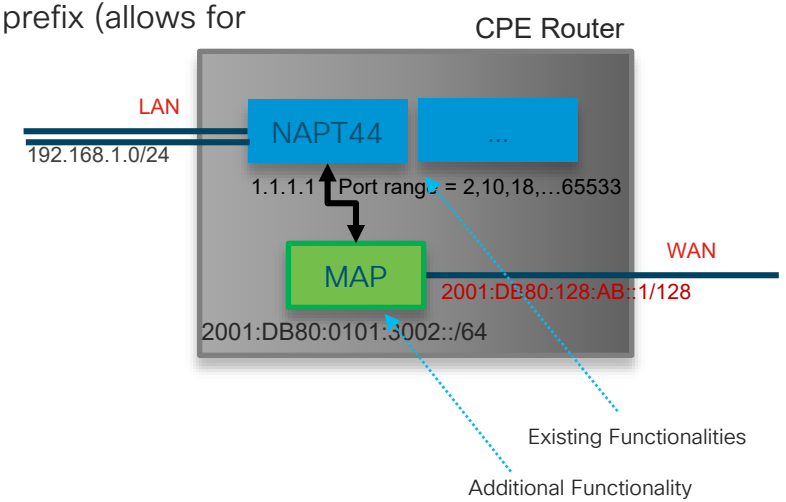

Supported on ASR9K,

MAP-E : Stateless 464 Encapsulation (RFC7597)

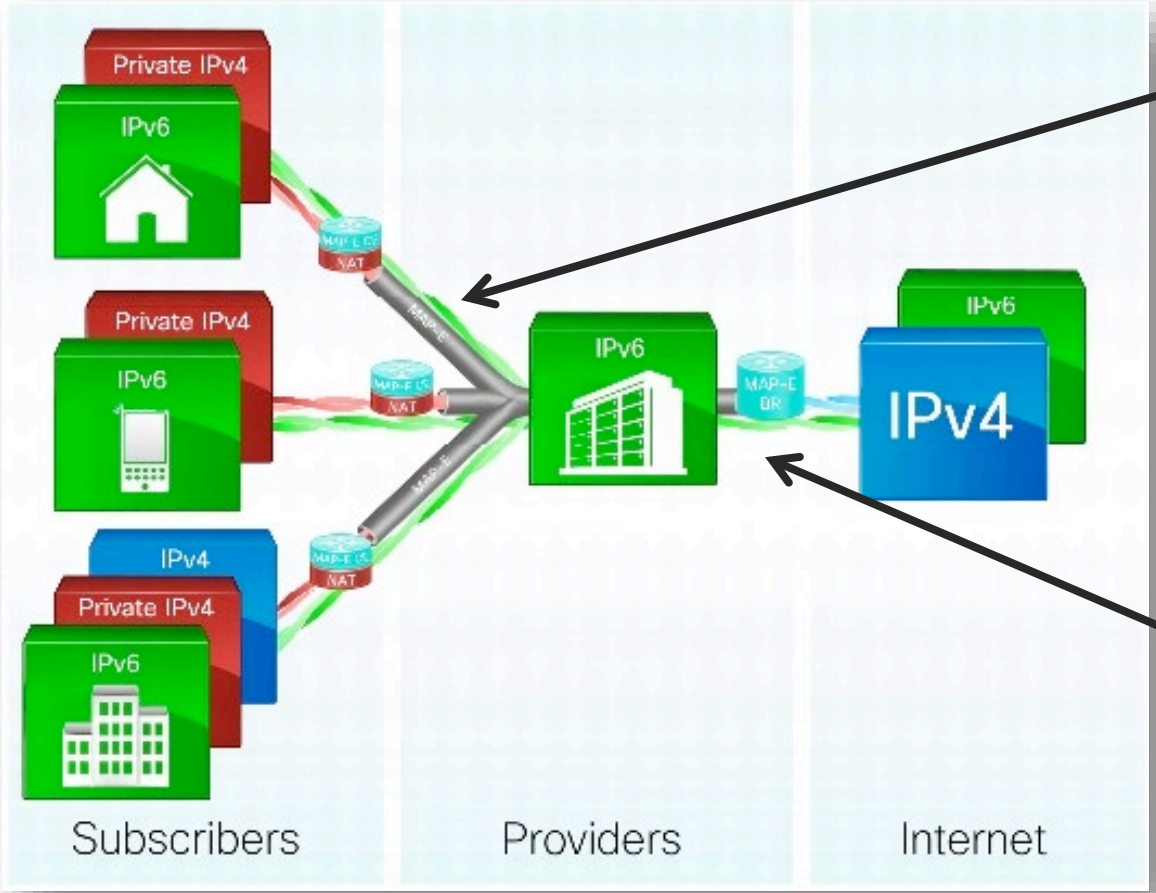

IPv4-over-IPv6

Stateless Tunneling function (on routers)

- No Stateful CGN-

MAP-T : Stateless 464 Translation (RFC7599)

Supported on ASR9K, ASR1K

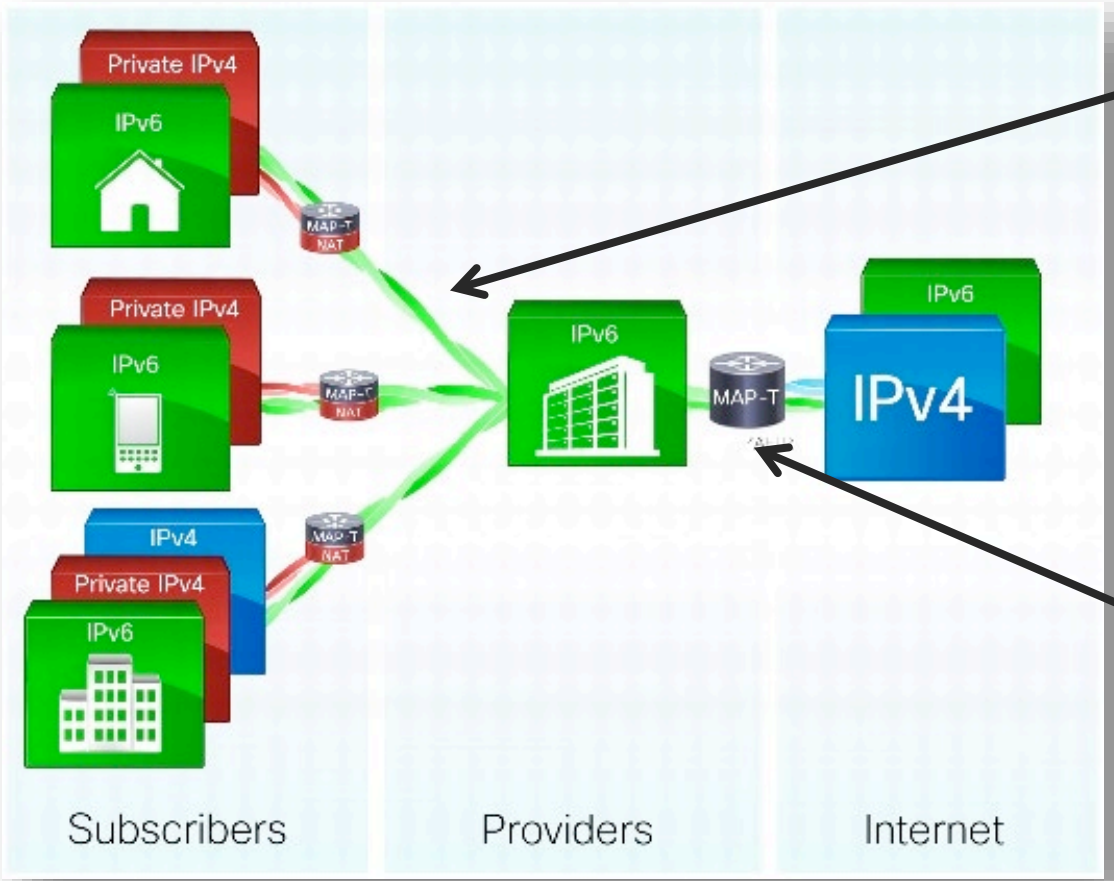

Native IPv6

Stateless 64 translation function (on routers)

- No Stateful CGN -

MAP

- **Advantages:**
  - Leverages IPv6 in the network
  - No CGN inside SP network
  - No need for NAT Logging (DHCP logging as usual)
  - No need for ALGs
  - No need for Stateful NAT64/DNS64

- **Disadvantages:**
  - Dependency on CPE router
  - Any application hardcoding any port# might not work without UPnPv2 support

MAP Design – Simplify Domain Addressing

http://map46.cisco.com/

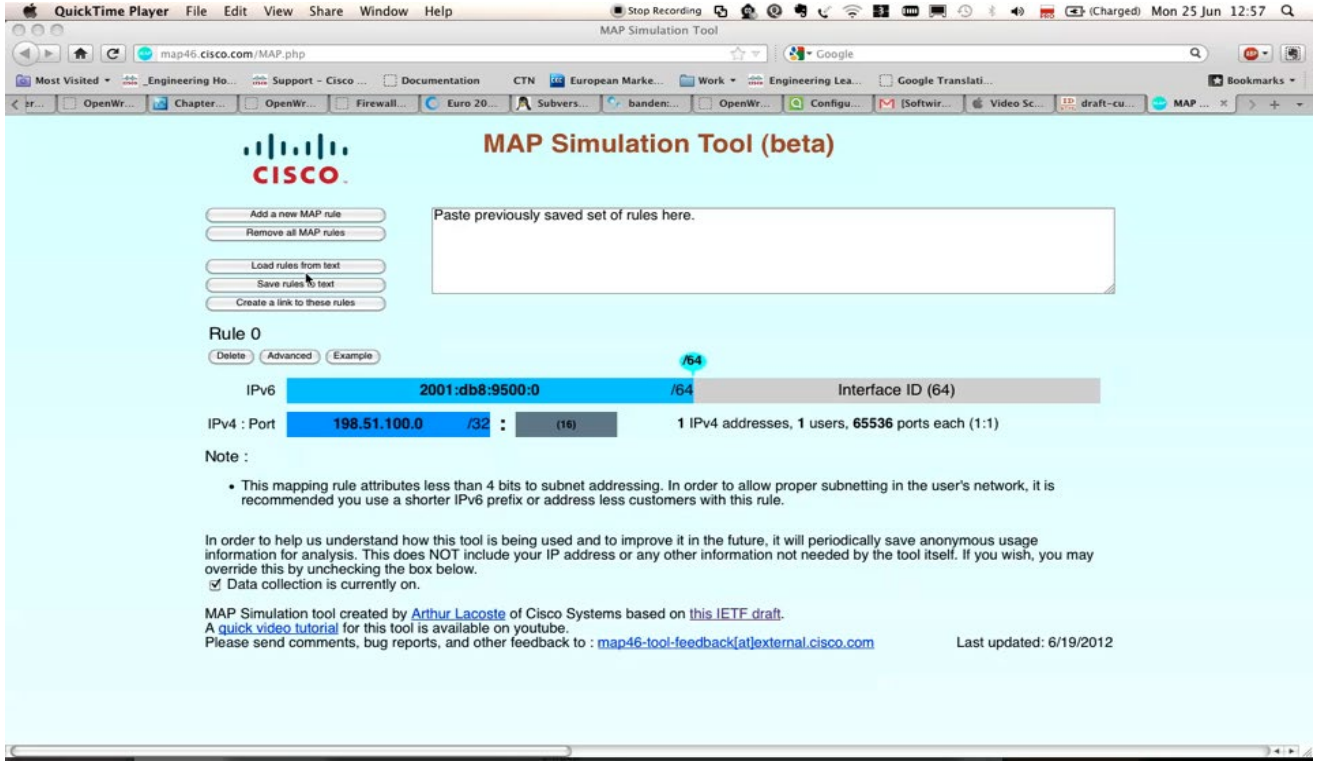

Agenda

- Introduction
- Transition Technologies - Summary
- Transition Technologies - Overview
  - Single-Stack IPv4 – Obtain more IPv4
  - Single-Stack IPv4 – CGN 44, 6rd
  - Dual Stack – Impact (& Happy Eyeballs)
  - Single-Stack IPv6 – DS-Lite, MAP-T/E
  - Single-Stack IPv6 – CGN 64
- IPv4 Address Sharing - Impact
- Conclusion

Single-Stack IPv6 (IPv6-only)

Disable IPv4

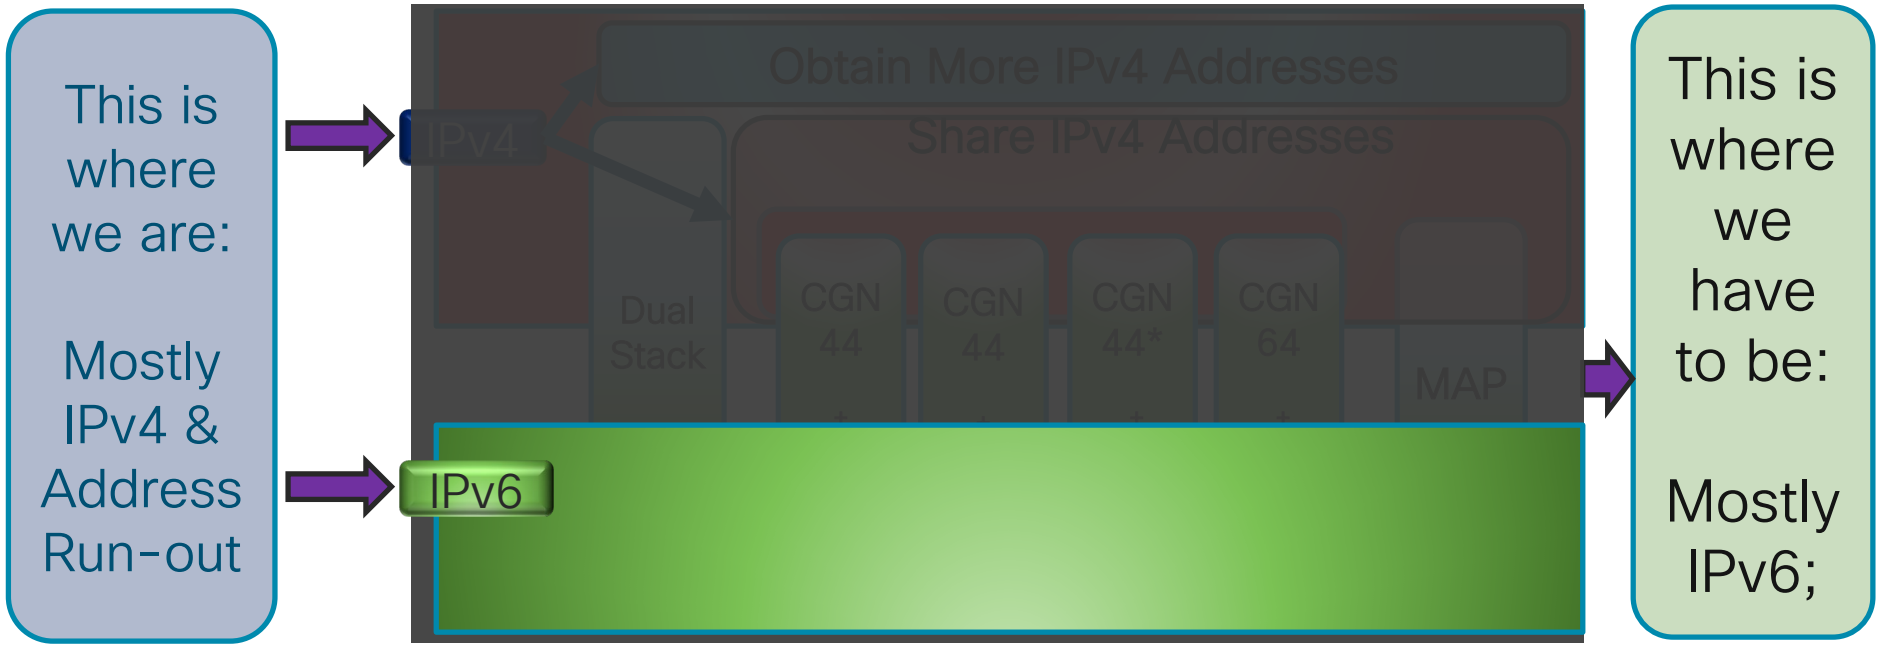

While Client-side apps (mobile or desktop) got IPv6-only support, few Server-side e.g. FaceTime, iMessage, iCloud etc. still need to catch up... ☹

Hence, the short-term need for NAT64...

Single-Stack IPv6 (IPv6-only) with CG NAT64

Disable IPv4

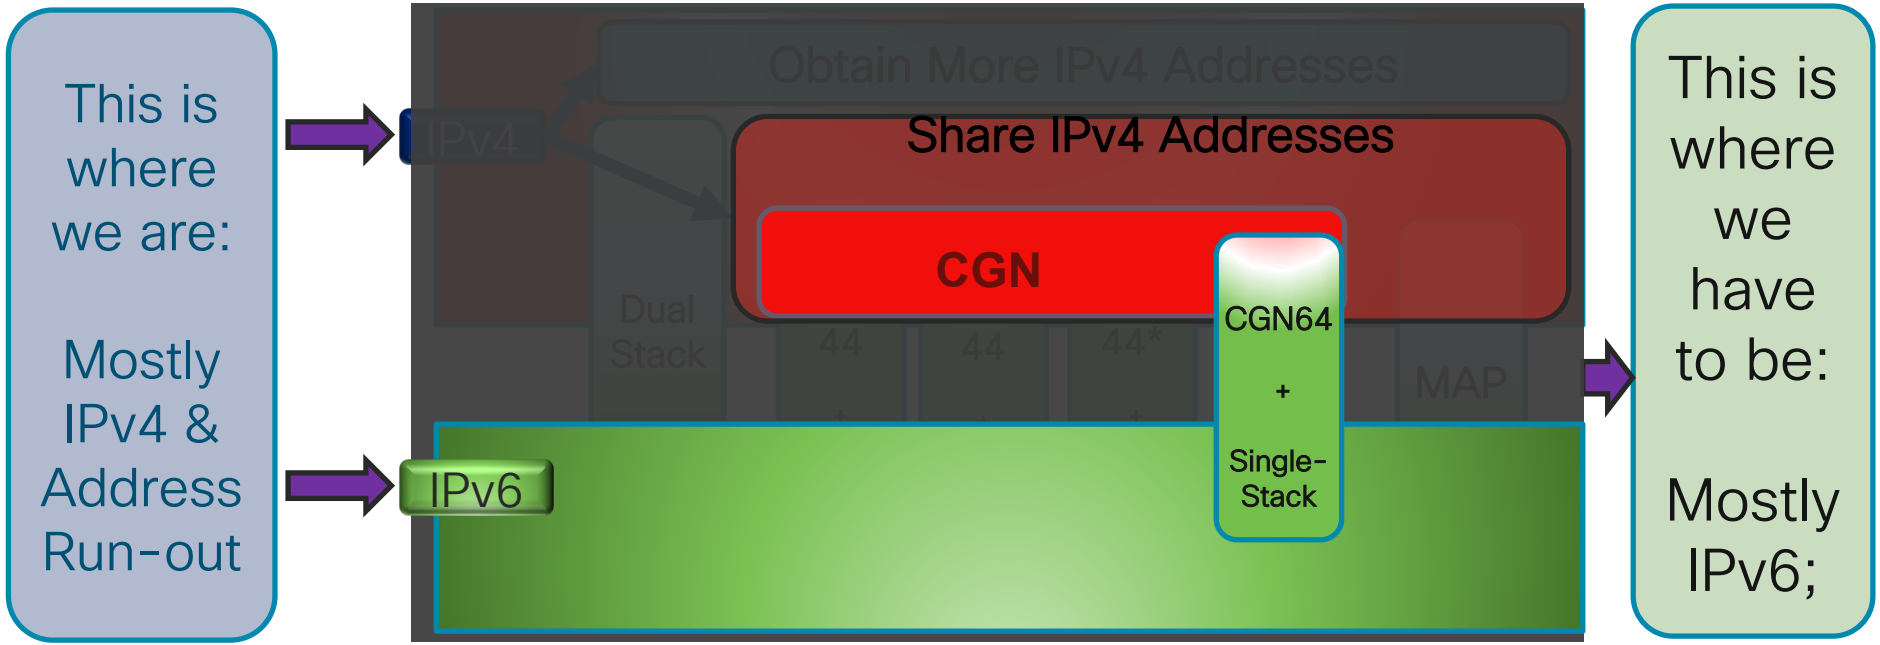

Single-Stack IPv6 (IPv6-only) with CG NAT64

NAT64 supported on ASR9K, ASR1K, CRS

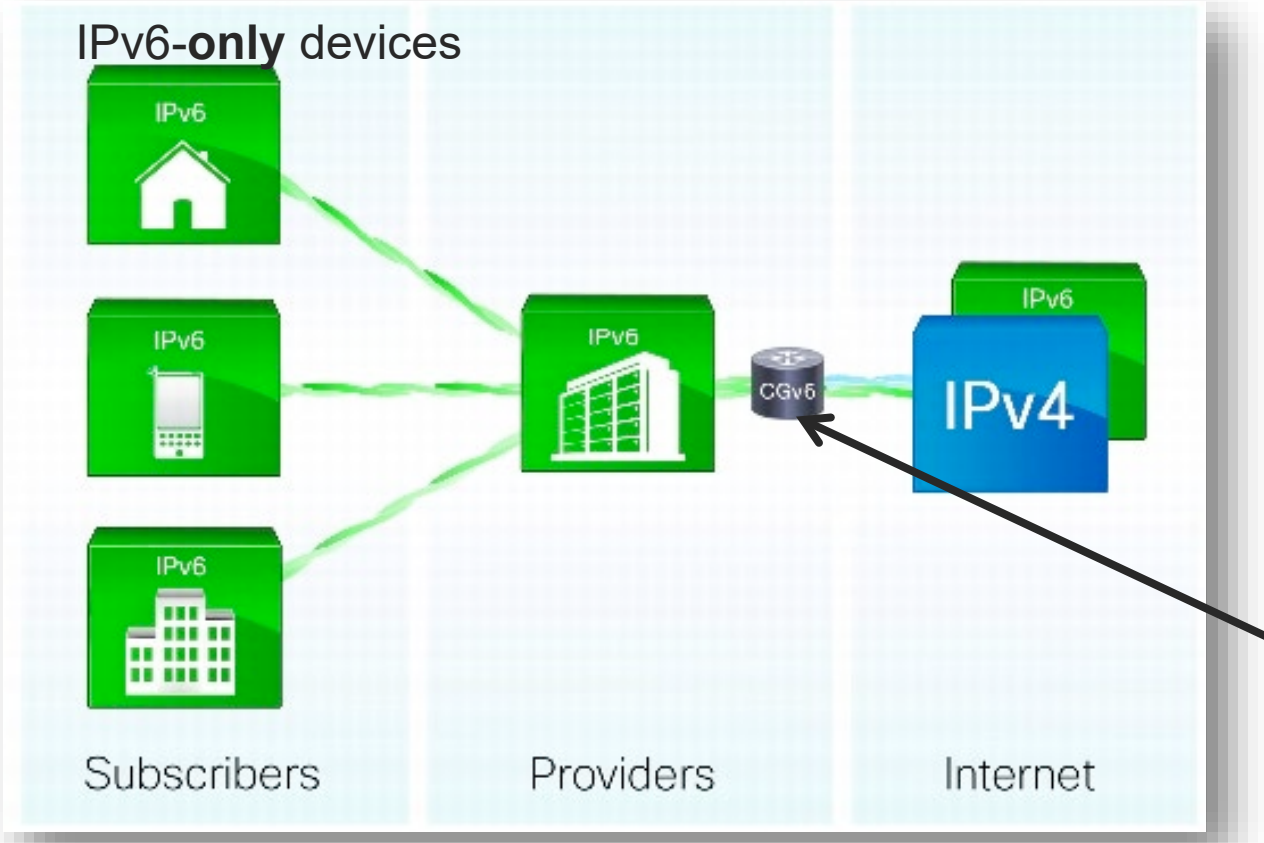

Stateless or Stateful NAT64 function for IPv4 reachability (on routers etc.)

NAT64 – Stateful

Supported on ASR9K, ASR1K, CRS

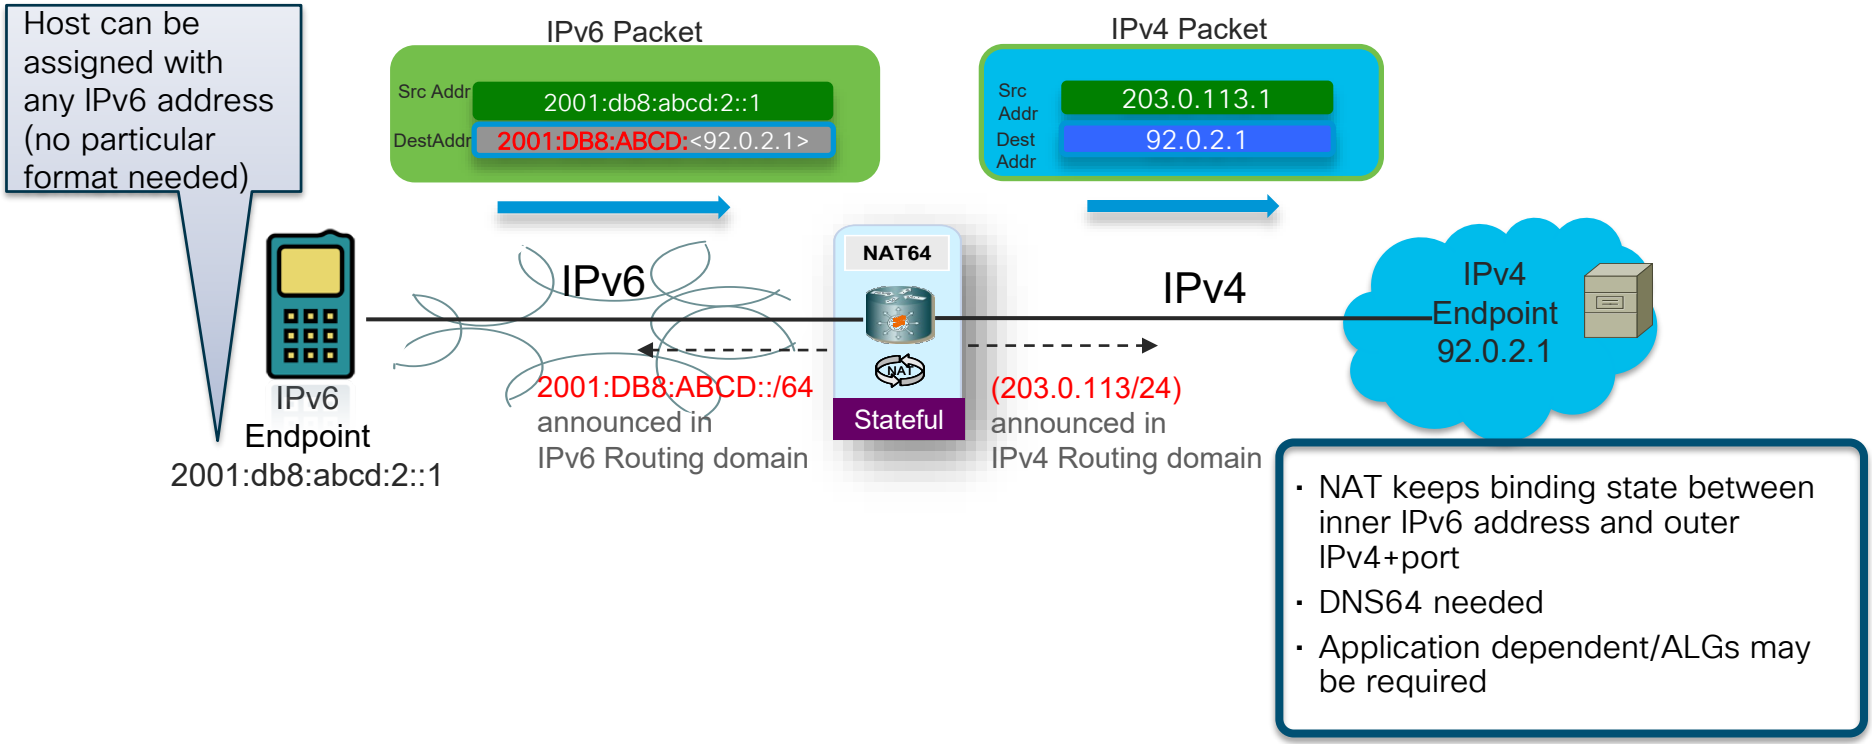

- NAT keeps binding state between inner IPv6 address and outer IPv4+port
- DNS64 needed
- Application dependent/ALGs may be required

NAT64 – Stateless

Supported on ASR9K, ASR1K, CRS

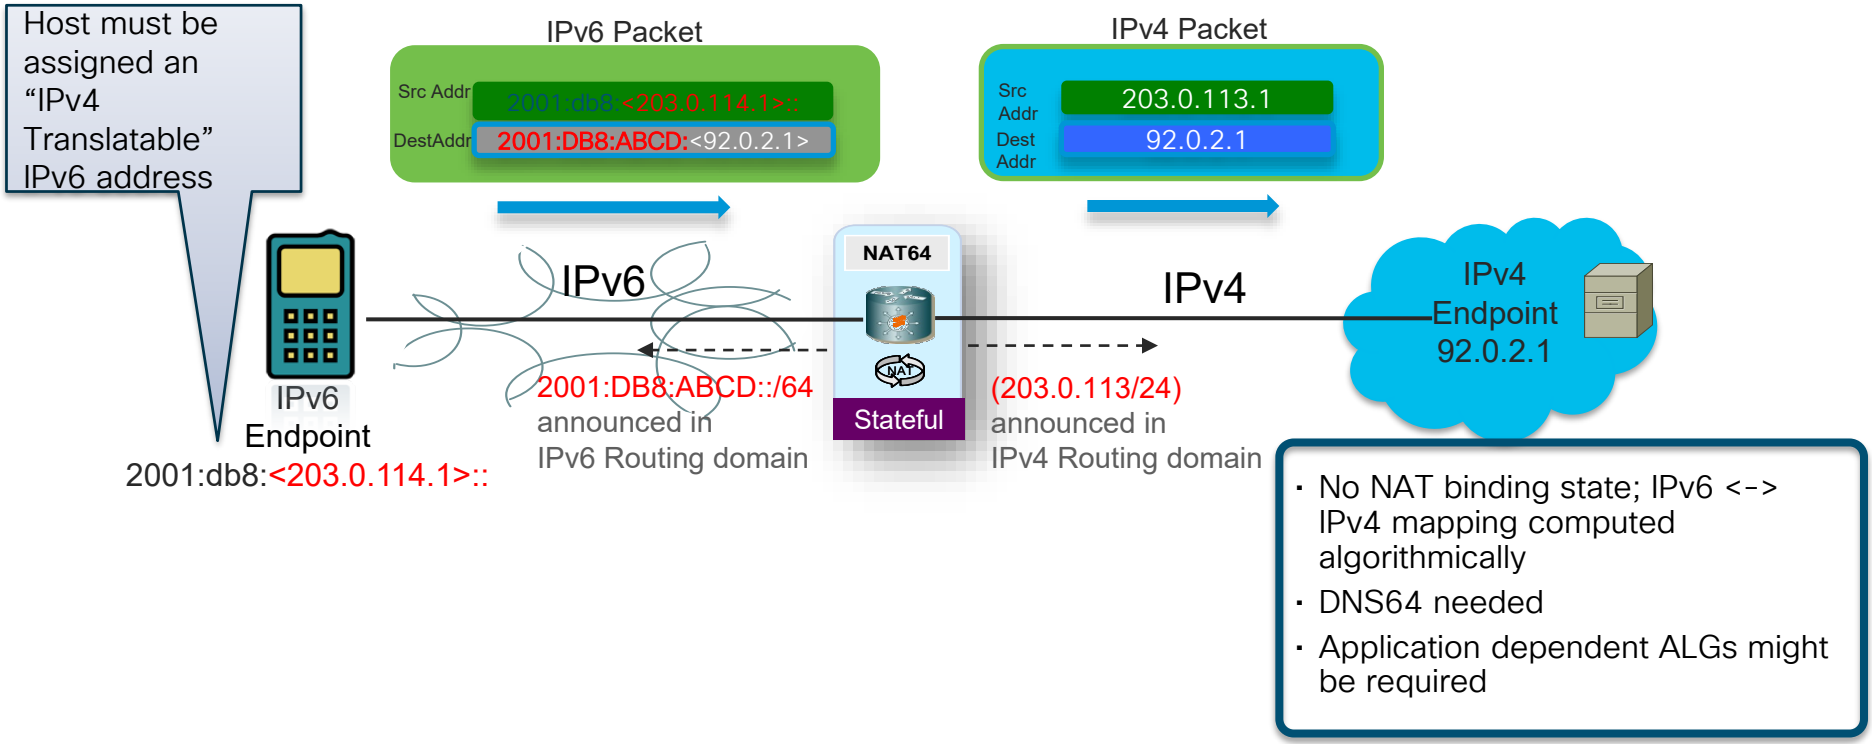

NAT64 – Stateful vs. Stateless

Stateful

- 1:N translation
- “NAPT”
- TCP, UDP, ICMP
- Shares IPv4 addresses

Stateless

- 1:1 translation
- “NAT”
- Any protocol
- No IP4 address sharing (or saving; see note1)
  - Just like dual-stack

Note 1:MAP however **does** save IPv4 addresses by combining NAT64 with NAT44

Note 2: IPv6-only DC using Stateless64 : RFC7755

DNS64 is usually needed with NAT64

- NAT64 translator is useful only if the traffic can come to it
  - IP addresses of IPv6 packets must be formulated accordingly
- DNS64 provides conversion of an IPv4 address into an IPv6 address
  - AAAA record is made up from A record (only if upstream AAAA not present) using IPv6 prefix of NAT64 translator (e.g. 2001:DB8:ABCD::)

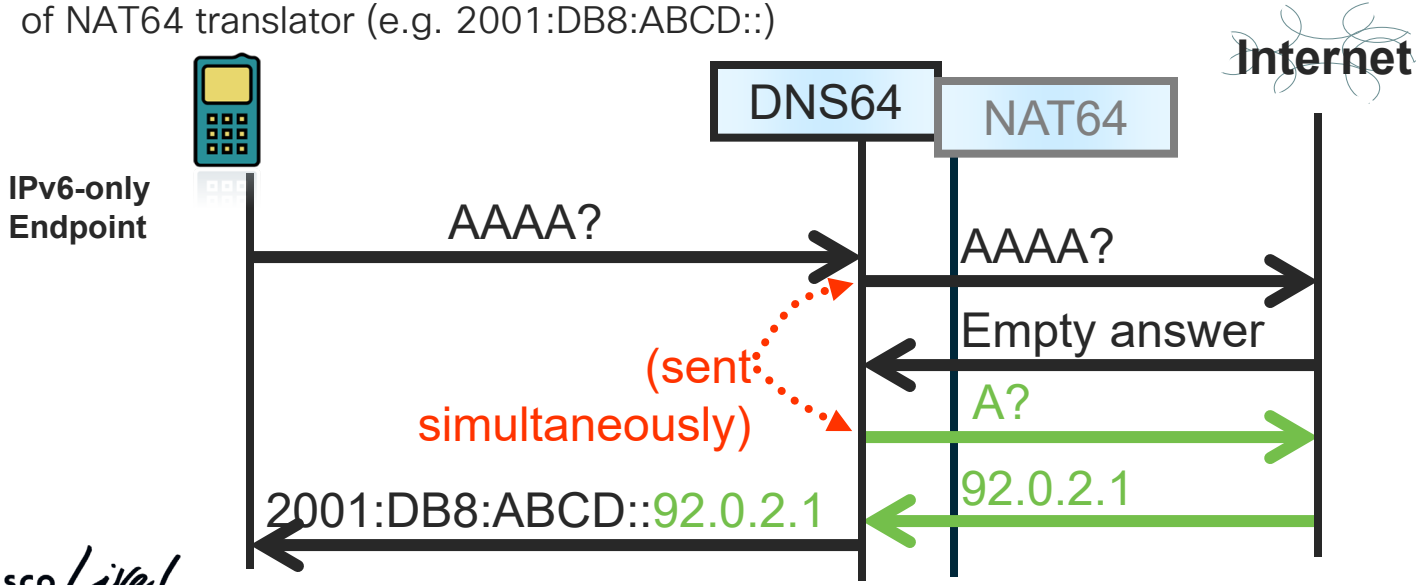

DNS64 – Watch out

- Works for applications that do DNS queries
  - http://www.example.com
  - IMAP, connecting to XMPP servers, etc.
- Works with DNSSEC (note [1])

- Doesn't work for applications that don't do DNS queries or use IP address literals
  - http://1.2.3.4
  - SIP, RTSP, H.323, XMPP peer to peer, etc.
- Doesn't work well if Application-level proxy for IP address literals (HTTP proxy) is used
  - Learn NAT64's prefix, RFC 7050
- NAT46/BIH (Bump In the Host), RFC6535
- 464XLAT (RFC6877)

[1] https://blog.apnic.net/2016/06/09/lets-talk-ipv6-dns64-dnssec/

464XLAT = Stateless + Stateful Better Together ☺ RFC6877

- Some applications may not work with IPv6-only
  - Apple- No issue; Android – Listed here*, but way out-of-date
- 464 translation helps most of those IPv4-only applications
  - Endpoint does “Stateless NAT64”;
  - Network does “Stateful NAT64”
- ***Benefit***: *Network can move to IPv6-only while allowing for any IPv4-only apps and ensuring seamless customer experience*

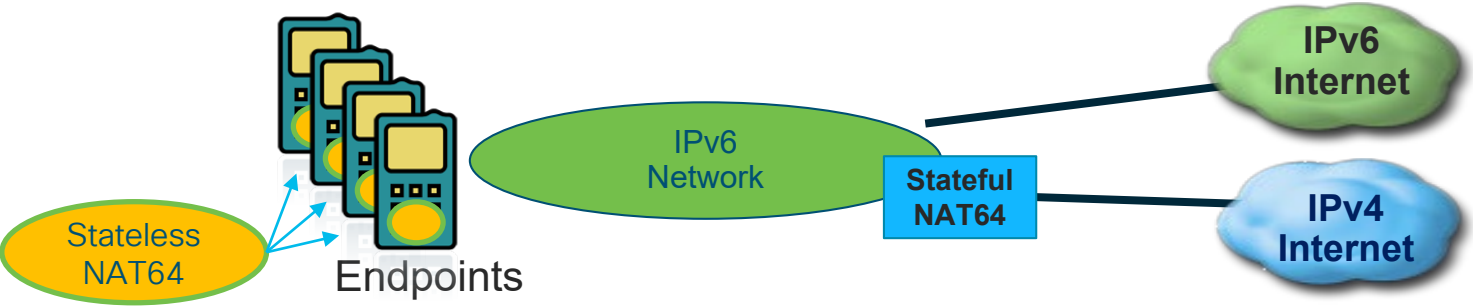

Note: The usefulness of XLAT may continue to subside, given apple mandate for apps to work with IPv6-only since 2016, as well as Cloud Providers enabling IPv6-only support

(exception: tethering)

* http://tinyurl.com/nat64-breakage

IPv6-Only Networks with XLAT 464

PLAT on ASR9K, ASR1K, CRS

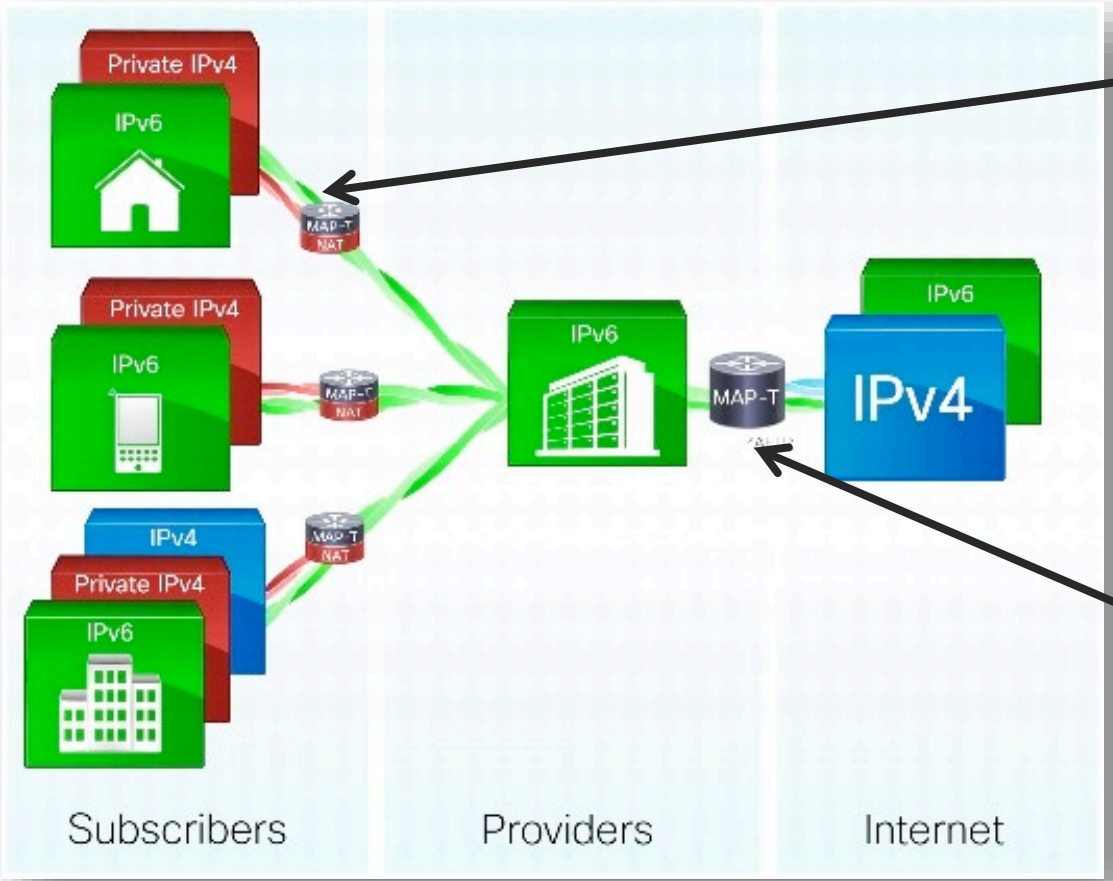

Stateless NAT64 function (on UE, CPE etc.)

Note: XLAT 464 is opposite of MAP-T

Stateful NAT64 function in the network (on routers etc.)

Many NAT64 Scenarios
#1 has been in focus for now

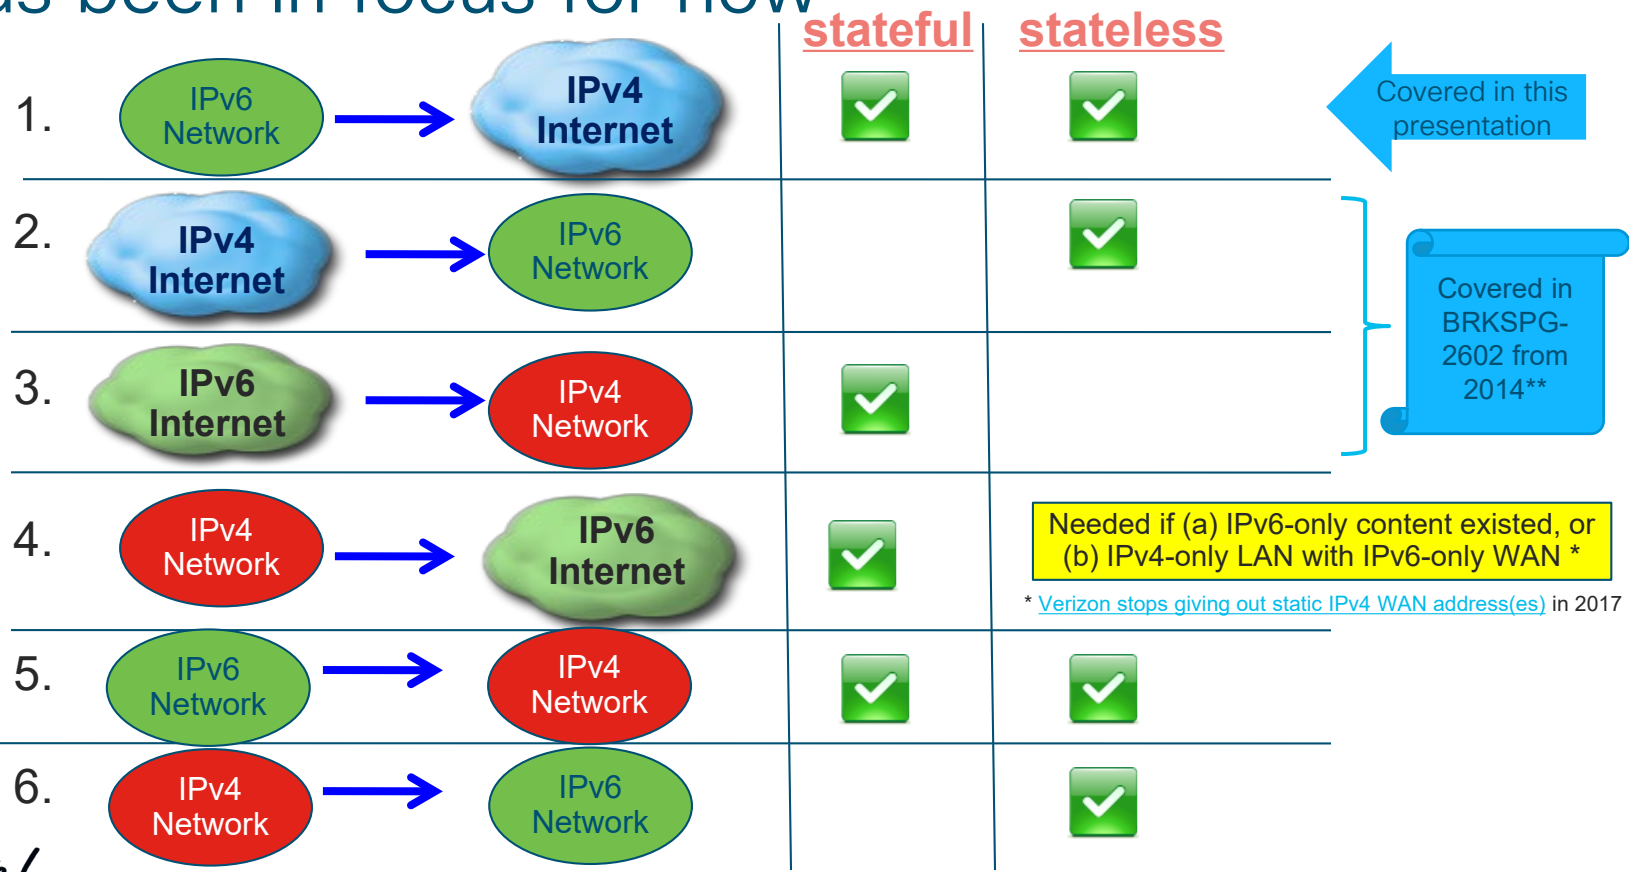

Agenda

- Introduction
- Transition Technologies - Summary
- Transition Technologies - Overview
  - Single-Stack IPv4 – Obtain more IPv4
  - Single-Stack IPv4 – CGN 44, 6rd
  - Dual Stack – Impact (& Happy Eyeballs)
  - Single-Stack IPv6 – DS-Lite, MAP-T/E
  - Single-Stack IPv6 – CGN 64
- IPv4 Address Sharing - Impact
- Conclusion

IP Address Sharing

Reputation matters..

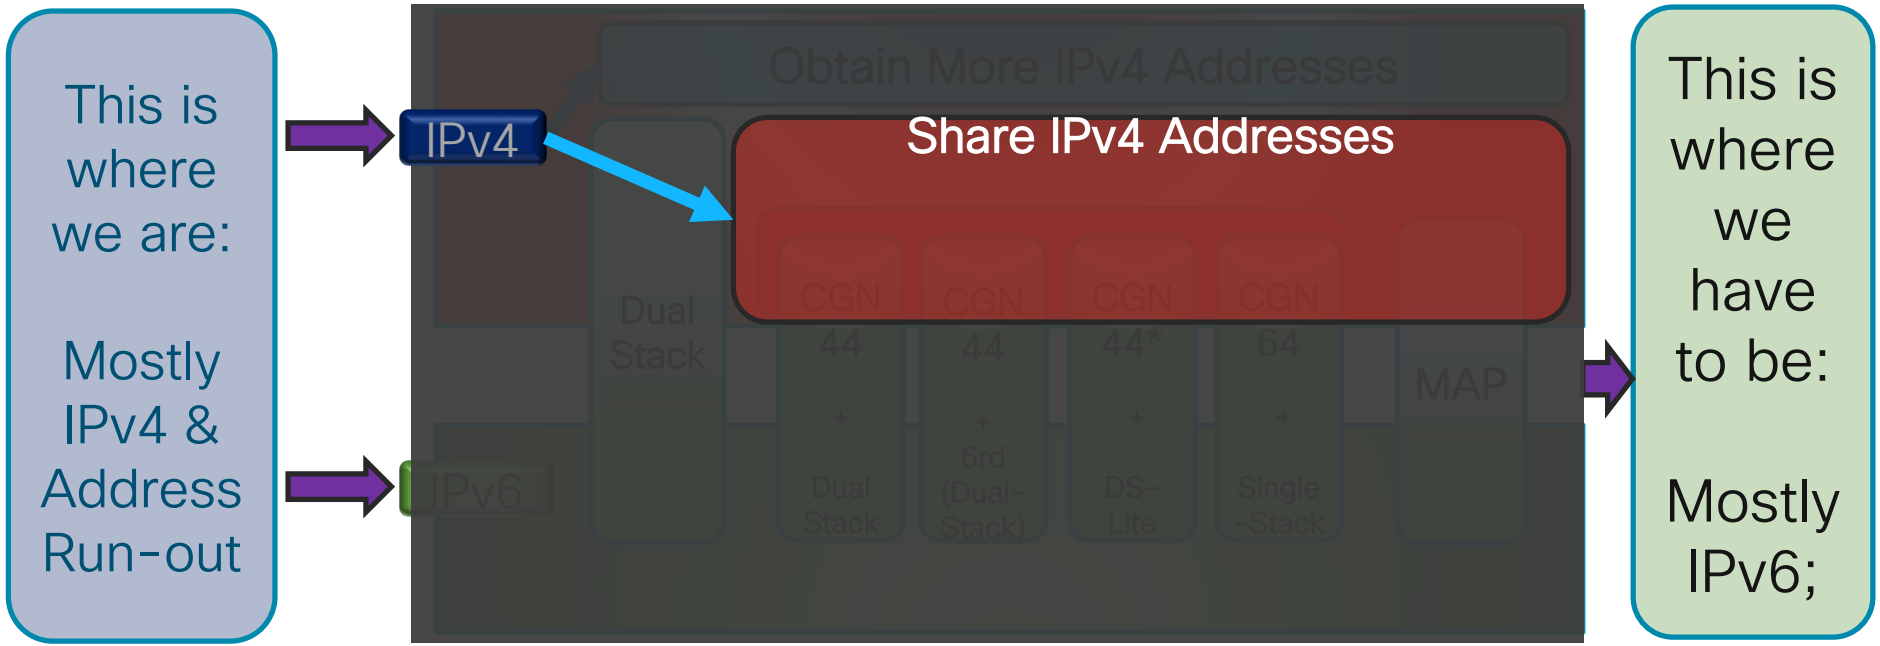

IP Address Sharing : Watch out for IP Reputation

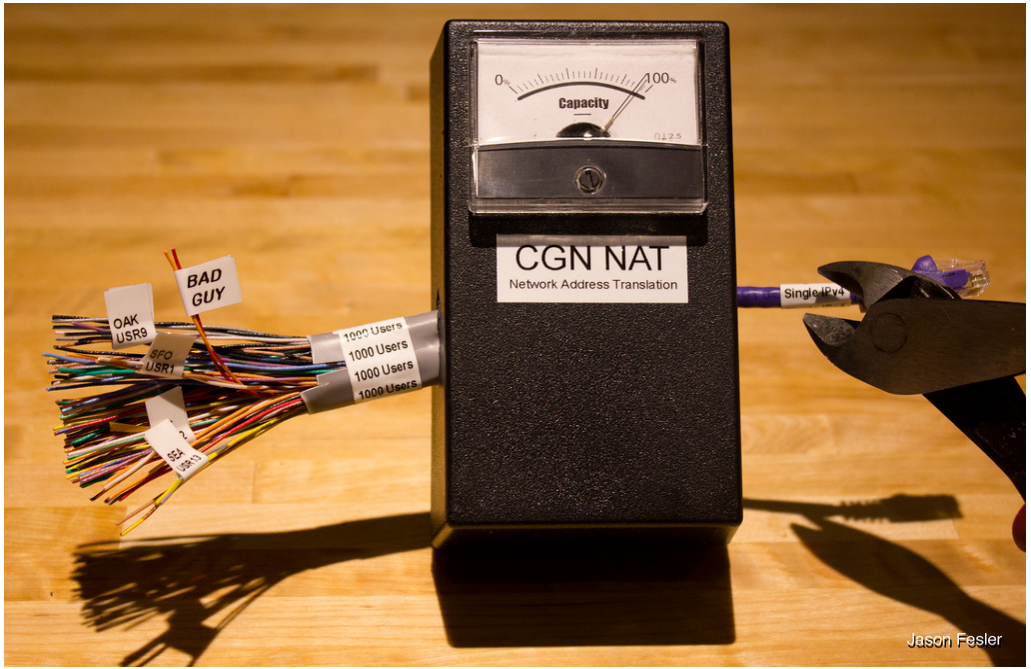

Image source: Jason Fesler, Yahoo!

- Reputation is calculated based on number of things..
  - e.g. Geo Location vs. Owner vs. ..
- Reputation impact varies depending on the use-case
  - e.g. CPEs vs. mailserver vs. content source vs. gaming server vs..

IP Address Sharing: Watch out for IP Reputation (1/2)

- Reputation based on IPv4 address
  - Shared IP address = shared suffering
- Workaround: Distinguish subscribers (sharing IP address, or not sharing)
  - draft-ietf-intarea-nat-reveal-analysis
  - draft-wing-nat-reveal-option
- Server logs currently only contain IPv4 address
  - Servers logs need to include source port number, recommended by RFC6302
- Best Solution – have users and content providers use IPv6!

IP Address Sharing: Watch out for IP Reputation (2/2)

- Affects NATs, as everyone knows
  - NAT44 (CGN44): a big NAT operated by an ISP, enterprise, or University
  - NAT444 (subscriber's NAT44 + ISP's CGN44)
  - NAT64 (CGN64)
  - DS-Lite (called "AFTR" = Modified CGN44)

- Also affects non-CGN architectures!
  - MAP (Mapped Address and Port)
  - Conceptually, a CGN with (some) fixed ports
  - Address + Port, SD-NAT, Deterministic NAT

Agenda

- Introduction
- Transition Technologies - Summary
- Transition Technologies - Overview
  - Single-Stack IPv4 – Obtain more IPv4
  - Single-Stack IPv4 – CGN 44, 6rd
  - Dual Stack – Impact (& Happy Eyeballs)
  - Single-Stack IPv6 – DS-Lite, MAP-T/E
  - Single-Stack IPv6 – CGN 64
- IPv4 Address Sharing - Impact
- Conclusion

Conclusion

More stateless, More IPv6, the better..

* Allows both arbitrary and algorithmic mapping
** Changes needed if IPv6 is not supported by existing CPE

	Options	CPE LAN IPv4 or IPv6	CPE WAN IPv4 or IPv6	Tunnel or Translate?	In-network "State"?	Arbitrary IP addressing of CPE?	Extra CPE features?
0	Single-Stack	IPv4	IPv4	-NA-	-NA-	Yes	No
1	Single-Stack	IPv4	IPv4	Translate	Yes (CGN44)	Yes	No
2	Dual-Stack	IPv4 + IPv6	IPv4+IPv6	-NA-	-NA-	Yes	No**
3	Dual-Stack	IPv4 + IPv6	IPv4+IPv6	Translate	Yes (CGN44)	Yes	No**
4	DS-Lite	IPv4 + IPv6	IPv6	Both	Yes (CGN44)	Yes	Yes
5	6rd	IPv4 + IPv6	IPv4	Tunnel	No	No	Yes
6	6rd + CGN	IPv4 + IPv6	IPv4	Both	Yes (CGN44)	No	Yes
7	MAP	IPv4 + IPv6	IPv6	Either	No	Yes*	Yes
8	Single-Stack	IPv6	IPv6	Translate	Yes (CGN64)	Yes	Yes|No

Conclusion
Drive for (Stateless) Simplicity...be Careful

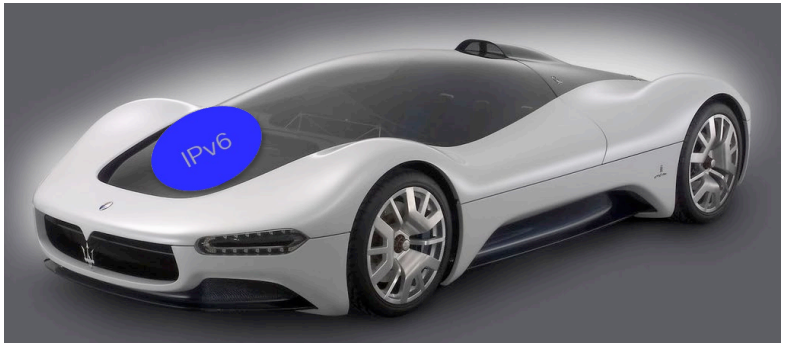

Complete your online session survey

- Please complete your session survey after each session. Your feedback is very important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (starting on Thursday) to receive your Cisco Live t-shirt.
- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Content Catalog on ciscolive.com/emea.

Cisco Live sessions will be available for viewing on demand after the event at ciscolive.com.

cisco Live!

Continue your education

Demos in the Cisco campus

Walk-in labs

Meet the engineer 1:1 meetings

Related sessions

Thank you

You make **possible**

Additional Slides

More IPv6 Sessions

When	Session	Title
29 Jan 2019 / 11:00	BRKIP6-2191	IPv6: The Protocol
29 Jan 2019 / 14:15	LABSPG-3122	Advanced IPv6 Routing and services lab
29 Jan 2019 / 14:30	BRKIP6-2616	Beyond Dual-Stack: Using IPv6 like you've never imagined
30 Jan 2019 / 11:00	BRKSPG-2602	IPv4 Exhaustion: NAT and Transition to IPv6 for Service Providers
30 Jan 2019 / 14:30	BRKIP6-2301	Intermediate - Enterprise IPv6 Deployment
31 Jan 2019 / 08:30	BRKRST-3304	Hitchhiker's Guide to Troubleshooting IPv6 - Advanced
31 Jan 2019 / 11:00	BRKRST-2619	IPv6 Deployment: Developing an IPv6 Addressing Plan and Deploying IPv6
31 Jan 2019 / 11:00	BRKSEC-3200	Advanced IPv6 Security Threats and Mitigation
31 Jan 2019 / 14:00	LTRIPV-2494	IPv6 Transformation Lab
31 Jan 2019 / 14:00	LABSPG-3122	Advanced IPv6 Routing and services lab
	LABIPV-2261	IPv6 planning, deployment and transition
	LABCRS-1000	Intro IPv6 Addressing and Routing Lab