

CS 43: Computer Networks

NAT, ICMP, IPv6

Kevin Webb

Swarthmore College

March 31, 2022

Today

- Internet Control Message Protocol (ICMP)
 - status / meta messages related to IP
- Network Address Translation (NAT)
 - Use one IP address for many devices
- IP version 6 (IPv6)

ICMP: Internet Control Message Protocol

- Used to communicate network information
 - “Control messages”, i.e., not data themselves
 - Error reporting
 - Unreachable host
 - Unreachable network
 - Unreachable port
 - TTL expired
 - Test connectivity
 - Echo request/response (ping)

ICMP: Internet Control Message Protocol


- Header:
 - 1-byte type
 - 1-byte code
 - 2-byte checksum
 - 4 bytes vary by type
- Sits above IP
 - Type 1 in IP header
 - Usually considered part of IP

<u>Type</u>	<u>Code</u>	<u>Description</u>
0	0	echo reply (ping)
3	0	dest. network unreachable
3	1	dest host unreachable
3	2	dest protocol unreachable
3	3	dest port unreachable
3	6	dest network unknown
3	7	dest host unknown
4	0	source quench (congestion control - not used)
8	0	echo request (ping)
9	0	route advertisement
10	0	router discovery
11	0	TTL expired
12	0	bad IP header

Ping Demo

Traceroute and ICMP


- Source sends sets of UDP segments (usually 3) to destination
 - first set has TTL =1
 - second set has TTL=2, etc.
 - unlikely port number
- When n th set of datagrams arrives to n th router:
 - router discards datagrams
 - and sends source ICMP messages (type 11, code 0)
 - ICMP messages includes name of router & IP address
- When ICMP messages arrives, source records RTTs



Traceroute and ICMP

stopping criteria:

- UDP segment eventually arrives at destination host
- destination returns ICMP “port unreachable” message (type 3, code 3)
- source stops



Traceroute Demo


```
6 Episode.IV (206.214.251.1) 68.642 ms 67.307 ms 67.005 ms
7 A.NEW.MOPE (206.214.251.6) 65.986 ms 68.502 ms 68.708 ms
8 It.is.a.period.of.civil.war (206.214.251.9) 67.067 ms 70.139 ms 66.52
9 Rebel.spaceships (206.214.251.14) 70.214 ms 70.192 ms 71.622 ms
10 striking.from.a.hidden.base (206.214.251.17) 71.427 ms 74.206 ms
11 have.won.their.first.victory (206.214.251.22) 71.665 ms 70.434 ms 7
12 against.the.evil.Galactic.Empire (206.214.251.25) 69.218 ms 70.621
13 During.the.battle (206.214.251.30) 69.059 ms 68.931 ms 69.981 ms
14 Rebel.spies.managed (206.214.251.33) 77.247 ms 72.757 ms 77.61
15 to.steal.secret.plans (206.214.251.38) 71.224 ms 71.164 ms 69.543
16 to.the.Empires.ultimate.weapon (206.214.251.41) 68.744 ms 68.824
17 the.DEATH.STAR (206.214.251.46) 72.316 ms 74.551 ms 66.354 ms
18 an.armored.space.station (206.214.251.49) 69.413 ms 70.334 ms 6
19 with.enough.power.to (206.214.251.54) 66.182 ms 66.627 ms 71.23
20 destroy.an.entire.planet (206.214.251.57) 71.926 ms 71.266 ms 70.
21 Pursued.by.the.Empires (206.214.251.62) 67.298 ms 65.956 ms 66.
22 sinister.agents (206.214.251.65) 65.020 ms 67.806 ms 70.508 ms
23 Princess.Leia.races.home (206.214.251.70) 68.894 ms 71.147 ms 71
24 aboard.her.starship (206.214.251.73) 72.130 ms 71.093 ms 74.026
25 custodian.of.the.stolen.plans (206.214.251.78) 68.568 ms 67.939 ms
26 that.can.save.her (206.214.251.81) 67.063 ms 69.874 ms 68.889 m
27 people.and.restore (206.214.251.86) 70.395 ms 70.144 ms
28 freedom.to.the.galaxy (206.214.251.89) 66.098 ms 65.432 ms
29 0-----0 (206.214.251.94) 75.931 ms 74.159 ms 80.012
30 0-----0 (206.214.251.97) 73.026 ms 73.403 ms 73.256
31 0-----0 (206.214.251.102) 83.602 ms 82.079 ms 70.743
32 0-----0 (206.214.251.105) 70.459 ms 69.403 ms 68.782 m
33 0-----0 (206.214.251.110) 68.516 ms 72.472 ms 71.811 ms
34 0-----0 (206.214.251.113) 69.056 ms 65.981 ms 68.202 ms
35 0-----0 (206.214.251.118) 66.790 ms 71.556 ms 74.292 ms
36 0-----0 (206.214.251.121) 68.286 ms 71.042 ms 71.587 ms
37 0-----0 (206.214.251.126) 72.702 ms 71.785 ms 72.442 ms
38 0-----0 (206.214.251.129) 78.143 ms 74.411 ms 72.828 ms
39 0-----0 (206.214.251.134) 69.692 ms 66.187 ms 67.369 ms
40 0-----0 (206.214.251.137) 69.184 ms 70.678 ms 67.445 ms
41 0-----0 (206.214.251.142) 70.383 ms 68.220 ms 67.543 ms
42 0-----0 (206.214.251.145) 67.593 ms 72.970 ms 73.220 ms
43 0----0 (206.214.251.150) 70.964 ms 69.082 ms 70.831 ms
44 0---0 (206.214.251.153) 73.856 ms 71.848 ms 70.311 ms
45 0--0 (206.214.251.158) 71.517 ms 69.204 ms 69.538 ms
46 0-0 (206.214.251.161) 68.076 ms 68.179 ms 67.620 ms
47 0-0 (206.214.251.166) 68.738 ms 70.518 ms 68.757 ms
48 00 (206.214.251.169) 68.281 ms 70.225 ms 74.811 ms
49 I (206.214.251.174) 70.203 ms 71.668 ms 71.672 ms
50 By.Ryan.Werber (206.214.251.177) 68.900 ms 71.461 ms 72.297 ms
51 When.CCIEs.Get.Bored (206.214.251.182) 75.816 ms 73.957 ms 71.333 ms
52 read.more.at.beaglenetworks.net (206.214.251.185) 70.254 ms 73.799 ms
```

Recall: IPv4 Addresses

- 32-bit number, must be *globally unique*
- $2^{32} \Rightarrow 4,294,967,296$ possible addresses
- How many do you have?

Address Scarcity

RIPE NCC Begins to Allocate IPv4 Address Space From the Last /8

14 Sep 2012

On Friday 14 September, 2012, the RIPE NCC, the Regional Internet Registry (RIR) for Europe, the Middle East and parts of Central Asia, distributed the last blocks of IPv4 address space from the available pool.

This means that we are now distributing IPv4 address space to Local Internet Registries (LIRs) from the last /8 according to section 5.6 of "IPv4 Address Allocation and Assignment Policies for the RIPE NCC Service Region".

This section states that an LIR may receive one /22 allocation (1,024 IPv4 addresses), even if they can justify a larger allocation. This /22 allocation will only be made to LIRs if they have already received an IPv6 allocation from an upstream LIR or the RIPE NCC. No new IPv4 Provider Independent (PI) space will be assigned.

It is now imperative that all stakeholders deploy IPv6 on their networks to ensure the continuity of their online operations and the future growth of the Internet.

[More information on IPv6 and its deployment, advice from experts and where to get training](#)

[More information on reaching the last /8](#)



OPINION

ARIN Finally Runs Out of IPv4 Addresses



IPv4 Address Cupboards are Bare in North America.

Network World | Sep 22, 2015 7:25 AM PT

RELATED



An insider's guide to the private IPv4 market

Techniques for Prolonging the Lifespan of IPv4



ARIN's registry and transfer policies can help bridge the gap from IPv4 to IPv6

on IDG Answers

If I buy a Chromebook and can't get to grips with OS can I convert to windows?

RELATED TOPICS

Internet

Cisco Subnet

IPv6

6 COMMENTS

It is often said, "the Internet is running out of phone numbers," as a way to express that the Internet is running out of IPv4 addresses, to those who are unfamiliar with Internet technologies. IPv4 addresses, like phone numbers are assigned hierarchically, and thus, have inherent inefficiency. The world's Internet population has been growing and the [number of Internet-connected devices continues to rise](#), with no end in sight. In the next week, the [American Registry for Internet Numbers](#) (ARIN) will have exhausted their supply of IPv4 addresses. The metaphorical IPv4 cupboards are bare. This long-predicted Internet historical event marks opening a new chapter of the Internet's evolution. However, it is somehow anti-climactic now that this date has arrived. The Internet will continue to operate, but all organizations must now accelerate their efforts to deploy IPv6.

INSIDER

Network jobs are hot; salaries expected to rise in 2016

Wireless network

ARIN IPv4 Address Exhaustion

The [Internet Assigned Numbers Authority](#) (IANA) delegates authority for Internet resources to the five RIRs that cover the world. The [American Registry for Internet Numbers](#) (ARIN) is the [Regional Internet Registry](#) (RIR) for the United States, Canada, the Caribbean, and North Atlantic islands. ARIN has been managing the assignment of IPv4 and IPv6 addresses and Autonomous System (AS) numbers for several decades. Each RIR has been managing their limited IPv4 address stores and going through their [various phases of exhaustion policies](#). ARIN has been in [Phase 4](#) of their IPv4 depletion plan for more than a year now. ARIN will soon announce that they have completely extinguished their supply of IPv4 addresses.

Seriously, we're done now. We're done

Exhausted with never-ending internet exhaustion

By [Kieren McCarthy](#) in [San Francisco](#) 15 Feb 2017 at 23:07 214  SHARE ▼



You may have heard this before, but we are really, really running out of public IPv4 addresses.

This week, the regional internet registry responsible for Latin America and the Caribbean, LACNIC, [announced](#) it has moved to "phase 3" of its plan to dispense with the remaining network addresses, meaning that only companies that have not received any IPv4 space are eligible. There is no phase 4.

That means LACNIC is down to its [last 4,698,112 public IPv4 addresses](#) (although that may increase as it recovers a little bit of space over time).

Another Year of the Transition to IPv6



By **Geoff Huston**

Author & Chief Scientist at APNIC


February 20, 2022 | Views: 5,865 | Comments: 2

I bet that nobody believed in 1992 that thirty years later, we'd still be discussing the state of the transition to IPv6! In 1992 we were discussing what to do about the forthcoming address crunch in IPv4, and having come to terms with the inevitable prospect that the silicon industry was going to outpace the capacity of the IPv4 address pool in a couple of years, we needed to do something quickly. We decided to adopt a new protocol, IP version 6, a couple of years later, and in December 1995, the IETF published RFC 1883, the specification of IPv6.

Private Addresses

- Defined in RFC 1918:
 - 10.0.0.0/8 (16,777,216 hosts)
 - 172.16.0.0/12 (1,048,576 hosts)
 - 192.168.0.0/16 (65536 hosts)
- These addresses should *never* be routed.
 - Anyone can use them.
 - Often adopted for use with NAT.

NAT: Network Address Translation



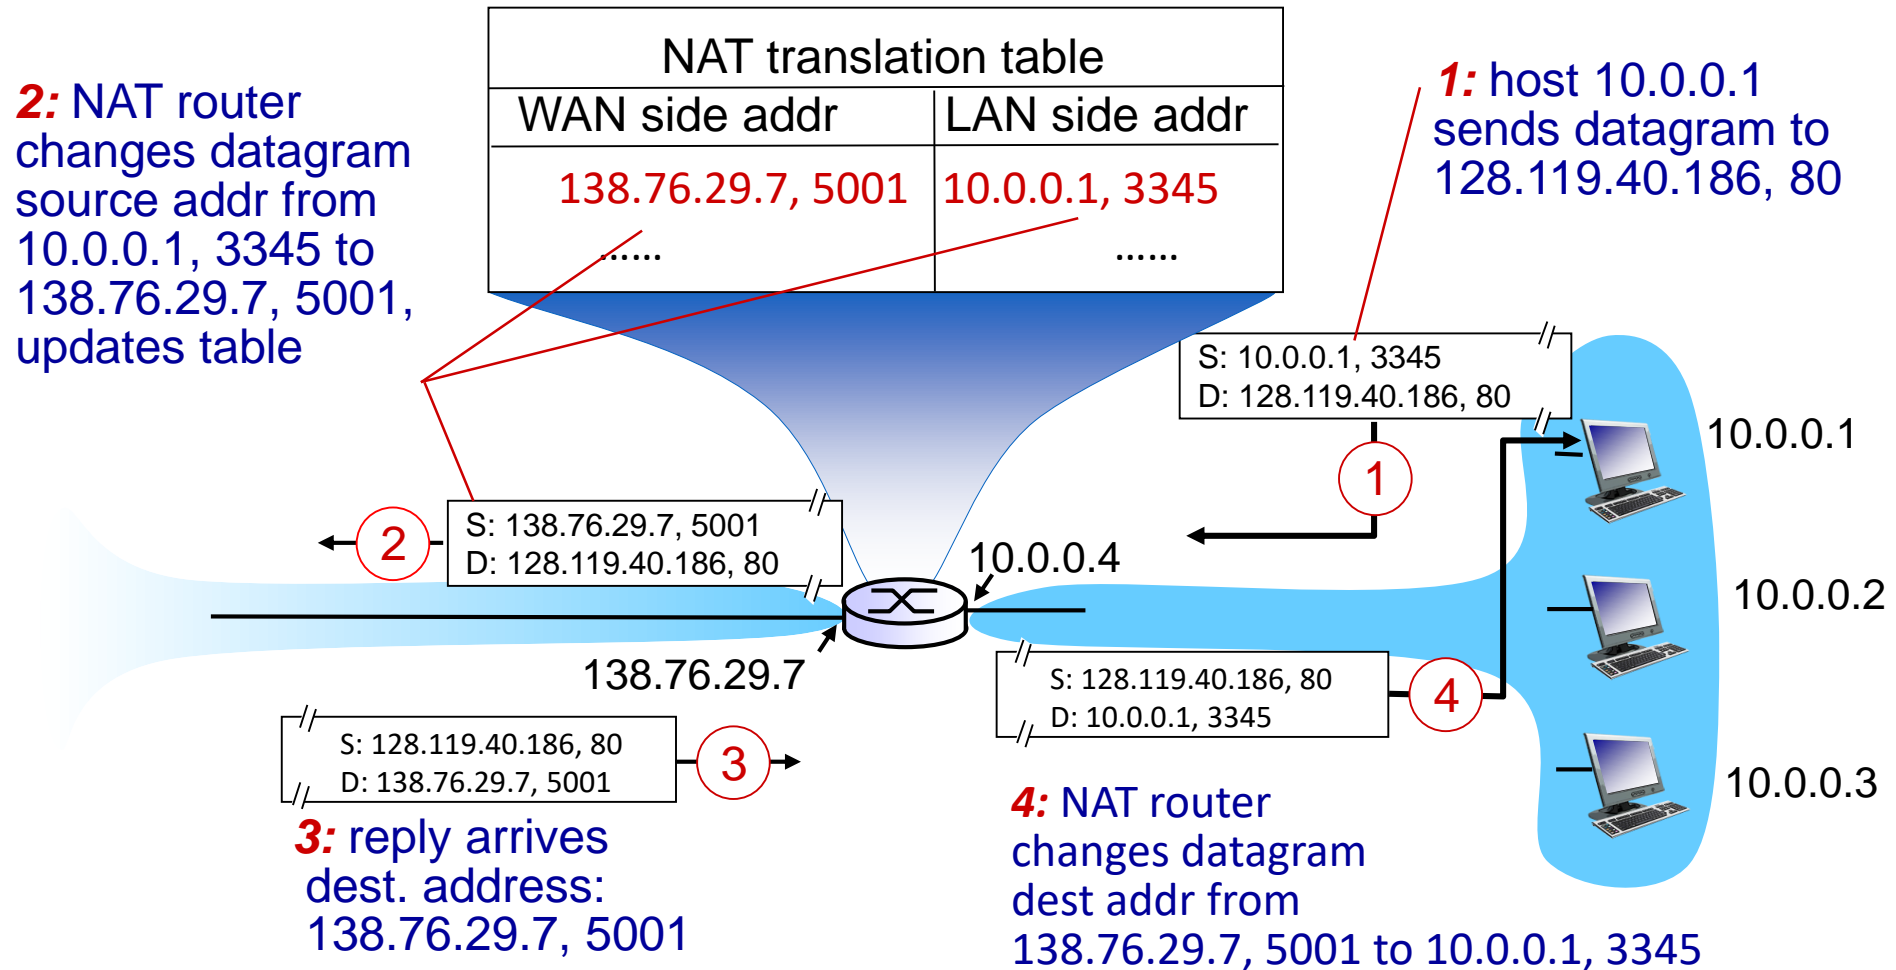
all datagrams *leaving* local network have *same* single source NAT IP address: 138.76.29.7, different source port numbers

datagrams with source or destination in this network have 10.0.0/24 address for source, destination (as usual)

Implementing NAT

- *Outgoing datagrams: replace* (source IP address, port #) of every outgoing datagram to (NAT IP address, new port #)
 - remote clients/servers will respond using (NAT IP address, new port #) as destination address
- *Record (in NAT translation table)* every (source IP address, port #) to (NAT IP address, new port #) translation pair
- *Incoming datagrams: replace* (NAT IP address, new port #) in dest fields of every incoming datagram with corresponding (source IP address, port #) stored in NAT table

NAT: network address translation



NAT: network address translation

NAT translation table	
WAN side addr	LAN side addr
138.76.29.7, 5001	10.0.0.1, 3345

Neither the sender nor receiver need to know that NAT is happening...

// D: 138.76.29.7, 5001

4: NAT router changes datagram dest addr from 138.76.29.7, 5001 to 10.0.0.1, 3345

10.0.0.1

10.0.0.2

10.0.0.3

NAT Advantages

- Organizations need fewer IP addresses from their ISP.
 - With a 16-bit port field, we can put 65535 connections behind one external IP address!
- Organizations can change internal network IPs without having to change outside world IPs.
- Devices inside the local network are not explicitly addressable or visible to the outside world.

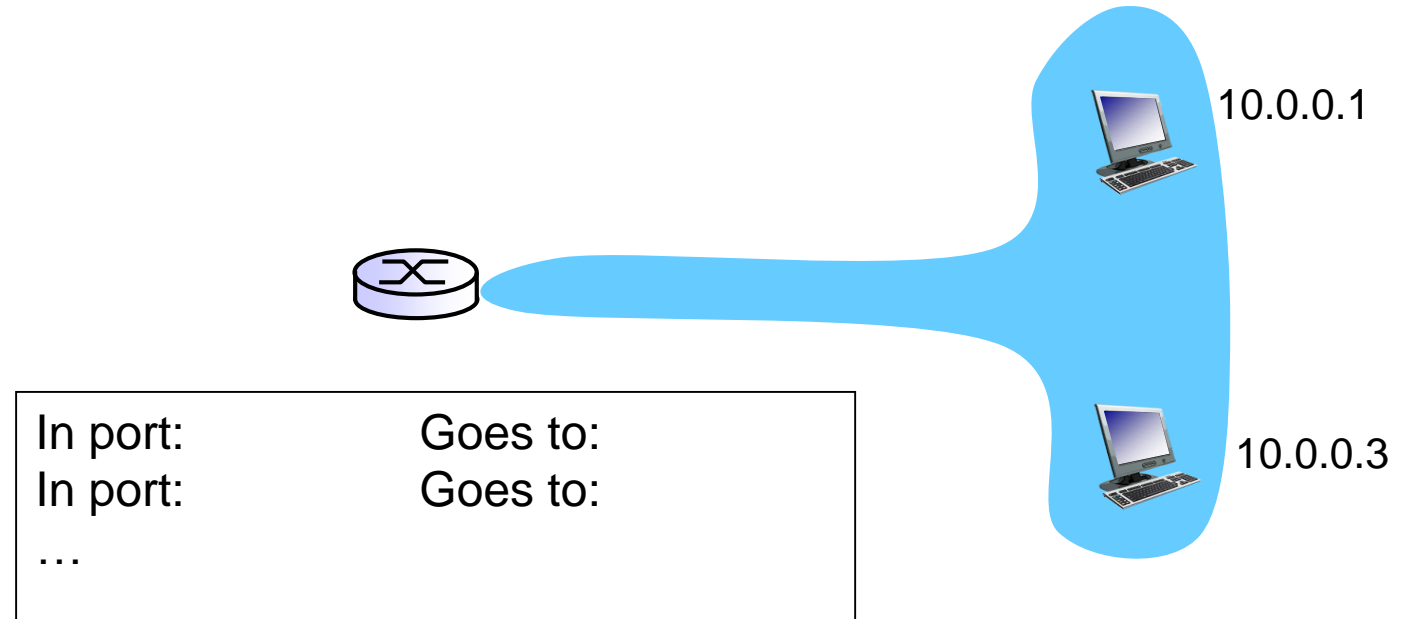
Do you agree that this is an advantage?

Devices inside the local network are not explicitly addressable or visible to the outside world.

- A. This is an advantage.
- B. This is a disadvantage.

Port Forwarding

- What if we wanted to run a web server on both these hosts?



How do we feel about NAT?

- A. NAT is great! It conserves IP addresses and makes it harder to reach non-public machines.
- B. NAT is mostly good, but has a few negative features. No big deal.
- C. NAT is mostly bad, but in some cases, it's a necessary evil.
- D. NAT is an abomination that violates the end to end principle, and we should not use it!

IPv6

- *Initial motivation:* 32-bit address space soon to be completely allocated, any day now™.
- Additional motivation:
 - header format helps speed processing/forwarding
 - header changes to facilitate QoS

IPv6 datagram format:

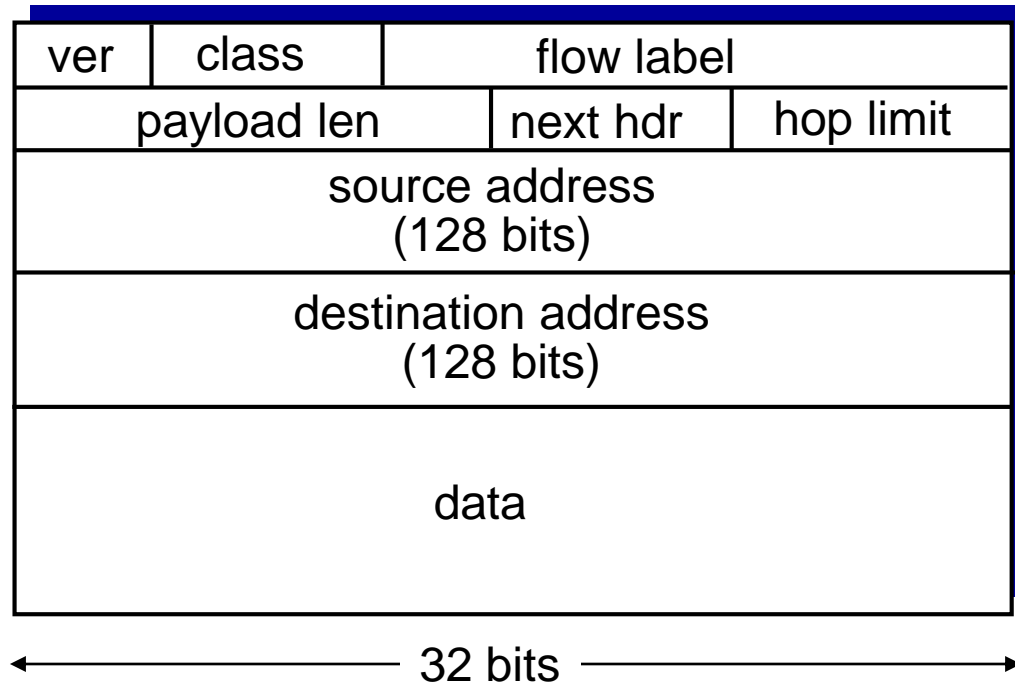
- fixed-length 40 byte header
- no fragmentation allowed

IPv6 datagram format

class: identify class/priority of packet

flow label: identify datagrams in same “flow.”
(purpose of “flow” not well defined).

next header: identify upper layer protocol for data



32-bits -> unique addresses:
4,294,967,296

64-bits -> unique addresses:
18,446,744,073,709,551,616

128-bits -> unique addresses:
340,282,366,920,938,463,463,374,607,431,768,211,456

Other changes from IPv4

- *checksum*: removed entirely to reduce processing time at each hop
- *options*: allowed, but outside of header, indicated by “Next Header” field
- *ICMPv6*: new version of ICMP
 - additional message types, e.g. “Packet Too Big”
 - multicast group management functions

IPv6 (vs. IPv4)

- Simpler, faster, better
- How much traffic on the Internet is IPv6?
- Why!?

IPv6 celebrates its 20th birthday by reaching 10 percent deployment

All I want for my birthday is a new IP header.

ILJITSCH VAN BEIJNUM - 1/3/2016, 12:00 PM

Twenty years ago this month, RFC 1883 was published: [Internet Protocol, Version 6 \(IPv6\) Specification](#). So what's an Internet Protocol, and what's wrong with the previous five versions? And if version 6 is so great, why has it only been adopted by half a percent of the Internet's users each year over the past two decades?

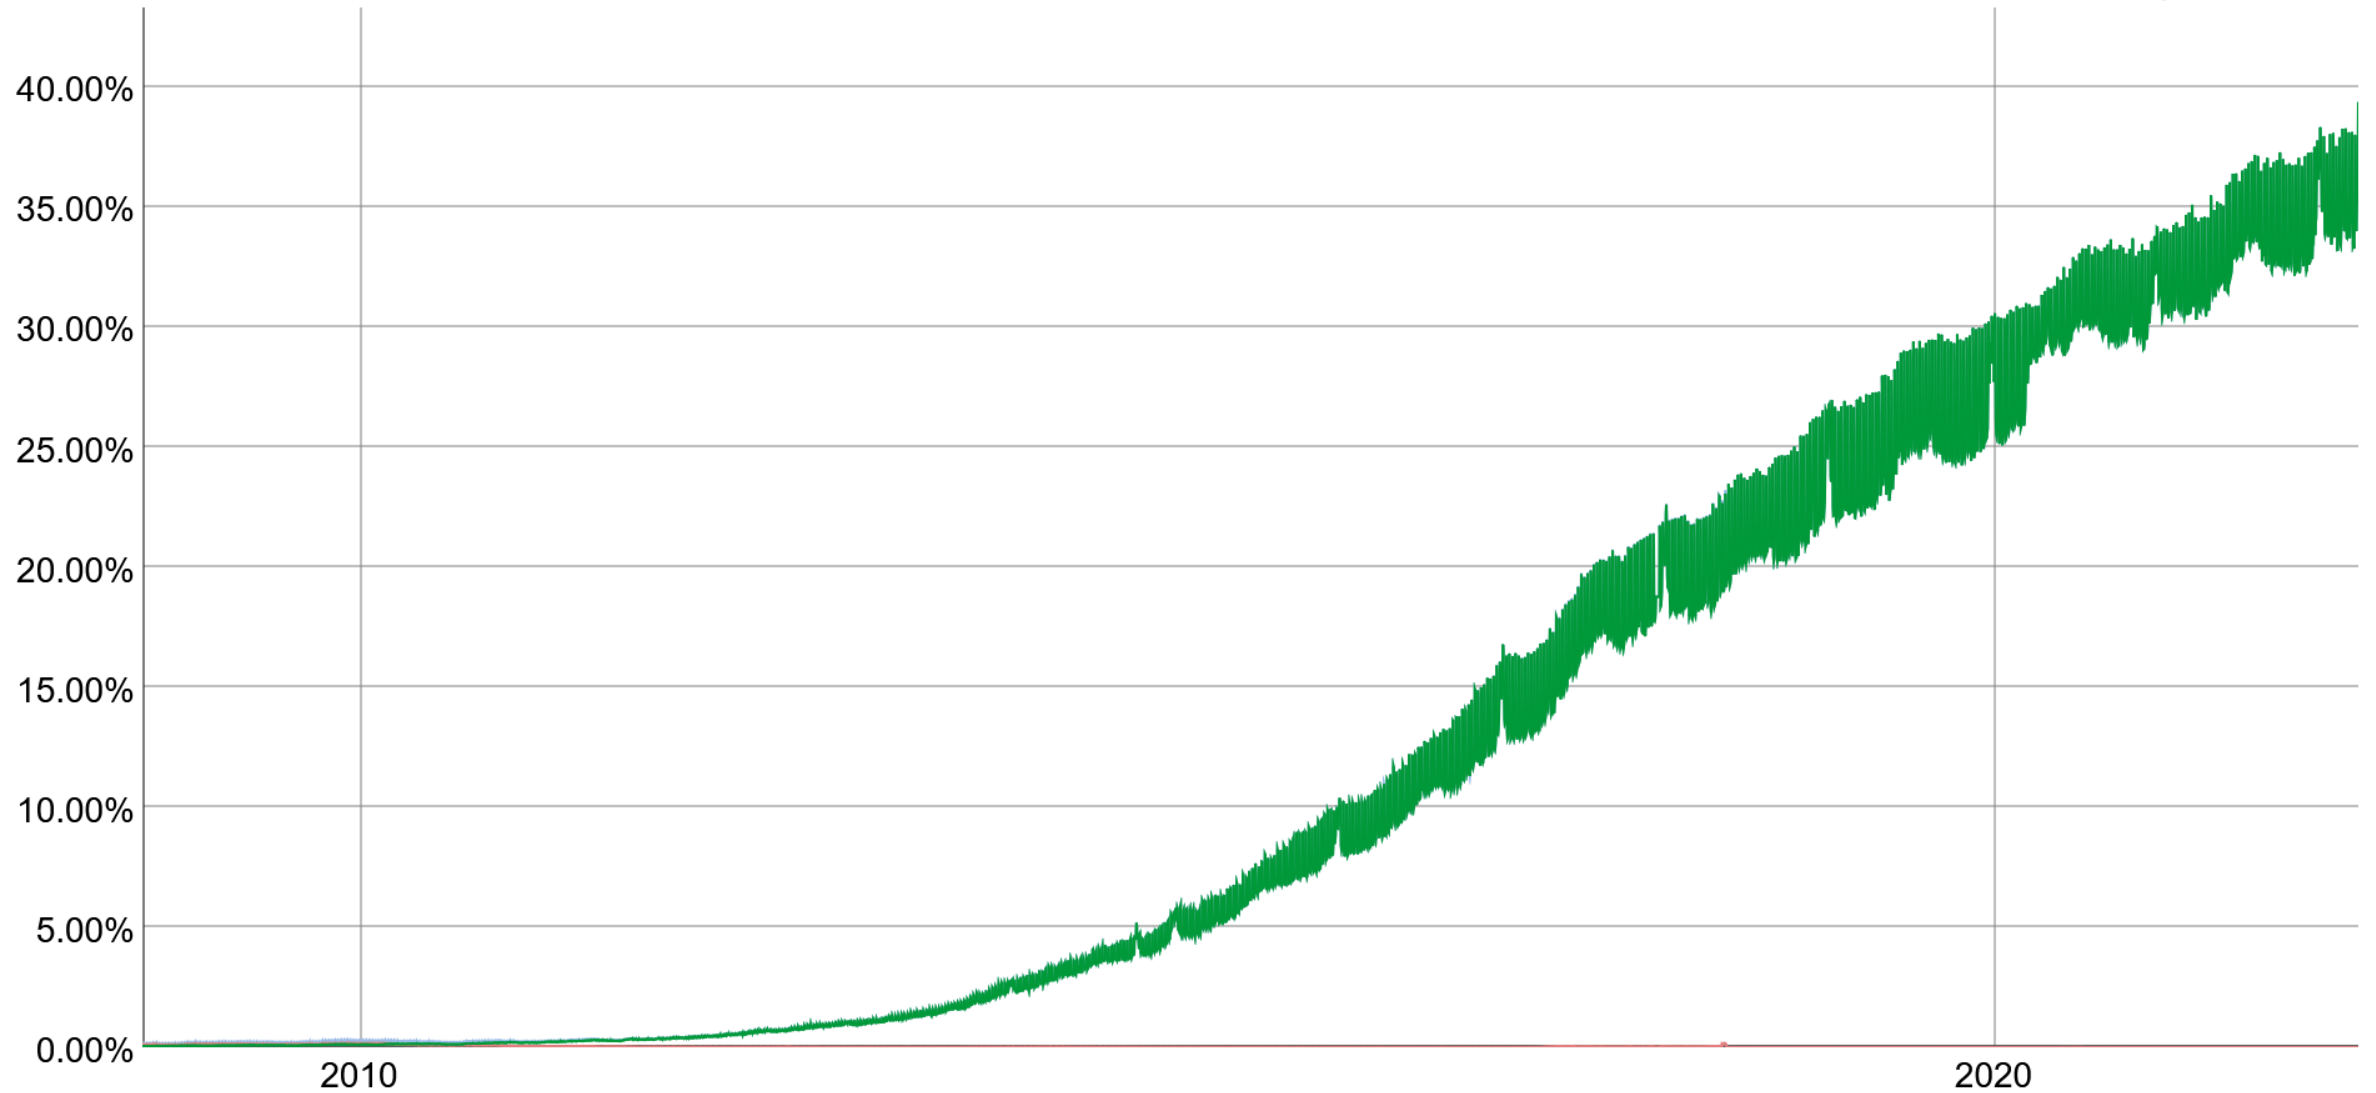
10 percent!

First the good news. According to Google's statistics, on December 26, the world reached 9.98 percent IPv6 deployment, up from just under 6 percent a year earlier. Google measures IPv6 deployment by having a small fraction of their users execute a Javascript program that tests whether the computer in question can load URLs over IPv6. During weekends, a tenth of Google's users are able to do this, but during weekdays it's less than 8 percent. Apparently more people have IPv6 available at home than at work.

IPv6 Adoption

We are continuously measuring the availability of IPv6 connectivity among Google users. The graph shows the percentage of users that access Google over IPv6.

Native: 39.37% **6to4/Teredo: 0.00%** **Total IPv6: 39.37%** | Mar 26, 2022



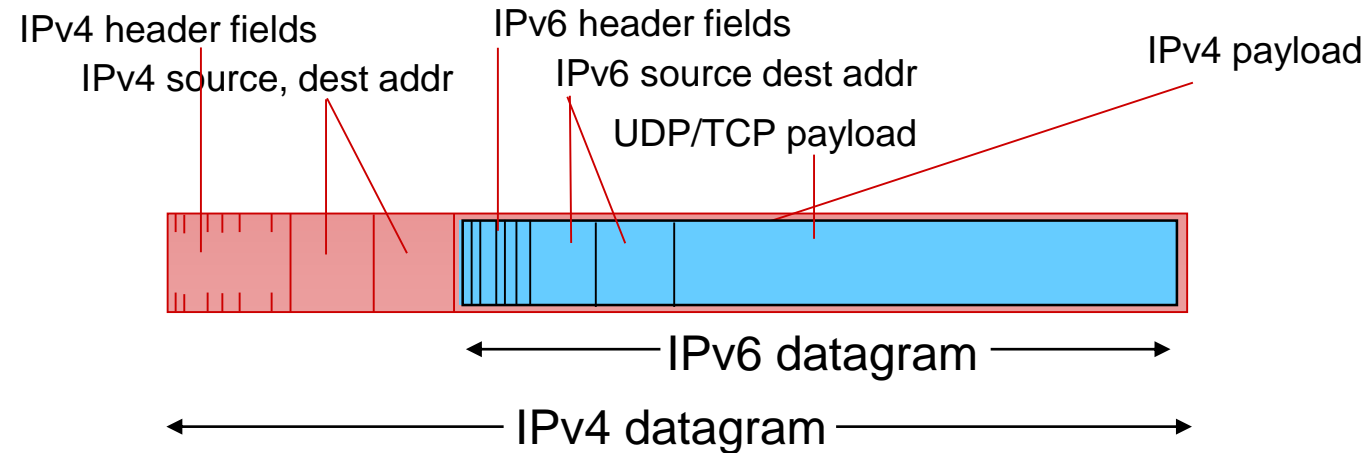
Source: <https://www.google.com/intl/en/ipv6/statistics.html>

Transitioning to IPv6


- Option 1: “Flag day”
 - How do we get *everyone* on the Internet to agree?
 - Whose authority to decide when?
 - Can you imagine how much would break?
- Option 2: Slow transition
 - Some hosts/routers speak both versions
 - Must have some way to deal with those who don't
 - Lack of incentive to switch

Tunneling

- IPv6 datagram carried as *payload* in IPv4 datagram among IPv4 routers



Tunneling



Summary

- Lots of mechanisms surrounding IP
 - NAT: translate routable IP address to multiple private IP addresses using ports/header rewriting
 - ICMP: small status messages, usually report errors
- IPv6: solution to IP address scarcity
 - adoption improving, but still < 50%