Advanced Computer Networking

CYBR 230 – Jeff Shafer – University of the Pacific

IPv6

IP Versions

Version	Description
0-3	Unused: Development versions of IP
4	**Current network-layer protocol**
5	Unused: Experimental stream protocol – ST
6	**New network-layer protocol (1996)**
7-9	Unused: Experimental protocols – TP/IX, PIP, TUBA
10-15	Not allocated

Motivation for IPv6: Scarcity! (Of IP addresses…)

Why Replace IPv4?

- The problem
  - IPv4 has ~4.3 billion addresses
  - World has ~6.6 billion people!
    - How many internet-capable devices per person?
- IP address exhaustion
  - Internet will not "collapse", but new devices / networks will not be able to join(*)
- When? **NOW!** Final rate of consumption was one /8 block (16.78 million addresses) per month
  - Feb 1st, 2011 – Five final /8 blocks handed out to Regional Internet Registries (RIRs)
  - RIR supply ran out within months

(*) Except via address translation…

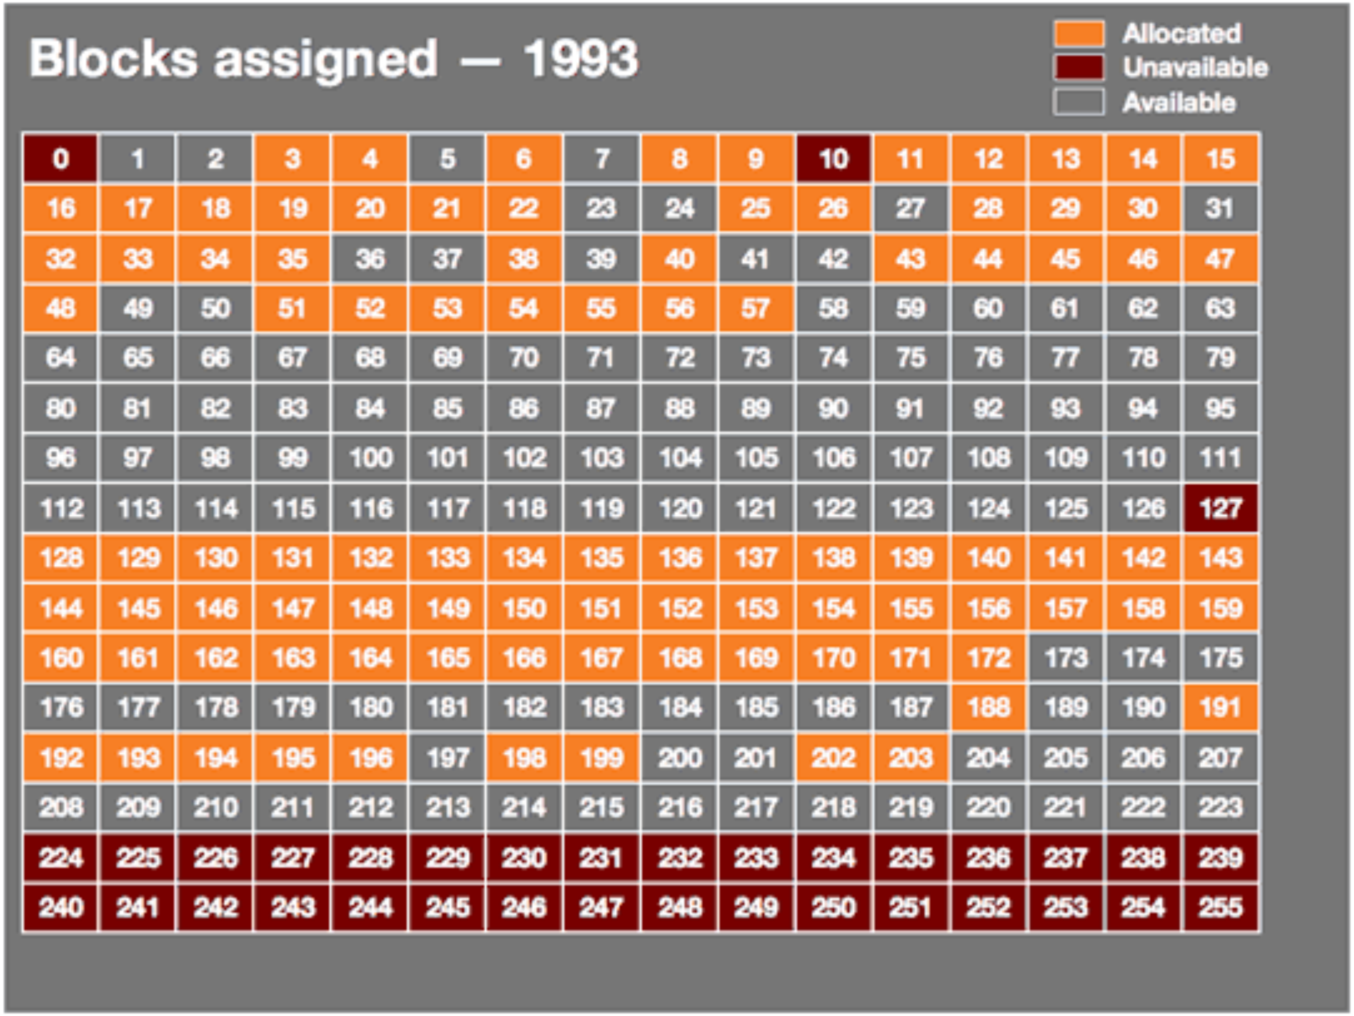

Figure from http://arstechnica.com/articles/paedia/IPv6.ars

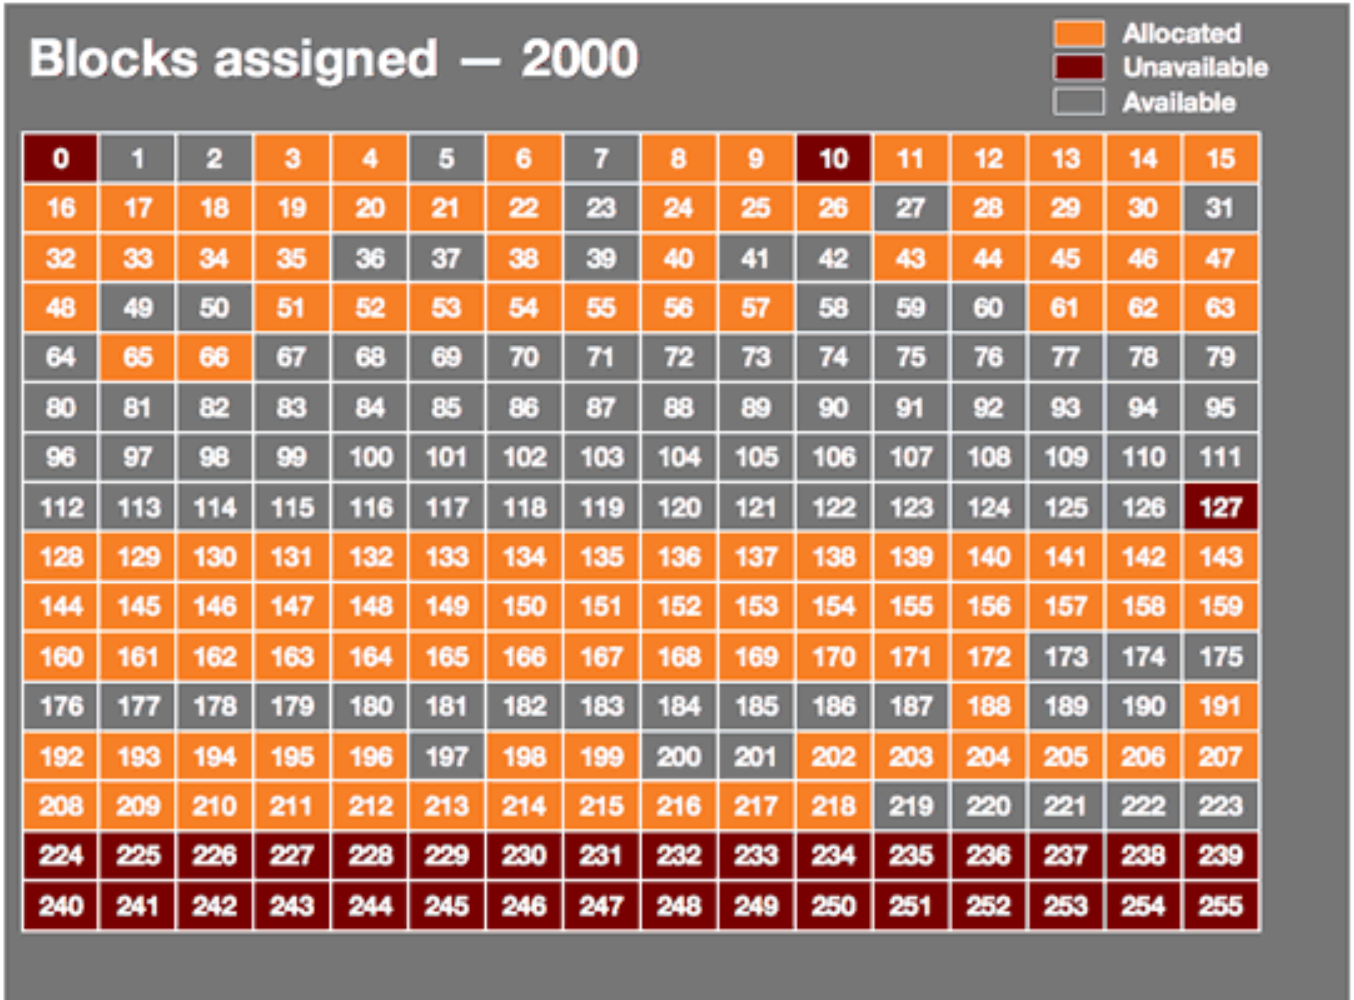

Figure from http://arstechnica.com/articles/paedia/IPv6.ars

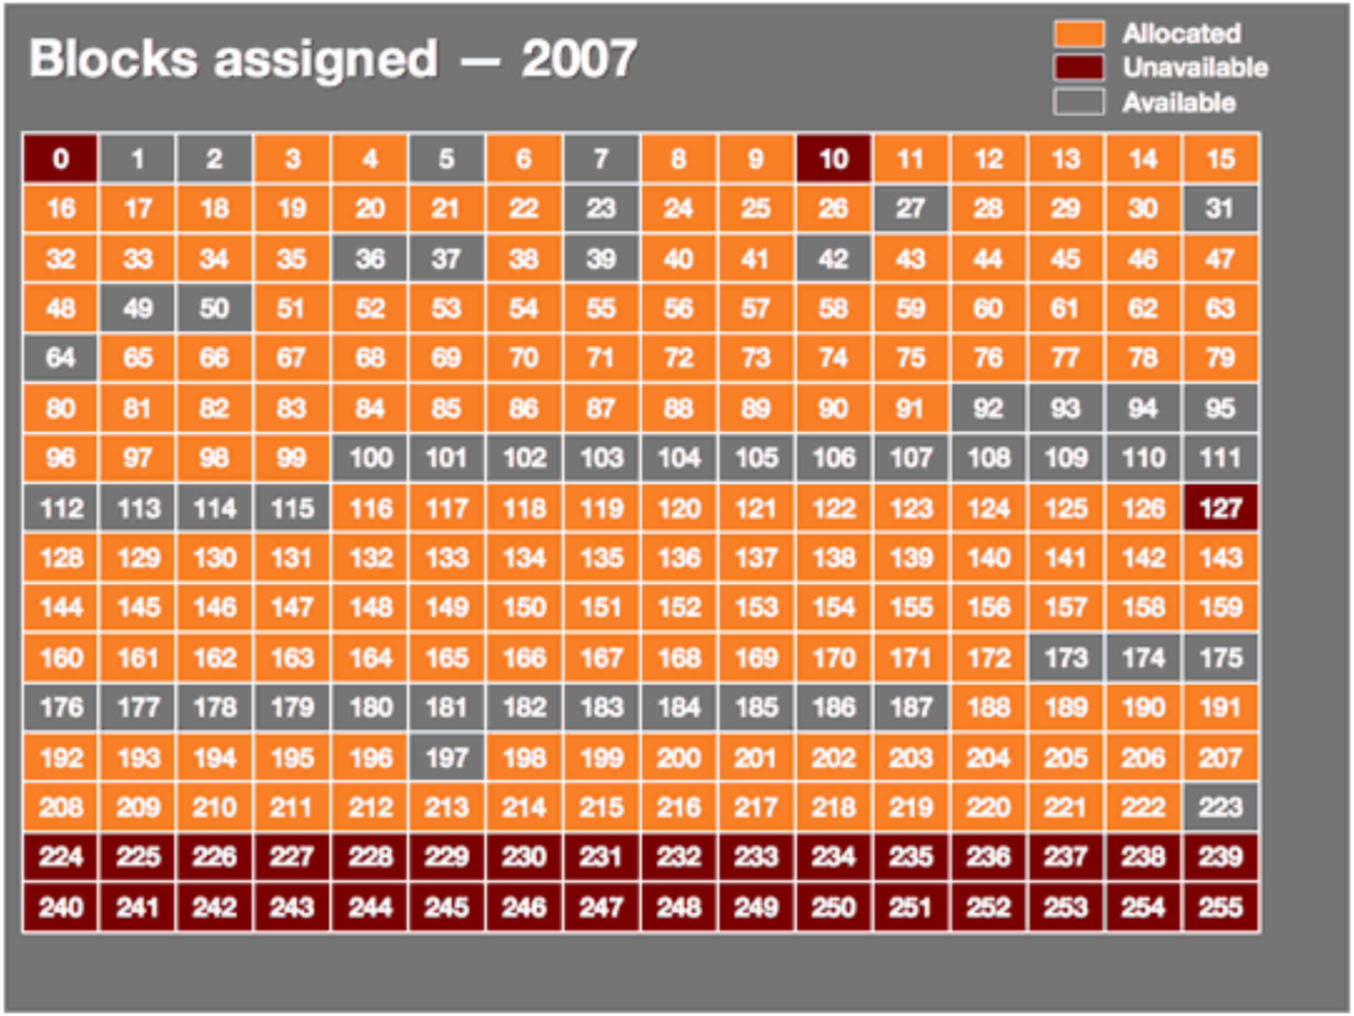

Figure from http://arstechnica.com/articles/paedia/IPv6.ars

IPv4 address space as of January 2010

Used Free
Unusable

http://arstechnica.com/tech-policy/news/2010/01/90-of-ipv4-address-space-used-ipv6-move-looking-messy.ars

IPv4 address space as of October 18, 2010

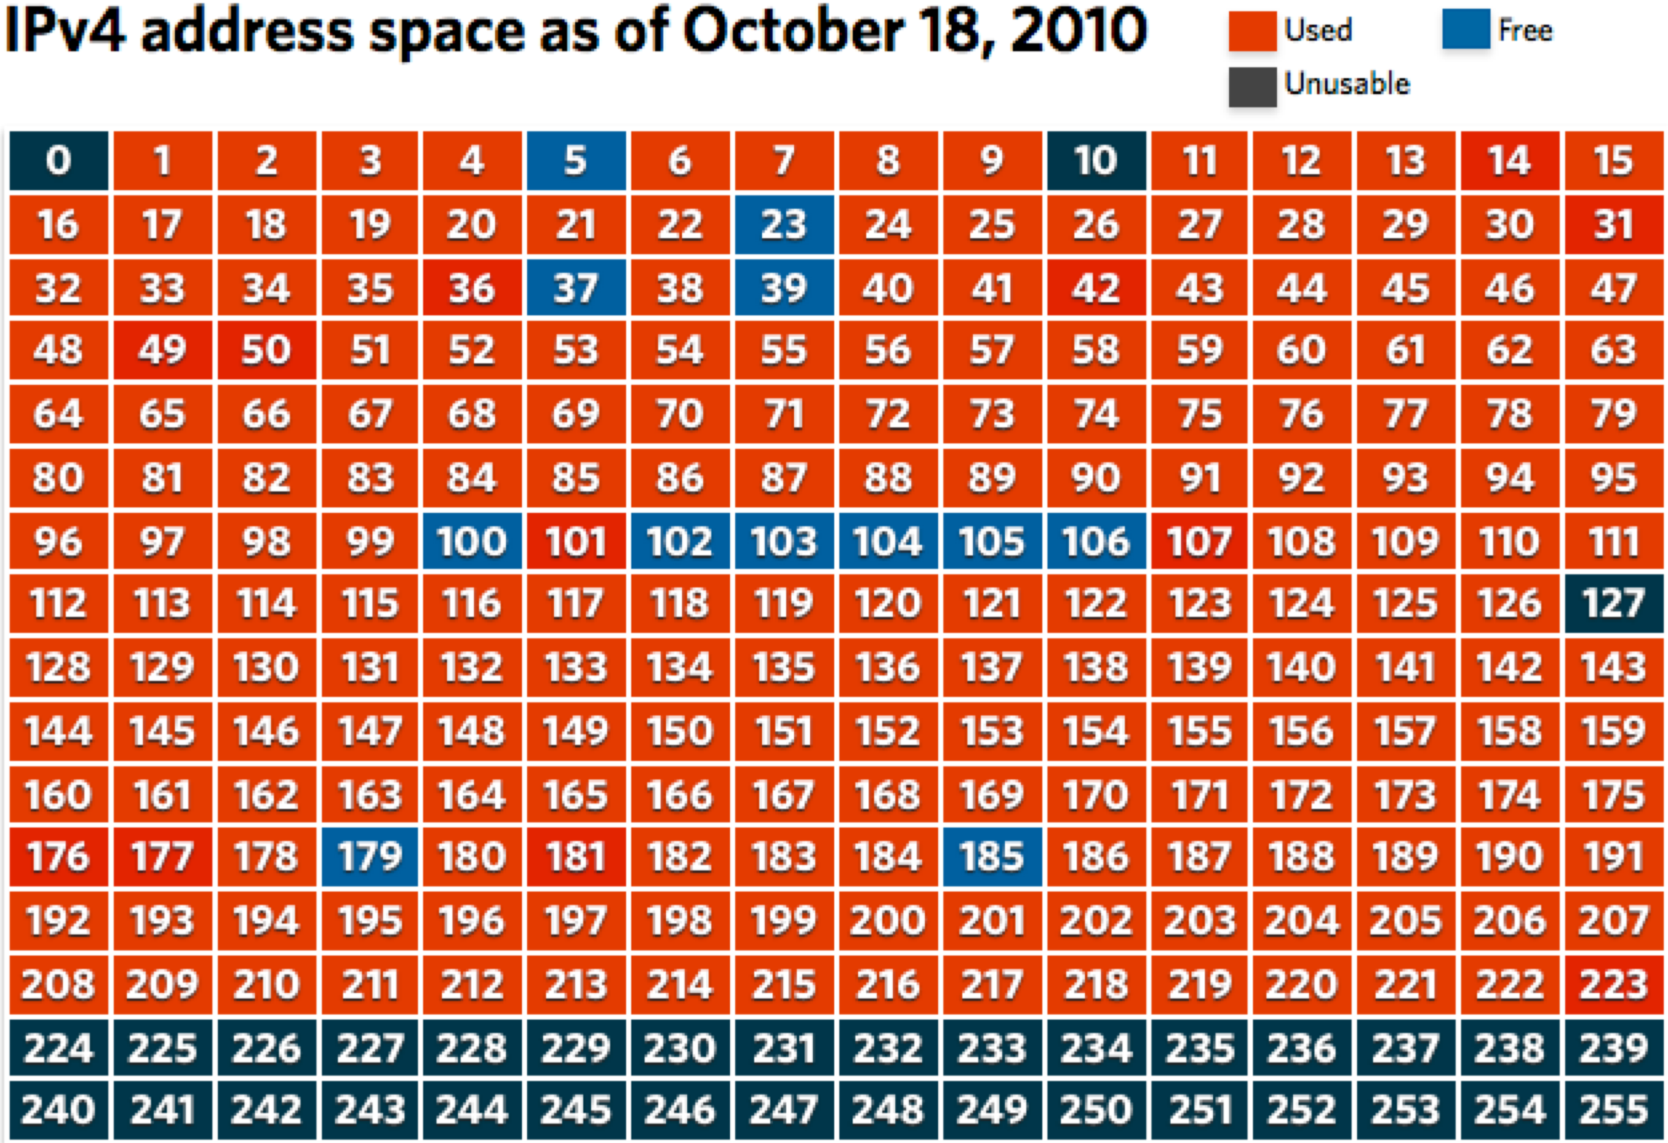

IPv4 address space as of February 2011

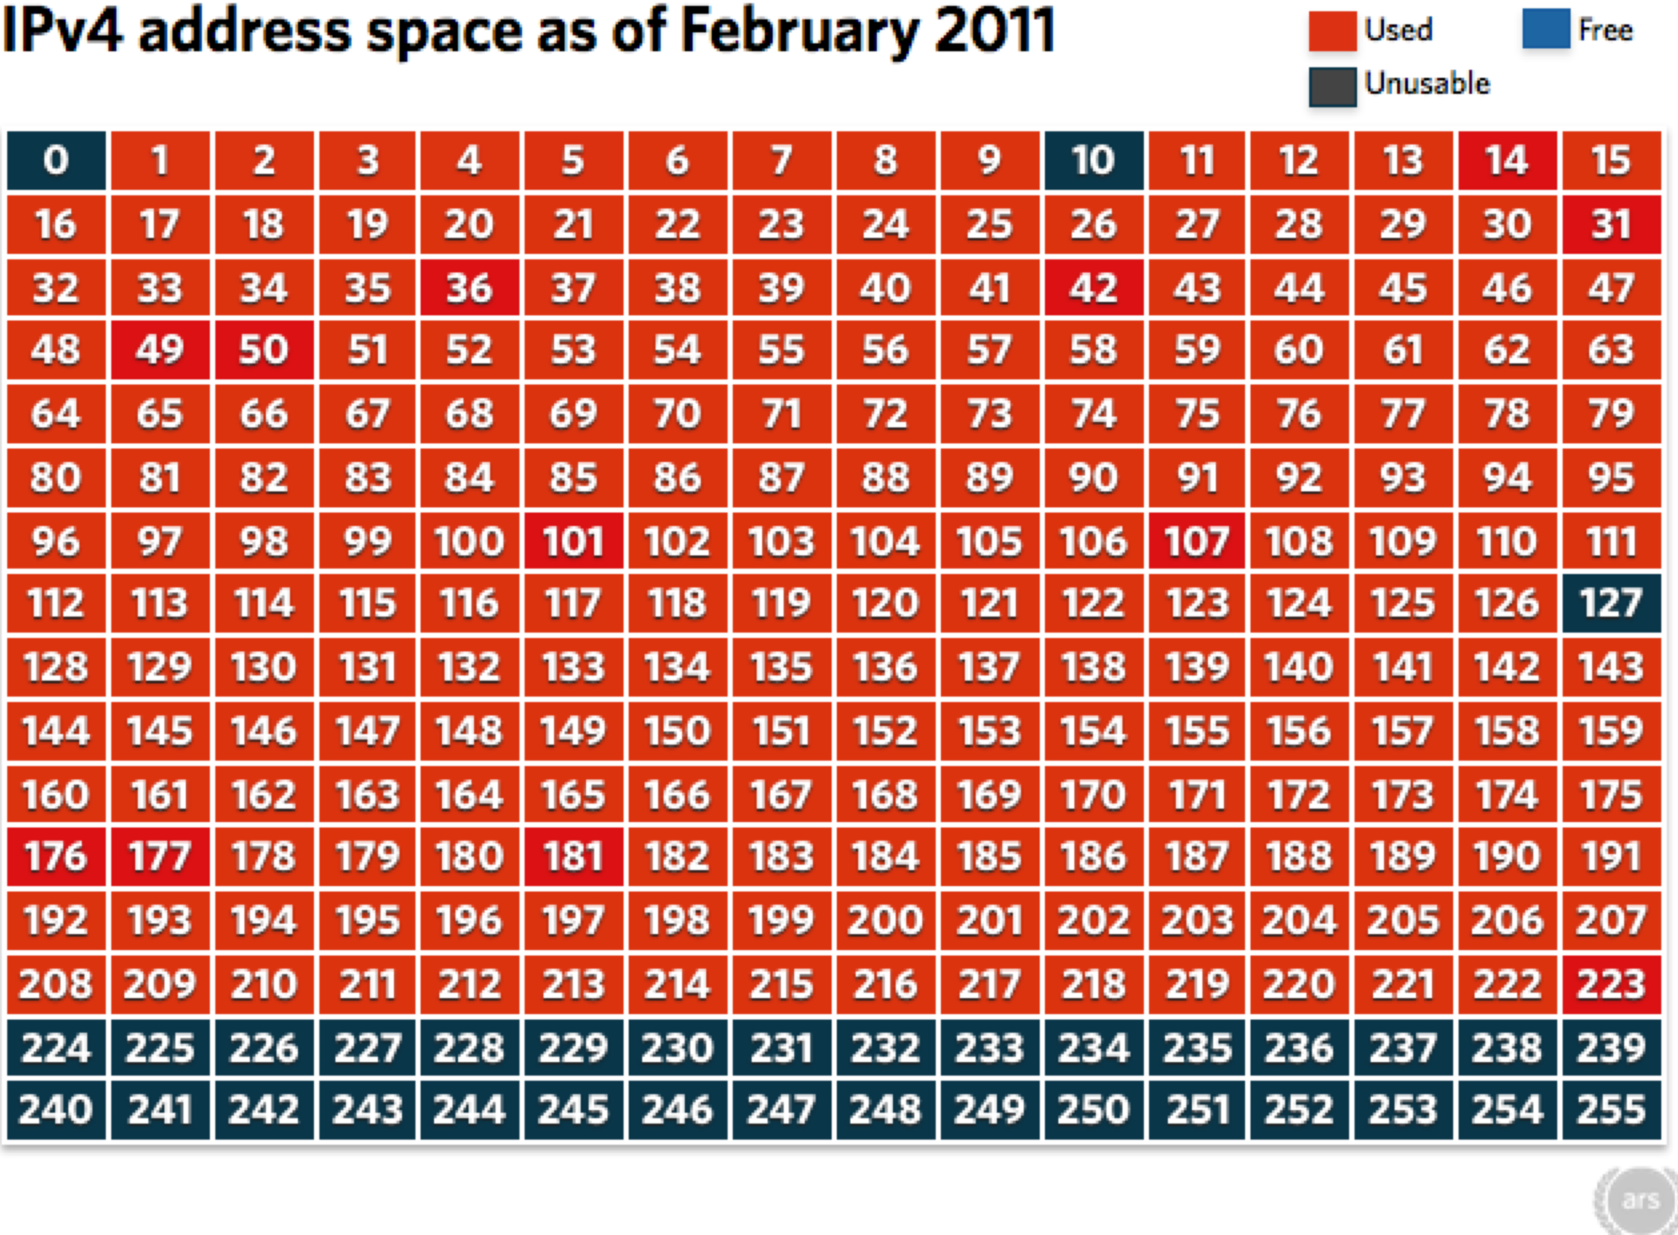

http://arstechnica.com/tech-policy/news/2011/02/river-of-ipv4-addresses-officially-runs-dry.ars

IPv4 Address Space

- Unavailable Addresses
  - 10.x – Private Addresses
    - Along with 192.168.x and 172.16.x to 172.31.x
  - 127.x – Local Loopback Addresses
    - Why an entire /8?
  - 224.x to 239.x — Multicast groups
  - 240.x to 254.x — Reserved for “future use”
    - Waste of address space
    - Impossible to re-use today because most IP software flags these addresses as invalid
  - 91 entities with entire class A’s (Govt, IBM, GE, HP, MIT, …)
- Current Allocation
  - http://www.iana.org/assignments/ipv4-address-space

Comparison – IPv4 vs IPv6

	IPv4	IPv6
Deployed	1981 *[RFC 791]*	1999 *[RFC ~~2460~~, 8200]*
Address Size	32-bit number	128-bit number
Address Format	Dotted Decimal Notation: `192.149.252.76`	Hexadecimal Notation: `3FFE:F200:0234:AB00:0123:4567:8901:ABCD`
Prefix Notation	`192.149.0.0/24`	`3FFE:F200:0234::/48`
Number of Addresses	2^{32} = ~4,294,967,296 (~4 billion)	**2^{128} = ~340,282,366, 920,938,463,463,374, 607,431,768,211,456**

https://biotech.law.lsu.edu/blog/ipv4_ipv6.pdf (ARIN Fact Sheet)

IPv6 Address Notation

- 128 bits – 8 groups of 4 hex digits
  - `2001:0db8:85a3:08d3:1319:8a2e:0370:7334`
- "User friendly!" "Easy to remember!"
- "Helpful" Shortcuts:
  - Omit leading zeros in a group (`0005:0db8:`… is equivalent to `5:db8:`…)
  - Collapse groups of all-zeros with :: (`2001:0000:0000:0000:0000:8a2e:0370:7334` is equivalent to `2001::8a2e:0370:7334`)

But we couldn't just stop with a larger address space....

IPv4 vs IPv6 - Differences

IPv6 is *not* just IPv4 with 128-bit long addresses…

It's a different network protocol that should be configured (and secured) separately, but runs over the same physical wires.

"Dual Stack"

IPv4 vs IPv6 - Similarities

- Datagram
  - Each packet is individually routed
  - Packets may be fragmented or duplicated
- Connectionless
  - No guarantee of delivery in sequence
- Unreliable
  - No guarantee of delivery
  - No guarantee of integrity of data
- Best effort
  - Only drop packets when necessary
  - No time guarantee for delivery

IPv4 vs IPv6 – Address Length

- Address Length
  - IPv4 – 32 bits (2^{32} = ~4 billion)
  - IPv6 – 128 bits (2^{128}= ~340 trillion, trillion, trillion)
- Standard subnet size in IPv6: 2^{64}
  - Upper 64 bits: Subnet address
  - Lower 64 bits: Devices within subnet
- **With such a large address space, no need to use all possible addresses**

IPv4 vs IPv6 - Autoconfiguration

- How does a new host determine its IP address?
  - IPv4: DHCP (Dynamic Host Configuration Protocol)
  - IPv6: Stateless Address Autoconfiguration (SLAAC)
- Stateless IP address autoconfiguration via Neighbor Discovery (part of ICMPv6)
  - Host sends out multicast *Router Solicitation*
  - Router(s) reply with a *Router Advertisement* message
    - Includes router subnet and gateway information
  - Host constructs its own IP address(es)
    - Upper 64 bits – From router(s), globally unique
    - Lower 64 bits – Client generates based on MAC address (or picks *randomly* for SLAAC Privacy Extensions)
  - Host verifies uniqueness via *Neighbor Solicitation*
- DHCPv6 is <u>also</u> a thing! (Alternate method, helps with DNS)

IPv6 – Special Addresses

- Loopback: `::1`
- Link Local
  - Scope limited to single network segment / link
  - Application: Network configuration, device discovery
  - Address matches `fe80::/10`
- Site local
  - Scope limited to single organization (similar to private IPv4 addresses)
  - Purpose: Each organization can randomly pick their own address instead of everyone using same range of private IPv4 addresses
  - Address matches `fc00::/7`
- Multicast
  - Address matches `ff00::/8`

IPv4 vs IPv6 – Fragmentation

- **IPv6 Fragmentation only done by transmitting host**
- Supported by an optional header
  - Design assumption that fragmentation will be less common in the future
- Routers never fragment a packet
  - Drop packets that are too large
  - Send ICMP error back to host
  - **Simplifies router design**
- Host should use Path MTU Discovery (PMTUD) to select correct (maximum) packet size

IPv4 vs IPv6 – Router Overhead

- Simplified packet processing for routers
- Simplified Header Format
  - Infrequently used fields are moved to optional header extensions
- No Header Checksum in IPv6
  - Easier for routers – No need to update checksum after decrementing TTL
  - Reliability maintained by link-level (Ethernet) and transport-layer (TCP, UDP) error checking

IPv4 Datagram

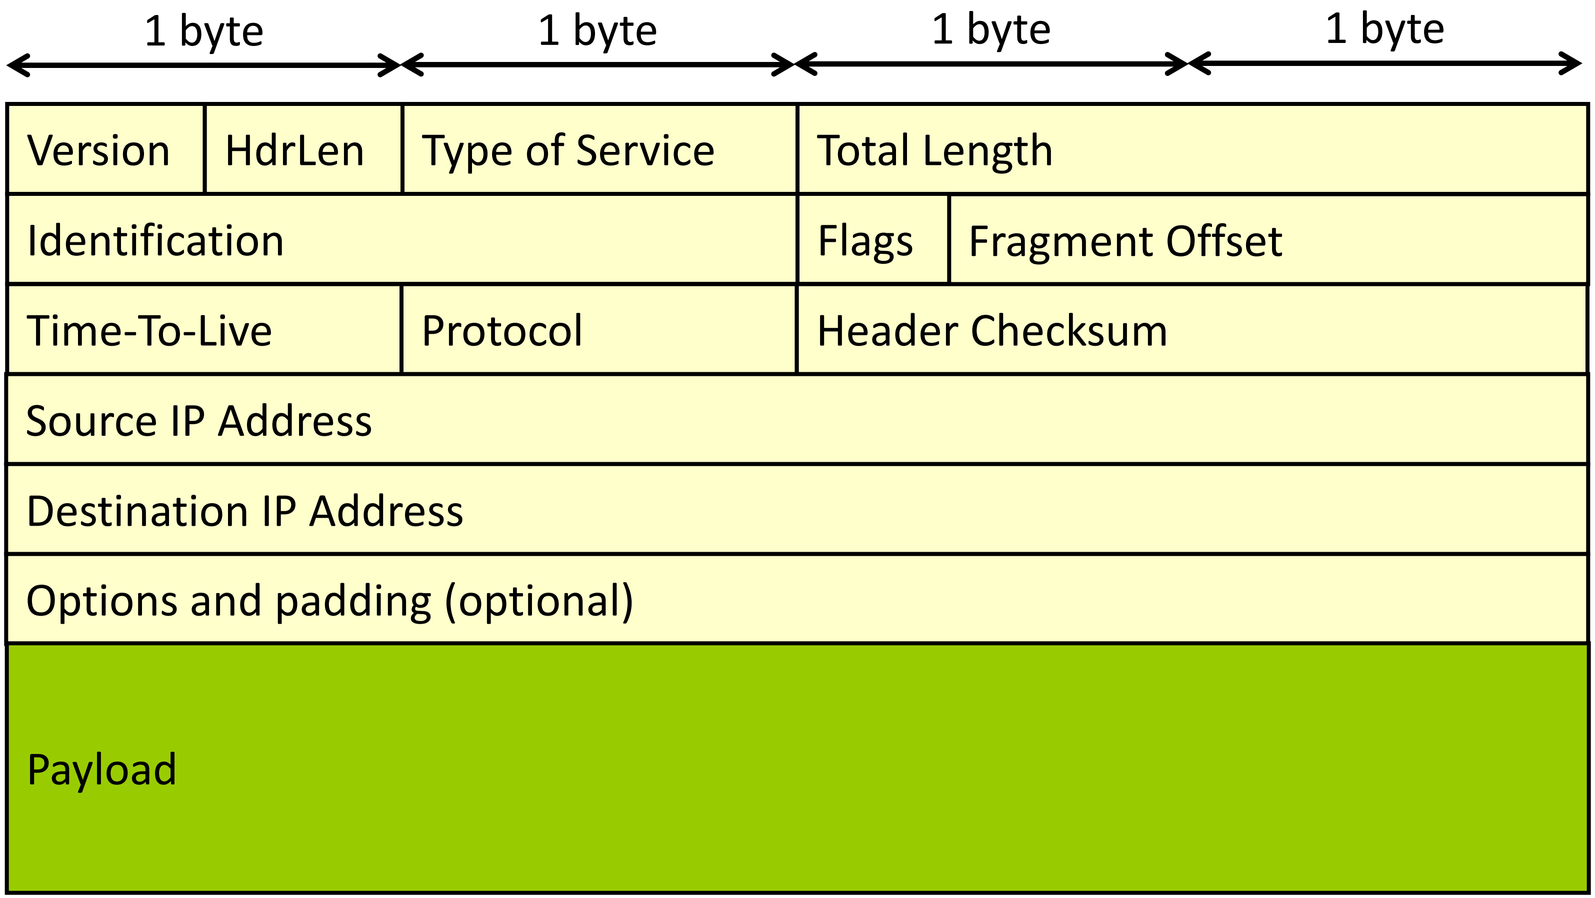

IPv6 Datagram (Base Header)

<table>
<tr><th colspan="2">1 byte</th><th colspan="2">1 byte</th><th colspan="2">1 byte</th><th colspan="2">1 byte</th></tr>
<tr><td>Version</td><td colspan="2">Traffic Class</td><td colspan="5">Flow Label</td></tr>
<tr><td colspan="4">Payload Length</td><td colspan="2">Next Header</td><td colspan="2">Hop Limit</td></tr>
<tr><td colspan="8">Source IP Address (128 bits)</td></tr>
<tr><td colspan="8">Destination IP Address (128 bits)</td></tr>
<tr><td colspan="8">Payload</td></tr>
</table>

IPv6 Datagram (Base Header)

- Version (4 bits)
- Traffic Class (8 bits)
  - Differentiated Services (DS) field
  - Explicit Congestion Notification field
  - *Can* be used by routers to prioritize traffic or decide what to drop during congestion
- Flow Label (20 bits)
  - Identifies stream of packets
  - *Can* be used by routers to avoid sending a single flow across multiple outbound paths (which could result in re-ordering at arrival)

IPv6 Datagram with Options

- Can append multiple headers

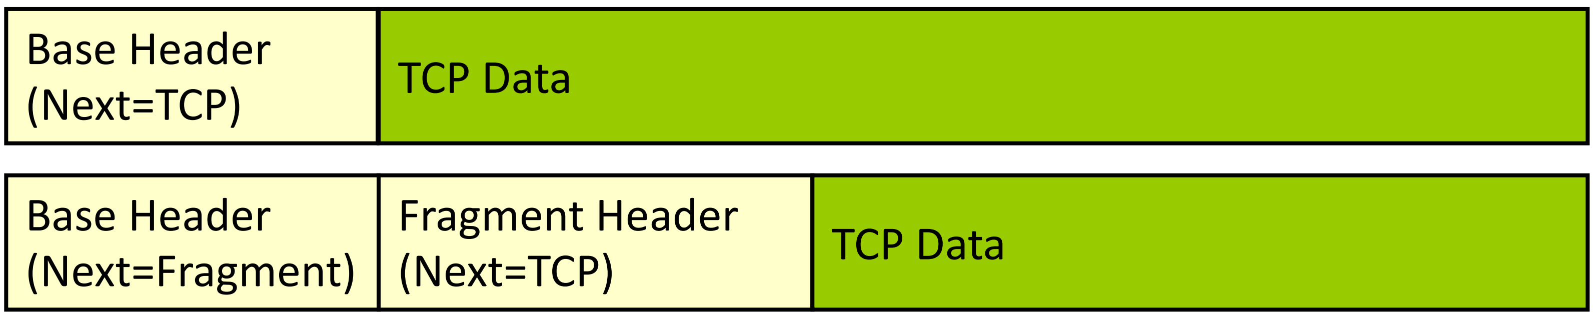

- Common additional headers
  - Fragmentation
  - Routing (allows source to specify preferred route)
  - Authentication Header (part of IPsec – verifies source)
  - Encapsulating Security Payload (part of IPsec – carries encrypted payload)

IPv6 – Routing

- **How can having bigger IP addresses (128 bits) make routing easier?**
  - Larger address space allows more intelligent network organization
    - Addresses match physical network organization
    - Collapse routing table entries
- Many competing proposals on the details…
- Basic constraint on all designs
  - Use upper 64 bits for routing
  - Use lower 64 bits for interface ID (clients pick this randomly or based on MAC address)

IPv6 – Routing

- **Besides the address layout, how does IPv6 make routing easier?**
  - No checksum calculation
  - No fragmentation
  - Infrequently used headers are optional
- **How does IPv6 make routing harder?**
  - Forwarding table entries 2-4 times larger
  - Need to route both IPv4 and IPv6 for the foreseeable future

IPv6 – Security

- What are the security implications of having a huge (sparse) address space?
  - Security through obscurity(?)
  - Blind random address scanning by worms is ineffective
    - Unlike in IPv4, which can be scanned in **5 minutes (!!)** over a 10GbE link: https://zmap.io/
  - Targeted scanning works great, however…
    - Listen to P2P networks?
    - Listen to internal routing protocols? (OSPF, etc…)
    - Use Neighbor Discovery on infected host?
    - Snoop through host configuration and log files on infected host?
    - https://www.usenix.org/system/files/login/articles/920/bellovin.pdf

IPv6 – Security (IPsec)

- Security – IPSec support ~~required~~ optional in IPv6
  - IPSec encrypts each IP packet independently
  - Was originally required but dropped because not all devices (e.g. embedded) could support it
- IPsec features
  - Data encryption – Data cannot be read or modified
  - Host authentication
  - Anti-replay – Captured packets cannot be reused by an attacker
- What are the strengths and weaknesses of putting security at the IP layer? (Doesn't SSL work fine?)
  - Security is independent of higher layers (either applications or protocols like TCP/UDP)
  - Encryption overhead is incurred per-packet (high!)

Deployment

- Why should I deploy IPv6 today?
  - My customers can reach anywhere on the Internet today
  - Google, Facebook, Twitter, etc… will always be reachable
  - Only new applications / users will suffer
- How do I deploy IPv6?
  - Flip a switch across the internet?
- Legacy routers may not be upgradeable
  - Hardware implementations are fixed
  - Software implementations may be insufficiently capable (either incapable or only at low performance)
- Islands of IPv6 in the sea of IPv4
  - Dual network stacks support both IPv4 and IPv6
  - Tunnel IPv6 across IPv4 networks
- Need to upgrade other systems
  - DHCP (SLAAC vs DHCPv6)
  - DNS
    - Starting adding IPv6 addresses to root nameservers in 2008
    - All 13 of 13 root nameservers are IPv6 accessible now
  - Firewalls, traffic shapers, etc.

IPv6 – Original Plan

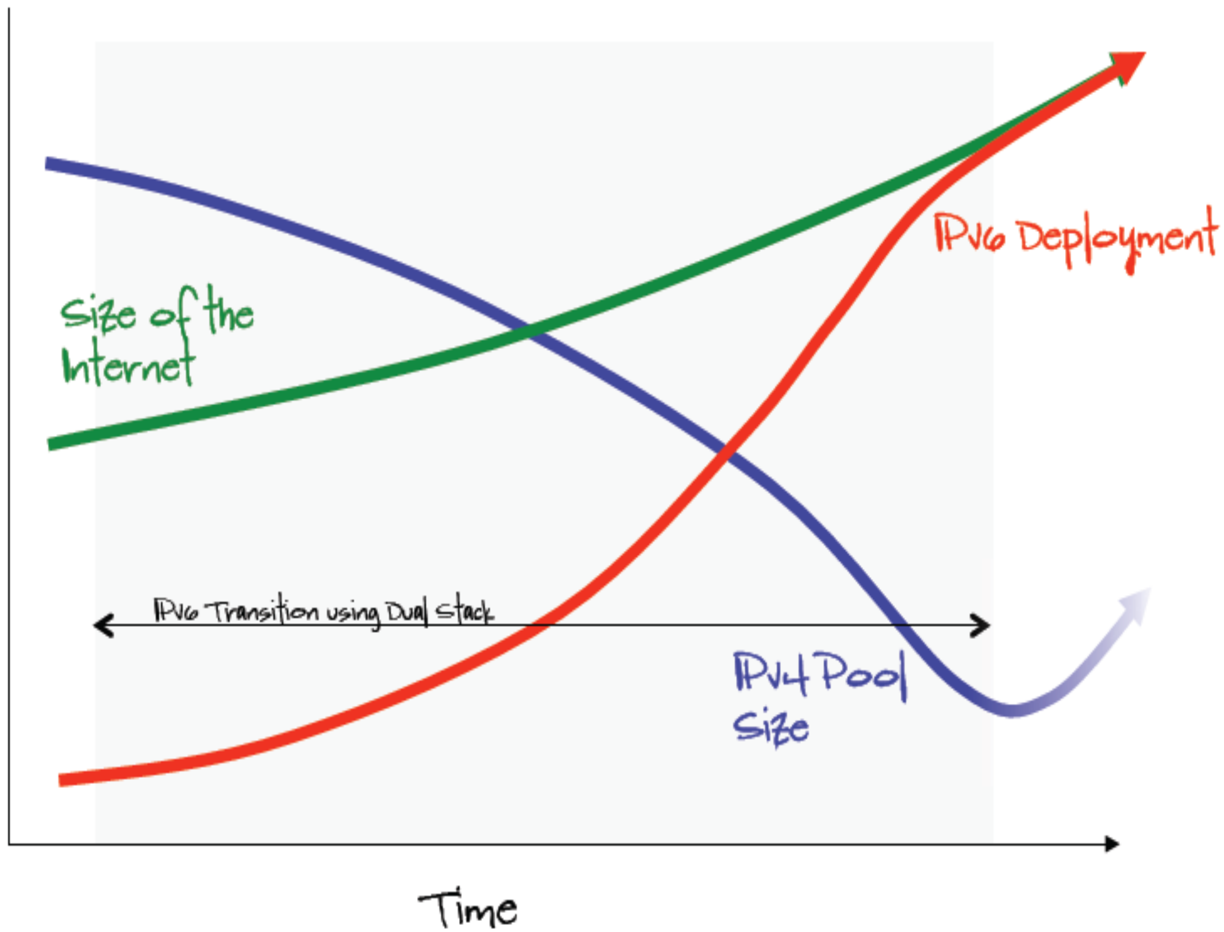

Figure from http://www.potaroo.net/presentations/2008-11-17-ipv6-failure.pdf

IPv6 – Current Status

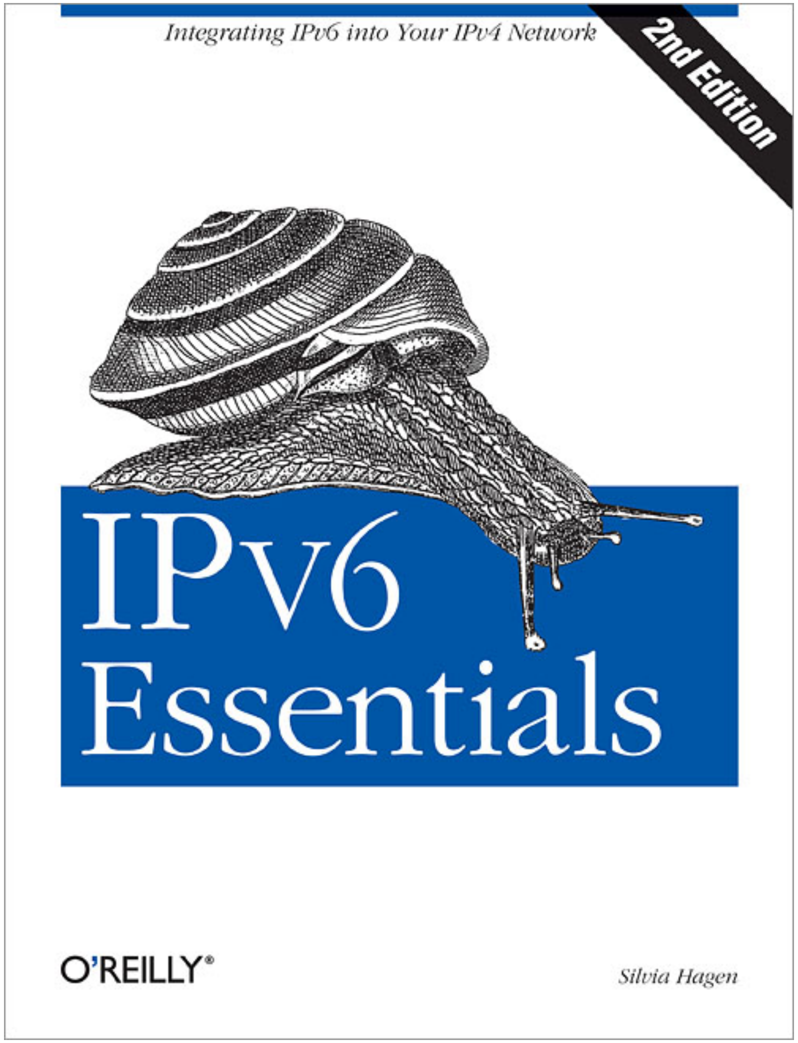

IPv6 – The New "Plan" (?)

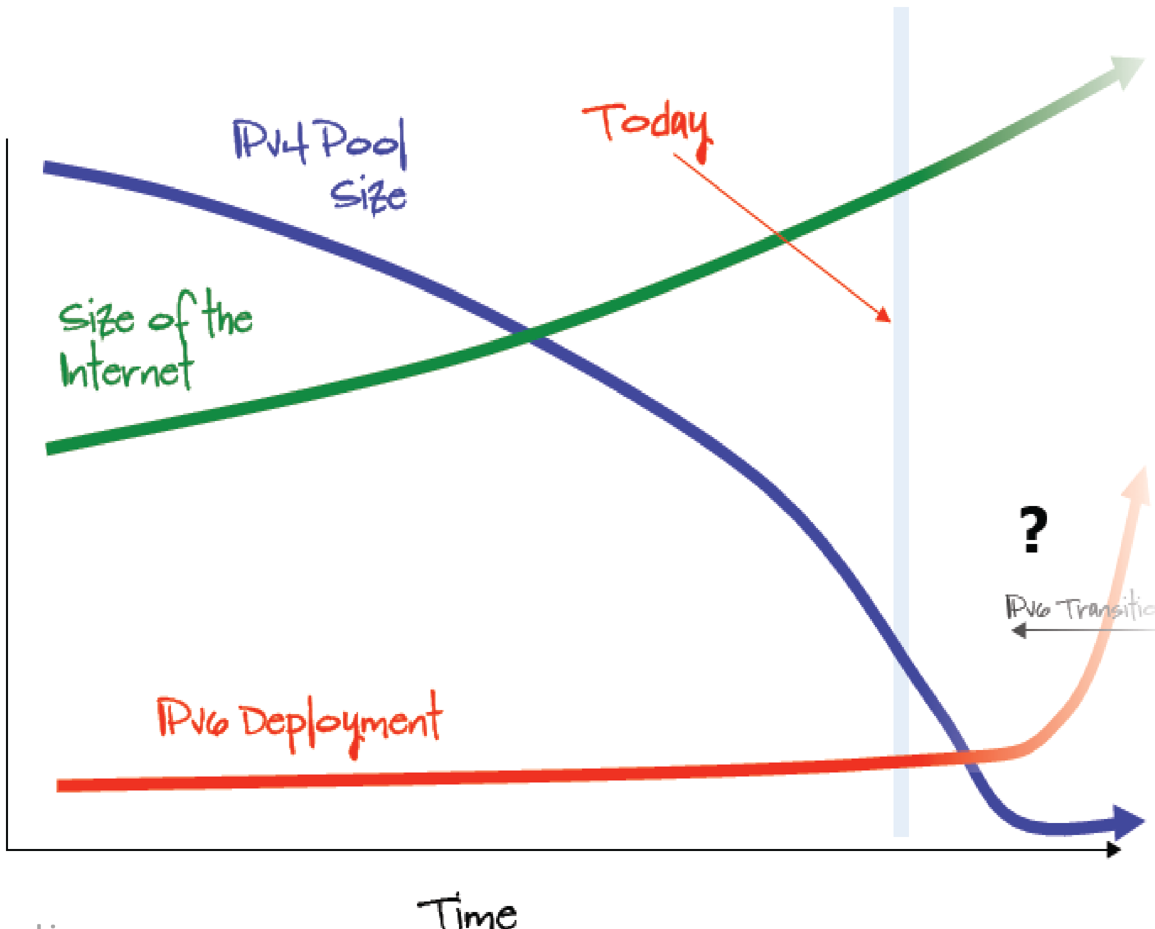

Figure from http://www.potaroo.net/presentations/2008-11-17-ipv6-failure.pdf

IPv6 Adoption | Per-Country IPv6 adoption

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.

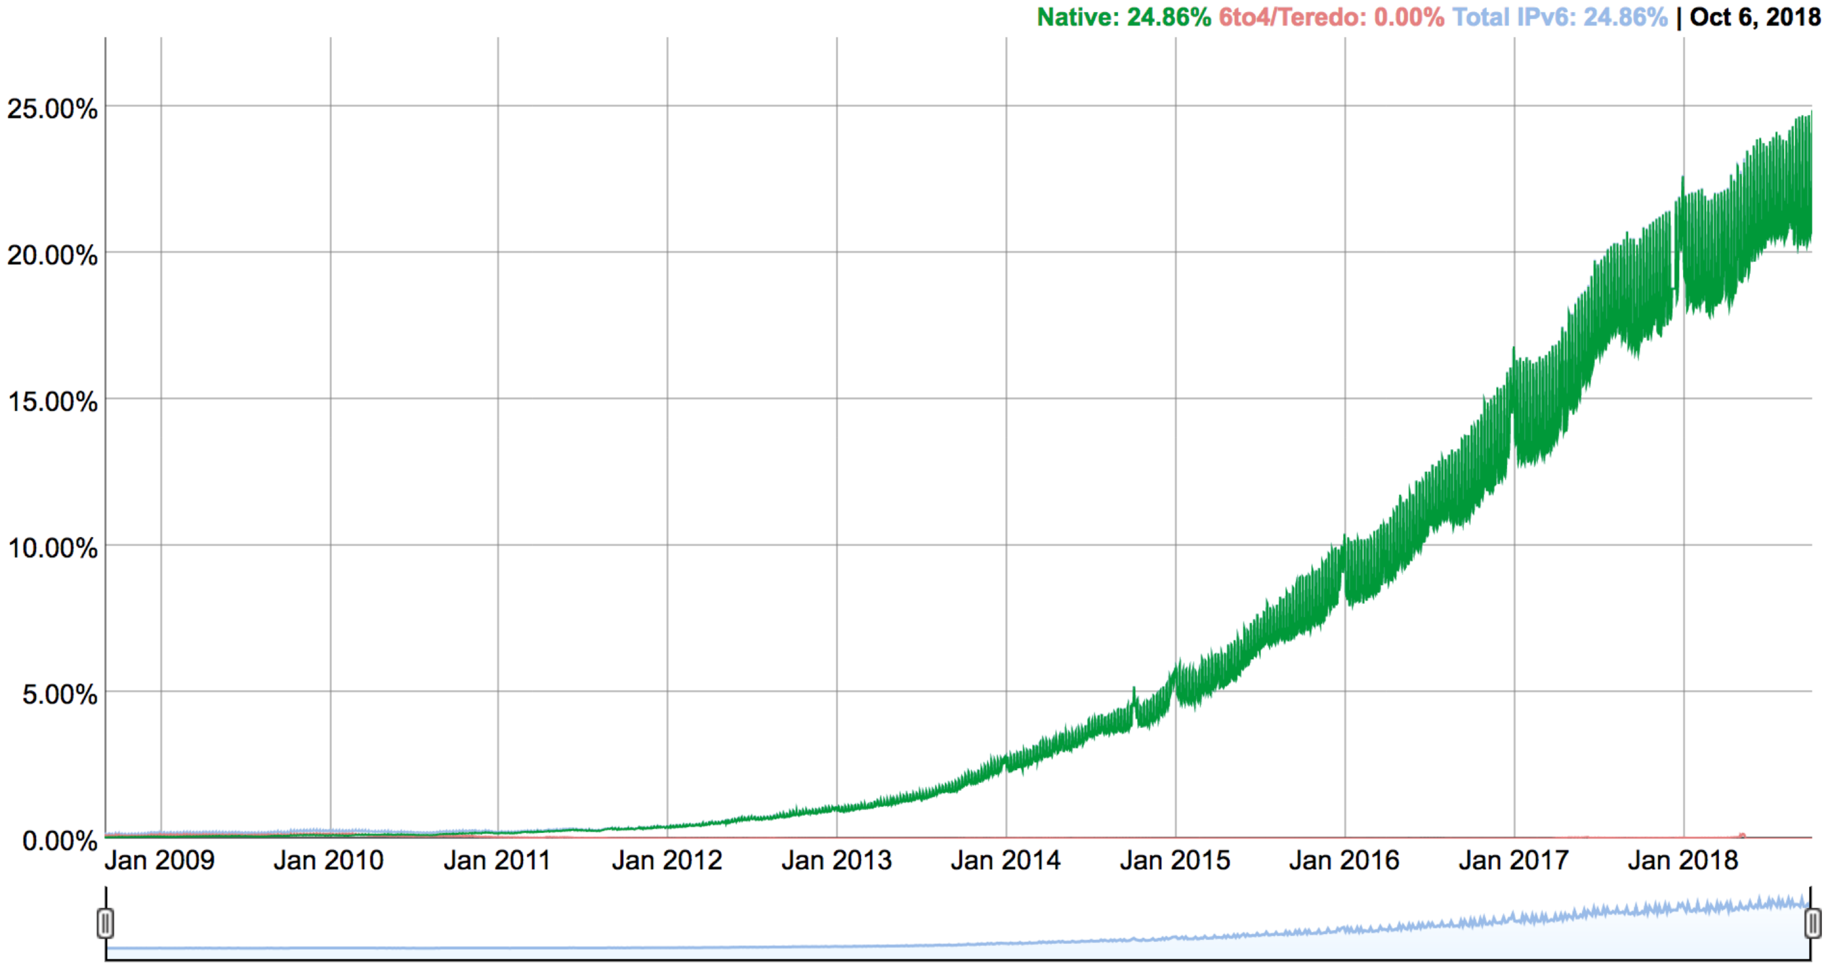

https://www.google.com/intl/en/ipv6/statistics.html

IPv6 – Failure is an Option

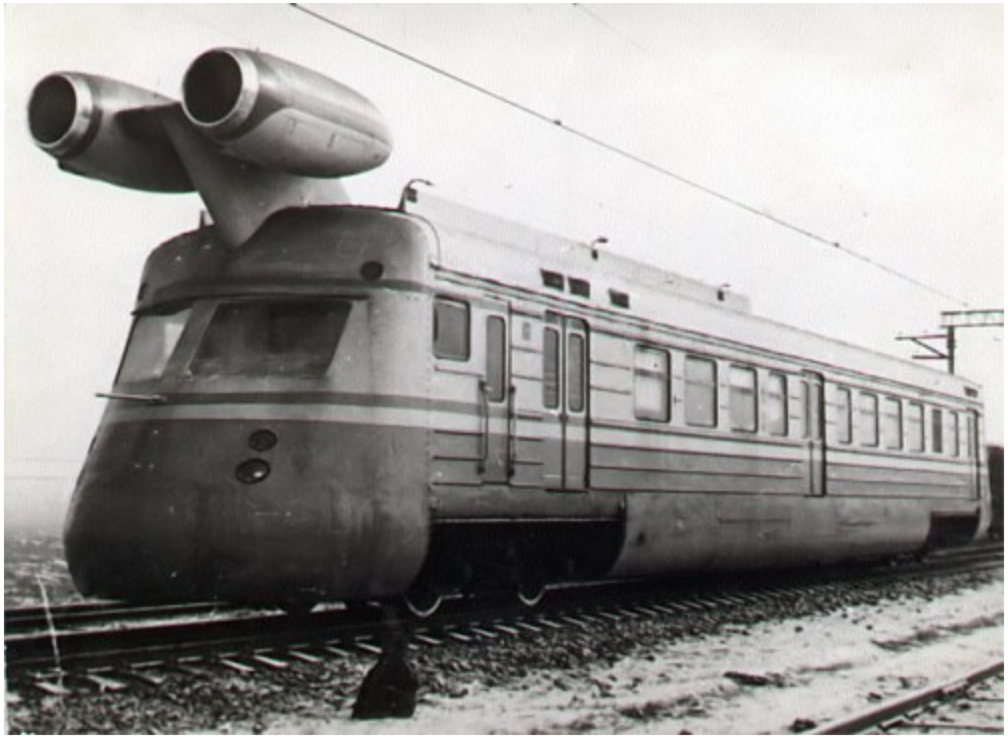

Is this IPv6?

IPv6 – Failure is an Option

- What happens if IPv6 "fails"?
  - Failure is defined as anything less than a complete migration from IPv4 to IPv6
  - Do we stop allowing new hosts to connect to the internet?
- What about using NAT?
  - Observation: Only 5-20% of assigned IPs are actually used by hosts.
  - Solution: Use lots of NAT to reclaim unused addressed
- What happens if this works, and we build "carrier-grade" NAT everywhere?
  - No more end-to-end connectivity?
  - Need coordination with ISP to deploy new services?
  - New opportunities for ISPs to filter traffic and charge for services?

http://www.potaroo.net/presentations/2008-11-17-ipv6-failure.pdf

IPv6 – Failure is an Option

- If an organization deploys NAT extensively, how can you get them to give up the reclaimed addresses?
- IP Address Marketplace
  - Can we create a marketplace? (Currently "forbidden" to sell IP blocks)
  - Imagine: *"For Sale: One Lightly-Used IP Block (only used by grandma to check email on Sunday)"*
  - Same problems as buying a used car:
    - Does the person selling the IP block actually "own" it?
    - What is the condition of the IP block? (If used for spam or illicit activity, IP block may be in blacklists worldwide)

http://www.potaroo.net/presentations/2008-11-17-ipv6-failure.pdf