

IPv6 Fundamentals

Leaving the chains of IPv4 behind

Carsten Strotmann

CREATED: 2025-01-27 MON 08:46

Agenda

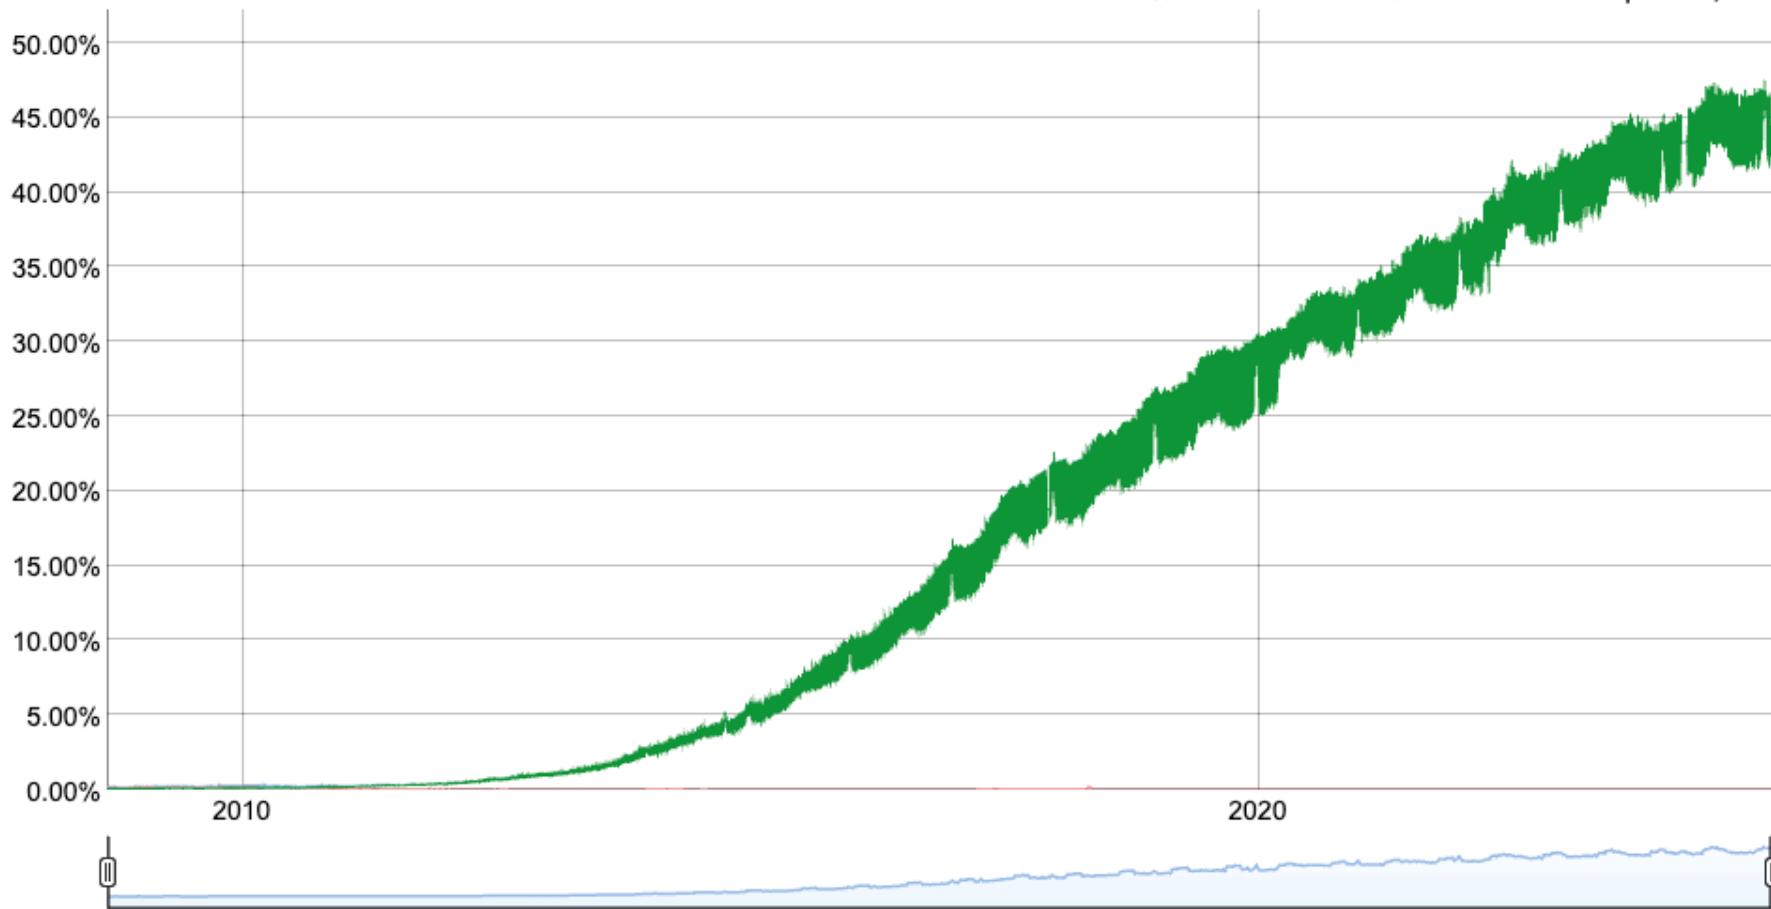
- IPv6 in 2025
- IPv6 challenges or "why don't we have it everywhere?"
- IPv6 Addressing

IPv6 in 2025

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.

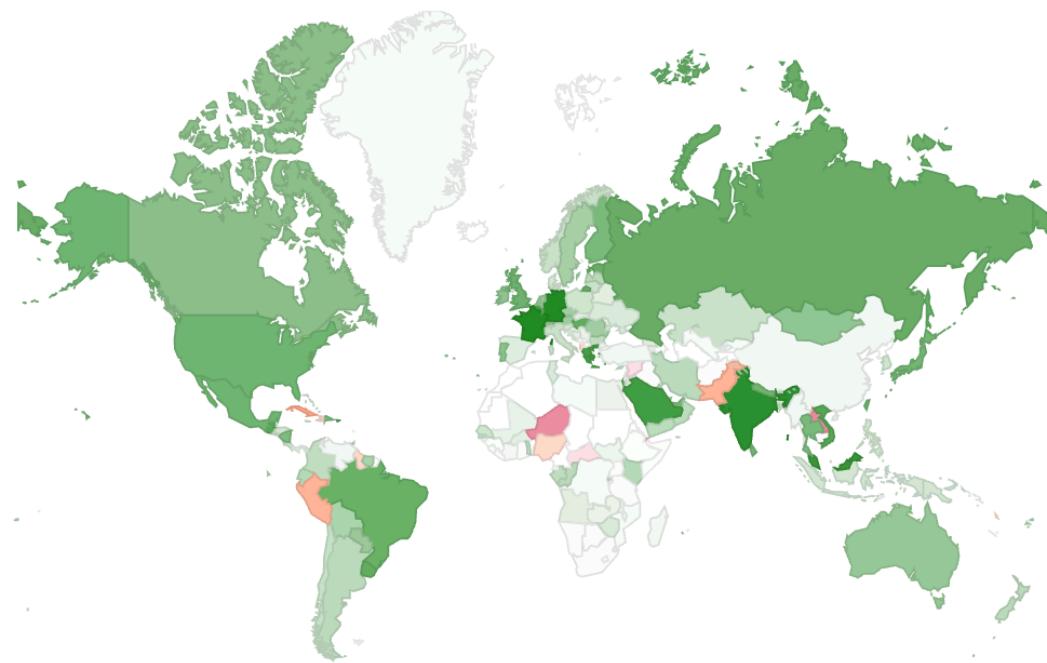
Native: 44.43% 6to4/Teredo: 0.00% Total IPv6: 44.43% | Jan 24, 2025



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

IPv6 by country as seen by Google

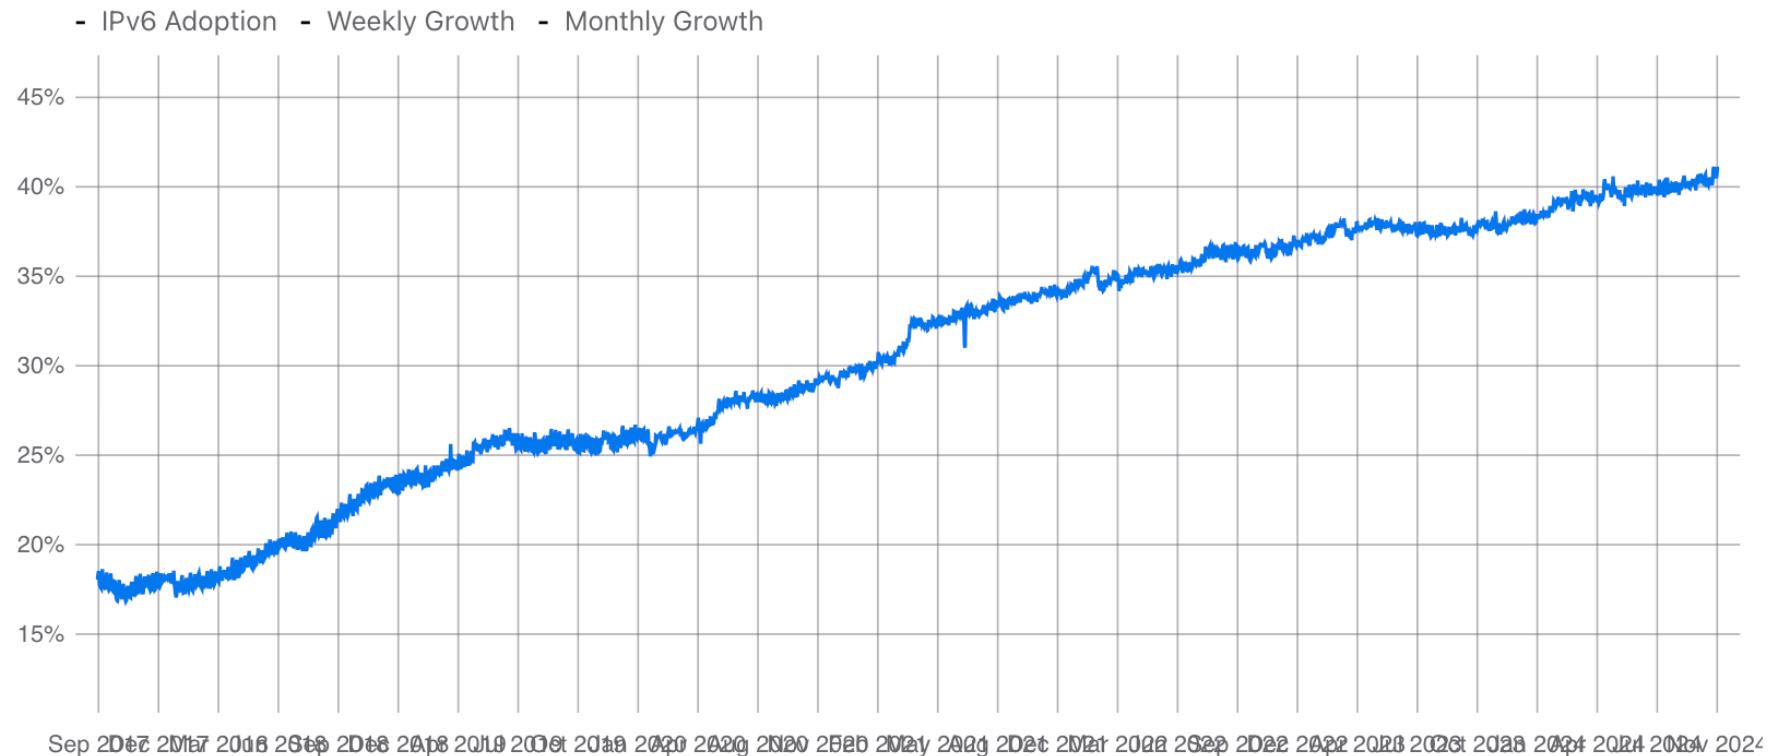
Per-Country IPv6 adoption



IPv6 as seen by Facebook

IPv6 Adoption

Export All ▾



Anatomy of IPv6 projects

IPv6 challenges

- IPv6 is as much a human challenge as it is a technical challenge
 - People are opposed to change
 - IPv6 changes may *break* things in the network

IPv6 challenges

- IPv6 cannot be a "side project"
 - It must have management attention and must be 100% backed by management

IPv6 challenges

- What will happen when introducing IPv6 to a network
 - There will be friction
 - There will be errors and breakage
 - Administrators/developers must have a safe environment to be able to innovate with IPv6

IPv6 Benefits

IPv6 Benefits

- Flexibility
- Modern IETF protocols building on IPv6 (but not IPv4)
- Lean network management and planning

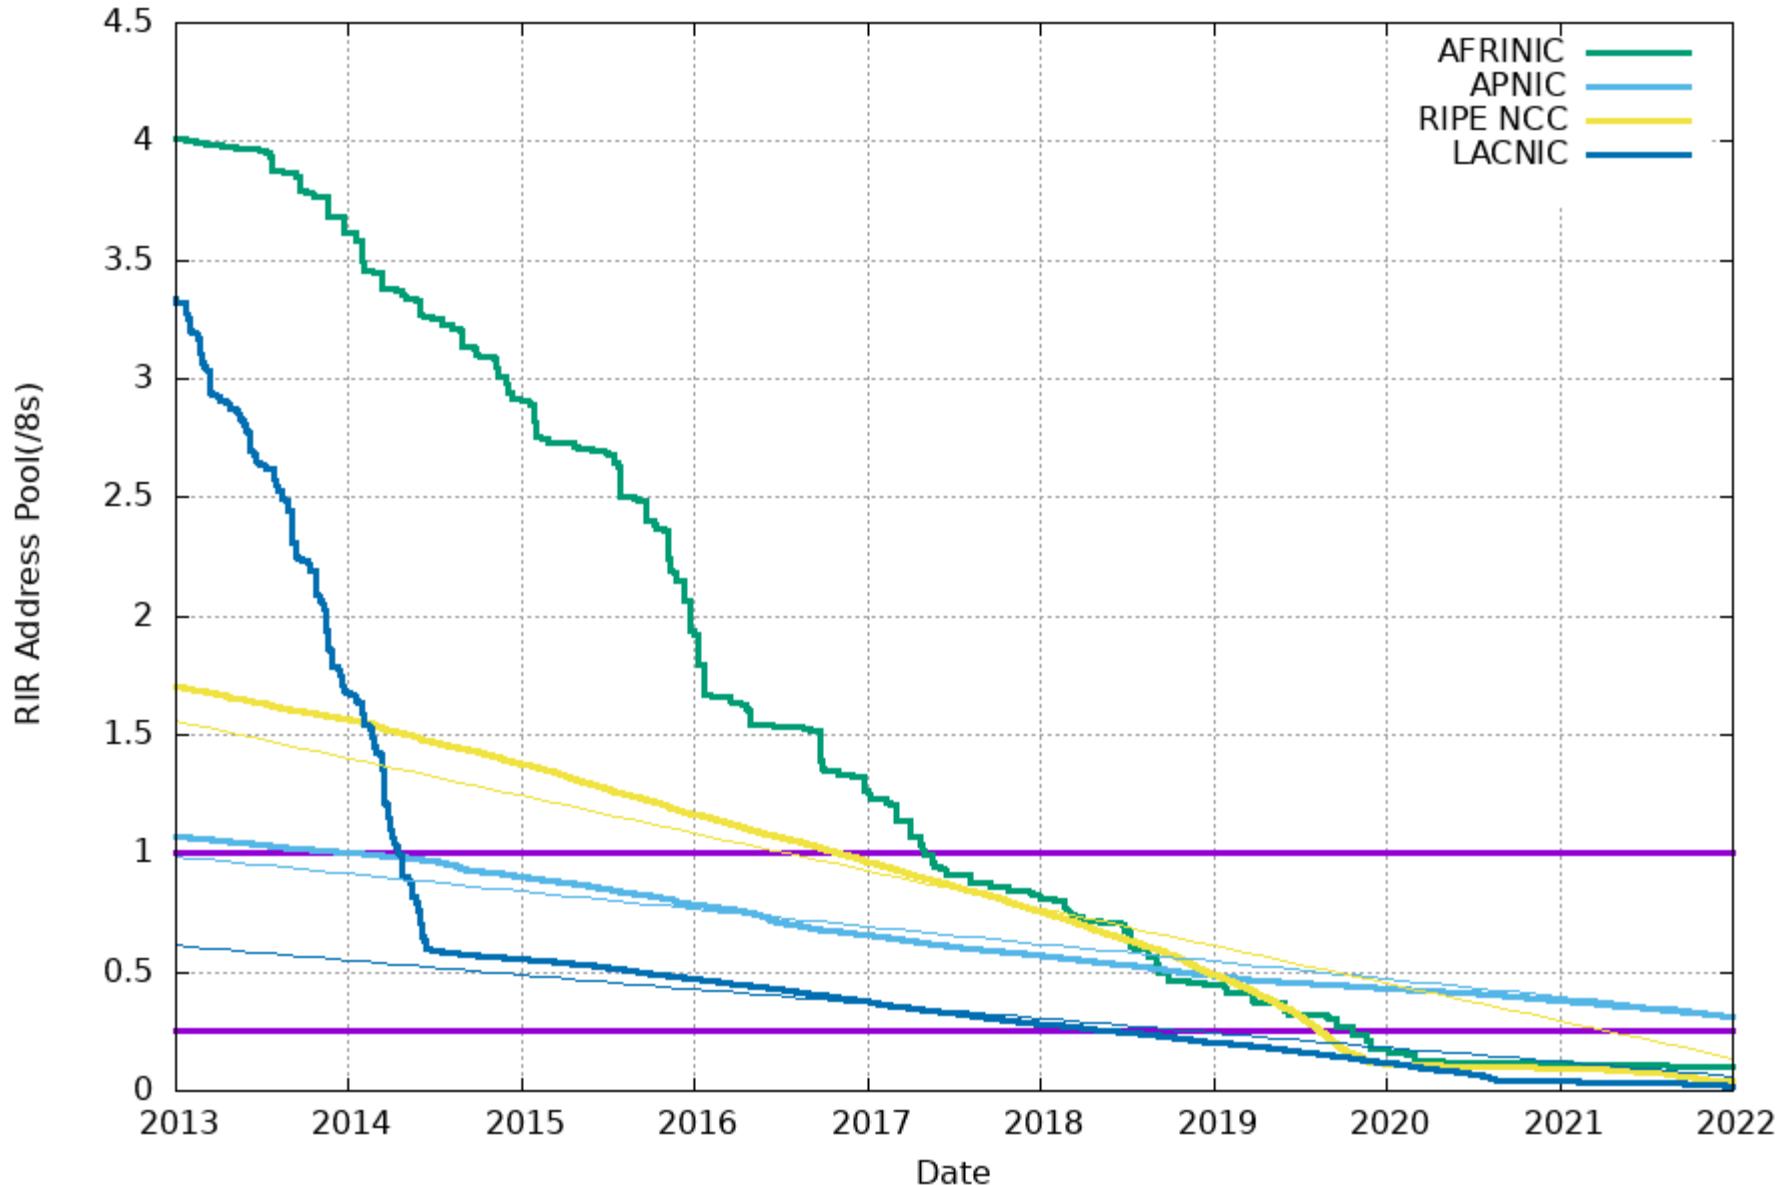
IPv6 Benefits

- Less load on Firewalls (no NAT for UDP based video-conferencing)
- Simpler firewall rules -> less complexity
- Better latency (XBox gaming ...)
- Stability (IPv4-as-a-service is getting brittle)
- Less costs per IP addresses (IPv4 = ~ 35-60 US\$ / IPv6 = 0 US\$)

IPv4 in 2025

- The last IPv4 blocks were given out to the RIRs by IANA on February 1st, 2011
 - ARIN run out of IPv4 addresses in 2015
 - RIPE run out of IPv4 addresses in 2020
 - LACNIC run out of IPv4 addresses in 2020
 - APNIC and AFRINIC are low on the last /8 block

RIR IPv4 Address Run-Down Model



December 2024 IPv4 Auction Sales Report

IPv4.GLOBAL

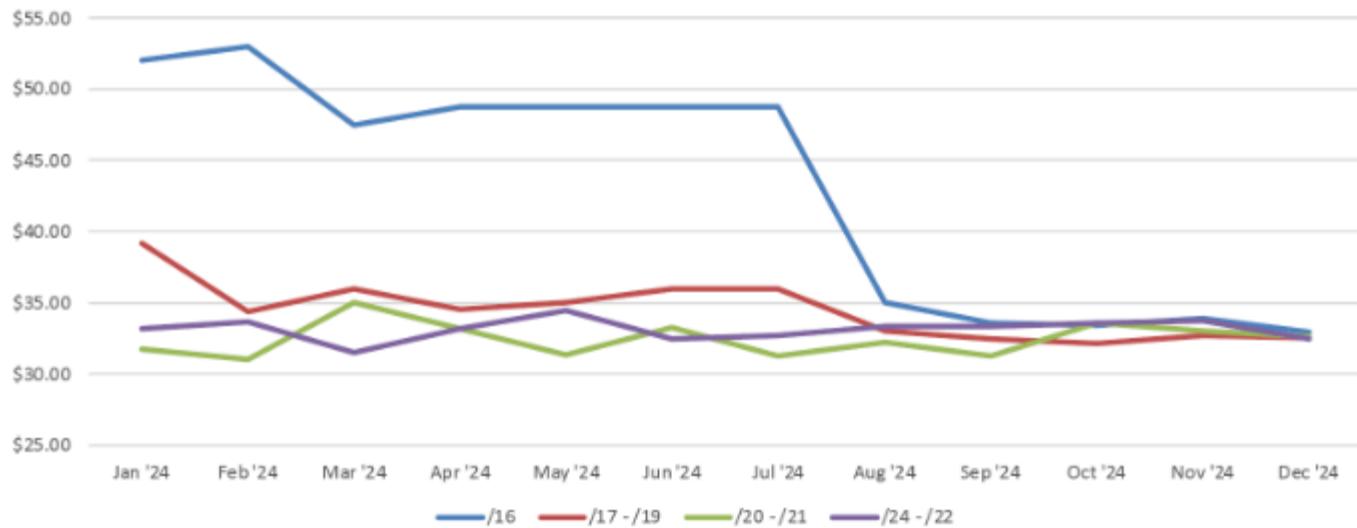
By  **Hilco.**
Streambank.

DECEMBER 2024

IPv4 Sales Data

auctions.ipv4.global/prior-sales

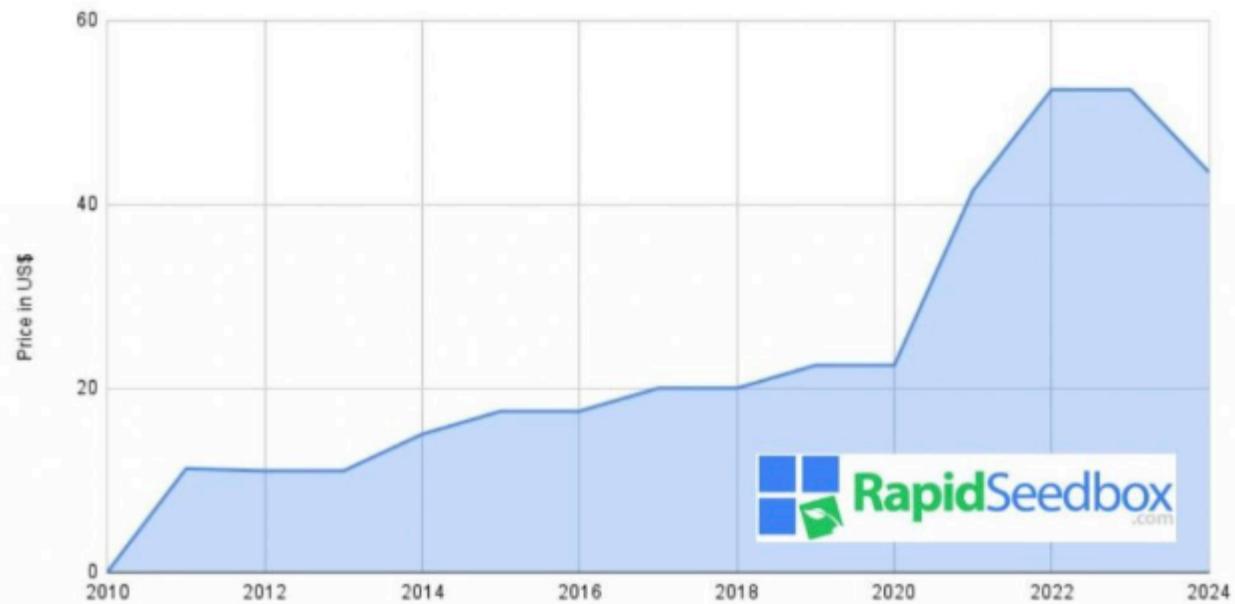
MONTHLY AVERAGE \$ PER ADDRESS



Report includes a mix of online and private sales. Does not include addresses from LACNIC. Some blocks may be part of a bundle with a negotiated price. Prices are reported as of date of agreement, not transfer.

IPv4 address costs

Average Purchase Price Per IP Address



IPv6 trends

- WiFi: One IPv6 /64 subnet per host -> natural user separation on cheap access points
- Datacenter: One IPv6 address per service endpoint:
 - No IP + Port tuple anymore
 - IPv6 overlay networks with security boundaries

IPv6 History

IPv6 History



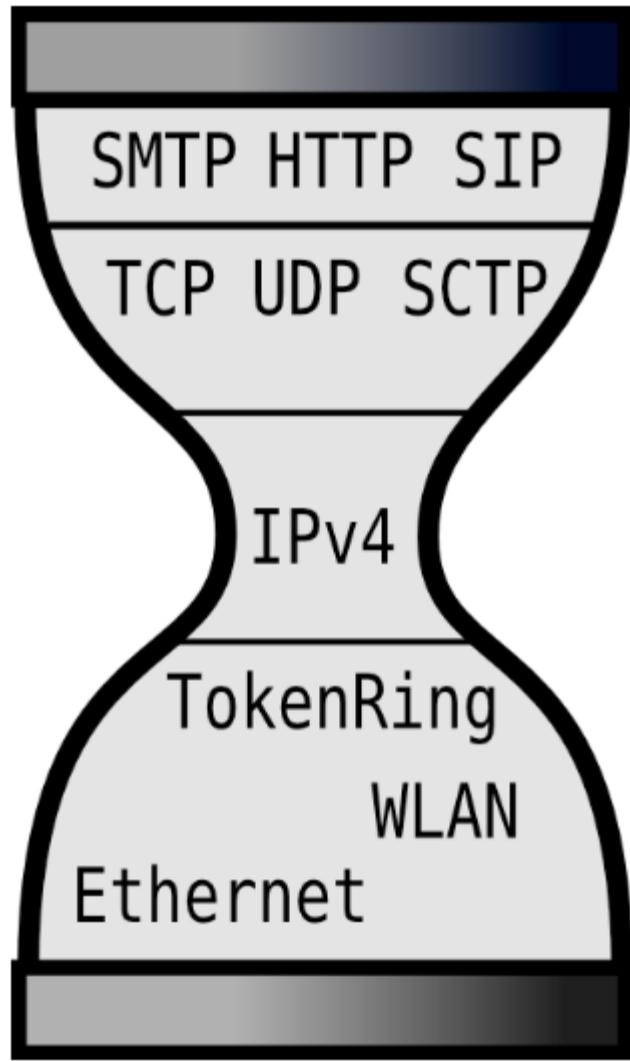
RFC 6540 (April 2012) IPv6 Support Required for All IP-Capable Nodes

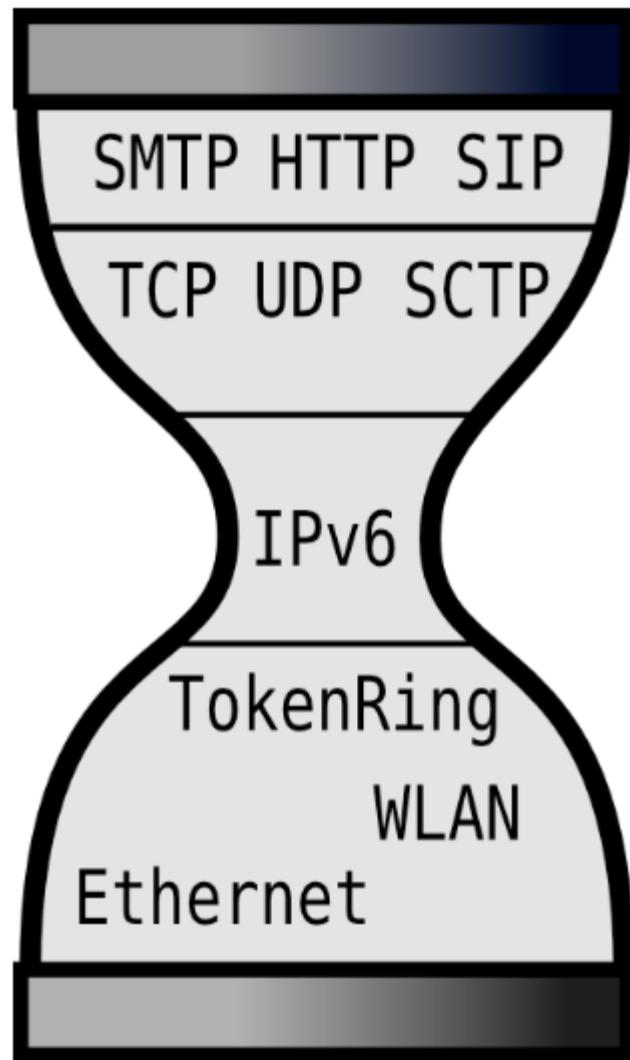
- New IP implementations must support IPv6.
- Updates to current IP implementations should support IPv6.
- IPv6 support must be equivalent or better in quality and functionality when compared to IPv4 support in a new or updated IP implementation.

RFC 6540 (April 2012) IPv6 Support Required for All IP-Capable Nodes

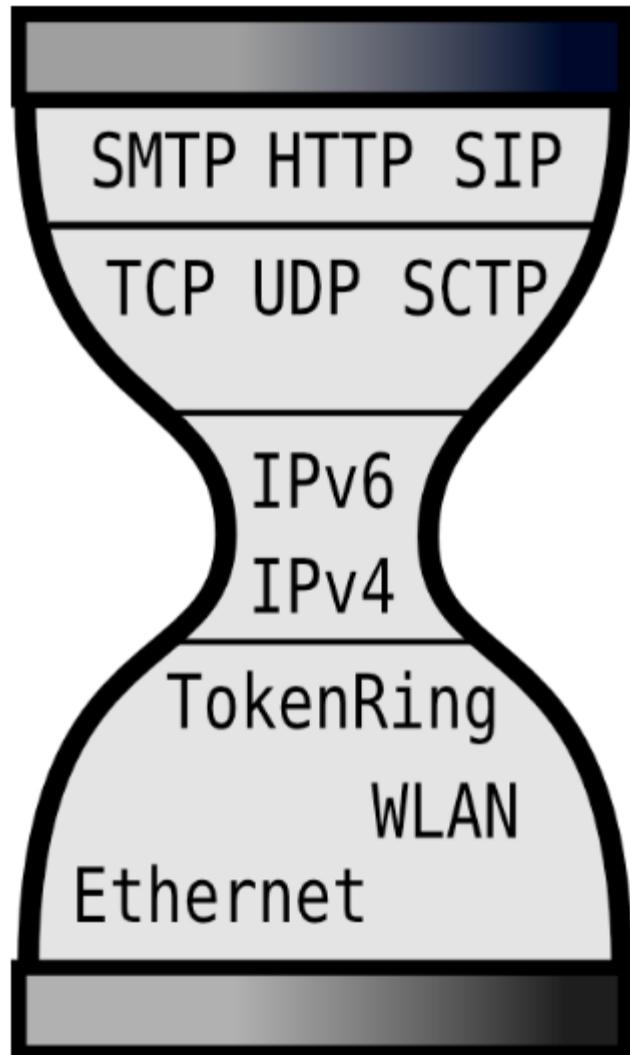
- New and updated IP networking implementations should support IPv4 and IPv6 coexistence (dual-stack), but must not require IPv4 for proper and complete function.
- Implementer are encouraged to update existing hardware and software to enable IPv6 wherever technically feasible.

IPv6 Basics

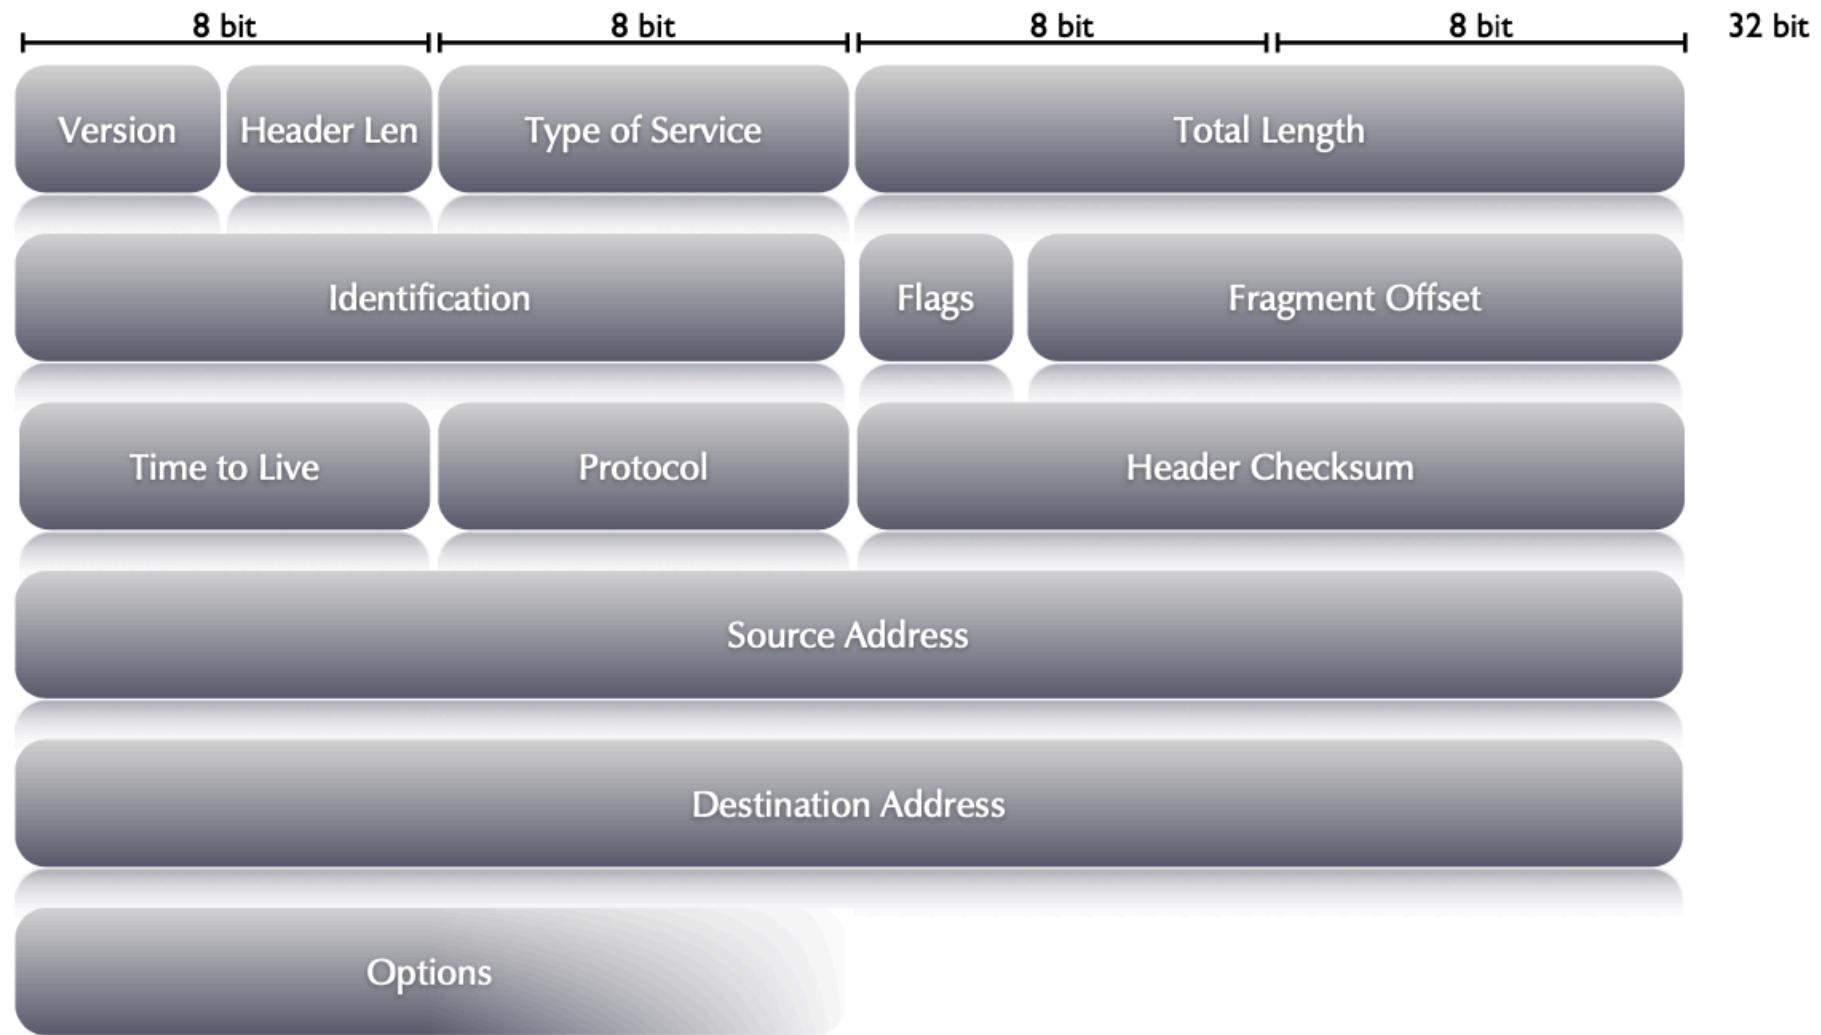




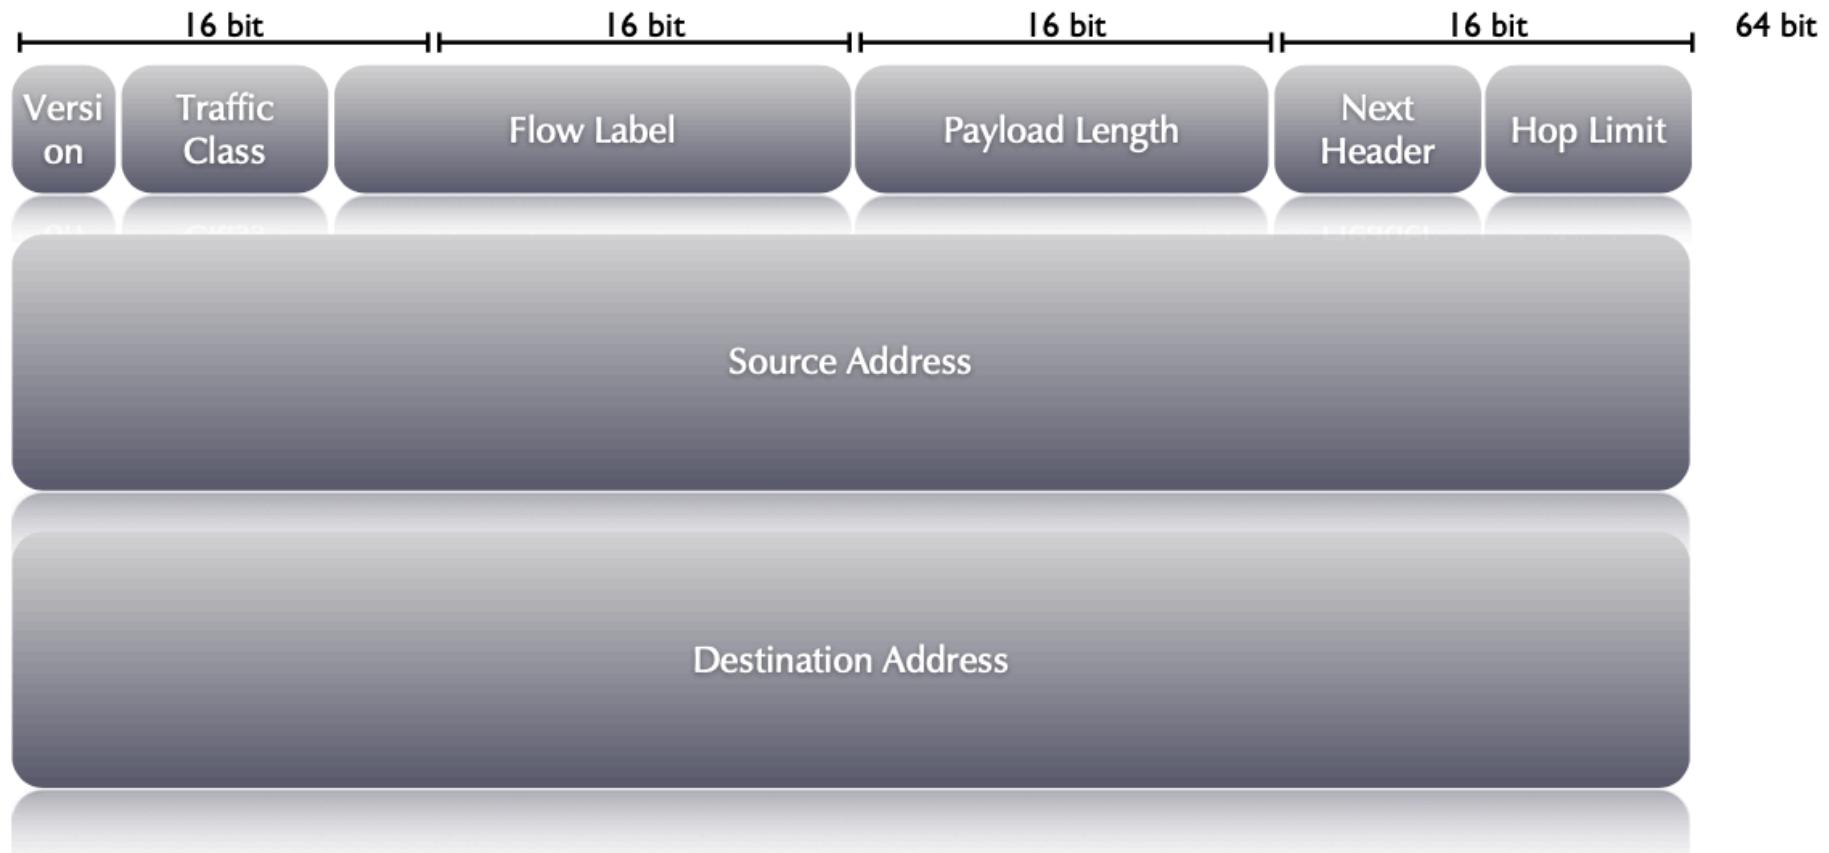
TCP/IP Hourglass for Dual-Stack



The IPv4 Packet



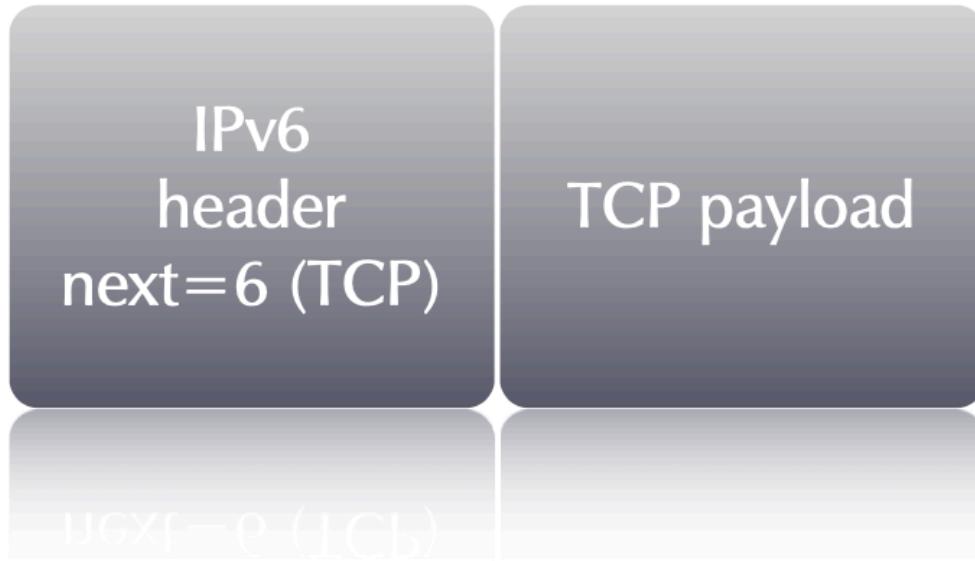
The IPv6 Packet



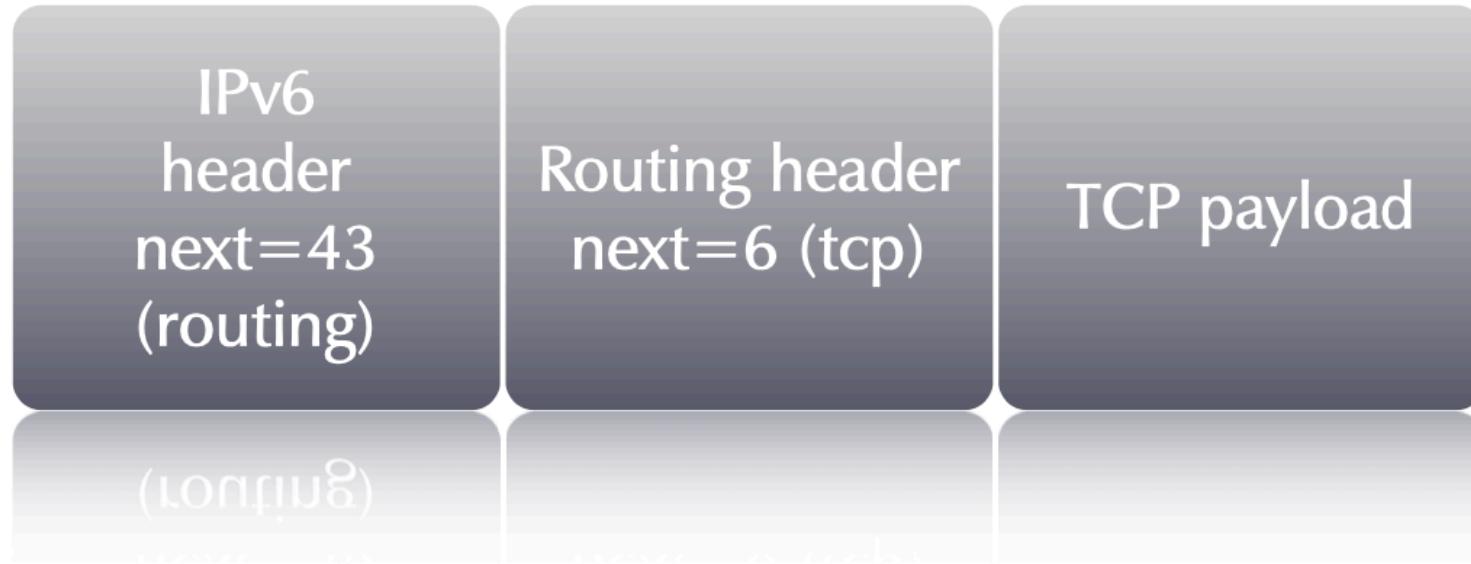
IPv6 Header

- IPv6 uses extension headers instead of fixed options
 - headers are chained one after another
 - *next header* field points to the next extension headers
 - or to the upper level protocol

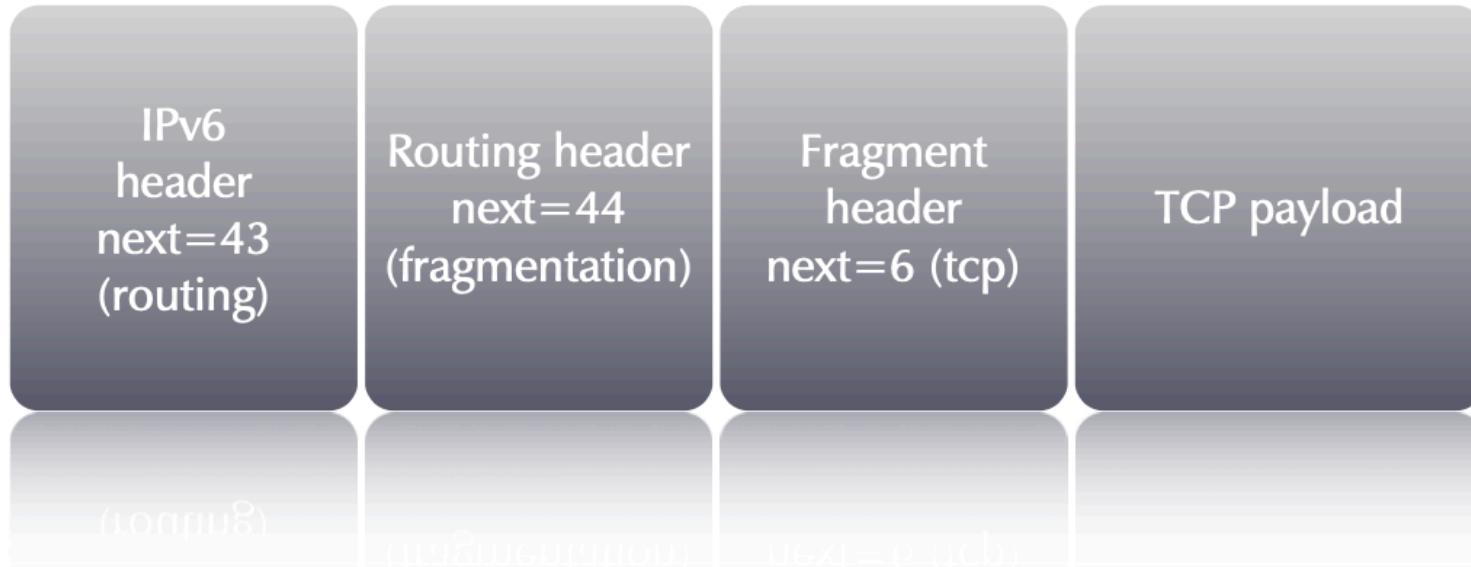
IPv6 Header Example 1



IPv6 Header Example 2



IPv6 Header Example 3



Rules for extension Headers

- The following rules apply to extension headers:
 - Each extension header should not appear more than once, with the exception of the Destination Options header.
 - The Hop-by-Hop Options header should only appear once.
 - The Hop-by-Hop Options header should be the first header in the list because it might be examined by every node along the path.

Rules for extension Headers (2)

- The Destination Options header should appear at most twice (before a Routing header and before the upper-layer header).
- The Destination Options header should be the last header in the list, if it is used at all.
- The Fragment header should not appear more than once and should not be combined with the Jumbo Payload Hop-by-Hop option

Rules for extension Headers (3)

- The entire IPv6 header chain must be contained in the first fragment of a packet
- Other than the path MTU constraints, there are no other limits to the number of IPv6 EHs that may be present in a packet
- The only way for a node to obtain the upper-layer protocol type or find the upper-layer protocol header is to parse and process the entire IPv6 header chain, in sequence, starting from the mandatory IPv6 header until the last header in the IPv6 header chain is found

Order of Extension Header

Order	Header
1	IPv6 Header
2	Hop-by-Hop
3	Destination Options
4	Routing Header
5	Fragment Header
6	Authentication Header
7	ESP Header
8	Destination Options
9	Upper Layer Payload

Operational Implications of Extension Headers

- ↳ RFC 9098 - Operational Implications of IPv6 Packets with Extension Headers
 - Packets with EHs are intentionally dropped in the public Internet in some circumstances

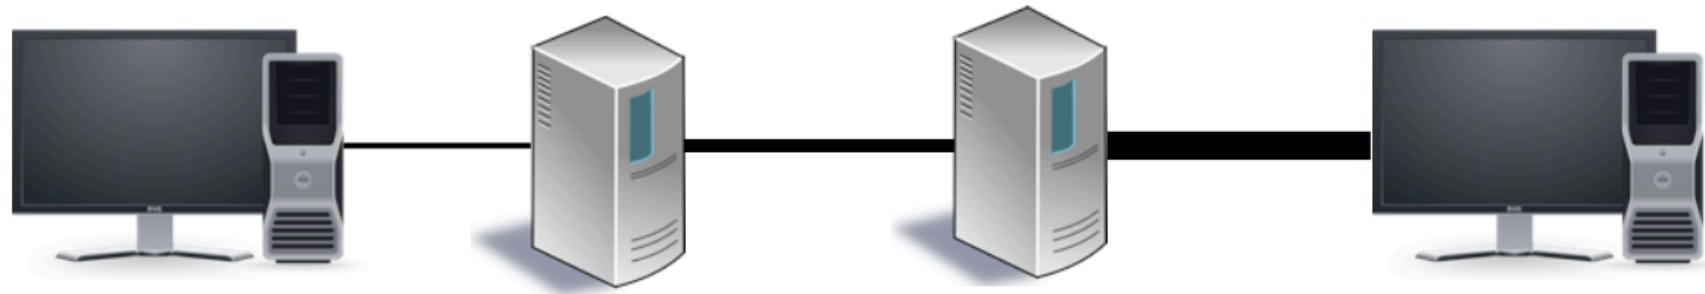
Path MTU Detection (PMTUD)

- Path Maximum Transmission Unit Detection
 - IPv6 Nodes have a minimum MTU of 1280 byte
 - only the sender is allowed to fragment a datagram

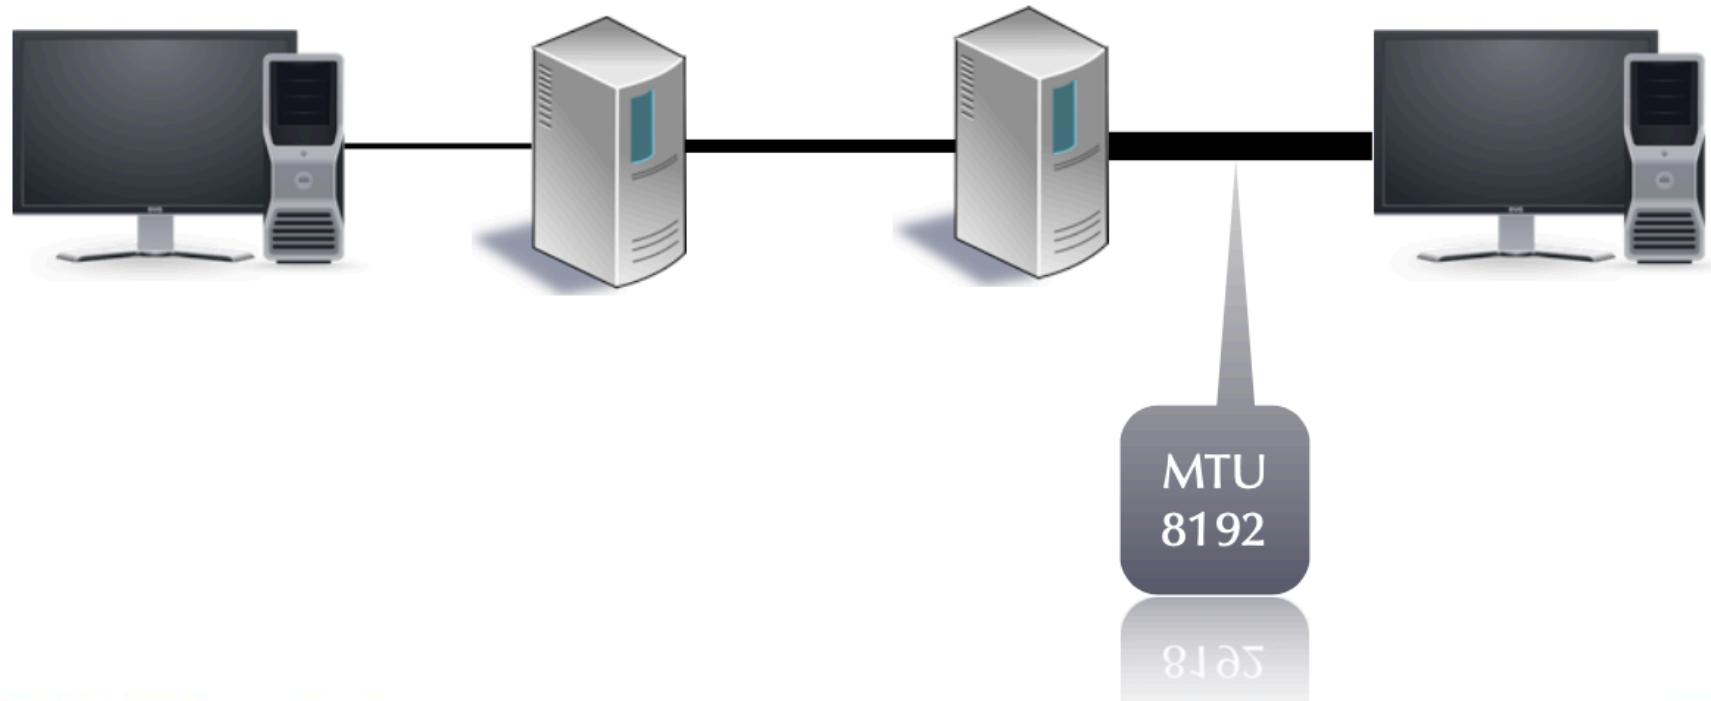
Path MTU Detection (PMTUD)

- if a device on the way cannot handle the size of a packet
 - it will drop the packet
 - and will inform the sender of MTU conflict and the maximum MTU (using ICMPv6 *Packet too big* message)
 - if ICMPv6 *Packet too big* messages are blocked, IPv6 is broken!

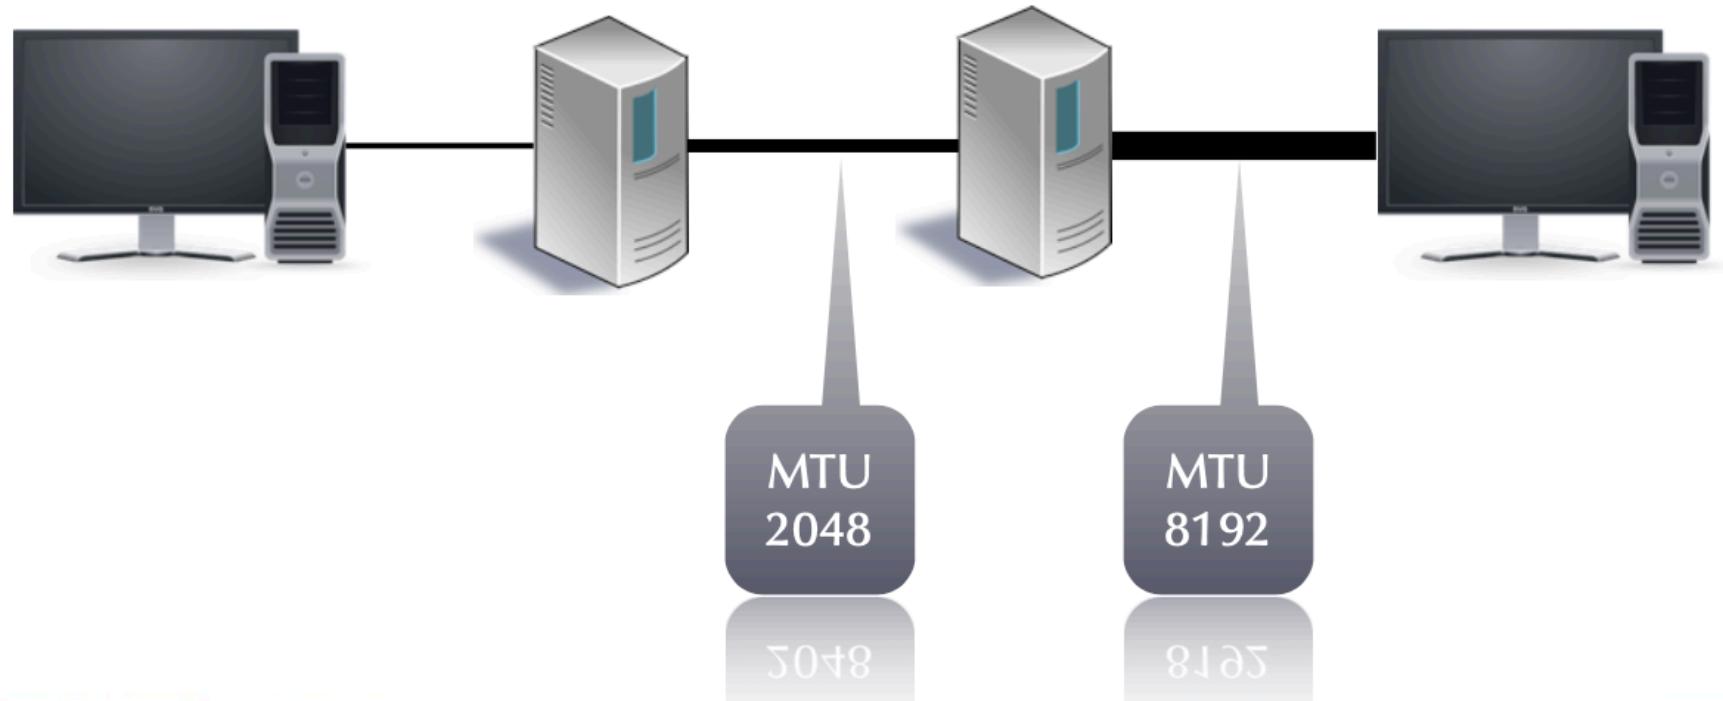
PMTUD (1)



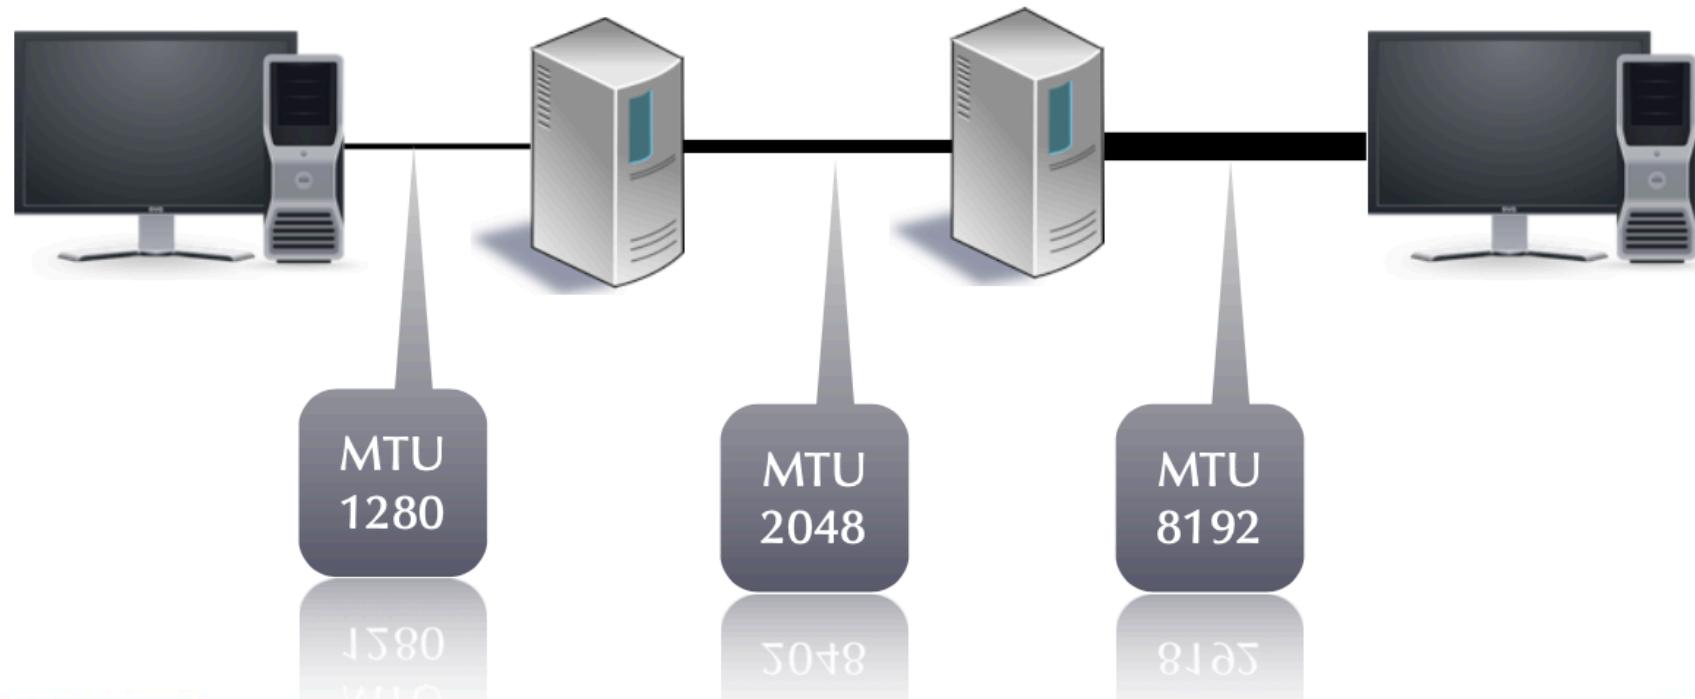
PMTUD (2)



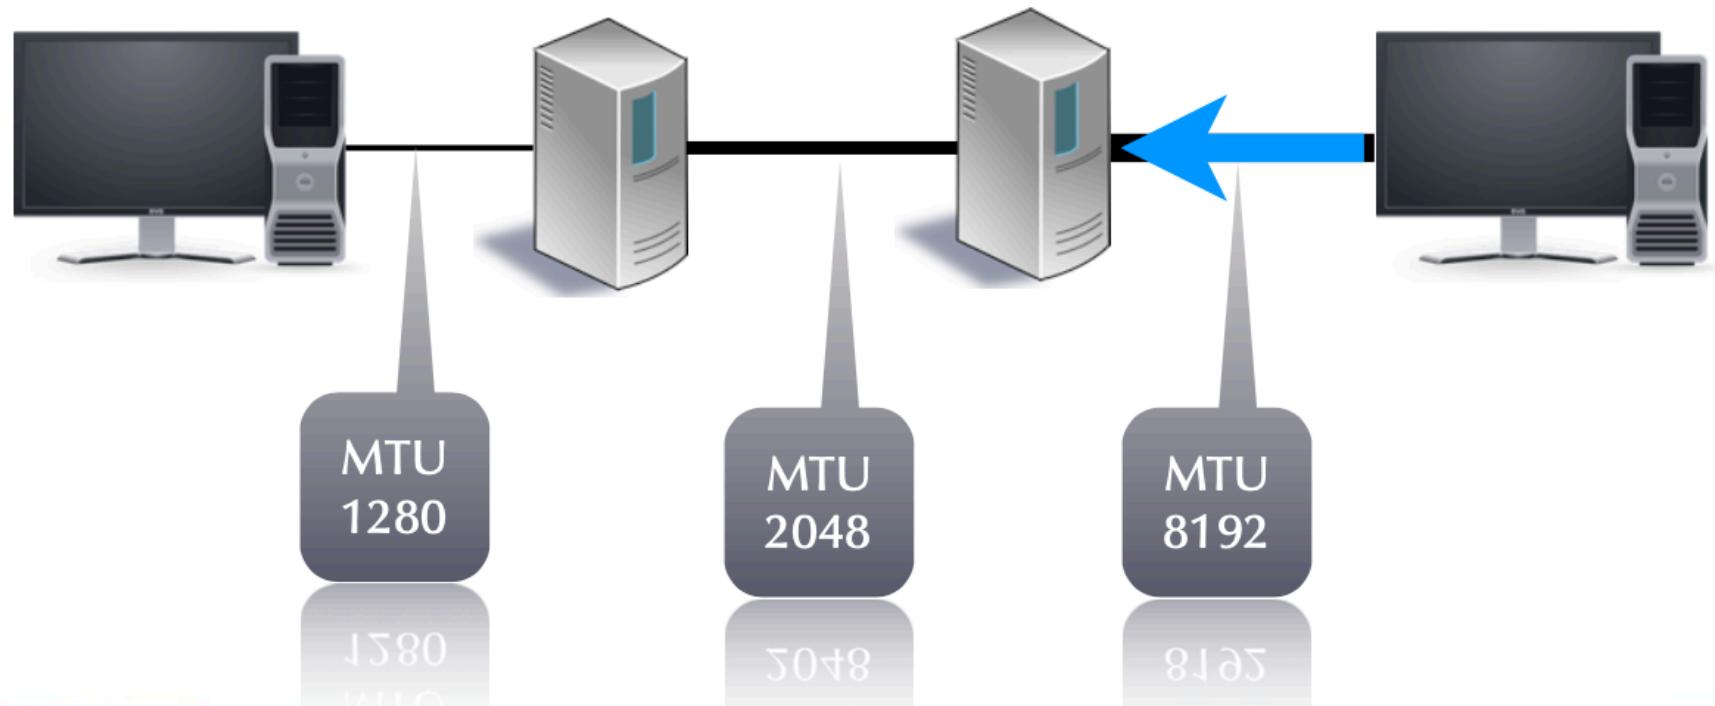
PMTUD (3)



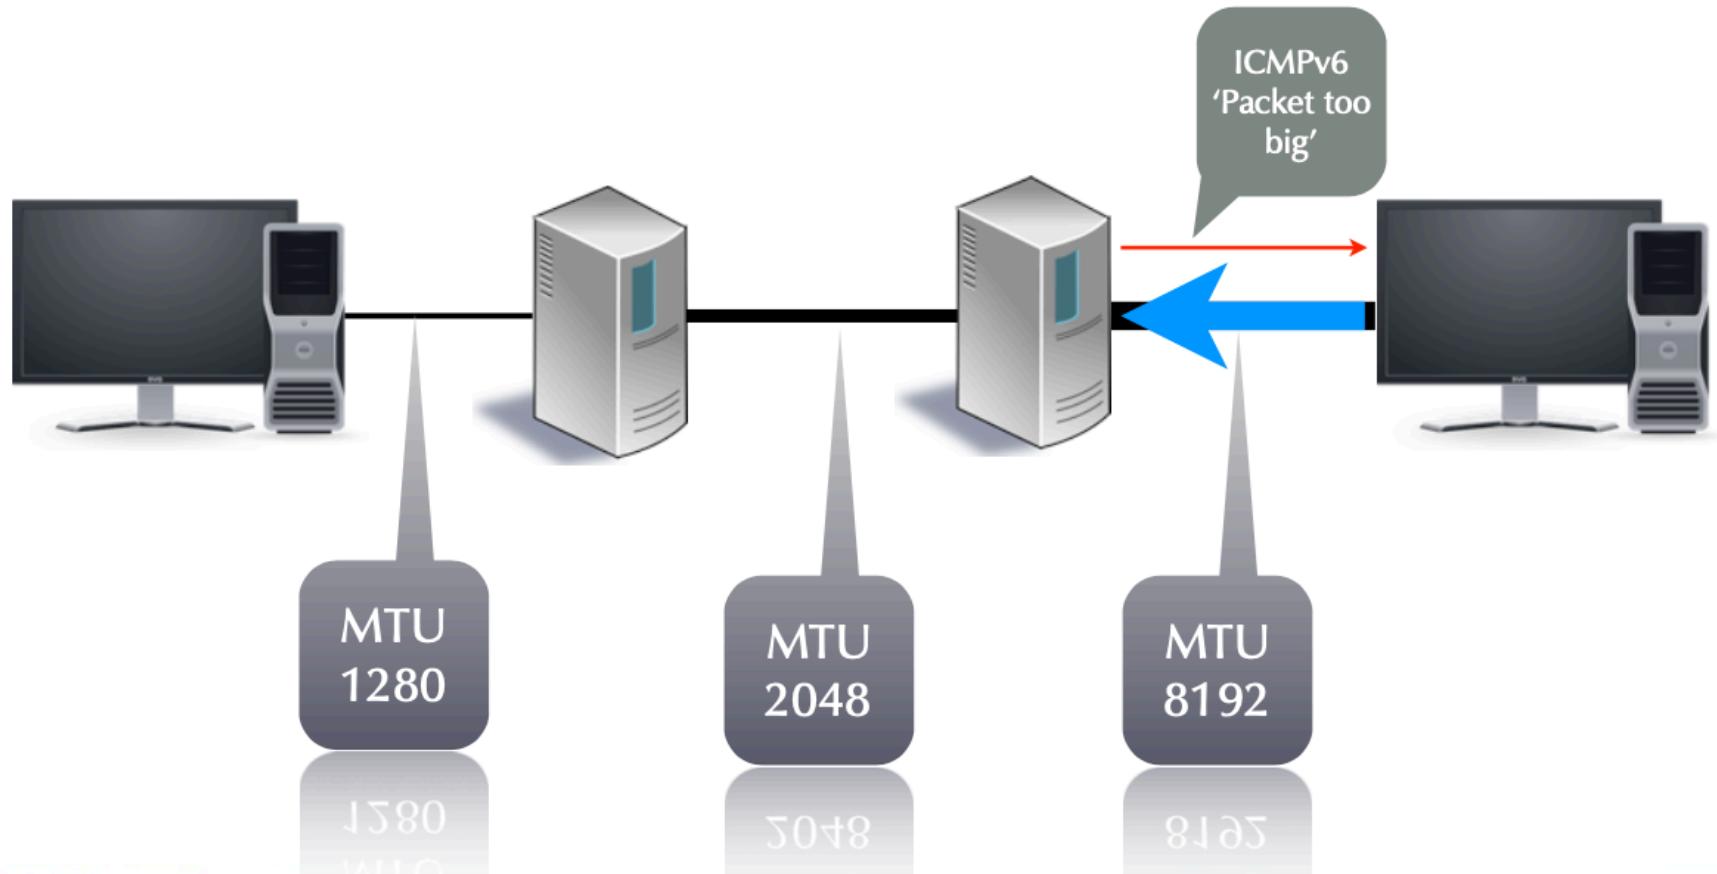
PMTUD (4)



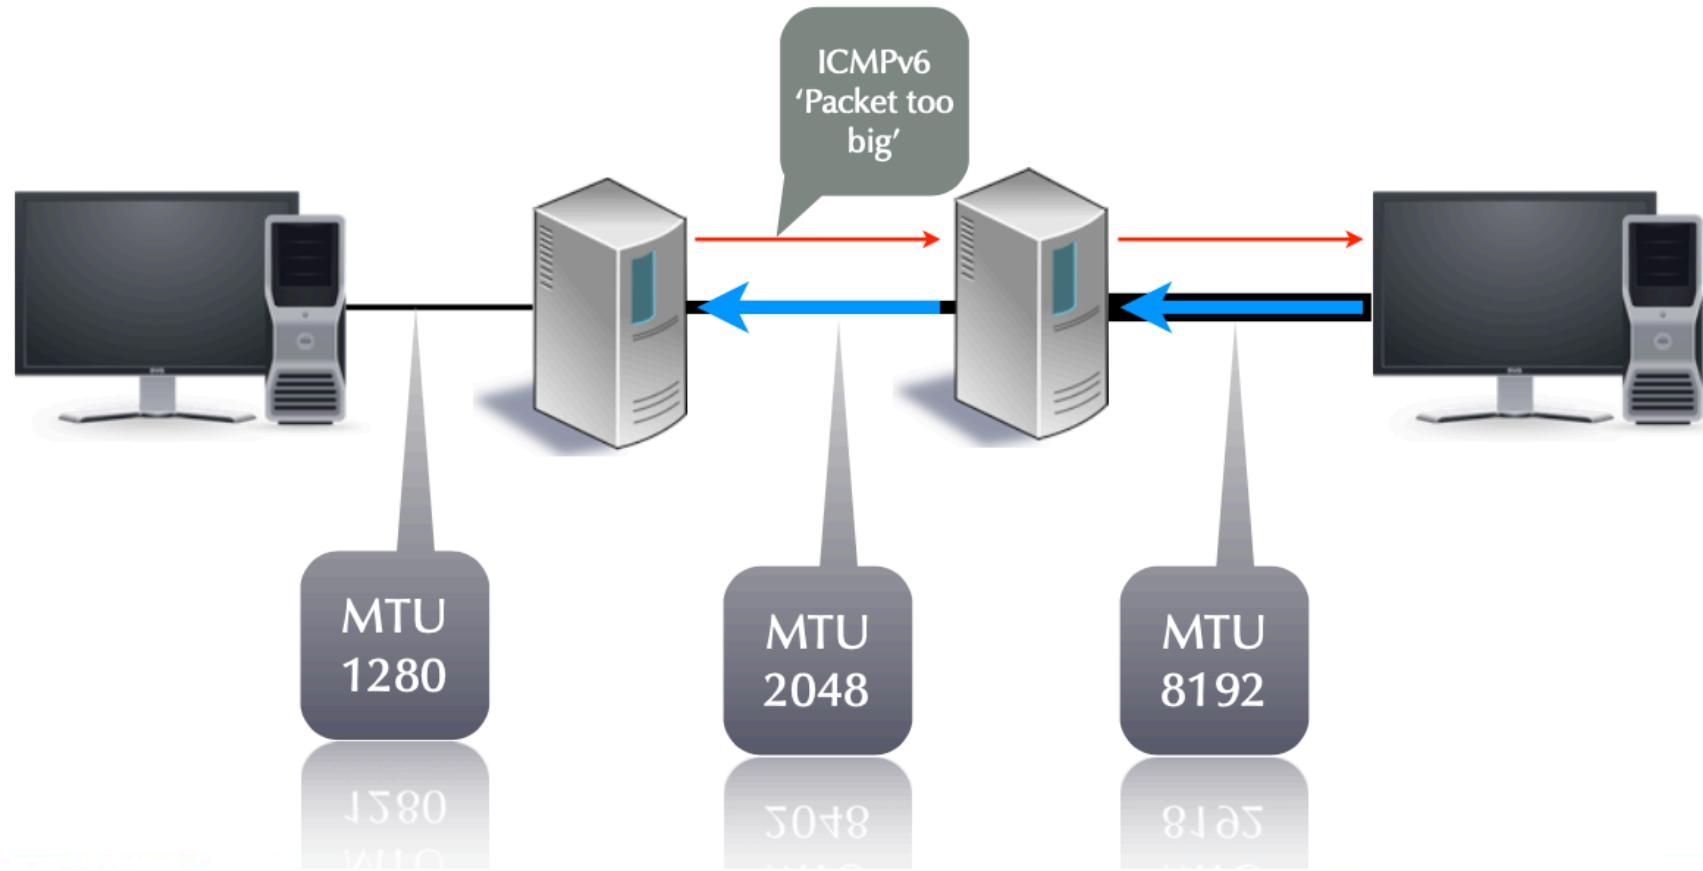
PMTUD (5)



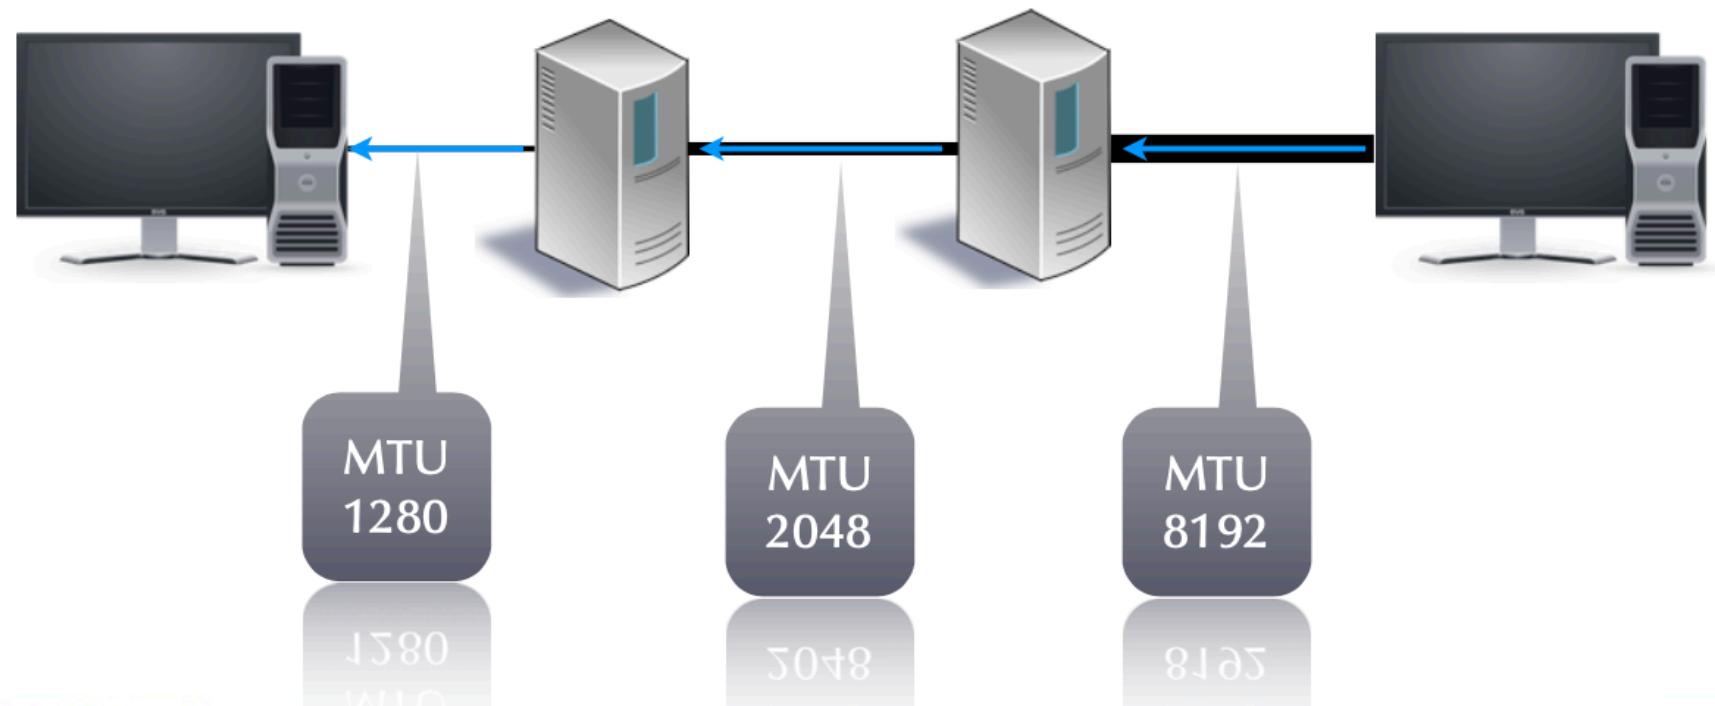
PMTUD (6)



PMTUD (7)



PMTUD (8)



IPv6 Addresses

IPv6 Addresses

- IPv6 Addresses are 128bit long (16 bytes)
 - IPv4 had 32bit per Address
 - per Bit, the Address space doubles
- IPv6 has 2^{128} Addresses (340 billion billion billion billion or 3.4×10^{38} or 340 undecillion)

IPv6 Address Format

- IPv6 Addresses are written using 32 hexadecimal digits
 - the digits are arranged in 8 groups of four digits, separated by colon :
 - the display rules for IPv6 addresses have been redefined in RFC 5952

IPv6 Address Format

- A full IPv6 Address:

2001:0db8:1c01:0029:021a:5f4f:fe9a:5aab

IPv6 Address Abbreviations

- leading zeros of every group of 4 hex-digits can be omitted:
2001:0db8:1c01:0029:021a:5f4f:fe9a:5aab →
2001:db8:1c01:29:21a:5f4f:fe9a:5aab

IPv6 Address Abbreviations

- *one* sequence of all zero groups can be replaced by a pair of colons: 2001:db8:100:0:0:0:541:fa20 → 2001:db8:100::541:fa20

IPv6 Address Abbreviations

- IPv6 Address Prefixes:
 - 2001:db8:100::/64
 - 2001:db8::/32
 - 2002::/16
- The Prefix defines the number of important bits in this network from the start

IPv6 prefix and host-id

- Prefix and Host-ID

2001:db8:100:200:abba:cafe:23:422



Prefix

Host-ID

Types of IPv6 Addresses

- Unicast (one to one communication)
- Multicast (one to many communication)
- Anycast (one to some communication)

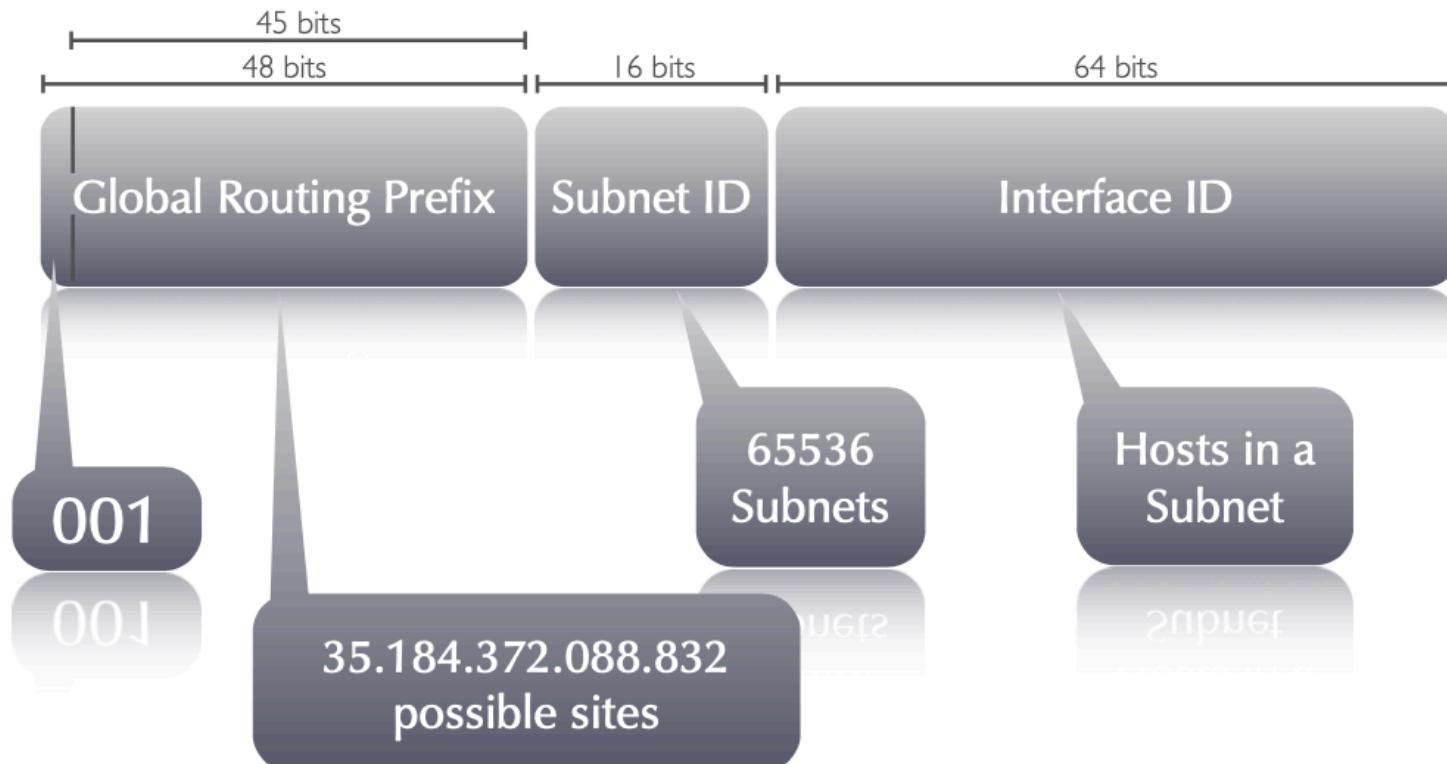
Types of IPv6 Unicast Addresses

- Global Unicast Addresses (usable in the whole Internet)
- Link-Local Addresses (only valid in a single subnet)
- Site-Local Addresses (deprecated, originally for a *site*)
- Unique-Local Addresses (like private IP addresses, but statistically unique)
- Special Addresses (see next slide)
- Transition Addresses (IPv4 to IPv6 Transition Protocols)

IPv6 Special Addresses

- Unspecified Address (all zeros): ::
- Loopback Address: ::1

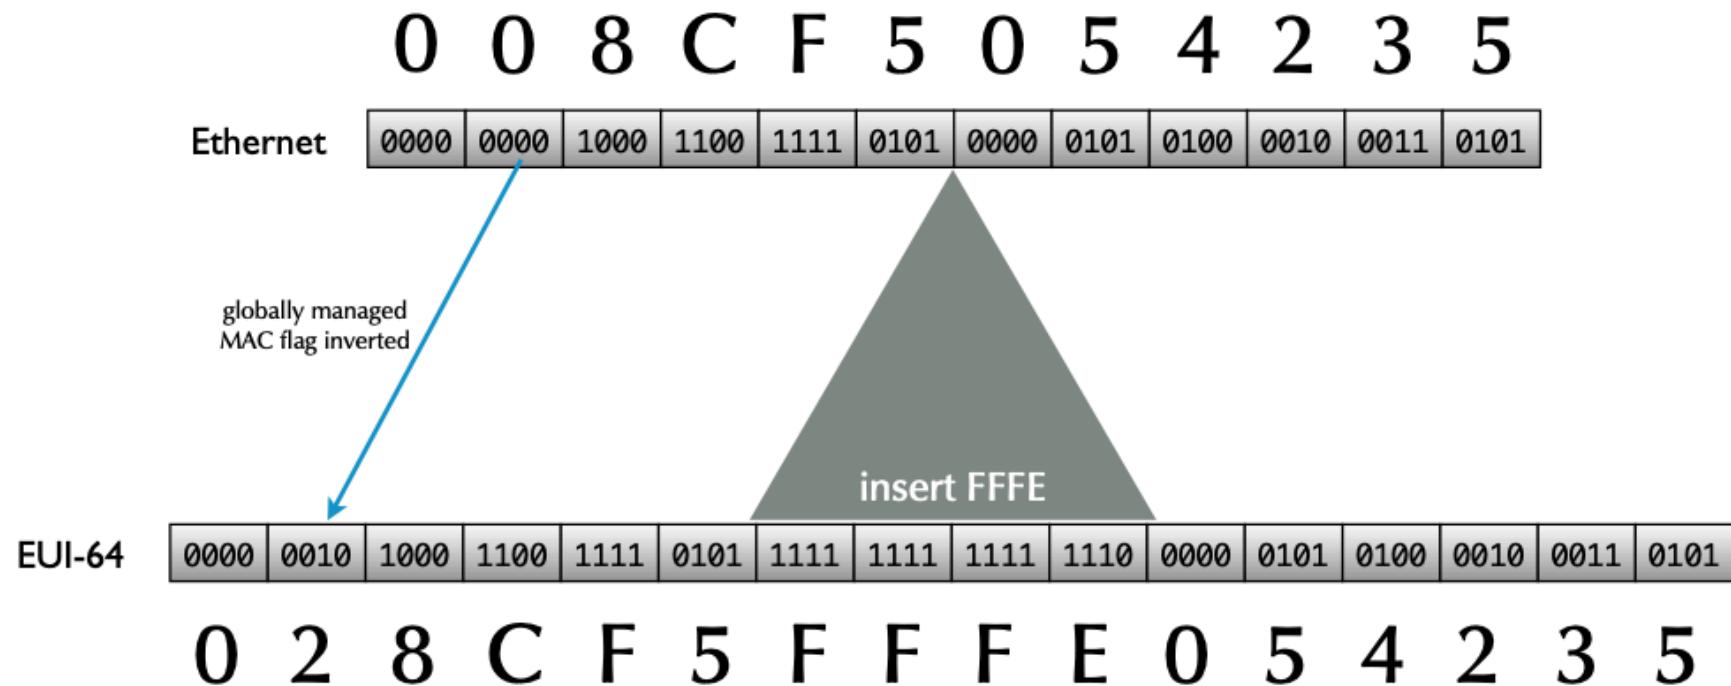
Global Unicast Addresses



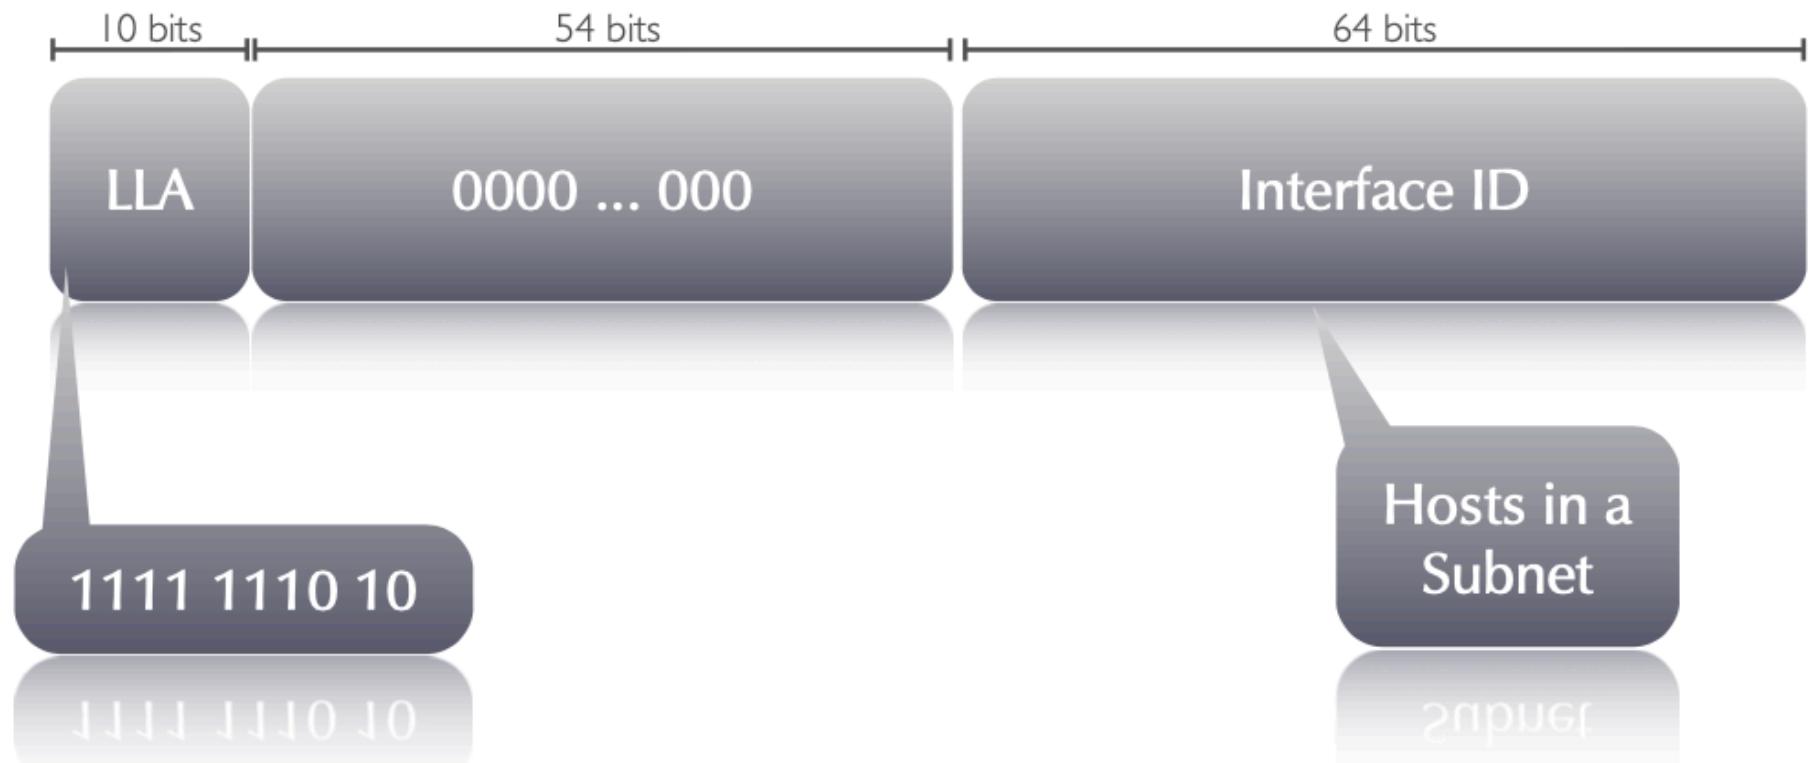
IPv6 Interface ID

- The Interface ID can be build from
 - an Interface EUI-64 ID (7th bit inverted, used for stable addresses on server machines)
 - MAC Address (48bit) (7th bit inverted and 0xffffe inserted after 3rd byte of the MAC address)
 - based on the network (stable only for a single network)
 - random (privacy)
 - manual

MAC-Address to Interface ID



Link-Local-Addresses



Link-Local-Addresses

- Link-Local addresses are used for communication of hosts on a single link
 - link-local addresses are always configured if IPv6 is enabled
 - used for Neighborhood Discovery and local communication

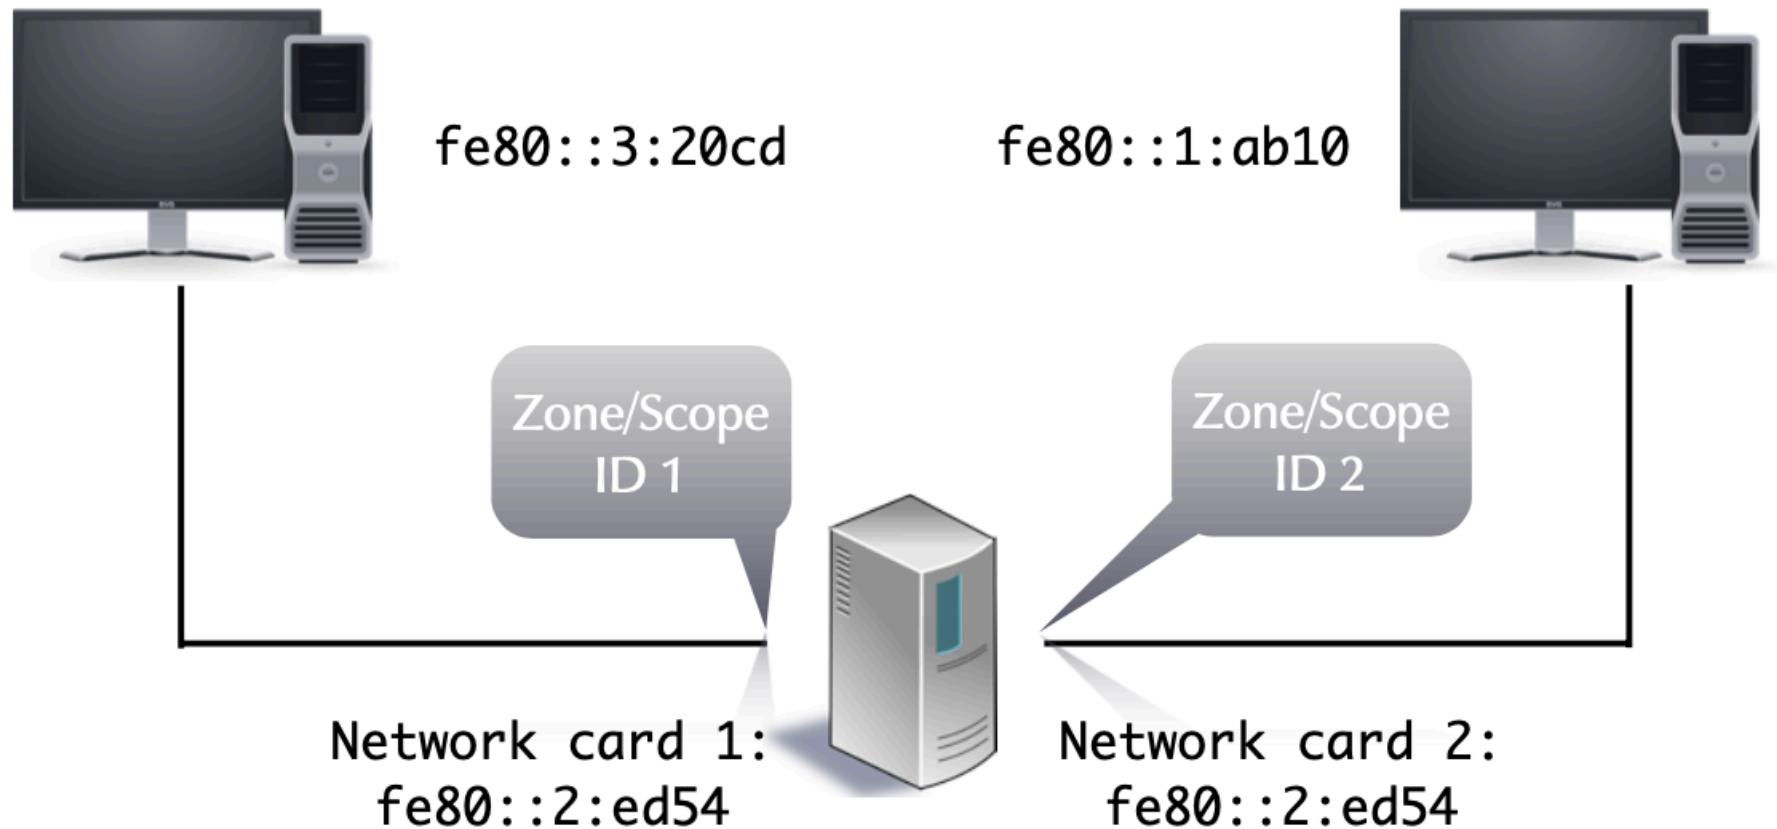
Link-Local-Addresses

- Prefix for Link-Local Addresses is `fe80::/10`
 - IPv6 router never forward link-local addresses
 - similar to IPv4 APIPA (automatic private IP addressing, `169.254.0.0/16`)

Zone IDs

- Link-Local addresses can be reused on every link
 - a machine with more than one link can reach different machines with the same link-local IPv6 address

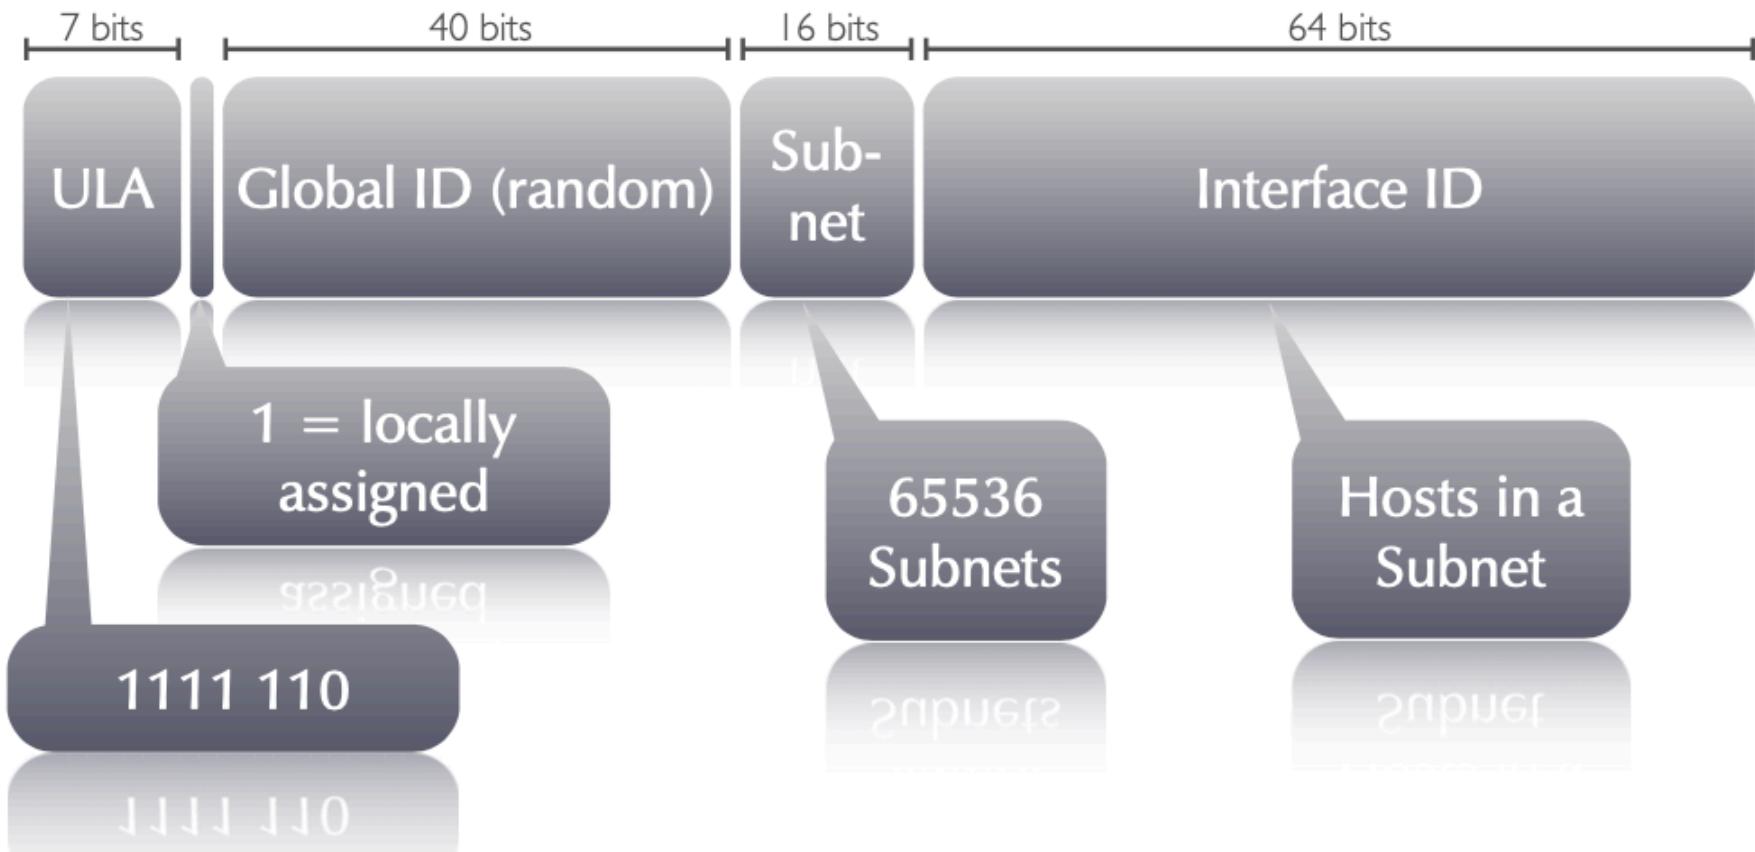
Zone IDs



Zone/Scope ID

```
en1: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
      inet6 fe80::226:bbff:fe05:fa4f%en1 prefixlen 64 scopeid 0x5
      inet 10.122.104.112 netmask 0xffffffff80 broadcast 10.122.104.127
        ether 00:26:bb:05:fa:4f
      media: autoselect status: active
      supported media: autoselect
```

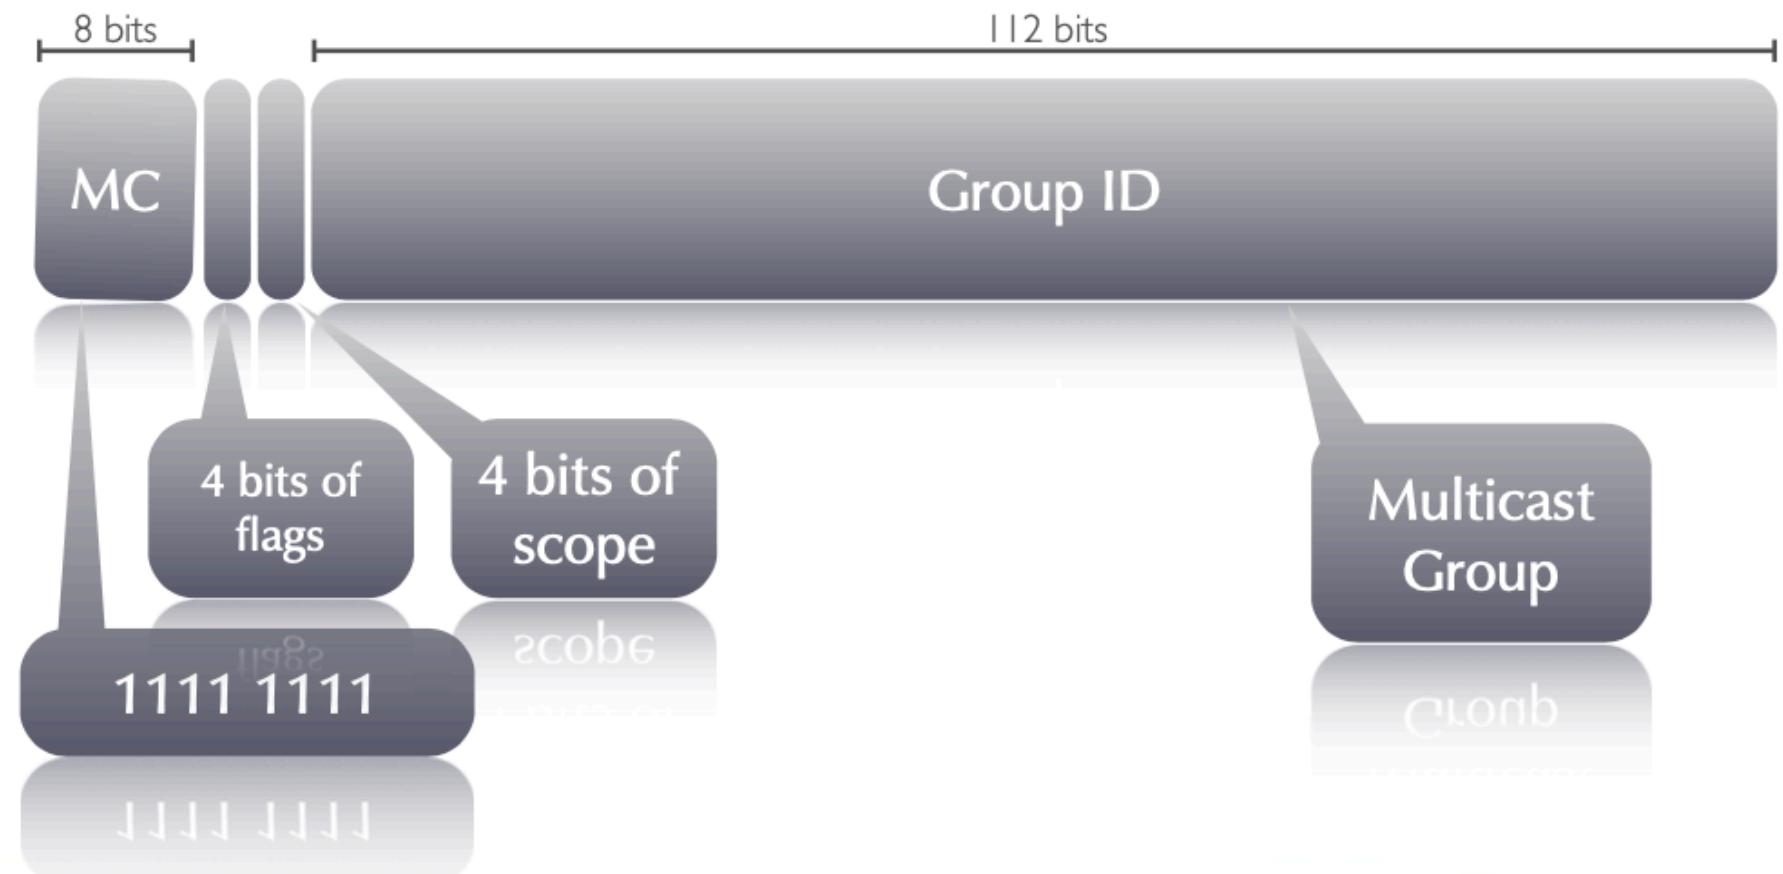
Unique-Local-Addresses



Unique-Local-Addresses

- Unique local addresses are private to an organization, yet statistically unique
 - low probability of address clashes if two sites with ULA merge
 - the L flag indicates local assignment (default)
- ULA Addresses have the prefix `fd00::/8`
 - ULAs have a global scope (not scope/zone ID) •will not be routed in the Internet

Multicast Addresses



Multicast Addresses

- multicast addresses always starts with ff
 - IPv6 uses multicast for internal services
 - no broadcasts anymore, only multicast!

Multicast Address Flags

BIT USAGE

0	Transient Flag: 0 = permanent address (IANA assigned, well-known) 1 = transient address
1	Prefix flag (RFC 3306)
2	Rendezvous Point (RFC 3956)
3	undefined

Multicast Address Scope

HEX-VALUE	SCOPE
1	Interface local
2	Link-Local
4	Admin-Local
5	Site-Local
8	Organization Local
E	Global

Node Local Scope

ADDRESS USAGE

ff01::1 all nodes address

ff01::2 all routers address

ff01::fb mDNSv6

Link-Local Scope

ADDRESS	USAGE
ff02::1	all nodes address
ff02::2	all routers address
ff02::fb	mDNSv6
ff02::f	UPnP Devices
ff02::1:1	Link Name
ff02::1:2	all DHCP agents
ff02::1:3	LLMNR

ADDRESS

USAGE

ff02::101

network time protocol (ntp)

ff02::1:ffXX:XXXX

Solicited-Node Address

ff02:0:0:0:0:2:ff00::/104

Node Information Queries [RFC 4620](#)

Site-Local/Variable Scope

ADDRESS USAGE

ff05::2 all routers address

ff05::fb mDNSv6

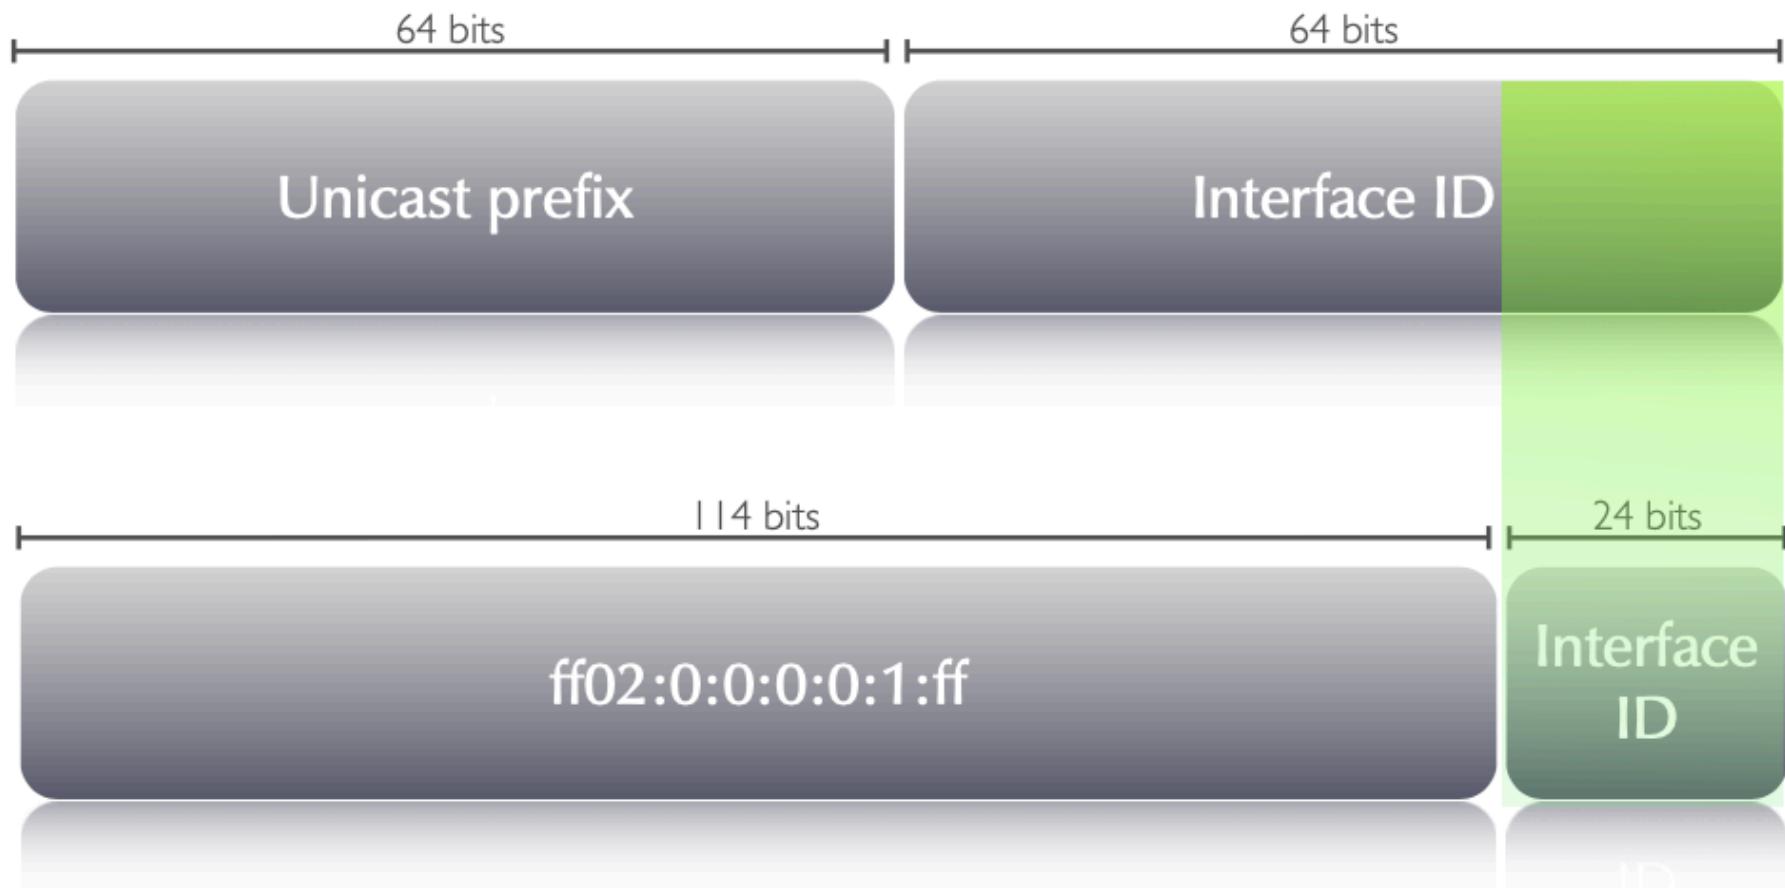
ff05::1:3 all DHCPv6 Servers

ff0x::101 network time protocol (NTP)

Solicited-Node Address

- the Solicited node address is used for link-layer address resolution
 - was ARP in IPv4
- the solicited-node multicast address is build from the prefix `ff02::1:ff00:0/104` and the last 24 bits of the unicast IPv6 address to be resolved

Solicited-Node Address



Transition Addresses

- IPv4 compatible Addresses $0:0:0:0:0:0:w.x.y.z$
($::w.x.y.z$)
 - deprecated by RFC4291
- IPv4 mapped Addresses $0:0:0:0:0:ffff:w.x.y.z$
($::ffff:w.x.y.z$)

Transition Addresses

- 6to4 Addresses: 2002:WWXX:YYZZ:<subnetid>:<interfaceID>
- ISATAP Addresses <64bitprefix>:0:5efe:w.x.y.z

special Addresses

- Original IPv6 address prefix for documentation
2001:db8::/32 ([↪RFC 3849](#))
- Additional IPv6 address prefix for documentation 3fff::/20
([↪RFC 9637](#))

Questions?

