FRANKFURT UNIVERSITY OF APPLIED SCIENCES

COMPUTER NETWORKS

Network Layer - Internet Protocol

Prof. Dr. Oliver Hahm

2024-12-19

AGENDA

- Addressing
- IPv4 Networks
- IPv6 Networks
- Packet Structure
- ICMP
- Address Autoconfiguration

THE NARROW WAIST OF THE INTERNET

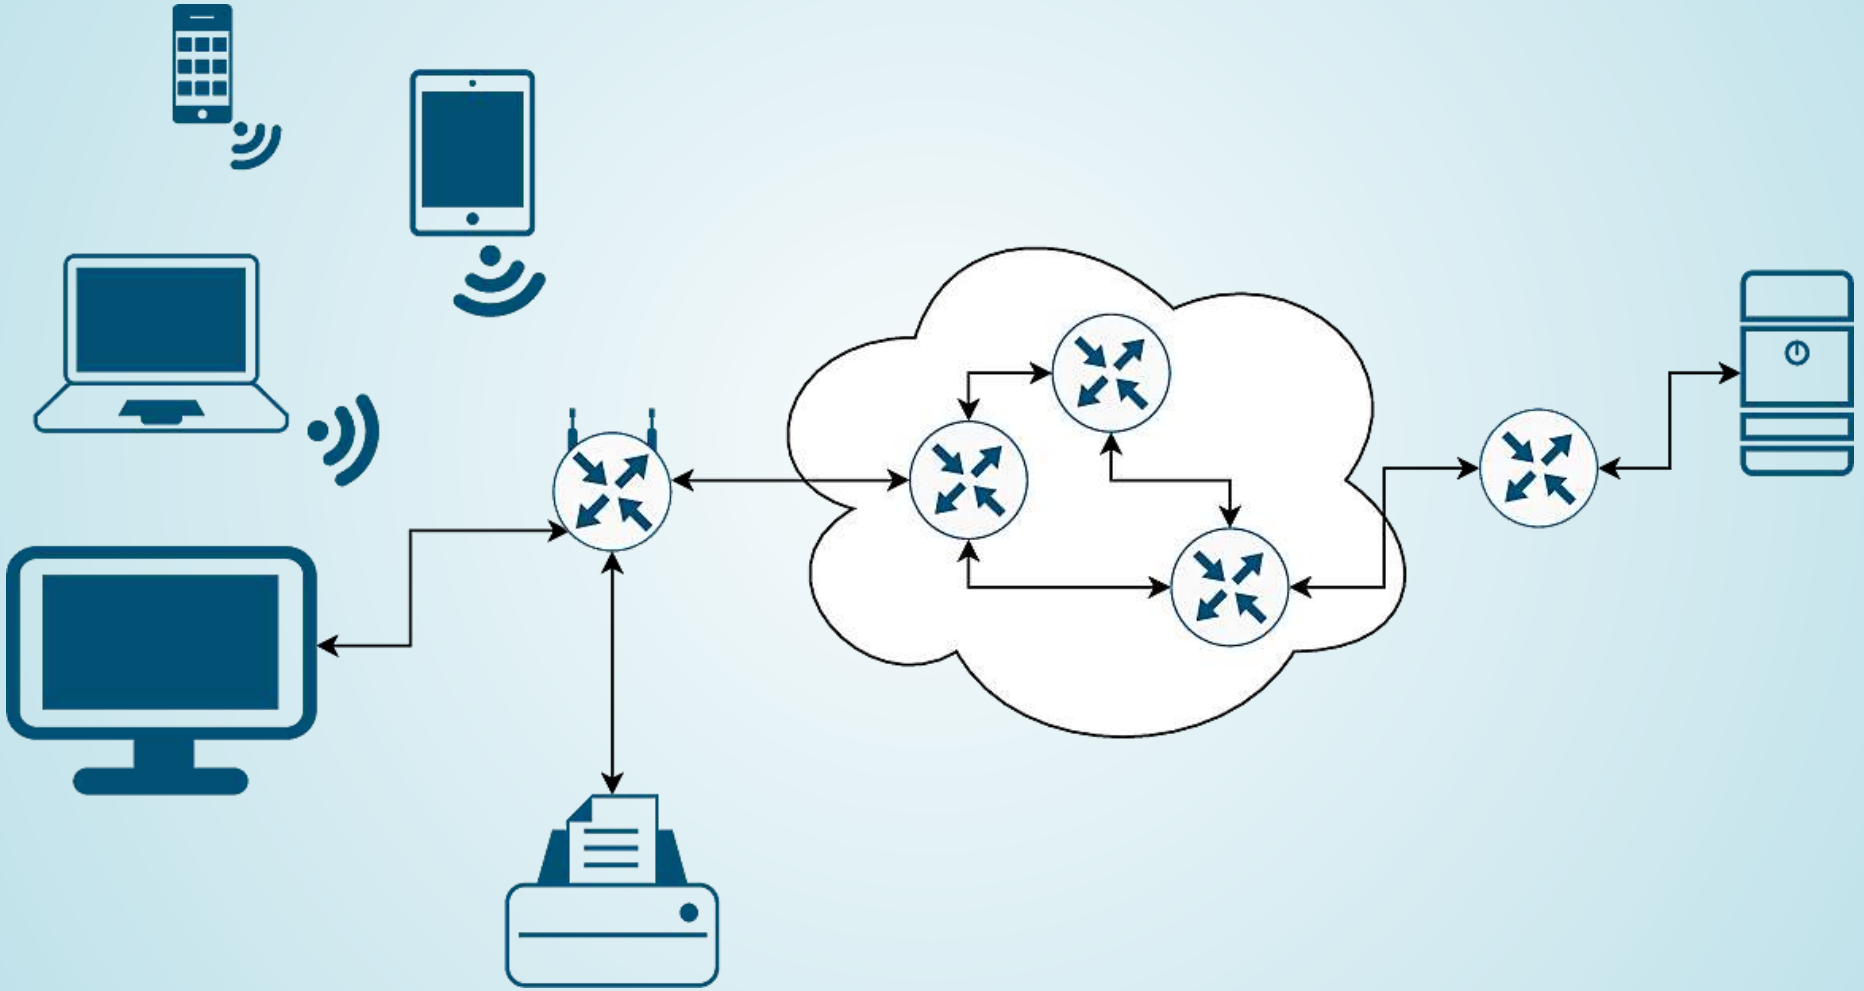

ADDRESSING

COMPUTER NETWORK ADDRESSES

- How many MAC address does your computer have?
- How many IP address does your computer have?

PURPOSE AND FORMAT

ADDRESSING IN THE NETWORK LAYER

- Physical addresses ($\rightarrow$ MAC addresses) are bound to a device
  $\Rightarrow$ it is impossible to maintain a logical hierarchy or replace hosts in a transparent manner
- Logical addresses are required, which are independent from the specific hardware
  - Logical addressing separates the logical position within the network from a physical device

Address Assignment

For local networks manual address assignment is typically not desired, hence mechanisms for *address autoconfiguration* are required.

FORMAT OF IP ADDRESSES

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

- IPv4 addresses have a length of 32 bits (4 bytes)
  - Thus, the address space contains $2^{32} = 4,294,967,296$ possible addresses
- IPv6 addresses have a length of 128 bits (16 bytes)
  - Thus, the address space contains $2^{128} = 3.4 * 10^{38}$ possible addresses

Address space = amount of all valid network identifiers

- The usual representation of IPv4 uses the dot-decimal notation
  e.g., `198.51.100.23` [1]
- The usual representation of IPv6 uses the hexadectets (*quad-nibbles*) seperated by colons
  e.g., `2001:0db8:0000:0000:0000:ff00:0042:8329` [2]

1. See RFC 5737
2. See RFC 3849

IPV4 NETWORKS

NETWORK IDENTIFIER AND HOST IDENTIFIER

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

- The 32 bits of an IPv4 address are split into
  - Network identifier (network ID)
  - Host identifier (host ID)
- All hosts with an identical **network ID** are in the same network
- How many bits are used for the **network ID** differs
  - The fewer bits are used for the **network ID**, the more bits remain for the **host ID**
  - $\Rightarrow$ The more hosts a network may comprise

(SUB)NETMASKS

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

- For specifying the network size a (sub-)netmask is required
  - All hosts in a network have a netmask assigned
  - A network may be further divided by using the first bits of the **host** as subnet identifier - this process is called subnetting
- Structure of the netmask:
  - 1-bits indicate, which part of the address space is used for network IDs
  - 0-bits indicate, which part of the address space is used for host IDs

Two **host IDs** are reserved
(i.e., cannot be assigned to network devices):

- Network Address: All host ID bits are set to 0 $\Rightarrow$ reserved to identify the network itself
- Broadcast Address: All host ID bits are set to 1 $\Rightarrow$ reserved for the broadcast address

NETMASK EXAMPLES

Network Mask: 255.255.0.0

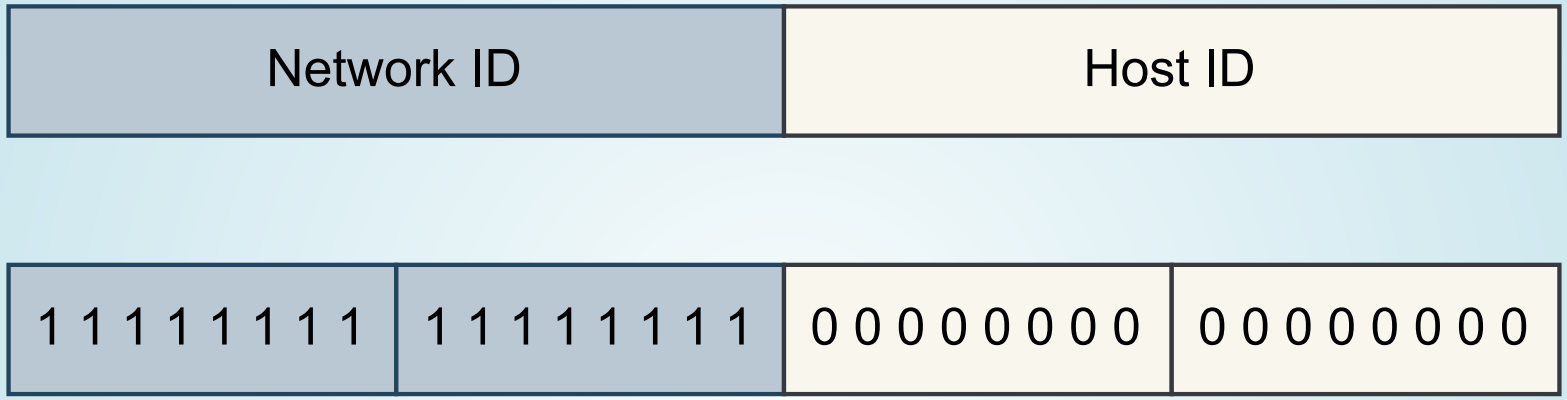

Network Mask: 255.255.255.240

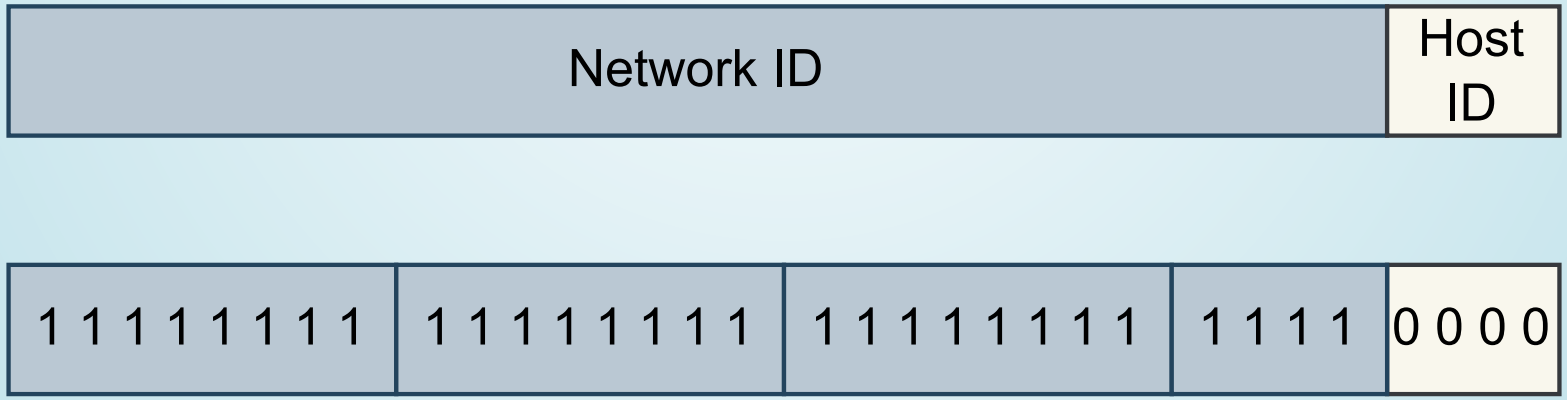

ADDRESS CLASSES

- Originally, IPv4 addresses were categorized into classes from A to C
  - Additionally, the classes D and E for special purposes existed

Class	Prefix	Address range			Network ID	Host ID
A	`0`	0.0.0.0	-	127.255.255.255	7 bits	24 bits
B	`10`	128.0.0.0	-	191.255.255.255	14 bits	16 bits
C	`110`	192.0.0.0	-	223.255.255.255	21 bits	8 bits
D	`1110`	224.0.0.0	-	239.255.255.255	—	—
E	`1111`	240.0.0.0	-	255.255.255.255	—	—

- $2^7 = 128$ **class A** networks with a maximum of $2^{24} = 16,777,216$ host addresses each
- $2^{14} = 16,384$ **class B** networks with a maximum of $2^{16} = 65,536$ host addresses each
- $2^{21} = 2,097,152$ **class C** networks with a maximum of $2^8 = 256$ host addresses each
- Class D contains multicast addresses
- Class E is reserved for future purposes and experiments

DRAWBACK OF ADDRESS CLASSES

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

- The original intention was to identify physical networks in an unique way via the network ID
- Drawbacks of Address Classes:
  - It is impossible to dynamically adjust them
  - Many addresses are wasted
    - A **class C** network with 2 devices wastes 253 addresses
    - The address space of **class C** networks is quite small
    - A **class B** network with 256 devices wastes $>$ 64,000 addresses
    - Only 128 **class A** networks exist
    - Migrating multiple devices to a different network class is complex task
- **Solution:** Logical networks are divided into subnets
  - 1993: Introduction of the Classless Interdomain Routing (CIDR)

SYNTAX OF THE CLASSLESS INTERDOMAIN ROUTING (CIDR)

- According CIDR IP address ranges are represented by this notation:
  `Network address/mask bits`
  - The number of mask bits indicates the number of 1-bits (prefix) in the subnet mask
- The table shows the possible splits of a **class C** network into subnets

Mask bits (prefix)	/24	/25	/26	/27	/28	/29	/30	/31	/32
Subnet mask	0	128	192	224	240	248	252	254	255
Subnet bits	0	1	2	3	4	5	6	7	8
Subnets IDs	1	2	4	8	16	32	64	128	256
Host bits	8	7	6	5	4	3	2	1	0
Host IDs	256	128	64	32	16	8	4	2	—
Hosts (maximum)	254	126	62	30	14	6	2	0	—

SUBNETTING EXAMPLE

Class B IP address

Subnet mask (255.255.248.0)

A part of the hosts IP address includes the subnet identifier

SUBNETS AND ROUTING

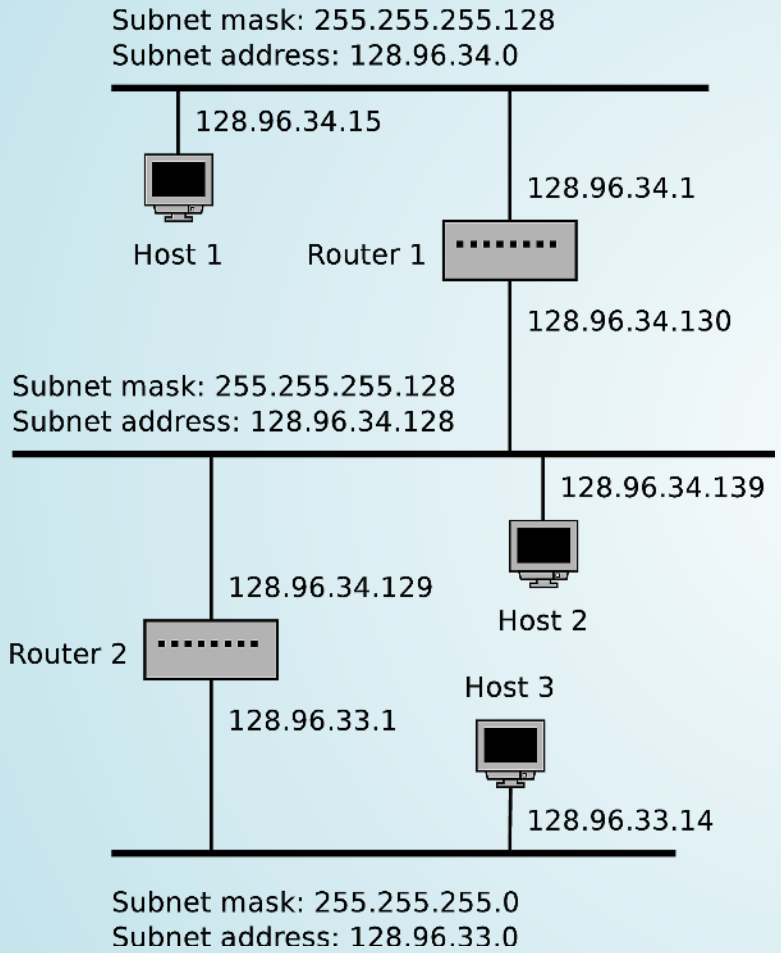

Source: Computernetzwerke. Peterson and Davie. dpunkt (2000)

- All hosts inside the same subnet have the same subnet mask
- If a host wants to transmit a packet, it performs a logical AND operation for its own subnet mask and the destination IP address
  - If the result is equal to the subnet address of the sender, the sender learns that the destination is inside the same subnet
  - If the result does not match the subnet address of the sender, the packet must be transmitted to a router, which forwards it to another subnet

PRIVATE NETWORKS

PRIVATE IPV4 ADDRESS SPACES

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Address space:	10.0.0.0 to 10.255.255.255
CIDR notation:	10.0.0.0/8
Number of addresses:	2^{24} = 16,777,216

Address space:	172.16.0.0 to 172.31.255.255
CIDR notation:	172.16.0.0/12
Number of addresses:	2^{20} = 1,048,576

Address space:	192.168.0.0 to 192.168.255.255
CIDR notation:	192.168.0.0/16
Number of addresses:	2^{16} = 65,536

NETWORK ADDRESS TRANSLATION (NAT)

NETWORK ADDRESS TRANSLATION (NAT)

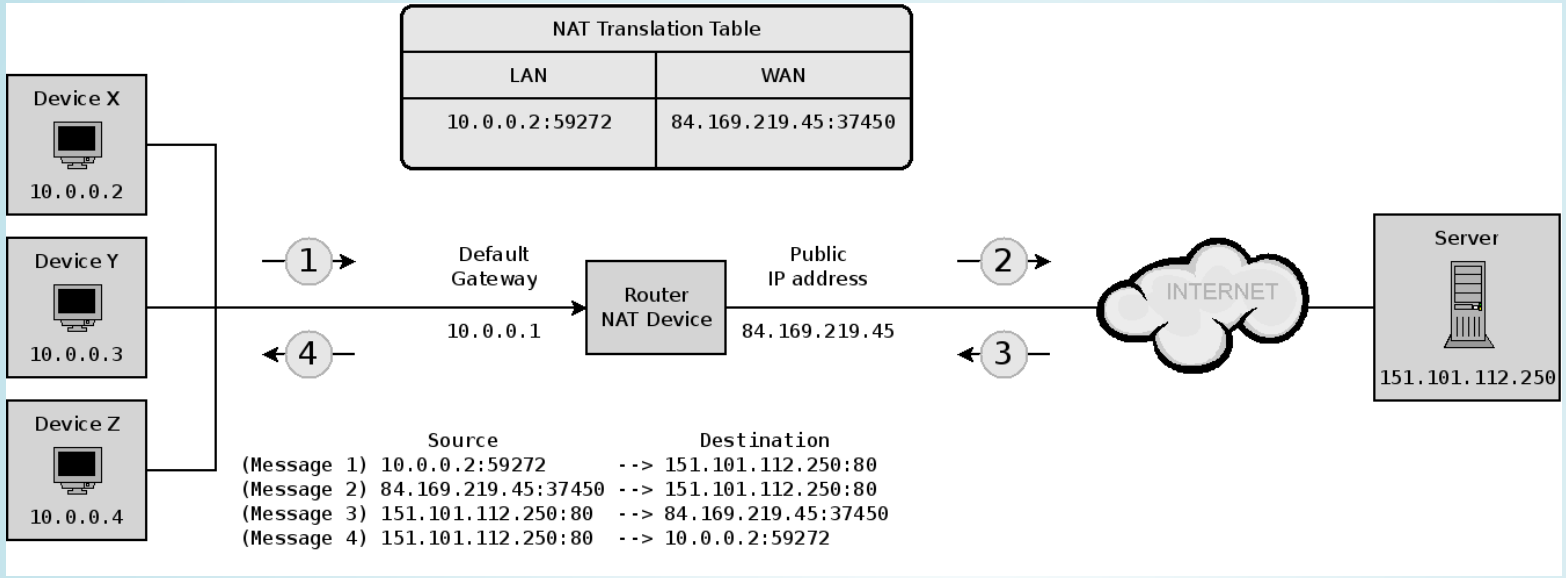

FRAGMENTATION

PACKET FRAGMENTATION

- If a network device does not receive all fragments of an IP packet within a certain period of time (a few seconds), the network device discards all received fragments
- Routers can split IP packets into smaller fragments, if the MTU makes this necessary and it is not prohibited in the packets
- But only the receiver can assemble fragments, none of the routers along the path

IPV6 NETWORKS

A "NEW" INTERNET PROTOCOL

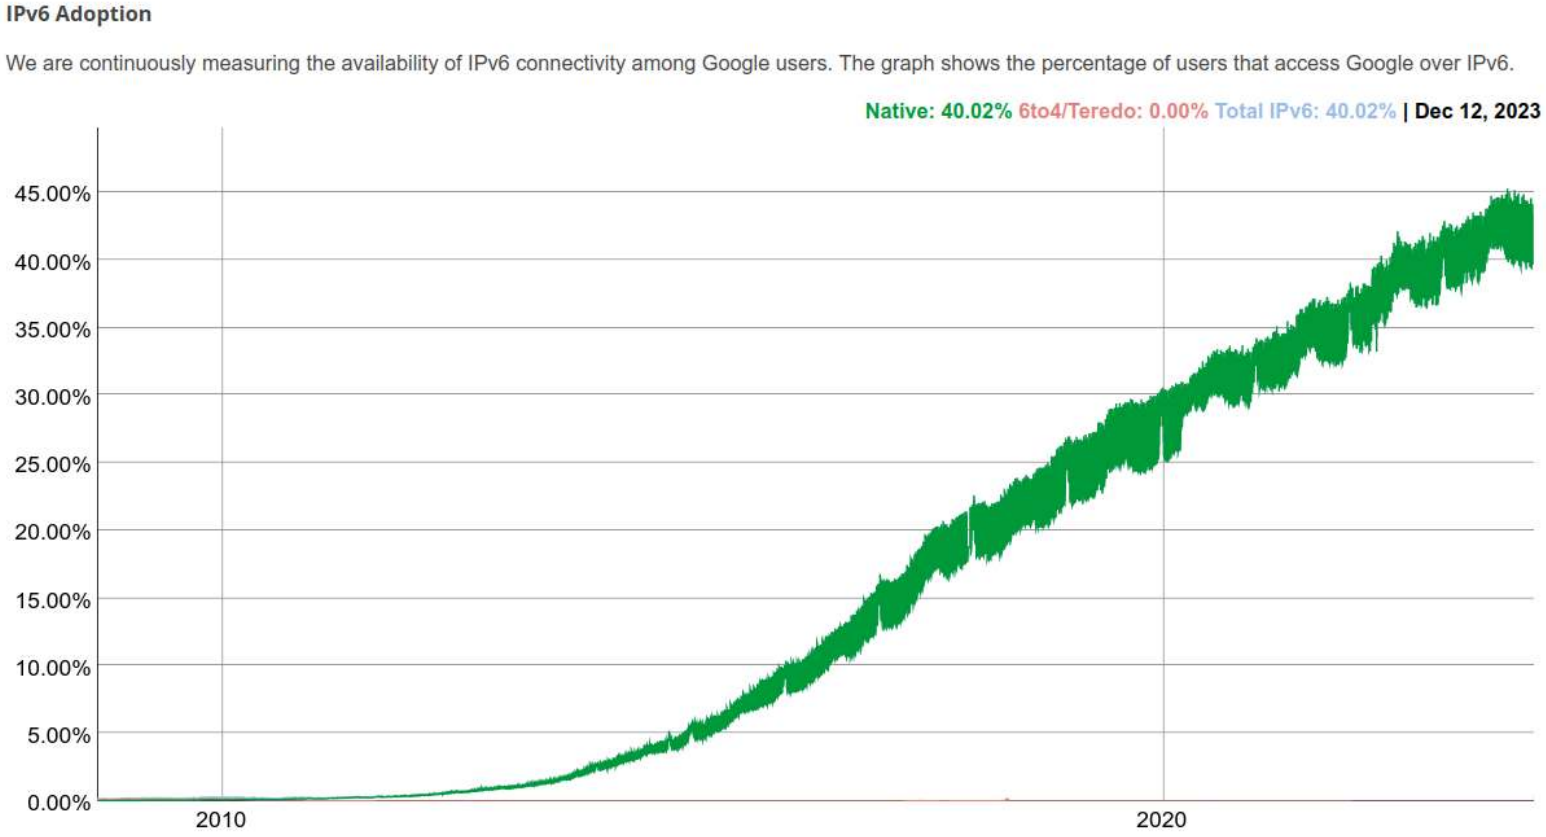

Source: Google IPv6 Stats

IPV6 IMPROVEMENTS

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

- Simplified format
  - Lean header with a fixed size plus optional next headers with a standardized format
  - No checksum, no fragmentation
- Improved Support for mobile applications
  - Improved support for multicast and anycast
  - Support for mobile devices

REPRESENTATION OF IPV6 ADDRESSES

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Notation of IPv6 addresses (URLs)

- IPv6 addresses are enclosed in square brackets
- Port numbers are appended outside the brackets
  `http://[2001:500:1::803f:235]:8080/`
- This prevents the port number from being interpreted as part of the IPv6 address

STRUCTURE OF IPV6 ADDRESSES

IPv6 addresses consist of two parts

64 Bits	64 Bits
Network Prefix	Interface Identifier
2001:638:208:ef34	:0:ff:fe00:65

1. Prefix (Network Prefix)
   - Identifies the network
2. Interface identifier (Interface ID)
   - Identifies a network device in a network
   - Can be automatically computed, manually set, or assigned via DHCPv6
   - If the **interface identifier** is computed the MAC address may be used:
     - **EUI-48 MAC addresses** are first converted into into a 64-bit address $\Longrightarrow$ modified EUI-64 address format

What might be problematical about using the MAC address as part of a (global) IP address?

PRIVACY

- Using the MAC address (even in a modified) form as part of the IP address makes a host globally identifiable
- In order to prevent this, the IETF has proposed privacy extensions
  - RFC 4941 describes a mechanism where the **interface identifier** is changing over time
  - RFC 7217 describes a mechanism where the **interface identifier** is derived from a stable secret

IPV6 ADDRESS TYPES

Described in RFC 4291

- Unicast
  - `fc00::/7 (1111 110)` $\Longrightarrow$ Unique local address, may be routed only in private networks.
  - `fe80::/10 (1111 1110 10)` $\Longrightarrow$ Link local addresses, may not be routed.
  - `::1/128 (0000..1)` $\Longrightarrow$ Loopback address
  - `::/128 (0000..0)}` $\Longrightarrow$ Unspecified
  - Anything else $\Longrightarrow$ Global unicast address (GUA), e.g., `2000::/3` (`2000...` until `3fff...`)
- Multicast
  - $\Longrightarrow$ Multicast addresses.
    (No explicit broadcast addresses, but multicast groups for *all nodes* (`ff01::1` and `ff02::1`) and *all routers* (`ff01::2`, `ff02::2` and `ff05::2`).
- Anycast

  $\Longrightarrow$ from Unicast address range

STRUCTURE OF IPV6 NETWORKS

- IPv6 networks are specified in CIDR notation
  - The address of a single device sometimes has `/128` attached
  - An example is the loopback address of IPv6: `::1/128`
    - All bits – except the last one – have value `0`
      (For IPv4, the loopback address is: `127.0.0.1`)
  - Internet Providers (ISPs) or operators of large networks get the first 32 or 48 bits assigned from a Regional Internet Registry (RIR)
    - The ISPs or network operators split this address space into subnets
    - End users usually get a `/64` or even a `/56` network assigned

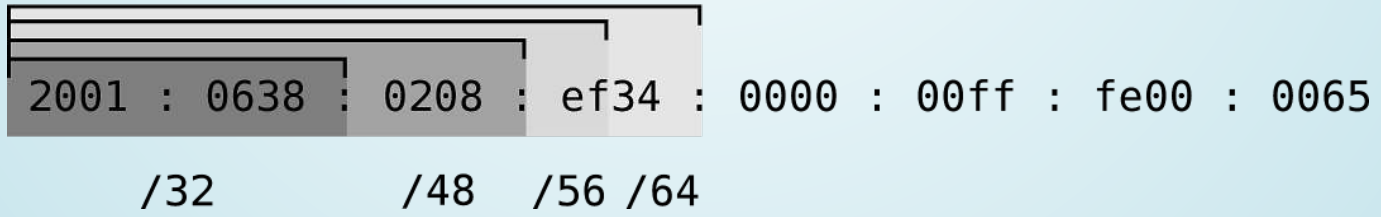

EMBEDDING IPV4 ADDRESSES INTO IPV6 (*IPV4 MAPPED*)

- The IPv4 address may be represented in hexadecimal or decimal notation

Example

IPv4 address:	`131.246.107.35`
IPv6 address:	`0:0:0:0:0:FFFF:83F6:6B23`
Shorter notation:	`::FFFF:83F6:6B23`

PACKET STRUCTURE

HOW TO DESIGN THE PACKET FORMAT?

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Which information do you expect in the packet format?

IPV4 PACKET STRUCTURE

STRUCTURE OF IPV4 PACKETS

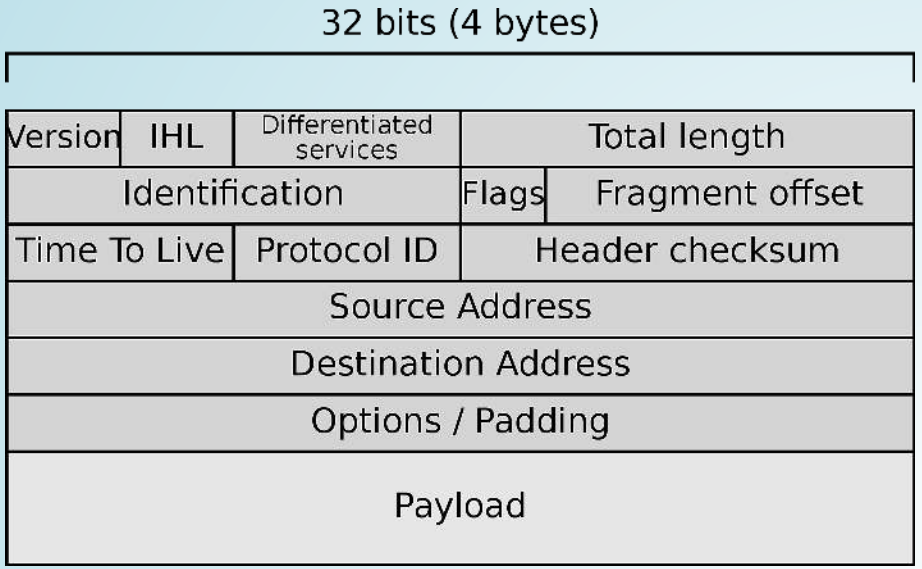

- The payload field contains the data from the Transport Layer

IPV6 PACKET STRUCTURE

WHAT IS DIFFERENT IN IPV6?

Do you expect the packet header to be longer or shorter compared to IPv4?

STRUCTURE OF IPV6 PACKETS: DESIGN

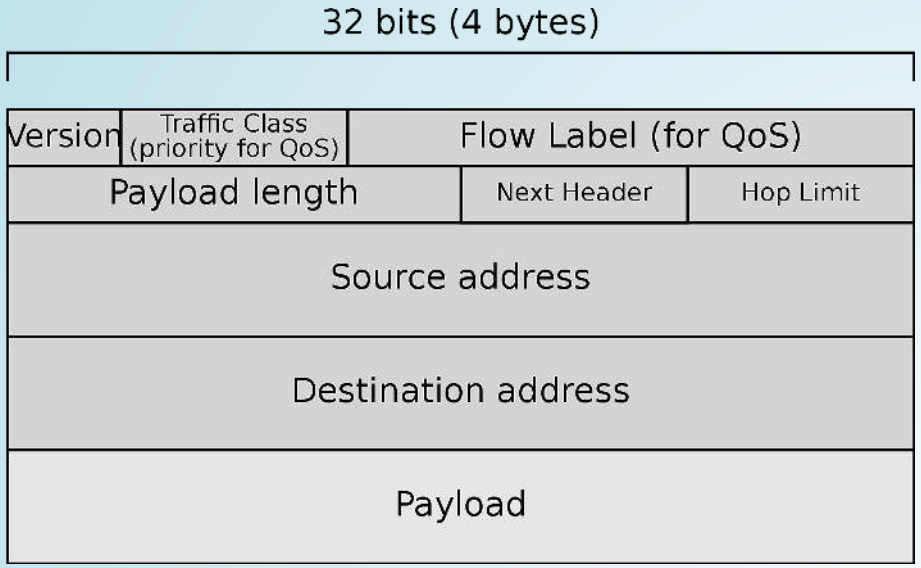

The size of the **IPv6 header** is fixed (320 bits $\Longrightarrow$ **40 bytes**)

- **Simplified package structure**, but simple option to add additional (new) features with a chain of extension headers
- No IHL, fragmentation fields, checksum, options, and padding

STRUCTURE OF IPV6 PACKETS

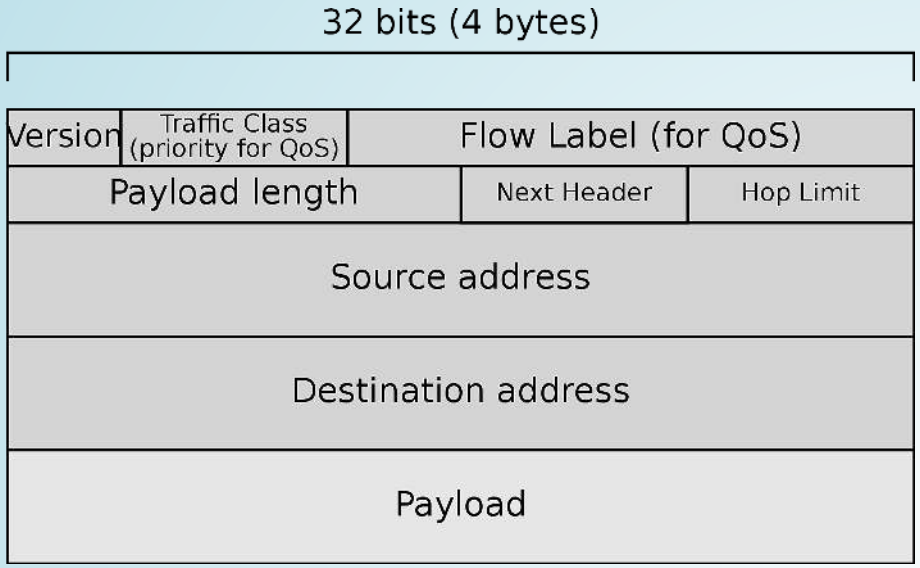

- The hop limit replaces the TTL field of IPv4
- Source and destination addresses keep their meaning
- After the address either the data from the transport layer or an extension header follows

ICMP

WHAT IS MISSING?

- How can we inform a sender about an error?
- How can we verify connectivity?
- How can we find the way a packet takes through the Internet?

THE ROLE OF ICMP

- The Internet Control Message Protocol (ICMP) is used for the exchange of…
  - diagnostic,
  - control, and
  - error messages
- ICMP is a component (*sub-protocol*) of IP
  - but it is treated as a separate protocol
- ICMPv4 is used for IPv4 networks, ICMPv6 is the corresponding protocol for IPv6 networks

USE CASES FOR ICMP

- All routers and terminal devices can handle ICMP
- Typical situations where ICMP is used:
  - A router discards an IP packet, because it does not know how to forward it
  - Not all fragments of an IP packet arrives at the destination
  - The destination of an IP packet cannot be reached, because the Time To Live (TTL) has expired
- ICMP specifies different sorts of messages, which can be send by a router as response to provide diagnostic information
- If an ICMP packet cannot be delivered, no further action is done

The most prominent example The *ping* command uses ICMP messages.

ICMP MESSAGE STRUCTURE

- Th type
- T ype

- The table contains some type-code combinations of ICMP messages

Type	Name of type	Code	Description
0	Echo reply	0	Echo reply (reply for ping)
3	Destination unreachable	0	Destination network unreachable
		1	Destination host unreachable
		2	Destination protocol unreachable
		3	Destination port unreachable
		4	Fragmentation required, but forbidden by the IP packet's flags
		13	Firewall at destination site rejects the IP packet
5	Redirect	0	Redirect Datagram for the Network (or subnet)
		1	Redirect Datagram for the Host
8	Echo Request	0	Echo request (ping)
11	Time Exceeded	0	TTL (Time To Live) expired
		1	Fragment reassembly time exceeded

ICMP Types and Codes

The original set of ICMP type and code values are defined in RFC 792, but multiple have been marked as deprecated in RFC 6633 and RFC 6918. A full list can be found at the IANA.

EXAMPLE OF USING ICMP: traceroute

```
$ traceroute -q 1 wikipedia.de
traceroute to wikipedia.de (134.119.24.29), 30 hops max, 60 byte packets
1  fritz.box (10.0.0.1)  1.834 ms
2  p3e9bf6a1.dip0.t-ipconnect.de (62.155.246.161)  8.975 ms
3  217.5.109.50 (217.5.109.50)  9.804 ms
4  ae0.cr-polaris.fra1.bb.godaddy.com (80.157.204.146)  9.095 ms
5  ae0.fra10-cr-antares.bb.gdinf.net (87.230.115.1)  11.711 ms
6  ae2.cgn1-cr-nashira.bb.gdinf.net (87.230.114.4)  13.878 ms
7  ae0.100.sr-jake.cgn1.dcnet-emea.godaddy.com (87.230.114.222)  13.551 ms
8  wikipedia.de (134.119.24.29)  15.150 ms
```

ADDRESS AUTOCONFIGURATION

AVOID MANUAL CONFIGURATION

- Why do we want avoid manual configuration?
- In which cases is manual (static) configuration preferable?
- How can we do automatic assignment?

REVERSE ADDRESS RESOLUTION PROTOCOL (RARP.JPG)

- Upon booting a network interface has no IP address assigned
- Manual address configuration is not desirable in many scenarios
- With the help of Reverse ARP, well-known hardware addresses are assigned to IP addresses, and recorded on a RARP server
- **Problem**: RARP requests are not passed on by routers, therefore a RARP server must be set up in each local network

RARP is obsolete. Replaced by DHCP (more modern and feature-rich).

DYNAMIC HOST CONFIGURATION PROTOCOL

- A host that needs … R packet
- A DHCP server rep… which contains an IPv4 address
- Additionally it may … outer, DNS server…
  → DHCP can be u…
- The assigned addr… st be renewed after expiration)
- In each subnet a D… uch a message on to the DHCP serve…

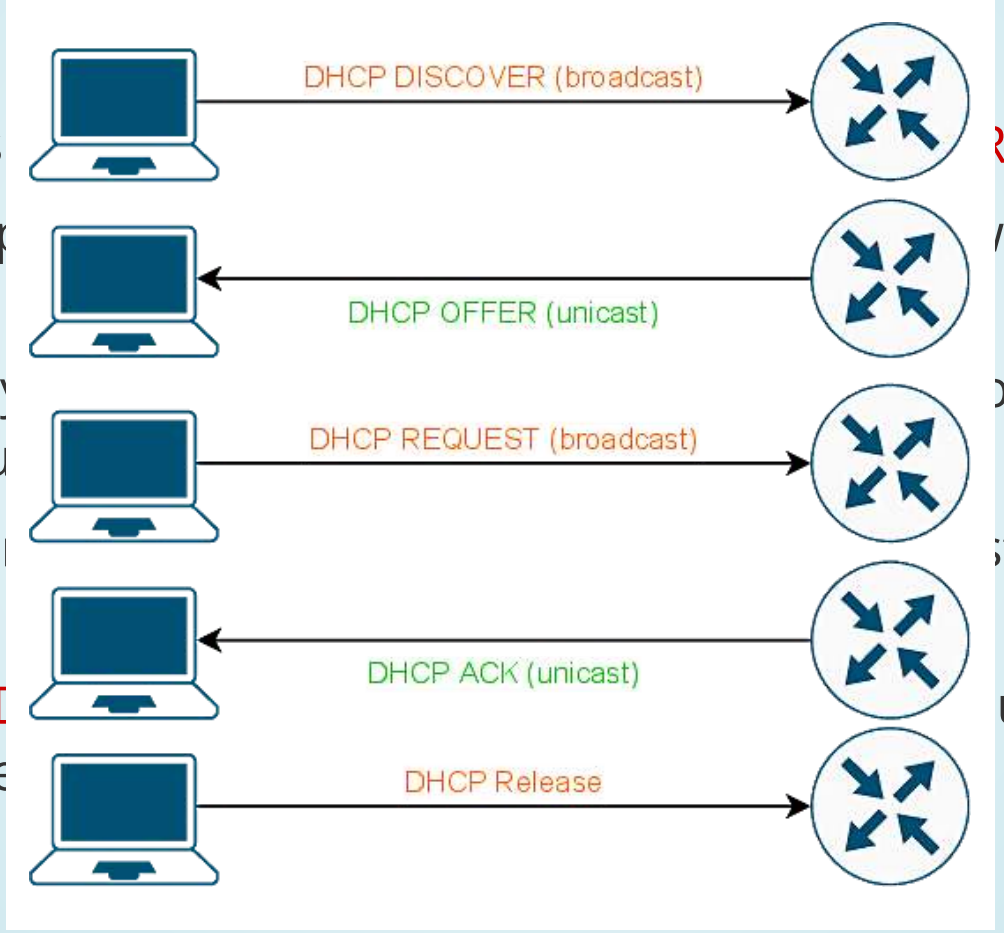

LINK-LOCAL ADDRESSES

- Link-local addresses are valid inside a local physical network
- IPv4 uses the prefix `169.254.0.0/16`, IPv6 uses the prefix `fe80::/10` for link-local addresses
- Are not guaranteed to be unique beyond their network segment, i.e., not globally routable
- In IPv4 the host ID is initially randomized, in IPv6 it can be derived from the MAC address
- A mechanism for Duplicate Address Detection (DAD) is mandatory [1]
- A link-local address can serve as a temporary solution until a globally routable or private address becomes available

1. In IPv4 ARP can be used for this purpose

STATELESS AUTO ADDRESS CONFIGURATION (SLAAC)

- SLAAC is specified for **IPv6** in RFC 2462
- Functioning of **SLAAC**
  - A host generates a tentative link-local address
  - DAD: The host sends a Neighbor Solicitation (NS) with the chosen IP address as destination address
  - If no host responds to the NS with an Neighbor Advertisement (NA) it keeps this address
  - Router solicitations (RS) or Router Advertisements (RAs) are used to find the responsible router for the network
  - The RA contains the network prefix which is used to determine a routable IP address

SUMMARY

You should now be able to answer the following questions:

- Why do we need logical addresses?
- How does an IPv4 address look like and which information does it contain?
- What is a subnet?
- Why do we need a new Internet Protocol?
- What happens in NAT network?
- What is the purpose of ICMP?
- How can IP address be configured automatically?