Network Layer

After examining the transport layer, we can move down to the **network layer**. The transport layer provides logical communication between application processes, while the network layer is responsible for addressing and routing packets between hosts and different networks through the most appropriate available path.

It is important to specify that, in its primary operation, a router mainly uses the first three layers of the network stack:

- Physical layer
- Link layer
- Network layer

A router does not normally need to interpret the application content of the communication, such as a web page, an email or a file. Its task is more limited: it receives a packet through one network interface, reads the network-layer information—particularly the destination address—and decides which interface should be used to forward it.

To perform this task, a router requires:

- The **physical layer**, to transmit and receive signals through the communication medium.
- The **link layer**, to communicate with devices directly connected to the same local network.
- The **network layer**, to read addresses and select a path towards the destination.

The upper layers—the transport and application layers—are mainly required by end hosts because this is where the communicating programs are located, such as browsers, web servers, messaging applications and email clients.

Note: Some modern routers can also inspect information from higher layers, such as port numbers and application protocols, for security and traffic-management purposes.

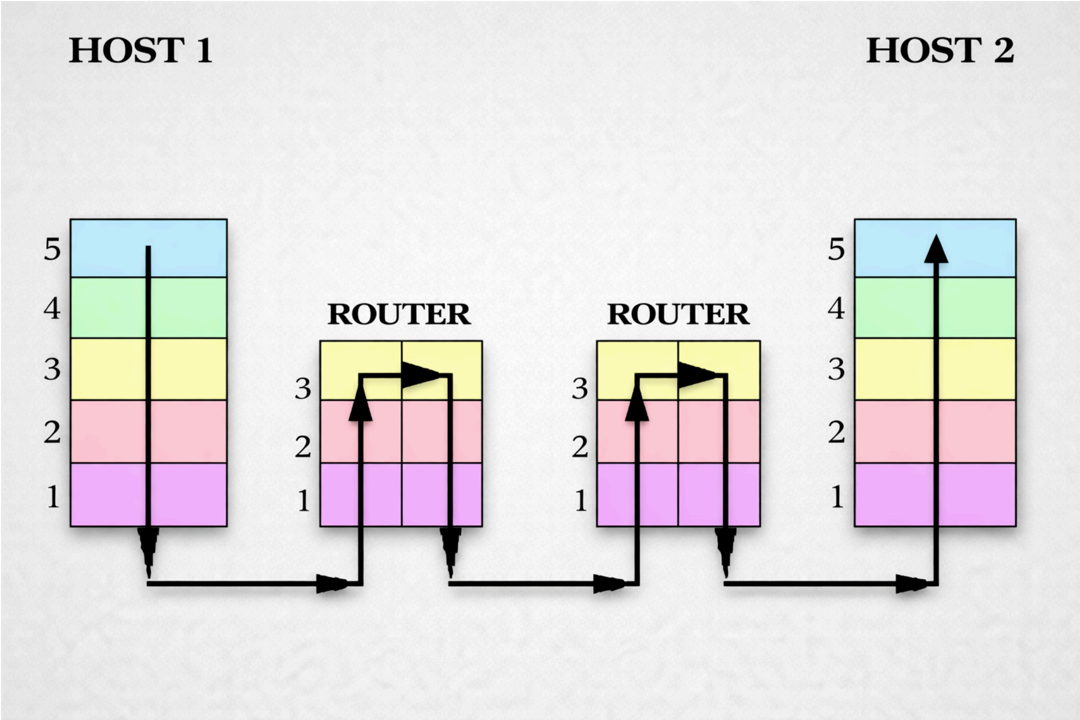

The Internet is based on a connectionless and unreliable protocol called the **Internet Protocol**, abbreviated as **IP**.

IP provides a **best-effort packet service**. This means that it does not guarantee packet delivery, delivery time, arrival order or a particular quality of service. When these guarantees are required, they are provided by higher layers, particularly the transport layer.

The fact that IP is best effort does not mean that it is useless or poorly designed. It means that IP performs a specific task: it attempts to deliver packets to their destination. Reliability, reordering and retransmission can be handled by other protocols, such as TCP, or directly by the application.

The main tasks of IP are **addressing and routing between heterogeneous networks**. Within local networks, communication instead uses local addressing mechanisms, usually based on MAC addresses.

For different networks to communicate with one another, they need a shared addressing system. This system is provided by the IP protocol.

Every interface connected to an IP network must have an **IP address**, meaning a logical identifier used to reach it through the network.

The IP address identifies the sender and recipient of a packet even when they are located on different networks.

An IP address therefore has a different role from a MAC address:

- A **MAC address** is used for local communication within the same physical or local network.

- An **IP address** is used to route packets between different networks.

When a packet must reach a distant destination, it passes through several routers. Each router reads the destination IP address and decides which network or next router should receive the packet.

This procedure is called **routing**. Routing is the collection of rules and mechanisms that allows packets to find a path through different networks until they reach their destination.

The data unit created by the IP protocol is called an **IP packet**.

Example
When a computer on a home network needs to reach a web server on the Internet, the IP packet contains both the sender's IP address and the recipient's IP address.

The home router forwards the packet to the ISP's network. From there, other routers read the destination IP address and decide, step by step, which path the packet should follow.

The Internet is often described as a **TCP/IP network** because it is primarily based on these two fundamental protocols:

- **IP** addresses and routes packets from one device to another through the network.
- **TCP** makes communication reliable by checking that data arrives correctly and in the proper order.

The term **TCP/IP** therefore refers to the collection of protocols that allows computers to communicate through the Internet and many modern local networks.

IP Address

An **IP address**, short for **Internet Protocol address**, is a numerical identifier assigned to a network interface connected to an IP network.

It identifies the interface and allows data to be routed correctly to and from the device.

There are two main versions of IP addresses:

- **IPv4 — Internet Protocol version 4**
- **IPv6 — Internet Protocol version 6**

Historically, IPv4 has been the most widely used version. It uses 32-bit addresses. As the Internet grew, the available IPv4 address space was progressively exhausted.

IPv6 was introduced to address this problem. It uses 128-bit addresses and therefore provides an enormously larger address space.

It is important to specify that an IP address is assigned to a network interface rather than directly to the host itself.

For example:

- A host using only one network interface, such as Wi-Fi, will normally have one IP address for that interface.
- A router connecting multiple networks has multiple network interfaces and therefore multiple IP addresses. A home router, for example, usually has at least one interface towards the local network and another towards the ISP's network.

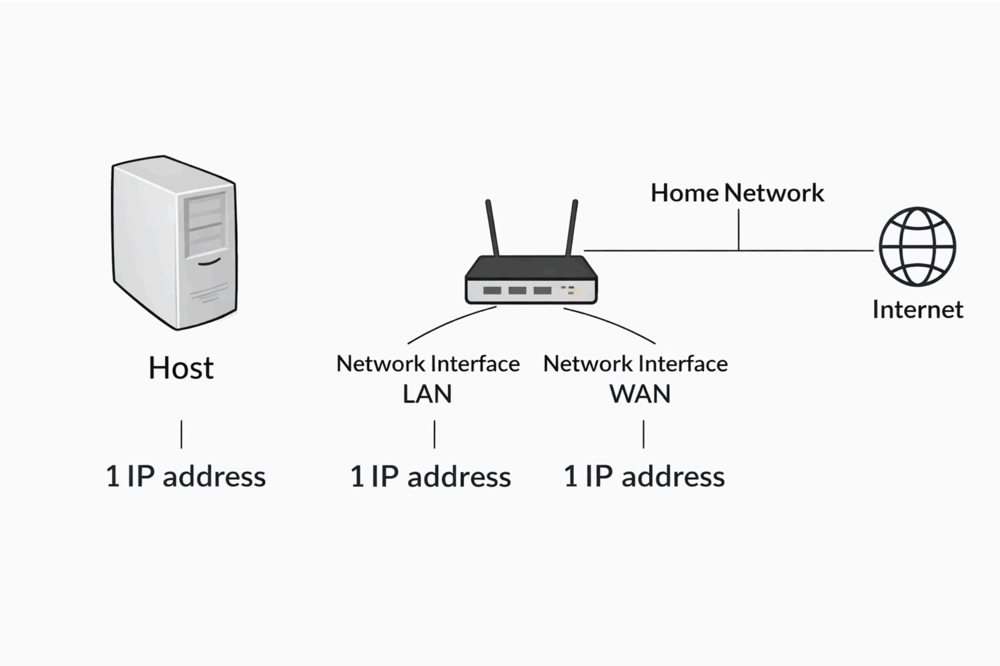

Why are IP addresses necessary when MAC addresses already exist?

IP addresses are required for global routing, while MAC addresses are generally limited to local networks.

The main reasons are:

- **Scalability and organisation:** IP addresses have a hierarchical structure, making them suitable for organising and routing traffic across very large networks such as the Internet. MAC addresses are usually assigned to individual interfaces by manufacturers and do not provide the same hierarchical structure.
- **Independence from the physical network:** IP can operate over different technologies, including Ethernet, Wi-Fi, optical fibre and mobile networks.
- **Flexibility:** IP addresses can be assigned dynamically, simplifying the management of networks where devices frequently join and leave. MAC addresses are generally assigned more permanently to network interfaces, although they can sometimes be changed or randomised.

IP addresses can be:

- Static
- Dynamic

Static IP Address

A **static IP address** is configured permanently on a device or assigned on a fixed basis by a network administrator or ISP.

This means that the interface normally uses the same IP address whenever it connects to the network.

Static IP addresses are often used for servers, critical network equipment and devices that must always be reachable through the same address.

Dynamic IP Address

A **dynamic IP address** is assigned temporarily to a device when it connects to a network.

Dynamic addresses are commonly managed by a **DHCP server**, using the **Dynamic Host Configuration Protocol**. The server automatically assigns IP addresses from a pool of available addresses.

The address may change when the device reconnects to the network or after a particular period of time.

Dynamic IP addresses are often used for home devices and ordinary users because they allow the available addresses to be used more efficiently. When a device is no longer connected, its address can return to the pool and be assigned to another device.

Some websites allow users to view their current public IP address. When the displayed address changes on different days, it is usually because the ISP has assigned a dynamic public IP address.

Public IP addresses are allocated in blocks by **Regional Internet Registries**, or **RIRs**. ISPs receive address blocks and distribute addresses to their customers, including businesses, organisations and home users.

To optimise the use of the available addresses, ISPs often assign dynamic IP addresses.

IPV4

An **IPv4 address** consists of 32 bits, or four bytes, divided into four groups of 8 bits. The groups are separated by dots.

For example:

`11001001.00100100.10101111.00001111`

Each byte can be converted into decimal form to make the address easier to read.

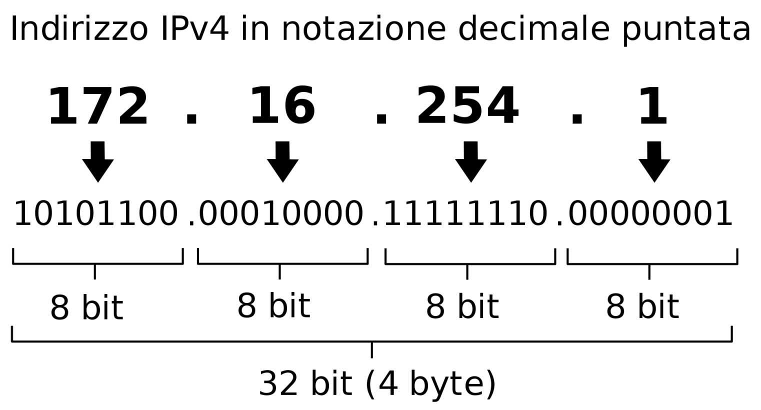

An 8-bit byte can represent values between 0 and 255.

An IPv4 address is therefore normally written as four decimal numbers separated by dots, for example:

`192.168.1.10`

Because IPv4 addresses contain 32 bits, the theoretical address space contains:

2^{32} addresses, approximately 4.3 billion combinations.

This initially seemed like an enormous number. However, as the Internet grew, the available public IPv4 addresses became insufficient.

Today, each person may use many connected devices, including smartphones, computers, tablets, smart televisions and Internet of Things devices.

Private addresses and address-translation techniques were introduced to reduce the consumption of public addresses, while IPv6 was developed to provide a much larger address space.

Private and Public IP Addresses

A **private IP address** is assigned to a device within a local network, such as a home or business network. It is not directly routable through the public Internet.

A **public IP address** is routable through the Internet. It can identify a router or network interface within the global network and may be reached by other Internet hosts when firewall rules and network configurations permit it.

Private IP addresses operate within local networks.

When a device needs to communicate with the Internet, the router can translate its private IP address into a public IP address.

This allows multiple devices on the same local network to share a single public IP address instead of requiring a separate public address for every device.

Using private IP addresses provides the following advantages:

- **Saving public IP addresses:** Public IPv4 addresses are limited. Using private addresses within a LAN reduces the number of public addresses required for a single network. This is particularly useful when many devices communicate mainly within the local network and do not need to be reached directly from the Internet.
- **Isolation:** Private addresses also provide a degree of isolation because internal devices are not directly routable from the public Internet.

As a result, an ISP will often assign a single public IPv4 address to a customer even when dozens of devices are connected to that customer's network.

To summarise:

- The customer receives a public IP address from the ISP, which is often dynamic.
- The devices within the customer's network use private IP addresses. When they communicate with the Internet, the router translates their communications so that they use the public IP address.

Example

Imagine a home network containing several devices, including computers, smartphones, printers and smart televisions. All these devices connect to a Wi-Fi router to access the Internet.

The router assigns a private IP address to each device, for example:

- `192.168.1.2`
- `192.168.1.3`
- `192.168.1.4`

When one of these devices accesses the Internet, the router translates the device's private address into the public address assigned by the ISP.

This public IP address is the address seen by Internet servers when the device sends requests.

A practical example:

- A computer on the home network attempts to access a website.
- The computer uses its private IP address, such as `192.168.1.2`, to communicate with the router.

- The router translates that private address into the public address assigned by the ISP, such as `203.45.67.89`.
- The request reaches the web server with the public source address `203.45.67.89`.
- The server sends its response to the router, which uses its translation information to forward the response to the correct private address inside the LAN.

How does the router translate a private address into a public address and perform the reverse operation?

NAT, or **Network Address Translation**, is a technique that allows multiple devices within a LAN to share a public IP address when connecting to the Internet.

When a device in the LAN sends a request using its private IP address, the router performs the following operations:

- It replaces the private source IP address with the router's public IP address because private addresses cannot be routed through the public Internet.
- It may also replace the source port with a temporary port chosen by the router to distinguish between different communications.
- It stores the translation in a **NAT table**, associating the new public port with the original private IP address and port.
- It forwards the packet to the destination server using the public address and the newly selected port.

When the server replies, it sends the packet to the router's public IP address and the port selected by NAT.

The router then:

- Consults its NAT table to determine which internal host initiated the communication.
- Replaces the public destination address and port with the internal host's original private address and port.
- Forwards the response to the correct device within the LAN.

Note: In everyday language, this process is often simply called NAT. More precisely, the version that also translates port numbers is called **PAT**, or **Port Address Translation**, and is also known as **NAPT**. This is the technique commonly used in home networks.

Example

Imagine the following situation:

- Your computer has the private IP address `192.168.1.100`.
- You want to access a website with the public address `93.184.216.34` on port `80`.
- Your router has the public IP address `203.0.113.1`.

1. Sending the Request

Your computer sends a packet to `93.184.216.34` on port `80`.

Because the computer cannot communicate through the public Internet using a private source address, the router modifies the packet through NAT/PAT:

- It changes the source from `192.168.1.100:54321` to `203.0.113.1:60001`.
- Port `54321` was the source port selected by the computer.
- Port `60001` is a temporary public port selected by the router.

The router stores the translation in its NAT table:

Public port	Internal IP address	Internal port
60001	192.168.1.100	54321

2. Receiving the Response

- The server replies to `203.0.113.1:60001`.
- The router checks its NAT table and finds that port `60001` corresponds to `192.168.1.100:54321`.
- It changes the packet's destination from `203.0.113.1:60001` to `192.168.1.100:54321`.
- It forwards the packet to your computer.

IPV6

IPv6 is a newer version of the Internet Protocol, introduced primarily to solve the shortage of IPv4 addresses.

The most evident difference is the size of the address space:

- IPv4 uses 32-bit addresses.
- IPv6 uses 128-bit addresses.

This difference is also visible in the address format:

- IPv4 address: `203.0.120.195`
- IPv6 address: `2001:0620:0000:0000:0211:24FF:FE80:C12C`

The 128 bits of an IPv6 address are written in hexadecimal form to improve readability.

The theoretical address spaces are:

- **IPv4 address space:** 2^{32} addresses.
- **IPv6 address space:** 2^{128} addresses.

The difference becomes clearer through comparison. IPv4 provides approximately 4.3 billion addresses, which is not even enough to assign one public address to every connected device used by every person on Earth.

IPv6 provides such a vast address space that, theoretically, an enormous number of unique addresses could be assigned to every person and object on the planet.

The 128 bits of an IPv6 address are divided into eight blocks of 16 bits.

Each 16-bit block is written using four hexadecimal characters, and colons are used to separate the blocks.

Example
2001:0620:0000:0000:0211:24FF:FE80:C12C

IPv6 addresses can be shortened to make them easier to write:

- Leading zeroes within each block can be omitted. For example, `0620` becomes `620`, while `0211` becomes `211`.
- A block containing only zeroes, such as `0000`, can be written as `0`.

Example
2001:0620:0000:0000:0211:24FF:FE80:C12C
becomes:
2001:620:0:0:211:24FF:FE80:C12C

IPv6 addresses can be compressed further. One consecutive sequence of all-zero blocks may be replaced with the symbol `::`.

The `::` compression can be used only once within an IPv6 address. If it were used more than once, it would be impossible to determine how many zero blocks had been omitted from each location.

Esempio
2001:0620:0000:0000:0211:24FF:FE80:C12C
becomes:
2001:620:0:0:211:24FF:FE80:C12C
and can then become:
2001:620::211:24FF:FE80:C12C

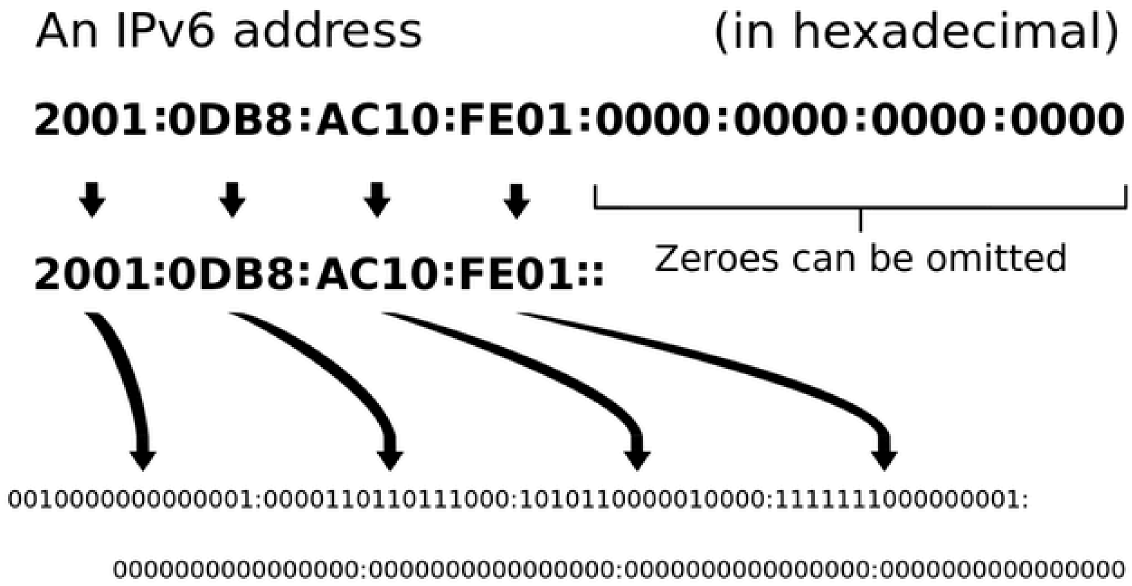

Because IPv6 provides an enormous number of available addresses, NAT is less necessary for allowing multiple devices to share a single public address.

One of the objectives of IPv6 is to allow many devices to have globally unique addresses.

IPv6 also provides addresses intended for internal use within local networks. These are called **Unique Local Addresses**, or **ULAs**.

ULAs are not intended to be routed through the global Internet.

They allow devices to use stable IPv6 addresses internally without making those addresses globally routable. For example, a printer may need to communicate with devices on its local LAN without needing to communicate directly with the entire Internet.

Although IPv6 solves many of the limitations of IPv4, IPv4 remains widely used.

The main reason is that the Internet was built and developed around IPv4 for many years. A very large number of networks, devices, servers, programs and configurations continue to support it.

A complete transition to IPv6 requires time, technical upgrades and compatibility between different systems.

IPv4 and IPv6 therefore coexist today. Many networks continue to use IPv4, often together with NAT and private addresses, while IPv6 is progressively introduced to support the continued growth and modernisation of the Internet.

Routing

After examining IP addresses, we can consider how a packet reaches a distant network.

Routing is the process of determining the best available path through which data should travel from a source to a destination, either within a network or between different networks.

The selected route is based on the information available to the routers, the configured rules and the routing metrics being used.

Routers maintain **routing tables** containing information about reachable networks and the next step required to reach them.

This next step may be:

- Another router, called the **next hop**.
- A network directly connected to the router.

Routes can be learned and updated dynamically by routers in response to changes in network topology or network conditions.

This allows routing tables to adapt when the network changes.

Once routing has determined the best available path, **forwarding** is the concrete operation performed for each individual packet.

The router:

1. Receives the packet.
2. Reads its destination IP address.
3. Consults its routing table.
4. Forwards the packet through the correct outgoing interface.

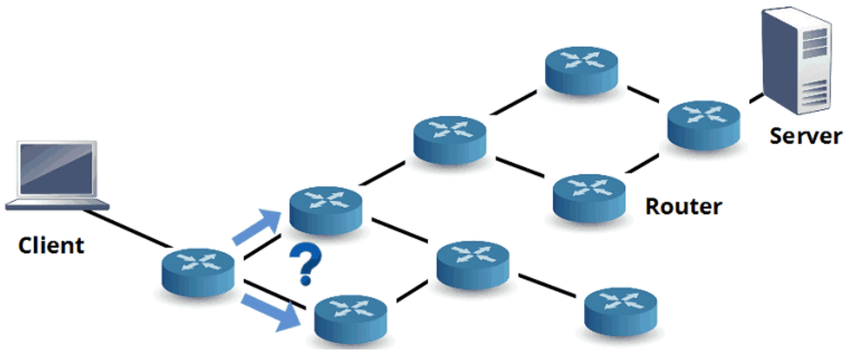

The forwarding process works as follows:

- **Packet reception:** The router receives a packet through one of its network interfaces. The packet may come from a directly connected network or from another router.
- **Destination-address analysis:** The router examines the destination IP address contained in the packet header. This identifies the packet's final destination.
- **Routing-table lookup:** The router consults its routing table, which contains information about reachable networks and the routes used to reach them. It searches for an entry matching the destination address.
- **Route selection:** Three main situations are possible:

- If the router finds a specific matching entry, it selects the best route according to the information contained in the table.
- If no specific route is found, the router checks whether a **default route** has been configured. If so, the packet is forwarded through that route, often towards a router connected to a larger network such as the ISP's network.
- If the router has no matching route and no default route, the packet is discarded.

- **Packet forwarding:** The router forwards the packet along the selected route. This may involve sending it to another router or directly to the destination host when the destination is located on a directly connected network.

Routing tables may be configured manually by a network administrator or updated automatically through **routing protocols**.

Routing protocols allow routers to exchange information about reachable networks and adapt to changes in the network.