

Network layer

Written by [Febogamedev](#)

After examining the transport layer, we can move down to the network layer. The transport layer handles logical communication between application processes, while the network layer addresses and routes packets between hosts and different networks along the most appropriate network path (Baker, 1995; Postel, 1981).

It is important to note that, in its primary operation, a router mainly uses the first three layers of the network stack: the physical, link, and network layers.


This is because the router does not 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 a network interface, reads the network-layer information—particularly the destination address—and decides through which interface to forward it.

To perform this task, the router needs:

- 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 the path toward the destination.

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

Note: some modern routers can also inspect higher-layer information, such as ports and application protocols, for security and management purposes.



The Internet is fundamentally based on an unreliable, connectionless protocol: **the Internet Protocol**, abbreviated as **IP** (Postel, 1981).

It is a **best-effort packet protocol**. This means that the service provides no guarantees regarding delivery, the time taken between transmission and reception, or the quality and integrity of the service during transmission. As we have seen, these quality attributes are implemented at 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: attempting to deliver packets to their destination. When reliability, reordering, or retransmission is required, these functions are handled by other protocols, such as TCP, or directly by the application (Postel, 1981).

As stated above, the main task of IP is addressing and routing between heterogeneous subnetworks, which use their own local addressing system, typically based on MAC addresses (Postel, 1981).

Different networks require a common addressing system in order to communicate with one another. This system is provided by the **IP protocol**.

Every device connected to an **IP network**—a network that uses the IP protocol—must have an IP address, which is a logical identifier used to reach it over the network (Hinden & Deering, 2006; Postel, 1981).

An IP address identifies the sender and receiver 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, by contrast, 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 set of rules and mechanisms that allows packets to find a path across different networks to their destination (Baker, 1995).

The packet 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 the sender's IP address and the destination IP address.

The home router forwards the packet to the operator'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 also called a **TCP/IP network** because it is mainly based on two fundamental protocols, TCP and IP (Deering & Hinden, 2017; Postel, 1981):

- *IP* handles packet addressing and routes packets from one device to another across the network.
- *TCP* makes communication reliable by ensuring that data arrives correctly and in the proper order.

The name TCP/IP therefore refers to the suite of protocols that allows computers to communicate over the Internet and many modern local networks.

IP address

An **IP address**, short for Internet Protocol address, is a unique number assigned to each device connected to the Internet (Hinden & Deering, 2006; Postel, 1981).

This address uniquely identifies a network interface on the Internet and enables data to be routed correctly to and from the device.

IP addresses have two main versions:


- **IPv4 (Version 4).**
- **IPv6 (Version 6).**

Historically, IPv4 has been the most widely used version and uses 32-bit addresses. As the Internet grew, the IPv4 address space was progressively exhausted (Internet Corporation for Assigned Names and Numbers [ICANN], 2011).

IPv6 was introduced for this reason. It uses 128-bit addresses and provides a vastly larger address space (Deering & Hinden, 2017).

It is important to note that an IP address refers not to the host itself but to one of its network interfaces. For example (Hinden & Deering, 2006):

- A host with only one network interface, such as Wi-Fi only, will have a single IP address.
- Routers that connect multiple links have several network interfaces and therefore several IP addresses. A home router usually has at least two: one facing the local network and one facing the operator's network.



Why do IP addresses exist when MAC addresses already exist?

IP addresses are required for global routing, whereas MAC addresses are limited to the local network.

The main reasons are:

- Scalability and organization: IP addresses are organized hierarchically, making them scalable and suitable for large networks such as the Internet. MAC addresses, by contrast, are assigned to individual devices by manufacturers and do not follow a hierarchical structure, making them less suitable for organizing a large-scale network (Fuller & Li, 2006).
- Independence from the physical network: the IP protocol can operate over different technologies, such as Ethernet, Wi-Fi, fiber, and mobile networks.
- Flexibility: IP addresses can be assigned dynamically, simplifying the management of networks in which devices join and leave. MAC addresses are assigned statically by device manufacturers (Droms, 1997).


Public IP addresses are allocated in blocks by **Regional Internet Registries (RIRs)**. ISPs obtain these blocks and distribute the addresses to customers such as businesses, organizations, and home users (Internet Assigned Numbers Authority [IANA], n.d.). To optimize the use of available addresses, ISPs often assign dynamic IP addresses..

A domain name can be associated with one or more *IP addresses*, and similarly, an *IP address* can be associated with one or more domains.

- The simplest relationship is one-to-one, where a domain corresponds to a single IP address, for example `www.example.com` → `192.0.2.10`.
- There is also a one-to-many relationship, where the same domain is associated with multiple IP addresses: this configuration can be used to distribute traffic across multiple servers and improve availability and performance.
- A many-to-one relationship is also possible, where several different domains point to the same IP address. This is very common in hosting services, where a single server hosts many websites. DNS therefore makes it possible to manage flexible associations between domain names and IP addresses, which do not necessarily have a unique one-to-one correspondence.

When a domain is associated with multiple *IP addresses*, *DNS* does not directly route the traffic; instead, it returns one or more available addresses to the device.

The selection can be made in different ways: the *DNS* server can alternate the returned addresses, as in **DNS round-robin**, where the server changes the order of the returned IP addresses cyclically to distribute requests, or it can select them according to criteria such as the user's geographic location, server load, or server availability. Once the *DNS* response has been received, the user's device establishes a connection to one of the returned *IP addresses*.



IP addresses can be:

- Static.
- Dynamic.

Static IP Address

A static IP address is permanently configured on a device or assigned on a fixed basis by a network administrator or an Internet Service Provider (Droms, 1997).

This type of address remains fixed, meaning that the device uses the same IP address whenever it connects to the network. Static IP addresses are often used for servers, critical network devices, and devices that must always be reachable through the same address.

Dynamic IP Address

A dynamic IP address is assigned temporarily to a device whenever it connects to the network.

Dynamic IP addresses are managed by a DHCP (Dynamic Host Configuration Protocol) server, which automatically assigns addresses to devices from a pool of available IP addresses (Droms, 1997).

These addresses may change whenever the device reconnects to the network or after a certain period of time.


Dynamic IP addresses are often used for home devices and ordinary users because they allow available addresses to be used more efficiently. When a device is not connected, its address can return to the pool and be assigned to another device.

IPv4

An IPv4 address consists of 32 bits, or four bytes, divided into four groups of 8 bits. Each group is separated by a dot, for example 11001001.00100100.10101111.00001111 (Postel, 1981).

Each of these four bytes can be converted into decimal form to make the address easier to read.

IPv4 address in dotted decimal notation



Each byte ranges from 0 to 255 because there are $2^8 = 256$, which is the number of combinations available with 8 bits.

An IPv4 address is therefore written as four decimal numbers separated by dots, for example 192.168.1.10.

Because IPv4 addresses contain 32 bits, the number of possible combinations is 2^{32} , approximately 4 billion (Postel, 1981).

At first, this number seemed enormous, but as the Internet grew, the available public IPv4 addresses became insufficient.

Today, each person may use many connected devices, such as smartphones, computers, tablets, smart TVs, and IoT devices.

Solutions such as private addresses were introduced to reduce the consumption of public addresses, while IPv6 was developed to provide a much larger address space (Deering & Hinden, 2017; Rekhter et al., 1996).

In-depth

IP addresses are suitable for very large networks such as the Internet because they have a hierarchical structure, which makes it possible to organize devices into networks and subnets.

An IP address, in fact, does not simply identify a device, but also contains information about the network to which it belongs.

This allows routers to manage traffic without having to know every device connected to the Internet individually.

For example, imagine that an operator owns an entire network identified by the prefix 203.0.112.0/20.

Where:

- The /20 part indicates that the first 20 bits of the address identify the network, while the remaining bits can be used to distinguish addresses within it. In this way, the prefix represents not a single address, but an entire block of IP addresses.

Within this network, there may be thousands of devices, each with its own IP address.

An external router, however, does not need to store a separate route for each of them: it can simply know that all addresses belonging to that prefix must be sent toward that operator.

The process can therefore take place progressively:

Internet → ISP → subnet → device

As the packet gets closer to its destination, routers use increasingly specific information.

This principle is known as route aggregation and makes it possible to greatly reduce the amount of information that routers need to store.

MAC addresses work differently.

An address such as: 3C:52:82:A1:43:11

identifies a network interface, but does not indicate which network it is located on. Two devices with similar *MAC addresses* may be located in completely different parts of the world. A *MAC address* therefore does not have a structure that allows routers to easily group devices according to their location within the network.

If the Internet used *MAC addresses* directly for routing, routers would theoretically have to maintain much more detailed information about individual devices and their locations.

With *IP addresses*, by contrast, entire groups of devices can be represented using only a few network prefixes.

A useful analogy is the postal system. A postal address is organized hierarchically:

Italy → Lombardy → Milan → Via Roma → number 15

A national postal sorting center does not need to know the exact location of an individual house: initially, it only needs to know that a letter addressed to Milan must be sent toward Milan.

The subsequent sorting centers will then progressively identify the area, the street, and finally the street number.

A *MAC address*, by contrast, would be more similar to a personal identification code with no geographical information. Knowing the code would make it possible to identify the recipient, but not where the letter should be sent.

For this reason, the two types of addresses have different and complementary functions: an *IP address* makes it possible to locate a destination within a hierarchical network and route packets on a large scale, while a *MAC address* is mainly used to identify network interfaces and deliver packets within the local link.

Private and public address

A private IP address is assigned to a device within a local network, such as a home or business network, and cannot be accessed or routed directly over the Internet (Rekhter et al., 1996).

A public IP address is routable over the Internet. It can identify a device or router on the global network and can be reached by other Internet hosts when firewall rules and configuration permit it.

Private IP addresses operate within local networks.

When a device needs to communicate over the Internet, the router translates its private IP address into a public IP address. This allows several devices on the same local network to share one public IP address instead of requiring a different public address for each device (Srisuresh & Egevang, 2001).

Private addresses therefore provide the following advantages:

- Conservation of public IP addresses: public IP addresses are limited and costly. Using private IP addresses within a local network reduces the need to obtain many public IP addresses for a single subnet. This is particularly useful when many devices are used only locally and do not require a direct Internet connection.
- Security: private addresses also provide a degree of isolation because internal devices are not directly reachable from the Internet.

Thanks to this technique, ISPs usually assign only one public IP address to each customer, even when dozens of devices connect to the Internet.

To summarize:

- As a user, I receive a public IP address—probably a dynamic one—from my ISP.
- My devices instead use several private IP addresses, which are translated into my public IP address when they communicate over the Internet.

Example

Imagine a home network containing several devices, including computers, mobile phones, printers, and smart TVs. All these devices connect to a Wi-Fi router to access the Internet. In this situation, the router assigns each device a private IP address from the private address range, for example 192.168.1.2, 192.168.1.3, and so on.

When one of these devices accesses the Internet, the router translates the device's private IP address into the single public IP address assigned by the ISP. This public IP address is what Internet servers see when the device sends requests.

Practical Example:

- A computer on your home network attempts to access a website.
- The computer uses its private IP address, for example 192.168.1.2, to communicate with the router.
- The router translates the computer's private IP address into the public IP address assigned by the ISP, for example 203.45.67.89.
- The computer's request reaches the website server with the public IP address 203.45.67.89.
- The server sends its response to the router, which uses the stored information to translate the public address back into the correct private address and forward the response to the computer on the local network.

How exactly does the router translate a private IP address into a public one, and vice versa?

Network Address Translation (NAT) is a technique that allows several devices on a LAN to share one public IP address when connecting to the Internet (Srisuresh & Egevang, 2001; Srisuresh & Holdrege, 1999).

When a device on the LAN sends an Internet request using its private IP address, the router's NAT operates as follows (Srisuresh & Egevang, 2001):

- It replaces the private IP address with the router's public IP address because private addresses cannot be routed over the Internet.

- It also replaces the source port with a temporary port selected by the router to distinguish the different connections.
- It stores this information in a NAT table, associating the new public port with the original connection—the host's private IP address and private port.
- It forwards the packet to the destination server using the public IP address and the newly generated port.

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

At this point:

- The router checks the NAT table to determine which internal host initiated the connection.
- It translates the destination address by replacing the public IP address and port with the host's original private IP address and private port.
- It forwards the response to the correct device within the LAN.

Note:

In everyday language, this mechanism is often simply called NAT. More precisely, however, the advanced form described here is called PAT or NAPT, and it is commonly used in home networks (Srisuresh & Holdrege, 1999).

Example

Imagine the following situation:

- You are using a PC with the private IP address 192.168.1.100.
- You want to access a website with the public IP address 93.184.216.34 on port 80.
- Your router has been assigned the public IP address 203.0.113.1.

1. Sending the Request

Your PC sends a packet to 93.184.216.34 on port 80, but it cannot access the Internet using a private IP address. Therefore the router modifies the packet through NAT:

- It changes the sender from 192.168.1.100 with port 54321 to 203.0.113.1 with port 60001. The choice is arbitrary: 54321 was the PC's port, while 60001 is a port selected by the router.
- It stores the translation in a NAT table:

Public Port	Internal IP Address	Internal Port
60001	192.168.1.100	54321

2. Receiving the Response

- The example.com 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 on port 54321.
- It changes the packet destination from 203.0.113.1 on port 60001 to 192.168.1.100 on port 54321.

- It forwards the packet to your PC.

IPv6

IPv6 is a more recent version of the IP protocol, introduced primarily to solve the shortage of IPv4 addresses (Deering & Hinden, 2017).

The most evident difference is the larger address space: IPv4 uses 32-bit addresses, whereas IPv6 uses 128-bit addresses.

The difference is immediately visible in the address format:

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

With 128 bits, written in hexadecimal form for readability, IPv6 can generate far more unique IP addresses than a 32-bit system (Hinden & Deering, 2006).

- IPv4 address space: $32 \text{ bit} = 2^{32}$ di indirizzi $\approx 4,3 \text{ billions of addresses}$
- IPv6 address space: $128 \text{ bit} = 2^{128}$ di indirizzi $\approx 340 \text{ sextillions of addresses}$

The scale of this difference becomes clearer through comparison. IPv4's address space of approximately 4.3 billion addresses is not enough to give every person on Earth a unique address, while a 128-bit system could theoretically assign an IP address to every grain of sand on the planet.

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

Using hexadecimal notation, each 16-bit block can be represented by four digits or letters, with colons used as separators.

Example

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

A shorter notation is commonly used to make IPv6 addresses easier to write (Kawamura & Kawashima, 2010):

- Leading zeros in each block can be omitted. For example, 0620 can become 620, and 0211 can become 211.
- A block consisting entirely of zeros, such as 0000, can be written simply 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 by replacing one consecutive sequence of all-zero hexadecimal blocks with the symbol "::" (Kawamura & Kawashima, 2010).

The :: compression symbol can be used only once within an IPv6 address. If it appeared more than once, it would be impossible to determine how many zero blocks had been omitted at each position (Kawamura & Kawashima, 2010).

Example

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

becomes

2001:620:0:0:211:24FF:FE80:C12C

becomes

2001:620::211:24FF:FE80:C12C

An IPv6 address (in hexadecimal)

2001:0DB8:AC10:FE01:0000:0000:0000:0000

↓ ↓ ↓ ↓ |—————|

2001:0DB8:AC10:FE01:: Zeroes can be omitted

0010000000000001:0000110110111000:1010110000010000:1111111000000001:
0000000000000000:0000000000000000:0000000000000000:0000000000000000

Because IPv6 provides an enormous number of available addresses, NAT becomes less necessary as a method for allowing many devices to share the same public address. The goal of IPv6 is to allow many devices to have a unique global address.

IPv6 also provides addresses intended for internal use on local networks, called **Unique Local Addresses (ULAs)** (Hinden & Haberman, 2005).

These addresses are not intended to be routed over the global Internet.

ULAs provide stable IPv6 addresses that can be used only within an internal network without exposing the devices to the global Internet. For example, a printer needs to communicate with its local LAN, not with the entire Internet.

Although IPv6 solves many of IPv4's limitations, IPv4 addresses are still widely used.

The main reason is that the Internet was created and developed for many years using IPv4, so a large number of networks, devices, servers, software systems, and configurations still support it.

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

For this reason, IPv4 and IPv6 coexist today. Many networks still use IPv4, often with techniques such as NAT and private addresses, while IPv6 is being introduced gradually to support broader and more modern network growth (Deering & Hinden, 2017).

Routing

After examining IP addresses, we can ask how a packet manages to reach a distant network.


Routing is the process of determining the optimal path—the most suitable path according to the available information and configured rules—through which data is directed from a source to a destination within one network or across different networks (Baker, 1995).

Routers maintain **routing tables** containing information about reachable networks and the next step required to reach them. This step may be another router, called the next hop, or a network directly connected to the router (Baker, 1995).

Routers can learn and update routes dynamically in response to changes in network topology or network conditions.

This allows routing tables to adapt to changes in the network.

After routing has determined the optimal route, **forwarding** is the actual operation in which the router takes an individual packet, reads its destination IP address, consults the routing table, and sends it through the correct outgoing interface. This process enables the packet to reach its final destination through the network using the best available path (Baker, 1995).




The forwarding process in a router works as follows:

- Packet reception: forwarding begins when the router receives a packet through one of its network interfaces. The packet may come from a network directly connected to the router or may have been forwarded by another router.
- Destination-address analysis: the router examines the destination IP address contained in the packet header. This address identifies the packet's final receiver.
- Routing-table lookup: the router consults its routing table, which contains information about available networks and the routes used to reach them. It searches for an entry matching the packet's destination address.
- Optimal-route selection: the lookup can produce three possible outcomes:
 - If the router finds a routing-table entry matching the destination address, it selects the optimal route according to the information in that entry.

- If no specific route is found, the router checks whether a **default route** is configured. If one exists, it forwards the packet to the router specified as the default path, often a router connected to a larger network such as the Internet operator's network.
- If the router has no default route and cannot find a specific route for the destination address, the packet is discarded.
- Packet forwarding: the router sends the packet along the selected route. This may involve forwarding it to another router or directly to the destination host when it belongs to a known local network.

Routing tables can be configured manually by a network administrator or updated automatically through **routing protocols**. These protocols allow routers to exchange information about reachable networks and adapt to network changes (Baker, 1995).



Supplementary materials

Key points

- **The network layer** is responsible for addressing and routing packets between hosts, including hosts belonging to different networks, while the transport layer manages logical communication between application processes.
- IRouters operate mainly at the physical, link, and network layers: they receive a packet, read the destination network address, and decide which interface should be used to forward it.
- The Internet uses the **IP protocol**, a connectionless, best-effort protocol: it attempts to deliver packets, but does not guarantee delivery, order, or arrival times.
- **IP** and **MAC** serve different and complementary functions: a **MAC address** is mainly used for local communication, while an **IP address** enables routing between different networks.
- An **IP address** identifies a network interface, not a device in general. A router with multiple interfaces therefore has multiple **IP addresses**.
- **IP addresses** are suitable for the Internet because they have a hierarchical structure. Network prefixes, such as **/20**, make it possible to represent entire blocks of addresses and aggregate routes, preventing routers from having to know every individual device.
- A domain name and an **IP address** do not necessarily have a one-to-one relationship: one domain can be associated with multiple IP addresses, and multiple domains can use the same IP address. **DNS** can also return different addresses to distribute traffic.
- IP addresses can be **static** or **dynamic**. Dynamic addresses are normally assigned automatically through **DHCP**.
- **IPv4** uses 32-bit addresses, normally written as four decimal numbers ranging from 0 to 255, and provides approximately 4.3 billion addresses.
- The shortage of **IPv4** addresses has encouraged the use of private addresses and **NAT**. Devices on a **LAN** can use private **IP addresses** and share a single public **IP address** to communicate with the Internet.
- Home networks normally use **PAT/NAPT**, commonly referred to simply as **NAT**: the router distinguishes connections from different devices by combining **IP addresses** and port numbers and storing these associations in a **NAT table**.
- **IPv6** uses 128-bit addresses, providing an enormously larger address space than **IPv4**. Addresses are written in hexadecimal and can be shortened by removing leading zeros and using **::** for consecutive sequences of zero-valued blocks.
- With **IPv6**, **NAT** becomes less necessary because, in theory, the available address space allows globally unique addresses to be assigned to a very large number of devices. **IPv4** and **IPv6**, however, continue to coexist.
- **Routing** determines the path to be used to reach a destination, while forwarding is the actual operation through which a router sends an individual packet through the correct interface.

- Routers use **routing tables**, which specify reachable networks and their corresponding next hops. If no specific route exists, they can use a **default route**; if no default route exists either, the packet is discarded.
- Routing tables can be configured manually or updated through **routing protocols**, allowing routers to adapt to changes in the network.

Exercises

What does the router have to do?

A computer on a home network needs to send a packet to a web server on the Internet. The packet passes through the home router and then through several routers in the ISP's network.

Answer the following questions:

- Explain what role the network layer performs during this communication.
- Indicate which information a router mainly needs to read in order to decide where to forward the packet.
- Explain why, in its primary operation, the router does not need to interpret the content of the requested web page.
- Indicate which three layers of the stack are mainly used by the router and what function each one performs.
- Explain why the transport and application layers are particularly important in the end hosts.
- Describe what it means that the IP protocol provides a best-effort service.
- If a packet is lost along the path, explain why IP is not necessarily responsible for retransmitting it.
- Finally, explain why the fact that IP does not guarantee delivery does not mean that the protocol is poorly designed.

IP or MAC?

A school has two different LANs connected through a router.

Computer A is located in LAN 1 and needs to communicate with computer B, which is located in LAN 2.

Both computers have a MAC address and an IP address.

Analyze the communication.

- Explain the function of the MAC address.
- Explain the function of the IP address.
- Indicate which of the two addresses is required to route the packet between the two LANs.
- Explain why the MAC address alone would not be suitable for organizing routing on a global network such as the Internet.
- Explain in what sense IP addresses are independent of the physical technology being used, such as Ethernet, Wi-Fi, fiber, or mobile networks.
- The router has one interface connected to LAN 1 and another connected to LAN 2. Explain why it must have at least two IP addresses.

- A student states: "Every computer always has exactly one IP address." Evaluate this statement and correct it using the concept of a network interface.

Assigning addresses on a network

A company has:

- 80 employee computers.
- 30 company smartphones.
- 6 network printers.
- 3 servers that must always be reachable at the same address.
- Some devices that frequently join and leave the network.

The administrator must decide which devices should use static IP addresses and which should use dynamic IP addresses.

Answer the following questions:

- Identify which devices could benefit most from a static IP address and justify your choice.
- Identify which devices could normally use dynamic addresses.
- Explain the function of a DHCP server.
- Describe what can happen to a device's dynamic IP address when the device disconnects from the network.
- Explain why dynamic assignment makes it possible to use the available pool of addresses more efficiently.
- A company server continuously changes its IP address. Explain what practical problem this could cause.
- Propose a general address-assignment strategy for the network described above and justify your choices.

A home with only one public IP address

A family has the following devices:

- A computer with the private IP address 192.168.1.10.
- A smartphone with the private IP address 192.168.1.11.
- A smart TV with the private IP address 192.168.1.12.
- The router has the public IP address 203.0.113.1.

The computer opens a connection to a web server using source port 52341. The router temporarily translates the communication using public port 60001.

Analyze what happens.

- Explain the difference between a private IP address and a public IP address.
- Explain why 192.168.1.10 cannot be used directly for routing on the Internet.
- Describe what the router modifies when the computer sends the packet to the Internet.
- Indicate which information the router must store in its NAT table.
- The server replies to 203.0.113.1:60001. Explain how the router can determine that the response must be forwarded to 192.168.1.10:52341.
- Explain why the smartphone and smart TV can use the same public IP address at the same time without mixing up their communications.

- Explain why the technique described in the handout is more precisely called PAT/NAPT, even though in common usage it is often simply referred to as NAT.
- Indicate two advantages of using private addresses discussed in the handout.

IPv4 and IPv6

Consider the following two addresses:

- IPv4: 172.16.254.1.
- IPv6: 2001:0620:0000:0000:0211:24FF:FE80:C12C.

Answer the following questions:

- Indicate how many bits an IPv4 address uses and how many groups it is divided into.
- Explain why each decimal number in an IPv4 address can have a value between 0 and 255.
- Indicate how many bits IPv6 uses.
- Explain why IPv6 provides an enormously larger address space than IPv4.
- Rewrite the IPv6 address by removing unnecessary leading zeros from each block.
- Compress the address further by correctly using the :: symbol.
- Explain why the :: symbol can be used only once in the same address.
- Explain why IPv6 makes it less necessary to use NAT to allow many devices to share the same public address.
- Explain what Unique Local Addresses (ULAs) are and why they can still be useful in an IPv6 network.
- Explain why IPv4 continues to be widely used despite the existence of IPv6.

Routing or forwarding?

A router receives a packet addressed to a remote network.

Its routing table contains the following information:

- Network A is directly connected.
- Network B is reachable through Router 2.
- Network C is reachable through Router 3.

For all other destinations, a default route is configured through Router 4.

Answer the following questions:

- Explain what is meant by routing.
- Explain what is meant by forwarding.
- Indicate the function of the routing table.
- Explain what is meant by next hop.
- If the packet is addressed to network B, describe what the router must do.
- If the packet is addressed to a network that is not explicitly listed in the table, explain what happens because of the default route.
- Imagine that the default route is also removed and the router does not know any path to the destination. Explain what happens to the packet.
- Explain why routing and forwarding are not the same operation, even though they work together.
- Imagine that the link to Router 2 stops working and that the router uses dynamic routing protocols. Explain how the routing table could adapt to the new situation.

- Finally, reconstruct the entire procedure: packet reception → reading the destination IP address → consulting the routing table → selecting the route → forwarding through the correct interface.

Bibliography

- Baker, F. (1995). *Requirements for IP version 4 routers* (RFC 1812). RFC Editor. <https://doi.org/10.17487/RFC1812>
- Deering, S., & Hinden, R. (2017). *Internet protocol, version 6 (IPv6) specification* (RFC 8200). RFC Editor. <https://doi.org/10.17487/RFC8200>
- Droms, R. (1997). *Dynamic host configuration protocol* (RFC 2131). RFC Editor. <https://doi.org/10.17487/RFC2131>
- Fuller, V., & Li, T. (2006). *Classless inter-domain routing (CIDR): The Internet address assignment and aggregation plan* (RFC 4632). RFC Editor. <https://doi.org/10.17487/RFC4632>
- Hinden, R., & Deering, S. (2006). *IP version 6 addressing architecture* (RFC 4291). RFC Editor. <https://doi.org/10.17487/RFC4291>
- Hinden, R., & Haberman, B. (2005). *Unique local IPv6 unicast addresses* (RFC 4193). RFC Editor. <https://doi.org/10.17487/RFC4193>
- Internet Assigned Numbers Authority. (n.d.). *Number resources*. Retrieved August 17, 2026, from <https://www.iana.org/numbers>
- Internet Corporation for Assigned Names and Numbers. (2011, February 3). *Available pool of unallocated IPv4 Internet addresses now completely emptied*. <https://www.icann.org/en/news/releases/release-03feb11-en.pdf>
- Kawamura, S., & Kawashima, M. (2010). *A recommendation for IPv6 address text representation* (RFC 5952). RFC Editor. <https://doi.org/10.17487/RFC5952>
- Postel, J. (1981). *Internet protocol* (RFC 791). RFC Editor. <https://doi.org/10.17487/RFC0791>
- Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., & Lear, E. (1996). *Address allocation for private internets* (RFC 1918). RFC Editor. <https://doi.org/10.17487/RFC1918>
- Srisuresh, P., & Egevang, K. (2001). *Traditional IP network address translator (Traditional NAT)* (RFC 3022). RFC Editor. <https://doi.org/10.17487/RFC3022>
- Srisuresh, P., & Holdrege, M. (1999). *IP network address translator (NAT) terminology and considerations* (RFC 2663). RFC Editor. <https://doi.org/10.17487/RFC2663>

Further reading

- James F. Kurose and Keith W. Ross, Computer Networking: A Top-Down Approach, 9th ed., Pearson, 2026 — Provides a clear and comprehensive treatment of the network layer, including IPv4 and IPv6 addressing, subnetting, DHCP, NAT, forwarding, routing algorithms, routing protocols, and the distinction between the data plane and the control plane.
- Andrew S. Tanenbaum, Nick Feamster, and David J. Wetherall, Computer Networks, 6th ed., Pearson, 2022 — Offers a systematic explanation of the network layer, with detailed coverage of IP addressing, IPv4 and IPv6, routing algorithms, packet

forwarding, internetworking, congestion, and the mechanisms that allow heterogeneous networks to communicate.

- Behrouz A. Forouzan, Data Communications and Networking with TCP/IP Protocol Suite, 6th ed., McGraw Hill, 2022 — Dedicates separate chapters to network-layer data transfer and packet routing, making it particularly useful for exploring IPv4 and IPv6 addressing, subnetting, address translation, forwarding, routing tables, and the mechanisms used to move packets between networks.
- Douglas E. Comer, Internetworking with TCP/IP, Volume One, 6th ed., Pearson, 2014 — Provides an in-depth treatment of the Internet Protocol and TCP/IP architecture, including IPv4 and IPv6 addressing, address resolution, subnetting, CIDR, datagram forwarding, routing, DHCP, NAT, and the principles that make large-scale internetworking possible.
- Wendell Odom, CCNA 200-301 Official Cert Guide, Volume 1, 2nd ed., Cisco Press, 2024 — Provides a practical approach to IPv4 addressing and subnetting, router configuration, static routes, routing tables, packet forwarding, and IPv6, making it especially useful for connecting theoretical network-layer concepts with the configuration and operation of real routers.

Editorial notes

In producing these materials, I have used generative artificial intelligence tools to support the writing process, particularly to improve the form of the text, reorganize content, refine wording, and speed up certain editorial tasks.

As I work independently on the production of these materials, I try to automate any activities that can reasonably be automated, so that I can devote more time to research, design, and content development.

However, artificial intelligence does not determine the content of this work: the choice of topics, structure, ideas, interpretations, examples, and teaching approach are developed by me. AI is therefore used primarily as a tool to support production and formal revision, while authorship and responsibility for the design and content remain mine.

Unless otherwise indicated, this material is distributed under the Creative Commons Attribution–NonCommercial–ShareAlike 4.0 International license (CC BY-NC-SA 4.0).

You are therefore free to share, redistribute, adapt, and create derivative works based on this material, provided that appropriate credit is given, the material is not used for commercial purposes, and any modified or derivative versions are distributed under the same license.

