ISP Architecture

Written by [Febogamedev]

After examining what a WAN is and the role played by ISPs, we can take a closer look at what happens when a user connects to the Internet.
The connection is not made directly to the "entire Internet". It first passes through the user's operator network, which is connected to other networks in turn.

The image below shows what actually happens when a user connects to the Internet:

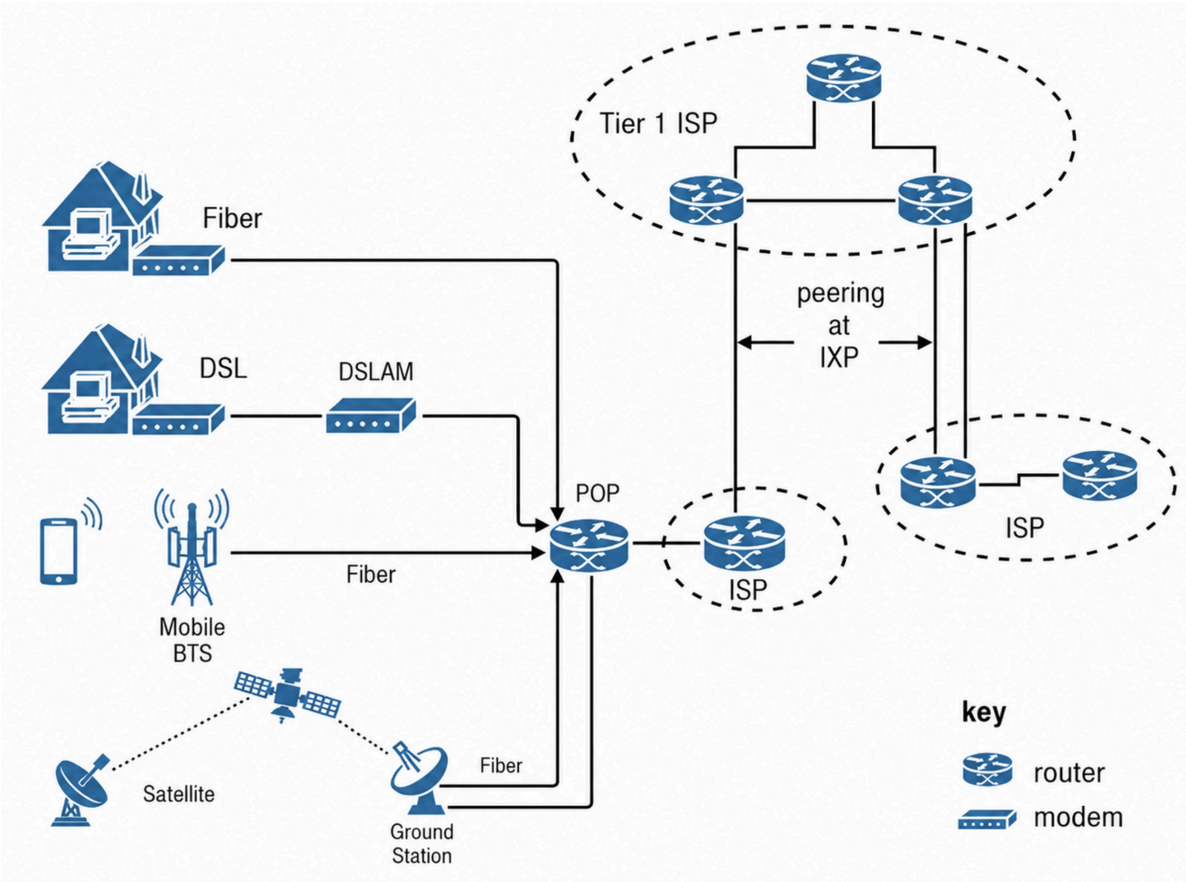

To connect to the Internet, a computer connects to an ISP, which enables it to exchange packets with all other hosts that can be reached through the Internet (Kurose & Ross, 2026). Common methods for connecting to an ISP include telephone or copper lines, fiber-optic connections, mobile networks, and, in some cases, radio or satellite links (Kurose & Ross, 2026).

The following sections examine telephone lines and fiber-optic connections:

Telephone Line
DSL (Digital Subscriber Line) uses a telephone line to transmit digital data.
In a DSL connection, the computer is connected to a device called a **DSL modem**, which converts digital data into signals suitable for transmission over a twisted-pair telephone line.

At the other end, a device called a **DSLAM** (Digital Subscriber Line Access Multiplexer) receives these signals and passes them into the operator's network (International Telecommunication Union [ITU], 1999; Kurose & Ross, 2026).

In practice, the home modem sends data through a copper twisted-pair telephone line. The line reaches a DSLAM, usually located in a telephone exchange or street cabinet. The DSLAM receives the lines of many subscribers simultaneously, keeps each subscriber's traffic separate, and forwards it to the ISP's network through a higher-capacity connection, often using fiber optics (Kurose & Ross, 2026).

This method is limited by the last mile of the connection.
In telecommunications, the term "**last mile**" refers to the section of cable connecting telephone exchanges to end users, namely the part of the telephone network known as the access network.
Connection performance decreases in this section because the last mile uses a much slower system—the telephone line—than the technology used within the ISP's network, namely fiber optics (ITU, 1999; Kurose & Ross, 2026).

Fiber Optics
For faster Internet access, fiber-optic cables can be installed all the way to homes, since fiber optics is much faster than a telephone line (ITU, 2008; Kurose & Ross, 2026).
This is achieved through configurations called **FTTH** (Fiber to the Home) or **FTTC** (Fiber to the Cabinet).

- With FTTC technology, the fiber-optic cable reaches a cabinet located a short distance from residential or commercial buildings. From the cabinet, the signal is transmitted to the customer using a different technology that relies on the existing copper cable for the final section of the connection. The last-mile limitation therefore remains, but its effect is reduced.
- FTTH technology provides better performance and more stable connections than FTTC because the fiber-optic cable reaches the user's home directly. The final section is therefore made of fiber rather than copper, greatly reducing the bottleneck typical of connections based on twisted-pair telephone lines (ITU, 2008; Kurose & Ross, 2026).

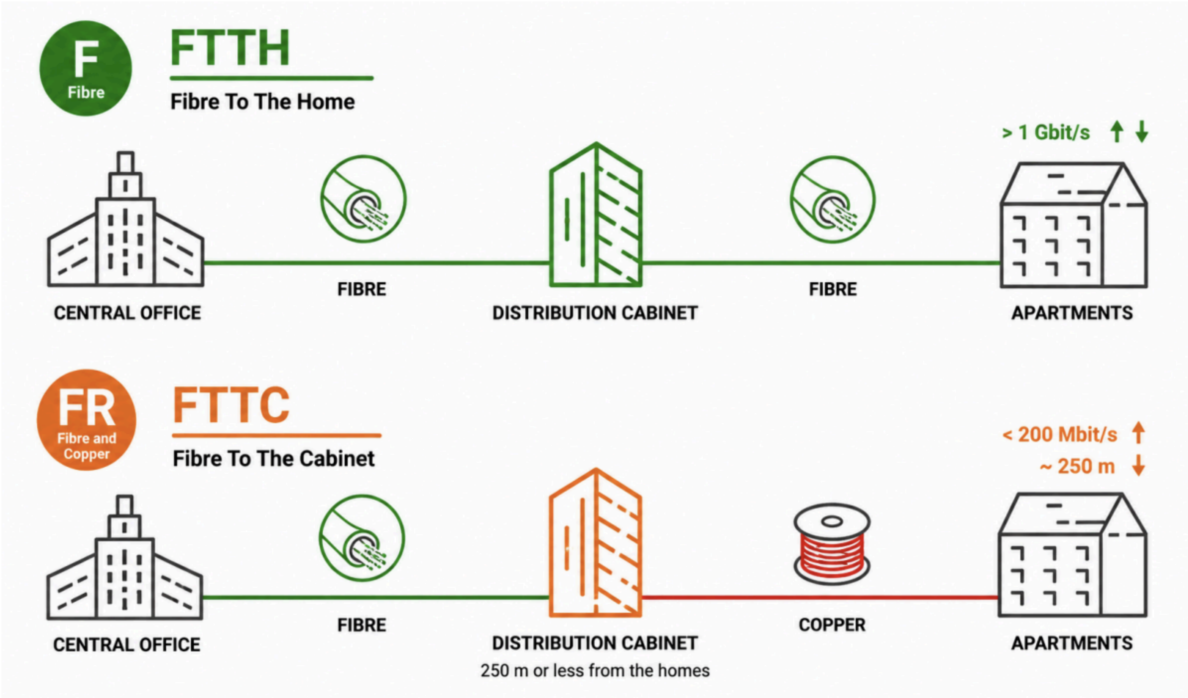

At this point, the data packets can be forwarded by the ISP (Kurose & Ross, 2026).
The point at which packets enter the ISP's network is called a **Point of Presence (POP)**.
ISP networks may operate at a regional, national, or international level.
The ISP's backbone consists of long-distance transmission links that interconnect POPs in the different cities served by the ISP Kurose & Ross, 2026).

Example
Suppose you live in Palermo. Your ISP has a PoP in Palermo connected to its national backbone by fiber.
When you browse the Internet, your traffic enters the local PoP, where it is processed and then routed through the national or international network.

When a packet must reach a destination on the Internet, it may follow different routes (Kurose & Ross, 2026; Rekhter et al., 2006).

- If the destination is within the same ISP's network, the packet can be routed directly through the operator's network.
- If the destination is on a different network, the packet must pass through an interconnection between networks. ISPs can exchange traffic through **peering** agreements, in which two networks connect directly to exchange data (Center for Applied Internet Data Analysis [CAIDA], n.d.; Kurose & Ross, 2026).

Many of these exchanges take place at Internet **Exchange Points, or IXPs**: physical locations where ISPs, corporate networks, and major content providers interconnect to exchange traffic more efficiently (Euro-IX, n.d.; Kurose & Ross, 2026).

Example

IXPs are not located everywhere. They are mainly found at major strategic connection hubs: cities with extensive network infrastructure through which national and international operators pass.
In practice, an IXP is a “marketplace” where ISPs, corporate networks, and content providers such as Google, Meta, and Netflix exchange traffic to reduce costs and latency—the delay before data reaches its destination.
Italy has several Internet Exchange Points, including MIX in Milan, NaMeX in Rome, and TOP-IX in Turin (MIX, n.d.; NaMeX, n.d.; TOP-IX, n.d.).

When an ISP cannot reach a destination through its own connections or peering agreements, it can purchase IP transit from another operator called a **transit provider**.
A transit provider is an ISP that sells connectivity to other ISPs or customers. It allows data traffic to be routed through its network to reach destinations that the customer’s network could not otherwise access directly (CAIDA, n.d.; Kurose & Ross, 2026).
This gives the customer organization access to the Internet without requiring it to negotiate peering agreements with many different networks.
The connectivity provided by a transit provider is called **IP transit**.

ISP Hierarchies

Now that we understand how an individual ISP works, we can examine the overall structure.
Not all ISPs are the same. They have different levels of coverage, independence, and decision-making power, creating a logical and commercial hierarchy.
This structure forms an ISP hierarchy divided into levels called tiers (CAIDA, n.d.; Kurose & Ross, 2026).
Understanding this hierarchy is essential for explaining how the Internet works on a global scale, who operates its infrastructure, and how data travels from one side of the world to the other.
The ISP hierarchy is divided into three main levels:

- Tier 1.
- Tier 2.
- Tier 3.

Tier 3 – Local or Retail ISPs
Tier 3 ISPs are the providers closest to end users.

- They do not have an extensive network of their own.
- They purchase connectivity from Tier 2 ISPs—or, in some cases, directly from Tier 1 ISPs—which act as transit providers.
- They provide Internet access services to households, schools, offices, and small businesses.
- They operate at a local or regional level.
- They do not participate directly in global peering (CAIDA, n.d.; Kurose & Ross, 2026).

<u>Example</u>
Tier 3 providers may be small local operators, wireless providers serving mountainous areas, or regional operators in underserved locations.

Tier 2 – Regional Providers
Tier 2 ISPs are intermediate providers: they operate their own networks but cannot reach the entire Internet independently.

- They have local peering agreements with other ISPs at the same level.
- They pay one or more Tier 1 providers for IP transit, allowing them to send data to networks they cannot reach directly. In this case, the Tier 1 ISPs act as transit providers for the Tier 2 ISPs.
- They often operate at a national or continental level.
- They provide connectivity to Tier 3 ISPs or large companies (CAIDA, n.d.; Kurose & Ross, 2026).

<u>Example</u>
An individual user normally connects through a Tier 2 ISP.
Suppose you live in Milan or Palermo and have a home Internet connection provided by Fastweb, Vodafone, or WindTre.
These operators have their own national networks, regional backbones, peering connections with other providers, and data centers.
However, they do not operate an independent global network, so they must purchase transit from Tier 1 providers to reach other countries or continents. They are therefore classified as Tier 2 ISPs.
As a user, you therefore connect to a Tier 2 ISP, which relies on Tier 1 ISPs to reach the rest of the world.

A Tier 2 ISP may act as a transit provider for other Tier 2 ISPs for several reasons:

- <u>Economic reasons</u>: a Tier 2 provider may offer more competitive prices than a Tier 1 provider in certain regions.
- <u>Technical reasons</u>: it may provide better service quality along specific routes.
- <u>Redundancy</u>: using multiple transit providers, including Tier 1 and Tier 2 networks, improves resilience and service continuity (Kurose & Ross, 2026).

<u>Example</u>
A Tier 2 provider with a strong network in Europe could supply transit to another Tier 2 provider that needs efficient access to that region, even if the second provider is already connected to a Tier 1 network.
Specific example:

- ISP A is a Tier 2 provider in Italy with a strong national network and connections to two Tier 1 providers.
- ISP B is another Tier 2 provider located in Spain.

- ISP B can purchase IP transit from ISP A, for example to improve routing toward Italy or to obtain a backup route.

Tier 1 – Global Backbone Operators
At the top of the hierarchy are large international backbones containing thousands of routers connected by fiber-optic links.
These ISPs, known as Tier 1 providers, do not pay for transit and operate international backbones that carry global traffic between continents through submarine cables crossing the oceans.
A Tier 1 operator can reach every network on the Internet without purchasing IP transit, thanks to its own connections and peering agreements with other Tier 1 providers.
In practice, Tier 1 providers play a fundamental role in connecting major networks and continents.
A Tier 1 ISP has peering agreements with all other Tier 1 providers. This means they exchange traffic without paying one another, based on a balanced and reciprocal exchange of data (CAIDA, n.d.; Kurose & Ross, 2026).

Supplementary materials

Key points

- The user connects to the Internet through their **ISP’s network**.
- Technologies such as **DSL**, **FTTC**, and **FTTH** form the access network and the last mile.
- Traffic enters the operator’s network through a **Point of Presence (POP)** and travels across its backbone.
- ISPs exchange traffic through **peering**, **Internet Exchange Points (IXPs)**, or by purchasing **IP transit**.
- ISPs are classified as **Tier 1**, **Tier 2**, or **Tier 3** according to the size of their networks and their connections with other operators.

Exercises

From the Computer to the ISP Network
Marco uses a DSL connection. His computer sends a request to a website on the Internet. The home connection uses a telephone copper pair and, in Marco’s area, users’ lines reach a DSLAM installed in a street cabinet.

Answer the following questions:

- Explain what function the DSL modem performs when Marco sends data.
- Describe the path followed by the data from the modem to the DSLAM.
- Explain what function the DSLAM performs.
- Explain why the DSLAM can receive lines from many users at the same time without mixing up their traffic.
- Indicate how the data can be forwarded from the DSLAM to the ISP network.
- Explain why the copper connection between the home and the operator’s infrastructure can represent a bottleneck.
- Identify which part of the path is referred to as the last mile.

FTTC or FTTH?
Two families live in the same neighborhood.

- <u>Family A</u>: uses an FTTC connection. Fiber reaches the street cabinet, while the final section to the home uses copper.
- <u>Family B</u>: uses an FTTH connection. Fiber reaches the home directly.

Compare the two solutions.

- Reconstruct the physical path followed by the connection in the FTTC solution.
- Reconstruct the physical path followed by the connection in the FTTH solution.
- Identify which of the two solutions still includes a final copper section.
- Explain why FTTC reduces the last-mile problem but does not eliminate it.
- Explain why FTTH can provide higher performance and a more stable connection.
- An operator wants to reduce the last-mile bottleneck as much as possible. Indicate which of the two technologies it should prioritize and justify your choice.

From the POP to the Backbone
A user living in Palermo uses a national ISP. Their traffic enters the operator's network through a POP in Palermo.
The ISP also has POPs in Rome, Milan, and Turin, connected through its national backbone.
Analyze the situation.

- Explain what a Point of Presence (POP) is.
- Describe what happens when the user's traffic reaches the Palermo POP.
- Explain what is meant by an ISP backbone.
- Describe the function performed by the backbone in connecting POPs located in different cities.
- Imagine that the destination is within the same ISP network. Explain why the traffic could remain entirely within the operator's network.
- Imagine instead that the destination belongs to another network. Explain why the traffic will eventually have to leave the ISP network.
- Reconstruct the conceptual path: user → access network → POP → backbone → destination.

Two ISPs Need to Exchange Traffic
ISP A and ISP B have many customers who frequently communicate with one another.
At present, traffic between the two networks follows an indirect path. The two operators are considering connecting directly at an Internet Exchange Point.

Answer the following questions:

- Explain what is meant by peering.
- Explain what function an IXP can perform in the situation described.
- Indicate which types of organizations, in addition to ISPs, can connect to an IXP according to the learning material.
- Explain why directly connecting the two networks can make traffic exchange more efficient.
- Explain why a more direct path can help reduce latency.
- A large content provider wants to exchange traffic directly with several ISPs. Explain why an IXP can be a strategic location for establishing these interconnections.
- Explain why an IXP should not be confused with an ISP that simply sells a standard Internet connection to end users.

When Peering Is Not Enough
A small ISP can directly reach some networks through peering agreements, but it does not have enough connections to reach every destination on the Internet.

The operator therefore decides to purchase a service from a larger ISP.
Analyze the solution.

- Explain what is meant by IP transit.
- Explain what function the transit provider performs.
- Describe what happens when a customer of the small ISP wants to reach a network that is not accessible through the available direct connections or peering agreements.
- Explain why purchasing transit allows the small ISP to avoid establishing direct agreements with every network on the Internet.

- Compare the role of peering with that of IP transit.
- Imagine that the ISP uses two different transit providers. Explain what possible advantage related to service continuity is described in the learning material.
- A second provider offers a particularly efficient path to a specific region. Explain why an ISP might choose to purchase transit from this operator as well.

Reconstructing the ISP Hierarchy

Consider three operators:

- ISP A: operates only in a few provinces, sells connections to households and small businesses, and purchases connectivity from larger operators.
- ISP B: owns a national network, has peering agreements with several networks, provides connectivity to some smaller ISPs, but must purchase transit to reach part of the Internet.
- ISP C: owns a large international backbone, can reach all Internet networks through its own infrastructure and peering agreements with other major operators, and does not purchase IP transit.

Analyze the three cases.

- Classify ISP A, ISP B, and ISP C as Tier 1, Tier 2, or Tier 3 respectively, and justify each choice.
- Explain why ISP A must purchase connectivity from other operators.
- Explain why ISP B can simultaneously purchase transit from a larger operator and sell it to a smaller operator.
- Explain which characteristic mainly distinguishes ISP C from the other two with regard to IP transit.
- Describe a possible path followed by a packet sent by a customer of ISP A to a very distant destination, indicating which operators it might pass through.
- Explain why the Tier classification is useful for understanding relationships between operators, but does not constitute a formal standard with always rigid boundaries.
- Finally, reconstruct the complete path of a packet by connecting the main concepts covered in the learning material: user → access network → POP → ISP backbone → peering or IP transit → other networks → destination.

Bibliography

Center for Applied Internet Data Analysis. (n.d.). *AS relationships*. Retrieved August 17, 2026, from https://www.caida.org/catalog/datasets/as-relationships/

Euro-IX. (n.d.). *Join Euro-IX*. Retrieved August 17, 2026, from https://www.euro-ix.net/en/forixps/join-euro-ix/

International Telecommunication Union. (1999). *Asymmetric digital subscriber line (ADSL) transceivers (Recommendation ITU-T G.992.1)*. https://www.itu.int/rec/T-REC-G.992.1

International Telecommunication Union. (2008). *Gigabit-capable passive optical networks (GPON): General characteristics (Recommendation ITU-T G.984.1)*. https://www.itu.int/rec/T-REC-G.984.1

Kurose, J. F., & Ross, K. W. (2026). *Computer networking: A top-down approach* (9th ed.). Pearson. https://www.pearson.com/en-us/subject-catalog/p/computer-networking-a-top-down-approach/P200000013385/9780135415603
MIX. (n.d.). *Chi siamo*. Retrieved August 17, 2026, from https://www.mix-it.net/chi-siamo/
NaMeX. (n.d.). *What we do*. Retrieved August 17, 2026, from https://www.namex.it/what-we-do/
Rekhter, Y., Li, T., & Hares, S. (2006). *A Border Gateway Protocol 4 (BGP-4) (RFC 4271)*. RFC Editor. https://doi.org/10.17487/RFC4271
TOP-IX. (n.d.). *Cos'è TOP-IX*. Retrieved August 17, 2026, from https://top-ix.org/consorzio/

Further reading

- James F. Kurose and Keith W. Ross, Computer Networking: A Top-Down Approach, 9th ed., Pearson, 2026 — Provides a clear introduction to Internet access networks and the Internet core, connecting technologies such as DSL and fiber access with ISPs, packet switching, routing, backbone networks, and the interconnection of different networks.
- Andrew S. Tanenbaum, Nick Feamster, and David J. Wetherall, Computer Networks, 6th ed., Pearson, 2022 — Offers a systematic treatment of broadband access networks, transmission technologies, routers, wide-area networks, Internet architecture, and the mechanisms through which independent networks are interconnected.
- William B. Norton, The Internet Peering Playbook: Connecting to the Core of the Internet, 2014 ed., DrPeering Press — Focuses specifically on Internet peering, IP transit, Internet Exchange Points, public and private peering, ISP interconnection strategies, and the economic and technical relationships that shape the core of the Internet.
- Iljitsch van Beijnum, BGP, O'Reilly Media, 2002 — Provides a more technical introduction to interdomain routing and explains how ISPs use BGP to exchange routing information, including Internet topology, transit and peering, ISP classification, multihoming, route selection, and connections to Internet Exchange Points.
- Gerd Keiser, FTTX Concepts and Applications, Wiley, 2006 — Provides a specialized introduction to fiber access networks, covering optical-fiber communication, Passive Optical Networks (PONs), FTTX architectures, and fiber-to-the-premises deployments, making it particularly useful for exploring the access-network and FTTH topics in greater depth.

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.