

What are Computer Networks?


We use the internet everyday, but it isn't often – if ever – that we need to consider how it works. That's because the protocols powering the internet were intentionally designed to operate as invisibly as possible. Built around a common language, networks enable devices to communicate with each other and share resources. By standardizing how computers talk with each other, we have expanded the scale of networks over time seeking to achieve a global cloud infrastructure.

“How we are at the small scale is how we are at the large scale. The patterns of the universe repeat at scale. There is a structural echo that suggests two things: one, that there are shapes and patterns fundamental to our universe, and two, that what we practice at a small scale can reverberate to the largest scale.

— adrienne maree brown

Connected Communities

Computer networks consist of nodes – which are devices that are seeking to communicate – as well as the links between them. Under some circumstances, nodes will connect to other nearby nodes and create a mesh that data can traverse while seeking its destination.



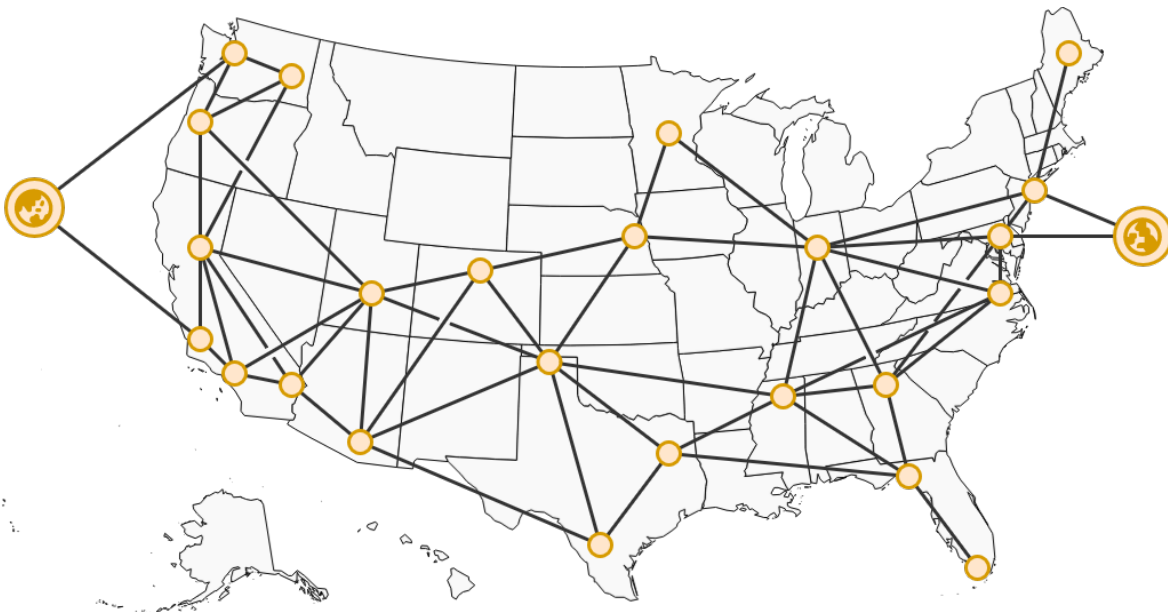
Edge networks form around bottlenecks that may arise in rural and disconnected areas. Functionally, this means that every person trying to connect to the World Wide Web must share a single Internet route – like an overly congested highway. This can increase the cost of internet

service, as well as effect the overall speed, availability and reliability.

The ways that communities connect to the Internet relies entirely on what is available within their physical geography. While in a major city, there may be many options for access to fast and reliable fiber Internet. Trying to connect to the Internet from rural Alaska, communities may find themselves restricted to a satellite connection.

These "digital deserts" can arise along geological boundaries – such as mountains or islands. More importantly, marginalized areas – such as Black, Indigenous and Hispanic communities – are not always offered equal or adequate Internet access. During 2021, it was estimated that over 42 million Americans do not have access to terrestrial broadband Internet – with 4 million in Texas alone.

What reasons – geological, political and social – do you think contribute to "digital deserts" without internet access?




Networks, similar to the rest of computing, leans heavily on abstractions that enable people (and network engineers) to comprehend the infrastructure required to power telecommunications at this scale. In order to build a global infrastructure, digital technologies have created a stratified system that simplifies data shared in between these layers.

This foresight during the creative process has worked to simplify our relationship with technology. You don't need to understand electrical engineering to build a computer system from parts you bought off-the-shelf. Similarly, data isn't concerned about routing its own path across the internet and only follows the one assigned to it.


Scales of Connection

Scale is a foundational way these systems are abstracted when trying to classify them. This is important to consider because different networks may have unique requirements. Your home network only needs to juggle a handful of people's data, while a college campus will be handling much more traffic from people located around the globe.



Personal Area Network

While your phone is connected to your headphones through BlueTooth, you are creating a Personal Area Network. These, as the name implies, operate on a smaller and more intimate scale. PANs utilize wired and wireless technologies to connect to each other, either through a hub device – like a cellphone or laptop – or directly to each other.



Local Area Network

A Local Area Network contains all of the nodes and links within a limited (often architectural or regional) area. This includes desktop, television and console devices attached by cable, as well as other devices connected wirelessly.

This could be as small as your home or some larger contained area – like a college campus or corporate headquarters. These institutions must subscribe to Internet service – just on a larger scale. They may have hundreds of interconnected wireless routers blanketing a mesh network over a large physical area.

Universities can have several campuses and corporations may have branch offices at different scales. A Virtual Private Network creates a private tunnel connecting two geographically separated LANs into one that is accessible by both locations. This enables devices over vast distances to communicate as if they were nearby. This can be accomplished invisibly through hardwired infrastructure, as well as on a device-by-device basis by connecting to a VPN server using the appropriate credentials.

When connecting computers to a wired network, there are a few devices that can incorporate physical cables to facilitate links between nodes. Judging solely by appearance, it can be hard to tell them apart.



Modem



Router



Repeater Hub



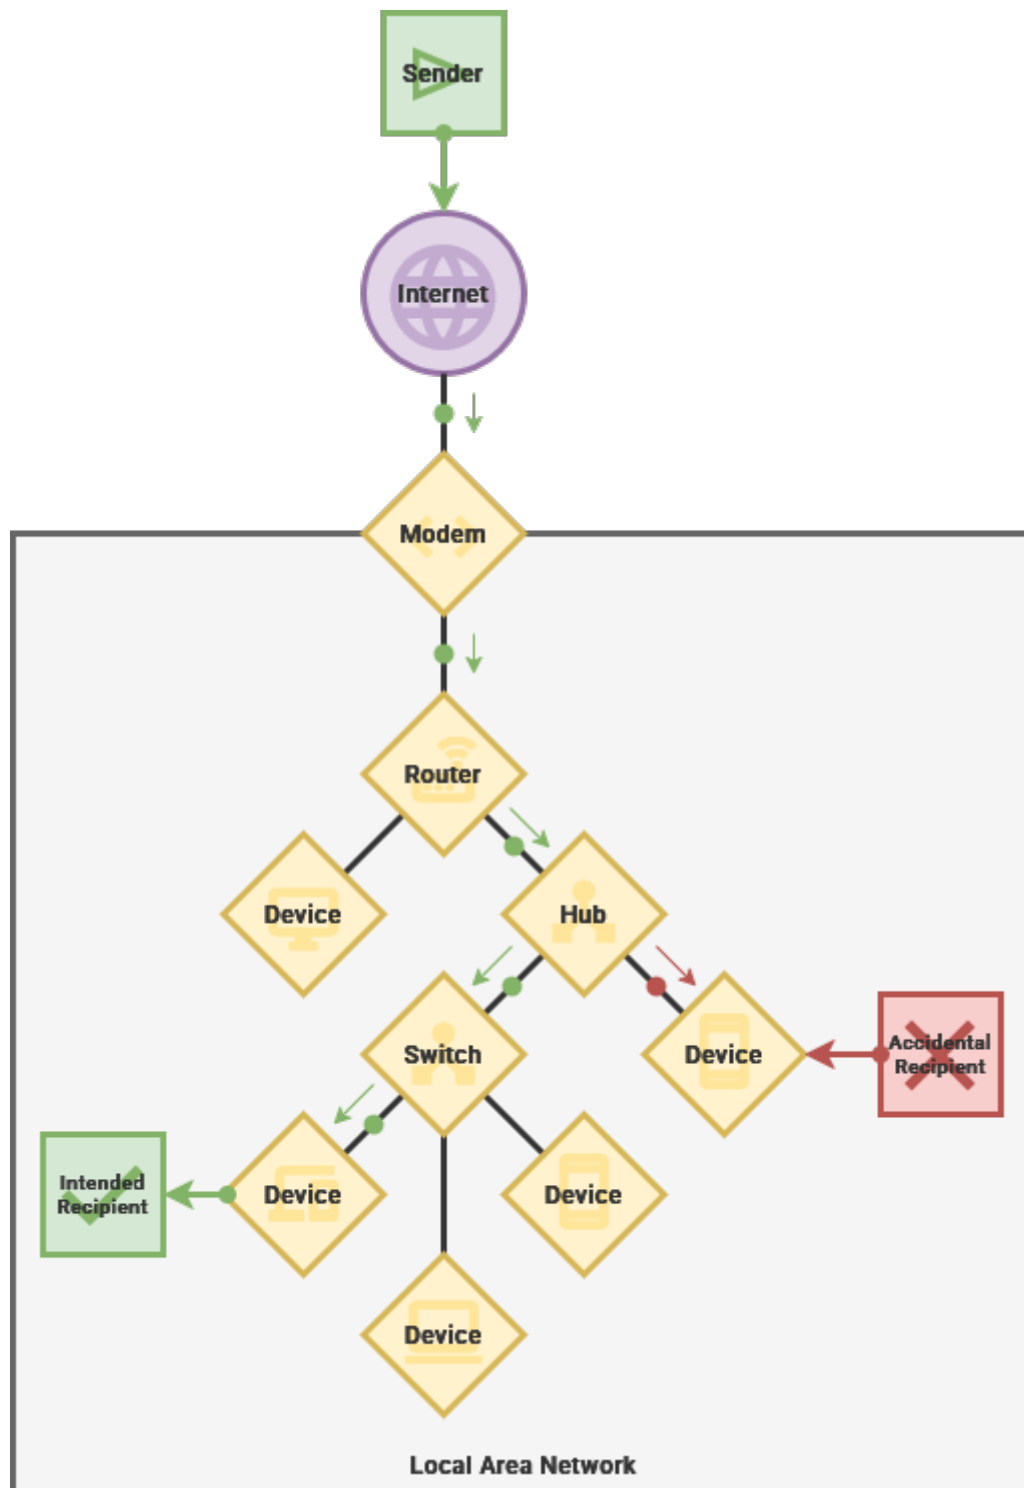
Switching Hub

Each Local Area Network has a modem responsible for transmitting data to and from an Internet Service Provider. This hardware is used to modulate – or translate – data into a signal that can be sent along a physical cable, radio wave or other connection.

The router connects to the modem and orchestrates communication between all the devices connected to it. While connected, each device is assigned a Private IP Address – a unique identification number on that network. This allows devices to quickly and intentionally exchange information over your network, even if there is no access to the outside of World Wide Web.

Three Private IP address ranges have been reserved for LAN networks: *192.168.68.100*, *172.16.0.0*, and *10.0.0.100*.

Ethernet is a standard for enabling network device communication over a wire. Ethernet cables are given a category designation – with higher categories meeting the performance requirements of data centers. Modern routers often incorporate wireless connectivity through the Wi-Fi standard – which turns data in radio waves that can be transmitted to devices through their wireless radio.



A repeater hub can connect many devices to the network on once, but will openly broadcasts all data it receives to every device connected to it. They can be cost-effective because of their simple design, but they greatly increase the potential for data sniffing – or the data being intercepted by someone other than the intended recipient.


What could result from the wrong device accidentally receiving it's data?

On the other hand, a switching hub behaves more intelligently by only sending data to it's intended recipient. This requires electronics to process the information being transmitted through it, but

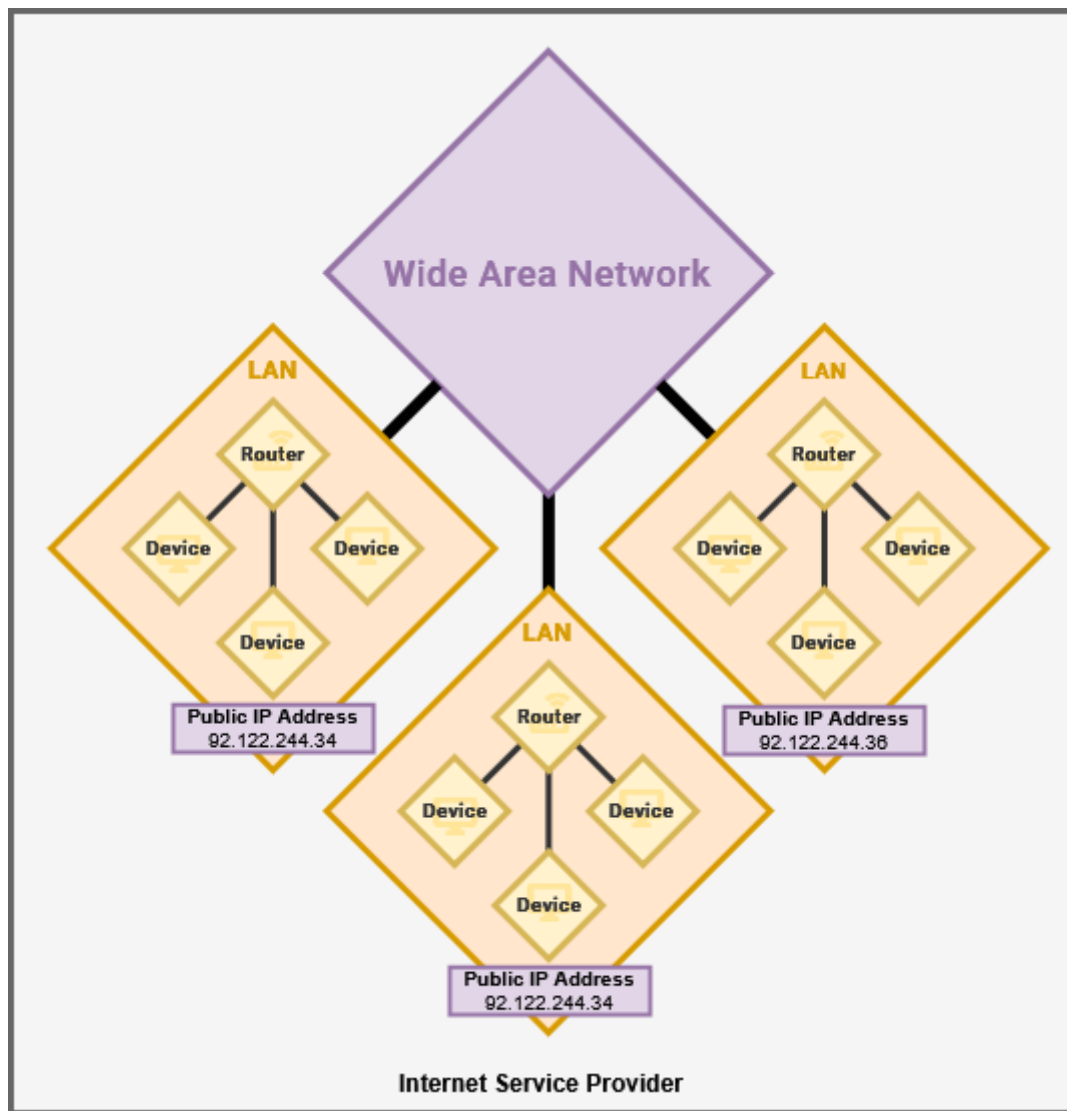
ultimately results in better reliability and security.

Wide Area Network

Your modem, acting as the gateway to the internet, is also assigned a Public IP Address. Similar to a phone number or street address, this is how networks find each other over the vast worldwide internet infrastructure. Whenever a device on your network contacts the World Wide Web, the router on the uses Network Address Translation to automatically convert data packets between your device's Private IP Address and the Public IP Address of your modem. This technique allows your LAN devices to access the Internet without the internet being able to access your LAN devices.




These disparate Local Area Networks – such as your home, your neighbors, city, county and state – are conglomerated together into a Wide Area Network or WAN.




Distributed Applications

The modern internet, as we know it, predominately operates within the client-server model. This means that one computer – a server – is used to respond to the requests of other computers – known as clients. Perhaps the most well known example of the client-server model is the modern World Wide Web.



Through a Web browser, we can navigate to a server using a graphical interface and enter a URL – such as example.com. This is more specifically known as a domain name and points towards the address of a Web server on the open internet. By leveraging the HTTP protocol, we can request data from a server and receive the response back in the form of an interactive website.



When you enter example.com into the browser's address bar, it needs to be translated into an IP address for our computer to connect to. The Domain Name System enables anyone in the world to know where to locate the web server over the World Wide Web. The predecessor to this infrastructure, which acts as the "phone book" of the entire Internet, was pioneered by Elizabeth Feinler.

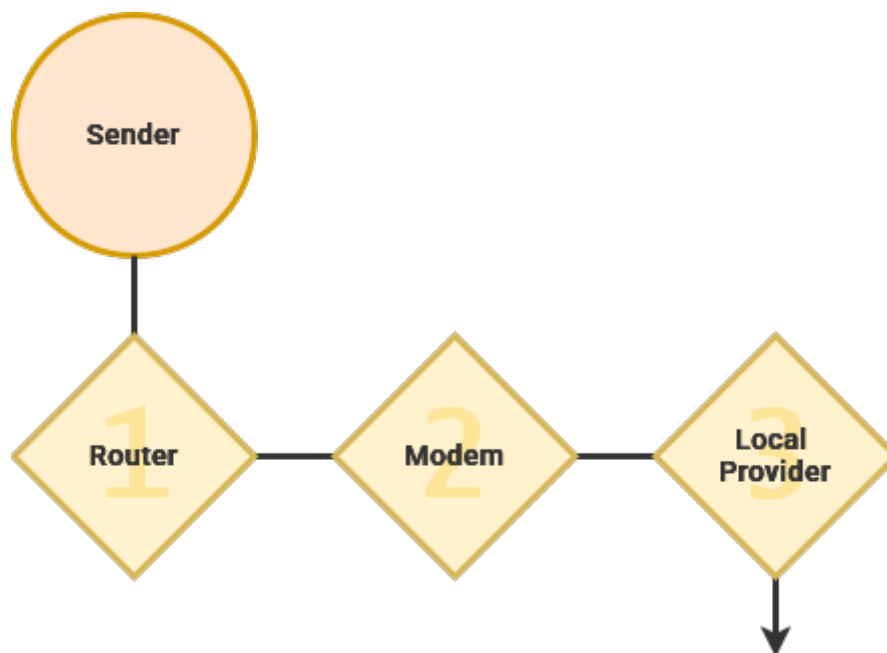
For example.com, the IP is "92.122.244.34".

While still less common for consumers, the peer-to-peer model is becoming more popular. These behave similarly to the mesh networks that allow ISPs to transmit data around the globe through interconnected networks. Within this network structure, each peer has the ability to act as both a server and a client to share data in a more efficient way. Each peer has the same privileges and power, creating decentralized networks – such as BitTorrent, OwnCloud and social media like the Fediverse or Bluesky.


The Path of Data

Data may need to travel vast distances to get from it's origin to final destination. This can include multiple internet service providers and connection types – ranging from physical cables to wireless connections. Carrying data around the globe can include anything from vast underwater cable networks to satellite relays in geostationary orbit.


Whenever data is transmitted over a network, it is first broken into small "packets" by your computer. These are transmitted to the network router and across the infrastructure laid by your Internet Service Provider to a local hub and, possibly, a regional or central hub.



Your Provider contracts access to middle-mile and long-haul providers specializing in quickly transmitting data across a geographical region. These carry data from outlying areas into major metropolitan areas where, if necessary, it can be sent around the globe. The Eastern US shares many undersea cable connections with Western Europe, just like the Western US connects to China and Japan.



Through the Border Gateway Protocol, data can find a route across this patchwork of autonomous and independently-owned network systems. This process relies on the mutual agreement between ISPs that every network system will act as a neutral peer to all other networks ensuring that messages will always be passed along towards its destination. If these data packets contain any erroneous or fabricated metadata, they will likely get lost during this exchange process.

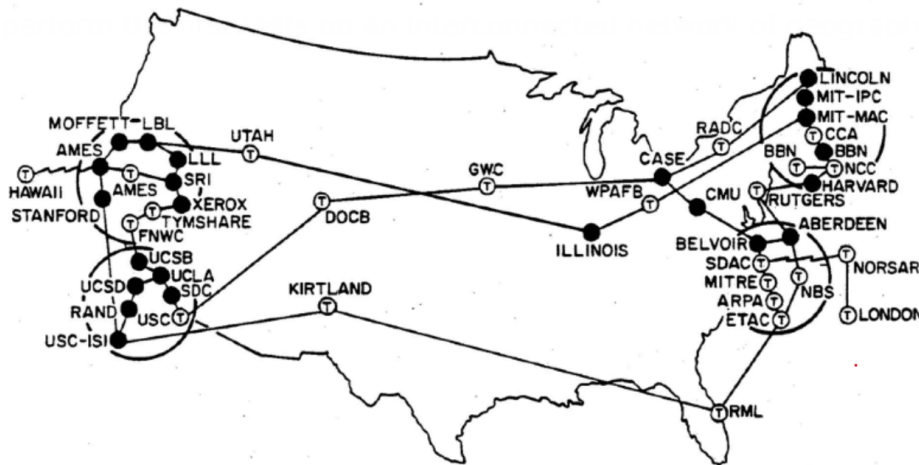


From here, packets will take the most direct route to its destination. Performing the same process in reverse, data transits through middle-mile and long-haul providers, before filtering through regional and local internet infrastructures. Finally, the data enters the intended router before being delivered to its destination.

When the receiver wants to send a response back to the original sender, it must repeat this entire process again. Modern software systems often implement mechanisms that will remember the quickest connection between two points. Web browsers, for example, are built on top of open protocols that enable two computers to create a persistent connection that can be reused for transporting data.

Open Standards

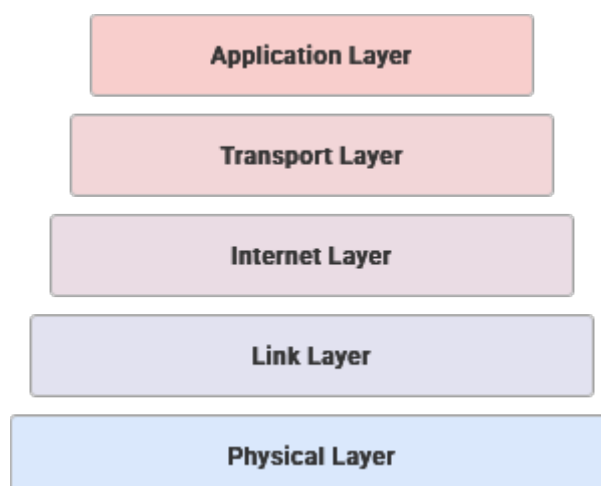
By the 1960s, computers had begun gaining traction and the US Department of Defense created the first Wide Area Network by connecting research campuses across the US with those in the United Kingdom and Norway. Before the concept of the World Wide Web, this network was used to connect locally isolated computer systems.



We quickly realized that we

would need to create tools, protocol and standards to ensure that these networking systems didn't become too fragmented. for us to communicate amidst rapid technological advancement. By the late 80s, globally expanding internet networks were increasingly relying on the TCP/IP protocol standards – more commonly known as the Internet Protocol Suite.

Simplified, this protocol standardized access to the Internet by creating functional layers that depend on each other while operating in parallel and communicating through a common language. The five layers of the Internet Protocol Suite are:



Each layer defines the requirements and expectations ensuring that disparate devices could mutually agree on a language to communicate with. At each level, we are ensuring that all nodes on the network are speaking a common language. By being able to effectively communicate, these layers can operate independently and only communicate tasks to the next layer as needed.

memory

Physical Layer

This level handles how a data signal is encoded and transmitted between hardware components using a range of connection types – such as electrical, optical or wireless links. Within this layer, each individual network-connected device is identified by a unique MAC address – with over 281 trillion possibilities.

Cable

Link Layer

This level defines the functional and procedural methods that are used to transmit data between nodes across their link. Within this layer, we are restricted to connections between nodes that are connected physically – such as through Ethernet and Wi-Fi. When necessary, a Virtual LAN can be used to segment a monolithic physical network into smaller virtual ones that operate in isolation.

lan

Internet Layer

This level covers the methods and specifications for transmitting data between intermediate routers along its path from origin to destination. When transmitting data across the world wide web, public IP addresses are used to route a path.

swap_horiz

Transport Layer

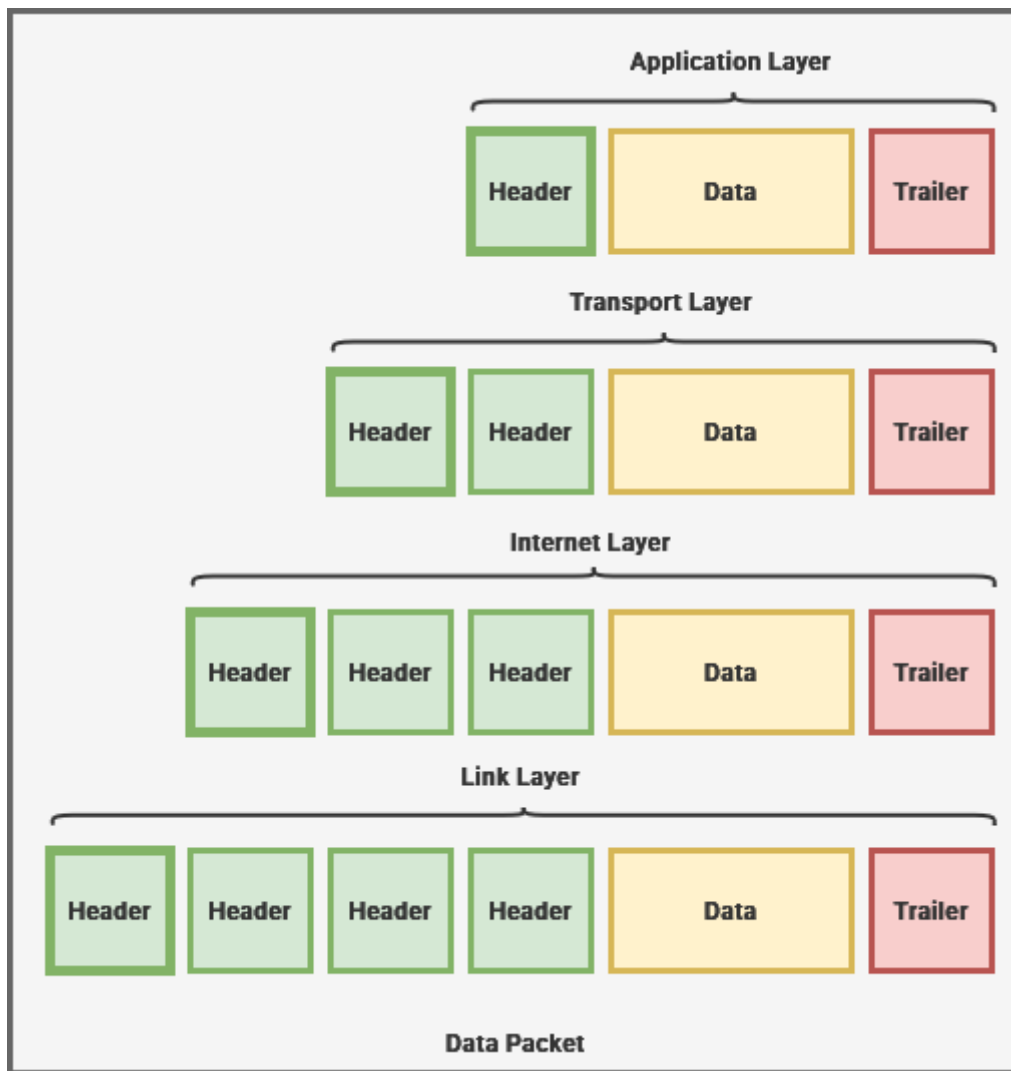
This level is responsible for ensuring that nodes can reliably communicate across this vast interconnected network. This is accomplished through an array of protocols that handle the security, reliability and flow logistics of data differently. Within this layer, any data to be sent over the network must be broken into smaller segments to facilitate smooth transmission.

api

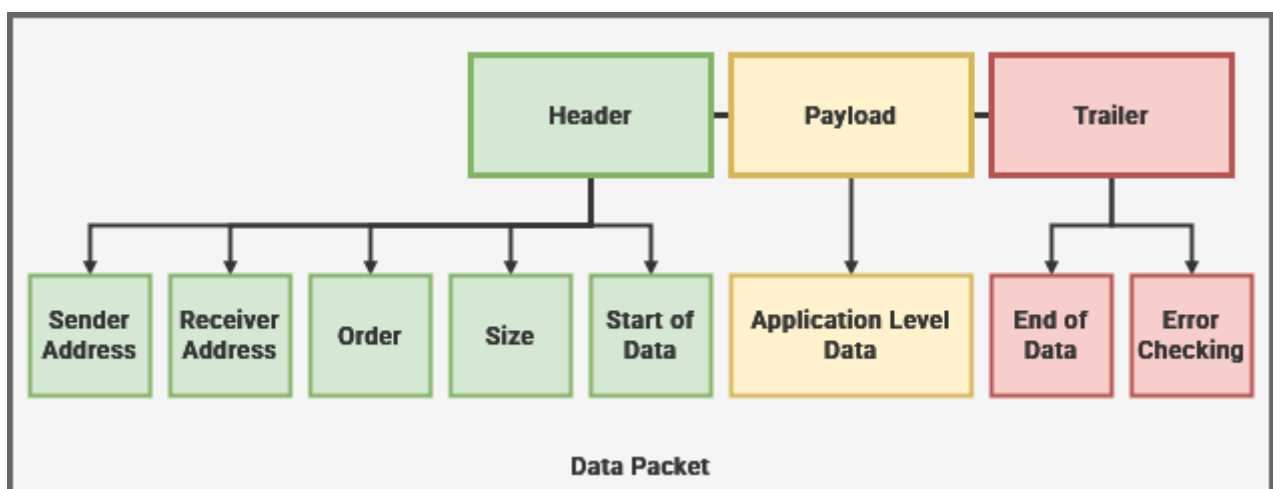
Application Layer

This is the top-most layer of the Internet where familiar services exist – such as the World Wide Web, e-mail, remote desktop connections and file transfers. This is the layer that people interact with most often and data transfer happens almost invisibly.

While passing between layers, "metadata" encapsulates the data – or "payload" – being sent over the Internet. This process assists navigation across the globe and ensures that it arrives in one piece without errors.



These "headers" detail how this data relates with other data, as well as the path taken and route still in progress. The "footer" will often contain a checksum – or abstracted data based on a mathematical equation – that can be compared to the received data and verify there aren't any errors.



Before data can be sent anywhere, the applications being used – such as between a web browser and web server – must agree on a method of communication. While passing down from the top-most Application Layer to the Transmission layer, software developers can choose different

protocols that define how the data will be transmitted over the open internet. During this process, data gets broken down into thousands of tiny "packets" that will need to be put back together on the other end.

Created alongside the internet in the early 60s, Transmission Control Protocol – more commonly known as TCP – is still one of the most widely used protocols. This protocol was built to ensure that all data will arrive without errors and in the exact order it was sent – all important properties for smooth web surfing, resilient email services and reliable data transfers. When using TCP, both devices must agree to "handshake" – or mutually agree on how they will be creating a link between them.

After the connection is established, each minuscule packet of information will be sent one after the other, ensuring each packet was received correctly before moving on to the next. When data becomes lost or corrupted, the sender is alerted so that it can be sent again. This process is extremely reliable, but can be slow because it requires more upfront coordination and processing power.

On the World Wide Web, we depend on TLS – or Transport Layer Security – to serve us private and secure websites using HTTPS through a TCP connection.




By contrast, User Datagram Protocol – or simply UDP – does not require a "handshake" to create an active connection. Instead, data packets can be sent directly to the server from a client device without needing explicit approval first. This requires that the server is always available to accept incoming connections without prior warning.

The server will respond to this request by also sending data without negotiation. Some applications – such as media services and video games – may begin to stream data by

broadcasting packets in quick succession. Using this protocol, there is no guarantee that the packet will be received in any particular order or timeframe. Even when a server sends packets in order, they may be received out of order because of the underlying network.

Neither the server nor the client know when to expect data and, by extension, neither know when data intended for them never arrives. While the server and client will be alerted if the data was corrupted along the way, the application needs to send it again. This protocol is preferable for time-sensitive applications, online video games and media streaming services.



When hosting digital services for a network, ports enable a single physical server to create a dedicated sub-addresses for multiple running multiple applications. This allows each port to use different transmission protocols – such as TCP and UDP. Unlike a physical hardware port, network ports reside entirely in the virtual space and only exist within your operating system while it's loaded.

Ports use a numeric identifier ranging from 0 to 65535.

If an IP address were to a building address, ports would be apartment numbers within that building. After receiving internet traffic, a computer will route data one final time to the applicable port. Services we use everyday are tied to a specific port on our computer. Often, these are defined through open standards that purposely reserve ports for specific services to reduce confusion.

| Port Number | Protocol | Use |
|-------------|----------|----------------|
| 53 | DNS | Web Browsing |
| 80 | HTTP | Web Browsing |
| 443 | HTTPS | Web Browsing |
| 20 | FTP | File Transfers |
| 22 | SSH | Remote Access |
| 25 | SMTP | E-Mail |

| Port Number | Protocol | Use |
|-------------|----------|--------|
| 110 | POP3 | E-Mail |
| 220 | IMAP | E-Mail |

Revision #56
Created 14 March 2025 01:50:23 by metaphorraccoon
Updated 27 April 2025 19:43:54 by metaphorraccoon