The Complete Guide to Networking & Routing Protocols

From AppleTalk to Fiber Optics: A Plain-English Journey Through the History and Technology of Computer Networks

---

Date	February 25, 2026
Audience	Non-technical readers, managers, and curious minds
Scope	LAN/WAN, routing protocols, switching, VLANs, ISPs, fiber

Table of Contents

1. Introduction: What Are Routing Protocols?

Every time you open a web page, send a text message, or stream a video, data travels across a vast web of interconnected networks. But how does that data know where to go? Routing protocols are the answer. They are sets of rules that routers—the traffic directors of the internet—use to determine the best path for your data to travel from its source to its destination.

Think of routing protocols as the GPS navigation system for the internet. Just as GPS considers traffic, road closures, and distance to find the fastest route for your car, routing protocols evaluate network conditions—speed, congestion, number of stops—to find the fastest path for your data.

Over the decades, these protocols have evolved from simple, broadcast-based systems suitable for a handful of computers in a single room, to incredibly sophisticated algorithms that manage billions of routes across the global internet. This guide tells that story in plain English, covering the technologies, the history, and the future of networking.

Diagram 1.1 — The Basic Idea of Routing

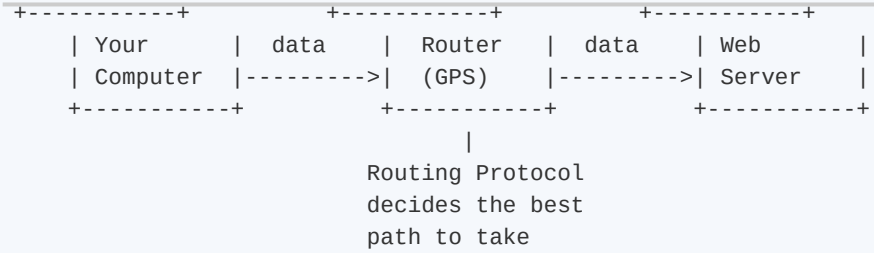

Diagram 1.1: A router sits between your computer and the destination. The routing protocol is the intelligence that tells the router which direction to send your data.

2. Network Fundamentals

What Is a Network?

A network is any group of computers or devices connected so they can exchange information. The network in your home—your laptop, phone, smart TV, and printer all talking to your Wi-Fi router—is a network. The global internet is also a network, just an unimaginably larger one.

IP Addresses: The Mailing Addresses of the Internet

Every device on a network needs a unique address so other devices can find it. This address is called an IP address (Internet Protocol address). The most common format, IPv4, looks like four numbers separated by dots: **192.168.1.25**. Each number ranges from 0 to 255. Think of it exactly like a mailing address: the first numbers narrow down the region (network), and the last numbers pinpoint the exact building (device).

Because the world has run out of IPv4 addresses (about 4.3 billion combinations), a newer system called **IPv6** is gradually taking over. IPv6 addresses are much longer and look like **2001:0db8:85a3::8a2e:0370:7334**, providing trillions upon trillions of unique addresses.

MAC Addresses: The Hardware Fingerprint

While IP addresses identify devices on a network level, every network interface card (the hardware that connects a computer to a network) has a permanent, factory-assigned identifier called a MAC address (Media Access Control). It looks like **00:1A:2B:3C:4D:5E**. MAC addresses work at the local level—switches use them to deliver data to the right device on a local network.

Diagram 2.1 — IP vs. MAC Addresses

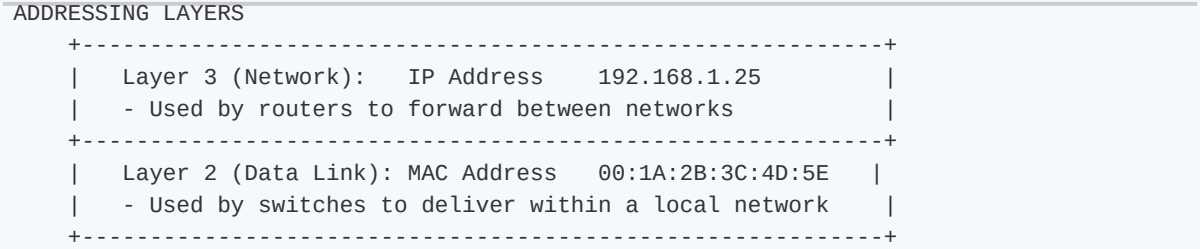

Diagram 2.1: IP addresses are used for routing between networks; MAC addresses are used for delivery within a single local network.

Packets: Data Broken Into Pieces

When you send a file across a network, it doesn't travel as one big chunk. Instead, it's broken into small pieces called **packets**. Each packet carries a portion of the data plus addressing information (source IP, destination IP). Packets may take different paths and arrive out of order—the receiving computer reassembles them. This design makes networks resilient: if one path fails, packets can take another.

Routers vs. Switches vs. Hubs

Three types of devices are commonly mentioned in networking:

Device	What It Does	Analogy
Hub	Broadcasts data to ALL connected devices	A megaphone in a room—everyone hears everything
Switch	Sends data only to the intended device using MAC addresses	A mail sorter—delivers letters to the right mailbox
Router	Forwards data between different networks using IP addresses	A post office—routes mail between cities

Diagram 2.2 — Hub vs. Switch vs. Router

```
HUB: sends to everyone          SWITCH: sends to one
    +----+----+----+----+           +----+----+----+----+
    | PC | PC | PC | PC |           | PC | PC | PC | PC |
    +--+-+--+-+--+-+--+-+           +--+-+----+----+--+-+
       |    |    |    |                 |              |
       +----+----+----+                 +-- data ------+
       all get the data                  only recipient gets it

    ROUTER: connects different networks
    +-----------+          +-----------+
    | Network A |---[R]---| Network B |
    +-----------+          +-----------+
```

Diagram 2.2: Hubs flood data everywhere (wasteful); switches are smart and send data only where needed; routers connect separate networks together.

3. LANs vs. WANs: Two Different Worlds

LAN — Local Area Network

A LAN covers a small geographic area—typically a single home, office floor, or building. All devices on a LAN are physically close and usually connected by Ethernet cables or Wi-Fi. LANs are owned and managed by the organization using them. Because distances are short, LANs are fast (typically 1 Gbps or more today) and inexpensive to build.

Common LAN activities include sharing printers, accessing local file servers, and communicating between workstations. The dominant LAN technology since the 1990s has been Ethernet, standardized as IEEE 802.3.

WAN — Wide Area Network

A WAN spans large geographic distances—connecting offices in different cities, states, or countries. The internet itself is the largest WAN. WANs typically lease connections from telecommunications companies (ISPs) because laying long-distance cables is enormously expensive. WANs are slower than LANs, more complex, and more costly, but they make global communication possible.

WAN technologies have included Frame Relay, ATM (Asynchronous Transfer Mode), MPLS (Multiprotocol Label Switching), and modern SD-WAN (Software-Defined WAN). Each generation improved speed, reliability, and flexibility.

MAN — Metropolitan Area Network

Between LANs and WANs sits the MAN, covering a city or campus. Universities, city governments, and cable TV systems often operate MANs. They use many of the same technologies as WANs but over shorter distances.

Diagram 3.1 — LAN vs. MAN vs. WAN Scale Comparison

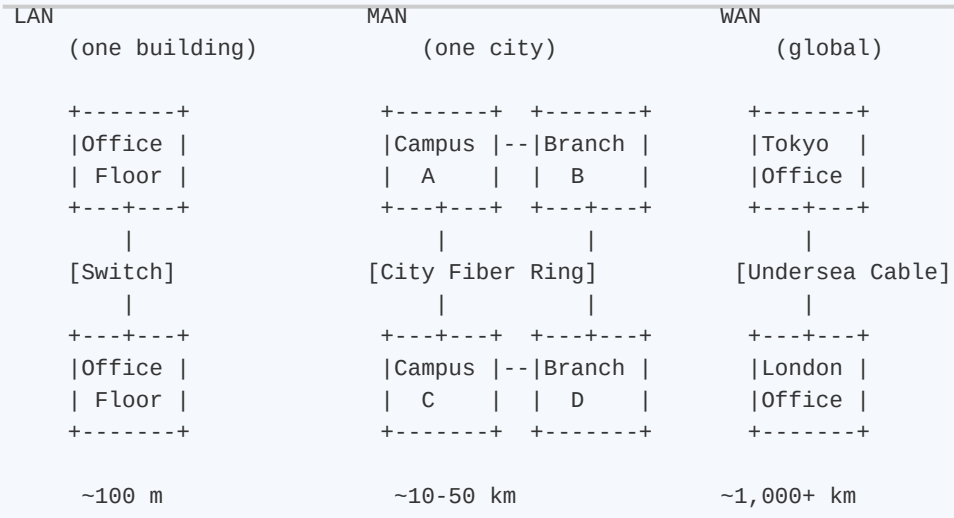

Diagram 3.1: Networks range from small (LAN) to city-wide (MAN) to global (WAN).

Characteristic	LAN	MAN	WAN
Distance	Up to ~1 km	~5-50 km	100+ km to global
Speed	1-100 Gbps	1-10 Gbps	50 Mbps - 100 Gbps
Ownership	Private	Private / shared	Leased from ISPs
Cost	Low	Medium	High
Latency	Very low (<1 ms)	Low (1-5 ms)	Variable (5-200 ms)
Example Tech	Ethernet, Wi-Fi	Metro Ethernet	MPLS, SD-WAN, Fiber

4. The OSI Model: Layers of Networking

To understand networking, engineers use a layered model called the **OSI Model** (Open Systems Interconnection). Think of it like building a house: you need a foundation, walls, wiring, and a roof—each layer serves a purpose and relies on the layer below it. The OSI model has seven layers:

Layer	Name	What It Does	Everyday Analogy
7	Application	User-facing apps (web, email)	The letter you write
6	Presentation	Data formatting, encryption	Translating the letter to the right language
5	Session	Manages connections between apps	Starting and ending a phone call
4	Transport	Reliable delivery (TCP/UDP)	Certified mail vs. regular mail
3	Network	Routing between networks (IP, OSPF)	The postal system choosing a route
2	Data Link	Local delivery (Ethernet, MAC)	The mail carrier on your street
1	Physical	Cables, signals, voltages	The roads and trucks

Routing protocols operate at **Layer 3** (Network). Switches operate at **Layer 2** (Data Link). Understanding which layer a technology works at helps explain its role and limitations.

Diagram 4.1 — Data Encapsulation Across the OSI Layers

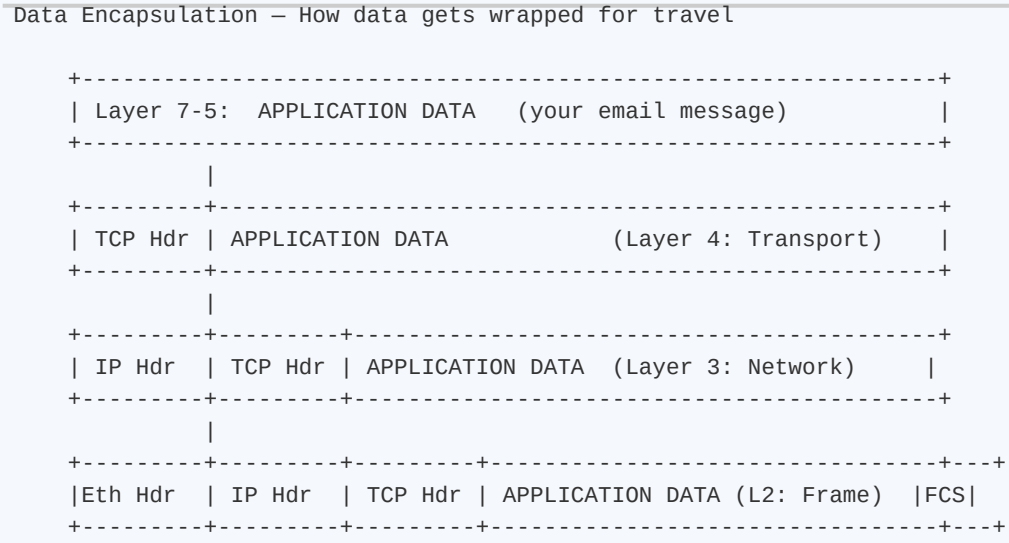

Diagram 4.1: As data moves down the layers, each layer adds its own header (addressing information). The receiving device peels off headers in reverse order.

5. Ethernet and the Birth of the LAN

Ethernet is the most important LAN technology ever created. Invented at Xerox PARC in 1973 by Robert Metcalfe and David Boggs, it was formally standardized as IEEE 802.3 in 1983. Ethernet defines how devices share a common communication medium—originally a single coaxial cable, later twisted-pair copper wires, and now fiber optics.

How Ethernet Works

Early Ethernet used a method called **CSMA/CD** (Carrier Sense Multiple Access with Collision Detection). In simple terms: before sending data, a device listens to see if anyone else is talking. If the line is clear, it transmits. If two devices transmit at the same time (a collision), both stop, wait a random amount of time, and try again. Modern switched Ethernet has largely eliminated collisions because each device gets its own dedicated connection to the switch.

Ethernet Generations

Standard	Speed	Year	Cable Type	Max Distance
10BASE5	10 Mbps	1983	Thick coax ("yellow cable")	500 m
10BASE2	10 Mbps	1985	Thin coax ("cheapernet")	185 m
10BASE-T	10 Mbps	1990	Cat 3 twisted pair	100 m
100BASE-TX	100 Mbps	1995	Cat 5 twisted pair	100 m
1000BASE-T	1 Gbps	1999	Cat 5e/6 twisted pair	100 m
10GBASE-T	10 Gbps	2006	Cat 6a twisted pair	100 m
25GBASE-SR	25 Gbps	2016	Multimode fiber	100 m
100GBASE-SR4	100 Gbps	2010	Multimode fiber	100 m
400GBASE	400 Gbps	2018	Single-mode fiber	10 km+

Notice the trend: speeds have increased from 10 Mbps to 400 Gbps—a 40,000x improvement—while the technology shifted from bulky coaxial cable to thin fiber optics.

Diagram 5.1 — Ethernet Evolution: Bus to Star Topology

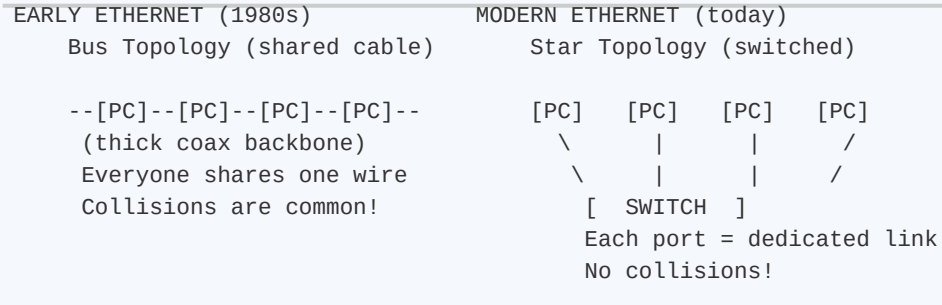

Diagram 5.1: Early Ethernet shared a single cable (collisions were inevitable). Modern Ethernet gives each device its own switched connection.

6. Switching: The Intelligent Bridge

A network switch is one of the most important devices in modern networking. While a hub blindly sends data to every port, a switch learns which devices are connected to which ports and sends data only where it needs to go. This is called **Layer 2 switching** because switches operate at Layer 2 of the OSI model, using MAC addresses.

How a Switch Learns

When a switch is first powered on, it knows nothing. As devices send data, the switch reads the source MAC address of each frame and records which port it came from in a **MAC address table** (also called a CAM table). Within seconds, the switch builds a complete map of which device is on which port. This process is entirely automatic.

Diagram 6.1 — Switch MAC Address Table

```
SWITCH MAC ADDRESS TABLE
    +------+-------------------+--------+
    | Port |    MAC Address    | VLAN   |
    +------+-------------------+--------+
    |  1   | 00:1A:2B:3C:4D:01|  10    |
    |  2   | 00:1A:2B:3C:4D:02|  10    |
    |  3   | 00:1A:2B:3C:4D:03|  20    |
    |  4   | 00:1A:2B:3C:4D:04|  20    |
    +------+-------------------+--------+

    When data arrives for MAC ...4D:03, the switch
    sends it ONLY out port 3.  Other ports are undisturbed.
```

Layer 2 vs. Layer 3 Switching

Traditional switches work at Layer 2 (MAC addresses). However, modern **Layer 3 switches** can also read IP addresses and make routing decisions, blurring the line between switches and routers. Layer 3 switches are common in large enterprise networks where high-speed routing between VLANs is needed.

Spanning Tree Protocol (STP)

When multiple switches are connected together for redundancy, loops can form—data circling endlessly between switches. **Spanning Tree Protocol (STP)**, invented by Radia Perlman in 1985, prevents this by automatically blocking redundant paths and re-enabling them only if a primary path fails. Think of it as a traffic system that closes some roads to prevent cars from driving in circles, but reopens them if the main road is blocked.

Diagram 6.2 — Spanning Tree Protocol Prevents Loops

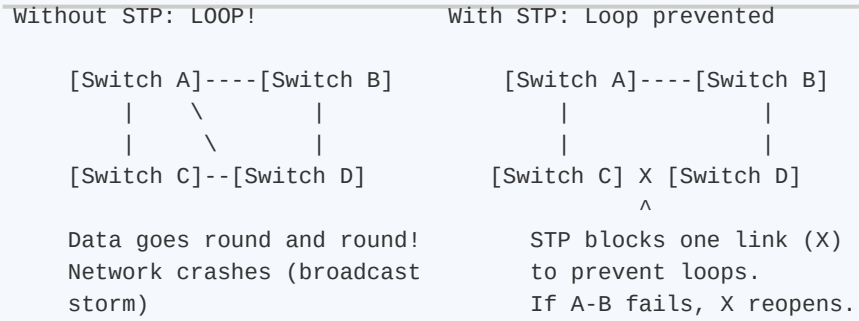

7. VLANs: Virtual Networks Inside a Network

Imagine you have one large office building with a single network, but you want the Accounting department and the Engineering department to be on separate networks for security and performance reasons. In the old days, you'd need separate physical switches and cables. **VLANs (Virtual Local Area Networks)** solve this by letting you create multiple separate networks on the same physical switch—no extra hardware required.

How VLANs Work

A VLAN is defined by assigning switch ports to a numbered group. Ports in VLAN 10 can only talk to other ports in VLAN 10; ports in VLAN 20 can only talk to ports in VLAN 20. Traffic between VLANs must pass through a router or Layer 3 switch (called **inter-VLAN routing**). VLANs are defined in the IEEE 802.1Q standard, which adds a small 4-byte tag to each Ethernet frame identifying which VLAN it belongs to.

Diagram 7.1 — VLANs: One Switch, Three Networks

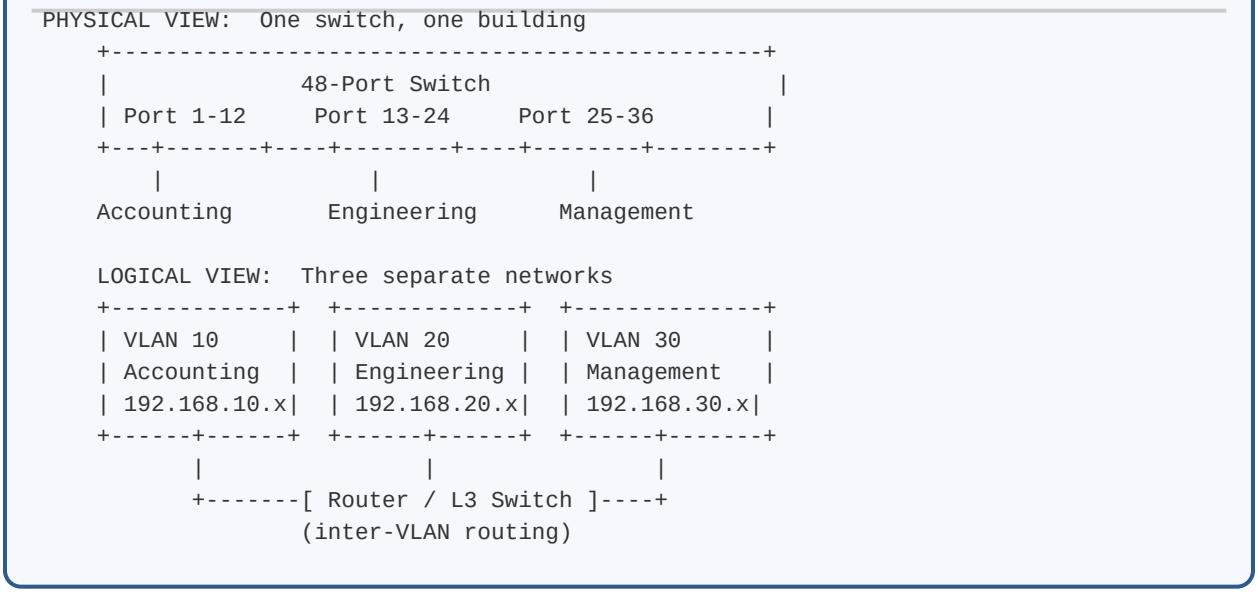

Diagram 7.1: VLANs let one physical switch act as multiple logical switches. Traffic between VLANs must be routed.

Trunk Ports

When VLANs span multiple switches, you need a way to carry traffic for all VLANs over a single cable between switches. This is called a **trunk link**. Trunk ports use 802.1Q tags to label each frame with its VLAN number. The receiving switch reads the tag and sends the frame to the correct VLAN. Access ports (regular ports) connect to end devices and carry traffic for a single VLAN—no tagging needed.

Diagram 7.2 — Access Ports vs. Trunk Ports

8. QinQ (802.1ad): VLAN Stacking for Service Providers

Standard 802.1Q supports up to 4,094 VLANs—plenty for a single organization, but not enough for a large Internet Service Provider (ISP) that serves thousands of customers, each with their own VLANs. **QinQ** (also called 802.1ad, VLAN stacking, or double-tagging) solves this by adding a second VLAN tag on top of the first.

How QinQ Works

When a customer's traffic enters the ISP's network, the ISP adds an outer VLAN tag (called the **S-Tag** or Service Tag) around the customer's existing inner VLAN tag (called the **C-Tag** or Customer Tag). This means Customer A's VLAN 10 and Customer B's VLAN 10 are kept completely separate—the outer tag distinguishes them.

Diagram 8.1 — QinQ Double Tagging

```
CUSTOMER A's frame:                CUSTOMER B's frame:
    [Dst][Src][C-Tag: VLAN 10][Data]   [Dst][Src][C-Tag: VLAN 10][Data]

                    ISP adds S-Tag (outer tag)
                             |
                             v
    INSIDE ISP NETWORK:
    Customer A: [Dst][Src][S-Tag:100][C-Tag:10][Data]   --> S-Tag 100
    Customer B: [Dst][Src][S-Tag:200][C-Tag:10][Data]   --> S-Tag 200

    Both customers use VLAN 10 internally, but the ISP
    keeps them separate using different S-Tags.
    Total VLANs possible: 4,094 x 4,094 = ~16 million!
```

Diagram 8.1: QinQ lets ISPs carry thousands of customers' VLAN traffic over a single infrastructure without conflicts.

Where QinQ Is Used

QinQ is primarily used by ISPs and large carriers to provide Ethernet services across their networks. Metro Ethernet services (connecting multiple business locations within a city) frequently rely on QinQ. It's also common in data center interconnects where multiple tenants share infrastructure.

9. Subnetting: Dividing Networks Into Pieces

As networks grow, they need to be divided into smaller, more manageable segments—just like a large city is divided into neighborhoods, each with its own zip code. This division is called **subnetting**.

The Basics of IP Addressing

An IPv4 address like **192.168.1.25** is actually a 32-bit binary number. Each of the four numbers (called octets) represents 8 bits. The address is divided into two parts: the **network portion** (which network is this?) and the **host portion** (which device on that network?). A **subnet mask** tells devices where the dividing line falls.

Diagram 9.1 — Anatomy of an IP Address and Subnet Mask

```
IP Address:   192 . 168 . 1   . 25
    In binary:    11000000.10101000.00000001.00011001

    Subnet Mask:  255 . 255 . 255 . 0     (also written as /24)
    In binary:    11111111.11111111.11111111.00000000
                  |<--- Network Part --->|  |<-Host->|

    Network:      192.168.1.0       (the 'neighborhood')
    Host:          .25              (the specific 'house')
    Usable hosts:  .1 through .254  (254 devices)
    Broadcast:     .255             (message to everyone)
```

Common Subnet Sizes

CIDR Notation	Subnet Mask	Usable Hosts	Typical Use
/8	255.0.0.0	~16 million	Giant networks (rare)
/16	255.255.0.0	~65,000	Large organizations
/24	255.255.255.0	254	Typical office/home network
/25	255.255.255.128	126	Medium department
/26	255.255.255.192	62	Small department
/27	255.255.255.224	30	Small workgroup
/28	255.255.255.240	14	Point-to-point or tiny segment
/30	255.255.255.252	2	Router-to-router link

Diagram 9.2 — Subnetting a /24 Into Four /26 Subnets

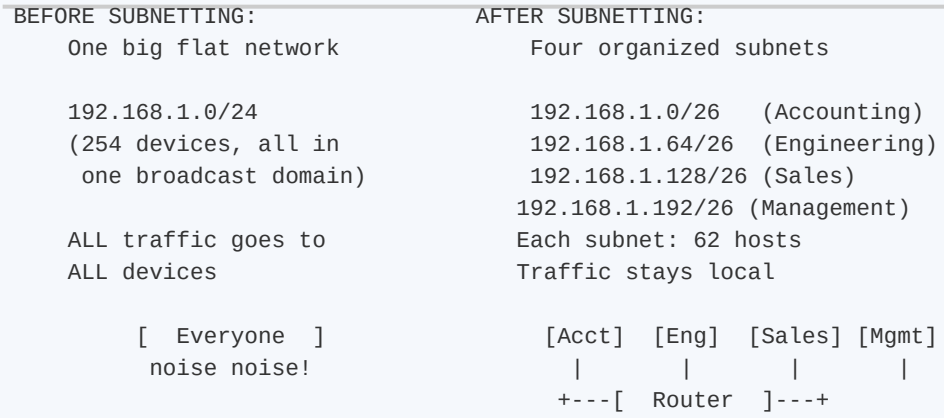

Diagram 9.2: Subnetting breaks a large flat network into smaller, more efficient segments.

10. The History of Routing Protocols

The story of routing protocols mirrors the story of computing itself: from isolated, proprietary islands to a connected, open world.

The 1960s-1970s: ARPANET and the Dawn of Routing

The very first computer network, ARPANET (1969), connected just four universities. Its routing was primitive—essentially hard-coded paths. As the network grew, the need for dynamic routing became obvious. Early research led to distance-vector algorithms (the basis for RIP) and link-state algorithms (the basis for OSPF). These two fundamental approaches still underpin all modern routing protocols.

The 1980s: The Proprietary Era

In the 1980s, networking exploded—but each vendor built its own ecosystem. Apple shipped AppleTalk with every Macintosh. Novell dominated office networking with IPX/SPX. IBM pushed Token Ring. DECnet, Banyan VINES, and XNS also competed. Each system had its own addressing, its own routing, and its own protocols. Networks couldn't easily talk to each other. It was like having five different phone systems, each incompatible with the others.

The 1990s: TCP/IP Wins

The rise of the internet forced a reckoning. TCP/IP—originally designed for ARPANET—became the universal standard. One by one, proprietary protocols were abandoned in favor of IP. Cisco rose to dominance selling routers that spoke IP. RIP (standardized in 1988), OSPF (1989), and EIGRP (1992) emerged as the primary routing protocols for IP networks.

The 2000s-Present: The Internet Age

Today, virtually all networks run IP. The routing protocol landscape has stabilized: OSPF and EIGRP handle internal routing for organizations, while BGP (Border Gateway Protocol) handles routing between organizations and across the global internet. Software-defined networking (SDN) is the newest frontier, allowing routing to be controlled by centralized software rather than configured on individual routers.

Diagram 10.1 — Routing Protocol Timeline

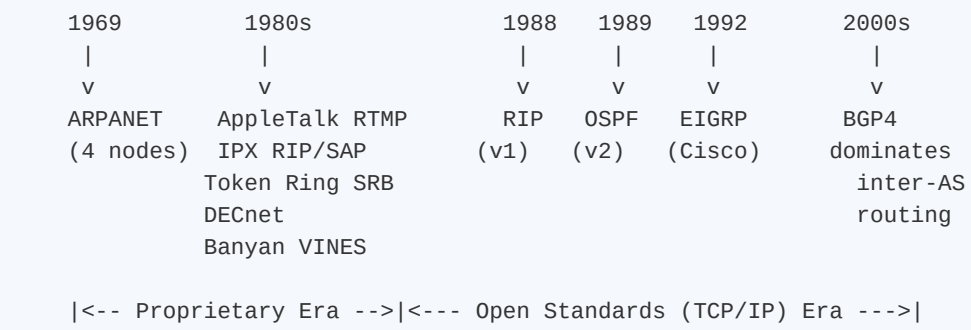

11. Legacy Protocols: AppleTalk, IPX/SPX, and Token Ring

AppleTalk (1984-2009)

When Apple released the Macintosh in 1984, it included a built-in networking system called AppleTalk. This was revolutionary: for the first time, regular people could connect computers together without an IT department. Plug in a cable, and Macs would automatically discover each other and share printers and files.

AppleTalk used its own routing protocol called **RTMP (Routing Table Maintenance Protocol)**. RTMP was a distance-vector protocol—similar in concept to RIP. Each router broadcast its entire routing table every 10 seconds. This worked well for small offices but created enormous traffic on larger networks. AppleTalk also used **NBP (Name Binding Protocol)** for device discovery and **ZIP (Zone Information Protocol)** to organize devices into named groups called zones.

AppleTalk peaked in the late 1980s and early 1990s, particularly in education and publishing. Apple began transitioning to TCP/IP in the late 1990s and finally removed AppleTalk support from Mac OS X Snow Leopard in 2009.

Diagram 11.1 — A Typical AppleTalk Network

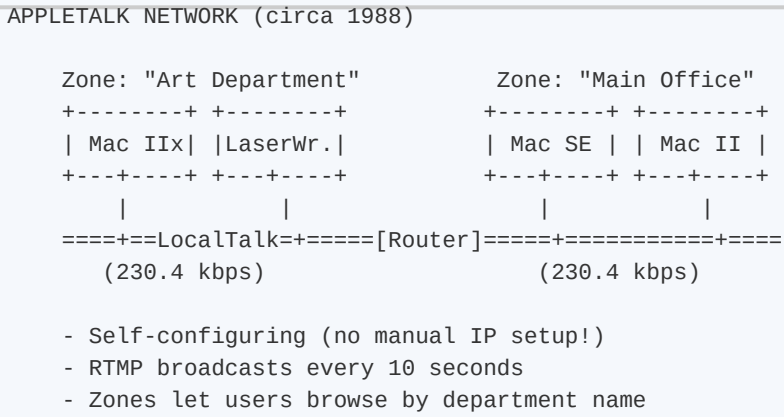

IPX/SPX (1983-~2005)

Novell's **NetWare** was the dominant network operating system in the late 1980s and 1990s. It used its own protocol suite: **IPX (Internetwork Packet Exchange)** for network-layer addressing and **SPX (Sequenced Packet Exchange)** for reliable delivery. If TCP/IP is like the modern postal system, IPX/SPX was a parallel postal system that only worked with Novell's mail carriers.

IPX used **RIP** (its own version, not the TCP/IP RIP) and **SAP (Service Advertising Protocol)** for routing. SAP was notable: servers would broadcast their available services ("I'm a file server!" "I'm a print server!") every 60 seconds. On large networks, SAP traffic could consume significant bandwidth. Novell later

introduced **NLSP (NetWare Link Services Protocol)**, a link-state protocol similar to OSPF, but by then TCP/IP had already won.

At its peak, Novell NetWare held over 60% of the network operating system market. The rise of Windows NT and TCP/IP eroded its position, and Novell eventually adopted TCP/IP as its primary transport. NetWare was discontinued in 2010.

Token Ring (1985-~2005)

IBM's Token Ring (IEEE 802.5) was a LAN technology that competed directly with Ethernet. While Ethernet used a "talk and hope for the best" approach (CSMA/CD), Token Ring was orderly: a special frame called a **token** circulated around a ring of computers. Only the computer holding the token could transmit data. When finished, it released the token to the next computer.

Token Ring ran at 4 Mbps initially and later 16 Mbps. It used **Source Route Bridging (SRB)** for connecting multiple rings—the sending computer discovered the path through the network and embedded the route in each frame. Token Ring was considered more predictable than Ethernet under heavy load because it guaranteed each device would get a turn. However, it was significantly more expensive: Token Ring network cards cost $500-$1,000 vs. $50-$100 for Ethernet. Ethernet's lower cost and simpler architecture ultimately won the market.

Diagram 11.2 — Token Ring vs. Ethernet Operation

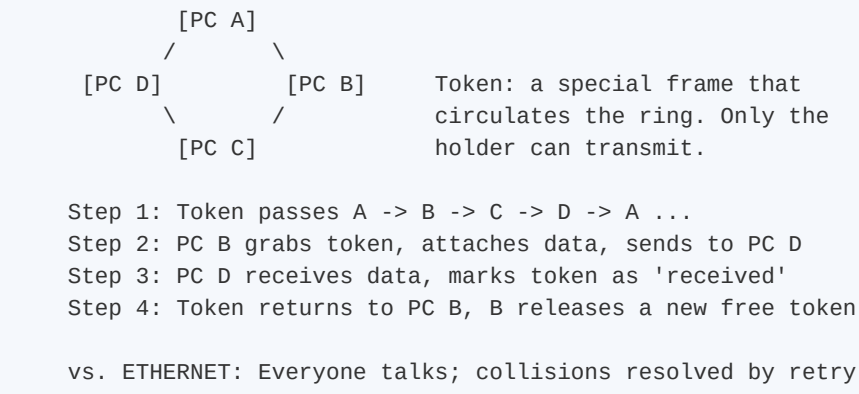

12. RIP: The First IP Routing Protocol

RIP (Routing Information Protocol) is the granddaddy of IP routing protocols. Its roots trace back to the routing algorithm in Xerox's XNS (Xerox Network Systems) from the late 1970s. RIP was formally standardized for TCP/IP in **RFC 1058 (1988)**. Despite its age, RIP is still used in small networks and remains an essential learning tool.

How RIP Works: Distance-Vector Routing

RIP is a **distance-vector** protocol. Each router maintains a table of known networks and their distance (measured in hops—the number of routers a packet must cross). Every 30 seconds, each router broadcasts its entire routing table to its neighbors. When a router learns about a shorter path, it updates its table. The maximum hop count is 15; a distance of 16 means "unreachable."

Diagram 12.1 — RIP Hop Count and Convergence

```
RIP HOP COUNT EXAMPLE

    [Net A]---[R1]---[R2]---[R3]---[Net B]

    R1 says: 'Net A is 1 hop, Net B is 3 hops'
    R2 says: 'Net A is 2 hops, Net B is 2 hops'
    R3 says: 'Net A is 3 hops, Net B is 1 hop'

    If R2-R3 link fails:
    [Net A]---[R1]---[R2]   X   [R3]---[Net B]
    R2 now says: 'Net B is unreachable (16 hops)'
    RIP slowly converges (can take minutes!)
```

Feature	RIP v1 (1988)	RIP v2 (1998)
Addressing	Classful only	Classless (CIDR supported)
Updates	Broadcast	Multicast (224.0.0.9)
Authentication	None	MD5 supported
Max hops	15	15
Update interval	30 seconds	30 seconds

RIP's Strengths and Weaknesses

RIP's great strength is simplicity—it's easy to configure and understand. Its weaknesses are significant: slow convergence (it can take minutes to react to changes), bandwidth waste (broadcasting full tables every 30 seconds), and the 15-hop limit. For any network larger than a handful of routers, OSPF or EIGRP is a far better choice.

13. EIGRP: Cisco's Intelligent Protocol

EIGRP (Enhanced Interior Gateway Routing Protocol) was developed by Cisco Systems in **1992** as an evolution of the older IGRP. For many years it was Cisco-proprietary, meaning only Cisco routers could use it. In 2013, Cisco published EIGRP as an informational RFC, but it remains most common in Cisco-centric environments.

How EIGRP Works: Advanced Distance-Vector

EIGRP is sometimes called an **advanced distance-vector** or **hybrid** protocol because it combines features of both distance-vector and link-state approaches. Instead of just counting hops, EIGRP uses a composite metric based on **bandwidth**, **delay**, **reliability**, and **load**. This means it can choose a slower but less congested path over a fast but overloaded one.

EIGRP uses the **DUAL (Diffusing Update Algorithm)** to ensure loop-free routing and extremely fast convergence. When a route fails, EIGRP can often switch to a backup route instantly because it pre-computes alternative paths. It only sends updates when something changes (no periodic broadcasts like RIP), saving bandwidth.

Diagram 13.1 — EIGRP Metric and Feasible Successor

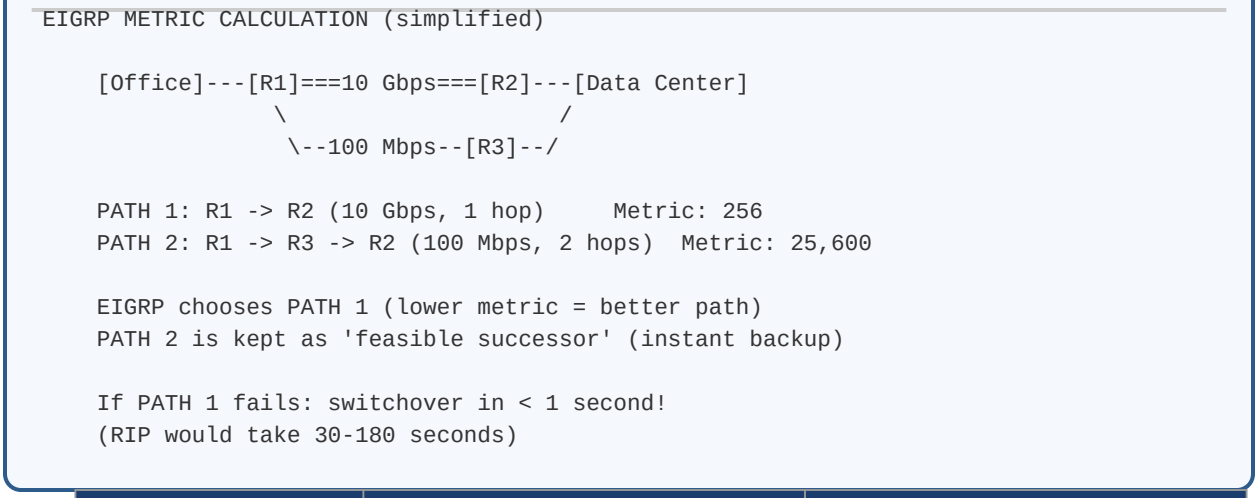

Feature	EIGRP	RIP v2
Metric	Bandwidth + Delay (composite)	Hop count only
Max network size	No hard limit	15 hops
Convergence	Sub-second	Minutes
Updates	Only on changes	Every 30 seconds
Backup routes	Pre-computed (instant failover)	None
Vendor	Cisco (primarily)	Open standard

14. OSPF: The Industry Standard

OSPF (Open Shortest Path First) is the most widely used interior routing protocol in the world. Developed by the IETF (Internet Engineering Task Force) and first standardized in **RFC 1131 (1989)**, OSPF was designed from the ground up to handle large, complex networks. The current version, OSPFv2, is defined in **RFC 2328 (1998)**. OSPFv3 extends the protocol to support IPv6.

How OSPF Works: Link-State Routing

Unlike RIP (which only knows distances), OSPF is a **link-state** protocol—every router builds a complete map of the entire network. Each router advertises information about its directly connected links (neighbors, costs, status) using **LSAs (Link-State Advertisements)**. All routers collect these LSAs into a **Link-State Database (LSDB)**, which is identical across the network. Each router then independently runs **Dijkstra's Shortest Path First algorithm** on this database to calculate the best route to every destination.

Diagram 14.1 — OSPF Shortest Path Calculation

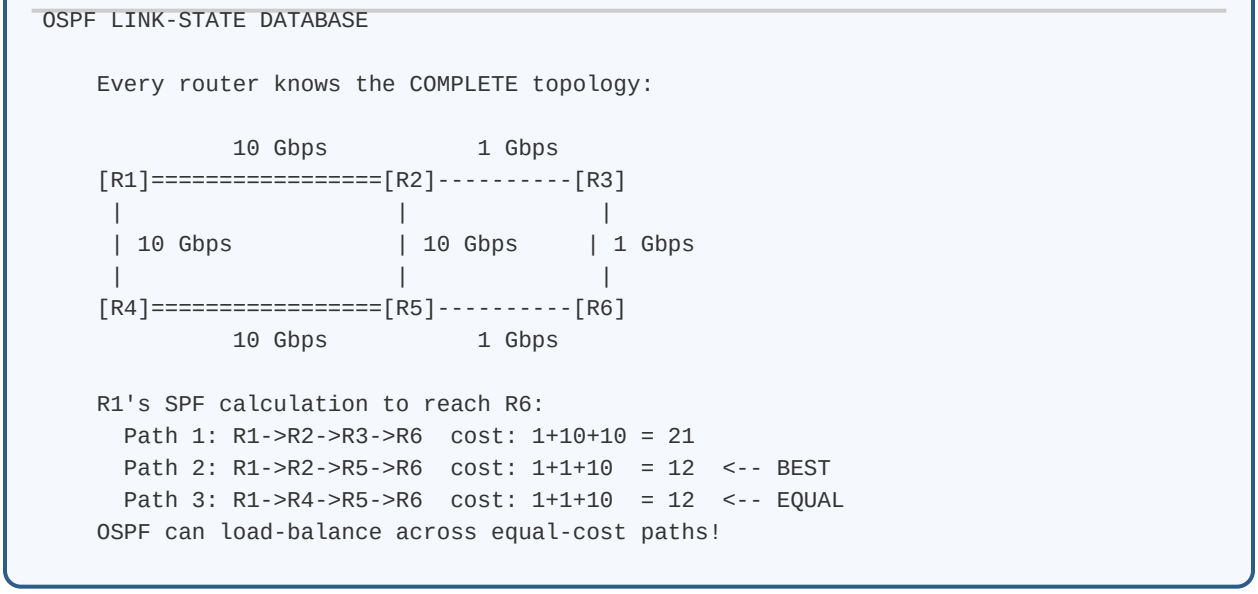

OSPF Areas: Dividing Large Networks

In very large networks, having every router know every link would create too much overhead. OSPF solves this with **areas**. An OSPF network is divided into areas, each maintaining its own link-state database. **Area 0** (the backbone) connects all other areas. Routers at area boundaries (**ABRs—Area Border Routers**) summarize routing information between areas, reducing the amount of data each router must process.

Diagram 14.2 — OSPF Multi-Area Design

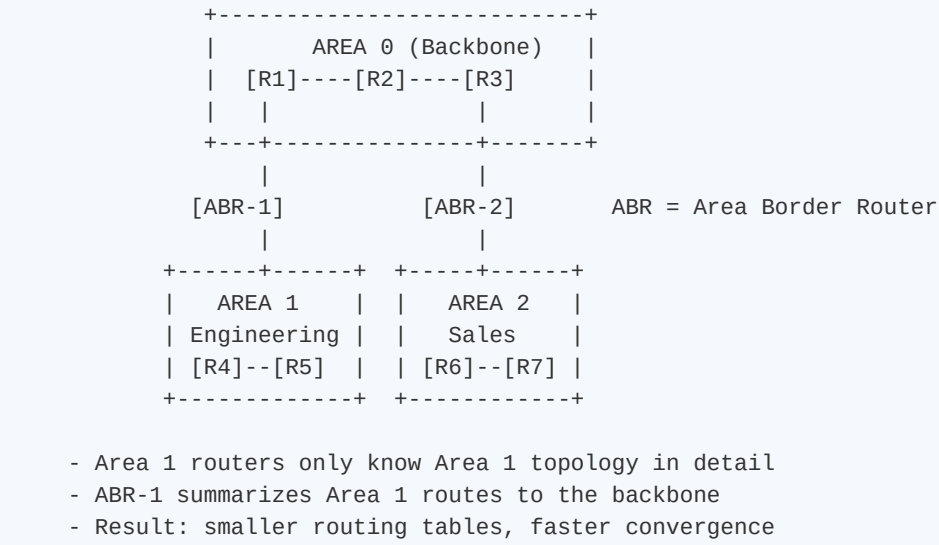

Diagram 14.2: OSPF areas reduce complexity in large networks by limiting the scope of detailed topology information.

Feature	OSPF	Why It Matters
Algorithm	Dijkstra (SPF)	Guarantees shortest path
Metric	Cost (based on bandwidth)	Prefers faster links
Convergence	Very fast (seconds)	Quick recovery from failures
Scalability	Excellent (areas)	Supports thousands of routers
Standard	Open (IETF RFC 2328)	Works on any vendor's equipment
Equal-cost paths	Supported	Load balancing across links
Authentication	MD5 / SHA supported	Prevents rogue routers

15. BGP: The Protocol That Runs the Internet

While OSPF and EIGRP route traffic *within* an organization, **BGP (Border Gateway Protocol)** routes traffic *between* organizations—between ISPs, data centers, and the global internet. BGP is literally the protocol that holds the internet together.

Autonomous Systems

The internet is divided into **Autonomous Systems (AS)**—independently operated networks, each identified by a unique AS number. Your ISP is an AS. Google, Amazon, and Netflix each have their own AS numbers. BGP is how these autonomous systems exchange routing information: "I can reach these networks; here is the path to get there through me."

Diagram 15.1 — BGP Connects Autonomous Systems

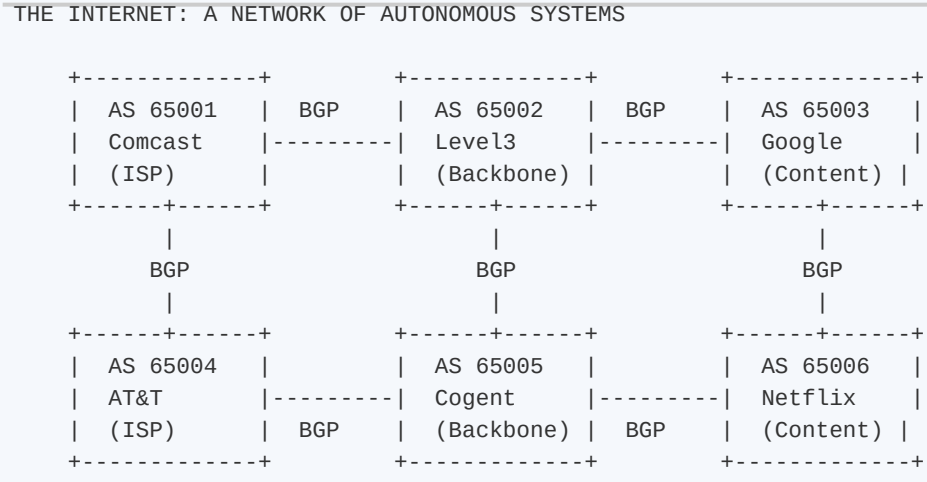

BGP is a **path-vector** protocol. Instead of just knowing the distance to a network, BGP knows the complete list of autonomous systems a packet must traverse. This allows BGP to make policy-based decisions: an ISP might prefer to send traffic through a partner rather than a competitor, even if the competitor's path is shorter.

The current version, **BGP-4**, has been the backbone of internet routing since 1994. As of today, the global BGP routing table contains over **1 million routes**. A single BGP misconfiguration can (and has) taken down large portions of the internet—it is arguably the most critical protocol in all of networking.

16. Internet Service Providers and How the Internet Works

The internet is not one single network—it's a network of networks, connected together by ISPs (Internet Service Providers) at various levels.

The ISP Hierarchy

Tier	Description	Examples
Tier 1	Global backbone providers. They peer with each other for free (settlement-free peering) and reach every part of the internet without paying anyone.	Lumen (CenturyLink), Cogent, NTT, Arelion (Telia)
Tier 2	Regional providers. They pay Tier 1 providers for some routes (transit) and peer freely for others.	Comcast, AT&T, Deutsche Telekom, Vodafone
Tier 3	Local ISPs. They buy all their internet access from Tier 1 or Tier 2 providers (transit customers).	Local cable companies, regional DSL providers, small fiber companies

Diagram 16.1 — The Path of Data Through the ISP Hierarchy

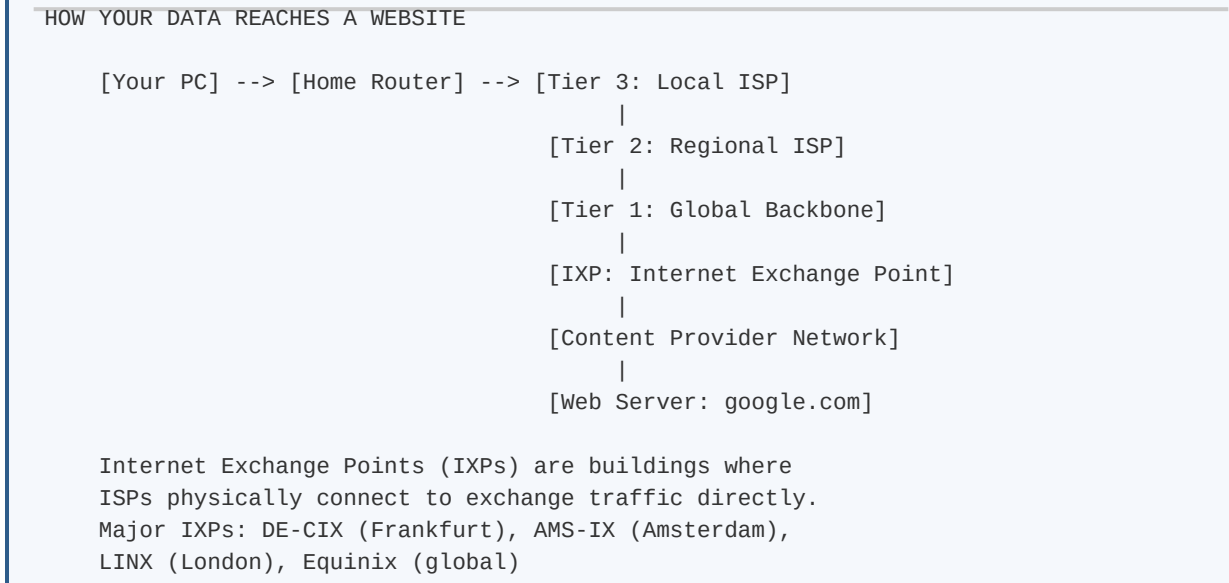

Peering vs. Transit

Peering is when two ISPs agree to exchange traffic directly, usually for free (settlement-free). This benefits both parties because it keeps traffic local and fast. **Transit** is when a smaller ISP pays a larger one for access to the wider internet. Most Tier 3 ISPs buy transit from Tier 2 providers, who in turn may buy transit from Tier 1 providers.

Content Delivery Networks (CDNs)

Companies like Netflix, Google, and Amazon don't want their content to travel across half the world. Instead, they place servers called **CDN nodes** inside or near ISPs' data centers. When you stream a Netflix movie, it likely comes from a server just miles away, not from Netflix's headquarters. This reduces latency, saves bandwidth for ISPs, and improves your experience.

17. Connection Speeds Through the Decades

The history of networking is a story of exponentially increasing speeds. To appreciate how far we've come, let's walk through the major milestones.

Understanding the Units

Baud measures signal changes per second (not always equal to bits per second). **bps (bits per second)** measures actual data throughput. **Kbps** = thousands of bits/second. **Mbps** = millions. **Gbps** = billions. To convert to familiar file sizes: divide by 8 (1 byte = 8 bits). So 100 Mbps = about 12.5 megabytes per second.

Speed	Era	Technology	Time to Download 1 GB File	Real-World Feel
300 baud	1960s-70s	Acoustic coupler modem	~309 days	Watching paint dry—character by character
9,600 baud	1970s-80s	External modem (Hayes Smartmodem)	~9.6 days	A page of text every few seconds
14.4 Kbps	1991	V.32bis modem	~6.4 days	Could barely load a simple web page
28.8 Kbps	1994	V.34 modem	~3.2 days	The early web: text with small images
56 Kbps	1996	V.90 dial-up modem	~1.6 days	Remember the modem screech? Loading an MP3 took 10+ minutes
1.5 Mbps	1999	T1 line / early DSL	~1.5 hours	Business internet. A revelation!
10 Mbps	1983	10BASE5 Ethernet (LAN only)	~13 min	First LAN speed standard
100 Mbps	1995	Fast Ethernet (LAN/broadband)	~80 sec	The standard office network for a decade
1 Gbps	1999	Gigabit Ethernet	~8 sec	Download an HD movie in seconds
10 Gbps	2002	10G Ethernet (data centers)	~0.8 sec	Enterprise and cloud computing
40 Gbps	2010	40G Ethernet (data centers)	~0.2 sec	Spine-leaf data center fabrics
100 Gbps	2010	100G Ethernet (backbone/DC)	~0.08 sec	ISP backbone and hyperscale DCs
400 Gbps	2018	400G Ethernet (cutting edge)	~0.02 sec	Next-gen data center interconnects

Diagram 17.1 — Network Speed Growth Over 50 Years

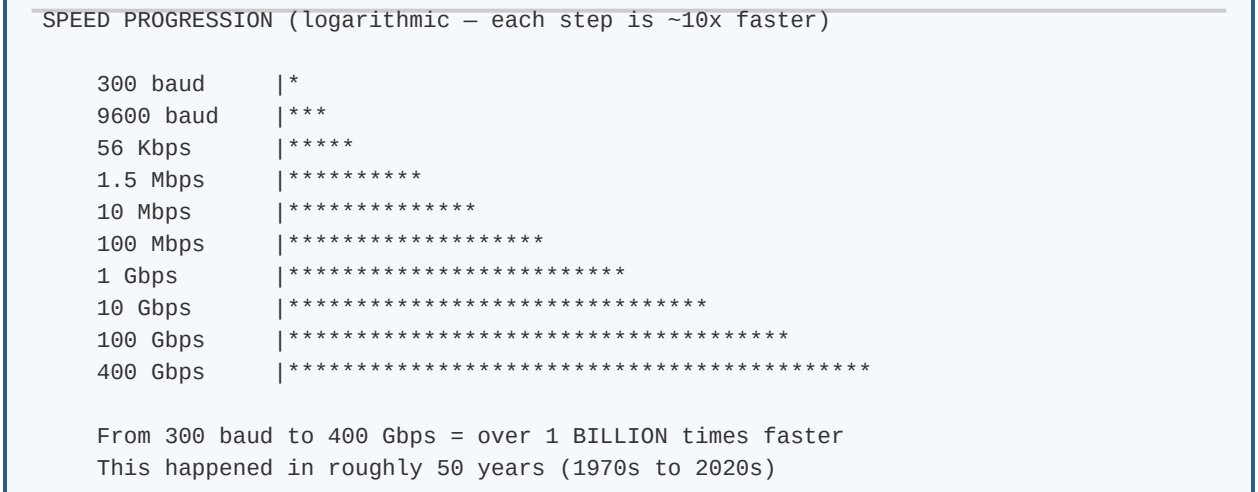

The Dial-Up Experience

For anyone who grew up with broadband, it's hard to imagine the dial-up era. At 56 Kbps, a single MP3 song (5 MB) took about 12 minutes to download. Web pages were designed to be tiny—images were heavily compressed, and sites like Google became popular partly because their homepage was just a search box (fast to load). Streaming video was unthinkable. You'd hear the famous modem handshake screech every time you connected, and picking up the phone would disconnect the internet.

The Broadband Revolution

DSL and cable modems (1-10 Mbps) in the early 2000s changed everything. Suddenly, music could be streamed. YouTube launched in 2005 because broadband had finally made video streaming viable. Netflix transitioned from mailing DVDs to streaming in 2007. Each speed increase unlocked new applications that were previously impossible.

18. Fiber Optics: The Undisputed Future

If there is one technology that will define the future of networking, it is **fiber optics**. While copper cables transmit data as electrical signals, fiber optic cables use **pulses of light** traveling through strands of ultra-pure glass thinner than a human hair. The advantages are so overwhelming that fiber is replacing copper everywhere—from undersea intercontinental cables to the last mile into your home.

How Fiber Optics Work

A fiber optic cable contains one or more glass fibers, each about 125 micrometers in diameter (slightly thicker than a human hair). A laser or LED at one end converts electrical data into pulses of light. The light bounces along the inside of the glass fiber through a principle called **total internal reflection**—the light literally bounces off the walls of the fiber, trapped inside, traveling at about 200,000 km/s (two-thirds the speed of light in vacuum).

Diagram 18.1 — Single-Mode vs. Multi-Mode Fiber

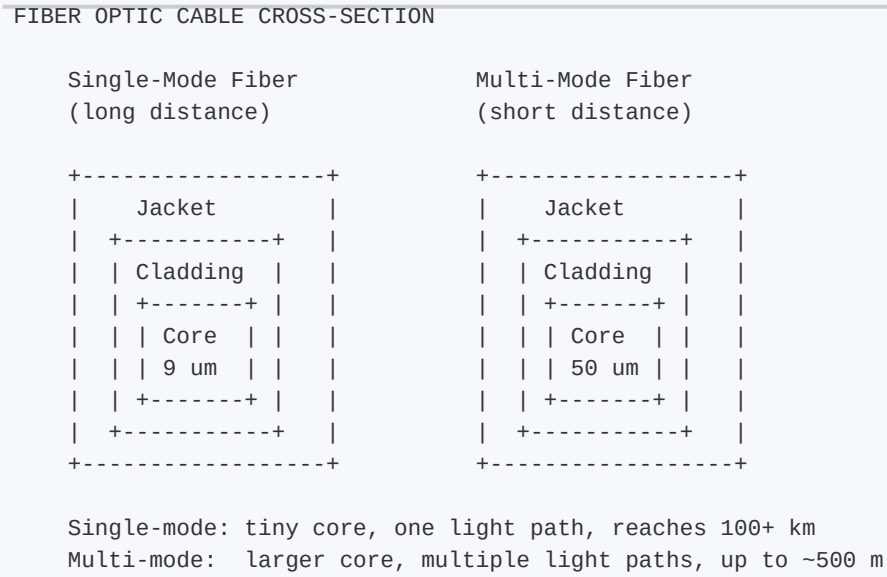

Why Fiber Is Superior to Copper in Every Way

Characteristic	Copper (Cat 6)	Fiber (Single-Mode)	Winner
Max speed	10 Gbps	400+ Gbps (and growing)	Fiber
Max distance	100 meters	100+ kilometers	Fiber (1000x)
Electromagnetic interference	Susceptible	Completely immune	Fiber
Signal loss (attenuation)	High (needs repeaters)	Very low	Fiber

Bandwidth capacity	Limited by copper physics	Virtually unlimited (WDM multiplexing)	Fiber
Weight	Heavy (copper is dense)	Very light (glass)	Fiber
Security	Can be tapped electrically	Extremely hard to tap (light leaks are detectable)	Fiber
Durability	Corrodes, affected by moisture and temperature	Glass does not corrode; immune to weather	Fiber
Cost (cable)	Low	Medium (dropping fast)	Copper (for now)
Cost (total ownership)	Higher long-term (more maintenance, cooling)	Lower long-term (less power, less space)	Fiber

Wavelength Division Multiplexing (WDM)

One of fiber's most remarkable capabilities is **WDM (Wavelength Division Multiplexing)**. Different colors (wavelengths) of light can travel through the same fiber simultaneously without interfering with each other—like multiple lanes on a highway. Dense WDM (DWDM) can carry **80 or more channels** on a single fiber, each at 100 Gbps or more. This means a single fiber pair can carry **8+ terabits per second**. No copper technology comes even remotely close.

Diagram 18.2 — WDM: Multiplying Fiber Capacity

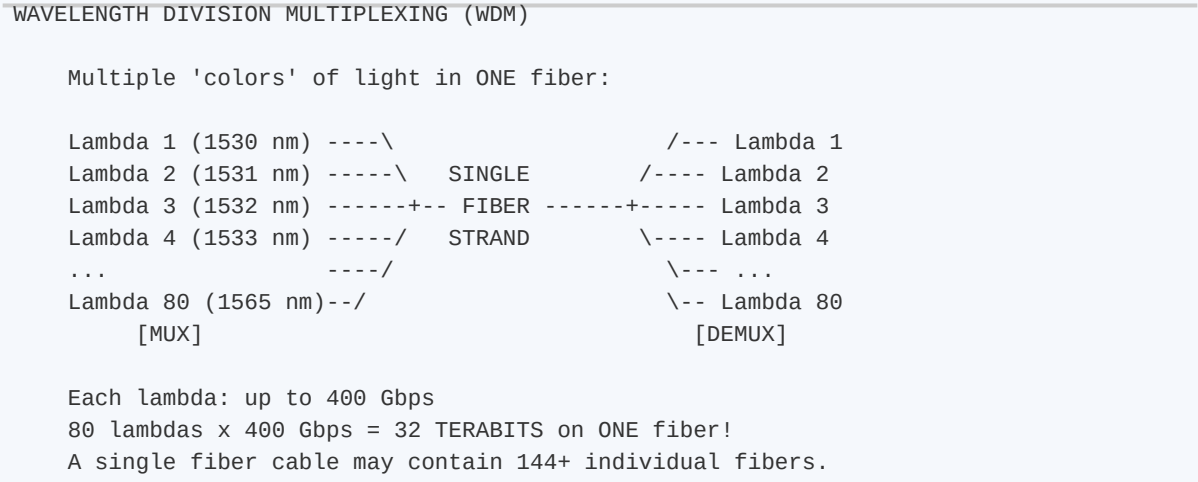

Fiber to the Home (FTTH)

For decades, fiber was limited to long-distance and enterprise connections because the cost of running fiber to individual homes was prohibitive. That's changing rapidly. **FTTH (Fiber to the Home)** deployments are accelerating worldwide. Technologies like **GPON (Gigabit Passive Optical Network)** and **XGS-PON (10G Symmetric PON)** allow a single fiber from the ISP to be split among 32-128 homes, making deployment economical.

Countries leading in FTTH include South Korea (95%+ fiber penetration), Japan, UAE, Singapore, and the Scandinavian countries. In the US, companies like AT&T;, Google Fiber, and Verizon FiOS are expanding fiber rapidly. Once installed, fiber lasts 25-50 years and can be upgraded to higher speeds simply by changing the equipment at both ends—the glass itself doesn't need replacement.

Undersea Fiber Cables

The global internet is connected by approximately 550 active undersea fiber optic cables spanning over 1.4 million kilometers. These cables carry **99% of intercontinental data traffic** (satellites carry less than 1%). Modern undersea cables like Google's "Grace Hopper" (connecting the US, UK, and Spain) carry over 340 Tbps. These cables are engineering marvels—laid on the ocean floor by specialized ships, sometimes at depths of 8,000 meters.

The Future Belongs to Fiber

Fiber optics is not just better than copper—it's in a completely different league. Copper has reached its physical limits: 10 Gbps over 100 meters is about as good as it gets. Fiber's limits are nowhere in sight: researchers have achieved 1 petabit per second (1,000 terabits) on a single fiber in laboratory conditions. As 5G, cloud computing, AI, virtual reality, and IoT drive ever-increasing bandwidth demands, fiber is the only technology that can keep up. The question is not whether fiber will replace copper, but how quickly.

19. Putting It All Together: A Packet's Journey

Let's trace the path of a single packet as it travels from your laptop to a web server across the country, using everything we've learned.

Diagram 19.1 — Complete Packet Journey Across the Internet

```
YOUR LAPTOP (192.168.10.50, VLAN 10)
         |
         | (1) Laptop creates packet: Src=192.168.10.50,
         |     Dst=93.184.216.34 (web server)
         |
    [Office Switch] (Layer 2)
         | (2) Switch reads MAC address, forwards frame
         |     within VLAN 10 to the router port
         |
    [Office Router] (Layer 3, runs OSPF)
         | (3) Router looks up destination in OSPF routing table.
         |     Best path: send to ISP via fiber uplink.
         |
    [ISP Edge Router] (Tier 3 ISP, runs BGP)
         | (4) ISP's BGP table shows the destination is in
         |     AS 65002. Best path: via Tier 2 provider.
         |
    [Tier 2 ISP Router] (BGP)
         | (5) Forwards via Tier 1 backbone.
         |     Travels over fiber at ~200,000 km/s.
         |
    [Tier 1 Backbone Router] (BGP)
         | (6) Crosses undersea fiber cable or long-haul fiber.
         |
    [Destination ISP / CDN Edge]
         | (7) Enters destination network.
         |
    [Web Server: 93.184.216.34]
         | (8) Server receives packet, sends response back.
         |     The whole round trip: ~20-100 milliseconds!
```

Diagram 19.1: A packet's journey involves Layer 2 switching (MAC addresses, VLANs), Layer 3 routing (OSPF inside the organization, BGP between ISPs), and physical transmission over fiber.

This entire journey—through switches, routers, multiple ISPs, across fiber optic cables, and back again—typically takes less than 100 milliseconds. The protocols described in this guide (Ethernet, VLANs, OSPF, BGP) are all working simultaneously, at every step, to make this possible. It is one of humanity's most remarkable engineering achievements, yet it happens billions of times per second, invisibly, all around the world.

20. Glossary of Key Terms

Term	Definition
ARP	Address Resolution Protocol — maps IP addresses to MAC addresses on a LAN
AS	Autonomous System — an independently managed network with its own routing policy
BGP	Border Gateway Protocol — routes traffic between autonomous systems (the internet backbone)
Baud	Signal changes per second on a communication line; in early modems, roughly equal to bps
Broadcast	A message sent to all devices on a network segment
CAM Table	Content Addressable Memory table — a switch's mapping of MAC addresses to ports
CDN	Content Delivery Network — servers placed near users to speed up content delivery
CIDR	Classless Inter-Domain Routing — flexible IP address allocation (e.g., /24, /26)
CSMA/CD	Carrier Sense Multiple Access / Collision Detection — Ethernet's original access method
DWDM	Dense Wavelength Division Multiplexing — many light channels on one fiber
EIGRP	Enhanced Interior Gateway Routing Protocol — Cisco's advanced routing protocol
Ethernet	The dominant LAN technology, standardized as IEEE 802.3
FTTH	Fiber to the Home — fiber optic cable delivered directly to residences
GPON	Gigabit Passive Optical Network — technology for sharing fiber among many homes
Hop	One step between routers on a network path
IP	Internet Protocol — the Layer 3 addressing scheme used by all internet traffic
IPX/SPX	Novell's proprietary protocol suite (obsolete)
ISP	Internet Service Provider — a company that provides internet access
IXP	Internet Exchange Point — physical location where ISPs connect and exchange traffic
LAN	Local Area Network — a network covering a small area (building/campus)
LSA	Link-State Advertisement — OSPF messages describing a router's connections
MAC	Media Access Control address — a hardware identifier burned into network interfaces
MPLS	Multiprotocol Label Switching — WAN technology using labels for fast forwarding
OSPF	Open Shortest Path First — the industry-standard link-state routing protocol
Packet	A unit of data with headers containing addressing information
QinQ	Double VLAN tagging (802.1ad) used by service providers
RIP	Routing Information Protocol — the simplest IP routing protocol (hop-count based)
Router	A device that forwards packets between different networks using IP addresses
SD-WAN	Software-Defined WAN — WAN managed by centralized software controllers
STP	Spanning Tree Protocol — prevents loops in switched networks
Subnet	A subdivision of an IP network, created using a subnet mask
Switch	A device that forwards frames within a LAN using MAC addresses
TCP/IP	Transmission Control Protocol / Internet Protocol — the protocol suite of the internet

Token Ring	IBM's LAN technology using token passing (obsolete)
Trunk	A switch port carrying traffic for multiple VLANs using 802.1Q tags
VLAN	Virtual LAN — a logical network segmentation within a physical switch
WAN	Wide Area Network — a network spanning large geographic distances
WDM	Wavelength Division Multiplexing — sending multiple light colors through one fiber

End of Document

This guide was generated to provide a comprehensive, accessible overview of networking technologies from the earliest dial-up connections to modern fiber optics.

Generated: February 25, 2026