

Networking Infrastructure

Reading time: ~10 min · Prereqs: Lessons 01–07 · *Snapshot 2026 — physical standards evolve slowly; SD-WAN and observability tooling moves faster.*

Why this matters to you

Every digital thing you have ever used — a web page, a Slack message, a video call, a payment going through — travelled over a network at least once. Networking is the infrastructure underneath all other infrastructure, and yet it's the layer most non-technical people treat as magic. That opacity is a problem. When a product manager says "it's slow," when a cloud bill triples for no obvious reason, when a security team says "we found lateral movement in the network" — understanding what a network actually is turns those phrases from noise into something actionable. This lesson is the mental model: the road analogy, the three very different worlds networks live in, and the mechanisms — VPN, SD-WAN, monitoring — that engineers use to manage them. No configuration, no commands. Just the map.

1. Streets, highways, and addressing

The cleanest analogy for a network is a road system. **Data travels in packets** — small, numbered envelopes each carrying a fragment of the message. The packet knows where it came from and where it's going. **Routers** are post offices: they look at the destination address on each envelope and decide which road leads toward it. No single router needs to know the full route; it only needs to know the next hop, and the next post office takes over from there.

The physical medium underneath can be copper (Ethernet cables, the Cat5/6/7 patch leads you see in offices), glass (**fiber optic**, which carries pulses of light and can run for hundreds of kilometres without signal degradation), radio waves (**Wi-Fi** in the 2.4, 5, and 6 GHz bands inside buildings), cellular radio (**4G/5G** via towers), or microwave beams bounced off satellites (**Starlink** and similar). The physics differ enormously — fiber is dense and fast; satellite adds 500–600 ms of delay because the signal

physically travels to orbit and back — but the addressing system on top is the same everywhere.

Every device has two kinds of address. A **MAC address** is burned into the network card at the factory: a 48-bit serial number that looks like `3C:22:FB:4A:01:0D`. It works only on your local network segment — the equivalent of a furniture serial number that is only meaningful in the building where the furniture lives. An **IP address** is assigned by the network and can change. Your home router gives every device a *private* IP address in the reserved ranges like `192.168.x.x` or `10.x.x.x`. Your internet provider gives your router a single *public* IP that the rest of the internet sees. The mechanism that maps many private addresses to one public IP is called **NAT** (Network Address Translation) — think of it as an apartment building with one street address and an internal mailroom that routes parcels to the right flat.

2. Three network worlds: SOHO, Enterprise, Data Center

The physics of networking are the same at every scale, but the engineering discipline required grows by orders of magnitude. It is useful to separate the world into three contexts.

SOHO (Small Office / Home Office) is the living-room router most people know. One box does everything: it is the router, the Wi-Fi access point, the firewall, and usually the DHCP server that assigns IP addresses. Top speed is typically 1 Gbps on the wired side, maybe 2–6 Gbps on Wi-Fi 6E. If it breaks, you reboot it. The entire network fits in the palm of your hand.


Enterprise networks cover an office campus or a company's national footprint. Here, dedicated switches connect hundreds or thousands of devices. The network is carved into **VLANs** (Virtual Local Area Networks) — logical segments that keep, say, the finance team's traffic separated from the guest Wi-Fi. Redundant uplinks ensure that a single cable cut does not take down a floor. A dedicated team writes change-control tickets before touching anything, because a misconfiguration in a core switch can silence an entire building.

Data-center networks are a different species. A large hyperscaler like Meta or Google runs tens of thousands of servers per building, connected at **100 Gbps or 400 Gbps** per link. The topology is typically *spine-leaf*: every leaf switch (serving a rack of servers) connects to every spine switch, so no path through the fabric is more than two

hops. Redundancy is N+2: two complete failures can happen simultaneously without affecting service. Links, switches, power supplies, and cooling circuits are all doubled or tripled. A data-center network engineer writes code — configurations are pushed via APIs and version-controlled like software — not clicked through a browser UI.

3. The three worlds, one diagram

Here is a side-by-side sketch of the three network worlds: their physical scale, the devices involved, and the gap in management complexity. The leftmost column is the living-room router. The rightmost is a data-center spine-leaf fabric.



4. Going private through public pipes: VPN and anonymizers

A **VPN** (Virtual Private Network) is an encrypted tunnel punched through a public network so that you can act as if you were physically inside a different, private one. Think of it as a pneumatic tube inside the postal system: your envelope goes into a sealed, opaque tube at one end, travels invisibly through the public mail system, and emerges at the other end as if it had never left the building. Nobody watching the public mail can read the contents — or even tell what building it came from.

Corporate VPNs (WireGuard, Cisco AnyConnect, Palo Alto GlobalProtect) solve a specific, legitimate problem: a remote worker in a coffee shop needs to reach internal company resources that are not exposed to the public internet. The VPN client on their laptop creates an encrypted connection back to the company's VPN gateway; from that point on the laptop looks as if it is sitting on the office LAN. This is also why IT

departments can enforce security policies on remote workers — once inside the tunnel, the traffic flows through the same firewalls and monitoring tools as in the office.

Consumer "VPN" services (NordVPN, ExpressVPN, Mullvad, and dozens of others) use the same technology for a different goal: hiding your IP address from the websites you visit. Your traffic exits not from your home, but from the VPN vendor's server in whichever country you chose. This is useful for bypassing geographic content restrictions, and it does prevent your ISP from seeing your browsing history. But it does not make you anonymous: the VPN vendor now sees everything your ISP used to see. You have moved trust, not eliminated it. The marketing phrase "military-grade encryption" is technically correct (AES-256 is the same cipher used by governments) and practically meaningless for most consumer use-cases where the threat model does not include governments decrypting traffic.

Tor (The Onion Router) is a different architecture designed for stronger anonymity. Your traffic passes through three volunteer-operated *relay nodes* worldwide, each of which strips one layer of encryption — like peeling an onion. The entry node knows who you are but not your destination. The exit node knows your destination but not who you are. No single node sees the full picture. The trade-off is speed: three relays, often on different continents, add 200–600 ms of latency. Tor Browser is a modified Firefox pre-configured to use the Tor network; browsing ipinfo.io from it shows a random exit node's IP in a country you may never have visited.

5. SD-WAN: software-defined networks for enterprises

Until the mid-2010s, companies with multiple office locations connected them with **MPLS** (Multiprotocol Label Switching) circuits — private, carrier-managed links guaranteed to be fast and reliable. The catch: expensive, slow to provision (weeks to months), and inflexible. A new office in a city meant waiting for the carrier to run a circuit.

SD-WAN (Software-Defined Wide Area Network) replaced MPLS with intelligence layered on top of ordinary broadband internet links. An SD-WAN device at each office constantly measures the quality of all available paths (ISP A, ISP B, cellular 5G backup) and routes traffic intelligently: video calls over the lowest-latency path, bulk backups over the cheapest path, financial transactions over the most reliable. A new branch office needs only a box that is shipped pre-configured; it calls home to the SD-WAN

cloud controller, downloads its policy, and is live in minutes. Companies like Cisco (Viptela), VMware (VeloCloud), and Palo Alto (Prisma SD-WAN) dominate the market as of 2026. Cost savings of 60–80% over MPLS are common, which is why the enterprise WAN market flipped quickly.

6. Watching the pipes: monitoring and administration

A network that nobody is watching is a network that fails silently. Network operations (NetOps) is a continuous practice of measurement, alerting, and planned maintenance.

The oldest tool in the kit is **SNMP** (Simple Network Management Protocol), invented in 1988 and still running in almost every enterprise network in 2026. Each network device (switch, router, firewall) runs an SNMP agent that reports metrics — port up/down status, error counts, CPU load, bandwidth utilisation — to a central **NMS** (Network Management System). Think of SNMP as each post office in the road network sending a hourly status report to a head-office dashboard.

NetFlow (Cisco) and its open-source equivalent **sFlow** take a different approach: instead of health metrics, they sample actual traffic. Every five minutes you get a report that says "these 20 IP pairs exchanged these volumes of traffic on these ports." That is how a network engineer discovers that a server is sending 40 GB of data to an IP address in an unexpected country at 2 AM — or that a developer's test job is accidentally saturating the office uplink.

Modern observability stacks layer dashboards (Grafana), alerting (Prometheus / Alertmanager), and log aggregation (Elasticsearch or Loki) on top of these feeds. When an alert fires, the on-call network engineer can open a dashboard, immediately see which link or device is misbehaving, correlate it with a recent change, and escalate or fix. This is not magic; it is years of work building the instrumentation. The networks that have it are dramatically easier to operate than the ones that don't.

What people argue about

- **VPN as privacy vs VPN as security theater.** Security professionals generally agree that consumer VPNs are oversold: they protect one threat model (ISP visibility) while doing little for others (browser fingerprinting, tracking pixels, account-level tracking). The counter-argument is that moving trust away from

your ISP — which may log and sell browsing data — to a vendor whose business model depends on not being caught logging is still a meaningful improvement for many users.

- **SD-WAN vs SASE.** As of 2026 the industry is actively debating whether SD-WAN is a stepping stone or a destination. **SASE** (Secure Access Service Edge) merges SD-WAN networking with cloud-delivered security (firewall-as-a-service, zero-trust access). Some vendors say SASE makes standalone SD-WAN obsolete; others say SASE is SD-WAN with a marketing rebrand and a larger bill.
- **IPv4 exhaustion vs IPv6 adoption.** The internet ran out of IPv4 addresses years ago; the fix is IPv6, which has astronomical address space. Adoption is slow because it requires updating every device and software stack. Most networks as of 2026 run "dual-stack" — both protocols simultaneously — in an awkward middle period that has lasted two decades longer than anyone planned.

Reading on

- **Computer Networking: A Top-Down Approach** by Kurose & Ross — the standard university textbook, surprisingly readable for the first three chapters even without a CS background. Chapters 1–2 cover the material in this lesson at greater depth.
- **Cloudflare Learning Center** — cloudflare.com/learning. Free. Cloudflare publishes extremely clear explainers on every networking topic: DNS, TLS, DDoS, VPNs, SD-WAN. Start with "What is a VPN?" and "What is SD-WAN?" — both are accurate and ad-free.
- **The Tor Project** — torproject.org/about/history. The project's own history page explains the design rationale plainly and is written for non-technical readers.
- **Julia Evans's "How do computers get IP addresses?"** — jvns.ca. Evans is remarkable at making low-level networking concepts feel approachable. Any of her networking zines or blog posts are worth an hour.

LAST-REVIEWED: 2026-06-02