Internet & Telecommunications: How the World Connects

Figures current as of October 2026

How information travels, who owns the networks that carry it, and who controls them. The course covers the internet's layers and its history; packets, addresses, and protocols; the fiber backbone and the exchanges where networks meet; the submarine cables that carry almost all traffic between continents; broadband at home; mobile networks and the airwaves they use; satellites; the cloud and its data centers; who is still offline, and why; who governs the internet; and the hacks, shutdowns, and splits that threaten it. A dated watch list closes the course.

MODULE 1 — What the Internet Is

Goal: Understand what the internet is made of, who owns each part, and how big it has become.

A network of networks that agree on common rules; its layers, from cables and radio to the applications; who owns each layer; how many people use it and how much data they move; and what the cloud physically is.

MODULE 2 — From Telegraph to Smartphone

Goal: Follow the networks from Morse's first line to the smartphone.

The telegraph and the first Atlantic cables; the telephone and the Bell System's monopoly, broken up in 1984; ARPANET and the protocols that became the internet; the web; the commercial internet, the 1996 telecommunications law, and the fiber glut of the dot-com bust; and the iPhone.

MODULE 3 — How Data Travels

Goal: Understand how a message crosses the internet, from addresses to routes.

Packets and addresses, and the slow move to a larger address system; the rules computers use to talk, to look up names, and to encrypt; how networks announce routes to each other, and how that fails; and why delay often matters more than speed.

MODULE 4 — The Backbone

Goal: Learn how fiber carries data and how networks connect and pay one another, and weigh whether content companies should pay the carriers.

Light in glass and the long-haul networks; transit and peering; the exchange points where networks meet; the content delivery networks that bring video close to viewers; and Europe's argument over a "fair share" fee.

MODULE 5 — Under the Sea

Goal: Understand the submarine cables that carry traffic between continents, and weigh whether they are protected well enough.

How cables are laid, landed, and repaired; the 694 systems in service or planned, and the content companies now building them; faults, the shortage of repair ships, the Red Sea and Baltic cuts, and the Gulf projects halted by the 2026 war.

MODULE 6 — The Last Mile

Goal: Follow broadband into the home and what it costs, and weigh whether American broadband is overpriced.

Copper, cable, fiber, and fixed wireless, and the Wi-Fi inside the home; the spread of fiber; the American market after the mergers of 2026; speeds in America and the reach of broadband in Europe, China, and India; American prices against the world's; and what a monthly bill contains.

MODULE 7 — Mobile Networks

Goal: Understand how a cellular network works and how carriers make money, and weigh whether 5G has paid off.

Phones, towers, and the network behind them; the generations from 1G to 5G; subscriptions, coverage, and traffic; the carriers' business; the equipment makers and the tower companies; and 6G, expected about 2030.

MODULE 8 — The Airwaves

Goal: Understand radio spectrum and how governments divide and sell it.

Why some frequencies travel far and others carry more; licenses, auctions, and the unlicensed bands that Wi-Fi uses; America's 2025 law restoring spectrum auctions; the sale of EchoStar's spectrum to AT&T; and SpaceX; and the world radio conference of 2027.

MODULE 9 — Satellites

Goal: Compare satellite broadband with the networks on the ground, and weigh whether satellites can replace fiber for rural homes.

High orbits and low orbits, and why delay differs; Starlink and its rivals, from Amazon's network to China's; satellites that reach phones directly; the crowding of low orbit; and satellites in war.

MODULE 10 — The Cloud and Data Centers

Goal: Understand the cloud, the data centers behind it, and why outages spread, and weigh whether concentration has made the internet fragile.

What the cloud is; the three companies that sell most of it; the electricity and money the data centers consume; the outages of 2024 and 2025; and the drone strikes on cloud data centers in the Gulf in March 2026.

MODULE 11 — Who Is Offline

Goal: Measure who is not online and why, and weigh whether cost now matters more than coverage.

The 2.2 billion people offline; the gap between coverage and use; the price of data in poor countries; the rural and gender gaps; and America's programs, from the end of its broadband subsidy to the rural build-out split between fiber and satellite.

MODULE 12 — Who Governs the Internet

Goal: Understand who makes the internet's rules, and weigh whether states should have a larger role.

The bodies that run names and addresses and write the standards; the UN's telecommunications agency and its conference of November 2026; the Internet Governance Forum; America's Federal Communications Commission, net neutrality, and the Universal Service Fund; and Europe's new network law.

MODULE 13 — Networks Under Attack

Goal: Understand how networks are hacked, switched off, and split, and weigh whether the internet is breaking into national networks.

The hacking of American phone carriers; the bans on Chinese equipment; government shutdowns, which set a record in 2025; Iran's wartime blackout and Russia's blocks on messaging apps in 2026; and the jamming of navigation signals in the Gulf.

MODULE 14 — Where Things Stand, and What to Watch

Goal: Close with the argument, the contradictions, and a dated watch list.

The course's argument on a page; the contradictions no single module states; what to watch from October 2026 on; and a last word.

A note on the sources: counts of people online and of prices for access come mainly from the International Telecommunication Union (ITU), the United Nations agency for telecommunications, and from the GSMA, the association of mobile operators. Mobile figures come also from Ericsson's *Mobility Report*, speeds from Ookla's speed tests, and American figures from the Federal Communications Commission (FCC), the National Telecommunications and Information Administration (NTIA), and the companies' own reports to investors.

Submarine cables are counted by TeleGeography, a research firm, and the cloud market by Synergy Research Group. Traffic, encryption, and outages come from Cloudflare Radar and other measurement services that watch the internet from outside, and government shutdowns from Access Now, a digital-rights group. Where governments, companies, and armed groups make claims about attacks, sabotage, and shutdowns that cannot always be checked, the course says who claims what.

Speeds are given in megabits per second (Mbps), gigabits per second (Gbps), or terabits per second (Tbps), and volumes of data in gigabytes (GB); a byte is 8 bits, so bits measure speed and bytes measure amounts. Delay is measured in milliseconds (ms), thousandths of a second. One abbreviation recurs: artificial intelligence (AI). The Federal Reserve is America's central bank. Each figure names its source and date, because counts differ by source: the ITU and other counters disagree about the number of people online by hundreds of millions. Amounts in other currencies are converted at the Federal Reserve's exchange rates for October 16, 2026. Where the field argues, over fees for content, the protection of cables, broadband prices, 5G, satellites, the cloud, the digital divide, governance, and the internet's splitting, the course makes each case as its advocates make it, gives the evidence, and then states its own reading once.

Key Numbers, as of October 2026

The measurements the course relies on

Measure	Figure	Period	Source and release
People online	About 6.0 billion, 74% of the world's people; 2.2 billion offline, 96% of them in low- and middle-income countries	2025	ITU, *Facts and Figures 2025*, November 17, 2025
Covered but not connected	3.1 billion people live under a mobile broadband signal but do not use mobile internet; 270 million live beyond any mobile broadband signal (the ITU counts about 312 million without mobile broadband coverage)	2025	GSMA, *State of Mobile Internet Connectivity 2026*, September 2026; ITU
The price of access	The median entry-level mobile data plan costs 1.4% of average monthly income per person, against a target of 2%; in low-income countries, 2 GB a month costs 7.4%	2025; the low-income figure for 2024	ITU price data, 2025
Traffic	Internet traffic grew 19% in 2025; the average American home used 711 GB a month in the second quarter of 2026	2025-26	Cloudflare, *Radar 2025 Year in Review*; OpenVault
Speeds	Median download speed worldwide: 109 Mbps on mobile networks and 130 Mbps on fixed broadband	August 2026	Ookla, Speedtest Global Index
Mobile	8.84 billion mobile subscriptions at the end of 2025; 5G subscriptions passed 3.1 billion at the end of March 2026, and 5G networks covered 60% of the world's people; the average smartphone used 22 GB a month in 2025	2025-26	Ericsson, *Mobility Report*, June 2026
China's 5G	4.838 million 5G base stations and 1.204 billion 5G subscriptions	End of 2025	Ministry of Industry and Information Technology (MIIT), via Xinhua
Fiber in America	Fiber passes 84.6 million unique homes, more than 60% of households; a record 11.8 million homes were passed in 2025	Data to September 2025	Fiber Broadband Association, with RVA, a research firm
Home broadband by radio	Fixed wireless customers: T-Mobile 8.45 million on December 31, 2025, its last reported count; Verizon 6.21 million on June 30, 2026; AT&T 2.94 million on September 30, 2026	2025-26	Company reports
America's mergers	Verizon closed its purchase of Frontier on January 20, 2026, AT&T its purchase of Lumen's consumer fiber business on February 2, and Charter its merger with Cox on August 20	2026	Company releases
Submarine cables	694 cable systems, in service and planned, and 1,893 landing stations; more than 1.5 million km of cable in service	March 2, 2026	TeleGeography
Cable faults and repairs	150-200 faults a year, 70-80% of them accidental, mostly from fishing gear and anchors; about 20 ships work mainly on repairs, and a repair took 51 days on average in 2024	Since 2014; 2024	International Cable Protection Committee (ICPC), via Reuters; testimony to the U.S.-China Economic and Security Review Commission, 2026
Exchange points	A record peak of 35.9 Tbps at IX.br São Paulo, which TELETIME, a Brazilian trade publication, calls the world's busiest exchange; 21.47 Tbps at DE-CIX Frankfurt	June 25, 2026; October 2026	IX.br; DE-CIX

Measure	Figure	Period	Source and release
Addresses	For the first time, more than half of users (50.10%) reached Google over IPv6, the newer address system, on March 28, 2026; an address in the older IPv4 system sold for $20.04 on average	March 2026; first half of 2026	Google; IPv4Center
Encryption	95.1% of web requests to Cloudflare used HTTPS, the encrypted form of the web's protocol	Third quarter of 2026	Cloudflare Radar
The cloud	Cloud infrastructure services earned $143.4 billion, 43% more than a year earlier; Amazon Web Services (AWS) held 28% of the market, Microsoft 20%, and Google 15%	Second quarter of 2026	Synergy Research Group, July 30, 2026
Data centers' electricity	About 485 terawatt-hours (TWh) in 2025, about 1.5% of the world's electricity, and about 950 TWh by 2030; American data centers used 4.4% of America's electricity in 2023	2023-30	International Energy Agency (IEA), April 16, 2026; Lawrence Berkeley National Laboratory, December 2024
Capital spending	Five large technology companies spent more than $400 billion in 2025; the four largest spenders, Amazon, Microsoft, Alphabet, and Meta, guided to about $720-745 billion for 2026	2025-26	IEA, April 2026; July 2026 earnings calls
Satellites	Starlink: 13.6 million consumer subscribers in 174 markets on September 30, 2026, and about 10,200 working satellites on June 30. Amazon Leo: 396 satellites launched by July 2, 2026	2026	SpaceX, August 4 and October 20, 2026; Amazon
Spectrum	AT&T paid EchoStar about $23 billion for about 50 megahertz (MHz) of spectrum, roughly $1.40 per MHz for each person covered; the FCC must finish auctioning the upper C-band by July 4, 2027	2025-27	EchoStar filings; FCC
Universal service	The FCC proposed a charge of 42.0% of carriers' interstate and international end-user revenue for the fourth quarter of 2026 to fund America's Universal Service Fund; the Supreme Court upheld the fund 6-3 on June 27, 2025	2025-26	FCC, September 2026; *FCC v. Consumers' Research*
Rural broadband money	The Broadband Equity, Access, and Deployment (BEAD) program awarded $18.19 billion for 3.79 million homes and businesses: 62.9% get fiber, 21.9% low Earth orbit (LEO) satellite, 13.3% fixed wireless, and 1.9% cable	August 28, 2026	NTIA data, via Broadband Breakfast
Shutdowns	At least 313 government internet shutdowns in 52 countries in 2025, the most since counting began in 2016; Myanmar ordered 95 and India 65. Iran's wartime blackout lasted 88 days, from February 28 to late May 2026	2025-26	Access Now, March 31, 2026; Cloudflare Radar
Hacking the carriers	Salt Typhoon, a hacking campaign attributed to China's state, reached at least 200 American companies and targets in 80 countries	Stated August 27, 2025	Federal Bureau of Investigation (FBI)
Jamming	978,000 events of jamming of satellite navigation, such as the Global Positioning System (GPS), affected ships worldwide in the first quarter of 2026, 98% of them in the Gulf	First quarter of 2026	Windward, a maritime-data firm

Counts differ by source because they measure different things: the ITU counts people who used the internet in the past three months, the GSMA counts users of mobile internet, and Ookla's medians come from tests that users choose to run. Speeds are in bits per second and volumes in bytes; a byte is 8 bits. Company figures are the companies' own.

Who's Who

The people who run networks and their rules, as of October 25, 2026

Regulators and governments

Name	Who they are	Why they matter to networks
Brendan Carr	Chairman of the FCC since January 2025	Leads the regulator that approved EchoStar's spectrum sales to AT&T and SpaceX in May 2026, and that adopted the rules for auctioning the upper C-band on July 22, 2026
Olivia Trusty	FCC commissioner, a Republican, since June 2025	With Carr, voted 2-1 in November 2025 to rescind a January 2025 ruling that had read America's wiretap law as requiring carriers to secure their networks
Anna Gomez	FCC commissioner, the only Democrat, since September 2023	Dissented from that vote; her term expired on June 30, 2026, but she may serve until the end of 2027 unless a successor is confirmed. Two of the commission's five seats are empty
Arielle Roth	Administrator of the NTIA since 2025	Approved the last of the 56 state and territory plans for the BEAD program on August 25, 2026, and decides what happens to its leftover money
Howard Lutnick	US Secretary of Commerce since February 2025	Oversees the NTIA; told senators in February 2026 that returning BEAD's leftover money to the Treasury "is not the plan"
Henna Virkkunen	Executive Vice-President of the European Commission for Tech Sovereignty, Security and Democracy since December 2024	Presented the EU's Digital Networks Act on January 21, 2026, without a mandatory fee on content companies
Doreen Bogdan-Martin	Secretary-General of the ITU since January 2023, the first woman in the post	The only candidate for Secretary-General at the ITU's Plenipotentiary Conference in Doha, which opens on November 9, 2026
Kurtis Lindqvist	President and chief executive of the Internet Corporation for Assigned Names and Numbers (ICANN) since December 2024	Leads the nonprofit that sets policy for the internet's domain names and coordinates its addresses
Maksut Shadayev	Russia's Minister of Digital Development, Communications and Mass Media since January 2020	His ministry introduced, in August 2025, a "whitelist" of services deemed essential that stay reachable during mobile shutdowns
Andrei Lipov	Head of Roskomnadzor, Russia's communications regulator, since 2020; age 56	His agency slowed Telegram in February 2026, and a nationwide ban on Telegram took effect on April 10, 2026
Li Lecheng	China's Minister of Industry and Information Technology since April 2025; age 61	Regulates China's carriers, which ran 4.838 million 5G base stations at the end of 2025 by his ministry's count
Zhuang Rongwen	Director of the Cyberspace Administration of China since 2018	Runs the regulator of China's online content and data; China's amended Cybersecurity Law took effect on January 1, 2026
Sattar Hashemi	Iran's Minister of Information and Communications Technology since August 2024; born 1976	President Masoud Pezeshkian ordered his ministry on May 25, 2026 to restore international access after the blackout that began on February 28

Carriers and equipment makers

Name	Who they are	Why they matter to networks
John Stankey	Chief executive of AT&T since July 2020	AT&T added a record 769,000 internet customers in the third quarter of 2026, and its fiber reached more than 39 million locations
Dan Schulman	Chief executive of Verizon since October 2025	Closed Verizon's purchase of Frontier, a fiber provider, on January 20, 2026; Verizon reports its third quarter on October 26
Srini Gopalan	Chief executive of T-Mobile US since November 2025	Runs America's largest provider of fixed wireless home broadband
Brian Roberts and Mike Cavanagh	Comcast's chief executive since 2002 and chairman since 2004 (Roberts); co-chief executive with Roberts since January 2026 (Cavanagh)	Run Comcast, America's second-largest broadband provider after Charter's merger with Cox
Chris Winfrey	President and chief executive of Charter Communications since December 2022	Closed Charter's purchase of Cox on August 20, 2026, creating America's largest broadband provider, with more than 35 million subscribers
Chen Zhongyue	Chairman of China Mobile since December 22, 2025	China Mobile had about 1.011 billion mobile customers on June 30, 2026, 687 million of them on 5G
Margherita Della Valle	Chief executive of Vodafone Group since 2023	Vodafone has a contract with AST SpaceMobile, running to 2034, for satellite service to ordinary phones
Timotheus Höttges	Chief executive of Deutsche Telekom since January 2014	The appeals court in Düsseldorf ruled on February 10, 2026 that a Meta unit owed Deutsche Telekom more than €30 million for carrying traffic
Akash Ambani	Chairman of Reliance Jio Infocomm since June 2022	Jio, India's largest carrier, had 503.6 million wireless subscribers in June 2026, by the national regulator's count
Börje Ekholm	President and chief executive of Ericsson since January 2017	Ericsson was the second-largest maker of mobile radio equipment in 2025 by Dell'Oro's estimate, and publishes the *Mobility Report*
Justin Hotard	President and chief executive of Nokia since April 2025	Nokia ranked third in mobile radio equipment in 2025 and gained share (Dell'Oro)
Ren Zhengfei	Founder (1987) and chief executive of Huawei	Huawei was the largest maker of mobile radio equipment in 2025 (Dell'Oro), though America bars its gear and several European governments restrict it

Satellites, the cloud, and the internet's builders

Name	Who they are	Why they matter to networks
Elon Musk	Founder (2002) and chief executive of SpaceX, which runs Starlink	Controls the largest satellite network; in October 2026 he said decisions on the use of Starshield, its military version, rest with the US government
Gwynne Shotwell	President and chief operating officer of SpaceX since 2008	Said in August 2026 that SpaceX was integrating 65 MHz of spectrum bought from EchoStar for service to ordinary phones
Rajeev Badyal	Vice president of Amazon Leo, Amazon's low-orbit satellite network, formerly Project Kuiper	Amazon Leo had launched 396 production satellites by July 2, 2026, but had named no date for general service by October 21

Name	Who they are	Why they matter to networks
Abel Avellan	Founder (2017), chairman, and chief executive of AST SpaceMobile	Builds large satellites that connect to ordinary phones; one, BlueBird 7, was lost after its launch on April 19, 2026
Jean-François Fallacher	Chief executive of Eutelsat, owner of the OneWeb low-orbit network, since June 2025	Runs Europe's main rival to Starlink
Matt Garman	Chief executive of AWS since June 2024	AWS held 28% of cloud infrastructure services in the second quarter of 2026; its Fastnet transatlantic cable is due to enter service in 2028
Satya Nadella	Chief executive of Microsoft since February 2014	Microsoft held 20% of cloud infrastructure services in the second quarter of 2026 (Synergy)
Thomas Kurian	Chief executive of Google Cloud since 2019	Google held 15% of cloud infrastructure services in the second quarter of 2026 (Synergy)
Matthew Prince	Co-founder (2009) and chief executive of Cloudflare	Cloudflare carries a large share of web traffic and measures the internet through Cloudflare Radar; its outage of November 18, 2025 took many websites offline
Roman Danyliw	Chair of the Internet Engineering Task Force (IETF) since March 2023, reselected for a term from March 2026	Leads the open body that writes the internet's core protocols
Tim Berners-Lee	Inventor of the World Wide Web	Proposed the web at CERN, the European particle-physics laboratory, in March 1989, and founded the World Wide Web Consortium (W3C) in 1994
Vint Cerf	Co-designer, with Bob Kahn, of the internet's core protocols, TCP/IP; Google's chief internet evangelist	Their paper of May 1974 set out the design that ARPANET, the American military's research network, adopted on January 1, 1983

Terms

The vocabulary of networks

Term	Meaning
Bit, byte, and bandwidth	A bit is a single 0 or 1; a byte is 8 bits. Bandwidth is how many bits a link can carry each second, in megabits (Mbps), gigabits (Gbps), or terabits per second (Tbps); volumes of data are counted in bytes, such as gigabytes (GB)
Latency	The time data takes to cross a network, in milliseconds (ms). Round-trip time, there and back, is what a user notices when a page loads or a call lags; distance, the speed of light in glass, and congestion set it
Packet	A small piece of a message, usually no more than 1,500 bytes, that carries the addresses of its sender and receiver; packets may travel by different routes, and the receiving computer puts them back in order
Internet service provider (ISP)	A company that sells connections to the internet to homes, businesses, or other networks
Protocol	An agreed set of rules for exchanging data. The Internet Protocol (IP) addresses and forwards packets; the Transmission Control Protocol (TCP) makes delivery reliable; the User Datagram Protocol (UDP) is faster but unchecked; QUIC, a newer protocol first built by Google, combines speed, reliability, and encryption
IP address, IPv4 and IPv6	The number that identifies a device or network on the internet. IPv4 addresses are 32 bits long, allowing about 4.3 billion; IPv6 addresses are 128 bits long, enough never to run out
Domain Name System (DNS)	The internet's directory, which turns names such as example.com into IP addresses; at its top sits the root zone, served by 13 named root servers through about 2,000 copies, or instances
Encryption and HTTPS	Encryption scrambles data so that only the intended receiver can read it; HTTPS is the web's protocol with encryption added, shown by the padlock in a browser
Router	A computer that reads each packet's destination address and passes it on toward the next network on its way
Autonomous system and the Border Gateway Protocol (BGP)	An autonomous system is a network run under one routing policy, such as an ISP's or a cloud company's, with its own number; BGP is how these networks tell one another which addresses they can reach
Transit and peering	Transit is paid access to the whole internet through a larger network; peering is two networks exchanging their own customers' traffic directly, usually without payment
Internet exchange point (IXP)	A building, or a group of them, where many networks connect to shared switches to exchange traffic by peering
Content delivery network (CDN)	A system of servers that keep copies of popular content, such as video, close to users, so that it travels a short way
Optical fiber and the backbone	Fiber is a strand of glass about as thin as a hair that carries data as pulses of light, many colors at once; the backbone is the high-capacity fiber that links cities and countries
Submarine cable and landing station	A cable laid on the seabed, with up to about two dozen fiber pairs inside and repeaters every 50 to 100 km to boost the light, powered from the shore; a landing station is where it comes ashore
The last mile	The final link from a network to a home or business, and usually the most expensive part to build
Digital subscriber line (DSL), cable, and fiber to the home	DSL sends data over telephone copper; cable broadband uses television's coaxial cable and the DOCSIS (Data Over Cable Service Interface Specification) standard; fiber to the home (FTTH), or to the premises (FTTP), runs glass all the way
Passive optical network (PON)	The usual way to build fiber to homes: one fiber from the operator's building is divided by unpowered splitters among 32 or 64 homes, which share its capacity
Fixed wireless access (FWA)	Home broadband sent by radio from a mobile network's masts to a receiver at the house
Wi-Fi	The standards for short-range wireless networks inside a home or building, which use radio bands that need no license

Term	Meaning
Passings and take-up	A home is "passed" when a network runs close enough to connect it; take-up is the share of passed homes that subscribe
Spectrum and hertz	The range of radio frequencies, measured in hertz, cycles per second: megahertz (MHz) and gigahertz (GHz). Low bands, below 1 GHz, travel far and through walls; mid bands, from 1 GHz to about 7 GHz, balance range and capacity; high bands, above 24 GHz, carry the most and travel the least
MHz-pop	The usual unit for comparing spectrum prices: what a license costs, divided by its width in megahertz and by the number of people in the area it covers
Licensed and unlicensed spectrum	Licensed bands are sold or assigned for one holder's exclusive use, as mobile carriers' are; unlicensed bands, such as Wi-Fi's, are open to any device that keeps to the rules on power
Mobile generations: 2G to 6G	Successive standards for mobile networks, about a decade apart: 2G brought digital calls and texts, 3G mobile data, 4G, also called LTE (Long-Term Evolution), fast mobile broadband, and 5G more capacity and shorter delays; 6G is expected about 2030
Radio access network (RAN), core, and backhaul	The RAN is the masts, antennas, and radios that link phones to the network; the core is the central system that tracks users, connects calls, and links to the internet; backhaul carries traffic between the two, by fiber or microwave
Geostationary orbit (GEO) and low Earth orbit (LEO)	A GEO satellite circles 35,786 km above the equator once a day, so it seems fixed in the sky; LEO satellites fly between a few hundred and 2,000 km up, so signals travel far less distance, but thousands of satellites are needed for constant coverage
Direct to device	Satellite service that reaches ordinary phones without a dish, using spectrum shared with or bought from mobile carriers
Cloud computing and hyperscalers	Cloud computing is renting computing power, storage, and software from a provider's data centers over the network, paid for by use; the hyperscalers are the few companies that run it at the largest scale, led by AWS, Microsoft, and Google
Cloud region and availability zone	A region is a cluster of data centers in one metropolitan area; an availability zone is one or more data centers within a region with their own power and networks, so that one can fail without taking down the rest
Net neutrality	The principle that internet providers should treat lawful traffic equally, without blocking it, slowing it, or charging for priority
Universal service	Programs that subsidize phone and internet service in rural areas, for poor households, and for schools, libraries, and rural clinics; America's Universal Service Fund is paid for by a charge on telecommunications companies' revenue that is passed on to customers
Multistakeholder governance	The model under which the internet's technical rules are set in open bodies where companies, engineers, academics, civil society, and governments all take part, rather than by treaty among states alone
Internet shutdown	A deliberate disruption of internet or mobile data by a government, in a region or a whole country, or the blocking of particular services, such as messaging apps
Splinternet	The idea that the internet is breaking apart into national networks with their own rules, filters, and standards
Jamming and spoofing	Jamming drowns out satellite navigation signals, such as those of GPS, with noise; spoofing sends false signals that make a receiver report the wrong position

Module 1 — What the Internet Is

Layers, Owners, and Scale

Recommended Reading / Reference

Book: Andrew Blum, *Tubes: A Journey to the Center of the Internet* (2012), a reporter's tour of the buildings, exchange points, cable landings and data centers where the internet physically lives

Reference: ITU, *Facts and Figures 2025* (November 17, 2025), the standard yearly count of who uses the internet, and Cloudflare Radar, a free public dashboard of the traffic that one of the largest networks carries

1a. A Network of Networks

The internet is not one network but tens of thousands of them, each with its own owner, that have agreed to carry traffic for one another. A university's campus network, a mobile carrier, a cable company, Google's network and the network in a small office are all separate. What makes them one internet is a shared set of protocols, the agreed rules for how data is formatted, addressed and passed along. The central one is the Internet Protocol (IP), which gives every connected device an address and defines how data moves toward it.

Engineers count the internet's independent networks as autonomous systems (ASes), each a network run under one routing policy, with its own number, that tells its neighbors which addresses it can deliver traffic to (Module 3 explains how). A routing report presented to network operators counted 76,870 autonomous systems in the global routing table, the shared list of routes between networks, on May 22, 2025. Of these, 11,066 passed on traffic for other networks, and the other 65,804 handled only their own.

No one runs the whole. The internet has no central computer, no master switch and no single owner, so no one can turn all of it off, although a government can cut off its own country (Module 13). Each network decides which others to connect with and on what terms, mostly through private contracts between companies (Module 4 explains who pays whom). The few things that must be unique worldwide, such as addresses and domain names, are coordinated by nonprofit bodies, and the technical standards are written in open meetings that anyone may join (Module 12).

Three things often mistaken for the internet run on top of it. The World Wide Web is one application, a system of linked pages first proposed in 1989 (Module 2), and email, video calls, online games and payment networks all use the internet without using the web. The cloud is a set of computers in data centers that companies rent out (section 1f). A mobile network is a separate system, built for telephone calls, that now carries internet traffic as well (Module 7).

Nor is the internet the same everywhere. The protocols are common to all, but what a user can reach depends on where the user is, because governments can order their networks to block sites and services. China filters its connections to the rest of the world, Russia and Iran have built the means to cut theirs off, and whether these national controls are splitting the internet into separate networks is a question Module 13 takes up.

1b. The Layers

The internet's designers split the work of moving data into layers. Each layer does one job, relies on the layer below and serves the layer above, and each can change without the others having to. Most textbooks teach five layers, from the physical cable at the bottom to the program a person uses at the top. An older reference model, published by the International Organization for Standardization (ISO) in 1984, has seven, but the five-layer version matches the internet as it was built.

The physical layer carries raw bits, the ones and zeros of digital data, as pulses of light in glass fiber, electrical signals in copper wire, or radio waves through the air and up to satellites (Modules 4 to 9 cover each). The link layer moves data across a single hop, such as from a laptop to a home Wi-Fi router or between two machines on the same cable. Ethernet, the standard for wired local networks, and Wi-Fi are both link-layer standards.

The network layer is where the internet proper begins. Data travels in packets, small parcels each labeled with a source and a destination address, and IP defines those labels. Routers, specialized computers that read each packet's destination and choose the next link toward it, pass every packet from network to network until it arrives. The transport layer sits above and manages the conversation between two devices. The Transmission Control Protocol (TCP) checks that every packet arrives and puts them back in order, while the User Datagram Protocol (UDP) sends without checking, which suits live voice, where a late packet is useless (Module 3).

The application layer holds the programs people use and the protocols they speak to one another. These include the web's Hypertext Transfer Protocol (HTTP), email, video streaming, and the DNS, which turns names such as example.org into numerical addresses (Module 3). An application does not need to know whether its packets travel over fiber, Wi-Fi or a satellite, and a cable does not need to know whether it carries a film or a bank transfer.

1c. How the Layers Fit Together

Each layer wraps the data handed down from the layer above in an envelope of its own, a step engineers call encapsulation. A web page leaves its server as HTTP, and TCP wraps it with numbers that let the receiver put the pieces in order. IP then adds the addresses, and Ethernet or Wi-Fi adds the details of the next hop. Each router along the way removes and replaces only the outermost wrapping, reads the IP address and passes the packet on, and only the receiving device unwraps it all the way to the page.

In a rough analogy with the post, the physical layer is the roads, rails and planes, the network layer is the address on the envelope and the sorting offices that read it, and the transport layer is the choice between a tracked parcel and an ordinary letter. The analogy fails at one point that matters, because the internet cuts each message into many small packets, which may travel by different routes and are put back together only at the far end.

The layered design explains how the internet grew without a central plan. Because IP is the one layer every network must share, anyone could build a new application on top of it without asking permission, as the web's inventor did and as the video services did later. Anyone could also add a new physical medium underneath, as Wi-Fi and mobile data did, provided it carried IP packets. Engineers call the related principle "end to end": the network moves packets and leaves the intelligence to the devices at either end.

The same design explains where change happens. The top and bottom layers have changed many times, from email to streaming and from copper to fiber and radio, while the narrow middle has changed slowly.

The version of IP that still carries most of the internet's traffic dates from the early 1980s, and its successor has taken decades to spread (Module 3).

1d. Who Owns Each Layer

Ownership follows the layers, and the table gives the main owners of each. The lower layers cost the most to build, and companies own most of them. They earn their money by selling access, so their business is the monthly subscription and the wholesale contract between networks. The upper layers are where most of the internet's money is made, from advertising, subscriptions, online shopping and rented computing.

Layer	What it does	Main owners or operators	Examples	Covered in
Physical	Carries bits as light, electricity or radio	Telephone and cable companies, tower and satellite firms, undersea-cable consortia, content companies, some governments	Fiber to homes; mobile masts; Starlink; transatlantic cables	Modules 4-9
Link	Moves data across one hop	The same network owners; households, for home Wi-Fi	Ethernet; Wi-Fi; mobile radio standards	Modules 6-8
Network	Addresses packets and routes them between networks	Each network for its own routers; addresses coordinated by the Internet Corporation for Assigned Names and Numbers (ICANN) and five regional registries	IP; BGP (Border Gateway Protocol), the routing protocol between networks	Modules 3, 4, 12
Transport	Manages each conversation: order, loss, pace	No owner: open standards that run on users' devices and on servers	TCP; UDP; QUIC, a newer protocol	Module 3
Application	The services people use	Companies and individuals; the standards themselves are open	The web; email; Netflix; cloud services	Modules 3, 10

This course's summary of the main owners; the modules named give the detail.

The largest content and cloud companies have moved down the layers. Google, Meta, Amazon and Microsoft began as applications and now own data centers, private fiber networks and stakes in undersea cables, so much of their traffic never crosses a carrier's public network (Module 5 gives their share of the traffic that crosses oceans).

A growing part of the physical layer belongs to firms that sell space rather than service. Tower companies own masts and rent room on them to several carriers at once (Module 7). Colocation companies own data centers and rent space, power and connections to other firms' servers, and exchange points, many of them run by associations of their members, give networks a place to meet and swap traffic (Module 4). These neutral owners let rivals share the most expensive parts of the system.

Governments own the radio spectrum in nearly every country and license it to carriers (Module 8). Some also own networks outright: China's three carriers, China Mobile, China Telecom and China Unicom, are controlled by the state, as are national carriers in many countries of Africa, Asia and the Middle East. Elsewhere governments act on networks they do not own, approving mergers, setting rules on privacy and content, paying for rural lines (Module 11) and, in some countries, ordering networks shut (Module 13).

The protocol layer belongs to no one, because TCP/IP, HTTP and the DNS are open standards, which anyone may build into a product without a license or a fee. Engineers from companies and universities write them in open bodies such as the Internet Engineering Task Force (IETF). The internet's unique identifiers are its addresses and its top-level domain names, the last part of a name, such as .com, and ICANN and the five regional registries coordinate them (Module 12 covers all of these bodies).

This mix is why the course says the internet is run more by agreements than by governments. A packet from a phone in Lagos to a server in Frankfurt may cross a mobile carrier's network, an undersea cable owned by a consortium, an exchange point run by a nonprofit association, and a cloud company's private network. No law requires any of them to carry it, and each does so under a contract or under the custom of free exchange between networks that Module 4 describes.

1e. How Big It Is

About 6.0 billion people used the internet in 2025, 74% of the world's people, by the count of the ITU, the United Nations agency for telecommunications (*Facts and Figures 2025*, November 17, 2025). The ITU counts people who used it from any place and on any device, estimated from national surveys. The number rose by more than 240 million in a year, and the ITU revised its estimate for 2024 from 5.5 billion to 5.8 billion, so figures from its earlier editions should not be compared with the new ones.

Other counts differ: DataReportal, a compiler of digital statistics, puts 2025's users at about 5.6 billion, 400 million fewer than the ITU, because it draws on different sources and estimates. This course uses the ITU's figure throughout, as the official series, compiled from governments' own data. Who the remaining quarter of the world's people are, and why they are offline, is the subject of Module 11.

Traffic grows faster than the number of users. Cloudflare measured global internet traffic 19% higher in early December 2025 than in a baseline week in mid-January, comparing average daily volumes and leaving out automated "bot" traffic (Cloudflare Radar, December 15, 2025). Most of the growth came from August onward, and because Cloudflare carries traffic for a large share of the world's websites, its view is broad but not complete. AppLogic Networks, formerly Sandvine, which sells traffic-management systems to carriers, put global traffic above 33 exabytes a day in its 2024 report, where an exabyte is a billion gigabytes.

Video is the largest part of that traffic. In AppLogic's 2025 report video remained the largest category of traffic by volume, and live sport set the peaks, running 30-40% above normal levels on the busiest days. The most recent public breakdown this course found, Sandvine's of October 2018, put video at 57.7% of the traffic flowing down to users worldwide, with Netflix alone at 15%. New uses are arriving on top: the same 2025 report found AI assistants in use by 7% of households on fixed networks and 4% of mobile users.

The average American home with broadband used 711 gigabytes (GB) of data a month in the second quarter of 2026, counting downloads and uploads together, by the count of OpenVault, a firm that measures usage on the networks of the internet service providers it works with (*Broadband Insights*, July 2026). That comes to about 23.7 GB a day, far more than a smartphone uses on its mobile plan (Module 7), and most of it is video, as the worked example shows.

Two units run through this course, and they are easy to confuse. Speed is measured in bits per second, as megabits per second (Mbps) or gigabits per second (Gbps), while volume is measured in bytes, as gigabytes. A byte is eight bits, so a line running at 8 Mbps moves one megabyte, a million bytes, each second. A month's use is a volume and a connection's capacity is a speed, and the two answer different questions: how much a home uses, and how much it can use at once.

1f. What "the Cloud" Physically Is

"The cloud" is a name for computers in someone else's building. When a phone backs up its photos, or a company runs its website or its accounts "in the cloud", the data sits on servers in data centers. A data

center is a building, often the size of a warehouse, filled with racks of computers, with its own power supply, backup generators and cooling, and connected to the internet by many fibers.

Cloud computing means renting that capacity instead of owning it: computing time, storage and software, paid for by the hour or by the gigabyte. The largest providers group their data centers into regions in particular metropolitan areas, each divided into zones with separate power and networks, so that a service can survive the loss of a building (Module 10). A customer chooses the region where its data lives, and with it the country whose laws apply to that data.

Not every server is in a rented cloud. Many companies, banks and governments still run data centers of their own, and others place their own machines in colocation buildings, paying for space, power and connections. The cloud companies' advantage is scale: they buy computers, electricity and fiber in bulk, and they let a customer rent a thousand machines for an hour and then hand them back.

So the cloud has an address: each region sits in a real place, chosen for reliable power, land, fiber routes and tax rules, and its buildings can fail, flood, burn or be attacked like any other. Module 10 covers the providers, their market shares, the electricity their data centers use, the outages of 2025 and the strikes on data centers in the Gulf in 2026.

The cloud also reaches out toward its users through content delivery networks, which keep copies of popular films, pages and software updates on servers inside or next to internet service providers, so that most requests are answered close to home (Module 4). A video that seems to come from "the cloud" usually comes from a server in or near the viewer's own city.

1g. The Course's Thread

The internet looks like a cloud, and the word invites the idea that it floats free of the ground. This course argues the opposite: the internet is a physical system of cables, radio, satellites and data centers. Companies own most of it, and it runs more by agreements between networks than by the decisions of governments.

Its weak points are physical and political, and 2025 and 2026 tested them. Undersea cables were cut in the Red Sea and the Baltic, where courts are still weighing whether the Baltic cuts were deliberate (Module 5). Faults at a few cloud and network companies took large parts of the internet down for hours in late 2025, and drones struck Amazon data centers in the Gulf in March 2026, by Amazon's account (Module 10). Governments ordered a record number of internet shutdowns in 2025, by the count of one rights group (Module 13).

The largest gap is no longer coverage but cost. Most people who do not use the internet live within reach of a mobile signal, and the price of handsets and data, along with a lack of skills, keeps them offline (Module 11). For the next billion users the question is less where to build networks than how to make using them affordable.

The modules follow the layers, with Modules 2 and 3 on the history and the protocols. Modules 4 to 10 work up from fiber and undersea cables through home broadband, mobile networks, spectrum and satellites to the cloud. Modules 11 to 13 turn to who is left out, who governs, and who attacks the networks, and Module 14 sets out what to watch.

Worked example — One Household’s Month of Data

Take an illustrative household of three in a 30-day month, and build its use from published per-hour rates and stated assumptions. Netflix publishes its rates per device: at its highest setting a stream uses up to 3 GB an hour in high definition (HD) and up to 7 GB an hour in ultra-high definition, also called 4K (Netflix Help Center, read October 25, 2026). Zoom, the video-call company, publishes the bits per second a group call uses at 720p, a common HD picture: 2.6 Mbps sent and 1.8 Mbps received (Zoom developer documentation, 2026). Everything else here is an assumption. Step 1, 4K streaming on the living-room television: 1.5 hours a day is 45 hours a month, and 45 × 7 GB = 315 GB. Step 2, HD streaming on other screens: 2.5 hours a day is 75 hours, and 75 × 3 GB = 225 GB. Step 3, video calls: one adult works from home with 1.5 hours of group calls on each of 22 working days, or 33 hours. Converting bits to bytes, 2.6 + 1.8 = 4.4 megabits a second, times 3,600 seconds, is 15,840 megabits an hour. Divided by 8, that is 1,980 megabytes, or 1.98 GB an hour, so 33 hours use 65.3 GB, of which 38.6 GB is sent. Step 4, browsing, social media and music on phones and laptops over Wi-Fi: an assumed 2 GB a day, or 60 GB. Step 5, software: an assumed 25 GB game update and 15 GB of phone and computer updates, 40 GB. The month comes to 315 + 225 + 65.3 + 60 + 40 = 705.3 GB, within 1% of OpenVault’s average of 711 GB. Streaming is 540 GB of it, 76.6%, against 9.3% for calls, 8.5% for browsing and 5.7% for updates. The total is sensitive to one choice: watching the 4K hour and a half in HD instead saves 4 GB an hour, or 180 GB, and cuts the month to 525.3 GB. Netflix’s rates are maximums, so the streaming lines are upper estimates for those hours. Last, the units: 711 GB a month is 5,688 gigabits, which spread over the month’s 2,592,000 seconds averages only 2.19 Mbps. A single 4K stream, at 7 GB an hour, runs at 15.6 Mbps while it plays, so a home’s line must be sized for its evening peak, not for its monthly average.

Module 1 Key Points

What the Internet Is

• The internet is tens of thousands of separate networks; a routing report counted 76,870 autonomous systems on May 22, 2025.

• About 6.0 billion people used the internet in 2025, 74% of the world (ITU); DataReportal counts 5.6 billion.

• Five layers: physical, link, network, transport, application. Companies own the lower layers; the protocols belong to no one.

• Cloudflare measured global traffic, excluding bots, up 19% during 2025.

• An average American home used 711 GB a month in the second quarter of 2026 (OpenVault); in the worked example, streaming is 77%.

• The internet’s weak points are physical and political, and its largest gap is now cost, not coverage.

Module 1 Quiz

Check Your Understanding

1. How many people used the internet in 2025, by the ITU’s count?

A) About 6.0 billion
B) About 5.6 billion, the ITU’s figure
C) About 5.5 billion, as it first estimated
D) About 8.1 billion, nearly everyone alive, since mobile signals now reach almost every inhabited place

2. What does the network layer do?

A) Carries light in glass
B) Checks that data arrives
C) Addresses and routes packets
D) Runs the programs that people use

3. Which layer of the internet belongs to no one?

A) The physical cables under the sea
B) The protocols, open standards that anyone may use without a license or a fee
C) The radio spectrum, which carriers buy
D) The data centers that host the cloud

4. How much data did an average American home use in the second quarter of 2026, by OpenVault’s count?

A) 22 GB a month
B) 58.2 GB a month
C) About 1,100 GB a month
D) 711 GB a month

5. What does the course say is now the largest gap in getting people online?

A) Coverage, since most offline people live beyond any mobile signal
B) Speed of connections
C) Electricity supply
D) Cost

Module 2 — From Telegraph to Smartphone

A Short History of Networks

Recommended Reading / Reference

Book: Katie Hafner and Matthew Lyon, *Where Wizards Stay Up Late: The Origins of the Internet* (1996), the standard narrative of how the ARPANET, the network the internet grew from, was built and by whom, drawn from interviews with the engineers who built it

Reference: Robert H. Zakon, “Hobbes’ Internet Timeline” (Request for Comments 2235, November 1997), a dated record of the network’s early growth and the source of this module’s host counts

2a. The Telegraph, 1844-1866

Electrical networks began with the telegraph, which sent messages as coded pulses of current along a wire. On March 3, 1843, Congress gave Samuel Morse $30,000 to build a test line. On May 24, 1844 he sent “What hath God wrought” from the Capitol in Washington to Baltimore, over 38 miles of wire strung along the Baltimore and Ohio Railroad (US Senate). Messages were spelled out in Morse code, letters made of short and long pulses, and read by trained operators at each end.

The line’s ownership set a pattern for American networks: Congress added $8,000 in 1845 to run it under the Post Office, and in 1847 the Post Office leased it to private investors. Private companies built the American telegraph from then on, and by the late 1860s one of them, Western Union, dominated it. Most European governments took the other path, and Britain’s Post Office took over its country’s telegraph companies in 1870. State administrations of the post, the telegraph and later the telephone ran most of Europe’s networks until they were privatized in the 1980s and 1990s.

Borders forced the first international agreements on networks, because lines that crossed frontiers needed common rules for codes, charges and the handing over of messages. In 1865 twenty European states founded the International Telegraph Union in Paris to set them, and it is today’s ITU, the oldest agency in the United Nations system (Module 12 covers its work now).

The ocean was harder to cross, and before 1858 news crossed the Atlantic by ship, a voyage of about ten days. The first transatlantic telegraph cable, between Ireland and Newfoundland, carried its first official message on August 16, 1858. A 98-word greeting from Queen Victoria to President James Buchanan took about 16 hours to send, and the cable failed within about three weeks, after its electrician, Wildman Whitehouse, applied up to 2,000 volts and damaged its insulation.

A lasting link came eight years later, when the steamship *Great Eastern*, then the largest ship afloat, landed a new cable at Heart’s Content, Newfoundland, on July 27, 1866. In early September a second line, made by recovering and splicing a cable lost the year before, reached the same bay. The 1866 cable sent about eight words a minute, about 80 times as fast as the first. Telegraph cables then spread along the world’s main trade routes, and Module 5 covers the fiber cables that carry the internet across the same oceans today.

2b. The Telephone and the Bell System, 1876-1984

The telephone carried the voice itself, and Alexander Graham Bell received his patent for it on March 7, 1876 and made his first call three days later, on March 10. The telephone network worked by circuit switching, in which a call held a continuous path of wire between the two phones for as long as it lasted. Operators at switchboards made the connections at first, and automatic exchanges took over the work later (Module 3 compares this method with the internet's).

Bell's company, the American Telephone and Telegraph Company (AT&T), grew into the Bell System, which made the equipment, ran the long-distance lines and owned most local telephone companies. In 1913, facing antitrust action by the Justice Department, it agreed to let independent telephone companies connect to its long-distance network and to stop buying them up, a deal known as the Kingsbury Commitment. The Communications Act of 1934 created the FCC to regulate interstate telephone service, and for the next half-century AT&T ran most American telephone service as a regulated monopoly.

For decades the government policed the monopoly rather than breaking it up. In 1956 a consent decree, a settlement approved by a court, ended an antitrust case filed in 1949. It kept AT&T to regulated communications, barring it from selling computers, and required it to license its patents to other companies. Those patents included the transistor, invented at its Bell Laboratories in 1947 and the basis of modern electronics. In 1968 the FCC's Carterfone decision let customers attach equipment not made by the Bell System to the network, and households would later use that freedom to connect modems, devices that send data as sound over a phone line.

The end came through the courts, after the Justice Department sued AT&T again on November 20, 1974. The two sides settled on January 8, 1982, and Judge Harold Greene approved the settlement, with changes, that August. On January 1, 1984, the Bell System was broken up: AT&T kept long distance, its equipment maker Western Electric and Bell Labs, and local service passed to seven regional companies: Ameritech, Bell Atlantic, BellSouth, NYNEX, Pacific Telesis, Southwestern Bell and US West.

The breakup opened long distance and telephone equipment to competition, and rival carriers such as MCI, a long-distance company that had begun with microwave relay links, and Sprint grew in the space it made. The regional companies did not stay apart, and after the 1996 Telecommunications Act (section 2e), Bell Atlantic merged with NYNEX and then with GTE, a large independent telephone company, to form Verizon in 2000. SBC Communications, the renamed Southwestern Bell, bought Pacific Telesis, Ameritech, AT&T itself in 2005 and BellSouth in 2006, and took the AT&T name. By the end of 2006, six of the seven regional companies had become two.

2c. ARPANET and TCP/IP, 1961-1990

The internet grew from a different idea about how to share a line. Packet switching cuts data into small blocks that share the lines with everyone else's traffic and are put back together at the far end, instead of reserving a circuit for each conversation. Leonard Kleinrock published the first paper on its theory in July 1961. Paul Baran at the RAND Corporation, a research institute working for the US Air Force, and Donald Davies at Britain's National Physical Laboratory developed the idea separately in the 1960s, and Davies gave the blocks the name they still carry, packets.

The first large network built on the idea was the ARPANET, paid for by the Advanced Research Projects Agency (ARPA) of the US Department of Defense. It was built so that researchers at universities and laboratories funded by ARPA could share scarce and expensive computers. A common story holds that it was designed to survive a nuclear war. Histories written by its builders, including the Internet Society's,

reject that story, because Baran's studies of a network that could survive an attack were separate work at RAND.

In September 1969 the contractor Bolt Beranek and Newman (BBN) installed the network's first switch at the University of California, Los Angeles (UCLA), and connected the first host computer (Leiner and others, 1997). A host is a computer attached to the network, as opposed to the switches that move traffic between hosts. By the end of 1969 four sites were linked: UCLA, the Stanford Research Institute, the University of California at Santa Barbara and the University of Utah. In 1971 Ray Tomlinson of BBN wrote the first program to send email between computers on the network, and chose the @ sign to separate a user's name from the name of the computer.

The ARPANET was one network, and the next problem was joining different ones, including the radio and satellite networks that ARPA was also paying for. Vinton Cerf and Robert Kahn published a design for doing so in 1974, which became the Transmission Control Protocol and Internet Protocol (TCP/IP). Its central choice was to keep the network simple and to put the work of checking and reordering data in the computers at each end. On January 1, 1983, the ARPANET switched from its original protocol to TCP/IP, and the change let it be split into a research network and a separate military network (MILNET).

Names and a wider network followed, beginning with the DNS, which was introduced in 1984 and replaced a single list of host names kept at the Stanford Research Institute with a directory spread across many servers (Module 3 explains how it works). In 1986 the National Science Foundation (NSF) created NSFNET, a backbone, or long-distance trunk network, that linked its supercomputer centers and the regional networks of universities at an initial speed of 56 kilobits per second (kbps). The NSF spent $200 million on the program from 1986 to 1995, and its acceptable-use policy barred traffic "not in support of research and education" (Leiner and others). The ARPANET itself was retired in 1990, its work taken over by NSFNET and the regional networks.

2d. The Web, 1989-1994

The internet of the 1980s served mostly researchers, through programs that needed typed commands, and the World Wide Web opened it to everyone else. Tim Berners-Lee, a British scientist at CERN, the European particle-physics laboratory near Geneva, wrote his first proposal for it in March 1989 and a formal one with Robert Cailliau in November 1990. By Christmas 1990 he had the first web server and browser running on a NeXT computer, at the address info.cern.ch.

The web joined three simple ideas, the first being that every page has an address, the Uniform Resource Locator (URL). Pages are written in the Hypertext Markup Language (HTML), which lets any page link to any other, and they are fetched with the Hypertext Transfer Protocol (HTTP). Berners-Lee announced the system publicly on August 6, 1991, in a message to an online discussion group, alt.hypertext, and invited others to build on it.

Two events in 1993 made it spread, the first on April 30, when CERN put the web's software into the public domain so that anyone could use and build on it without paying. Early that year the National Center for Supercomputing Applications (NCSA) at the University of Illinois released Mosaic, a browser that showed pictures alongside text and was soon available for Windows and Macintosh computers. By late 1993 there were more than 500 web servers, and the web carried about 1% of internet traffic (CERN).

In October 1994 Berners-Lee founded the World Wide Web Consortium (W3C) at the Massachusetts Institute of Technology, with support from CERN, the European Commission and the Defense Advanced Research Projects Agency (DARPA), as ARPA had been renamed, and the W3C still develops the web's

standards (Module 12). By the end of 1994 the web had 10,000 servers, 2,000 of them commercial (CERN), and businesses had begun to treat it as a place to sell.

2e. Going Commercial, 1991-1998

The internet became a business in the first half of the 1990s. Companies built networks of their own to sell access, and they linked them at exchange points outside NSFNET, where its ban on commercial traffic did not apply. NSF's policy of handing the network to private firms culminated in April 1995 with the end of its funding for the NSFNET backbone (Leiner and others). From then on, the long-distance lines of the American internet belonged to companies.

The first boom began that year, when Netscape, whose browser was written by former members of the Mosaic team, sold shares to the public on August 9, 1995. The price of its shares roughly doubled on the first day of trading. Dial-up access, through a modem over a home phone line at speeds of up to 56 kbps by the late 1990s, brought millions of households online through providers such as America Online.

Congress rewrote the law of 1934 in the Telecommunications Act, signed on February 8, 1996. It aimed to bring competition to local telephone service, let the regional Bell companies into long distance once their local markets were open, and loosened the rules on how many broadcasting stations one company could own and on what cable companies could charge. Its Section 230 shields online services from liability for most of what their users post. The same section declared it American policy to preserve the internet's free market "unfettered by Federal or State regulation", a line later cited in the argument over net neutrality (Module 12).

The last piece of central coordination moved out of direct government hands as well. The internet's addresses and its top-level domain names, such as .com, had been managed under US government contracts. In 1998 the Commerce Department began handing that job to a new California nonprofit, the Internet Corporation for Assigned Names and Numbers (ICANN), and Module 12 explains how ICANN works and how American oversight of it ended.

2f. Boom and Bust, 1996-2002

Money poured into the internet and the networks under it in the late 1990s. By a widely cited estimate, telecommunications companies invested more than $500 billion, most of it borrowed, in fiber, switches and wireless networks in the five years after the 1996 Act. Investors bid up the shares of internet companies, and the Nasdaq Composite, the stock index where most of them traded, closed at a peak of 5,048.62 on March 10, 2000.

Capacity was built faster than demand grew, and the price of carrying data fell. The Nasdaq fell 78%, to 1,114.11 on October 9, 2002, and by common estimates about $5 trillion of market value had been lost by the end of that year. Carriers that had borrowed to build went bankrupt, among them Global Crossing, which had laid undersea cables and filed for bankruptcy protection on January 28, 2002.

The largest failure was WorldCom, the second-largest American long-distance carrier and the owner of one of the biggest internet backbones. On June 25, 2002 it disclosed that $3.8 billion of operating expenses had been booked as if they were investments, and late on Sunday, July 21, 2002 it filed for bankruptcy protection. With $107 billion of assets and $41 billion of debt, it was the largest bankruptcy in American history to that date (CNN Money). Its former chief executive, Bernard Ebbers, was later convicted of fraud.

The bust left the networks behind, because the fiber laid in the boom stayed in the ground, much of it unused at first, and new owners bought it for a fraction of what it had cost to build. The traffic of the next decade, from broadband homes and then from video, grew into that capacity, and the bust became the standard warning for later investment booms in networks and data centers (Module 10).

Homes moved off dial-up in the same years. By the early 2000s, digital subscriber lines (DSL), which send data over telephone wires at frequencies above the voice, and cable modems, which use the cable-television network, were replacing modems that tied up the phone line. Both were always on and many times faster, and both reused wires that the telephone and cable companies already owned, which gave those companies their hold on home broadband (Module 6 explains the technologies and the market).

2g. The Smartphone, 1999-2008

Phones reached the internet before smartphones did, first in Japan, where the mobile arm of Nippon Telegraph and Telephone (NTT), NTT Docomo, launched its i-mode service in February 1999. It let millions of subscribers read email and simple web pages on their handsets. Research In Motion's BlackBerry made email on the move a habit of office workers in the early 2000s.

The iPhone changed what a phone was for when Steve Jobs, Apple's chief executive, announced it at the Macworld conference in San Francisco on January 9, 2007, as a phone, a music player and an internet device in one, with a touch screen in place of a keyboard. It went on sale in the United States on June 29, 2007, at $499 and $599, only on the network of Cingular, which was then being renamed AT&T. It reached the internet over Wi-Fi or over a slow second-generation mobile data service (Apple).

The step that turned smartphones into internet machines came a year later. Apple opened its App Store on July 10, 2008, letting outside developers sell programs for the phone, and the first phone running Google's Android software went on sale in October 2008. Google gave Android away as open-source software, whose code anyone may use and change, and it later became the system on most of the world's smartphones. The mobile networks those phones needed are the subject of Module 7, which traces them from the first generation to the fifth.

The table traces the network's growth through these stages, with counts of hosts, computers with their own address on the network, in the early years and estimates of people later. The host counts were compiled from surveys of the address system. The two measures differ, because in the 1990s one host might serve many users, while today one person may use several devices.

Date	Stage	Hosts or users	Source
End of 1969	The ARPANET's first sites linked	4 sites	Leiner and others; Zakon
August 1981	Before the switch to TCP/IP	213 hosts	Zakon
October 1984	The DNS introduced	1,024 hosts	Zakon
December 1987	NSFNET in service	28,174 hosts	Zakon
July 1989	Berners-Lee's web proposal (March)	130,000 hosts	Zakon
July 1992	A year after the web's announcement	992,000 hosts	Zakon
End of 1994	The W3C founded (October)	10 million web users	CERN
July 1995	NSFNET backbone defunded (April)	6,642,000 hosts	Zakon
1995	Netscape's share sale (August)	About 39 million users, 0.7% of the world	DataReportal series (secondary)
July 1997	After the 1996 Act	19,540,000 hosts	Zakon

Date	Stage	Hosts or users	Source
1999	The height of the boom	More than 150 million users	Computer Industry Almanac (secondary)
2000	The Nasdaq’s peak (March)	About 361 million users	DataReportal series (secondary)
2005	Two years before the iPhone	About 1.0 billion users, 15% of the world	DataReportal series (secondary)

Zakon is Robert H. Zakon, “Hobbes’ Internet Timeline,” Request for Comments 2235 (November 1997), which compiled surveys of the domain name system; Leiner and others is “The Past and Future History of the Internet,” *Communications of the ACM* (February 1997). Web users from CERN’s history of the web. Users for 1995, 2000 and 2005 from a chart citing DataReportal, and for 1999 from the Computer Industry Almanac as reported by the University of Missouri-St. Louis (March 2000), all secondary; Pingdom (March 31, 2011) gives “less than 40 million” for 1995.

Worked example — From 39 Million Users to 6 Billion

Two starting figures are available: CERN counted 10 million web users at the end of 1994, and a series compiled from DataReportal gives 39.2 million internet users for 1995, which Pingdom’s “less than 40 million” supports. This example starts from the 1995 figure, because it counts the same thing as the ITU’s figure for 2025, people using the internet, while CERN’s counts users of the web, a smaller group in a year when the web was new. The end point is the ITU’s 6.0 billion users in 2025 (Module 1). Step 1, the growth multiple: 6,000 million ÷ 39.2 million = 153.1. Step 2, the period: 2025 − 1995 = 30 years. Step 3, the compound annual growth rate, the steady yearly rate that would produce that multiple: 153.1 raised to the power 1/30 is 1.183, so users grew 18.3% a year on average. Step 4, the doubling time: at 18.3% a year, users doubled every ln 2 ÷ ln 1.183 = 4.1 years, where ln is the natural logarithm, a little over seven doublings in 30 years. Starting instead from CERN’s 10 million at the end of 1994 gives a multiple of 600 over 31 years, 22.9% a year and a doubling every 3.4 years, so the choice of starting point moves the answer by 4.6 percentage points. Step 5 shows that the average hides a slowdown: from 1995 to 2005, when the same series reached 1.0 billion, growth averaged 38.3% a year, a doubling every 2.1 years. From 2005 to 2015, when it reached 3.0 billion, growth averaged 11.6%, a doubling every 6.3 years. From 2015 to the ITU’s 6.0 billion in 2025 it averaged 7.2%, a doubling every 10 years, or 6.4% if the series’ own 5.6 billion for 2025 is used. Step 6, the ceiling: 6.0 billion is 74% of the world’s people, which implies about 8.1 billion people in all. At 7.2% a year, and if the population stayed at 8.1 billion, an assumption, users would outnumber it in about 4.3 years, which cannot happen, so growth must keep slowing as the people still offline prove harder and costlier to reach (Module 11).

Module 2 Key Points

From Telegraph to Smartphone

• Morse’s line of May 24, 1844 was built with public money and leased to private investors in 1847, setting America’s pattern of private networks.

• The Bell System was broken up on January 1, 1984 into AT&T and seven regional companies; by 2006 six had merged into two.

• The ARPANET’s first switch was installed at UCLA in September 1969, and the network switched to TCP/IP on January 1, 1983.

• Berners-Lee proposed the web in March 1989, and CERN put it in the public domain on April 30, 1993.

• Telecom companies invested more than $500 billion after the 1996 Act; WorldCom’s bankruptcy of July 21, 2002 was the largest to that date.

• In the worked example users grew 18.3% a year from 1995 to 2025, but only 7.2% a year in the last decade.

Module 2 Quiz

Check Your Understanding

1. What did the breakup of the Bell System on January 1, 1984 do?

A) It moved local telephone service to seven regional companies.

B) It put AT&T's long-distance lines under the FCC as a public agency.

C) It ended the 1956 consent decree and let AT&T keep its local telephone companies while it entered computing.

D) It sold Bell Labs to the government.

2. Why was the ARPANET built, according to the histories written by its builders?

A) It was built to survive a nuclear war.

B) It was built so that researchers funded by ARPA could share scarce and expensive computers.

C) It was built to sell internet access to businesses.

D) It was built to replace the Bell System's telephone lines.

3. What happened on January 1, 1983?

A) CERN put the web into the public domain.

B) The NSF stopped funding its backbone.

C) The ARPANET switched from its original protocol to TCP/IP.

D) The Bell System was broken up into AT&T and seven regional companies.

4. What did WorldCom do on July 21, 2002?

A) It merged with MCI to become the second-largest long-distance carrier.

B) It laid the first transatlantic fiber cable between New York and London.

C) It filed for bankruptcy.

D) It disclosed that $3.8 billion of operating expenses had been booked as investments.

5. In the worked example, how fast did internet users grow on average from 1995 to 2025?

A) About 18.3% a year

B) About 7.2% a year

C) About 38% a year

D) About 153 times a year

Module 3 — How Data Travels

Packets, Addresses, Protocols, and Routes

Recommended Reading / Reference

Book: James F. Kurose and Keith W. Ross, *Computer Networking: A Top-Down Approach* (9th edition, 2025), one of the most widely used university introductions, which explains each protocol from the application down to the wire

Reference: Geoff Huston's yearly routing reports for APNIC, the Asia-Pacific address registry ("BGP in 2025," February 4, 2026), and Cloudflare Radar's public figures on the protocols in use

3a. Packets, Not Circuits

The telephone network moved calls by circuit switching, in which placing a call reserved a path through the exchanges between the two phones for the whole call, whether anyone spoke or not (Module 2 tells its history). That guaranteed the quality of the call but wasted capacity, since much of a conversation is silence and most computer traffic comes in short bursts.

The internet uses packet switching instead, cutting a sender's data into packets, each carrying a header, a label with the source and destination addresses and a few other fields, followed by a slice of the data itself. On most networks a packet holds at most about 1,500 bytes, where a byte is eight bits, so a photo of 3 megabytes (MB) travels as about 2,000 packets. Each packet travels on its own, passed from router to router, and the receiving device puts them back together in order.

A router is a computer built to forward packets. It reads each packet's destination address, looks it up in its table of routes, and sends the packet out on the link that leads toward that destination. To forward a packet a router needs to read only its header, and it keeps no record of the conversation the packet belongs to, which is why routers can be fast and the network can stay simple.

Packet switching shares lines efficiently, because thousands of conversations interleave on the same fiber and capacity is used only when someone sends. It is also resilient, since routers can steer packets around a broken link. The cost is that the internet guarantees nothing. When a link is busy, packets wait in queues inside routers, which adds delay, and when a queue is full they are dropped. Engineers call this "best effort" delivery, and the protocols in the rest of this module are largely about making it reliable enough to use.

3b. Addresses: IPv4, Its Market, and IPv6

Every device that sends or receives packets needs an IP address. The version used since 1983, IPv4, writes each address as 32 binary digits, usually shown as four numbers from 0 to 255, such as 192.0.2.1. Thirty-two bits allow 2 to the 32nd power, or 4,294,967,296, addresses, fewer than the number of people alive and far fewer than the phones, computers and other devices now connected. Addresses are handed out in blocks, and the Internet Assigned Numbers Authority (IANA) allocates large blocks to five regional internet registries, which pass them on to carriers, companies and universities in their regions (Module 12 covers these bodies). IANA gave out its last blocks of IPv4 addresses on February 3, 2011, and the registries' own stocks have since run out or been rationed. Since then a network that needs more IPv4 addresses has mostly had to buy them from one with addresses to spare.

Two things kept the internet growing on IPv4. The first is network address translation (NAT), which lets a home router share one public address among every device in the house, giving each a private address that works only inside the home. Carriers do the same on a much larger scale, putting many customers behind each public address. The method works, but it hides devices from one another, which complicates video calls, games and anything else that needs a device to accept an incoming connection.

The second is a market, in which IPv4Center, a broker, recorded 596 transactions moving 5,016,064 addresses for about $58.46 million in the first half of 2026 (IPv4Center, July 2, 2026). The average price per address across those deals was $20.04, and the median $20.00, down 35.7% from an average of $31.15 a year earlier and 60-67% below the peak of $50-60 in 2021-22. Because large blocks sell for less per address, dividing the total value by the total number of addresses gives only $11.65.

The lasting answer is IPv6, the next version of the protocol, defined in the 1990s. Its addresses are 128 bits long, which allows about 3.4×10^{38} of them, enough to give every device its own public address with no need to share. IPv6 cannot talk directly to IPv4, so networks must run both side by side during a transition that has already lasted decades. Large websites and carriers switched IPv6 on permanently on June 6, 2012, a date the industry called World IPv6 Launch.

The transition passed a milestone on March 28, 2026, when 50.10% of the users reaching Google did so over IPv6, the first day above half, after a previous high of 49.56% on June 21, 2025 (Google’s figures, as reported by the Internet Society’s Pulse, April 21, 2026). Other measures come out lower because they count something else. Cloudflare Radar found 40.1% of the requests reaching Cloudflare’s network using IPv6 (reported April 17, 2026). Google counts the users who connect to it, while Cloudflare counts requests, which include traffic from servers and automated programs as well as from people.

Adoption varies widely by country, and by Google’s count in April 2026 about 73% of users in France and 72% in India reached it over IPv6, against 65% in Saudi Arabia, 17% in Italy, 10% in Spain and 4% in Egypt. Where the largest carriers have switched IPv6 on, most of their customers use it without knowing, because phones and home routers choose it automatically when it is offered.

3c. TCP, UDP, and QUIC

IP moves packets but promises nothing about them: they may arrive late, out of order, twice or not at all. Transport protocols, running in the computers at each end rather than in the network, decide what to do about that. The two oldest, from the internet’s first years, take opposite approaches.

The Transmission Control Protocol (TCP) makes delivery reliable. Before sending any data it sets up a connection with a three-message handshake, which costs one round trip, the time for a message to reach the other end and a reply to come back. It numbers the data it sends, the receiver acknowledges what arrives, and the sender resends whatever goes missing. TCP also slows down when packets are lost, on the assumption that the network is congested, and this restraint by every sender keeps shared lines from jamming. Reliability has a cost, because TCP delivers everything in order, so one lost packet holds up all the data behind it until a fresh copy arrives, a delay engineers call head-of-line blocking. For a file download that costs a fraction of a second, but for a live call it means a frozen picture, because data that arrives late is useless.

The User Datagram Protocol (UDP) does almost nothing, sending each packet with no handshake, no numbering and no resending. That suits applications that prefer speed to completeness, such as voice and video calls, online games and most DNS lookups (section 3d), which deal with the occasional lost packet themselves.

QUIC, a newer transport protocol that runs on top of UDP, does TCP's job in a different way. Google designed it, and the Internet Engineering Task Force (IETF) published it as a standard in May 2021 (Module 12 covers the IETF). QUIC builds encryption in, so that setting up both the connection and its encryption takes one round trip instead of the two that TCP and encryption need separately. It also carries each request in a stream of its own, so a lost packet holds up only the stream it belongs to.

The web's own protocol, the Hypertext Transfer Protocol (HTTP), has three versions in use. HTTP/1.1, from the late 1990s, handles one request at a time on each connection. HTTP/2, standardized in 2015, sends many requests at once over a single TCP connection, and HTTP/3 does the same over QUIC. The table gives their shares of the requests reaching Cloudflare, along with the shares encrypted.

Requests reaching Cloudflare's network	Share	Period
Using HTTP/2	52.46%	Third quarter of 2026
Using HTTP/1.x (versions 1.0 and 1.1)	27.37%	Third quarter of 2026
Using HTTP/3, over QUIC	20.17%	Third quarter of 2026
Encrypted, using HTTPS (HTTP inside an encrypted connection, section 3e)	95.10%	Third quarter of 2026
Human traffic encrypted	98.60%	July 2026
Automated (bot) traffic encrypted	87.80%	July 2026

Cloudflare Radar data as compiled by TechnologyChecker.io (April 9, 2026, updated October 1, 2026); third-quarter shares are provisional. They describe traffic reaching Cloudflare, which serves a large share of websites but not all of them.

Counts of adoption depend on what is counted. W3Techs, which surveys websites, found 38.8% of sites supporting HTTP/3 in April 2026, more than the share of requests that used it. A browser usually reaches a site over HTTP/2 first and switches to HTTP/3 only once it learns the site offers it, and some networks block the UDP traffic that QUIC needs. Meta said in October 2020 that more than 75% of its own internet traffic already ran on QUIC.

3d. DNS: The Internet's Directory

People use names and routers use numbers, and the DNS translates one into the other, so that a browser asked for www.example.org must first learn that name's IP address. The system is a hierarchy, with the root at the top, which knows the servers for each top-level domain, the last part of a name, such as .org, .com or .uk. Each top-level domain's servers in turn know which servers are authoritative, that is, hold the official records, for every name registered under it.

A lookup runs in steps, starting when the device asks a recursive resolver, a server run by its internet service provider or by a public service such as Cloudflare's 1.1.1.1 or Google's 8.8.8.8. If the resolver does not have the answer, it asks a root server where to find .org, asks a .org server where to find example.org, and asks that domain's authoritative server for the address. Each answer comes with a time to live, the number of seconds it may be reused, so a resolver that has looked up a popular name recently answers at once from memory.

The root is served by 13 named identities, lettered A to M and run by 12 organizations, since Verisign runs both A and J. The number 13 dates from an early limit on the size of a DNS reply. Each identity is not one machine but many copies, called instances, in different cities, which share an address through anycast: every copy announces the same address, and routing delivers each query to the nearest one. By the operators' published count there were 2,003 instances on August 11, 2026, and together they answered about 1.6 million queries a second in July 2026 (the root operators' figures, as reported in August 2026).

Caching keeps most traffic away from the root, because resolvers remember where .com and .org are for up to two days. The load falls instead on resolvers and on authoritative servers. Cloudflare alone said it was answering about 67 million DNS queries a second at the end of 2025, as a resolver for users and as the authoritative server for its customers' domains (Cloudflare Radar, December 15, 2025).

Because nearly every connection starts with a lookup, the DNS is a point of control and of failure. Governments order resolvers to give false answers for blocked sites (Module 13), and errors in DNS records have caused some of the largest outages (Module 10). Lookups were long sent unencrypted, so anyone on the path could see which sites a user visited. DNS over HTTPS, standardized in 2018, wraps lookups in the same encryption as web pages (section 3e).

3e. Encryption: HTTPS and the Post-Quantum Switch

Packets cross many networks owned by others, so anyone along the path could read them unless they are encrypted. Transport Layer Security (TLS) encrypts a connection between two programs, typically a browser and a web server, and HTTPS is HTTP carried inside TLS. The networks in between can still see the addresses at each end, the timing and the volume, and often the name of the site, but not the pages, passwords or messages.

Before encrypting, the two sides run a handshake, in which the server proves its identity with a certificate, a digital document signed by a certificate authority that browsers trust, and the two sides agree on a secret key for the session. In TLS 1.3, the version published in August 2018, the handshake takes one round trip, half the two that the version before it needed.

Encryption is now the norm, as the table in section 3c shows, but it was not in the web's first two decades, when most pages were sent unencrypted and browsers marked only shopping and banking pages as secure. The change came in the late 2010s, as browsers began warning users about unencrypted sites and free certificates became available to every site.

The next change is already under way. A large enough quantum computer, a machine that computes with the rules of quantum physics, could break the methods used today to agree on session keys. No such machine was known to exist as of October 2026, but an adversary could record encrypted traffic now and decrypt it once one does, a threat that security engineers call "harvest now, decrypt later". Post-quantum key agreement uses mathematical problems thought to resist quantum computers, under standards that the US National Institute of Standards and Technology (NIST) published in August 2024.

The share of human web traffic reaching Cloudflare that used post-quantum key agreement nearly doubled in 2025, from 29% at the start of the year to 52% in early December (Cloudflare Radar, December 15, 2025). More than a quarter of requests from iPhones, running Apple's iOS, used it by then. Because both ends must support a method before a connection can use it, these shares rise when the makers of browsers and operating systems switch it on by default, not when users choose it.

3f. Routing and BGP

Inside its own network, an operator routes packets however it likes. Between networks, the roughly 77,000 autonomous systems (ASes) of Module 1, routing runs on the Border Gateway Protocol (BGP). Each network announces to its neighbors the blocks of addresses it can deliver to, called prefixes and written like 1.1.1.0/24, which means the 256 addresses whose first 24 bits match. Neighbors pass the announcements on, adding their own number to the path, until every network has heard a route to every prefix.

Each router then chooses among the routes it has heard according to its owner's policy, which usually favors routes through customers that pay over routes through providers that charge (Module 4). When a router forwards a packet, the most specific matching route wins, the one covering the fewest addresses, whatever the policy. That rule lets networks steer traffic precisely, and it is also what lets mistakes and attacks spread.

The table of routes keeps growing, and Geoff Huston, chief scientist of APNIC, reported that the IPv4 table seen by every network that carries the full set of routes passed 1 million entries in 2025, and that the IPv6 table grew from 201,000 entries to 222,000 ("BGP in 2025," February 4, 2026). Every router at the core of the internet must hold the full table in fast, expensive memory, so its growth is a cost that every large network pays.

BGP was first specified in 1989 for a small community whose members trusted one another, and it still believes what it is told. A network can announce addresses that belong to someone else, by mistake or on purpose, and its neighbors will often pass the claim on. A leak is the mistaken passing on of routes that a network should have kept to itself, and a hijack is the announcement of another network's addresses as one's own.

Two cases show the effect, the first in February 2008, when Pakistan Telecom, ordered by its government to block YouTube inside Pakistan, announced a more specific route to YouTube's addresses to send local requests nowhere. It passed the route to its international provider, which spread it worldwide, and YouTube was cut off for a large part of the internet (Doug Madory of Kentik, a network-monitoring firm, September 1, 2023). The second came on June 27, 2024, when a Brazilian carrier, Eletronet, began announcing Cloudflare's public DNS address, 1.1.1.1, as its own at 18:51 Coordinated Universal Time (UTC), and a second Brazilian network leaked the block around it a minute later. By Cloudflare's account, users in about 300 networks in 70 countries were affected until the leak was fixed at 02:28 UTC the next day (BleepingComputer, July 5, 2024).

The main defense is the Resource Public Key Infrastructure (RPKI), a system of digitally signed records in which an address holder states which networks may announce its addresses. A network that checks these records can drop a route whose origin is not authorized, although RPKI cannot catch a leak that keeps the right origin. Not every network checks, so the protection is partial, and Cloudflare reported only "continued improvement" in routing security in 2025, without figures.

3g. Latency, Bandwidth, and Jitter

Speed in everyday talk means bandwidth, the rate at which a connection moves data, in megabits or gigabits per second (Mbps or Gbps). Latency is delay, the time a packet takes to arrive, usually measured as the round-trip time to a server and back in milliseconds (ms). Jitter is the variation in that delay from one packet to the next, loss is the share of packets that never arrive, and each of the four limits different things.

Latency has three main sources, and distance sets a floor, because light in fiber travels at a fixed speed and a packet to another continent must cover thousands of kilometers each way (Module 4 gives the physics). Each router along the way adds a little processing time, but the largest and most variable share often comes from queues, as packets wait behind others on a busy link, in a home Wi-Fi router or on a mobile network's radio link. Equipment with oversized queues can add hundreds of milliseconds under load, a problem engineers call bufferbloat.

Latency, not speed, limits many of the things people do online, because much of what the internet does is conversation, in which one side must wait for the other's reply before it can continue. A web page needs a

DNS lookup, a connection, an encryption handshake and then requests for its parts, each costing at least one round trip (worked example). A faster line shortens only the time spent moving data. Once a line can deliver a page's data quickly, more speed changes little, while every millisecond cut from the round trip still helps.

Other uses also depend on delay or its variation more than on bandwidth. A video call needs only a few megabits per second, but delay makes people talk over each other, and jitter forces the receiving device to hold packets in a buffer to smooth the flow, which adds delay of its own (Module 10 traces a call). Online games, the remote control of machines and financial trading depend on milliseconds, and some traders pay for the shortest routes between exchanges (Module 4). Downloads, backups and ultra-high-definition (4K) video are the uses that need bandwidth.

This is why many of the largest gains of the past decade came from cutting round trips rather than adding speed. Content delivery networks (CDNs) put servers near users (Module 4), TLS 1.3 cut the encryption handshake to one round trip, and QUIC merged it with the connection. It is also why satellites in low orbit, much closer to the ground than geostationary ones, can serve uses that the older satellites could not (Module 9).

Worked example — The Half-Second Before a Web Page Appears

Follow a first visit to a page on a server on another continent, with these illustrative assumptions: a round trip of 10 ms to the user's DNS resolver, and 80 ms from the resolver to the domain's authoritative server, with the root and .org answers already cached. The round trip from the user to the web server is 100 ms, the server takes 50 ms to prepare the page, the page's main file is 100 kilobytes (thousands of bytes), and the user's line runs at 100 Mbps. Over HTTP/2, which runs on TCP with TLS 1.3, the steps run one after another. Step 1, the DNS lookup: 10 + 80 = 90 ms. Step 2, the TCP handshake: one round trip to the server, 100 ms, so 190 ms have passed. Step 3, the TLS handshake: one more round trip, reaching 290 ms. Step 4, the request and the start of the response: one more round trip plus 50 ms at the server, reaching 440 ms. Step 5, the transfer: 100 kilobytes is 800 kilobits, which at 100 Mbps takes 8 ms if the server sends it all at once (an assumption: TCP and QUIC start slowly, adding round trips), so the page's first file is complete at 448 ms. Over HTTP/3, QUIC replaces Steps 2 and 3 with a single handshake of one round trip, and the same file is complete at 90 + 100 + 100 + 50 + 8 = 348 ms, a saving of 100 ms, or 22.3%. QUIC saves the round trip because it sets up the connection and the encryption in the same exchange, while TCP must finish its own handshake before TLS can begin. With the older TLS 1.2, whose handshake took two round trips, the HTTP/2 page would have taken 548 ms. Now test the claim that latency matters more than speed. Doubling the line to 200 Mbps cuts the transfer from 8 ms to 4 ms and the HTTP/2 total to 444 ms, a saving of 4 ms. Halving the round trip to the server to 50 ms, by using a server on the user's own continent, cuts the HTTP/2 total to 298 ms and the HTTP/3 total to 248 ms. In the HTTP/2 case, 87.1% of the 448 ms is spent waiting on the lookup and the round trips, and 1.8% on moving the data. A real page then fetches style sheets, scripts and images, each needing further round trips, so the full load takes longer, but the lesson holds: round trips, not megabits, set the wait.

Module 3 Key Points

How Data Travels

• Data travels in packets of up to about 1,500 bytes, each forwarded separately, and the internet delivers them on a best-effort basis only.

• On March 28, 2026, 50.10% of users reached Google over IPv6, the first day above half; Cloudflare counted about 40% of requests.

• An IPv4 address averaged $20.04 in the first half of 2026, down 35.7% in a year (IPv4Center).

• The DNS root has 13 identities run by 12 operators, served by 2,003 instances in August 2026.

• About 95% of requests to Cloudflare were encrypted in the third quarter of 2026, and about 20% used HTTP/3.

• In the worked example QUIC saves one 100-millisecond round trip, while doubling the line’s speed saves only 4 milliseconds.

Module 3 Quiz

Check Your Understanding

1. What share of users reached Google over IPv6 on March 28, 2026?

A) About 40%
B) 50.10%, the first day above half
C) 95.10%, the same share as encrypted requests to Cloudflare
D) About 74%, the share of the world's people who use the internet at all

2. Why does HTTP/3 over QUIC save a round trip when a connection opens?

A) Because QUIC sends packets faster than TCP does.
B) Because QUIC skips encryption on the first request.
C) Because QUIC keeps DNS answers in the browser.
D) Because QUIC sets up the connection and its encryption in one exchange, where TCP and TLS need one each.

3. How is the DNS root served?

A) By one computer in Virginia
B) By 13 servers, one per operator
C) By 13 identities run by 12 operators, through about 2,000 instances
D) By ICANN alone, from a single data center whose address every resolver keeps in memory

4. What does BGP do?

A) Encrypts web pages
B) Tells networks which addresses each can reach
C) Turns names into addresses
D) Resends lost packets

5. In the worked example, what does doubling the line's speed from 100 to 200 Mbps save?

A) It saves 50 ms.
B) It halves the load time, from 448 ms to 224 ms.
C) It saves 100 ms, as much as moving from HTTP/2 to HTTP/3.
D) It saves only 4 ms of 448 ms.

Module 4 — The Backbone

Fiber, Exchanges, and Who Pays Whom

Recommended Reading / Reference

Book: William B. Norton, *The 2014 Internet Peering Playbook: Connecting to the Core of the Internet* (DrPeering Press, 2014 edition), a practical guide to how and why networks agree to exchange traffic, by a co-founder of Equinix, the data-center and exchange company, and a former chair of the North American Network Operators' Group

Reference: TeleGeography's research on transit prices, and the public ranking of the internet's autonomous systems (AS), the separate networks it is made of, that the Center for Applied Internet Data Analysis (CAIDA) at the University of California, San Diego publishes as AS Rank

4a. Light in Glass

Almost every long-distance link on the internet is a strand of glass. An optical fiber has a core of very pure glass inside a layer of glass that bends light less, so light sent down the core reflects off the boundary and stays inside, even around gentle bends. A laser at one end encodes bits in the light, and a receiver at the other end reads them. Copper wire and radio carry data too, but glass loses far less of the signal over distance and carries far more of it.

Light in glass travels at about 200,000 km a second, about two-thirds of its speed in a vacuum, 299,792 km a second. The glass slows it by its refractive index, the ratio of light's speed in a vacuum to its speed in a material, which is about 1.46 in standard fiber. That speed sets a floor on delay that no equipment can remove: each 1,000 km of fiber adds 5 milliseconds (ms) one way. Module 3 explains why delay, more than raw speed, limits calls and games, and the worked example applies the floor to New York and London.

Long links use infrared light at wavelengths of around 1,550 nanometers, or billionths of a meter, where glass absorbs the least. Even there the light fades, because standard fiber loses about 0.2 decibels a km, where a decibel is a measure of power ratios and 3 decibels is a halving, so half the light is gone every 15 km. A typical design therefore places an amplifier about every 80 km (standard values, as used in an engineering exercise of the Karlsruhe Institute of Technology). These erbium-doped fiber amplifiers, short lengths of fiber laced with the element erbium and energized by a laser, boost every wavelength at once without turning the light into electricity.

Each fiber carries many signals at once, one on each wavelength, or color, a method called wavelength-division multiplexing (WDM). Coherent optics, which encode bits in the phase of the light wave as well as in its strength, have multiplied what each wavelength carries. Ciena's WaveLogic 6 Extreme, sold from the first half of 2024, sends up to 1.6 terabits per second (Tbps) on one wavelength over metro distances and 800 gigabits per second (Gbps) over the longest links, and Nokia's ICE7 sends up to 1.2 Tbps (company figures). Nokia says its equipment can carry close to 100 Tbps on a single fiber pair, one fiber for each direction, by using wider bands of wavelengths (product page, updated June 11, 2026).

Phil McKinney, chief executive of CableLabs, the cable industry's research body, put the figure for one strand at 50,000 gigabits per second, or 50 Tbps, at an industry conference in 2025 (Fiber Broadband Association, April 2, 2026). Capacity therefore grows mostly by lighting more wavelengths on fibers already in the ground, and by better equipment at their ends, rather than by laying new cable. That is a main reason why the price of carrying traffic keeps falling (section 4f).

4b. Long-Haul Networks on Land

A long-haul network links cities hundreds or thousands of km apart. On land its cables run in ducts along railways, highways and pipelines, where rights of way already exist, and they meet in buildings where carriers and data centers connect to one another. Telephone companies such as AT&T and Verizon own such routes. So do wholesale carriers such as Lumen and Zayo, which sell capacity to other networks, and the largest cloud companies, which build their own (Module 10).

Carriers sell long-haul capacity in three forms. Dark fiber is a strand leased with no equipment on it, which the buyer lights itself. A wavelength is one channel between two points, lit by the seller. IP transit is the service of carrying traffic to and from any address on the internet, priced by capacity (sections 4c and 4f). Coherent optics can carry 800 Gbps over 3,000 km or more before the signal must be turned back into electricity and regenerated, by Nokia's account.

Distance on the ground is longer than distance on the map. Cables follow roads, rails and coasts, and every detour adds delay at 5 ms per 1,000 km of glass. For most uses a few ms do not matter. For automated trading between financial centers they do, and the financial trading community values the most direct routes (The Broadcast Bridge, 2015), as the worked example shows for the Atlantic.

4c. Transit, Peering, and the Tier 1 Networks

The internet is made of independent networks called autonomous systems (AS), each run under one routing policy and identified by its own number. They tell one another which addresses they can reach using the Border Gateway Protocol (BGP), which Module 3 describes. Every network must find a way to reach every address, and it can do so in one of two ways, a choice that decides who pays.

Transit is a paid service, in which a network buys from a larger one the right to send traffic to, and receive traffic from, the whole internet, at a price per Mbps (megabits per second) of capacity per month. A small ISP or a hosting company usually buys transit from two or more providers, so that the failure of one does not cut it off.

Peering is an exchange of traffic between two networks' own customers, and only that: neither gives the other a route to the rest of the internet. Most peering is settlement-free, meaning that no money changes hands, because both sides gain about equally. Two networks peer either over a direct cable between their routers, called a private network interconnect, or across an exchange point (section 4d). When one side sends far more than it receives, or is much smaller, the other may ask to be paid, an arrangement called paid peering.

A Tier 1 network is one that reaches every other network without buying transit from anyone, because it peers with all the other Tier 1 networks. There is no official list, and the label is partly a marketing claim. Size is usually measured by the customer cone, the set of networks that reach the internet through a given network. On CAIDA's AS Rank, read on October 25, 2026, the largest is autonomous system 3356 (AS3356), run by Lumen and formerly by Level 3, which buys transit from no one, peers with 75 networks and has 6,506 direct customers, with 52,292 networks in its customer cone.

Rankings differ by method: bgp.tools, another public ranking, lists Lumen, Cogent, GTT, Tata Communications of India and NTT of Japan, through their American-registered networks, as the five networks registered in the United States with the largest customer cones (read the same day). Below the Tier 1 networks sit regional carriers that buy some transit and peer where they can, and at the edge sit the

access networks that serve homes (Module 6). The content and cloud companies sit outside this old hierarchy, a change section 4e explains.

4d. Internet Exchange Points

An internet exchange point (IXP) is a place where many networks connect to a shared switch and exchange traffic with any of the others that agree, instead of running a separate cable to each. A network that joins can peer with hundreds of others through a single port. Some exchanges are run by associations of their members or by nonprofit bodies, and others by companies such as DE-CIX of Frankfurt.

Exchanges keep local traffic local: without one, two networks in the same city may hand traffic to each other through a transit provider in another city or country, paying for the privilege and adding delay. Brazil shows the effect at scale. IX.br, run by the Brazilian Network Information Center (NIC.br), the nonprofit that manages the country's .br domain, operates exchanges in 39 metropolitan areas. The table gives the records set at the largest exchanges in 2026, with the number of networks that meet at each.

Exchange	Record peak	When, and what drove it	Networks connected
IX.br São Paulo	35.9 terabits per second (Tbps); TELETIME calls it the world's largest exchange by traffic	June 25, 2026, a World Cup day with six matches streamed online	Not given separately
IX.br, all 39 metropolitan areas	55 Tbps	The same day	About 3,800 autonomous systems, with about 7,000 interconnections
DE-CIX Frankfurt	21.47 Tbps, its first time above 21 Tbps	Champions League evenings in October 2026, 20 minutes after kick-off	Almost 1,100 networks
DE-CIX, all its exchanges	32.10 Tbps	The same evenings	More than 4,600 networks in more than 100 countries

TELETIME (São Paulo), July 22, 2026; IT Brief UK, October 14, 2026, reporting DE-CIX's figures. Peaks are the highest traffic measured at an instant, not averages. The two operators measure their own traffic.

Live football set both records, because matches streamed to millions of homes at once produce sharper peaks than any other use, and an exchange must be built for those minutes, not for the average. DE-CIX said in October 2026 that traffic at Frankfurt had risen by more than 160% since the start of 2020, and Ivo Ivanov, its chief executive, noted that Frankfurt had first crossed 20 Tbps "just three and a half weeks" before the record.

Exchanges are growing fastest outside the old hubs. In its report for 2025, DE-CIX counted more than 4,300 connected networks, up 25% in a year, and 220 Tbps of connected capacity, up 40%. Peak traffic at its Mumbai exchange rose 63% to 2.76 Tbps, at UAE-IX in Dubai it passed 1 Tbps for the first time, and at New York and Dallas it was approaching 2 Tbps (DE-CIX, May 12, 2026). The company's revenue for the year was €70.9 million ($79.6 million), small beside the traffic it carries, because exchanges charge their members for ports rather than for each bit.

Exchanges are where the agreements of section 4c become physical. The networks that meet at one bring their own routers into the same building, connect them to the exchange's switch, and then decide one by one whom to peer with. The exchange itself carries no traffic on its own account and buys no transit for its members.

4e. Content Delivery Networks

A content delivery network (CDN) keeps copies of popular content on servers placed inside or next to the networks that serve users, so most requests are answered nearby rather than from a distant origin. Delay falls, and the long-haul and transit links carry each video once instead of once for every viewer. CDNs steer each user to a nearby server through DNS answers that differ by location, or through anycast, in which many sites announce the same address and routing delivers each user to the nearest one.

Some CDNs sell the service to others. Akamai, founded in 1998, delivers content for paying customers from servers placed in networks around the world. Cloudflare, which sells delivery and security services, says its network spanned 335 cities in more than 125 countries at the end of 2024. It says 13,000 networks connected to it directly and that it was within 50 ms of 95% of the world's internet-connected population (company figures).

The largest streaming companies built their own. Netflix's Open Connect puts Netflix's servers inside ISPs and at exchange points. In March 2016 Netflix said the servers sat in close to 1,000 locations, that Open Connect delivered 100% of its video, and that close to 90% of its traffic went over direct connections to residential ISPs. One server's output rose from 8 Gbps in 2012 to more than 90 Gbps in 2016. Netflix fills its servers during off-peak hours, so that copying films does not compete with watching them, and says connections to Open Connect are always free.

Google does the same through Google Global Cache, servers that it places inside ISPs to deliver YouTube and its other services. The cloud companies also run delivery networks of their own, tied to their data centers (Module 10).

These caches flattened the old hierarchy: content once flowed from a hosting company through a transit provider and a Tier 1 network down to an ISP, and each handoff was paid. Now more of it travels from a content company's network straight into the ISP, over peering or from a cache inside the ISP, so less of it crosses paid transit. The ISP saves on transit and the content company saves on delivery. But the ISP must still build and run the line to every home (Module 6), and that is where the argument over payment in section 4g begins.

4f. Who Pays Whom

Money follows traffic in four directions. Households and businesses pay their ISP for access, and content companies pay data centers, clouds and CDNs to store and deliver their content, or build their own. Networks that cannot reach every address through peering buy transit from those that can. And sometimes a large sender pays an access network for delivery straight into it, under paid peering.

Transit keeps getting cheaper. TeleGeography, a research firm, found that in the second quarter of 2026 the lowest offers in the most competitive markets were $0.03 per Mbps a month for a port of 100-gigabit Ethernet (100 GigE) and $0.07 for a 10 GigE port, excluding local access and installation (report of September 16, 2026). At $0.03, transit for a full 100 Gbps would cost $3,000 a month. Prices for 100 GigE fell 17% a year overall, compounded, from the second quarter of 2023 to the second quarter of 2026; city by city, they fell 11-13% a year in London, New York and Singapore, 11% in São Paulo and 35% in Mumbai.

TeleGeography gives two reasons for the falls. In the global hubs, new transmission equipment drives the cuts, while in Africa, Latin America and Asia, new ocean cables with more fibers (Module 5) bring regional and global carriers into competition, and local peering and local content mean that less of the price reflects the cost of crossing an ocean. Demand is moving to 400 GigE ports, bought mainly for the largest cloud

companies and AI traffic in the United States and Europe, at an average of 3.5 times the monthly charge for 100 GigE.

A rough scale shows where the cost of the internet lies. Suppose one home draws 5 Mbps at the evening peak, an illustrative figure that Module 5's worked example derives. Carrying that peak over transit at $0.03 per Mbps would cost about 15 cents a month, against the tens of dollars a month a home pays for access (Module 6). The large costs lie in the access networks that reach every address, and the dispute of section 4g is over who should pay for them.

4g. Should Content Companies Pay Network Operators? The Cases as Their Advocates Make Them

The case for a fee: Deutsche Telekom, other large European operators and the GSMA, the mobile industry's association, argue that a handful of large traffic generators send most of the traffic that forces operators to upgrade, while the operators alone pay for the upgrades. A study the operators commissioned put the infrastructure cost driven by such traffic at up to €36-40 billion ($40.4-44.9 billion) a year, as the Body of European Regulators for Electronic Communications (BEREC) reported in October 2022. Dollar figures converted from euros and won here use the Federal Reserve's rates for October 16, 2026.

Paid interconnection, they add, is a normal commercial deal, and German courts agree. Deutsche Telekom sued Edge Network Services, a Meta subsidiary that kept sending traffic through 24 private interconnection points at seven locations after its paid contract ended in 2020. A court in Cologne awarded Telekom more than €20 million ($22.5 million) in 2024. On February 10, 2026 the Higher Regional Court in Düsseldorf upheld the claim at more than €30 million ($33.7 million), finding that Meta's unit had accepted a paid contract by continuing to use the links (Steptoe, a law firm). Korea shows that content companies can pay. Netflix's traffic on SK Broadband's network reached 1.2 Tbps in September 2021, when the Korean carrier was claiming network fees from it (The Register, 2023). Netflix reportedly paid SK Broadband 40 billion won (about $30.0 million) when the two settled (Jeong and Ali, *International Journal of Communication*, 2026).

The case against a fee: content companies, their trade association the Computer & Communications Industry Association (CCIA), the Internet Society and consumer groups answer that the traffic flows because the operators' own customers ask for it, and those customers already pay for access. Content companies already pay for delivery through caches inside ISPs (section 4e), long-haul networks and ocean cables (Module 5). BEREC itself concluded that "there is no evidence of 'free-riding'," that rising traffic does "not directly lead to significant incremental costs" when compared with total network costs, and that the market for exchanging traffic is "largely competitively driven."

BEREC also restated its finding of 2012 that a model in which the sending network pays "might be of significant harm to the internet ecosystem." Korea, the critics say, shows the harm. A rule of 2016 made Korean ISPs pay one another for the traffic they send, where they had exchanged it free before. Twitch, the live-streaming service, closed its Korean service on February 27, 2024, saying its network fees there were "ten times more expensive than in most other countries" (Dan Clancy, its chief executive, December 2023).

What the evidence says: the European Commission left a mandatory fee out of the Digital Networks Act it proposed in January 2026 (Module 12). Henna Virkkunen, its executive vice-president for tech sovereignty since December 1, 2024, said the commercial relationships between operators and content companies "are the industry standard" for exchanging traffic. In place of a fee the proposal offers a voluntary conciliation procedure, run with national regulators and BEREC, which the CCIA warns could harden into binding fees.

The operators' cost estimate comes from a study they paid for, which BEREC presented as their claim and did not validate. BEREC added that fixed access networks are "to the largest extent not traffic-sensitive," so their cost depends on the homes they reach more than on what those homes download. Falling transit prices (section 4f) point to a competitive market for carrying traffic. Korea, the one country that makes networks pay to send, has produced lawsuits and an exit. Netflix's case there ran from April 2020 to September 18, 2023 and ended in a partnership with undisclosed terms, and no study this course found measures a gain in investment from the Korean rules. Courts, as in Düsseldorf, enforce paid deals case by case under ordinary contract law.

The course's reading, stated once: the evidence does not support a fee set by law. Carrying traffic between networks is cheap and getting cheaper, the costly access network is paid for by the households it serves, and the one country that charges senders has produced disputes rather than measured gains. Where a large sender uses an operator's links on terms it has accepted, as Meta's unit did in Germany, contract law already makes it pay. The open question is whether the conciliation procedure becomes a fee in practice as the proposal passes through the European Parliament and the Council.

Worked example — New York to London and Back, by Glass and by Satellite

Inputs: New York (40.71° north, 74.01° west) and London (51.51° north, 0.13° west), on an Earth of mean radius 6,371 km; the speeds of light in fiber and in a vacuum (section 4a), the second applying to radio through space; and a geostationary satellite 35,786 km above the equator (Module 9 explains the orbit). The benchmark is Hibernia Express, a cable opened in 2015, whose tested round trip between the two cities was under 58.95 ms (Equinix, November 3, 2015). Hibernia's builder said that was more than 5 ms faster than existing cables. The satellite's position over the Atlantic, at 30° west, is an illustrative assumption. Step 1: the shortest path over the Earth's surface between the two cities, the great circle, is 5,570 km. Step 2: in glass along that path, one way takes 5,570 ÷ 200,000 = 0.02785 seconds, or 27.85 ms, so a round trip takes 55.7 ms, which no fiber can beat. Step 3: Hibernia's 58.95 ms is 3.25 ms more. If all of it were distance, the cable's path would be 58.95 ÷ 2 × 200,000 ÷ 1,000 = 5,895 km one way, 325 km or 5.8% longer than the great circle, and in practice part of the 3.25 ms goes to equipment. Step 4: the satellite is 39,051 km from New York and 39,088 km from London, standing about 25° above the horizon in both. One way, up and down, is 78,139 km, which takes 78,139 ÷ 299,792 = 0.2606 seconds, or 260.6 ms. Step 5: a request and its reply make four such legs, 521.3 ms in all, 8.8 times Hibernia's round trip, before any processing on the ground. No geostationary hop can take less than about half a second, as Module 9's worked example shows. Result: the best Atlantic cable comes within 6% of the limit set by physics, while a geostationary hop adds about half a second to every exchange. That is too slow for natural conversation or fast games (Module 3), and it is one reason why cables carry nearly all traffic between continents (Module 5).

Module 4 Key Points

The Backbone

• Light in fiber travels about 200,000 km a second, so each 1,000 km of glass adds 5 ms of delay one way.

• One fiber pair can carry close to 100 Tbps, Nokia says, because each fiber carries many wavelengths at once.

• Transit, paid access to the whole internet, cost as little as $0.03 per Mbps a month in mid-2026 (TeleGeography).

• IX.br São Paulo set the exchange record, 35.9 Tbps, on a World Cup day, June 25, 2026.

• Caches inside ISPs carry most video: Netflix said in 2016 that Open Connect delivered 100% of its video.

• BEREC found “no evidence of ‘free-riding’” in 2022, and the EU’s 2026 proposal left out a mandatory fee.

Module 4 Quiz

Check Your Understanding

1. About how fast does light travel through an optical fiber?

A) About 200,000 km a second
B) At 299,792 km a second, the speed of light in a vacuum
C) About 100,000 km a second, a third of light's speed in a vacuum
D) About 250,000 km a second in modern glass

2. What is the difference between transit and peering?

A) Transit carries traffic within one country, and peering carries it between countries.
B) Peering is the paid service, and transit is the free exchange between equals.
C) Transit is a paid service that reaches the whole internet, while peering swaps only the two networks' own customers' traffic, often for free.
D) Transit runs over long-haul fiber, while peering runs over the lines to homes.

3. What drove the record traffic peaks at the São Paulo and Frankfurt exchanges in 2026?

A) Streamed football matches drove them.
B) Nightly backups between cloud data centers drove them.
C) Software updates pushed to millions of phones at once drove them.
D) The training of AI models in nearby data centers drove them.

4. How did content delivery networks change who pays whom?

A) They made Tier 1 transit essential
B) They added a fee for every gigabyte delivered to a home, which content companies now pay to the Tier 1 networks that carry their video
C) They moved video onto satellites
D) More content flows straight into ISPs, so less crosses paid transit

5. What is the course's reading on whether content companies should pay network operators?

A) A fee set by law
B) No legal fee; contracts already cover it
C) A fee for the largest senders only, set by BEREC
D) Every content company should pay operators in proportion to its share of their traffic, as Korea requires of its ISPs

Module 5 — Under the Sea

Submarine Cables and Their Risks

Recommended Reading / Reference

Book: Nicole Starosielski, *The Undersea Network* (Duke University Press, 2015), a media scholar's account of how the cable routes, the landing towns and the people who lay and guard the cables shaped the network, centered on the Pacific

Reference: TeleGeography's Submarine Cable Map (2026 edition, March 2, 2026) and its free online lists of cables and landing stations, and the International Cable Protection Committee, the industry body that counts cable faults

5a. What a Submarine Cable Is

A submarine cable is a line of optical fibers, which carry data as light (Module 4), laid on or under the seabed between landing points on two coasts. For most of an ocean crossing it is about as wide as a garden hose, while each fiber is roughly the diameter of a human hair (TeleGeography). Most of the bulk is protection and power: steel wires for strength, a copper layer that carries electricity, and plastic insulation around a small number of fiber pairs, one fiber for each direction. TeleGeography describes 16-24 fiber pairs as typical for new cables (as compiled by VoxBooster, 2026).

Light fades over distance, so at intervals of tens of km the cable passes through a repeater, a pressure-proof housing for the amplifiers that boost the signal. The repeaters run on direct current fed down the copper layer from power equipment at the landing stations, so a cable needs electricity at both ends as well as light. A branching unit splits a cable toward a third landing point, so that one system can serve several countries.

Laying a cable begins with a survey of the seabed along the chosen route, which avoids steep slopes, busy fishing grounds and other cables where it can. A cable ship carries the cable coiled in tanks and pays it out as it sails. In deep water the cable lies directly on the ocean floor. Near shore, where anchors and trawl nets reach the bottom, it carries extra layers of armor and is buried under the seabed, often by a plow towed behind the ship. For Hibernia Express, the Atlantic cable of 2015 in Module 4's worked example, the installers buried heavily armored cable for nearly 2,000 km (The Broadcast Bridge, 2015).

At each end the cable comes ashore to a landing station, a building that holds the power equipment and the terminals that send and receive the light, and that connects the cable to networks on land. Cables are engineered for a design life of at least 25 years, but owners often retire them sooner, because newer cables carry far more data at a lower cost per bit (TeleGeography, 2026). Ships can recover a retired cable and re-lay it along a new path.

5b. The Network in 2026

TeleGeography's 2026 Submarine Cable Map, published on March 2, 2026, shows 694 cable systems, counting both those in service and those planned, and 1,893 landing stations. Its estimate of the cable in service in early 2026 is more than 1.5 million km, enough to circle the Earth at the equator more than 37 times.

Cables carry more than 99% of international data traffic, the ITU said on November 29, 2024. A Reuters explainer of April 28, 2026 gave the figure as a share of “the world’s internet traffic,” a wider measure than the ITU’s. TeleGeography says satellites carry a very small share, and that cables carry far more data at a fraction of satellites’ cost (Module 9 compares the two). Capacity has grown with each generation: Hibernia Express was designed in 2015 for 53 terabits per second (Tbps) (The Broadcast Bridge), and Amazon’s Fastnet (section 5c) for more than 320 Tbps.

The busiest routes cross the North Atlantic and the Pacific and link the countries of Asia, and a few narrow passages carry much of the traffic between regions. Many cables between Europe and Asia pass through Egypt and the Red Sea, which is why the cuts there in 2024 and 2025 (section 5e) slowed traffic in several countries at once. The Strait of Hormuz matters less. The International Cable Protection Committee (ICPC), the industry body that counts faults, said in May 2026 that traffic through it is under 1% of global international bandwidth, and that backup systems would limit the damage from a cut there (as reported by Europe in Motion).

5c. Who Builds Them Now

Telephone companies built and owned the cables for most of their history, usually in consortia that shared the cost and the capacity among many carriers. The largest content and cloud companies have since joined the consortia and then led the building. By the end of 2024, content providers used about 75% of the world’s used international bandwidth, and at least 80% on the routes across the Atlantic and the Pacific and within Asia (TeleGeography figures, as compiled by VoxBooster). TeleGeography names Google, Meta, Microsoft and Amazon as the major investors in new cable.

They build for their own traffic. A cloud company moves data between its own data centers on different continents, and it wants capacity it controls, on routes it chooses, without waiting for a consortium to agree. By the count of Mappr, a mapping publisher, Google held stakes in 34 cable systems in July 2026. Microsoft has invested with partners, as in MAREA, which it built with Meta and Telxius, the infrastructure arm of Spain’s Telefónica, between Virginia and Bilbao.

Meta confirmed Project Waterworth on February 14, 2025: a cable of 50,000 km touching five continents, with landings in the United States, Brazil, India and South Africa, and 24 fiber pairs. Much of it will lie in deep water, at depths of up to 7,000 meters, beyond the reach of anchors and trawl nets. Meta called it a “multi-billion dollar, multi-year investment” and tied it to rolling out AI services. TeleGeography said it would be the first cable wholly owned by Meta, which was then part-owner of 16 others (TechCrunch).

Amazon announced Fastnet on November 4, 2025. It is a transatlantic cable of Amazon Web Services (AWS), Amazon’s cloud business, from Maryland to County Cork in Ireland, landing away from the usual cable corridors, and AWS says it will enter service in 2028. It has extra layers of steel armor near shore and an optical switching branching unit that can send traffic to landings added later (Amazon, updated September 17, 2026).

TeleGeography expects the building to continue. Alan Mauldin, its research director, forecast on July 7, 2026 that spending on new cables would average $5 billion a year from 2026 to 2035. New cables would add more than 1,010,000 km in that decade while more than 450,000 km of older cable may be retired, and the routes within Asia alone would take $11 billion.

5d. Faults and Repairs

Cables break about 150-200 times a year, a count that held steady from 2014 to 2025 even as the total length of cable grew, according to the ICPC (in the same Reuters report). The ICPC and other experts attribute 70-80% of faults to accidental human activity, mainly fishing gear and ship anchors. TeleGeography adds that earthquakes and other natural events cause some, that failures of the equipment under water are less common, and that deliberate sabotage is very rare. Fish bites caused no faults at all from 2007 to 2014.

A single fault rarely cuts a country off, because most traffic has more than one route. When one cable fails, its owners move the traffic onto others while the damaged one is repaired. Places served by only one or two cables are the exception, and for them a single break can cut most international traffic until a ship arrives.

A repair starts on shore, where the owners test the line to estimate where the fault lies. A repair ship then sails to the site. In water shallower than 1,500 meters it uses a remotely operated vehicle, an underwater robot, and the work typically takes 1-3 days once the ship is there. In deep water the crew drags a grapnel, a hooked tool, to lift the cable to the surface, which takes a week or more, and then splices in a new section and lowers the cable back.

The waiting takes longer than the work. Rear Admiral Mike Brookes, commander of the US Office of Naval Intelligence, told the US-China Economic and Security Review Commission in 2026 that a repair took 51 days on average in 2024. The figure counts the wait for permits to work in a country’s waters, for a ship to become free and for good weather. Owners arrange repairs in three ways: consortium zones, in which the owners in a region share ships kept on standby; private zones, in which an owner pays for ships for its own systems; and deals made as needed.

About 65 cable ships work on fiber-optic cables worldwide, and about 20 of them mainly on repairs, by the same testimony, which counted at least six new cable ships on order. The fleet is old. TeleGeography and Infra-Analytics, a consultancy, estimated on June 30, 2025 that about two-thirds of cable maintenance ships will reach the end of their service life by 2040, while the length of cable in the oceans grows a net 48%. They put the cost of keeping repair service at today’s level at about $3 billion, for 15 replacement ships and five more.

5e. The Incidents of 2024-2026

In the Red Sea, the Rubymar, a cargo ship struck by missiles fired by Yemen’s Houthi movement in mid-February 2024, was abandoned with its anchor down and drifted for almost two weeks before it sank. Three cables in the area, Seacom, AAE-1 (Asia-Africa-Europe 1) and EIG (Europe India Gateway), were cut on February 24, 2024. John Kirby, spokesman for the US National Security Council, said in March 2024 that they were cut “most likely by an anchor dragging from the Rubymar,” and Ryan Wopschall, general manager of the ICPC, said it was “generally accepted” that the ship had dropped its anchor when it was hit.

HGC Global Communications, a Hong Kong carrier that uses the three cables, estimated in March 2024 that they carried about 25% of the region’s traffic. The repairs waited on permits from the Houthi authorities in Sana’a, and the cable ship Niwa, owned by E-Marine of the United Arab Emirates, completed them in July 2024 (DatacenterDynamics, July 25, 2024).

On September 6, 2025, two more cables, SMW4 (South East Asia-Middle East-Western Europe 4) and IMEWE (India-Middle East-Western Europe), failed near Jeddah in Saudi Arabia, as NetBlocks, a monitor of internet connectivity, recorded. Service in Saudi Arabia, Pakistan, the United Arab Emirates and India

slowed, and Microsoft reported higher delays on its Azure cloud. The cause has not been confirmed. Omantel, Oman's main carrier, said it moved traffic to other routes with no effect on its customers.

In the Baltic Sea the question is intent. On December 25, 2024, the Eagle S, an oil tanker flagged in the Cook Islands that had just left a Russian port, dragged its anchor across the Gulf of Finland, damaging the Estlink 2 power cable between Finland and Estonia and four other cables (The Record). Finnish prosecutors charged three of its senior officers in August 2025. In October 2025 the Helsinki District Court dismissed the case, finding that Finland lacked jurisdiction.

On August 27, 2026, the Helsinki Court of Appeal ruled unanimously that Finland does have jurisdiction, because the damage and its effects on Finland's power and telecom supply occurred there, and sent the case back for trial. It held that the events were not a "maritime accident" under the UN Convention on the Law of the Sea (UNCLOS). The first anchor drop might have been accidental, the court said, but the crew then told Finnish authorities that both anchors were raised while the ship dragged one for more than three hours.

The deadline to seek leave to appeal to Finland's Supreme Court was October 26, 2026, and whether anyone had appealed was not known on October 25, 2026, the day before. The three officers have left Finland. Fingrid and Elering, the Finnish and Estonian grid operators that own Estlink 2, claim about €105 million ($117.9 million) from them (The Record), at the Federal Reserve's rate for October 16, 2026, the rate used for euros throughout this module.

Finland stopped a second ship, the Fitburg, on December 31, 2025. Prosecutors say the cargo ship, flagged in St Vincent and the Grenadines and sailing from Russia to Israel, dragged a damaged anchor at least 130 km across the seabed, cut two telecom cables between Finland and Estonia and tried to damage eight more connections. Its Russian captain and Azerbaijani bosun, who deny the charges, went on trial in the Helsinki District Court on September 8, 2026, with prosecutors seeking at least two and a half years in prison (SubTel Forum). Most of the dragging, prosecutors say, happened in Estonia's exclusive economic zone (EEZ), the waters out to 200 nautical miles where a coastal state controls resources, and the defense argues that this puts the case outside Finland's authority. No verdict had been reported as of October 2026.

Who claims what differs. The Atlas Institute for International Affairs, a research group, counted at least 11 cables cut or damaged by Russia's "shadow fleet" of aging, opaquely owned ships since October 2023 (April 12, 2026). But officials from several countries around the North and Baltic seas say their governments increasingly believe the breaks were accidental and not directed by the Kremlin (The Record, August 2026). No court has ruled on intent in either Finnish case.

Taiwan has the clearest verdict. On February 25, 2025, the Hong Tai 58, a Togo-registered freighter with a Chinese crew, cut the Taiwan-Penghu No. 3 cable of Chunghwa Telecom in a no-anchoring zone off Tainan, after lingering in the area for three days. On June 12, 2025, the Tainan District Court sentenced its captain, a Chinese national surnamed Wang, to three years in prison for damaging the cable intentionally, finding that he had anchored on purpose and let the ship drift. He had said it was negligence (Central News Agency). Taiwan recorded 36 cases of cable damage in its waters from 2019 to 2023 (DatacenterDynamics, June 13, 2025).

The Iran war of 2026, whose effects on Iran's own networks Module 13 covers, put the Gulf's cables at risk without, so far, cutting one. In early March 2026 the cable ship Ile de Batz stopped work on a Europe-Asia cable through the Gulf, and at least two other cable projects were paused indefinitely (Bloomberg, May 13, 2026). Maintenance in the Gulf relies on E-Marine's five ships, and only one of them had been inside the Gulf since late February (InCyber, June 15, 2026).

In May 2026 the Islamic Revolutionary Guard Corps (IRGC) declared that foreign cable operators must obtain Iranian permits and pay “protection fees” to keep cables on the Iranian seabed, and repair operators issued force majeure notices for the Gulf, invoking contract clauses that excuse work made impossible by events beyond their control (Windward, a maritime-intelligence firm, May 18, 2026). Only two systems, FALCON and Gulf Bridge International, pass through Iranian waters. No cut of a Gulf cable had been reported in the sources this course read, as of late October 2026.

5f. Laws, Navies and Money

The law of the sea protects cables on paper. UNCLOS gives every state the right to lay cables on the high seas and the continental shelf, and its Article 113 requires each state to make it an offense for ships flying its flag, or people under its jurisdiction, to break a cable willfully or through culpable negligence. That leaves enforcement beyond a coastal state’s territorial waters mostly to the flag state, often a small registry far away, which is why the Finnish cases turned on jurisdiction. Governments added new measures in 2025-26.

Who	Measure	When	What it does
European Union	Action Plan on Cable Security	February 2025	Four strands: prevention, detection, response and recovery, and deterrence; proposes a reserve fleet of cable vessels for the EU
European Union	Funding from the Connecting Europe Facility, and a Cable Security Toolbox	Announced February 2026	€347 million ($389.7 million). In 2026, €60 million for repair equipment and €20 million for monitoring tools; in 2026-27, €267 million for priority projects. A separate €20 million call is for faster repairs, with equipment kept at ports and shipyards and a first pilot in the Baltic. Also 13 Cable Projects of European Interest, built in stages over 15 years
North Atlantic Treaty Organization (NATO)	Baltic Sentry	From January 14, 2025	Allied patrols in the Baltic Sea to deter and detect damage to undersea infrastructure
United States	FCC rules for cable landing licenses	Adopted August 7, 2025; in force from November 26, 2025, with the rest from July 8, 2026	Applicants owned or controlled by a “foreign adversary,” such as China or Russia, and new landings in those countries, are presumed to be denied; equipment on the FCC’s Covered List of security threats is barred; licensees must certify security plans

European Commission, EU Blue Economy Observatory, February 11, 2026; WorldPorts and the Atlas Institute for the Action Plan and Baltic Sentry; Federal Register, October 27, 2025, and the FCC’s public notice on the rules’ effective date.

Most of the money for cables and their upkeep still comes from their owners, through construction budgets and maintenance agreements. Public money began to flow only in 2025-26, through the measures in the table, and the Atlas Institute has proposed that the EU fund a repair fleet of its own from 2028. Taiwan’s prosecution and Finland’s two cases test whether criminal law can deter damage at sea.

5g. Are Cables Protected Well Enough, and Who Should Pay? The Cases as Their Advocates Make Them

The case that protection falls short: security officials, NATO and analysts such as the Atlas Institute argue that the risks have grown faster than the defenses. Repair capacity is thin and aging (section 5d), so several deliberate cuts at once could leave a region waiting weeks for ships. Ships linked to Russia have dragged anchors across Baltic cables again and again, prosecutions have stalled over jurisdiction, and a

court in Taiwan has found a captain guilty of cutting a cable on purpose. War freezes repair: in the Gulf the IRGC claimed the right to license and charge cable owners, and repair crews stood down. Owners on their own, these advocates argue, will not pay for spare ships that sit idle in peacetime, or for naval patrols.

The case that protection is broadly adequate: the cable industry and many engineers answer that the network is built to fail safely. Faults have held steady while the length of cable grew, and most are accidents that better charts, deeper burial and rules for fishing reduce. Routes overlap, so traffic moves to other cables, as Omantel's did in 2025, and new cables add routes every year. They add that several European governments now think the Baltic breaks were accidents, and that no navy can watch every km of the world's cables. Owners already pay for repair through their maintenance agreements, and they can pay for more ships when the need grows.

What the evidence says: the record of accidents supports the second side. Faults per year have not risen as the cable grew, so the risk per km has fallen, and the Red Sea cuts of 2024 were traced to a drifting wreck, not a targeted attack. The record of repair supports the first: an average of 51 days in 2024, and two-thirds of the maintenance ships at the end of their lives by 2040, with a costly renewal needed just to keep service level. On deliberate damage the evidence is mixed and unsettled. Taiwan's court found intent in one case, Finland's courts have so far ruled only on jurisdiction, and in the Gulf a threat stopped work without a cut.

The course's reading, stated once: cables are well protected against accidents and poorly protected against the two risks that grew in 2024-26, slow repair and deliberate or reckless damage in contested waters. Redundancy handles single faults, but not several cuts at once in a region with few routes and fewer ships. Owners should keep paying for the cables and their repair agreements, because they earn the revenue. The gap is public. Spare repair ships for crises, patrols, and laws that reach ships beyond territorial waters protect a service whose loss falls on whole economies, so governments have the stronger case to pay for them.

Worked example — How Many Homes One New Cable Could Serve at the Evening Peak

Inputs: Fastnet's design capacity, more than 320 Tbps (section 5b), split into channels of 400 gigabits per second (Gbps), a common size for one wavelength (an assumption for the arithmetic). An average American home used 711 gigabytes (GB) of data a month (OpenVault, Module 1), of which downloads are 11.2 parts in 12.2 (OpenVault's ratio of downloads to uploads, Module 6). A month has 30.4 days, and the evening peak is taken as 2.5 times the round-the-clock average, an assumption, since no published busy-hour figure was found. Step 1: 320 Tbps is 320,000 Gbps, or 800 channels of 400 Gbps. Step 2: downloads are 711 × 11.2 ÷ 12.2 = 652.7 GB a month. In bits that is 652.7 × 8 = 5,222 gigabits, spread over 30.4 × 86,400 = 2,626,560 seconds, an average of 1.99 megabits per second (Mbps). Step 3: at the evening peak a home draws 2.5 × 1.99 = about 5.0 Mbps. Step 4: one 400 Gbps channel serves 400,000 ÷ 5 = 80,000 homes at the peak, and the whole cable 800 × 80,000 = 64 million homes, or about 51 million if links are kept below 80% of capacity for safety. Step 5: DE-CIX Frankfurt's record peak of 21.47 Tbps (Module 4) is, at 5 Mbps a home, the evening load of about 4.3 million homes, and Fastnet's design capacity is 14.9 times that peak. Result: in principle one cable could carry the evening downloads of about half of America's 123.75 million home broadband subscriptions (Module 6). In practice it will not, because a home's evening video mostly comes from caches close to it (Module 4). A transatlantic cable like Fastnet carries traffic between data centers, which is why a cloud company built it.

Module 5 Key Points

Under the Sea

• Cables carry more than 99% of international data traffic (ITU); TeleGeography’s 2026 map shows 694 systems, in service and planned.

• Content providers such as Google and Meta used about 75% of used international bandwidth at the end of 2024 (TeleGeography).

• Cables suffer 150-200 faults a year, 70-80% of them accidents with fishing gear and anchors (ICPC).

• The average cable repair took 51 days in 2024, counting the wait for permits, ships and weather.

• Finland’s Court of Appeal revived the Eagle S case on August 27, 2026; no court has ruled on intent in Finland.

• No Gulf cable cut was reported by October 25, 2026, though the IRGC demanded permits and “protection fees” in May.

Module 5 Quiz

Check Your Understanding

1. What causes most faults in submarine cables, by the ICPC's count?

A) Accidents cause most of them, chiefly fishing gear and anchors.
B) Sharks bite through most of them.
C) Deliberate sabotage by state-linked ships causes most of them, as courts have found.
D) Failures of repeaters and other equipment under the sea cause most of them.

2. How long did a cable repair take on average in 2024, counting the wait for permits and ships?

A) One to three days
B) About six months
C) Ten to twenty days
D) 51 days

3. Who used about 75% of the world's used international bandwidth at the end of 2024?

A) Telephone carriers
B) Content providers such as Google and Meta
C) National governments
D) Banks and financial traders, whose orders between New York, London and Tokyo fill most transatlantic and transpacific cables

4. Where did the Eagle S case stand in October 2026?

A) The officers were convicted of sabotage in 2025.
B) Finland closed the case for good.
C) An appeals court had ruled that Finland has jurisdiction and sent the case back for trial, and no court had ruled on intent.
D) Finland handed the officers to Russia, whose court acquitted them.

5. What is the course's reading on whether cables are protected well enough?

A) Not against slow repair and deliberate damage
B) Fully, thanks to redundancy
C) Yes, if owners pay more
D) No, and the owners alone should pay for patrols, spare ships and new laws, since the cables earn them their revenue

Module 6 — The Last Mile

Broadband at Home and What It Costs

Recommended Reading / Reference

Book: Susan Crawford, *Fiber: The Coming Tech Revolution—and Why America Might Miss It* (Yale University Press, 2018), a Harvard law professor's case that fiber to every home is basic infrastructure and that the United States has left its building to a few private companies

Reference: The FCC's National Broadband Map and its annual broadband report under Section 706 of the Telecommunications Act; OpenVault's quarterly *Broadband Insights*; and Ookla's Speedtest reports of measured speeds

6a. The Access Technologies

The last mile is the link between a home and its ISP's network. It is the costliest part of the internet to build, because it must reach every address, and in the United States five technologies provide it: copper telephone lines, cable television's coaxial lines, fiber, radio from cell towers, and satellites.

Digital subscriber line (DSL) service sends data over the copper pairs that once carried only phone calls, from a telephone exchange or a cabinet on the street. Its speed falls with the length of the copper, so homes far from the cabinet get little. Carriers are retiring it: AT&T aims to decommission the large majority of its copper network by the end of 2029, although the approvals it needs could delay that (AT&T's third-quarter report for 2026). The EU's plan for its own copper networks is in Module 12.

Cable broadband runs over hybrid fiber-coaxial networks: fiber to a node that serves a neighborhood, then coaxial cable shared by the homes on that node. The standard it uses, the Data Over Cable Service Interface Specification (DOCSIS), delivers gigabit downloads in its version 3.1, but uploads have long been the weak side, because most of a cable's frequencies are set aside for traffic toward the home. OpenVault, which measures traffic on the networks of the ISPs it works with, put the average upstream speed on cable networks at 47.8 megabits per second (Mbps) in the second quarter of 2026, up 33.7% in a year.

Fiber to the home runs glass all the way to the house. Most such networks are passive optical networks (PON), in which one fiber from the ISP's equipment is split by unpowered splitters among dozens of homes that share its capacity. The older standard, gigabit PON (GPON), shares 2.5 gigabits per second (Gbps) toward the homes and 1.25 Gbps back. The newer 10-gigabit symmetric PON (XGS-PON) shares 10 Gbps in each direction, under the ITU's standards G.984 and G.9807.1. Upgrading a fiber network means changing the equipment at its ends, not the glass.

Fixed wireless access (FWA) uses a mobile network's radio to reach a receiver in the home, usually by 5G or 4G from an existing cell site (Module 7 explains the radio). It needs no new line to the house, so a carrier can sell it wherever its towers have spare capacity, and the carriers limit sign-ups where they do not. Satellites reach homes that no line or tower does, and Module 9 covers them.

Inside the home the connection ends at a box, a modem or an optical terminal, usually combined with a Wi-Fi router, and from there most devices connect by radio. Wi-Fi's speed falls with distance and with every wall and floor, and all the devices on a network share the router's airtime, so a phone in a far room may measure a fraction of the plan's speed. Households answer with mesh systems, several linked routers spread through the house, and with cables for computers and televisions that stay in one place. The radio bands Wi-Fi uses are covered in Module 8.

6b. Fiber’s Spread

Fiber is replacing copper and competing with cable. RVA, a research firm, counted 84.6 million unique US homes marketed with fiber when it closed its annual count in September 2025, over 60% of households, in its survey for the Fiber Broadband Association (FBA), the industry’s trade group (FBA, April 15, 2026). Counting each provider’s passing separately, so that a home with two fiber providers counts twice, the total came to 98.3 million passings, and builders passed 11.8 million homes in the year, a record (FBA, December 16, 2025). Canada has gone further: RVA counted 12.4 million unique Canadian homes with access to fiber, nearly 75% of homes, of which 6.3 million were connected.

Most homes with fiber have one fiber provider: 16% have more than one fiber passing, 15% two and 1% three or four. Of the homes that can buy fiber, 46.5% do, counted by unique passings, a share called the take rate, or take-up (FBA, April 2026). The FBA expects fiber to overtake cable as the main US broadband platform as early as 2028, and it has set goals for 2030 of 90% of homes passed and 50% connected, both its own projections. Competition raises take-up: where a second fiber provider arrives, the share of homes buying fiber from either rises to about 61%, and RVA counted about 60 million homes still waiting for their first fiber passing (FBA, December 2025).

The telephone companies are the main builders. AT&T reached more than 39 million consumer and business locations with fiber by September 30, 2026, and says it will pass 40 million by the end of 2026 and more than 60 million by the end of 2030. It counted 13.31 million fiber customers at the end of September 2026 (as reported by Cord Cutters News). Verizon, after buying Frontier (section 6d), reached more than 30 million homes and businesses with fiber and planned at least 2.0 million new passings in 2026 (Verizon, January 30, 2026). T-Mobile builds through joint ventures with investment firms: Lumos with EQT, and Metronet with KKR (T-Mobile).

Homes are sending more data, which favors fiber’s equal speeds in both directions. OpenVault counted average upstream use of 58.2 gigabytes (GB) a month per subscriber in the second quarter of 2026, up 19.2% in a year, with fiber subscribers sending 94.1 GB. The ratio of downloads to uploads fell from 16.6 to 1 in the second quarter of 2019 to 11.2 to 1, and OpenVault’s report warned that the return path, “long treated as secondary,” must be freed from performance constraints. Cable companies are responding: Charter said in July 2026 that it is evolving its network to offer symmetrical and multi-gigabit speeds.

6c. Fixed Wireless’s Rise

Fixed wireless is the fastest-growing kind of American home broadband, and the three national mobile carriers sell it. T-Mobile, the largest provider, had 8.45 million 5G broadband customers at the end of 2025, after adding 1.879 million in the year (its fourth-quarter report). It stopped reporting subscriber counts in 2026. Its chief financial officer, Peter Osvaldik, said only that it added broadband customers, on fiber and fixed wireless together, “in the upper 400,000 range” in the second quarter (Broadband Breakfast, July 23, 2026). Verizon had 6.21 million fixed wireless connections on June 30, 2026, after adding 193,000 in the second quarter (Verizon, July 24, 2026). AT&T had 2.94 million on its Internet Air service at the end of September 2026 (as reported by Cord Cutters News).

Ookla, which runs the Speedtest app, reported that the three carriers’ fixed wireless connections passed 17 million in the first quarter of 2026. That was 13.74% of the 123.75 million US home broadband subscriptions in the Census Bureau’s American Community Survey, as summarized by the technology blog of the Institute of Electrical and Electronics Engineers (IEEE) Communications Society in August 2026. About 70% of fixed

wireless speed tests in the second quarter came from urban users, a sign that the service now competes in cities, though Ookla notes that test counts are not a direct measure of subscribers.

The service is improving. The FCC's broadband report of August 2026 found that fixed wireless download speeds rose 25.1% and upload speeds 36.9% between late 2024 and early 2026. It also found that in 23 states more than half of all locations could get fixed wireless at 100/20 Mbps, meaning at least 100 Mbps down and 20 Mbps up, the FCC's benchmark (as reported by CSI Magazine, August 17, 2026).

The carriers sell it on price and simplicity. A survey of 12 large ISPs by InternetProviders.ai, an affiliate-marketing site and so a weak source, listed T-Mobile Home Internet and Verizon 5G Home with no rise in price after a promotion and no equipment fee in early 2026 (March 16, 2026). It found equipment fees of $5-15 a month and step-ups of 30-50% when the cable companies' promotions end. The trade-off is capacity: the service shares each cell site's radio with phones, and its medians trail fiber's, above all for uploads (section 6e).

6d. The American Market After the Deals of 2026

Three deals closed in 2026 and reshaped the market. Verizon completed its purchase of Frontier Communications, a telephone company with a large fiber network, on January 20, 2026, paying about $9.8 billion in cash and taking on about $12.9 billion of debt (Verizon's second-quarter report for 2026). AT&T completed its $5.75 billion purchase of the consumer fiber business of Lumen Technologies on February 2, 2026, adding about 1 million fiber customers and more than 4 million fiber locations in 11 states. AT&T kept the customers and has agreed to sell a controlling stake in the network itself, which it placed in a unit named Forged Fiber 37.

The largest was Charter Communications' $34.5 billion purchase of Cox Communications, which closed on August 20, 2026. The FCC approved it first, on February 27, 2026 as the Atlanta Journal-Constitution reported (Light Reading dated the approval March 2026), and the California Public Utilities Commission approved it last, about a week before the closing. California's conditions included $275 million of network upgrades and a $20 plan for low-income households.

Broadband Breakfast counted more than 35 million broadband subscribers and more than 70 million passings for the combined company, and Light Reading about 38 million customers in 45 states, both at the closing. Cox Enterprises, the family company that owned Cox, holds about 26% of Charter, and its chief executive, Alex Taylor, became chairman of Charter's board. By about August 2027, a year after the closing, the company will take the name Cox Communications while selling under Charter's Spectrum brand. Smaller deals continued around it. WOW!, a cable company, was taken private by DigitalBridge and Crestview, two investment firms, for about $1.5 billion including debt, closing on December 31, 2025, and Verizon bought Starry, which sells fixed wireless to apartment buildings and had about 100,000 customers, in early 2026. The FCC said in August 2026 that it had approved $130 billion of deals since January 2025, at least $50 billion of them in cable, fiber and wireline (as reported by CSI Magazine).

The deals came as cable lost ground. Comcast ended September 2026 with 28.343 million residential broadband customers after losing 144,000 in the third quarter, and said its domestic broadband revenue fell 4.7% from a year earlier on "lower average rates" and fewer customers (Comcast, October 22, 2026). Charter, before Cox, had 29.4 million internet customers on June 30, 2026, after losing 172,000 in the second quarter, and its residential internet revenue fell 3.2% from a year earlier, to $5.8 billion (Charter, July 24, 2026). Chris Winfrey, its chief executive, said in the same report that it operates "in a competitive environment across all of our products."

AT&T, by contrast, added 769,000 internet customers in the third quarter of 2026, a record for the company, with 442,000 on fiber and 327,000 on fixed wireless (AT&T, October 21, 2026). The cable companies answer with price guarantees, cheaper promotions and mobile phone service sold with broadband. Comcast launched a five-year price guarantee on April 15, 2025, and its national plans of June 2025 include a free mobile line for a year (Comcast). It added 368,000 mobile lines in the third quarter of 2026, reaching 10.555 million.

6e. Speeds, Here and Abroad

Speeds are best compared as medians, the middle result of many tests, which a few very fast or very slow lines cannot skew. Ookla's Speedtest Global Index put the US median fixed download at 306.86 Mbps in May 2026, ninth in the world (as reported by SlashGear), and the global median at 129.68 Mbps in August 2026, with Singapore fastest at 462.50 Mbps (as reported by VietNamNet). These are tests that users choose to run, not a random sample, and Wi-Fi inside the home can hold them down. Quality varies by state: in the first half of 2026, 74.7% of Speedtest users in Connecticut reached at least 100/20 Mbps, against 43.5% in Alaska, the lowest (Ookla, as reported by TelecomLead, September 15, 2026). Ookla's tests also show how far the technologies differ, above all in upload speed.

Service	Technology	Median download	Median upload	Period
AT&T Fiber	Fiber	374.75 Mbps	320.65 Mbps	First half of 2026
T-Mobile Home Internet	Fixed wireless (5G)	222.72 Mbps	18.12 Mbps	Second quarter of 2026
AT&T Internet Air	Fixed wireless	160.43 Mbps	10.06 Mbps	Second quarter of 2026
Verizon 5G Home Internet	Fixed wireless (5G)	126.53 Mbps	12.20 Mbps	Second quarter of 2026

Ookla Speedtest Intelligence data, as reported by the IEEE Communications Society's technology blog, July 23 and August 27, 2026. Ookla named AT&T Fiber the fastest US fixed network in the first half of 2026. Verizon's figures exclude Starry, a fixed-wireless provider it bought. These are medians of tests that users ran, not the speeds of the plans they bought.

Fiber's median upload is 18 to 32 times the fixed wireless medians in the table, which matters for video calls, backups to the cloud and sending large files. The fixed wireless downloads, by contrast, are enough for several high-definition video streams at once (Module 1 gives the rates streaming uses).

Europe has more fiber coverage. The European Commission found that very-high-capacity networks (VHCN), the gigabit-capable lines it counts, reached 85.6% of EU households in mid-2025, and fiber to the premises (FTTP) 74.1% (State of the Digital Decade report, June 17, 2026). Across the 38 members of the Organisation for Economic Co-operation and Development (OECD), fiber made up 47% of fixed broadband subscriptions at the end of 2024, against 28% in 2019, and more than 80% in Korea, Iceland, Spain, Finland and Lithuania. Fixed wireless made up 6% and satellites 1% (OECD figures, as reported by RCR Wireless News, November 3, 2025). The Commission said that fiber still lags its goal of gigabit connections for all by 2030, and that it will review its targets in 2027.

China built the largest fiber network. Its Ministry of Industry and Information Technology (MIIT) counted more than 698 million fixed broadband users at the end of April 2026. Of these, 667 million (95.5%) had connections of 100 Mbps or more and 253 million (36.2%) had gigabit service, and China had more than 32 million ports of 10-gigabit PON (as reported by TelecomLead, June 29, 2026).

India connects mainly by phone. The Telecom Regulatory Authority of India (TRAI) counted 1,021.43 million mobile broadband subscriptions at the end of June 2026, a count that includes connections between machines, against 47.71 million wired broadband connections and 18.36 million fixed wireless ones, 12.94

million of them over 5G (TRAI, July 28, 2026). Fixed connections of both kinds made up 6.6% of India's 1,001.50 million broadband subscriptions, leaving out connections between machines.

6f. Is American Broadband Uncompetitive and Overpriced? The Cases as Their Advocates Make Them

The case that it is: consumer advocates point first to price. Broadband Genie, a British price-comparison site, ranked the United States 167th of 214 countries by median monthly price, at about $80, in a study of 2,631 fixed-line plans published on April 17, 2026. It found Northern America the second-most expensive region, averaging $98.40, after Polynesia. The advocates point to promotions that end, such as Comcast's step-up after one year (worked example) and the cable step-ups in section 6c.

They point, too, to choice. Only 43.4% of Americans could choose among three or more fixed terrestrial, or land-based, services at 100/20 Mbps in June 2025 (FCC, August 2026). Most homes with fiber have just one fiber provider (section 6b), and where cable is the only wired line, a step-up is hard to escape.

The case that it is not: USTelecom, the telephone companies' trade association, and the FCC under its chairman, Brendan Carr (Module 12), point to value and to change. USTelecom's 2026 Broadband Pricing Index (May 19, 2026) found the inflation-adjusted price of the most popular tiers, 100-940 Mbps, down 43.6% since 2014 and 6.0% in 2025. It found gigabit plans down 48.9% since 2016 and the real price per Mbps down more than 85% since 2014, while prices before inflation held essentially flat. USTelecom says one American household in three now subscribes to a gigabit plan, and American median speeds rank ninth in the world (section 6e). The number of Americans without access to land-based service at 100/20 Mbps fell by about 23% from June 2024 to June 2025, the FCC found.

Counting satellite services as well, 77% of Americans could choose among three or more services at 100/20 Mbps, and in August 2026 the FCC found broadband being deployed to all Americans "in a reasonable and timely fashion." Competition shows in the results, they add: cable is losing customers, AT&T is adding them at a record pace, and Comcast's revenue per customer is falling (section 6d).

What the evidence says: the headline comparison and the industry index measure different things. Broadband Genie compares the monthly prices of plans, not prices per Mbps or as a share of income, and American plans are faster than most. USTelecom's index adjusts for inflation and speed, but an industry group publishes it, and its page does not set out its data sources. The market evidence is stronger than either. Comcast and Charter lost customers while AT&T added a record number, and Comcast's broadband revenue fell on lower average rates, which is what competition does to prices.

On choice, the answer depends on what counts, since the FCC's two figures differ only in whether satellites are included. Ookla's medians show fixed wireless to be a substitute for fiber's downloads but not for its uploads (section 6e), and OECD figures show fixed broadband prices across its members falling 20% from 2020 to 2025, so falling prices are not peculiar to the United States (as reported by RCR Wireless News).

The course's reading, stated once: American broadband is expensive at the headline and more competitive than its critics allow. Cable's hold on the market has given way, in most places, to rivalry with fiber and fixed wireless, and the falling revenue per customer at the largest cable company is the clearest sign that prices are under pressure. The problem that remains is narrower. A home with one wired provider and a weak wireless signal pays the full price after the promotion, with no alternative, which is why the FCC's count of land-based choices matters more than its count with satellites.

Worked example — A Broadband Bill, from the Promotion to What a Customer Pays

Inputs: Comcast's national prices for its 500 Mbps plan, announced on June 26, 2025: $55 a month with a one-year price guarantee, $70 with a five-year guarantee, and an "everyday" price of $85, with unlimited data and the Wi-Fi gateway included. The release does not state taxes or installation fees. A gateway fee of $14 a month, Comcast's charge on plans that bill equipment separately, comes from the InternetProviders.ai survey, an affiliate site and so a weak source. Comcast's results give $6,241 million of domestic broadband revenue in the third quarter of 2026, and its residential customers numbered 28.486 million on June 30 (second-quarter release) and the September figure given in section 6d. Step 1: after the first year the price rises from $55 to $85, by $30 or 54.5%. Step 2: two years on the one-year plan cost 12 × $55 + 12 × $85 = $660 + $1,020 = $1,680, an average of $70 a month, the same as the five-year price. Step 3: over five years the one-year plan costs $660 + 4 × $1,020 = $4,740, against 60 × $70 = $4,200 with the five-year guarantee, $540 more. Step 4: on a plan that charges for the gateway, the everyday bill is $85 + $14 = $99, and the gateway costs $168 a year. Step 5: what Comcast actually collects. Customers averaged (28.486 + 28.343) ÷ 2 = 28.4145 million in the quarter, so revenue per customer per month was $6,241 million ÷ 3 ÷ 28.4145 million = $73.21. Result: the average customer pays 33.1% more than the $55 promotion and 13.9% less than the $85 everyday price, a blend of customers on promotions, guarantees, bundles and full prices, and on slower and faster plans. The $73.21 is an average, not a price anyone is quoted, and the advertised $55 is the price of a first year, not of the years after it.

Module 6 Key Points

The Last Mile

• Fiber reached 84.6 million unique US homes, over 60% of households, when RVA closed its 2025 count for the Fiber Broadband Association.

• Fixed wireless from T-Mobile, Verizon and AT&T passed 17 million US connections in early 2026 (Ookla).

• Charter closed its $34.5 billion purchase of Cox on August 20, 2026, after the Verizon-Frontier and AT&T-Lumen deals closed.

• Comcast lost 144,000 broadband customers in the third quarter of 2026, while AT&T added a record 769,000.

• Comcast collected about $73.21 a month per broadband customer that quarter, against its $55 one-year promotion.

• Fiber to the premises reached 74.1% of EU households in mid-2025 (European Commission).

Module 6 Quiz

Check Your Understanding

1. How many unique US homes were marketed with fiber when RVA closed its count in September 2025?

A) About 123.75 million
B) 84.6 million
C) More than 39 million homes
D) About 17 million, mostly in cities, since rural homes are served by fixed wireless and satellites instead

2. What happened to Comcast's broadband customers in the third quarter of 2026?

A) It gained a record 769,000 of them.
B) Its count was unchanged, as its price guarantees kept customers from leaving.
C) It added customers for the first time in years, helped by its five-year price guarantee.
D) It lost 144,000 of them, leaving 28.343 million.

3. In the worked example, about how much did Comcast collect per residential broadband customer per month in the third quarter of 2026?

A) $55, its promotional price
B) $85, its everyday price
C) About $73.21, between its $55 promotion and its $85 everyday price
D) $99, its everyday price plus the gateway fee

4. Which company had the most US fixed wireless customers at the end of 2025?

A) Verizon Communications
B) T-Mobile
C) AT&T Internet Air
D) SpaceX's Starlink, whose satellites reach more homes than any carrier's towers

5. What is the course's reading on American broadband prices and competition?

A) Pricey at the headline, but more competitive than critics say
B) Cheap by world standards
C) Overpriced, because cable still faces no rival in most of the country, so the FCC should cap prices
D) Fairly priced everywhere, since fixed wireless and satellites now reach every home

Module 7 — Mobile Networks

From 1G to 5G, and the Business of Carriers

Recommended Reading / Reference

Book: Jon Agar, *Constant Touch: A Global History of the Mobile Phone* (Icon Books, 2003; revised and updated edition, 2013), a historian of technology's short account of how the mobile phone and its networks spread around the world

Reference: Ericsson, *Ericsson Mobility Report*, published twice a year, whose June 2026 edition is the source of most subscription and traffic figures here, and the GSMA, the mobile operators' trade association, whose research arm publishes *The Mobile Economy 2026* (February 2026)

7a. How a Cellular Network Works

A mobile network links a phone to the internet in four stages. The phone exchanges radio signals with a base station: the antennas, radios and processing equipment mounted on a tower, a rooftop or a lamppost. A link called backhaul carries the traffic from that cell site to the operator's core network, the central computers that connect it to the internet and to other phones. Engineers call the radio side, from phones to base stations, the radio access network (RAN).

A base station's antennas divide the area around it into sectors, each serving a cell. Reusing the same frequencies from cell to cell lets a limited range of radio spectrum serve a whole country (Module 8 explains spectrum). Older networks kept a frequency a few cells away from its next use, while 4G and 5G cells mostly use the same channel side by side and coordinate to limit interference. Large towers cover wide rural cells, while small cells, low-power base stations on lampposts and walls, add capacity in crowded streets and stations. When a phone moves, the network passes its connection from one cell to the next, a step called handover. Each phone identifies its account with a subscriber identity module (SIM), a small chip, or its software equivalent, that holds the subscriber's identity.

Backhaul runs over fiber where it is available and over point-to-point microwave radio where it is not. The GSMA claims that more than half of 5G backhaul now runs over radio links in the E-band, at 71-76 and 81-86 gigahertz (GHz), or billions of cycles a second (GSMA presentation, February 12, 2026). The core network checks each subscriber's identity, tracks which cell each phone is in, sets up its data sessions, records use for billing and hands traffic to the internet. In 5G the core is designed as software for general-purpose servers rather than as dedicated hardware.

Everyone in a cell shares its capacity. That capacity is the product of three things. The first is the width of the spectrum the operator uses there, measured in megahertz (MHz), millions of cycles a second. The second is spectral efficiency, the number of bits a second that each hertz carries, which depends on the radio technology and on the quality of the signal. The third is the number of cells. An operator that needs more capacity can buy spectrum, upgrade its radios or build more and smaller cells, and the worked example at the end of this module counts how many people one site can serve.

7b. Five Generations, and What Each Changed

A generation of mobile technology is a set of standards that equipment makers and operators build to, so that phones work on any network that follows them. The 3rd Generation Partnership Project (3GPP), a joint body of standards organizations from Asia, Europe and North America founded to write 3G, now writes the

standards the whole world uses (Module 12 explains how standards are made). Each generation must meet minimum requirements that the ITU calls International Mobile Telecommunications (IMT). A new generation has arrived about once a decade (table).

Generation	First commercial service	Main standard	The ITU's peak-speed target
1G	NTT (Nippon Telegraph and Telephone), Japan, 1979, first in Tokyo; Ameritech, United States, 1983	Several national analog systems	None
2G	Radiolinja, Finland, 1991	Global System for Mobile Communications (GSM), among others	None
3G	NTT Docomo, Japan, October 1, 2001	Two rival standards, one of them written by 3GPP	At least 144 kilobits per second (kbps)
4G	TeliaSonera, Stockholm and Oslo, December 14, 2009	Long-Term Evolution (LTE)	100 megabits per second (Mbps) on the move; 1 gigabit per second (Gbps) standing still
5G	South Korea's three operators and Verizon, April 3, 2019	5G New Radio, first specified in 3GPP's Release 15 (December 2017)	20 Gbps down; 10 Gbps up

Launch dates as given in Wikipedia's articles on each generation (read October 25, 2026). Targets are the ITU's minimum peak rates: IMT-2000 for 3G, IMT-Advanced (2008) for 4G and IMT-2020 for 5G. The requirements for 6G are called IMT-2030 (section 7f).

The first generation sent voice as an analog radio signal, much as a broadcast station does, so calls could not be encrypted and each call took a frequency channel to itself. The second generation made the signal digital. Digital speech could be compressed and encrypted, several calls could share one channel, and the same network could carry text messages and, with later upgrades, slow data. The European GSM standard spread from Finland to most of the world.

3G was designed for data, so that phones could browse the web and make video calls. 4G made the network run entirely on the Internet Protocol (IP). Voice became one more stream of data, carried as voice over LTE, and the radio split each wide channel into many narrow subcarriers, which made it efficient at high speeds. 4G was the network of the smartphone boom, whose history is in Module 2.

5G's new radio allows much wider channels than 4G's (Module 8 gives the limits by band), and its antennas have many elements that aim beams at individual phones, so one site carries far more traffic. 5G also opened the millimeter-wave bands above 24 GHz to mobile service, though their signals travel only short distances. Its design aims at lower delay as well as higher speed, which matters for video calls, games and machines that respond in real time.

Old generations are switched off so that their spectrum can be reused. In the United States, AT&T shut its 3G network on February 22, 2022, T-Mobile closed Sprint's on March 31 and its own on July 1, and Verizon followed on December 31. Parts of Europe are keeping 2G longer: Finland's operators must run GSM until the end of 2029, and Britain has set 2033 for switching its 2G off. In the first quarter of 2026, by the count of Ericsson, the Swedish equipment maker, 5G subscriptions rose by 162 million while 4G fell by 77 million, 3G by 18 million and 2G by 23 million.

7c. Standalone 5G and What It Adds

Most operators launched 5G in its non-standalone form: new 5G radios attached to their existing 4G core network and controlled through 4G signals. That raised speeds quickly and cheaply, but the network behind the radios was still 4G. Standalone 5G pairs the new radios with a new 5G core. T-Mobile US launched the first nationwide standalone network, in 2020.

The standalone core is what makes 5G's other features possible. It supports network slicing, which divides one physical network into virtual networks with guaranteed speed or delay for particular customers, such as a broadcaster sending live video from a stadium. It also lowers delay and improves the uplink, the traffic sent from phones to the network. About 390 operators had launched commercial 5G by June 2026, and more than 90 of them had launched standalone 5G (Ericsson).

Counts differ by definition. Dell'Oro Group, a market-research firm, counted 72 standalone 5G core networks in service around November 2025, a count of networks built rather than of operators selling service. Paid features are still few: Ericsson counted 84 commercial offers built on slicing or similar guarantees in June 2026, up from 65 in November 2025. The next step, 5G-Advanced (3GPP's Release 18), adds AI tools for running networks and more precise location, and 12 operators had launched it commercially by late 2025 (Dell'Oro).

7d. The State of Mobile in 2026

There were about 8.84 billion mobile subscriptions at the end of 2025, by Ericsson's count, more than one for every person on Earth, because a subscription is an active SIM and many people have more than one. Counting people instead, GSMA Intelligence, the GSMA's research arm, put unique mobile subscribers at 5.8 billion in 2025, about 70% of the world's population. Ericsson counted 7.47 billion smartphone subscriptions and forecasts 9.54 billion subscriptions of all kinds in 2031.

5G subscriptions reached 2.94 billion at the end of 2025 and 3.1 billion at the end of March 2026, while 4G subscriptions were just under 4.6 billion (*Ericsson Mobility Report*, June 16, 2026). Ericsson forecasts 6.4 billion 5G subscriptions at the end of 2031, two-thirds of the total. Take-up is uneven. At the end of 2025, 5G was 79% of North American subscriptions, about 370 million, and 60% in North East Asia, but 35% in the region of India, Nepal and Bhutan, where that share still came to 430 million.

Estimates of coverage differ by source. Ericsson calculates that 5G networks covered 60% of the world's people at the end of 2025, and 50% outside mainland China. The ITU's *Facts and Figures 2025* (November 17, 2025) puts 5G coverage at 55% of the world's population in 2025. The ITU builds its figure from data that operators and regulators report country by country, while Ericsson uses its own estimates, so the five-point gap is a difference of method, not a change over time.

Mobile networks carried 146 exabytes (EB) of data a month at the end of 2025, an exabyte being a billion gigabytes (GB). Including fixed wireless access (FWA), home broadband delivered over the mobile network, they carried 203 EB (Module 6 gives America's FWA subscriber counts). Traffic in the first quarter of 2026 was 22% higher than a year earlier. The average smartphone used 22 GB a month worldwide in 2025, from 5 GB in Sub-Saharan Africa to 25 GB in North America and 37 GB in India, Nepal and Bhutan, and Ericsson forecasts a world average of 42 GB in 2031. The world's median mobile download speed was 109.05 megabits per second in August 2026, by Ookla's Speedtest Global Index, from tests that users choose to run (as reported by VietNamNet, October 14, 2026).

Uplink traffic grew faster than downlink at 43 of the 55 operators Ericsson measured. Machines are counted separately: Ericsson put cellular connections for the internet of things (IoT), such as meters, trackers and cars, at about 4.5 billion at the end of 2025 and forecasts 7.8 billion in 2031. China is the largest 5G market. Its Ministry of Industry and Information Technology (MIIT) counted 4.838 million 5G base stations at the end of 2025, 34.4 for every 10,000 people, and 1.204 billion 5G subscriptions out of 1.827 billion mobile subscriptions (Xinhua, China's state news agency, January 29, 2026). By the GSMA's count, China held more than 40% of the world's 5G connections (June 2026).

7e. The Carriers' Business

Mobile operators, also called carriers, earn most of their money selling access, and their revenue has not grown with the traffic they carry. China's telecom industry earned 1.75 trillion yuan ($261.7 billion) in 2025, up 0.7% (Xinhua), in a year when it added about 583,000 5G base stations (this course's arithmetic from the MIIT's counts). The GSMA's 2026 report describes "slowing connectivity revenues" and operators looking to AI and to services for businesses for growth. Dollar figures converted from other currencies in this module use the Federal Reserve's rates for October 16, 2026.

Spending has followed. Capital spending by telecom operators worldwide rose 5% in the first half of 2026, by Dell'Oro's count, but the firm expects it to grow 0-1% a year from 2025 to 2030. Capital intensity, capital spending as a share of revenue, peaked at 18% in 2022, during the first 5G builds. Dell'Oro expects it to fall to about 14% by 2028 and then rise modestly as early 6G investment begins (October 6, 2026).

The largest carriers count their customers in different ways, so their figures are not strictly comparable (table). China Mobile, controlled by the Chinese state, is the largest in the table by customers and by revenue. Reliance Jio leads India, where its average revenue per user (ARPU), revenue divided by customers, was 215.6 rupees ($2.27) a month in April-June 2026 (company report, July 17, 2026). America Movil, based in Mexico, and Vodafone, based in Britain, run networks in many countries, while most carriers serve one.

Carrier	Main markets	Mobile customers (date)	A revenue figure (period)
China Mobile	China	About 1.011 billion, 687 million of them on 5G (June 30, 2026)	$152.0 billion (2025)
Reliance Jio	India	503.6 million wireless subscribers (June 2026, the regulator's count)	Revenue per user only (see text)
America Movil	Mexico and Latin America	332.4 million wireless lines (September 30, 2026)	250.6 billion pesos ($13.5 billion), third quarter of 2026
Verizon	United States	147.0 million wireless retail connections (June 30, 2026)	$138.2 billion (2025)
AT&T	United States	109.8 million retail wireless subscribers (June 30, 2026)	$125.6 billion (2025)
T-Mobile US	United States	34.7 million postpaid accounts (June 30, 2026); it reports accounts, not subscribers	$19.0 billion of service revenue, up 9%, second quarter of 2026
Vodafone	Europe, Turkey and Africa	No group-wide total in its annual results	€40.5 billion ($45.5 billion), up 8.0%, year to March 31, 2026

Customer counts: company figures compiled by TelecomLead, a telecom trade publication (September 29, 2026), except America Movil's (its report of October 20, 2026) and Jio's, from the Telecom Regulatory Authority of India (TRAI). Revenue for 2025: TelecomLead's own dollar figures (April 21, 2026). America Movil, T-Mobile and Vodafone: their reports of October 20, July 23 and May 12, 2026. Verizon counts connections, AT&T subscribers, T-Mobile accounts and America Movil lines.

Towers are often owned by separate companies that rent space on each one to several carriers, so the number of tenants per site, the tenancy ratio, drives a tower company's profit. The largest, China Tower, had 2.172 million sites and 3.871 million tenants on June 30, 2026, a ratio of 1.78. American Tower had 149,686 sites at the end of 2025, Cellnex, Europe's largest, 114,542 in the first quarter of 2026, and India's Indus Towers 267,611 towers on June 30, 2026. Crown Castle owns about 40,000 American towers. Their revenue is mostly rent: China Tower took in 48.69 billion yuan ($7.28 billion) in the first half of 2026, 35.26 billion of it from towers, and Cellnex €1.994 billion ($2.24 billion), with 1.63 tenants per site. American Tower expects to spend $1.8-1.91 billion and to build 1,700-2,300 new sites in 2026 (company figures compiled by TelecomLead, September 14 and October 21, 2026).

7f. Equipment Makers, Open RAN, and 6G

Five companies sell almost all of the world's mobile radio equipment. Huawei, Ericsson, Nokia, ZTE (Zhongxing Telecommunication Equipment) and Samsung, ranked in that order for 2025, held 96% of radio access network sales in the first three quarters of 2025, against 95% in 2024. Huawei and Ericsson alone held nearly two-thirds (Dell'Oro, as reported by Fierce Network). Huawei and ZTE are Chinese, Ericsson is Swedish, Nokia Finnish and Samsung South Korean, and Module 13 covers the restrictions on the two Chinese makers in America and Europe.

The market has shrunk. Sales of radio equipment lost about a fifth of their value between 2022 and 2024, as operators finished their first 5G builds, and Dell'Oro expected them to be flat in 2026 and to grow about 1% a year from 2025 to 2030 (February 17, 2026). In its ranking for 2025, Huawei and Nokia gained share, Ericsson and Samsung held steady and ZTE lost ground, and Dell'Oro judged the market's concentration to be at a 10-year high.

Open radio access networks (Open RAN) are the industry's attempt to widen the field. The approach splits a base station into a radio unit, a distributed unit and a centralized unit, with published interfaces between them. An operator can then buy each part from a different supplier and run some of them as software on ordinary servers. Open RAN's growth in 2025 was "at the low end" of a target of 5-10%, said Stefan Pongratz, a vice president of Dell'Oro, on December 31, 2025. Private networks, mobile networks that factories, ports and campuses run for themselves, are another hoped-for market, but Dell'Oro reported on October 20, 2026 that sales of radio equipment for them grew only at mid-single-digit rates in the first half of 2026.

6G is on 3GPP's calendar. On June 10, 2026, at a meeting in Singapore, 3GPP agreed the timeline for Release 21, its first set of 6G standards. The service requirements are to be frozen, or fixed so that manufacturers can build to them, in March 2027, the system design in June 2028, the detailed functional specifications in December 2028 and the final message formats in March 2029. Ericsson and the GSMA both expect the first commercial 6G services around 2030.

The forecasts for 6G start small. Ericsson expects 180 million 6G subscriptions at the end of 2031, and the GSMA expects more than 5 billion 6G connections by 2040, about half of all connections. What 6G will add is still being studied in 3GPP's Release 20. Its spectrum, including the bands to be discussed at the World Radiocommunication Conference of 2027, is covered in Module 8.

7g. Has 5G Paid Off, and Is 6G Needed? The Cases as Their Advocates Make Them

The case that 5G has not paid off: analysts and investors who follow the operators argue that 5G added costs without adding revenue. Operators bought spectrum and radios and then cut back, and sales of equipment shrank (section 7f). China built the largest 5G network while its telecom revenue barely moved (section 7e). Most operators have not launched the standalone core that paid features need, and the features on sale are few (section 7c). Low-orbit satellites now compete for rural customers (Module 9). On this view, 6G risks repeating the cycle: another round of spending for speeds that most users do not need and will not pay more for.

The case that it has: equipment makers and many operators answer that 5G was built to carry growth, and it does. 5G carried 48% of mobile data at the end of 2025, by Ericsson's count, though it held a third of subscriptions (this course's arithmetic), and Ericsson forecasts 85% by 2031. Fixed wireless gave operators a new business: Ericsson counted 185 million FWA connections at the end of 2025, and 71% of the operators that offer FWA offer it over 5G. Operators also keep paying for mid-band capacity. AT&T says that the spectrum it bought from EchoStar (Module 8), used on 23,000 sites under a special FCC authorization, made 5G downloads up to 80% faster. On 6G, Ericsson's scenario has uplink traffic tripling or more from 2025 to 2031 "if additional AI traffic materializes," and the GSMA argues that the 2030s will need new spectrum and a new standard (Module 8).

What the evidence says: the traffic and the money point in different directions. If revenue grows about 1% a year, an assumption close to China's 0.7%, while traffic grows about 22%, the revenue each gigabyte brings in falls by about 17% a year (this course's arithmetic). Operators must therefore cut their cost per gigabyte at least that fast to stand still, and 5G's added capacity is how they do it. The evidence does not support the promise of large new revenue: paid features built on slicing number in the dozens worldwide, and capital spending is falling as a share of revenue (sections 7c and 7e). For 6G, the evidence so far is a timetable rather than a business case. The standards will be ready by 2029, the makers' own forecasts for its first years are small, and Dell'Oro expects only a modest rise in spending (sections 7e and 7f).

The course's reading, stated once: 5G has paid off as capacity, not as a new business. It carries the growth in traffic at a cost per gigabyte that operators could not otherwise reach, and it created fixed wireless, but the new revenue its promoters promised has not come, and operators now spend a smaller share of their revenue on their networks than in 2022. 6G should be judged the same way, as a capacity upgrade to be built when traffic requires it rather than the start of a new spending cycle. The operators' own capital plans already treat it that way.

Worked example — How Many People Can One 5G Site Serve?

Take one mid-band 5G site. It uses a 100 MHz channel, the widest that 5G allows in mid band (Module 8), and three antenna sectors, a common design taken here as an assumption. Its average spectral efficiency is an assumed 5 bits a second per hertz, with 3 and 8 as low and high cases, since real values depend on signal quality, antennas and load. Demand is Ericsson's world average of 22 GB per smartphone a month in 2025, and the busiest hour is assumed to carry 10% of a day's traffic. Step 1, capacity: one sector carries 100 million hertz × 5 bits = 500 megabits per second (Mbps), and three sectors carry 1.5 gigabits per second (Gbps); the low and high cases give 0.9 and 2.4 Gbps. Step 2, demand per smartphone: 22 GB a month is 0.733 GB a day, and 10% of that is 0.0733 GB, or 586.7 megabits, in the busy hour. Spread over its 3,600 seconds, that is 163 kilobits per second (kbps). Step 3, typical users: 1.5 Gbps ÷ 163 kbps = about 9,200 smartphones at the busy-hour average, or 5,500 to 14,700 across the cases. Step 4, users at a given speed at the same moment: 1,500 Mbps ÷ 25 Mbps = 60 users streaming at 25 Mbps, or 15 at 100 Mbps, and the cases give 36 to 96 at 25 Mbps. Step 5, fixed wireless: a home using 711 GB a month, the American average in the second quarter of 2026 (Module 1), uses 711 ÷ 22 = 32.3 times a smartphone's data. At the same 10% busy-hour share it needs 5.27 Mbps in that hour, close to Module 5's estimate of 5.0 Mbps of downloads alone at the evening peak, so the middle-case site could serve about 285 such homes and nothing else. The result: one site serves thousands of typical phones but only dozens of people at video speeds at once, and each home on fixed wireless takes the room of about 32 phones. That arithmetic is why operators keep adding spectrum and cells, and why fixed wireless competes with phones for the same capacity.

Module 7 Key Points

Mobile Networks

• 5G subscriptions reached 3.1 billion at the end of March 2026, out of about 8.84 billion mobile subscriptions (Ericsson).

• 5G covered 60% of the world’s people at the end of 2025 by Ericsson’s estimate, and 55% by the ITU’s.

• Five makers held 96% of radio equipment sales in the first three quarters of 2025; Huawei and Ericsson alone held nearly two-thirds (Dell’Oro).

• Operators’ capital spending peaked at 18% of revenue in 2022; Dell’Oro expects about 14% by 2028.

• 3GPP plans to freeze 6G’s detailed specifications in December 2028, with first commercial service expected around 2030.

• In the worked example one mid-band site serves about 9,200 typical smartphones but only 60 users at 25 Mbps at once.

Module 7 Quiz

Check Your Understanding

1. Why do Ericsson's and the ITU's estimates of 5G coverage in 2025, 60% and 55%, differ?

A) 5G coverage fell during 2025
B) The ITU counts only standalone 5G, while Ericsson also counts 5G radios attached to 4G cores
C) Ericsson leaves out China
D) They use different methods and data

2. What does standalone 5G add to the non-standalone form that most operators launched first?

A) Wider channels
B) Satellite links to phones
C) A new 5G core, which makes network slicing, lower delay and a better uplink possible
D) A faster 4G core network

3. What did the worked example find about one mid-band 5G site?

A) It serves thousands of typical smartphones but only dozens of users at video speeds at once
B) It serves 9,200 users at 25 Mbps at once
C) It needs a 400 MHz channel
D) It can serve only a few hundred smartphones

4. By Dell'Oro's count, what share of radio access network sales did the five largest makers hold in the first three quarters of 2025?

A) About two-thirds
B) About half, with Open RAN suppliers taking a growing share of the rest each year since 2022
C) About 96%
D) 80%

5. Under the timeline 3GPP agreed on June 10, 2026, when are the detailed functional specifications for 6G to be frozen?

A) March 2027, with the service requirements
B) December 2028
C) Around 2030, when the first services start
D) June 2026, at the meeting in Singapore

Module 8 — The Airwaves

Spectrum, Auctions, and Wi-Fi

Recommended Reading / Reference

Book: Thomas W. Hazlett, *The Political Spectrum: The Tumultuous Liberation of Wireless Technology, from Herbert Hoover to the Smartphone* (2017), an economist and former chief economist of the FCC on how American spectrum policy moved from assignment by officials toward markets

Reference: The FCC's auction pages (fcc.gov/auctions), with the rules and results of every American spectrum auction, and the Congressional Research Service's report *Spectrum Provisions in P.L. 119-21, the FY2025 Reconciliation Law* (February 25, 2026)

8a. What Spectrum Is

Radio waves are electromagnetic waves, like light but at lower frequencies. A wave's frequency is the number of times it oscillates each second, measured in hertz: a megahertz (MHz) is a million cycles a second and a gigahertz (GHz) a billion. The radio spectrum is the range of frequencies used to communicate, and every service, from broadcast radio to radar, satellite links and mobile phones, uses its own slices of it.

A wave's frequency sets its wavelength, the distance between its peaks, which equals the speed of light divided by the frequency. A 600 MHz signal has a wavelength of half a meter, a 3.5 GHz signal one of 8.6 centimeters, and a 28 GHz signal one of about 1.1 centimeters (this course's arithmetic). Lower frequencies travel farther, bend around obstacles and pass through walls more easily, so a few towers can cover a wide rural area. A wave bends around obstacles not much larger than its wavelength, so a half-meter wave slips past trees and corners that block a centimeter wave, and walls, leaves and rain absorb higher frequencies more strongly. Higher frequencies also fade faster with distance, but they are available in much wider blocks, and the width of a block sets how much data it can carry. Antennas scale with wavelength too, which is why low-band antennas on towers are large, while a millimeter-wave panel the size of a book can hold dozens of antenna elements that aim beams.

Operators therefore sort the bands into three tiers (table). Low bands provide coverage, mid bands balance coverage and capacity, and high bands, the millimeter-wave frequencies, carry the most data over the shortest distances. A cell's capacity grows with the width of its spectrum, the efficiency of its radios and the number of cells reusing the same frequencies, so operators can buy spectrum or build more cells as partial substitutes (Module 7 works through one site). Mid band is where 5G's extra capacity comes from, and it is still being built: Ericsson estimates that 5G on mid-band frequencies covered about 45% of the world's people outside mainland China at the end of 2025.

Tier	Frequencies	Wavelength	Strengths	Weaknesses
Low band	Below 1 GHz	More than 30 cm	Long reach; passes through walls	Little spectrum, so narrow channels
Mid band	About 1 to 7.125 GHz	About 4.2 to 30 cm	Reach and capacity together; 5G channels up to 100 MHz wide	Often has to be cleared of satellite, radar or other users first
High band (millimeter wave)	24.25 to 71 GHz	About 4.2 to 12.4 mm	Channels up to 400 MHz wide, or 2 GHz above 52.6 GHz; the most capacity	Short range; blocked by walls and leaves

Wavelengths are this course's arithmetic (the speed of light divided by the frequency). Channel widths are the limits in the 5G standard written by the 3rd Generation Partnership Project (3GPP), whose mid-band range ends at 7.125 GHz. Sources draw the edges of the mid band differently, and some treat the part above 6 GHz separately.

Mobile licenses come in two shapes. Older bands are paired: one block carries traffic from the network to phones and a matching block, some distance away in frequency, carries traffic back. Most newer bands are unpaired, a single block in which the network switches between sending and receiving many times a second, which lets it give more of the block to downloads, since phones download far more than they send. SpaceX's purchase in section 8e included both kinds.

Two transmitters on the same frequency in the same place interfere with each other, so spectrum has to be divided by frequency, by place and sometimes by time. A license gives its holder the exclusive right to transmit in a band of a stated width, over a stated area, for a stated term. The airwaves that carry mobile traffic are thus a public resource, divided by treaties and governments and then sold or lent to companies, which is why most of this module is about rules rather than radio.

8b. Who Decides

Spectrum is managed at two levels. Internationally, the Radio Regulations, a treaty that the ITU administers, allocate bands to services such as mobile, broadcasting, satellite and radar in each of three world regions (Europe, Africa and the Middle East; the Americas; and Asia and the Pacific), so that neighboring countries do not interfere and equipment can be sold worldwide. World Radiocommunication Conferences revise them every three to four years (section 8f), and the ITU's Radio Regulations Board hears complaints of interference between countries (Module 9 describes one). Module 12 explains the ITU itself.

Within each country a regulator assigns bands to users. In the United States the FCC manages spectrum for commercial, state and local use, while the NTIA manages the federal government's, so moving spectrum from military or other federal users to carriers needs both. Countries assign mobile licenses by auction, by administrative decision or by renewal for a fee, and many have used auctions since the 1990s.

Auctions put a market price on spectrum. The FCC held auctions from 1994 until its authority to hold them lapsed on March 9, 2023, "for the first time in its 30-year history," as the Congressional Research Service (CRS), Congress's research agency, put it in a report of February 25, 2026. The law of July 4, 2025 restored that authority (section 8d). Prices are compared in dollars per MHz-pop: the price divided by the number of megahertz and by the number of people the licenses cover. The measure lets a regional license for a narrow band be compared with a national license for a wide one (worked example).

Most recent FCC auctions run in rounds over the internet. Bidders bid on blocks of spectrum in each license area, the price rises after each round wherever demand exceeds the blocks on offer, and bidding ends when no new bids arrive. Licenses come with build-out obligations, deadlines by which the holder must cover a stated share of the population or area, which is why the FCC's approval of AT&T's purchase dealt with build-out dates (section 8e).

8c. Unlicensed Bands, Wi-Fi, and Sharing

Not all spectrum is licensed. In unlicensed bands anyone may transmit with certified equipment that stays within power limits, and devices share the band by listening before they transmit. No user is protected: each device must accept interference from the others. Wi-Fi, Bluetooth and many other devices use the unlicensed bands at 2.4 GHz and 5 GHz and, in the United States since 2020, at 6 GHz. Because phones switch to Wi-Fi at home and at work, unlicensed bands also take traffic off mobile networks.

The FCC adopted rules in 2020 opening all 1,200 MHz of the 6 GHz band, from 5.925 to 7.125 GHz, to unlicensed use (CRS). The upper part, from 6.425 to 7.125 GHz, is now contested between Wi-Fi and mobile networks. The GSMA, the mobile operators' trade association, claims that countries covering "over 80% of the global population" support using it for mobile (February 2026), and China's Ministry of Industry and Information Technology approved it for 6G technology trials on May 8, 2026 (GSMA, June 2026). Wi-Fi equipment that uses the band is sold as Wi-Fi 6E and Wi-Fi 7. The CRS lists reallocating all or part of the band for licensed use "to align with international allocations" as one option for Congress. Some versions of America's 2025 spectrum law would have kept the band from auction, but the law as enacted excludes only two other bands (section 8d).

The Citizens Broadband Radio Service (CBRS) tests a third model, sharing. Its 150 MHz, from 3.55 to 3.7 GHz, has three tiers. Federal users, chiefly the Navy, have first call. Holders of priority licenses, auctioned in August 2020 for about $4.5 billion, come second, and anyone else may use what is left at low power. Spectrum access systems, databases run by private companies, tell each radio which channels it may use at each moment (CRS).

CBRS is contested too. Under a plan reported in May 2025 and initiated by AT&T, its general-access users would move to the lower 3 GHz band and CBRS would be auctioned again for high-power 5G. The lower 3 GHz band, 3.1-3.45 GHz, is one of the two bands the 2025 law kept out of its pipeline (section 8d). The Wireless Internet Service Providers Association (WISPA), whose members use the band, warned that changing its power levels or license areas "would strand investment, chill innovation, and impose costs." In September 2025, 10 Republican senators urged Brendan Carr, the FCC's chairman (Module 12), to preserve existing CBRS operations, and the 2025 law did not exclude the band from reallocation (CRS).

8d. America's Pipeline After the Law of July 4, 2025

The law signed on July 4, 2025, Public Law 119-21, known as the One Big Beautiful Bill Act, restored the FCC's auction authority through September 30, 2034. It set a pipeline of 800 MHz for commercial wireless use: 500 MHz between 1.3 and 10.5 GHz that the NTIA must identify among federal bands, and 300 MHz that the FCC must auction. At least 100 MHz of the FCC's auctions must come from the upper C-band, 3.98 to 4.2 GHz, and be completed by July 4, 2027 (CRS).

The federal part runs on a longer clock. The NTIA must identify at least 200 MHz by July 2027 and the rest by July 2029, and the FCC must auction at least 200 MHz of that spectrum by July 2029 and the remainder by July 2033. Two bands, 3.1-3.45 GHz and 7.4-8.4 GHz, are excluded. Federal bands are slow to clear because the agencies using them must move or share, and the law lets auction proceeds pay those relocation and sharing costs. The Congressional Budget Office estimates that the auctions will raise "over $85 billion" in the fiscal years 2025 to 2034, all of it for reducing the deficit once costs are covered (CRS). Auction 113's $3.572 billion, below, was about 4% of that estimate (this course's arithmetic), so most of the money depends on the upper C-band and on federal spectrum not yet identified.

The upper C-band comes first. The FCC proposed in November 2025 to auction 100 MHz of it, with the option of going up to 180 MHz (CRS). Satellite operators use the C-band to deliver television programming to cable and broadcast companies, and the FCC's order of July 22, 2026 (FCC 26-46), adopted with one commissioner dissenting in part, will auction 160 MHz of it, from 3.98 to 4.14 GHz, as eight blocks of 20 MHz licensed by region for 15 years. A guard band from 4.14 to 4.16 GHz will separate the new licenses from satellite service, which keeps 4.16 to 4.2 GHz. The auctioned band will join the C-band licenses sold in 2021 in a single 440 MHz band from 3.7 to 4.14 GHz.

Moving the satellites costs money. The order sets a $6.3 billion pool of incentive payments, of which SES, a Luxembourg satellite operator, would receive $5.607 billion, Eutelsat $504 million and Telesat $189 million. It also estimates rebates of $3.83-5.71 billion to retrofit radio altimeters, the instruments that measure an aircraft's height above the ground on frequencies just above the band. The operators' transition plans are due on November 5, 2026, the main transition in the 75 largest license areas by December 30, 2030, and the final one by June 30, 2031 (Davis Wright Tremaine, a law firm, July 2026). A candidate to run the clearinghouse that will manage the payments is to be identified by December 15, 2026, and bidding in the auction, numbered 115, is tentatively scheduled to start on April 27, 2027 (FCC, July 24, 2026).

The FCC's first auction in four years, and its first under the 2025 law, ran from June 2 to June 23, 2026. Auction 113 offered 200 licenses for Advanced Wireless Services (AWS-3) spectrum, a name unrelated to Amazon Web Services, in the 1695-1710, 1755-1780 and 2155-2180 MHz bands, that the FCC was selling for the second time, after two companies tied to EchoStar returned them. EchoStar had warned in its filing for the first quarter of 2026 that it might owe up to about $2.921 billion in re-auction payments, which the FCC charges when licenses resell for less than the original bids. On June 26, 2026 the FCC confirmed that those payments had been satisfied (EchoStar's quarterly report for the second quarter). Seventeen qualified bidders took part, and winning bids totaled $3.572 billion after 72 rounds (FCC, June 23 and September 4, 2026). "After years on the sidelines, FCC auctions are finally back," Carr said, and Module 13 explains what the proceeds paid for.

8e. EchoStar Sells Its Spectrum

The largest spectrum deals of 2025-26 were private sales, not auctions. EchoStar, which owns the Boost Mobile brand, sold most of the spectrum under its 5G network to AT&T and SpaceX. The sales ended its attempt to run a national network of its own. It has reported "no customer traffic on our 5G Network" since November 15, 2025 (EchoStar's quarterly report on Form 10-Q for the first quarter of 2026), and it had 7.38 million wireless subscribers on June 30, 2026, now served over AT&T's network (EchoStar, August 3, 2026).

AT&T agreed on August 25, 2025 to buy EchoStar's licenses at 3.45-3.55 GHz and 600 MHz, about 30 MHz and 20 MHz nationwide, for $22.65 billion in cash, a price usually reported as $23 billion. The FCC's wireless and space bureaus approved the sale around May 12-13, 2026. They required AT&T to build out the 600 MHz spectrum faster, refused to extend its deadline to build on the 3.45 GHz licenses, and had EchoStar put $2.4 billion in escrow (Mobile World Live, May 13, 2026). AT&T completed the purchase in July 2026 (Cercle Finance, July 29, 2026).

SpaceX agreed on September 7, 2025 to buy 50 MHz of AWS-4 and H-block spectrum, two bands near 2 GHz, for $17 billion in cash and SpaceX shares, and to pay about $2 billion of interest on EchoStar's debt through November 30, 2027 (DatacenterDynamics, May 13, 2026). An amendment of November 5, 2025 added up to 15 MHz of unpaired AWS-3 spectrum for $2.6 billion in SpaceX stock. EchoStar's quarterly report says the total therefore "increased from $17 billion to approximately $20 billion," of which up to $11 billion is SpaceX stock valued at $212 a share.

The FCC approved the SpaceX purchase in May 2026, with waivers that allow the spectrum to be used on the ground, from satellites or both, and the spectrum is expected to change hands on or about November 30, 2027 (EchoStar 10-Q). The two sales left three companies, AT&T, T-Mobile and Verizon, with national mobile networks, and gave SpaceX spectrum for service from satellites straight to phones (Module 9). Carr presents such private sales as a second route to more spectrum alongside auctions. He said on May 13, 2026 that the two routes together would release about 300 MHz of low- and mid-band spectrum by the end

of 2027. For satellites, the FCC says it processed 3,418 applications in 2025, 21% more than in 2024, and freed more than 20,000 MHz of spectrum for their use (FCC, December 23, 2025).

8f. Europe, India, and the Conference of 2027

Prices depend on the band. Britain's regulator, Ofcom, sold millimeter-wave licenses on October 24, 2025: the country's three mobile operators each won 800 MHz at 26 GHz and 1 GHz at 40 GHz for £13 million ($17.3 million), or £39 million ($52.0 million) for 5,400 MHz in all, with no further payments to choose frequencies (Ofcom). The low price reflects the short reach of millimeter-wave signals. Dollar figures converted from other currencies in this module use the Federal Reserve's rates for October 16, 2026.

India's auction of June 2024 drew little demand, two years after its operators had bought 51.2 GHz of spectrum in August 2022. The government offered 10,523 MHz worth 96,238 crore rupees ($10.13 billion) at its starting prices, a crore being 10 million. The auction ended within two days, having sold 141.4 MHz, about 1.3% of the spectrum on offer, for 11,340.78 crore rupees ($1.19 billion). Bharti Airtel bought 97 MHz for 6,856.76 crore rupees, Vodafone Idea 30 MHz and Reliance Jio 14.4 MHz (Free Press Journal, June 26, 2024).

The next global decisions come at the World Radiocommunication Conference of 2027 (WRC-27), to be held in Shanghai from October 18 to November 12, 2027 (ITU). Under agenda item 1.7 it will consider identifying bands for International Mobile Telecommunications (IMT), the ITU's term for mobile broadband: 4.4-4.8 GHz everywhere outside the Americas, parts of 7.125-8.4 GHz in some regions or worldwide, and 14.8-15.35 GHz worldwide. Agenda item 1.13 studies new satellite allocations in mobile bands, so that satellites can connect directly to ordinary phones (Caribbean Telecommunications Union briefing, 2025; Module 9). Other items cover the E-band radio links that carry much mobile backhaul (Module 7), communications on and around the Moon, and the protection of radio astronomy from satellite networks.

The ITU also sets the requirements that a 6G technology must meet, called IMT-2030. Its working party on mobile systems, meeting in Geneva from May 27 to June 5, 2026, worked on the draft requirements, evaluation criteria and submission templates for candidate technologies, the step through which 3GPP's 6G standards will be recognized (Module 7 gives 3GPP's timetable). In Europe, the Digital Networks Act that the European Commission proposed in January 2026 would lengthen spectrum licenses and set more of their terms at the level of the European Union (Module 12).

The GSMA supports identifying 4.4-4.8 GHz and 7.125-8.4 GHz for mobile. It claims that each country will need 2-3 GHz of mid-band spectrum for mobile in 2035-2040, and 2.5-4 GHz in the busiest markets, against about 2 GHz available by 2030, with 6G channels two to four times as wide as 5G's widest in mid band. It wants satellite service in mobile bands to be secondary to networks on the ground and to protect them. Brazil has proposed identifying 7.125-8.4 GHz in the Americas, while America's 2025 law kept 7.4-8.4 GHz out of its pipeline (section 8d).

Worked example — What AT&T Paid for EchoStar's Spectrum

The inputs are AT&T's price of $23.0 billion as usually reported, or $22.65 billion in cash on EchoStar's filing; 50 MHz of spectrum; and 325 million people covered, the population that CreditSights, a credit-research firm, used for nationwide licenses (September 5, 2025). The 15-year period and the 6% cost of money below are illustrative assumptions. Step 1, MHz-pops: 50 MHz × 325 million people = 16.25 billion MHz-pops. Step 2, price per MHz-pop: $23.0 billion ÷ 16.25 billion = $1.42, or $1.39 on the cash price, matching CreditSights' estimate of about $1.40. Step 3, a comparison: the FCC's C-band auction of 2021, Auction 107, raised $81.2 billion in gross bids for 280 MHz of mid band (FCC, January 15, 2021; Fierce, February 17, 2021). On the same population that is 280 × 325 million = 91.0 billion MHz-pops, or $0.89 per MHz-pop. Those licenses cover only the 48 contiguous states, so their true figure is slightly higher, and their buyers also paid the satellite operators' moving costs, which this leaves out. Even so, AT&T paid about 59% more per MHz-pop. Step 4, per person: $23.0 billion ÷ 325 million = $70.77 for every American. Step 5, per month: spread over 15 years with no interest, that is $4.72 a year, or 39 cents a month. Step 6, with financing: at an assumed 6% a year, paying off $70.77 over 15 years takes $7.29 a year, or 61 cents a month. Step 7, per customer: divided instead among AT&T's 109.8 million retail wireless subscribers (Module 7), the price is $209.47 each, or $1.80 a month over 15 years at 6%. The result: a $23 billion price comes to 39 to 61 cents a month for each American, and about $1.80 a month for each AT&T wireless subscriber, because spectrum is used by hundreds of millions of people for many years.

Module 8 Key Points

The Airwaves

• Low bands reach far; high bands carry more but travel short distances; mid band, about 1 to 7.125 GHz, balances the two.

• The law of July 4, 2025 restored FCC auctions through September 30, 2034 and set an 800 MHz pipeline.

• Auction 113, the FCC’s first in four years, raised $3.572 billion in winning bids in June 2026.

• AT&T paid about $1.42 per MHz-pop for EchoStar’s spectrum, about 39 cents per American a month over 15 years.

• SpaceX’s purchase from EchoStar rose from $17 billion to about $20 billion when 15 MHz of AWS-3 spectrum was added.

• Bidding for 160 MHz of upper C-band is tentatively set for April 27, 2027; WRC-27 opens in Shanghai on October 18, 2027.

Module 8 Quiz

Check Your Understanding

1. Why can a few towers cover a wide rural area with low-band spectrum?

A) Low bands offer the widest channels
B) Low bands are unlicensed
C) Low-band signals are relayed by satellites over long distances
D) Low frequencies travel farther and pass through walls

2. When did the FCC's general authority to auction spectrum lapse?

A) March 9, 2023
B) July 4, 2025, when the reconciliation law passed
C) September 30, 2034
D) June 23, 2026, at the end of Auction 113

3. How does the Citizens Broadband Radio Service share its 150 MHz?

A) It is sold as exclusive national licenses, like other mobile bands.
B) It is fully unlicensed, like Wi-Fi's 5 GHz band.
C) Satellite operators have first call, and phones use the rest.
D) Federal users come first, priority licensees second, and anyone else may use what is left.

4. How much did Auction 113 raise in winning bids in June 2026?

A) $81.2 billion
B) $23 billion
C) $3.572 billion
D) Over $85 billion

5. In the worked example, about what did AT&T pay per MHz-pop for EchoStar's spectrum?

A) $0.89
B) About $1.42
C) About $70.77
D) About $209.47

Module 9 — Satellites

Starlink, Its Rivals, and the Crowding of Orbit

Recommended Reading / Reference

Book: Ashlee Vance, *When the Heavens Went on Sale: The Misfits and Geniuses Racing to Put Space Within Reach* (2023), a technology reporter's account of the start-ups, among them Planet Labs and Rocket Lab, that made satellites and rockets small and cheap enough to fill low orbit

Reference: Jonathan McDowell's *General Catalog of Artificial Space Objects* (planet4589.org), an astronomer's independent record of every satellite launched, and SpaceX's reports to the FCC every six months on Starlink's maneuvers to avoid collisions

9a. Geostationary Against Low Orbit

A satellite stays up because it falls around the Earth as fast as the surface curves away beneath it, and the higher the orbit, the longer each circuit takes. At 35,786 kilometers (km) above the equator a circuit takes 23 hours and 56 minutes, the time the Earth takes to turn once, so a satellite there stays over one spot and a fixed dish can point at it. This is geostationary orbit (GEO). From there a satellite can see about 42% of the Earth's surface (this course's geometry), so three of them can cover nearly all of the inhabited world.

The cost of that height is delay. A request from a home goes up to the satellite and down to a ground station, and the reply makes the same two legs back, so the signal travels at least four times 35,786 km. At the speed of light that takes about half a second, as the worked example shows. Streaming video, which loads ahead of what is being watched, barely notices. Video calls, online games and many business applications do.

Low Earth orbit (LEO), from a few hundred to about 2,000 km up, cuts the distance. Starlink's satellites fly at about 480-550 km, some 65 times closer to the ground than geostationary ones, and Eutelsat OneWeb's at about 1,200 km. A satellite at 550 km circles the Earth in about 96 minutes (this course's arithmetic) and is in view of any one home for only minutes, so continuous service needs thousands of satellites, with each home's connection passed from one to the next. Everyone under a satellite's beams shares its capacity, so the capacity over a given area, not the network's total, decides how many homes it can serve (section 9g).

Low orbits also clean themselves. The thin air at those heights slowly drags satellites down, so a dead satellite at 480 km falls back and burns up within a few years, while one in geostationary orbit stays up indefinitely. SpaceX said in a report to the FCC that it was lowering about 4,400 of its satellites from 550 km to 480 km during 2026, and it now expects each satellite to last about five years (Space Intel Report, January 13, 2026).

9b. Starlink in 2026

Starlink, owned by SpaceX, is the largest satellite network ever built. On June 30, 2026 it had roughly 10,200 operational satellites, about 9,600 of them for broadband and the rest serving phones directly, about 12 million subscribers in 167 markets, and about 800 terabits per second (Tbps) of network capacity, by SpaceX's accounts of its second quarter (Nasdaq, August 4, 2026; Subscription Insider, August 5, 2026). Its third-quarter report of October 20, 2026 counted 13.6 million consumer subscribers and coverage in 174 markets on September 30, and about 11,100 satellites in orbit, a count that includes some approved for

deorbiting (Broadband Breakfast; TeslaNorth, October 20, 2026). Trackers that count satellites themselves give slightly different totals, and counts of about 8,000 satellites still in circulation are out of date.

Growth has been fast. Starlink had 2.3 million subscribers at the end of 2023, 4.4 million at the end of 2024 and 10.3 million on March 31, 2026, according to the prospectus that SpaceX filed with the Securities and Exchange Commission (SEC) in May 2026, before its initial public offering of shares (Broadband Breakfast, May 22, 2026). It added more than 1.7 million consumer subscribers in the second quarter of 2026 (Subscription Insider, August 5, 2026) and 1.6 million in the third.

Average revenue per user (ARPU), revenue divided by subscribers, has fallen as subscribers multiplied: from $99 a month in 2023 to $85 in the second quarter of 2025 and $66 in the first and second quarters of 2026, before rising 5% to $69 in the third, after Starlink raised its prices (TeslaNorth, October 20, 2026). Revenue has still grown fast. SpaceX's connectivity business, mostly Starlink, earned $11.4 billion in 2025, up 50% and 60% of SpaceX's revenue of $18.7 billion, and $4.3 billion in the second quarter of 2026, up 66% from a year earlier. Of that quarter's revenue, about $2.5 billion came from consumers and about $1.8 billion from businesses and governments, whose spending had more than doubled, up 108% (Nasdaq; Subscription Insider). In the third quarter it earned $5.2 billion, up 69%, with $3.1 billion from consumers and $2.1 billion from businesses and governments, up 91% (SpaceX, October 20, 2026, as reported by Broadband Breakfast and TeslaNorth).

SpaceX plans to add capacity with third-generation satellites (V3), each designed for 1 Tbps of downloads, which its prospectus said would begin launching on its Starship rocket in the second half of 2026; a Starship flight in the third quarter carried the first 26 (SpaceX, October 20, 2026). Elon Musk, SpaceX's chief executive, said on the call in August 2026 that a V3 satellite would carry about ten times as much as the current version. The more capacity each launch adds, the faster the network can grow and the lower its cost per subscriber.

9c. The Rivals

Amazon's network, called Project Kuiper until Amazon announced in November 2025 that it was renaming it Amazon Leo, is Starlink's main Western rival (table). The FCC's authorization of July 2020 required half of its satellites to be in orbit by July 30, 2026. Amazon asked for a 24-month extension, and in June 2026 the FCC instead granted a conditional waiver, under which satellites launched after the deadline lose priority in their radio bands until March 2028 (KeepTrack, a satellite-tracking site, October 21, 2026).

Amazon says production is no longer what holds it back: its factory in Kirkland, Washington, can build five satellites a day, and "now we can easily make 27 a week," said Rajeev Badyal, vice president of Amazon Leo, in October 2026. Launches have lagged, and none was listed between July 2 and October 21, 2026. "The plan is to launch initial service in 2026," Chris Weber, its vice president of business and product, said after the network missed a mid-2026 target (Light Reading, October 12, 2026). Amazon has announced three terminals, rated at up to 100 megabits per second (Mbps), 400 Mbps and 1 gigabit per second of downloads, and a South African partner, Herotel, plans to offer the service in 2027 (KeepTrack; Wikipedia).

Network	Owner	Satellites launched	Planned size	Service, October 2026
Amazon Leo	Amazon (United States)	396 production satellites and 2 prototypes, since deorbited, by July 2, 2026; 392 objects in orbit in October	3,236 authorized by the FCC in July 2020	Terminals for business customers since November 2025; no date named for general service

Network	Owner	Satellites launched	Planned size	Service, October 2026
Eutelsat OneWeb	Eutelsat (France)	652 operational on October 20, 2024, the latest count found	Not confirmed in the sources read	Sold wholesale through operators and governments
Guowang	China SatNet, a state company	About 177 by June 17, 2026	More than 13,000	No public service reported in the sources read
Qianfan (Spacesail)	Shanghai Spacesail Technologies	258 by September 15, 2026	More than 15,000, with 1,296 in its first phase	Missed its goal of 648 satellites in orbit by the end of 2025

Amazon Leo: Light Reading (October 12, 2026) and KeepTrack (updated October 21, 2026). OneWeb, Guowang and Qianfan: launch records compiled in Wikipedia's articles on each (read October 25, 2026); no 2026 count for OneWeb was found.

Amazon is also buying spectrum for phones. On April 14, 2026 it agreed to buy Globalstar, a satellite operator, for $11.57 billion, or $90 a share, for its satellites and radio spectrum, subject to regulators' approval (SDxCentral). Globalstar carries the emergency texting service of Apple's iPhones (section 9d).

Eutelsat OneWeb sells capacity wholesale, through telecom operators and governments, rather than to households, and it is Europe's main low-orbit network. Jean-François Fallacher has been chief executive of Eutelsat, its owner, since June 1, 2025. China is building two state-backed networks, Guowang and Qianfan, also called Spacesail, which together plan more than 28,000 satellites, nearly three times Starlink's fleet in mid-2026 (table).

Geostationary operators remain large. SES, a Luxembourg company, completed its purchase of Intelsat in July 2025, creating a fleet of about 90 geostationary satellites and nearly 30 in medium Earth orbit, between low and geostationary orbit. The combined company had pro forma revenue of €3.7 billion ($4.16 billion) and contract backlog above €8 billion ($8.98 billion) (Advanced Television, July 17, 2025), and SES stands to receive most of the C-band payments described in Module 8. Viasat, an American company, also runs broadband satellites in geostationary orbit. Dollar figures converted from euros in this module use the Federal Reserve's rate for October 16, 2026.

9d. Satellites Direct to Phones

The newest market connects satellites straight to ordinary phones, with no dish. T-Mobile US and Starlink launched T-Satellite, a texting service, on July 23, 2025. It has since added some apps, such as WhatsApp and Google Maps, and costs $10 a month on some plans, while voice calls were still in development in July 2026. Customers of AT&T and Verizon can use it with a compatible phone, texts to emergency services are free for all T-Mobile customers, and it works for roaming customers in Canada and New Zealand (satelliteinternet.com, July 2, 2026). Starlink had launched more than 650 satellites for this service by January 2026, and SpaceX's prospectus reported 7.4 million devices using it each month in about 30 countries.

A phone's small antenna and low power make the link weak, so each satellite carries little data to phones, and service so far is mostly texts and light apps. SpaceX bought spectrum from EchoStar to change that (Module 8). Gwynne Shotwell, SpaceX's president and chief operating officer, said in August 2026 that it was integrating 65 megahertz (MHz) of EchoStar spectrum, and SpaceX plans to launch next-generation satellites for phones in 2027 and to expand mobile service by the end of 2027.

AST SpaceMobile, an American company, is taking a different route: large satellites designed to give phones broadband rather than texts. It reported a peak of 98.9 Mbps from an orbiting satellite to an unmodified smartphone over international waters, and the FCC has authorized commercial service in the

United States using up to 248 satellites (company report, May 11, 2026). One satellite, BlueBird 7, was lost after a New Glenn rocket left it in too low an orbit on April 19, 2026, and six more, BlueBirds 8 to 13, went up within 50 days by early August, giving AST 13 spacecraft in orbit. AST's partners include AT&T, Verizon and Vodafone. It reported revenue of $31.5 million in the second quarter of 2026 and kept its forecast of $150-200 million for the year (company report, August 10, 2026).

Apple has offered emergency texting by satellite on iPhones since 2022 through Globalstar, which reserves 85% of its network capacity for Apple. Apple owns 20% of Globalstar, a stake bought in November 2024 for about $400 million, and agreed to provide $1.1 billion to expand the service. Under Amazon's purchase, Amazon Leo is to power the satellite services of the iPhone and the Apple Watch, emergency texting included (SDxCentral). Which satellites may use mobile bands, and on what terms, is on the agenda of the world radio conference of 2027 (Module 8).

9e. Crowding and Debris

Low orbit is filling up. Space.com counted about 16,000 operational spacecraft in mid-2026 (July 15, 2026), so SpaceX's count of roughly 10,200 working Starlink satellites was about 64% of everything operating in orbit (this course's arithmetic). Every one of them must avoid the others and the debris of earlier launches.

SpaceX's reports to the FCC show how often. Its satellites made 207,152 maneuvers to avoid collisions from December 2025 to May 2026, against 148,696 in the previous six months and 144,404 in the six months before that (Space.com, July 15, 2026; Space Intel Report, January 13, 2026). The year to May 2026 thus brought 355,848, about 35 for each of the roughly 10,200 satellites working in June 2026 (this course's arithmetic). SpaceX moves a satellite whenever its estimated chance of a collision exceeds 3 in 10 million, against a common industry threshold of 1 in 10,000.

SpaceX reported that one Chinese satellite, Honghu-2, alone forced 1,143 maneuvers, and it complains that some operators, including Chinese and Russian ones, do not share their satellites' precise orbits. Hugh Lewis, a professor of astronautics at the University of Birmingham, projects one million cumulative Starlink maneuvers by about June 2027 (Space.com).

Debris adds to the risk. When the upper stage of the rocket that launched Qianfan's first satellites broke up in August 2024, US Space Command counted more than 300 pieces large enough to track. Each fragment is a new hazard that crosses the orbits of working satellites, which is one reason operators fly low enough for dead satellites to fall back quickly (section 9a).

Orbit is also a contest over radio frequencies. ITU rules dating from the 1990s limit the interference that low-orbit systems may cause to geostationary ones, and the world radio conference of 2027 will consider such limits in one band, 17.3-17.7 gigahertz (GHz), along with satellite service to phones (Module 8).

9f. Satellites in War

Starlink became a military network in Ukraine. SpaceX switched it on there in 2022 at the request of Ukraine's government, and Ukraine's forces came to depend on it for front-line communications and drone operations. Russian forces used smuggled terminals too, as Ukraine said in February 2024. In February 2026 SpaceX cut off terminals in Ukraine that were not on an approved list, which reports said disrupted Russian military communications.

In October 2026 a cap appeared on Starshield, the military version of Starlink that Ukraine uses to steer interceptor drones against the jet-powered Shahed attack drones that Russia launches. The connection cut

out whenever a terminal moved faster than 160 km an hour (100 miles an hour), far slower than the drones it chases (Associated Press, October 16, 2026). Who imposed it was disputed. Elon Musk said that the US government owns and controls Starshield and that decisions about its use rest with the US military and ultimately with the president. President Donald Trump said on October 16, “it wasn’t me, and I don’t think it was Elon.” Press accounts conflict: Reuters, citing US officials, reported on October 15 that the White House had ordered the cap, and The Wall Street Journal, as reported by the Kyiv Post on October 21, that Musk had.

Serhii Beskrestnov, a defense technology adviser to President Volodymyr Zelenskyy, called the cap “not only a technical issue and not only a business issue, but also a political one,” and said Ukraine was negotiating with the Pentagon, SpaceX and the US government. He said Ukraine wanted either the cap lifted or a separate, costlier Starshield service for aircraft without one, and he called Rassvet, a Russian satellite network that already gave Russia windows of coverage over Ukraine of 30 to 40 minutes, “danger No. 1” among the technological challenges Russia poses. The two episodes of 2026 show that decisions made by one company, or by the government it answers to, can cut off or restrict the communications of armies in the field.

In Iran, Starlink worked against the government. During Iran’s internet shutdowns of 2026 (Module 13), smuggled terminals kept some Iranians online: at least 50,000 were in the country by January 2026, by an estimate that the British Broadcasting Corporation (BBC) reported in May 2026. Iran made possessing one a crime, and the BBC and Al Jazeera report different maximum prison terms. Mohammad Amin Aghamiri, secretary of Iran’s Supreme National Cyberspace Council, has called Starlink a “tool of war.” Iran has pressed the ITU’s Radio Regulations Board since 2023 to stop Starlink transmissions licensed by the United States and Norway, and it summoned Norway’s ambassador on October 4, 2026 (Al Jazeera, October 14, 2026).

9g. Can Satellites Replace Fiber and Towers for Rural Homes? The Cases as Their Advocates Make Them

The case that satellites can replace them: SpaceX and other advocates of low orbit argue that a satellite connection reaches any home within days, with no trenches, poles or permits, and the NTIA’s 2025 restructuring of America’s rural broadband program let satellites compete for its money on equal terms (Module 11). Starlink’s growth shows that people choose it where they have the choice (section 9b), and its capacity rises with each generation of satellites. Delay is no longer a weakness at low orbit (worked example). Subsidizing fiber to the most remote homes, they add, costs several times as much per location as a satellite connection (Module 11 gives the bids), so the same money connects more homes, sooner.

The case that they can only fill gaps: advocates of fiber and of mobile networks argue that a satellite’s capacity is shared by everyone beneath it, so a network that suits scattered homes cannot serve a town or a growing suburb. A satellite at 480 km lasts about five years, so the network must be rebuilt continually, while fiber in the ground lasts for decades. Ericsson, the Swedish equipment maker, forecasts about 33 million satellite broadband subscriptions worldwide in 2031, up from about 10 million at the end of 2025, against 350 million fixed wireless connections. The GSMA, the mobile operators’ trade association, says the effect of satellites that connect directly to phones on bringing offline people online is “likely to be small.”

What the evidence says: delay is no longer the obstacle, but capacity per area is. The worked example shows that Starlink’s network as a whole has room for its subscribers, yet one satellite’s average capacity would be used up by the homes of a small city at a typical evening load, and its beams must serve everyone else beneath it too. Independent measurement has found the limits before. In August 2022 the FCC rejected

SpaceX’s claim to $885.5 million from the Rural Digital Opportunity Fund for 642,925 locations in 35 states, citing Ookla’s tests, whose median speeds fell below the program’s required 100 Mbps down and 20 Mbps up (FCC, August 10, 2022; Fierce Network). America’s subsidy decisions of 2025-26, made after satellites had grown, still gave most rural locations to fiber and a minority to satellites (Module 11).

The course’s reading, stated once: low-orbit satellites fill gaps, and the gaps are large. They are the right answer for the most remote homes, where a fiber line would cost several times as much and serve few people, but they cannot replace fiber and towers where homes cluster, because the capacity over each area is shared and set by the satellites overhead. A five-year satellite life also makes their cost recurring. The sound test for each location is the cost of meeting its expected demand over a decade or more, which favors satellites where homes are few and far apart and fiber where they are not.

Worked example — Half a Second Against a Few Milliseconds

The inputs are the speed of light, 299,792 km a second; the Earth’s mean radius, 6,371 km; Starlink’s 550 km shell; and a path in which a request goes from the home to the satellite and down to a ground station, and the reply returns the same way, four legs in all. The dish’s angle of 40° above the horizon, the 20 milliseconds (ms) allowed for processing and routing, and the evening load per home are illustrative assumptions. Step 1, the geostationary height: by Kepler’s law, an orbit whose period equals one turn of the Earth, 86,164 seconds, has a radius of 42,164 km, which less the Earth’s equatorial radius of 6,378 km is 35,786 km. Step 2, the shortest delay: with the satellite straight overhead, four legs are 4 × 35,786 = 143,144 km, which light covers in 477.5 ms. Step 3, a real dish: at 40° above the horizon the satellite is 37,778 km away, so four legs take 504.1 ms, and with 20 ms of processing the round trip is about 524 ms. Step 4, Starlink: at 550 km the overhead path of four legs is 2,200 km, or 7.3 ms, and at 40° the satellite is 812 km away, so four legs take 10.8 ms, or about 31 ms with the same 20 ms added. Low orbit cuts the delay about 17-fold. Step 5, capacity per subscriber: SpaceX’s 800 Tbps divided among its 12.0 million subscribers, both counted on June 30, 2026, the last date it gave its capacity, is 66.7 Mbps each, if every subscriber were online at once and all of the capacity were over them, which it is not, since much of it passes over oceans. Step 6, capacity per satellite: 800 Tbps across 9,600 broadband satellites averages 83.3 gigabits per second (Gbps), about 56 times the middle-case 5G site in Module 7. Step 7, homes per satellite: a home using 711 GB a month, the American average in the second quarter of 2026 (Module 1), with an assumed 10% of each day’s traffic in its busiest hour, needs 5.27 Mbps in that hour, as Module 7’s worked example computes. One satellite’s average capacity therefore serves about 15,800 such homes at once, if all of its beams were over them. The result: low orbit solves the delay of geostationary satellites but not the sharing of capacity, which is set by how many satellites are overhead.

Module 9 Key Points

Satellites

• A geostationary round trip takes at least 477.5 ms by geometry; Starlink’s low orbit cuts the delay about 17-fold.

• Starlink had about 10,200 operational satellites on June 30, 2026, and 13.6 million consumer subscribers on September 30 (SpaceX).

• Starlink’s average revenue per user fell from $99 a month in 2023 to $66 in 2026, while revenue kept rising.

• Amazon Leo had launched 396 production satellites by July 2, 2026, and still planned initial service in 2026.

• Starlink made 207,152 collision-avoidance maneuvers from December 2025 to May 2026, up 39% on the previous six months.

• Satellites fill large gaps but cannot replace fiber where homes cluster, because capacity over each area is shared.

Module 9 Quiz

Check Your Understanding

1. Why does a request and reply through a geostationary satellite take about half a second?

A) Geostationary satellites process each signal slowly

B) The signal is carried partly by fiber

C) The signal must travel at least four times 35,786 km, up and down for the request and again for the reply

D) Geostationary satellites pass each signal through several other satellites before it lands

2. How many subscribers did Starlink have on June 30, 2026, by SpaceX's account?

A) About 8,000

B) 10.3 million

C) About 4.4 million, the count SpaceX gave in its prospectus for the end of 2024

D) About 12 million

3. What did the FCC do after Amazon asked for more time to meet its July 30, 2026 deadline?

A) Revoked its license

B) Granted a conditional waiver

C) Granted the 24-month extension

D) Reassigned its spectrum to SpaceX

4. In the worked example, how much capacity would each Starlink subscriber have if all were online at once?

A) About 67 Mbps

B) About 800 Tbps

C) About 5 Mbps, the evening load of a home that uses 711 GB a month

D) About 83.3 Gbps

5. What is the course's reading on whether low-orbit satellites can replace fiber and towers for rural homes?

A) They will replace fiber for most homes within a decade.

B) They are too slow for video calls at any height.

C) Their delay rules them out, whatever their capacity.

D) They fill gaps: right for the most remote homes, but shared capacity per area limits them where homes cluster.

Module 10 — The Cloud and Data Centers

Hyperscalers, Power, and Outages

Recommended Reading / Reference

Book: Tung-Hui Hu, *A Prehistory of the Cloud* (2015), a media scholar's account of the cloud's physical and political history, from the time-sharing computers of the 1960s to fiber laid along railroad lines

Reference: Synergy Research Group's quarterly estimates of the cloud market, and the post-incident reports that Amazon Web Services, Microsoft and Cloudflare publish after major outages

10a. What the Cloud Is

"The cloud" is computing rented from someone else's data centers. A company that once bought its own servers, the computers that run websites and store data, now rents computing power, storage and software and pays by the hour or by use. The largest providers are Amazon Web Services (AWS), Microsoft's Azure and Google Cloud. The industry calls them hyperscalers, because each builds and runs data centers at a scale no customer could match.

What is rented comes in layers, and the first, infrastructure, is raw computing, storage and network capacity, usually as virtual servers, slices of a physical machine that behave like separate computers. Platforms add tools that the provider runs, such as databases and AI services. Software means finished applications, such as email, used through a browser. Synergy Research Group, a market-research firm, counts the first two, with hosted private clouds, as cloud infrastructure services, its measure of the market (section 10b).

Renting turns a large purchase into a running cost and lets a company add capacity in minutes rather than months. Its price is dependence, because software written for one provider's tools is costly to move to another, and a customer shares that provider's failures (section 10e).

Providers group their data centers into regions and availability zones. A region is a set of data centers in one metropolitan area, such as northern Virginia or Frankfurt. An availability zone is one or more data centers with their own power, cooling and network links. AWS says the zones in a region are many kilometers apart but within 100 kilometers (km) of each other, and it counted 124 zones in 39 regions on a page updated on September 17, 2026.

The design lets a customer survive the loss of a building. A service run in two or three zones keeps working if one zone fails, and one copied to a second region can survive the loss of a whole region. Many customers do neither, because each extra copy costs money. AWS's oldest region, us-east-1 in northern Virginia, hosts so many companies' systems that its outage of October 20, 2025 stopped services used around the world (section 10e).

Data centers reach users through the providers' own fiber networks, through exchange points where networks meet, and through content delivery networks that keep copies of popular files near users (Module 4). AWS says it runs nearly 20 million km of its own fiber on land and under the sea (September 2026). The cloud companies are also now the main builders of new ocean cables (Module 5).

10b. The Market

Cloud infrastructure services earned $143.4 billion worldwide in the second quarter of 2026, by Synergy Research Group's estimate of July 30, 2026. That was 43% more than a year earlier, the fastest growth in eight years, and the eleventh quarter in a row in which growth sped up. Revenue over the twelve months to June was about $500 billion.

Three companies take most of the money: Synergy put Amazon's share at 28%, Microsoft's at 20% and Google's at a record 15%, together 63%. A year earlier Amazon had held 30% and Google 13% (Synergy, as reported by Shattered.io in September 2026). On these shares Amazon's cloud revenue grew by about a third while the market grew 43%, and Google's grew by about two-thirds (this course's arithmetic).

The fastest growth is among smaller providers that rent out AI computing. Synergy named Oracle, CoreWeave, Crusoe, Nebius and Nscale among the fastest-growing second tier. It estimated that cloud services for generative AI, systems that produce text, images and code, grew 165% in the year to the second quarter. John Dinsdale, Synergy's chief analyst, said on July 30, 2026 that "AI technology has lit a fire under the cloud market."

The companies' own figures differ from Synergy's, because Amazon, Microsoft and Google each report revenue for a cloud division whose contents they define, including software and other items, so their totals and growth rates cannot be compared directly. Market shares in this course are Synergy's, which measure every provider the same way, and its figures for the second quarter of 2026 are the latest available.

10c. Inside a Data Center

A data center is a building of computers and the equipment that keeps them running. Servers sit in racks, metal frames about two meters tall, joined by network switches and fiber to each other and to the outside world. Power arrives from the grid through substations and transformers. Batteries carry the load for the seconds or minutes that diesel generators need to start if the grid fails.

Nearly all the electricity a computer uses ends up as heat, so cooling is the other half of the building. Older halls blow chilled air through the racks, while newer ones increasingly pipe liquid to plates on the chips, because the graphics processing units (GPUs) that run AI draw far more power per rack than ordinary servers. Cooling by evaporating water saves electricity but consumes water.

Data centers used about 485 terawatt-hours (TWh) of electricity in 2025, each a billion kilowatt-hours, or about 1.5% of the world's electricity, by figures from the International Energy Agency (IEA) of April 16, 2026. Their use rose 17% in 2025, against 3% for all electricity demand. Centers built mainly for AI used about 155 TWh (IEA figures compiled by Our World in Data, July 20, 2026).

The IEA expects data centers to use about 950 TWh by 2030, about 3% of the world's electricity, with centers built for AI about half of it. The United States used about 45% of the world's data-center electricity in 2024 (IEA, April 2025). Lawrence Berkeley National Laboratory estimated in December 2024 that American data centers used 176 TWh in 2023, 4.4% of the country's electricity, and would use 6.7-12% by 2028. The width of that range reflects different assumptions about how many AI chips are shipped and how hard they run.

Builders look for power, land, fiber and nearby customers. The IEA reported in June 2026 that 19 of the 21 large AI data centers under construction in the United States lie within 65 km of a city of more than 100,000 people, so their demand falls on grids that also serve large populations. It also counted a pipeline of 47 gigawatts (GW) of gas-fired generation planned at data-center sites themselves.

10d. The Spending, and the Networks It Needs

A few companies pay for most of the build-out: five large technology companies spent more than $400 billion of capital in 2025 (IEA, April 2026). In earnings calls in July 2026, the four largest spenders, Amazon, Microsoft, Alphabet (Google's parent) and Meta, guided to about $720-745 billion for 2026, and about $800 billion with Oracle. Microsoft and Alphabet report third-quarter results on October 28, 2026 and Amazon on October 29, and those reports will update the plans.

The IEA said in June 2026 that these companies' capital spending now exceeds the world's investment in oil and gas production. The last great wave of network building, the fiber boom of the late 1990s, ended in a bust (Module 2).

AI changes the network inside the building first: training a large model splits the work across thousands of GPUs that must exchange partial results throughout the job, so they are joined by very fast links, many of them fiber. In such a building most traffic runs between machines inside it rather than out to users. Answering users' requests, called inference, moves far less data per request and can run closer to the users.

AI also adds traffic between data centers, because companies spread training across campuses that may be hundreds of kilometers apart and copy models and data between regions, all of which needs long-distance fiber. Zayo, an American fiber company, was investing $4 billion in more than 5,000 route miles (about 8,000 km) of new long-haul routes for AI data centers (Data Center Frontier, March 20, 2025). In January 2026 Meta and Corning announced a multiyear deal worth up to $6 billion for Corning's fiber, cable and connectors for Meta's American data centers.

The same demand reaches the ocean floor and the exchange points: Meta has tied its planned cable around five continents to its AI plans (Module 5), and exchange operators credit AI for part of their traffic growth (Module 4).

10e. When the Cloud Fails

Four outages between July 2024 and November 2025 show how concentrated systems fail (table). Each began with a routine change that the company itself made, which automation spread to many machines before anyone could stop it, and none was an attack.

Date	Whose system	How long, on the company's timeline	What users lost
July 19, 2024	CrowdStrike's Falcon security software on Windows computers	The faulty update was live for 78 minutes, from 04:09 to 05:27 Coordinated Universal Time (UTC); repairs took days	Airlines, banks, hospitals and broadcasters lost the use of crashed computers
October 20, 2025	AWS, us-east-1 region (northern Virginia)	About 14.5 hours, from 11:48 p.m. Pacific time on October 19 to 2:20 p.m. on October 20	Slack, Atlassian, Snapchat and many other services, and AWS's own services in the region
October 29, 2025	Microsoft, Azure Front Door	More than eight hours, from about 15:45 UTC	Microsoft 365, Xbox, Minecraft, the Azure Portal and customers' websites
November 18, 2025	Cloudflare's network	Failures from 11:20 UTC; main impact fixed by 14:30 and all services by 17:06	Websites and apps behind Cloudflare, and its own dashboard and logins

CrowdStrike's technical note of July 20, 2024 and SecurityWeek; AWS's post-event summary (October 2025); Microsoft's statements as reported by The Stack, October 29-30, 2025; Cloudflare's post-mortem as summarized by iLert. Pacific time was seven hours behind UTC in October 2025.

The first was not a cloud failure, but it shows the same risk. CrowdStrike, an American security company, sent its Falcon software a routine update to the files that tell it what to look for. A logic error in the update crashed the computers that received it, and Microsoft estimated on July 20, 2024 that about 8.5 million Windows devices were hit, under 1% of all Windows machines. Because a crashed computer could not download a fix, many had to be repaired by hand, one at a time.

At AWS the fault lay in automation that maintains the DNS records for DynamoDB, its database service. Two copies of the automation ran out of step: a delayed one applied an old plan over a newer one, and a clean-up step then deleted it, leaving the database's address in the region with an empty record. AWS's post-event summary calls the cause "a latent race condition," a fault that appears only when two processes run in an unlucky order. DNS was restored in under three hours, but systems that depended on the database, including the one that launches new virtual servers, took until the afternoon. AWS has since switched off that automation worldwide.

Azure Front Door is Microsoft's global system for receiving web requests, for its own services and its customers' sites, and passing them on. Microsoft blamed "an inadvertent tenant configuration change," a settings change for one customer's account, and said "a software defect" had let the change "bypass safety validations." It froze all further changes, rolled the whole network back to the last configuration known to work, and brought servers back in stages to avoid overloading them.

At Cloudflare, which sits in front of many websites to speed and protect them (Module 4), the outage began with a change to who could read one of its databases. The change made a routine query return its rows twice, so a file used by its bot-detection system roughly doubled in size and passed a hard limit in the software that handles its traffic. That software failed across the network, recovering and failing again as good and bad copies of the file went out every five minutes. Engineers first suspected a distributed denial-of-service (DDoS) attack, a flood of junk traffic. Matthew Prince, Cloudflare's co-founder and chief executive, called it the company's worst outage since 2019 in its report of November 18, 2025.

Each company published a detailed account within days, and each fix changed how updates are tested and released. The outages ended in hours, but each stopped thousands of other companies' services at the same moment, which is the question of section 10g.

10f. War Reaches the Data Center

The Iran war, which began on February 28, 2026 with American and Israeli strikes on Iran, brought drone strikes on cloud data centers (Module 13 covers the war's effects on networks). On March 1-3 drones hit AWS sites in two Gulf states, in strikes that press reports attributed to Iran. AWS said two of its facilities in the United Arab Emirates (UAE), in its ME-CENTRAL-1 region, were "directly struck," and a strike near a facility in Bahrain, in its ME-SOUTH-1 region, damaged equipment.

The damage was physical: Amazon said the strikes "caused structural damage" and "disrupted power delivery to our infrastructure" (DatacenterDynamics), and the sites suffered fires, water damage from sprinklers and power cuts (The Register, March 2, 2026). Its storage service returned high rates of failed requests, and AWS urged customers to move their workloads to other regions. Later in March, drone activity disrupted the Bahrain region a second time, and Iran's military had threatened the region's digital and energy infrastructure if the bombing of Iran continued (DatacenterDynamics).

Two months later the UAE region had not recovered: on April 30, 2026 AWS said it was "currently unable to reliably support customer applications" there and that restoring the region "is expected to take several months." It told customers to "migrate all accessible resources to other Regions" and to restore the rest from

backups kept elsewhere, and it paused billing in both regions (Developing Telecoms, May 4, 2026). In September 2026 AWS said it could not restore several sites in Bahrain and one in the UAE, and that it had “exhausted every option” to recover some data held in Bahrain (Al Jazeera, October 13, 2026).

The strikes tested the design of section 10a. Availability zones guard against a fire or flood in one building, but two of the UAE region’s zones were hit in the same attack, and AWS builds each region with at least three. Customers with copies in another region could switch to them, and those that relied on the region alone could not. The Gulf states’ plans to become centers of AI computing rest on data centers and cables inside a war zone (Module 5 covers the Gulf’s cables).

10g. Has Concentration Made the Internet More Fragile? The Cases as Their Advocates Make Them

The case that concentration has made it more fragile: its advocates, among them network engineers and financial regulators, argue that the internet has gained single points of failure. Three companies hold almost two-thirds of cloud infrastructure (section 10b), and services that look independent often run in the same few regions. Within 30 days in late 2025, a race condition at AWS, a configuration change at Microsoft and a doubled file at Cloudflare each stopped services worldwide. One provider’s failure can spread to another: Google Cloud’s outage of June 12, 2025 also took down Cloudflare services that relied on it (TechCrunch). Concentration also gives an enemy a target, as the drones in the Gulf showed.

Europe’s financial regulators act on this view: on November 18, 2025 the European Supervisory Authorities, the EU’s regulators of banking, insurance and securities, named 19 technology companies as critical providers to Europe’s financial firms, AWS, Google, Microsoft and Oracle among them. The designation puts them under direct oversight under the Digital Operational Resilience Act (DORA), which has applied since January 17, 2025.

The case that it has not: the providers and many engineers answer that a few large clouds are more reliable than thousands of company server rooms. They spend more on spare power, networks and security than most customers could, and they publish detailed reports on their failures. Their outages end in hours, and customers that spread their services across zones and regions keep working through them. Most large disruptions are not cloud failures at all: almost half of the major disruptions that Cloudflare counted in 2025 were shutdowns ordered by governments (Module 13). The internet’s core, its routing system and DNS root servers, remains spread across thousands of independent networks and sites (Module 3).

What the evidence says: outages are becoming rarer per site but more concentrated in providers. The Uptime Institute, which researches and certifies data centers, reported on May 13, 2026 that outage rates per site had fallen for the fifth year in a row, though more slowly, and that about 1 in 10 operators called their latest outage serious or severe. In nine years of tracking publicly reported outages, it found that third-party providers caused about two-thirds of them: cloud and internet companies, telecom carriers, and colocation firms, which rent space in data centers. Andy Lawrence, executive director of its research arm, Uptime Intelligence, said with the report that “digital infrastructure is remarkably resilient.” The outages of 2025 show the other side: each lasted hours, not days, but each stopped thousands of companies at once. No study found for this course measures whether the internet as a whole fails more often than it did.

The course’s reading, stated once: concentration has made the internet’s failures larger and more correlated, not more frequent. Outages per site are falling, but when one large provider fails, thousands of services fail together, as 2025 showed three times in a month. The triggers were the providers’ own changes, which slower, staged releases and better testing can catch, and each provider changed its

methods afterward. War is a risk that availability zones were not built for, and only copies in another region or with another provider guard against it. The fragility is real but specific, and it falls hardest on customers that rely on a single region.

Worked example — A Video Call from New York to London

Follow one moment of a group video call at 720p, a high-definition picture 720 lines tall, from a laptop in New York to one in London. Every delay below except the ocean crossing is an illustrative assumption, a typical range rather than a measurement. Step 1, capture and encoding: the camera takes a frame every 33 milliseconds (ms) at 30 frames a second, and the software compresses sound and picture in about 20-40 ms. Step 2, packets: the audio is cut into 20 ms pieces, a common setting of Opus, a widely used codec, or program that compresses sound, and filling each piece adds 20 ms. The pieces travel by UDP (User Datagram Protocol), which sends packets without waiting for replies (Module 3). At the 2.6 megabits per second (Mbps) that Zoom gives for sending video in a 720p group call (Module 1), with an assumed 1,100 bytes of video in each packet, the laptop sends about 300 video and 50 audio packets a second, or about 1.2 million an hour. Step 3, the home and its line: Wi-Fi to the router takes about 2-10 ms, and the access network to the provider's first router about 2 ms on fiber and up to about 30 ms on cable or 4G. Step 4, the providers' networks: an exchange or a direct link to the call company's network in the same city adds about 1-5 ms, and the media server, a computer in the call company's data center that forwards each person's streams to the others, another 1-5 ms. Step 5, the ocean: half of Hibernia Express's tested round trip of 58.95 ms, or 29.5 ms (Module 4). Step 6, the far side: London's provider, access network and Wi-Fi add about 5-40 ms, a jitter buffer holds packets for 20-60 ms so that uneven arrivals play smoothly, and decoding and display take 10-20 ms. Total: about 110 ms at best and 260 ms at worst, one way. The ITU's standard G.114 advises keeping one-way delay under about 150 ms for natural conversation, so the call meets it only when the access networks and buffers are quick. The ocean is a ninth to a quarter of the total, so most of the delay is in the two homes and the software. Data: an hour of this call uses 1.98 gigabytes (GB), as Module 1's worked example computes from Zoom's rates, along with what such calls add to a household's month. Zoom's bandwidth page, as read in October 2026, gives 1.2 Mbps each way for a one-to-one call at 720p, or 1.08 GB an hour, and 3.8 Mbps up and 3.0 Mbps down at 1080p, or 3.06 GB.

Module 10 Key Points

The Cloud and Data Centers

• Cloud infrastructure services earned $143.4 billion in the second quarter of 2026, up 43%; Amazon held 28% (Synergy).

• Data centers used about 485 TWh in 2025, about 1.5% of the world’s electricity, by the IEA’s estimate.

• The four largest spenders guided to about $720-745 billion of capital spending for 2026 in July 2026.

• AWS’s outage of October 20, 2025 began when its own automation left a DNS record empty; it lasted 14.5 hours.

• Drones struck AWS data centers in the UAE and Bahrain on March 1-3, 2026; restoring the UAE region would take months.

• Outages per site fell for a fifth year (Uptime Institute), but one large provider’s failure stops thousands of services at once.

Module 10 Quiz

Check Your Understanding

1. By Synergy Research Group's estimate, how much did cloud infrastructure services earn worldwide in the second quarter of 2026?

A) $143.4 billion, up 43%
B) About $500 billion in that one quarter
C) About $100 billion, unchanged from a year earlier as growth stalled
D) $720-745 billion, the sum the four largest companies guided to

2. What caused AWS's outage of October 20, 2025, by AWS's own account?

A) A drone strike on its northern Virginia data centers
B) A race condition in automation that left a DNS record empty
C) A cyberattack
D) A faulty update from CrowdStrike that crashed the servers running its database service

3. What did the Uptime Institute report in May 2026 about outages?

A) They doubled after the cloud outages of 2025
B) They rose for a fifth year
C) Outage rates per site fell for a fifth year
D) Cloud and internet companies caused only a small share of publicly reported outages, far fewer than companies' own server rooms

4. What did AWS tell customers on April 30, 2026 about its region in the UAE?

A) Fully restored in a week
B) That it would close the region for good
C) That only one of its data centers had been hit
D) That it could not reliably support customer applications there, and that restoring it would take several months

5. What is the course's reading on whether concentration in a few clouds has made the internet more fragile?

A) No: the outages of 2025 were fixed in hours
B) Yes: outages are now more frequent
C) Its failures are larger and more correlated, not more frequent
D) Yes, because a single provider now carries most of the world's internet traffic

Module 11 — Who Is Offline

Access, Affordability, and Public Money

Recommended Reading / Reference

Book: Jan van Dijk, *The Digital Divide* (2020), a Dutch communication scientist's synthesis of research on who lacks access, skills and benefits online, and why the divide outlasts the spread of devices

Reference: The ITU's *Facts and Figures 2025* (November 17, 2025), and the *State of Mobile Internet Connectivity 2026* of GSMA, the mobile operators' association (September 2026), the two main annual counts of who is offline and why

11a. Who Is Offline

About 2.2 billion people did not use the internet in 2025, by the ITU's *Facts and Figures 2025*, published on November 17, 2025. The ITU counts as a user anyone who went online, on any device and in any place, in the previous three months, so the offline had not used it at all in that time. They are the 26% of the world's people not counted among the users in Module 1, and 96% of them live in low- and middle-income countries.

Where a person lives decides most of it. In high-income countries 94% of people were online in 2025, against 23% in low-income countries. Europe, the Americas and the former Soviet states had 88-93% of their people online, Asia-Pacific 77% and the Arab States 70%. Africa was the least connected region, at 36%, and the least developed countries, the UN's list of the poorest, stood at 34%. In cities 85% of people were online, against 58% in the countryside, and in low-income countries only 14% of rural people were.

Within countries the gaps follow sex and age. Worldwide 77% of men and 71% of women were online, about 280 million more men than women, and the ITU's parity score, women's rate divided by men's, has stayed at 0.92 since 2019. Young people lead: 82% of those aged 15-24 were online, against 72% of everyone else, and in low-income countries the young were 1.9 times as likely to be online as everyone else there.

The totals move with each edition: when the ITU gathered new survey data it revised its 2024 estimate of the offline population from 2.6 billion to 2.3 billion, the other side of the revision to its count of users that Module 1 describes.

11b. Covered but Not Connected

Most offline people live where a signal already reaches. The GSMA, the mobile operators' association, measures two gaps in its report of September 2026, with data for 2025. Its coverage gap is the people who live outside any mobile broadband network of the third generation (3G) or later: 270 million, or 3% of the world's people. Its usage gap is the people who live under such a network but do not use mobile internet: 3.1 billion, or 38%.

The usage gap is thus about 11 times the coverage gap, and more than 90% of the 3.4 billion people who do not use mobile internet live in covered areas. The coverage gap shrank by about 40 million people in 2025. In Sub-Saharan Africa the usage gap was 66% of the population and the coverage gap 9%, so only 25% of people there used mobile internet. The GSMA counts 4.8 billion people, 59%, as connected on a device they own or mainly use, and another 650 million, 8%, who use mobile internet only on someone else's and so are counted within the usage gap.

The ITU measures coverage differently and finds a larger gap: about 312 million people, 4%, lived without mobile broadband coverage in 2025, almost half of them in Africa. The ITU builds its estimate from network data that operators and regulators report, while the GSMA uses the coverage model of its research arm, GSMA Intelligence, so the gap of 42 million between the two is a matter of method. The ITU's own figures point the same way as the GSMA's: 96% of people were covered by 3G or better but 74% were online (Module 1), leaving about 1.8 billion covered but offline (this course's arithmetic).

Coverage also differs in quality: in low-income countries 4G reached 56% of people in 2025 and only 38% of rural people, and 19% of rural people in the least developed countries had no mobile network at all (ITU). Average download speeds in high-income countries were about 3.5 times those in low- and middle-income countries in 2025 (GSMA). In Sub-Saharan Africa half of mobile subscribers still used 3G phones or feature phones, basic handsets with little or no internet use, against 13% worldwide (GSMA).

11c. Why Covered People Stay Offline

The GSMA asks non-users why, and for people who have never used mobile internet the barriers it ranks highest are the cost of a handset, then poor literacy and digital skills, then fears for safety and security, such as fraud and harassment. For people who use it but then stop, safety comes first, then the cost of data, then a poor connection. Its surveys also ask about relevance, whether there is content and service worth having in a person's language.

Handsets are the first barrier: the cheapest internet-capable handset cost about 15% of average monthly income in low- and middle-income countries in 2025, 44% for the poorest fifth of people, and 76% for the poorest fifth in Sub-Saharan Africa (GSMA). About 82% of people aged 10 and over owned a mobile phone in 2025, but only 53% in low-income countries (ITU). In Africa 66% owned a phone while 36% were online, a gap the ITU puts at 30 points.

In 2026 handsets became more expensive: the GSMA reports that memory-chip prices more than doubled between the third quarter of 2025 and the first quarter of 2026 and rose another 80-90% in the second quarter. It projects that smartphone shipments will fall 14% in 2026, and those of phones under $100, the ones most first-time users can afford, 36%.

Even those online use less: a typical user in a high-income country generates nearly 8 times the mobile data of one in a low-income country (ITU). In low- and middle-income countries women are 12% less likely than men to use mobile internet, about 200 million fewer women, and rural adults are 28% less likely than urban ones (GSMA).

11d. What It Costs

Affordability has an agreed yardstick: in 2018 the Broadband Commission for Sustainable Development, a UN-backed body, set a target that entry-level broadband in low- and middle-income countries should cost less than 2% of monthly gross national income (GNI) per person by 2025, replacing an earlier 5% target. GNI per person is a country's total income divided by its population, so the target is a share of average income, not of a poor person's. The Alliance for Affordable Internet, a campaign group, states the same idea as "1 for 2": 1 gigabyte (GB) of mobile data for 2% of average monthly income.

By the ITU's count for 2025, the median price of an entry-level mobile data plan worldwide fell to 1.4% of monthly GNI per person, from 1.5%, while entry-level fixed broadband held at 2.5%. Mobile plans met the 2% target in 130 of 205 economies and fixed plans in 88 of 195, so about one economy in three still missed

it for mobile. Entry-level mobile broadband prices fell by about half between 2013 and 2025, after adjusting for inflation and local price levels (ITU, November 17, 2025).

The poorest countries are far from the target (table). In nine of ten low-income economies 5 GB of mobile data cost more than 10% of monthly income, and for the poorest 40% of earners in low-income and least developed countries the burden is about 20%. Users in low-income countries spend about 22 times the share of income on mobile broadband that users in high-income countries spend, and fixed broadband, where it is offered there, costs more than a quarter of average income (ITU). Spending on phones and internet takes 2-3% of household budgets in rich countries and up to about 9% in poorer ones, by the ITU's *Global Connectivity Report 2025*.

Mobile data allowance	Median price in low-income countries, share of monthly GNI per person	Times the 2% target	That share of $97.92 a month, the World Bank's low-income ceiling (illustrative)
1 GB a month	5.5%	2.8	$5.39
2 GB a month	7.4%	3.7	$7.25
5 GB a month	12.3%	6.2	$12.04
10 GB a month	20.4%	10.2	$19.98
The 2% target	2.0%	1.0	$1.96

ITU, affordability report on 2024 data (published 2025); medians across low-income countries. The last column applies each share to the World Bank's ceiling for a low-income country, $1,175 of GNI per person a year in its classification of July 1, 2026; it is this course's arithmetic, not a measured price.

The GSMA's measure looks better but points the same way. Across low- and middle-income countries 1 GB cost 1.1% of average monthly income in 2025, a five-year low, yet 47 of the 133 such countries, 35%, had still not met the 2% target. For the poorest 40% of people, 80 of the 133, or 60%, missed it, and a 5 GB plan cost a median 1.8% of income overall but 6.3% for the poorest fifth.

11e. America: Covered but Not Subscribed

The United States shows the same gap at a higher income. By the FCC's 2026 report on broadband deployment, 96.9% of Americans could buy fixed broadband of at least 100 megabits per second (Mbps) down and 20 Mbps up delivered over land, by wire or fixed wireless rather than by satellite, in data to mid-2025 (as reported in the trade press on August 17, 2026). Adoption follows income instead: in its survey for 2025 the Pew Research Center found that 54% of adults in households earning under $30,000 a year had home broadband, down from 57% in 2023 and 2024, against 94% of those earning $100,000 or more (as reported by Telecompetitor, January 23, 2026). About a third of those lower-income adults, 34%, relied on a smartphone with no home broadband, against 4% of those earning $100,000 or more.

For about two and a half years the federal government paid part of the bill. The Affordable Connectivity Program (ACP), funded with $14.2 billion, gave eligible low-income households $30 a month toward internet service, or $75 on tribal lands, and a one-time discount of up to $100 on a device. More than 23 million households enrolled. Congress did not add money, and the program ended on June 1, 2024 when its funds ran out.

Its end shows what price does: the Brattle Group, an economics consultancy, estimated that about 5 million households dropped home internet altogether and many more moved to cheaper, slower plans (as reported by Broadband Breakfast, May 31, 2025). In a survey by the National Lifeline Association, a trade group of providers that serve low-income households, nearly 40% of former ACP households said they had cut back

on food to pay for internet. Proposals to revive the program failed, among them an amendment blocked in Congress on May 14, 2025.

What remains is patchy: New York requires providers to offer a $15 plan to low-income households, and California's utilities commission made a $20 low-income plan a condition of Charter's purchase of Cox (Module 6). The FCC's Lifeline program, paid for by the Universal Service Fund (Module 12), still gives qualifying households $9.25 a month toward a broadband plan.

11f. $42.45 Billion for Coverage

The largest public program pays for networks, not bills. The Infrastructure Investment and Jobs Act of November 2021 gave the Broadband Equity, Access, and Deployment (BEAD) program $42.45 billion, run by NTIA, to build networks to homes and businesses that lack adequate service. Unserved locations, with less than 25 Mbps down and 3 Mbps up, came first, then underserved ones, below 100 Mbps and 20 Mbps. A location is an address on the FCC's broadband map that could take service, and each of the 56 states and territories draws up a plan that NTIA must approve.

NTIA rewrote the rules on June 6, 2025. Its restructuring dropped the earlier preference for fiber and made the program neutral between technologies, so that any network delivering at least 100 Mbps down and 20 Mbps up with a delay of 100 milliseconds (ms) or less could count as a priority project. It ordered a new round of bidding, called "Benefit of the Bargain," in which the bid needing the least BEAD money generally won, with other factors weighed only among bids within 15% of the cheapest. It also dropped requirements on climate resilience and labor, and it let providers set the price of the low-cost plan they must offer poorer households.

NTIA approved the last of the 56 plans, Illinois's, on August 25, 2026 (Fierce Network). By Broadband Breakfast's tally of NTIA's award data on August 28, 2026, the plans spend $18.19 billion to reach 3.79 million locations, and the table gives the split by technology. Satellite's share of the money is far smaller than its share of the locations, because a satellite award pays a provider to reserve capacity rather than to build.

Technology	Share of BEAD locations, all 56 approved plans	Average bid per location in Tennessee's Benefit of the Bargain round
Fiber	62.9%	$4,903
Low-orbit satellite	21.9%	$886
Fixed wireless access, by radio from a nearby tower to an antenna at the home	13.3%	$1,289
Cable (fiber and coaxial)	1.9%	$6,305

Shares of locations: Broadband Breakfast's tally of NTIA award data, August 28, 2026; dollar shares by technology were not published in the sources read. Tennessee: Broadband Breakfast's analysis of the bids, July 30, 2025, before awards were made.

The restructuring changed the states' plans. Louisiana's revised plan, which NTIA approved in November 2025, covers 127,108 locations for about $498.2 million of its $1.355 billion allocation, with fiber at 81% of the locations against 95% in its first plan (Wireless Estimator, February 11, 2026). In Illinois, fiber's share fell from 74.8% in December 2025 to 68.1% in the approved plan of August 2026, while fixed wireless rose from 9.8% to 23.6% and satellite fell from 15.1% to 8.1% (Fierce Network).

The satellite share is shrinking: in July 2026 NTIA asked states to remove up to half of their satellite locations, saying the FCC's newest map showed many were already served by a wired or fixed wireless network or were not buildings that need broadband, such as sheds (Broadband Breakfast, July 14, 2026).

Alex Karras and Michael Santorelli of New York Law School estimated that 35-42% of the roughly 900,000 satellite locations could go, cutting satellite grants from more than $1 billion to $602-671 million (Broadband Breakfast, July 16, 2026). In draft plans SpaceX was set to receive nearly $739 million and Amazon Leo, Amazon's low-orbit service (Module 9), more than $312 million. These awards are separate from the $885.5 million that the FCC provisionally awarded SpaceX in its 2020 Rural Digital Opportunity Fund auction, which the FCC rejected in August 2022 (Module 9). The argument over whether satellites or fiber should serve the last homes is in Module 9.

Less than half of the $42.45 billion went to construction awards. NTIA put the money left after approving the plans at about $21 billion, a figure it restated in its memo of September 3, 2026, while Broadband Breakfast's own count of August 28 was about $22.5 billion. The law allows non-deployment uses such as digital-skills training and devices, and the states want to decide them. Howard Lutnick, the Secretary of Commerce, told senators that returning the money to the Treasury "is not the plan" (reported February 11, 2026). Two federal decisions narrowed the states' choices. An executive order signed by President Trump on December 11, 2025 told the Commerce Department to withhold non-deployment money, as far as federal law allows, from states with "onerous" AI laws. NTIA's memo of September 3 sent part of the savings to a new round of bidding for rural locations still unserved after defaults, misreported coverage or map changes, a process it said could take about nine months. Guidance on other uses, first expected in March 2026 and then promised for summer 2026, was still pending in October 2026, and Gigi Sohn, executive director of the American Association for Public Broadband, called the September memo "partial guidance."

11g. Is the Divide Now About Cost and Skills, Not Coverage? The Cases as Their Advocates Make Them

The case that cost and skills now matter most: the GSMA, the ITU's work on affordability and the Alliance for Affordable Internet argue that the networks are mostly built and the task now is getting people onto them. Far more people live under a signal without using it than live beyond one (section 11b), and the barriers non-users name first are handsets, skills and safety. Mobile data still misses the 2% target in about one economy in three, and the GSMA estimates that closing the usage gap could add $3.5 trillion to the economic output of low- and middle-income countries between 2023 and 2030. In America they point to near-universal fixed coverage beside low adoption among the poor, and to the households that dropped service when the ACP ended.

The case that coverage still matters: advocates for rural and low-income areas argue, citing the ITU's own coverage figures, that a signal on a map is not service a person can use. In low-income countries most rural people have no 4G, and in the least developed countries many rural people have no mobile network at all (section 11b). Hundreds of millions of people, almost half of them in Africa, lack mobile broadband entirely, and speeds in poorer countries are a fraction of those in rich ones. In America, BEAD exists because millions of homes and businesses still lacked adequate service (section 11f).

What the evidence says: the two sides are counting different people. Counted by heads, the evidence favors cost and skills: people under a signal who do not use it outnumber those beyond any signal by about 6 to 1 on the ITU's figures and about 11 to 1 on the GSMA's (section 11b). Real prices have halved since 2013 while use has grown (section 11d). The end of the ACP gives a test of price: in Pew's survey for 2025, the first full year after it ended, adoption among low-income adults was 3 points lower than in 2024, while among the highest earners it was 2 points higher. Counted by places, coverage still binds, because the uncovered are concentrated in rural Africa and the least developed countries, where the networks that exist

are often 3G. The GSMA's barrier rankings come from what non-users say, which shows their stated reasons rather than what would change their behavior.

The course's reading, stated once: for the world as a whole, the divide is now mainly about cost and skills, by a margin of several to one in the number of people affected. Coverage of usable quality still binds for the rural poor in Africa and the least developed countries, so the two explanations describe different people rather than compete. Money spent on networks where a signal already exists reaches few new users, while the price of handsets and data and the teaching of skills decide whether the billions under a signal come online. Whether satellites or fiber should reach the last uncovered homes is a separate question (Module 9).

Worked example — BEAD's Price per Location, and a Starlink Household's Bill

Step 1, the national average: $18.19 billion ÷ 3.79 million locations = $4,799 of BEAD money per location (section 11f). Step 2, fiber against satellite in one state: in Tennessee's bidding the average fiber bid of $4,903 was 5.53 times the average low-orbit satellite bid of $886, and fixed wireless, at $1,289, was 1.45 times it (table, section 11f). Step 3, two extremes: had every location gone to satellite at Tennessee's price, BEAD would have spent 3.79 million × $886 = $3.36 billion, and had every one gone to fiber at Tennessee's price, 3.79 million × $4,903 = $18.58 billion. Step 4, fiber's implied national price, an illustration rather than a published figure: if satellite, fixed wireless and cable cost what they did in Tennessee, the national mix implies fiber at ($4,799 − 21.9% × $886 − 13.3% × $1,289 − 1.9% × $6,305) ÷ 62.9% = $6,859 per location. Step 5, the household's side: BEAD pays for the network, but the household pays the monthly bill. Starlink's American residential plans were listed at $55, $85 and $130 a month by speed tier from June 2026, before taxes and equipment, up from $50, $80 and $120 (Benton Institute, May 18, 2026; satelliteinternet.com, May 28, 2026), so ten years of service cost $6,600, $10,200 or $15,600. Tennessee's $886 satellite bid is 8.7% of the ten-year bill on the $85 plan. Step 6, a ten-year total: assume, for illustration, that a fiber household pays $75 a month, close to the average revenue per home-broadband customer in Module 6, or $9,000 over ten years. Subsidy plus bills then come to $4,903 + $9,000 = $13,903 for fiber, against $886 + $10,200 = $11,086 for satellite on the $85 plan and $16,486 on the $130 plan. Satellite stays cheaper over ten years unless its household pays more than ($13,903 − $886) ÷ 120 months = $108.48 a month. The result: satellite's up-front cost is 5.5 times lower, but over ten years fiber's total is 25% higher than satellite's on the $85 plan, while satellite's is 19% higher than fiber's on the $130 plan. The comparison turns on the household's bill, and fiber's case rests on the decades after that, the subject of Module 9.

Module 11 Key Points

Who Is Offline

• About 2.2 billion people were offline in 2025, 96% of them in low- and middle-income countries (ITU).

• GSMA counts 3.1 billion people under a mobile broadband signal who do not use mobile internet, about 11 times the uncovered.

• The median entry-level mobile data plan cost 1.4% of monthly income per person worldwide in 2025 (ITU).

• The Affordable Connectivity Program ended on June 1, 2024; low-income home broadband adoption fell to 54% in 2025 (Pew).

• BEAD’s approved plans spend $18.19 billion on 3.79 million locations; about $21 billion of its $42.45 billion is left over (NTIA).

• Tennessee’s fiber bids cost 5.5 times satellite’s per location, but ten years of household bills can reverse the comparison.

Module 11 Quiz

Check Your Understanding

1. How does the GSMA's usage gap compare with its coverage gap for 2025?

A) The usage gap, 3.1 billion people, is about 11 times the coverage gap.

B) The two gaps are about equal, at roughly 1.5 billion people each.

C) The coverage gap is larger, because 3.1 billion people live where no mobile broadband signal reaches them.

D) The coverage gap is about 312 million and the usage gap about 270 million, so the two are now nearly the same size.

2. By the ITU's count for 2025, what did the median entry-level mobile data plan cost worldwide, as a share of monthly income per person?

A) About 2.5%

B) 1.4%

C) 7.4% of income

D) Just over 20%

3. What happened to the Affordable Connectivity Program?

A) Congress made it permanent in 2025 at $30 a month.

B) It ran out of money and ended on June 1, 2024.

C) It became part of BEAD.

D) It was folded into the Lifeline program, which now pays $30 a month to 23 million households.

4. Under NTIA's restructuring of June 6, 2025, which bid generally won a BEAD award?

A) The one using fiber wherever fiber could be built

B) The fastest network, whatever its cost to build

C) The one offering the most jobs and the strictest labor standards, as judged by each state

D) The one needing the least BEAD money

5. What is the course's reading on whether the digital divide is now mainly about cost and skills rather than coverage?

A) Mainly yes for the world as a whole, but coverage of usable quality still binds for the rural poor

B) No: coverage remains the main gap

C) Yes: coverage is now universal

D) Neither can be judged, because the ITU and GSMA disagree

Module 12 — Who Governs the Internet

Standards Bodies, Regulators, and the Rules

Recommended Reading / Reference

Book: Laura DeNardis, *The Global War for Internet Governance* (2014), a scholar's account of who controls the internet's names, numbers, standards and points of control, and why each has become a contest between companies, engineers and states

Reference: The FCC's quarterly public notices on the universal service contribution factor and the filings of the Universal Service Administrative Company, the source of the fund's figures here, and the ITU's pages for its 2026 conference in Doha

12a. No One Runs It, but Someone Governs Each Part

No government or company runs the internet as a whole, but someone sets the rules for each of its parts. Private companies own most of the cables, towers and data centers, and national regulators license them (Module 1 maps who owns each layer). The shared core that lets separate networks work as one is coordinated by a few nonprofit bodies, most of them open to anyone who wants to take part. Their supporters call this model "multistakeholder": companies, engineers, civil-society groups and governments share the decisions, and no one of them controls the outcome.

Two computers cannot share an address and two websites cannot share a name, so someone must keep the master lists. That job belongs to the Internet Assigned Numbers Authority (IANA), which maintains the root zone, the top of the DNS that the root servers publish (Module 3), along with the global pools of IP addresses and the registries of numbers that protocols use. Since 2016 IANA has been run by Public Technical Identifiers (PTI), an affiliate of the Internet Corporation for Assigned Names and Numbers (ICANN), a California nonprofit set up in 1998 (Module 2).

ICANN makes the policy for domain names: it decides which top-level domains exist, from .com to the newer endings added since 2013, and it accredits the companies that sell names. Its policies come from councils of registries, registrars, network operators, businesses and civil-society groups, and governments advise its board but do not vote.

Until 2016 the US government held the last word. ICANN performed the IANA functions under a contract with NTIA, which signed off on every change to the root zone. NTIA announced on March 14, 2014 that it would give up that role if the community designed a replacement that kept the system stable and open. It let the contract expire on September 30, 2016, and from October 1, 2016 the functions passed to PTI, overseen by the communities that use the names, numbers and protocol registries.

IP addresses reach networks through five regional internet registries (RIRs), member-owned nonprofits that allocate blocks of addresses to ISPs and other organizations in their regions. The African Network Information Centre (AFRINIC) serves Africa, and the other four serve Asia and the Pacific, North America, Latin America and the Caribbean, and Europe with the Middle East and Central Asia. Each sets its address policies in open meetings of its members, mostly network operators.

AFRINIC, based in Mauritius, shows what happens when that self-governance breaks down. In 2020 it accused Cloud Innovation, a company holding 6.2 million of its addresses, of breaking its membership agreement, and years of lawsuits followed. Disputed elections left AFRINIC without a valid board from 2022, and a court placed it in receivership in September 2023. A board election in June 2025 was annulled over

allegations of forged proxy votes, and on July 16, 2025 ICANN warned that it might reassess AFRINIC's status as a registry.

The repair came from the members, with a court and a government standing by. A re-run election in September 2025 filled eight board seats, after the government of Mauritius had ordered a Supreme Court judge to investigate AFRINIC and Cloud Innovation had petitioned to wind it up. In May 2026 the Supreme Court of Mauritius let ICANN join that case to argue that the addresses AFRINIC hands out "are not assets of AFRINIC", and the litigation was still going on in September 2026 (court filings and press reports).

12b. How Standards Are Made

The protocols that every network speaks are written by the Internet Engineering Task Force (IETF), which has no members, no votes and no fee to join. Anyone may join its mailing lists and working groups, which also meet in person three times a year. Since its first meeting, in January 1986, it has published its work in the Requests for Comments (RFCs), a numbered series begun on the ARPANET in 1969, some of which become internet standards, such as the encryption that protects web traffic (Module 3).

Its decision rule comes from a 1992 talk by David Clark, a computer scientist at the Massachusetts Institute of Technology: "We reject: kings, presidents and voting. We believe in: rough consensus and running code." A working group's chair judges whether objections have been answered, not whether a majority agrees. A proposal is expected to work in real software before it becomes a standard. Companies employ most participants, but people take part as individuals, so a company cannot buy a seat or a vote.

The World Wide Web Consortium (W3C), which Tim Berners-Lee founded (Module 2), sets standards for web pages and browsers, among them the accessibility guidelines that many governments require public websites to meet. Unlike the IETF, it charges organizations to join, and its standards pass through public drafts and review by members. It has been an independent nonprofit since January 2023.

Mobile standards come from the 3rd Generation Partnership Project (3GPP), formed in December 1998 by standards bodies from Europe, the US, Japan and South Korea and later joined by China's and India's, for seven in all. Its members, mostly equipment makers and carriers, agree on proposals by consensus in technical groups. It issues its work in numbered releases, each a fixed set of features that phones and networks can be built to (Module 7 gives the timetable for 6G). The ITU sets the minimum performance a technology must reach to count as a new generation (Module 7).

Open processes keep any one government from controlling the result, but they favor those who can pay for engineers' time and travel, so large companies and rich countries write most standards. Governments have a formal vote only at the ITU, one reason some want it to have a larger role (section 12g).

12c. The ITU, the UN, and a Forum Made Permanent

The International Telecommunication Union, which began as the International Telegraph Union (Module 2), is now the UN's agency for telecommunications, with 194 member states, each with a vote. Its radiocommunication sector runs the Radio Regulations, the treaty that divides spectrum and satellite orbits among countries (Module 8). Its standardization sector writes telecom standards, and its development sector runs programs in poorer countries and publishes the counts of who is online (Module 11).

Its supreme body, the Plenipotentiary Conference, meets every four years to set its course and elect its leaders, from the Secretary-General and the directors of its three bureaus to the 48 members of its Council and the 12 of its Radio Regulations Board. The 2026 conference meets in Doha, Qatar, from November 9 to

27, 2026, and its secret ballots begin on November 12 with the Secretary-General. Doreen Bogdan-Martin of the US, Secretary-General since January 2023, is the only candidate for her post, while the races for deputy and for the radiocommunication and development bureaus are contested (ITU, October 15, 2026). In 2022 she had beaten Rashid Ismailov of Russia with 139 of 172 votes.

Proposals before the conference cover space and orbits, submarine cables and network resilience, and the International Telecommunication Regulations, a treaty on how international traffic is exchanged and paid for (ITU briefing, October 9, 2026). The treaty's last revision, in Dubai in December 2012, split the membership over whether it should reach internet policy: 89 countries signed and 55 did not, among them the US, the UK, Canada and Australia (European Digital Rights, a campaign group, December 19, 2012).

The UN's main forum for the internet is the Internet Governance Forum (IGF), set up after the World Summit on the Information Society (WSIS) of 2003 and 2005. It negotiates nothing and binds no one: governments, companies, engineers and campaigners meet once a year to discuss policy. Its mandate needed renewal at intervals until December 17, 2025, when the UN General Assembly, reviewing the summit 20 years on (WSIS+20), made it a permanent UN forum. The text passed without a vote after governments had objected to 55 of its 127 paragraphs, by the account of DENIC, the German domain registry.

Russia and some developing countries had sought a "governmental track" in the IGF that would negotiate recommendations, and the final text only "encourages government dialogue with all stakeholders participating." The US objected to references to the UN's Sustainable Development Goals, gender and climate, and Russia and Iran to references to human rights. The next IGF meets in Nairobi, Kenya, on December 14-18, 2026, the first since it became permanent (Media Rights Agenda, June 18, 2026), and the next WSIS review is set for 2035.

12d. The FCC under Brendan Carr, and Net Neutrality

In the US, the Federal Communications Commission licenses spectrum and satellites, regulates telephone and cable companies, reviews their mergers and runs the Universal Service Fund (section 12e). Its five commissioners are nominated by the president and confirmed by the Senate, with no more than three from one party. Brendan Carr, a commissioner since 2017, has been chairman since President Trump designated him on January 20, 2025. In September 2026 only three seats were filled: besides Carr, Olivia Trusty, a Republican confirmed on June 18, 2025, and Anna Gomez, a Democrat whose term expired on June 30, 2026 and who may stay until the end of 2027 unless a successor is confirmed. A nominee for a vacant seat, Danielle Thumann Severs, had had no Senate hearing by September 25 (Broadband Breakfast).

Carr's program has three parts: cutting rules, opening spectrum for auction (Module 8) and national security (Module 13). Under a review he named "Delete, Delete, Delete", the FCC had removed or proposed removing 1,108 rules, 134,928 words and 312 pages of the Code of Federal Regulations by January 14, 2026, and it had closed a record 2,048 inactive proceedings, by Carr's account to Congress. It eased the retirement of copper telephone lines and sped up access to utility poles for new cables. It also investigated the diversity programs of Verizon, Comcast and other companies (Broadband Breakfast, April 29, 2025).

The FCC's distance from the White House has narrowed. An executive order of February 18, 2025 made independent regulators, the FCC among them, submit significant rules to the White House for review. On June 29, 2026, in Trump v. Slaughter, the Supreme Court ruled 6-3 that the president may remove members of the Federal Trade Commission at will, overruling a 1935 precedent. The ruling did not name the FCC, but its test, whether an agency exercises executive power, appears to reach one that grants licenses and enforces rules, though the Court left the limits of its holding open.

Net neutrality is the principle that an ISP should carry all lawful traffic alike, without blocking or slowing particular sites or charging them for faster delivery. The rules have changed with each change of party, because the law leaves open whether broadband is a “telecommunications service”, which the FCC may regulate much as it regulates telephone lines, or an “information service”, which it largely may not. The FCC imposed rules in 2015, repealed them in 2017 and restored them in 2024.

On January 2, 2025, the Sixth Circuit Court of Appeals struck down the 2024 rules in Ohio Telecom Association v. FCC, holding that broadband is an information service, so the FCC had no authority to impose them. It relied on the Supreme Court’s Loper Bright decision of 2024, which ended courts’ deference to an agency’s reading of an unclear law. Carr told Congress in January 2026 that “a federal court invalidated the prior FCC’s plan to expand government control of the Internet”. The rules that remain are the states’: California’s and Washington’s laws are still in force (Perkins Coie, a law firm, 2025).

The evidence on what such rules do is mixed. C. A. Hooton found that the US rule changes of 2010, 2015 and 2017 had no measurable effect on the investment of listed telecom companies (*Telecommunications Policy*, 2019). The Phoenix Center, an advocacy research group, claimed in 2017 that the 2015 rules cut the industry’s annual investment by about a quarter, roughly $30 billion a year. A 2022 study of 32 countries of the Organisation for Economic Co-operation and Development, 2000 to 2021, by Wolfgang Briglauer and three co-authors found “a significant and strong negative impact on fiber investments”. None of the three measured what the rules did to what consumers pay or receive.

12e. The Universal Service Fund

The Universal Service Fund pays to extend telephone and broadband service where the market alone would not. The 1996 Telecommunications Act (Module 2) wrote the goal into law, and a nonprofit, the Universal Service Administrative Company (USAC), runs the fund’s four programs under the FCC’s direction (table).

Program	What it pays for	Requirement for October to December 2026
High Cost	Support for carriers serving rural and remote areas, now mostly for broadband	$1,078.77 million
Lifeline	Monthly discounts on phone or broadband service for low-income households	$186.79 million
E-Rate (schools and libraries)	Discounts on connections and internal networks for schools and libraries	$696.85 million, less $69 million of unused funds
Rural Health Care	Connections for rural hospitals and clinics	$188.78 million, less $56 million of unused funds

USAC’s projections, filed July 31, 2026; each includes administrative costs and corrections for earlier quarters. The FCC applied the unused funds (Benton Institute, September 14, 2026).

The money comes from telecommunications companies, which pay a share of their interstate and international revenue from end users, a share the FCC calls the contribution factor and resets every quarter. Most pass the charge to customers as a line on the phone bill. Revenue from broadband internet access is not assessed, so as customers moved from long-distance calls to data, the base the factor applies to shrank, while the programs, which now pay mostly for broadband, did not.

The FCC proposed a factor of 38.8% for July to September 2026 and, on September 14, one of 42.0% for October to December. By its practice a proposed factor takes effect unless the Commission acts within 14 days of its release (FCC public notice, June 12, 2026). Proposals to widen the base, to broadband subscriptions or to large online companies, have been debated for years, and Carr argued in 2021, as a

commissioner, that large technology companies should pay in. For the fourth quarter of 2026 the base was still telecommunications revenue alone.

The fund survived a constitutional challenge. Consumers' Research, an advocacy group, argued that Congress had handed the FCC, and the FCC had handed USAC, an unlimited power to tax. On June 27, 2025, in FCC v. Consumers' Research, the Supreme Court upheld the scheme 6-3, reversing the Fifth Circuit Court of Appeals, which had struck it down in July 2024. Justice Elena Kagan's opinion held that Congress had given the FCC enough guidance and that the FCC, not USAC, makes the decisions. Justice Neil Gorsuch, joined by Justices Clarence Thomas and Samuel Alito, dissented.

12f. Europe's Network Rules, and the Fight over Data

The EU sets most of its telecom rules in common, and on January 21, 2026 the European Commission proposed to rewrite them as a Digital Networks Act (DNA). It would replace the European Electronic Communications Code, a 2018 directive that each member state wrote into its own law, with a regulation that applies directly in all 27. It would also absorb the rules for the Body of European Regulators for Electronic Communications (BEREC), the EU's spectrum program and parts of its net-neutrality regulation. Payments from large content companies to network operators were left out as a mandatory fee (Module 4). Its main parts, as summarized by the law firm Gleiss Lutz, are these:

- A single passport: an operator could serve the whole EU after notifying one member state, which would confirm within a week.
- Longer spectrum licenses, of about 40 years; summaries of the text differ on whether that is a minimum or a maximum.
- One EU-wide authorization for satellite services, replacing national ones within 36 months.
- An end to copper: national plans for fiber due by October 31, 2029, and copper networks closed across the EU by December 31, 2035, earlier where fiber covers at least 95% of premises.
- A preparedness plan for networks in a crisis, which BEREC would draw up within 12 months of the act taking effect.

The proposal needs the approval of the European Parliament and of the Council, where member states sit. A Council progress report of June 2, 2026 recorded concerns that the single passport would invite "forum shopping", in which operators choose the most lenient regulator, and about leaving many details to later rules and BEREC guidelines, and it set no date for a common position. In Parliament, Michał Kobosko, a Polish member of the Renew group, leads the work for the industry committee.

A narrower law already applies. The Gigabit Infrastructure Act, in general application since November 12, 2025, makes authorities decide on permits for fiber and 5G within four months and makes network owners and public bodies share their ducts and poles on fair terms (Arthur Cox, a law firm, September 15, 2025). New and renovated buildings must be ready for fiber from February 12, 2026 (Digital Policy Alert).

The other contest is over data. Data localization laws require that data about a country's people, or certain kinds of data, be stored or processed inside the country. Russia has required since 2015 that personal data on its citizens be kept on servers in Russia, and China's Cybersecurity Law of 2017 requires operators of "critical information infrastructure" to store personal data in China. The EU's General Data Protection Regulation (GDPR), applied since 2018, takes a different route: personal data may leave the EU only for countries, or under contracts, that protect it. Freedom House, a US-based monitor, wrote in *Freedom on the Net 2025* that "expanding data localization requirements are contributing to the global decline of internet freedom", a decline it recorded for a 15th consecutive year.

Trade rules on data weakened in 2026. The World Trade Organization's pledge of 1998 against customs duties on electronic transmissions, such as downloads, lapsed when its ministerial conference in Yaoundé, Cameroon, ended on March 30, 2026 without renewing it (Bloomberg). Sixty-six members covering about 70% of world trade pledged on their own not to impose such duties, but India and the US did not sign (Deloitte, April 1, 2026). From May 8, 2026 the US and 18 other members, among them Japan and South Korea, made the same pledge among themselves (fibre2fashion, May 14, 2026).

12g. Multistakeholder Governance or a Larger Role for States? The Cases as Their Advocates Make Them

The case for the multistakeholder model: its advocates, among them the Internet Society (ISOC), ICANN and the US and EU governments, argue that the people who build and run networks are best placed to coordinate them, and that the record shows it. The root zone and the address registries have run without a government contract since 2016, and open standards let any network join without asking permission. Giving states votes over the core, they argue, would hand governments that censor at home a say over the network everyone uses. They point to WSIS+20, which reaffirmed the model by consensus, and to the ITU election of 2022 (section 12c).

The case for a larger role for states: its advocates, among them Russia, China, Iran and many of the developing countries of the Group of 77 (G77), argue that governments answer to their citizens for security, crime and public policy, while companies and engineers answer to no electorate. The open forums, they say, are dominated by Western companies and technical communities that poorer countries cannot afford to staff. In the WSIS+20 talks the G77 and China called a new financing mechanism for digital development a red line. States, they note, can make binding rules: the UN Convention against Cybercrime, whose negotiation Russia proposed, opened for signature in Hanoi on October 25, 2025. Critics of the model add AFRINIC's crisis as proof that self-governance can fail.

What the evidence says: in the forums, the multistakeholder side prevailed in 2025-26 (section 12c). The state-made instrument has been slow. The cybercrime convention had 79 signatories by July 2026, the US not among them, but only three ratifications by May 2026, against the 40 it needs before it can take effect 90 days later. Rights groups and companies oppose it: the Cybersecurity Tech Accord, whose members include Microsoft, Meta and Cisco, called it a "surveillance treaty". AFRINIC's case cuts both ways: a court and a government stepped in, but the repair was an election by its members (section 12a). Inside borders the trend runs toward state control: shutdowns reached a record in 2025 (Module 13), and data localization is spreading. No study read for this course compares outcomes under the two models.

The course's reading, stated once: the question has a different answer at each layer. For the shared core of names, numbers and protocols, the multistakeholder model has worked: the core stayed single and stable through the 2016 handover and the AFRINIC crisis, and no state-made body has offered a working replacement. For what crosses each country's borders, states already govern, and no forum will take that power back. The model should be judged on keeping the core open and single, which it has done, not on protecting users from their own governments, which it cannot do. A larger role for the ITU would add votes over the layer that works and none over the layer where users lose access.

Worked example — How the 42.0% Is Set, and How It Reaches a Phone Bill

The FCC sets the factor so that one quarter's contributions pay for one quarter's programs. The programs' requirements for October to December 2026 are in the table in section 12e, and USAC projected $6,909,924,575 of interstate and international end-user telecommunications revenue from 3,191 contributors (filing of September 1, 2026). The method is the one set out in the FCC's notice for the third quarter (June 12, 2026). Step 1: $1,078.77 million + $186.79 million + $696.85 million + $188.78 million = $2,151.19 million, less $125 million of unused funds, leaves a requirement of $2,026.19 million. Step 2: the FCC subtracts the requirement from the revenue, because the revenue carriers report includes the charges they bill to recover their contributions, and then takes off 1% for contributions it does not expect to collect: ($6,909.92 million − $2,026.19 million) × 0.99 = $4,834.90 million. Step 3: the factor is $2,026.19 million ÷ $4,834.90 million = 41.91%, which the FCC rounds up to the next tenth of a point, 42.0%, the proposed figure. The same arithmetic on the third quarter's published inputs gives 38.77%, rounded up to 38.8%. Between the two quarters the requirement rose 1.2% while the revenue base fell 4.3%, and the factor rose 3.2 points. Step 4: if the base fell by another 4.3% and spending held, the factor would reach 44.7% (an illustration, not a forecast). Step 5: the bill. Take a mobile plan with $40 a month of voice and text charges, an illustrative assumption, since data is not assessed. The FCC lets mobile carriers treat 37.1% of such revenue as interstate, a "safe harbor" it set in 2006, so $14.84 is assessable. A carrier may not bill more than the factor times that amount, so the line on the bill can be at most 42.0% × $14.84 = $6.23 a month, or $74.79 a year, 15.6% on top of the $40. Step 6: four quarters like this one would come to $8.10 billion a year, a rough scale for the fund rather than a budget.

Module 12 Key Points

Who Governs the Internet

• US control of the internet’s master lists of names and numbers ended on October 1, 2016; a multistakeholder community now oversees them.

• The UN made the Internet Governance Forum permanent on December 17, 2025; it next meets in Nairobi on December 14-18, 2026.

• The ITU’s conference in Doha, November 9-27, 2026, elects its leaders; Doreen Bogdan-Martin is the only candidate for Secretary-General.

• A US appeals court struck down federal net-neutrality rules on January 2, 2025; state laws such as California’s still apply.

• The Universal Service Fund’s proposed charge reached 42.0% for late 2026 as its revenue base shrank; the Supreme Court upheld the fund 6-3.

• The EU’s Digital Networks Act, proposed January 21, 2026, would close copper networks by 2035 and let operators serve the EU on one notification.

Module 12 Quiz

Check Your Understanding

1. What changed on October 1, 2016?

A) ICANN was founded to take over the domain name system
B) The five regional registries were set up
C) The IANA functions passed to PTI, ending US government oversight
D) The ITU took charge of the root zone and the address pools under a treaty its member states signed in Dubai

2. What did the Supreme Court decide in FCC v. Consumers' Research on June 27, 2025?

A) It struck down the fund as an unlimited tax that only Congress may levy.
B) It upheld the fund's contribution scheme, 6-3.
C) It ordered the FCC to add broadband revenue to the fund's base.
D) It ruled that the president may remove FCC commissioners at will, overruling a 1935 precedent.

3. On what ground did the Sixth Circuit strike down the FCC's 2024 net-neutrality rules?

A) Broadband is an information service, so the FCC lacked authority
B) The rules cut fiber investment
C) Congress had repealed them
D) States had their own laws

4. Why did the Universal Service Fund's factor rise to a proposed 42.0% for late 2026?

A) Lifeline benefits were doubled
B) Broadband was added to the base
C) A court ruling
D) Its revenue base shrank

5. What is the course's reading on multistakeholder governance against a larger role for states?

A) States should govern every layer of the internet, since only states can make rules that bind everyone.
B) The ITU should take over the root zone so that every state gets a vote.
C) The multistakeholder model works for the core, but states already control their borders.
D) Neither model works, so the internet needs a new treaty body.

Module 13 — Networks Under Attack

Hacks, Shutdowns, and the Splinternet

Recommended Reading / Reference

Book: Andy Greenberg, *Sandworm: A New Era of Cyberwar and the Hunt for the Kremlin's Most Dangerous Hackers* (2019), a reporter's account of the Russian military hackers who cut power in Ukraine and released NotPetya, the attack the White House called the most costly in history

Reference: Access Now's annual #KeepItOn reports and Cloudflare Radar's quarterly summaries of internet disruptions, the main counts of shutdowns used here

13a. Inside the Carriers: Salt Typhoon and Volt Typhoon

The most damaging attacks on networks in 2024-26 did not cut them. The attackers lived inside the routers and management systems of the carriers themselves, often for years. Microsoft gives the hacking groups it links to China the family name "Typhoon", and two of them matter most here. Salt Typhoon spied on telephone carriers, and Volt Typhoon, by the US government's account, placed itself inside power, water and telecom networks so that it could disrupt them in a crisis.

The table gives the main disclosures about Salt Typhoon in order. Several disclosures came first from reporters, and US agencies, chiefly the Federal Bureau of Investigation (FBI) and the Cybersecurity and Infrastructure Security Agency (CISA), confirmed them.

Date	What was disclosed	Source
September 26, 2024	A group linked to China had broken into US internet providers	*The Wall Street Journal*
October 5, 2024	The systems that carriers use to carry out court-ordered wiretaps had been reached	*The Wall Street Journal*
November 13, 2024	The group had stolen call records, the communications of a small number of people in government and politics, and data on law-enforcement requests	FBI and CISA joint statement
December 27, 2024	A ninth US telecom company had been hit; carriers named in reports include Verizon, AT&T, T-Mobile, Charter, Lumen, Consolidated and Windstream	The White House; press reports
January 17, 2025	Sanctions on Sichuan Juxinhe Network Technology, a Chinese company, for direct involvement	US Treasury
January 7, 2026	Intrusions into the systems of several House committees in December 2025 were attributed to the group	*Financial Times*
February 3, 2026	A senator, Maria Cantwell, asked for a hearing with the chief executives of AT&T and Verizon, saying the carriers had not provided assessments of the breach by Mandiant, a security firm	Senator Cantwell's letter

Dates are those of publication or announcement.

On August 27, 2025, agencies of 13 countries, led by the National Security Agency (NSA), the FBI and CISA, set out how the group works. It got in through known flaws in equipment from Cisco, Ivanti and Palo Alto Networks, some of them years old, and then changed the settings of routers at the core and edges of carriers' networks to keep its access. It used the trusted links between carriers to move from one network to the next. The advisory named three Chinese companies that, it said, provide cyber products and services to units of the People's Liberation Army (PLA) and the Ministry of State Security (MSS), China's intelligence

services. China's foreign ministry called it false information driven by political motives and accused the US of getting other countries to "smear and frame" China (CP24, a Canadian news channel, August 28, 2025).

How large the campaign was depends on what is counted. Brett Leatherman of the FBI's cyber division said that the group had hacked at least 200 American companies, in a campaign that reached 80 countries (*The Washington Post*, August 27, 2025). The FBI also described more than 600 organizations in 80 countries as its targets (*The Register*, August 28, 2025). The two figures appear to measure different things: breaches of American companies, and organizations targeted worldwide.

The FCC's response changed with its chair. In January 2025, under its previous chair, Jessica Rosenworcel, it ruled that a 1994 wiretap law, the Communications Assistance for Law Enforcement Act (CALEA), obliges carriers to secure their networks against unlawful access, and it proposed requiring cybersecurity plans. On November 20, 2025, it rescinded the ruling by 2-1, with Brendan Carr and Olivia Trusty in favor. Carr called the ruling "unlawful" and "ineffective" and cited commitments the carriers had made instead: faster patching, tighter access controls and more sharing of threat information. Anna Gomez, the dissenting commissioner, said: "That is governing by hope rather than by duty."

Volt Typhoon is a different kind of threat. Microsoft first warned of it in May 2023. On February 7, 2024, CISA, the NSA and the FBI said "with high confidence" that the group, which they called state-sponsored by China, was positioning itself inside the computer networks of communications, energy, transport and water companies, ready to move into the systems that run physical equipment and disrupt them in a crisis. In some victims it had kept its foothold "for at least five years". The Justice Department said in January 2024 that it had removed the group's malware from hundreds of hijacked small-office routers. China's government has denied the accusations against both groups.

13b. The Huawei and ZTE Question

Huawei and ZTE (Zhongxing Telecommunication Equipment), two of the five largest makers of radio-network equipment (Module 7), are barred or restricted in the US and much of Europe. The governments that restrict them argue that Chinese law can compel companies to assist state intelligence work, so equipment at the heart of a network could be used to spy on it or shut it down. Huawei says its equipment poses no security risk.

The US moved first. The Secure and Trusted Communications Networks Act, signed in March 2020, set up a program to pay small carriers, most of them rural, to tear out Huawei and ZTE equipment and replace it, known as "rip and replace". The FCC stopped approving new Huawei and ZTE equipment in November 2022. Congress appropriated $1.9 billion, about $3.08 billion short of the costs the FCC approved, so the FCC paid only 39.5% of each carrier's approved costs, by its own count, until a law of December 23, 2024 let it borrow the gap from the Treasury.

The FCC repaid the loan in full on September 4, 2026, from the proceeds of its auction of AWS-3 (Advanced Wireless Services) spectrum licenses (Module 8), "earlier than anticipated", in Carr's words. The work itself is slower. Of 126 funded projects, 53 (42%) were complete by the FCC's report of June 2026, up from 13 in December 2025, and the deadline of May 8, 2026 passed with about 30 extensions of three to six months. Carriers blamed the supply chain (35% of participants), labor (15%), weather (15%) and permits (11%) (Broadband Breakfast and Light Reading, June 17-18, 2026).

Europe has moved more slowly. The EU's 5G Security Toolbox of 2020 left the decisions to member states, and ten of them have banned or restricted Huawei, by one public tally. Germany agreed with its three operators in July 2024 that critical Huawei and ZTE components must leave their 5G core networks by the

end of 2026 and their management systems by the end of 2029 (RCR Wireless, July 12, 2024). The UK barred purchases of new Huawei 5G equipment after 2020 and requires its removal by the end of 2027.

On January 20, 2026, the European Commission proposed a revised Cybersecurity Act that would let the EU assess “high-risk” suppliers and order them phased out, without naming any country, within 36 months for mobile networks. Henna Virkkunen, the Commission’s executive vice-president for tech sovereignty, security and democracy, said the voluntary toolbox had “not gone far enough”. Member states struck the fixed deadline in favor of case-by-case assessment, and the text still needed agreement with the Parliament (Reuters via Techzine, September 30, 2026).

Timotheus Höttges, chief executive of Deutsche Telekom, and 16 other executives warned that phasing out “high-risk” suppliers could cost Europe’s telecom sector up to €40 billion ($44.9 billion at the Federal Reserve’s rate for October 16, 2026), money they say would otherwise go into fiber and 5G (Techzine). Removing two of the five main suppliers also leaves operators fewer choices in a market that is already concentrated (Module 7). The cost figure is the executives’ own estimate, while governments assess the security risk in classified terms and have not quantified it in public.

13c. Shutdowns by the Numbers

An internet shutdown is a deliberate cut or restriction of internet access, ordered by a government, for a country, a region or particular services. Access Now, a digital-rights group that has tracked them since 2016, counted at least 313 in 52 countries in 2025, the most it has recorded, against 304 in 2024 and 289 in 2023 (report of March 31, 2026). Myanmar had 95 and India 65, together half the total, and Pakistan had 20. Since 2016, people in 100 countries have lived through at least one.

Conflict was the largest cause, with 125 shutdowns in 14 countries during wars and armed unrest. Protests led to 64, elections to 12 and school exams to 11, cuts that governments justify as a way to stop cheating. Access Now counted blocks of particular platforms separately, 94 in 40 countries, and it recorded 14 shutdowns of low-orbit satellite service in seven countries, up from four in 2024 (Module 9 covers satellites in war). When 2026 began, 75 shutdowns in 33 countries were still in force.

Other trackers count differently. Top10VPN, a site that reviews virtual private networks (VPNs), services that hide and reroute a user’s traffic, counted 212 major government-imposed outages in 28 countries in 2025. They lasted 120,095 hours in all, up 70%, and cost an estimated $19.7 billion, up 156% (Top10VPN, January 5, 2026). Its count includes social-media blocks and slowed connections as well as blackouts. Cloudflare, a network company whose Radar service publishes what its traffic shows, counted 174 major disruptions of all kinds in 2025, almost half of them shutdowns ordered by governments (December 15, 2025). Access Now counts each order, so a region’s repeated cuts or an exam season adds many entries, while Top10VPN weights each outage by its hours and Cloudflare sees only what its own traffic shows.

Governments cut access in several ways. The bluntest is to make networks stop announcing their addresses through the Border Gateway Protocol (BGP), the system networks use to tell each other which addresses they can reach (Module 3), so that the country drops off the world’s routing tables. Others leave the routes up and filter the traffic instead, letting through only approved users and sites, a list known as a whitelist. Others cut only mobile service, or slow chosen services until they are unusable. Monitors see each method differently: the Internet Outage Detection and Analysis project (IODA) at Georgia Tech watches routes and probes addresses, Kentik measures traffic for the networks that use its software, Cloudflare sees requests to the sites it serves, and NetBlocks runs its own measurements.

13d. Iran’s Two Blackouts

Iran cut itself off twice in 2026. The first blackout began on January 8, during nationwide protests over the economy and a collapsing currency, and lasted until January 27. Trackers placed its start between 16:30 and 18:45 Coordinated Universal Time (UTC), depending on how each measures (Cloudflare; IODA and Kentik with two Iranian digital-rights groups, January 14, 2026; net4people, a research forum). Iran’s networks withdrew 98.5% of their announced addresses in IPv6, the newer version of the Internet Protocol, but kept most of their older IPv4 routes so that whitelisted users could stay online. IODA found only about 3% of the addresses it probed responding.

The second began on February 28, 2026, within hours of the American and Israeli strikes that started the war. Cloudflare Radar saw Iran’s traffic fall at about 10:30 local time to below 1% of normal, while Iran’s address space stayed announced, a sign of heavy filtering with whitelists rather than withdrawn routes. Approved users kept access through subscriber identity module (SIM) cards issued to them that bypass the filters, known as “white SIM cards”. NetBlocks called it the longest nationwide shutdown on record on April 5. From April 20 the authorities restored access to favored groups, a tiered system that, according to press reports, allies of the Islamic Revolutionary Guard Corps (IRGC) backed and officials close to President Masoud Pezeshkian opposed.

On May 25, Pezeshkian ordered the communications ministry, led by Sattar Hashemi since August 2024, to restore international access “to its pre-January status”, and access returned in part from May 26 (Al-Monitor, May 26, 2026). Cloudflare counted an 88-day shutdown. Traffic came back at about 40% of its earlier level on May 27, peaked as high as 90% and settled near 59% (Cloudflare, July 28, 2026). Mohammad-Reza Aref, the first vice president, called the restoration “the first step toward free and regulated access to cyberspace”.

The status it returned to was already filtered. In October 2026 most Iranians used a domestic internet that blocks most international services and thousands of websites, and they reached the rest through VPNs (Al Jazeera, October 14, 2026). Fatemeh Mohajerani, the government’s spokesperson, said that “decision-making bodies are sometimes forced to bring down internet speeds”, citing protection against cyberattacks. Mohammad Amin Aghamiri, secretary of the Supreme National Cyberspace Council, backed a tiered model that would keep international access for selected groups, such as traders, researchers and journalists, during future restrictions, and he warned that more restrictions could come if the war intensified. The sources this course read report no new nationwide blackout after fighting resumed on August 31, 2026.

13e. Russia, China, and the Election Blackouts

Russia has built the legal and technical means to run its own internet. Amendments of 2019, known as the “sovereign internet” law, require providers to install state-controlled equipment that can filter, slow or block traffic, run by Roskomnadzor, the communications regulator led by Andrei Lipov since 2020. They also created a national domain name system that the state can use in place of the global one. YouTube has been slowed since July 2024.

In 2025-26 the state turned these tools on foreign messaging apps. Roskomnadzor said on August 13, 2025 that it was restricting calls on WhatsApp and Telegram, citing fraud, and on February 10, 2026 it confirmed that it had “slowed down” Telegram (The Record, February 11, 2026). On February 12 it removed WhatsApp’s main domain from the national domain name system, which WhatsApp called an attempt to block it fully. Telegram was blocked nationwide from April 10 (Institute for Security and Development Policy,

a Swedish research institute, July 6, 2026), though Pavel Durov, its founder, said about 65 million people in Russia still used it daily, many of them through VPNs (Kyiv Post, April 10, 2026).

The state's alternative is Max, a messenger built by VK, a Russian internet company. Since September 2025 it must be preinstalled on devices sold in Russia, public-sector employees and students are required to use it, and it works only with Russian or Belarusian SIM cards. Critics fear it could be used for surveillance. Mobile internet has also been cut region by region since May 2025, officially to stop Ukrainian drones, and in August 2025 the digital ministry of Maksut Shadayev introduced a whitelist of services deemed essential that stay reachable during cuts. On September 27, 2025, mobile shutdowns hit 54 regions and whitelists applied in 30 (Na Svyazi, a monitoring project, via Global Voices). More than 400 VPNs were blocked by mid-January 2026, 70% more than in the autumn of 2025 (the same institute), and Top10VPN put Russia's 2025 losses from 57 shutdowns at $11.9 billion, more than half of the world total.

China's Great Firewall (GFW), the system that filters traffic at the borders of China's networks, keeps out most foreign platforms and blocks many tools for getting around it. It can misfire in public. On August 20, 2025, for about 74 minutes after midnight Beijing time, it sent forged reset signals, which tell both ends to drop a connection, to every connection on port 443, the channel used for encrypted web traffic, from a device that matched no known part of the system (GFW Report, a research group). China also tightened its rules: the Cyberspace Administration of China published incident-reporting measures on November 1, 2025, and an amended Cybersecurity Law took effect on January 1, 2026.

China's methods are also for sale. On September 11, 2025, more than 500 gigabytes (GB) of internal documents and source code leaked from Geedge Networks, a company linked to Fang Binxing, the firewall's chief designer, and from a laboratory of the Chinese Academy of Sciences. GFW Report verified the files. They describe "Tiangou", a commercial system that reports on the leak call a "Great Firewall in a box", able to inspect traffic, detect VPNs and log whole sessions. It was deployed in Myanmar, in 26 data centers whose dashboards monitored 81 million connections at once, and was reportedly sold to Pakistan, Ethiopia and Kazakhstan.

Elsewhere, shutdowns around elections and exams went on into 2026, and Russia's strikes on Ukraine's power system cut connectivity there (table).

Dates	Country	What happened	Measured by
October 29 to November 3, 2025	Tanzania	Election shutdown, one day and then five more; traffic more than 70% below normal	Cloudflare
January 7-8 and January 26, 2026	Ukraine	Russian strikes on energy infrastructure cut traffic in the Dnipropetrovsk and Kharkiv regions by about half	Cloudflare
January 13-17, 2026	Uganda	Mobile operators ordered offline around the presidential election, an order the Uganda Communications Commission defended; domestic traffic at the Uganda Internet Exchange Point fell from about 72 gigabits per second (Gbps) to 1 Gbps	Cloudflare; exchange data
About 60 hours around March 15, 2026	Republic of Congo	Nationwide election shutdown	Cloudflare
April 13-23, 2026	Sudan	Ten exam shutdowns of about 3.5 hours each	Cloudflare
June 2, 11 and 28, 2026	Iraq	Exam shutdowns of about 90 minutes each	Cloudflare

Cloudflare Radar's 2025 Year in Review (December 15, 2025) and its disruption summaries for the first quarter of 2026 (as reported by *The Register* and SDxCentral, April 28-29, 2026) and the second quarter (July 28, 2026).

13f. Jamming the Signals from Orbit

Networks also depend on signals they do not control. Global navigation satellite systems (GNSS), such as the American Global Positioning System (GPS), Europe's Galileo and China's BeiDou, broadcast time and position from orbit. Ships report their positions through the Automatic Identification System (AIS) using these signals, and many mobile networks take their timing from them. Jamming drowns the signals in noise, and spoofing fakes them so that receivers compute a false position.

The 2026 war made the Gulf the most jammed region on Earth. Windward, a maritime-data company, counted 978,000 jamming events worldwide in the first quarter of 2026, 959,627 of them (98%) in the Gulf, against 7,630 in the Black Sea. More than 1,100 vessels were affected in the first 24 hours after the strikes of February 28, and Windward found at least 35 clusters of jamming between February 28 and March 2. Spoofed ships appeared on its maps at airports, at a nuclear power plant and inland near Bandar Abbas. Windward did not say who was jamming, and the war also disrupted cable work in the Gulf (Module 5).

13g. Is the Internet Splitting into National Networks? The Cases as Their Advocates Make Them

The case that it is splitting: its advocates, among them digital-rights groups such as Access Now and researchers who measure censorship, point to what 2026 showed. Iran ran two nationwide blackouts and kept a national network for approved users (section 13d). Russia pulled a foreign messenger out of its national name system, banned another, built a state alternative and keeps whitelists of approved services for its regional cuts (section 13e). China's filtering is now sold abroad as a product, and shutdowns of satellite service, the last way around a national cut, are rising (section 13c). Even the World Trade Organization's pledge against customs duties on electronic transmissions lapsed (Module 12). Each step, they argue, turns a country's internet into a gated network that the state opens at its discretion, a "splinternet" of national systems.

The case that it is overstated: its advocates, among them engineers and researchers such as Nick Merrill of the University of California, Berkeley, argue that the internet's core is still one: a single root zone, one set of address registries and common protocols (Module 12). Countries that cut themselves off come back, because staying off costs too much, and Iran restored access after its wartime blackout. Circumvention is mass behavior: about 41% of Russian internet users relied on VPNs by 2025 (the Center for European Policy Analysis, citing Cybernews), and Telegram kept tens of millions of daily users in Russia after its ban. Merrill found fragmentation spread across many layers rather than moving in one direction, and China's model rare even among the countries of its Belt and Road Initiative (Techdirt, July 13, 2020).

What the evidence says: fragmentation is measurable and growing at three layers: the applications people may use, the content they may see, and the state's control of the gates. The core of protocols and identifiers has not split, since no country runs a rival root zone or address registry that others use. Full national disconnection remains rare and costly, and Iran's was partly reversed. Partial disconnection is getting cheaper and more common: whitelists let a state keep banks, officials and approved users online while cutting everyone else, and Iran and Russia have proposed or introduced them as standing arrangements (sections 13d and 13e). Freedom House recorded a 15th year of decline in online freedom (Module 12).

The course's reading, stated once: the splinternet is real but layered. The internet is not splitting into separate networks at its core, and 2026 shows why: a country that leaves the shared core pays heavily and comes back. What is spreading is the gated national network, in which the state decides who reaches the global internet and when. Whitelists make that model cheaper than a full blackout, so the split to watch is

political rather than technical: between countries where access is the default and countries where it is a permission.

Worked example — What Iran’s Blackouts Cost, and What the Estimates Leave Out

Shutdown costs are modeled, not measured. Top10VPN multiplies each shutdown’s hours by a cost per hour from the Cost of Shutdown Tool (COST) of NetBlocks, which draws on World Bank, ITU, Eurostat and US Census data on the size of each economy’s online activity. Partial outages are scaled by the market share of the networks cut, regional ones by the region’s share of output and platform blocks by their users. Slowing service to second-generation (2G) mobile speeds counts as a blackout, and blocks of more than 365 days are left out as permanent censorship. The inputs are these. NetBlocks put Iran’s wartime blackout at about $37 million a day, and Sattar Hashemi, the communications minister, put it at $35.7 million (MEA Tech Watch, a technology news site, April 1, 2026). The wartime blackout lasted 88 days and the January one 19. Iran’s output was $119.7 billion in the quarter to June 2025 (CEIC, a data company, converting official figures at the central bank’s average market rate). Step 1: 88 days × $37 million = $3.26 billion, or $3.14 billion at the minister’s rate. NetBlocks’ running total of $1.8 billion on April 16, after 48 days, implies $37.5 million a day, close to its rate. Step 2: the January blackout adds 19 × $35.7-37 million = $0.68-0.70 billion, for $3.82-3.96 billion over 107 days, or 35.9% of the 298 days of 2026 to October 25. Step 3: per hour, $37 million ÷ 24 = $1.54 million, 9.4 times Top10VPN’s 2025 world average of $164,000 an hour ($19.7 billion ÷ 120,095 hours), because the average mixes regional cuts and platform blocks with national blackouts. Iran’s two blackouts alone equal 19-20% of the world’s 2025 total. Step 4: output was $119.7 billion ÷ 91 days = $1.32 billion a day, so the estimates assume that a near-total blackout removes 2.7-2.8% of each day’s output. Iran has several exchange rates, and at a weaker one its dollar output would be smaller and that share larger. What the method leaves out: it treats business delayed as business lost, and it ignores longer costs, such as firms and skilled workers that leave and investment not made. It counts nothing for news, emergency calls, health care or rights. It also misses the months of partial service after restoration (section 13d), which no tracker counts as a shutdown.

Module 13 Key Points

Networks Under Attack

• The FBI said Salt Typhoon hacked at least 200 US companies; the FCC rescinded its carrier-security ruling on November 20, 2025, by 2-1.

• The FCC repaid its $3.08 billion rip-and-replace loan on September 4, 2026, but only 53 of 126 projects were complete in June.

• Access Now counted a record 313 shutdowns in 52 countries in 2025; Myanmar and India accounted for half.

• Iran’s wartime blackout lasted 88 days from February 28, 2026; traffic then settled near 59% of its earlier level.

• Russia removed WhatsApp’s domain from its national name system in February 2026, blocked Telegram from April 10, and requires Max on new devices.

• The course’s reading: the splinternet is real but layered; the core stays single while whitelists make gated national networks cheaper.

Module 13 Quiz

Check Your Understanding

1. What did the FCC do on November 20, 2025?

A) It ordered carriers to replace all Huawei and ZTE equipment by May 8, 2026.
B) It rescinded its January 2025 wiretap-law ruling, by 2-1.
C) It fined AT&T for Salt Typhoon.
D) It adopted binding cybersecurity rules for every carrier after the Salt Typhoon breaches.

2. How many American companies did the FBI say Salt Typhoon had hacked, as of August 27, 2025?

A) At least 200, in a campaign reaching 80 countries.
B) Nine, all of them telephone carriers, as the White House said in December 2024.
C) More than 600, all of them inside the United States, by the FBI's count.
D) About a million, counting every customer whose call records were taken.

3. How many internet shutdowns did Access Now count in 2025?

A) 174, almost half ordered by governments
B) 212, in 28 countries
C) At least 313, in 52 countries
D) 304, a record

4. How long did Iran's wartime blackout of 2026 last, by Cloudflare's count?

A) 19 days
B) 107 days
C) Six days, in June
D) 88 days, from February 28 to late May

5. What is the course's reading on the "splinternet"?

A) Already split into national internets
B) Overstated; little has changed
C) Real but layered: the core is still one
D) A split driven by new protocols that national networks adopt

Module 14 — Where Things Stand, and What to Watch

The Argument, the Contradictions, and the Calendar

Recommended Reading / Reference

Book: Tim Wu, *The Master Switch: The Rise and Fall of Information Empires* (2010), a Columbia law professor's history of how the telephone, radio, film and cable television each moved from open competition to control by a few firms, and his argument about whether the internet will follow

Reference: The measurements that date the watch list below: the ITU's yearly *Facts and Figures*, Cloudflare Radar and the Internet Outage Detection and Analysis project (IODA) for traffic and shutdowns, and TeleGeography's cable map

14a. The Argument on One Page

The course has made one argument from thirteen directions. The internet looks like a cloud, but it is a physical system of glass, copper, radio, satellites and data centers, which companies own and which agreements, more than governments, run. About 6.0 billion people used it in 2025, by the ITU's count (Module 1), through about 77,000 separate networks that share one set of protocols and one system of names and addresses (Modules 1 and 3). Its traffic crosses oceans in cables, 694 systems on TeleGeography's 2026 map, counting those planned (Module 5), and ends in data centers that used about 485 TWh of electricity in 2025 (Module 10).

Ownership has moved toward the companies that once only used the networks. Content and cloud companies accounted for about 75% of the world's used international bandwidth at the end of 2024 (Module 5), they bring their video into ISPs through caches of their own (Module 4), and three of them held 63% of the cloud market in the second quarter of 2026 (Module 10). The carriers kept the costliest part, the last mile to every home (Module 6), and their revenue has not grown with the traffic they carry (Module 7). The shared core belongs to no one: the root zone, the address registries and the open standards stayed single through the end of American oversight in 2016 and through the crisis at AFRINIC, Africa's address registry (Module 12).

Its weak points are physical and political, and 2024-26 tested both. Cables were cut in the Red Sea in 2024, most likely by a drifting wreck's anchor, and in the Baltic, where courts have not settled whether the damage was deliberate, and a cable repair took 51 days on average in 2024 (Module 5). Within 30 days in late 2025, routine changes at AWS, Microsoft and Cloudflare each stopped services around the world, and in March 2026 drones struck AWS data centers in the Gulf, in strikes that press reports attributed to Iran (Module 10). Governments ordered at least 313 shutdowns in 2025, by Access Now's count, and Iran's wartime blackout lasted 88 days (Module 13).

The largest gap is cost, not coverage. Of the 2.2 billion people offline in 2025, about 1.8 billion lived under a mobile broadband signal, on the ITU's figures, and the GSMA counts 3.1 billion people under one who do not use mobile internet, against 270 million with no coverage (Module 11). Non-users name the price of a handset as the first barrier, and in low-income countries 2 GB of data a month cost a median 7.4% of average income in 2024. America shows the same pattern at a higher income: 96.9% of Americans could buy land-based broadband at 100/20 Mbps, but only 54% of adults in households earning under $30,000 a year had broadband at home in 2025 (Module 11).

Carrying data keeps getting cheaper, while reaching each home stays dear. The lowest transit offers cost $0.03 per Mbps a month in mid-2026, after prices for 100-gigabit ports fell 17% a year from 2023 (Module 4), and one new transatlantic cable could carry the evening downloads of tens of millions of homes (Module 5). The cost sits at the edges. America's rural program pays about $4,799 per location it connects (Module 11), a mid-band 5G site serves only dozens of people at video speeds at the same moment (Module 7), and a satellite's capacity is shared by everyone beneath it (Module 9). Most of the course's contested questions are, at bottom, arguments over who pays for the edges.

14b. Five Contradictions

Each module states its own findings. Read together, they produce tensions that no single module states.

Built to route around damage, run through a few places. Packet switching and the Border Gateway Protocol send traffic around a broken link (Module 3), and ocean cables back one another up, so most cable faults cut no one off (Module 5). Yet the services people use have gathered in a few providers and regions, and one empty DNS record in AWS's northern Virginia region stopped services around the world for about 14.5 hours in October 2025 (Module 10). The protocols kept working. What failed was the commercial layer above them, where redundancy is each customer's choice and costs money.

The networks' largest customers are among their largest builders. Europe's large carriers argue that content companies should pay for the traffic they send (Module 4). The same companies accounted for about 75% of the world's used international bandwidth at the end of 2024, increasingly on cables they own (Module 5), and their caches inside ISPs cut the carriers' transit bills (Module 4). BEREC found that rising traffic does not directly lead to significant extra costs compared with the cost of the networks as a whole, so the argument is over who pays for the last mile, not over who pays to carry the traffic (Module 4).

Money for networks, not for bills. The usage gap is about 11 times the coverage gap by the GSMA's count, and the course reads the divide as mainly one of cost and skills (Module 11). Yet America's largest program, BEAD, spends $18.19 billion to build networks to 3.79 million locations, while the program that paid part of the bills of more than 23 million low-income households ended on June 1, 2024, when its money ran out. The Brattle Group estimated that about 5 million households then dropped home internet altogether (Module 11).

A billion connections, and 65 shutdowns. India had 1,001.5 million broadband subscriptions in June 2026, all but 6.6% of them mobile (Module 6), and smartphones in the region of India, Nepal and Bhutan used 37 GB a month in 2025, more than North America's 25 GB (Module 7). India also ordered 65 shutdowns in 2025, second only to Myanmar, by Access Now's count (Module 13). A government can build mass access and keep the power to withdraw it, which is why the course measures openness by whether access is a default or a permission, not by how many people a country connects (Module 13).

The satellites that get around the state answer to a company, and to a government. Low-orbit satellites reach homes that no line or tower reaches, and smuggled Starlink terminals kept some Iranians online during the shutdowns of 2026 (Module 9). Governments have answered: shutdowns of satellite service rose from four in 2024 to 14 in 2025 (Module 13). Yet one company's satellites are about 64% of the spacecraft operating in orbit, and the speed cap that cut Ukraine's Starshield links on fast drones in October 2026 showed that its decisions, or those of the government it answers to, can restrict an army's communications in the field (Module 9).

14c. What to Watch

The dates below are the ones that would confirm or overturn the course's readings. Each entry says what the event would show, not what the course has already reported about it.

When	What	Why it matters
October 26, 2026	Verizon's third-quarter results; the deadline to seek leave to appeal the Eagle S ruling to Finland's Supreme Court	Verizon's newest counts of fixed wireless and fiber customers (Module 6); whether the first Baltic cable case returns to trial in Helsinki (Module 5)
October 28-29, 2026	Third-quarter results: Microsoft, Alphabet and T-Mobile on October 28, Amazon on October 29	Whether the largest spenders raise their 2026 capital spending above the $720-745 billion they guided to in July (Module 10), and how many broadband customers T-Mobile, which no longer reports its fixed wireless count, adds (Module 6)
October 29, 2026	The FCC's October open meeting, with proposals on more mid-band spectrum and on satellite service to ordinary phones	How fast the FCC widens the spectrum pipeline (Module 8) and opens mobile bands to satellites (Module 9)
November 5, 2026	Satellite operators' plans for clearing the upper C-band are due at the FCC	The first step toward the auction that must finish by July 4, 2027 (Module 8)
November 6, 2026	Charter's third-quarter results, its first since buying Cox	Whether America's largest broadband provider keeps losing customers to fiber and fixed wireless (Module 6)
November 9-27, 2026	The ITU's Plenipotentiary Conference in Doha, with ballots from November 12, the Secretary-General's first	Whether states win a larger role for the ITU, and what it decides on cables, orbits and the treaty on international traffic (Module 12)
November 2026 (expected)	The ITU's *Facts and Figures 2026*; Ericsson's November *Mobility Report*	New counts of users, of the offline and of the price of access (Modules 1 and 11), and of 5G subscriptions (Module 7)
December 14-18, 2026	The Internet Governance Forum in Nairobi, the first since the UN made it permanent	How the multistakeholder model works in its permanent form (Module 12)
Mid-December 2026 (expected)	Cloudflare Radar's review of 2026	Traffic growth, the shares of encrypted and post-quantum traffic, and the year's disruptions (Modules 1, 3 and 13)
End of 2026	Germany's deadline for critical Huawei and ZTE components to leave its 5G core networks; AT&T's target of 40 million fiber locations	How fast Europe removes Chinese equipment (Module 13), and how fast fiber spreads (Module 6)
Early 2027 (expected)	Access Now's count of 2026's shutdowns; TeleGeography's 2027 cable map	Whether shutdowns set another record in the year of Iran's two blackouts (Module 13), and how fast the cable network grows (Module 5)
March 2027	3GPP freezes the service requirements of Release 21, its first set of 6G standards	The first fixed step toward 6G, expected in service around 2030 (Module 7)
April 27, 2027	Bidding is tentatively set to begin in Auction 115, for 160 MHz of the upper C-band	The first band of the 2025 law's 800 MHz pipeline, and a test of the over $85 billion that Congress's budget office expects auctions to raise by 2034 (Module 8)
October 18-November 12, 2027	WRC-27, the World Radiocommunication Conference, in Shanghai	Bands for future mobile networks, and the rules for satellites that connect directly to phones (Modules 8 and 9)
End of 2027	Britain's deadline for removing Huawei's 5G equipment; the latest date the FCC's only Democrat may serve as a holdover	Europe's second test of removing Chinese equipment (Module 13), and whether the FCC is left with commissioners of one party (Module 12)
On or about November 30, 2027	SpaceX's purchase of EchoStar's spectrum is due to close	Whether satellite service to ordinary phones grows beyond texts and light apps (Modules 8 and 9)

When	What	Why it matters
2028	Amazon’s Fastnet cable is due to enter service; in December, 3GPP freezes the detailed 6G specifications	More ocean capacity owned by a cloud company (Module 5), and the standard that 6G equipment will be built to (Module 7)
October 31, 2029, and December 31, 2035	Under the EU’s Digital Networks Act as proposed: national fiber plans due, then copper networks closed	Whether Europe retires copper by law, as AT&T plans to retire most of its own by the end of 2029 (Modules 6 and 12)
No date set	NTIA’s guidance on BEAD’s leftover money; the Fitburg verdict; the EU Council’s position on the Digital Networks Act; Amazon Leo’s general service	Where about $21 billion goes (Module 11); the first ruling on intent in a Baltic cable case (Module 5); Europe’s network law (Module 12); Starlink’s main Western rival (Module 9)

Dates are as announced by October 25, 2026. Entries marked expected follow the timing of earlier editions and had not been announced. The Digital Networks Act is a proposal, so its dates would change if the European Parliament and the Council amend it.

14d. The Last Word

The history in Module 2 begins with 38 miles of wire between Washington and Baltimore, and the course ends with a network of about 77,000 networks that no one runs. Each stage of its history turned on three questions: who builds the lines, who pays for them, and who may cut them. The telegraph and the telephone answered with private monopolies in America and state administrations in Europe. The internet answered with agreements: open protocols, free exchange between networks and private contracts, under which it grew from about 39 million users in 1995 to about 6.0 billion in 2025. Those agreements still hold at the core, and the pressure has moved to the edges: the last mile that poor households cannot afford, the few providers and regions where the services people use now sit, and the gates that governments build to decide who reaches the rest of the world.

A reader who follows the internet from here will learn more from a few measurements than from announcements. The ITU’s yearly count shows who is online and what access costs, Cloudflare Radar and IODA show traffic and shutdowns as they happen, and TeleGeography counts the cables and the time repairs take. The quarterly subscriber counts show whether fiber and fixed wireless keep taking customers from cable, and Ookla’s medians show the speeds people actually get. Each is dated and sourced, and each could prove the course’s readings wrong.

Module 14 Key Points

Where Things Stand, and What to Watch

• The internet is a physical system of cables, radio, satellites and data centers, mostly company-owned and run more by agreements than by governments.

• Its weak points are physical and political: cable cuts, cloud outages, drone strikes and a record 313 shutdowns tested them in 2025-26.

• The largest gap is cost, not coverage: 3.1 billion people live under a mobile signal they do not use (GSMA).

• Carrying data keeps getting cheaper; reaching each home, by fiber, radio or satellite, is where the cost and the arguments lie.

• The shared core of names, numbers and protocols stayed single; what is spreading is the gated national network.

Module 14 Quiz

Check Your Understanding

1. What does the course argue the internet is?

A) It is one network that the US government ran until it handed control to ICANN in 2016.

B) It is a physical system that companies own and agreements run.

C) It is a weightless system of software.

D) It is a set of national networks that governments own and link to one another by treaties negotiated at the ITU.

2. What tension does the module draw from Europe's argument over fees for content companies?

A) Carriers own every ocean cable

B) Content companies refuse all peering

C) Regulators already set fees by law

D) Those asked to pay are among the largest builders

3. Which American policy does the module set against the course's reading on the digital divide?

A) Building networks while a bill subsidy ended

B) Capping broadband prices wherever only one wired provider serves a home, as the FCC voted to do in 2026

C) Spending more of BEAD on satellites than on fiber

D) Ending the net-neutrality rules, which the course says keeps low-income households offline

4. What does the module say the speed cap on Starshield in October 2026 showed?

A) It showed that Starlink's satellites cannot follow fast drones.

B) It showed that Russia had jammed Starlink across all of Ukraine.

C) It showed that geostationary satellites are too slow for drones, as the delay over any orbit above 2,000 km would predict.

D) It showed that one company's or one government's decisions can restrict an army's communications.

5. Where does the course say the pressure on the internet has moved?

A) To the root servers

B) To the protocols

C) To the edges: the last mile, a few providers, and national gates

D) To the ocean floor

Answer Sheet

Answers to every module's Check Your Understanding quiz.

Module	Answers
Module 1 Quiz	1. A 2. C 3. B 4. D 5. D
Module 2 Quiz	1. A 2. B 3. C 4. C 5. A
Module 3 Quiz	1. B 2. D 3. C 4. B 5. D
Module 4 Quiz	1. A 2. C 3. A 4. D 5. B
Module 5 Quiz	1. A 2. D 3. B 4. C 5. A
Module 6 Quiz	1. B 2. D 3. C 4. B 5. A
Module 7 Quiz	1. D 2. C 3. A 4. C 5. B
Module 8 Quiz	1. D 2. A 3. D 4. C 5. B
Module 9 Quiz	1. C 2. D 3. B 4. A 5. D
Module 10 Quiz	1. A 2. B 3. C 4. D 5. C
Module 11 Quiz	1. A 2. B 3. B 4. D 5. A
Module 12 Quiz	1. C 2. B 3. A 4. D 5. C
Module 13 Quiz	1. B 2. A 3. C 4. D 5. C
Module 14 Quiz	1. B 2. D 3. A 4. D 5. C