TeleGeography

The State of the Network

2026 Edition

A comprehensive analysis of the global telecommunications landscape, focusing on infrastructure and market shifts.

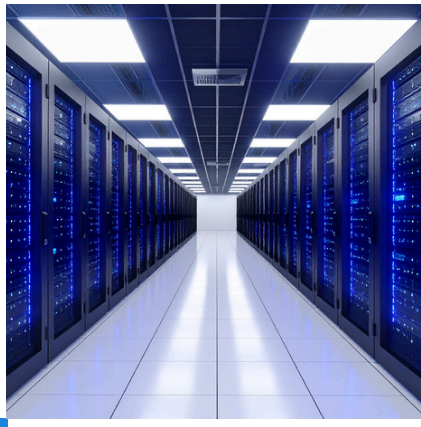

About Us

TeleGeography is a telecommunications data provider known for independent analysis. Trusted since 1989, our research is vital to telecom service providers, large enterprises, government agencies, and other research firms. Our data and analysis platforms deliver comprehensive, unbiased research and actionable insights.

Learn more about our work.

01 The State of Things

An overview of this resource and how you can get the most from these insights.

02 Transport Networks

Transport Networks Research Service assesses the global telecom capacity market and evaluates factors shaping long-term demand and price movement.

03 IP Networks

In IP Networks Research Service, we delve into the core meaning of our extensive internet capacity and traffic datasets.

04 Data Centers

Data Center Research Service analyzes an unprecedented shift fueled by global regulatory changes and the rise of AI.

05 Cloud & WAN

We're seeing a slowdown in new cloud region deployment. Cloud and WAN Research Service explores, with insight on wide area network development.

06 Voice

The downward slide continues for international voice traffic. International Voice Report details the latest.

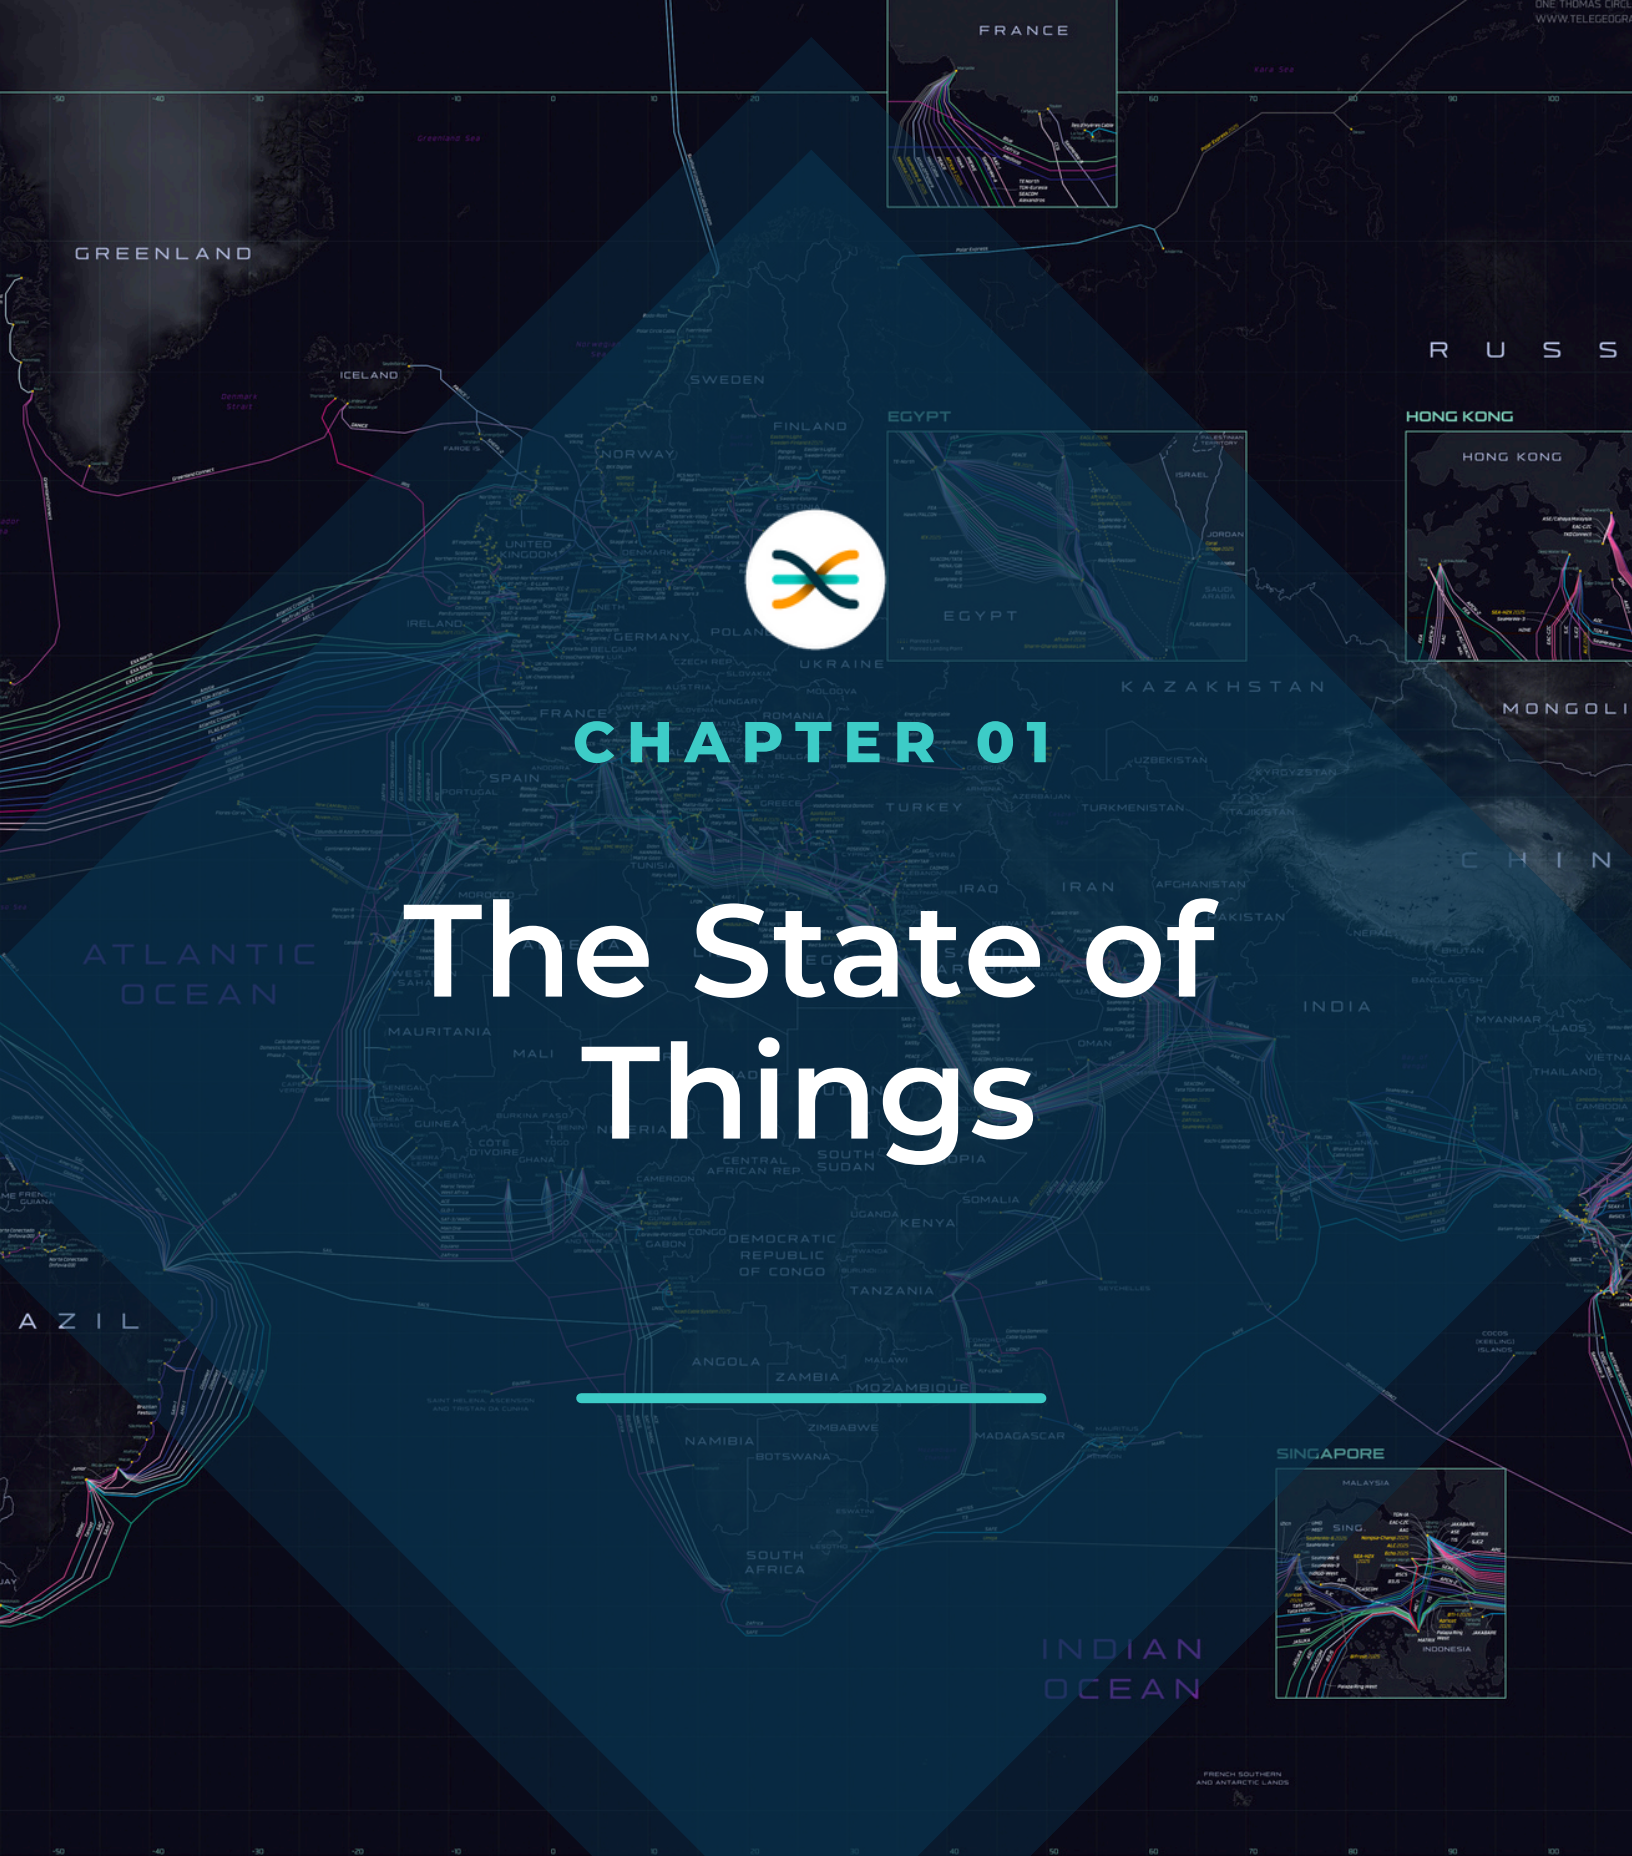

CHAPTER 01

The State of Things

Your Guide to Global Telecom

Who uses this independent data and analysis?

Our work empowers organizations in diverse industries around the globe, including:

- Telecommunications service providers
- Large enterprises
- Government agencies

Welcome to TeleGeography's annual comprehensive analysis of the global telecommunications landscape. This report focuses on infrastructure and the current market dynamics and recent industry shifts.

In this 2026 edition, we cover:

CABLE TALK

We highlight a massive surge in submarine cable investment, projected to exceed $14 billion through 2027, to meet rising bandwidth needs. Content and cloud providers like Google and Meta now dominate international capacity, significantly outpacing traditional internet backbone providers.

PRICE CHECK

The report also notes that while bandwidth demand and traffic continue to grow, market prices for capacity and IP transit are steadily declining.

BEHIND THE HEADLINES

We address how geopolitical tensions and the rise of artificial intelligence are reshaping network routing and data center strategies.

Using This Resource

DATA FOR YOUR BIGGEST DECISIONS

There's no one way to leverage TeleGeography data and insight. As we review our top findings from the last year, shedding light on complex market dynamics, there are four key areas supported by our work.

Guiding Infrastructure Investment and Development

This document is essential for organizations planning capital expenditures, as it tracks where physical infrastructure is being built and where it's needed. For example, it highlights that the aggregate cost of new undersea cable construction is expected to reach over $14 billion between 2025 and 2027, driven largely by content providers like Google and Meta rather than traditional internet backbone providers. By identifying specific high-growth regions, investors can target underserved markets rather than saturated ones.

Benchmarking Pricing and Costs

For enterprises and carriers negotiating contracts, TeleGeography research provides vital leverage by revealing global pricing trends. It offers granular data on cost erosion and illuminates regional disparities that affect total cost of ownership. For instance, cross-connect prices in U.S. markets are roughly three times higher than in European markets. This helps companies accurately budget for connectivity and decide whether to prioritize specific hubs, based on colocation and interconnection costs

Navigating Geopolitical Risk

The report correlates geopolitical tensions with network vulnerability, detailing how certain conflicts have created permitting challenges and physical threats to cable repair vessels, leading to delays. It also sheds light on regulatory hurdles, such as difficulty obtaining permits in the South China Sea, which is forcing operators to design new routes to bypass problematic waters. This intelligence supports network architects in building diversity into their systems to avoid single points of failure.

Adapting to Technological Disruptions

Finally, this research helps the industry prepare for massive technological shifts, specifically the impact of Artificial Intelligence (AI) and the transition to cloud infrastructure. For instance, in this e-book you'll read about an AI-driven "power scarcity" crisis, with wait times for data center power extending to four years or longer.

CHAPTER 02

Transport Networks Research Service

Learn more about
Transport Networks Research Service

Findings

To many people, the concepts of global network infrastructure and bandwidth markets are difficult to grasp. But those who follow this sector understand that it's a fundamental building block of the global economy. As with other areas of industry, the capacity market can struggle managing growth and uncertainty. Our *Transport Networks Research Service* assesses the state of the global telecom capacity market and evaluates the factors that shape long-term demand and price movements. We look at market conditions on both a global level and on a regional level, focusing on critical submarine cable routes.

Features of this service include:

- Detailed profiles of over 460 network service providers and more than 670 submarine cable systems, also downloadable in Excel.
- A new bandwidth demand search interface with historical data back to 2005 for countries, subregion, regions, and major submarine cable routes. Forecasted data are now available to 2031.
- A regularly-updated table of Planned Submarine Cables, containing ownership, landings, cost, and other critical information, and divided by route deployment.
- A network database, that allows customers to find carriers and submarine cables by location.
- Latency data from network operators for a variety of long-haul terrestrial and submarine routes.
- The Submarine Cable Fault Database features the latest publicly-disclosed cable fault information.
- An enhanced Submarine Cable Map including tools to search by cable owner, a feature not available in our free map.

Demand Trends

Global bandwidth demand continues to climb, albeit at a steadily slowing rate. As recently as 2020, year-on-year demand had increased 45%, but the pace decelerated to 29% in 2024. Still, this represents a steady 32% CAGR and a tripling of demand between 2020 and 2024, with demand now surpassing 6.4 Pbps.

FIGURE 2

Used International Bandwidth Growth by Region

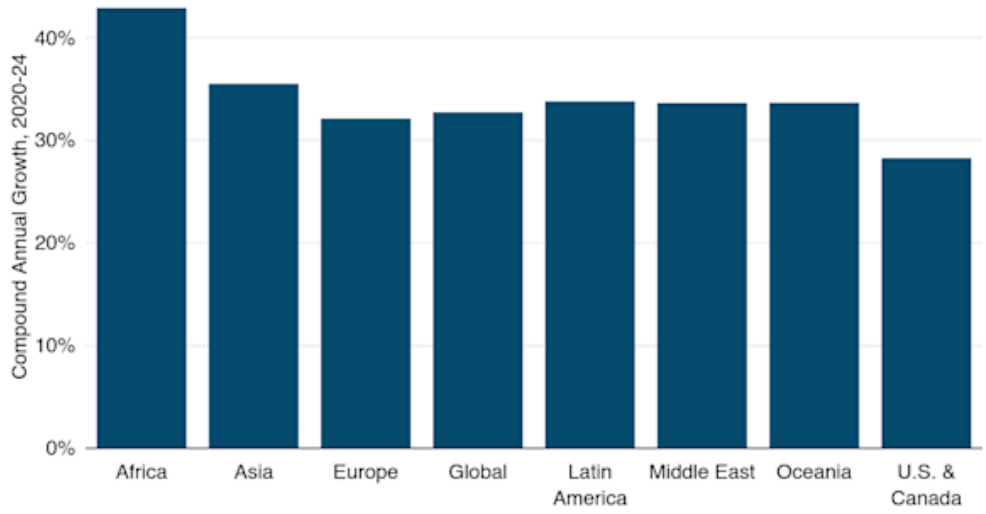

Source: TeleGeography

FIGURE 1

Worldwide International Bandwidth Growth

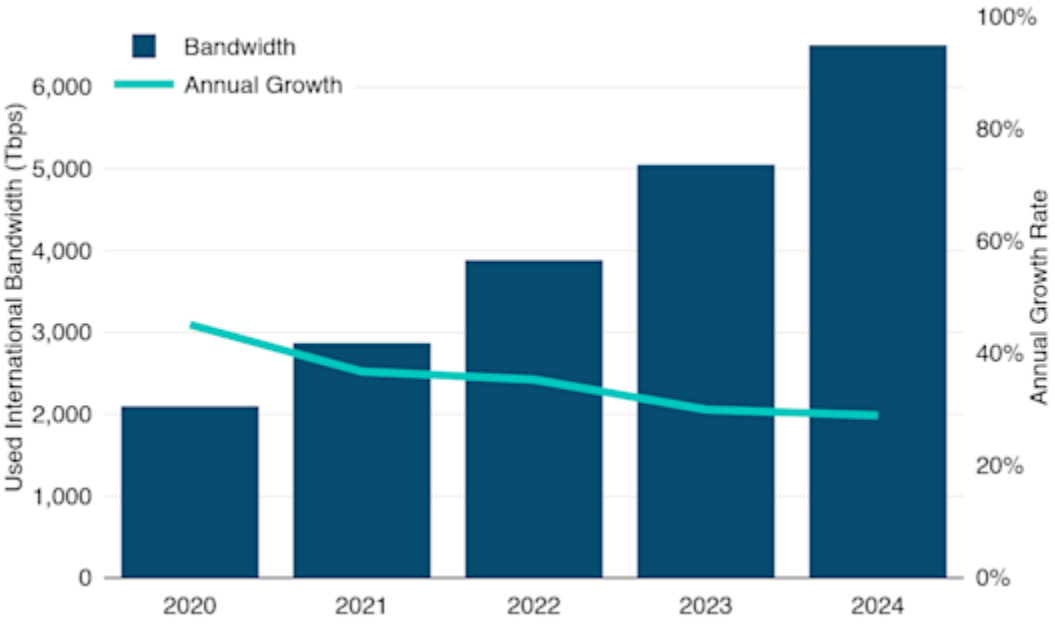

Source: TeleGeography

Now take a look at the bar chart below. On a regional level, most parts of the world have seen very comparable growth at about 32-35% CAGR since 2020. Africa outpaces other regions with more than 40% CAGR demand growth, and the mature U.S. market falls slightly below the average with about 29% growth in bandwidth demand.

In 2024, content & cloud networks accounted for almost **three quarters** of all bandwidth demand.

Learn more about
Transport Networks Research Service

The Role of Content Providers

Content and cloud providers—most specifically a handful of companies like Google, Meta, Microsoft, and Amazon—overwhelmingly account for most of the world's international bandwidth usage. As recently as 2016, internet backbone providers accounted for the majority of demand. Not anymore. In 2024, content and cloud networks accounted for almost three quarters of all bandwidth demand.

Take a look at the 100% bar chart below. You can see here that, while content providers' share of total demand varies by route, it's becoming dominant just about everywhere. On the massive trans-Atlantic, trans-Pacific, and intra-Asian routes, these networks account for at least 80% of bandwidth demand.

FIGURE 3

Share of Used Bandwidth by Category for Major Routes

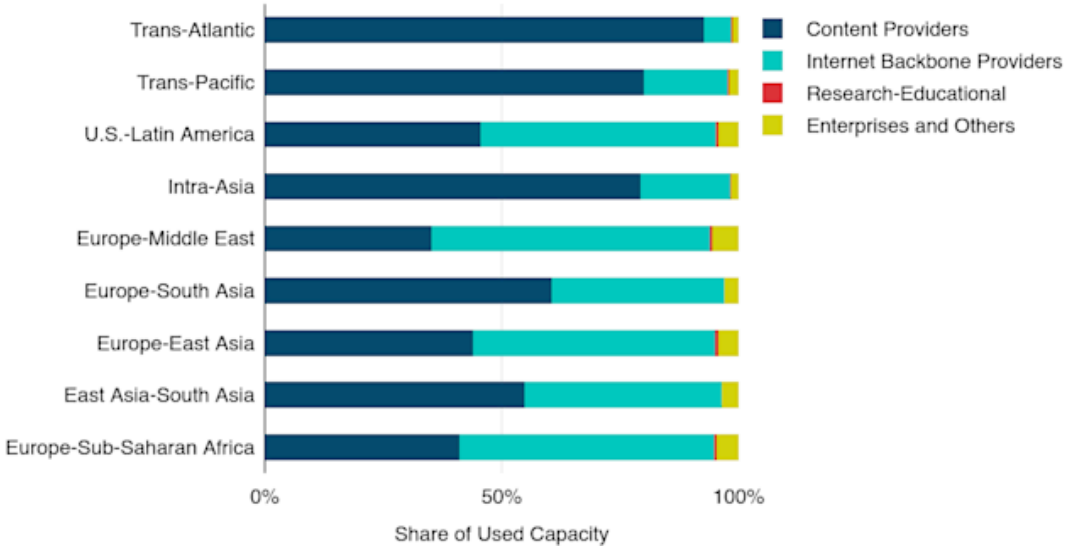

Notes: Data shows used bandwidth as of year-end 2024.

Source: TeleGeography

A few smaller routes like Europe-East Asia and Europe-Sub-Saharan Africa are still driven by internet backbone demand, but the dynamic is shifting even in these locations. This is clearly illustrated in the bar chart below, which shows that content provider demand growth is outpacing that of other network sources on all of the route groupings we track. Take a look at Latin America, where content demand growth is almost double that of other sources. Within the next year, content providers will account for the outright majority of demand on U.S.-Latin American routes.

FIGURE 4

Content Providers versus Others Bandwidth Growth by Region

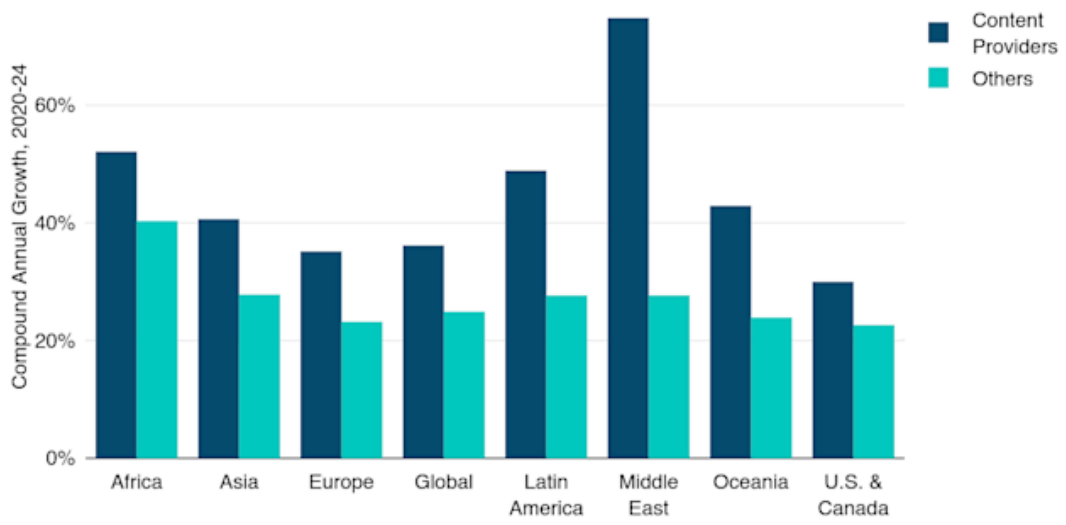

Source: TeleGeography

Meeting Demand Requirements

To keep pace with ever-increasing bandwidth demand, a steady stream of investment has driven tremendous growth in subsea cable infrastructure. Aggregate cost of new construction over the past nine years has averaged over $2 billion per year. With the exception of a few anomalous years, we haven't seen this level of investment on subsea cable infrastructure since 2000-2001—and it's not letting up. The value of new submarine cables planned to enter service from 2025-2027 could to reach over $14 billion.

FIGURE 5

Construction Cost of Submarine Cables

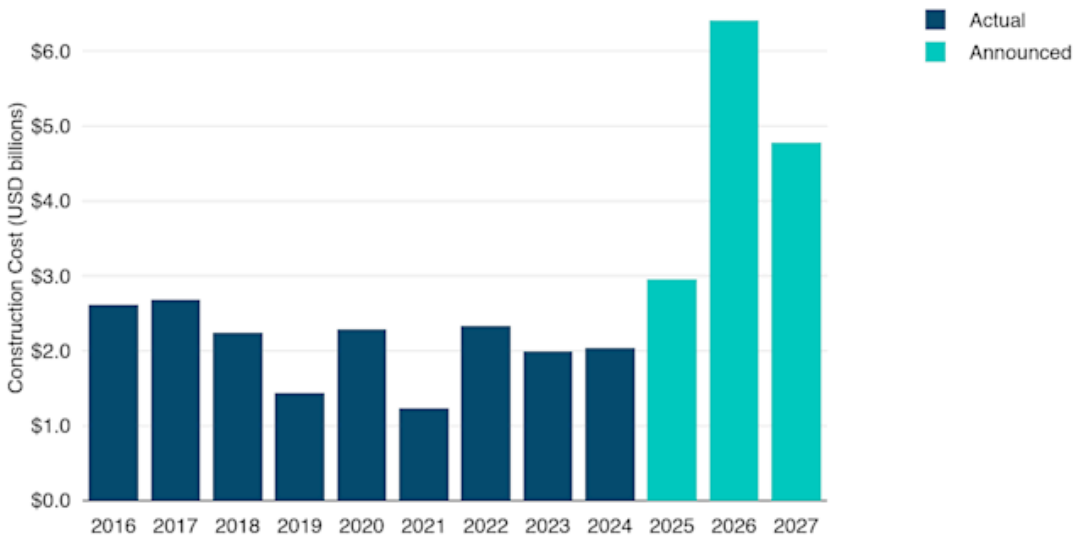

Notes: Total construction costs of all international and domestic submarine cables entering service in designated years. Construction costs exclude the cost of subsequent capacity upgrades and annual operational costs. 2025-2027 construction costs based on announced contract values and TeleGeography estimates. Not all planned cables may be constructed.

Source: TeleGeography

How is this substantial investment in subsea infrastructure being deployed regionally? In the past three years, no single route stands out versus the others, although intra-Asian routes have had the most extensive investment with $1.2 billion in new cables, and Latin American routes have seen the least with about $200 million in new investment.

Look at the coming three years in the horizontal bar chart below. We see a significant surge in new cable investment across every route we track. The biggest surge by far is expected in the trans-Pacific, where an aggregate of over $3 billion in spend will be driven by multiple Google and Meta-led cables along with several others.

FIGURE 6

Construction Cost of New Submarine Cables Entering Service by Region

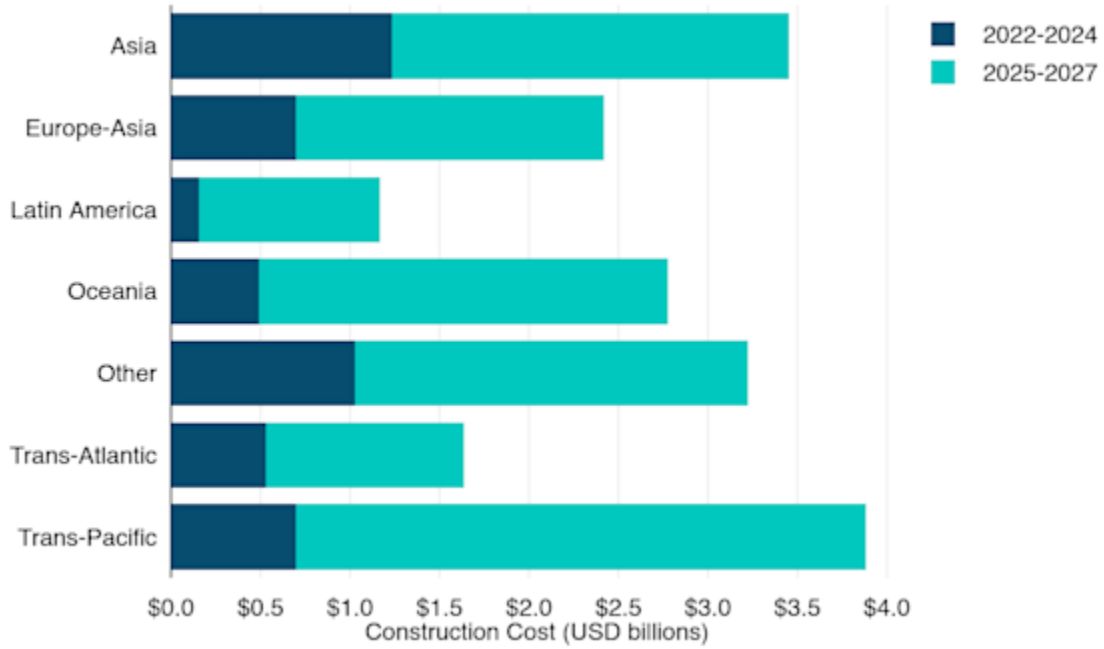

Notes: Construction costs based on the year that the cable entered service. Construction costs exclude the cost of subsequent capacity upgrades and annual operating costs. 2025-2027 construction costs based on announced contract values and TeleGeography estimates. Not all planned cables may be constructed.

Source: TeleGeography

Pricing

Capacity demand has been doubling nearly every three years on many routes. As a result, investment in new submarine cables has surged. In some instances, delays in cable completion (for reasons ranging from geopolitics to supply chain disruptions) has slowed price erosion to single digits. In others, new supply has created intense price competition. But as always seems to be the case, new projects that take advantage of the latest technology impact price. Wavelength prices overall in 2024 continued their steady decline. Between Q4 2021 and Q4 2024, weighted median 100 Gbps wavelength prices across the key global routes below decreased an average of 11% compounded annually.

How these big new investments are impacting each region varies. The figure below maps out the 100 Gbps wavelength price in the dark blue columns and the CAGR price decline in the red circles. What are we seeing?

FIGURE 7

Weighted Median 100 Gbps Wavelength Prices & CAGR Price Decline on Global Routes

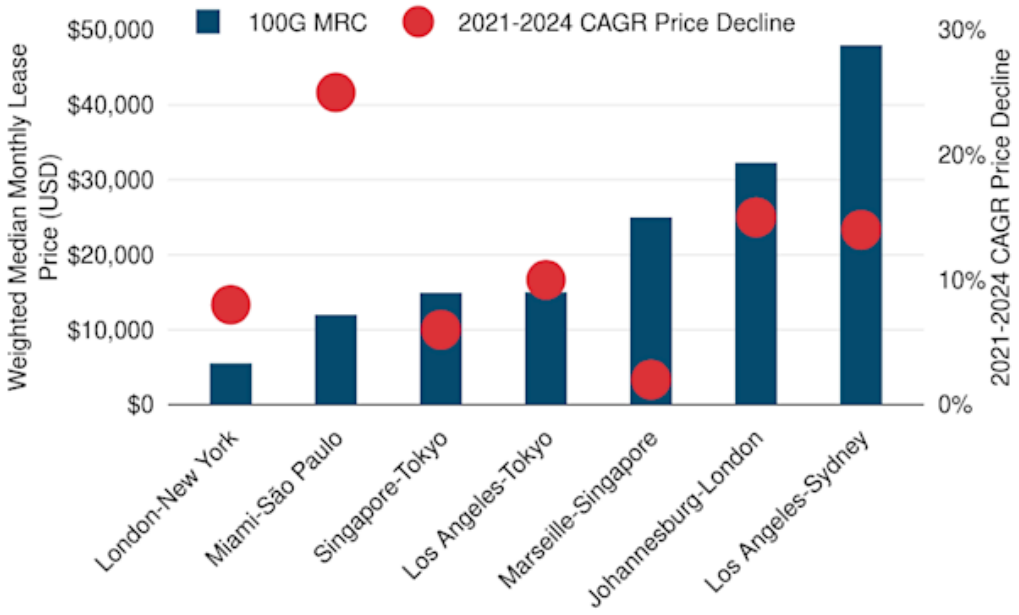

Notes: Each column represents the weighted median monthly lease price for an unprotected 100 Gbps wavelength on the listed route. Circles represent the percentage decline of the weighted median price calculated as a three year compound annual growth rate. Prices are in USD and exclude local access and installation fees.

Source: TeleGeography

On routes with more ample supply, price erosion has returned to form as new high capacity cables enter service. For example, the U.S.-Latin America route continues to fall at a brisk pace, a reflection of the imminent launch of Firmina, diverse fiber pair ownership, and upgrades to existing systems. Between 2021 and 2024, 100 Gbps wavelength prices on Miami-São Paulo decreased 25% compounded annually, to $12,000 per month. Similarly, Johannesburg-London, which saw a massive influx of new capacity from the launch of Equiano, recorded an annual 15% price drop for 100 Gbps wavelengths over the past three years. In Q4 2024, weighted median 100 Gbps wavelength prices on the core Africa-Europe route were $32,272 per month. That's 5.9 times the price of London-New York, compared to 7.5 times more expensive just three years ago.

On other routes, planned cables are just starting to come to fruition and geopolitical hurdles remain. As a result, price erosion is still sluggish. Routes connecting to (and within) Asia are a prime example of this. Between 2021 and 2024 100 Gbps prices on Singapore-Tokyo and Marseille-Singapore decreased just 6% and 2% annually from 2021-2024, respectively. This is beginning to change, though. With delayed systems starting to enter service at the end of 2024, these routes are about to receive a substantial amount of new supply and competitors over the next few years, likely setting the stage for increasing price erosion.

Outlook

Numerous factors are shaping the trajectory of the global bandwidth market on the global, regional, and local levels. The individual analysis section of the *Transport Networks Research Service* delve into detail on developments shaping each of the route groupings and regions we track. But here, let's focus on a few of the most widely-impactful trends shaping long-haul capacity demand and pricing.

Content Providers Expanding Frontiers

Expect to see content providers increasingly drive long-haul network development across all global routes.

Content providers initially focused their route development on major inter-hub routes. This was reflected in heavy investment on trans-Atlantic, trans-Pacific, and intra-Asian routes. They've since expanded their footprints dramatically. Their investments now include new systems linking all global regions including Africa, the Middle East, India, Oceania, and Latin America.

In 2024, Google announced its Pacific Connect Initiative, a collection of cable systems criss-crossing the world's largest ocean. The earliest of these systems will enter service in 2026, with others coming online as they are completed. Although this initiative does not represent Google's only recent cable investments, it represents a massive chunk of capacity and capital. The cost of these systems is well over $1 billion.

Meta's Project Waterworth is another vast and ambitious plan to reach new markets. This group of planned cables will form a loop around the globe, though detailed landing and routing has not yet been made available.

Locations of Data Centers and Power Optimization

Network optimization will drive long-haul bandwidth growth across routes linking data centers.

Operators are beginning to redistribute computing workloads across data centers to optimize power consumption and processing capacity, a strategy known as "spatial temporal load shifting." Google, for example, leverages its global network to shift workloads based on the real-time availability of renewable energy. As other companies adopt similar strategies, locations with abundant renewable energy options will likely become prime targets for future data center and network infrastructure development. These will add to bandwidth demand on a growing intermeshed network of global data center-to-data center routes that prioritize both access to end users and access to compute and storage resources.

Geopolitical Concerns

Geopolitical tensions will directly affect the timing and location of subsea cable deployments.

While geopolitical concerns have always played a role in determining which companies deploy long-haul networks (and in which locations), several recent developments are reshaping network deployment trends.

The Red Sea currently faces major problems. Competing Yemeni territorial claims have created a permitting nightmare for cable laying vessels attempting to operate there, and Houthi attacks on shipping lanes have manifested a physical threat to operations. The massive 2Africa cable, which is largely laid, has yet to realize deployment of a large length of cable that would lie in Yemeni waters. Other cables planned to traverse the Red Sea in coming years include IEX, Africa-1, Raman, and SeaMeWe-6, which are also awaiting deployment.

Cable repair operations have been threatened in the Red Sea as well. The simultaneous cable faults on AAE-1, EIG, and SEACOM/TGN-Eurasia off the coast of Yemen in late February 2024 resulted in 5 months of delay before maintenance vessels were able to conduct repairs due to the threat of Houthi attacks on shipping. These problems in the southern Red Sea are spurring efforts to develop terrestrial bypass solutions from the Saudi Arabian Red Sea coast across to the U.A.E. and Oman.

Beyond the Red Sea, cable builders are finding it increasingly difficult to receive Chinese permits for new cable deployment in the South China Sea. The ADC intra-Asian cable finally entered service at the end of 2024, followed by the SJC-2 cable in July 2025 after multi-year delays. (SJC-2 was originally planned to be in service in Q4 2020!)

These challenges of laying cables in the South China Sea are leading operators to develop new routes.

The builders of the planned Apricot cable aim to avoid this problem by laying the cable from Japan to Singapore via a route to the east of the Philippines. In addition, U.S. government opposition to direct China-to-U.S. cables has boosted the development of several cables from Southeast Asia to the U.S. These include Echo and Bifrost.

Artificial Intelligence (AI)

AI will contribute to shifts in traffic patterns—not just locally—but also in the long-haul.

There's hardly a hotter topic in the network world right now than AI. Just as the broader economy braces for the massive but unknown impact of this technological development, we have to ask what kind of effect AI will have on the transport network. In the U.S., terrestrial fiber providers Lumen and Zayo have both cited AI-driven demand as justification for laying new fiber. Lumen has touted big deals with Google, Meta, Amazon, and Microsoft for metro and long-haul fiber.

The impact of AI on network traffic extends far beyond the networks of major hyperscalers. AI-driven traffic is increasingly interwoven with traditional traffic streams carried across all carrier networks. New AI-focused companies may emerge that seek to build their own network infrastructure.

In conclusion, while it's clear that AI will certainly influence long-haul network demand, the precise magnitude and nature of this impact remain uncertain. The rapidly evolving model training techniques and inference deployment strategies make it challenging to predict what the future will hold.

CHAPTER 03

IP Networks Research Service

Learn more about
IP Networks Research Service

Findings

In our *IP Networks Research Service*, we delve into the core meaning of our extensive internet capacity and traffic datasets. We also explore the factors that influence IP transit pricing and highlight the significant roles played by individual backbone operators. While international internet bandwidth and traffic growth had shown a gradual slowing in previous years, the last couple of years have revealed a leveling out in annual growth. By any measure, bandwidth and traffic growth continue to be brisk. Globally, IP transit prices are still on a downward trend, though it's important to note that substantial regional differences in prices persist.

Internet Traffic and Capacity

Global internet bandwidth increased by 23% in 2025, maintaining its consistent pace of steady growth. Total international bandwidth now stands at an impressive 1,835 Tbps, representing a 4-year compound annual growth rate (CAGR) of 24%. Although the pace of growth has slowed ever so slightly in recent years, bandwidth has still more than doubled since 2021.

Naturally, capacity growth varies across different regions. Once again, Africa led the way with the most rapid growth in international internet bandwidth, expanding at a compound annual rate of 38% between 2021 and 2025. The Middle East followed, growing at a 27% compound annual rate over the same period.

FIGURE 1

International Internet Bandwidth Growth by Region

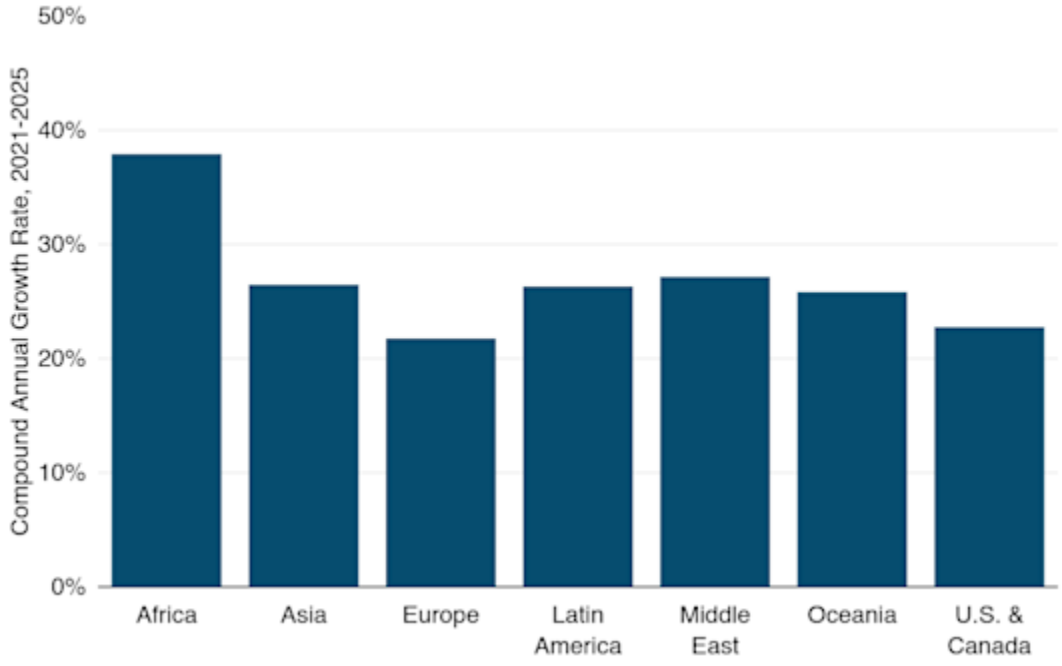

Notes: Data as of mid-year.

Source: TeleGeography

International internet traffic growth largely mirrors that of internet bandwidth. Average international and peak international internet traffic increased at compounded annual rates of 24% and 23%, respectively between 2021 and 2025—essentially identical to the compounded annual growth rate in bandwidth over the same period.

Prices

Providers' shift to predominantly 100 Gbps internet backbones continues to reduce the average cost of carrying traffic and enables profitability at lower prices. As a result, price erosion remains the universal norm, reflecting the introduction of competition into new markets and the response of more expensive carriers to lower prices.

Trends in the IP transit market generally mirror regional trends of the transport market, which is driven by new infrastructure and technologies that enable greater scale and efficiency. Because IP transit relies on these high-capacity links, when transport prices fall we usually see IP transit prices follow. And while IP transit prices continue to decline around the world, the larger story in the most established hubs has been the bottoming out of reported low prices, as a growing number of carriers position themselves at the lower end of the price range.

Across the cities included in the figure below, 100 GigE prices fell 12% compounded annually from Q2 2022 to Q2 2025. In Q2 2025, the lowest 100 GigE prices on offer in the most competitive markets remained steady at $0.05 per Mbps per month. The lowest for 10 GigE also held at $0.07 per Mbps per month.

While 10 GigE remains a relevant increment of IP transit, particularly in more emerging markets, its share of the transaction mix continues to yield to 100 GigE. In 2025, providers indicated that a majority of their sales mix in key U.S., European, and Asian hubs were now 100 GigE. On average, across the cities noted, the Monthly Recurring Charge (MRC) for a 100 GigE port was 6.6 times the MRC for a 10 GigE port.

FIGURE 2

Weighted Median 100 GigE IP Transit Port Prices & CAGR Price Decline in Key Global Cities

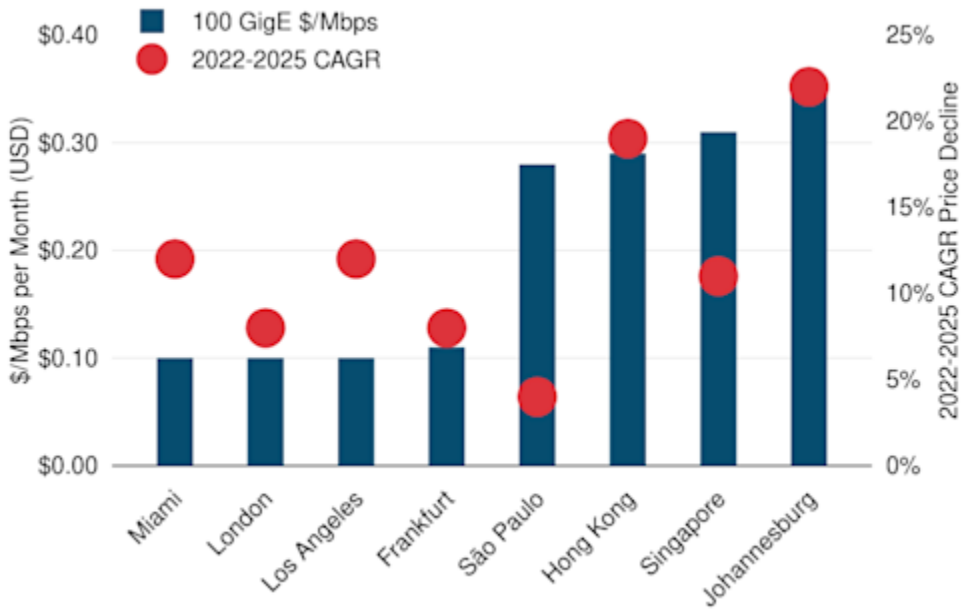

Notes: Each column represents the weighted median monthly price per Mbps for an IP transit port in the listed city. Circles represent the percentage decline of the weighted median price calculated as a three year compound annual growth rate. Prices are in USD and exclude local access and installation fees. 100 Gigabit Ethernet (100 GigE) = 100,000 Mbps.

Source: TeleGeography

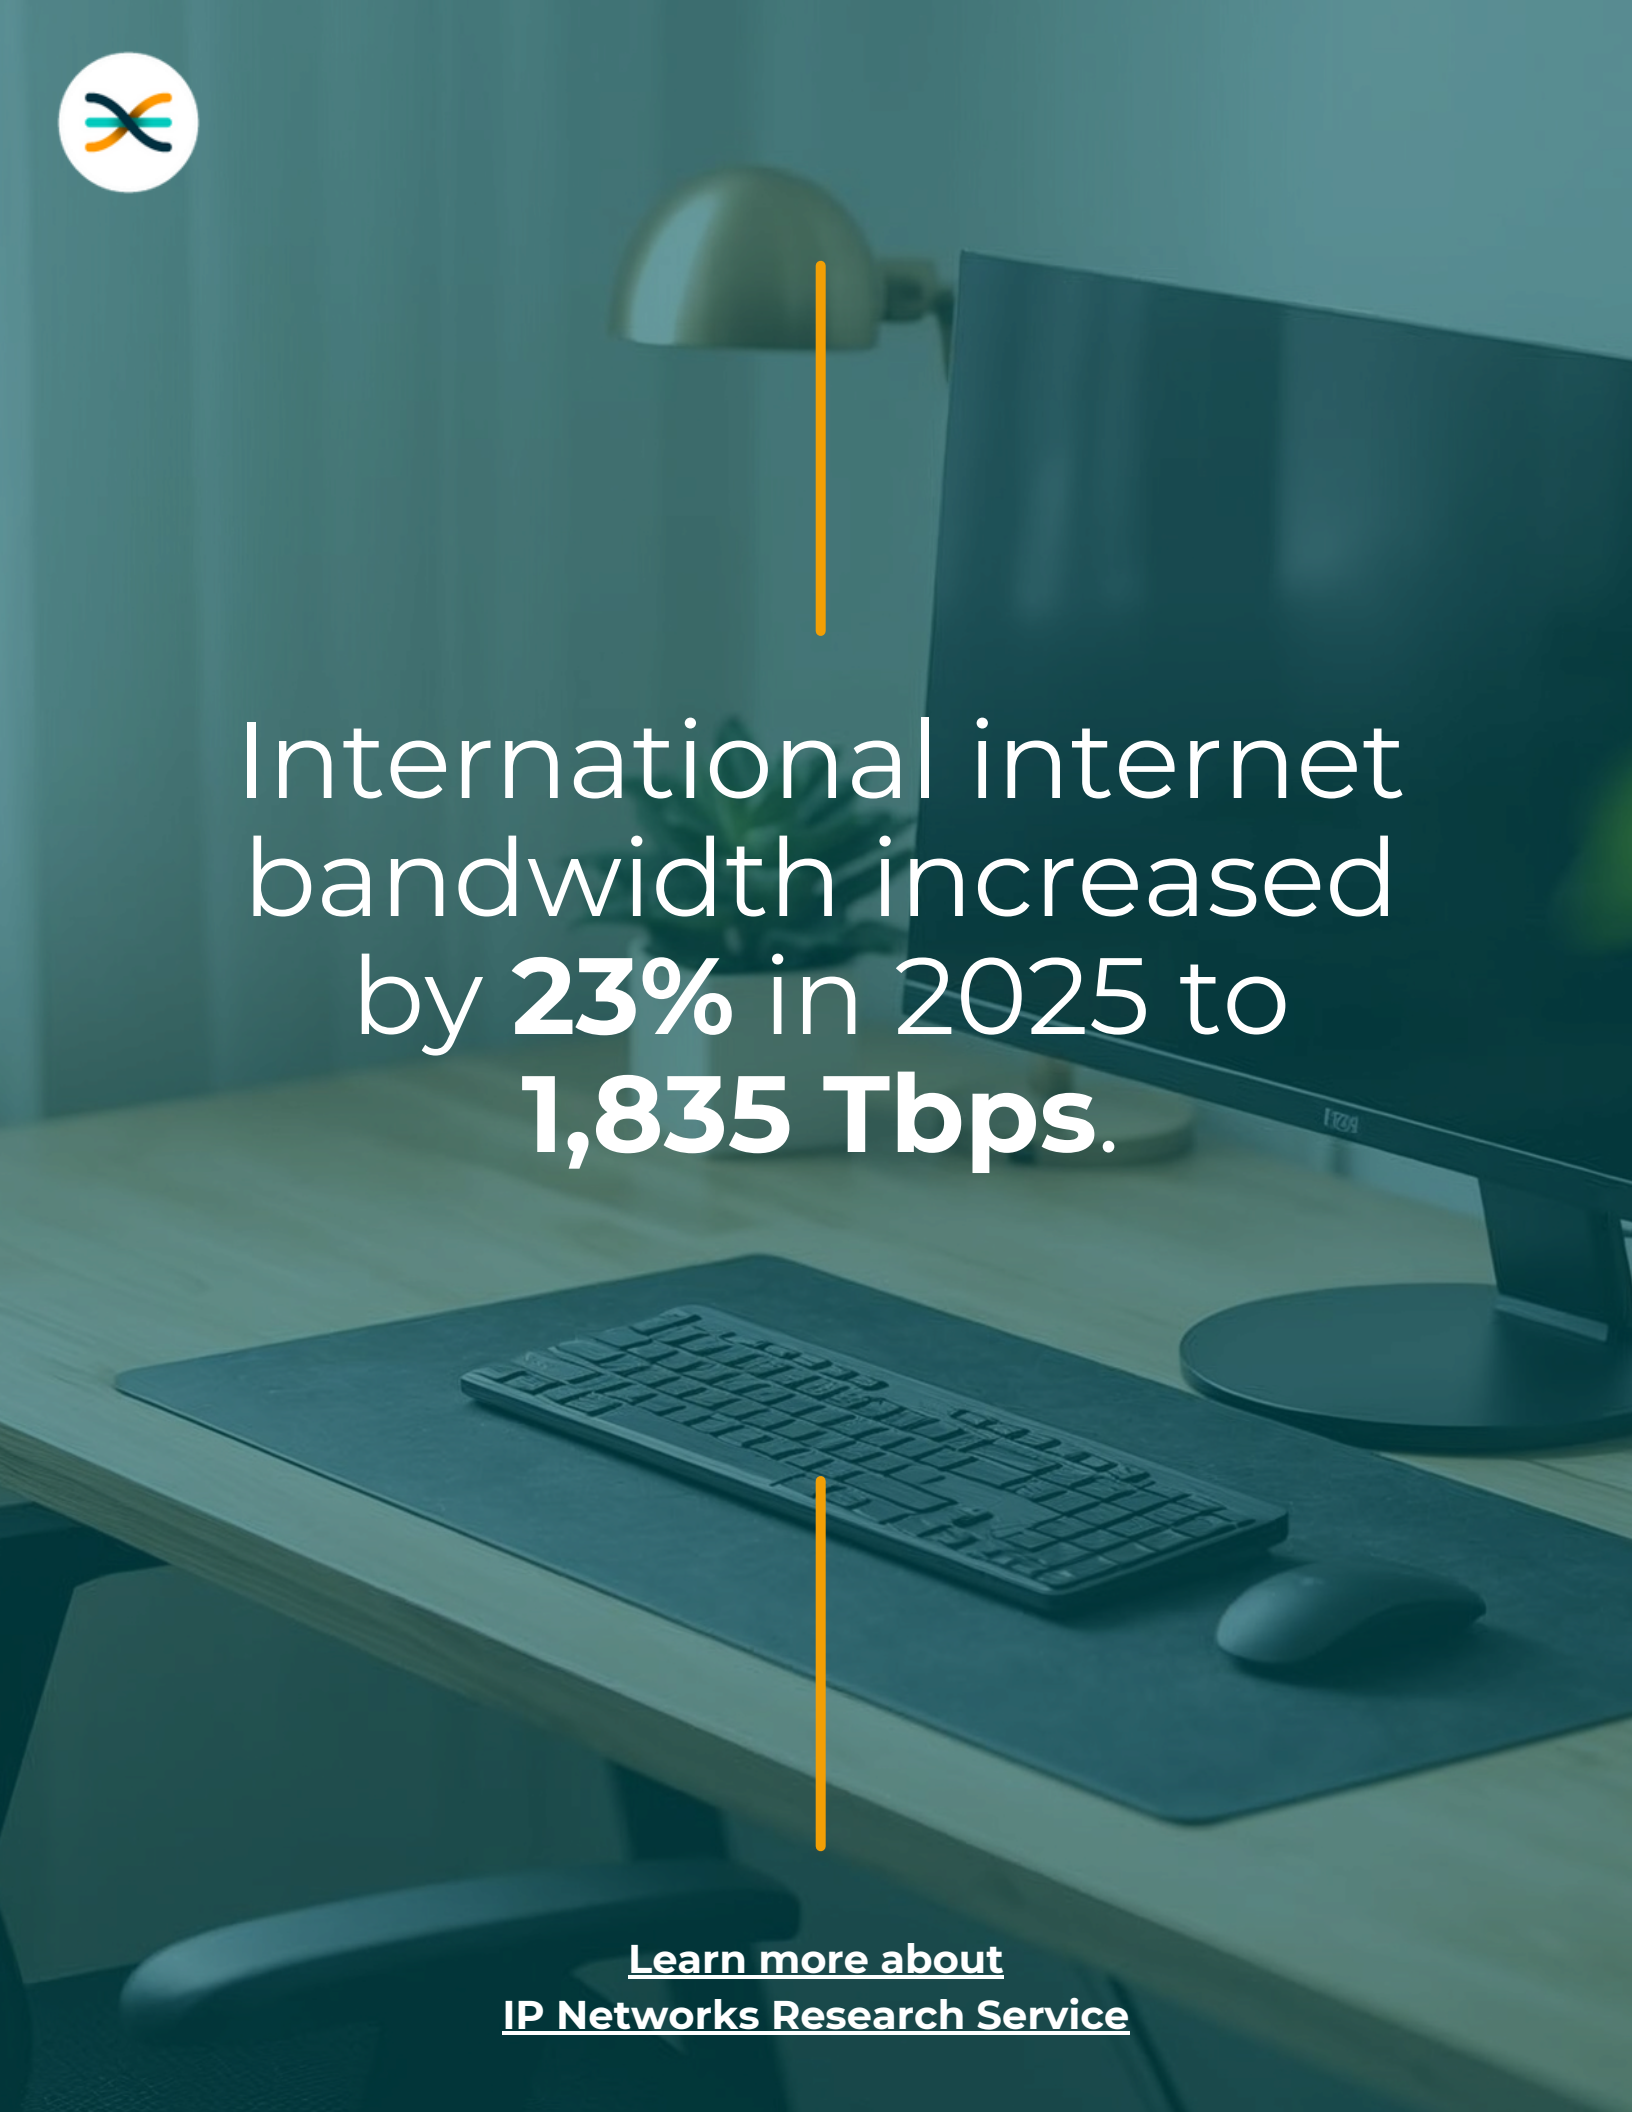

International internet bandwidth increased by **23%** in 2025 to **1,835 Tbps**.

Learn more about IP Networks Research Service

The activation of large-scale subsea cable systems drove the greatest price declines in historically expensive developing markets like Africa and South Asia. These high capacity systems reduce unit costs, fostering intense competition. Increased local peering and content localization further improve network performance and reduce the subsea cost component of IP transit.

In developed markets, carriers are in the process of rolling out 400 GigE services across their network and customer demand has started to materialize. At the moment, providers report that the service represents a very small portion of their sales mix (think single digits) and is mainly limited to the largest global hubs in Europe and the U.S. Across key cities in the U.S. and Europe, 400 GigE prices range from $0.08 to $0.09 per Mbps. That's an average of 3.3 times the price of a 100 GigE port across key cities. As networks scale to meet increasing demand, price erosion for IP transit remains a certainty in nearly all markets for the foreseeable future.

Provider Connectivity

Our rankings of provider connectivity includes analysis based on BGP routing tables, which govern how packets are delivered to their destinations across myriad networks as defined by autonomous system numbers (ASNs). Every network must rely on other networks to reach parts of the internet that it does not itself serve; there is no such thing as a ubiquitous internet backbone provider.

If you want a single, simple number to identify the best-connected provider in the world, you may come away disappointed. There are several ways to measure connectivity, and each highlights different strengths and weaknesses of a provider's presence. One basic metric is to count the number of unique Autonomous Systems (AS) to which a backbone provider connects, while filtering out internal company connections. The results are presented in the table below.

Hurricane Electric has experienced consistent gains, and now ranks as the clear number one in terms of connections. Cogent has also experienced steady growth. Lumen and Hurricane Electric had swapped the top spot back and forth for several years. Lumen (the rebranded CenturyLink) experienced huge gains a few years ago when the company bought Level3. Since then, the number of ASNs connected to Lumen has stagnated.

FIGURE 3
Number of Connections for Selected Providers

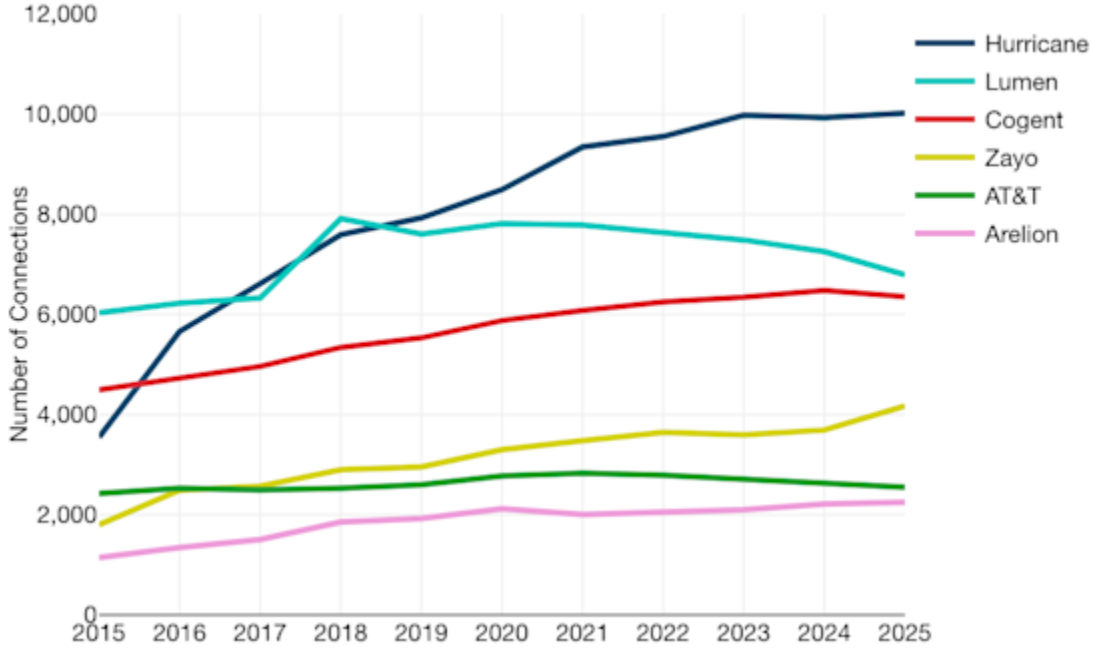

Notes: Data shows the number of connections to other ASNs. The line indicating Lumen's number of connections reflect Level 3 (parent ASN 3356) rather than Lumen (formerly parent ASN 209) prior to 2018.

Source: TeleGeography

In addition to examining overall number of connections, we also used our analysis of BGP routing tables to look at the "reach" (a measure of the number of IP addresses an upstream ASN has been given access to from downstream ASNs) and "share" (which compares an upstream provider's reach to all other upstream providers of a downstream ASN.) The results of this analysis paint a different picture. In some cases, an ISP might end up high-ranked in terms of number of connections but low-ranked in terms of share or reach when the number of IP addresses passed from its customers is relatively small.

Finally, to focus on which backbone providers best serve the end-user ISP market and corporations, we compare upstream provider connections to downstream broadband ISPs, calculated the top providers to Fortune 500 companies, and examined connectivity to specific industry sectors such as hosting, medical, and finance.

Outlook

The combined effects of new internet-enabled devices, growing broadband penetration in developing markets, higher broadband access rates, and bandwidth-intensive applications will continue to fuel strong internet traffic growth. While end-user traffic requirements will continue to rise, not all of this demand will translate directly into the need for new long-haul capacity. A variety of factors shape how the global internet will develop in coming years:

- **Artificial Intelligence (AI).** This is the most hyped demand driver in recent years, but its impact on international internet capacity is not entirely clear. The impact of AI will not solely be felt in the networks of the large cloud providers who are involved in AI, such as Amazon, Google, Meta, and Microsoft. Many companies are offering GPUs-as-a-Service (GPUaaS) which allows anyone access to GPUs to train their own models or use for inference. As users deploy AI models and inference clusters using these GPUs, the traffic will end up running through the networks of many operators, not just those of the major cloud providers.

- **Bypassing the public internet.** The largest content providers have long operated massive networks. These companies continue to experience more rapid growth than internet backbones, and they are expanding into new locations. Many other companies, such as cloud service providers, CDNs, and even some data center operators, operate private backbones that bypass the public internet. As a result, a rising share of international traffic is carried by these networks.

- **International requirements of new applications.** While there's little doubt that enhanced end-user access bandwidth and new applications will create large traffic flows, it is not clear how much of this growth will require the use of international links. In the near-term, the increased reliance on direct connections to content providers and the use of caching will continue to have a localizing effect on traffic patterns and dampen international internet traffic growth.

- **IP Transit Price Erosion.** International transport unit costs underlay IP transit pricing. As new international networks are deployed, operational and construction costs are distributed over more fiber pairs and more active capacity, making each packet less expensive to carry. We already see a major shift from 10 GigE requirements to 100 GigE requirements, and expect that 400 GigE will emerge in two to three years as a significant part of the market. The introduction of new international infrastructure also creates opportunities for more regional localization of content and less dependence on distant hubs. As emerging markets grow in scale, they too will benefit from economies of scale, even if only through cheaper transport to internet hubs.

CHAPTER 04

Data Center Research Service

Learn more about
Data Center Research Service

Findings

The data center market continues to wrestle with an unprecedented shift that started—publicly at least—six years ago with the announced moratoria on new development in Singapore and Frankfurt. Those regulatory interventions shone a light on the challenges the industry faces in securing ample power. The following years proved that data centers would need much more power than had been realized, at a time when regulators and community organizations were growing increasingly critical of the sector's existing demand.

The acceleration of artificial intelligence (AI) is taking this problem to a whole new level. The data center industry has rallied to implement innovative power and and cooling solutions while securing new energy partnerships. Meanwhile, utilities struggle to update power grids in the face of growing public opposition to new energy infrastructure.

Let's delve into some of the findings from our 2025 *Data Center Research Service* update.

Market Connectivity Score

First, we'll take a brief look at a few of the markets highlighted in our Market Connectivity Score (MCS) tool. For those who subscribe to the DCRS, be sure to browse the latest updates in this interactive platform, which pulls from our wide array of research areas to evaluate the health of more than 3,000 metropolitan markets worldwide. The MCS can be found in the Data section.

London: As of October 2025, London ranks as the strongest interconnection market in the world. London is particularly dominant in the carrier category. No other market has more direct connections to international cities or a larger ecosystem of transport and WAN service providers. On the infrastructure side, London is a top 5 data center and IX market. It’s high ranking as a connectivity hub is bolstered by its demographics. A global powerhouse economy, the city is home to more Fortune 5000 companies than almost any other market.

Singapore: Singapore is indisputably one of the world’s strongest interconnection hubs, but it's not without weaknesses. The first market to introduce a moratorium on new data center development in 2019, its government has struggled with increasing concerns over the power used by the local data center sector. It may surprise some to see that the MCS ranks Singapore not just among the largest, but also among the *fastest-growing*, markets in Asia. The reason is the city-state’s subsea sector. Not only does it have a massive concentration of cables already, but more new cables are slated to land there in the near future than anywhere else in the world.

New York: The New York metropolitan area is a top 10 global market for connectivity, and second only to Washington (Northern Virginia) as the highest ranked U.S. market as of this writing. Although New York doesn't stand out in any single broad MCS category, the city

is strong enough in several categories to rank highly overall. Where New York shines most is its macroeconomic factors. The market is second only to Tokyo for the number of global 5000 companies headquartered there, and its population is greater than nearly any other major communications market. Interestingly, New York also ranks higher than other U.S. markets for remoteness from other hubs. As emerging markets vie to draw market share from larger hubs, remoteness from other centers of gravity can be critical.

Power Scarcity

Formally and informally, operators indicate that markets across the globe are seeing increased delays in securing power for new builds. While 1-2 year wait times were typical to secure power, lags now typically extend for 4 years or much longer. Existing generation and transmission resources are proving insufficient to meet the surge in demand for cloud and AI computing.

One key and growing response to the crisis has been to deploy new data centers closer to the source of power generation.

AWS has proactively taken this approach in a few different markets, most recently with its purchase of the Cumulus data center campus that's colocated with the Susquehanna Steam Electric Power Station in Salem Township, PA.

Google has penned an agreement with Intersect Power and TPG Rise Climate to build solar power generation alongside new data centers.

Microsoft has partnered with Constellation Energy to revive part of the Three Mile Island nuclear power plant (now dubbed Crane Clean Energy Center) to fuel data centers.

Other operators are building out their own energy infrastructure. Texas Tech University System and Fermi America are planning an ambitious energy campus in the Texas panhandle to fuel large-scale AI data center development.

The data center industry is anxiously awaiting the commercial availability of small modular reactors (SMRs) as a potential source of ample power generation. Such technology would allow for nuclear reactors of around 300 MW to be deployed at data center campuses, providing dedicated power to the sites. Early development of SMRs is already underway in several countries.

Dominion Power is exploring the feasibility of building a SMR at the North Anna Power Station, where AWS is investing in a new campus.

Google has partnered with Kairos Power to supply SMRs at several of its sites by 2035.

Equinix and Switch signed purchase agreements for SMR capacity from Oklo.

In mid-2025, Equinix announced additional agreements related to SMR deployments with Radiant, Rolls Royce, and Stellaria.

AI's Effect

Compounding the existing constraints on the data center market, we've reached a major inflection point in the development of artificial intelligence (AI), and more specifically in the development of generative AI. AI will require many changes within data centers. The most immediate concerns include provisioning higher density cooling and higher capacity interconnection.

In response to rapidly increasing demand for scalable AI capacity, a new breed of "neocloud" operators is emerging. These are companies that are centered on using GPU clusters to conduct AI and other high-performance computing workloads. Two entities in particular are at the forefront of providing data center infrastructure for neocloud applications—CoreWeave and Stargate.

Coreweave offers cloud-based AI workload access at more than 30 zones across the U.S. and

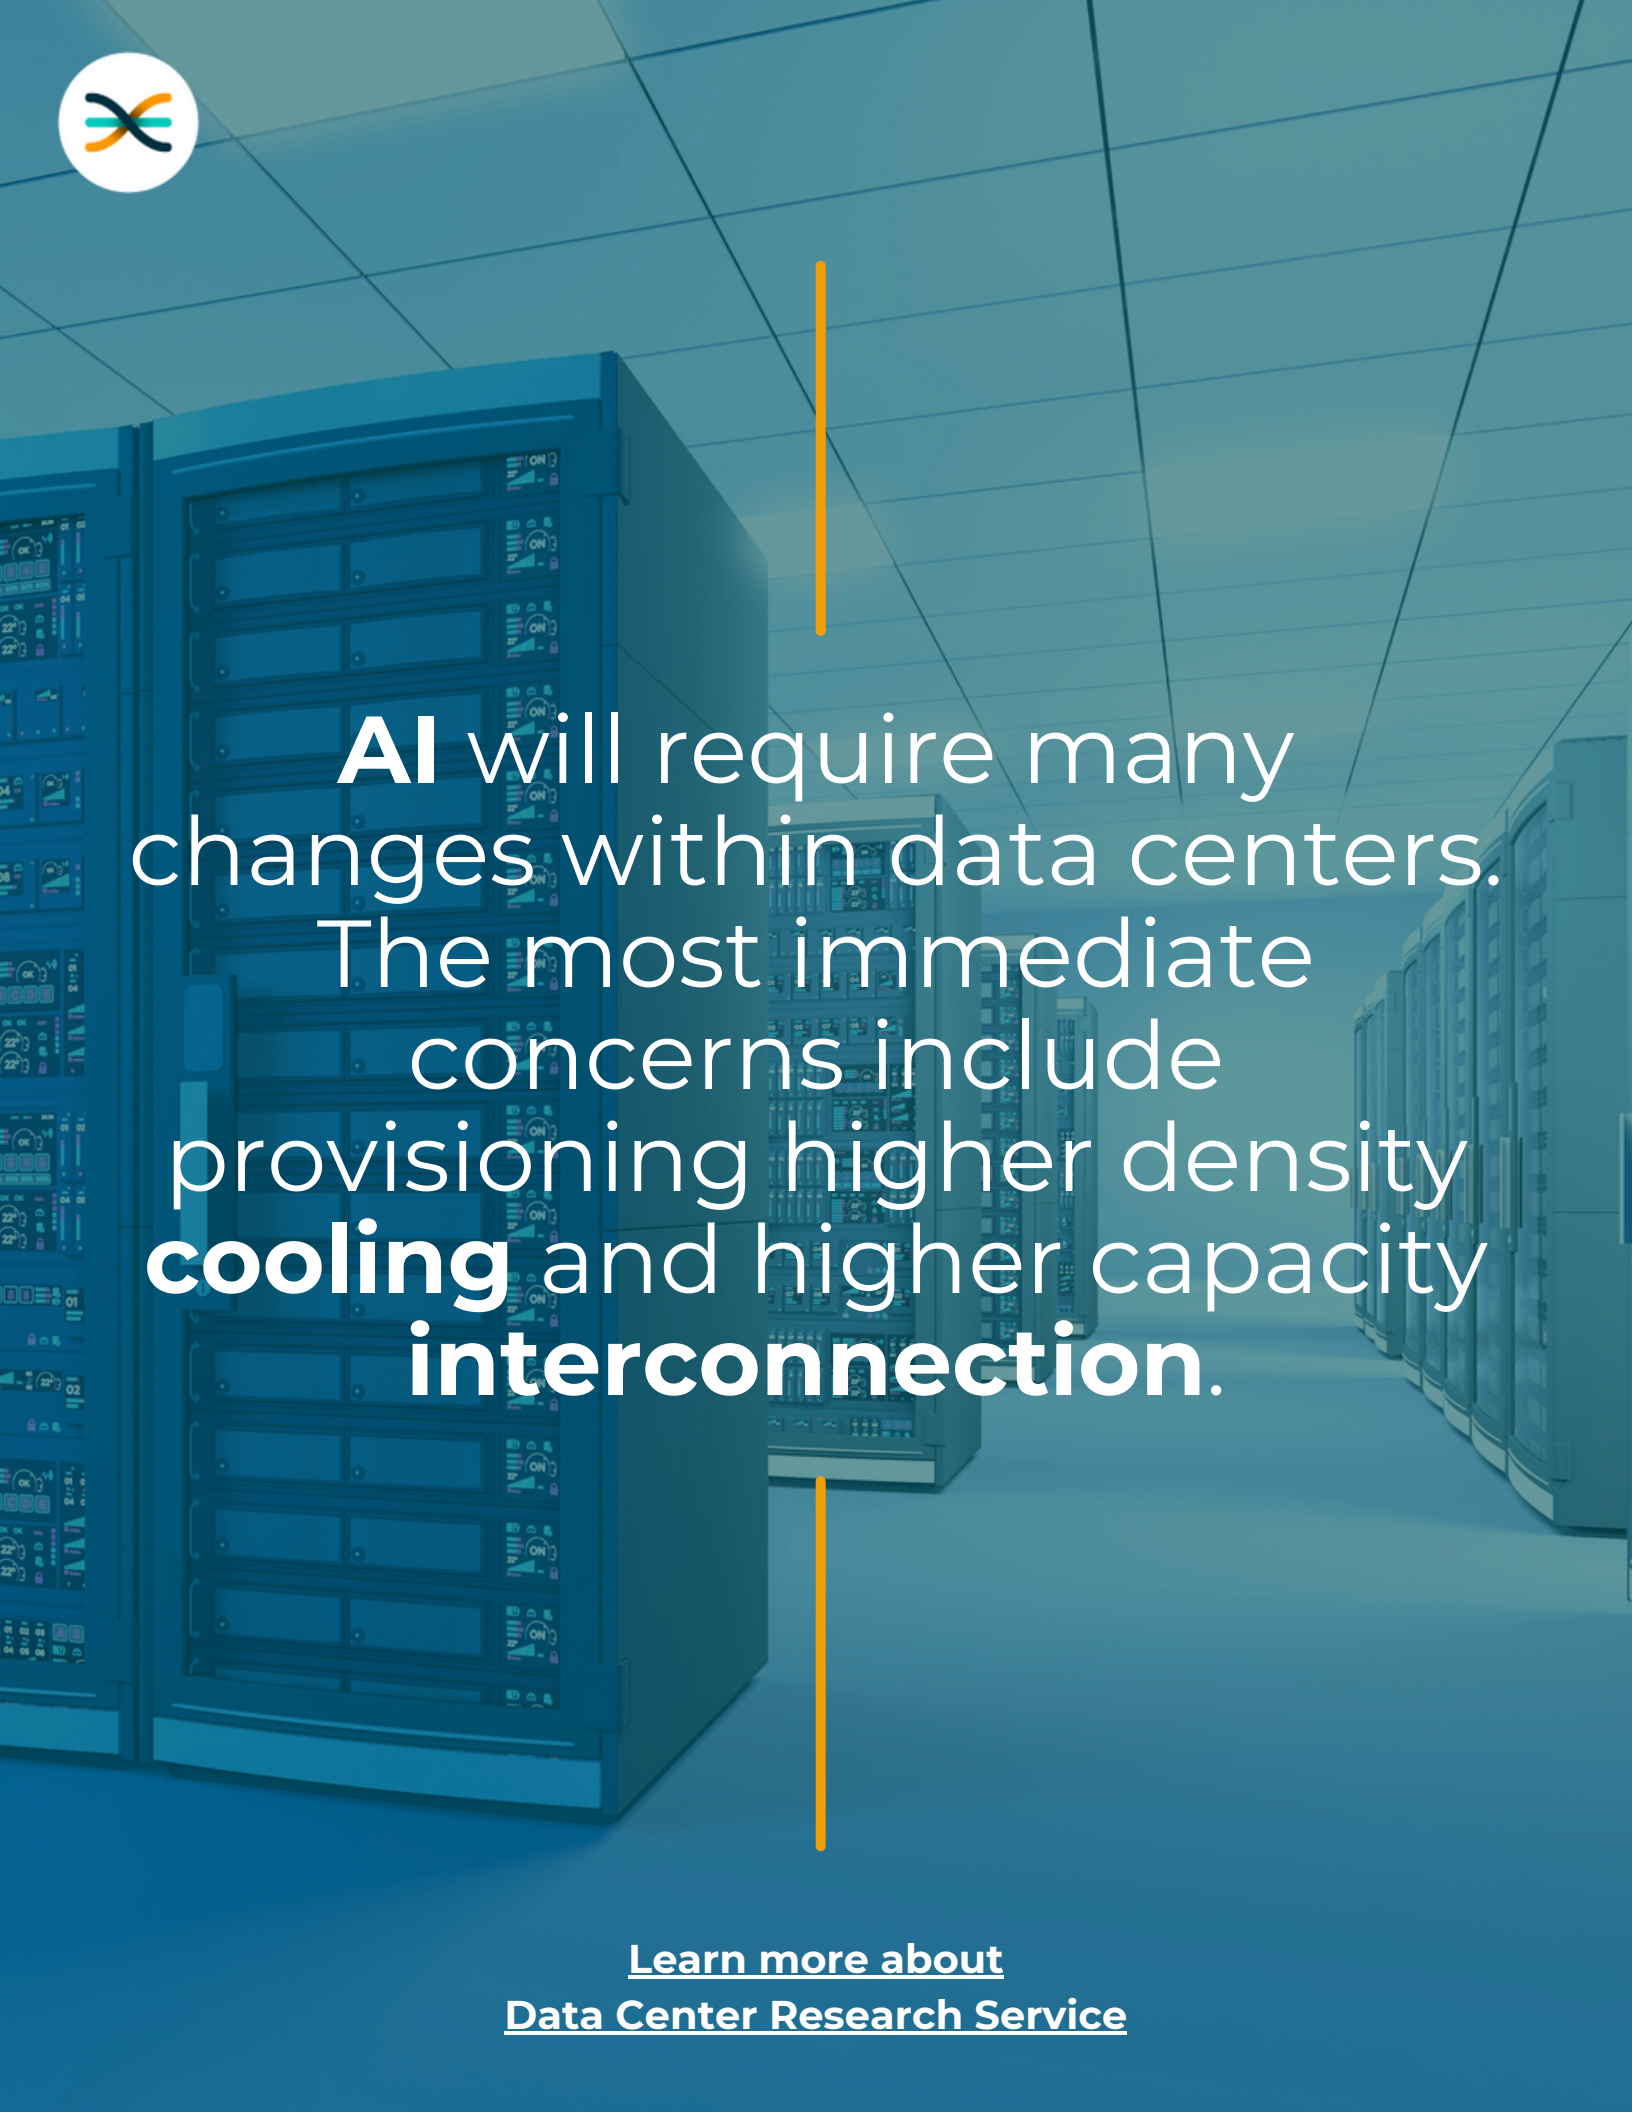

AI will require many changes within data centers. The most immediate concerns include provisioning higher density **cooling** and higher capacity **interconnection**.

Learn more about
Data Center Research Service

Europe.

Stargate—which is spearheaded by OpenAI in partnership with SoftBank, Crusoe, Oracle, and other partners—has extensive pipeline plans at several locations in Texas, and at other U.S. and international nodes.

As of 2025, numerous colocation operators tout AI-readiness and openness in accommodating advanced applications. But being "AI-ready" is tough to execute on the scale that's demanded, especially as effective targets for AI readiness continue to rapidly move.

Critical challenges have to be addressed to effectively deploy AI. These include boosting power, reconfiguring PDU block and rack designs, and providing space to accommodate liquid cooling infrastructure.

These changes need to be made while also providing proper support for existing customers in increasingly hybridized environments.

Data Center Developments

The Washington metropolitan area—or more specifically Northern Virginia (NoVA)—dominates as the world's largest data center market. As of 2025, we estimate its gross commercial data center capacity at nearly 25 million square feet of operational capacity.

As of 2025, London edges out Beijing as the third-largest commercial data center market, with 13.5 million square feet of operational capacity. London metro's capacity supports a vast array of network, cloud, and peering infrastructure, making it not only one of the largest data center markets, but one of the largest interconnection hubs on the planet.

As of 2025, U.S. metro areas account for 5 of the 10 largest markets we track. In Europe, only London and Frankfurt make that list.

Digital Realty and Equinix have greater scale and geographical diversity than any other operator. Both control over 30 million square feet of operational data center capacity and have significant presence in every global region. NTT Global Data Centers, the international arm of NTT Data, also has a strong presence across multiple regions, but Equinix is almost twice as large as this provider.

More than 60 commercial data center providers control more than 1 million square feet of operational capacity each. While few compare to the behemoths that top this list, many are growing rapidly, are flush with new investment, and are critical players in global interconnection development.

A handful of cloud and content providers generate unparalleled demand for compute, storage, and traffic. These few providers are effectively primary drivers of global network and data center investment. In the last year alone (between August 2024 and 2025), Amazon, Google, Meta, and Microsoft deployed 16 new data centers globally. Accelerated growth is expected in the near term with at least 55 new hyperscale data center and cloud region locations in active development just among these 4 companies.

Kolkata surged to the top of the list of large fast-growth markets with new deployments over the past year by CtrlS, STT, and NTT GDC. Johor Bahru and Jakarta predictably remain on the list of large and still fast-growing Asian markets, due to their growing status as alternatives to Singaporean growth.

We estimate that there are 26 markets that have more than 1 million square feet of operational commercial data center capacity and 4-year CAGR of at least 10%. Seven of those markets are in Asia, 6 are in Europe, and 5 are in the U.S. Even more extraordinary—4 of the Asian markets are in India alone. Chennai, Kolkata, and Mumbai have all seen commercial data center capacity growth of 20% or more since 2021, and New Delhi has exceeded 13% growth. In Europe, Warsaw, Berlin, and Milan are all seeing robust growth of at least 15% CAGR.

Many markets, small and large, see low growth over time and periods of no new data center builds going into operation. Amsterdam, Los Angeles, New York, and Singapore all have new data center capacity growth of less than 4% CAGR since 2021. In the cases of Amsterdam and Singapore, this is directly related to regulatory and geopolitical conditions.

Digital Realty, Vantage, NTT GDC, and Equinix have been particularly prolific in launching new data center capacity over the past two years. Digital Realty has opened more than 20 new data centers in markets like France, the U.S., Japan, South Africa, Greece, and Brazil to name a few. Vantage has rolled out more than 15 new sites—many on existing campuses—in markets across North America and Europe, alongside a few in Asia (Cyberjaya and Taipei). NTT GDC has deployed new data centers in the same regions as Vantage, with a particular focus on India, where 7 of its 15 newest data centers are located. Among the more than 25 new data centers Equinix has deployed in the past two years, several are in markets that are new for the company.

Among operators tracked in our database, nearly 500 commercial data center sites are known to be in the construction pipeline right now. While this development is spread across the globe, the region comprising the U.S. and Canada outpaces other regions with the largest percentages of new site deployments, followed closely by Europe.

Power

We estimate that, as of 2025, commercial operators in the top ten data center markets by MW consume about 16 gigawatts (GW) of power. That's enough power to generate electricity for roughly 12 million homes—or, in this case, only about 1,000 commercial data centers!

Only 22% of data center sites reporting are able to provision high-density aisles that exceed 20 kW per rack. This is troubling, especially when considering the fact that AI applications will require density levels in the range of 45 to 80 kW per rack—far in excess of traditional standards of high density.

Connectivity

The most common companies across global colocation facilities are Cogent, AT&T, Lumen, Zayo, and Verizon. With the sale of Lumen's EMEA and Latin American assets to Colt and Cirion, its portfolio has now shrunk to include almost exclusively U.S. assets. Carriers like China Telecom, NTT, Cirion, Vocus, du, and MTN Group are ubiquitous in their respective regions.

By our estimates, CoreSite's LA1 data center in Los Angeles is the most carrier-dense colocation site in the world. Equinix's Kleyerstraße 90 site in Frankfurt and its SG1 site in Singapore are also central nodes of international internet connectivity.

DE-CIX Frankfurt, ix.br São Paulo, and AMS-IX Amsterdam are among the largest exchanges in the world, each exceeding 14 Tbps of peak traffic volumes. Each year, new peering exchanges come online in almost every region of the globe. 2025 is no exception with new openings in Africa, Asia, Europe, Latin America, Oceania, and North America.

Pricing

Current Trends

Singapore is routinely among the most expensive markets we track. With the new licensing regime in the city-state, supply has become very scarce, with new data center supply growing at just 2.2% CAGR in square footage between 2021 and 2025. Median colocation rates there have been at or above $600 per kilowatt since H2 2023.

In Europe, London prices have crept upward, holding at $450 or higher since H1 2023. Frankfurt rates hit a near-term peak of $500 in early 2023 and have since declined to $450. Rates in capacity-constrained Amsterdam have steadily climbed upward to reach levels comparable to London and

Frankfurt.

Among the markets included in the H2 2025 update are four Latin American metros: São Paulo, Santiago, Santiago de Querétaro, and Bogotá. These high-demand markets currently see median rates comparable to those seen in other global hubs. The latter three markets report medians between $400 and $460 per kilowatt. You can find the full available rates for these and all other markets under "Pricing Data" in the H2 2025 pricing data deliverable.

Despite sometimes dramatic price fluctuations between periods, U.S. markets almost always have higher cross connect rates than European markets. As of H2 2025, the average U.S. cross connect price across our market sample is 3 times as high as the average European rate. Cross connect prices for Asian metros generally fall between the rates of U.S. and European markets.

In our study, we model TCO for colocation rates assuming the average monthly cost of a cabinet with either one or five fiber cross connects. Among markets covered in our H2 2025 pricing update, European cities come out more costly than U.S. markets, and Singapore soars in price above the others.

In U.S. markets, where a single cross connect can count for as much as 20% of TCO, going to five cross connects means a 70% jump in TCO. It also means that those cross connects can count for at least half of TCO. In European markets, conversely, the low cost of cross connects relative to base colocation results in a more modest 15-30% hike in TCO when moving from one to five cross connects.

Expectations

When asked how much they expect prices to change in the next 12 months, contributors in H2 2025 anticipated moderate near-term inflation for both base colocation and cross connect rates. Respondents across the market sample expected colocation prices to rise around 3-5%, but expectations for inflation in U.S. markets were higher. Universal expectations for cross-connect rate increases were more modest at 0-3%.

Regional power delivery difficulties, government actions, and lags in wholesale contract renewals will continue to complicate the relationship between industrial consumer rates and the power prices passed to data center customers. In the U.S., data centers' outsized impact on industrial electricity demand growth will likely drive up rates for colocation customers.

The biggest ongoing concern for data center operators remains the availability of power and space to develop. The scramble for resources needed by power-hungry AI clusters will more directly affect wholesale and hyperscale pricing in the near term, but this capacity constraint is bound to hit the colocation market as well.

CHAPTER 05

Cloud & WAN Research Service

Learn more about
Cloud & WAN Research Service

Findings

Cloud Geography 2024 Developments

Globally, 13 new cloud regions were launched in 2024. This marks a significant slowdown in new cloud region deployment (26 were launched in 2023.) There hasn't been a drop below 20 new regions in a year since 2017, and the 2024 figures represents just a third of the number of regions launched at the peak in 2019. Microsoft opened the most regions (4) followed by Google, Oracle, Huawei and OVH with 2 new regions each, and AWS with a single region. Alibaba, IBM and Tencent did not launch any regions in 2024. Regionally, Asia led last year with 3 new regions, Africa, Latin America and Europe, US & Canada launched 2 each, Oceania and the Middle East launched a single region each.

Asia added the most new cloud regions in 2024—a mere 3. This is the same number as in 2023, and down from 8 the year before. Oceania launched 1 new region, down from 2 the previous year. Outside of Brazil, Latin America had no cloud regions until recently. In the past two years the region has grown at a fast pace. The region tied for second place with Africa in terms of new regions launched. Continuing its momentum from 2023, 2 additional regions launched in Mexico (2 had opened in 2023). This brings the country's total region count up to 7. Africa remains the most underserved region of the world—until recently only South Africa contained any cloud regions. Two regions were launched in 2024, including the first region outside of South Africa—in Egypt.

Europe added the most new cloud regions in 2023, with 9 new regions launched. Most of these regions opened in Western Europe, and were meant both to address continuing demand, and to offer a new breed of data sovereign regions. But in 2024 only 2 regions launched. Generally an underserved market for cloud services, the Middle East has been active in cloud infrastructure development the past couple years, but not in 2024. After launching 6 new regions in 2023, only a single region went live in 2024—Oracle's single zoned region in Riyadh. Finally, like Europe, the U.S. & Canada only saw two regions launched in 2024—by Microsoft in the U.S. and OVH in Canada. The U.S. & Canada is a very mature market in terms of the number of cloud regions.

Local Zones

AWS first announced a new element to their cloud network infrastructure in early 2022, the *local zone*. With increased focus on moving cloud services closer to end-users, the local zone provides low-latency access to applications running closer to customers. Google offers a similar service with their dedicated interconnect service, an option to connect to a low-latency colocation facility that provide latencies of less than 5 milliseconds between the workloads in a specified region and the on-ramp location. OVH launched the same local zone service in the first quarter of 2024 and Microsoft has

joined in as well with their extended zone service launched the summer of 2024.

FIGURE 1
Local Zone and Low Latency Locations

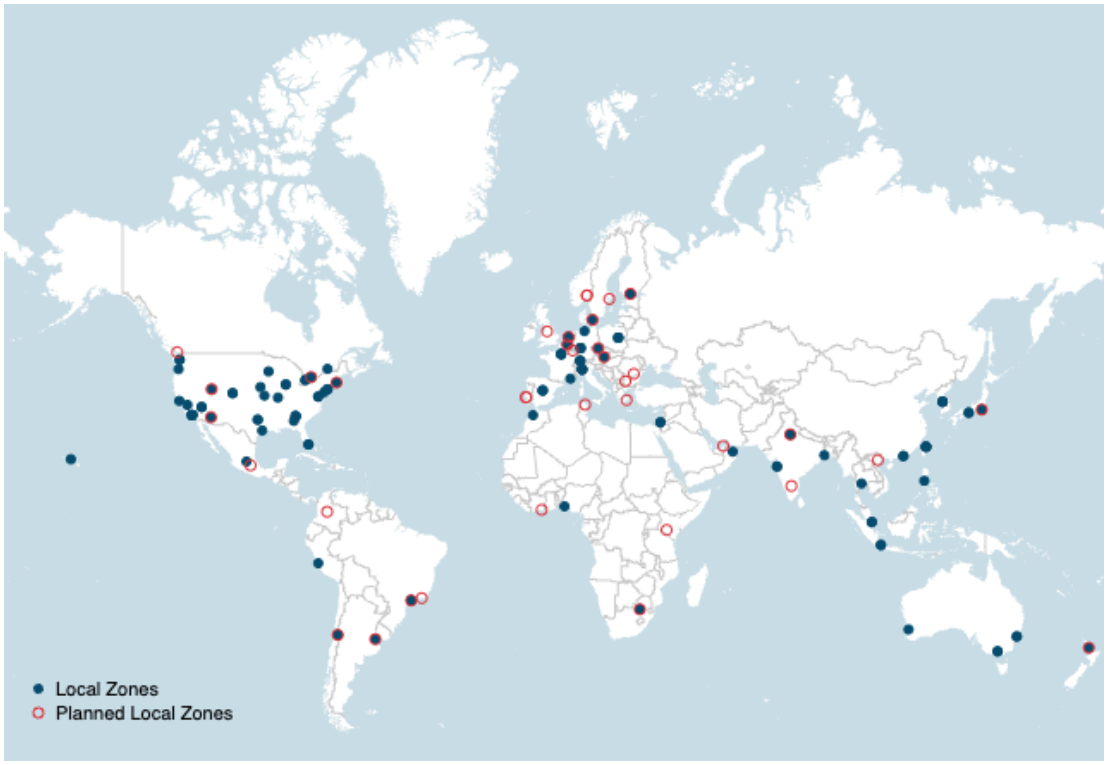

Notes: Data only include IaaS cloud providers AWS, Google Cloud, OVH Cloud and Microsoft Azure. Data as of Q1 2025.

Source: TeleGeography

AWS' local zones differ from Google's low-latency facilities in that they are often launched in facilities far from the zone's parent region, in effect becoming a sort of mini-cloud region for markets that may not have a standard cloud region. Take, for example, their local zones in Africa. AWS currently has a single cloud region, located in South Africa. AWS has launched local zones in Nairobi and Lagos. These zones will serve a select set of cloud services locally in addition to being linked to the parent region in South Africa. OVH's local zone service closely resembles AWS' launching new local zone, small scale regions, in markets where they do not currently have regions.

AWS has launched 34 local zones so far—17 in the U.S. and an additional 17 outside the U.S. The company has more than 14 zones planned to launch over the next year or so. Meanwhile Google offers low-latency colocation facilities in more than 70 markets where they also offer their dedicated interconnection service. OVH launched 19 local zones, primarily in the Europe and the U.S. in 2024 and have another 22 planned for 2025. Microsoft's single extended zone is live in Los Angeles, with a second planned for Perth, Australia.

Cloud Connectivity Services

Cloud services have become a critical component of many enterprises' data management. How enterprises reach the cloud service providers' data centers has become an important issue. Traditionally, the plain old internet sufficed. But there's more than one way to skin a cat. Companies seeking better performance may peer with cloud service providers (CSPs), either through their network service provider (NSP) or directly with the CSP if the company has an autonomous system number (ASN) and meets the CSP’s peering requirements. For better security, companies may instead

choose to connect via IPSec VPNs, tunneling through the public internet.

Still other companies may have high-capacity requirements and business-critical applications in the cloud. For these businesses, cloud services cannot be left susceptible to the performance of the public internet. For them, cloud service providers (CSPs) and their carrier and colocation partners offer *dedicated links* to CSP networks. These links effectively extend an enterprise's network into the cloud provider's network, thus bypassing the public internet.

FIGURE 2

Global Cloud Data Center and On-Ramp Locations

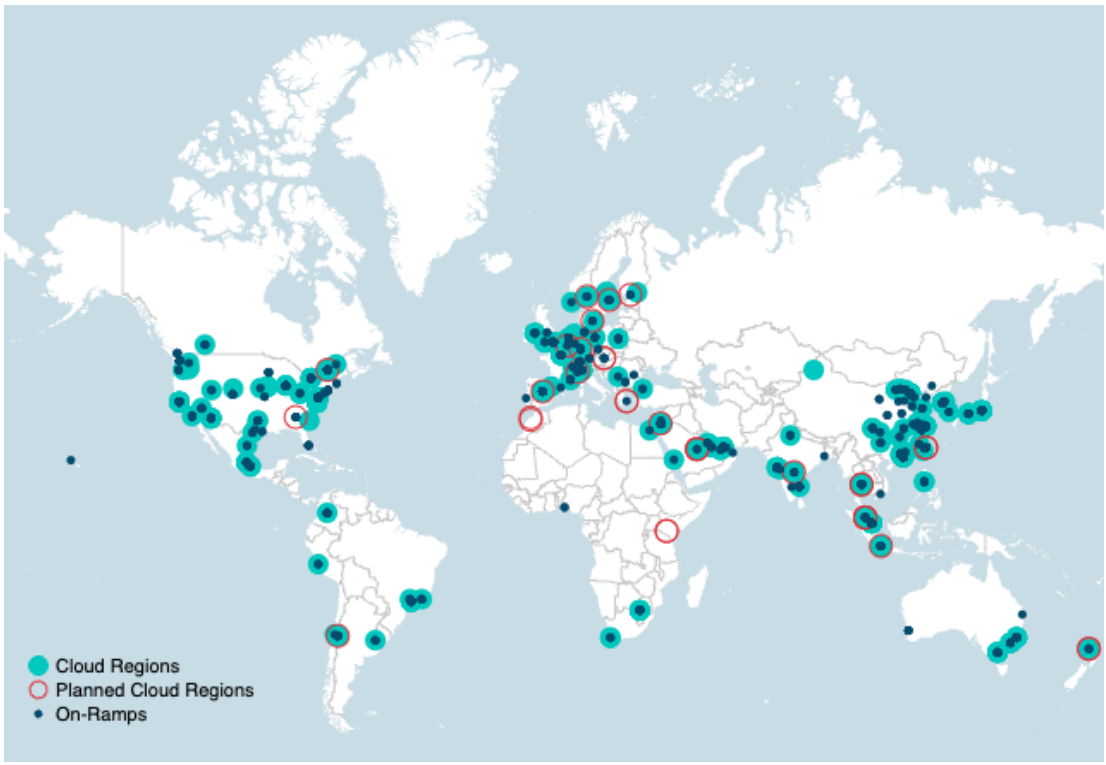

Notes: Data only include IaaS cloud providers from Alibaba, AWS, Google Cloud, Huawei Cloud, IBM, Microsoft Azure, Oracle Cloud, OVH Cloud and Tencent Cloud. Data as of Q1 2025.

Source: TeleGeography

Enterprise network managers have a wide array of service providers to choose from for a dedicated cloud connection service. While enterprises can set up a link directly with the cloud provider, more frequently a third-party (think a carrier, colocation provider, or connectivity specialist) is used. Selection of a provider often depends on the location of the enterprise WAN in relation to the cloud providers' zones or data centers. If a company has routers located within the same colocation facility as the cloud provider, it can often work directly with the cloud service provider to facilitate the direct connection between the networks.

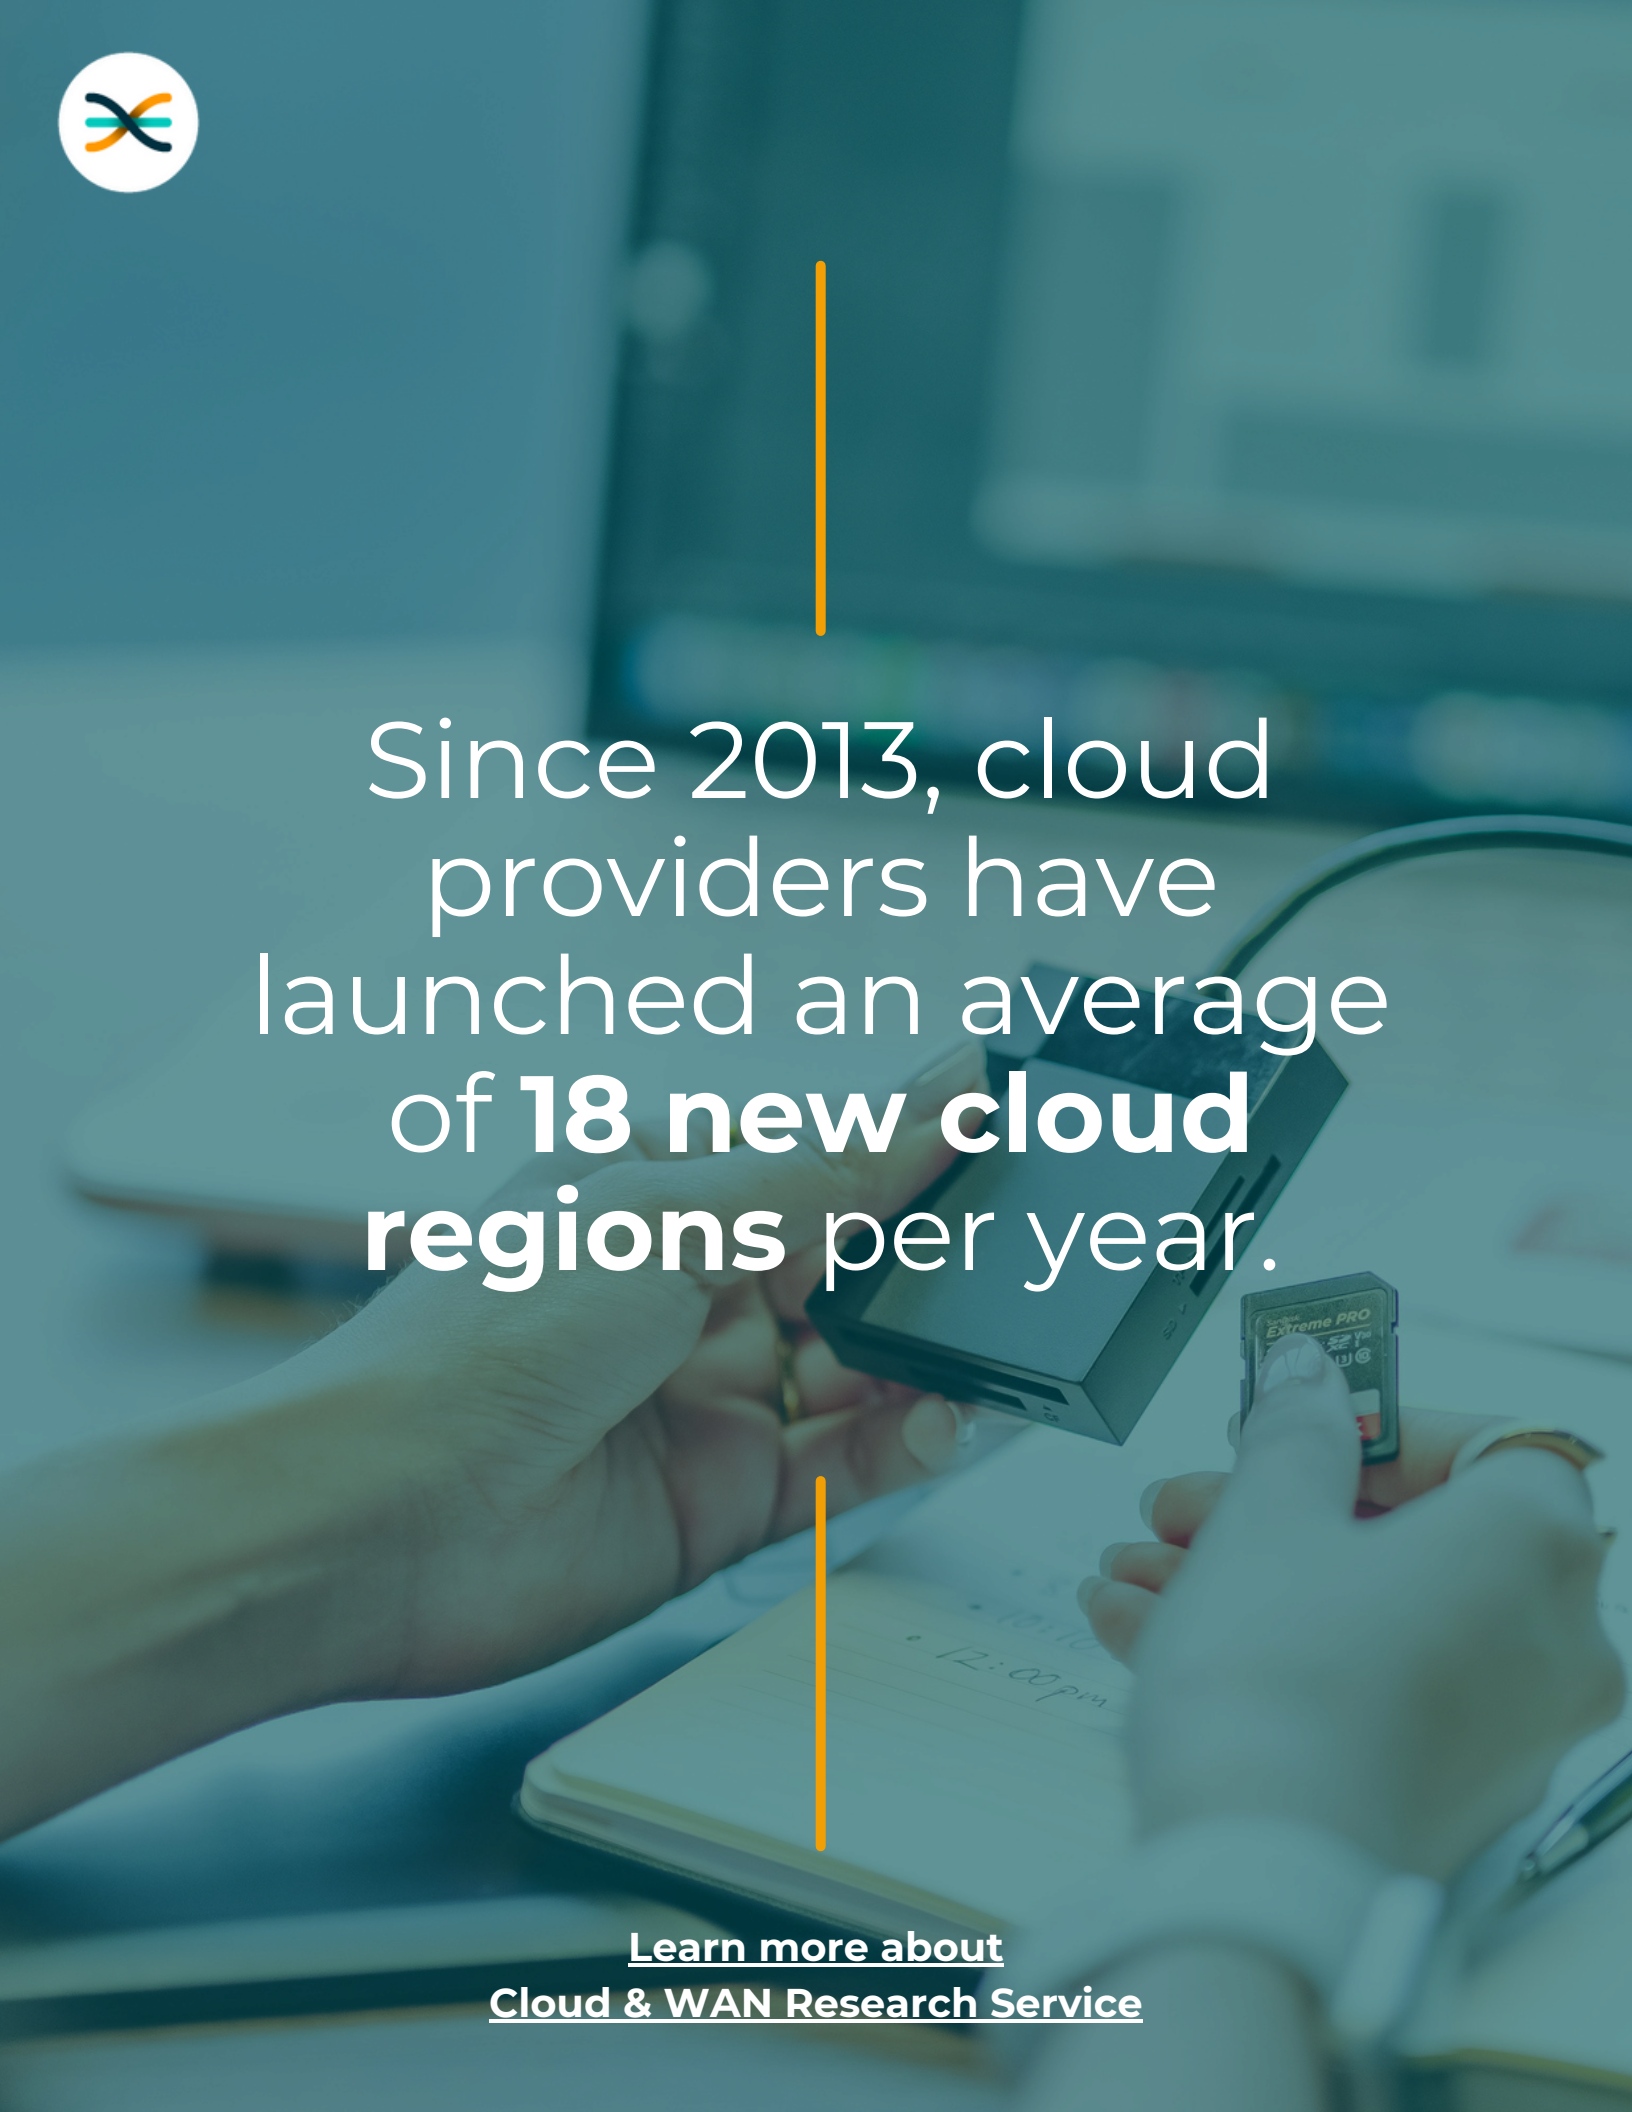

Since 2013, cloud providers have launched an average of **18 new cloud regions** per year.

Learn more about
Cloud & WAN Research Service

FIGURE 3
Cloud On-Ramps by Top Metros

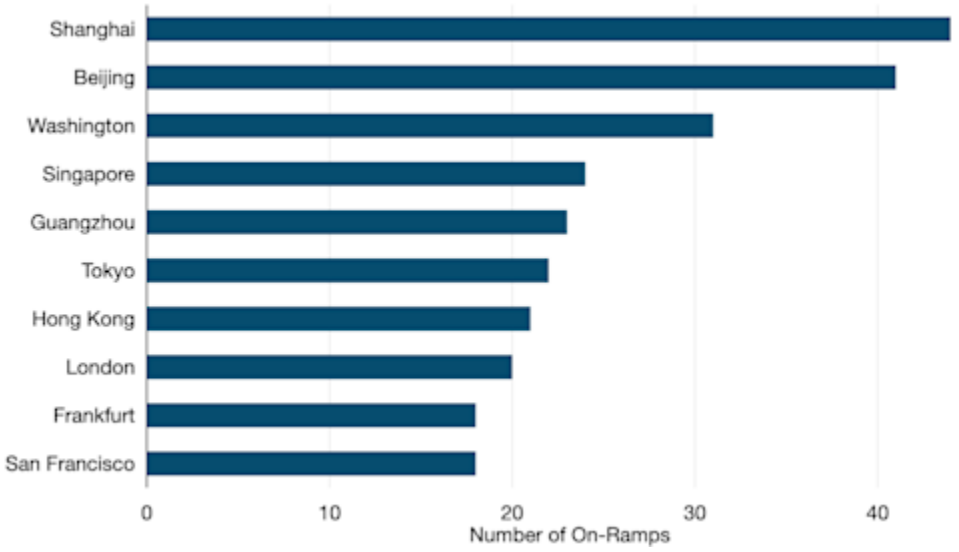

Notes: Data represent the number of dedicated connection on-ramp offerings by cloud service providers at colocation facilities. Data include available information from Alibaba, AWS, Google Cloud, IBM, Microsoft Azure, Oracle Cloud and Tencent Cloud. Data as of Q1 2025.

Source: TeleGeography

With a total count that nears 280, Asia is home to the most in-service cloud zones. Europe follows suit with over 130 in-service zones. Together, these two regions account for 63% of the world's cloud data centers. The remainder are housed in the United States & Canada (20%), Latin America (6%), the Middle East (6%), Oceania (4%), and Africa (2%). At the country level, China and the United States are the clear leaders with close to 160 availability zones for China and 130 for U.S & Canada. Japan, Australia and India round out the top 5, but are home to only between 20 and 30 zones each.

FIGURE 4

Service Availability Zones by Region

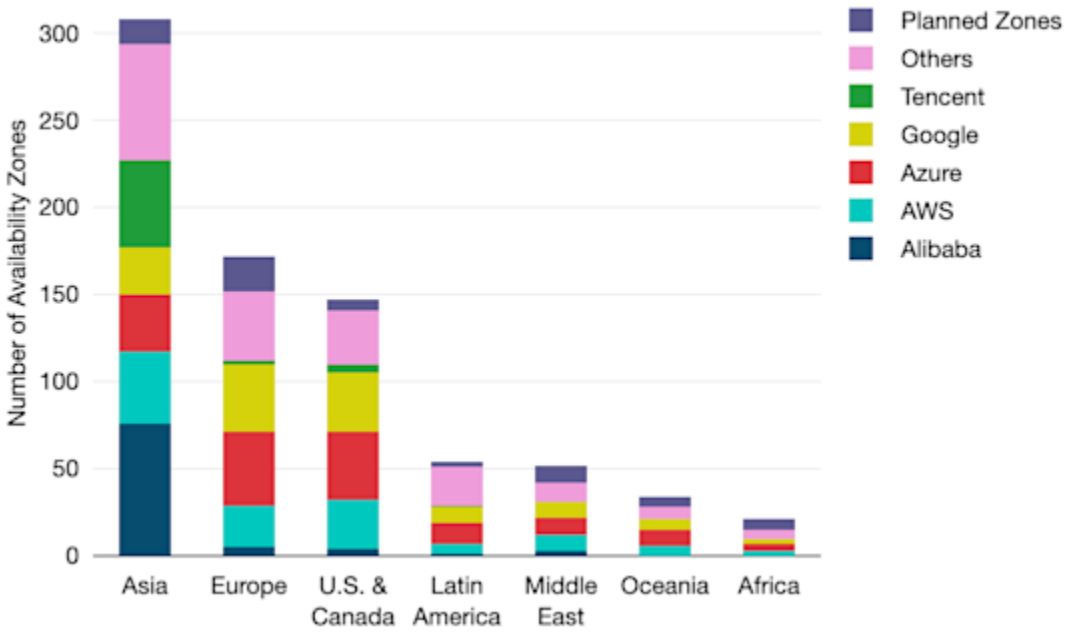

Notes: Data represent only infrastructure as a service (IaaS) providers and does not include software as a service (SaaS) providers. Data only include the data center information from Alibaba, AWS, Google Cloud, IBM, Microsoft Azure, Oracle Cloud, OVH Cloud and Tencent Cloud. Zones are defined as discrete data centers or data center sites. Azure regions that do not have zones are counted as one zone in this figure. Data as of Q1 2025.

Source: TeleGeography

Since 2013, cloud providers have launched an average of 18 new cloud regions per year. In 2019, Oracle joined the fray, launching 12 new cloud regions. Among all providers, a whopping 45 new regions were added in 2019. Early 2020 looked equally promising, with cloud providers on track to launch as many or more regions than the year prior. Alas, COVID-19 struck, stifling these ambitions. Nonetheless, this rampant expansion continued to pick up pace soon after as cloud operators successfully launched 23 and 27 new regions in 2020 and 2021 respectively. The year 2022 ended with a similar number of data centers with 23 new regions launched. A further 23 were added in 2023.

The pace seems to be picking up again with more than 35 planned regions for the next year or two. Azure leads the pack, contributing 17 new cloud regions in the pipe. Google, AWS and Oracle are also on the bandwagon, announcing plans for 8, 6, and 4 additional new regions respectively.

FIGURE 5
Planned Cloud Zones

Notes: Data include available information from Alibaba, AWS, Google Cloud, IBM, Microsoft Azure, Oracle Cloud and Tencent Cloud. Circles are scaled according to the number of zones per metro. Data as of Q1 2025.

Source: TeleGeography

WAN Services Coverage

The geographic coverage of carriers' enterprise network services varies significantly. Not every carrier connects to every city in their customers' networks, and not all services are available everywhere. When narrowing down the universe of potential suppliers, enterprises must first consider how their geographic requirements overlap a potential service provider's physical network. They then must determine if the specific data services they require are enabled at each of the service providers' PoPs. This analysis examines carrier network connectivity and service availability from a geographic perspective.

Global Business Center Product Comparison

Layer 3 MPLS IP VPN remains the most common enterprise-wide area network product across the key 165 business center metros. In these metro areas, carriers offer over 3,000 offerings of this service. Ethernet over MPLS is the second most common service in these locations, with over 2,600 offerings, and DIA was third, with over 2,300 offerings. EVPN is offered over 2,000 times in these metros, and DWDM is offered over 1,700 times.

FIGURE 6

Number of Offerings of Eight Enterprise Products across 165 Key Business Centers

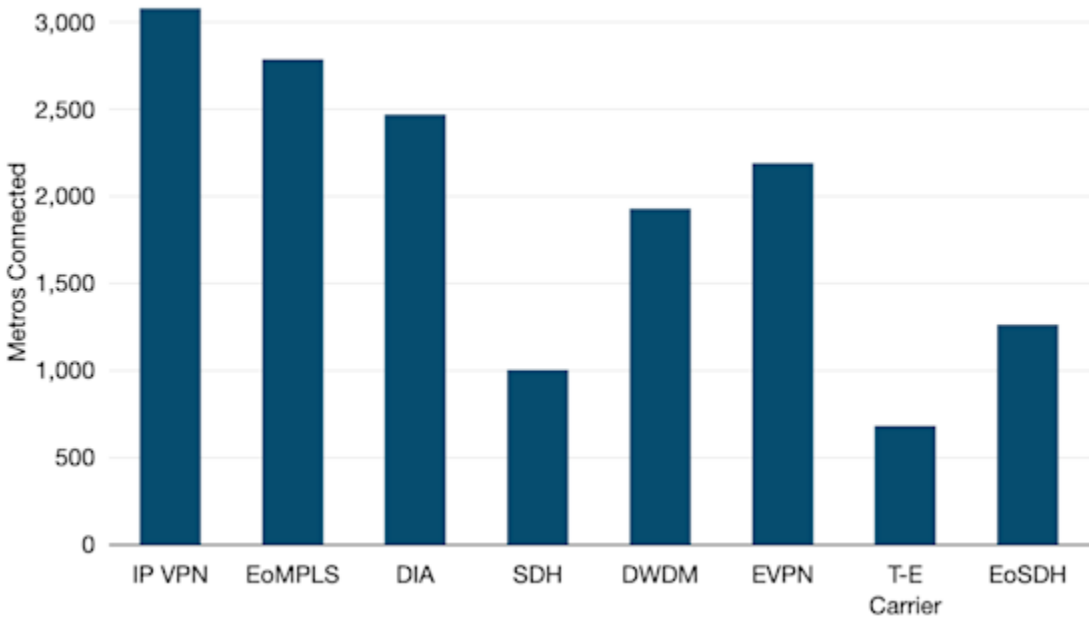

Notes: Total number of offerings that participating carriers indicated they could offer each of the listed eight services in 165 business centers. Data as of Q4 2024.

Source: TeleGeography

Product Geographic Comparison

Each of these five enterprise products (plus SDH, EoSDH, and T-E services) are most widely available in Europe, followed by the U.S. and Canada, and then Asia.

FIGURE 7

Number of Offerings of Top Services by Region in Key Metros

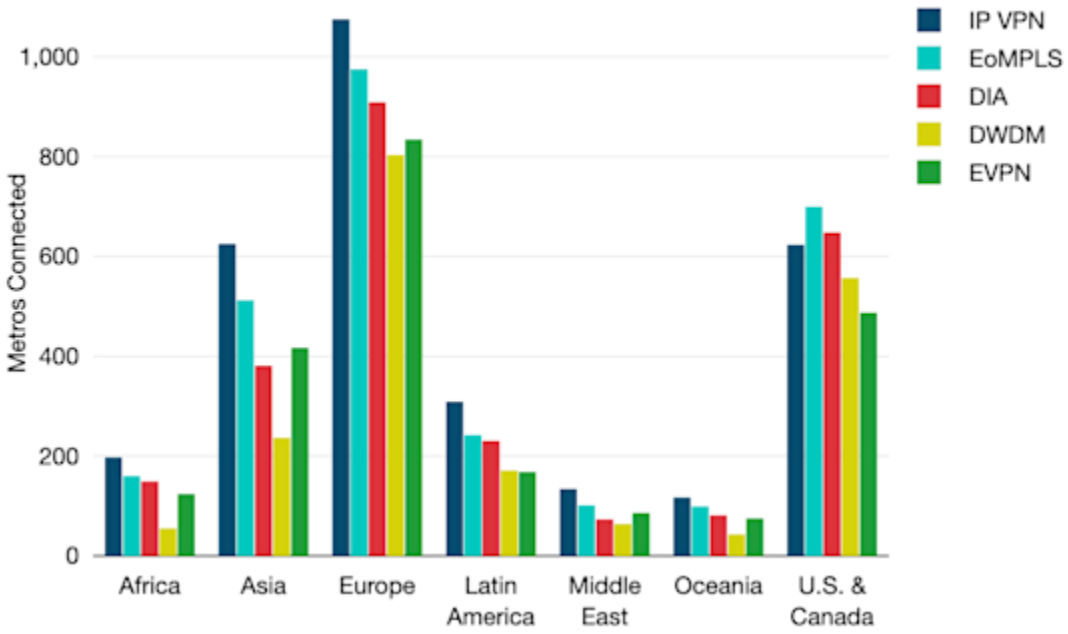

Notes: Each bar represents the total number of sites for each product across 165 key business centers. Data as of Q4 2024.

Source: TeleGeography

Global Business Center Provider Comparison

When sourcing a WAN, enterprises can keep it simple by relying on one primary global service provider, or they can work with many carriers to get the lowest prices in each region. In fact, according to our most recent WAN Manager Survey, a healthy minority—27%—of companies sourced their MPLS from a single global provider. A far smaller share—15%—source their DIA from a single provider. Enterprises should therefore have a strong command of the provider landscape in different regions of the world. This is part of the reasoning behind this report's expansion beyond the original 165 metros to include *all* carrier locations. It is also useful to understand which carriers can cover the key global business centers where the vast majority of global offices are located. This section focuses on which carriers can cover these major metros.

BT is the most widespread IP VPN provider across the 165 business centers, covering more than 110 metros. The remainder of the top ten includes a roster of well known providers: Verizon, Orange Business Services, AT&T, Vodafone, etc.

FIGURE 8

Top International Service Providers by Number of IP VPN Metros Within 165 Key Global Business Centers

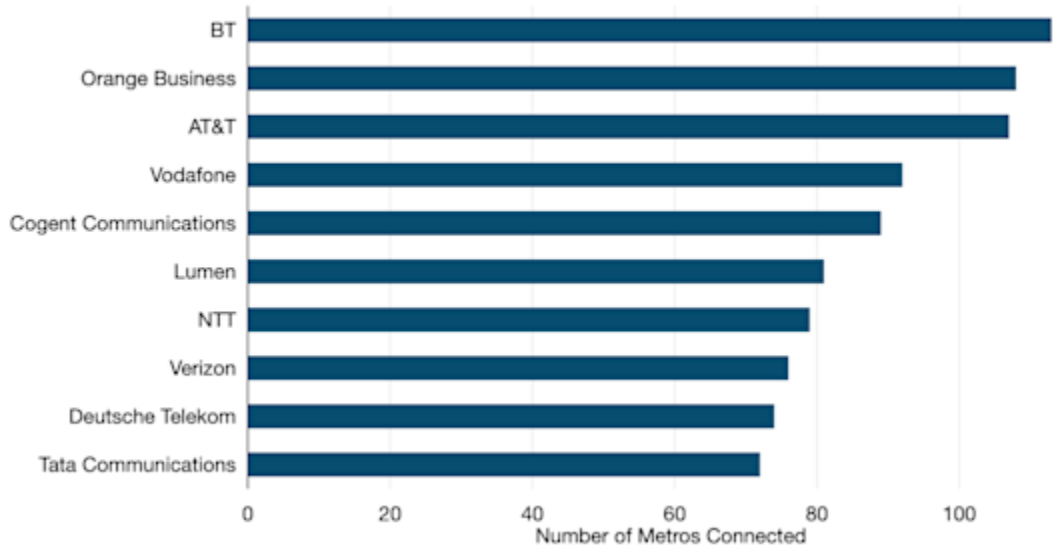

Notes: Each bar represents the number of cities among the key 165 business center metros in which the listed carrier can offer IP VPN. Data as of Q4 2024.

Source: TeleGeography

Content Delivery Networks

At TeleGeography we track the cache locations (more than 2,300) by city (more than 500) of the major CDN companies (over 20). Akamai is the largest CDN company by number of sites, close to 350, with Cloudflare in second with just under 300 and CDNetworks a distant third with just under 200 cluster sites. Geographically, Asia has the highest number of cache sites (over 700), Europe is a distant second with over 500 locations followed by the U.S. and Canada with over 450.

FIGURE 9
Content Delivery Network PoP Map

Notes: Data as of Q1 2025.

Source: TeleGeography

So are content delivery networks (CDNs) cloud networks? Not necessarily, but most of the top cloud providers offer CDN services on their cloud network like AWS, Google, and Microsoft. So then are CDN providers cloud providers? These days the lines are blurring. CDN providers are offering cloud services and have been adding traditional cloud services like compute, storage or networking to the product lists. For example, Akamai—the dominant player and paragon of a CDN provider as well as the bellwether for the CDN industry—began to offer cloud services in earnest after its acquisition of cloud computing company Linode in 2022.

Why start selling cloud services? The CDN industry has in a troubled state as of late. Rising capex, declining margins, and stiff competition has had its effect on some major providers the past two years— Stackpath shut down, Edgio declared bankruptcy, and Lumen exited the CDN market. Akamai still dominates the CDN market but has shifted its focus to offering new services like security and now infrastructure as a service as CDN profits have dipped.

Introduction: Sizing the WAN Market

What is the annual value of the global enterprise wide area network (WAN) market? While no one could provide a precise, objective number, we have created a model to give a range of plausible values for this market. We did this using proprietary TeleGeography data to build a market size based not on vendor reporting, but on enterprise network demand and pricing data we have collected over decades of primary research. This is a "bottom-up" approach to market sizing rather than "top down". Through this approach, we have determined some estimates for the size of the global enterprise WAN market broken out by subregions, countries, and specific network products.

Key Findings

- The global enterprise WAN market is potentially worth about half a trillion dollars.
- MPLS, while on the decline, is still a huge contributor to WAN revenue representing about one third of the market.
- Local access services are worth another third of the market, just ahead of DIA.
- East Asia is the biggest regional enterprise WAN market in the world, driven largely by China.
- The U.S. is still the single largest country for WAN revenue, representing about one quarter of global WAN revenue.
- The top four countries alone accounted for more than half of the global WAN market.

Global Enterprise WAN Market Size

So, how large is the global enterprise WAN market? We ran several versions of our model to account for uncertainty in our assumptions and data. The primary model we will focus this analysis on is the "median-median" model. This model combines the *median* site count based on our assumptions about the distribution of site counts among all global enterprises, and the calculated *median* price across all cities within a given country for each product and port size in the model. The other versions of the model we ran were based on including lower and higher numbers of potential sites, and the low and high prices for various services.

FIGURE 1

Range of Total Market Sizes (USD Billions, 2025)

Site Count Statistic	Total Site Count	Price Statistic	Market Size (USD Billions)
Median	13,203,750	Low	$227.60
Low	10,563,000	Median	$403.88
Median	13,203,750	Median	$504.84
High	15,844,500	Median	$602.17
Median	13,203,750	High	$1,523.87

Notes: The table demonstrates the total site count and annual market size for each model in billions of USD.

Source: TeleGeography

Our best estimate of the **global enterprise WAN market is $505 billion**. This tracks with the general estimate of the global telecom market being worth about $2 trillion, given that enterprise is just one segment, leaving out all consumer, small business, wholesale, wireless, and other telecom services. Of course, there will always be uncertainty around these numbers, which is why we have presented results from various model runs. We doubt very much that our high price model paints a realistic picture of the market, so the $1.5 trillion value is unlikely to match reality, but other combinations of different site counts and price ranges certainly seem like plausible alternatives.

Global Product Revenue

Our primary goal with this model, however, isn't just to come up with a singular global market size, but rather to compare contributions of the key WAN products and geographies. We look at revenue for MPLS, DIA, business broadband, local access, and SD-WAN to more closely examine revenue sources both by product and subregion or country. Globally, we see massive differences between the revenue contributions of these WAN products and geographies.

FIGURE 2

Global Market Size of WAN Products (2025)

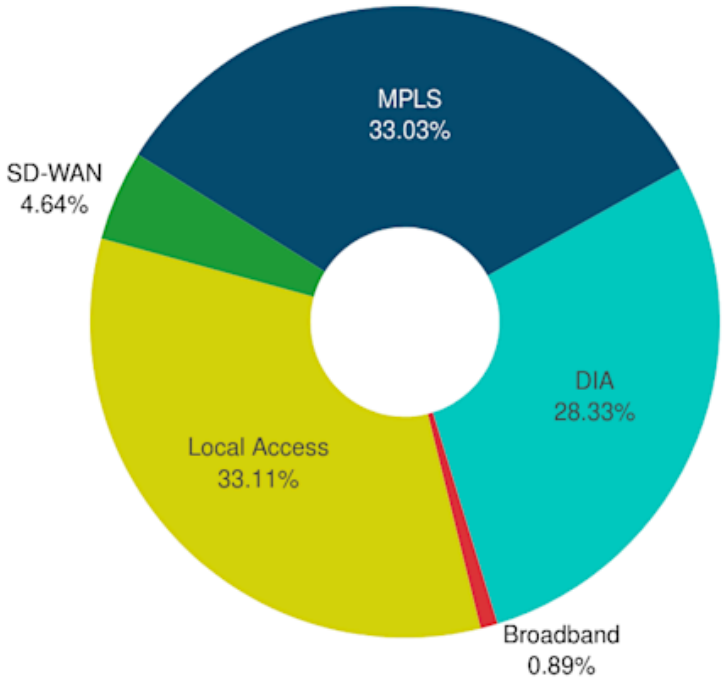

Notes: Each section represents the percentage contribution of the listed product to the total median-median model run global market size.

Source: TeleGeography

- MPLS and local access are very similar, both contributing about one third of the total market revenue.
- DIA is not far behind with a bit more than one quarter of global revenue.
- SD-WAN is well behind at about 5%, and business broadband–usually sold at much lower price points–is only about 1% of global WAN revenue.

Market Size by Geography

The key reason we used a bottom-up approach was to be able to break out our results by geography and product. In this version of the model, we included nearly every political division for which we have GDP and population data, which comes out to 231 individual countries (which exceeds the 193 sovereign states recognized by the U.N., including as separate countries various dependent territories or other minor states).

Subregional Market Sizes

Economic development, corporate presence, and bandwidth demand are far from uniform across our subregional geographic divisions, so it should be no surprise that WAN revenue is wildly different across the globe. The below chart details the percentage that each subregion contributed to the global market size of roughly half a trillion dollars.

FIGURE 3

Subregional Contribution to Global WAN Market Size (2025)

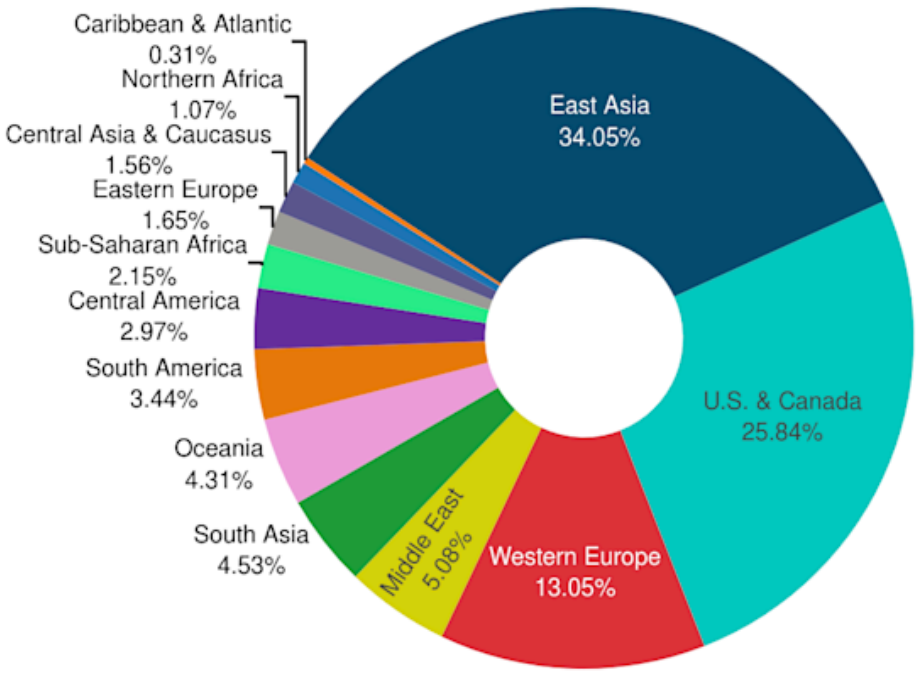

Notes: Each section represents the percentage that each subregion contributed to the global market size for the median-median model run.

Source: TeleGeography

- East Asia is the biggest enterprise WAN market in the world, contributing roughly one third of global revenue.
- The U.S. & Canada subregion contributed about one quarter. However, this is notable as this subregion is only two countries, compared to 19 in East Asia.
- Western Europe, despite accounting for 25% of enterprise sites in the model, contributed only 13% of global revenue. This is due mostly to that region having consistently the lowest telecom prices in the world.

East Asia is the biggest enterprise WAN market in the world, contributing roughly **one third** of global revenue.

Learn more about
Cloud & WAN Research Service

- Most other subregions had very small revenue contributions of around a few percent. The Caribbean and Atlantic region was the smallest, with less than 1%.

Country Market Sizes

As we explained above, we first distributed sites by subregion, then by individual country within each subregion. But we also reallocated specific bandwidths according to different demand and pricing in each subregion. For some subregions, like Western Europe or the U.S. & Canada, demand is fairly uniform across different countries. But in others, such as Oceania, and East Asia, we further distributed port/circuit sizes by country to account for how different bandwidth demand is between countries in those regions. This allocation, along with wide variance in service pricing, leads to massively different revenue numbers for most products at the country level.

There is a massive difference in market size between countries. The upshot is that a few countries have an outsized contribution, and there is a long tail of countries with very small markets.

FIGURE 4

Total Annual Revenue by Country (2025)

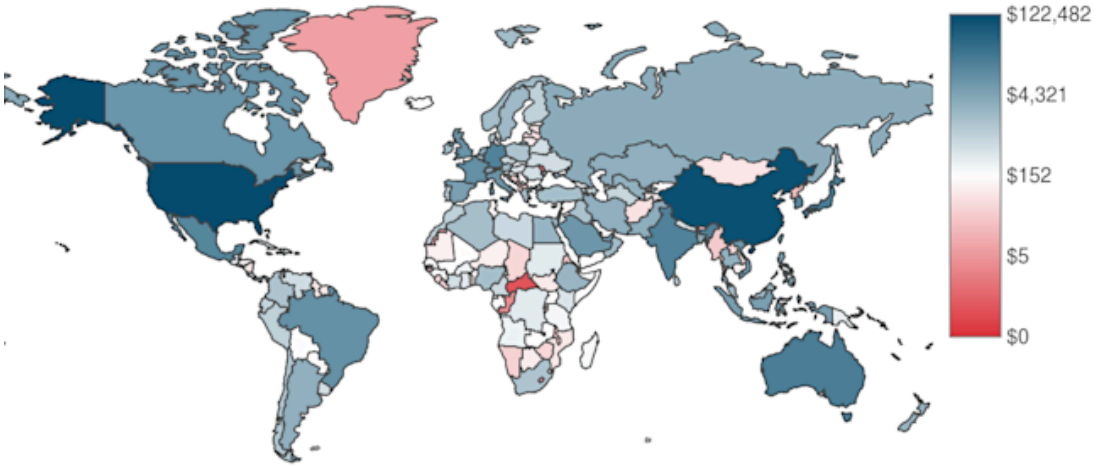

Notes: The map displays the total enterprise WAN revenue for each country across all products included in the model. This is for the median-median model run. Prices are shown in millions of US dollars.

Source: TeleGeography

- The U.S. and China clearly standout as massive revenue generators for enterprise WAN services.
- Western Europe contributes a large amount as a subregion, but many individual countries within the subregion contribute a fairly small amount.
- We have a long tail of smaller economic markets that together account for collectively very little of global WAN revenue.

Conclusion

While never a static market, the global enterprise networks market has been in a state of flux for many years and that won't likely abate anytime soon. The move to cloud, the adoption of SD-WAN, and the integration of more affordable internet are still impacting WANs, let alone the now emergent trends of NaaS, automated ordering and provisioning, AI-driven bandwidth demand boost, and automated network management. It is more important than ever for industry players to understand the size and shape of the enterprise connectivity market. This model and report seeks to offer guidance on where the current market is for underlay and overlay services, and thus inform how carriers, MSPs, and other technology vendors invest in and prepare for the future. To that effect, we will also soon follow up with a forecast report looking at plausible future network services allocations and pricing.

For the here and now though there are several key trends worth taking note of:

- MPLS is still massively important to enterprise WAN revenue.
- Local access varies considerably in price and thus revenue around the world.
- DIA may be more prominent and demand for DIA at higher bandwidths, but lower prices mean it is still not the core of underlay revenue.
- SD-WAN is a must-have for most enterprise networks, but still pales in comparison to network services in terms of revenue.

Business broadband will likely grow in popularity, particularly as more high-speed plans are available, but it won't be able to replace even a small fraction of traditional WAN revenue.

CHAPTER 06

International Voice Report

Learn more about
International Voice Report

Findings

The downward slide continues. The year 2014 represents the peak for international voice traffic. International call minutes declined the following year, for first time since the Great Depression—and it's been downhill ever since. The slump in voice traffic has turned into a rout, as carriers' traffic fell by 6.2% in 2019 and by a further 7.2% in 2020. The COVID-19 pandemic spurred a short-term rally in international call volumes in early 2020, but things pretty much returned to the new normal. Traffic fell a further 6.0% in 2021, by 5.8% in 2022, and by an astonishing 7.8% in 2023. This dreary trend continues unabated, with international voice traffic dropping by another 7.0% in 2024.

FIGURE 1

Circling the Drain?

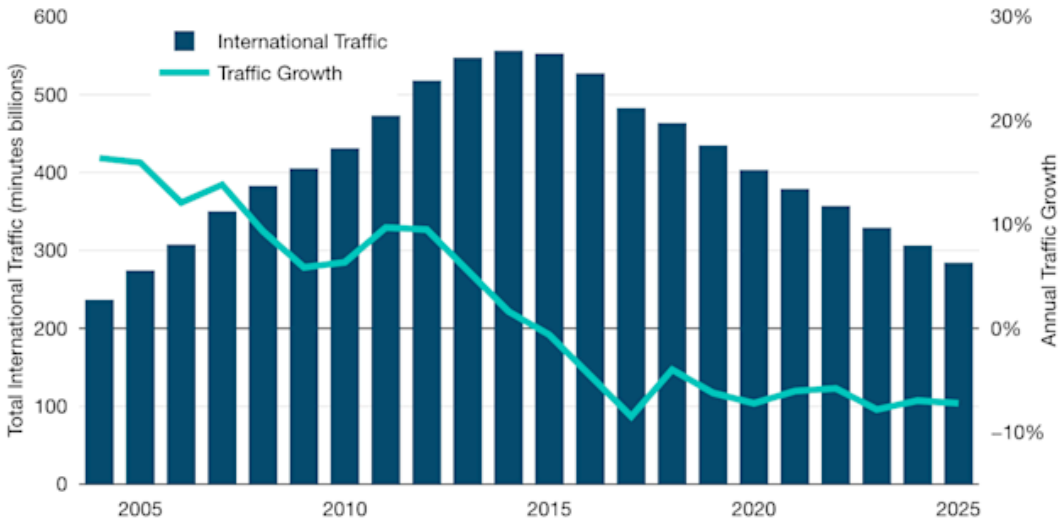

Notes: Data for 2024 are projections.

Source: TeleGeography

The OTT Effect

The new-ish market dynamic—social calling that replaced business communications as the primary driver of ILD usage—fueled a long era of international call traffic growth that began in the 1990s. In 1990, U.S. international call prices averaged over one dollar per minute(!) and business users accounted for 67% of ILD revenue. A wave of market liberalization in the subsequent decade brought new market entrants, causing prices to tumble, and making international calling ever more affordable to consumers. In the early 2000s, the introduction of low-cost prepaid phones made it possible for billions of people in developing countries to obtain their own phones, and to keep in touch with friends and family abroad easily. Call volumes soared, and by 2015, calls to mobile phones in developing countries accounted for 48% of global ILD traffic.

The transition to mobile and social calling drove a 20-year boom in voice traffic, but has also left the industry uniquely vulnerable to the rise of mobile social media. While Skype was the dominant communications application for computers, a veritable menagerie of smartphone-based communications applications, such as WhatsApp, Facebook Messenger, WeChat (Weixin), Viber, Line, KakaoTalk, and Apple's FaceTime now pose a greater threat. We believe that WhatsApp had about 2.5 billion monthly active users in 2025, with Facebook Messenger topping 1 billion. WeChat had about 1.3 billion active users at the same time. TeleGeography estimates that seven OTT communications applications—WhatsApp, Facebook Messenger, WeChat, QQ, Viber, Line, and KakaoTalk—combined had roughly 6 billion monthly users in September 2025. These estimates exclude other apps, such as Apple's FaceTime, Google Meets, and Zoom.

It's hard to pin precise numbers on the volume of international OTT communications. However, a simple thought experiment helps to illuminate its likely scale. Between 1983 and 2007, international phone traffic grew at a compounded annual growth rate (CAGR) of 15%, and traffic grew an even faster 21% CAGR between 1927 and 1983. It's hard to believe then that the recent decline in traffic means that people have lost interest in communicating with friends and family abroad. Rather, it suggests that they are turning to other means of keeping in touch.

TeleGeography has fairly reliable estimates of Skype's traffic through 2013, when the company carried 214 billion minutes of on-net (Skype-to-Skype) international traffic. Telcos terminated 547 billion minutes of international traffic in 2013, and OTT plus carrier traffic totaled 761 billion minutes. We had been assuming that total (carrier plus OTT) demand for international communications was growing at a relatively modest 13% annually since 2013. For this analysis we further assume that texting, video, and email have quite recently eaten into overall voice growth, and that this annual growth gradually glided down to 0% by 2023. A bit of arithmetic would tell us that the combined volume of carrier and OTT international traffic would have expanded to 1.6 trillion minutes in 2023. (The figure Where Did the Minutes Go?: The OTT Effect illustrates this nicely). Traditional carrier traffic has slumped, but OTT traffic has risen to fill the void. This calculation suggests that cross-border OTT traffic overtook international carrier traffic in 2016, and would exceed 1.3 trillion minutes in 2025, dwarfing the 284 billion minutes of carrier traffic projected by TeleGeography.

Wholesale carriers terminated approximately **217 billion minutes** of traffic in 2024, down 8% from 2023.

Learn more about International Voice Report

FIGURE 2

Where Did the Minutes Go?: The OTT Effect

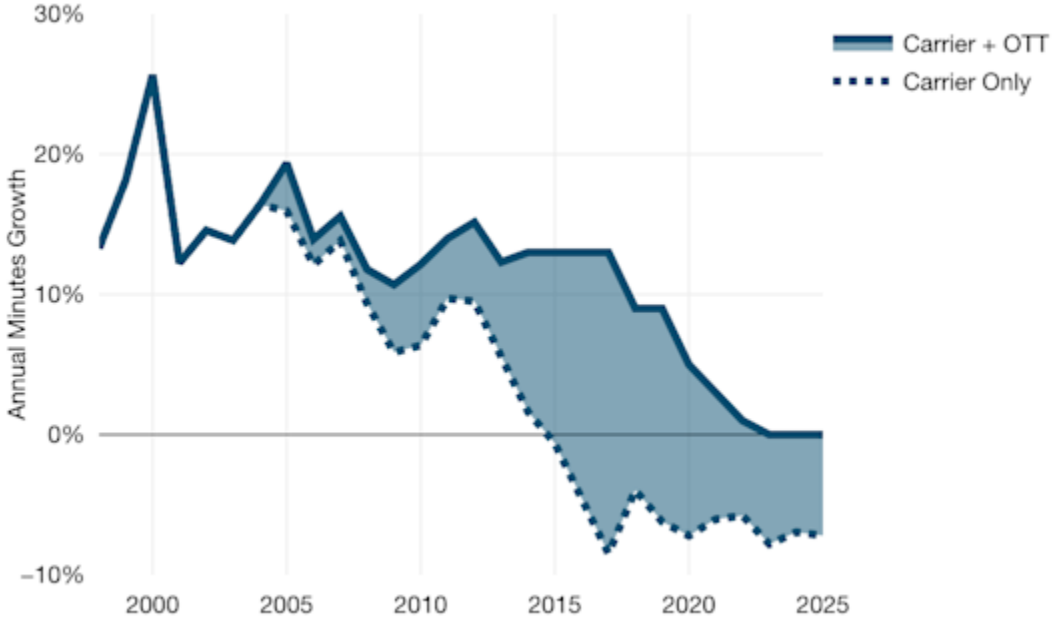

Notes: OTT traffic reflects in-app cross-border traffic only, and excludes calls originated on apps but terminated to the PSTN.

Source: TeleGeography

International Wholesale Services

Many retail service providers, such as mobile operators, MVNOs, and cable broadband providers, rely heavily on wholesale carriers to transport and terminate their customers' international calls. Wholesale carriers terminated approximately 217 billion minutes of traffic in 2024, down 8% from 2023. Wholesale traffic declined at an average rate of 1.5% per year over the past ten years, compared to a -6% CAGR for overall traffic. Wholesale carriers terminated more than two-thirds (71%) of international traffic in 2024, essentially unchanged from the year before.

Traffic to mobile phones in emerging markets has spurred expansion in wholesalers' share of the overall market. In 2024, wholesale carriers terminated over 86% of traffic to South America. In contrast, wholesale carriers terminated only 55% of traffic to western Europe.

FIGURE 3

Wholesale Revenues and Traffic by Market Type, 2014 and 2024

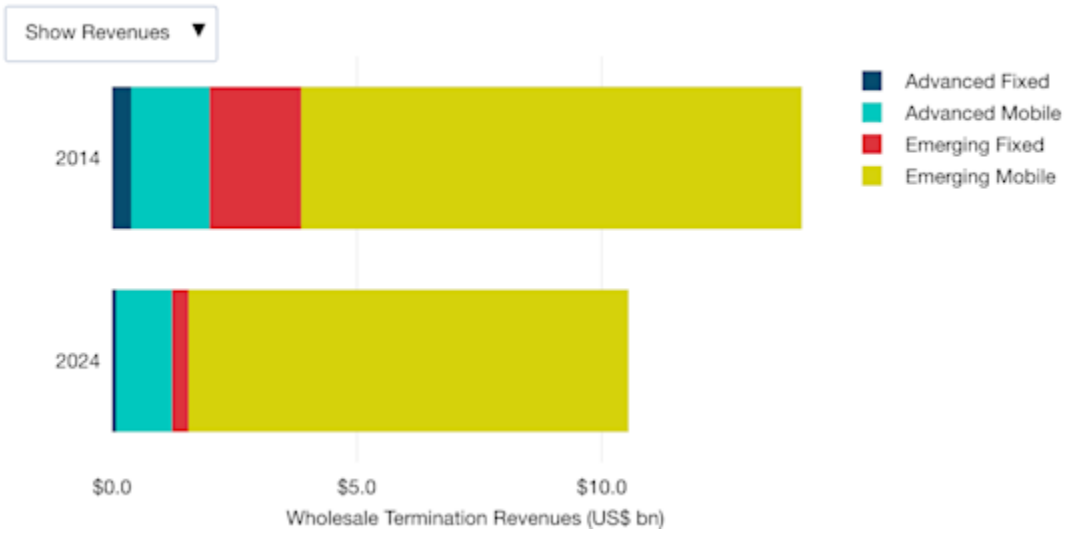

Notes: Advanced markets are defined as those of the upper 30th percentile of countries as ranked by GDP per capita in 2024. Emerging markets are those of countries in the lower 70th percentile. World and regional totals exclude countries with extremely low traffic volumes.

Source: TeleGeography

Wholesale revenues are down 25% from ten years ago. Let’s take a moment to dive in a bit and see what's going on. A few years ago, one could say that traffic to mobile phones in emerging markets had driven international wholesale market growth. No longer. Toggle between the "Show Revenues" and "Show Traffic" options in the "Wholesale Revenues and Traffic by Market Type, 2014 and 2024" figure. It's plain as day that revenues to fixed lines in both advanced and emerging economies has collapsed. At the same time, you'll see that revenues from calls to mobiles in all markets has shrunk slightly over the past ten years, to about $10.1 billion.

Who's carrying this traffic? Take a look at the figure “Traffic Volumes of Major Carriers.” It shows a ranking of some of the world’s largest international carriers. The top 8 operators in the figure carried 41% of all global traffic in 2024, about 125 billion minutes. Only one of the eight largest carriers in the world—BTS—terminated more traffic in 2024 than in 2023. Chew on that for a bit.

FIGURE 4

Traffic Volumes of Major Carriers

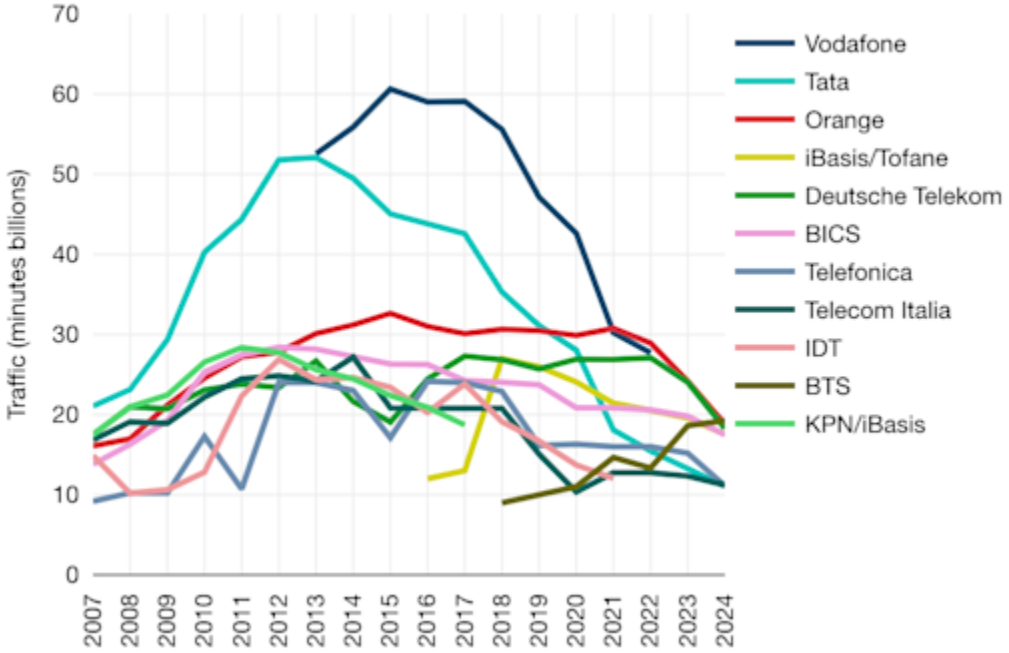

Notes: Pre-2013 data for Vodafone, which reflect the combined traffic of Vodafone and the former Cable & Wireless Communications, are not available.

Source: TeleGeography

Prices & Revenues

Retail ILD call revenues have slowly withered in recent years. So, too, has ILD's contribution to overall carrier revenues.

Let's look back a few years. In 2014, retail international call revenues (revenues that exclude wholesale revenues and termination payments) generated $96 billion. During that year, wireline, broadband, and wireless services, in total, generated $1.4 trillion. Thus, ILD accounted for 7.1% of total revenues in 2014.

In 2024, ILD accounts for only 3.3% of total carrier revenues.

For the mobile market, outgoing ILD revenues as a share of overall wireless revenues had remained relatively static; they had even increased from 2010 to 2012. Since then, international mobile revenues have followed the same downward trajectory as fixed ILD revenue trends. In both the fixed and mobile sectors, ILD calls account for a noticeably smaller share of overall carrier revenues than they did a few years ago.

FIGURE 5

Global Retail Revenues from International Calls

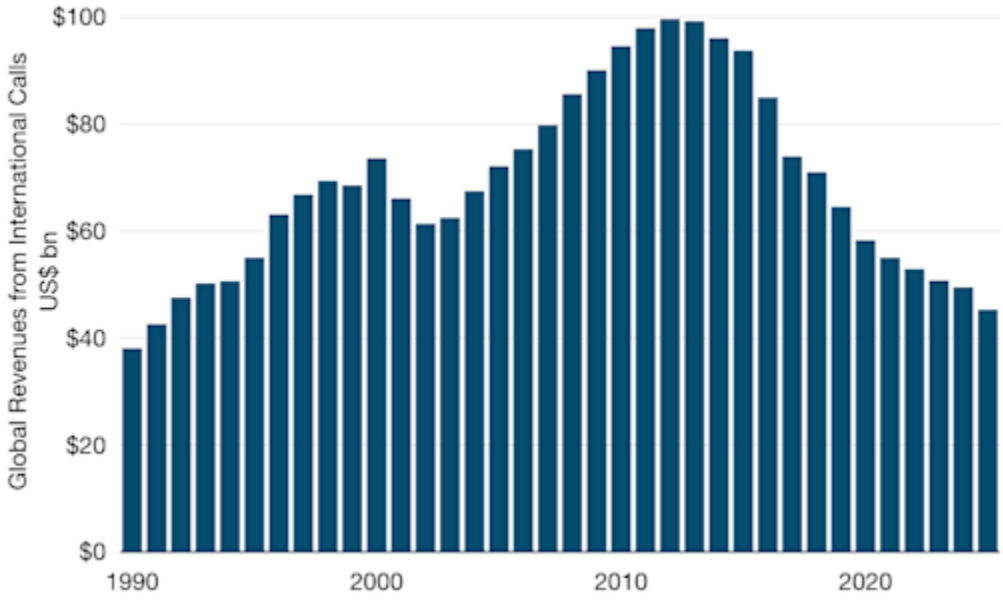

Notes: Data measure retail revenues on outgoing international calls; totals do not include revenue from wholesale services or incoming international traffic termination. Data for 2025 are projections.

Source: TeleGeography

Contact Us

Since 1989, we've delivered the most comprehensive data and analysis in the global telecom industry.

Learn more about our platforms and services.

TeleGeography Headquarters
One Thomas Circle, N.W.
Suite 360
Washington, DC 20005

+1 (202) 741 0020

info@telegeography.com