TeleGeography

Introduction 1
Transport Networks 3
IP Networks 7
Data Centers 11
Cloud and WAN 17
International Voice 21
Glossary 24
Research Catalog 26

The State of the Network

2025 Edition

Sponsored by

telecomegypt

Brought to You By Real Data

Hello, gentle readers, and welcome to the **2025 State of the Network Report**—our eighth edition.

The TeleGeography team spent the last year compiling market surveys, building algorithms, and analyzing emerging trends, and now we're ready to share the highlights.

Think of this e-book as your 2025 update on the telecommunications industry.

As always, we extract the major global bandwidth headlines, take a snapshot of the global internet, peruse the latest in data centers, check in on the cloud, and finish with an update from the voice market.

If you're interested in topics like subsea cable activity, cloud geography developments, the effect of AI on long-haul capacity requirements, and data center power scarcity, you're in the right place.

We'll leave you to it. Thanks, as always, for checking out our research.

— The TeleGeography Team

P.S. If these insights are up your alley, we've got lots more in our full suite of research apps.

The Experts Behind This E-book

Lane Burdette
Senior Research Analyst
Connect on LinkedIn

Paul Brodsky
Senior Research Manager
Connect on LinkedIn

Patrick Christian
Senior Research Manager
Connect on LinkedIn

Jon Hjembo
Senior Research Manager
Connect on LinkedIn

Alan Mauldin
Research Director
Connect on LinkedIn

Tim Stronge
Chief Research Officer
Connect on LinkedIn

Marvin Tan
Senior Research Analyst
Connect on LinkedIn

Juan Velandia
Senior Research Analyst
Connect on LinkedIn

Recent Presentations

Watch & Download

- 2025 Telecom Workshop
- AI and Submarine Cables
- How To Measure Connectivity
- 2024 Telecom Workshop

TRANSPORT NETWORKS

Living up to Our Potential (Capacity)

Those who follow the telecommunications space know that global network infrastructure and bandwidth markets are among the most fundamental building blocks of the global economy.

As with other areas of this industry, the capacity market sees growth, struggle, uncertainty, and advancement. 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.

Demand Trends

Worldwide bandwidth demand continues to grow at a steady pace. Annual demand growth has decelerated slowly, but aggregate demand more than tripled between 2019 and 2023 to reach an eye-popping 5 Pbps.

On a regional level, most parts of the world have seen comparable growth at about 35-40% CAGR since 2019. The markets that stand out are Africa, where capacity growth is still surging at a nearly 50% CAGR, and the U.S. & Canada, where market maturity has slowed demand to around 30% CAGR.

The Role of Content Providers

Content and cloud providers—most specifically a handful of companies like Google, Meta, Microsoft, and Amazon—are firmly entrenched as the biggest users of network capacity globally. As recently as 2016, internet backbone providers accounted for the majority of demand. Not anymore. As of 2023, content and cloud networks accounted for more than 70% of all bandwidth usage.

This demand scale is not the same on every route that we track at TeleGeography.

Construction Cost of Submarine Cables

Construction costs in USD billions

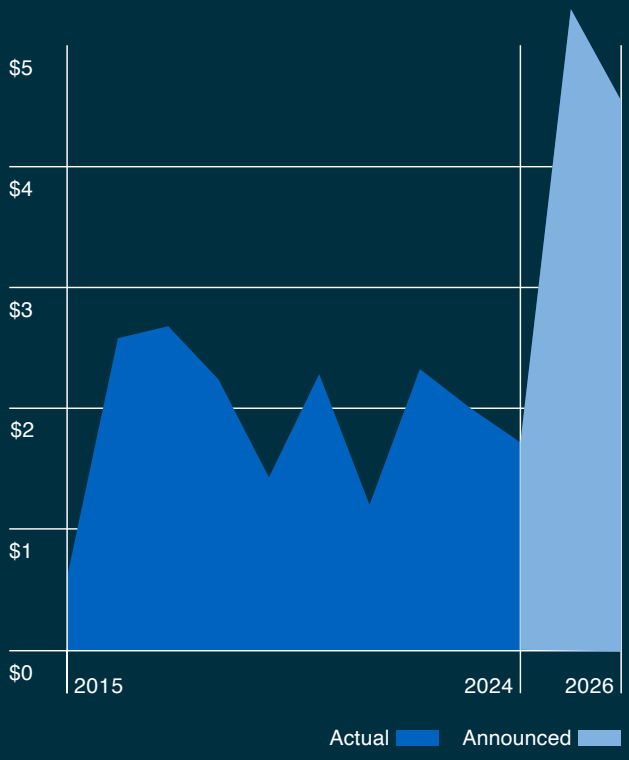

Notes: Total construction costs of all international and domestic submarine cables entering service in designated years. Construction costs exclude the cost of subsequent capacity upgrades and annual operational costs. 2024-2026 construction costs based on announced contract values and TeleGeography estimates. Not all planned cables may be constructed.

On some of the biggest subsea routes like the trans-Atlantic, trans-Pacific, and intra-Asia, content providers account for the vast majority of demand (80% or more). Other routes remain more carrier-driven, such as the U.S.-Latin America routes, and routes connecting Europe to Africa, Asia, and the Middle East.

Why the contrast? Content providers are focused on connecting data centers across different zones. Due to the concentration of service delivery in major Asian, European, and U.S. markets, core routes connecting these regions are of highest priority to the content providers.

That said, content provider demand is rapidly growing everywhere and outpaces demand growth even on routes where carriers continue to drive overall capacity usage. As might be expected, content provider demand growth is fastest in regions where carriers are still dominant, like Africa, Latin America, and the Middle East. But there's no part of the globe where content demand growth isn't outpacing that of internet backbone providers.

Meeting Demand Requirements

Investment in new submarine cables has surged in recent years. Despite some fluctuations, new cable investment has averaged over $2 billion per year in the past 8 years. The value of new submarine cables entering service from 2024-2026 is forecasted to reach over $10 billion.

The amount of potential capacity the newest generation of cables will provide is incredible. Several major routes will see their potential capacity more than double once new cables are completed.

Pricing

We've witnessed an unprecedented slowing of bandwidth price erosion globally over the past few years, as network investors, carriers, and customers grappled with delays in new network projects, existing system upgrades, and rising equipment costs. For the first time, many customers found themselves asking if prices were actually increasing and when price declines would return to typical levels. While the supply chain constraints and card shortages that spurred this trend have resolved themselves, the geopolitical issues that contributed, as well, have not. As a result, recent price trends vary dramatically by region. Price erosion has started to accelerate on some routes, but remains stalled on others.

Looking at weighted median 100 Gbps wavelength price trends on major international routes—between 2020 and 2023—weighted median 100 Gbps wavelength prices decreased an average of 13% compounded annually. That's compared to 18% over the prior three years (2017-2020).

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, still feeling the effects of new cables, diverse fiber pair ownership, and upgrades to existing systems. The route is also anticipating the launch of the Firmina cable in 2025. Between 2020 and 2023, 100 Gbps wavelength prices

on Miami-São Paulo decreased 28% compounded annually. That has exceeded the 22% decrease between 2017 and 2020. Johannesburg-London, which saw a massive influx of new capacity from the launch of Equiano in 2023, recorded an annual 22% price drop for 100 Gbps wavelengths over the past three years.

On routes with continued delays in new supply, price erosion is still stalled. For example, 100 Gbps prices on Marseille-Singapore and Hong Kong-Singapore decreased just 2% and 6% annually from 2020-2023. That's compared to 28% and 16% annually from 2017-2020.

The Europe-Asia and intra-Asia routes have been especially impacted by recent delays in supply. While card shortages have resolved themselves, geopolitical issues play a key role in delays of new systems on both routes. Until new systems start coming into service this year, available inventory is going for 2023 prices or potentially higher.

Outlook

What's fueling the changes we see in the global bandwidth market? Let's look at some key trends shaping long-haul capacity demand and pricing.

Geopolitical Concerns

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

The Red Sea is the focus of major problems. Even before the recent spate of rebel attacks on commercial shipping vessels, the Yemeni civil war created permitting headaches. Cable laying vessels require permits to enter a country's territorial waters. When two different entities claim the same swath of sea, the situation becomes complicated.

The near simultaneous cable faults on three cables off the coast of Yemen in February 2024 created a major challenge. Given the location of these faults, it is uncertain when maintenance vessels may be able to conduct repairs due to Houthi rebel attacks.

Subsea cable activity is also geopolitically challenging in the South China Sea. Cable builders find it increasingly difficult to receive Chinese permits for cable deployment in this region. Builders of the planned Apricot cable hope to avoid this issue by linking Japan to Singapore via the east side 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.

Artificial Intelligence (AI)

There's hardly a hotter topic in the network world right now than AI. 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 market. The answer? It depends.

Model training. If a given model is pulling data from all over the world, it will increase demand for long-haul bandwidth,

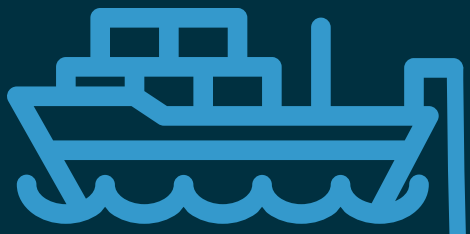

Google and, to a lesser degree, Meta are increasingly deploying their own subsea cables. This approach allows them to control the system design and landings, but also to move swiftly and avoid potential delays from working with partners.

but if data used is stored locally in the same campus or data center, it won't. Data residency requirements will also factor into the use of long-haul networks for model training.

Training clusters to inference clusters data transfer. If models are largely trained in more remote locations with cheaper and more abundant power sources, long-haul capacity will be needed to transfer data from training to inference locations. The size and frequency of training model updates will also factor into this aspect of demand. Of course, if training needs to be located closer to inference zones, long haul capacity won't come into play as much.

Inference. The location of inference clusters will also be critical in determining the ultimate effect of AI on long haul capacity requirements. If inference is deployed in zones close to end users, it may not impact long-haul demand as much. If the model requires data pulls from multiple locations, though, it could increase long haul demand.

These are just a few of the many factors to consider. As AI technology rapidly evolves, we expect that long-haul demand will grow, even if metro area demand increases more rapidly.

Private Cable Expansion

Google and, to a lesser degree, Meta are increasingly deploying their own subsea cables. This approach allows them to control the system design and landings, but also to move swiftly and avoid potential delays from working with partners. Google is the owner of 15 active and planned private cables, with many more in development—particularly in the Pacific. Meta is planning the Anjana trans-Atlantic cable, its first proprietary cable.

However, even on a private cable, content providers are not the only users. Some portion of the fiber pairs are typically available for sale or swaps. In the case of swaps, a content provider would swap fiber pairs in exchange for fiber pairs on other cables, for landing rights in a country, or for terrestrial backhaul within a country.

Content providers are not just swapping fiber pairs, but also selling them. Google is also selling IRUs for whole and partial fiber pairs on the company's private cables. Presumably, Meta would also make fiber pairs available for sale on Anjana.

Google's private cable investment is spread globally, but especially focused in the Pacific, where the company's Pacific Connect Initiative is crisscrossing the ocean with multiple cables. Notably, Google is receiving some financial support from the U.S. and Australian governments for these projects, largely to enhance connectivity to islands in the South Pacific.

Completely private content provider cables are unlikely to become the dominant model on every route. In some countries, they may not possess the legal or regulatory authority to land cables themselves and may need to rely on carriers or specialist companies who can operate a cable on their behalf. Also, few content providers currently have enough bandwidth demand to justify investing alone in a new cable.

IP NETWORKS

The Future of Global Internet Expansion

In our IP Networks Research Service, we analyze the meaning of our robust internet capacity and traffic data sets. We also discuss factors impacting IP transit pricing, and the role individual backbone operators play. Based on primary survey data gathered from dozens of regional and global network operators around the world, we conclude that COVID-related expansion of internet traffic and bandwidth was largely a one-off phenomenon, and that the trends we had been observing in recent years have reasserted themselves. International internet bandwidth and traffic growth have been gradually slowing in recent years, but they remain brisk. IP transit price declines continue globally, but significant regional differences in prices remain.

Internet Traffic and Capacity

Global internet bandwidth rose by 22% in 2024, continuing its slow but steady decline over the past few years. Total international bandwidth now stands at 1,479 Tbps, representing a 4-year CAGR of 25%. Clearly, the pace of growth has been slowing recently. Still, bandwidth has more than doubled since 2020.

Capacity growth varies across regions. Once again, Africa experienced the most rapid growth of international internet bandwidth, growing at a compound annual rate of 41% between 2020 and 2024. Asia is a distant second, rising at a 28% compound annual rate over the same period.

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 25%, respectively, between 2020 and 2024—about the same as the 25% compounded annual growth rate in bandwidth over the same period.

International Internet Bandwidth Growth By Region

Compound annual growth, 2020–24

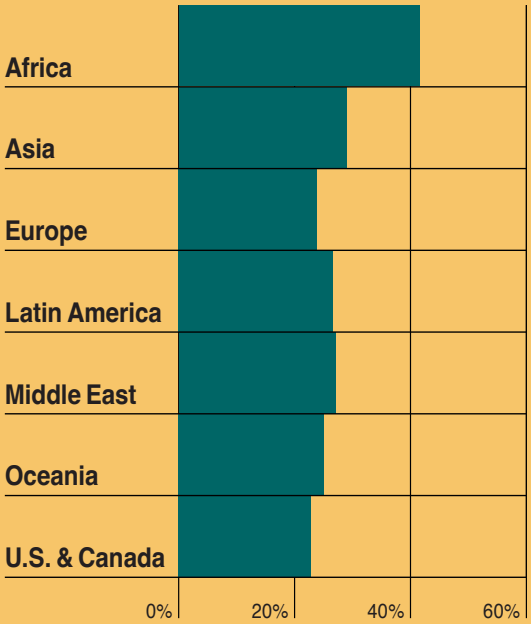

Notes: Data as of mid-year.

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. It reflects the introduction of competition into new markets and the response of more expensive carriers to lower prices. Trends in the IP transit market generally follow regional trends of the transport market. And while some have suggested that price erosion may slow as a result of recent inflation and supply chain constraints (as it has in the wavelength market), we have not seen this trend make its way into the IP transit market.

Across a range of markets, 10 GigE prices fell 13% compounded annually from Q2 2021 to Q2 2024. A comparable sample of 100 GigE port prices fell 15% over the same period. In Q2 2024, the lowest 10 GigE prices on offer in the most competitive markets were at the brink of $0.07 per Mbps per month. The lowest for 100 GigE was $0.05 per Mbps per month.

The sharper decline in 100 GigE reflects increased availability and competition at port size and the advanced maturity of 10 GigE. 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 2024, providers indicated that a majority of their sales mix in key U.S., European, and Asian hubs were now 100 GigE. On average, across seven cities—London, Miami, Frankfurt, Los Angeles, Hong Kong, Singapore, and São Paulo—the Monthly Recurring Charge (MRC) for a 100 GigE port was 6.8 times the MRC for a 10 GigE port.

Prices in developing markets that activated large scale subsea systems saw the greatest change. As 100 Gbps wavelength prices between Europe and South Africa fell 38% post the activation of the Equiano cable, 10 GigE IP transit prices in South Africa decreased 45% over the past year. This clearly supports a correlation between network scale, competition, and price.

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 $0.07 to $0.08 per Mbps. That's an average of 2.8 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 include 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.

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.

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 compared 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. 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

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.

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.

DATA CENTERS

A Power Struggle

The data center market is in the midst of an unprecedented shift that started—publicly at least—five 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 sufficient power. Since then, it turned out that data centers would need much more power than had been realized—just at a time when regulators and community organizations were growing increasingly critical of the sector's existing demand.

The accelerated commercialization of AI in 2023 took this problem to a whole new level. The data center industry has rallied to implement innovative power and cooling solutions while securing new energy partnerships. Meanwhile, utilities struggle to update power grids in the face of growing public opposition.

Let's delve into some findings from our Data Center Research Service.

Market Connectivity Score

The Market Connectivity Score (MCS) is our answer to the question "Where is the next big hub going to be?", and pulls from our wide array of research areas to evaluate the health of more than 3,000 metropolitan markets worldwide. Here is a brief look at a few of the markets highlighted in the MCS.

Frankfurt. Long-established as a regional center of networking, peering, and data center development, Frankfurt tops the MCS list as the strongest connectivity market in the world. The home of the world's largest data center-neutral IX, DE-CIX, Frankfurt has the world's strongest peering ecosystem (both by numbers of IX platforms present and numbers of ASNs connected) and strongest level of cloud

Frankfurt

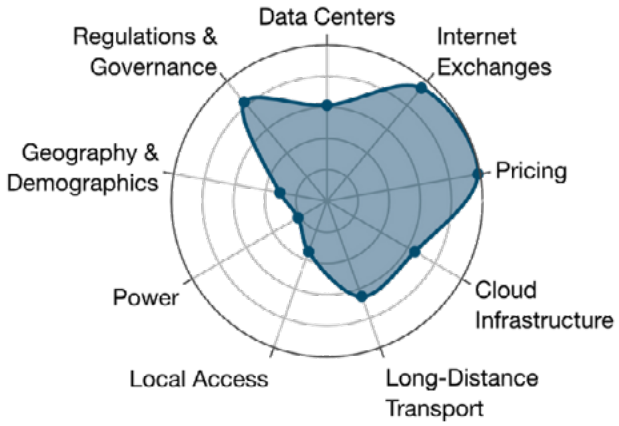

Singapore

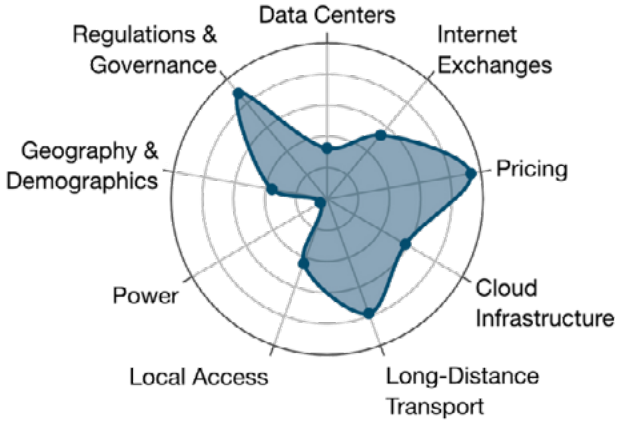

São Paulo

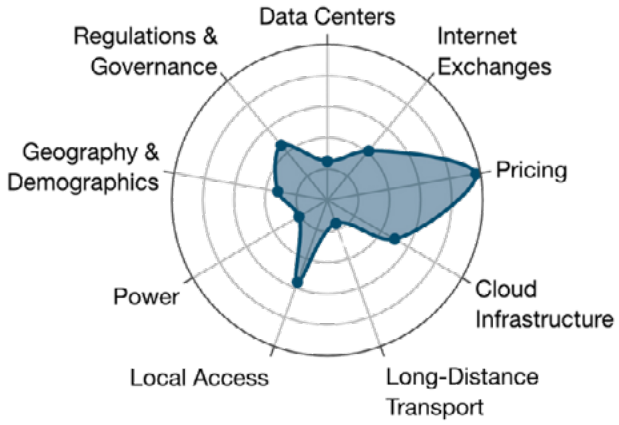

Notes: Data as of Q1 2025.

development. No other city has more connected international internet capacity or number of transport carriers serving the 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.

São Paulo. São Paulo is a critical market for Latin American connectivity, ranking highest not only in Brazil but in the entire region. Among other metrics, it benefits from favorable IP Transit pricing, with Miami-São Paulo remaining one of the most competitive routes to or from Latin America. São Paulo also hosts an incredible number of internet exchanges alongside a high number of ASNs connecting to those exchanges.

Power Scarcity

Several critical hubs have undergone major power scarcity crises, but utility problems are global and not confined to only a few major markets. Formally and informally, operators indicate that markets across the globe are seeing increased delays in securing power for new builds. 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.

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.

AI's Effect

Compounding the existing constraints on the data center market, we've reached a major inflection point in the development of AI, and more specifically in the development of generative AI.

AI work consists of two basic phases: training and inference (the work for which the model is trained). Each has a different impact on the data center market.

Training. In the training phase, workloads don't need to be located close to where the application will ultimately be run, so they can be conducted outside of core markets. The main requirements are that significant power, space, and GPU capacity are available. For this reason, AI training demand creates significant opportunity for data center operators in secondary and more remote markets.

Inference. In the inference phase, compute has to be done close to end users and is much more latency sensitive. At this point, workloads must shift to network and cloud-dense locations, creating further demand in highly concentrated markets. But in addition, inference loads will also need to be distributed closer to edge markets to reach more end users.

AI will require many changes within data centers. The most immediate concerns include the provisioning of higher density cooling and higher capacity interconnection. Standard air-cooling systems can't support GPU-based, power-intensive AI applications, so liquid cooling solutions will become far more prevalent within data centers, particularly where AI applications are being run.

Many data centers will be unable to accommodate such computational advances without extremely difficult and expensive retrofits. Many new AI-ready data centers will be in demand at a time when the data center market is already under regulatory scrutiny and power is becoming scarce. The full effects of AI on data center markets are as yet unknown. However, it is clear that these technologies—after years of development and impressive results—are here to stay.

Data Center Developments

The Washington metropolitan area—or more specifically Northern Virginia (NoVA)—dominates as the world's largest data center market. As of 2024, we estimate its gross commercial data center capacity at more than 24 million square feet of operational capacity.

Despite an impressive 16 to 17 million square feet of capacity, Beijing remains a particularly closed ecosystem with heavy restrictions on international investment and carrier presence. China Telecom, ChinData, GDS, and several other sizable operators have deployed extensively across the metro area, most heavily in Beijing itself and in Langfang, approximately 40 miles from the capital.

Asian and U.S. metro areas account for 8 of the 10 largest data center markets. In Europe, only London and Frankfurt make the list of largest markets.

We estimate that, as of 2024, colocation operators in the top ten data center markets by MW consume about 13 GW of power.

That's enough power to generate electricity for roughly 10 million homes—or, in this case, only about 1,000 commercial data centers!

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, the third-largest global data center operator, has a heavily concentrated footprint which is smaller than Equinix's but still twice as big as the next-largest 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.

By our tally, no large market comes close to the 4-year growth seen in Johor Bahru. The promise that pent-up demand for access to Singapore would spill over into Malaysia has come true. JB's commercial data center market growth has exceeded 40% CAGR since 2020.

We estimate that there are 24 metropolitan markets that have more than 1 million square feet of operational data center capacity and 4-year CAGR of at least 10%. Nine of those 24 markets are in Asia and seven are in Europe. Even more extraordinary—four of the Asian markets are in India alone. Mumbai, New Delhi, Hyderabad, and Chennai have all seen commercial data center capacity growth of 15% or more since 2020. In Europe, Warsaw, Berlin, and Madrid are all seeing robust growth.

Many markets, small and large, see low growth over time and periods of no new data center builds going into operation. Amsterdam, Dallas, New York, Hong Kong, and Singapore all have new data center capacity growth of less than 5% CAGR since 2020.

Digital Realty and Equinix dominate the scale and geographical diversity of recent market growth, as they do overall capacity. But Vantage and STT's group of companies have also built out more than 4 million square feet of new data center space in the past two years. STT's new builds have launched across Southeast Asia and India, though the pace of GDS rollouts in China has largely driven the overall growth for the STT group. Vantage's new data center launches since mid-2022 have been truly global, spanning core U.S. markets, Canada, several German locations, Australia, Taiwan, South Africa, and more.

Among operators tracked in our database, nearly 400 data center sites are known to be in the construction pipeline right now. While this development is spread across global regions, Europe and Asia outpace other regions with the largest percentages of new site deployments.

Power

We estimate that, as of 2024, colocation operators in the top ten data center markets by MW consume about 13 GW of power. That's enough power to generate electricity for roughly 10 million homes—or, in this case, only about 1,000 commercial data centers!

Only 17% 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

Lumen, Cogent, Zayo, Verizon, and AT&T are the most prominent carriers across global facilities. These five operators are especially widespread in the North America and Europe. Elsewhere, providers like 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 TELEHOUSE's London Docklands campus are also central nodes of international internet connectivity.

We continue to see new peering exchanges coming online across the globe in both established and developing markets. Recent deployments are geographically dispersed, with new IXs coming online in almost every region of the globe each year between 2020 and 2024.

Pricing

Current Trends

Starting in 2022 and continuing into 2024, anticipated price inflation has become a reality in the colocation market—at least in Europe and Asia. Between H2 2021 and H1 2024, the average price per kilowatt for colocation in our market sampling increased between 20% and 40% in Europe and Asia, respectively. Colocation prices have largely leveled off since H1 2024 as wholesale electricity rates have stabilized and dropped significantly since their 2022 peak.

Colocation prices in the U.S. have remained largely stable throughout the period of volatility elsewhere. Persistent "local turf wars" and vacancy issues among some operators have essentially kept prices flat. More critically, U.S. markets have not been as exposed as European markets to fluctuations in gas and coal prices.

Singapore is routinely among the most expensive markets we track. With the new licensing regime in the city-state, supply has become incredibly scarce. As a result, median rates surge 20% year-on-year to exceed $600 per kilowatt in H1 2024. At the end of 2023, Singapore's median rate reached an all-time high and has only decreased slightly since then. In Europe, London prices have crept upward, breaching $500 per kilowatt in the last half of 2023, and remaining there this cycle.

Regardless of the predictability of median price changes across markets, the average reported price range (low to high) for base colocation rates within markets can vary. The average range across markets covered in the H1 2024 update cycle was just above $200 per kilowatt between low and high quotes. Some markets had variance as little as $100, while others had differences of more than $400.

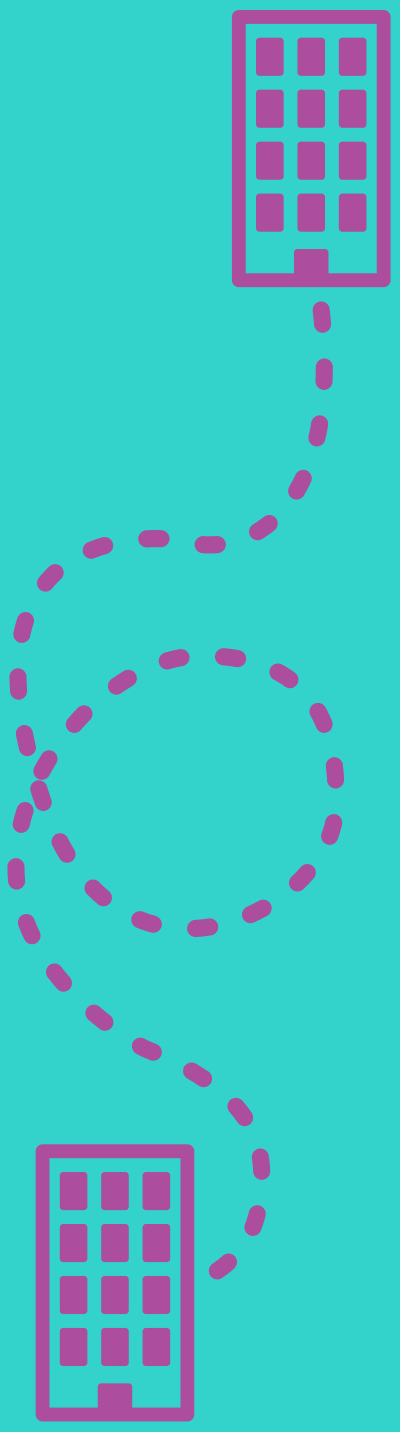

In the eleven years that we've tracked data center pricing, we've observed a consistent gap among North American, Asian, and European cross connect prices—with the prices in North America the highest. The disparity between the U.S. and Europe narrowed leading into H2 2020 but has since widened again, generally hovering around 2x for the past two years.

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. The average TCO in European markets when one cross connect is assumed ($2,000) was about 50% higher than that in North American markets. The gap between average TCO in Europe and the U.S. grew dramatically between 2022 and 2023 (having previously been closer to 15%). This was largely due to the sharp upturn in European prices per kilowatt, contrasted with the continued, muted response to macroeconomic conditions among many U.S. operators.

When five cross connects were assumed in our TCO model, the difference in TCO between these two regions became essentially non-existent. The drastic increase in base colocation pricing seen in European markets over the past few cycles counter-balanced the U.S.'s high cross connect rates.

Expectations

It appears that as of H1 2024, operators are no longer expecting drastic colocation price changes in the near term. Lower, stabilized wholesale electricity rates seem to support the sense that colocation prices will settle down.

Here are a few general trends to watch as we move into 2024.

Falling European electricity rates do not preclude continued volatility for colocation pricing. Complications in power delivery specific to individual countries, government actions, and lags in wholesale contract renewals will be among many factors complicating the relationship between spot rates and the ultimate power prices passed down to data center operators and their customers.

The biggest ongoing concern in the data center market will be the availability of power and space to develop further in key markets. This challenge will be compounded by the rapid growth of generative AI applications. While the scramble for resources needed by power-hungry AI clusters will more directly affect the wholesale side of the market in the near term, the overall effect of data center capacity constraint is bound to hit the colocation market as well.

CLOUD AND WAN

Rolling Out New Regions

The world of WAN services can seem like the Wild West to even the savviest of WAN managers. Like Gary Cooper in High Noon, we try to bring some order to this world with our Cloud and WAN Research Service. We detail cloud connectivity offerings and cloud geographies, as well as international wide area networking (WAN) services of more than 250 service providers. This analysis examines the evolution of WAN services and architecture, geographic coverage, and pricing. We also cover cloud connectivity services (dedicated connections) with profiles and analyses of the major public IaaS cloud service providers and colocation providers that offer cloud on-ramp services.

Cloud Geography Developments

Globally, 23 new cloud regions were launched in 2023. This is on par for the last four years—24 were added in 2022, 27 in 2021, and 25 in 2020. This is down from the peak in 2019 when 41 regions came online. Oracle took the lead in terms of total regions launched with seven new single zone regions. Google and Huawei were next, with both opening four new regions. Microsoft opened three. AWS and Alibaba added two new regions each. IBM rounded out the pack, opening a single region.

Europe added the most new cloud regions in 2023, with eight being launched (up from seven in 2022). The generally underserved Mideast market has been active in cloud infrastructure development the past couple years. Six new cloud regions were launched in 2023, second only to Europe. Similar to the Mideast, Latin America had no cloud regions (except in Brazil) until recently. In the past two years, LATAM has been growing at a fast pace, adding five new regions. New cloud regions continue to be added in Asia and Oceania, though at a less intense pace and in only two countries in 2023. Three new regions launched in China, and two launched

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.

in Australia. No new regions were added in Africa or the U.S. and Canada in 2023, but Google launched a region in South Africa the first month of 2024 and Oracle announced plans for a region in Kenya. U.S. has three more regions in the pipe for 2024.

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.

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.

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 five, but are home to only between 20 and 30 zones each.

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 eight, six, and four additional new regions respectively.

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.

Product Geographic Comparison

IP VPN, EoMPLS, DIA, DWDM, EVPN, SDH, EoSDH, and T-E services are all most widely available in Europe, followed by the U.S. and Canada, and then Asia.

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

Global Cloud Data Center and On-Ramp Locations

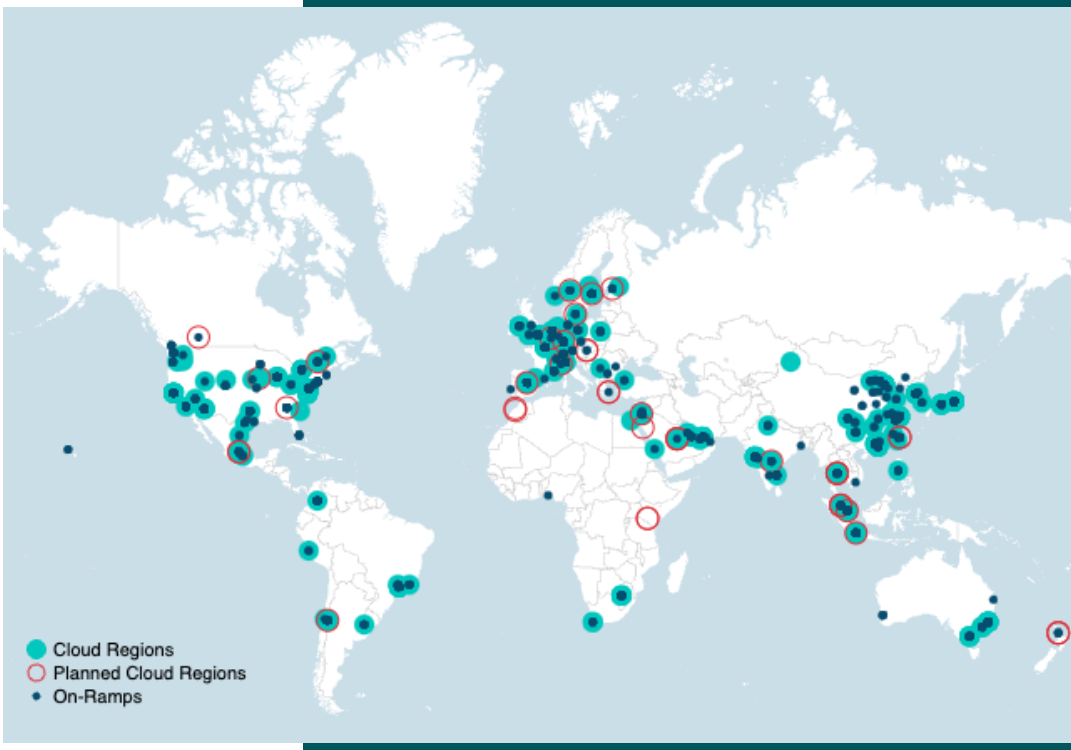

Notes: Data only include IaaS cloud providers from Alibaba, AWS, Google Cloud, Huawei Cloud, IBM, Microsoft Azure, Oracle Cloud, and Tencent Cloud. Data as of Q3 2024.

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.

SASE

The enterprise WAN has significantly changed over the past few years. Many corporate IT infrastructure teams have reduced or eliminated their use of private MPLS networks and replaced them with SD-WAN-enabled hybrid networks, utilizing more cost-effective internet connections. Likewise, corporate data traffic is often now connecting to cloud-based XaaS products rather than on-premises data centers. These new configurations put pressure on legacy “castle and moat” network security postures. In the past three years, Secure Access Service Edge (SASE) has emerged as the network security framework many enterprises are adopting to absorb these changes and remain safe.

SASE is a network security framework that leverages multiple security services and SD-WAN to address challenges introduced by the cloud-based hybrid WAN. There are roughly a dozen SASE products that focus on improving the user-to-app experience while reducing costs and complexity. The Zero Trust Security (ZTS) approach predates the SASE framework, and was effectively subsumed into it. ZTS represents the idea to “never trust, always verify” and is made up of security services that emerged to identify and manage network and application access for specific users and devices.

Key SASE Findings

- The most common SASE products being offered are DDoS protection, Secure Web Gateway (SWG), and Firewall as a Service (FWaaS).
- Multi-Factor Authentication (MFA) is the most common ZTS product.
- The most prominent pricing strategy across all products is to charge per user.
- Palo Alto is the most used security channel partner for carriers/MSPs across all products and services.

INTERNATIONAL VOICE

I'm Goin' Down, Down, Down

The downward slide continues. The year 2014 represents the peak for international voice traffic. International call minutes declined the following year, for the first time since the Great Depression—and it's been downhill ever since.

The slump in voice traffic has turned into a rout. According to our International Voice Report, 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 6.1% in 2021, by 5.8% in 2022, and by an astonishing 7.8% in 2023.

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 telephones, 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

When we compare top international carriers, we note that the top eight operators carried 42% of all global traffic in 2023, about 139 billion minutes.

None of the eight largest carriers in the world terminated more traffic in 2023 than in 2022.

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 2024, with Facebook Messenger topping 1.3 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 2024. These estimates exclude other apps, such as Apple's FaceTime, Google Meet, and Skype (the latter two of which have over 1 billion downloads from Google's App Store).

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.

Traditional carrier traffic has slumped, but OTT traffic has risen to fill the void. Our calculations suggest that cross-border OTT traffic overtook international carrier traffic in 2016, and would exceed 1.3 trillion minutes in 2024, dwarfing the 305 billion minutes of carrier traffic we projected.

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 236 billion minutes of traffic in 2023, down 8% from 2022. Wholesale traffic declined at an average rate of 1% per year over the past ten years, compared to a -2% CAGR for overall traffic. Wholesale carriers terminated nearly three-fourths (72%) of international traffic in 2023, the same as the year before.

Traffic to mobile phones in emerging markets has spurred

expansion in wholesalers' share of the overall market. In 2022, wholesale carriers terminated over 87% of traffic to South America. In contrast, wholesale carriers terminated only 56% of traffic to Western Europe.

Wholesale revenues are down 11% 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. Revenues to fixed lines in both advanced and emerging economies have collapsed. At the same time, revenues from calls to mobiles in all markets have stagnated over the past ten years, at around $11.5 billion.

Who's carrying all this traffic? When we compare top international carriers, we note that the top eight operators carried 42% of all global traffic in 2023, about 139 billion minutes. None of the eight largest carriers in the world terminated more traffic in 2023 than in 2022. Chew on that for a bit.

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 2023, ILD accounts for only 3.1% 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.

Global Retail Revenues from International Calls

USD billions

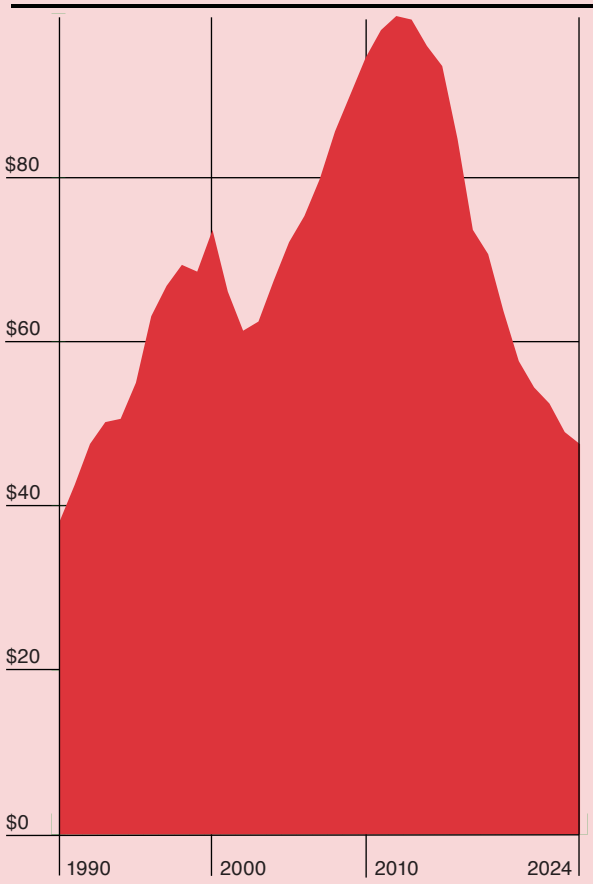

Notes: Data measure retail revenues on outgoing international calls; totals do not include revenue from wholesale services or incoming international traffic termination. Data for 2024 are projections.

Glossary

Addressable Wholesale Capacity—The amount of capacity that wholesale operators are able to sell in the form of managed bandwidth services.

Autonomous System (AS)—Organizes data about IP addresses that are accessible through its network and announces that data across other networks using standardized BGP routing tables.

Autonomous System Number (ASN)—A unique id number that a network must have in order to appear in the global routing tables.

Average Traffic—The sum of all traffic across a link in one month, divided by the number of seconds in the month.

Bandwidth—A measure of information-carrying capacity on a communications channel. May also be referred to as "capacity."

Bandwidth Demand—See Used bandwidth.

Bit—A binary unit of information that can have either of two values, 0 or 1.

Bit Rate—The amount of capacity transmitted by a single wavelength.

Border Gateway Protocol (BGP)—A standardized gateway protocol that exchanges routing information among autonomous systems on the internet.

Channel—Transmission path for a telecommunications signal.

Colocation—The lease of space to house transmission equipment at the same physical location of a carrier or ISP.

Compound Annual Growth Rate (CAGR)—This typically refers to the change in price over a given period of time.

Content Providers—One of the four components of used bandwidth. Includes networks deployed by operators such as Google, Facebook, Microsoft, Amazon, Apple, as well as content delivery networks and many others.

Cross-connect—A physical cable interconnecting equipment (servers, switches, routers) in a data center.

Ethernet—A protocol originally used most frequently in local area networks. Despite its local network origins, Ethernet is a common bandwidth product on long-haul submarine cables.

Fiber Pair—Submarine telecommunications cables contain strands of fiber optic cable. Light is transmitted uni-directionally on fibers; thus, a bi-directional circuit requires a pair of fibers.

High Density—Rack space designated for cabinets with servers that draw more power than standard. We categorize cabinets with 10 kW density or higher as high-density.

Hub Markets—The most critical converging points of global network interconnection. Markets with the most international bandwidth and the largest

interconnection facilities.

Internet Backbone Providers—One of the four components of used bandwidth. Includes the carriers that operate layer 3 IP backbones.

Internet Bandwidth—Refers to the capacity, not average or peak traffic, deployed by internet backbone providers.

Internet Exchange (IX)—A physical location where networks come together to connect and exchange traffic with each other.

Latency—The time it takes for a signal to traverse fiber.

Lit Capacity—The amount of bandwidth available for use on a submarine cable.

Mobile Virtual Network Operator (MVNO)—A wireless communications services provider that doesn't own the network infrastructure it uses to provide services to its customers.

Packet—Generic term for a bundle of data, organized in a specific way for transmission. Consists of the data to be transmitted and certain control information, including the destination address.

Peak Traffic—The 95th percentile of traffic across a link in one month. This is calculated by dividing one month's traffic into five-minute increments, ranking the traffic levels of each increment, and removing the top 5%.

Peering—A practice that allows networks to exchange traffic. The actual exchange of traffic via peering relationships can either be a private transaction between a few operators, or through public arrangements via an internet exchange.

Potential Capacity—The theoretical maximum capacity that a cable could handle with current technology. Often referred to as design capacity.

Purchased Bandwidth—The total of used bandwidth and purchased but unused bandwidth.

Rack Density—The amount of power drawn by servers.

Route Diversity—The need for users of submarine cables to acquire capacity on multiple geographically diverse paths.

Secondary Markets—Markets that are not as large as global hubs but are significant interconnection points on a sub-regional level.

Site Density—The ratio of facility power to data center floor space.

Submarine Cable—A group of optical fiber strands bundled with electrical cabling inside a protective sheath. Cables are laid directly on top of the ocean floor, but are typically buried underneath the sea floor near land, in shallow water, and in areas heavily used by fishing industry.

Upgrade—The installation of additional wavelengths on existing lit fibers or the lighting of previously unlit fiber pairs.

Used Bandwidth—The sum of all capacity deployed by Internet backbone providers, content providers, research and education networks, and enterprises and others. Also referred to as used capacity.

Wavelength—A bandwidth sales product of a single wavelength (usually at a capacity of 10 Gbps or 100 Gbps) on fiber-optic systems employing DWDM.

Research Catalog

Cloud and WAN Research Service

This tool profiles international WAN services offered by 180 providers and analyzes trends in VPN, Ethernet, DIA, and IPL availability and pricing, as well as cloud connectivity services. This unique subscription is also home to:

- **SD-WAN Research**
  The only product that catalogs and analyzes the SD-WAN market so you can find the right fit.

- **WAN Manager Survey**
  This special survey report is a treasure trove of analysis based on the experiences of WAN managers whose day-to-day role covers designing, sourcing, and managing U.S. national, regional, and global corporate wide area computer networks.

- **WAN Market Size Report**
  This vital report presents individual market sizes for key elements of the corporate network broken out by geography.

- **SASE Research**
  This new section analyzes network security offerings and how they are being offered.

Data Center Research Service

A comprehensive online guide for understanding data centers, network storage, and the nature of interconnection.

GlobalComms Database

The most complete source of data about the wireless, broadband, and fixed-line telecom markets.

i3forum Insights

A user-driven voice benchmarking tool for i3forum consortium members; powered by TeleGeography.

International Voice Report

The most comprehensive source of data on international long-distance carriers, traffic, prices, and revenues.

IP Networks Forecast Service

Detailed historical data and forecasts of IP transit service volumes, prices, and revenues by country and region.

IP Networks Research Service

The most complete source of data and analysis about international internet capacity, traffic, service providers, ASN connectivity, and pricing.

Network Pricing Database

A unique database made up of 10 modules that correspond to our 10 network pricing data sets, all of which are available individually.

- **Business Broadband**
  An extensive database of broadband service providers, plans, and prices.

- **Dedicated Internet Access**
  TeleGeography's database of dedicated internet access price benchmarks for corporate and retail customers.

- **Ethernet Over MPLS**
  This database presents information on prices connected to Layer 2, point-to-point Ethernet private line transport service delivered over an MPLS mesh.

- **Ethernet Over SDH or SONET**
  In this module, we track long-haul city-to-city routes between major global business centers.

- **Ethernet VPN**
  TeleGeography's database of layer 2 Ethernet VPN or VPLS services targeted at mid-market/enterprise customers.

- **IP Transit**
  A database of wholesale internet access price quotes by port speed and committed data rate from more than 30 carriers in over 100 cities around the world.

- **Local Access**
  A database of global local access prices, reflecting actual transaction prices paid by carriers for leased private lines and Ethernet circuits.

- **MPLS VPN**
  TeleGeography's price benchmark tracks VPN port and capacity charges at capacity increments between 128 Kbps and 10 GigE.

- **TDM**
  TeleGeography experts routinely survey facilities-based service providers that offer point-to-point private line TDM. Both domestic and international routes are covered in our list of tracked and surveyed routes.

- **Wavelengths**
  In this module, we focus on long-haul city-to-city routes between major global business centers.

Transport Networks Forecast Service

Detailed forecasts of international bandwidth supply, demand, prices, and revenues, updated quarterly.

Transport Networks Research Service

The most complete source of data and analysis for long-haul networks and the undersea cable market.

WAN Cost Benchmark

Provides tailored end-to-end price benchmarks for enterprise wide area networks, based on the client's specified site locations and service requirements.

WAN Geography Benchmark

A WAN Geography benchmark is your personalized cloud and WAN compass. This bespoke tool helps users optimize their network architecture for the cloud.