TeleGeography

The State of the Network

2024 Edition

Sponsored by

telecomegypt

Brought to You By Real Data

Another lap around the sun, a new TeleGeography State of the Network Report—a tradition we've kept since 2017.

No seven-year itch for us, though; we're just getting started. As the world of telecommunications continues to turn, each edition brings even more to report on than the last.

If you're new here, think of this e-book as our 2024 check-in on all things telecom.

After compiling the data and analysis our team dutifully captured throughout 2023, 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.

This report is just the tip of the iceberg (you'll find much more within our full suite of research apps), but it's a great sampling of our core data sets.

What can you expect to read about in our latest report?

For starters, investment is surging to meet demand across all global routes. Many global networks are returning to more typical rates of utilization. Another major disruptive component has emerged in the data center market. And the pace of cloud region expansion seems to be picking up again.

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

— The TeleGeography Team

The Experts Behind This E-book

Lane Burdette
Research Analyst
Connect on LinkedIn

Paul Brodsky
Senior Analyst
Connect on LinkedIn

Patrick Christian
Principal Analyst
Connect on LinkedIn

Jon Hjembo
Senior Manager
Connect on LinkedIn

Alan Mauldin
Research Director
Connect on LinkedIn

Tim Stronge
VP of Research
Connect on LinkedIn

Marvin Tan
Research Analyst
Connect on LinkedIn

Juan Velandia
Research Analyst
Connect on LinkedIn

Recent Presentations

Watch & Download

- 2024 TeleGeography Telecom Workshop
- The Subsea Cold War
- Mythbusters IV: Resurrection
- 2023 TeleGeography Telecom Workshop

TRANSPORT NETWORKS

As Demand Grows, Uncertainty Follows

For much of the world, COVID is largely seen in the rear-view mirror, along with the COVID-driven bump in bandwidth deployments. The bandwidth market now continues merrily along as demand grows across nearly all networks. Our Transport Networks Research Service assesses the state of the global telecom transport industry and evaluates the factors that shape long-term demand growth and price erosion. We assess market conditions on both a global level and on a regional level, focusing on critical submarine cable routes.

Demand Trends

By any measure, the global bandwidth market is thriving. International bandwidth demand has nearly doubled from 2020 to 2022, and has now reached 3.9 Pbps (petabits per second).

Let's break this demand growth down to a more granular level. If we consider used international bandwidth growth by region, two observations jump out. The first is that demand growth has been strongest on links connected to Africa, which experienced a compound annual growth rate of over 50% between 2018 and 2022. The second is that growth in the rest of the world remains strong. Even Latin America saw a 36% compounded annual growth rate over the last five years. While trailing the pack, keep in mind that this annual growth rate implies a doubling of bandwidth every 27 months.

The Role of Content Providers

Who's gobbling up all this international capacity? Historically, it's been carrier networks, provisioning public internet services. As the internet has evolved, major content and cloud service providers—in particular Google, Meta, Amazon, and Microsoft—have become the main sources of demand. Companies like these are the dominant users of international bandwidth, accounting for 71% of all used international capacity in 2022.

The capacity requirements for companies such as these vary in scale and by route. Content providers prioritize the need to link their data centers and major interconnection points. As such, they often deploy massive amounts of capacity on core routes, while focusing much less than traditional carriers do on secondary long-haul routes. To get a sense of this contrast, note that in 2022, content providers accounted for 92% of used capacity on the trans-Atlantic route but just 31% on the Europe-East Asia route.

While the share of content provider capacity on some routes may be much lower than on others, the growth in their demand across all routes has been relentless. A comparison of content providers' international capacity demand growth compared to that of all other networks reveals a stark contrast. Across every region, content providers added capacity at a compound annual rate of at least 41% between 2018 and 2022, compared to a rate no higher than 44% for all the others.

Meeting Demand Requirements

Demand for international bandwidth is nearly doubling every two years. To meet this demand, companies are investing in existing networks and in new infrastructure. The lit capacity on major submarine cable routes continues to soar, keeping pace with demand. Between 2018 and 2022, lit capacity tripled on several routes. The pace of growth was the most rapid on the trans-Atlantic route, where lit capacity increased over 3-fold between 2018 and 2022.

Submarine cable operators are lighting additional

capacity on existing systems. Not only that, but new systems are coming online across all routes. The year 2016 initiated a period of significant global investment in the sector. Cables with a combined construction cost of $8.8 billion entered service between 2018 and 2022, and every major subsea route saw new cables deployed during this timeframe. Investment is expected to surge across all global routes. Based on publicly-announced planned cables, over $11 billion worth of new cables are expected to enter service between 2023 and 2025.

Pricing

Prices continue to decline (somewhat), but the biggest story recently has been how the pace of price erosion compares to previous years. For many key global routes, it has been notably slower—a reflection of different levels of market maturity and delays in supply due to geopolitical challenges and global supply chain issues. Capacity upgrades, which historically took 6-12 weeks from order to installation, rose to 50+ weeks for some vendors.

While this improved over the course of 2022, delays are anticipated to continue throughout 2023. Looking at weighted median 100 Gbps wavelength price trends on major international routes—between 2019 and 2022—weighted median 100 Gbps wavelength prices decreased an average of 13% compounded annually. That's compared to 23% over the prior three years (2016-2019). Trends do, of course, vary by market.

On routes with more ample supply, we see higher rates of price erosion. For example, the U.S.-Latin America route continues to fall at a brisk pace, still feeling the effects of new cables and upgrades to existing systems. While price erosion on Miami-São Paulo was certainly less over the past three years than the historical trend, it is still above the range of 15-20% annual price erosion that we tend to see on most key global routes.

In comparison, on routes with continued delays in new supply, price erosion has stalled. Marseille-Singapore and Hong Kong-Singapore are key examples of this. Wavelength prices on both routes are already extremely competitive and don't have as much room to fall, but the

Construction Cost of Submarine Cables

Construction costs in USD billions

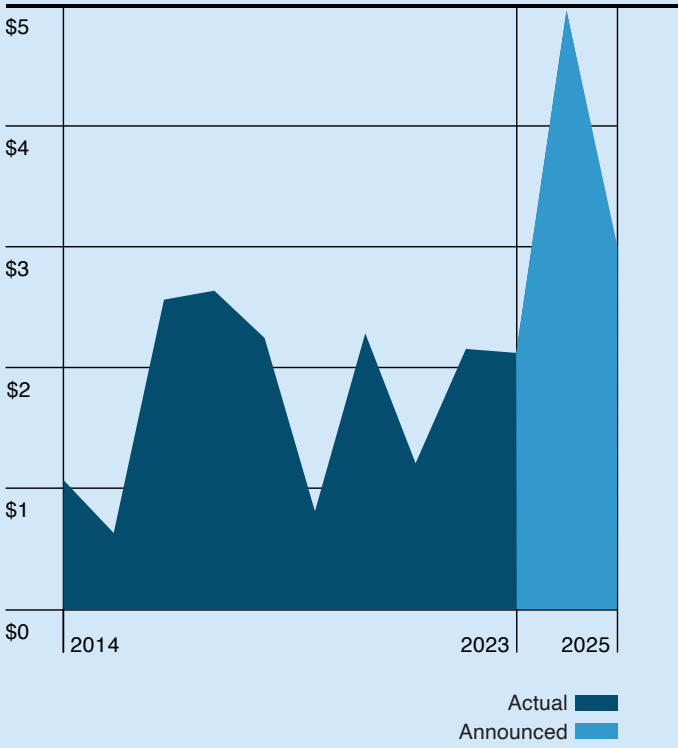

Notes: Total construction costs of all international and domestic submarine cables entering service in designated years. Construction costs exclude the cost of subsequent capacity upgrades and annual operational costs. 2023-2025 construction costs based on announced contract values and TeleGeography estimates. Not all planned cables may be constructed.

Europe-Asia and intra-Asia routes have also been especially impacted by recent delays in supply and for the time being available inventory is going for 2022 prices or potentially higher.

Outlook

What does the future hold for the global bandwidth market? The two most predictable trends are persistent demand growth and price erosion. Beyond that, operators will have to navigate the major uncertainties of an evolving sector. Here are a few of the key trends, among many, that will affect the long-haul capacity market in the coming years.

Rising Utilization

The most fundamental driver for new cable construction is the limited availability of potential capacity. On the surface, this issue may not appear important on major cable routes, where the percentage of potential capacity that is lit has only recently exceeded 50%. However, demand continues to rise at an exponential rate and could soon lead to capacity exhaustion without new cable investment.

Even with the introduction of many new cables and the ability for older cables to accommodate more capacity, the growth of potential capacity has failed to outpace that of lit capacity. If we consider the percentage of potential capacity that is lit on major submarine cable routes, we'll see that it has begun to rise.

Looking at the lit share of potential capacity is not the only way to measure utilization. In fact, the availability of fiber pairs is emerging as a key metric on routes where content providers are involved. Thus, when gauging potential supply on a route, it's important to bear in mind not just how much unlit capacity remains but whether unlit fiber pairs are available as well.

Uncertain Growth for Content Providers

Content providers' international capacity has grown at a rapid rate in recent years, but how long can this last? The recent layoffs at major content providers and declining stock prices have created some questions

about these companies' network investment. Thus far, these issues do not appear to have a material impact on the bandwidth demand growth forecasted by these companies.

Most network planners in these companies focus on meeting expected growth for a 2- to 3-year planning horizon. In our discussions with content providers, all of them have indicated challenges in forecasting their longer-term demand requirements. None of them foresee a decline in demand and continue to anticipate the need for future cable investments. A few aspects that influence growth rates include the following.

Maturing networks. The law of large numbers dictates that a large entity growing rapidly cannot maintain that pace of growth forever. We are certainly seeing evidence of this on major routes. For example, across the Atlantic, annual growth for content providers had been in excess of 80% but has now dipped below 30%. This is a typical pattern for networks as they mature. Even with slowing cumulative growth rates, the incremental volume of bandwidth added each year is still massive. So while global content provider bandwidth growth slowed to "only" 39% in 2022, this still equates to an incremental increase of 783 Tbps.

Artificial Intelligence (AI). The most frequently cited future application that will drive demand is AI. Google, Meta, Microsoft and Amazon have all invested heavily in their own AI models which will increase demands on their network infrastructure. Microsoft's infrastructure is also supporting OpenAI, the company behind ChatGPT. While AI models require substantial compute power, the degree to which AI will impact international bandwidth demand remains unclear.

Multiple product lines and users. Content providers' bandwidth demand comes from a large number of services within each company. In the case of Google, there is search, YouTube, maps, cloud, and many more. It's also worth noting that the bandwidth demand for Google Cloud, AWS, and Microsoft Azure isn't related to these companies' internal demand, but rather has to do with enterprises' implementation and usage of their cloud platforms.

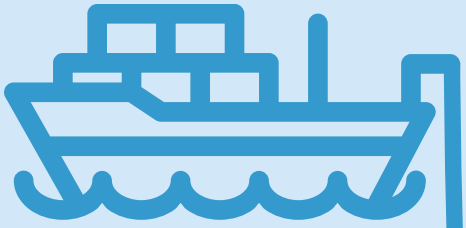

While geopolitical concerns have always played a role in determining which companies deploy long-haul networks and where they do so, several recent developments are reshaping network deployment trends.

Timing of new cables. In recent years, major content provider investments have reduced reliance on carriers and focused on securing enough wholly-owned fiber pairs to achieve sufficient route diversity. Increasingly, new capacity is added largely through the introduction of new cable systems. Thus, annual capacity growth rates observed on some routes could appear lumpy as they are largely influenced by when new submarine cables enter service.

Supply Limitations

The global shortage of chips is continuing to lead to some delays in network upgrades. These issues are improving but may not be fully resolved until 2024. However, other supply side factors could throttle the pace of demand growth in the longer term. There is a limit to how many new submarine cables can be added each year. Cable factories can only produce so many kilometers of cable a year. In addition, there are a limited number of cable laying ships and experienced crews to engage in marine installation. Increasing factory size, the number of installation vessels, and crews will certainly occur, but it takes several years for these measures to be implemented.

Geopolitical Concerns

While geopolitical concerns have always played a role in determining which companies deploy long-haul networks and where they do so, several recent developments are reshaping network deployment trends. In one example, thawing relations between Israel and other Middle Eastern countries has allowed the potential for systems connecting Europe, the Middle East, and Asia to transit across Israel. Several planned projects, including the Blue and Raman cables, hope to capitalize on this opportunity.

In contrast, cable builders find it increasingly difficult to receive Chinese permits for cable deployment in the South China Sea. Operators of the planned Apricot cable hope to avoid this problem by building a cable from Japan to Singapore that runs to the east side of the Philippines. In addition, U.S. government opposition to direct China-to-U.S. cables has encouraged the development of several cables from Southeast Asia to

the U.S. These include Echo, Bifrost, ACC-1, and Hawaiki Nui.

The Europe-Asia route has also been impacted by contemporary geopolitics. China Telecom and China Mobile opted to leave the SeaMeWe-6 consortium cable when American-supplier SubCom was selected as the supplier instead of Chinese-supplier HMN Tech. As a result, the Chinese carriers that left SeaMeWe-6 along with other carriers in Europe and the Middle East are rumored to be planning another cable called Europe-Middle East-Asia (EMA) that HMN Tech would build. The precise landing points and expected activation date are not yet available.

Wholesale Market Challenges

The rapid expansion of major content providers' networks has caused a shift in the global wholesale market. Google, Microsoft, Meta, and Amazon are investing in new submarine cable systems and purchasing fiber pairs. This removes huge sources of demand from the addressable wholesale market. On the other hand, it drives scale to establish new submarine cable systems and lower overall unit costs.

Many submarine cable business models actually rely on this capital injection, allocating fiber and network shares to the largest consumers to cover initial investment costs, then selling remaining shares of system capacity as managed wholesale bandwidth. Unit cost savings of large investments are a great incentive to investment for operators, but they don't want to be left with too much excess bandwidth. It's often a race to offload wholesale capacity before a new generation of lower-cost supply emerges. The carriers most likely to succeed are those with massive internal demand and less dependence on wholesale market revenues.

Both content and carrier network operators are reckoning with massive bandwidth demand growth, driven by new applications and greater penetration into emerging markets. The sheer growth in supply will drive lower unit costs for bandwidth. In the face of price erosion, the challenge for wholesale operators is to carve out profitable niches where demand trumps competition.

IP NETWORKS

Have We Reached Homeostasis?

Three years after the COVID-19 pandemic struck, the internet appears to have achieved a state of normalcy. After a tumultuous 2020, in which the COVID-19 pandemic caused internet traffic patterns to shift and volumes to surge, network operators have returned to the business of adding bandwidth and engineering their traffic in a more measured manner.

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 hard 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 had 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 23% in 2023, continuing to fall from the pandemic-generated bump of 2020. Total international bandwidth now stands at 1,217

Tbps, representing a 4-year CAGR of 28%. COVID bump aside, the pace of growth has been slowing. Still, we do see a near tripling of bandwidth since 2019.

Strong capacity growth is visible across regions. Once again, Africa experienced the most rapid growth of international internet bandwidth, growing at a compound annual rate of 44% between 2019 and 2023. Asia is a distant second, rising at a 32% compound annual rate over the same period.

International internet traffic growth largely mirrors that of internet bandwidth. Both average and peak international internet traffic increased at a compound annual rate of 30% between 2019 and 2023—slightly above the 28% compounded annual growth rate in bandwidth over the same period. All of the stay-at-home activity associated with COVID-19 resulted in a spike in traffic from 2019-2020. The return to more normal usage patterns over the last couple of years has resulted in a substantial drop in average and peak traffic growth. Average traffic growth dropped from 46% between 2019-2020 to 23% between 2022-2023, while peak traffic growth dropped from 45% to 21% over the same time period.

This return to normalcy can be seen across regions of the world. With the initial rapid traffic growth due to COVID-19 continuing to wane in 2023, many global networks appear to have started to return to more typical rates of utilization. Global average and peak utilization rates were essentially unchanged from the year before at about 26% and 44% percent, respectively, in 2022.

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

International Internet Bandwidth Growth By Region

Compound annual growth, 2019–23

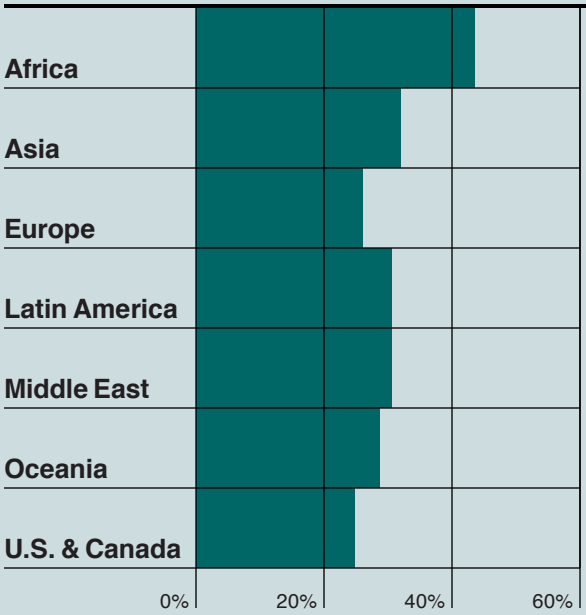

Notes: Data as of mid-year.

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 2020 to Q2 2023. A comparable sample of 100 GigE port prices fell 16% over the same period.

The sharper decline in 100 GigE reflects the advanced maturity of 10 GigE, as well as more carriers offering it—resulting in greater competition. 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 2023, 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, São Paulo, Hong Kong, and Singapore—the Monthly Recurring Charge (MRC) for a 100 GigE port was 6.9 times the MRC for a 10 GigE port. Operators are poised to adopt 400 GigE IP transit ports as the next fundamental upgrade from multiple 100 GigE ports.

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 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:

Post-COVID-19 growth trajectory. Initial evidence suggests that the spike in the rate of bandwidth and traffic growth in 2020 from the pandemic was a one-time event and we have returned to more traditional rates of growth.

IP Transit Price Erosion. International transport unit

Number of Connections for Selected Providers

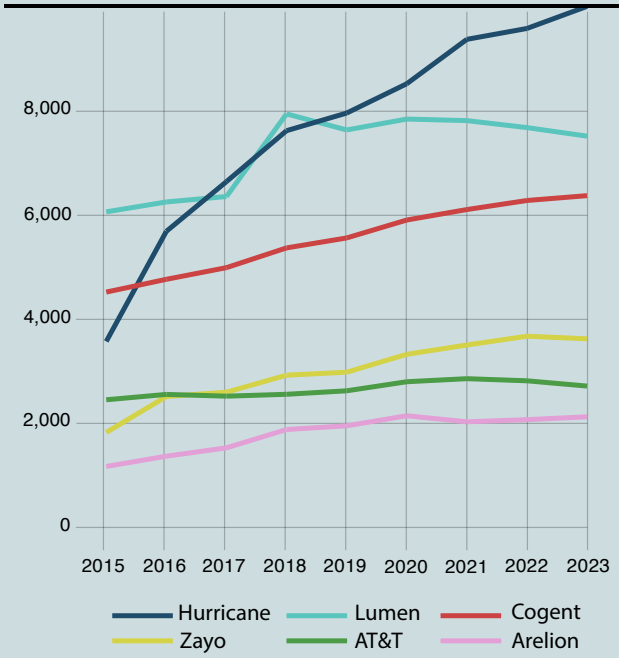

Notes: Data shows the number of connections to other ASNs. The line indicating Lumen's number of connections reflect Level 3 (parent ASN 3356) rather than Lumen (formerly parent ASN 209) prior to 2018.

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.

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.

International versus domestic. While there's little doubt that enhanced end-user access bandwidth and new applications will create large traffic flows, the challenge for operators will be to understand 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.

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, are also building their own private backbones that bypass the public internet. As a result, a rising share of international traffic may be carried by these networks.

Artificial Intelligence (AI). This is the most hyped demand driver in recent years, but its impact on international internet capacity is not entirely clear. A large amount of AI-driven demand is likely to be carried over the private networks of Google, Microsoft, Amazon, and Meta. Microsoft's infrastructure is also supporting OpenAI, the company behind ChatGPT.

DATA CENTERS

Can’t Catch a Break

Recently, we’ve been closely monitoring the intensifying pressures of insatiable demand and supply constraints in key data center markets. Supply constraints have come in the form of both short-term and long-term challenges. In the short term, supply chain disruptions have hindered development timelines. On the long-term side, regulators and utility providers have begun taking a hard look at the data center sector and how to grow it sustainably going forward. In some cases, these entities have severely disrupted development during the interim period.

None of these challenges have been resolved. And as we move through 2023, another major disruptive component has been added to the mix—the accelerated growth of generative Artificial Intelligence (AI). AI will have profound effects on both data center demand and on how data centers will be designed to accommodate vastly more sophisticated operations moving forward. In our Data Center Research Service, we highlight these current obstacles facing the industry.

Of course, the current pains will ultimately produce positive changes. For one, development across a wider distribution of geographic locations could ease constraints on power and space in hub markets. It’s also possible that price volatility in the electricity market could spur an even greater focus on the use of

Largest Retail Operators by Gross Floor Space

Million square feet, 2023

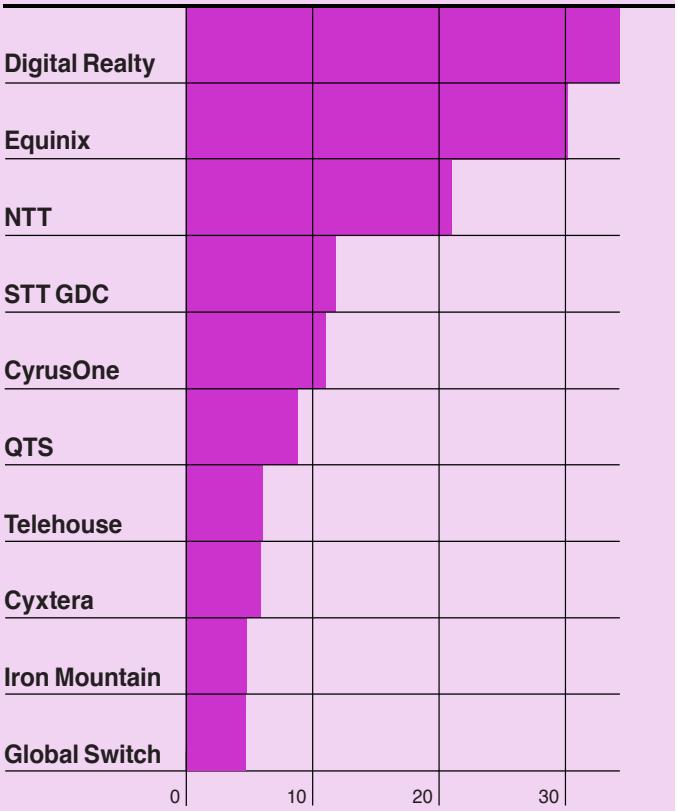

energy-efficient equipment. And, ultimately, these disruptions could drive development of sustainable practices across the data center value chain (e.g., liquid-cooled servers, recycling waste heat, use of renewable energy generation, deployment of onsite generation, gray water cooling, and other solutions).

In the meantime, we continue to see rapid expansion of data center and interconnection market infrastructure across the globe, both in core and developing markets. Network, data center, cloud, and internet exchange operators continue to work together to build new and more widely distributed interconnection nodes.

Let's take a look at the most recent findings from our data center research.

Capacity

Data Center Developments

By our 2023 estimates, The Washington metropolitan area—or more specifically Northern Virginia (NoVA)—dominates as the world's largest data center market. With more than 22 million square feet of operational capacity, NoVA is 30% larger than the next-biggest data center hub, Tokyo.

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.

Digital Realty and Equinix have much greater scale and geographical diversity than all of their competitors. Each of these two operators controls at least 30 million square feet of operational data center capacity. And both of them have significant footprints across every global region. NTT has a footprint about 30% smaller than that of Equinix but nearly twice as big as the next-largest provider.

More than 50 commercial data center providers control over 1 million square feet of operational capacity—each. While few are nearly as large as the behemoths on top of this list, many are growing rapidly, are flush with new investment, and are immensely critical players in the development of the global interconnection market.

Looking at sizable markets with at least 1 million square feet of operational capacity, least 20 of these are growing at 10% CAGR or more. If we include smaller, dynamic markets in that mix, the number reaching that level at least doubles.

A few markets are particularly notable for both their extensive deployed capacity and their surging growth in the past five years: the Latin American markets of Santiago, Campinas, and Santiago de Queretaro; Berlin in Europe; and Mumbai in South Asia. Each has at least 2 million square feet of operational data center capacity and has grown between 24-36% CAGR since 2019. All of these markets have much more capacity in the pipeline, too.

It's also worth highlighting a couple of markets that have seen very low growth. Hong Kong, which is at the center of geopolitical and network deployment challenges, has only grown 2% CAGR over the past five years. Amsterdam and Washington (NoVA) have seen compound growth of about 5% in new capacity. Growth in these markets will further contract as restrictions on new development are imposed in the Netherlands and as the NoVA market waits pensively for new power transmission capacity to come online.

Digital Realty and Equinix have continually led the market in the amount of new capacity deployed and in the geographic diversity of those investments. Digital Realty's new site capacity deployed between August 2021 and August 2023 was spread evenly across five continents. Equinix's focused heavily on Europe and Asia.

Vantage has outpaced even Equinix—at least in sheer gross capacity growth—over the past two years. Since August 2021, the company has rapidly built out large-scale facilities from North America to Europe and further afield in South Africa and Australia.

Among the operators tracked in our database, more than 350 data center sites are known to be in the pipeline right now. While this construction is spread across global regions, Asia outpaces other regions with the largest percentage of new deployments.

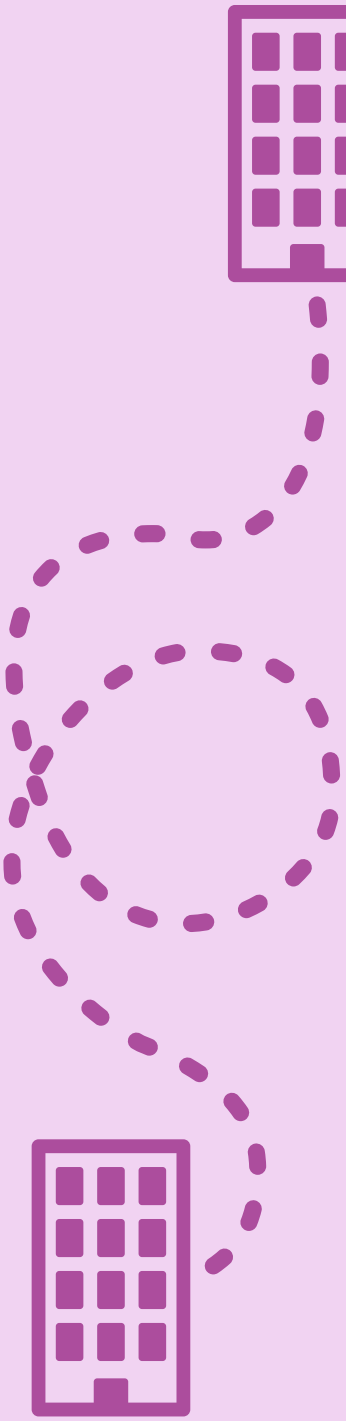

Not all of these data centers will be deployed as soon as hoped. NoVA has nearly a dozen sites in development, but none of those sites will be deployed until power becomes available. Dublin has six, but currently numerous projects are being rejected by the market's local government. On the other hand, Johor Bahru, Kuala Lumpur, and Jakarta all have a half dozen or more data centers in the immediate pipeline, and these are highly likely to be operational within the next few years.

Power

We estimate that, as of 2023, colocation operators in the top ten data center markets consume about 12 gigawatts (GW) of power. That's enough power to generate electricity for roughly 9 million homes—or, in this case, only about 1,000 commercial data center facilities!

Only 11% 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 U.S. & Canada. Operators like Tata, NTT, and China Telecom are ubiquitous in data centers throughout Asia and far beyond; Vodafone, Deutsche Telekom, and euNetworks are heavily represented in European data centers and in other regions; Telefonica, Embratel, Oi, and Flo Networks are among the carriers offering extensive connectivity in Latin American data centers.

By our estimates, SUNeVision's MEGA-i data center in Hong Kong is the most carrier-dense colocation site in the world, though Coresite LA1 (better known as One Wilshire) rivals that position. 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 2019 and 2023. Notable launches in the past year have included international operator deployments in India, Southeast Asia, Africa, and the Nordics along with multiple localized operator IX launches in Italy. A steady stream of new exchanges is slated to come online, most imminently in Europe, Latin America, South Asia, and the Middle East.

Pricing

Current Trends

Starting in 2022 and continuing through 2023, expectations of price inflation became reality in the colocation market—at least in Europe and Asia. Average prices per kilowatt for colocation in our market sampling increased between 40% and 50% respectively over the two years between H2 2021 and H2 2023. In the U.S., despite ongoing expectations that prices will start to rise, ongoing "local turf wars" and vacancy issues among some operators continued to artificially drive prices downward.

In Singapore—the most expensive colocation market we track—supply has become incredibly scarce due to the city-state's new licensing regime. As a result, median rates surged 30% year-on-year to exceed $660 per kilowatt in H2 2023. In our ten years of tracking colocation pricing, we've never seen Singapore's rates this high. Frankfurt has typically been among our most expensive markets as well, although there was a slight decrease in our observed median rate this cycle.

With the sole exception of the New York metro, the U.S. market registered more affordable median colocation rates than all European and Asian markets in our most recent survey.

Between H1 2022 and H1 2023, U.S. cross connect rates rose substantially before correcting back downward in the most recent reporting period. Although utility rates have no direct impact on cross connect

pricing, their inflationary effect on the colocation market is so strong that some operators had increased cross connect prices in order to distribute increased fees across contracted services. In contrast to U.S. markets, European and Asian cross connect prices generally remained steady.

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 ($1,960) was about 60% 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 nonexistent. The drastic increase in base colocation pricing seen in European markets over the past few cycles counterbalanced the U.S.'s high cross connect rates.

On the metro level, Singapore remained untouched as the most expensive market in our entire survey in both the one and five cross connect TCO models. This was unsurprising considering the fact that the median base colocation price was far higher than all other metros surveyed at over $660 per kilowatt and that the median cross connect price in Singapore was also rather high at $180 per month.

Expectations

Operators continued to anticipate price inflation across the market looking forward to 2024, but reported expectations were mixed in H2 2023. Across our full sampling of metros, the median responses ranged from no change to 15% expected inflation in colocation rates for the coming year, with the global average resting around 10%.

Respondents continued to expect cross connect rates to rise as well, though not nearly as much as base

colocation rates. Across our sampling of markets, operators indicated that prices could rise a further 5% over the course of the coming year.

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

Inflation in European electricity costs has settled down significantly following 2022's surge. This development does not preclude continued volatility. 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.

U.S. markets are not immune to the inflationary pressures of the global market. Localized competition will continue to temper the effect, but the surge in cross connect rates we saw earlier in 2023 will likely be followed by some increases in colocation rates.

The biggest ongoing concern in the data center market will be the availability of power and space to develop further in key markets. As regulators and utilities continue to push sustainability goals for the industry, inflationary pressure will continue in markets like Singapore and Frankfurt.

The biggest ongoing concern in the data center market will be the availability of power and space to develop further in key markets.

CLOUD AND WAN

So Many Ways To Get There

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 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 of over 250, Asia is home to the most in-service cloud zones. The United States and Canada follows suit with over 125 zones. Together, these two regions account for about 65% of the world's cloud data centers. The remainder are housed in Europe (19%), Latin America (5%), Oceania (4%), the Middle East (4%), and Africa (2%). At the country level, China and the United States are the clear leaders with close to 140 availability zones for China and 110 for the U.S. Japan, Australia and India round out the top 5, 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

With a total count of over 250, Asia is home to the most in-service cloud zones. The United States and Canada follows suit with over 125 zones. Together, these two regions account for about 65% of the world's cloud data centers.

a similar number of data centers with 23 new regions launched.

The pace seems to be picking up again, with more than 35 planned regions for 2023. If you add in regions planned for 2024 and beyond, there are currently plans to launch close to 50 new cloud regions. Azure leads the pack, contributing 20 new cloud regions. Google, AWS and Oracle are also on the bandwagon, announcing plans for 10, eight, and seven additional new regions respectively. Rounding out the pack, Huawei has plans for two new regions.

WAN Pricing Trends

Trends Across Key Business Centers

MPLS

There is no denying that the prevalence of MPLS in the WAN has decreased as enterprises continue to move to hybrid network designs. In 2022, respondents to our WAN Manager Survey reported that they employed MPLS at 51% of their network sites. That's down from a reported high of 82% in 2018. It is worth noting that this is the first year that MPLS usage saw some stability after the past several years of downward trends. While its role in the WAN is diminishing, MPLS remains a critical component of many enterprise networks, particularly at sites with stricter security or higher SLA requirements. Providers are responding to the perception that MPLS is expensive. Prices continue to decline across geographic regions as providers look to position the service more competitively. Keep in mind, however, that individual prices and price trends vary by market and provider.

Overall, MPLS prices remain highest in developing or remote markets, such as Johannesburg, Mumbai, and São Paulo, where international Layer 1 connectivity is expensive and fewer service providers have PoPs. Markets that are major connectivity hubs and where international capacity is cheap, such as London, New York, and Hong Kong, are the least expensive. Competition reflects the fact that most carriers offering any international service tend to have PoPs in these cities.

DIA vs. MPLS

An optimized WAN routes traffic over the most cost-effective link that supports application performance. Where much of an end-user's traffic is bound to the internet anyway, carrying it over MPLS from the customer premise to a gateway is not only expensive, but also impacts performance. The most common "hybrid" WAN combines MPLS for mission-critical traffic that can't be run over the public internet, with DIA for traffic destined to the internet. This is particularly true where a local breakout will improve the performance of SaaS applications and support the volumes of general internet traffic most companies generate. In most cases, the question is not whether to opt for MPLS or DIA, but rather, what is the appropriate size of each connection—and, when upgrading a site's capacity, where can bandwidth be added most cost effectively?

DIA is universally less expensive than MPLS. In Q4 2022, 10 Mbps DIA connections in key cities were an average of 34% less expensive than a comparable Best Efforts MPLS port. Individual premiums vary dramatically. In New York and Singapore, DIA ports are just 27% and 23% less expensive than MPLS, respectively. In São Paulo and Mumbai, DIA ports are 46% and 45% times less expensive than MPLS, respectively. And in Johannesburg, a 10 Mbps DIA port is a staggering 60% less expensive than MPLS.

Business Broadband vs. DIA vs. MPLS

Business broadband delivers the most cost-effective site connectivity in a hybrid WAN. If we compare the 100 Mbps monthly price of best efforts MPLS, DIA, and business broadband across geographies, business broadband is by far the least expensive option.

On average, across ten major markets, the price for best efforts MPLS is a shocking 63 times the price of broadband. In markets such as New York and Singapore, where broadband prices are a bit higher, a 100 Mbps MPLS port was just six and five times more expensive than a comparable broadband connection, respectively. In markets where broadband prices are low or MPLS remains expensive, the difference can be much larger. For example, a 100 Mbps MPLS port in

Mumbai was 341 times more than a broadband connection.

While DIA is a more affordable option in comparison with MPLS, the average price multiple is still 28 times the average price of broadband. Price gaps again are lowest in Singapore and New York at just two and four times more expensive. Mumbai once again reported the largest difference, with a 100 Mbps DIA port coming in at 133 times more than a comparable broadband connection. With more and more traffic destined for cloud applications, why not take advantage of business broadband? Particularly if cloud on-ramps are in close proximity to users.

SD-WAN and the Hybrid WAN for Cost Optimization

SD-WAN is one tool that assists enterprise customers to integrate internet services into a hybrid WAN. To provide insight into how incorporating SD-WAN into the corporate WAN can impact total network spend, it is useful to look at how these costs apply to a specific network. Looking at the total cost of the overlay and its impact on a network's total cost of ownership (TCO) affords the most apples-to-apples comparison between service providers—particularly with a number of pricing models currently in the market.

To do this level of analysis we created a hypothetical network based on our median WAN Cost Benchmark customer, along with some input from our WAN Manager Survey that queried IT infrastructure managers from around the world about their network configurations. The resulting hypothetical network is comprised of 150 sites spread across major international business centers.

Over the past year, the cost of both the unmanaged and managed SD-WAN overlays in this network scenario have decreased. Thirty-three percent for the unmanaged solution and 17% for the managed solution. As a result, in 2022, the unmanaged SD-WAN overlay contributed just 11% to the network TCO, while the managed solution contributed 27% to the network TCO.

Overall, even after investing in the cost of an SD-WAN overlay, if enterprise customers are able to remove some of their MPLS and integrate DIA or broadband, they may be able to achieve some real network savings or increased network capacity while staying within their existing budget.

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 2,500 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 1,800 times.

Global Cloud Data Center and On-Ramp Locations

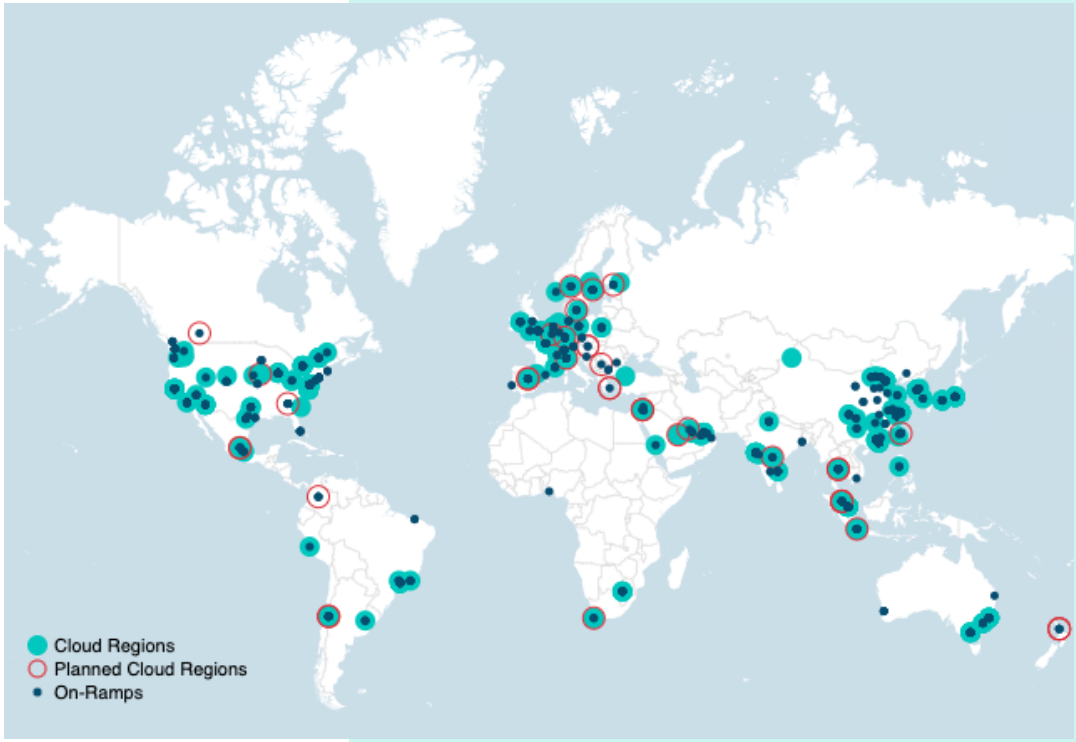

Notes: Data only include IaaS cloud providers from Alibaba, AWS, Google Cloud, Huawei Cloud, IBM, Microsoft Azure, Oracle Cloud, and Tencent Cloud. Circle size reflects number of on-ramps in a given city. Data as of Q1 2023.

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.

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.

The enterprise WAN market is in a state of flux. Cloud computing, the migration of the data center away from corporate premises, local internet breakouts, and the introduction of SD-WAN have significantly disrupted the way multinational corporations design and source their networks. To assist with this task, we created the global enterprise WAN Market Size Report. Based largely on the WAN Manager Survey, which we have published annually since 2018, this report relies on real metrics and data and adds to our assumptions about how these data play out across the globe and into the future.

WAN Market Size

The 2021 global WAN market for the largest multinational enterprises was worth a median value of $59.2 billion according to our model. The range of potential market sizes was $53 to $70 billion. This number includes the key elements of corporate network connectivity: (1) MPLS port charges, (2) DIA port charges, (3) local access charges, (4) business broadband connections, and (5) SD-WAN equipment and encrypted throughput charges.

We also looked at the global market shares of specific connectivity products, and found that:

Access loops to MPLS ports made up the single largest category of global WAN revenue at 31%, which was $18.3 billion.

MPLS was nearly as large at $17 billion or 29%. Together MPLS and MPLS access loops constituted 60% of the global WAN market. This is a key finding given that MPLS revenue is extremely likely to decline over the coming years.

DIA port sizes skewed larger than those for MPLS, so although they are generally slightly cheaper than MPLS ports, these larger ports made up a similar portion of global WAN revenue at $16.3 billion or 28%.

East Asia dominated the global market in WAN revenue in our median model run at 41%—a dollar value of $24 billion. Our model assumes this region to be only 15% of global WAN sites, so the large revenue comes primarily from consistently high prices in parts of the region. The U.S. & Canada was the next largest revenue contributor at $9.8 billion or 17%. This is despite the fact that the region is assumed to represent 30% of global WAN sites.

INTERNATIONAL VOICE

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 full-scale retreat. According to our International Voice Report, carriers' traffic fell by 4.0% in 2018, 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 these standards, 2022 was actually not a bad year, as traffic fell by "only" 5.8%.

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 applications, such as WhatsApp, Facebook Messenger, WeChat (Weixin), Viber, Line, KakaoTalk, and Apple's FaceTime, now pose a greater threat. WhatsApp had about 2.5 billion monthly active users in 2022, with Facebook Messenger topping 1.3 billion. WeChat reported 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 2023. These estimates exclude other apps, such as Apple's FaceTime, Google Hangouts, 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 interna-

When we compare top international carriers, we note that the top nine operators carried nearly half of all global traffic in 2022. That's about 178 billion minutes.

tional traffic in 2013, and OTT plus carrier traffic totaled 761 billion minutes. If we assume that total international (carrier plus OTT) traffic has continued to grow at a relatively modest 13% annually since 2013 (with a drop to 9% in 2018 due to texting, video, and email), the combined volume of carrier and OTT international traffic would have expanded to 1.8 trillion minutes in 2021, and to almost 1.9 trillion minutes in 2022. 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 near 2.1 trillion minutes in 2023, dwarfing the 337 billion minutes of carrier traffic projected by 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 257 billion minutes of traffic in 2022, down 5% from 2021. 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 2022, up from 70% 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 Sub-Saharan Africa and South America. In contrast, wholesale carriers terminated only 56% of traffic to Western Europe.

Wholesale revenues have changed only marginally from ten years ago. But let's take a moment to look under the hood. Over the past decade, traffic to mobile phones in emerging markets has driven international wholesale market growth. As a portion of overall wholesale carrier revenues, calls to advanced economies shrank, as did revenues from calls to fixed lines in emerging markets.

Who's carrying all this traffic? When we compare top international carriers, we note that the top nine operators carried nearly half of all global traffic in 2022. That's about 178 billion minutes. Among the nine largest carriers in the world, only one terminated more traffic in 2022 than in 2021.

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 2013, retail international call revenues (revenues that exclude wholesale revenues and termination payments) generated $99 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 2013.

In 2023, ILD accounts for only 3.5% 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

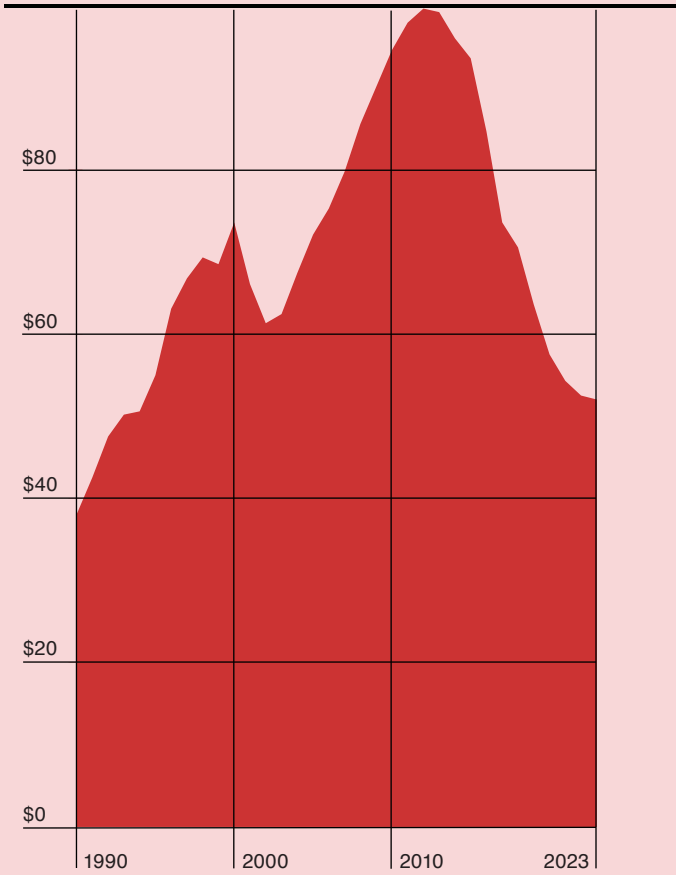

Notes: Data measure retail revenues on outgoing international calls; totals do not include revenue from wholesale services or incoming international traffic termination. Data for 2023 are projections.

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.