TeleGeography

The State of the Network

2023 Edition

Select chapters brought to you by

orange™

i3 forum

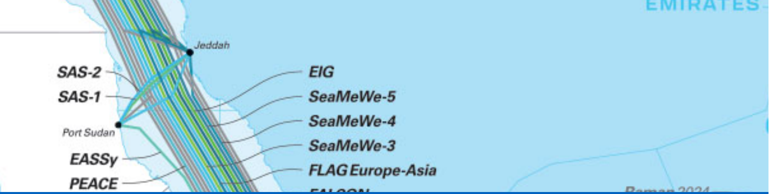

Time to Talk Telecom

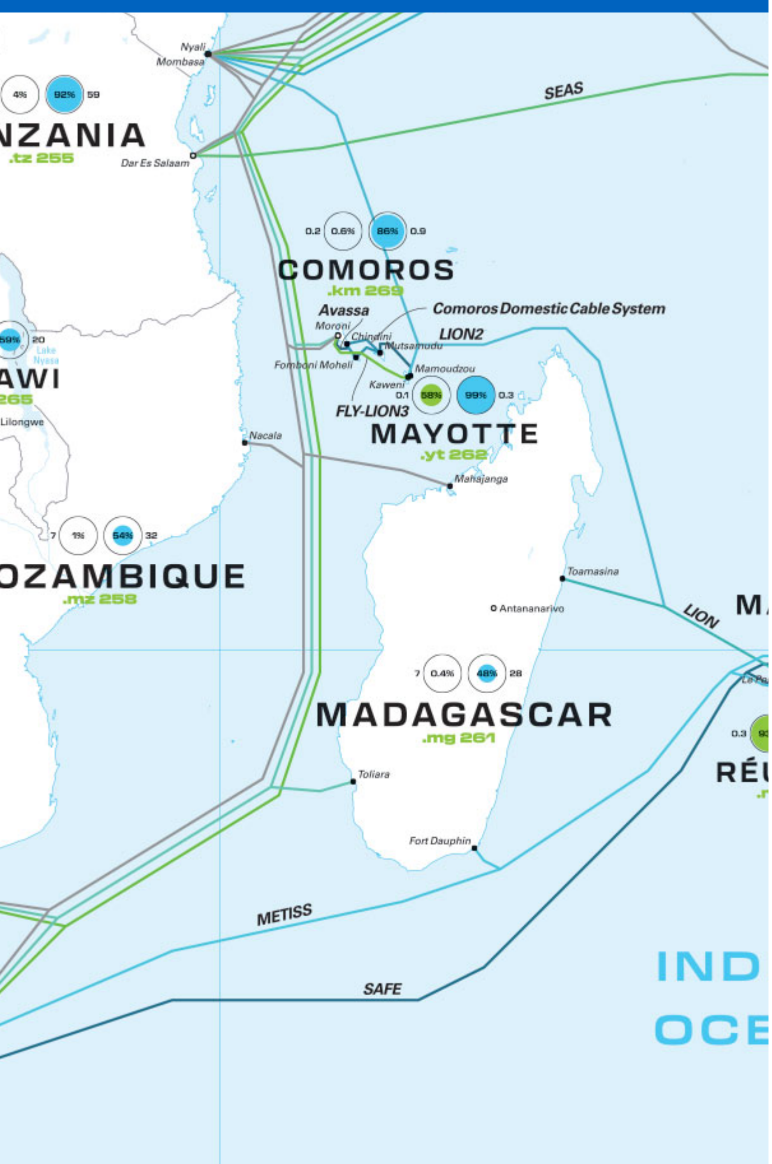

What can we say about our 2023 State of the Network Report?

Let's start with the basics.

For our first-time readers: this is our annual check-in on all things telecom—our sixth annual effort! It's informed by another year of data collection and analysis from TeleGeography's larger research portfolio. (To get specific: this intel was collected throughout 2022. You can find more of it within our full suite of research apps.)

Per usual, 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.

But what's new and different? What's the data telling us in 2023?

Here's a taste of what you'll find in this year's rundown:

- Demand for international bandwidth is nearly doubling every two years.
- 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.
- Critically necessary plans are underway to pursue more sustainable development in the data center sector.
- Close to 50 new cloud regions are expected to launch over the next couple of years.

Taken together, this information pieces together what we'll be talking about in the year ahead, each data point another patch in the quilt that is 2023.

Thanks, as always, for downloading and exploring 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

- 2023 TeleGeography Telecom Workshop
- Economic Indicators and the Future Submarine Cable Industry
- Regional Overview: Latin America Market Trends

GLOBAL BANDWIDTH

Still in High Demand

The COVID-driven bump in bandwidth deployments has largely subsided, but demand continues to grow across nearly all networks. Our Transport Networks Research Service assesses the state of the global telecom transport network 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 route markets.

Demand Trends

If demand is the key factor in assessing the health of the global bandwidth market, then the market is thriving. Between 2019 and 2021 alone, international bandwidth used by global networks doubled to reach nearly 2,900 Tbps.

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 52% between 2017 and 2021. The second is that growth in the rest of the world remains strong. Even Latin America saw a 31% 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 two and a half years.

The Role of Content Providers

Who's driving all this demand growth for international capacity? Historically, it's been carrier networks, provisioning public internet services. More recently a handful of major content and cloud service providers—namely Google, Meta (formerly Facebook), Amazon, and Microsoft—have become the primary sources of demand. These companies are the dominant users of international bandwidth, accounting for 69% of all used international capacity in 2021.

But their capacity requirements vary extensively by route. Content providers prioritize the need to link their data centers and major interconnection points. As such, they often take tremendous 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 2021, content providers accounted for 92% of used capacity on the trans-Atlantic route but just 21% 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 51% between 2017 and 2021, compared to a rate no higher than 45% 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 2017 and 2021, lit capacity more than tripled on many routes. The pace of growth was the most rapid on the trans-Atlantic route, where lit capacity increased nearly four-fold between 2017 and 2021.

In addition to lighting new capacity, 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 $9.2 billion entered service between 2017 and 2021, and every major subsea route saw new cables deployed during this timeframe. Investment is expected to continue across all global routes. Based on publicly announced planned cables, over $10 billion worth of new cables are expected to enter service between 2022 and 2024.

Pricing

Bandwidth prices on primary global routes continue their downward trajectory. Looking at weighted median 100 Gbps wavelength price trends on major international routes—between 2018 and 2021—weighted median 100 Gbps wavelength prices decreased an average of 12% compounded annually.

Prices for 100 Gbps on the core Los Angeles-Tokyo and Hong Kong-Singapore routes fell the least, just 7% compounded annually since 2018. By comparison, the weighted median 100 Gbps price on Miami-São Paulo fell 17%.

Prices continue to decline, but the bigger story this year 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 difficulties. Capacity upgrades, which historically took 6-12 weeks from order to installation, have risen to 30-35 weeks. And these delays are anticipated to continue for the next 12-18 months.

Across the Los Angeles-Tokyo, London-New York, Hong Kong-Singapore, London-Singapore, and Miami-São Paulo routes, 100 Gbps wavelength prices decreased an average of 12% between 2018 and 2021. That's compared to 20% over the past five years. On routes with ample supply, like London-New York, price erosion has been largely consistent with historical trends. In comparison, on routes with limited supply like Los Angeles-Toyko, price erosion has significantly slowed.

Construction Cost of Submarine Cables

Construction costs in USD billions

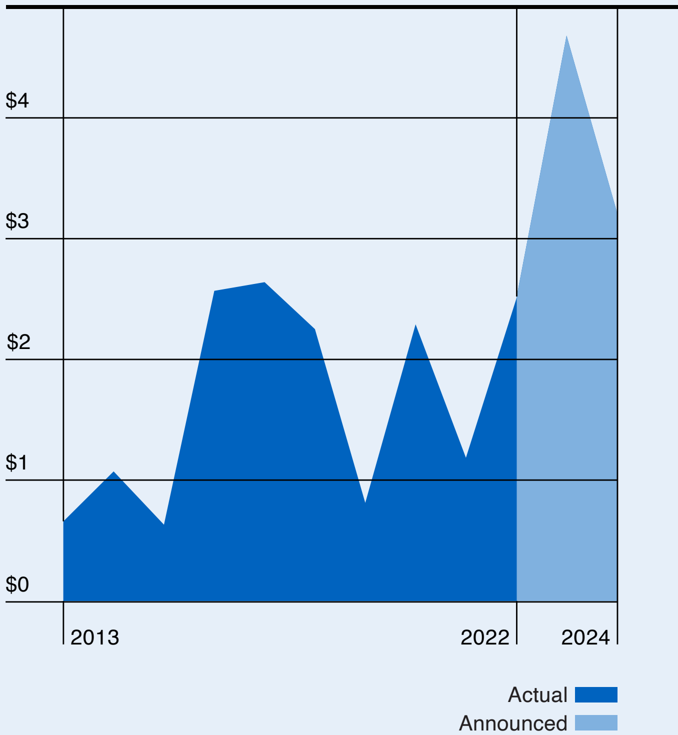

Notes: Total construction costs of all international and domestic submarine cables entering service in designated years. Construction costs exclude the cost of subsequent capacity upgrades and annual operational costs. 2022-2024 construction costs based on announced contract values and TeleGeography estimates. Not all planned cables may be constructed.

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 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 30%.

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? Most network planners in these companies focus on meeting expected growth for a two- to three-year planning horizon. In our discussions with content providers, all of them have indicated challenges in forecasting their longer-term demand requirements.

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. 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 2021, this equates to an annual increase of 563 Tbps.

New applications. Artificial intelligence and virtual reality are most frequently cited as future applications that will drive demand. The degree to which these will impact international demand remains unclear.

Multiple product lines and users. Content providers’ bandwidth demand comes from a large number of services within their companies. 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 on enterprises’ implementation and usage of these cloud platforms.

Timing of new cables. In recent years, major content provider investments have reduced reliance on carriers and have 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.

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. Thawing relations between Israel and other Middle Eastern countries have 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.

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.

In contrast, cable builders are finding it increasingly difficult to receive permits from China to deploy cables in the South China Sea. Operators of the planned Apricot cable hope to avoid this issue 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 boosted the development of several cables from Southeast Asia to the U.S. including Echo, Bifrost, and Hawaiki Nui.

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. 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.

GLOBAL INTERNET

The Dust Is Settling

The internet continues its return to normal—however you define that. After a tumultuous 2020, in which the COVID-19 pandemic caused internet traffic patterns to shift and volumes to surge, network operators 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.

”

Global internet bandwidth rose by 28% in 2022, continuing the return to “normal” from the pandemic-generated bump of 2020.

Total international bandwidth now stands at 997 Tbps, representing a four-year CAGR of 29%.

Internet Traffic and Capacity

Global internet bandwidth rose by 28% in 2022, continuing the return to “normal” from the pandemic-generated bump of 2020. Total international bandwidth now stands at 997 Tbps, representing a four-year CAGR of 29%. COVID bump aside, the pace of growth has been slowing. Still, we do see a near tripling of bandwidth since 2018.

Strong capacity growth is visible across regions. Africa experienced the most rapid growth of international internet bandwidth, growing at a compound annual rate of 44% between 2018 and 2022. Asia came in second, rising at a 35% compound annual rate during the same period.

International internet bandwidth growth largely mirrors that of internet traffic. Average and peak international internet traffic increased at a compound annual rate of 30% between 2018 and 2022—slightly above the 29% 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. As one may expect, the return to more normal usage patterns has resulted in a substantial drop in average and peak traffic for 2021-2022. Average traffic growth dropped from 47% between 2019-2020 to 29% between 2021-2022, while peak traffic growth dropped from 46% to 28% 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 2022, 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 26% and 45% 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 in 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 16% compounded annually from Q2 2019 to Q2 2022. A comparable sample of 100 GigE port prices fell 25% over the same period.

The sharper decline in 100 GigE reflects 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 2022, providers indicated that a majority of their sales mix in key U.S. and European hubs are now 100 GigE.

Customers with the highest traffic commitments receive the best price. IP transit transactions, which are expressed as unit price per Mbps, are lowest for full port allocation. In Q2 2022, the lowest 10 GigE prices on offer were at the brink of $0.09 per Mbps per month. The lowest for 100 GigE were $0.06 per Mbps per month.

On average, across the cities noted, the Monthly Recurring Charge (MRC) for a 100 GigE port is 6.7 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.

International Internet Bandwidth Growth By Region

Compound annual growth, 2018–22

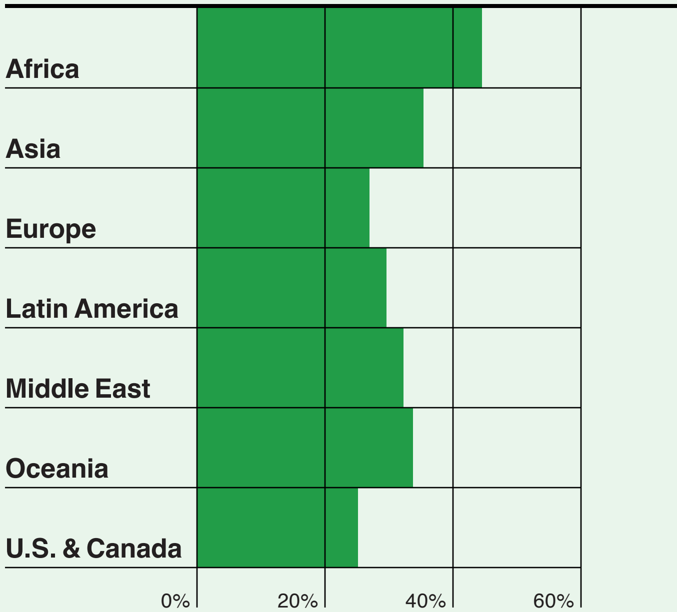

Number of Connections for Selected Providers

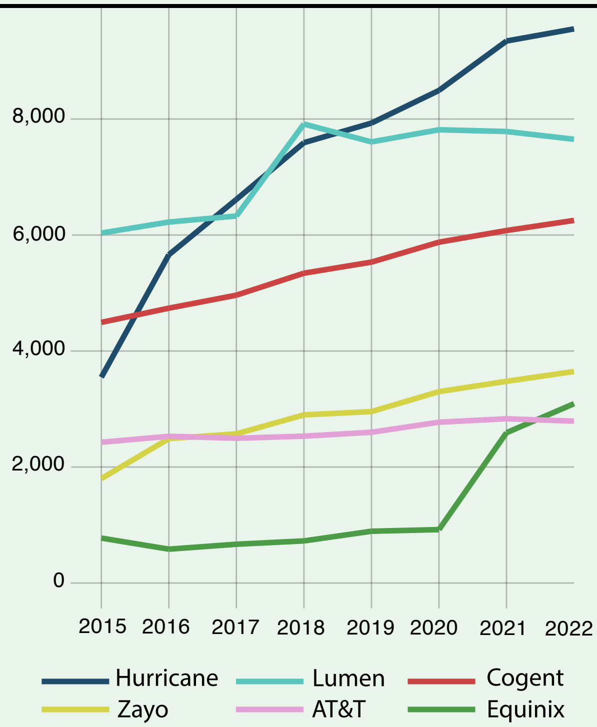

Notes: Data show the number of connections to other ASNs. The line indicating Lumen's number of connections reflect Level 3 (parent ASN 3356) rather than Lumen (formerly parent ASN 209) prior to 2018.

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. Some members of the internet old guard, such as AT&T, have begun to slump in the rankings. Data center operator Equinix has now exceeded AT&T's number of observable AS connections, for the first time ever.

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:

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 largely returned to more traditional rates of growth. Operators we spoke to indicated they no longer see the pandemic leading to upward adjustments to their demand forecasts.

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 requirements 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.

DATA CENTERS

Meeting Major Uncertainties

The data center market has reached a challenging juncture where power usage and industry sustainability have become an immediate and widespread concern. Some of the world's biggest interconnection markets—including Frankfurt, Singapore, Amsterdam, and Northern Virginia—have faced unprecedented disruptions in capacity expansion due to critically important regulator evaluations of the sector's impact. Or, in the case of Virginia, due to a lack of power transmission capacity. These markets and others are reaching inflection points where future development will never again follow the same trajectory as it has in the past.

While critically necessary plans are underway to pursue more sustainable development in the data center sector, an even more acute short-term problem faces the industry. The war in Ukraine is throwing the European wholesale electricity market into chaos. Ultimately, that situation highlights the urgent need to reduce dependence on fossils (like gas), but in the short term, operators just want access to any power available.

Of course, the current pains will ultimately produce positive changes. For one, development across a wider distribution of geographical 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 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.

Capacity

Data Center Developments

By our 2022 estimates, Asian and U.S. metro areas account for 13 of the 15 largest retail colocation markets in the world. In Europe, only London and Frankfurt make the top 15 ranking.

A number of sizable regional edge markets have cropped up around the globe in recent years. APAC markets like Mumbai, Seoul, and Sydney already have at least 2 million square feet of retail colocation space and have continued to grow more than 10% CAGR over the past five years. These locations and many others have become focal points for regional connectivity, with investments from both local and major international colocation providers.

Among operators that primarily derive revenue from retail colocation leases, Equinix's footprint is 50% larger than that of NTT. Among wholesale data center operators, Digital Realty's capacity is 300% bigger than that of its next-largest competitor, STT, at 33.4 million gross square feet.

While most operators concentrate development in a single region, a small group of operators has heavily distributed infrastructure investment across the globe. Among these operators are Digital Realty, Equinix, NTT, STT, and Global Switch. STT far outpaces other operators in recent capacity development, with at least

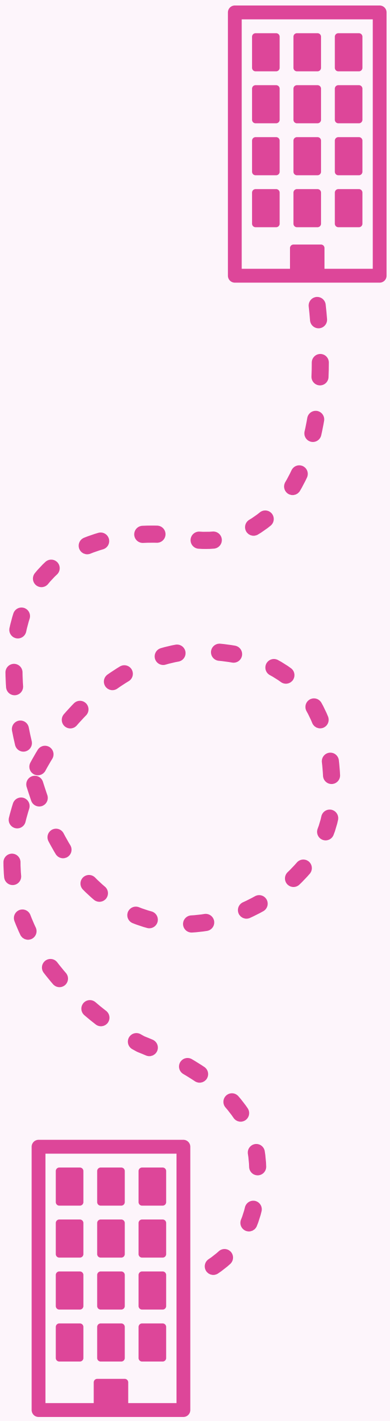

7.5 million square feet of new gross data center space deployed between 2020 and 2022. This includes several sites in India, London, Singapore, and Thailand, along with at least 50 GDS sites across China.

NTT Global Data Centers has pursued aggressive multi-regional growth as well with new sites across the U.S., Europe, and India, as well as Jakarta.

Recent acquisitions since 2020 have fueled continued expansion for Digital Realty and Equinix. For Digital Realty, these include Interxion and Teraco. For Equinix, these include sites from Axtel, Bell Canada Enterprise, GPX, MainOne, and Entel Chile.

Among the operators tracked in our database, more than 250 data center sites are known to be in the pipeline right now. While this construction is spread across global regions, Asia far outpaces other regions with the largest percentage of new deployments.

The U.S., China, India, and Germany top the list of countries with the most new data center deployments scheduled in the next few years. Sites are dispersed across a vast number of individual metropolitan markets, but Northern Virginia, Mumbai, Frankfurt, Langfang, and Shanghai are particularly notable for having between 7 and 12 sites in the near pipeline.

Among the proprietary data center operators tracked in the Data Center Research Service, all are rapidly expanding into new markets. Collectively, Meta (formerly Facebook), Microsoft, Google, and Amazon have deployed 17 new data centers globally (many of which come in the form of cloud service availability zones). Their growth is expected to accelerate over the near term with at least 50 more proprietary sites and cloud region deployments in the immediate pipeline.

Meta currently operates twelve proprietary data center campuses. Several of these campuses including Altoona, Prineville, Sarpy, and Los Lunas are currently operating at least 4 million square feet of data center capacity each. The company has eight more campuses in the pipeline. All but one of these (Singapore) are based in the U.S.

Power

We estimate that, as of 2022, retail colocation operators in the top eight global data center markets consume about 4.2 gigawatts (GW) of power. That's enough power to generate electricity for roughly 3.4 million homes—or in this case, only about 400 retail colocation facilities.

High-density power provisioning has become a priority in the data center industry. Most operators (82%) report that their site density levels exceed 100 watts per square foot (W/sq ft) and fully 40% provide site density levels exceeding 150 W/sq ft.

Operators at most sites (70% of those reporting) support only density levels of up to 10 kilowatts per rack (kW/rack). The share of sites offering the highest density levels exceeding 20 kW/rack is just 12%.

The average site density levels in San Francisco (Silicon Valley), Sydney, London, and Frankfurt exceed 280 W/sqft. This puts their average density levels into the very highest range that we track. On the other end of the spectrum for major markets, average supportable site density levels in Chicago and Singapore are half that of Silicon Valley.

As of 2022, our survey indicates that most sites don't operate at a very low PUE level. A significant minority of sites (40%) operate below 1.5, but that percentage hasn't shifted over the past five years.

Connectivity

Lumen (formerly CenturyLink), Verizon, and Zayo are the most prominent carriers across global facilities. These three operators are especially widespread in the U.S. & Canada. AT&T and Cogent are also ubiquitous in U.S. & Canada facilities, while Colt is heavily represented in European data centers. Operators like China Telecom, Tata, and NTT cover data centers throughout Asia and far beyond; MTN is heavily concentrated in Africa; and Telefonica, Oi, and Embratel are among the carriers offering extensive connectivity in Latin American sites.

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.

Recent IX deployments have been geographically dispersed, with new regional exchanges notably coming online in almost every region of the globe each year between 2018 and 2022. A steady stream of new exchanges are also slated to come online soon, most imminently in Europe, South Asia, and the Middle East.

Pricing

Individual Pricing Components

In the colocation market—unlike bandwidth markets—long term price movements tend to be quite modest. However, as of 2022, we are seeing prices start to rise significantly. In Europe and Asia, prices increased about 15% over the course of the year and are expected to go much higher in 2023. But what's going on with our sampling of U.S. markets? As observed in H1 2022, ongoing “local turf wars” and vacancy issues among some operators continue to artificially drive prices downward. This is expected to change in 2023 and take a strong inflationary turn.

Singapore and Frankfurt are routinely among the most expensive markets we track. H2 2022 was no exception, with median monthly colocation rates in both markets hovering around $500 per kilowatt. U.S. colocation rates have yet to follow the inflationary trend that we're beginning to see in Europe and Asia, and generally declined or remained stable in H2 2022. Each U.S. market covered registered cheaper median colocation rates than all European and Asian markets in the survey.

The price per kilowatt is often discounted at higher density levels. Our global sampling in H2 2022 averaged 7% lower rates per kilowatt for 10-kilowatt cabinets than for standard 4-kilowatt cabinets.

In the nine years that we've tracked data center pricing, we've observed a consistent gap among North American, Asian, and European cross-connect prices. The disparity between the U.S. and Europe narrowed leading into H2 2020 and has generally remained at about 2-3x ever since. That observation continues to play out across metropolitan markets. Nearly all of the median fiber cross-connect rates reported in key U.S. markets for H2 2022 ranged between $200-$275 per month, higher than the medians in all other global markets.

Total Cost Model

Regional differences in base prices per kilowatt and the costs of cross-connects contribute directly to differences in average TCO. In a comparative model assuming one cross-connect for a cabinet, average TCO was about 55% higher in European markets than U.S. markets. The gap between average TCO in Europe and the U.S. grew dramatically between H1 and H2 (having previously been closer to 15%). In Europe, the base cost of the 4-kilowatt cabinet accounted for nearly all of the TCO (93%). In North America, the base rate accounted for about 80% of the total cost. The cost of a fiber cross-connect accounted for the rest, and had a far more significant impact on North American TCO.

When five cross-connects are assumed, the North American average TCO exceeded the European average by nearly 10%. The difference in price movement between the regions was stark. In moving from one to five cross-connects, the average North American price jumped 80%, while it moved just 30% in Europe. The Asian average TCO, meanwhile, jumped 40% to exceed $2,600 per month.

On the metro level, Singapore remained untouched as the most expensive market in our entire survey—unsurprising considering the fact that base colocation prices were far higher than all metros surveyed. Due to their relatively expensive cross-connect fees, every U.S. metro covered in our H2 2022 survey saw average TCO rates jump at least 60% when moving from one to five cross connects in a cabinet.

”

In the colocation market—unlike bandwidth markets—long term price movements tend to be quite modest.

However, as of 2022, we are seeing prices start to rise significantly.

Largest Retail Operators by Gross Floor Space

Million square feet, 2022

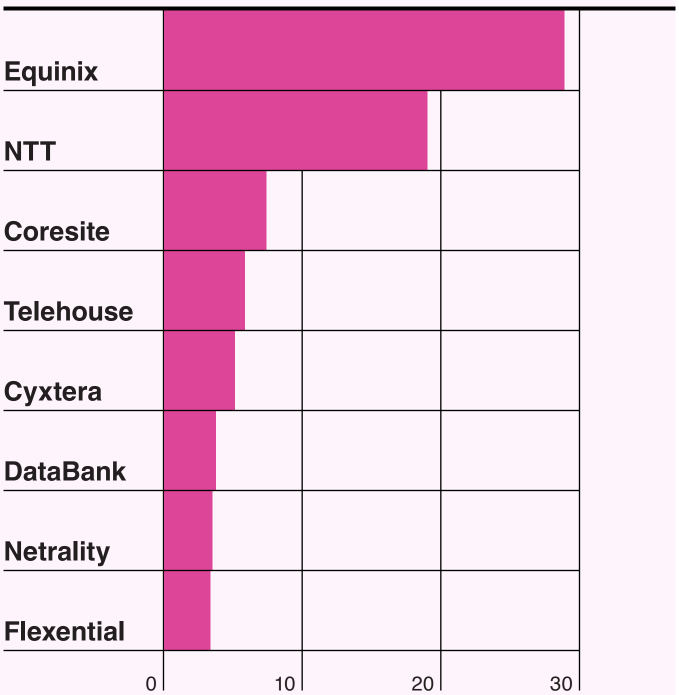

Elsewhere, only Seoul saw a similar increase. All other markets saw increases of 40% or less.

Price Trends

As we start to see the effects of capacity scarcity and high energy costs translating to increasing colocation prices, operators universally expect price inflation throughout the market in 2023. As of H2 2022, numerous operators, particularly in European hubs, indicated that colocation rates would rise dramatically (in some cases as much as 20% or even 40%) in the next 12 months.

Ultimately, pricing reflects response to global macroeconomic indicators, local market dynamics, and unique factors noted at the outset of this section. But it certainly appears that at the start of 2023, the reality of inflation in the colocation market has caught up to the expectations.

CLOUD

Expanding Horizons

The world of WAN services can seem like the Wild West to even the savviest of WAN managers. 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 200 service providers. We 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 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.

Asia is home to the most in-service cloud regions with more than 75 regions. Europe and the United States & Canada each feature almost 50. Together, these three regions account for more than 80% of the world's cloud data centers. Oceania, Latin America, and the Middle East each house slightly more than 10 cloud regions. Africa currently has 4.

Since 2013, cloud providers have launched an average of 15 new cloud regions per year. In 2019, Oracle joined the fray, launching 12 new cloud regions. Among all providers, 37 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 26 new regions in 2020 and 2021, respectively.

There are currently plans to launch close to 50 new cloud regions over the next couple of years. Azure leads the pack, contributing 19 new cloud regions. Google, Oracle, and AWS are on the bandwagon, announcing plans for 13, 8, and 6 additional new regions, respectively. Alibaba just recently launched two new regions in Korea and Thailand.

WAN Cost Overview

A network's total cost reflects how all the services mentioned above are integrated, the network geography, the capacities and traffic prioritization required, and the cost of local access. In addition, contractual matters such as volume and term-length influence total cost.

The following overview provides a breakdown of different elements that comprise the total cost of a global WAN.

Geography. The location of network sites affects cost. Prices are lowest in the most developed markets where competition is well-established and the most modern technologies are densely deployed. The regions with the lowest cost for WAN services are U.S. & Canada, Europe, and parts of Asia-Pacific. Latin America and Africa have fewer competing service providers, benefit less from economies of scale, and lag in deployment of technologies such as Ethernet. As a result prices tend to be higher in these regions.

Technology. Transport technology comprises the main component of the underlying cost of the network. MPLS IP VPN is the most expensive service per Mbps. However, the bandwidth needs are much lower than for other technologies because this transport type is reserved for the most business-critical traffic. EVPN is more expensive than IP VPN at low capacities but much less expensive at higher capacities. EPLs are less expensive per Mbps for higher capacity links than legacy SDH/SONET circuits. DIA is the lowest cost technology and is the most attractive option for offloading non-sensitive traffic to the public internet. We'll get back to this price comparison in a minute.

Traffic Prioritization. MPLS service gives network managers the ability to separate traffic streams by category with CoS designations. Typically there are at least three CoS levels. The highest, CoS1, is reserved for the most sensitive business-critical applications like video or voice calls. For prioritized traffic, most carriers charge a premium that can increase port prices by 20% even in developed markets.

Local Access. Connecting off-site locations to the core network comprises a significant portion of total WAN cost. Prices vary considerably around the globe, and even within the same metro areas. However, developed countries generally report the lowest prices and highest capacities, while circuits in less developed and less competitive countries are lower in capacity and higher in price.

Since 2013, cloud providers have launched an average of 15 new cloud regions per year.

WAN Pricing Trends

Trends Across Key Business Centers

In each pricing review, we examine pricing trends around Best Efforts MPLS and DIA ports as well as business broadband across key global business centers. While its role in the WAN is diminishing, MPLS remains a critical component of many enterprise networks and prices for the service continue to decline across geographic regions.

Overall, MPLS prices remain highest in developing or remote markets, such as Johannesburg, Mumbai and São Paulo, where international layer 1 connectivity is most 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.

Orienting networks towards greater cloud utilization generally requires additional bandwidth at each site. So, how much more expensive is it to increase site capacity in different geographies? Budgeting for these increases requires some metrics. Somewhat surprisingly, increasing capacity 5 times, from 20 Mbps to 100 Mbps, only increases the MRC by a multiple of 2.0 to 2.7.

For MPLS connections, the other consideration is the cost of dedicated local access to connect the customer premise to the carrier network edge. How much local access contributes to a site's total cost varies significantly by market. Local access contributes as much as 62% and 45% of the site cost in New York and London, respectively. It adds as little as 8% in Singapore. Network managers need to consider the cost contribution that local access will make to a site's total price.

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, it can impact performance. The most common "hybrid" WAN combines MPLS for mission critical traffic that can't be run over the public internet, and DIA for traffic destined to the internet. Particularly 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 so much 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 key business centers, at 10 Mbps, Best Efforts MPLS connections are an average of 1.7 times more expensive than DIA. Let's look at each individual market, however, because the individual premium varies dramatically. In Johannesburg and Mumbai, MPLS ports are 3.2 and 2.4 times more expensive than DIA, respectively. In New York and London, MPLS ports are 1.1 and 1.2 times more expensive than DIA, respectively. The price difference seems to be greatest in markets where private, international connectivity is expensive, or where regulation for international carriers is more challenging. Some carriers' MPLS and DIA pricing plans have converged as providers look to maintain their MPLS customer base. But by and large we continue to see a fairly significant difference between MPLS and DIA pricing globally.

Global Cloud Data Center and On-Ramp Locations

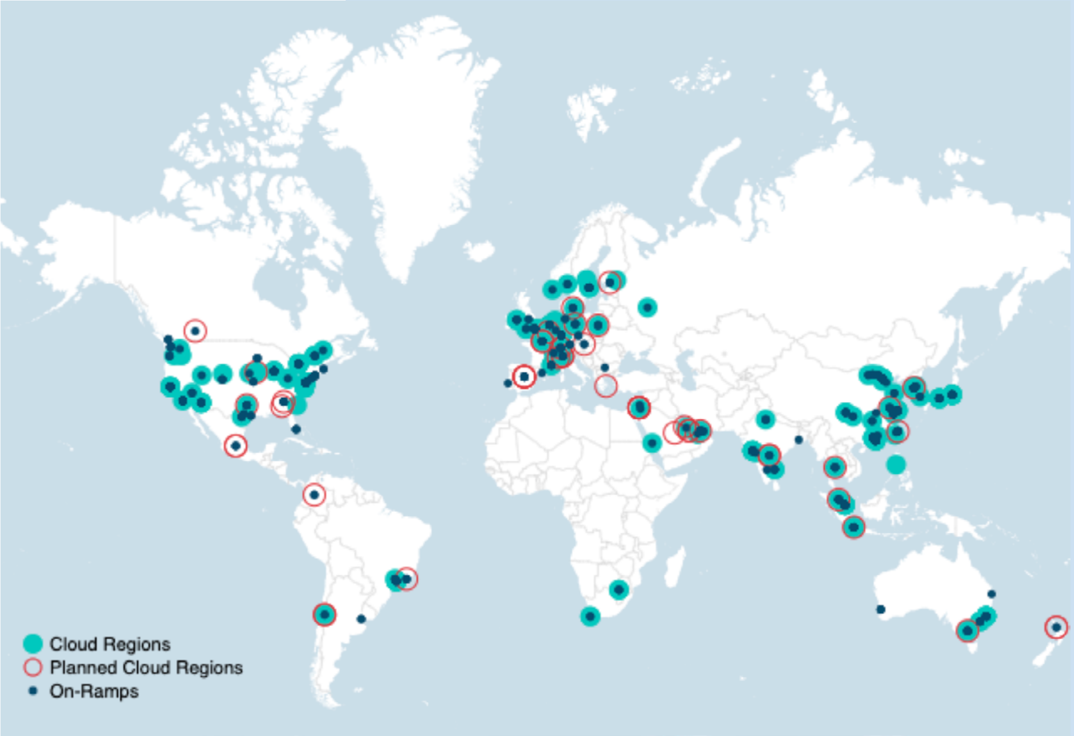

Notes: Data include IaaS cloud provider data center information from Alibaba, AWS, Google Cloud, IBM Cloud, Microsoft Azure, Oracle Cloud, and Tencent Cloud. Data as of Q1 2022.

Business Broadband vs. DIA vs. MPLS

Business broadband is a cost effective alternative to MPLS as part of 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 the markets included below, the average price for best efforts MPLS is a shocking 32 times the average price of broadband. While DIA is a more affordable option in comparison with MPLS, the average price multiple is still 11 times the average price of broadband. With more and more traffic destined for cloud applications, why not take advantage of business broadband.

Broadband service prices are a fraction of MPLS IP VPN or DIA, but broadband is not without its drawbacks. (While corporations can source broadband through an aggregator or carrier, this analysis looks at sourcing directly from the local ISPs.) Broadband services are often contended, may or may not be symmetric, and do not typically include quality of service guarantees like DIA or CoS like MPLS. Further, the price of business broadband varies greatly based on these factors as well as geography, technology, and local regulations. Broadband often reflects ability to pay in the local market, so the quality of service for international traffic may be affected by more contention.

It is a cost-efficient connectivity option for businesses and branch offices where guaranteed speed, security, and service levels are not absolutely necessary, or where other data infrastructure is not available. And SD-WAN overlay often can mitigate many of the performance and security issues associated with broadband. As faster, higher-quality broadband services become more available around the world, broadband will continue to serve enterprise needs.

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 Product Comparison

Carriers reported service availability in 5,244 cities around the globe in 2021. No surprise, MPLS-based IP VPNs are most widely available around the world, particularly so in Europe and the U.S. and Canada.

We also keep on hand a cohort of 165 important business hubs around the world (think Tokyo, Frankfurt, New York). We analyze the geographic reach among carriers for this cohort of important cities as well.

The list of leading WAN service providers includes many well-known names. BT, Verizon, AT&T, Lumen and Orange are the largest providers of IP VPN, EoMPLS, DIA, and EVPN.

Enterprises can find a wide selection of providers in some of the world's most important business centers. There are more than 80 providers of IP VPN in London, 70 in Frankfurt, and 60 in Hong Kong. These cities—and others—tend to be home to high numbers of service providers for other popular WAN servi

VOICE

Incoming Call

You might know that 2015 marked a turning point in the international voice market—the first time since the Great Depression that international call traffic declined, even if only by one half percent. It's been a race downhill ever since.

Carriers' traffic fell by 4% in 2018 and by 6% in 2019 and a further 7% in 2020. The COVID-19 pandemic spurred a short-term rally in international call volumes in early 2020, but things returned to the new normal. Traffic fell a further 6% in 2021, similar to the two previous years.

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 65% 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 billion monthly active users in 2021, with Facebook Messenger topping 1.3 billion. WeChat reported about 1.2 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 nearly 6 billion monthly users in September 2021. These estimates exclude apps for which directly comparable data is unavailable, including 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 international 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

”

When we compare top international carriers, we note that the top eight operators carried nearly half of all global traffic in 2021. That’s about 177 billion minutes.

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 1.6 trillion minutes in 2020, and to almost 1.8 trillion minutes in 2021. 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 1.6 trillion minutes in 2022, dwarfing the 359 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 272 billion minutes of traffic in 2021, down 4% from 2020. 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 2021, up from 70% the year before.

Traffic to mobile phones in emerging markets has spurred expansion in wholesalers’ share of the overall market. In 2021, wholesale carriers terminated over 86% of traffic to Sub-Saharan Africa, and over 87% to South America. In contrast, wholesale carriers terminated only 56% of traffic to western Europe.

Wholesale revenues have changed only marginally from ten years ago. Wholesale ILD generated $8.8 billion in 2011 and $10.7 billion in 2021. 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 eight operators carried nearly half of all global traffic in 2021.

That's about 177 billion minutes. Among the eight largest carriers in the world, only one terminated more traffic in 2021 than in 2020.

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 2022, ILD accounts for only 3.3% of total carrier revenues.

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

Global Retail Revenues from International Calls

USD billions

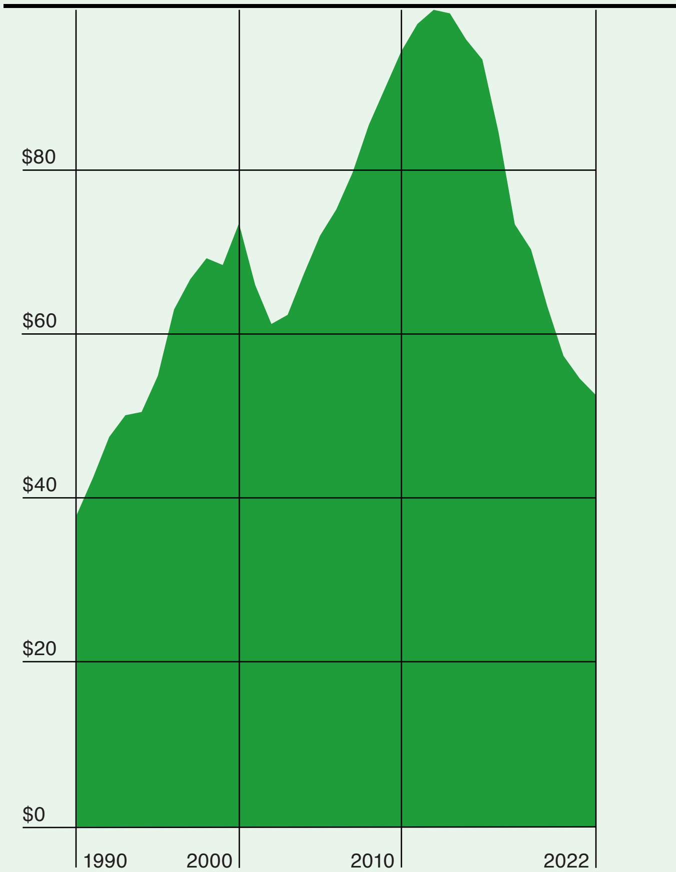

Notes: Data measure retail revenues on outgoing international calls; totals do not include revenue from wholesale services or incoming international traffic termination. Data for 2022 are projections.

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 Service
  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.

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.