TeleGeography

The State of the Network

2022 EDITION

Washington DC, USA

TeleGeography
One Thomas Circle, N.W.
Suite 360
Washington, DC 20005
United States
Tel: +1 (202) 741 0020
Fax: +1 (202) 741 0021

United Kingdom

Primetrica Limited/
TeleGeography
The Station Master's House
New North Road
Exeter, Devon. EX4 4HF
United Kingdom
Tel: +44 (0) 1392 493626
Fax: +44 (0) 1392 420126

Singapore

TeleGeography Asia Pte Ltd
78B Boat Quay
Singapore 049866
Tel: +65 6812-5797
Fax: +65 6812 5889

Telecom in ‘22

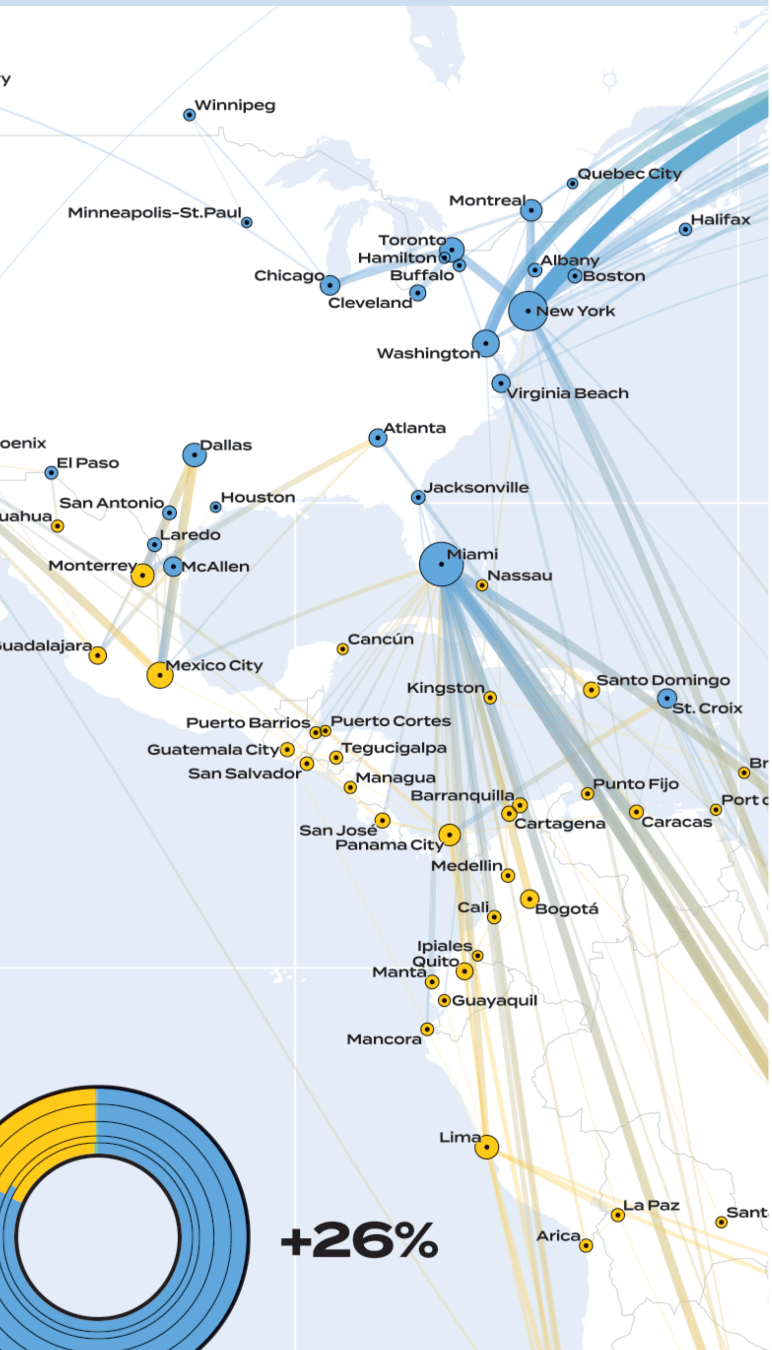

Believe it or not, with 2022 comes our *fifth* State of the Network Report.

To our new readers: welcome. If this is your first State of the Network rodeo, think of this e-book as an annual telecom check-in, informed by another year of data collection and analysis from TeleGeography’s larger research portfolio. As 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.

There’s no other way to put it. This particular State of the Network Report is a weird one.

Why? Well, due to the lingering effects of the COVID-19 pandemic, 2021 saw its fair share of uncertainty. And at the same time, 2021 may be remembered as the year that the internet returned to normal—however you define that.

Yes, working and learning from home dramatically altered traffic patterns. But global internet bandwidth rose by 29% in 2021, which we consider a return to “normal” over the previous year’s COVID-driven surge of 34%.

At the same time, content and carrier network operators continue to reckon with massive bandwidth demand growth driven by new applications and greater penetration into emerging markets. Indeed, strong capacity growth is visible across regions. (Africa experienced the most rapid growth of international internet bandwidth, growing at a compound annual rate of 45% between 2017 and 2021.)

But more on all this later.

As always, this analysis was informed by TeleGeography data. This intel was collected throughout 2021 and you can find more of it within our full suite of research apps.

Thanks for your continued readership and interest in our work. Enjoy this year’s report.

— The TeleGeography Team

The Experts Behind This Ebook

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

Anahí Rebatta
Senior Analyst
Connect on LinkedIn

Tim Stronge
VP of Research
Connect on LinkedIn

Recent Presentations

Watch & Download

- 2022 TeleGeography Telecom Workshop
- The Mystery of International Bandwidth
- Connectivity in Latin America (En Español)
- How to Improve Your IP Traffic Performance and Security

GLOBAL BANDWIDTH

A Very Demanding Industry

The ongoing crisis sparked by the global outbreak of COVID-19 has amplified the international bandwidth market's critical role in keeping the world connected and moving forward. Working and learning from home have dramatically altered traffic patterns, yet the global telecommunications network has proven remarkably resilient in the face of changes.

Bandwidth demand has accelerated across nearly all networks. Operators have felt this increase most acutely in the access networks, but all parts of the global network have been impacted. Many network operators have accelerated their capacity rollouts to stay ahead of demand, and local ISPs have increased caching capabilities to reduce reliance on international links. After an initial spike in traffic with the onset of the pandemic, many operators indicated a stabilizing of demand growth and a return to typical growth rates, building from those elevated levels.

Meanwhile, life—and business—goes on. On the commercial side, operators race to keep revenue margins ahead of eroding prices, while bandwidth demand and supply continue to grow across global routes.

Our Global Bandwidth Research Service assesses the state of the global telecom transport network industry, evaluates the factors that shape long-term demand growth and price erosion, and provides some thoughts on the impact of COVID-19 on the industry. 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 2018 and 2020 alone, international bandwidth used by global networks more than doubled to exceed 2,000 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 56% between 2016 and 2020. The second is that growth in the most developed markets in the world—Europe and U.S. & Canada—wasn't far behind. While mature markets typically grow more slowly than developing markets, that's not really the case when it comes to global bandwidth demand.

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, Facebook, Amazon, and Microsoft—have become the primary sources of demand. As of 2020, these companies are the dominant users of international bandwidth, accounting for two-thirds of all used international capacity.

But their capacity requirements vary extensively by route. Content providers' top priority in their international network planning is 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 2020, content providers accounted for 91% of used capacity on the trans-Atlantic route but just 12% 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 in the following figure reveals a stark contrast. Across six of the world's seven regions, content providers added capacity at a compound annual rate of at least 62% between 2016 and 2020, compared to a rate no higher than 49% for others.

Meeting Demand Requirements

Demand for international bandwidth is more than 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 2016 and 2020, 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 five-fold between 2016 and 2020.

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 $10.9 billion entered service between 2016 and 2020, 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 $8 billion worth of new cables are expected to enter service between 2021 and 2023.

Pricing

The network faced a host of new demands and challenges in 2020, but many of the key trends that characterized the wholesale market prior to the pandemic held true. Demand growth remains robust. And while the pace of price erosion moderated in many markets over the past year, prices still continued to decline. Looking at weighted median 100 Gbps wavelength price trends on major international routes—between 2017 and 2020—weighted median 100 Gbps wavelength prices decreased an average of 16% compounded annually. The average decline across these routes slowed to 6% between 2019 and 2020, a reflection

Construction Cost of Submarine Cables

Construction costs in USD billions

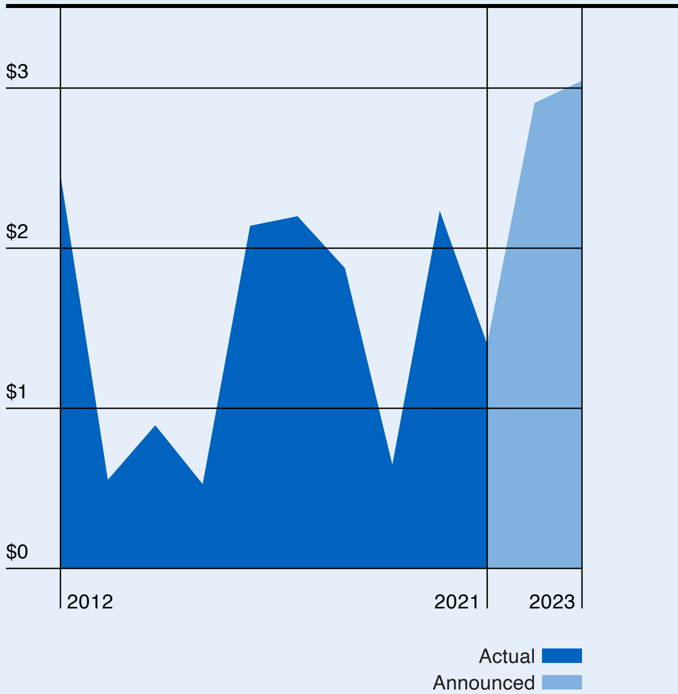

Notes: Total construction costs of all international and domestic submarine cables entering service in designated years. Construction costs exclude the cost of subsequent capacity upgrades and annual operational costs. 2021-2023 construction costs based on announced contract values and TeleGeography estimates. Not all planned cables may be constructed.

of different levels of market maturity as well as delays in supply due to both geopolitical challenges and supply chain issues stemming from COVID-19.

Prices for 100 Gbps on the core Los Angeles-Tokyo and London-New York routes fell the least, just 10% and 13% compounded annually since 2017. Wavelength prices on both routes are already extremely competitive and don't have as much room to fall. But for the Pacific, this was a notable slowdown from previous years. In comparison, the U.S.-Latin America route is still feeling the effects of new cables and upgrades to existing systems after years of scant competition. The weighted median 100 Gbps price on Miami-São Paulo fell the most, 22%. And while the trend lines might look a little flat on some of the other routes in between, prices are still eroding at a steady clip across most global routes.

As 100 Gbps price erosion has outpaced that of the more mature 10 Gbps, it has compressed price multiples between the two services. And this has made it much more economical for customers to upgrade their networks in anticipation of future growth—particularly in 2020 as regional providers reported increased demand for higher capacities as a result of the COVID-19 pandemic.

In Q4 2020, carriers priced 100 Gbps wavelengths an average of 4.3 times higher than 10 Gbps for 10 times the capacity. That's down from 6.4 times more in 2015.

Multiples vary by route, corresponding with regional price differences. Shorter, intra-regional terrestrial links exhibit lower price multiples than longer, transoceanic subsea connections.

We also tend to see low multiples where 100 Gbps adoption is strong, since this drives down unit cost, or in markets where 10 Gbps prices are still high. London–New York has the lowest multiple amongst the subsea routes featured here for two reasons: 100 Gbps prices are low on this route, and sales of 10 Gbps circuits have stalled, resulting in more stable prices.

The same goes for Los Angeles-Tokyo. And while the price multiple on Hong Kong-Singapore might look high (at 6.3), carriers who offer both 10 Gbps and 100 Gbps service on the route reported an average price multiple of 4.7. This falls much closer to other core subsea routes and speaks to the sizable 100 Gbps market on the key intra-Asia connection.

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 and a global pandemic. 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.

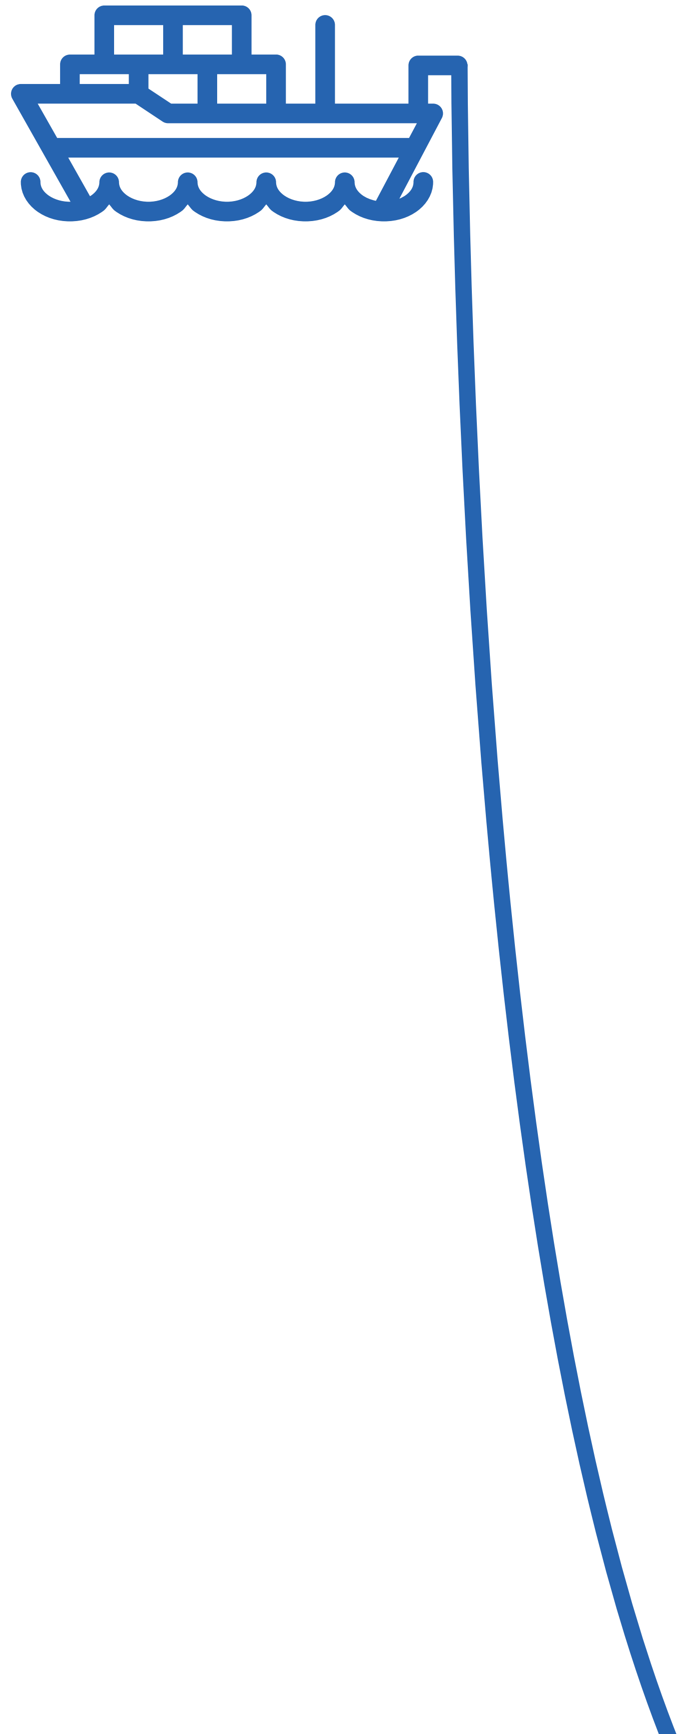

Many older cables laid in the late 1990s and early 2000s may soon become candidates for retirement. This process has already started with the retirements of two trans-Atlantic cables, TAT-14 and Columbus-III, in December 2020, while the PAN-AM cable in Latin America is planned to retire soon.

A few aspects that influence growth rates include the following.

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.

Looming Cable Retirements

Cables are engineered to have a minimum design life of 25 years, but what really matters is the *economic* life. The economic life depends on a cable's revenue exceeding the costs. If the costs of operating a cable continually exceed the revenues, an operator may consider retiring the cable. This could happen well before a cable runs of out capacity. Many older cables laid in the late 1990s and early 2000s may soon become candidates for retirement. This process has already started with the retirements of two trans-Atlantic cables, TAT-14 and Columbus-III, in December 2020, while the PAN-AM cable in Latin America is planned to retire soon.

Addressing the Shannon Limit

Transmission technologies continue to advance, further increasing bit rates. At some point, the industry will face a major challenge as it approaches the Shannon Limit—the theoretical limit of channel capacity given a specified channel bandwidth and signal-to-noise ratio (SNR).

The industry is tackling this problem by taking a multipronged approach. A few of the major strategies include increasing the number of fiber pairs, introducing multi-core fiber, and continuing to introduce more powerful processors. The concept of Spatial Division Multiplexing (SDM) has emerged at the forefront of strategies for increasing subsea cable system throughput moving forward. SDM simply refers to the use of an increased number of paths in a cable (either more fiber pairs or more cores per fiber pair).

Wholesale Market Challenges

The rapid expansion of major content providers' networks has caused a shift in the global wholesale market. Google, Microsoft, Facebook, 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 unrelenting price erosion, the challenge for wholesale operators is to carve out profitable niches where demand trumps competition.

GLOBAL INTERNET

Internet Traffic Goes Full Steam Ahead

The year 2021 may be remembered as the year that the internet returned to normal—however one may choose to 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 Global Internet Geography 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 29% in 2021, a return to "normal" over the previous year's COVID-driven surge of 34%. Total international bandwidth now stands at 786 Tbps, representing a four-year CAGR of 29%. The pace of growth had been slowing, but we still see a near tripling of bandwidth since 2017.

Strong capacity growth is visible across regions. Africa experienced the most rapid growth of international internet bandwidth, growing at a compound annual rate of 45% between 2017 and 2021. Oceania sits just behind Africa, rising at a 38% compound annual rate during the same period.

International internet bandwidth growth largely mirrors that of internet capacity. Average and peak international internet traffic increased at a compound annual rate of 29% between 2017 and 2021—matching the 29% compounded annual growth rate in bandwidth over the same period. All of the stay-at-home activity associated with COVID-19 last year 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 2020-2021. Average traffic growth dropped from 48% between 2019-2020 to 23% between 2020-2021, while peak traffic growth dropped from 46% to 26% 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 waning in 2021, many global networks appear to have started to return to more typical rates of utilization. Global average and peak utilization rates declined slightly to 26% and 45% percent, respectively, in 2021.

Prices

Now that internet backbone operators have adapted their networks to accommodate changes in traffic flows, they've resumed a more measured approach to capacity planning and network upgrades in 2021. Price trends have resumed their downward trajectory and regional characteristics accordingly.

International Internet Bandwidth Growth By Region

Compound annual growth, 2017–21

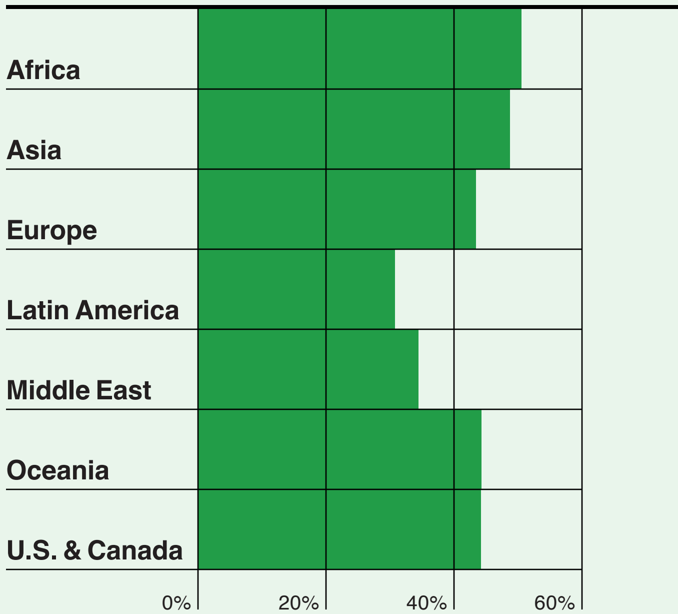

"

Global internet bandwidth rose by 29% in 2021, a return to "normal" over the previous year's COVID-driven surge of 34%.

Total international bandwidth now stands at 786 Tbps, representing a four-year CAGR of 29%.

Across a range of markets, 10 GigE prices fell 18% compounded annually from Q2 2018 to Q2 2021. A comparable sample of 100 GigE port prices fell 30% 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. Most internet backbone operators have 100 GigE deployed, many with multi-100 GigE transactions. Following early speculation that the next increment of port capacity might jump to 1 Tbps, operators are poised to adopt 400 GigE IP transit ports as the next fundamental upgrade from multiple 100 GigE ports.

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 2021, 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.

Price erosion for 100 GigE ports in the cities above has exceeded that noted for 10 GigE ports significantly, 30% versus 18%, attributed to more carriers offering it and more competition. On average, across the cities noted, the Monthly Recurring Charge (MRC) for a 100 GigE port is just over seven times the MRC for a 10 GigE port.

Provider Connectivity

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

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

We've seen little change among the top providers based on this ranking system. Hurricane Electric and Lumen have swapped the top spot for several years. Hurricane edged out then-Level 3 in 2017 as the best-ranked ISP in terms of overall connections, but the Lumen (at that time CenturyLink) merger with Level 3 moved the combined entity back to the top in 2018. Hurricane Electric maintained its lead in 2021.

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

10 Highest Capacity International Internet Hub Cities

International internet bandwidth, 2021

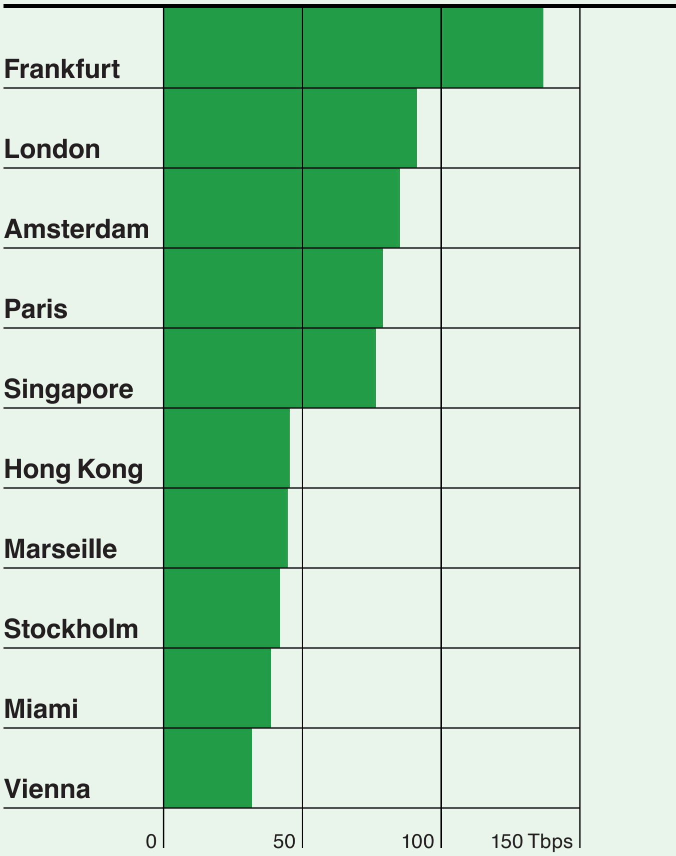

and we have largely returned to more traditional rates of growth. Operators we spoke to indicated they no longer see the pandemic leading to upward adjustments to their demand forecasts.

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 emerge in 2021 and comprise a substantive proportion of the market within three years. 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

New Investors, Same Goal

At TeleGeography, we've observed a recent inflection point in the data center—and broader interconnection—market. While traditional hub markets have seen continual growth, numerous secondary markets and metropolitan nodes on the frontier of network development have seen tremendous new investment by both local and international operators. This investment has involved a mix of network, data center, cloud, and internet exchange operators, working together to build new and more widely distributed interconnection nodes. A confluence of factors has contributed to the trend.

Critically, we're seeing a growth in bandwidth-intensive cloud-based applications that require proximity to end-users, low-latency, and edge computing resources.

High costs can also play a role. Some hub markets are highly competitive when it comes to colocation and network prices, and others are exorbitantly expensive.

Geopolitical and regulatory issues have also increasingly played a role in pushing new network investment beyond traditional hubs. Since 2019, Amsterdam, Frankfurt, and Singapore have all experienced moratoria on the development of new data center sites, as they address land and power constraints. Hong Kong now faces long-term uncertainty about China's role in local governance.

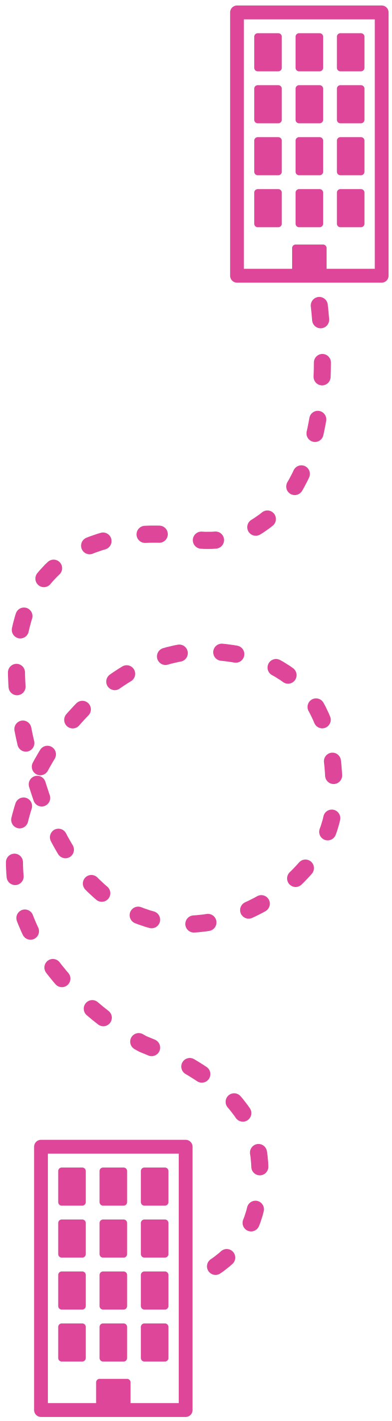

Investment is also chasing profitability in underserved areas. As traditional markets mature and profits erode in the face of increasingly intense competition, operators are looking for first-mover advantages in increasingly nascent markets. This motivation has pushed investment into increasingly volatile markets over the past few years.

Pricing

Individual Pricing Components

As of H1 2021, the European average price per kilowatt for a 4-kilowatt colocation cabinet was about 33% higher than the North American rate. The most recent period saw a bit of adjustment—all attributable to U.S. markets. In H1 2021, we saw universal downward adjustments in U.S. median colocation rates, averaging nearly 15% lower than rates seen in H2 2020. Operator price adjustments did play a role here, but sampling changes were a bigger driver in the lower average.

Among mature global markets, Singapore and Hong Kong are always among the priciest in our survey, but major network convergence points in Europe and North America are also costly places to rent server space. Frankfurt, Hong Kong, and Singapore all registered median rates of at least $424 per kilowatt.

Reported per-kilowatt rates for high-density cabinets (cabinets with 10-kilowatt density) averaged 13% lower than for standard 4-kilowatt cabinets. In Asian markets, the discount was negligible, while in U.S. markets, it reached nearly 20%.

When observing large-scale retail leases (100 kilowatts), we generally see discounts relative to standard leases, but the size of those discounts varies by market and region. Asian markets registered only small discounts, and as with high-density, operators in Hong Kong reported higher prices per kilowatt for larger-scale leases. Discounts per kilowatt averaged more than 15% in European markets and nearly reached 25% in U.S. markets.

The average price multiple for a North American fiber cross-connect was just 2.2 times the average European rate, having narrowed from multiples as high as four times or more as recently as H1 2017. European rates have risen at a steady clip, and averaged nearly $130 per cross-connect as of H1 2021. In Asia, cross-connect rates fall between the European and North American averages but move much closer to the North American rates.

Historically, operators in North America have charged more for fiber cross-connects than for Ethernet, whereas European operators typically charged more for Ethernet cross-connects. Now, most European operators have largely swung in the direction of discounting Ethernet cross-connect fees relative to the cost of fiber cross-connects, with the exceptions of those in Frankfurt and Amsterdam.

Total Cost Model

Regional differences in base prices per kilowatt and the costs of cross-connects contribute directly to differences in average TCO. Among the markets covered in our H1 2021 pricing update, the average TCO in European markets when one cross-connect is assumed was $1,719, about 24% higher than that the $1,383 price in North American markets. The average Asian TCO was close to $2,000. Hong Kong, Singapore, and Frankfurt are among the more expensive markets in the survey, with average total costs of at least $2,000 per month.

When five cross-connects are assumed, the North American average TCO exceeded the European average by a modest 8%, but the difference in price movement between the regions was stark. The average North American price jumped more than 70%, while it moved just 29% in Europe. On the metro level, Hong Kong remained untouched as the most expensive market in our entire survey–unsurprising considering the fact that both its base and cross-connect prices are among the highest of all metros surveyed. New York, Sydney, and Singapore comprised an expensive cluster averaging at least $2,700 per month.

Price Trends

Among operators across all markets, our H1 2021 survey indicated that base colocation rates could rise slightly in the next 12 months. Notably, in Frankfurt, which is already one of the priciest markets in the world, there was some indication that prices could rise between 5%-10% in the near term.

In the Asia region, there has been geopolitical pressure that could ultimately reduce demand growth in Hong Kong, and there are ongoing regulatory controls on data center development in Singapore that could increase demand in an already-constrained market. In Europe, regulatory controls are being implemented on new data center developments in two of the largest hubs—Amsterdam and now Frankfurt. It would come as no surprise if prices in these markets rise as supply becomes more scarce.

And yet, we still see stable growth rates in pricing historically and modest expectations for change going forward.

Capacity and Providers

Market Capacity and Growth

Tokyo remains the largest retail colocation market in the world, with 13.8 million square feet of gross retail capacity reported in 2021—but over 30% of that is accounted for by the various entities of the NTT Group.

The second and third largest markets— Hong Kong and Washington (NoVA)—each have at least a 45% smaller data center footprint than Tokyo.

A number of sizable regional markets have cropped up around the globe in recent years. Montreal, Brussels, Johannesburg, Madrid, Moscow, and Mumbai are particularly noteworthy as markets with at least 1 million square feet of retail colocation space that have also seen rapid growth in the last five years—each with between 11% and 28% gross capacity CAGR.

Long-term gross capacity growth across markets tends to be modest in both large and smaller markets. Between 2017 and 2021, the median compound annual growth rate among a sampling of 95 markets highlighted in the study was just 2%.

Vacancy

Among the metros with sufficient reporting samples, Stockholm, Brussels, and Dallas had relatively high space availability exceeding 30% as of 2021. Dallas was among the markets reporting higher vacancy in our 2020 study as well.

On the opposite end of the spectrum, respondents indicated that fitted colocation capacity in Kiev and Taipei was at least 85% filled.

Average market vacancy among all markets with sufficient reporting hovered around 23%, lower than the average 30% observed across our full sample a year ago.

Providers

With a footprint that's more than 50% larger than its next-biggest competitor NTT, Equinix has more than doubled its gross data center footprint in just five years to reach 27 million square feet of capacity.

While Equinix has increased its global footprint by about 70% since 2017, the next largest operator, NTT, has grown by about 24% in the same time period to reach 18 million square feet.

On the wholesale side, Digital Realty's capacity is 2.5 times as large as that of its next-largest competitor, CyrusOne, at 26 million gross square feet. This excludes Interxion capacity (which we still count as retail capacity) but it includes about 1.7 million square feet of capacity in Latin America from its Ascenty operating unit.

The STT Group of companies (consisting of ST Telemedia Singapore, STT GDC India, GDS Services, STT GDC Thailand, and Virtus) now reaches 10 million gross square feet of capacity. GDS Services has been developing new hyperscale sites across China at a blistering rate, adding 26 sites since September 2019 alone and developing at least 15 more in the near-term pipeline as of August 2021.

Among the operators tracked in our database, at least 200 data center sites are known to be in the pipeline right now. While this construction is spread across global regions, Asia and Europe far outpace other regions including North America with the largest percentages of new deployments.

Data center operators are investing both in edge and core markets for future development. Retail operators are doubling down in large markets like London, Frankfurt,

> **Long-term gross capacity growth across markets tends to be modest in both large and smaller markets. Between 2017 and 2021, the median compound annual growth rate among a sampling of 95 markets highlighted in the study was just 2%.**

Largest Retail Operators by Gross Floor Space

Million square feet, 2021

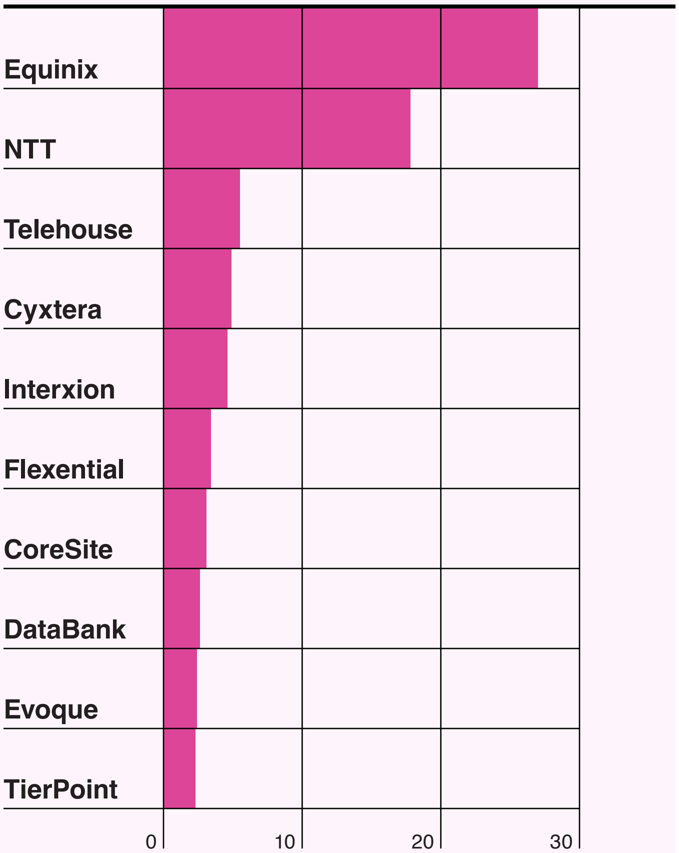

and Singapore—but smaller markets like Berlin and Mumbai are well-represented too.

Likewise, wholesale construction runs the gamut from the very largest markets like Washington, Singapore, and London, to secondary Chinese markets and other secondary locations like Bangkok and Santiago de Queretaro.

Proprietary Data Centers

Among the proprietary data center operators tracked in the Data Center Research Service, all are rapidly expanding into new markets. Collectively, Facebook, Microsoft, Google, and Amazon have deployed 13 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 46 more proprietary sites and cloud region deployments in the immediate pipeline.

Facebook alone currently operates 12 proprietary data center campuses with 19.1 million square feet of operational capacity and room for further growth. That's up more than 44% from their reported operational capacity just one year ago. The company has six more campuses in the pipeline.

Power

We estimate that as of 2021, retail colocation operators in the top eight global data center markets consume about 5 gigawatts (GW) of power. That's enough power to generate electricity for roughly 1.5 million homes—or in this case, only about 440 retail colocation facilities.

Despite increased interest in high-density service provisioning, reported density levels haven't shifted much. At the highest levels we track, only about 20% of sites currently provision site density levels exceeding 200 W/sqft, and that proportion hasn't dramatically shifted in at least the last seven years.

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 11%.

The average site density levels in Dallas and San Francisco (Silicon Valley) exceed 250 W/sq ft. This puts their average density levels into the very highest range that we track. Dallas also has a far above-average supportable rack density level of 20 kW/rack. On the other end of the spectrum for major markets, London and Paris have much lower average site and rack density levels of less than 130 W/sqft and under 10 kW/rack.

As of 2021, 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 four years.

Connectivity

As in the previous few years, 2021 respondents indicated that Lumen (formerly CenturyLink), Verizon, and Zayo are the most prominent carriers in their facilities. These three operators are especially widespread in North America. AT&T and Cogent are also common in North American facilities, while Colt, GTT, and BT are heavily represented in European data centers. Telstra, China Telecom, Tata, and NTT are among the most ubiquitous carriers across Asian sites; 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 Equinix's Kleyerstraße 90 site in Frankfurt rivals that position. TELEHOUSE's London Docklands campus, CoreSite's One Wilshire carrier hotel (624 South Grand), and Equinix's Ashburn 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. In four of the past five years, more than 20 new peering platforms cropped up. Recent deployments have been geographically dispersed, too, with new regional IXs notably coming online in almost every region of the globe each year between 2017 and 2020.

CLOUD

Coming to a Region Near You: The Cloud

You can't talk about demand in the network world without considering cloud infrastructure. Our Cloud and WAN Infrastructure research has taught us as much.

Consider this. 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, 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.

Geography of Cloud Connections

The cloud is not ubiquitous. Dedicated connections to infrastructure as a service (IaaS) providers generally map closely to where those providers have deployed data centers. There are two separate sets of data center groupings relevant to dedicated interconnect:

CSP data centers. These locations are the sites where enterprise data are stored and processed in “the cloud.” More often than not, cloud providers’ data centers are separate from colocation facilities, and are housed in private buildings exclusively for the use of the CSP.

On-Ramps. These locations are the points of interconnection between enterprises and cloud service providers or network to network interconnections (NNIs). These are often located at colocation facilities. Colo facilities are sometimes called third-party data centers or multi-tenant data centers but, to avoid confusion with cloud providers’ private data centers, we will refer to them as “colocation facilities.”

An enterprise or an enterprise’s network service provider partner is responsible for bringing its data traffic to an on-ramp. A CSP takes responsibility for hauling traffic between the dedicated connection location and its own data centers.

CSP Data Centers (Availability Zones and Regions)

The cloud network infrastructure of AWS, Microsoft Azure, and Google Cloud is organized around geopolitical regions and zones in proximity to clusters of data centers. Terminology and architecture vary somewhat between the three providers:

Zones. Availability zones typically consist of multiple data centers in a geographic region or on a campus.

Regions. Regions are territories (generally, a metro area) that comprise one or multiple zones. All AWS and Google regions contain multiple zones; some of Azure’s regions hold just one zone.

”

AWS launched the first cloud region in 2006. Microsoft launched its first Azure region in 2008. Alibaba launched its first in 2011. The pace of new deployments has surged since then.

Cloud network infrastructure is designed to provide redundancy, availability, and fault tolerance. Data centers that make up a zone are physically separated from each other. They may not be separated by great distance, and may even be in the same site, but they are completely isolated from each other in terms of power source, cooling, and network connectivity. Individual data centers, availability zones, and geopolitical regions are all configured to operate independently. Special zones (or "national clouds") are isolated from other regions, often for regulatory compliance or other legal purposes. These are typically private cloud platforms for governments like AWS' GovClouds in the U.S., or for specific countries like Azure's national clouds in China and Germany.

The Roll Out of Regions

AWS launched the first cloud region in 2006. Microsoft launched its first Azure region in 2008. Alibaba launched its first in 2011. The pace of new deployments has surged since then. By 2014 AWS had launched 10 regions. In that year alone Microsoft launched 10 new regions for a total of 18 overall. The following year, Google launched its first cloud regions (four, in fact) and Alibaba launched its first cloud region outside of China.

From 2006 to 2013 the four major cloud providers launched on average a bit over 2 regions per year. Since then? On average cloud providers have launched a whopping 15 new cloud regions per year. In 2019, Oracle joined the fray, launching 12 new cloud regions. Among all providers, 28 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, cloud operators managed to launch 21 new regions in 2020 and 18 more in 2021.

Planned Regions

There are currently plans to launch 49 new cloud regions over the next couple of years. Azure leads the pack with plans to launch 21 new regions. Google has announced plans for 10 additional regions, AWS 8 and Oracle has 7 in the pipeline. Alibaba has just announced its plans for two new regions in South Korea and Thailand.

Of the 49 new planned regions, the lion's share (16) will arise in Europe—Madrid, Milan and Paris will have two to three new regions each. The Middle East is growing in importance on the cloud map, with six planned additions. Israel and Qatar will both see two new regions.

In Latin America, Mexico will see its first two regions, both located in Querétaro; and Chile will see its second region in Santiago. Also, a new region is planned to launch in Brazil— in Rio de Janeiro. Eight new regions are planned for Asia, both in developing hubs like Bangkok, Kuala Lumpur, and Jakarta; as well as in more established hubsin Beijin, Hyderabad, and Seoul.

Three new regions are planned in Oceania, where New Zealand will see its first region located in Auckland. There is one new region planned for Africa in Johannesburg to add to the two already active. Finally, four new regions are planned in the United States and Canada.

Global Cloud Data Center and On-Ramp Locations

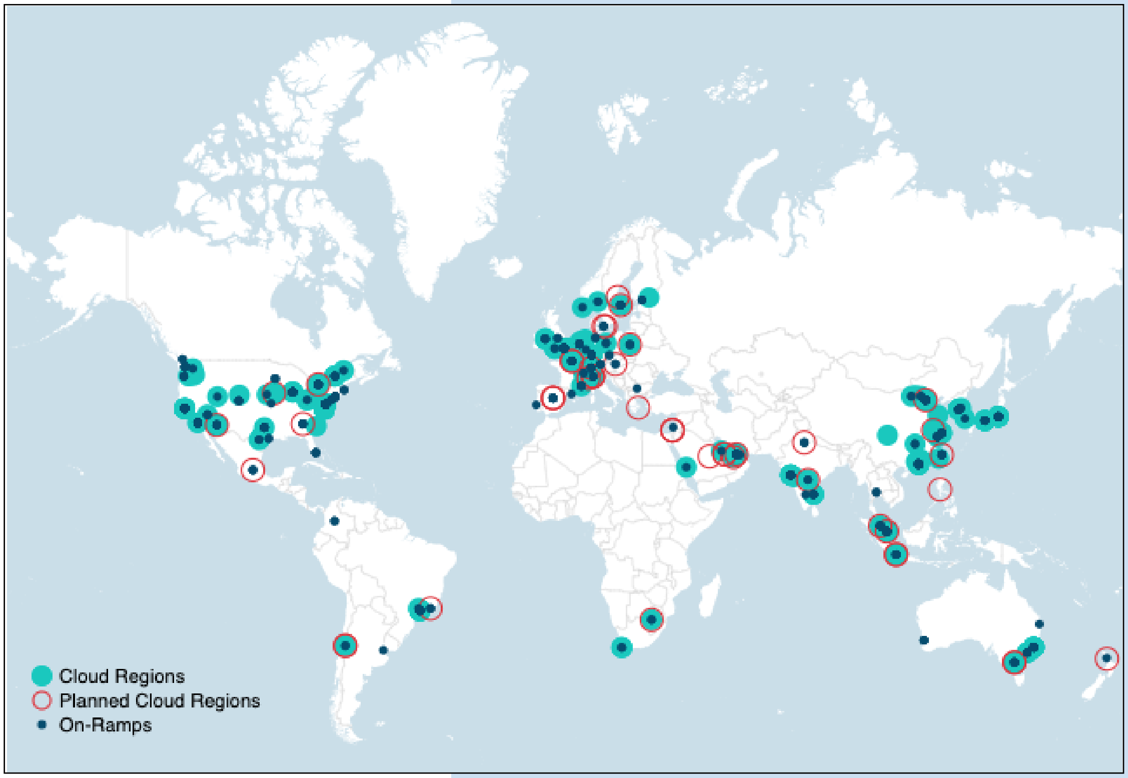

Notes: Data only include IaaS cloud providers data center information from Alibaba, AWS, Google Cloud, IBM, Microsoft Azure, and Oracle Cloud. Data as of Q1 2021.

VOICE

One Big Dropped Call

Still falling, after all of these years. 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, as the slump in voice traffic has turned into a rout. Carriers' traffic declined by 9% in 2017 and 4% in 2018 and a further 6% in 2019. 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 7% in 2020, slightly faster than the two previous years.

Impact of COVID-19

So, was there any effect of COVID-19 on the international voice market? The short, big picture answer is: "No." A more detailed analysis suggests that the answer is actually: "Yes, but not that much." We got an early glimpse of the impact of the pandemic late in 2020 when we surveyed a number of international operators on the market's reaction to COVID. Slightly more than half of the operators reported that they had seen a jump in international call volumes as the pandemic tightened its grip in March (but nearly a third still saw a drop in traffic compared to the year before). The bump in traffic was short-lived, however. Only one carrier (of 24) reported that traffic levels remained elevated by the end of the year. And when the final numbers came in for 2020, global traffic had continued to drop at an even faster rate than in 2019. But, who knows? Maybe it would have been an even bigger drop without the early COVID-effect surge.

One specific area of the voice market where we believe the pandemic may have had an impact is the mobile-originated share of international traffic. The mobile-originated share of traffic dropped for the first time ever in 2020, albeit by a very small amount. (From 62.4% in 2019 to 62.0% in 2020.) The upward trend seemed irreversible ever since mobile subs passed fixed subs in 2002, and then mobile-originated traffic surpassed fixed-originated traffic in 2013. Still, next year we anticipate that the upward trajectory of mobile-originated traffic share will return.

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

It's hard to believe 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.

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 (with surprisingly not much growth from 2020). 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 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 1.4 trillion minutes in 2021, dwarfing the 375 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 285 billion minutes of traffic in 2020, down 7% from 2019. Turns out, even though wholesale traffic declined in 2020, over the last ten years wholesale traffic grew at a compounded annual rate of 1% while overall traffic dropped by 2% per annum. Wholesale carriers terminated more than two-thirds (70%) of international traffic in 2020.

Traffic to mobile phones in emerging markets has spurred expansion of the wholesale market, and that demand continues to drive wholesale's growth. In 2020, wholesale carriers terminated over 85% of traffic to Sub-Saharan Africa and South America. In contrast, wholesale carriers terminated only 54% of traffic to western Europe.

Wholesale revenues have changed only marginally from ten years ago. Wholesale ILD generated $16.1 billion in 2019 and $14.5 billion in 2020. But let's take a moment to look under the hood. The geographic sources of this revenue have changed substantially.

Revenues on calls to Europe grew 61% between 2013 and 2020. That's largely due to traffic growth to European mobile phones. The U.S. & Canada regions provides a mirror image. Revenues on calls to U.S. & Canada fell by 61%.

Over the past decade, traffic to mobile phones in emerging markets has driven international wholesale market growth. Revenues from calls to mobiles in emerging markets noticeably increased from $8.1 billion in 2010 to $11.3 billion in 2020. As a portion of overall wholesale carrier revenues, calls to advanced economies shrank, as did revenues from calls to fixed lines in emerging markets.

Wholesale revenues are bolstered by a select set of low-traffic routes with stubbornly high prices. For example, the France to Tunisia route accounts for just 0.3% of international traffic, but, at $0.51 per minute, it provides

Global Retail Revenues from International Calls

USD billions

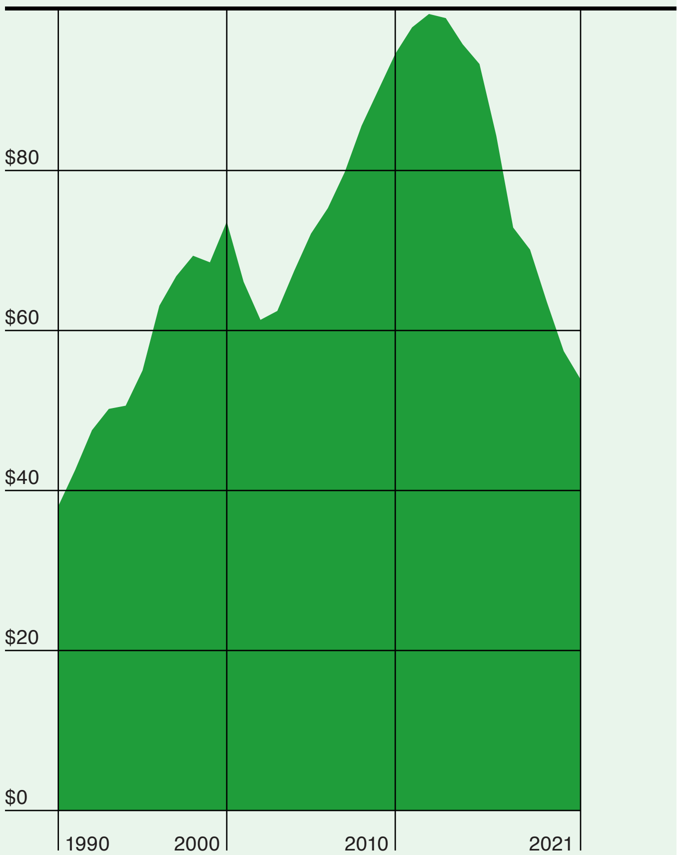

Notes: Data measure retail revenues on outgoing international calls; totals do not include revenue from wholesale services or incoming international traffic termination. Data for 2021 are projections.

3.6% of all revenues Thanks to low termination prices in Mexico, the U.S.-Mexico route serves as a converse example: that massive route represents 7% of all international traffic in the world, but only 0.2% of wholesale carrier revenues.

Who's carrying all this traffic? When we compare top international carriers, we note that the top 10 operators carried *over half* of all global traffic in 2020. That's about 220 billion minutes. Among the eight largest carriers in the world, only one terminated more traffic in 2020 than in 2019—and just barely.

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 2021, ILD accounts for only 3.7% 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.

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

Business Broadband Pricing Data
An extensive database of broadband service providers, plans, and prices.

Cloud and WAN Infrastructure
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.

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

Dedicated Internet Access Pricing Data
TeleGeography's database of dedicated internet access price benchmarks for corporate and retail customers

Ethernet Over MPLS Pricing Data
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 Pricing Data
In this module, we track long-haul city-to-city routes between major global business centers.

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

Global Bandwidth Forecast Service
Detailed forecasts of international bandwidth supply, demand, prices, and revenues, updated quarterly.

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

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

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

GlobalComms Forecast Service
Wireless, broadband, and wireline market metrics and forecasts by country and region.

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

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

IP Transit Pricing Data
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 Pricing Data
A database of global local access prices, reflecting actual transaction prices paid by carriers for leased private lines and Ethernet circuits.

MPLS VPN Pricing Data

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

SD-WAN Research Service

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

TDM Pricing Data

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.

TeleGeography Report and Database

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

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.

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.

Wavelengths Pricing Data

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