Executive Summary

Executive Summary

When TeleGeography released the first edition of *International Bandwidth* in April 2000, many investors regarded bandwidth as the new crude oil—the raw material of the new economy and a source of future riches. Since the crash of tech stocks, however, venture capital and bank funding for grand new network schemes have dried up. Bankruptcies and consolidations have replaced IPOs as the staple of financial headlines.

Does this melt-down make any real difference to bandwidth buyers? Yes and no. The uncertainty surrounding the sector makes it doubly important that buyers—especially those buying long-term—conduct a rudimentary health-check on suppliers. Also, comparisons to the energy industry notwithstanding, bandwidth is no commodity. Prices vary greatly across regions. Disparate terms and hidden costs keep bandwidth contracts far from standard and negotiations drawn out.

Despite the financial busts, the same forces that drove the industry in 1999 and 2000 continue to propel it in 2001. In particular, the combined effects of dazzling technical advances and much stronger competition continue to push prices down fast—at 50 percent or more each year on many key routes. New technology such as optical switching could disrupt existing product sets, just as it did when dense wavelength division multiplexing (DWDM) emerged. Finally, despite the current gloom, bandwidth suppliers continue to deploy networks.

Decisions about what to buy and how to buy it depend, as they always have, on a detailed understanding of market dynamics—in supply, demand, product sets, and technologies. This report provides a road map for buyers seeking to answer the hardest question of all: what happens next?

Fiber Optic Networks: Markets

Analysis of markets for fiber optic network bandwidth suggests that the investment boom has by no means played itself out. Between 1998 and 2000, new submarine cables increased trans-Pacific bandwidth from 14 Gbps to 244 Gbps. That achievement pales in comparison with what is yet to come. Once fully upgraded, cables now under construction could add an astounding 17 terabits per second (17,000 Gbps) to the trans-Pacific route. Terrestrial networks boast even greater upgrade capabilities. A handful of pan-European networks hold the potential to carry 1 petabit per second (1,000,000 Gbps) of traffic between cities—equivalent to 2 Mbps for every person in Europe. Because operators can upgrade these new systems gradually, the effects from the huge capital infusions during the late 1990s will not be fully felt for many years to come.

A booming stock market enabled financing for multiple network build-outs in the late 1990s. However, it was technology that fueled the greatest increases. Whereas early fiber optic systems used a single wavelength to transport information, modern networks employ Dense Wavelength Division (DWDM) to make use of many different light frequencies within a cable. Such advances mean that, with all wavelengths lit, bandwidth on a single network currently under construction could dwarf the combined existing capacity of all the other cables on the same route (see Figure 1. Giant Steps: Upgradable Capacity in Perspective).

Despite the well-publicized problems of some builders, many are still adequately funded. Global Crossing, Interoute, Level 3, Telia, and TyCom all appear to have ample resources to continue their builds, although more conservative financial management may slow the pace of deployment.

Compared to the astounding increases in lit and potential bandwidth, the actual market structure of the industry has remained relatively stable in the twelve months following the release of *International Bandwidth 2000*. No major new supplier has emerged since TyCom's entry into the bandwidth provisioning market in early 2000. However, many capacity providers are moving into related markets. Companies that initially supplied only wholesale bandwidth have added new services such as colocation, IP transit, switched minutes, last-mile facilities, and content distribution.

The massive upsurge of bandwidth in many parts of world has sparked extraordinary price declines; it also has engendered a bewildering array of new products for buyers. The Fiber Optic Networks chapter of this report sorts out these market options. First, an analysis of technical advances and regulatory changes examines the factors transforming the industry. Next, bandwidth products at different levels of the network value chain—from dark fiber to fully provisioned capacity—are compared. Finally, the chapter compares the costs and benefits of the financial arrangements (direct ownership, lifetime IRUs, short and long leases) available for acquiring bandwidth.

Optical Revolutions

Dense wavelength division multiplexing, the key development in fiber optic technology of the 1990s, will continue to wring more bandwidth from a single pair of optical fibers over the next several years. Other emerging technologies continue to widen the boundaries of the possible. In 2001, some suppliers will deploy networks that exploit a new optical "window"

Figure 1. Giant Steps: Upgradable Capacity in Perspective

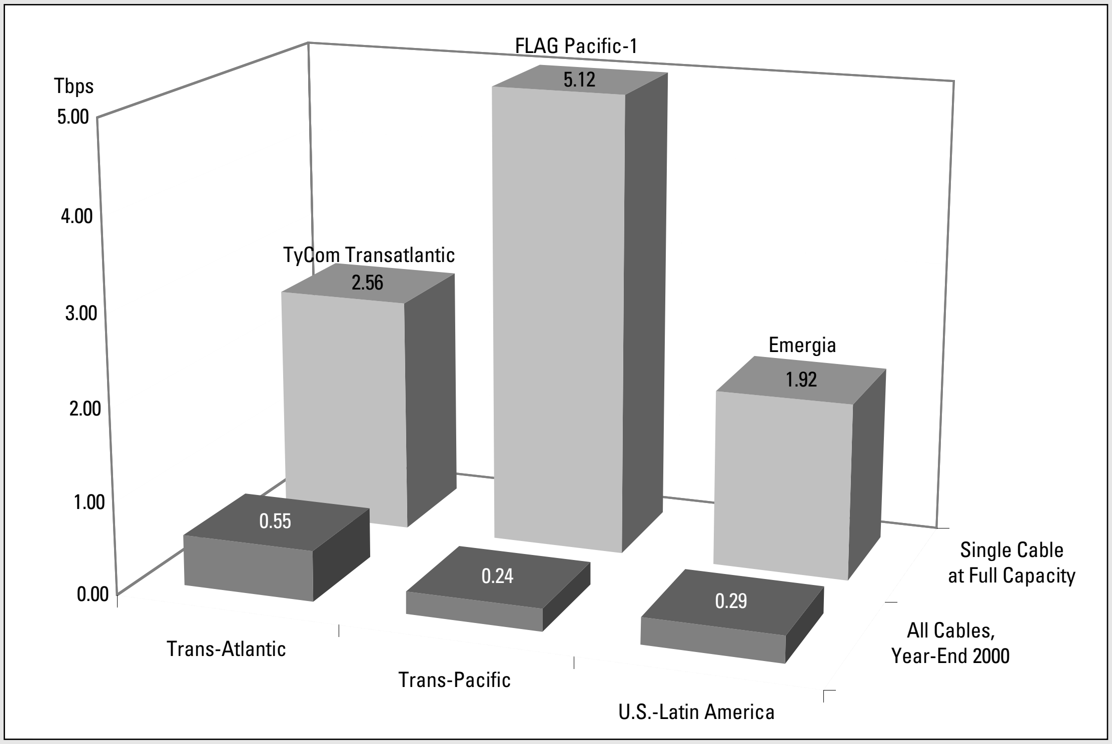

Source: TeleGeography research

at 1600 nanometers, the so-called L band. By 2002, the first 40 Gbps wavelength system should be in place. Meanwhile, improvements in both laser and fiber technology will enable designers to squeeze more signals into each window. Current networks space wavelengths 100 or 200 GHz apart; those to be deployed this year and next will space wavelengths 50 or even 25 GHz apart. Ultimately, these new technologies will allow operators to transmit up to 320 wavelengths on a single fiber pair.

Just as DWDM revolutionized fiber optics in the 1990s, optical switching could dominate the new decade. Switching tasks cannot easily be accomplished in the optical domain because light signals and photons are much more difficult to manipulate than electrical signals and electrons. Hence, at every point in a network where signals must be separated and routed discretely—that is, at every single node—the light signal must be converted into an electrical signal, processed, and then converted back into an optical signal. This continual opto-electronic conversion comes at a price. Each nodal point requires a complete set of multiplexing and demultiplexing equipment, which is very expensive—once purchased, it must be housed, maintained, managed, and powered. Moreover, traditional technologies were not designed to handle the very rapid increases in bandwidth or the continual reconfiguration of networks brought on by the Internet. Every time a new wavelength is required, skilled network designers must study network configuration maps, assign paths, and send out technicians to manually set up connections at every site. A shift to all-optical networks could introduce vast savings in both time and money; it also could enable new bandwidth products featuring rapid provisioning and network flexibility—a far cry from the often cumbersome arrangements still in place.

In addition to exploring the possibilities of state-of-the-art fiber optic technology, this year's report offers a more complete picture of the underlying cost of existing networks. While construction costs dominate initial expenses, especially in terrestrial networks, the high cost of DWDM equipment accounts for a larger proportion of ongoing costs. Optimistically, one can expect the costs of electronic equipment to fall with improvements in the price-performance ratio of electronic componentry. DWDM and related technologies likely will continue to reduce the costs of bandwidth for at least the next two to three years. The main technological challenges will emerge from long-standing difficulties in applying optical techniques to network management, network architecture, switching, and provisioning.

Prices in Fiber Optic Networks

Several factors are conspiring to drive down prices. First, the one-off impact of competition has pulled formerly inflated prices toward cost in Northern and Western Europe, the U.S., and on trans-Atlantic routes (see Figure 2. Transcoceanic Annual Lease Rates, 1998-2001). The impact of new competition will be felt in other regions over the next year or two. Second, very large annual increases in supply precipitate a decline in underlying cost per megabit. Third, corporate financial distress—and the corresponding pressure on bandwidth sales managers to generate immediate cash flow—has resulted in "fire sales" that have driven prices on some routes to cost or even below.

These three effects have combined to reduce prices at a record pace, especially at capacities of 155 Mbps and above. Ultimately, certain fixed underlying costs dictate a "floor" price for any particular circuit speed on any particular route—the so-called "cost plus" model. Differences across regions, products, and stages of development help to define the price trajectory. Either way, the price per megabit will continue to fall rapidly for the foreseeable future as users trade up to higher capacities and sellers pass on savings from technology and economies of scale. Capacity in public telecommunications networks typically increases in fourfold increments (e.g., from 155 Mbps to 622 Mbps, and from 622 Mbps to 2.5 Gbps), while prices tend to only double between these capacities.

Pricing, however, is not often straightforward; the real-world circumstances of buyers complicate the negotiation process. Large buyers often purchase networks, not specific routes. Some costs, such as operations and maintenance (O&M) costs for submarine capacity, are hidden or may not be fully appreciated. Equipment or break-out points may also be charged. Buyers must account for last mile connections and colocation. However complex, accurately projecting the decline in pricing is absolutely vital, especially on long contracts which can leave buyers holding an expensive, wasted asset.

Satellite Systems

Unlike developments in fiber optic networks, satellite technology has advanced in evolutionary steps rather than in revolutionary leaps. There was no satellite equivalent of DWDM;

Figure 2. Transoceanic Annual Lease Rates, 1998-2001

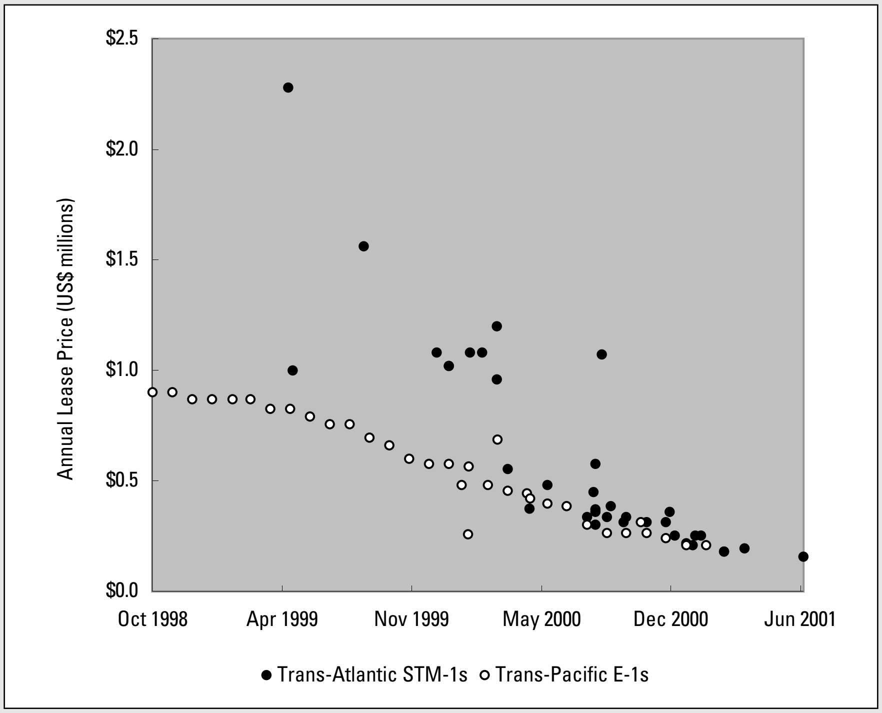

Note: Trans-Atlantic prices reflect New York-London STM-1 (155 Mbps) lease rates; trans-Pacific prices reflect Los Angeles-Tokyo E-1 (2 Mbps) lease rates.

the combined throughput of all satellites worldwide is now less than that of the TAT-14 undersea cable.

The role of satellites has evolved to meet the changing circumstances of the market. In wealthier countries, satellites are primarily used for broadcasting. Satellites still provide point-to-point services in areas where there is no fiber or where the cost of fiber is prohibitively expensive. Nevertheless, in recent years it has become clear that on many point-to-point routes, the satellite era is over.

The chapter on satellite systems examines the roles played by communications satellites in the international bandwidth market. It offers an overview of satellite communications technology, followed by a discussion of the primary applications for which satellites are used. The chapter also provides analysis of supply, looking first at satellite system operators, and then at the total available amount of satellite bandwidth. The analysis concludes with an overview of satellite capacity pricing trends.

Network Ends

A straw poll on recent long-haul network bids suggests that last mile connections add about 30 percent to the cost of an end-to-end connection. Fortunately, some cities—like London, Paris, and New York—have more than five metropolitan area networks (MANs) to choose from. Many others, especially outside the U.S., have only the incumbent carrier to look to for the “local tail” connection. Even in cities with new metro network builds, however, achieving a similar level of competition to that of the long distance environment is a long-term challenge. But any growth in metro networking is more than welcome for purchasers, who continue to be plagued by delays and high prices in the last mile.

The options for hosting your equipment—whether it be telephone switches or storage arrays—are also multiplying. But the exuberant atmosphere that once surrounded colocation is quickly evaporating. Many ambitious build-out plans have been scaled back, or even abandoned. The colocation business is at an uncertain crossroads, and the Network Ends section of this report will help sort out the options and avoid the pitfalls of network interconnection and local delivery.

The Demand Paradox

The new market pessimism has forced bandwidth providers to return to a question that has haunted the industry for several years: is there, or will there be, a capacity glut? Resolving this quandary requires an accurate gauge of bandwidth demand. Most efforts have focused on bottom-up approaches to modeling demand: gather a wide range of statistics on end-user traffic (e.g., number of users, peak user flows, bandwidth per application), then run them through a set of assumptions to see how much bandwidth to provision. To test the accuracy of this approach, the results of the predictive models are measured against TeleGeography’s historical inventory of actual network deployment. A comparison suggests that most models chronically overestimate the capacity actually deployed in networks. In particular, Internet bandwidth has historically been far smaller than most models have predicted.

Even more troubling—at least at first glance—is a comparison of projected network deployment against the future supply of lit, available bandwidth. Not a single model predicts that bandwidth provisioned in networks will amount to even half of the lit capacity across the Atlantic and Pacific (see Figure 3. Projected Network Deployment versus Supply, 1998-2005). This huge gap provides a clue that projections based on network provisioning requirements tell only half the story; purchased but unused capacity also figures strongly in the bandwidth demand equation. Service providers tend to purchase excess capacity for a range of reasons, including network redundancy and long provisioning times. This excess buying behavior is bad news for models that focus strictly on network traffic requirements to predict future bandwidth demand. For investors of companies that depend on bandwidth sales, though, it is a most welcome phenomenon.

Figure 3. Projected Network Deployment versus Supply, 1998-2005

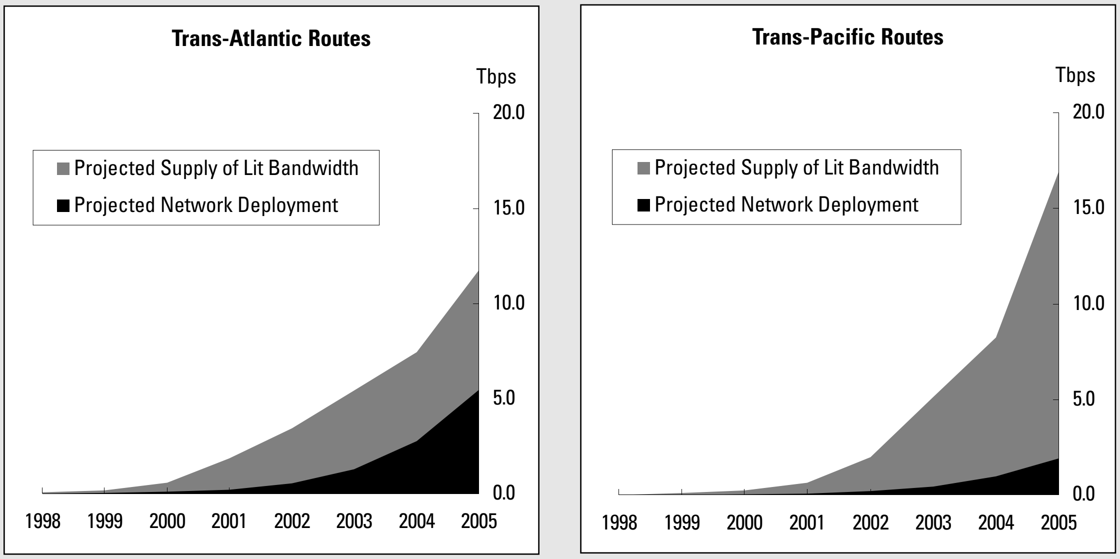

Note: Network utilization data reflect projections from TeleGeography medium growth scenario outlined on page 138. Projected supply is based on reported initial capacity of systems during the year they are scheduled to be ready for service and assumes that cables will reach their potential capacities in five years through upgrades of equal increments in each year. Projected supply figures include only cable systems announced by early 2001, and thus may underestimate future supply of available bandwidth.

Source: TeleGeography research

Buying & Trading Bandwidth

The supply of bandwidth continues to increase at exponential rates. Prices on most routes are plummeting. And, although bandwidth demand is more robust than many models predict, such slack exists on some routes that many carriers and ISPs are acquiring bandwidth far in excess of current needs. Yet, bandwidth markets are not a buyer's paradise. Perils and pitfalls abound: hidden expenses, contractual snafus, and service failures can prove expensive. The chapter on buying bandwidth focuses on strategies for negotiating the lowest price possible, avoiding contractual and regulatory entanglements, and guaranteeing quality and availability.

The emergence of market intermediaries seeking to simplify and automate the sale and purchase bandwidth could prove a valuable tool for bandwidth buyers. The biggest benefit offered by bandwidth exchanges is greater pricing transparency, which helps buyers to identify the "going rate" for a specific type of circuit on a particular route. Service offerings can vary greatly among the different types of bandwidth exchanges and brokers, however. The chapter on buying and trading bandwidth concludes with a comparison of the products, operational models, and fee structures of these new market intermediaries.

Network Profiles

The first half of *International Bandwidth 2001* focuses on the supply and demand trends shaping the industry. The last 225 pages in the report provide detailed information on the actual networks. Profiles are presented by network type:

1. Submarine cable systems: profiles of 69 international cables ready for service by January 2002 with a capacity of at least ten Gbps.

2. Terrestrial networks—Americas: profiles of 35 fiber optic networks in the U.S., Canada, Mexico, and South America offering major bandwidth (OC-3 and above) sales to carriers and ISPs.

3. Terrestrial networks—Europe: profiles of 27 pan-European networks specializing in major bandwidth (STM-1 and above) offerings to purchasers.

4. Satellite systems: profiles of 22 satellite operators that collectively account for more than 80 percent of worldwide transponder capacity.

Each network profile includes contact information, network details, service offerings, and a map of connectivity.

Conclusion

Even as people continue to drive cars and burn oil in a recession, so do they continue to make phone calls, send email and search Web pages, all the while consuming bandwidth. Despite the current gloom, bandwidth supply and demand will continue to expand.

Simultaneous build-out of multiple new networks will translate to over-supply of capacity on some routes for many years to come. The need to realize revenue during a bear market puts sellers under further pressure. As a result, international bandwidth will remain a buyer's market for some time—especially for those buying in volume.

Still, even buyers' markets present risk. Purchasers must take great care when contracting with suppliers in an industry suffering significant instability. Also, with transmission and switching technologies leaping forward and costs falling fast, the possibility of investing in soon-to-be overpriced bandwidth remains a risk. More than ever, this is a market in which the principle *caveat emptor* applies.

