Matching Bandwidth to the Requirements of the Internet

By Peter Sevcik and Rebecca Wetzel
July 2001

Until the wakeup call of early 2001, prevailing wisdom was to meet insatiable bandwidth demand with indiscriminate bandwidth provisioning. Sadly, many carriers accelerated their demise by following this call. Others though bruised, survive and are attempting to apply painful lessons as they plan for the future. For carriers to chart a successful course, they need to prepare for the as-yet undefined customer needs that will materialize during the next five years. A carrier must navigate through a wide range of possible traffic scenarios and architecture choices, all the while insuring that it remains competitive.

Although building too much bandwidth capacity was an expensive mistake in the recent past, it may well be the best strategy for the future. Internet demand grew largely based on the growth of the browser-enabled user population. Although this continues to be a large source of traffic growth, it has now become predictable. The future will be much less predictable as completely new and different traffic profiles hit the Internet.

To absorb the sudden impacts of new application classes, Internet core service providers should accumulate bandwidth pools starting at about 40 Gbps. Such capacity buffers, however, must cost much less than the previous build-out. In order to manage costs and realize the best return on their investment, service providers must make the highest use of their fiber spectrum through spectral efficiency, and take advantage of every possible savings to provision and manage bandwidth to customers.

A lynchpin to success will be to grow bandwidth flexibly and sensibly rather than inflexibly and indiscriminately as in the past. With unpredictability the only constant, savvy carriers will make flexibility, cost effectiveness and customer responsiveness competitive tools. Successful carriers will prepare to meet the demands of multiple scenarios, and not lock themselves into solutions that satisfy only one or two outcomes. This report presents four possible scenarios and explores the traffic implications of each.

Alternative Growth Scenarios

Depending on whom you ask, different sets of forces will reign supreme in shaping demand for Internet backbone capacity. No one can yet say which of these competing forces will in fact dominate. Will content or connectivity be king? Will VPNs rule? Will new models for linking users with content emerge? As unsettling as it may be, it is vital for carriers to delve into alternative scenarios such as these and consider the dramatic impacts they will have on their network architectures.

Scenario A: Content is King

You've seen the headline "Content Is King," and many content owners and storage vendors are convinced that they will soon occupy the center of the network universe. The content is king scenario spawns a hub and spoke traffic pattern. At the hubs of the traffic flows are a limited number of Web sites with massive storage capabilities. The spokes reach out to millions of distributed users concentrated within major population centers.

In this model, although users have thousands of Web site destinations to choose from, most eyes are drawn to content at fewer than 50 sites. This is analogous to the availability of dozens or even hundreds of TV stations via cable or satellite, of which, only a handful attract the lion's share of viewers. The infrastructure needed to support

NetForecast Report 5054

this traffic flow consists of roughly 35 Web hosting facilities connected by massive bandwidth capacity to about 35 metropolitan areas.

In this scenario, traffic grows with the evolution of Web-based applications, and as applications and content become more compelling, users stay on the network longer. Over time, applications such as streaming media flourish and grow, and traffic becomes more broadcast in nature. As bandwidth becomes more plentiful and cheap, the cost to serve content from the source becomes roughly equivalent to the cost to serve from a local cache. Despite changes in traffic volume and bandwidth economics, however, the fundamental hub and spoke nature of the traffic flow remains consistent.

Scenario B: Connectivity is King

Connectivity is as much of a contender for the position of "king" as content. Both businesses and consumers have a longer history of paying much more for connectivity than for content. Connectivity-based traffic can grow from a variety of sources.

Broad adoption of peer-to-peer applications would automatically raise connectivity to prominence. Peer-to-peer applications are emerging to serve both business and consumer markets, and future peer-to-peer applications are likely to encompass such activities as business partners sharing product design information and consumers sharing files. Telephone companies and alternative service providers are adding traffic to this scenario by wholesale migration of telephone traffic to the Internet using voice-over-IP (VoIP) technology.

In the connectivity-is-king scenario, peer-to-peer usage starts out emulating traditional telephone traffic, with traffic distribution following population demographics. At the outset, 80 percent of traffic remains local, reflecting personal and business relationships strongly influenced by geography. Over time, however, as businesses and consumers take advantage of the network to transcend geographic limitations, the peer-to-peer model diverges from the telephony example, and the 80/20 ratio of local to long-haul traffic begins to change – ultimately reversing. Peer-to-peer traffic tends to be unpredictable, with traffic peaks and valleys reflecting changing business needs and consumer whims.

As it grows, long-haul peer-to-peer traffic continues to reflect population distribution. Major metropolitan areas such as New York, San Francisco and Dallas outstrip the likes of Topeka and Spokane in day-to-day traffic volume. This said, however, events can precipitate sudden traffic shifts to demographic backwaters. Traffic spikes to demographically out-of-the-way places are likely to seriously challenge performance where circuits are not sized to comfortably accommodate such loads.

With connectivity as king, traffic growth remains modest until a video version of Napster bursts upon the scene, inflicting 'peer-to-peer pressure' on the network through precipitous traffic growth within and between major population centers.

Scenario C: Applications Delivered via "Network Palettes"

A scenario just now coming into focus is the delivery of applications and content to businesses and consumers via "network palettes [1]." A network palette is an emerging application platform comprised of technologies that enable a fast, personalized and secure user experience. A network palette includes capabilities such as replication of applications, databases and content along with their acceleration and personalization.

A network palette employs thousands of loosely coupled machines, each operating a small part of an application or database. XML dominates this scene in which service-point machines interact to keep a palette of information fresh and properly synchronized across many locations. Data coherency among the service points varies as application usage and user behavior change.

A user's individual experience consists of a composite "painted" from data elements, or "colors" from one or more palettes. PCs and PDA devices are supported via wire or wireless links with persistent connections to the palettes. Mobile users are handed off to specialized palette edge service points similar to cellular telephone network operation. However, businesses maintain persistent connections to each other's servers day and night. Data is constantly refreshed among machines to be used by machines. Humans see or use a very small portion of the machine-to-machine traffic. Machines, not people, drive traffic, and the resulting traffic patterns are relatively consistent.

For example, an airline must order everything from toilet paper to engines parts. But rather than join a number of business exchanges, the airline connects to all of its suppliers via a network palette. From the network palette, the airline can easily locate any needed item, and continuously monitor price and availability. Specialized search tools are active on the airline's behalf around the clock gathering information on items of interest. The search tools maintain persistent connections to suppliers, and deduce short- and long-term patterns in pricing and availability, resulting in the most favorable terms with the least expenditure of effort.

This scenario produces dramatically different traffic demands, where trunks per network grow quickly and then level off once a limited number of network palettes are built. Bandwidth per trunk grows modestly and keeps growing as more data is added to the network palettes. The total bandwidth required is the highest in this scenario as torrents of machine-to-machine traffic are unleashed on the core.

Scenario D: VPNs Rule

As the Internet matures, it is likely to attract an enormous pool of traffic from private leased line and frame relay networks to virtual private networks (VPNs), which share common infrastructure. VPN technology supports private networks much more cost-effectively. In this scenario, frame relay service constitutes the first phase of a widespread migration from totally dedicated point-to-point lines to totally shared VPN infrastructure for intra-corporate and inter-corporate data traffic.

In the "VPNs Rule" scenario, businesses wish to take advantage of the lower cost of shared infrastructure, but do not trust their traffic to the Internet, turning instead to interconnect to their partners via secure tunnels on the Internet for their communications.

Many diverse locations serve as traffic end points, and traffic grows slowly over time. There is plenty of bandwidth headroom for areas with abundant and cheap transport, but little to none where transport is scarce and expensive. All in all, life for a carrier and its private line customers continues to look remarkably similar to the way it looks today from a bandwidth point of view. However, the bandwidth is now supplied over the more shared infrastructure of the Internet.

How the Scenarios Shape the Internet Core

The Internet core is made up of several carriers that provide services to ISPs and enterprises who in turn perform access, hosting, and/or edge services. These customer groups provide first and often multiple levels of traffic aggregation before they connect to the long-haul Internet core. The Internet core carriers then switch the aggregated traffic onto their very high capacity long distance trunks. Despite the traffic aggregation of the feeder markets, each of the four scenarios will have a dramatic impact on the number of such trunks required and the bandwidth of each trunk as shown in Figures 1 and 2.

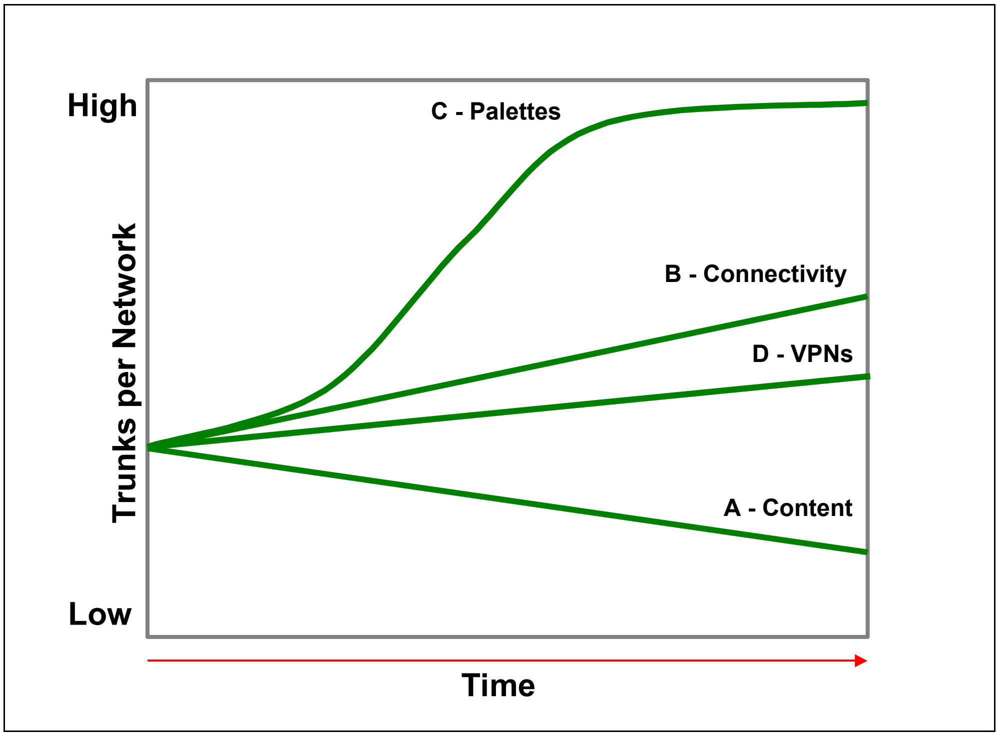

Figure 1 – Scenario Implications on Core Trunks

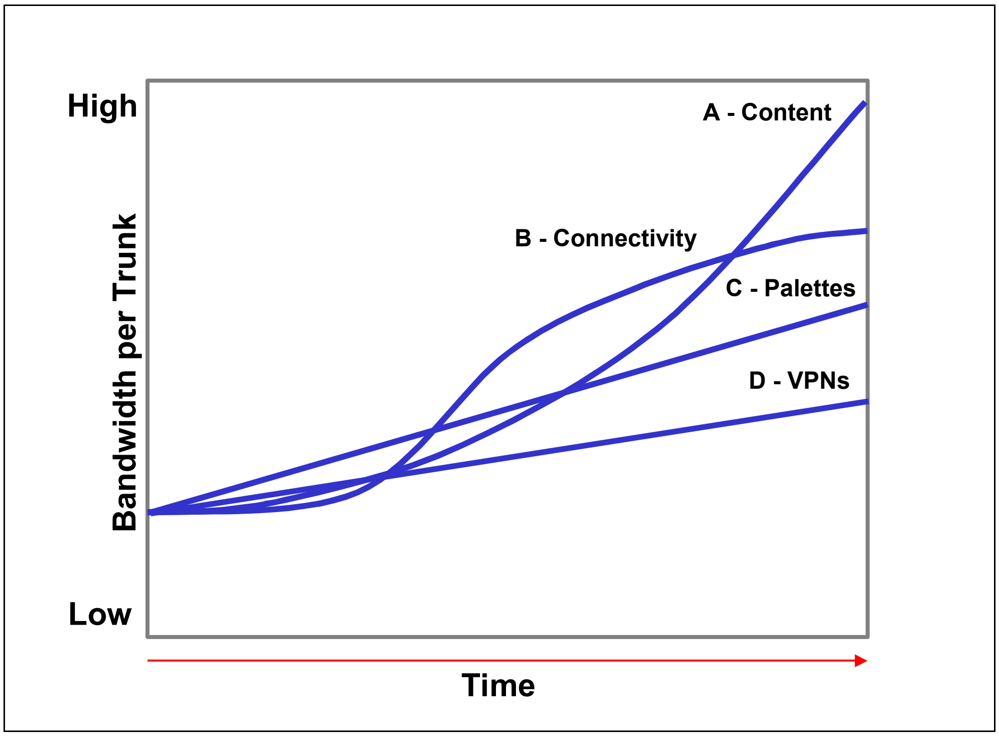

Figure 2 – Scenario Implications on Core Bandwidth

Figures 1 and 2 show vastly different networks will be needed depending on which scenario is in ascendancy at any time. Bandwidth may be drastically over or under provisioned by a carrier that does not react in time. Even the locations where the bandwidth terminates will change. Understanding and reacting to the fundamental market drivers of bandwidth will be a clear competitive advantage.

Alternative Bandwidth Delivery Architectures

Some say that none of the above scenarios pose problems because there is such a fiber glut that no conceivable scenario will dent it for years. Granted, there is a temporary Internet core bandwidth glut [2]. But contrary to the current conventional wisdom of Wall Street analysts, there is no fiber glut. Long-haul telecommunications conduit and fiber are untapped bandwidth reserves. No one has ever called the oil reserves of the Middle East or Alaska a fuel glut. In both cases a tremendous amount of work and investment must be made to transform the reserves into a useful commodity. This work takes time, and cannot be performed fast enough to accommodate a sudden change in demand for the finished product.

Fiber is transformed into a useful commodity by lighting it with wave division multiplexing (WDM) technology and adding switching equipment to create bandwidth a user can call upon as needed. Recent dramatic changes in transmission efficiencies have spawned alternative views of how bandwidth will be provisioned and delivered.

The Future Role of Bandwidth

One view is that bandwidth is becoming such an inexpensive commodity that it can be wasted in order to save on some other more expensive part of the architecture. The proponents of this view use the analogy of the computer chip. Once computers were expensive and had to be shared among users. Today they are inexpensive and are dedicated to each user. Computer innovation, fueled by Moore's Law, is putting a chip in almost every man-made item. The proponents of this analogy conclude that since bandwidth is becoming inexpensive, it is inevitable that networks will supply dedicated bandwidth to each user and some day to each man-made item.

This analogy is, however, flawed. The microchip ushered in an age in which computing bandwidth (mega instructions per second, MIPS) became irrelevant, and innovation focused on ubiquity. Everyone owns many chips, most of which are hidden in products. No one cares how many MIPS of computing bandwidth each chip has.

The appropriate analogy between the chip and network revolutions is that innovation will give users and then "things" ubiquitous access to the Internet, and bandwidth will not matter. Users will no longer care or even know the data rate at which they are connected, provided the access rate is fast enough. They will only care *that* they are connected. Therefore, the appropriate analogy between chip and network technologies applies to access not bandwidth.

Shared vs. Dedicated Bandwidth (aka Packets vs. Circuits)

The misapplication of the chip analogy also leads to the false conclusion that users will have an inexpensive dedicated connection between themselves and the destination with which they are communicating. Such a connection needs a fixed modest bandwidth sufficient to perform the task of the connection. Since clearly there will be millions of connections, networks should be built with many circuits that can be switched into place as needed. Circuit switches establish dedicated bandwidth while packet switches manage shared bandwidth.

The circuit switching (connections) versus packet switching (datagrams) debate has raged many times in past years [3]. Circuit switching proponents have proposed the telephone switch, X.25 switching and ATM switching. Packet switching proponents solidly converged on the Internet Protocol (IP) as the de facto international standard of packet switching. In the mean time the old telephone switch network is migrating to an IP infrastructure and X.25 and ATM are falling by the wayside.

The fact is that users and devices will connect to the Internet with packets not circuits, obviating the need to provide millions of circuits to each user connection. It is already too late and impractical to provide a dedicated connection to a user session. For example, a typical Web page loads by providing information from about a dozen sources. Packet switching provides access to the many sources without the overhead of a connection setup. Each source supplies its data within fractions of a second using a shared bandwidth pool.

The only other reason to build a long-haul network with many small circuits is to support the enterprise private line market. However, that market is already transitioning to the shared model of frame relay followed by the shared VPN model. The private line market will always migrate to the lowest cost approach. Analysis [4] shows a progressive acceleration of cost declines by SONET, then ATM and finally Ethernet technology networks. Ethernet is a clear long-term low-cost winner as a transmission service. Ethernet wins this battle precisely because it is packet-based and takes advantage of shared bandwidth pools.

Adapting to Unpredictability

Rather than a single scenario establishing dominance, the most likely outcome is a kaleidoscope composed of all of the above. To complicate matters, the relative supremacy of each scenario is likely to change over time. In the short term, content may be king, with network palettes and VPNs ramping up slowly. Peer-to-peer applications are likely to pop up sporadically as a wild card, possibly to fade into oblivion.

This unpredictability makes planning a daunting task for any carrier. In the old days, it was possible to optimize and tune a network for a single purpose. But going forward, those carriers able to successfully accommodate a mix of scenarios will remain standing, while those pinioned by inflexible networks will join a formidable list of casualties.

Most of today’s networks are designed around the premise that the future will require bigger versions of the networks of the past. Unfortunately, yesterday’s networks are not flexible enough to respond to a mix of the scenarios described above. Bigger may sometimes be better, but not when you’re building the Titanic. You need a network able to survive an “iceberg” scenario.

Put simply, carriers must adapt to survive. But what does it mean to adapt? Among other things, it means that carriers must be able to:

- Maintain a strategic over-build of bandwidth
- Implement high spectral efficiency
- Keep system costs as low as possible
- Quickly change capacity deployment to meet shifts in demand
- Keep core backbone speeds ahead of the feeds from the network edge

- Incorporate fewer but more versatile boxes
- Sell flexible capacity not static wavelengths

Although no mean feat, all of this is possible. However, these goals must be viewed in the context of a realistic future vision of the Internet.

Maintaining High Spectral Efficiency

Spectral efficiency is defined as the total data bandwidth divided by the available transmission bandwidth. A fixed guard band of unused bandwidth must isolate each trunk or wavelength within a WDM spectrum. The approach of supplying many small trunks creates the need to maintain many guard bands of unused bandwidth between each trunk (low efficiency) while the approach of a few large trunks reduces the number of guard bands of unused bandwidth (high efficiency). Providing many small bandwidth circuits in the pursuit of a circuit-centric strategy will clearly lead to low spectral efficiency.

Users will connect by packets to sharing a metropolitan bandwidth pool and then be aggregated for connectivity into the long-haul carriers. Therefore, long-haul carriers will continue to receive traffic in ever increasing aggregations of users and bandwidth load. These large access pipes connecting to the long-haul networks must in their own right be trunked onto yet larger long-distance trunks.

It is these trunks that are described in Figures 1 and 2 above. Each of these trunks is sufficiently large to justify wavelengths in the tens of Gbps, starting at 40 Gbps and going faster as edge rates increase. The future of long-haul carrier flexibility lies in its ability to create, deploy and re-deploy such high capacity trunks out of the few strands of fiber that will be lit each year.

For the purposes of traffic engineering these trunks may in fact use some form of virtual circuit for assigning a wavelength to a traffic bundle between major geographic regions. The emerging method for this form of engineering is MPLS and its Lambda derivatives. These very high capacity virtual trunks are again taking bandwidth out of a shared pool of total bandwidth available in the lit fiber plant of a carrier.

High spectral efficiency systems will allow carriers to economically build large pools of bandwidth which can be filled by packet switches. Low spectral efficiency systems will put a carrier at a significant economic disadvantage. Low spectral efficiency will force the carrier to mine more of the fiber reserve in order to provide an equivalent amount of bandwidth in use. It is only bandwidth in use that generates revenue.

Again, just like the more competitive oil company can get more gallons of gasoline and other products from each barrel of oil, so must the carriers get as much total bandwidth in use out of a fiber strand. The only way to get high spectral efficiency is to deploy large bandwidth pools that the packet switches can fill. Dedicating bandwidth to any single user or even a single access network would be extremely inefficient. Carriers do not need thousands of unsold small pipes.

Keeping Capacity Ahead of Shifts in Demand

Carriers must have the flexibility to put the right amount of capacity where needed, when needed to accommodate unexpected traffic spikes and demand shifts due to changing usage and growth scenarios. The primary customer of the long-haul carriers will be companies and metropolitan carriers that aggregate traffic. The size of the customer circuit will continue to be larger. The bursty nature of the packet traffic on these large access circuits will magnify the shifting demands of the scenarios described above.

Each trunk is the fundamental pool of bandwidth for which access users are contending. The average trunk must therefore operate at speeds significantly higher than the highest access line rates. With today's edge technology interfacing to the core at rates as high as 10 Gbps, carriers must have trunk speeds at 40 Gbps as a starting point in order to keep up with the growth of access bandwidth.

Failure to do this could cause network performance to degrade catastrophically and without warning, and/or catch carriers without sufficient capacity to provide existing customers with bandwidth upgrades and provision new customers. Not only will this stifle growth, it will prompt customers to switch service providers.

Providing Capacity, Not Wavelengths

Customers want to buy capacity from carriers – and they want the flexibility to buy that capacity in the increments they need. Unfortunately legacy optical systems lock carriers into selling fixed wavelengths – not flexible capacity – because in the past a carrier could sell only an entire wavelength, not a portion of it. The good news for carriers is that this limitation is being lifted. Advanced optical technologies will allow carriers to sell customers capacity in the true bandwidth increments they want, not by the wavelength. Carriers need to move in this direction in order to adapt to customers' needs.

The counter approach of selling many small wavelengths not only puts the carrier at an economic disadvantage due to poor spectral efficiency but it also puts the carrier at a market disadvantage since it is selling the wrong service relative to market demand.

Summary

To stay afloat during these evolutionary times, a carrier needs more than just a massive fiber network. An impressive network that cannot adapt to the changing demands of customers and their networked applications will sink as surely as the Titanic. Carriers need to ponder a variety of future scenarios and build a network that will adapt to a range of outcomes to become unsinkable. Recent advances in optical network technology should be applied to the Internet now in order to prepare for a variety of future bandwidth needs. The most important attribute of such a nimble carrier will be its ability to operate at a high level of spectral efficiency. Carriers must not follow the false promise of equipment vendors that promote many dedicated circuits at the cost of low spectral efficiency and the disastrous consequence of being out of step with market needs.
¤

References

1. Sevcik, "The New Platform – The Network Palette," Business Communications Review, May 2001.
2. Sevcik, "Internet Bandwidth Supply and Demand – It's Time for Accountability," Business Communications Review, January 2001.
3. Sevcik, "Why Circuit Switching is Doomed," Business Communications Review, September 1997.
4. Sevcik, "A Contender With Staying Power," Business Communications Review, September 2000.