Why Moving to a Pure Optical Network Makes Sense … Today

Abstract

This paper provides a basic understanding of communications networks and an overview of today's core networks. It also highlights the issues preventing existing fiber networks from taking full advantage of their capacity. Finally, it outlines the advantages of a pure optical network.

Introduction

Internet traffic is the driving force behind the phenomenal growth rate of communication networks. Statistics show that Internet traffic in North America has reached 350,000 terabytes per month, the equivalent of over 9,500 years of continuous video, with continuing growth through the decade. One is left to wonder if bandwidth demand will ever be satisfied.

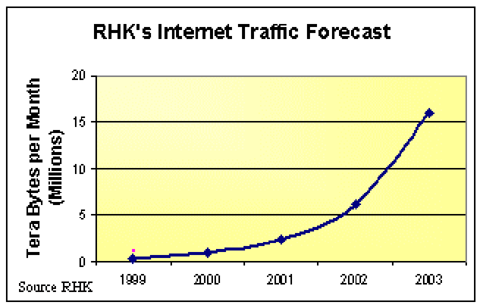

For network operators, customer demand for data-intensive services does not necessarily guarantee soaring revenue growth and profit margins. To get a substantial piece of the "Internet pie," operators must figure out ways to increase margins while meeting future market demands. The solution is clear: ensure that the underlying network infrastructure can keep pace with and capitalize on the opportunity growth. But when one considers that today's infrastructure was engineered for voice-based traffic, keeping pace with the explosion in data traffic may not be so simple after all.

Technology advances such as Dense Wavelength Division Multiplexing (DWDM) address fiber capacity issues, but networks are still experiencing crippling bottlenecks at the fiber termination points and engineers are plagued with a host of network management challenges. To solve these problems, network operators have to constantly add more equipment and replace existing network elements -- approaches that end up being costly and inefficient.

To fully realize the economic potential of the Internet, service providers need a new approach for optical transport, bandwidth management and networking -- a new type of network that can support data-centric IP traffic, handle bandwidth demand and scale to accommodate rapid expansion due to new and existing service growth. In short, carriers need to move to pure optical networks to resolve these issues.

A Closer Look at Today's Network

The design of the current public switched telephone network (PSTN) was based on consumer demand where 80 percent of all calls originated and terminated in the same local area and only 20 percent of the traffic traveled over long distance networks. Backbone networks did not get overtaxed, and the long distance portion of the network generated the lion's share of operators' profits. The Internet turns this business model on its head, with the majority of traffic occurring over long-haul networks at rock bottom rates of $19.00 per month for unlimited "dedicated" bandwidth. With voice services expected to grow a mere three percent, networks designed to carry voice must now try to adapt to the new world of communications in packets of data.

Designed for voice-based traffic and predictable traffic patterns, the SONET standard allowed telecom network providers to build interoperable optical networking elements to transport synchronous, asynchronous or non-hierarchical digital signals in SONET format over optical fiber. The introduction of SONET resulted in reliable, efficient and affordable bandwidth, and demand soon outpaced the original design intent of SONET's capacities and initial systems exhausted.

To cope, more and more SONET terminals were added to the network, creating stacked or layered ring configurations resulting in a proliferation of network elements. But how do you communicate between rings?

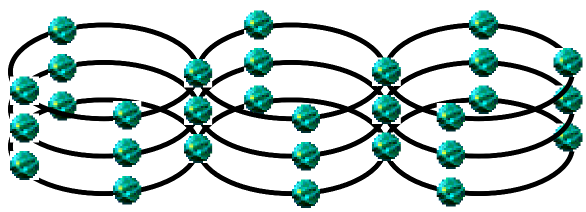

Figure 1: Stacked SONET Rings

To maneuver through multiple network layers, signals must be manipulated by Digital Cross Connects (DCS), ATM switches and Routers -- all elements that add latency and cost to signal transport. In effect, all optical signals must be dis-aggregated and re-aggregated and entirely processed at each network node, regardless of the traffic's destination. This adds cost complexity, bottlenecks and the potential for failure.

Functional for voice-based traffic and predictable traffic patterns, the underlying SONET-based architecture is its own worst enemy. It continuously manipulates signals as they pass through the network, converting optical streams to electrical pulses and then changing them back again to their optical form for routing to the next destination. It is this electrical regeneration of optical signals and multi-layer network structure that is at the core of the problem.

Finally, as it became apparent that SONET networks were suitable for voice and high priority data traffic, not low revenue best-effort Internet traffic, planners started designing dual networks -- one for Internet data and one for the rest of the traffic. For various and obvious reasons, maintaining two discrete networks does not make economic sense. At every turn, costs have skyrocketed in an attempt to make SONET adapt to a data-centric traffic environment. Industry analysts have estimated the cost of revamping the underlying PSTN infrastructure at $1 trillion over the next 20 years.

DWDM Solves One Problem - But Creates Others

Considered one of the most significant networking technology advancements, DWDM enabled operators to increase network capacity by an order of magnitude without the need for additional fiber. Increasing capacity on the same amount of fiber was viewed as the best way to save operators money and time.

However, by increasing the number of wavelengths per fiber, DWDM put an even greater strain on each network node. Now, each node must electronically process and regenerate many more wavelengths. So while DWDM solved the fiber capacity problem, it didn't change the nature of the bandwidth management problem -- and in fact, created a network management nightmare. By increasing capacity, DWDM inadvertently forced operators to add additional equipment for routing and switching signals and further exacerbated inherent problems with the SONET network.

Another Way to Describe the Problem

The growth in communications infrastructure due to the Internet can be compared to the rise of the railroads over 100 years ago. The network fiber backbone of today is similar to a system of railroad tracks crossing the nation. Each individual track (fiber) must carry passenger trains (real-time traffic), perishable commodities (web interactions), and freight that is not time sensitive (email). Each track goes from one city to another, and cities exchange freight connected by multiple tracks. When a train arrives at a city along its route, it stops to drop off passengers and freight destined for that city. Simultaneously, the train picks up new passengers and freight destined for cities further down the line. The rest of the train's payload, the passengers and commodities continuing on, must wait until this exchange process is completed before all elements can resume the journey.

Continuing with the train analogy, the advance in optical networking through DWDM now allows more than one train (wavelength) to occupy a given track at any time. This has greatly increased the capacity of traffic between cities, but the technology used in the connecting points or "switch yards" has not kept pace. Each train added requires multiple new connections to existing tracks, configured to allow its payload to leave on any outbound track. This ability in an optical network is referred to as non-blocking switching.

In the train model of what is happening in today's Internet, when a train of railcars with passengers and freight arrives at a switchyard, the cars themselves are not switched to another track for the next leg of their journey. Instead, each car is unloaded and then loaded to a different car for the trip across the switchyard. At the other side of the yard the cars are again unloaded and loaded on a third car that will travel to the next destination city. In the optical world this is the equivalent to every OC-48 or 192 being unpacked to its smallest component (STS-1 or SDH-1) and then repacked on a different Optical Carrier stream. This is the complexity that each optical to electrical and back to optical conversion adds to managing the network.

The ideal solution to this problem would be to have all the trains going from a particular city in the network to another on a single, high-speed track. When trains arrived at an intermediate point the cars would be switched to a clear track for the next leg of their journey. As the capacity of each track expands to handle more trains at the same time and at higher speeds, the focus shifts from trying to completely fill each train to keeping the network moving as fast as possible. A pure optical network does just this for the Internet backbone.

Moving to a Pure Optical Network

Today service providers charge their customers based on specific transmission speeds, the electronically related measure of "bits per second", and are limited to selling only those speeds their networks are built for. As the capacity of the individual network fibers continues to increase, a new measurement of network usage, the wavelength, is required. In the move to a wavelength-based system, the billable unit will change to the number of wavelengths and length of use, offering network providers the ability to offer premium services that take full advantage of the available bandwidth.

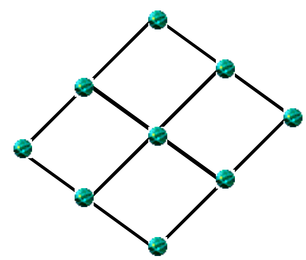

Figure 2: Mesh Network

The network that will allow for these new services will be a pure optical mesh network that operates independently of the information contained on individual wavelengths. Several breakthroughs in technology are required to make this all-optical network a reality, beginning with wavelength-routing and all-optical wavelength conversion.

The introduction of wavelength-routing and all-optical wavelength conversion will allow the network to utilize wavelengths on an individual fiber in the most efficient manner.. Bandwidth utilization in a pure optical network is defined as the ability to use any available wavelength connecting any two points in the network to create a new end-to-end connection on demand. A mesh network allows for the efficient utilization of all the bandwidths on each fiber, with turn-up time for a new wavelength reduced from months to less than a second.

In the near future, LANs will connect directly to the new optical mesh network - the Optical Internet. Photonic switching will integrate with DWDM transmission. There will be no OEO conversions or electrical regenerations for thousands of miles. Software will allow end-users to provision and manage their own services. Carriers will have the capability to profitably turn raw fiber-optic capacity into maximum usable bandwidth, on-demand anywhere in the network at the lowest cost per managed bit.

Extending the Bounds

Today's network infrastructure has reached the breaking point. Hoping that current technology can be made to accommodate Internet growth is a risky proposition. In each major market, there are several network providers in fierce competition to grab market share and meet customer demand. It is clear that today's public network has limits and that the application of band-aid solutions only increases network management complexity and overall costs.

It's time to evolve to a new measure of bandwidth—the wavelength—, and the pure optical mesh network that can fully utilize it.

Luxcore, an Atlanta-based company, is delivering on the vision of this network by building an optical internetworking system based on new photonic switching technology. By eliminating the need for electronic interfaces, Luxcore's **lambdaX** product line delivers a comprehensive and fully scalable solution that fulfills market needs today and in the future.