YIPES NAN

ENTERPRISE APPLICATIONS FOR NATIONAL ETHERNET

MICHAEL KENNEDY, PHD
NETWORK STRATEGY PARTNERS, LLC
MANAGEMENT CONSULTANTS TO THE NETWORKING INDUSTRY
www.nspllc.com
October, 2003

NETWORK
STRATEGY
PARTNERS

Network Strategy Partners, LLC (NSP) — management consultants to the networking industry — helps service providers, enterprises, and equipment vendors around the globe make strategic decisions, mitigate risk and affect change through custom consulting engagements. NSP's consulting includes go-to-market strategies, development of new service offers, pricing and bundling as well as infrastructure consulting. NSP's consultants are respected thought-leaders in the networking industry and influence its direction though confidential engagements for industry leaders and through public appearances and trade magazine articles. These interactions assure NSP's clients that they will be among the first to know the latest industry concepts and emerging technology trends. Each consulting engagement is uniquely structured—no forced methodologies or canned reports are employed. NSP's consultants' collective experience is derived from leading firms across a broad spectrum of professional disciplines including management consulting, engineering, marketing, financial analysis, and IT management. Contact NSP at www.nspllc.com.

Contents

Executive Summary

In this paper Network Strategy Partners describes a new form of wide-area Layer 2 connectivity—national Ethernet—and examines the economic and productivity benefits it delivers to enterprises that are using networks to create a single national workplace. Our analyses are taken from case studies and interviews with enterprises as diverse as financial services firms, universities and law firms. We found that by extending Ethernet nationwide and by building upon its original application as a Local Area Network protocol, and more recently as a Metro Area Network protocol, that workers are able to collaborate on projects whether they are sitting together around a conference table, working at home or located in an office thousands of miles away. This method of working yields many benefits including improved business continuity—important in a world increasingly concerned about terrorism—and the ability to bring the most knowledgeable people and up-to-date data to bear on every decision. We found, just as importantly, that the economics of national Ethernet, as compared to traditional Private Line services, make this new network-centric way of working more affordable and consequently within the reach of more enterprises and more of their sites, just as in metro area Ethernet networking.

National Ethernet services have been pioneered by SIAC—the technology subsidiary of the New York Stock Exchange and the American Stock Exchange, the major research universities that are building Internet2, and by Yipes. These national network designs support easily reconfigured and resized network links, low-cost scalable network interfaces, and network services that can interchangeably support any type of data traffic (Protocol). The designs feature national connectivity when and where it is needed, without wastefully nailing up unused bandwidth for years just in case it might be needed to accommodate unforeseen growth and/or rare catastrophic events, as is done with traditional Private Line services.

Ethernet's universal appeal as a LAN protocol is now being extended to national networking applications. For example, its connections being Ethernet are "Plug and Play." No specialized LAN/WAN equipment is required and the same interface is used for 1 Mbps connections as 1 Gbps connections. Furthermore, it supports all types of Layer 3 protocols at every interface—essential to building networks that are both resilient and affordable.

This report also drills down into the specific sources of cost savings of Yipes National Area Network (NAN) service as compared to traditional Private Line services. Three primary sources of cost savings are identified.

1. Scalability – NAN is scalable from 1 Mbps to 1 Gbps. Bandwidth can be increased without an onsite service call within a few minutes and not more than two hours. This ability to scale bandwidth on short notice gives network managers the ability to increase average circuit utilization by at least 50% over that achieved with traditional Private Line services. This is so because traditional

service contracts lock in capacity for three to five years and because there are very large bandwidth gaps between the DS1, DS3 and SONET data rates.

2. Lower Prices – Monthly recurring service charges are rarely more than those of traditional Private Line services and can be ½ to $^{1}/_{3}$ as much, even before accounting for the utilization effects (see above).

3. Declining Unit Prices – Traditional Private Line unit prices ($/Mbps) are flat with increasing bandwidth. No incentive is provided to scale up to additional capacity. In contrast, NAN unit prices decline with increasing capacity. This is a pass-through of the benefits Yipes realizes by employing optical Ethernet technology. Since communications budgets rarely increase as fast as bandwidth requirements, the NAN pricing structure makes expansion plans affordable.

As traditional early adopters of new technologies, both the Financial and Higher Education sectors have already embraced Yipes' approach of utilizing Ethernet for metro area connectivity and its logical extension into national area networking (NAN) by building Gigabit Ethernet networks of their own to serve their respective communities of interest. Yipes' NAN provides an identical service but uses a pricing structure that makes this approach affordable to a broad range of enterprises. It removes the financial risk of building the flexible high-speed national network enterprises need to support the new network-centric way of working. Network architects can build initial network configurations within existing budget limits, now being committed to traditional Private Line services. They also have the assurance that they can expand the capacity of the network as needed and with little advance planning. Large bandwidth can be supplied shortly before it is needed and not wastefully tied up in unused capacity for years in advance.

Introduction

New networking technology is changing the way work is done and helping enterprises respond to external forces including terrorism, environmental concerns and trade liberalization. The result is an increase in collaboration among workers throughout the nation and an increase in the amount and scope of information available to all workers, regardless of their location. These nationally oriented business models have taken hold in fields as diverse as education, finance and legal services.

Ethernet, originally a LAN protocol, and more recently developed as a Metro Area Network (MAN) service, now has been developed to provide national service. Figure 1 shows a conceptual view of national Ethernet service.

Figure 1

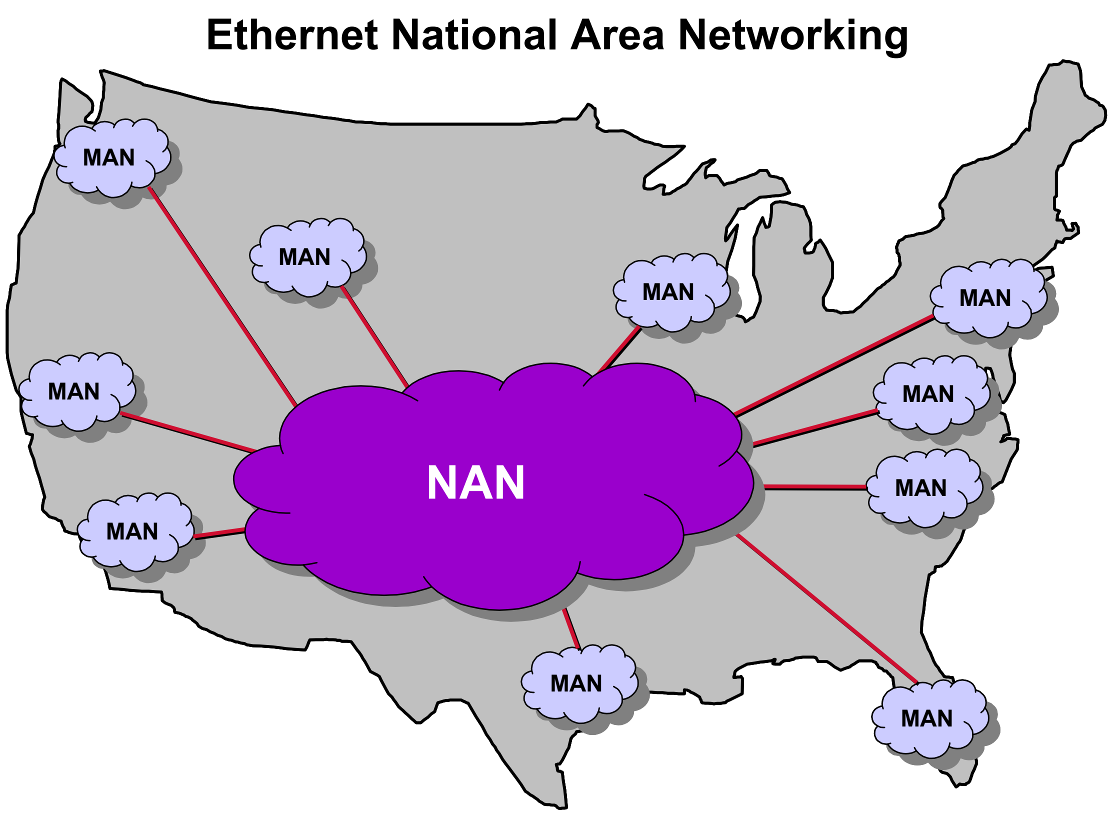

National Ethernet service extends the Metro Area Network concept wherein all workers within a single metro area appear to be working on a single enterprise LAN to the national level. The end-user's network experience appears to be identical when communicating with a colleague or computing resource at the next desk or in another metro area. This extends the successful LAN paradigm across the nation, including the instantaneous sharing of data among workers that facilitates group collaboration and sharing of enterprise resources such as databases, ERP systems and servers. This, in turn, increases the velocity of business by eliminating previously time-consuming handoffs across offices and regions. It also increases the scope of information an individual can access and can result in improving decision-making, reducing time and cost of obtaining information, and improving productivity, in that decision-makers can more easily compile facts on their own without the assistance of specialists or aides. Our research also shows that national Ethernet economics mimic those of Ethernet LANs, in which bandwidth becomes inexpensive and is no longer an obstacle to such network-centric operations.

Figure 2 provides a high-level view of the technology underlying national Ethernet networks.

Figure 2

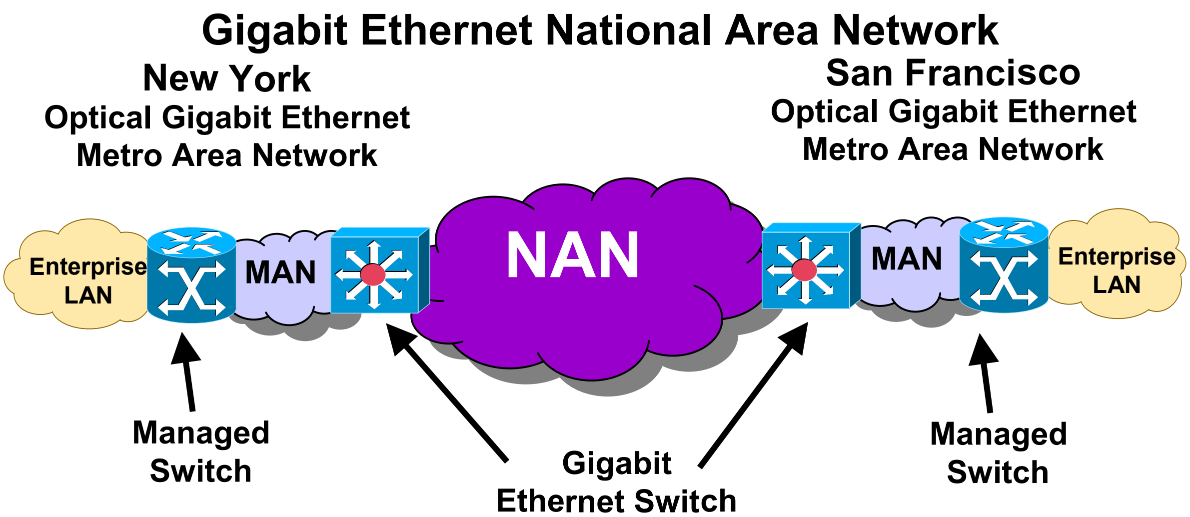

Optical Gigabit Ethernet switches link the enterprise's LAN to an Ethernet MAN and in turn to the national Ethernet service. The result is a single Gigabit optical switched Layer 2 network connecting all members of the enterprise. A worker in the New York metro area, for example, has Layer 2 Ethernet connectivity not only to all of his colleagues and computing resources in New York but all of those in San Francisco as well. To the worker, it all appears to be the same LAN. Similar "LAN-Like" connectivity is provided to all other served metro areas.

National Ethernet service provides the simplicity, granularity, scalability and ease of use required to support nationally dispersed work sites with timely and secure information. The resulting business benefits include:

- Improved business continuity
- Terrorism counter measures, such as work force dispersion
- Improved decision-making through the pooling of the talents of more people
- More timely data
- Equal dissemination of information throughout the nation

- Improved life styles by removing location from employment decisions
- A better environment through reduced work-related travel
- Better utilization of IT resources

This paper drills down into each of these business benefits and shows how, by extending Ethernet functionality to the nation, Ethernet is a more effective and cost-efficient means of delivering these benefits than traditional Private Line services.

Business Continuity

Concerns about business continuity were thrust upon all enterprises by the terrorist attacks of September 11th and their continuing aftermath. One obvious strategy to reduce risk exposure is to nationally disperse the work force and critical assets so that an attack at one point will not destroy the entire enterprise. The recent massive electric power blackout provided further evidence of the need for national, rather than regional or metro, risk-minimization strategies. IT architects, making their contributions, are redesigning their computing and networking infrastructures to eliminate single points-of-failure and provide equipment and network link redundancy and national diversity.

Traditional transport services use a costly brute force approach to create network link redundancy. Since such service cannot easily be resized or reconfigured to handle new protocols, designers are forced to size every link to meet maximum traffic loads and all conceivable Layer 3 protocols. Therefore, network resiliency improvements depend entirely upon how much extra budget can be allocated to oversized links in anticipation of rare events. Many designs, for example, include backup data centers used in the rare event of loss of the primary data center. Traditional network transport services would require that all connections into each data center be sized to carry the maximum amount of traffic of each type carried by the primary data center. Consequently, most of the transport capacity linking backup data centers is wasted since it is rarely, if ever, used.

Ethernet services, however, can be resized easily and carry all Layer 3 protocols. They provide a more resilient and robust network than traditional services, regardless of the available budget, because the entire budget is directed toward carrying live payload whether it is carried on the primary or secondary route. National Ethernet's flexibility also provides better protection against wholly anticipated events in that it can be easily reconfigured, whereas traditional solutions cannot.

Financial Services

Financial services companies—perhaps the hardest hit by the events of September 11th—began re-thinking their operating practices to be better prepared for future events and to reduce their risk-exposure through geographic-dispersion of business operations. SIAC (Security Industries Automation Corporation)—the technology subsidiary of the New York Stock Exchange and the American Stock Exchange, with responsibility for the design, development, implementation and operation of the exchanges' computer systems and communications networks—has developed an optical networking system, SFTI (Secure Financial Transaction Infrastructure) to meet the financial industry's needs for a secure data communications infrastructure. SFTI embodies the national Ethernet design concept.

SFTI (pronounced "safety") incorporates a number of design principles to insure that financial data traffic will continue to move uninterrupted, improving the industry's protection against possible threats including:

- No single point-of-failure.
  - Redundant equipment, connections, power supplies, and network links
  - Self-healing fiber-optic rings
- Connections are fungible.
  - Any working connection can carry the traffic of a broken connection.
  - All connections support all traffic types.
- Circuits do not run directly to SIAC Data Centers.
  - Users connect to two or more Access Centers via their carriers of choice
  - Access Centers and SIAC Data Centers are connected with geographically- and physically-diverse pathways
- Utilization of Access Centers provides the flexibility to add, move or change Data Center locations without disrupting communications.
- SFTI Access Centers and Data Centers provide clients with active management and support.

Figure 3 shows the use of NAN to provide a hypothetical financial services firm located in Houston connections into SFTI Access Centers located in Chicago and New York.

Figure 3

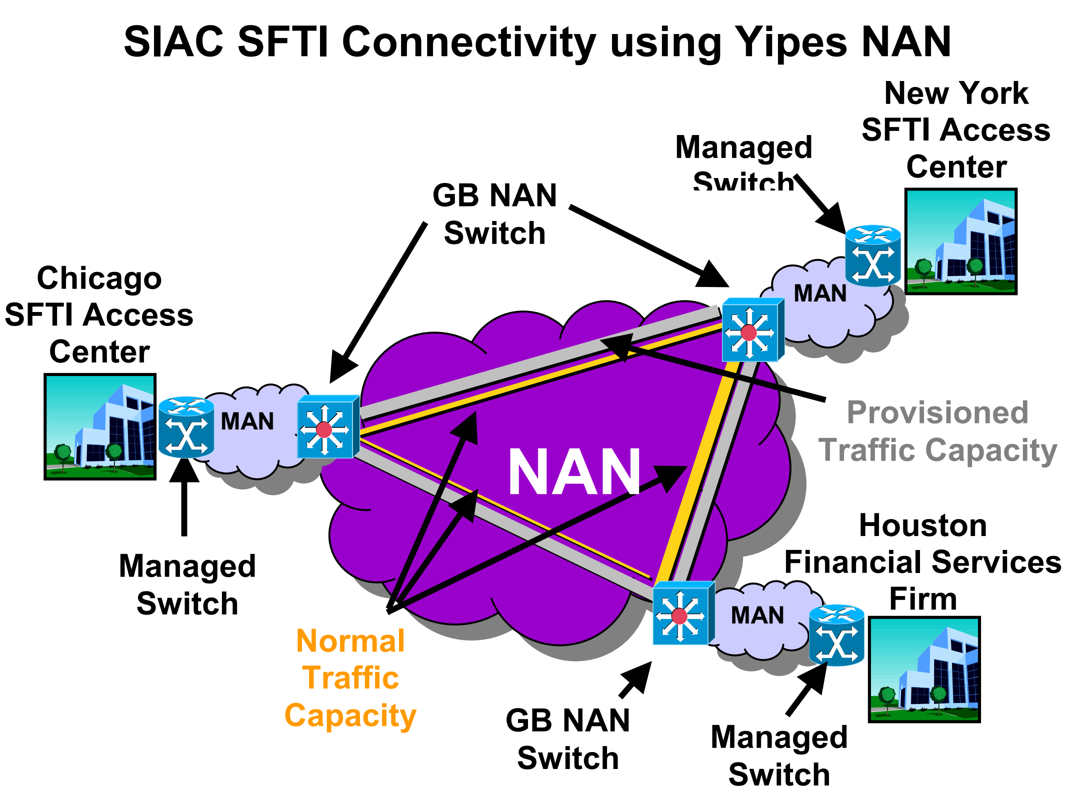

Yipes NAN provides the financial services firm located in Houston regionally-diverse, but affordable connections to SFTI Access Centers located in New York and Chicago. The diverse design is made affordable because the intercity links can be resized easily to meet potential emergency service requirements rather than being permanently nailed-up as traditional services would require. In this hypothetical example, more than 90% of normal traffic would flow between Houston and New York. The Houston to Chicago link would normally carry 10% of total traffic. Traditional network designs would require that the Houston to New York link be at least 10% underutilized while the Houston to Chicago link would be at least 90% underutilized so that each link can carry 100% of the traffic during an emergency. No such waste would be incurred using NAN because each link can be quickly resized in an emergency. SIAC refers to this re-use capability as fungibility, meaning that the circuits can be put to other use. The term is borrowed from finance where a fungible asset is one that is more valuable because it has multiple uses.

NAN also supports ring and/or mesh connectivity through its redundant optical links and Ethernet switches eliminating single points-of-failure. NAN provides further network resiliency by eliminating the need to configure complex router interfaces and plan for future network administration—essential to rapid redeployment after a disaster or terrorism attack. NAN interchangeably supports Layer 3 protocols such as IP, IPX, AppleTalk, and IP multicast as well as Layer 3-based routing protocols such as RIP, OSPF and EIGRP. Traditional Private Line services would require extensive re-working of routers and port cards to achieve comparable re-use. Finally, SFTI envisions flexible modular interchange, add and removal of any combination of Access Centers and Data Centers. NAN uses pervasively available Ethernet to support this flexibility strategy far more easily than do traditional Private Line services, with their fixed configurations and complicated interface specifications. NAN provides an Ethernet handoff to the SFTI Access Center. There is no need to procure space and install a router at the SFTI access node, as is required with traditional Private Line services.

Legal Services

Dispersion of data centers with a built-in capability for any one center to take on the duties of the others in an emergency is not limited to the high-end requirements of the financial services industry. A law firm, for example, with 18 offices in the United States and elsewhere had been operating its business with a star configuration with networked links homing on its New York City headquarters. The company-wide data applications were supported by a New York area data center while each regional office had a data center used for its local applications (See Figure 4).

Figure 4

Original Hub & Spoke Network Architecture

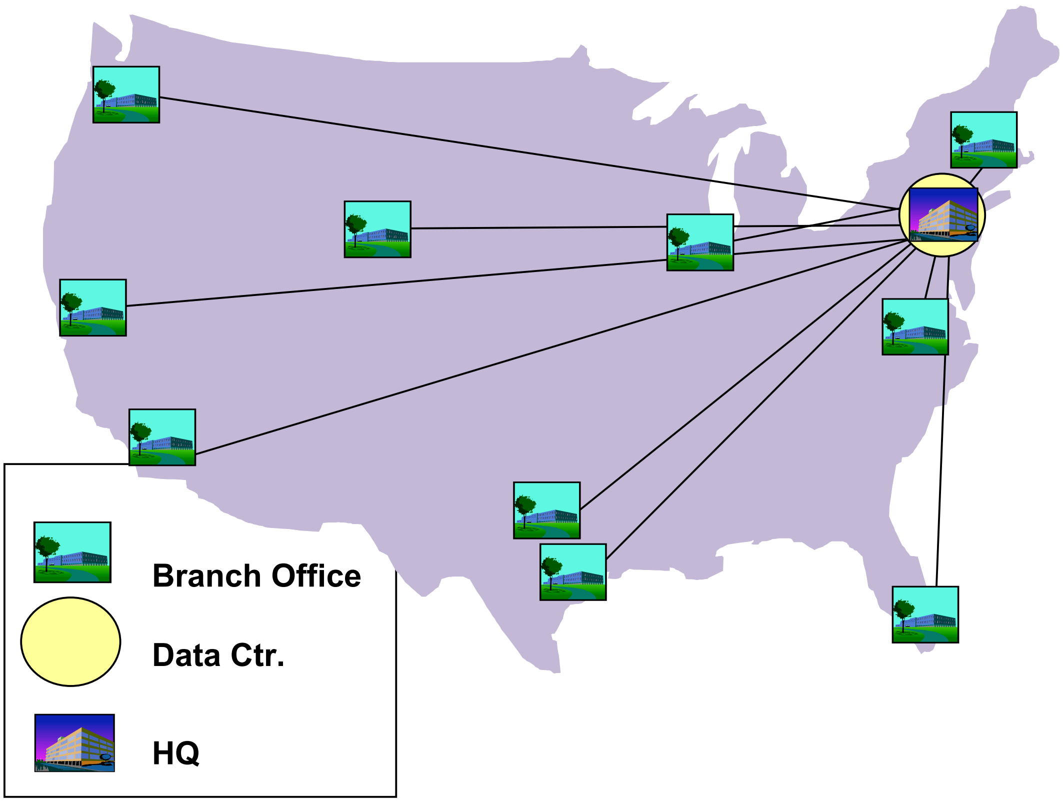

This hub and spoke configuration is typical of many enterprises since most communications are between the headquarters location and the branch offices. The September 11th attack, however, forced the law firm to rethink its network design and IT practices. The company redesigned its network and data center operations so that data processing is supported by four data centers—collocated at headquarters and three regional offices—each capable of performing all required processing activities. The headquarters and other offices are connected with nationally diverse network links, allowing for redundant access to at least two data centers from each office (See Figure 5).

Figure 5

New Resilient Network Architecture using NAN

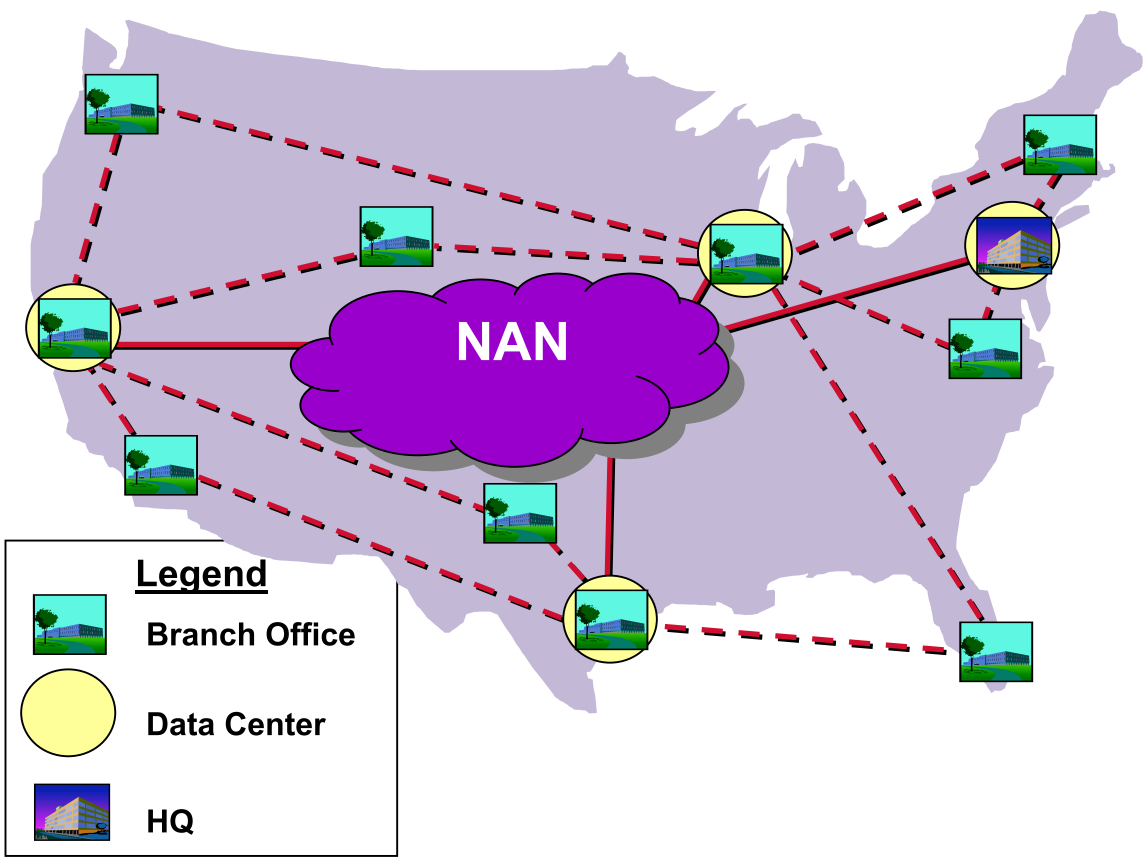

This design not only eliminates all single points-of-failure, but also saves IT costs. IT costs are reduced through the consolidation of computing activities to four data centers versus the original design that employed a data center at every location. NAN, through its low-cost and flexible Ethernet service, provides simple and easily reconfigurable support for this type of flexible and robust network design. Using NAN, no network link need be sized at more than 1 Mbps. However, should emergency reconfiguration be required any NAN link can quickly be scaled up to as appropriate.

Improved Decision Making

Nationally distributed business models have many other benefits beyond supporting business continuity initiatives. A simple, scalable and easy to use network design such as national Ethernet supports better decision making through its ability to link more people together with the right data to make better decisions. The economic benefit derived by bringing more people together with good data can be illustrated by comparing an eBay auction with a traditional one. Sellers benefit by being able to show their products to prospects worldwide—anyone on the Internet—rather than just those who show up at a physical auction. They also have a better chance to obtain a higher sale price with more bidders participating. Buyers also benefit by being given more products (and opening prices) to choose from which protects them from bidding above the market price. This benefit of using a network to bring the best knowledge and thinking to bear on a problem is applicable in many enterprises. Health Care providers, for example, are using networking technology to bring the best experts together within their entire enterprise to advance medical practices. In the absence of a national network, research is clustered within pockets of senior practitioners and their students grouped within each of several teaching hospitals (e.g., Boston, Baltimore, Houston).

Universities

The academic community has been especially active in exploring ways to use networking technology to improve the quality and reach of research and education. The Internet2 project has this as an explicit goal. Major U.S. research universities have articulated a common vision of what is needed in order to realize the next generation of network-based applications of importance to the higher education community. These include new modes of interactive collaboration and new ways of enhancing learning through IT, the integration of complex distributed digital library collections with academic programs, greater access to expensive specialized research facilities (e.g. accelerators and supercomputers) and life-long scholarly pursuits facilitated through ready access to learning materials from homes, offices, or anywhere convenient to the learner.

Table 1 lists several of the applications envisioned for Internet2.

Table 1

Internet2 Applications

Application	Description
Instructional Management System	Support student learning any-time/any-place/any-pace flexibly incorporating audio, video, and graphics. Support teachers' research and publishing activities.
Digital Libraries	Extend existing Internet-based digital libraries to incorporate images, video and audio in both stored and real-time formats.
Tele-Immersion	Allow individuals at different locations to share a single virtual environment. For example, participants could interact with each other virtually at a conference table, approximating what is possible in a physical conference room.
Virtual Laboratory	Enables a group of researchers located around the globe to work together on a common set of projects using a distributed problem-solving environment. Virtual laboratories have been proposed in many disciplines such as biology, astronomy, pharmacy, and materials science.
Tele-Medicine	Extends existing Tele-medicine capabilities with improved video quality and will allow the transfer of real-time data streams. For example, 3-D brain scans for remote consultation among brain surgeons.

National connectivity to and among major U.S. research universities has important social equity implications. These universities are heavily funded by Federal Grants but access to research results in the absence of a simple, flexible and affordable network would be limited to a select few able to work in the labs and attend the lectures at each of these universities. National Ethernet designs provide the simple, flexible and affordable national connectivity envisioned for the Internet2 project.

Internet2 employs the national Ethernet design concept. It is built around GigaPoPs, regional network interconnect points providing access to the inter-GigaPoP network for several Internet2 members (See Figure 6).

Figure 6

Internet2

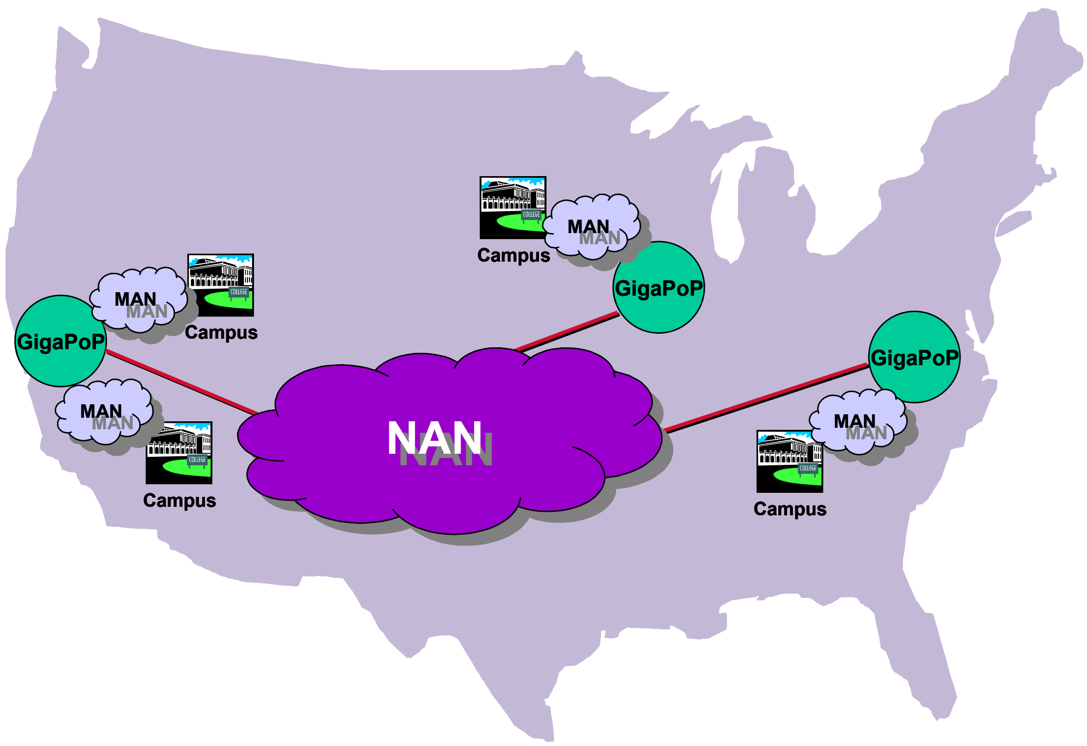

A GigaPoP is a secure and environmentally conditioned location that houses communications equipment connected with geographically diverse optical (and wireless) communications paths to maximize service reliability. (Note the design similarity to the financial industry's SFTI.) Bandwidth requirements are now in the 1-5 Gbps range and will extend to 10 Gbps. However, this design does not leave out students and faculty working at home or in small off-campus locations. Each GigaPoP acts as the hub of a Metro Area Network providing access from many medium and small sites onto the very high bandwidth Internet2 backbone.

National Ethernet is ideally suited to providing the national connections among the GigaPoPs. It can scale from 1 Mbps up to hundreds of Mbps in actual throughput. This provides flexible and affordable national connectivity essential to the broad reach and accessibility envisioned for Internet2. Traditional data transport services with fixed bandwidth allocations lack the flexibility needed to balance data throughput requirements with budget allocations making them a poor tool for increasing access to the best research and teaching.

Network Service Requirements

Network services needed to support working across the nation as if everyone is located in the same office embody three primary requirements.

1. Simplicity – Connections should be "Plug & Play." A simple interface should be employed suitable for very large installations down to very small ones. The service should support all types of Layer 3 protocols. Specialized designs or specialized interfaces should be avoided.

2. Scalability – The service should easily support changing bandwidth requirements from 1 Mbps up to the Gigabit range. Even under ordinary operating conditions traffic requirements are unpredictable. On one hand very inexpensive LAN connectivity even up to 1 Gbps is pushing up wide area network performance expectations while the wide-spread business decline, especially in the High-Tech sector, has made earlier demand projections look ludicrously optimistic. Scalability is even more important in support of business continuity and disaster recovery actions—any network link should be easily modified to assume the load if others are lost.

3. Affordability – The new network designs envision multi-megabit per second connections for very small sites with many bandwidth increments ranging nearly to a gigabit for large sites and GigaPoPs. Under traditional Private Line pricing schemes small sites lack the budget for multi-megabit services and even the wealthiest firms are resistant to Gigabit pricing. Furthermore, traditional Private Line service offers have big bandwidth and price gaps that most enterprises cannot afford to cross (For example, 1.5 Mbps, 45 Mbps and 155 Mbps.)

Ethernet meets these requirements much more effectively than traditional Private Line services. It is simpler to implement and manage then traditional Private Line services. No specialized WAN equipment (CSU/DSU, Router, or associated rack, power, and equipment space) is required to connect to the service. All of an enterprise's existing switches and routers are interchangeable with any growth or reconfiguration of the network topology (since they all have the same interfaces). The dual nature of existing Layer 2/3 switches exploits this simplicity. All Layer 3 protocols are supported. Furthermore, as a Layer 2 (virtual circuit) service Ethernet provides a familiar and proven migration from legacy Layer 2 Frame-Relay services. In contrast, substantial network reworking to meet diverse site requirements is required when using traditional Private Line services. Network designs, for example, are dramatically different for sites using T1 versus OC3 services.

Traditional services offerings also are difficult and costly to scale up to additional bandwidth. A service upgrade from DS3 (45 Mbps) to OC3 (155 Mbps), for example, is a major undertaking. At the customer premise, the CSU/DSU equipment must be replaced as well as the customer's WAN router port (or possibly the entire router). Within the carrier's central office, the fiber must be moved from the carrier's Digital Cross Connect system to a SONET port typically on another system. Once these physical changes are completed, network engineers must determine and assign an STS 3 (Or three STS 1) circuits through each cross-connect and SONET multiplexer from one end of the service to the other potentially involving many transport spans and cross-connect points. Should circuit capacity be inadequate at any point on this route other customer circuits may need to be similarly rearranged to accommodate the new service. Once the engineering work is completed, the commercial and accounting transactions are equally complex and cumbersome. OC3 is a different service offering than DS3 so the existing DS3 service contract must be terminated likely incurring an early termination fee. A new contract must be negotiated for the OC3 service. This usually involves at least a three-year term and substantial up-front non-recurring charges due to the cost and difficulty of provisioning services.

NAN Pricing Analysis

The preceding sections of this report described the benefits of national Ethernet designs including the financial industry's SFTI, research universities' Internet2, and Yipes NAN. This section takes a closer look at the pricing structure of NAN and shows how it makes the leading-edge capabilities such as those enjoyed by the New York Stock Exchange and the country's major research institutions available and affordable to a wide range of large as well as smaller enterprises.

Traditional Private Line services are offered at a small number of data rates (e.g., 1.5 Mbps, 45 Mbps and 155 Mbps) and usually require a three- to five-year contract commitment. NAN service, in contrast, is smoothly scalable in 1 Mbps increments from 1 Mbps to 1 Gbps. Bandwidth, furthermore can be modified in anywhere from a few minutes to a couple hours. The pay-as-you-go structure of NAN service pricing versus the three to five-year commitment required under traditional Private Line service tariffs provides a dramatic utilization improvement and cost advantage for users of NAN service. Figure 7 illustrates the source of this cost advantage.

Figure 7

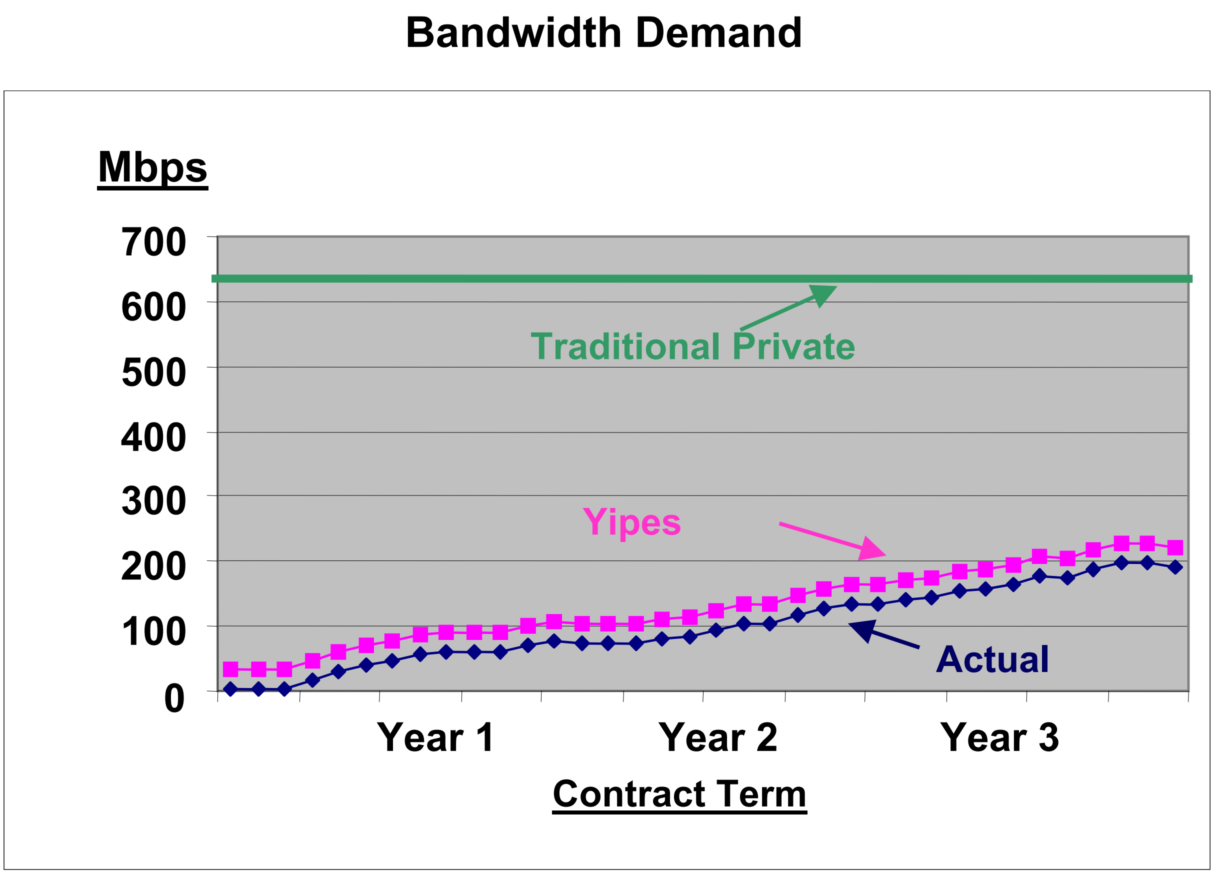

Figure 7 shows actual traffic and the amount of bandwidth that would be provisioned month-to-month for Yipes NAN and a traditional Private Line service over a three-year traditional Private Line service contract. Since the traditional Private Line service contract is for three years, the network manager must project his needs up to three years in advance. In this example actual usage will peak at 200 Mbps in the 10th month of the third year. Since the penalties of being caught short of bandwidth are very great orders are usually padded with some additional capacity. Therefore the three-year traditional Private Line services contract must be for at least 200 Mbps of bandwidth despite the fact that this capacity will not be reached until the end of the third year. The smallest data rate offered using traditional Private Line that will meet this minimum requirement is OC12 (622 Mbps). Consequently, the enterprise is locked into a service with huge excess bandwidth capacity for three years. No more than 32% of the available capacity will be used over the three-year contract period while average utilization will be 13%. In contrast, NAN service can be purchased with month-to-month bandwidth allocations. In this example, the monthly purchases are always 30 Mbps greater than monthly demand. Such close matching of actual demand to supplied bandwidth is feasible because forecasting mistakes can be corrected within an hour or two rather than being locked in for three years as is the case with traditional Private Line services. In this example, average NAN bandwidth over the three-year period is 110 Mbps and doesn't cross the 100 Mbps threshold until near the latter part of the second year. The actual financial savings from the use of scalable NAN service is even greater than implied from a bandwidth perspective. This is so, as we will now show, because NAN service is much less expensive (megabit for megabit) than traditional Private Line service.

Traditional Private Line service pricing consists of two price components—non-recurring and monthly recurring service charges. Traditional Private Line services customers also incur the cost of owning and operating CSU/DSU and routers needed to connect their LAN to the WAN. Since NAN requires no additional customer equipment for line termination and most bandwidth upgrades are done without an onsite visit, there is no cost of extra equipment and non-recurring charges are low and infrequent. Monthly recurring charges for NAN also are very low as compared to traditional Private Line service as we will now illustrate with several examples.

Figure 8 provides several examples of traditional Private Line service offerings using DS1 and DS3 Private Line services with comparably sized NAN Layer 2 virtual circuits. On-net pricing refers to service that uses a Yipes optical connection from the customer premise to the closest NAN PoP. Off-net service uses a third party leased line to connect the customer to the closest NAN PoP.

Figure 8

NAN Pricing Examples

- **NY-Chicago, 10 Mbps, Colo to On-net:**
  - ***Yipes NAN MRR = $1848***
  - ***DS3 range = ~$3800 - $9000***
- **Seattle-DC, 3 Mbps, off-net to off-net:**
  - ***Yipes NAN MRR = $3786***
  - ***2xDS1 range = ~$5000 - 12,000***
- **SF-Dallas, 30 Mbps, on-net to off-net:**
  - ***Yipes NAN MRR = $8086***
  - ***DS3 range = ~$9000 - $15,000***

NAN service is less expensive than comparably sized traditional Private Line offerings. When on-net, NAN costs about half as much as traditional Private Line. Even when off-net offerings are considered, monthly savings can be as much as 2/3 of the traditional Private Line price. The NAN service is rarely more expensive than traditional Private Line.

However, actual savings will be substantially more than these simple examples portray when the granularity of the NAN service is taken into account. Figure 9 illustrates the cost benefit of a granular service offering.

Figure 9

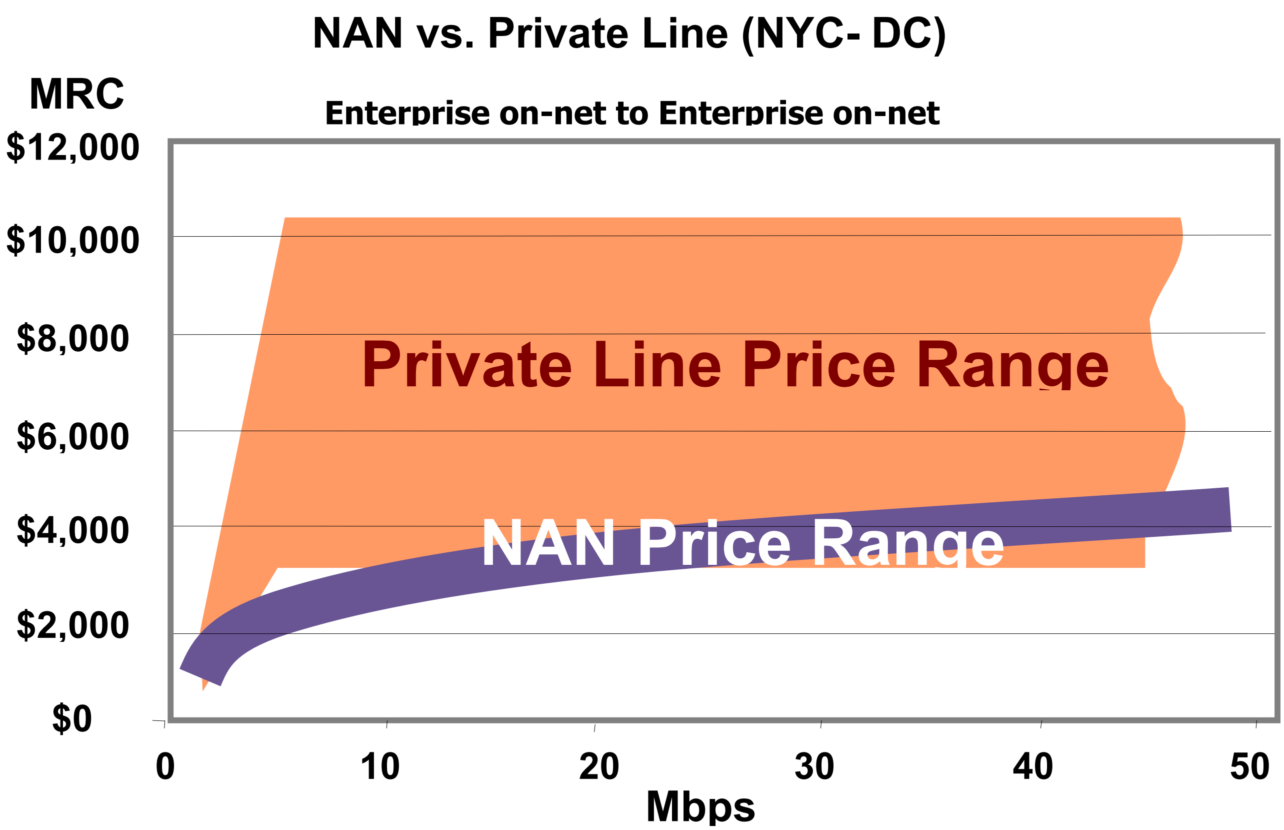

Figure 9 compares NAN pricing with T1 and DS3 traditional Private Line alternatives. The Private Line range is very broad reflecting the "cost plus" approach to Private Line pricing. The actual price for any circuit is specific to its physical build-out. Regardless of the price level, however, Private Line prices grow linearly to about 7-10 Mbps and then flatten out. This is due to their dependence on nailed up data rates—DS1 (1.5 Mbps) and DS3 (45 Mbps). Private Line prices rise with the addition of each additional T1 circuit as each additional circuit is priced the same as the first. Once bandwidth of between 7-10 Mbps is reached (4-7 T1 circuits) then it costs less to use a DS3 (45 Mbps) circuit. Customers, therefore, have no mid-priced option in the 10-45 Mbps range.

NAN pricing in contrast has declining unit cost with increasing bandwidth purchases. Its price curve flattens out with increasing bandwidth capacity—the incremental bandwidth price is declining. NAN pricing, in addition, is near or below the lowest price for Private Line service. NAN pricing is less than half of the highest price for Private Line service—even ¼ of the price in some instances. This pricing structure is supported and sustainable through Yipes' use of Ethernet switches and ports for service delivery. They are much less expensive from both a first and recurring cost perspective than the TDM and SONET designs employed by traditional Private Line service offerings. The high port capacity and software control of provisioned service underlies Yipes' ability to offer such attractive pricing.

NAN's pricing structure removes the financial risk of building the flexible high-speed national networks enterprises need to support their new network-centric way of working. Network architects can build initial network configurations within existing budget limits now being committed to traditional Private Line services with the assurance that they can expand the capacity of the network as needed and with little advance planning. Large bandwidth can be supplied shortly before it is needed not wastefully tied up in unused capacity years, if ever, before it is used.

Conclusion

Enterprises are adopting a new way of working that links all members of the firm tightly together at the same time that they are physically moving farther apart, either as part of business continuity initiatives or simply as life-style choices. New network architectures that provide flexible quantities of easily-configured (and re-configurable) bandwidth are being deployed in support of this new way of working. Yipes NAN service is an effective vehicle for constructing these new networks because it is simpler, more scalable and more affordable than traditional Private Line services.

Michael Kennedy, PhD

Michael Kennedy, co-founder of Network Strategy Partners, LLC, helps leading companies create and implement network industry strategies. Recently, he helped incumbent and new entrant service providers improve the profitability of their service offerings. Michael's work for equipment vendors is focused on helping them create

compelling business cases in support of their marketing efforts. He also works with end-users to develop network architectures, strategies and designs that meet corporate business goals.

Dr. Kennedy is a widely quoted authority on the dynamics of the communications industry. He has over 100 articles, speeches, and major conference presentations to his credit. Dr. Kennedy's bimonthly *Telecommunications Magazine* column focuses on the business applications of networking technology.

Michael has an extensive multi-disciplinary background in management consulting, telecommunications industry analysis, financial analysis and communications engineering gained over more than 30 years working in the industry. Prior to founding Network Strategy Partners, he was Director of Consulting Services for Strategic Networks, Boston, MA and lead Arthur D. Little, Inc.'s Cambridge, MA enterprise network consulting practice. Kennedy led Stamford, CT-based Gartner Group's telecommunications strategy advisory service prior to entering the management consulting profession. Michael's experience as a financial analyst was earned while providing equities research on the "sell-side" at Soundview Financial Corporation, Stamford, CT and at AT&T where he was responsible for financial analysis of the strategic plan. In addition, Dr. Kennedy held engineering positions at AT&T, as a Member of Technical Staff of Bell Labs, and at IBM.

Dr. Kennedy holds three degrees in electrical engineering, a Ph.D. from New York University, New York, NY, a master's degree from the Massachusetts Institute of Technology, Cambridge, MA, and a Bachelor of Science from the University of Akron, Akron, OH. He is a registered Professional Engineer in the State of Ohio.