SPECIAL REPORT

Corporate Networks in Transition:

Realizing the Value of Private Optical Networks

by Jim Metzler

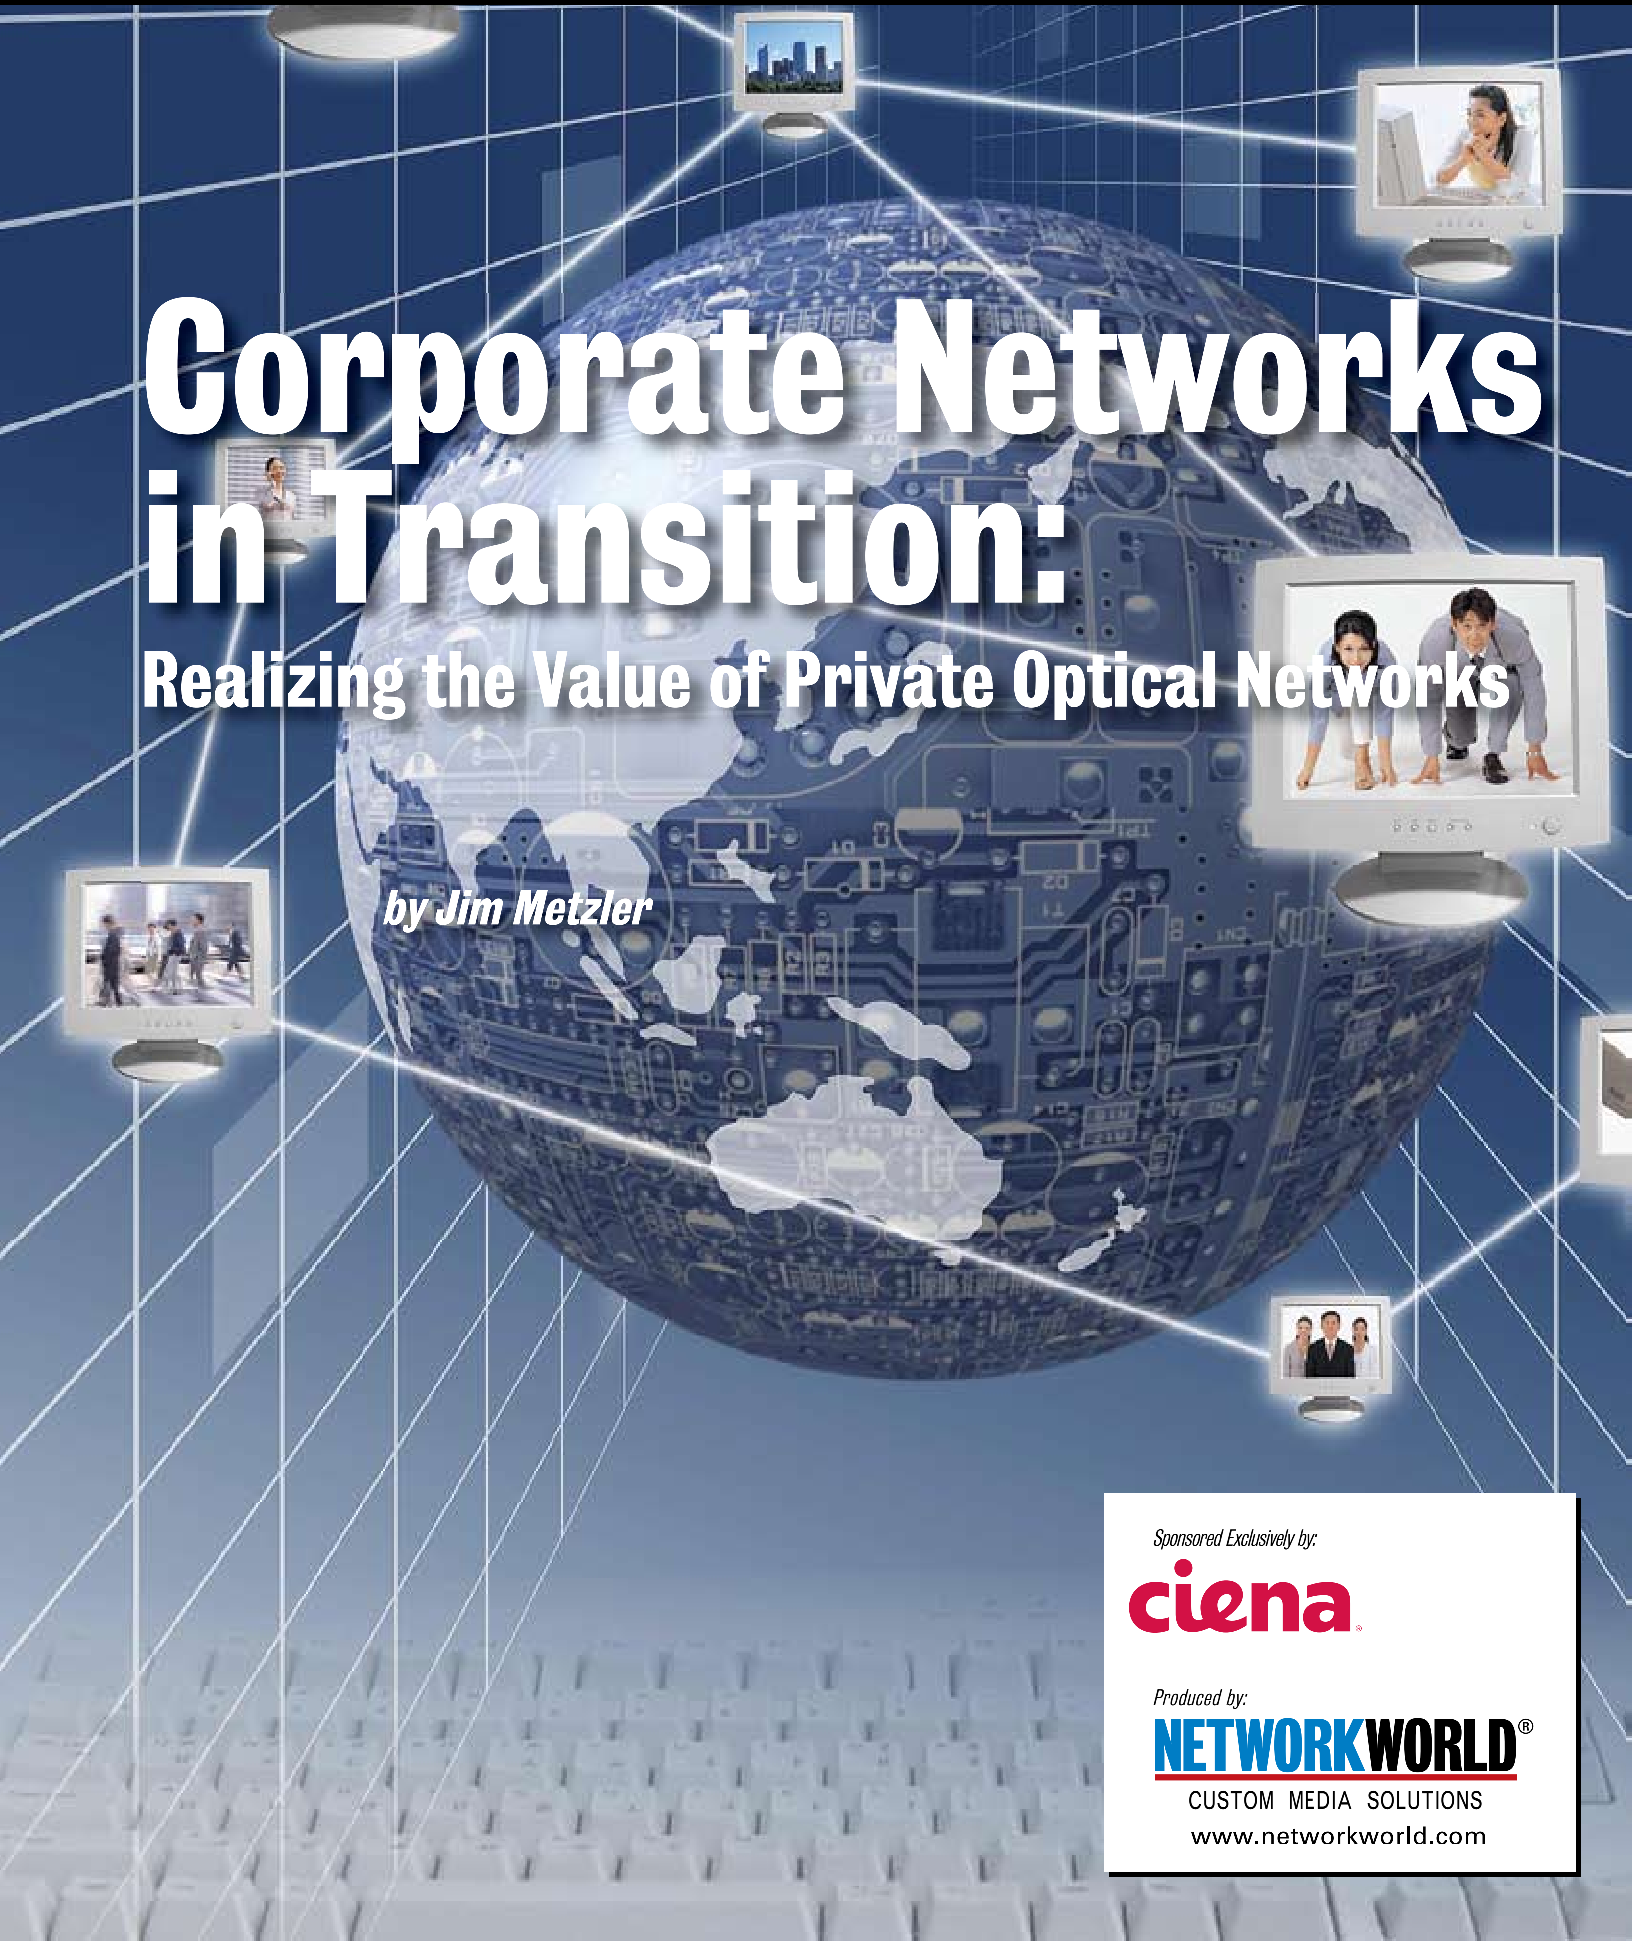

Sponsored Exclusively by:

ciena

Produced by:

NETWORKWORLD®
CUSTOM MEDIA SOLUTIONS
www.networkworld.com

Table of Contents

As recently as a few years ago, most enterprises understood the cost associated with wide area networking, but had less of an understanding of the business value of the WAN. That situation is changing. The enterprise WAN is going through a transition both in terms of its perceived business value, as well as how enterprise IT organizations architect their WAN.

For a transition to occur in terms of how IT organizations utilize WAN technologies, two fundamental factors have to be in place. One of these factors is that there must be broad-based demand that cannot be efficiently met by current technologies and practices. The second factor is that there must be alternative technologies and practices that can meet the demand more efficiently.

As will be described in the next section of this Special Report, the demand that the typical enterprise WAN must support is undergoing significant evolution and growth. This evolution and growth is caused by issues such as the increased interest in business continuity and disaster recovery, the continued issuance of government and industry regulations, the deployment of enterprise applications, and the continued movement to implement new forms of distributed computing.

To support these specific demands as well as other more general business requirements, the emerging enterprise WAN must have a set of key characteristics. For example, virtually all enterprises are under continued pressure to become increasingly agile. To enable business agility, one of the characteristics of the emerging enterprise WAN is that the WAN itself must become increasingly agile. Also, most enterprises run a non-stop operation. To support non-stop operations, the emerging enterprise WAN must be more highly available than the traditional WAN. The support of non-stop business operations also requires that IT organizations have increased control over the provisioning of additional WAN capacity to support unforeseen requirements in a timely fashion.

One way that many enterprises have responded to these new demands and business requirements is by deploying an optical network based on dark fiber or SONET. Throughout this Special Report, dark fiber is defined as optical fiber infrastructure (cabling and repeaters) that is currently in place but is not being used. Optical fiber conveys information in the form of light pulses so the 'dark' means no light pulses are being sent[1].

The use of optical networking has always been appealing because an optical network has the characteristics that were outlined above. However, the use of optical networking is becoming more appealing to enterprises in part because the availability of dark fiber is increasing. In addition, because the cost of dark fiber is dropping precipitously, there are a growing number of situations in which a compelling economic case can be made for deploying dark fiber.

---

[1] David Foxley, "Dark Fiber", 1999 TechTarget.com, Inc., 31 January 2002.

The increasing demands on the enterprise WAN

The volume of traffic that the enterprise WAN has to support continues to grow at an alarming rate. For example, one recent study[2] indicated that WAN traffic volumes were doubling every nine months.

One of the trends driving increased WAN traffic is ongoing merger and acquisition activity. For a variety of political and organizational reasons, merged organizations often maintain separate facilities. In addition, it is common for the acquiring company to want to give employees in the acquired company access to a set of core applications. All of the traffic that these applications generate has to transit the enterprise WAN.

This section of the Special Report will outline some of the other evolving trends that are placing continually higher traffic volumes on the enterprise WAN.

Business continuity and disaster recovery

The goal of business continuity is to ensure that key business processes continue to function after a disaster. A major component of business continuity is a disaster-recovery plan. The goal of a disaster-recovery plan is to ensure that key IT resources, such as data center resources, are functioning after a disaster. Interest in assuring business continuity and disaster recovery picked up notably after the Sept. 11 terrorist attacks and ratcheted up again after the recent hurricanes.

One common component of a data center disaster recovery plan is remote disk mirroring. Enterprises that implement remote disk mirroring continually back-up or mirror the contents of the disks at one of their data centers onto the disks at either one of the enterprise's other data centers, or at a site designed solely to support disaster recovery. Adding to the challenge of disk mirroring is that the size of the databases that need to be backed up is increasing at a dramatic rate.

While not as common as disk mirroring, another evolving trend is to run applications requiring high availability on separate server clusters in disparate data centers. This approach is most applicable in those situations in which the company already has an optical network in place to support disk mirroring and can use this network to also support the diverse server clusters with little or no additional cost.

There is no doubt that disk mirroring and diverse server clusters can add significant value as part of an IT organization's disaster recovery plan. There is also no doubt these techniques place significant extra traffic on the enterprise WAN that typically can not be met by use of traditional WAN technologies.

Industry and governmental regulations

The need for assuring business continuity and disaster recovery has also been driven by government regulations. For example, implementing a disaster-recovery plan is a U.S. government mandate for the financial services sector[3].

Over the last few years, a number of other government and industry regulations have also been implemented that have drastically impacted organizational behavior. These regulations include:

The Sarbanes-Oxley Act

This act requires management to make a written assertion stating their responsibility for establishing and maintaining an adequate control structure and procedures for financial reporting.

HIPAA

HIPAA (the Health Insurance Portability and Accountability Act) requires companies in the health-care industry to provide administrative simplification, security and privacy.

The Gramm-Leach-Bliley Act

This act requires companies to give consumers privacy notices that explain the institution's information-shar-

[2] Migrating the Cisco WAN to the Cisco 7600 Series Router, Cisco IT@Work Case Study

[3] "Sound Practices to Strengthen Resilience of the US Financial System", Board of Governors of the Federal Reserve System Docket No. R-1128, Office of the Comptroller of the Currency Docket No. 02-13 and Securities and Exchange Commission Release No. 34-46432; File No. S7-32-02.

ing practices and to give consumers the right to limit some of the sharing of its information.

Each of these acts requires that companies place a greater emphasis on assuring the accuracy, security and confidentiality of data. It is very difficult to accomplish these goals when there are multiple copies of the company's data on servers in branch offices. Hence, these regulations have influenced enterprise behavior by encouraging enterprises to consolidate their servers out of branch offices and into centralized data centers in an attempt to increase control over the data that is used to run their businesses. As a result of this consolidation, traffic that used to transit a high-capacity LAN in each branch office now transits the enterprise WAN.

Enterprise-wide application deployments

Over the last few years, most companies have come to regard their branch offices as a key touch point with their customers and hence critical to the success of their business. In part because of the growing importance of branch offices, many employees who used to reside in a headquarters facility now reside in branch offices.

Many companies are looking to fully leverage this customer touch point by deploying business critical applications out to branch offices. It is worth noting that many of these applications, such as VoIP, consume a significant amount of bandwidth and are notably delay sensitive.

However, as discussed in the preceding section, while an increasing number of employees are migrating to branch offices, enterprises have also been consolidating servers out of branch offices and into centralized data centers. As mentioned, one of the motivations to centralize servers is the desire to gain more control over the data that is used to run their businesses. Another factor driving server consolidation is the need to reduce the cost of the servers and the associated software licenses, maintenance, floor space, and heating and cooling costs.

While there are many compelling benefits of server consolidation, it places significant extra demands on the enterprise WAN. For example, many applications are written in a way that they perform well when they are run over a high capacity LAN. However, these same applications can perform poorly when they are run over the relatively low capacity WAN that connects branch offices to the corporate data center. Because enterprises do not want to experience application degradation, this means enterprise-wide application deployment requires the WAN to behave like the branch office LAN.

Evolution of distributed computing

In many instances, the phrase "distributed computing" is used interchangeably with the phrase "n-tier applications". N-tier applications are applications in which the functionality that had at one time been provided by a monolithic application is now decomposed into multiple tiers. Because these tiers are implemented on separate systems, n-tier applications typically make more demands on the network than do monolithic applications.

The first generation of n-tier applications was referred to as client-server applications. These two-tier applications were structured around the use of a PC that accessed an off-the-shelf database product using a language such as Structured Query Language (SQL). The information flow in these applications went from the PC to the database and back again over a company's private network.

The success of client-server computing led to the deployment of more sophisticated application architectures, such as three-tier applications. The typical three-tier application is comprised of a Web browser, Web server and a database server. The information flow in a three-tier application travels in a linear fashion from the Web browser to the Web server and to the database and then back again over the Internet using standard protocols such as HTTP. This form of application has become quite popular despite the addition of significant traffic to the WAN. In particular, applications that use HTTP consume five to 10 times as much WAN capacity as applications that use a thin client.

The next step in the development of distributed computing is the deployment of Web services-based applications. Web services refers to reusable software modules that encapsulate business functionality and which can be accessed over an IP-based network.

Figure 1: Web Services Based Applications

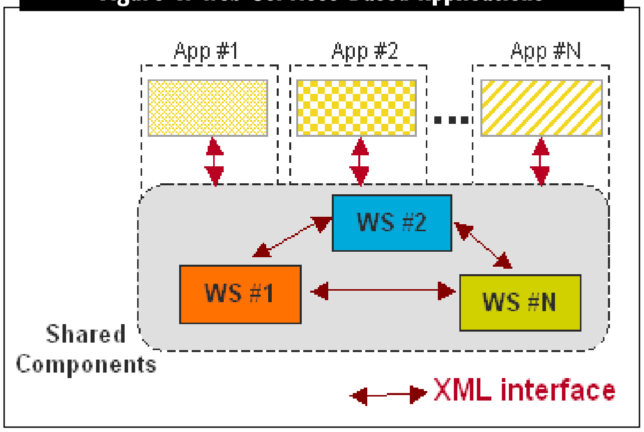

Figure 1 illustrates the relationship between applications and Web services. Specifically, a given application is comprised of a number of Web services that communicate among themselves using XML. In addition, a given Web service may be a part of multiple applications.

As an example, Table 2 contains a set of possible relationships between a company's applications and its Web services.

Table 1 : Relationship between Applications and Web Services

Application #	Web Services that Comprise the Application
1	WS #1, WS #5, WS #7, WS#17
2	WS #3, WS #5, WS #12, WS #19, WS # 25
3	WS #1, WS #5, WS #14, WS #20
4	WS #5, WS #11, WS #17

As shown in Table 2, Web service No. 17 is used in two applications — application No. 1 and application No. 4. Web service No. 5 is used by all four of the applications.

One of the primary reasons why Web services-based applications have become so popular is that they enable business agility. In a Web services-based environment, it is easy to create a service that is a network-enabled component representing business functionality that can be easily reused. As a result, IT organizations that adopt a Web services approach to applications development can modify business processes such as supply chain management and CRM in a fraction of the time and for a fraction of the cost than in a more traditional application environment.

However, Web services-based applications demand a highly available, high-capacity WAN in part because in most cases the servers that run the Web services are housed within:

- Multiple data centers owned by a given organization.
- Multiple data centers owned by different entities.

These application deployment options place tremendous demands on the WAN infrastructure over which they run. Applications that would have previously executed within a single data center now execute in multiple data centers. This results in a very significant amount of additional traffic placed on the WAN. For example, James Kobielus[4] estimates that XML requires two to 100 times more bits than non-XML binary transfers. However, in spite of this additional traffic and the latency that is inherent in a WAN, these Web services-based applications are expected to perform as if they were running over a LAN in a single data center.

[4] "Wrestling XML Down to Size", James Kobielus, BCR, December 2004

The need for a highly available, controlled, agile WAN

One of these emerging business requirements is that most enterprises run a non-stop operation. Given that the majority of companies run their operations over their network, if the enterprise WAN is not available, business operations are impacted. As a point of reference, a recent study[5] looked at the cost of network downtime on a variety of business operations. According to that study, the average hourly cost of network downtime ranges from $90,000 for a catalog sales operation to multiple millions for either credit card sales authorization or brokerage operations.

Table 2 translates WAN availability into unplanned downtime. Table 2 also translates unplanned downtime into the revenue that it impacts on the assumption that the enterprise books $3,000 a minute in revenue over its WAN.

Table 2 : WAN outages

WAN Availability	Annual Unplanned Downtime	Revenue Impact
99.0%	88 hours	$15,840,000
99.9%	8.8 hours	$1,584,000
99.99%	Less than 1 hour	$158,400
99.999%	Roughly 5 minutes	$15,840
99.9999%	Roughly 30 seconds	$1,584

As noted, if the enterprise WAN is not available, business operations and revenue are impacted. For example, customers who are unable to buy products and services will seek other suppliers, and sales are lost while the WAN is unavailable. However, in many cases, if this new supplier provides a high-quality product or service, the customer and their associated revenue are lost forever.

There are many other ways that a WAN outage impacts business. Employees in a branch office whose job calls for them to access a centralized application cannot do so, and they become non-productive. Customers attempting to contact the company with a question about a product or service become frustrated. In some cases, a business outage that was caused by a network outage gets reported on the front page of prominent publications such as the *Wall Street Journal*. That notoriety certainly diminishes the reputation of the company and can negatively impact the company's market capitalization.

Another key business trend that the enterprise WAN must support is the rapid pace and high demands related to business change. This means that not only does the WAN have to support mergers and acquisitions, but it must do so the day that the merger or acquisition is finalized. Analogously, the enterprise WAN has to support new facilities on the day the facility is activated, whether that new facility is a small branch office, a data center, or a manufacturing facility.

In addition, more than 50% of companies are currently making significant changes to one or more business processes. In the traditional application development environment, if a business manager wants to change a key business process it is the applications organization that is the roadblock. This follows because it typically takes the applications organization time to either develop a new application or integrate existing applications.

However, with the deployment of Web services-based applications, a business manager can modify a business

[5] The Effects of Network Downtime on Profits and Productivity, Performance Technologies, Inc. www.netcordia.com/news/TheCostOfNetworkDowntime.pdf

process and require little if any additional programming. As such, the applications organization will no longer be a roadblock preventing business managers from deploying new processes. But if the WAN is not capable of meeting the challenges that are associated with Web services applications, then the network organization becomes the roadblock.

To support these emerging business trends, what is needed is an enterprise WAN that is highly available, controlled and agile. As was pointed out by the New Paradigm Resources Group[6], networks based on dark fiber fit that description as they provide organizations the following advantages:

- Allows the organization to control the growth and development of its network by providing nearly unlimited bandwidth at no incremental cost.
- Enables stronger guarantees relative to network capacity, speed, flexibility, security and reliability.
- Provides greater control over the provisioning of telecom services.
- Lowers cost.

The availability and economics of optical networks

Over the last few years, many articles have been written that discuss the increasing use of dark fiber on the part of companies in various industries. One such article was authored by Marguerite Reardon[7]. In the article she notes that during the 1990s a wide variety of companies dug up city streets to put in fiber, leading to an excess of dark fiber that has resulted in falling prices. As a result of these conditions, Reardon concludes that "even many midsize companies can afford to take out long-term leases on dark fiber and buy the equipment to run their own network on it."

One company that has deployed a fiber network is Gannett, having built metro fiber networks in three cities in which it has newspapers. According to Gary Gunnerson[8], IT architect at Gannett, "Analysts have long said that it's too expensive and requires specially trained optics engineers to build and run these networks. I have found just the opposite to be true. The fiber and the equipment are so cheap now, and anyone who is familiar with IP networking gear can handle a short-distance optical network. I could teach them in half a day," he says. Gunnerson added that companies that spend $7,000 to $10,000 per month on telecommunications services should consider building their own fiber networks. He stated that by deploying its own fiber network, Gannett reduced its cost by at least 30%.

The Bank of America deployed a fiber network that allowed it to respond to the governmental regulations for data storage and replication requirements following the Sept. 11 terrorist attacks. According to Larry Schaeffer[9], senior vice president of network services at Bank of America, "Probably the most beneficial aspect of owning the network ourselves is the flexibility we have. We can turn up new services in a matter of hours or days instead of waiting months for a carrier to do it."

Network World Senior Editor Tim Greene[10] documented how the New York Presbyterian Hospital dropped its service provider in favor of implementing its own fiber-optic network. The situation that the New York Presbyterian Hospital faced was that it had exhausted its ATM based OC-3 network and was considering buying an OC-12 SONET network. However, the hospital realized that with the growing traffic demands that it had to support, the OC-12 network would be at exhaust by the time it was implemented.

New York Presbyterian Hospital decided to lease dark fiber from NYSERNET, a consortium that provides dark fiber to educational and research organizations in the New York area. NYSERNET is one of more than 100 entities that offer

6 Dark Fiber: Means to a Network, New Paradigm Resources Group, February 2002

7 Dark fiber: Businesses see the light, news.com.com/ Dark+fiber+Businesses+see+the+light/2100-1037_3-5557910.html

8 Ibid

9 Ibid

10 DWDM is the right RX for New York Presbyterian Hospital, http://www.networkworld.com/news/2005/051605-presbyterian.html

dark fiber services in the United States. In addition to consortiums like NYSERNET, this includes traditional service providers, alternative services providers, municipalities and utilities. The conventional wisdom is that the traditional service providers deliver the most reliable service. In point of fact, very often traditional service providers resell the services of other providers, many of whom provide very highly available services.

Akbar Kara is the chief technologist at the Texas Lonestar Education and Research Network. Kara, who is also the former director of core resources at the New York Presbyterian Hospital, commented that the cost of traditional network services increases as the speed of the service increases. For example, an OC-12 circuit costs more than an OC-3 circuit. However, with an optical network the cost of the service does not change as additional bandwidth is provisioned. Kara added that a customer-operated optical network provides virtually unlimited bandwidth and is capable of carrying storage and other network traffic. He also emphasized that when implemented in a ring architecture, an optical network is inherently resilient and affords the enterprise the greatest amount of flexibility.

The Denver Children's Hospital

Another organization that has deployed an optical network is The Children's Hospital in Denver. Albert Oriol is the hospital's IS program office director and data security officer. According to Oriol, the hospital has a dark fiber-based optical network between its two data centers, allowing them to split storage clusters across data centers. For additional disaster-recovery functionality, the storage-area network in each data center continually writes to a tape library in the other data center using the optical network for transport.

Oriol was bullish about the performance of the hospital's optical network. He stated that the network uptime "is as good as it gets" and that the quality of service has been fantastic. In addition, Oriol did not find that complexity was an issue relative to the ongoing management of an optical network. He commented that his organization has no more trouble managing an optical network than they do managing a TCP/IP network.

The hospital is in the process of deploying a new campus and is committed to using an optical network to service it. Because the geographic scope of the network would be greater than the current network, the hospital is leaning towards using a SONET service from a local provider.

Figure 2: SONET-based WAN

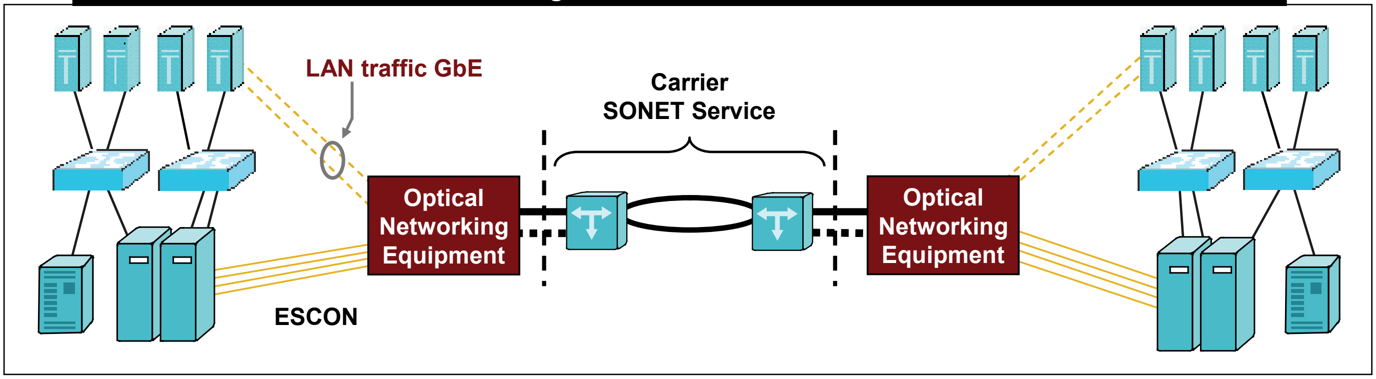

According to Leo Bodden[11], the current director of the hospitalSection 0 s network group, New York Presbyterian Hospital pays NYSERNET $542,000 per year as opposed to the $693,000 that it used to pay for its previous OC-3 network. To date, the hospital has lit 21 wavelengths, and they have the ability to add more bandwidth by merely increasing the speed of the lasers or adding more lasers.

It is important to realize that companies interested in deploying optical networks have more options than just dark fiber. Figure 2 illustrates an example of a fiber-optic network that many companies have deployed when the WAN spans a long distance. In this example, an IT organization acquires a SONET service from a carrier and deploys the equipment that would allow self management and provisioning of their own network.

The new WAN architecture

One of the enablers of the enterprise WANSection 0 s transition is that the robustness of Ethernet based transport networks has increased dramatically in the last few years. Another one of the enablers is the development of software-programmable hardware. This hardware enables cost-effective and timely modifications to WAN hardware.

One of the key issues that must be addressed by a network architecture is the role of routing. The traditional WAN makes heavy use of routing. However, while routed WANs have been around for a long time, they remain complex. In addition, for the type of requirements that were described earlier in this Special Report, a routed WAN can be very expensive.

Another shortfall of a routed WAN is that small increases in circuit utilization can lead to unacceptable levels of packet loss. Figure 3 depicts how the probability of packet loss increases as the circuit utilization increases. For example, if the circuit utilization increases from 35% to 50%, the probability of packet loss increases by a factor of 10, from 0.1% to 1%.

Figure 3 also includes some thresholds for packet loss that, when exceeded, the indicated application will not perform acceptably. For example, at 50% circuit utilization the probability of packet loss is 1%. If packet loss exceeds 1%, most users will find VoIP calls to be of unacceptable quality.

Figure 3 : Packet loss in a routed network

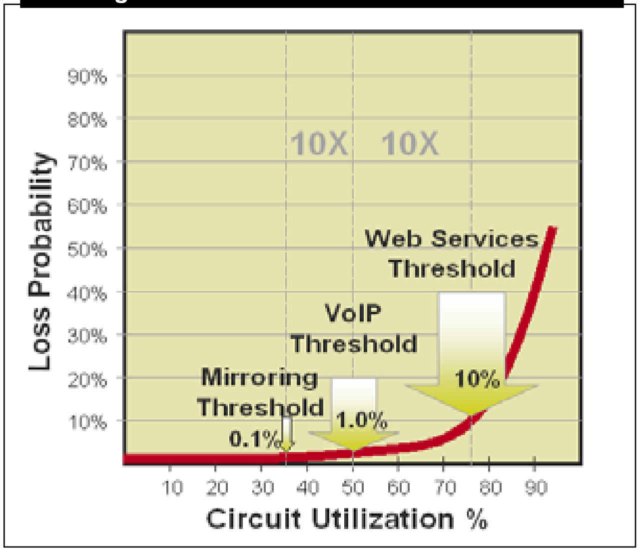

Another one of the key questions that must be answered by a network architecture is where to put the intelligence that is capable of performing extensive processing of packets. In a traditional WAN, this intelligence is embedded at the edge as well as throughout the core of the network. In particular, in the traditional WAN, the data path from the origin of the traffic to the destination is comprised of several routers, each of which inspect and do extensive processing of each packet as it transits the WAN. It is highly inefficient to perform extensive processing of the same packet several times. A more modern approach to network architecture is to position the intelligence at the edge of the network. Using this architectural approach, the packet is processed once at the edge of the network and then transmitted through the WAN with the absolute minimum of additional processing.

An example of this approach to network architecture is MPLS. In an MPLS network the first router in the data path assigns a label to each packet. A label is always 20 bits in length and is part of the 32-bit MPLS header. The routers

[11] Ibid

in the core of the network ignore the packetSection 0 s network layer header. Instead, when a packet arrives at one of these core routers, the only processing that occurs is to use the input port number and the label to perform a search of the forwarding table. When a match is determined, the forwarding component replaces the label and directs the packet to the outbound interface for transmission to the next hop in the data path.

One of the factors that has driven service providers to deploy MPLS in their networks is also relevant to the deployment of enterprise WANs. That factor is the desire on the part of service providers to minimize the complexity of their ongoing operations.

Performing extensive processing of packets is inherently complex. What this means from an architectural perspective is that the complexity of a network increases as the number of places in the network that perform extensive processing increases. As such, a network that only does extensive processing at the edge is less complex than a network that also performs extensive processing in the core. This added complexity has a direct impact on the IT organization. For example, the amount of time that it takes the IT organization to perform tasks such as troubleshooting increases notably if intelligence is distributed throughout the network. In addition, every time that either a new application gets deployed or an existing application gets modified, the IT organization has to tune the application to ensure that it runs well over the WAN. The difficulty of this task is increased dramatically if intelligence is distributed throughout the network.

To put complexity in context, the typical IT organization spends 70% of its resources on ongoing production and maintenance. This leaves a relatively small portion of the IT resources available to support new high-value initiatives. To combat this, most IT organizations are working to reduce the complexity of their networks. Implementing intelligence just at the edge of the network is a major step towards reducing complexity.

Another reason to avoid a network architecture that places intelligence in the core of the network is the blurring of the ISO seven-layer model. This approach binds the network layer with the middleware layer, which is highly risky as the major vendors have fundamentally different approaches to middleware and their products are still in a state of flux. As the middleware evolves, any work that was done to bind it to the network must be redone.

For example, consider MicrosoftSection 0 s .Net and SunSection 0 s Java 2 Platform Enterprise Edition (J2EE) both of which are intended to accelerate and simplify the development of business applications in general, and of Web services-based applications in particular. However, there are significant differences between the two approaches. J2EE supports just one language (Java) while .Net currently supports roughly 20 languages and the list is expanding. .Net only runs on Windows-based machines while Java runs on virtually any operating system. In the case of J2EE, the Java code is compiled into bytecode, which is executed in the Java Runtime Environment. .Net-supported languages get translated into Microsoft Intermediate Language, which the Common Language Runtime interprets and translates into a native executable[12].

[12] The great debate: .Net vs. J2EE, http://www.javaworld.com/javaworld/jw-03-2002/jw-0308-j2eenet.html

Table 3 summarizes the relative advantages of placing intelligence at the edge of the network vs. in the core.

Table 3 : Location of network intelligence

	Intelligence in the Core	Intelligence at the Edge
Core Switching	Layers 3 - 7	Layers 1 2
Prioritization	Packet contents	Header information
Processing Overhead	Significant	Minimal
Middleware	Tightly coupled	Decoupled
Impact on Applications	Increased complexity; More maintenance expense	Transparent
Security Risk	High	Low

Summary

The enterprise WAN is transitioning away from being merely a cost to be minimized, to a business asset that helps even the largest multi-site enterprises leverage agile IT for strategic advantage. One of the major factors driving this transition is the role that the WAN plays in supporting the ongoing spate of mergers and acquisitions, the increased importance of business continuity and disaster recovery, the continuing issuance of industry and governmental regulations, the deployment of enterprise-wide applications and the movement to a Web services model for application development. Other factors driving this transition include the need for a WAN with much higher availability than is provided by the traditional WAN, and one that can enable the deployment of new services in a more timely fashion.

Relative to being able to support emerging business requirements, the traditional IP-routed WAN has several shortcomings including complexity and cost. With a traditional IP-routed WAN, every increase in circuit capacity increases the monthly recurring charges. Given the rate of growth in WAN traffic, many companies have found this to be an urgent concern. Another shortcoming of an IP-routed WAN is packet loss. As was shown in this Special Report, a relatively small increase in circuit utilization can cause a 10-fold increase in packet loss and make business-critical applications such as VoIP perform unacceptably.

To transition the WAN to be highly available, controlled and agile, many companies are moving away from a traditional IP-routed WAN to an optical network. This transition is being driven in part by the fact that optical facilities are readily available and are increasingly cost effective. This transition is also being enabled by the fact that many companies are finding that optical networks are no more complex than a traditional IP-routed network.

The shift from the use of traditional network services to the use of an optical network fundamentally changes how IT organizations think about their WAN. When using traditional network services, each new application deployed consumes bandwidth and is a factor driving up the monthly cost. As a result, some companies are reluctant to add applications such as splitting server clusters into disparate data centers for disaster-recovery purposes. However, once an optical network is in place, applications like this can most likely be added without incurring any additional cost.

To successfully transition the WAN to support business initiatives, IT organizations need to develop a strategic plan that incorporates optical networking. This plan must include an analysis of the fiber options available, such as the use of dark fiber within metropolitan areas and SONET/SDH for longer distances. As part of a broader company-wide business plan, companies should acquire only a new facility if that facility is located in a site where it has access to one or more fiber routes.