

# DESIGNING NETWORKING SOLUTIONS FOR THE NINETIES: A NEW APPROACH

John A. Miller and Chung Ping Cheng

*John A. Miller is Executive Vice President of AT&T Paradyne, Largo, Florida; and Chung Ping Cheng is member of technical staff at AT&T Bell Laboratories, Middletown, New Jersey. Mr. Miller is directly responsible for the Data Networking Systems department, and manages the end-user DCE market, including product management and development for all AT&T Paradyne products. He has a B.S.E.E. from Gonzaga University, Spokane, Washington; and an M.S. and Ph.D. in electrical engineering from Stanford University, Palo Alto, California; and was AT&T-sponsored for the Senior Executive Program at MIT's Sloan School of Business, Cambridge, Massachusetts. He joined AT&T in 1966. Mr. Cheng is a systems engineer in the Accu-link® Networking Products Department, and is responsible for*

(continued on page 26)

In the last 10 years, information networking has become a major topic in telecommunications. The increasing need for flexible and efficient telecommunications, has made better networking solutions central to the communications manager's goals. This has become especially true since the introduction of high-performance, high-speed digital services, commonly known as T1. T1 tariffs were introduced in 1983, and allowed the network designer to digitally transmit information at 1.544 megabits per second (Mbps) over a wide geographic area. This paper describes the evolution of voice and data networking, and presents a new networking concept—using T1 technology—to migrate existing networks to a more robust and efficient architecture. Our focus is on using telecommunications services to create a core transport and switching network to move and manage information (i.e. voice, data, video, or image) over a wide area network (WAN).

## **Introduction: Historical Perspective**

In the early 1980s, analog private line (APL) and Dataphone® Digital Service (DDS) were the dominant network transmission vehicles. To accomplish data communication among geographically dispersed locations, modems and data service units (DSU) were designed to transport data through these facilities. Every circuit was either an individual point-to-point or multi-point connection. Thus, the concept of data networking was straightforward, but quite primitive by 1990's standards. By the mid-1980s, the T1 tariff had become more affordable and therefore useful in designing information networks. Many users found that by combining voice and data applications on the same T1 facilities, they could cost-justify implementing their own T1 backbone network. T1 networks were formed by connecting major sites with T1 facilities.

**Panel 1. Terms and Acronyms in This Paper**

APL	analog private line
ASDS	Accunet <sup>®</sup> Spectrum of Digital Services
BMS	bandwidth management service
BMS-E	bandwidth management service—extended
DCE	data-communications equipment
DDS	Dataphone <sup>®</sup> Digital Service
DSU	data service unit
CCR	customer controlled reconfiguration
CO	central office
CPE	customer premises equipment
IOC	inter-office channel
LSA	licensed space arrangement
M24	DS-0 multiplexing service
MUX	multiplexer
NMC	network management channel
NMS	network management system
PBX	private branch exchange
SDM	subrate data multiplexing
SDN	software defined network
WAN	wide area network

This was a pivotal time in information networking since T1 networking caused an evolution from individual private line designs to facility consolidation through T1.5 facilities such as Accunet<sup>®</sup> T1.5. Many vendors developed T1 multiplexers with sophisticated networking functions for operating a T1 backbone network. However, T1 backbone networking had a critical flaw: it did not address sites where T1 facilities could not be used for economic or availability reasons. Users were forced to maintain and manage their T1 backbone network separate from remote sites where subrate tail circuits had to be used.

In the late 1980s, several new networking services have been introduced by public carriers such as AT&T. These new Central Office (CO) based services let network designers follow a new concept called *Hybrid Networking*. The key strength of this solution is that it

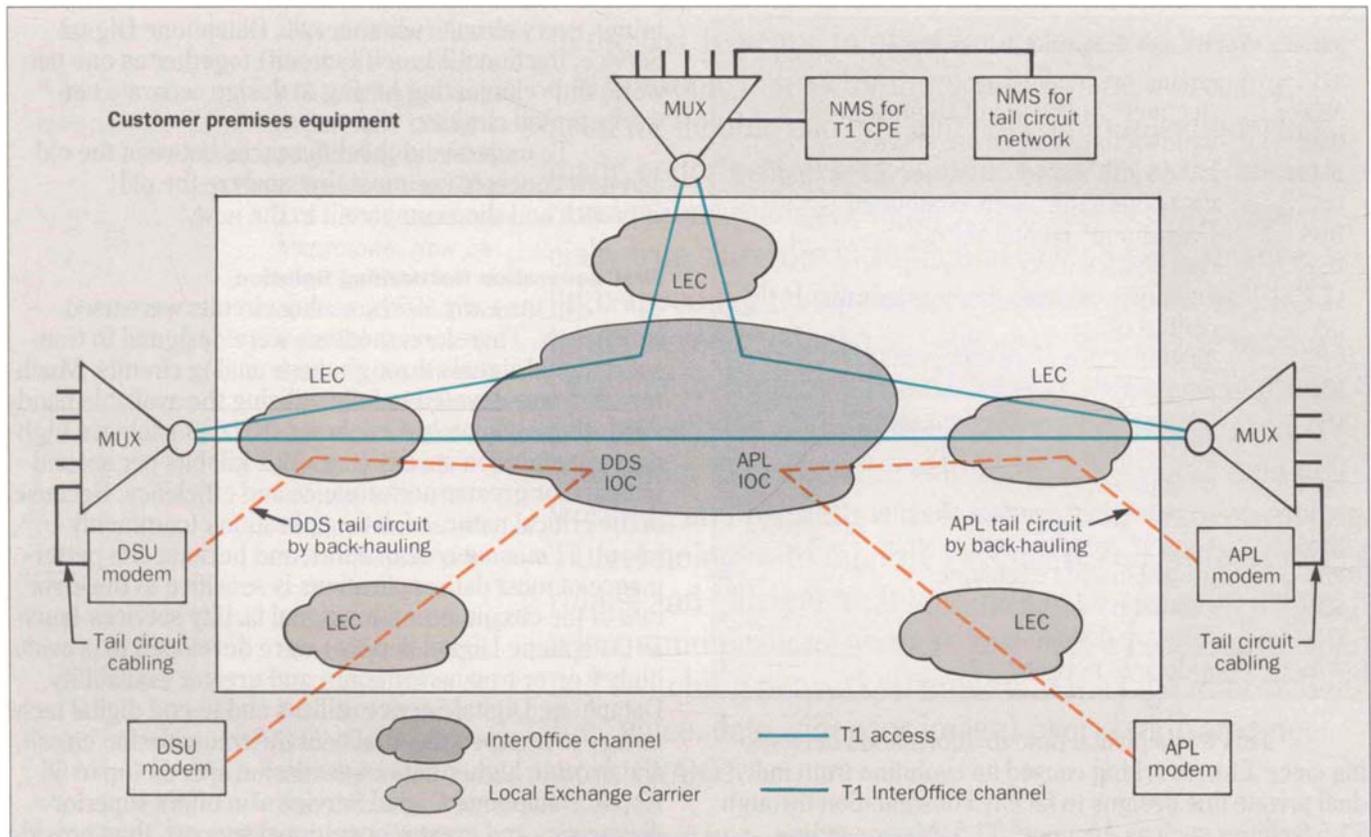
brings every circuit (whether APL, Dataphone Digital Service, fractional T1, or T1 circuit) together as one network, thus eliminating having to design separate networks for tail circuits.

To understand the differences between the old and new concepts, we must first analyze the old approach and then compare it to the new.

**First Generation Networking Solution**

In the early 1970s, analog circuits were used exclusively. Therefore, modems were designed to transport digital signals through these analog circuits. Much research was devoted to fully utilizing the available bandwidth [approximately 4 kilohertz (KHz)] to achieve high data transmission speeds [e.g., 19.2 kilobits per second (Kbps)] for greater performance and efficiency. Because of the critical nature of data applications (commonly known as *mission critical data*), and because the performance of most data applications is sensitive to the error rate of the circuit, premium digital facility services (such as Dataphone Digital Service) were developed to provide higher error-free performance and greater availability. Dataphone Digital Service utilizes end-to-end digital technology to improve the quality of the transmission circuit, and provide higher data transmission speeds (up to 56 Kbps). Dataphone Digital Service also offers superior diagnostics and greater operational support, thus providing a greater availability objective. Because Dataphone Digital Service was a premium service, the tariff was often two to three times higher than the APL tariff.

Both subrate services—APL and Dataphone Digital Service—began as point-to-point services, and users paid for them on a circuit-by-circuit basis. To reduce the facility charge, multi-point circuits were developed to allow multiple remote drops to share one subrate or 56K circuit. Because every circuit is independent, data networking was very simple. Integrating these services was done through the Network Management System (NMS) that provided centralized monitoring and control. To offer this functionality, the NMS needed a connection to

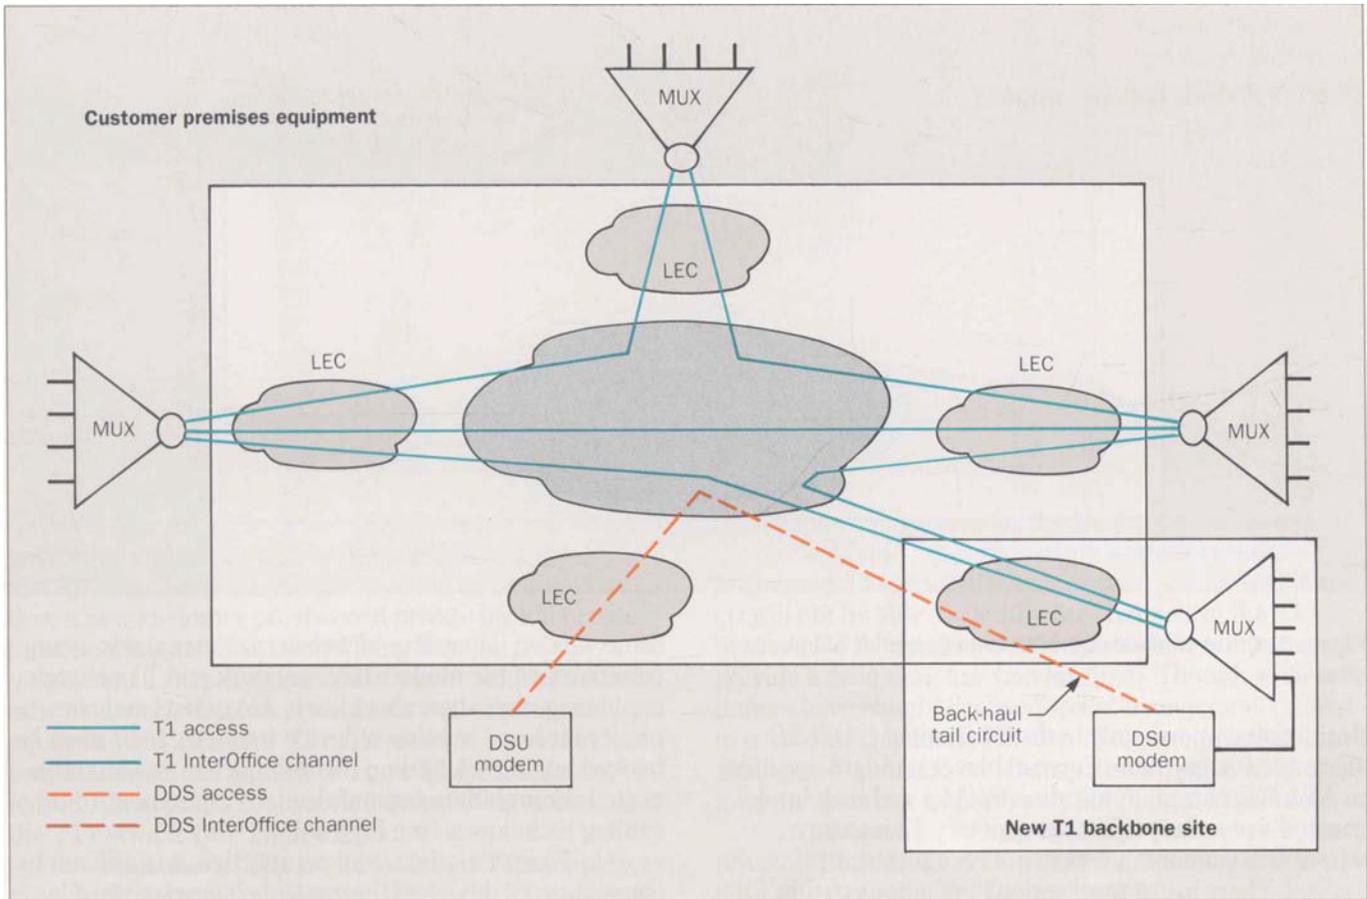


**Figure 1. Inefficiency of "backhauling" tail circuits**

every network device. Devices that had to be connected included: devices collocated with the NMS; devices at the remote end of the circuits; and devices that terminated circuits between two remote sites. Most vendors used a network management channel (NMC) for this connectivity. For APL, this was generally an independent, non-disruptive "side" channel. For Dataphone Digital Service, this either disrupted the primary channel or required higher bandwidth and more expensive (up to 20 percent more) local loops to offer non-disruptive network management functions.

### **Second Generation: T1 Backbone Networking**

By 1985, many users recognized the savings in a T1 facility. A point-to-point T1 facility could replace multiple analog or digital lines between two user sites. As a rule, the price of a T1 facility was lower than the price of no more than 10 digital or 20 analog lines. Typically, T1 networks were economically justified by the fact they replaced voice tie lines used for connecting PBX networks. At the same time, a T1 facility could easily replace subrate data circuits and still have enough bandwidth for future growth. The user could install new circuits easily and quickly by defining these circuits over their "bandwidth utility." Therefore, many users built proprietary T1

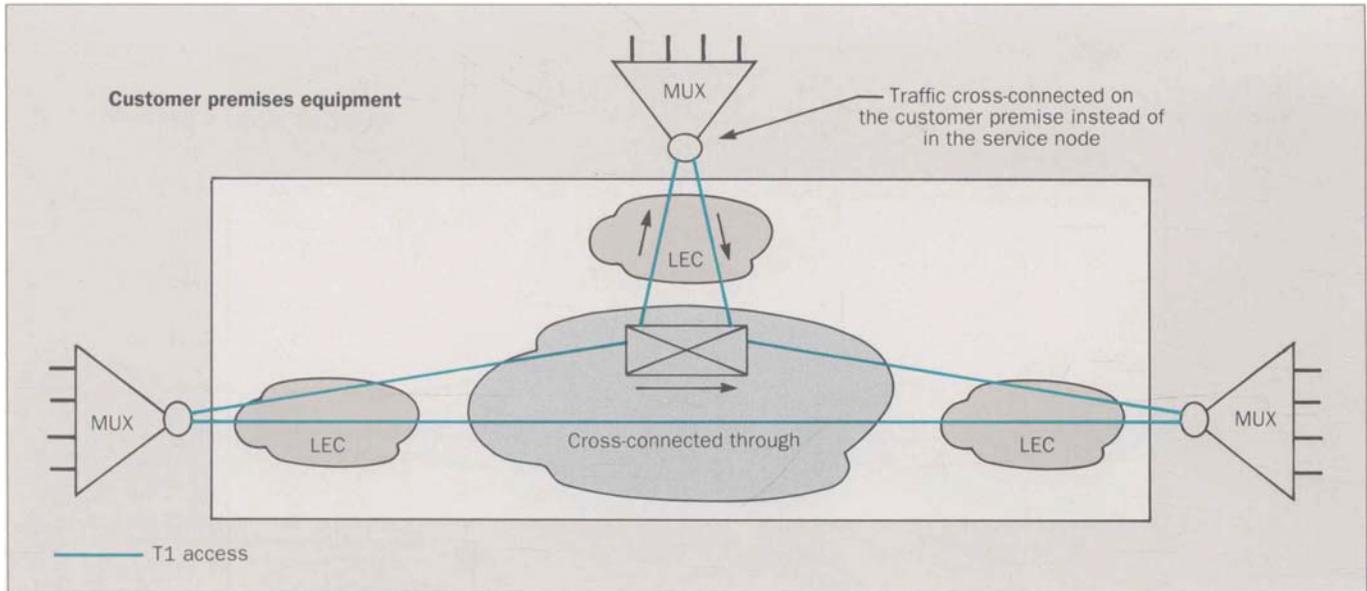


**Figure 2. Complexity of adding new sites using customer premise equipment and tail circuits**

backbone networks to connect their large locations. T1 backbone networking became the most popular telecommunications solution in the mid- to late 1980s.

At this time, there were no CO-based networking functions offered by the facilities providers. The only way users could build T1 backbone networks was to lease T1 facilities from the telephone company, buy expensive backbone T1 multiplexers from a customer

premise equipment (CPE) vendor, and become a “dedicated” private telephone company to serve their intracorporate communications needs. T1 backbone multiplexers were located at the customer premises, and supported all networking functions: efficient data multiplexing, voice compression, dynamic routing, automatic restoration, and network management functions such as single point of control, diagnostics, and performance monitoring. The T1 facility provided only an end-to-end physical layer connection between multiplexers. Because the COs did not support networking functions, many CPE vendors



**Figure 3. Cross-connected traffic is switched at backbone sites**

decided to conform only to the T1 framing [D4/ESF (Extended Superframe Format)] level standard specified in AT&T Technical Publication 62411,<sup>1</sup> and thus implemented a non-channelized proprietary T1 interface where DS-0 boundaries were not recognized.

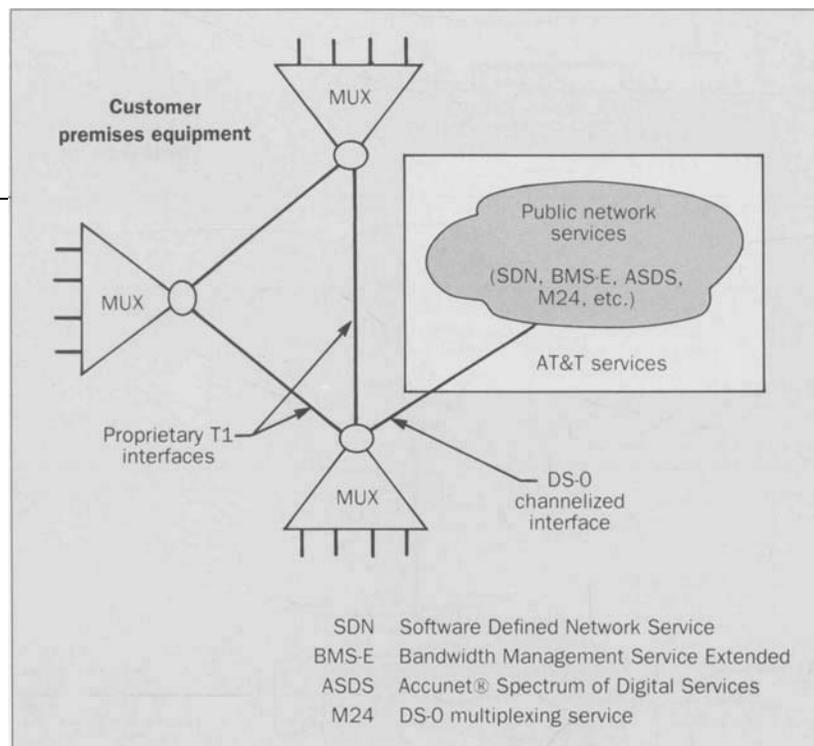
There are several serious limitations to this T1 backbone networking solution.

**Limitation 1.** The T1 backbone solution provides networking functionality only to locations on the T1 backbone network. For small sites where the traffic is not large enough to cost-justify T1 facilities, no T1 multiplexer will be used, and therefore, no networking functionality is available. The solution is to backhaul the traffic from the small sites to their nearest backbone site, thus creating a separate tail circuit network hanging off the backbone network. Unfortunately, the T1 multiplexer and the modem/DSU are usually not designed by the

same vendor. Therefore, network management is a concern because the modem/DSU network and T1 network are managed as separate entities. Also, costs and complexity increase because a device (modem/DSU) must be backed into a port card on the multiplexer and interconnected through data-communications equipment (DCE) cabling techniques (see Figure 1).

From a modem/DSU perspective, it is difficult to connect every device to the control channel of the modem/DSU network management system. From a T1 multiplexer perspective, there is no intelligent communication with the tail circuits. And from a system architecture perspective, reliability suffers because of increased single points of failure; availability suffers from reduced diagnostic capabilities, and the complexity of requiring multiple NMSS.

**Limitation 2.** Separate tail circuits lead to inflexible network configurations. Every tail circuit must be backhauled to the nearest backbone node. As shown in Figure 1, this configuration requires two local access



**Figure 4. Typical "gateway" depiction for hybrid networking**

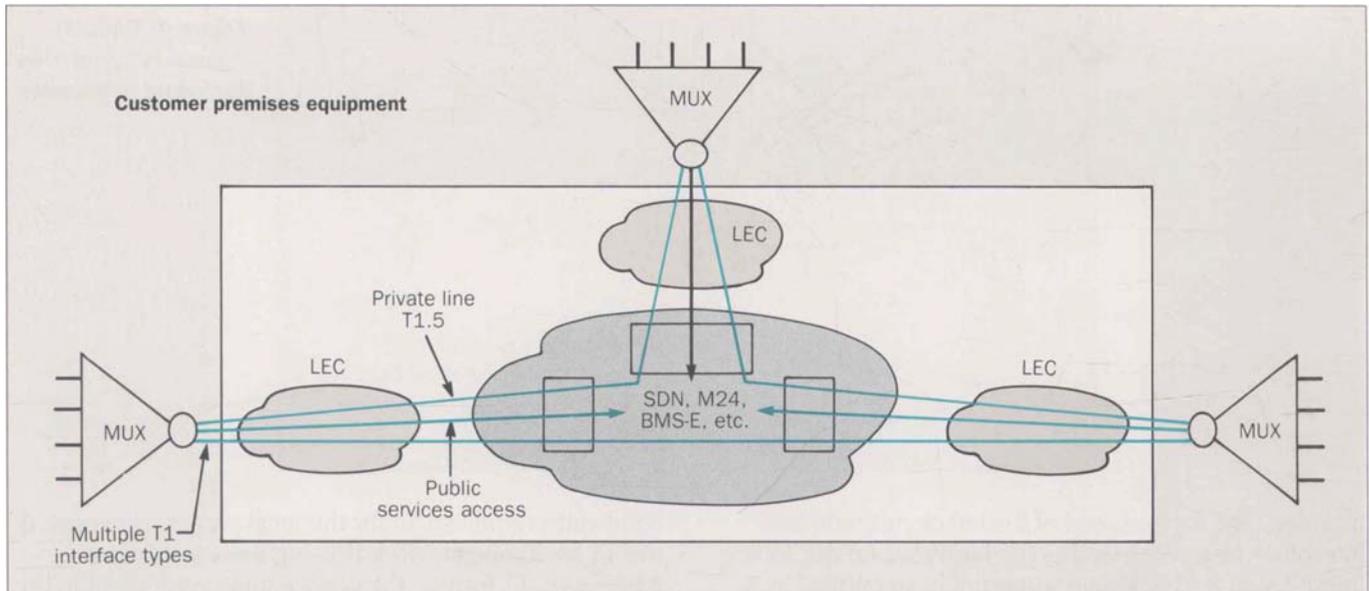
facilities (one for each end of the tail circuit) and one interoffice charge, as well as the hardware on the T1 multiplexer. If a backbone site is not in an optimal location, a new customer-provisioned private backbone site must be designated as the hubbing location for the nearby small sites. The customer will be limited to the office locations that house sophisticated communications equipment. Then a new T1 multiplexer is needed at that backbone site, and at least two new T1 facilities must be leased. These extra costs seriously limit the savings of the T1 network (see Figure 2).

**Limitation 3.** The premises-based T1 multiplexers provide the networking functionality in the T1 backbone network. The T1 multiplexer supports not only the traffic originating and terminating at that location, but also the traffic cross-connected through the node. The cross-connected traffic must enter the node on one T1 facility, and leave on another (see Figure 3). This means that cross-connected traffic requires twice as much bandwidth on the local access of that node.

A more natural place to perform traffic grooming is in the CO. A CO-based networking function can cross-connect the traffic without sending it on the local access to and from the premises multiplexer. This reduces the

bandwidth requirement for the local access. However, if the T1 backbone network is using *non-channelized proprietary* T1 format, the service node equipment in the CO will not be able to identify the channels in the T1 because DS-0 boundaries are kept, and therefore cannot provide a grooming function for them. Though some customers have sought licensed space arrangements (LSAs) to overcome this limitation, the level of negotiating complexity and technical issues may discourage both the customer and carrier from actively pursuing this solution.

**Limitation 4.** From a reliability point of view, the private T1 backbone networking solution may give a false sense of security. To provide automatic restoration capability, at least two different proprietary T1 paths must be at every backbone site. Because every T1 facility has three segments [one local access from each end point and one inter-office channel (IOC)], this suggests that every backbone node has at least two local access proprietary T1 lines. In most cases, these T1 lines are routed in the same cabling sheath; thus, it still creates a single point of failure. When the local access line fails, it disrupts not only traffic to that backbone node, but also traffic passing through it. This defeats the purpose of having two different paths to a node to avoid a single



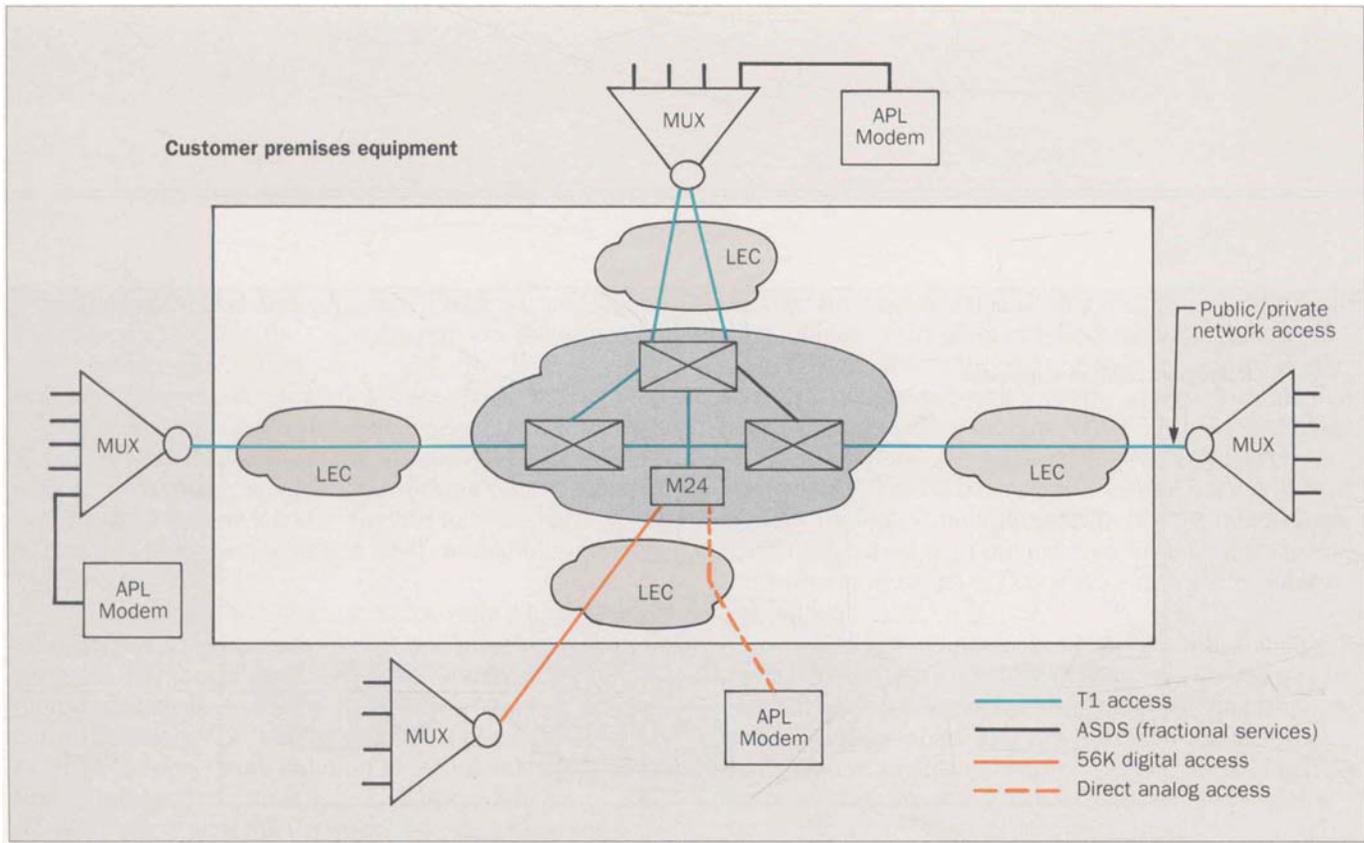
**Figure 5. Multiple T1 interfaces accessing public services**

point of failure. Unfortunately, the local access is also the segment where failures most often occur. Although total route diversity can be achieved on the local access, it is costly to implement, and is subject to availability.

**Limitation 5.** Besides needing information traffic between major user sites, many users also need to access network services like SDN (software defined network) and the Megacom<sup>®</sup> telecommunications service from those sites. Because most network services require a DS-0 channelized interface, a non-channelized proprietary T1 interface on the multiplexer cannot be used to access these services. Therefore, proprietary T1 multiplexers are forced to develop separate DS-0 channelized T1 interfaces. When the user must connect to another backbone node, a proprietary T1 interface is used. But when the user needs to connect to the public network, a DS-0 channelized T1 interface must be used (see Figure 4). The result of this approach is a less effective net-

work design, i.e., one with a potentially higher cost because both the DS-0 channelized and proprietary T1 interfaces are required.

One type of T1 interface accesses the private line T1 facility and uses a proprietary multiplexing format. Another type accesses other services by using a DS-0 channelized format. Figure 5 clearly shows different traffic is routed through the same service node location. The only difference between the two T1 facilities is that the traffic on the channelized T1 terminates in the public network, and the traffic on the proprietary T1 continues through it. Ideally, all traffic should be combined on the same access facility for maximum bandwidth flexibility and, more important, to avoid unnecessary leasing of additional T1 facilities. Proprietary architectures cannot be combined efficiently because two types of T1 interfaces are normally required. Moreover, because separate DS-0 channelized T1 interfaces from proprietary T1 multiplexers terminate directly in the CO, they typically do not support T1 multiplexer-based network management



functions. T1 multiplexer supervisory channels cannot ride over these specialized interfaces because they do not directly communicate with a T1 interface on the other end. This approach is usually promoted as a “gateway” to network services. However, it restricts network management and the overall efficiency of the network design.

**Limitation 6.** The final limitation of the T1 backbone networking approach is the T1 multiplexer’s proprietary nature. The T1 backbone network functions best only in a dedicated, private network. The networking capability of the proprietary T1 multiplexer decreases dramatically when the user wants to access network services. This limitation was acceptable when network based services were not available or needed. Over the past five years, many CO-based networking services were introduced, such as the M24 DS-0 multiplexing service, SDN, Megacom, BMS-E (Bandwidth Management Service-Extended), and ASDS (Accunet Spectrum of Digital Services). All these services support a more comprehensive networking approach and provide better cost economies for a complete networking solution. In the future, more

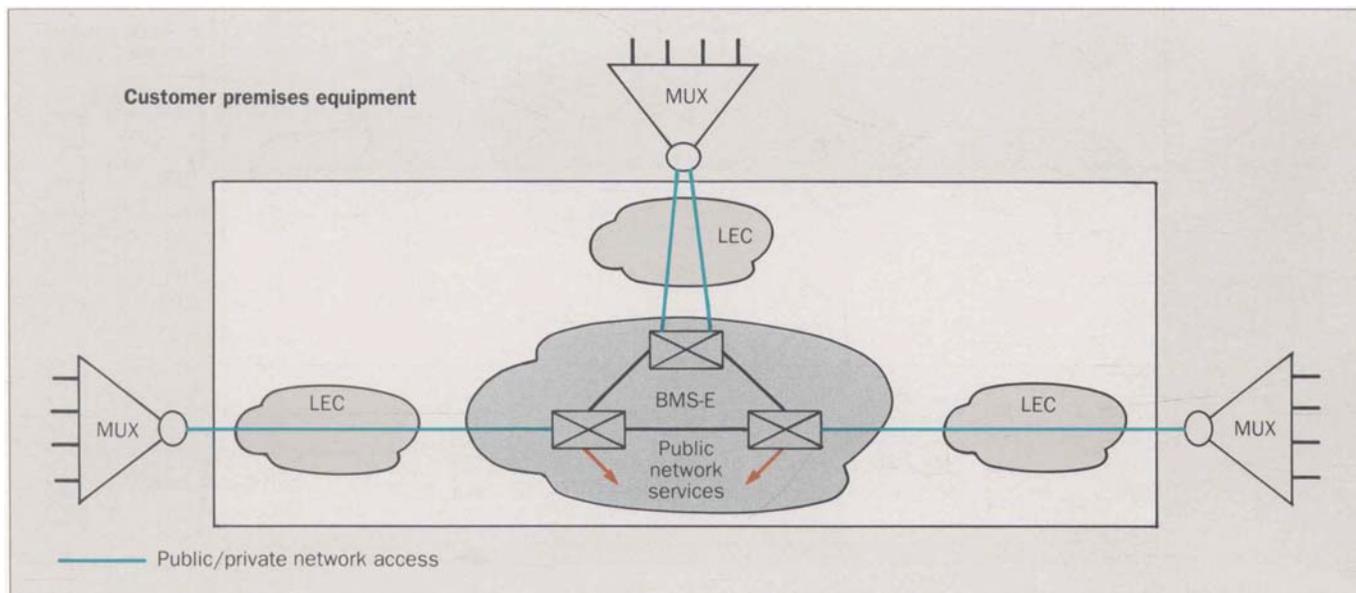
**Figure 6. Integrating public and private network access**

CO-based services will be introduced, and most will require DS-0 channelization compatibility. Therefore, the private T1 backbone networking concept is becoming obsolete. A hybrid networking concept combining both private and public networking capabilities is emerging because of lower overall network costs, improved reliability, and better network management.

#### **Hybrid Networking: The Design for the Nineties**

To completely realize the concepts of hybrid networking, a basic understanding of the T1 interface as defined by the CO, and contained in AT&T Technical Publication 62411, is required.

T1.5 represents the transmission of 1.544 Mbps. This rate equals 8,000 DS-1 frames (193 bits long) repeated each second, with each DS-1 frame containing one bit for framing overhead and 192 bits for user data. In a DS-0 channelized T1 interface, these 192 bits are further divided into 24 individual DS-0 channels. To



**Figure 7. Using access to BMS-E to perform networking**

support a channelized T1 interface, the T1 multiplexer must be able to handle each DS-0 independently and intelligently. This requires the multiplexer to perform networking functions under specific requirements, resulting in a more flexible and cost-effective network. Therefore, T1 multiplexers must adhere not only to the bit rate and frame level of compatibility, but also to the format level commonly known as DS-0 format.

With DS-0 level compatibility, many network services can be accessed, including M24, CCR (customer controlled reconfiguration), BMS-E, ASDS, SDN, and the Megacom system. Through these CO-based services, a more robust hybrid network configuration can be achieved. By efficiently combining CO- and premises-based networking capabilities, a hybrid network can result in better network design and lower costs by:

- Integrating public network access and private network access into one access facility to reduce local access

charges (see Figure 6).

- Using fractional T1 services to reduce IOC charges. Customers purchase T1 access circuits that carry the NxDS-0 fractional T1 circuit as well as other circuits that can be routed to services such as SDN (see Figure 6).
- Using M24 to bring 56 Kbps data circuits and analog voice circuits directly into the T1 network avoiding backhauling (see Figure 6).
- Using BMS (Bandwidth Manager Service), BMS-E, and CCR to provide fast provisioning and automatic restoration (see Figure 7).

When the T1 backbone networking concept is compared with the hybrid networking approach, the latter has the following advantages:

- It expands the network to include smaller sites with 56 Kbps data and analog circuits, thereby eliminating the need to treat those circuits separately.
- It increases network reliability by using the local access of a node to carry only the traffic terminating or

---

originating at that node. All pass-through traffic will be cross-connected in the CO, reducing access lines and recurring costs.

- It integrates public network access services (SDN, Megacom, BMS) and private network access services (i.e., private line traffic) onto the same T1 access facilities, while maintaining full networking functionality for the private network traffic (hybrid networking).
- It allows migration to future DS-0 compatible services (ISDN, SDDN).

Until now, the discussion has only addressed efficient methods to access public and private network services, and the ability to integrate modems and 56K digital circuits directly into the network by using network compatible services such as M24. What about efficiently extending the network solution to include user locations that require only a 19.2 Kbps, 9.6 Kbps, 4.8 Kbps, or 2.4 Kbps connection to the network? The proprietary networking solution *cannot directly connect these locations to the network*. They are handled as separate tail circuits that must be backhauled either to a 56 Kbps or T1 multiplexer site. In this case, the economic deficiencies and network management issues again present themselves.

One alternative is to use the CO-based M24 service to bring the site directly to the network through analog circuits. However, these circuits are treated as voice circuits. If an analog circuit is used to carry data traffic, the user bandwidth carried is typically between 1.2 Kbps and 19.2 Kbps, but because the M24 service only deals with DS-0's, it will still occupy a bandwidth of 64 Kbps in the network.

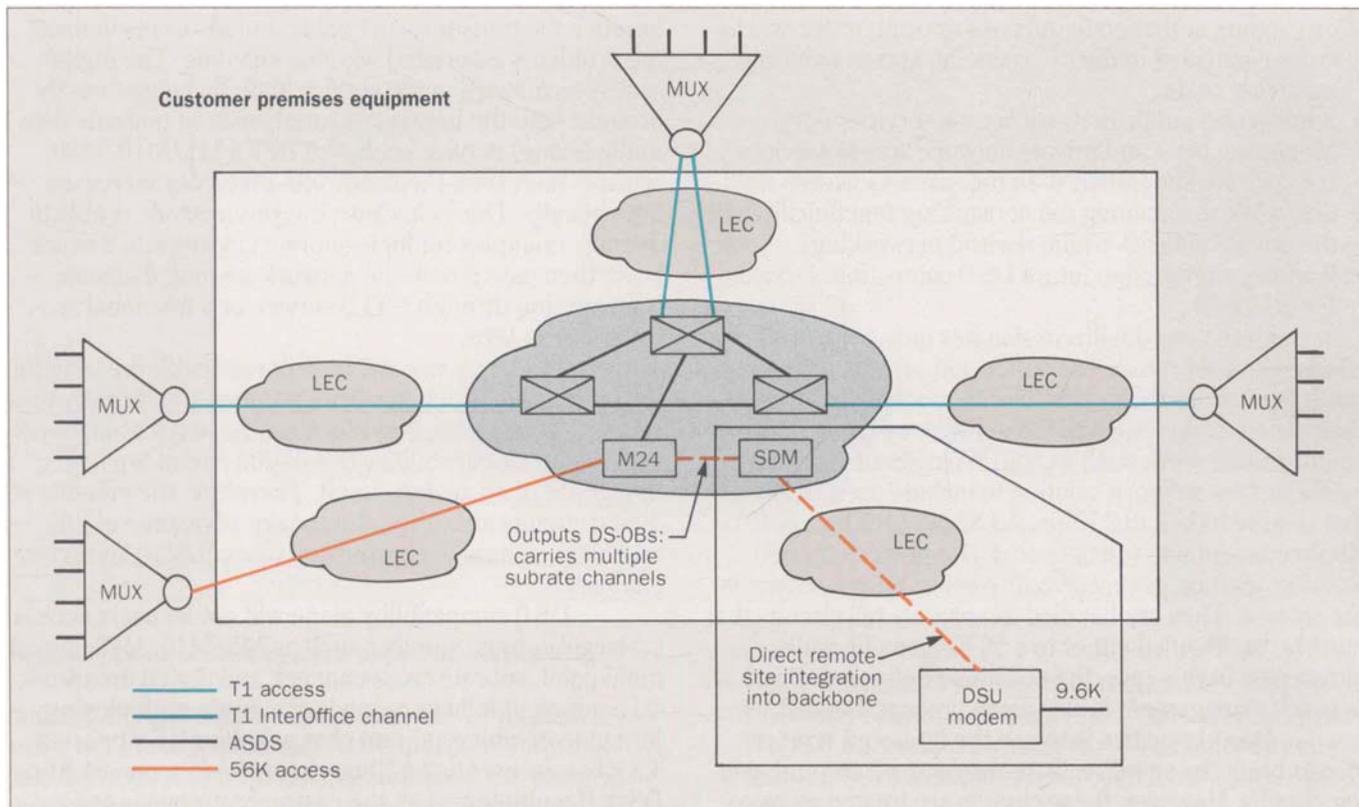
One way to address this inefficiency is to backhaul the analog circuit to the nearest multiplexer site, use a matching modem to demodulate the circuit back to the baseband data rate, and transport the circuit through the network as a digital data circuit (see Figure 1). But then the same backhauling and tail circuit issues arise. A more elegant solution is to migrate the analog access to digital access, and use public network services to directly integrate the sites into the network. This will not only

improve the transmission quality, but also will eliminate the problems associated with backhauling. The digital access terminates in the service node by being directly brought onto the network through an SDM (subrate data multiplexing) service such as AT&T's M5/M10/M20 service. With SDM, the bandwidth efficiency increases dramatically. This is because the service node is able to digitally multiplex multiple subrate circuits into a single DS-0, then pass it onto the network for long distance transmission through a T1.5 service or a fractional service such as ASDS.

The key economic benefit rests with the tariff for digital circuits, which has been reduced significantly to compete with analog circuits. Also, the service, support, and diagnostic capabilities of a digital circuit are better than those of an analog circuit. Therefore, the benefits of digital circuits are many. But to take advantage of this capability, a multiplexer must go beyond DS-0 level compatibility.

DS-0 compatibility alone will not let users access CO-based subrate services such as M5/M10/M20, digital multi-point, subrate cross-connect, and digital broadcast. CO equipment follows a standard subrate multiplexing format to combine subrate channels (five 9.6 Kbps, ten 4.8 Kbps, or twenty 2.4 Kbps channels) into one 64 Kbps DS-0. If multiplexers on the customer premises use a proprietary multiplexing format, the CO equipment will not be able to unbundle (i.e., demultiplex) the DS-0 into the original subrate circuits and provide networking services for them. This situation would make the premises equipment the only equipment that could switch subrate traffic to and from the customer site destinations. Once again, all subrate circuits would have to be backhauled to the premises multiplexers to gain access to those capabilities. This would result in lower reliability and higher bandwidth requirements on the local access facility.

To access CO-based subrate services, the premises multiplexer must follow the standard SDM format presented in AT&T Technical Reference 54075.<sup>2</sup> This level of compatibility will let the network designer

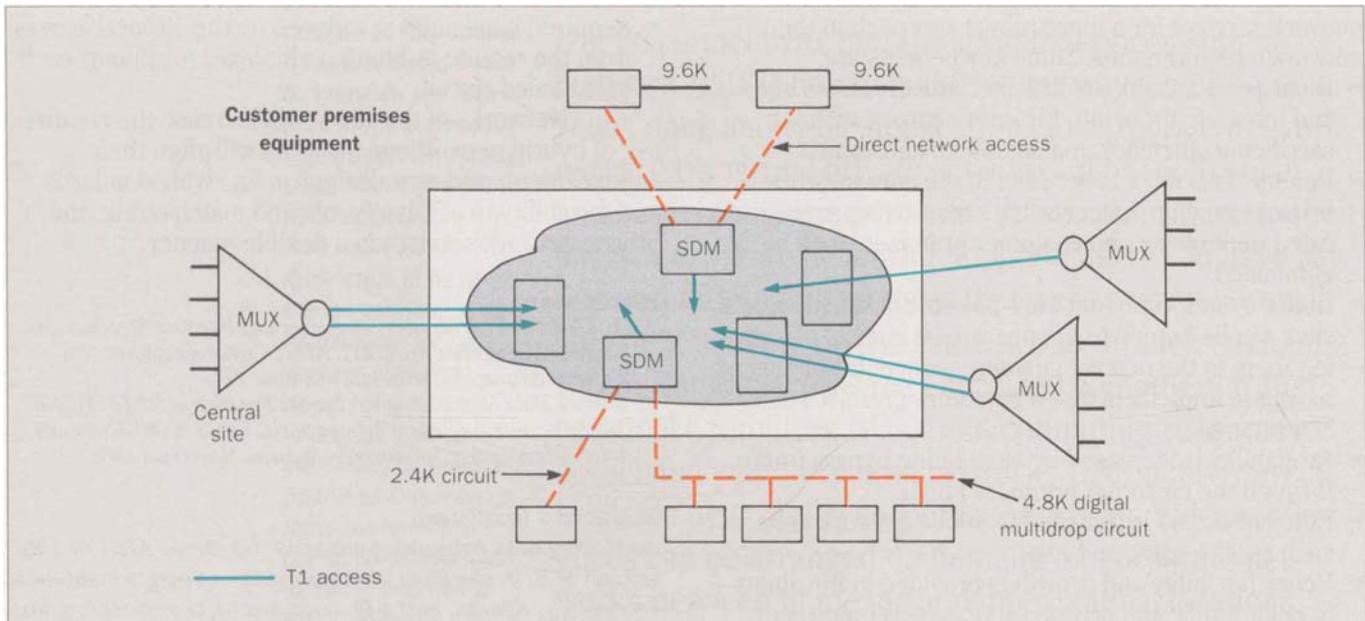


**Figure 8. Direct remote site integration through digital access**

integrate every user site into a single network, from the T1 backbone site to a small site with a single 2.4 Kbps data circuit. A T1 multiplexer at a central site would directly terminate traffic from a remote site with a single subrate access (e.g., 9.6 Kbps) by combining multiple subrates into DS-0B(s) [when a DS-0 uses the SDM format it is noted as a DS-0B], send them over a T1 facility to a CO where the service node equipment can identify each DS-0B and circuit in the DS-0B, perform the necessary networking grooming, and send them directly to

individual remote sites (see Figure 8).

SDM is an important tool for implementing an optimal hybrid network. It should be used where the requirements dictate direct subrate digital access. However, there may be other sites that dictate multiple subrate and voice circuits. The SDM format does not support voice applications and does not use bandwidth efficiently when mixing different data rates into one DS-0. Therefore, the SDM format should be used only for traffic terminating in the CO, or traffic needing CO networking functionality. An efficient proprietary multiplexing scheme is also needed to support traffic between premises T1 multiplexers. For optimum flexibility in



hybrid networking, the premises multiplexers must support both standard and proprietary multiplexing formats on the same T1 interface—without sacrificing network management functions—to achieve the best overall network design.

The advantage of this approach is that traffic from every user site, regardless of its speed or termination point, is included in the network. No tail circuits or backhauling are needed because all remote substrate access circuits will use CO-based multiplexing capability (M24, SDM) to connect to the T1 network. Sites requiring higher bandwidth use T1 or fractional T1 facilities (see Figure 9).

By being able to selectively combine proprietary multiplexer networking capabilities and public service offerings, the network design will have maximum flexibility and the best overall position for future growth. Depending on the traffic requirements at each site, an independent decision can be made whether to take

**Figure 9. Network design for the Nineties**

advantage of CO-based services or use the premises multiplexer capabilities. As traffic patterns change, the network configuration can be changed without major redesign. When more networking functions are introduced in the CO, the network can easily adopt to them and utilize the new capabilities.

#### **Summary**

Information networks have become vital business tools that directly influence a company's earnings. The changes caused by economic, political, technological, and environmental factors cause businesses to adjust to stay competitive. The information network must react quickly and efficiently to a business's requirements for new tactical and strategic moves that it needs to remain a force in its industry. Hybrid networking offers this flexibility through its ability to leverage public and private

network services for a more robust approach to the network's requirements. Some key benefits are:

- Businesses are able to directly connect remote sites and integrate them into the core network without sacrificing efficiency, management, and control.
- Remote sites are easily added to the core network because tandem concentrator sites—using sophisticated multiplexers on customer premises—can be eliminated.
- Better overall economics are possible because remote sites can be homed to the nearest CO instead of homing them to the nearest customer concentrator site, or having to implement a new customer premises concentrator site.
- Availability is increased by eliminating bypass traffic through the customer premises node.
- Efficient access is provided to public network services such as SDN, ASDS, and BMS.
- Better flexibility and growth is provided in the ability to tailor public and private services as requirements dictate.
- The requirement of one local access for every backhauled circuit is eliminated.
- Hardware on the remote hubbing multiplexer to support each backhauled circuit is eliminated.

- Required bandwidth is reduced on the T1 local access from the remote hubbing multiplexer to support each backhauled circuit.

Networks of the 1990s will increase the creative use of hybrid networking. Vendors will align their enhancements and new designs in line with standards and the ability to efficiently mix and match public and private network services in a flexible manner.

#### References

1. *ACCUNET® T1.5 Service: Description and Interface Specifications*, Technical Reference No. 62411, AT&T Customer Information Center, Indianapolis, Indiana, December 1988.
2. *Subrate Data Multiplexing—A Service Function of DATAPHONE® Digital Service*, Technical Reference No. 54075, AT&T Customer Information Center, Indianapolis, Indiana, November 1988.

Biographies (continued)

developing data networking products. He joined AT&T in 1985 with an M.S. in electrical engineering from Georgia Institute of Technology, Atlanta, and a Ph.D. in electrical engineering from Purdue University, West Lafayette, Indiana.

(Manuscript received April 8, 1991)