Internetworking LANs through WANs

In this report:

Synopsis

Editor's Note

The next logical step in local area network (LAN) technology evolution is interconnecting several LANs to form one large network—an internetwork or internet. This report examines the migration from LANs to WANs (wide area networks), an important trend in the computer/ communications industry today. It describes the protocols used on the internet, details the planning steps necessary to form and manage an internet, and highlights the products and services necessary to complete the task.

Report Highlights

LANs proliferated through the last decade. Their phenomenal growth initially resulted from their ability to interconnect terminals and workstations and share devices such as printers, file servers, and modems. LANs make economic sense. It costs less to provide fewer high-cost devices on a LAN, where they are shared by all network users, than to equip each workstation with its own modem, printer, and large disk drive. Users can share applications and files over LANs and communicate via electronic mail, thus, allowing organizations to finally achieve "true" office automation.

Once an organization has installed LANs and office automation technology in all of its offices, a natural extension is to link all corporate offices together. Using wide area networks (WANs), users can form an organization-wide enterprise internet. The enterprise internet provides the same functionality as an individual LAN, but on a larger scale—electronic mail, applications, and files can now move throughout the organization.

LAN-to-WAN Migration

Migration from a LAN or LANs to an enterprise internet is no small undertaking. Interconnecting LANs is often more challenging than building individual LANs, because the interconnections typically extend beyond a single premises and thus require

This report was developed exclusively for Datapro by L. Michael Sabo, Communications Network Consultant with SSDS, Inc. Mr. Sabo is currently developing a high-performance network architecture for interactive distributed real-time processing. Mr. Sabo holds a BS degree in Computer Science and an MS degree in Data Processing Management. He is a member of the Advisory Board for Datapro Network Management.

coordination between remote sites. Tuning an internet's performance is also difficult. A large portion of an internetwork is not owned by the user, unlike LANs. Regulated communications carriers provide the LAN interconnection circuits and, depending on the internet's scope, numerous carriers might be involved.

Users must carefully consider the networking protocols used on an internetwork. An internetwork might use any of several commercially available protocols, including IBM SNA, Digital DECnet, TCP/IP, OSI, and even Apple Appletalk. TCP/IP and the OSI protocols are not proprietary and, in a heterogeneous computing environment, should be considered over the proprietary protocols. Currently, TCP/IP is far more popular than OSI. TCP/IP protocols have been used longer and many more commercial products support internetworking via TCP/IP. Over the next several years, however, OSI will gain greater commercial acceptance.

Transmission Control Protocol/Internet Protocol (TCP/IP)

TCP/IP was originally developed for the U.S. military's ARPANET network. TCP/IP is actually a suite of protocols that is as successful in commercial applications as in military applications. TCP/IP works equally well over long haul facilities or LANs and is thus well suited to internetworking. It has been implemented on nearly every type of computer and represents a communications service's "lowest common denominator," providing peer-to-peer communications and projecting higher level services.

Transmission Control Protocol, as a formal transport layer, evolved from the original NCP transport layer of ARPANET, specifically addressing the problems of reliable communication through essentially unreliable subnetwork interfaces—such as linkages to datagram networks or packet radio environments. The Internet Protocol is the network layer of TCP/IP. It routes information through internets. TCP/IP, the "middle" of a full-layered protocol, seems an unlikely commercial success yet its popularity has grown steadily and is now the focus of a growing number of implementations.

TCP provides traditional virtual circuit, byte-stream-oriented communications services for programs (or users, through programs). Services include end-to-end flow control, error control, connection setup, and status exchange. The TCP, in other words, provides the communications that programs and users are accustomed to. Without it, programs could not be assured of send/receive order, data correctness, etc.

IP provides a datagram-oriented gateway service between subnetworks allowing hosts (typically using TCP) to access other hosts. IP does not enhance datagram reliability or integrity—it only lets them move from one subnetwork to another. IP also provides fragments and reassembles datagrams so that large IP datagrams can cross networks with small maximum packet sizes.

TCP is an OSI Transport Layer (Layer 4) protocol and provides several functions.

Data Transfer Support: provides a virtual circuit connecting the called and calling user, regardless of the lower level delivery system. (IP, in fact, is a datagram service.)

Error Checking: includes detecting lost, duplicate, and out-of-sequence information elements.

Flow Control: regulates the data rate from a sender so that the receiver's buffers are not filled, preventing the processing of incoming data.

Status and Synchronization Control: includes the ability to set up and break connections, mark significant points in the dialog, etc. This includes the ability to signal an unusual event (interrupt). The data flow for TCP is unlike other protocols in that it provides only one Transport Data Unit (TPDU) type, termed the TCP segment.

IP is an OSI Network Layer (Layer 3) function, normally responsible for routing and delivery. In a multi-subnet environment, an internet protocol is responsible for the following general communications areas.

Name Control and Translation: subnetwork members may not have the same naming conventions, and there may not be common naming control. IP assumes a internet-wide naming convention generated by combining a network ID assigned by the network administrator and a local host address.

Figure 1.
Remote Bridging

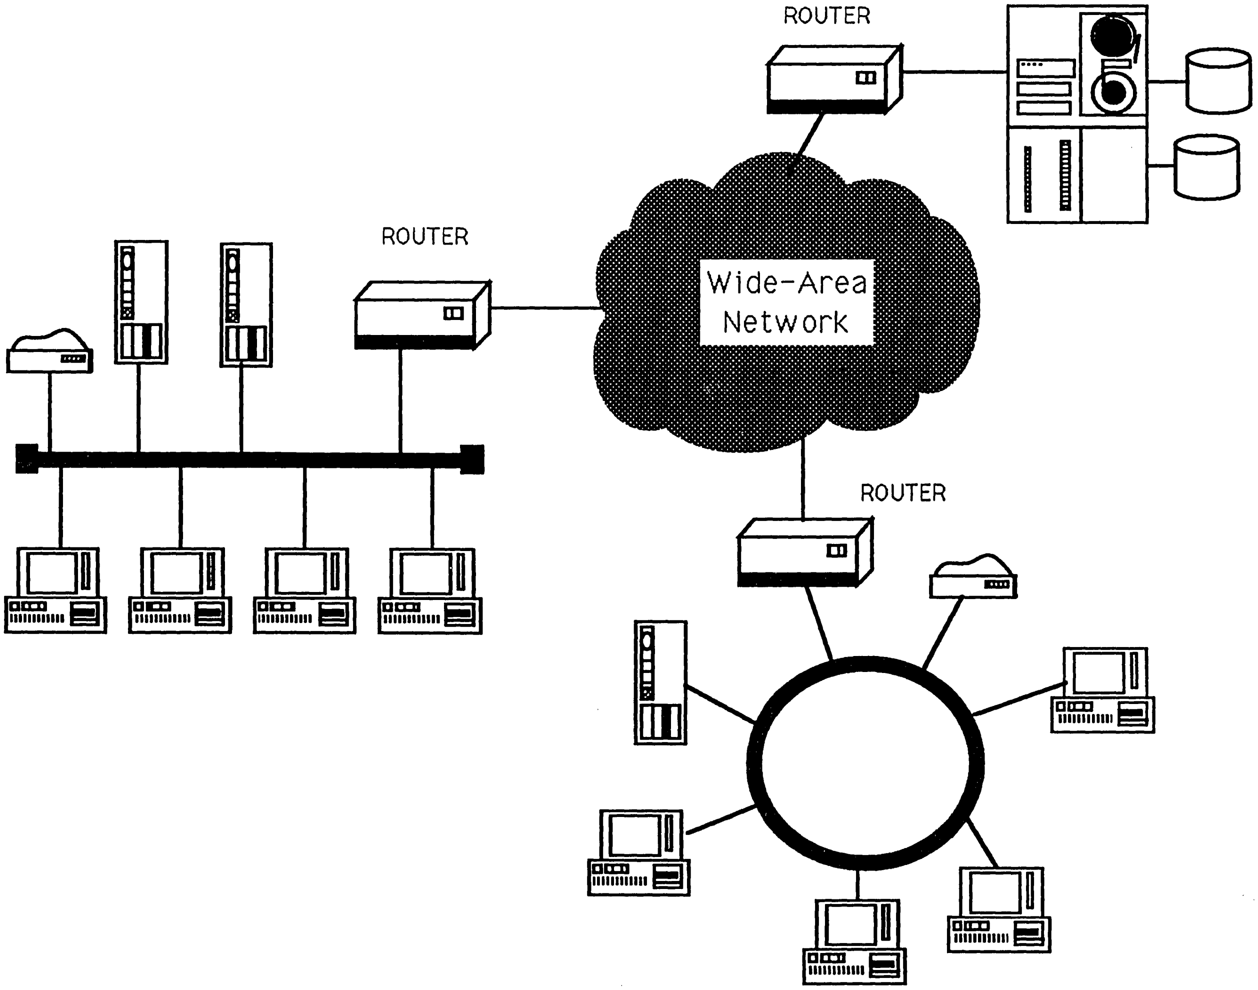

A relatively simple internetwork using remote bridging.

Status Translation and Communications: the result of internet operations must be communicated to the local user in an intelligible form. In addition, the operations that the user performs must return status conditions to alert the user to problems. IP provides four status messages: destination unreachable/invalid, timeout, parameter error, and redirect requested. Redirect requested means that another gateway has a shorter route to the destination.

Routing: IP provides Gateway-Gateway Protocol (GGP) message exchange to determine gateway and related host status. GGPs also verifies each interface's operability. "Probes" are sent to interfaces, and if several successive probes fail to return, the interface is marked down.

Management: IP allows operation center management of the internet environment, providing control and information gathering.

Fragmentation and Reassembly: IP routes large datagrams into networks that can't handle the datagrams' size via a technique called fragmentation. IP subdivides the original datagram into pieces small enough for transmission over the destination network and appends an additional header on the TCP information element. This header is a minimum of 20 bytes and can be 24 bytes or more.

Source and destination address are composed of a 1-byte network ID and a 3-byte host ID. The definition of subaddresses within a host is the TCP layer's responsibility.

When designing a TCP/IP internetwork, we strongly suggest that users request an official network number. While not mandatory, that simple step during the initial implementation can diminish future problems. To receive an official network number, write or call the DDN Network Information Center, SRI International, 333 Ravenswood Avenue, Menlo Park, CA 94025. The telephone number is (800) 235-3155. This service is free of charge and users are not required to be connected to the DDN to receive a number.

Open Systems Interconnection (OSI)

OSI's goal is much the same as that of TCP/IP—to provide an open and nonproprietary communications protocol suite. The OSI protocol suite grew out of the effort to define the now-popular seven-layer communications model.

Each layer in the OSI suite provides one or more services obtained using primitives. Each layer supports one or more protocols. Some layers, including the OSI network layer, have many protocols that support various networking requirements. These multiple protocols make internetworking in the OSI environment difficult. The two OSI network layer services are the connection-oriented network service (CONS) and the connectionless-mode network service (CLNS); they are incompatible.

These separate network layer services imply different transport layer services, which complicates the internetworking problem further. As an example, consider a LAN whose hosts use TP0/CONS for their transport and network layers. An internetwork is needed to another LAN, whose hosts support only TP4/CLNS for their transport and network layers. Because the network layers are not compatible a router cannot handle this situation. One solution is to allow only one network protocol within the enterprise internetwork. This restriction may not be feasible in all cases—however, in such situations, a transport layer store-and-forward relay is required. A transport layer relay essentially acts as a protocol converter, copying transport service primitives between each transport and network layer implementation.

Due to the above-mentioned problem with the OSI protocols, internetworking in the TCP/IP environment is often far easier.

Network Planning

Typically, LANs operate tens or hundreds of times faster than the WAN interconnecting them. This causes a severe throughput bottleneck and thus is a significant wide area internetworking issue that must be addressed.

To limit the bottleneck's effects, LAN managers should perform a requirements assessment to determine what devices are provided between LANs over the WAN. To determine required bandwidth, LAN managers must analyze each service as to the number of times the service will be used, and for what duration. If a special resource resides at remote site and significant traffic flows toward it, it might make sense to place the device at another location, or purchase another identical device at the local location.

Bandwidth requirements of the WAN interconnection circuits are dependent on a number of factors:

- response time requirements by application;
- amount of data traversing each circuit during peak utilization;
- amount of communications overhead; and
- type of internetworking equipment (routers or bridges).

Users must also consider the "highway effect." The highway effect means that demand for a new, high quality service will be greater than the aggregate demand for the service it replaces. Communications sizing should, therefore, consider reserve for the highway effect on top of growth reserve.

Systems Integrators

When developing an enterprise internetwork, users should consider hiring an experienced communications systems integrator. A good systems integrator knows the industry and the available products better than most internal communication groups. Communications systems integrators confront constant internetworking problems and can often handle these situations calmly and effectively. Equipment manufacturers cannot stay abreast of and obtain hands-on experience with *other* manufacturers' products. Systems integrators are an integral part of many successful internetworking projects.

Internetworking Equipment

The most popular internetworking devices are bridges and routers. Hybrid internetworking equipment, called brouters, are also appearing. Brouters provide both bridging and routing functions. Remote bridges are very simple to administer, while routers have more sophisticated filtering and security attributes.

Remote Bridges

Remote bridges link two or more geographically dispersed LANs over a wide area network, forming a LAN internetwork. A remote bridge operates as a store-and-forward device for LAN datagrams. With a bridge, stations connected to separate LANs can communicate as if they were both on the same LAN. Remote bridges, unlike local bridges, must operate in pairs to perform their internetworking function.

Remote bridges have two primary functions: filtering and forwarding. A remote bridge makes rapid decisions about a datagram's destination. It filters all packets on the LAN side and forwards only those addressed to remote locations.

Remote bridges learning or adaptive capabilities make routing information unnecessary. Each time a new sending or receiving address crosses the bridge, it is automatically added to the bridge's database. This lets users add or relocate devices on the internetwork without configuring a routing table.

A remote bridge uses data link layer addresses to make forwarding decisions. By observing the frames it receives, the remote bridge builds a database of all network devices that have transmitted. An entry in this database includes a physical destination address and the network generating the frame. All local traffic is filtered out, eliminating unnecessary traffic across the WAN. Since a remote bridge works exclusively at the MAC layer, it is not protocol specific.

The fundamental performance determinant of a remote bridge is available bandwidth.

Figure 2.
MAC-Layer Bridging

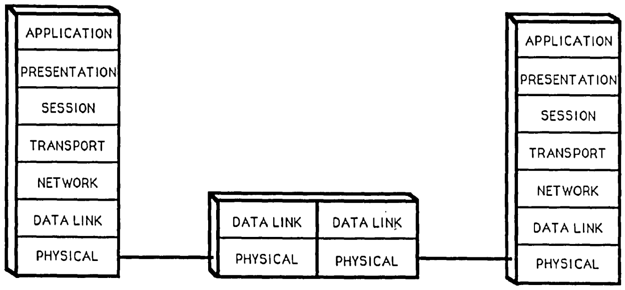

Remote bridges work exclusively from the MAC layer; they are not protocol specific.

Routers

Routers are somewhat more complex than bridges. Routers depend on the network layer address to perform the routing function. This is why routers are protocol-specific and can only route certain protocols.

Routers depend on routing protocols, such as Internal Gateway Protocol (IGP), to determine the most efficient path for packets to be forwarded to their destination. These protocols provide information such as shortest distance based on hop count or shortest delay.

Flow control is an integral routers function. If data comes in faster than a router can process it, the data must be queued. A router removes a packet from the input queue and examines its destination address. It then queries the routing table, which lists addresses of various nodes residing on the network. In addition, the routing table contains the paths between these nodes and how many hops (stops at intermediate routers) it will take to arrive at the destination. Most routers are configured to choose the best path based on the fewest hops.

Routing information is typically obtained and managed dynamically using gateway-to-gateway protocols. It should be noted that routers were once referred to as gateways. Gateways now are defined as devices that translate between protocols suites such as X.400 and the Simple Mail Transfer Protocol (SMTP) of the TCP/IP suite.

Routers that interconnect an enterprise internetwork to the large research internet must use the official routing protocol as defined in the Internet Request for Comments (RFC) 904. This routing protocol is termed an Exterior Gateway Protocol. Routers that operate within private internetworks

Figure 3.
Routing

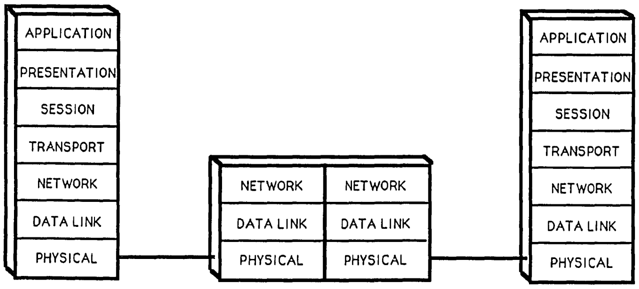

Routers depend on the network layer address to perform the routing function; they are protocol specific, and only certain protocols can be routed.

often use private routing protocols. These protocols are referred to as Interior Gateway Protocols (IGPs). Two popular IGPs are Routing Information Protocol (RIP) and HELLO.

RIP: requires routers to broadcast their routing tables to each of their neighbors periodically. This data contains reachability information, including hop count to specific destinations. It is used by the receiving router to make routing decisions.

HELLO: uses network delay as its routing metric instead of hop count. The HELLO protocol timestamps outgoing packets with the system clock value. Upon receipt of the packet, the receiver calculates the packet delay based on its system clock. Regular polling ensures system clock synchronization.

Remote Bridges versus Routers

The decision to bridge or route depends on protocols, network topology, and security requirements. Internetworking schemes often include a combination of bridges and routers. Due to remote bridges' simplicity, network managers should "test the waters" of internetworking initially with remote bridges. Additionally, remote bridges can forward packets much faster than routers because they do not process protocols past the data link layer. One of the strongest arguments for choosing routers over bridges is that bridges will copy broadcast packets until a "broadcast storm" overtakes the network. Routers are not susceptible to this effect, rejecting previously viewed broadcasts.

Routers can also provide flow control to reduce congestion and queueing delays. This capability is extremely important where low bandwidth circuits are used to interconnect LANs. Conversely, for smaller networks with high bandwidth circuits, remote bridges provide better performance.

Long Haul Facilities

One of the most critical determinations when designing a wide area internetwork is the necessary long haul circuit bandwidth. For applications where a few users will exchange minimal traffic, leased lines with 9.6K bps or 19.2K bps speeds will likely suffice. Internetworks typically require at least 56K bps, and many eventually migrate to T1.

56K bps Digital Service

It is possible to use conditioned leased lines providing 19.2K bps service as a long haul facility for extremely small internetworks, where sporadic electronic mail is the primary source of traffic. For an enterprise internetwork to provide a satisfactory degree of service, however, 56K bps is a practical minimum. 56K bps circuits are available as leased digital service or switched service. Switched service will accommodate those needing wide bandwidth relatively infrequently. Leased 56K bps digital service is more common. Carriers can provide 56K bps service over copper or fiber land lines, point-to-point microwave links, or satellite circuits.

T1 and Fractional T1 Service

T1 service is provided over leased facilities operating at 1.544M bps. T1 service can be acquired from many sources. Typically, local telcos provide service to their nearest central office. If the T1 circuit terminates inside the Local Access and Transport Area (LATA), only the local telco will be involved. If, however, users require interLATA T1 service, multiple telcos will be involved, complicating the process.

Fractional T1 falls between 56K bps and 1.544M bps service. With fractional T1, the telco still provides an entire T1 circuit to the local site; however, the user pays only for required bandwidth. Fractional T1 can be purchased in increments of 64K bps, up to the full 1.544M bps. Where available, it is an excellent alternative to multiple leased 56K bps circuits, when 56K bps is insufficient bandwidth and full T1 is not cost-effective.

X.25 Public Data Networks (PDNs)

Public data networks (PDNs) use packet switching technology to create a multipoint, or mesh topology. PDNs employ the X.25 protocol and can provide international service. PDN providers include Tymnet, Telenet, and AT&T. Because they employ a large number of packet switches and alternate routes between packet switches, PDNs are very reliable. A typical concern of PDNs is their end-to-end performance. The actual bandwidth is normally significantly less than the speed of the interface into the PDN.

To internetwork LANs via a PDN usually requires an X.25 gateway to convert LAN protocols to X.25.

Satellite

A satellite circuit consists of a transmitting earth station and a satellite with transponders that receive uplinked signals and amplify/rebroadcast them on a downlink to a receiving earth station.

Satellite solutions are attractive when extremely high bandwidth, low error rate, or wideband communications are required in a remote location where terrestrial service is not available.

Because communications satellites are parked in an orbit 22,300 miles above the earth, the round trip time from one earth station to another via the satellite is approximately ½ second. Unless this delay is compensated for, many protocols will believe the connection has been lost and end the session. Various techniques have been developed to compensate for this delay.

Essentially, satellite circuits provide a data pipe. As such, both bridges and routers will operate over these facilities.

Frame Relay with Fast Packet Switching

Frame relay is an emerging technology that will be employed for large internetworks of LANs. Internetworking equipment currently must provide a physical link interface for each long haul point-to-point circuit. As an internetwork grows, this situation can quickly become unmanageable. Using frame relay in conjunction with fast packet switching, only a single high-speed data link is required to attain complete interconnectivity, even if the network grows substantially. Fast packet, often referred to as Asynchronous Transfer Mode (ATM), is the basis for broadband ISDN and the IEEE 802.6 Metropolitan Area Network (MAN) standards.

Additionally, frame relay supports voice and data simultaneously. This technology is well suited to LAN interconnection due to its ability to handle bursty traffic patterns by allocating bandwidth as needed. Frame relay can be used with most commercially available routers and bridges to provide wide area bandwidth on demand. The frame relay standard is poised to overtake X.25 over the next several years, and is currently being standardized through a coordinated effort between CCITT and ANSI.

Network Management

Typically, LAN internetworks lack an end-to-end administration entity. This can make the administration of the internet very difficult. In large internets, special attention must be paid to network management. Often, finger-pointing between LAN administrators and long haul providers occurs because of lacking network management tools and an organization to handle the problem. In an internetworking situation, it is typical to find a LAN administrator who handles problems internal to the LAN and a WAN administrator who handles problems concerning the LAN interconnection facilities.

Administrators can use new integrated network management systems that support open standard network management protocols, such as the Simple Network Management Protocol (SNMP). Users with SNMP-based Network Management Stations (NMSs) can manage both LAN and WAN facilities.

SNMP is based on the manager-agent paradigm. An agent is software operating the network device. It responds to queries for information, which it gathers in its Management Information Base (MIB), from the manager. The manager software runs on the NMS. An SNMP manager can perform three operations with regards to agent interaction: it can get, get-next, or set specific variables in the agent's MIB. Correspondingly, an agent responds to the manager requests and, in the event of a catastrophic event, issues an unsolicited

Figure 4.
Long Haul Facilities

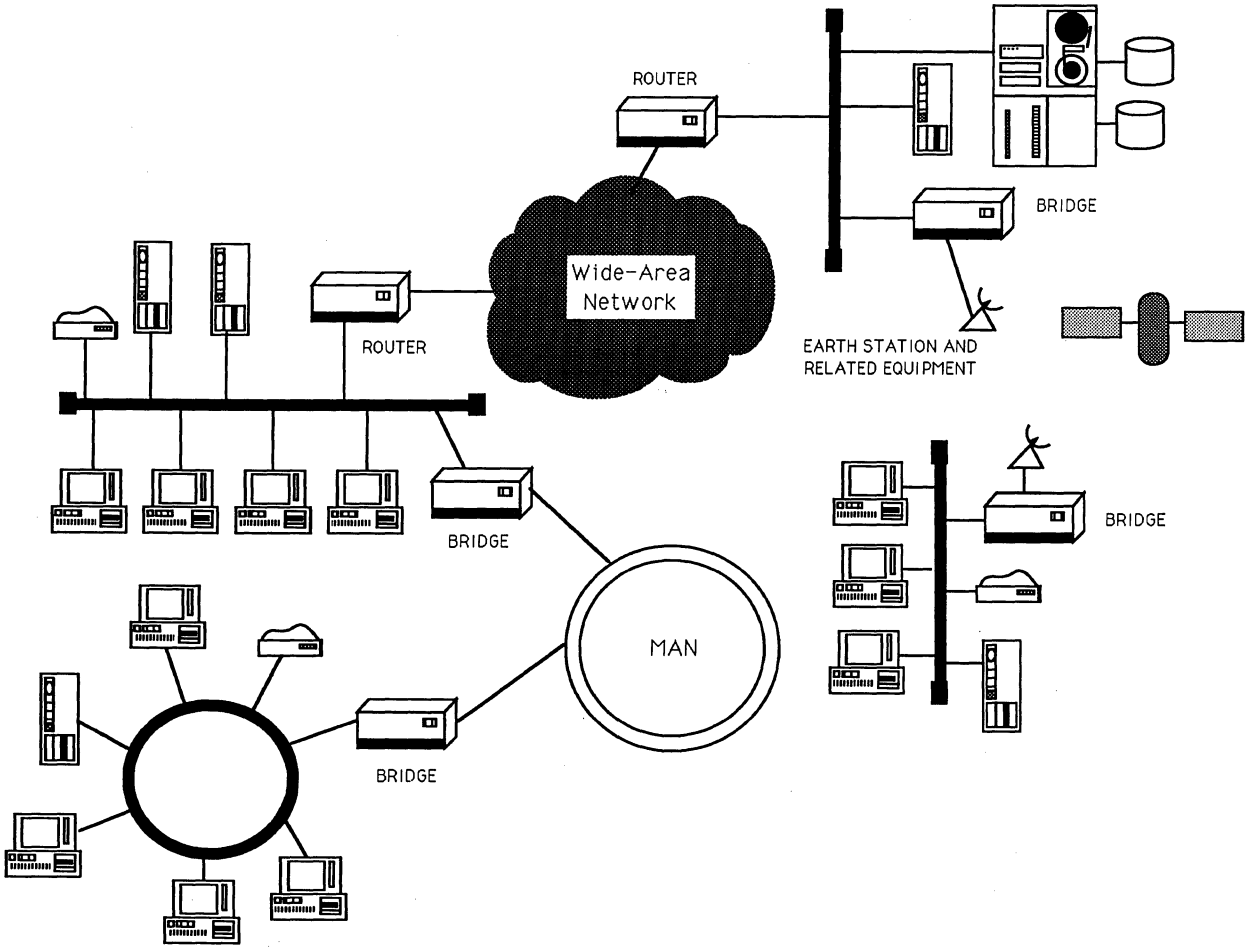

A complex internetwork with multiple types of long haul communications facilities, linked via bridges and routers.

trap message. SNMP is truly a simple network management standard, and for that reason it has found wide acceptance in the networking community. As internetworking migrates to OSI, however, the OSI network management solutions will eclipse SNMP.

OSI-Based Network Management

OSI network management is a robust set of standards and draft standards that address the entire scope of managing large open systems. In many respects, OSI-based network management and SNMP are similar; they both use the request/response, and are based upon the concept of an MIB. The OSI network management includes a richer set of service primitives, however, including request for an action to be taken, confirmed responses on demand, and the ability to create or delete managed objects.

OSI has specified the Common Management Information Protocol (CMIP) and Common Management Information Services (CMIS) as the international network management standard interfaces. CMIS, as defined by OSI, is the software services interface, whereas CMIP is the actual protocol mechanism implemented to access the CMIS services.

OSI is also standardizing system management functions in Specific Management Functional Areas (SMFAs), which include Configuration, Fault, Performance, Security, and Accounting SMFAs.

Trends

Increasingly, LANs are used as a logical computer bus. The aggregate of services over this logical computer bus forms a logical computer system. In this scenario, the overall performance of a distributed system is achieved through intense parallel processing, distributed through workstations and servers. However, users can begin to gain the benefits of centralized control. This trend will extend over wide area networks in the future, as long haul circuit bandwidth increases and becomes cost-effective. Network file servers, client hosts, and diskless server workstations form the foundation for this trend. This technology will drive enterprise internetworking to a new level. Users will enjoy identical interfaces for all systems, regardless of location on the network or system hardware.

Summary

Internetworked LANs will proliferate throughout the 1990s. All of the necessary communications components, including protocols, bridges, routers, and communications facilities, are available. TCP/IP is the most popular internetworking protocol; however, OSI is beginning to gain momentum and should reach high gear in the next several years. An internetwork presents many management problems not experienced in a strict LAN environment. We highly recommend an integrated network management concept to address these problems. ■