UNIFIED NETWORK MANAGEMENT FROM AT&T

Roberta S. Cohen, Hsin K. Kan, and Raymond J. Pennotti

Roberta S. Cohen *is head of the Unified Network Management Department of AT&T Bell Laboratories in Holmdel, New Jersey. Ms. Cohen joined AT&T in 1979. Her department is responsible for the Unified Network Management Architecture and network management protocol definition and specification. She has an A.B., an M.S., and a Ph.D. in sociology from the University of Illinois—Champaign-Urbana.* ***Hsin K. Kan*** *is head of the Data Network Systems Department of Bell Laboratories in the Red Hill facility, Middletown, New Jersey. Mr. Kan's department is responsible for the system architecture, development, and integration of new integrated UNMA products and services. He joined the company in 1976. Mr. Kan has a B.S.E.E. from National Chiao Tung University, and an S.M. in electrical*
(continued on page 136)

This article describes AT&T's *Unifed Network Management Architecture* (UNMA), a blueprint for common, end-to-end network management of complex data, voice, and combined networks using AT&T and other vendor's equipment and services. In this article, we introduce the target network management functional areas and present details of the overall architecture, including the *network management protocol* that serves as the standard interface among the interconnected management systems. We present examples of element management system functions, including configuration and fault management for modems and multiplexers, and we show how, by coordinating a network-wide configuration database with data from each element management system, a network management integrator can provide superior, end-to-end network management capabilities.

AT&T's Approach to Network Management

Communications networks in the U.S. today include three distinct networking domains: the customer's premises, the local exchange network, and the interexchange network. Customers have both voice and data equipment on premises; they use a local exchange carrier (LEC) for services and connections between local-access transport areas; and their communications cross the country through one or more interexchange networks. (Panel 1 is a list of terms and acronyms used in this paper.) Typically, these customer networks include equipment and services from many vendors.

To manage the equipment, services, and all other parts of the network, network operators rely on systems that we call *element management systems* (EMSs). Each EMS manages one or more network elements. Network elements can be equipment (such as modems, switches, and connecting media) or logical entities (such as packet ser-

Panel 1. Terms and Acronyms in This Paper

ACSE	association control service elements
ANSI	American National Standards Institute
AOE	application operating environment
APL	analog private line
ATR	automatic trouble reporting
CCITT	International Telegraph and Telephone Consultative Committee
CCR	commitment, concurrency, and recovery
CMISE	common management information service elements
DBU	dial back-up unit
DDD	direct distance dial
DIS	draft international standard (ISO)
DP	draft proposal (ISO)
DSU	data service unit
DTE	data terminating equipment
EMS	element management system
FA	facility fault alarm
FTAM	file transfer, access, and management
ISDN	Integrated Services Digital Network
ISO	International Organization for Standardization
LAN	local-area network
LEC	local exchange carrier
NMP	network management protocol
NPU	network polling unit
OSI	Open Systems Interconnection
PC	personal computer
ROSE	remote operations service elements
SQL	structured query language
UNMA	Unified Network Management Architecture

vices, circuits, and channels). AT&T and a variety of vendors offer network elements as well as EMSs. EMSs are found at the customer's premises, in the local exchange carrier environment, and inside the interexchange network. To unify and integrate network management across the networking domains and across these many EMSs, AT&T's UNMA uses a standardized *network management protocol* (NMP). The NMP provides a common means of communicating the information and instructions needed for integrated, end-to-end, multidomain network management among the systems that it manages.

Using NMP is not enough, however, to ensure integrated, end-to-end management of a network. The management systems themselves must provide integrated, unifying functionality. UNMA calls for an *integrating* network management system, a special kind of EMS, that acts as the hub for such functioning. (See Figure 1.) The integrating system provides its users with a *unified user interface* to the network management functions of the integrator as well as to the EMS with which it connects, and gives the network administrator value-added functions for end-to-end network management.

AT&T's commitment is to:

- Establish a network management architecture that integrates the management of multiple voice and data products and services.
- Provide end-to-end network management across customer premises equipment, LEC, and interexchange networks and services, including AT&T's and other vendors' equipment in all three domains.
- Build on the rich base of individual element management systems by including them in the architecture and adding an integrating system.
- Use international standards for network management interfaces between customer premises and regulated network services as well as for the interfaces between individual element management systems and the integrating system.
- Build on our operations experience to define an architecture based on high performance at a good price, a rich feature set, and intelligent analysis with expert systems.
- Establish priorities for features and supported configurations by working closely with key customers to build initial products and services.
- Work with vendors who are interested in meeting the standard interface.

Network Management Functions. The network man-

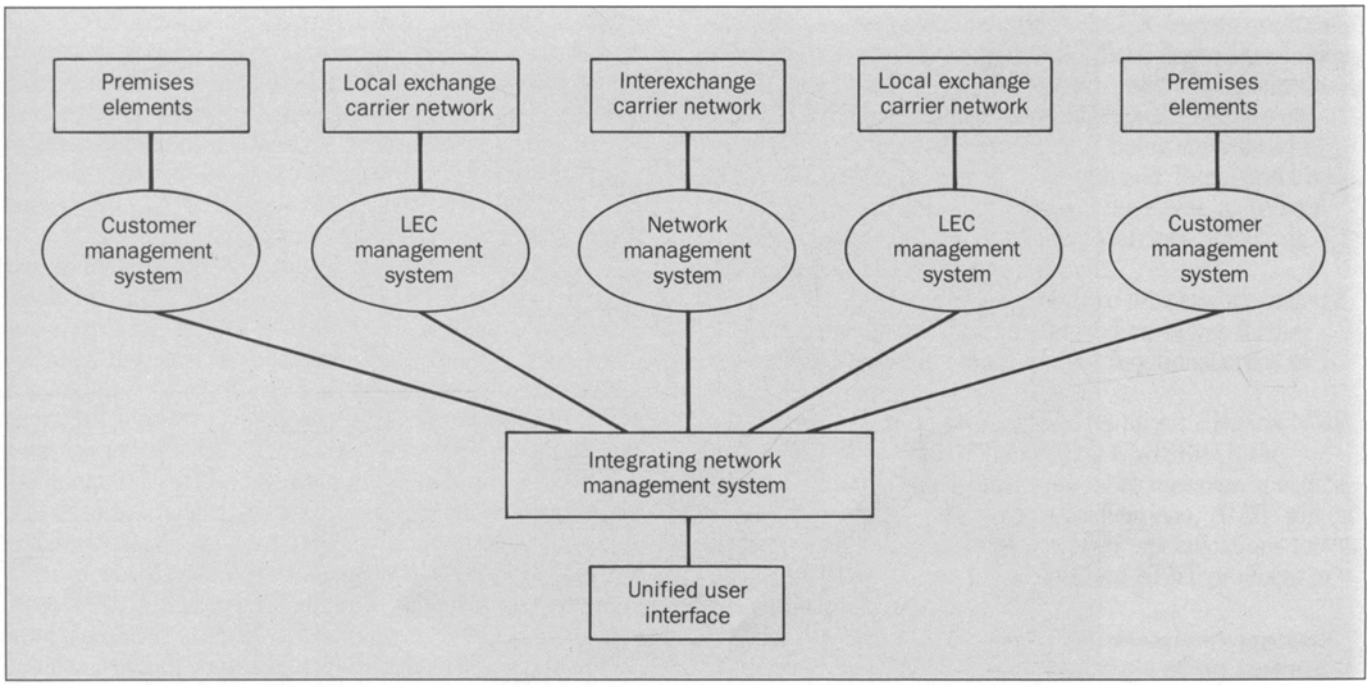

agement functions to be included in the UNMA can be grouped into six major categories:

- *Configuration and name management* includes inventory management, configuration management, change management and provisioning, and directory management. This means tracking assets; managing the active network configuration, including "adds and moves"; and managing the information needed to manage a network, such as organizational directories and vendor service contracts.
- *Fault management* includes network status supervision, end-to-end and segment testing, diagnosis of problems and repair, back-up and reconfiguration, and trouble tracking. These functions allow a network manager to identify a condition affecting service, isolate the problem through tests and other diagnostics, reconfigure around the problem while it is fixed, and track the trouble, on-line, with a historical status recording capability.

Figure 1. AT&T end-to-end network management.

- *Performance management* involves performance definition and monitoring as well as identifying trends and thresholds. These functions support the evaluation of network resources and overall performance by measuring network activity, usage, and availability.
- *Accounting management* supports budgeting, bill-back, and bill verification activities for network managers.
- *Security management* includes the activities needed to establish and maintain network access security and partition access by authorization level, function, geography, time of day, or other criteria.
- *Network planning* covers the areas of capacity, contingency, and strategic planning. These functions allow managers to improve system capacity and cost while planning for future events and growth.

Each network management function requires that

Figure 2. AT&T network management protocol.

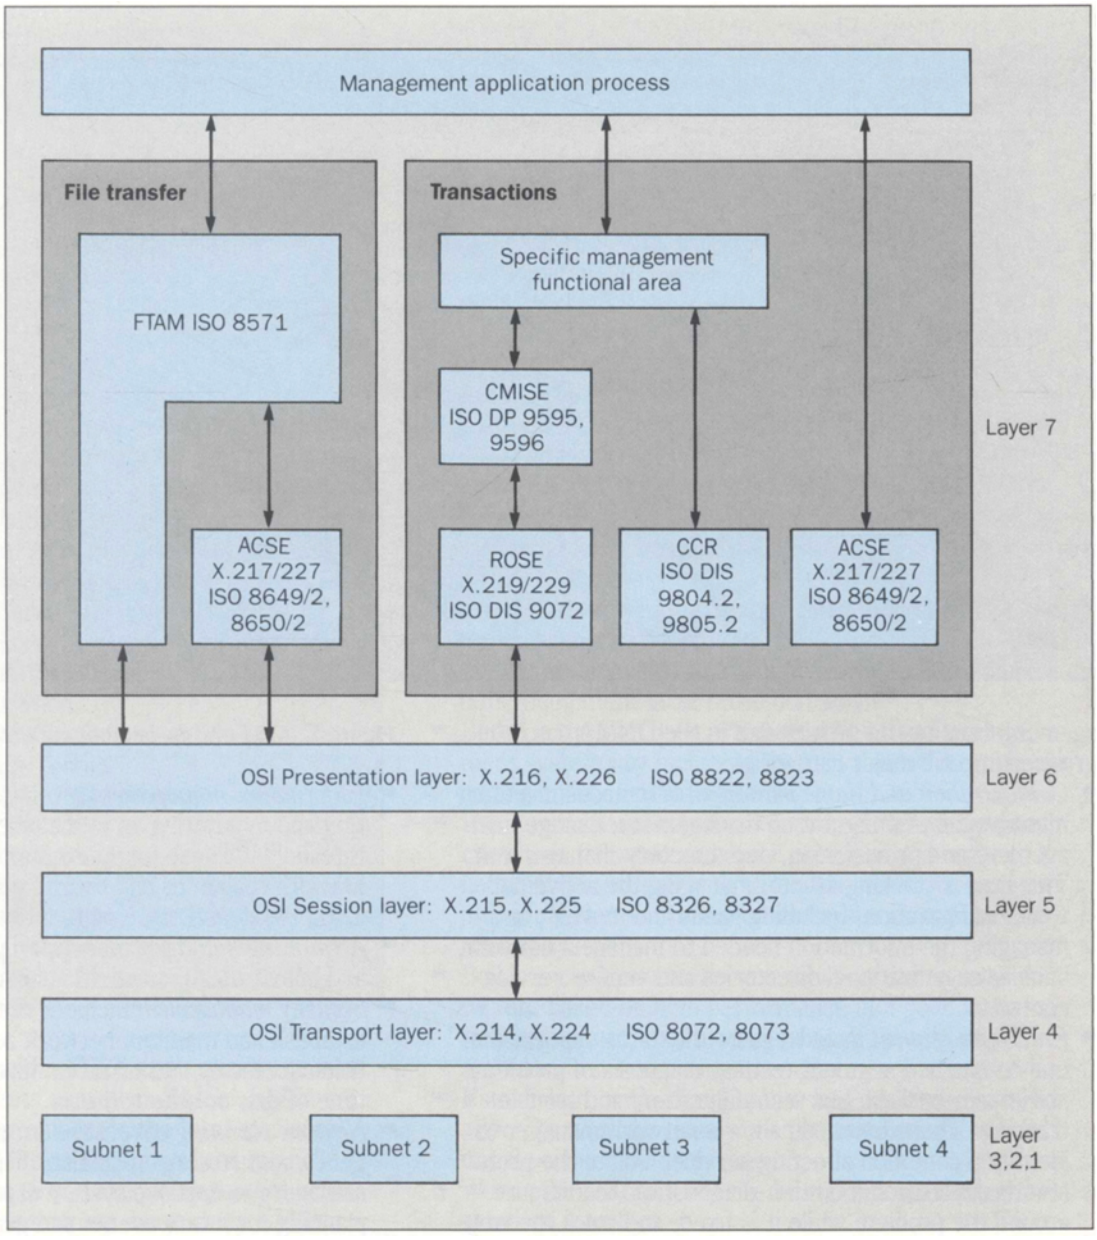

specific information or instructions be processed at the EMS and, possibly, flow between the EMS and the integrating system as well. Under UNMA, the messages needed to accomplish this, some of which have been defined and published by AT&T,[1,2] are the specific functional modules of the application layer protocol described below.

Fault management is the most sought-after function of the network management functions. Network managers must be able to relate network events, alerts, and alarms to the active network configuration. AT&T, in defining, refining, and implementing its UNMA, has focused first on the areas of fault and configuration management. Indeed, in April of 1988, AT&T published message set specifications as part of AT&T's NMP[1,2] that are compatible with the International Organization for Standardization Open Systems Interconnection Basic Reference Model.[3]

Standard Network Management Protocol. UNMA is intended to be an open architecture. The NMP interface from the EMS to the integrating system is based on emerging domestic and international standards of the ISO/OSI.

Protocol content. The basis for NMP is the OSI seven-layer stack of protocols shown in Figure 2. In this figure, the lower three layers—the physical, data link, and network layers—are represented by the multiple subnets shown at the bottom. The NMP is implemented in the higher layers of the seven-layer model and is independent of the specific implementation of the three lower layers as long as Layer 4, the transport layer, is properly implemented.

The OSI protocol stack is generally well defined within standards bodies. Specifications for the transport layer (Layer 4),[4] session layer (Layer 5),[5] and presentation layer (Layer 6)[6] are accepted standards. Within the application layer (Layer 7), there are numerous sublayers used for network management.[7] ACSE (association control service elements)[8] and FTAM (file transfer, access, and management)[9] are also final specifications. ROSE (remote operations service elements)[10] and CCR (commitment, concurrency, and recovery)[11] are official drafts that have undergone extensive review and are relatively stable. CMISE (common management information service elements)[12] is an unofficial draft that is being considered by the ISO, the International Telegraph and Telephone Consultative Committee (CCITT), and the ANSI-accredited T1M1 committee of the Exchange Carriers Standards Association. It may undergo change.

CMISE provides a structure for network management application messages that are themselves further defined in the specific management functional areas of Layer 7.

AT&T has also published its target *common* NMP protocol choices up to and including CMISE.[13] If the CCITT, ISO, or T1M1 final protocol or message specifications are different from these specifications, AT&T will, at an appropriate time, change the NMP specifications to conform to the standards and migrate AT&T products to the standard as needed.

Benefits of a standard network management protocol. There are advantages to using any standard for network management integration, and more extensive advantages when an international standard based on the OSI Basic Reference Model is used. These include the following:

- A common communications platform permits one network for OSI user applications and network management.
- Open naming conventions and standard data fields make management in a multivendor environment easier.
- Future OSI enhancements can be used to support network management.
- Network management expenses can be reduced through common, reliable specifications and software.
- There is increased demand for intelligent analysis and exploitation of expert system technology.

AT&T and its customers are building extensive local and national networks based on 802 LAN, X.25, and ISDN protocols. These networks support a wide range of business applications. Any of these networks can be used to transport

NMP as long as the standard Layer 3-to-Layer 4 OSI interface is available. Furthermore, Layers 4, 5, 6 and ACSE and FTAM in Layer 7 are application-independent and support network management as well as other transaction and file-oriented applications. There can be a significant savings when standard software is used to support all these needs because software development and maintenance expenses are reduced. There is also the reliability of a proven, well-used approach.

An open naming convention for network objects is needed to permit unique identification of elements in a multivendor environment. This naming convention has to be hierarchical, so that a registered code at the higher level allows independence across vendors and a structural approach at the lower levels gives an integrating system an understanding of the managed elements.[14] Furthermore, particular data fields must have a standard set of well-understood values. For example, the meaning of alarm severity values must be standard across vendor equipment to allow an integrating system to handle alarms properly. Standardized field values also allow a performance threshold to be applied across a network of multivendor equipment.

OSI layer functions are continually extended by experts in the standards bodies that understand the increasing complexity of application communication needs. Network management is just one example of a sophisticated application and can benefit from future enhancements to OSI. One near-term example is the ISO working group that is now defining a commitment, concurrency, and recovery (CCR) service to allow distributed database management. AT&T is already planning to take advantage of this work in our network management applications.

In the past, AT&T defined and developed a new machine-to-machine specification almost every time we needed a new interface between EMSs or an enhancement to an existing interface. NMP defines a generic message set that can be used for multiple products and services with enhancements provided as needed. We have already seen that multiple services can use the fault management specification by merely adding new fields and new field values. This message set will become more robust and stable as we gain experience in its use.

The integrating network management system will eventually use expert systems technology to reduce the need for expensive, highly trained network management technicians. Today AT&T EMSs have expert system capabilities that automatically analyze faults, initiate repairs, and clear alarms. To achieve these results in a complex multivendor network, all EMSs need to perform similar functions. For example, all EMSs used for fault management functions would have to be able to clear an alarm based on a designated command. NMP defines these functions and provides an impetus for the development of standard applications across EMSs.

The process of migrating an entire generation of products to the NMP is expected to take many years. The NMP application message specifications will continue to expand and will eventually include all the target network management functions identified earlier.

UNMA System Platform

UNMA is inherently a distributed architecture. Network management functions occur within EMSs and within an integrating network management system, with the actual distribution left to implementors at this time. To provide a common framework for UNMA applications development, a UNMA system platform for integrating systems and EMS development has been defined and developed. The platform offers a distributed execution environment for rapid and concurrent development of network management applications, and at the same time, provides a flexible system architecture to meet different customer and system needs.

Figure 3 shows the UNMA system platform architecture from an application designer's perspective. The platform uses the standard UNIX® System V capabilities and the UNIX system application operating environment (AOE). The platform has the following key components:

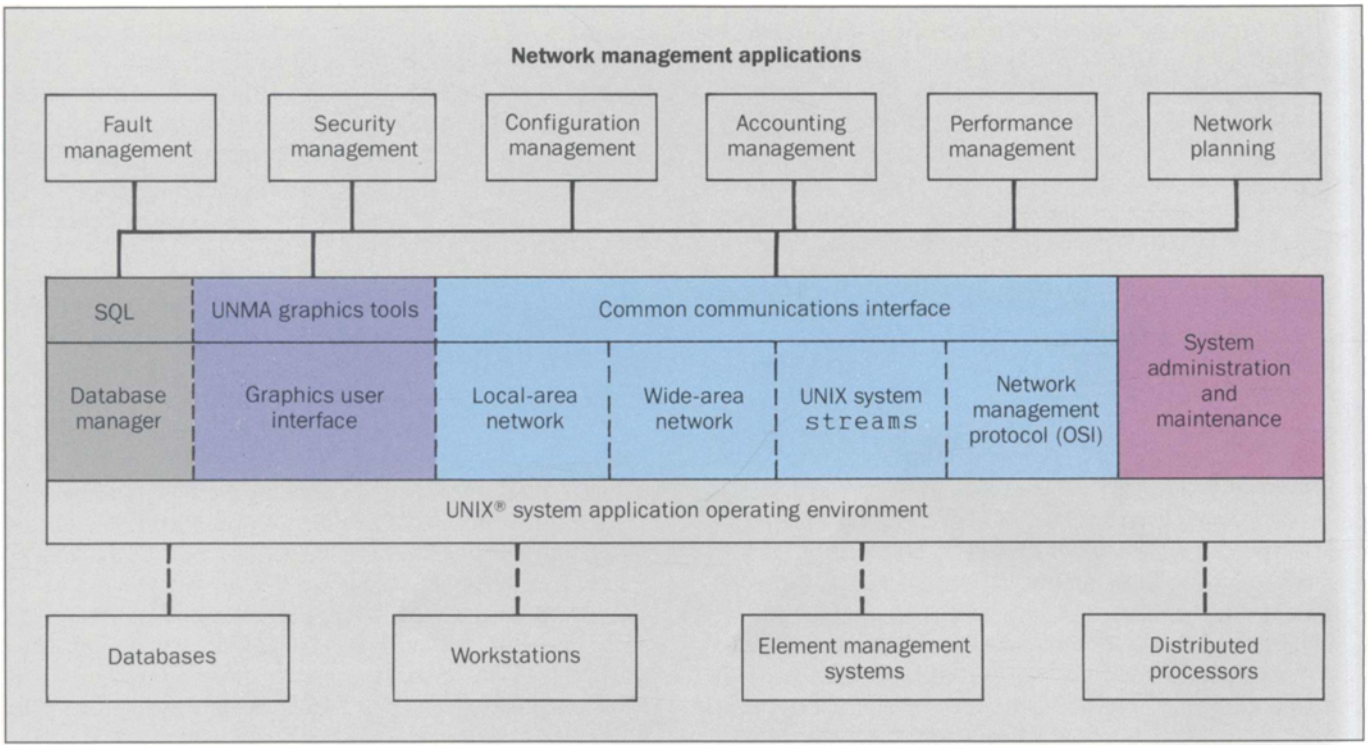

Figure 3. UNMA system platform architecture.

- *Database Manager.* The database manager is responsible for searching for and updating information in the database at the request of network management applications. Requests are in the form of SQL (structured query language) statements embedded in the application code. The database manager sees that the database is kept in a consistent state, backed up at appropriate intervals, easily recovered, and made available only to authorized client processes. For a fully distributed system, it also provides a mechanism for distributed transactions.
- *Graphics User Interface.* The user interface is responsible for taking display data from a network management application, formatting it for the UNMA workstation family, and displaying it in a window for the user. A set of tools designed for network management applications, such as for displaying a network configuration map, is also available to UNMA application designers. The user interface platform is based on XWINDOWS technology and enforces AT&T's network management "look and feel" user interface standards and, thus, achieves the goal of having a unified and distinctive UNMA user interface.
- *Communications.* The communications service is responsible for providing reliable communication between the various network management applications through a common interface that is independent of the underlying mechanisms used. It builds on an ISO-based NMP protocol stack plus the NMP message set for

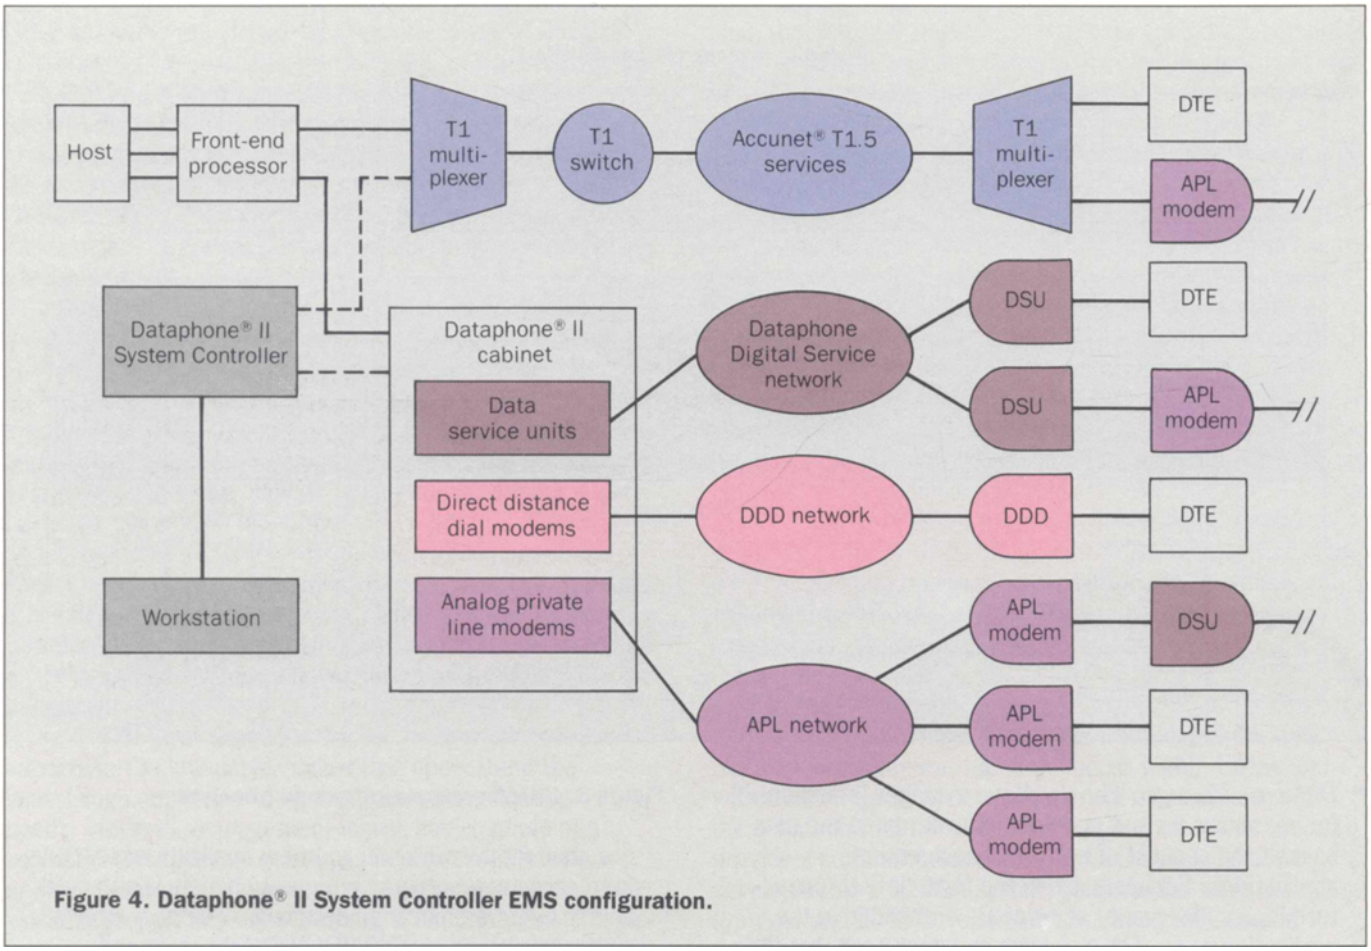

Figure 4. Dataphone® II System Controller EMS configuration.

intersystem communication (EMS and integrating system) and a local distributor capability for interprocessor communication within an integrating system (distributed processors and workstations). Routing of messages (including NMP messages) is performed by a directory service that translates the name supplied by the application into an appropriate address.

- *System administration and maintenance.* The system administration and maintenance component is used to initialize the system, bring it down gracefully, and respond to the variety of errors detected in the operation of the platform components, such as an overload condition. In addition, the system administration component can control or respond to the integrating system's own reconfiguration, which may occur on schedule or spontaneously in response to a problem. This component also synchronizes data across processors, checks platform subsystems, and writes and keeps the execution trace and error log files.

Security is an important function of system admin-

***Dataphone*® II System Controller**

The Dataphone II System Controller is an element management system that provides configuration and fault management functions to the AT&T family of data communications equipment products. Currently, this EMS supports the following products (see Figure 4.):

- The complete Dataphone II product family, which includes the 2000 series analog private line (APL) modems, the 2600 and 2700 series data service units (DSUs), the 839 series dial back-up units (DBUs), as well as rack-mounted direct distance dial (DDD) modems with diagnostic options.
- The series 700 family of multiplexers and switches.

Modems and Multiplexers. APL modems allow data terminating equipment (DTE) to transmit data over dedicated four-wire voice-grade lines. The digital data stream from the DTE is encoded into symbols that are transmitted over private lines as analog waveforms. When the quality of dedicated lines deteriorates to unacceptable levels, the user can restore service by placing switched calls using DBU units.

The series 700 family includes the Acculink™ T1 multiplexers and switches, 56-kilobit-per-second (Kb/s) time-division multiplexers, as well as statistical multiplexers and packet switches. T1 multiplexers and T1 switches connect DTE with the Accunet® T1.5 service. Both data and voice can be transported over the T1 links. In addition, T1 switches can interchange channels from different T1 links; and thus, traffic from a T1 node can be routed to different nodes. (For the purposes of this article, a T1 node is a generic term used to refer to either a T1 multiplexer or a T1 switch.)

The APL modems, DSUs, DBUs, and the T1 nodes all have network management functions built into them. However, some of the more advanced network management functions, in particular those that require a global view of the network, are provided by a centralized network management system like the system controller EMS.

System Controller Configuration. This EMS consists of a central processor, workstations, printers, and a network polling unit (NPU).

The central processor is from the AT&T 3B family of processors running the UNIX operating system. As many as 11 workstations from AT&T's PC family can be connected to the central processor. Menus and forms provide a user friendly interface for executing commands.

The NPU handles all communications functions between the central processor and the modem network. The NPU directly communicates to the local modems (i.e., those that are collocated with the EMS). Communication to remote modems is through a diagnostic network that uses "secondary channels" on the same private lines that transmit the primary data.

The system controller does not use the NPU for its connection to the multiplexer network. Instead, the central processor is directly connected to a collocated multiplexer or switch. Communications to other nodes are through a supervisory network interconnecting all the nodes using preallocated bandwidth on the T1 lines.

istration and maintenance. Certain users have certain levels of privilege, and must not be allowed to change or inspect data that is not in their area of responsibility. This is handled on a system-wide basis with password protection. Dedicated workstations, protocol encryption, and other conventional security measures can also be taken.

EMS Applications

The element management system, generally speaking, provides network management users with the functions and applications needed to monitor and/or control the network elements that they manage. Often these functions include configuration and fault management functions. This section will discuss those functions as exemplary of EMS functioning, using a particular EMS, the Dataphone® II System Controller, described in Figure 4 and the accompanying panel.

Configuration Management. Configuration management gives a timely and accurate representation of the network configuration (both physical and logical), including

the status of network resources. It also offers the user the ability to configure network resources (devices, facilities, and logical connections), including the scheduling of configuration changes. Finally, configuration management provides the user with methods for restoring the network configuration.

Configuration changes. The configuration of a network resource may change in a variety of ways:

- The network administrator may configure or reconfigure the resource in real time.
- The administrator may schedule a configuration change to occur at some time in the future.
- The resource may reconfigure itself in response to a fault condition.

To illustrate this point, we consider the T1 multiplexer application. Series 700 Acculink T1 multiplexers allow a user to segment T1 bandwidth into logical, end-to-end connections called channel groups. The channel group is analogous to a logical private line partitioned within a T1 facility. End-to-end voice and data circuits are then assigned bandwidth from the channel group.

An AT&T EMS allows a user to configure the parameters that characterize the channel group, such as its bandwidth, its route through the T1 network, and whether it should be automatically restored if the network fails. This EMS also lets the user assign voice and data circuits to channel groups and to configure the parameters that characterize the circuits.

The user may also schedule changes in the configuration of the channel group and the circuits that it carries. This function is available through the intelligence residing in the multiplexer itself, as well as through the EMS. For example, if an application only runs at specific times during the day, the user may want to commit T1 bandwidth to the application only when it is running. Bulk data transfer often is such an application—that is, scheduled to occur between two or more host computers during a network's off-peak hours. With the AT&T EMS for T1 multiplexers, the administrator can schedule the connection of the channel group for the period of time during which the file transfers occur. At other times, the channel group can be disconnected and T1 bandwidth can be available for other applications.

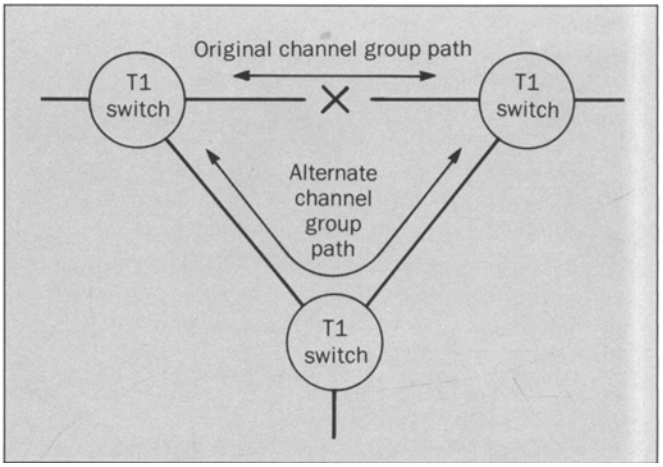

Figure 5. Rerouting of channel group after a facility failure.

The user may designate that a channel group be rerouted if a network fails. Although this option is set from the EMS, the rerouting function resides in the multiplexer. The network notifies the EMS of the status of channel group connections (routing failure, connected or disconnected). The EMS then allows the user to send inquiries to the network to determine the current route of a channel group. (See Figure 5.)

Configuration restoration. With this function, the user can restore the configuration of a network resource to a reference configuration. This capability is needed because of the rerouting of channel groups that occurred in the multiplexer application mentioned earlier. The restoration functionality resides in the multiplexer. As a result of the restoration, the route of the channel group returns to the reference configuration routing.

The role of the EMS is twofold. It monitors the reconfiguration that accompanies the restoration. It also provides commands that allow the user to acquire the cur-

rent route of the channel group, as well as the utilization of bandwidth on the T1 facilities that originally carried the channel group. The data returned by these commands help the user make the decision to restore the channel group to its original route (the reference state) or leave it in its current route.

Fault Management. Fault management is an important network management function. Although faults happen infrequently in a network, they are important because they can disrupt service. The process is as follows:

1. The EMS monitors the network for alarms, unusual status conditions, and events.
2. Transient alarms and those that are of no interest to the user are filtered out. The remaining alarms are displayed at the workstation and stored in a database so that management reports are available on demand. Based on parameters specified by the user, alarms can also trigger EMS actions.
3. Once a problem is detected in the network, users may choose to run the diagnostic tests and commands provided to help isolate and identify the problem.
4. Service is then restored, even if only on a temporary basis.
5. The problem is tracked until final repairs are made.

Alert monitoring. The EMS we use as our example, the Dataphone II System Controller, continually monitors the network for alarms, status reports, and events. Alarms are indications that some network resource is malfunctioning. Status reports may show that some network resource is not in its normal operational state (e.g., executing tests that disrupt data flow). Alarms and status reports are conditions that have a duration. Events, on the other hand, are transient. An example of an event is a change in a T1 node configuration.

All devices under the control of the Dataphone II System Controller can monitor their own health. However, because a device that fails may not be able to report the failure, a positive acknowledgment scheme (polling) is used. The Dataphone II System Controller requests a health report from each device at regular intervals. The device acknowledges the requests stating whether or not it is healthy. A "no response" means a failure has occurred either in the diagnostic communication network or in the device itself.

Apart from monitoring their own health, devices also monitor:

- Failures at the network interface (e.g., facility fault for APL modems and red alarm for the T1 nodes).
- Failures at the terminal interface (e.g., streaming terminal for APL modems and channel alarm for T1 nodes).
- Breakdown in communications over the management network (e.g., no response for APL modems and supervisory data link timeout alarm for T1 nodes).

Alarm processing and storage. Incoming alarms are displayed so that the user can take action. In addition, they are stored in a database from which information can be retrieved for management reports.

However, some alarms, such as a facility fault (FA) alarm for APL applications, are transient in nature. More than 80 percent of FA alarms last 2 minutes or less. Most users want to filter out transient alarms to display and/or store only those that exceed a certain time threshold.

Filtering can also be used to regulate the actions that are triggered by alarms. One example is the automatic trouble reporting (ATR) process that is described later.

Fault isolation. The network elements offer the user a collection of tests to isolate and quantify an alarm. The user can initiate these tests either by locally accessing the network elements or by using an EMS workstation. Tests can be classified based on those that disrupt the data flow and those that do not.

Again, in the example of the FA alarm for the APL modem, the FA alarm suggests that a circuit is of poor quality. The user can send a command to the modem to run a test that reports an estimate of the analog parameters of the channel without disrupting the data. If necessary, the user can perform additional tests and get accurate measurements of these analog parameters by temporarily suspending data flow. Also, the user can perform tests that measure bit error rate of the circuit.

Sometimes, communication between two pieces of equipment breaks down without an alarm from the network. This can be because of a fault in the terminal equipment or some undetected problem in the network (e.g., a broken cable). Both the APL modems and T1 nodes provide loopback tests to isolate these problems.

Restoring service. Faults in the network usually disrupt service. If the service is important, the user needs to restore the service as soon as possible, even with temporary measures.

More than 99 percent of the faults in APL modems are because of circuit failures. With DBU units, a user can initiate calls over the public switched network and replace the failed private line circuits until permanent repairs are made. Then, the user can restore the service back to the repaired lines.

T1 nodes are too critical to allow failures of the node. Therefore, all common logic of a T1 node is backed up by duplicate components that will automatically take over when the primary component fails. In addition, the T1 switches automatically reroute traffic if a T1 facility fails. This requires either that additional bandwidth be available or that a priority scheme exist to resolve contention. In both cases, good network design is necessary to ensure smooth operation.

Trouble tracking. Temporary measures to restore service usually have costs associated with them. For dial back-up, it is the charge of the calls. For rerouting, it is the possibility that low-priority traffic may be replaced by higher priority traffic that is being rerouted. Therefore, it is imperative to repair failed components as soon as possible. The system controller EMS we use as our example has a trouble ticket system to manage the repair process.

Because most failures are facility failures, close coordination with the maintenance center(s) of the common carrier(s) is important to ensure a speedy repair. The automatic trouble reporting feature of the EMS allows the user to report faults (both device and facility) to AT&T's maintenance centers automatically. The user can control this process by specifying that only alarms from certain devices and only those that last longer than a certain time period be reported.

External Interfaces. Under UNMA, an EMS supports an external interface to an integrating network management system to provide that system with data in support of fault and configuration management. In response to requests from an integrating system, the EMS will provide data on alerts and configuration changes.

UNMA-conforming EMSs also support remote terminal access from an integrating system, thereby allowing the user of an integrator to access the network management functionality of EMS directly.

Integrator Applications

As illustrated in the previous section using the management of modems and multiplexers as an example, each EMS under UNMA performs network management functions pertinent to the network elements and domains that it manages. However, in networks that span multiple domains and multiple vendors, the EMS applications alone *cannot* provide an end-to-end integrated view of the network. Some additional, value-added functions, *integrated management applications*, must be implemented in the UNMA integrating system. This section provides two examples of such value-added integrated functions.

Integrated Configuration Management. Under UNMA, each EMS has its own configuration database that includes the network elements that it manages. For example, the system controller EMS discussed in the previous section contains configuration data regarding the modems, multiplexers, and DBUs it manages. To manage a network in an integrated, end-to-end fashion, the integrating system requests downloads of configuration data from each EMS so that it can present the integrating system user with a complete network configuration such as the one depicted in Figure 6. When the network changes (for example, a T1 node is added, removed, or reconfigured), the EMS database is typically updated and the EMS transmits to the

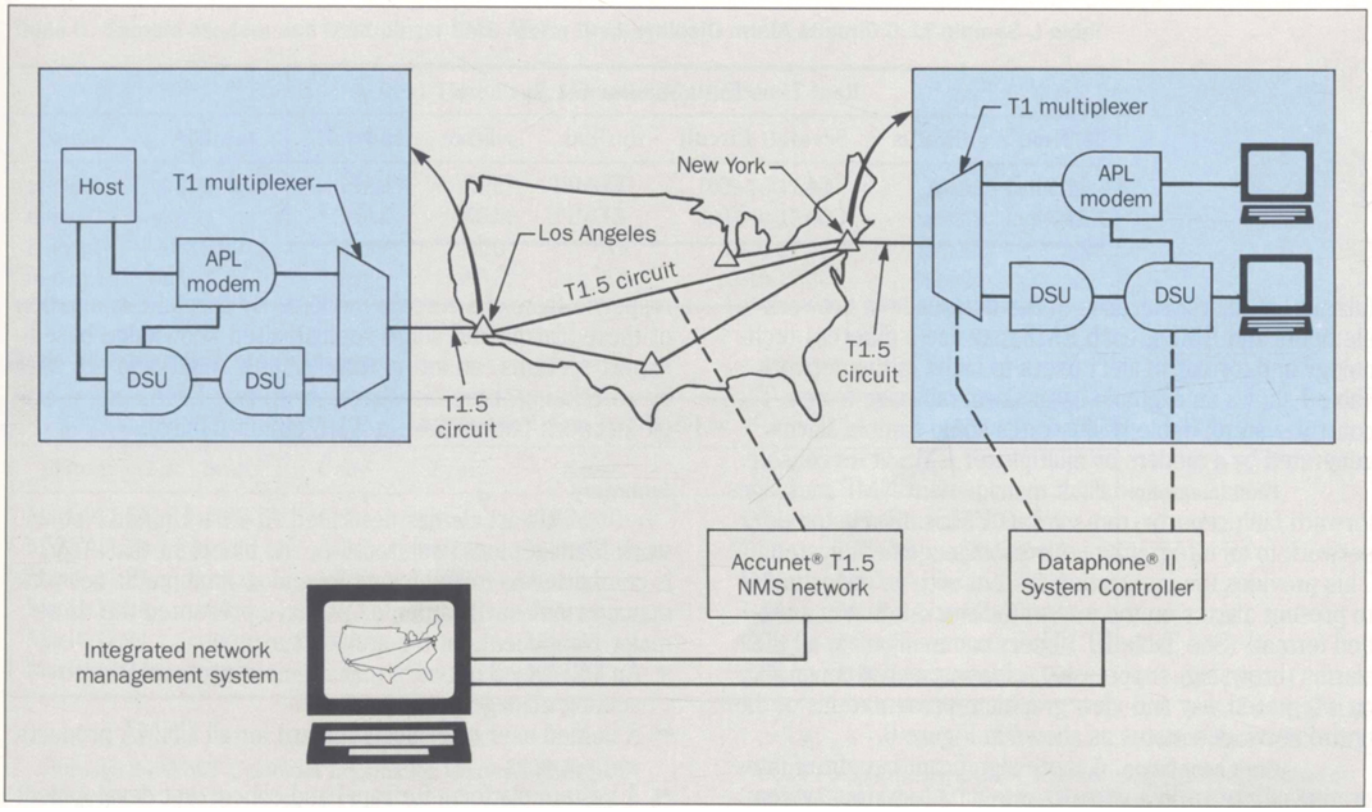

Figure 6. Sample network configuration.

integrating system an NMP configuration-data-update message reporting the change.

By requesting, downloading, and synchronizing the integrating system's configuration data with each EMS database, the integrator can present a user with a view of an end-to end, integrated network configuration that includes all the network elements in a customer network. Furthermore, by using the graphics tools provided by the UNMA platform, hierarchical configuration display of customer networks is also possible.

Figure 6 is an example of a simple network configuration display on a graphics workstation showing network connectivity and elements of T1.5 lines and T1 multiplexers, APL modems, and DSUs. An EMS may provide the equipment information for this view, while an Accunet T1.5 service EMS may keep the integrating system current regarding the T1.5 facilities used in the network. Using reported connectivity as a guide, the integrating system can build relatively complex network configurations with data from many EMSs.

Integrated Fault Management. A variety of element management systems can monitor and present faults

Table I. Sample T1.5 Circuits Alarm Display

Real-Time Fault Monitor (T1.5)						
Time	Status	Srvc	Circuit	Text	From	To
17:34pm	Y	CGA	001	COMPL	NY	LA1
17:34pm	R	CGA	002	COMPL	NY	LA1

(alarms) to users. Because of the difference in network elements and EMSs, each EMS may use a different technology and format to alert users to faults in the network. Table I shows an example of an alarm message from a T1.5 control system; Table II illustrates some sample alarms generated by a modem or multiplexer EMS at its console.

Fault integration. Fault management NMP messages forward fault reports from various EMSs in a customer's network to an integrating network management system. This provides the integrating system with an opportunity to present alarms on the integrated workstation in a unified format. (See Table III.) Users can manipulate all the alarms (browsing, suppressing, clearing, and archiving) in an integrated way and view graphic representations of the entire network's status as shown in Figure 6.

Fault correlation. A more significant benefit of integrating all alarms in a network into an integrating system is the possibility of providing a sophisticated intelligence-based expert system that can correlate the multiple alarms from various EMSs to a single network failure, and thus greatly reduce the cost of trouble-shooting the various individual alarms.

For example, in Figure 6, several alerts are generated by the premises equipment EMS and T1.5 network monitoring EMS when the T1.5 circuit between New York and Los Angeles goes down. Specifically, T1.5 monitoring at central office cross-connects transmits two alarms showing a loss of carrier group for both directions. The modem and multiplexer EMS generates four alarms, two from the T1 multiplexers and two from the modems, complaining about possible failure in the facility and no response from the remote modem. By careful examination of these alarms and some sophisticated knowledge-based expert systems, an integrating system could help the user by correlating these six alarms from two EMSs into a single network failure (i.e., a T1.5 circuit failure).

Summary

This article has described AT&T's Unified Network Management Architecture, the blueprint that AT&T is committed to follow for multivendor, multimedia network management environments. We have presented the three major components of the architecture:

- An OSI-based network management protocol for connecting management systems
- A unified user interface standard for all UNMA products and services
- A system platform for rapid and concurrent development of UNMA applications.

The architecture calls for a three-tier structure of network management, network elements, element management systems, and integrating management systems, with emphasis on capitalizing on and improving existing management applications.

Acknowledgments

We wish to acknowledge the contributions of Tom Chu, Anthony Longhitano, and Tom Walker to this article.

References

1. *AT&T Network Management Protocol—Fault Management Message Set Specification*, Publication No. TR54005, April 1988, available

Table II. Sample Modem and Multiplexer EMS Alarm Display

Real-Time Fault Monitor (D/M EMS)							
Name	**Address**	**Network**	**Srvc**	**Fault**	**Date**	**Status**	**ATR**
m-mod1	4/31/1	inms	APL	FA	03/10-17:34	nonfilt	na
m-mux1	m2/11	inms	MUX	IA	03/10-17:34	nonfilt	na
lb-mod1	4/031/1/01	inms	APL	NR	03/10-17:34	nonfilt	na
m-mux1	m2/11	inms	MUX	EA	03/10-17:34	nonfilt	na

Table III. Integrating Alarm Display

Integrator Alarm Log Window						
Time	**State**	**Code**	**Type**	**Name**	**EMS**	**Notes**
Mar10 17:34	failed	APL	wire	CIR-m-m1-x1	D/M	APL FA nonfilt ATR NA
Mar10 17:34	failed	APL	modem	lb-mod1	D/M	APL NR nonfilt ATR NA
Mar10 17:34	critical	MUX	MUX	m-mux1	D/M	MUX IA nonfilt ATR NA
Mar10 17:34	critical	MUX	T-1	CIR-m-mux1	D/M	MUX EA nonfilt ATR NA
Mar10 17:34	critical	CGA	T-1	NY-LA1	CCR	D/M T1 alarm circuit 002
Mar10 17:34	critical	CGA	T-1	NY-LA1	CCR	D/M T1 alarm circuit 001

through the AT&T Customer Information Center, Indianapolis, Indiana.
2. *AT&T Network Management Protocol—Configuration Management Message Set Specification*, Publication No. TR54006, April 1988, available through the AT&T Customer Information Center, Indianapolis, Indiana.
3. *Information Processing Systems—Open Systems Interconnection—Basic Reference Model*, International Organization for Standardization, Publication No. 7498, October 1984.
4. *Transport Service Definition for Open Systems Interconnection for CCITT Applications*, CCITT Recommendation X.214, International Telegraph and Telephone Consultative Committee and *Transport Protocol Specification for Open Systems Interconnection for CCITT Applications*, CCITT Recommendation X.224, International Telegraph and Telephone Consultative Committee (Red Book), October 1984.
5. *Session Service Definition for Open Systems Interconnection for CCITT Applications*, CCITT Recommendation X.215, and *Session Protocol Specification for Open Systems Interconnection for CCITT Applications*, CCITT Recommendation X.225, International Telegraph and Telephone Consultative Committee (Red Book), October 1984.
6. *Information Processing Systems—Open Systems Interconnection—Connection Oriented Presentation Service Definition*, ISO 8822, International Organization for Standardization, 1986, and *Information Processing Systems—Open Systems Interconnection—Connection Oriented Presentation Protocol Definition*, ISO 8823, International Organization for Standardization, June 1987.
7. *Information Processing Systems—Open Systems Interconnection—Application Layer Structure*, ISO DP 9594, International Organization for Standardization, March 10, 1987.
8. *Information Processing Systems—Open Systems Interconnection—Service Definition for Common Application Service Elements, Part 2: Association Control*, ISO 8649/2, International Organization for Standardization, 1986, and *Information Processing Systems—Open Systems Interconnection—Protocol Specification for Common Application Service Elements, Part 2: Association Control*, ISO 8650/2, International Organization for Standardization, 1986.

9. *Information Processing Systems—Open Systems Interconnection—File Transfer, Access and Management, Part 1: General Introduction, Part 2: Virtual Filestore Definition, Part 3: File Service Definition, Part 4: File Protocol Specification,* ISO 8571-1, International Organization for Standardization, 1987(E).
10. *Remote Operations: Model, Notation and Service Definition,* CCITT Recommendation X.219, International Telegraph and Telephone Consultative Committee, 1988 and *Remote Operations: Protocol Specification,* CCITT Recommendation X.229, International Telegraph and Telephone Consultative Committee, 1988.
11. *Information Processing— Open Systems Interconnection—Definition of Application Service Elements—Commitment, Concurrency and Recovery,* ISO DIS 9804.2, International Organization for Standardization and *Information Processing—Open Systems Interconnection—Specification of Protocols for Application Service Elements—Commitment, Concurrency and Recovery,* ISO DIS 9805.2, International Organization for Standardization, 1988.
12. *Information Processing Systems—Open Systems Interconnection—Management Information Service Definition,* ISO DP 9595/1,2, International Organization for Standardization and *Information Processing Systems—Open Systems Interconnection—Management Information Protocol Specification,* ISO DP 9596/1,2, International Organization for Standardization, August 1987.
13. *AT&T Network Management Protocol Specification—Transport through Application Layers,* Publication No. TR54004, January 1988, available through the AT&T Customer Information Center, Indianapolis, Indiana.
14. *AT&T Network Management Protocol—Data Modeling and Naming Framework,* Publication No. TR54007, April 1988, available through the AT&T Customer Information Center, Indianapolis, Indiana.

Biographies (continued)

engineering and an Sc.D. in computer science from the Massachusetts Institute of Technology. ***Raymond J. Pennotti*** *is head of the Data Systems Definition and Test Department of Bell Laboratories in Middletown, New Jersey and has been with AT&T since 1970. Mr. Pennotti's department provides systems engineering, system test, and field test for AT&T's data communications equipment, including analog private line and switched modems, Acculink™ multiplexers, and associated network management applications. Mr. Pennotti has a B.S.E.E. from Manhattan College and an M.S.E.E. and a Ph.D. in electrical engineering, both from the Polytechnic Institute of Brooklyn.*

(Manuscript received July 28, 1988)