# ♦ Network Management: Emerging Trends and Challenges

*Shervin Erfani, Victor B. Lawrence, Manu Malek, and Binay Sugla*

*Network management is arguably one of the most important challenges faced by the owners and operators of advanced networks and services. Lucent Technologies, as a major equipment vendor, and Bell Labs, as its R&D arm, have had extensive historical experience with designing network management systems for AT&T telecommunication networks and other associated data networks. This experience enables Lucent to provide unique differentiation to its products and also places it in an excellent position to address the network management needs and challenges of the emerging networks and services. This issue of the* Bell Labs Technical Journal *provides a sample of some of the many the activities, tools, technology trends, and designs currently under way in network management in various wings of Lucent. This paper, an introduction to the issue, discusses the evolution of network management, placing it in the context of underlying networking trends. In addition, it provides a brief introduction to the papers presented in this issue, which together provide a snapshot of Lucent's network management activities in support of an efficient, reliable, and flexible networking vision.*

## Introduction

Bell Labs has played a historical role in bringing reliability and dependability to the telecommunications infrastructure of the 20th century. A key component of this possibly unique capability is an understanding of the systems employed to manage these networks and services. The current challenge facing Bell Labs is to effectively harness this expertise to bring public switched telephone network (PSTN)-level reliability and dependability to data networks and services of today and tomorrow. This issue of the *Bell Labs Technical Journal* is devoted to bringing wide exposure to the ongoing work in network management in Bell Labs and Lucent Technologies business units.

This paper is organized as follows. In the section immediately below, we describe the evolution of network management and the role of network management functions. The section following it provides the key components of a network management system (NMS) and lays a foundation for understanding the papers in the issue. The next section, "The Network Environment: Then, Now, and in the Future," presents an overview of the current trends in networking and services that have strong implications for the requirements of a management system. Subsequently, "A Practical Vision for Managing Networks" delves into the management requirements induced by the networking trends mentioned in the previous section and thereby motivates the work described in this issue. The last major section, "A Snapshot of Trends and Techniques," provides a short summary of the papers in this issue.

## Evolution of Network Management

Telecommunication and services are in the process of revolutionary yet evolutionary changes due to transformation of the regulatory environment that in turn has given rise to rapid improvements in the underlying application, networking, computing, and

**Panel 1. Abbreviations, Acronyms, and Terms**

AAL—ATM adaptation layer
ADSL—asymmetric DSL
ATM—asynchronous transfer mode
B-ISDN—broadband ISDN
CATV—community antenna television
CIM—common information model
CMIP—common management information protocol
CMISE—common management information service element
CORBA*—Common Object Request Broker Architecture
DEN—directory enabled network
DiffServ—differentiated services
DSL—digital subscriber line
DMTF—Distributed Management Task Force
DWDM—dense wavelength division multiplexing
EMS—element management system
FCAPS—fault, configuration, accounting, performance, and security
GDMO—guidelines for the definition of managed objects
HDSL—high bit-rate DSL
HFC—hybrid fiber coax
IDL—Interface Definition Language*
IETF—Internet Engineering Task Force
IP—Internet protocol
ISDN—integrated services digital network
ITS—Internet telephony server
ITU-T—International Telecommunication Union–Telecommunication Standardization Sector
LAN—local area network
LDAP—lightweight directory access protocol
MCN—management communications network
MIB—management information base
MPLS—multiprotocol label switching
MUX—multiplexer
NMS—network management system
NMS—network management system
NNI—network-to-network interface
OAM&P—operations, administration, maintenance, and provisioning
ONU—optical network unit
ORB—object request broker
OS—operations system
PON—passive optical network
POTS—"plain old telephone service"
PSTN—public switched telephone network
QoS—quality of service
RFC—Request for Comments
RMON—remote monitoring
RSVP—resource reservation protocol
RTP—real time protocol
SDH—synchronous digital hierarchy
SLA—service level agreement
SMI—structure of management information
SMK—shared management knowledge
SMON—switch monitoring
SNMP—simple network management protocol
SONET—synchronous optical network
SS7—Signaling System 7
STM—synchronous transfer mode
SVC—switched virtual circuit
TCP—transfer control protocol
TDM—time division multiplexed/multiplexing
TDM—time division multiplexing/multiplexed
UDP—user datagram protocol
VDSL—very high bit-rate DSL
xDSL—any of various DSL modem technologies

transmission technologies. Lucent is supplying service providers with a variety of switching and transport technologies capable of delivering voice and high-speed data services. With these changes, network operators will soon require next-generation network management applications that deliver rich automated functionality capable of responding at high speeds, yet flexible enough to rapidly accommodate services that will be introduced.

The purpose of network management is the assignment and control of proper network resources, both hardware and software, to address service performance needs and the network's objectives. With the ever-increasing size and complexity of underlying networks and services, it has become impossible to carry out these functions without the support of automated tools. The progress in network management reflects the needs driven by these trends.

The evolution of network management is shown in **Figure 1**. Prior to the 1980s, network management was largely proprietary and, due to high development costs, it was considered affordable by only a few very large telecommunication carriers. Several developments in the 1980s caused a significant diffusion of

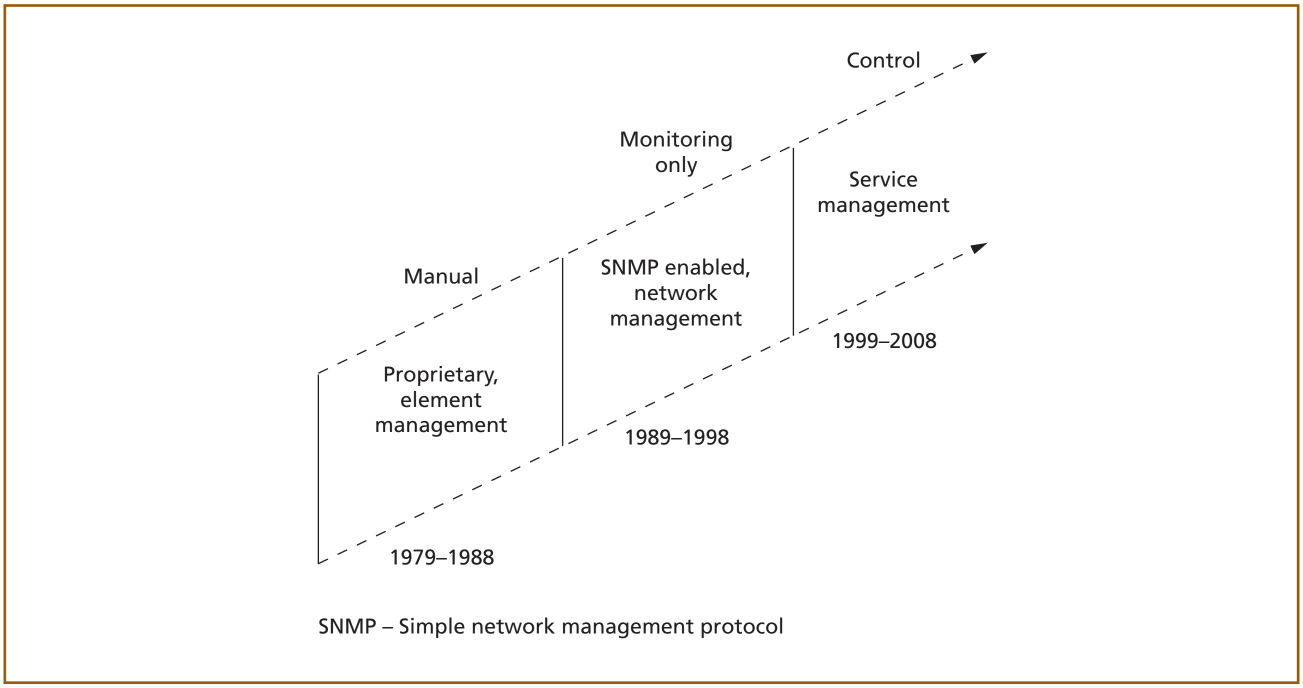


***Figure 1.***
***Evolution of network management.***

network management technologies. Standards like simple network management protocol (SNMP)[1] and common management information protocol (CMIP)/common management information service element (CMISE)[2,3] paved the way for a standardized method of communicating with the managed network elements, while platforms like OpenView* obviated the need to develop several core custom modules required in a network management system. By implementing a suite of integrated common software modules for topology discovery, map manipulation, drill-down viewing, and database support, these platforms reduced the cost of owning a network management system by an order of magnitude. However, the management functionality still remained in its primitive form, requiring network managers to observe and interpret raw variables manually. Given the number of variables per managed element and the typically large number of managed elements, there was a driving need for tools that provided higher levels of functionality at higher levels of abstraction.

Thus, the decade of the 1990s has been devoted to the realization that the only feasible way to manage large and increasingly complex networks is to devise innovative management applications for all aspects of management. Consequently, a number of automated management applications that perform advanced monitoring functions and attempt to provide an end-to-end network-service-level view of the network have been built. However, management at the network level is still insufficient, and there is a strong need to design and implement management applications that address the service level. The current focus on service level agreements (SLAs) is a precursor to this trend. (It may be noted that the AT&T long distance telephony network that was supported by Bell Labs had managed applications that operated at the service level and performed both monitoring and management applications that required control of the network elements. However, these applications were proprietary and assumed a single vendor solution.)

The general lack of effective, automated management applications has given rise to an important and fast-growing business opportunity—network management service providers, like Lucent's Netcare® Network Management Services.[4] These network man-

**Business management layer**
- Goal setting, finance, budgeting
- Planning, product definition

**Service management layer**
- Customer contact and interface
- Quality of service
- Interaction between services

**Network management layer**
- Connectivity among nodes
- Network control and coordination
- Network statistical log and events

**Element management layer**
- Control of subsets of network elements
- Gateway access to network elements
- Maintenance of statistical log and events

**Network element layer**
- Implementation of management commands
- Detection of problems

***Figure 2.***
***Network management layers.***

agement service providers serve a very useful function of offloading management tasks from enterprises and network providers who either lack the skilled staff to manage their own networks or would like to reduce the cost of management.

## Network Management Functions

Though this issue of the *Bell Labs Technical Journal* takes a broader, more popular view of the meaning of the term *network management* to include aspects of element and service management, network management is, strictly speaking, only one layer of the management systems hierarchy.

**Figure 2** shows this hierarchical, layered approach for partitioning and organizing management activities and tasks, which is aligned with management processes.[5] The layers recognize that the management of a telecommunications network service provider must deal with network elements, networks, services, and business issues. These five layers from the top down are:

- Business management layer,
- Service management layer,
- Network management layer,
- Element management layer, and
- Network element layer.

The layers are logical; they do not need to correlate with any physical implementation. The organization of management tasks is service-based in the sense that it will allow a specific service-management-layer item to become traceable to its corresponding required tasks at the network management layer all the way down to the network element layer. The philosophy is that the business requirements are the drivers for all management issues. A particular business goal dictates a specific set of requirements on the underlying service management layer, which in turn dictates the requirements for integration and interoperability on the underlying network management layer. At the lower layers, activities affect physical resources; at the higher layers, they affect relatively abstract entities and processes like SLAs. Thus, implementation occurs in the reverse order. The upper four layers distinguish management functions. The lowest layer (the network element layer) represents the management view provided by the network resources (either hardware or software).

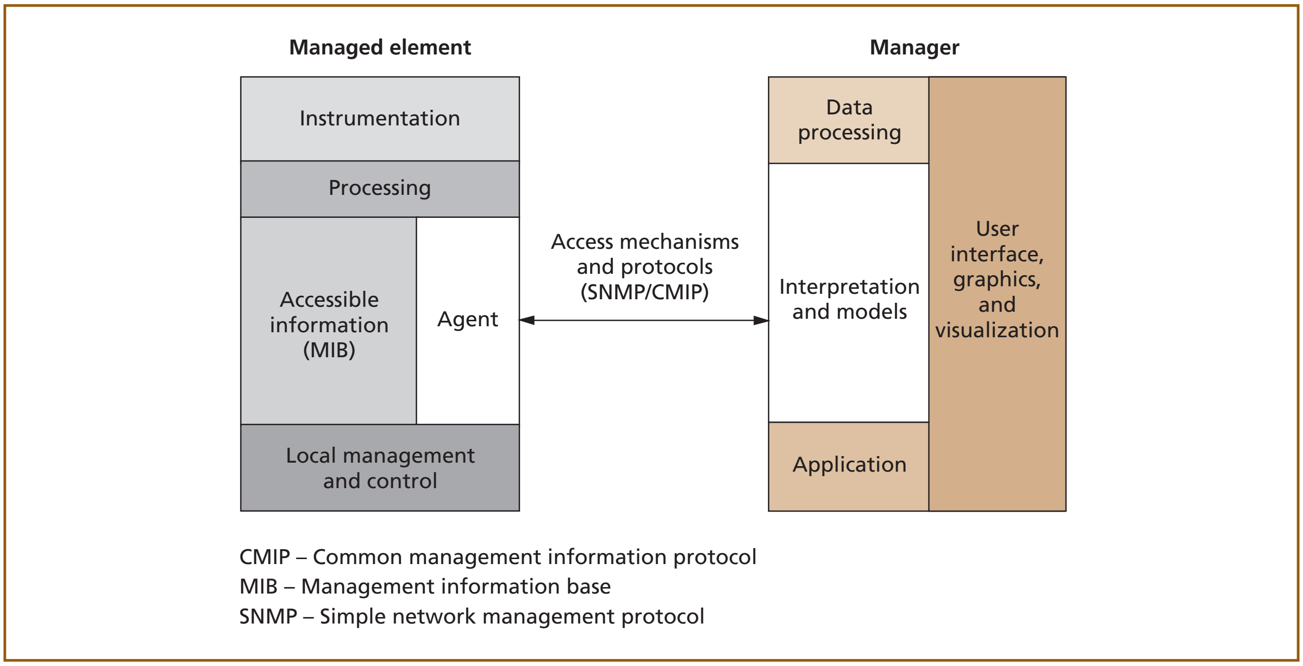


***Figure 3.***
***Manager-agent model.***

Similarly, the required tasks of each management layer can be categorized into the following five management functional areas: fault, configuration, accounting, performance, and security (FCAPS).[6] These management areas distinguish functions by the kinds of activities they perform. The management functional areas and organizational layers are orthogonal (mutually independent) classification schemes. For example, fault management, when applied to the network management layer, includes the functions that deal with analysis of alarms from all the network elements within a given network in order to localize the fault when functions at the element management layer have been unable to do so within a particular network element and/or limited subnet of network elements.

In practice, there can be many element management systems (EMSs) and network management systems (NMSs) for implementing management functions in each layer, which communicate with functions in other layers—typically the layers directly above or below the layer in question. The EMSs and NMSs manage a set of network elements or a subnet, respectively—that is, a *domain*, may have *inter-domain* communications.

The characteristics of inter-layer communications between management functions are:

- Lower layers are managed by upper layers.
- Each layer manages multiple occurrences of the layer directly below.
- Each layer uses network management functions/services of lower layers.
- An FCAPS application in a layer may communicate with other FCAPS applications in the same layer (for example, synchronous optical network (SONET) applications may talk to asynchronous transfer mode (ATM) applications).

This architectural picture has been hard to translate in practice due to the inherent technical difficulties and costs of development.

## Key Components of a Network Management System

The key constituents of a network management system are best understood by considering the relationship between the manager and the management agent residing on the management element, as shown in Figure 3. Given the diversity of managed elements like servers, routers, switches, and the wide variety of operating systems and programming interfaces, it is critical that the manager be able to communicate in a meaningful way with the managed element. This is done via the notion of an agent that not only can understand the communication protocol, but also can organize information in a consistent way so that both the manager and agent can refer to the same informational entity without ambiguity. Standardized communication protocols like SNMP or CMIP/CMISE along with associated information models—structure of management information (SMI) for SNMP and object-based models for CMIP/CMISE—are used for this purpose.

### Management Information and Knowledge Structure

Management information models are abstractions of the network resources. They define the structures and contents of management information that the management functions act upon. A management information model provides a common characterization of the network resources, enables multiple management functions to interact with each other, and supports different management functions. Management information is exchanged among management systems where management functions are implemented.

Attaining interoperability of network management systems and a common view of managed resources in a managed network environment requires that information models comply with standard models (or be able to map to standard models via proxy translations).[7] The management functions currently exchange management information by means of techniques defined in the ITU-T X.700 series of recommendations.[2,3,6,8-10] These recommendations incorporate the important object-oriented and manager/agent paradigms for information modeling.

The object-oriented approach for information exchange allows for grouping of data and the executable operations to be fully encapsulated into an object and object properties to be extended through "inheritance." The data can be manipulated only via the access operations provided in the object. The guidelines for the definition of managed objects (GDMO), ITU-T Recommendation X.722,[10] allow for a common data structure for managed objects in the managed and managing systems. Managed objects include their names, attributes, operations that can be performed on attributes, notifications that objects can emit, and behavioral descriptions of the objects. By using GDMO, any attribute, operation, notification, and behavior exposed by managed objects, as well as inheritance and containment relationships among managed objects and managed object classes, can be defined.

The manager/agent paradigm allows for hierarchical exchange of management information between functions (see Figure 3). A managing system assumes the role of "manager" for the purpose of issuing directives and receiving notifications, and a managed system assumes the role of "agent" for the purpose of carrying out directives and emitting notifications. This concept is well in line with the layered approach to the network management discipline. A system that manages a lower-level system may simultaneously be managed by a higher-level system, allowing for a cascaded management hierarchy.

The communicating management functions must share a common view of the available management information as well as a common understanding of the scope of management operations. This is referred to as shared management knowledge (SMK).[11] SMK comprises all data, information, and services that enable and support the management of a network. The most important concepts of SMK are:

- Integrated naming and addressing schemes for objects, which allow for both human and machine readability;
- Object classes, which allow for extensions to new services/technologies;
- Object attributes/services, which allow for vendor extensions and options; and
- Manager/agent/object relationships, which allow for data abstraction at upper layers.

The conceptual repository for management information is called the management information base (MIB).[10] A *MIB* is defined as a structured collection of managed object classes organized according to a management information model. Conceptually, specific domain MIBs can be constructed to include management information concerning data networks, telecommunication networks, and connected systems as well as application- and service-related databases. Each

MIB presents its own object groups defined for the corresponding application. For example, more than a hundred standards-based MIB modules are defined by a number of different Internet Engineering Task Force (IETF) Requests for Comments (RFCs).

### Management Information Exchange Protocols

Communicating functions in each management functional layer use network management protocols and message sets for transferring information between themselves and for triggering actions in other layers (usually, in adjacent layers). The manager and agent are linked by a network management protocol. For management of telecommunication networks, the management information exchange is based on CMIP. Nevertheless, because of its complexity and the slowness of industry deployment of CMIP chip implementations, alternatives to CMIP have arisen. Although CMIP can in theory be used as a general-purpose distributed computing facility, it is still immature in practice, lacking in software development support. As an alternative, Common Object Request Broker Architecture (CORBA*), which offers both a communication and computing facility, can be used.

The protocol used for management of data networks in a transfer control protocol (TCP)/Internet protocol (IP) environment is simple network management protocol (SNMP),[1] which is now expressing its legitimacy. From the telecommunications network point of view, CMIP provides a flexible template in communicating messages between manager and agent, which incorporates the needed aspects of the information model within a GDMO template. However, the convergence of voice and data networks requires blending the customer environment of SNMP-based management with the service provider's more complex environment. The complexity of telecommunications carrier networks requires complete management solutions to support the network elements, end-to-end layer networks, and subscriber-oriented services. In contrast, in an enterprise network environment, a total solution encompassing the network device management and high-level user service management can be achieved using the SNMP-based framework. The strength of SNMP is its simplicity. However, simplicity does not come without deficiencies. Because SNMP relies on user datagram protocol (UDP), which is a connectionless protocol, SNMP is itself connectionless. No ongoing connections are maintained between a manager and its agents.

Clearly, in an integrated network management environment, protocol conversions between SNMP and proprietary solutions or between SNMP and CMIP stacks are needed. The use of protocol adapters has long been promoted.

## The Network Environment: Then, Now, and in the Future

Traditionally, there have been two different market segments, voice and data, with differing characteristics and requirements. Two distinct networks, the PSTN and the Internet, have existed for voice and data traffic. For voice communications, the sequenced delivery of voice samples with latency less than a few hundred milliseconds is essential, but absolute data integrity is not necessary. Under these circumstances the guaranteed-bandwidth, circuit-switched infrastructure of PSTN was appropriate. For data networking, information integrity is paramount and the real-time transfer of this information is of less importance.

### The Development of Integrated Network Architectures

The wish to develop a single integrated network architecture capable of accommodating all types of services was first addressed with the development of the integrated services digital network (ISDN). ISDN, introduced in the early 1980's, was meant to provide customers with unified access to a range of digital services offered by carriers. The limitations of ISDN due to its channelized narrowband characteristics caused the introduction of broadband ISDN (B-ISDN), which required broadband switching and multiplexing techniques. By the late 1980s advances in technology made high-bandwidth information transport networks a realistic prospect. The technology and standard to provide this broadband switching and multiplexing capability was ATM, where fixed-size cells are used to transport information, preferably over high-speed and jitter-free synchronous fiber transport using optical network (SONET/SDH) standards.

### The Need for Higher Bandwidth and Differentiated Services

Deregulation and de-monopolization have caused a surge in the variety of services and information technologies, resulting in the most dynamic communication and computing network environment in history. The rapid growth of the Internet, as well as intranets and extranets, has caused a tremendous need for higher bandwidth both in the access and backbone networks. In addition, the relaxation of regulatory constraints is enabling network and service providers to offer various multimedia services with differing quality-of-service (QoS) categories.

Building the high-capacity optical backbone network of the future relies on innovative solutions incorporating cutting-edge technologies coupled with sound networking principles. Today's time division multiplexed (TDM)-based transport networks have been designed to provide an assured level of performance and reliability for the predominant voice and private-line services. SONET/SDH has been widely deployed in the current transport infrastructure, providing high-capacity transport, scalable to gigabit-per-second rates, with excellent jitter, wander, and error performance for 64 kb/s voice connections and private-line applications.

In contrast, the IP networks generally lack the means to guarantee high reliability and predictable performance. The best-effort service provided by most packet-based data networks, with unpredictable delay, jitter, and packet loss, is the price paid to achieve maximum link utilization through statistical multiplexing. To overcome the shortcomings of IP networks for carrying time-sensitive packets, such as voice, new protocols like real-time protocol (RTP), multiprotocol label switching (MPLS), resource reservation protocol (RSVP), and differentiated services (DiffServ) have been developed.

RTP is a framing protocol that uses the basic transport service of UDP but adds the required sequence numbers and timestamps to ensure ordered delivery and timely playback of real-time voice and video streams. RTP complies with ITU H.323 [12] for voice and video over packet networks; as a result, it has been widely adopted for multimedia teleconferencing. With MPLS, which is applicable to any network layer protocol, the mapping from packet headers to a bit stream is performed just once as the packet enters the network.[13] The stream to which the packet is assigned is encoded with a short fixed-length value known as a *label*. The label is used as an index into a routing table, which specifies the next hop and a new label. The old label is replaced with the new label and the packet is forwarded to its next hop. This is very similar to attaching an ATM virtual circuit label to an IP packet.

RSVP is an IP signaling protocol that enables the reservation of bandwidth to a particular IP destination and includes support for multicast applications.[14] RSVP is effective only if it is implemented end-to-end, including all intervening points. It ties up network resources and has scalability issues. Furthermore, only two service classes—controlled load and guaranteed service—are defined. DiffServ is proposed as a mechanism to enable certain users to receive a better service without the large-scale change necessitated by the widespread deployment of RSVP.

The surging demand for high-bandwidth and differentiated data services is challenging this dual architecture model of TDM-based transport and best-effort packet data networks.[15]

### Lucent's View of Network R/Evolution: Network of Networks

The Lucent vision is to provide a network of networks that will work together to deliver future services seamlessly and reliably. The elements of this vision are:

- *One network* that will be characterized by:
  - Coexistence and convergence of data networking and voice networking;
  - Optical backbone networking;
  - Multiple protocols and services support, including frame relay, IP, ATM, and virtual private networking;
  - A variety of broadband access mechanisms; and
  - Flexible bandwidth and service capabilities.
- *Open software platforms for network and service management* that will make possible:
  - A system that greatly simplifies operations by managing broadband pipes (rather than individual trunk groups);

- Lowering of OAM&P costs through a common packet backbone;
- Utilization of topology management and flow-through provisioning techniques;
- Creation (operationally) of a single network switch; and
- A new style of network management based on active directories and policy managers.

To satisfy new network requirements, Lucent has formulated a coherent strategic view of how a network satisfactorily delivers a full range of services to customers and proposed the Lucent Network Architecture—a data-centric optical transport network architecture[16]—as shown in **Figure 4**. Next-generation network architectures for cost-effective, reliable, and scalable evolution will employ both transport networking and enhanced service layers, working together in a complementary and interoperable fashion. Building a network capable of providing QoS controls requires accommodation of four protocols: Internet protocol (IP), asynchronous transfer mode (ATM), synchronous transfer mode (STM) such as SONET or SDH, and dense wavelength division multiplexing (DWDM).[16]

Transport networking enables the service layers to operate more effectively, freeing them from the constraints of physical topology and allowing them to focus on the challenge of meeting service requirements. Hence, optical transport networking will provide a unified, optimized layer of high-capacity, high-reliability bandwidth management, creating the so-called "optical data networking" solutions for higher-capacity data services with guaranteed quality.

In this architecture, integrated voice and data traffic is carried over the ATM/IP core network. Multiprotocol devices are connected via DWDM systems. The optical core will evolve over time to utilize optical add/drop multiplexers and cross-connect systems. Below is a brief description of the architectural elements in Figure 4.

**ATM core network.** Some public service providers are planning to use their ATM networks to support integrated voice and data traffic. ATM is a high-speed connection-oriented switching technology that is intended to support the concurrent transmission of all services—voice, video and data traffic—by segmenting their streams into small fixed-size of cells of 53 bytes. Voice traffic will be routed through an end-office switch to a tandem/toll switch, then carried through the constant bit rate of ATM adaptation layer 1 (CBR AAL-1) circuit emulation virtual circuits or variable-bit-rate real time of ATM adaptation layer 2 (VBR-rt AAL-2) virtual circuits in the ATM backbone network. In this architecture, a function converts between TDM and ATM voice traffic, with possible employment of voice compression and silence suppression. This function can be an integral part of either the tandem/toll switch or the ATM switch, or it can be implemented by a separate network element in the ATM network. In Figure 4, the 7R/E gateways, as part of the end offices, support a full range of packet and circuit connections to the core packet network. The lower half of Figure 4 indicates that the customer premises equipment is connected to an end office, that is, a 7R/E node, which in turn is connected to an ATM data edge MUX/switch, then connected to an ATM core switch. It is possible that the toll/tandem switch can be collocated with an ATM edge multiplexer (MUX)/switch. Thus, voice traffic may be aggregated with data traffic in the same ATM edge MUX/switch before being transported to the backbone network.

**IP core network.** Some Internet service providers are considering providing voice service over their IP backbone network. In this architecture, voice traffic from residential customers and/or businesses will be first routed through a 7R/E node, then carried through the IP backbone networks. In this case, the lower half of Figure 4 is the same as the IP network architecture supporting data services. Voice traffic is routed through a 7R/E node (an end office), encapsulated in RTP/UDP/IP, carried to the IP edge router, and then transferred through the IP core network.

**IN/SS7 network.** For the ATM backbone network, a signaling gateway performs the translation between Signaling System 7 (SS7) and ATM network-to-network interface (NNI) signaling protocols for the dynamic set up of switched virtual circuits (SVCs) for carrying voice traffic. The same signaling gateway will be able to control the dynamic allocation of network servers and

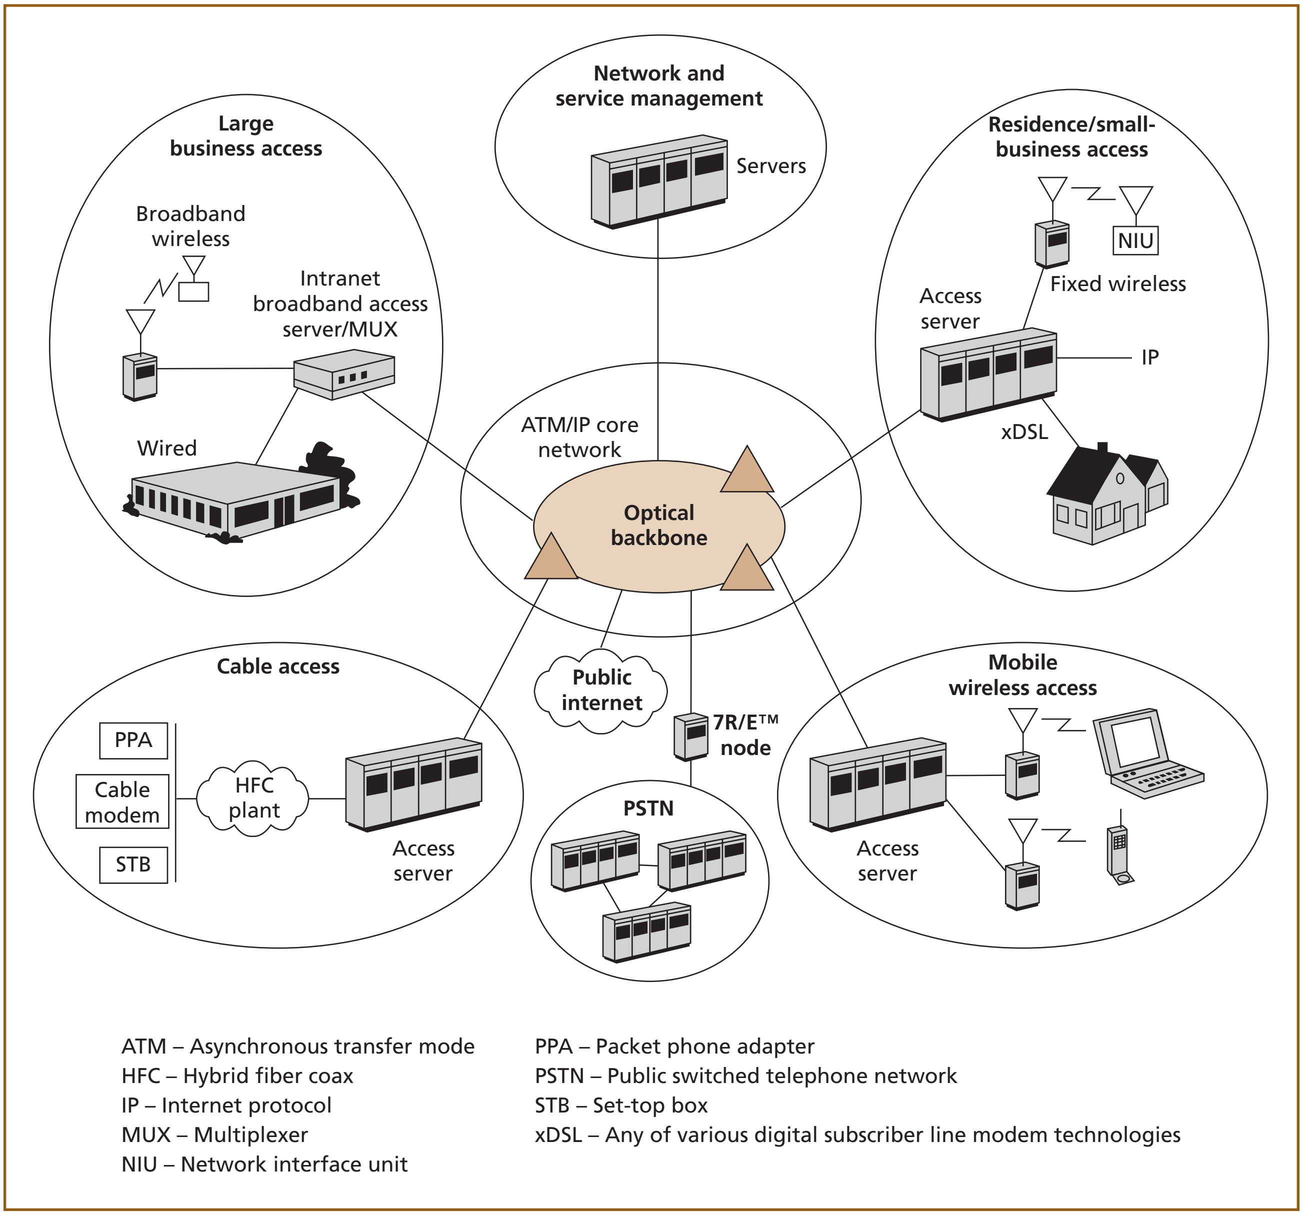


***Figure 4.***
***Lucent's Network Vision Architecture.***

resources. The signaling gateway can be part of the tandem/toll switch.

For the IP core network, a gatekeeper performs the translation between SS7 and the IP session management protocol to set up sessions dynamically for carrying the voice traffic. The gatekeeper, Internet telephony server (ITS) gateway, and edge router functions can be combined with the voice edge function performed by a 7R/E gateway node.

**Broadband access.** Various technologies are used for access to the backbone network by various types of end users. These range from traditional circuit-switched narrowband access to PSTN to broadband wireless access for mobile users. These technologies include xDSL, HFC, and the passive optical network (PON).

New access nodes are needed to provide integrated multi-service, multi-technology access to the backbone network. These edge nodes may need to perform conversion between circuit and packet domains. Some new telecommunications carriers will build packet-based networks from the ground up and therefore will not have to deal with integration of legacy voice/TDM networks. On the other hand, the vast majority of carriers must integrate their current networks with the emergent network infrastructures based on the optical layer.

A family of digital subscriber line (DSL) technologies, collectively referred to as xDSL, allows a few megabits per second to be carried over the existing copper plant (unshielded twisted pair) in a relatively fast and cost-effective way. The recently proposed ADSL-Lite runs at about 1.5 Mb/s, and asymmetric digital subscriber line (ADSL) at 6 Mb/s. Very high bit-rate digital subscriber line (VDSL) operates at higher data rates. Except for high bit-rate digital subscriber line (HDSL), xDSL is predominantly seen as an overlay technology, operating on the same pairs as narrowband services such as "plain old telephone service" (POTS) or ISDN. This is achieved by using the frequency spectrum beyond that used for POTS on the copper wire. The narrowband signal uses the low part of the spectrum, while the broadband signal uses the higher frequencies. The signals from the two are combined using a POTS splitter. At the customer premises, a POTS splitter allows these separate signals to be passed to an xDSL modem and to the telephones. Depending on the type of DSL and the bandwidths required by the customer, different deployment architectures are possible for this technology.

Another widespread system that is being used for access to high-rate digital services is the coaxial cable-based distribution used by cable television operators. A cable modem allows a few megabits per second of data to be carried over hybrid fiber coax (HFC) access systems. The frequency spectrum is split both at the customer site and at the cable headend, thereby allowing access to community antenna television (CATV), optical backbone, and optionally, PSTN. Amplifiers must be placed regularly along the cable to deliver satisfactory television signals to customers. Such systems have access to a majority of homes in the United States, but are not likely to be available for businesses that are not located in a residential area.

CATV systems are divided into 6 MHz channels in North America, corresponding to the analog television bandwidth. Each such channel can carry 30 or 40 Mb/s with good performance in the downstream direction (to the customer). However, in the upstream direction (from the customer) the performance is not nearly that satisfactory. In fact, most CATV systems today are strictly one way, downstream only. Upstream capacity is added by using the lower part of the frequency spectrum, below about 45 MHz. Aside from the smaller available bandwidth, the lower frequencies are more susceptible to interfering signals. Another problem associated with the shared use of the upstream channel is the cumulative addition of noise originating from all customer locations. For these reasons, only relatively low rates can be carried upstream, and the system is best used for asymmetric services, such as Internet access. Such services are available today.[17]

There is now great interest in using wireless technology to provide higher-rate digital services. The next generation of cellular service (the so-called "third generation") will provide digital service at rates up to 5 Mb/s. For residential and small business access, an architecture similar to digital cellular radio, but without handoffs, is being considered for providing multi-Mb/s access. Base stations can provide a coverage radius of 5 to 10 km, depending on the terrain. Operation in the 2 to 5 GHz region would be most likely. The use of steered beam antenna techniques should provide sufficient antenna gain to permit operation at approximately 1 watt of power at both the base station and remote terminal. Current research indicates that signal processing techniques show promise of improving the antenna gain from a multi-antenna array. The primary use of this system would be for Internet access as well as voice service. Introduction of this type of access service may be aided by the use of existing cellular physical sites. For large businesses, point-to-point access at up to 100 Mb/s is envisioned. Using higher frequencies, up to 28 GHz, allows the use of available wide spectral bandwidths, at the expense of shorter distance coverage due to rain attenuation. Distances of 2 to 3 km are feasible. The system can be upgraded to provide point-to-multipoint service using antenna steering or combining.[17]

Convergence of fixed and mobile access is concerned with the provisioning of network and service capabilities that are independent of access techniques.[18]

Optical access options exist as well. In the near term, a technology similar to the packet-based optical core, but with lower capacity, will provide metropolitan transport and business access to the core network described above. A promising technique to provide some degree of sharing the use of fiber among multiple customers is the PON. In the simplest configuration, only one wavelength is used in each direction. DWDM is a further enhancement to increase capacity. The end fibers in a PON may supply individual customers—particularly small- and medium-sized businesses—or may terminate in optical network units (ONUs) that serve small clusters of homes in a residential area. The latter case in frequently referred to as "fiber to the curb," where the individual homes may be accessed by wire pairs, possibly using VDSL technology. Passive coaxial cable distribution is another possibility. Because the number of homes served by each ONU is small (4 to 8), sharing of capacity is readily achieved. Ultimately, the ONUs in a PON may serve individual homes, achieving "fiber to the home." An issue in this case, other than cost, is supplying reliable power if the system is to provide primary telephone service.[17]

## A Practical Vision for Managing Networks

How is a global network made of diverse components and providing diverse services going to be managed? Realizing network management architectures involves a careful balancing of a web of tradeoffs among a set of critical factors to meet associated technical and networking challenges. Implicit in the balancing exercise is the ubiquitous cost/functionality tradeoff—that is, at what price functionality? Rather than a single solution for all applications, a range of network management architectures will arise as each market segment applies its unique priorities to make fundamental architecture tradeoffs. In this section we briefly describe the new network management requirements and the techniques to meet them.

### New Management Requirements

The objectives of the traditional NMS/OS have been to reduce operations complexity and cost by increasing flow-through. Reliability, security, and effective presentation techniques have always been desirable requirements. But the focus of traditional network management, both in carrier networks and enterprise networks, has generally been at the element level, relying on communications protocols and element MIBs. As networks grow in size and as services become more complex, scalability issues arise and management at the network level becomes more important.

The desire for defining and managing services quickly over the network infrastructure has created a requirement for more sophisticated network management techniques. To support services such as packet voice, virtual private networks, and multimedia services, network management systems must be capable of providing mechanisms for managing QoS and SLAs. Moreover, these systems must respond to customers, who more and more require the ability to view their logical subnetworks and to be billed commensurate with the provided QoS.

**Table I** shows some services and their corresponding network and management implications.

### Meeting Management Requirements

**Table II** shows some of the common practices in traditional network management and their corresponding shortcomings. To meet these challenges, the traditional network management platforms and techniques need to be modified and enhanced. Rapid introduction of new network elements, protocols, and applications require a highly flexible and responsive management architecture. Some elements of this architecture are described below.

**Distributed network management.** The traditional centralized architecture for management has reliability and scalability problems. The hierarchical management architecture somewhat ameliorates the situation, but this architecture still is not sufficiently flexible and responsive. Distributed management architectures solve the problems of reliability and scalability; however, they are generally more complex and present new challenges in terms of security, concurrent programming, and synchronization.

Supporting service requirements on an individual device basis is too cumbersome and not cost effective. To deal with the increased complexity of network

**Table I. Services and their network management implications.**

| Service | Network implications | Management implications |
|---|---|---|
| Voice, video, multimedia | • Latency of no more than 200 ms<br>• Bandwidth of 6 to 300 kb/s | • Bandwidth management<br>• QoS management |
| Electronic commerce | • Continuous availability<br>• Trust mechanisms | • Security management |
| Consumer lifeline services | • Continuous availability<br>• Low and high bandwidth<br>• Low cost | • Efficient configuration management<br>• Fault management |
| Network access (VPN, wireless, RADIUS) | • Bandwidth of at least 64 kb/s<br>• Security<br>• Security management | • Efficient configuration management<br>• Fault management |

VPN – Virtual private network
RADIUS – Remote authentication dial-in user service

**Table II. Practices in network management and corresponding shortcomings.**

| Practice | Shortcoming(s) |
|---|---|
| Use of centralized and hierarchical architecture | Slow response |
| No agent-agent communication | Restrictive |
| Scope fixed at element level | Static and inflexible, scalability problems |
| Predefined MIBs | Inflexible |
| Web-based interface | Failure to reduce complexity of network manager's job |
| No application-level QoS management | Outdated ways of managing virtual overlays |
| Application of policy-based management primarily to single points for authorization, access control, and security | Limited scope; must be extended to other areas of management |
| Shortage of management applications | Restrictive; no automated QoS management possible |

MIB – Management information base
QoS – Quality of service

management as a result of the growing variety of service requirements, the focus must be shifted from manager-agent communication protocols to information management and distribution. This has resulted in the growing popularity of CORBA-based distributed management[19] and directory-enabled networking (DEN) paradigms.[20]

CORBA provides a flexible platform for manager-agent and manager-manager communication. It contains mechanisms for communication between objects through an object request broker (ORB), an intermediary between client and server. The ORB has the capability of selecting the server that is best suited for satisfying the client's requirement, without the client being aware of details about the server. CORBA supports distributed object computing in a variety of programming languages. CORBA Interface Definition Language (IDL)* defines services offered by a particular distributed object. Being vendor independent, CORBA has gained tremendous popularity for management in a heterogeneous network/systems environment.

The directory enabled networks (DEN) vision is to provide network-enabled application information from a single X.500-based directory, using lightweight directory access protocol (LDAP) for access. The proposed information model is Common

Information Model (CIM) put out by the Distributed Management Task Force.[21] The CIM/DEN paradigm is gaining broad support in both carrier and enterprise network management.

**Policy-based management.** Network management techniques relying only on communications protocols and device-oriented data stores are not adequate for coordinating and configuring networks with complex layered services. Policy-based management provides a possible solution.[22,23] A *policy* is a formal representation of information affecting network/component behavior, but specified independent of components. The enterprise-wide directory advocated by DEN is well suited for policy-based management that allows coordinating and configuring networks with complex services. The policy applications, however, currently have a limited scope: their current focus is on QoS-based congestion management, user-access security management, and configuration management. The applications need to be extended to other areas of management. Policy-based management and related issues are covered in detail in one of the papers in this issue.[24]

**Intelligent agents in network management.** In both SNMP-based enterprise network management and CMIP-based carrier network management, management systems interact with management agents running on network elements. *Intelligent agents* are autonomous software entities capable of performing functions and making decisions on behalf of a user or another program. A client, such as an NMS, can delegate certain tasks to an intelligent agent, which will perform the task with minimal involvement from the client. Using intelligent agents on network elements reduces the processing burden on the NMS and the volume of management traffic. Intelligent agents can be used, for example, in *management by delegation*, where they execute programs asynchronously to perform functions like event correlation, freeing the NMS to perform other tasks. One of the papers in this special issue provides a tutorial on intelligent agents and how they may be utilized in network management.[25]

**Management of convergent networks.** Figure 5 shows a conceptual architecture for integrated multi-domain management. With packet voice gaining tremendous momentum in recent years, the challenge is to develop efficient ways of managing multi-domain networks. The network elements will include both TDM and packet components and interfaces, and the element managers must be able to manage both domains. Furthermore, an inter-domain management function is needed to reconcile the differences in paradigms used to manage the legacy TDM elements and the new multi-domain elements. In the service management layer, issues of end-to-end QoS management and SLA management need to be addressed.

## A Snapshot of Trends and Techniques

This issue of the *Bell Labs Technical Journal* outlines some trends, techniques, and applications to manage the evolving network. It describes how some of the network management challenges mentioned in the previous sections are being addressed, highlighting some network management activities that are taking place within Lucent. The papers presented in this issue fall into four major areas:

- Network management information models and applications,
- New techniques and paradigms,
- Management of large networks, and
- Management of diverse and converged networks.

The coverage of innovations here is by no means comprehensive or complete; it is only a sample of Lucent's activities in facing the challenges of network management.

### Network Management Information Models and Applications

Network performance monitoring is an important part of effective network management. Three papers in this issue address this important function for next-generation networks. "Real-Time Performance Monitoring and Anomaly Detection in the Internet: An Adaptive, Object-Driven, Mix-and-Match Approach" by Ho, Macey, and Hiller[26] describes a powerful framework capable of automatic and proactive fault and anomaly detection and performance management in IP networks and for IP services. This platform would facilitate the introduction and automatic management of emerging new IP applications.

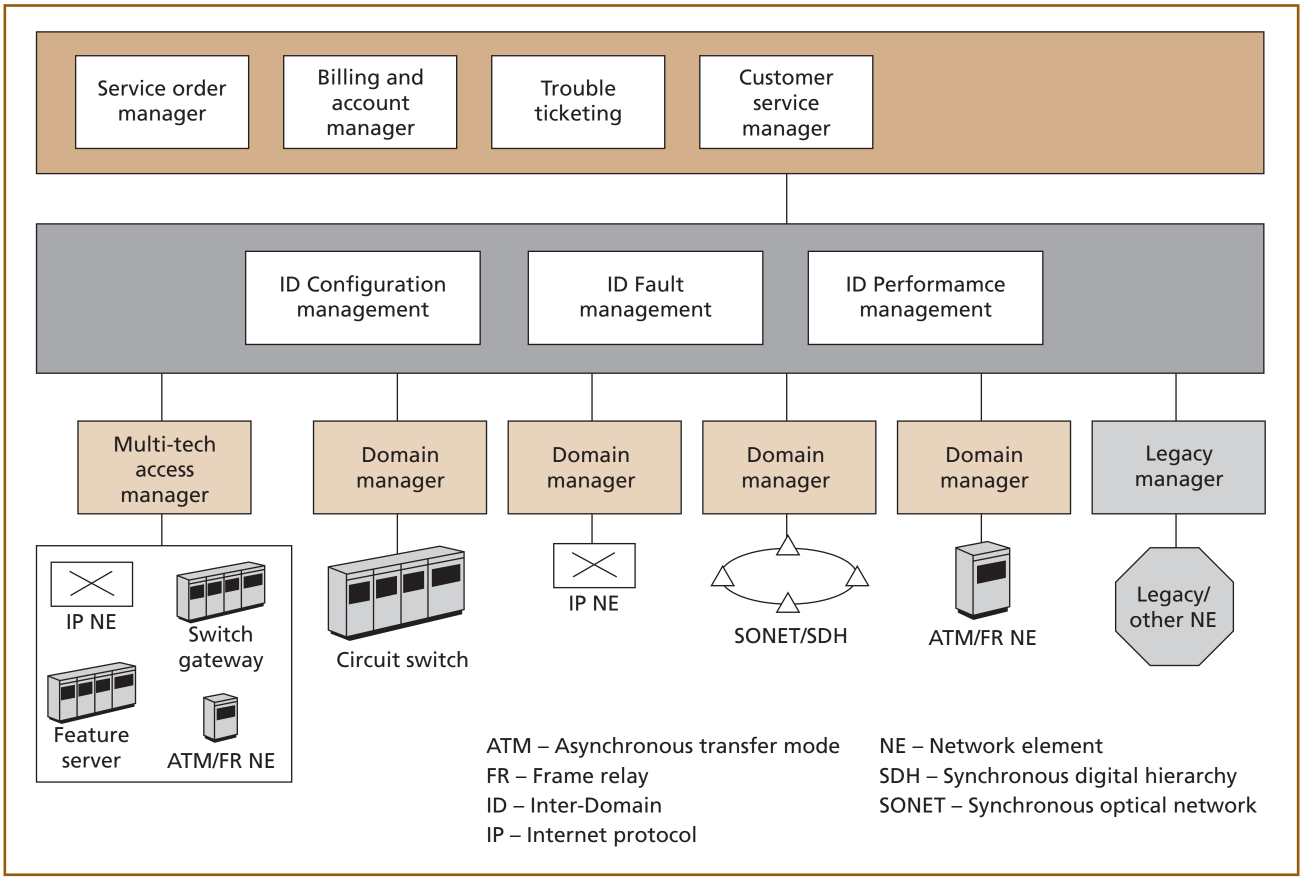


***Figure 5.***
***A conceptual architecture for integrated multi-domain management.***

"Switch Monitoring—The New Generation of Monitoring for Local Area Networks" by Romascanu and Zilbershtein[27] describes a method to organize data-collection reports in switches as well as extensions to handle layer-3 and above statistics at the switch level. This paper reviews the development of the current and future standards in the area of monitoring LAN switches. Because standard remote monitoring (RMON) probes are insufficient to monitor switched LANs, Lucent has developed new methods to perform switch monitoring (SMON).

Building network management applications whose scope is network-wide has been a largely unaccomplished goal of the communications industry. Two potentially complementary and hopefully converging efforts under way in the industry—the CIM and DEN initiatives—are now both under the auspices of the Distributed Management Task Force, as mentioned earlier in this paper. The former initiative defines object-modeling formalisms for specifying a network information model, as well as a set of base schemas. The latter defines a model schema that is intended for storage using LDAP-accessible directory software technology. The paper by Nissenbaum, "A Directory-Hosted Network Management Information Model: A CIM/DEN Implementation,"[28] describes a prototype of Lucent's implementation of a CIM/DEN system.

### New Techniques and Paradigms

Several papers in this issue deal with rather new paradigms in network management. Two papers focus on policy-based network management techniques. The paper "Policy-Based Management for IP Networks" by Stevens and Weiss[24] contrasts conventional approaches to network management with a new, more cohesive approach. The authors describe policy-based network management and techniques for implementing it. They also describe the challenges and opportunities in implementing effective policy-based network management. The paper "Policy-Based Network Load Management" by Hossain et al.[29] describes the design and implementation of the prototype of a policy-based load management system. The technique used in the prototype can intelligently manage incoming Internet service requests among a group of servers, resulting in QoS improvement.

Bieszczad et al., in "Management of Heterogeneous Networks with Intelligent Agents,"[25] provide a brief tutorial on intelligent agents and describe how using intelligent agents can address some of the shortcomings of the current approaches to network management and help in management of heterogeneous networks. The authors address ongoing Lucent work on an intelligent agent platform, LucINA, and its applications.

### Management of Large Networks

Unless properly engineered, network management of large SONET/SDH networks can severely stress the embedded data communications network (DCN), resulting in less responsive communication between the network manager and network elements. Budka et al. in their paper, "Engineering Large SONET/SDH Management Networks,"[30] explore the demands large networks place on the SONET/SDH operations network and the ways an operations network can be made robust. The authors present the characteristics of large SONET/SDH management networks and describe the engineering principles for improving their stability and performance. They also describe how these principles are incorporated during the design and testing of some Lucent product families.

The management of today's optical transport networks typically includes some or all of FCAPS management. Carrying this management information across the network requires the presence of a management communications network (MCN), also referred to as the DCN. The MCN itself needs FCAPS management. For instance, when an optical fiber cut is detected on a fiber carrying data at a rate of hundreds of gigabits per second, it is imperative that the detection of the fiber cut be reliably and quickly relayed to the management systems through the MCN. A well-managed MCN can give a better guarantee of doing this than an unmanaged MCN.

The availability of management for the MCN is very sparse in today's management systems, and "Managing the Management Communications Network in Optical Transport Systems" by Shankar and Chatterjee[31] is a right step toward addressing the problem. This paper describes algorithms for planning, designing, and analyzing the MCN. We envision this work as forming the basis for a suite of applications that can improve the management of the communications infrastructure in optical transport networks.

The software that supports telecommunications network elements encompasses a multitude of different functions. Besides the features and functionality that support service, call, and connection control, there is a vast amount of software that is required to support network management functionality. This, although frequently overlooked during initial system specification, often comprises the majority of the effort during software design and construction. "Reusable Management Frameworks for Third-Generation Wireless Networks" by Choy et al.[32] describes some work that has been done to analyze common management functionality using design pattern and framework technology. This paper describes how these assets are being used in the management of CORBA and SNMP-based third-generation wireless networks, supporting both voice and high-speed data services.

### Management of Diverse and Converged Networks

Rapid evolution of the telecommunications industry has resulted in the convergence of voice and data transport over many diverse architectures and technological domains such as TDM, ATM, SONET, SDH, frame relay, DWDM, and IP. Lucent is working to provide an inter-domain network management solution such that diverse technologies from multiple vendors can be managed in an integrated manner and the differences among the heterogeneous network equipment become transparent. "A Network Management System Solution for Diverse and Converged Voice and Data Networks" by Bagga et al.[33] gives an overview of the proposed Inter-Domain Management System Architecture.

Today's multiple network architectures supporting voice and data services can be consolidated using packet technologies. This should result in operational savings, but it introduces the challenges of managing multi-domain integrated networks. "Network Element Management Architecture for the 7R/E Trunk Access Gateway" by Whang, Jain, and Chan[34] addresses the challenges of managing an integrated gateway element and describes the system requirements, features, information models, manager/agent architectures, and major components of the element management system (EMS) for the 7R/E Trunk Access Gateway. The authors describe how the designed EMS meets the challenges.

Network management has traditionally been centralized and executed by a select group of technicians. Management of today's multi-service multi-vendor networks requires new architectures and techniques and often involves different platforms that need to be integrated. In "NetCare® Network Management Services—Managing Multi-Vendor Networks," Denison and Giovannetti[4] present the successful approach of NetCare Network Management Services to multi-vendor network management and describe how, through these services, the status of every component in the network is tracked, regardless of the type of network or the manufacturer.

In today's IP networks, network operators lack the tools to monitor use of network services by individual customers, to measure the type and quality of service these customers receive, and ultimately to establish SLAs and bill for the services rendered. Blott et al. in "User-Level Billing and Accounting in IP Networks,"[35] present an approach to accounting in IP networks based on a special-purpose network probe, which the authors term a *NetCounter*. The key functionality of a NetCounter is real-time, in-network correlation of network traffic with the individual users that generated it. This approach, adopted by Lucent, has several technical advantages over alternative approaches: NetCounter achieves substantial in-network aggregation, reducing the volume of usage data generated by between two and four orders of magnitude (when compared to flow-logging systems). In addition, it captures usage data at the level of individual users and records detailed, user- and service-specific performance metrics.

The network and computer vulnerabilities in Lucent, if left unidentified or unmitigated, will result in serious and costly compromises to Lucent's intellectual properties. Given the voluminous nodes and hosts in the Lucent intranet, it is not operationally feasible to scan the entire network. The final paper in this issue, "Managing Cyber Security Vulnerabilities in Large Networks" by Chang et al.,[36] describes a statistical sampling and analysis methodology combined with a network and host security discipline for developing Lucent's cyber security profile in an effective and efficient manner. Through root cause analyses, as explained in this paper, Lucent developed a focused plan for mitigating vulnerabilities effectively and efficiently. These patent-pending methodologies will become the enablers of cyber security management in a large networked environment.

### Other Related Lucent Activities

Management applications that control the network based on monitored information are rare. However, Bell Labs has recently built automated management applications that perform advanced monitoring and control functions on an end-to-end network/service level. The IP Network Configurator is an example of a novel management tool that configures a network taking into account the interdependencies between network elements.[37,38] In addition, the NetViz system provides a framework for constructing interactive visualization of complex hierarchical and time-varying networks.[39]

Novel approaches to management are required to meet the twin requirements of simplicity and more powerful functionality. Two unique Bell Labs solutions, Active Bell Labs Engine (ABLE)[40] and Extensible Network and Application Management Instrumentation (XNAMI),[41] achieve this goal by offering solutions that are simple, powerful, flexible, and capable of being implemented in a distributed setting.

## Conclusion

We have reviewed the evolution of network management as well as the essentials of network management and its functions and components. We have also highlighted the requirements for managing the network as it continues to evolve. We hope this cross-sectional view of network management activities will inspire continued excellence in network management research and development at Lucent Technologies.

### *Trademarks

CORBA is a registered trademark and Interface Definition Language (IDL) is a trademark of the Object Management Group.

OpenView is a registered trademark of Hewlett-Packard Company.

### References


1. W. Stallings, *SNMP, SNMPv2, SNMPv3, and RMON 1 and 2*, 3rd ed., Addison Wesley, Reading, Mass.,1999.
2. International Telecommunication Union, "Common Management Information Protocol Specification for CCITT Applications," Rec. X.711, 1991, http://www.itu.int/itudoc/itu-t/rec/x/index.html
3. International Telecommunication Union, "Common Management Information Service Definition for CCITT Applications," Rec. X.710, 1991, http://www.itu.int/itudoc/itu-t/rec/x/index.html
4. B. A. Denison and F. Giovannetti, "NetCare® Network Management Services—Managing Multi-Vendor Networks," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 225–236.
5. International Telecommunication Union, "TMN Management Functions," Rec. M.3400, 1993, http://www.itu.int/itudoc/itu-t/rec/m/index.html
6. International Telecommunication Union, "Management Framework for Open System Interconnection (OSI) for CCITT Applications," Rec. X.700, 1993, http://www.itu.int/itudoc/itu-t/rec/x/index.html
7. G. Lin, J. D. Price, and T. K. Srinivas, "Network Information Models and OneVision Architecture," *Bell Labs Tech. J.*, Vol.3, No.4, Oct.–Dec. 1998, pp. 208–221.
8. International Telecommunication Union, "Information Technology—Open Systems Interconnection—Systems Management Overview," Rec. X.701, 1992, http://www.itu.int/itudoc/itu-t/rec/x/index.html
9. International Telecommunication Union, "Information Technology—Open Systems Interconnections—Structure of Management Information—Part 1: Management Information Model," Rec. X.720/ISO 10165-1, 1992, http://www.itu.int/itudoc/itu-t/rec/x/index.html
10. International Telecommunication Union, "Information Technology—Open Systems Interconnections—Structure of Management Information: Guidelines for the Definition of Managed Objects," Rec. X.722/ISO 10165, 1992, http://www.itu.int/itudoc/itu-t/
11. International Telecommunication Union, "Principles for a Telecommunications Management Network," Rec. M.3010, 1993, http://www.itu.int/itudoc/itu-t/rec/m/index.html
12. International Telecommunication Union, "Packet-Based Multimedia Communications Systems, ITU-T Rec. H.323, Feb. 1998, http://www.itu-int/ITU-T/index.html
13. E. C. Rosen, A. Vishwanathan, and R. Callon, "MultiProtocol Label Switching Architecture," IETF Internet-Draft, IETF, Aug. 1999, http://ietf.org/internet-drafts/draft-ietf-mpls-arch-06.txt
14. R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, "Resource ReSerVation Protocol (RSVP) Version 1—Functional Specification," IETF RFC 2205, Sept. 1997, http://www.ietf.org/rfc/rfc2205.txt
15. R. C. Alferness, P. A. Bonenfant, C. J. Newton, K. A. Sparks, and E. L. Varma, "A Practical Vision for Optical Transport Networking," *Bell Labs Tech. J.*, Vol.4, No.1, Jan.–Feb. 1999, pp. 3–18.
16. D. C. Dowden, R. D. Gitlin, and R. L. Martin, "Next-Generation Networks," *Bell Labs Tech. J.*, Vol.3, No.4, Oct.–Dec. 1998, pp. 3–14.
17. K. G. August, V. B. Lawrence, and B. R. Saltzberg, "An Introduction to Future Communications Services and Access," *Bell Labs Tech. J.*, Vol.4, No.2, April-June 1999, pp. 3–20.
18. F. G. Harrison and S. R. Hearnden, "The Challenge to Realize Convergence of Fixed and Mobile Communications," *Electronics & Commun. Engineering J.*, Vol. 11, No. 3, June 1999, pp.164–168.
19. Object Management Group, *The Common Object Request Broker: Architecture and Specification*, December 1993. http://www.omg.org/corba/
20. Directory-Enabled Networking (DEN), http://www.dmtf.org/pres/rele/denfaq.html
21. W. Bumpus, "DMTF Expands Its Role in Developing Information Model Standards," *J. of Network and Systems Management*, Vol. 6, No. 3, Sept. 1998, pp. 357–360.
22. M. Sloman, "Policy Driven Management for Distributed Systems," *J. of Network and Systems Management*, Vol. 2, No. 4, 1994, pp. 333–360.
23. E. C. Lupu and M. Sloman, "Towards a Role-Based Framework for Distributed Systems

Management," *J. of Network and Systems Management*, Vol. 5, No. 1, 1997, pp. 5–30.
24. M. L. Stevens and W. J. Weiss, "Policy-Based Management for IP Networks," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 75–94.
25. A. Bieszczad, P. Biswas, W. Buga, M. Malek, and H. Tan, "Management of Heterogeneous Networks with Intelligent Agents," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 109–135.
26. L. L. Ho, C. J. Macey, and R. Hiller, "Real-Time Performance Monitoring and Anomaly Detection in the Internet: An Adaptive Object-Driven, Mix-and-Match Approach," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 23–41.
27. D. Romascanu and I. E. Zilbershtein, "Switch Monitoring—The New Generation of Monitoring for Local Area Networks," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 42–54.
28. H. D. Nissenbaum, "A Directory-Hosted Management Information Model: A CIM/DEN Implementation," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 55–74.
29. A. Hossain, H. F. Shu, C. R. Gasman, and R. A. Royer, "Policy-Based Network Load Management," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 95–108.
30. K. C. Budka, R. J. DiPasquale, B. W. A. Rijsman, and A. B. Sripad, "Engineering Large SONET/SDH Management Networks," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 136-154.
31. R. V. Shankar and A. K. Chatterjee, "Managing the Management Communications Network in Optical Transport Systems," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 155–170.
32. K.-C. Choy, M. S. Halkyard, S. Krishnamoorthy, and M. R. Panda, "Reusable Management Frameworks for Third-Generation Wireless Networks," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 171–189.
33. Y. S. Bagga, A. Chaudhury, C. Hird, T. Lee, W. Pretyka, D. Puseljic, S. Russo, R. Shafie-Khorasani, T. Tignor, and R. Ting, "A Network Management System Solution for Diverse and Converged Voice and Data Networks," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 190-203.
34. K. K. W. Whang, R. Jain, and P. Chan, "Network Element Management Architecture for the 7 R/E Trunk Access Gateway," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 204–224.
35. S. M. Blott, C. E. Martin, Y. J. Breitbart, J. C. Brustoloni, T. R. Gramaglia, H. F. Korth, D. M. Kristol, R. H. Liao, E. L. Scanlon, and A. Silberschatz, "User-Level Billing and Accounting in IP Networks," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 237-251.
36. E. S. Chang, A. K. Jain, D. M. Slade, and S. L. Tsao, "Managing Cyber Security Vulnerabilities in Large Networks," *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999, pp. 252-272.
37. B. Sugla and P. Krishnan, "IP Network Configurator: Advancing the State of the Art in IP Network Deployment and Operations," Lucent Technologies White Paper, November 1998.
38. http://www.lucent.com/OS/ipnc.html
39. T. He and S. Eick, "Constructing Interactive Network Visual Interfaces," *Bell Labs Tech. J.*, Vol. 3, No. 2, Apr.–June 1998, pp. 47–57.
40. D. Raz and Y. Shavitt, "An Active Network Approach for Efficient Network Management," The First Intl. Workshop on Active Networks (IWAN99), July 1999, Berlin, Germany, LNCS 1653, pp. 220–231.
41. A. John, K. Vanderveen, and B. Sugla, "An XML-based Framework for Dynamic SNMP MIB Extension," Proc. 10th IFIP/IEEE Intl. Workshop on Distributed Systems: Operations & Management (DSOM99) October 1999, Zurich, Switzerland, LNCS1700, pp. 107–120.

*(Manuscript approved January 2000)*

*SHERVIN ERFANI is a member of technical staff in the Data Communication Technologies Department at Bell Labs in Holmdel, New Jersey. He holds a combined B.S. and M.S. degree in electrical engineering from the University of Tehran in Iran and M.S. and Ph.D. degrees, also in electrical engineering, from Southern Methodist University in Dallas, Texas. Since joining Bell Labs, Dr. Erfani has been engaged in network management, service provisioning, network design and planning, and network security management. He has published more than 50 technical papers, holds a patent, and is the senior technical editor of the* Journal of Network and Systems Management. *Dr. Erfani also teaches courses at Stevens Institute of Technology in Hoboken, New Jersey.*

*VICTOR B. LAWRENCE is director of the Advanced*

*Communications Technology Center at Bell Labs in Holmdel, New Jersey. He holds both B.Sc. and Ph.D. degrees in electrical engineering from the University of London in the United Kingdom. At Bell Labs, he has held various assignments in the areas of signal processing, data communications, and exploratory development of transmission products and services. He is currently responsible for technology transfer, systems engineering, and exploratory development of multimedia systems over wire and wireless networks. An IEEE Fellow and a Bell Labs Fellow as well, Dr. Lawrence holds 17 patents and has published over 40 technical papers. For several of the papers, he has received special recognition. In addition, he is the author of a chapter in* Introduction to Digital Filtering, *co-editor of* Tutorials in Modern Communications, *and co-author of both* Intelligent Broadband Multimedia Networks *and* Engineering of Intelligent Communication Systems*.*

*MANU MALEK, a distinguished member of technical*

*staff in the Next-Generation Network Systems Engineering and CALA Network Planning Department at Bell Labs in Holmdel, New Jersey, holds a Ph.D. in electrical engineering and computer science from the University of California at Berkeley. Dr. Malek is currently working on operations and management planning for next-generation networks. He is the author, co-author, or editor of 7 books, the author or co-author of over 50 published technical papers, and a co-holder of 2 patents in multimedia and network management. An IEEE Fellow and an IEEE Communications Society Distinguished Lecturer, Dr. Malek is the founder and editor-in-chief of the* Journal of Network and Systems Management.

*BINAY SUGLA is currently department head of Network*

*and Services Management Research at Bell Labs in Holmdel, New Jersey. He joined Bell Labs after receiving a Ph.D. in electrical and computer engineering from the University of Massachusetts in Amherst. At Bell Labs, Dr. Sugla first designed and implemented CAPER, a visual parallel programming environment. He later worked on Network Flight Simulator, a real-time simulator of the dynamics of telephony network. For the past few years, he has been involved in the management of IP networks and services.* ◆