

◆ The Management Paradigm Shift: Challenges from Element Management to Service Management

Shervin Erfani, Victor B. Lawrence, and Manu Malek

In today's marketplace, the value of a network—whether public or private—is defined by the services it provides. This is a fundamental change from networks whose major value has been perceived as providing data communications. While operations support systems have served the needs of telecommunications providers for many years, major changes will be required in these systems to support this value shift. This issue of the Bell Labs Technical Journal, which emphasizes a management-paradigm-focused view of the network, provides a sample of many trends and challenges found within Lucent Technologies—from implementing element management systems to implementing service management systems. This paper, which presents a view of service management as managing the relationship between clients and applications in a heterogeneous network, serves as an introduction to the issue.

Introduction

As more low-cost and high-capacity communication between client and servers becomes available, the enterprises are becoming interested in implementing services that will help them streamline their processes, improve their applications, reduce their costs, and increase their revenue. At the same time, providing mere transport facilities is becoming less profitable for network providers. It has become evident that a larger part of the profit for network providers will be derived from the services they offer. To this end, Lucent Technologies has been exploring the underlying networking trends and the evolution of information processing technologies to provide a framework for the network full-service solution. This issue of the *Bell Labs Technical Journal* is devoted to the understanding of service management requirements and the corresponding components as related to the ongoing activities in Lucent business units. The issue is thus a follow-up to last year's issue on network management.¹

Lucent is committed to providing architectures for service management and a roadmap to achieve it. In this paper, we present a service management paradigm for the emerging network. We provide an end-to-end management view for the network infrastructure and the increasing variety of applications it supports for an enterprise's end-user community. We describe the network paradigm shift and the key components of this emerging networking environment. Having considered the environment, we then examine how well Lucent is positioned to support its customers. Lucent's position is shown to depend upon its many enabling technologies and products, some of which are described in the papers in this issue.

In the next section of this paper, we present an overview of the current trends in networking and services that have strong implications for the requirements of a service management platform. Subsequently, we address end-to-end network management, with emphasis on service management. We then

Panel 1. Abbreviations, Acronyms, and Terms

API—application programming interface
ASP—application service provider
ATM—asynchronous transfer mode
CCSP—CyberCarrier service provider
CMIP—common management information protocol
CMISE—common management information service element
CNM—customer network management
COPS—common open policy service
CSM—customer service management
DCSP—data center service provider
DSLAM—digital subscriber line access multiplexer
DSL—digital subscriber line
EMS—element management system
FCAPS—fault, configuration, accounting, performance, and security
GDMO—"Guidelines for the Definition of Managed Objects"
HFC—hybrid fiber-coaxial
IETF—Internet Engineering Task Force
IOP*—Internet Inter-Orb Protocol
ILEC—incumbent local exchange carrier
IP—Internet protocol
IPSS—IP service switch
ISDN—integrated services digital network
ITU—International Telecommunication Union
ITU-T—ITU Telecommunication Standardization Sector
LDAP—lightweight directory access protocol

MIB—management information base
MPLS—multiprotocol label switching
NMS—network management system
NSP—network service provider
OAM&P—operations, administration, maintenance, and provisioning
OMG—Object Management Group
OSI—Open Systems Interconnection
OSS—operations support system
PBX—private branch exchange
PCM—pulse code modulation
POTS—"plain old telephone service"
QoS—quality of service
RFC—Request for Comments
RFP—Request for Proposal
RTP—realtime transport protocol
SLA—service-level agreement
SMS—service management system
SNMP—simple network management protocol
SONET—synchronous optical network
TCP—transmission control protocol
TDM—time division multiplexed
TMN—Telecommunications Management Network
TOM—Telecom Operations Map
UDP—user datagram protocol
VPN—virtual private network
WAN—wide area network
WDM—wavelength division multiplexing
xDSL—any of various DSL technologies

report on the current status of key management technologies, which motivates the Lucent's activities described in this issue. Finally, we summarize the papers in this issue, which we hope will smooth the way for more advancement in upcoming technologies.

Service and Networking Trends

Today, most service providers still run two parallel networks—one for voice and one for data services. Applications are converging on top of these networks. Intelligent services in the voice network create additional value for service providers. As in the voice network, higher-layer services in the data network influence the lower layers of the network infrastructure.

The converged network of the future will be more feature rich than today's network, and it will be driven

by the customers' business needs. The so-called "optical communication networks for the next-generation Internet" will provide a client/server relationship among enterprise clients and various applications residing on servers in a number of data centers through a high-capacity heterogeneous network. The network will provide essential applications and services to the clients it supports.

This logical separation of the network and services will enable the applications to operate more effectively. It will free them from the constraints of physical topology and allow them to focus on the challenge of meeting service requirements, such as dynamic allocation of resources, easy addition of new services and/or applications, and efficient navigation to desired services.

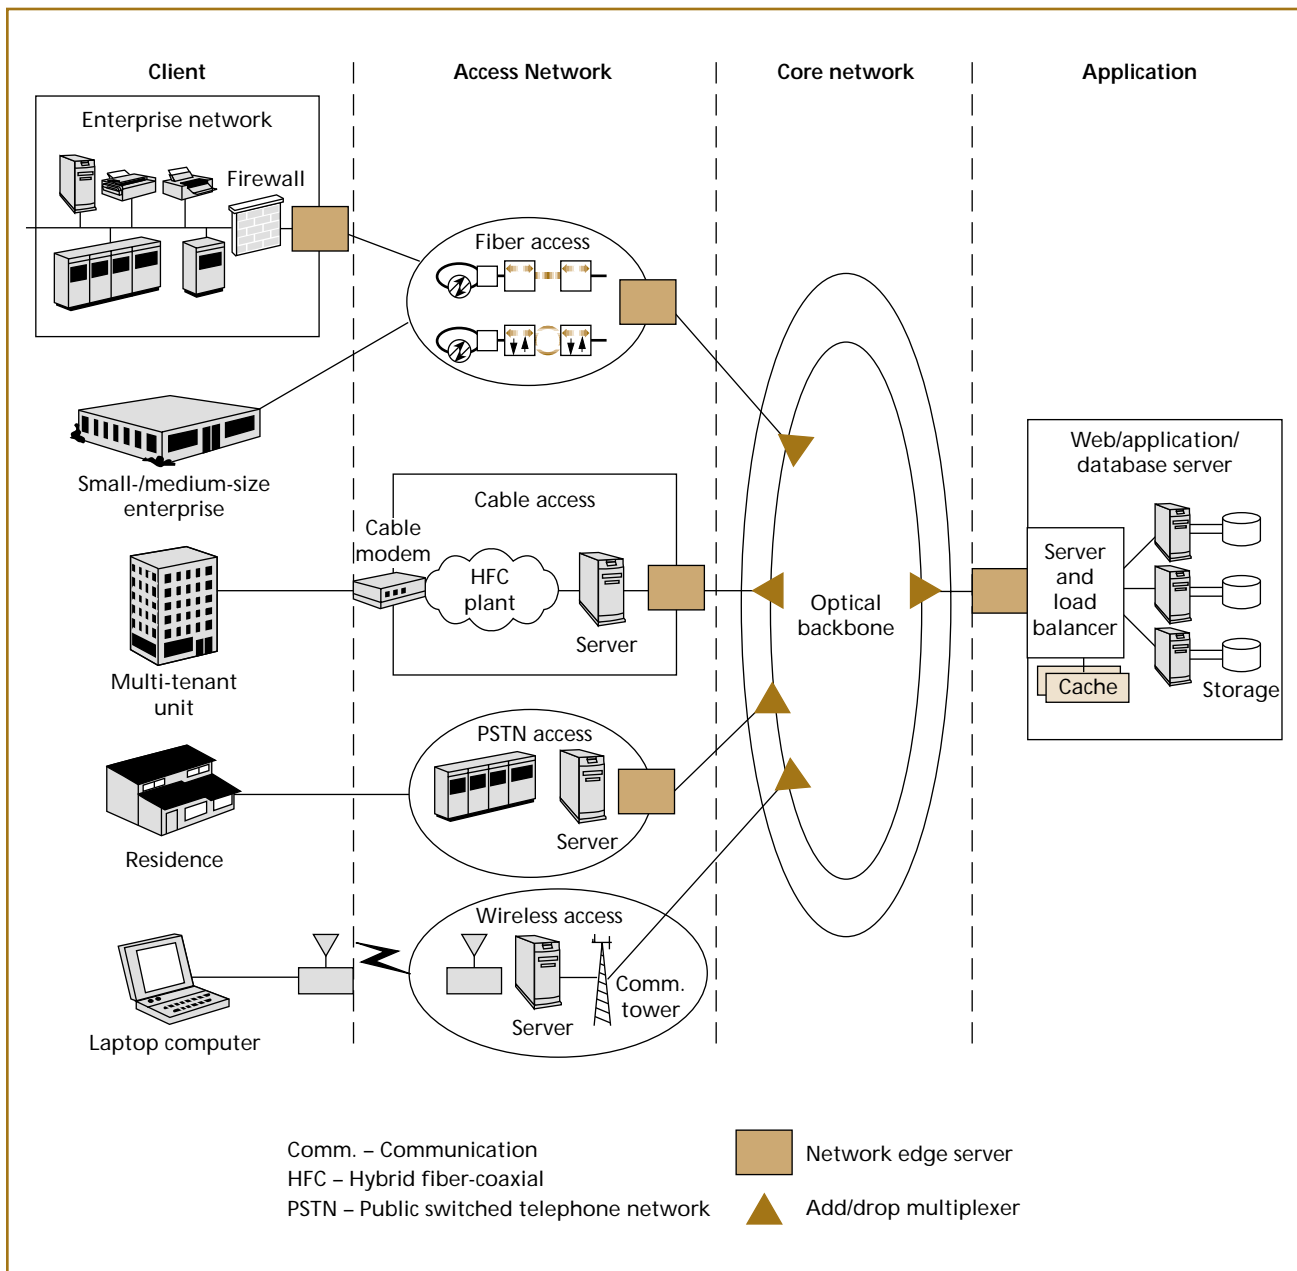


Figure 1.
The emerging network.

Figure 1 shows this view of the emerging network. In this architecture, the client/server traffic is carried over the core network. It is envisioned that this core network will be a high-speed network using optical transmission facilities with photonic switching technology. Various types of end users, who are the *clients* of applications, will be using different technologies for access to the backbone network. The value of

this network is that it enables almost latency-free communication between the clients and servers in different locations. The access networks provide integrated multi-service, multi-technology access to the backbone network. The edge nodes in the access networks perform a variety of functions, such as access security, conversion between protocols, and traffic load-balancing. The gateways between the access networks

and the core network perform similar functions as edge nodes. Client services, such as Web hosting services, e-mail, and video-on-demand services, are actually performed by a number of computing platforms residing on a number of different servers distributed in many data centers.

From the management point of view, the end-to-end network may be partitioned into four segments as shown in Figure 1: client segment, access network segment, core network segment, and application segment.

Client Segment

The client ranges from individual end-users or small- to medium-size enterprises to large corporations. Clients are the components and resources of an end user or an enterprise network that need to get connected to applications over another network. For teleworkers, Lucent has developed the Internet protocol service switch (IPSS)—which supports virtual private network (VPN) access solution—by partnering with SpringTide. There are many methods by which residential clients can be connected to the network. These include access servers and digital subscriber line access multiplexers (DSLAMs). Lucent's AnyMedia™ Access System is a next-generation digital loop carrier system that supports "plain old telephone service" (POTS) and the integrated services digital network (ISDN).

For small- and medium-size enterprises, there are Internet protocol (IP)-based private branch exchanges (PBXs), IP-enabled business communications and information systems, and a family of wide area network (WAN) access solutions. There are also servers and applications hosted by the enterprise. It is quite likely that these enterprise-hosted servers and applications need to interoperate with other servers in various locations.

Some client enterprises are considering providing voice service over their IP networks. In this case, voice traffic is routed through an access gateway, encapsulated in RTP/UDP/IP, carried to the IP edge router, and then transferred through the IP network.

Access Network Segment

Access networks are used to provide integrated multi-service, multi-technology access and to support higher bandwidth, remote data access, and ever-increasing mobility. The residential and enterprise users

may use various fixed or mobile access networks. For most residential users, the family of DSL technologies, collectively referred to as xDSL, allows a few Mb/s to be carried over the existing unshielded twisted-pair copper plant relatively fast and cost effectively.

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 Mb/s of data to be carried over hybrid fiber-coaxial (HFC) cable access systems. The frequency spectrum is split both at the customer site and at the cable head end, thereby allowing access to community antenna television (CATV), optical backbone and, optionally, public switched telephone network (PSTN).

Another compelling technology is wireless access, which is becoming a dominant method for accessing services on a network. 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 is being considered for providing multi-Mb/s access. Convergence of fixed and mobile access is concerned with provision of network and service capabilities that are independent of the access technique.²

Optical access options exist as well. In the near term, the optical core will repeat itself, although with lower capacity, to provide metropolitan transport and business access to the core network. As the optical networking technology moves into access networks and the cost of bandwidth goes down, the preferred way to carry IP traffic becomes IP over SONET or IP over Gigabit Ethernet.

Core Network Segment

The vision for the core network is seamless and reliable delivery of future services. This core network will provide:

- Coexistence and convergence of data and voice networking;
- Optical backbone networking;
- Support of multiple protocols and services, including frame relay, IP, asynchronous transfer mode (ATM), and virtual private networking;

- Support of a variety of broadband access technologies; and
- Flexible bandwidth and service capabilities.

The wavelength division multiplexing (WDM)-based photonic network technologies will be used for core networking, which can provide ample capacity and wavelength routing functions. An intelligent optical core using WDM will meet a carrier's needs for scale, provisioning, and restoration, while increasing overall efficiency. Dense wavelength division multiplexing (DWDM) is a further enhancement to increase capacity.

Two techniques can be used for transport of IP traffic over WDM:

- Using existing transfer modes (SONET/ATM/frame relay) between IP and WDM, and
- Using wavelength routing capability.

In the first approach, which is currently used, IP traffic is carried over the existing network transfer modes (SONET/ATM/frame relay) to provide flexible and granular access to bandwidth. However, this technique suffers from bandwidth inefficiencies and involves processing at each intermediate node. One remedy to this approach is to eliminate the IP layer altogether from the optical core network by setting up optical paths between ingress-egress edge nodes using the wavelength routing capability of the optical layer. Lucent's WaveStar™ LambdaRouter is using this technique to provide transparent transport of various traffic in a backbone network with terabit-per-second capacity.

Core network functionality is usually provided by one or more network service providers (NSPs). However, the emerging application service provider (ASP) market is triggering an unprecedented data center build out, which in turn has created a huge market for a variety of new kinds of NSPs.³ The business model may lead some NSPs to work with ASPs, either as a transport service, where both the ASPs and end users are charged, or as a reseller, where end users are charged for a rich suite of services with a single bill.

Application Segment

The separation of applications from core network and access technologies encourages application diversity—any ASP can supply various applications. Many

small- to medium-size enterprises leverage their networks by creating and/or expanding the services they already offer to the end users. Specifically, they accomplish this by:

- Expanding the reach and functionality of existing network services;
- Leveraging the supply chain by supporting valuable new services, applications, and businesses for suppliers and partners;
- Supporting large communities of interest; and
- Extending the influence of the enterprise brand.

It is important to distinguish between application services offered by the service providers to their customers and the solution and configuration choices they make in deploying their networks and data center infrastructures. An *application service* is a function or set of functions that is sellable and visible to clients. Applications are driven by business needs to provide features that incorporate voice, data, and video in a seamless, organized, and value-added manner.

Service providers will leverage the capabilities of the optical communication networks for the next-generation Internet by providing a broad array of applications to enterprises as network-hosted services and communications applications. For example, a Web portal service delivery provides cross-platform development advantages, enabling different types of endpoints to access common services from a generic Web browser without requiring custom software. These applications include “e” (electronic) and “i” (Internet/IP) services.⁴ The growth in these areas portends significant revenue opportunities and unprecedented growth (see **Panel 2**).

The key elements of an application data center are servers running e-business and servers running middleware software for load balancing and caching, IP address management, security, and policy-based management.⁵ In addition, there are storage solutions and high-performance routers to connect them all together. The function of an edge adaptor server is to offload the servers from TCP/IP processing, enable information to be shared between servers, and provide an optical signal that is compatible with the deployed WDM transport systems. This functionality alleviates

the major bottleneck of getting the information into and out of servers running various applications. Another layer of middleware is used to tie applications together and provide a communications-enabling set of application programming interfaces (APIs). The business objective of the application data centers is to integrate these technologies to provide economies of scale by serving multiple enterprises with a minimum amount of client-specific equipment and software.

Given a set of application services, the NSPs or ASPs will make decisions on how they wish to implement their solution. For example, they may choose an ATM connection-oriented infrastructure versus an IP-routed connectionless WAN infrastructure; or, based on their network architecture, they may select centralized versus distributed network management.

Lucent Network Realization

Currently, Lucent is supplying service providers with a variety of switching and transport technologies capable of delivering voice and data at high speed. To satisfy new network requirements, Lucent, having formulated a coherent strategic view of how a network satisfactorily delivers a full range of services to customers, has developed the Lucent Network Architecture—a data-centric optical transport network architecture.⁶ 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.

A typical Lucent network architecture for access to the Cyber Data Center with optical transport is shown in **Figure 2**.

End-to-End Management

The management method used for the network as shown in Figure 2 must allow rapid implementation of new services and support differentiated quality of service (QoS) and service-level agreements (SLAs). Network providers are making significant changes due to transformation of the regulatory environment, the growth of the Internet, and support of a diverse variety of multimedia traffic. Different service providers—namely, network service providers (NSPs), data center service providers (DCSPs), application service

Panel 2. Application Services

- **Communications Applications**
 - E-mail
 - Voice mail
 - Unified messaging
 - Videoconferencing
 - Workgroup applications/collaboration
 - Distributed call center
- **Business Applications**
 - Enterprise resource planning (ERP)
 - Payroll
 - Sales force automation
 - Web hosting
 - Internet-enabled applications
- **Content hosting**
 - Distance learning
 - Communities of interest
 - Entertainment
 - Network services
- **E-Commerce**
 - On-line auctions
 - Supply chain management
 - Person-on-line assistance
 - Transactions/ordering

providers (ASP), and CyberCarrier service providers (CCSPs)—are beginning to exploit their infrastructure to offer a wide variety of services with varied requirements and architectures.³ Clearly, the objective is efficient delivery of services. Management at the network level is insufficient, and there is a strong need to design and implement management systems that address the data centers and applications. The current focus on monitoring and managing SLAs is a precursor to this paradigm shift.

Service management is concerned with managing the services a network provides to customers. Service management includes functions such as order processing, customer trouble management, and billing. Network management, however, involves the activities that are responsible for the assignment and control of proper resources in the network, both hardware and software, to provide information movement among the end users and network nodes in a timely and cost-effective manner that meets the user's performance needs and the network's objectives.

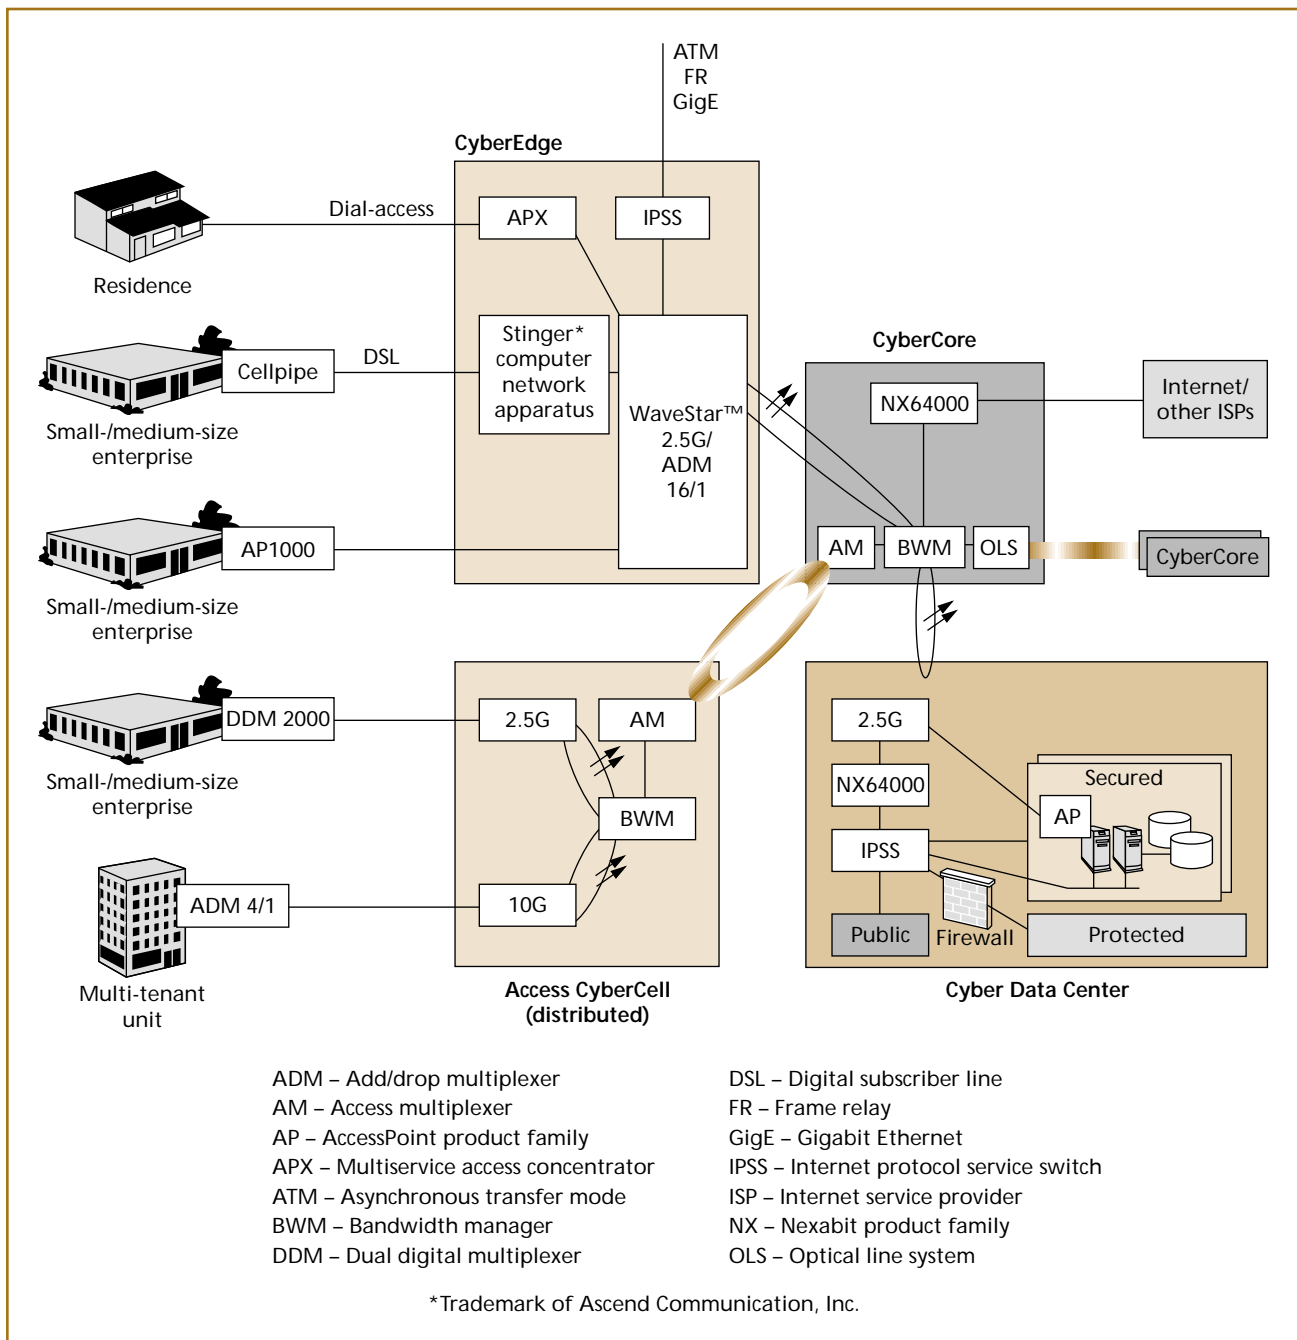


Figure 2.
A Lucent network architecture.

By this definition, service management may be considered as much an art as it is a science. As a science, it relies on computer science techniques to develop new computing tools and communications platforms for monitoring and control functions. As an art, it relies on human experts to make the final judge-

ment for implementing the tools and platforms and devising the corresponding processes. Lucent is combining the broad spectrum of expertise of its engineers and designers to address the management paradigm shift and deliver an efficient, reliable, and flexible management platform for its optical networking vision.

Supporting transparent application access and interoperability in the information network shown in Figure 1 requires integration between network management and service management. Managing services supported by any network requires the management functions to be interlinked. The partitioning proposed in Figure 1 breaks the end-to-end network up into four separate management domains: client, access network, core network, and application. **Figure 3** shows a corresponding high-level service management architecture. The basic architectural elements in Figure 3 are management layers, management protocols and standards-based interfaces, and service management and inter-domain network management MIBs. A brief description of these elements follows.

Management Layers

The architecture in Figure 3 shows four functional layers: service management, inter-domain network management, domain management, and domain. The management 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 inter-domain network management layer all the way down to the corresponding network element in a particular domain. The philosophy is that the service requirements are the drivers for all management issues. A particular set of requirements in the service management layer dictates the requirements for integration and interoperability on the underlying inter-domain network management layer. At the lower layers, activities affect physical resources; at the higher layers, they affect relatively abstract entities and processes like SLA management.

The required tasks of each management layer can be categorized into the five management functional areas of fault, configuration, accounting, performance, and security (FCAPS). Each layer is typically implemented with one or more deployable and interoperable operations support systems (OSSs) with well-defined interfaces.

Depending on the relative position of OSSs with respect to the management layers in Figure 3, they can be called service management system (SMS), network

management system (NMS), or element management system (EMS). In practice, there can be many EMSs for implementing management functions in each domain, which communicate with functions in other layers. For example, McKiou and Esposito in their paper in this issue, "OneLink Manager™ EMS for the 7R/E™ Switch,"⁷ describe the design of the EMS for the Lucent's 7R/E packet network element, which performs a complete FCAPS management function.

Service management layer. To meet the challenge, the traditional service management platforms and techniques need to be enhanced. To support services such as IP telephony, virtual private networks (VPNs), and multimedia services, service management systems must be capable of providing mechanisms for managing QoS and SLAs. Rapid introduction of new value-added services and applications requires highly flexible and responsive management functionality. To support services such as IP telephony, VPNs, and multimedia services, service management systems must be capable of providing mechanisms for managing QoS and SLAs.

Panel 3 lists general requirements for the service management platform and its OSS architecture from a set of Requests for Proposals (RFPs) Lucent has received for network-hosted services. "CyberCarrier Service and Network Management" by Brenner et al.³ provides a service/network management platform to address the requirements listed in Panel 3.

The objective of the service management layer is to consider all the functions a network provider requires to provide services to end customers. Service management deals with customer contact and interface, QoS, and interaction among services. The FCAPS-related service management functions, shown in Figure 3, are:

- *QoS management:* This is an important function in management of packet transport networks, particularly IP networks. This function is needed to provide performance guarantees to individual traffic streams or, optimally, aggregate requirements in a network that uses statistical multiplexing. QoS must be ensured regardless of the underlying transport technology. The network needs to support QoS-

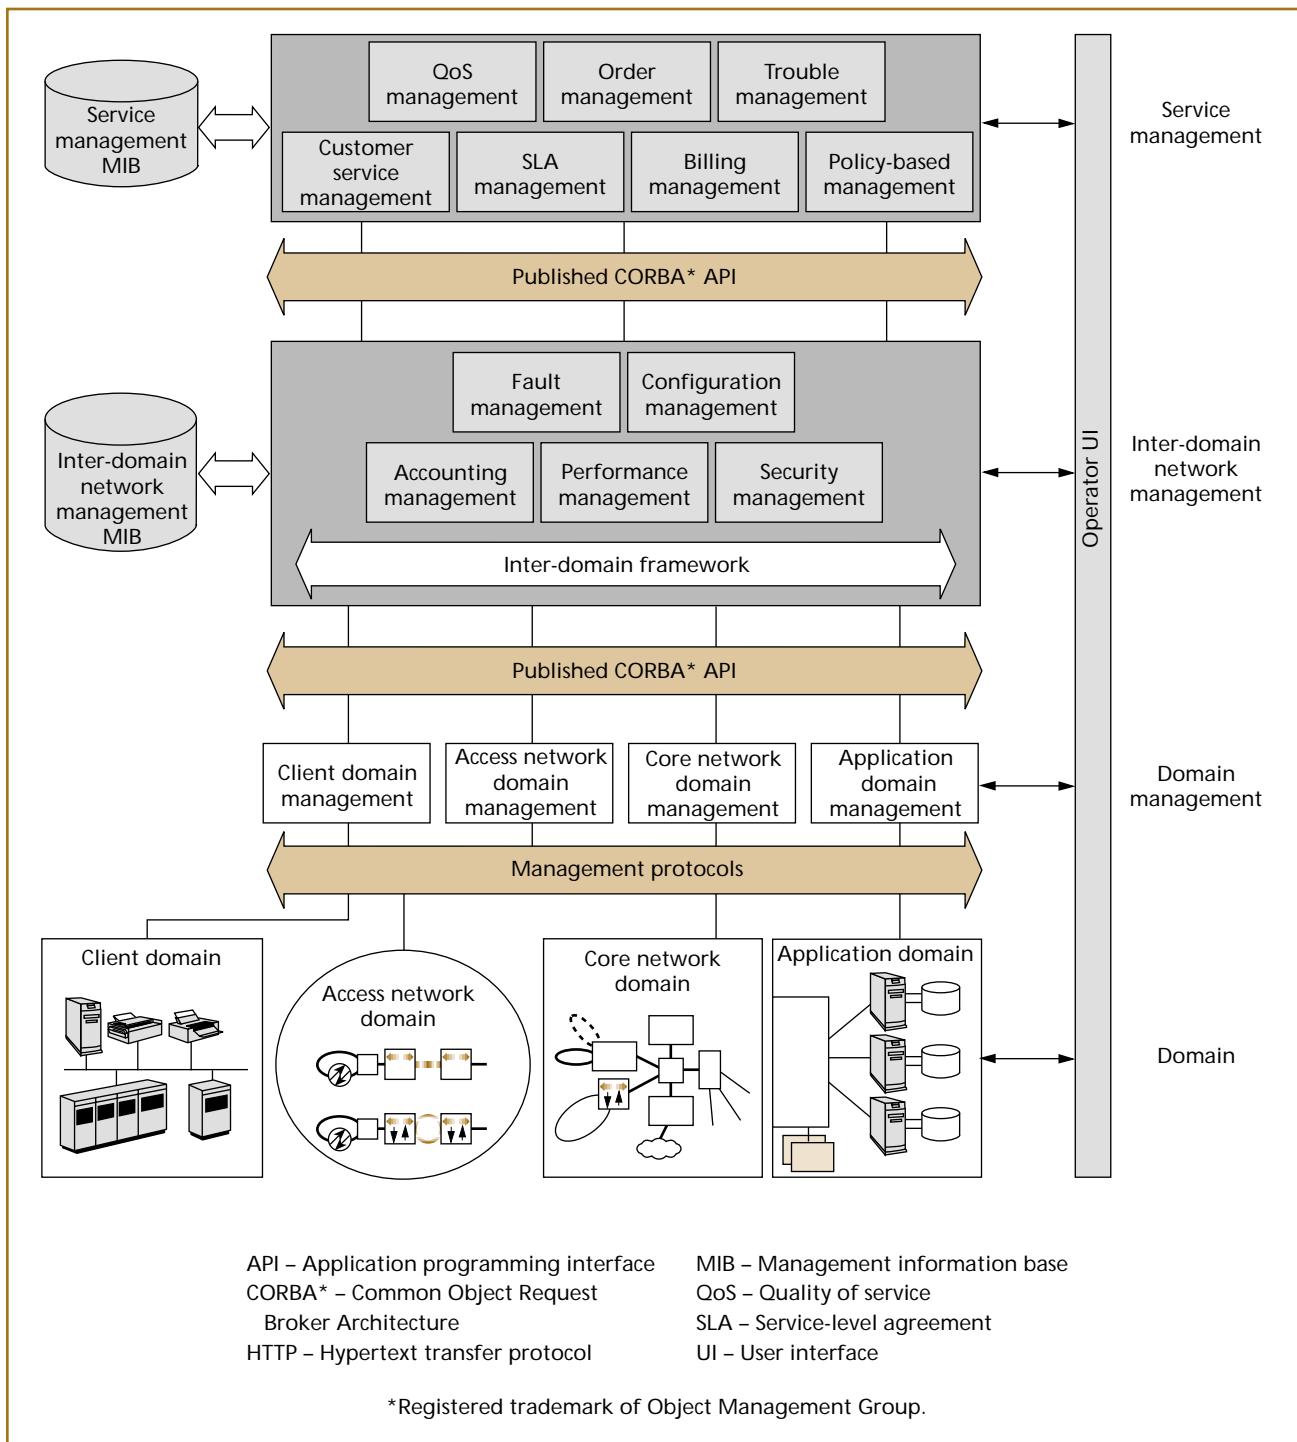


Figure 3.
High-level service management architecture.

capable switches and routers, mechanisms to select an end-to-end path, and mechanisms to reserve bandwidth along that path. Typical techniques are the Internet Engineering Task

Force's (IETF's) Diffserv and multiprotocol label switching (MPLS).⁸

- *Order management:* This function involves entering and managing customer orders for

Panel 3. General Requirements for a Service Management Platform

Structured architecture—Application hosting service management features shall be defined in a layered fashion allowing the customer to select the level of support and functionality desired, as well as stand-alone or integrated implementation of systems.

Complete management functions—Full fault, configuration, accounting, performance, and security (FCAPS) management functions are required at the service management layer.

Northbound interfaces—The management systems of the managed entities shall support a “TMN-like northbound interface” to the upper layer management functions.

Availability—The key components of management systems shall have high-availability to support customer-required (up to 99.999%) availability of service delivery components. The necessary fault tolerance and redundancy shall be provided.

Architectural scalability—The management platform shall be reliable, scalable, secure, and feature rich.

Modifiability—The management platform shall show how it can evolve to offer increased bandwidth, additional services and/or QoS for future network infrastructure/data center requirements.

Standards adherence—The management platform shall show the capability to support new and emerging technologies and the adherence to industry standards.

End-to-end solution—The management platform shall support an end-to-end service.

Remote management access—The management platform shall support Web-accessible remote access to network management functions for authorized network managers for monitoring and configuring the service quality.

Rapid deployment—The management platform shall support rapid deployment of management systems and services.

Security capability—The management platform shall support implementation of security management functions including application-level security and authentication as:

- Confidentiality and privacy
- Data integrity
- Authentication
- Access control and authorization
- Service availability and prevention for denial-of-service attacks
- Real-time intrusion detection and response
- Directory-enabled offerings or security policy rules
- Data center physical security (such as iris scanners or video surveillance)

continued on next page

one or many services, generating a single workflow, and distributing the workflow to downstream systems. The order management supports the business requirements of enterprises, access, Internet, and core network service providers through a set of standard APIs. A preferable method of customer access to the order management function is via the Web.

- *Trouble management*: This function is used to administer troubles experienced with services provided by the service provider to the users. The trouble management function is responsible for creating and exchanging trouble reports, monitoring the status of trouble reports, and

requesting escalation of the priority for resolution of the trouble. For end-to-end solution, the trouble management system needs to be database and platform independent.

- *Customer service management*: This function is for managing the interaction between the customer and the provider. The tasks include traditional customer care applications, but they also cover new applications such as access to the customer's SLA reports, on-line care, and Web-based customer network management.
- *SLA management*: The SLA specifies connectivity and performance level agreed to between an end user and a provider of service. SLA

Panel 3. General Requirements for a Service Management Platform *(continued)*

Billing, auditing, and reporting capabilities—The management platform shall handle a variety of pricing and billing methods that will best reflect the customer preferences and perception of value from a service. These value-based methods shall include:

- Activity-based billing (for example, per transaction, page view, and email-sent charges)
- Usage-based methods (for example, charges sensitive to time spent on an application, usage of disk space, or of server computing power)
- QoS-based methods (higher price for higher service level in terms such as percentage of service availability, bandwidth capacity, data and network redundancy, as well as customer refunds/credits in case of an SLA violation)
- Cross-service packages and usage/spending-based discounts and promotions
- Account for price difference between power

vs. regular users and usage/spending-based tier pricing (that is, lower cost for incremental user and cap pricing with variable fees over the cap per user, usage, transaction)

Quality-of-service management—The management platform shall handle a variety of QoS classes and standard methods such as:

- QoS-based routing methods and standards
- Multiple service classes
- Bandwidth-on-demand to support high-speed applications
- Broader, universal, and flexible connectivity

Load-balancing capability—The network management platform shall support an integrated set of management tools focused on capacity planning and load-balancing capability. The network management platform shall support distributing customers' traffic through load-balancing edge servers to achieve best performance or avoid unwanted threshold events.

management allows for setting thresholds and emitting threshold-crossing notifications. The tasks include being able to craft SLAs, design and provision the network to support them, and control the network to ensure the service levels are achieved. When SLA contract violations occur, the resolution process is determined based on the contract for the affected customers. SLA reports and/or notifications are sent to the customer via e-mail and/or to a Web-based customer service management system.

- *Billing management*: The billing functions include usage measurements, tariffing/pricing, collections and finance, and fraud prevention and counter-measures. Billing requires detailed usage data collection within the network. A mediation function collects usage data from network devices and systems and delivers that data to a back-office billing application. Mediation in IP networks requires substantial in-network aggregation.

- *Policy-based management*: Policy-based management is used where complicated analysis of events and very fast reaction times are required—as in SLA management, for example. A *policy* is a formal representation of information affecting network/component behavior, but specified independent of components. The policy-based management function interacts with the service QoS management function, security management, and inter-domain network configuration management. The network administrator expresses the high-level policies through a user interface. The policies are then interpreted by a router or server.⁹ Using lightweight directory access protocol (LDAP), policy servers retrieve appropriate policies from the directory, interpret them, and program the network devices using protocols such as common open policy service (COPS). Directories and LDAP are the primary means for storing and distributing network policy information.

Inter-domain network management. The objectives of traditional NMSs/OSSs have been to reduce operation complexity and cost by increasing flow-through, but the focus has generally been at the element level. As networks grow in size, technology becomes more diverse, services become more complex, scalability becomes a larger issue, and inter-domain management becomes more critical.

The inter-domain network management introduces a set of coordination and correlation functions into the network management layer that transcends any one domain. It aims to integrate and coordinate all the resources necessary to configure, monitor, test, analyze, evaluate, and control the end-to-end network so that service-level objectives are met at a reasonable cost. The driving forces for introducing this layer are:

- Managing diverse client, access, core, and application networks efficiently;
- Minimizing the overall network-monitoring and control functions;
- Relinquishing control of strategic assets to the client, access network, core network, and application managers (that is, domain managers);
- Minimizing the effect of multiple management protocols and standards; and
- Reducing the cost of network operations.

The inter-domain network management functions encompass FCAPS and may be implemented in a single system, or distributed among new or existing domain managers. Each domain manager needs to implement its network management functions and interwork with both peer domains and end users.

The inter-domain fault management collects faults across multiple domains and determines the root cause domain. It provides the craft personnel with a consolidated inter-domain view of alarms, potentially reducing the need for additional staff to manage multiple domains of access, core network, and data centers.

The inter-domain configuration management sets up the connections for services offered to subscribers and controls status of the network and data center as well. It enables single-point entry of provisioning requests—regardless of underlying domain—that will reduce overall data entry errors. It also offers end-to-end views of connections and their

underlying infrastructure, independent of domain.

Inter-domain accounting management includes collecting the data for usage of services and billing for the usage. The services may be network based, application based, or a combination of the two, and billing may be flat rate or usage based. The challenge is to gather a huge volume of packet-based usage data for individual users, map the flows to the users (in the face of dynamic assignment of IP addresses), and measure the individual user's service levels and thresholds. A solution to IP-based accounting is to use an in-network device that sits on the network between the user and the backbone and associates each packet with a user.¹⁰

In inter-domain performance management, functions are used to monitor the network so that preventive maintenance, capacity planning, and statistical analysis can be performed. The collected data is used to prevent potential degradation of resources so that the QoS promised to the customer can be maintained. The inter-domain performance management feature, in conjunction with customer service management (CSM) feature, allows customers to view the status of their own network services. Customers can access reports that contain information only about their own respective customer services.

Inter-domain security management tasks include both securing the management information and managing the network and information security.¹¹ Both require supporting security services, such as authentication and access control, privacy and confidentiality, data integrity, nonrepudiation, and denial-of-service prevention. These services are described in ITU Recommendation X.800.^{12,13}

Domain management. In the architecture shown in Figure 3 domain management systems provide an integrated set of network management functions that apply to their own domains. This approach has the advantages of:

- Reducing the cost of interfacing with multiple NMSs in other network segments or domains,
- Enabling the inter-domain systems to monitor and control domain resources for utmost effectiveness and productivity,
- Quickly deploying the domain NMS and EMS solutions, and

- Accommodating diverse technologies and facilitating growth and management in a cost-effective manner.

Conformant to the supported standards, the domain managers use ITU-T TMN/OSI's common management information protocol (CMIP), IETF's simple network management protocol (SNMP), and OMG's Common Object Request Broker Architecture (CORBA*) between the manager and managed systems.

Furthermore, the network management within the domain and between a domain and the corresponding inter-domain management can be agent based. Agents can act as middleware (any software entity that is interposed between a client and a server, a peer and other peers, or an application and a platform) to inter-link the domain managers with the inter-domain management layer, as discussed by Bieszczad et al.¹⁴ and others,¹⁵ who provide a detailed account of intelligent agents in telecommunications management.

Service Management Information Base

The service MIB shown in Figure 3 is a conceptual repository for all management information and parameters necessary for normal functioning of the integrated service management platform. A MIB is 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 IETF Requests for Comments (RFCs). OSI MIB is a collection of managed object classes defined in the special-purpose notations in "Guidelines for the Definition of Managed Objects" (GDMO).¹⁶

Attaining interoperability of 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). The

management functions currently exchange management information by means of techniques defined in the ITU-T X.700 series of recommendations.¹⁶ These recommendations incorporate the important object-oriented and manager/agent paradigms for information modeling. XML (Extensible Markup Language), standardized by the World Wide Web Consortium, is becoming popular as a mechanism for representing management data in a standard manner. For example, the Distributed Management Task Force (DMTF)¹⁷ announced version 1 of its XML encoding specification for encoding its Common Information Model (CIM) schema in XML in 1998.

Management Protocol and Interface Capabilities

Two critical aspects of communications among human operators and the management functions in the service management architecture in Figure 3 are protocols and interfaces. Network management protocols are defined as interactions between the managing systems and the managed entities. The management architecture needs a number of protocols to communicate with users, management domains, MIBs, and host operating systems. Protocols are not part of the functional architecture, because they are *means* to implement management commands rather than *tasks* to implement management functions. The interfaces are defined between a management system and a MIB, between one NMS and other NMSs and applications, and for host machines. Note that the functions of the service management platform should be accessible from any layer of the communication protocol that needs network management services. The management functions need to support secure multi-user remote administration for network managers and system administrators. The remote access to OSSs is performed preferably over a secure Web-based interface. Remote access to network elements is often achieved via Telnet.

The communication and distributed computing mechanism for the service management platform is achieved by CORBA, an industry standard.¹⁸ CORBA defines a framework for developing object-oriented distributed applications. It is a communication middleware that is not only platform independent, but also language independent. Moreover, CORBA is vendor independent and is thus the logical choice for manage-

ment systems that must operate in a heterogeneous environment.

CORBA plays two important roles in the management platform shown in Figure 3—namely, it serves as both a communication backbone for OSSs and an open application interface for service management and provisioning. All management layers are basically CORBA clients and servers. CORBA as a communication backbone helps to hide low-level communication complexity, simplify development, and enable scalability with the same software architecture. As an open API for external systems, CORBA allows an upper-layer OSS to connect to the lower-layer managed entity.

Lucent is adopting CORBA as the interface of choice for implementing inter-OSS interfaces and has been seeking to define multi-vendor standard interfaces based on this technology for ATM/frame relay and IP at the EMS-to-NMS and higher interfaces. Implementation of the management architecture in Figure 3 should support CORBA Internet Inter-ORB Protocol (IIOP*) and a published API for its upstream systems to communicate with downstream systems and their GUI servers. Nevertheless, communications between network elements and their corresponding domain managers can be thorough standardized management protocols like SNMP and CMIP/CMISE.^{19,20}

The CORBA-based architecture in Figure 3 can reuse existing MIB specifications using gateways between CORBA-based inter-domain network management and the existing CMIP/SNMP-based domain managers, if any.

Some Solutions and Ongoing Challenges

Realizing a network management architecture involves a careful balancing of a web of tradeoffs among a set of critical factors to meet the associated technical and networking challenges. 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. Once the management systems are in place, there are also the issues of creating processes and tools that can support the development of new techniques.

We have identified many challenges in end-to-end service/network management. The papers presented in this issue of the *Bell Labs Technical Journal*, which are summarized in the following sections, describe ongoing efforts from element management to service management within the Lucent community to address some of these challenges.

Managing Applications and Hybrid Network Elements

Silverman, Brenner, and Shannon, in “Toward a Vision for Network and Service Management,”²¹ offer a high-level application-driven framework for shifting from a voice-centric network environment to the data-centric distributed computing environment of today. This framework is inspired by the TeleManagement Forum’s Telecom Operations Map (TOM) model,²² which is gaining acceptance in industry. The authors elaborate on this paradigm shift and conclude that the management platform must conceal network complexity to enable customers to be more agile. In a companion paper, “Implementing a Management System Architecture Framework,” Goers and Brenner²³ consider the architectural framework offered by Silverman, Brenner, and Shannon and provide recommendations for an implementation that promotes commonality, simplicity, and flexibility.

A mixture of private and public services is currently provided through intranets, extranets, and the Internet. While the service management principles remain the same, a public service requires more controls and filters on the visibility and granularity of the service provider’s internal service state information. Lucent has recognized the importance of managing an end-to-end broadband network for access to applications and data centers. Brenner et al., in “CyberCarrier Service and Network Management,”³ discuss a management functional architecture that divides the problem of managing applications and hybrid network elements into tractable pieces. They provide an overview of the service management architecture of Lucent’s CyberCarrier Solution, which is a more detailed breakdown of the architecture presented in Figure 3.

It appears that the current state of the industry is not really offering SLA management tools, but rather SLA monitoring tools. The subtle difference is that most

of the current tools/systems allow one to track SLA thresholds according to selected metrics. However, none of them effectively manage the contractual agreement between the service subscriber and the service provider. Moreover, none of today's products allow for extensive customization through which a service provider can differentiate its own offering from its competitor's offering. Further, customers more and more require being able to view their logical subnetworks and be billed commensurate with the provided QoS.

With optical networking gaining tremendous momentum in recent years, the challenge is to develop efficient ways of managing multi-domain networks. The network elements will include both optical and packet components and interfaces, and the element managers have to cope with that. In addition, an inter-domain management function is needed to reconcile the differences in paradigms used to manage the legacy systems and the new optical elements. In "Hybrid Network Management," Epstein et al.²⁴ address integrated management of data and fiber-optics-based transport. The authors offer a CORBA-based solution that provides simplicity, scalability, and flexibility. They show that the CORBA-based management techniques combined with standard information models and protocols will enable multi-vendor network management solutions and encourage the use of the third-party applications.

Reducing the Number of Management Systems

For the converged network of today, the service management environment needs to be rich with applications developed across network components in a distributed architecture. With industry endorsed APIs, third-party tools will create services that were traditionally built as part of the telecommunications network management.

The OneLink Manager EMS is designed to allow rapid integration of disparate technologies and equipment from third-party vendors into coordinated OAM&P functions with a unified standards-based interface. McKiou and Esposito in their paper, "OneLink Manager™ EMS for the 7R/E™ Switch,"⁷ mentioned earlier, describe the design of this EMS for the Lucent's 7R/E Packet Network Solution. The OneLink Manager design is an acknowledgement by

Lucent that an environment of technological churn exists and that Lucent has adopted a new approach to OAM&P integration for its products.

Reengineering Existing Processes and Reducing the Operations Cost

As operations continue to evolve to support new services, carriers will need to reengineer existing processes to support these services and gain efficiency. Traditional transport management systems and data management systems use different operations philosophies to address their needs. Service providers continue to see the value of well-automated operations processes. Starting with a core time division multiplexed (TDM) network that can be managed with minimum FCAPS support and expanding to packet-based network services seems a logical path for most incumbent local exchange carriers (ILECs). However, a thorough operations cost analysis is needed before this path is taken. Lucent has developed cost models to analyze the trunk operations cost for the TDM and packet networks. Tsay's "Analysis of Service Operations Cost for TDM and Packet Tandem Networks"²⁵ sheds some light on this path. This paper shows that the network operations cost not only depends on the underlying technologies, but also on the network architectures deployed to support various services.

Integration with Legacy OSSs

New OSSs will need to integrate with existing systems to provide seamless customer service management and operations support. This would allow service providers to add the management applications they need whenever they need them. Eggert, Johnston, and Vaughan, in "The Flexent™ Element Management System—Using Web and Object Technologies to Manage a Wireless Network,"²⁶ describe an EMS for a wireless network that provides wireless communications services to Lucent's customers worldwide. They describe the architecture and features of the Flexent EMS and show how Lucent uses the Web and object technologies in its element management applications. The paper discusses the unique challenges of element management for the wireless network domain.

Wireless service providers are trying to use the ATM backbone network to carry their pulse code

modulation (PCM) streams. In wireless networks, ATM technology offers the opportunity for voice-data integration with QoS management, simplified and multi-service provisioning, and efficient facility consolidation. However, integrating ATM OSSs into a wireless service provider's operations environment is a challenging task. Schlaerth, in "Service and Network Management Strategies for ATM in Wireless Networks,"²⁷ discusses integration of management functions for the wireless and ATM networks. The conclusion is that ATM service management may be scaled by partitioning it to coincide with existing domains (as shown in Figure 1) and integrating customer network management (CNM) at the lowest level of domain management. This allows wireless service providers to introduce ATM without making expensive changes to their overlying network and service management systems.

Conclusion

We have outlined a functional architecture that addresses the needs of next-generation management systems. This architecture is based on the vision that the value of a network will be greatly enhanced by how easily services are provided. We have presented the essential functions of this service management platform. The fundamental issue for a coherent management platform capable of accommodating multiple technologies and multiple vendors is flexibility for manager-agent communications among various management domains. It has been argued that the flexibility could be achieved using CORBA-based interfaces among management systems. Increasingly, enterprise networks are using shared computing applications, requiring new management systems to support them. Furthermore, the service management platform should promote and drive inter-domain network management, application modularity, and adherence to standards. Service providers are concerned with the clean separation of management control and data so that the management platform can be extended to support new services.

The papers in this issue of the *Bell Labs Technical Journal* describe a wide range of technologies, methods, and strategies for managing the complexities in

service, network, and element management processes and systems. We hope you find this cross-sectional view of activities insightful and inspiring.

Acknowledgements

We would like to thank Michael Devlin and Ramdas Iyer for their insightful comments.

*Trademarks

CORBA and IIOP are registered trademarks of Object Management Group, Inc.

Stinger is a trademark of Ascend Communication, Inc.

References

1. "Network Management" issue, *Bell Labs Tech. J.*, Vol. 4, No. 4, Oct.–Dec. 1999.
2. 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.
3. M. R. Brenner, M. Chu, G. Gross, and M. Malek, "CyberCarrier Service and Network Management," *Bell Labs Tech. J.*, Vol. 5, No. 4, Oct.–Dec. 2000, pp. 44–62.
4. D. C. Dowden, K.F. Kocan, and J. Kozik, "The Future of Network-Provided Communications Services," *Bell Labs Tech. J.*, Vol. 5, No. 3, July–Sept. 2000, pp. 3–11.
5. 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.
6. 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.
7. K. W. McKiou and D. Esposito, "OneLink Manager™ EMS for the 7R/E™ Switch," *Bell Labs Tech. J.*, Vol. 5, No. 4, Oct.–Dec. 2000, pp. 80–96.
8. 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-mppls-arch-06.txt>>.
9. 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.
10. D. Mitra, K. E. Sahin, R. Sethi, and A. Silberschatz, "New Directions in Service Management" *Bell Labs Tech. J.*, Vol. 5, No. 1, Jan.–Mar. 2000, pp. 17–34.
11. L. G. Raman, *Fundamentals of Telecommunications Network Management*, IEEE Press, Piscataway, N.J., 1999, p. 31.
12. International Telecommunication Union,

- "Security Architecture for Open Systems Interconnection for CCITT Applications," Rec. X.800, 1991, <<http://www.itu.int/itudoc/itu-t/rec/x/index.html>>.
13. International Telecommunication Union, "Layer Two Security Service and Mechanisms for LANs," Rec. X.800 Amd. 1, 1996, <<http://www.itu.int/itudoc/itu-t/rec/x/index.html>>.
 14. 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.
 15. "Intelligent Agents for Telecommunications Management" issue, *J. of Network and Systems Management*, Vol. 8, No. 3, Sept. 2000.
 16. International Telecommunication Union, "Guidelines for the Definition of Managed Objects," ITU-T Rec. X.722, Geneva, Switzerland, Jan. 1992, <<http://www.itu.int/>>.
 17. <<http://www.dmtf.org>>.
 18. Object Management Group, "The Common Object Request Broker: Architecture and Specification," Revision 2.2, Feb. 1998, <<http://www.omg.org/corba/>>.
 19. International Organization for Standardization, "Common Management Information Protocol (CMIP)," ISO Rec. 9596, Nov. 1998.
 20. International Organization for Standardization, "Common Management Information Service Element (CMISE)," ISO Rec. 9595, Nov. 1998.
 21. K. S. Silverman, M. R. Brenner, and G. E. Shannon, "Toward a Vision for Network and Service Management," *Bell Labs Tech. J.*, Vol. 5, No. 4, Oct.-Dec. 2000, pp. 21-30.
 22. TeleManagement Forum, "Telecom Operations Map," Document GB910, Version 2.1, Mar. 2000, <<http://www.tmforum.org>>.
 23. W. C. Goers and M. S. Brenner, "Implementing a Management System Architecture Framework," *Bell Labs Tech. J.*, Vol. 5, No. 4, Oct.-Dec. 2000, pp. 31-43.
 24. H. I. Epstein, A. Asthana, S. A. Corum, and L. Mak, "Hybrid Network Management," *Bell Labs Tech. J.*, Vol. 5, No. 4, Oct.-Dec. 2000, pp. 63-79.
 25. D. Tsay, "Analysis of Service Operations Cost for TDM and Packet Tandem Networks," *Bell Labs Tech. J.*, Vol. 5, No. 4, Oct.-Dec. 2000, pp. 97-112.
 26. M. A. Eggert, R. A. Johnston and G. W. Vaughan, "The Flexent™ Element Management System—Using Web and Object Technologies to Manage a Wireless Network," *Bell Labs Tech. J.*, Vol. 5, No. 4, Oct.-Dec. 2000, pp. 113-125.
 27. J. P. Schlaerth, "Service and Network Manage-

ment Strategies for ATM in Wireless Networks," *Bell Labs Tech. J.*, Vol. 5, No. 4, Oct.-Dec. 2000, pp. 126-137.

(Manuscript approved February 2001)

SHERVIN ERFANI is a member of technical staff in the New Business Initiatives 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 Advanced Communications Technologies Vice President 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 Advanced Network Systems Engineering and CALA Network Planning, received a Ph.D. in electrical engineering and computer science from the University of California at Berkeley. Dr. Malek is currently working on next-generation service and network management architectures and implementations. He is the author, co-author, or editor of 7 books, and the



author or co-author of over 50 published technical papers. He is an IEEE Fellow, an IEEE Communications Society Distinguished Lecturer, and the founder and editor-in-chief of the Journal of Network and Systems Management. ♦