An OSI-Based Interoperability Architecture for Managing Hybrid Networks

Ian Sugarbroad

HYBRID NETWORKS IN TODAY'S ENVIRONMENT comprise a wide variety of equipment and services offered by different vendors and service providers. Managing these networks requires the utilization of an array of disparate (non-interworking) Network Management Systems (NMSs). Service providers—Inter-Exchange Carriers (IECs), Local Exchange Carriers (LECs), public packet networks, etc.—must maintain networks that physically consist of equipment provided by many vendors and logically require diverse signalling and transmission standards and protocols (X.25, SS7, PRA, BRA, SNA, etc.). Private networks that rely on the services provided by public networks are even more complicated. Designers of private networks utilize many different services provided by IECs, LECs, etc., and add further complexity by interworking these services with even more diverse equipment. The present solutions for managing these networks have severe limitations. The future solution being forged by the national and international standards organizations is an Open Systems Interconnection (OSI)-based interoperability architecture.

The Public Carrier Environment

This environment includes both LECs (telephone companies, or telcos) and IECs. Public carrier Network Management (NM) is provided by Operation Support Systems (OSSs)and by Network Operations Centers (NOCs).[1] OSSs collect data from the network elements and use this data to provide NM support to maintain and operate the internal telco network. This situation is illustrated in Figure 1.

The Hybrid Corporate Network Environment

A typical hybrid corporate network employs a private backbone network for its internal voice and data traffic, and a mixture of LEC and IEC services for overflow voice traffic. The private backbone network (voice) consists of mains (access switches) and nodes (tandem switches) linked by leased facilities. The backbone network may also employ T1 multiplexers for efficient bandwidth management. The IEC services include WATS services to access off-network locations and Virtual Private Network (VPN) services for overflow voice traffic. For

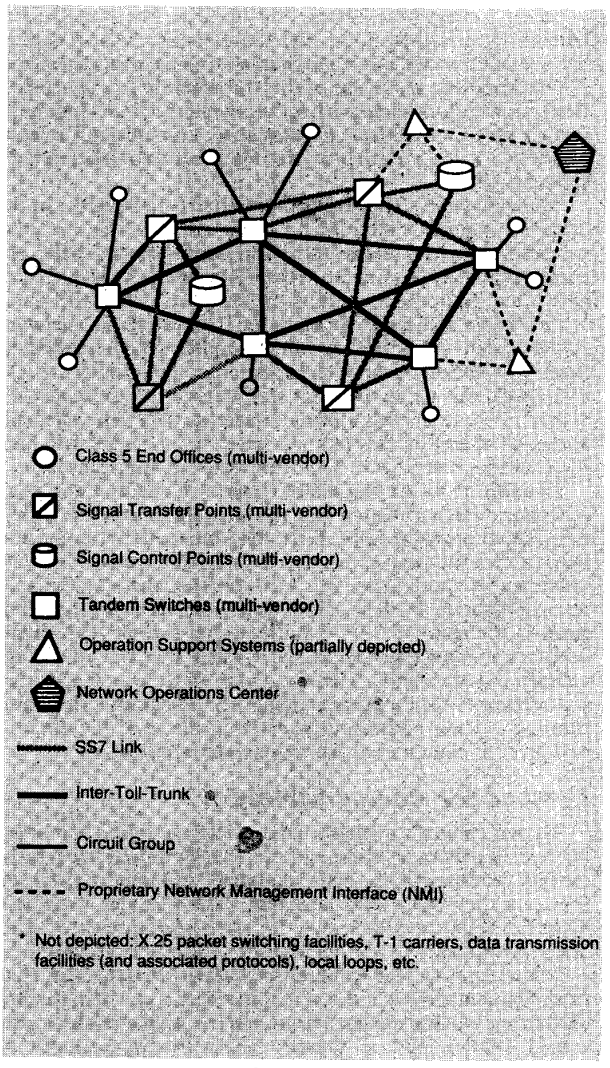

Fig. 1. Public carrier network.

[1]NOCs are served by hybrid OSSs that act as NM systems integrators. These OSS integration systems typically receive data from many subordinate OSSs for presentation to network managers within the NOC.

0163-6804/90/0003-0061 $01.00 © 1990 IEEE

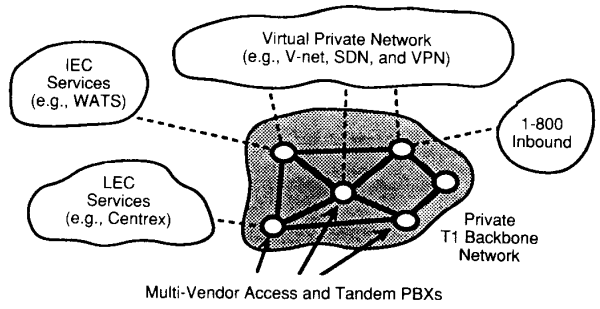

Fig. 2. Typical hybrid corporate network.

transaction-oriented data services, packet switching is employed, whereas bulk transfer and SNA networks employ circuit-switched and private-line-based data networks. A hybrid corporate network is thus an amalgam of equipment and NMSs from many vendors. The complexity of the typical hybrid corporate network is illustrated in Figure 2.

The Network Management Challenge

In the most general terms, NM involves the processing of Operation, Administration, and Maintenance (OAM) data to extract value-added information about the status and health of the network and make use of this information to take corrective actions. Thus, there are two components of NM—namely, surveillance and control. Furthermore, accounting, forecasting, network design, and provisioning are also considered to be elements of NM.

Managing hybrid networks is extremely challenging. In the telco environment, the NM functions are provided by OSSs maintained by the telcos. In hybrid corporate networks, the corporation may require the Private Branch Exchanges (PBXs) and other equipment to directly interface with third-party NMSs. In both cases, it is becoming imperative that network elements deliver OAM data using standard protocols and messages. In a hybrid network environment, there usually are several NMSs managing different segments of the network. The task of integrating NM information across different systems is a manual operation that requires network managers with high levels of expertise. Figure 3 illustrates the dilemma they face.

In addition to the formidable task of manually integrating real-time data, segmented NM (as illustrated in Figure 3) has many other disadvantages. First, the expense of obtaining and maintaining the varying systems is steep, and worse, it may be unknown in many instances. Also, because these systems do not interoperate, the loss of one system can leave critical blind spots in the network. Survivability of the overall NM system becomes an issue especially when systems are planned and developed independently. Finally, the job of the system administrator responsible for keeping all of the NM equipment up and running quickly becomes expensive in terms of the amount and level of expertise of the manpower required.

Demand for Interoperability

To overcome the difficult and unwieldy situation of managing hybrid networks, more and more corporations are requiring that vendors provide open standard OAM interfaces for the products they offer, and also that public service providers make OAM data available over the same open interfaces. This demand implies the eventual solution illustrated in Figure 4.

In this solution, the end user receives OAM data from Network Elements (NEs) and service providers (via OSSs) in standard formats using standard protocols (i.e., OSI standards, discussed below), and is thereby given the opportunity to integrate these data within a single system.

The OSI Standards

The purpose of OSI is to provide a common (nonproprietary) way of communicating between computers. OSI is defined in terms of standards and represents a tremendous undertaking by national and international standards organizations. The major contributors to the OSI suite of standards are the International Organization for Standardization (ISO) International Electrotechnical Committee (IEC) and the International Telephone and Telegraph Consultative Committee (CCITT), a division of the International Telecommunications Union (ITU), which is under the United Nations.

The ISO/IEC membership is made up from national standards bodies, including the American National Standards Institute (ANSI), the British Standards Institution (BSI), etc. CCITT members are the national Post, Telephone, and Telegraph administrations (PTTs). In countries where PTTs do not exist (e.g., the U.S.), an agency of the national government leads the delegation to the CCITT.

The ISO/IEC and the CCITT are relatively immune to market demands for standards. In the arena of NM, the demand is so strong for OSI-based NM that in 1988, a number of the major equipment vendors and service providers formed the OSI NM Forum (OSI/NMF) [1]. The Forum functions as a facilitator of OSI-based NM by making recommendations where standards do not exist, and by choosing appropriate standards for the varying tasks of managing hybrid networks.

Network Management System Management Functional Areas

The ISO/IEC has defined five categories of system management functional areas: fault management, accounting management, configuration management, performance management, and security management [2]. These functional areas are a collection of services that must work together across different systems to achieve NM interoperability. In more specific terms, these functional areas are defined below.

Fault Management

Fault management is the set of capabilities (services) that enables the detection, isolation, and correction of abnormal operation in the OSI environment. Faults cause systems to fail to meet their operational objectives. Faults manifest themselves as specific events (e.g., errors) in the operation of a system, and they may be persistent or transient. Fault management also includes the capabilities required to maintain and examine error logs; accept and act upon error messages; trace, isolate, and correct faults; and carry out test sequences.

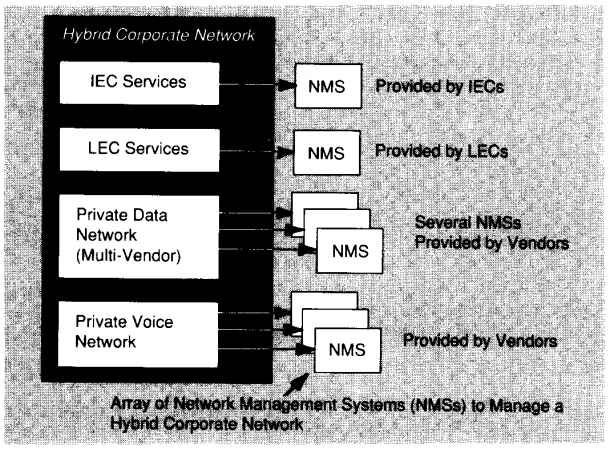

Fig. 3. Hybrid corporate network management.

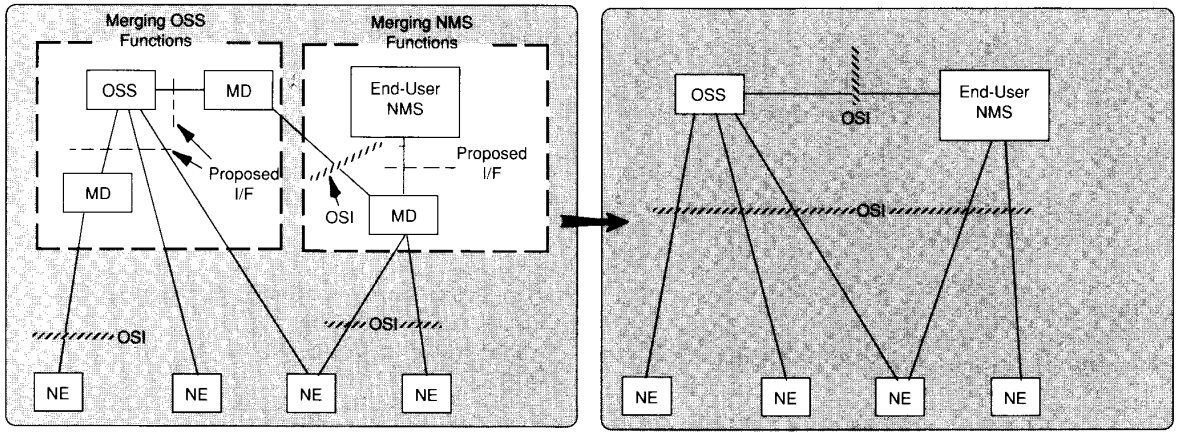

NE - Network Element
MD - Mediation Device
OSS - Operation Support System (Public)
NMS - Network Management System (Private)
OSI - Open Systems Interconnection (the standard that is the basis of network management interoperability)

Fig. 4. Proposed evolution plan for NM interoperability.

Accounting Management

This is the set of capabilities that enables charges to be established for the use of managed objects and costs to be identified for the use of those managed objects. Accounting management also includes the set of capabilities to inform users of costs incurred or resources consumed, enable accounting limits to be set for the use of managed objects, and enable costs to be combined where multiple managed objects are invoked to achieve a given communication objective.

Configuration Management

This is the set of capabilities that exercises control over, identifies, collects data from, and provides data to managed objects for the purpose of assisting in the continuous operation of interconnection services. Configuration and name management also includes the set of capabilities to set the parameters of the system; initialize and close down managed objects; collect data giving the open system state (both routinely and in recognition of a significant change of state); change the system configuration; and associate names with sets of managed objects.

Performance Management

This is the set of capabilities needed to evaluate the behavior of managed objects and the effectiveness of interconnection activities, gather statistical data, and maintain and examine logs of system state histories.

Security Management

Security management provides a set of capabilities to enable the control and distribution of information to various open systems for use in providing OSI security, and to report on the security provided and any security-relevant events that have occurred.

The Interoperability Architecture

Conceptual Architecture

The OSI/NMF operating environment consists of a group of systems known as Conformant Management Entities (CMEs) [1]. CMEs communicate with each other using the OSI/NMF protocol stack (OSI P-stack) to exchange OSI/NMF-defined messages ("M" messages). This information is exchanged for the purpose of NM. The interface over which two CMEs communicate is known as the P + M interface (illustrated in Figure 5).

A managed system (represented by one or more managed objects) is a physical or logical resource in the network that is made visible over the P + M interface via some CME and is subject to being managed in some capacity. Each CME has control of one or more managed objects, and the CMEs may themselves be represented by managed objects.

The Forum's policy is to use existing standards and draft standards from the international arena (mainly ISO/IEC and CCITT). The Forum will generate interim technical recommendations to fill gaps where they do not exist in order to achieve interoperability. These interim solutions will be reviewed and replaced as international standards emerge.

Physical Example

Consider a network that shares information over fully open and standard interfaces (see Figure 6). This example represents the operating environment that is the goal of the network management interoperability architecture.

This environment will promote the uninhibited exchange of NM messages and OAM data from NE to OSS, from NE to NMS, from OSS to NMS, from NMS to NMS, etc. As this type of interworking is achieved, the technical constraints of pairwise proprietary interfaces facing NM/OAM network integrators will be lifted.

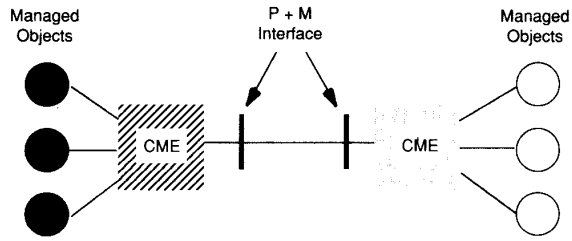

Fig. 5. OSI/NMF operating environment.

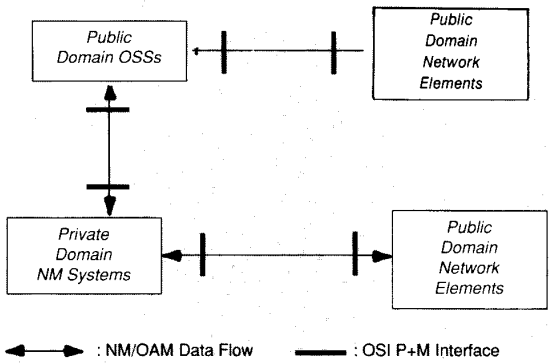

Fig. 6. Interoperability.

The OSI P-Stack

The OSI P-Stack is based on the seven-layer OSI model. Each layer is defined in terms of protocol standards produced by ISO/IEC (noted as ISO ####) and by the CCITT (noted as X.###). Figure 7 illustrates the standards that have been recommended by the OSI/NMF. It also represents how the management specific service elements (the "Ms" in the "P+M" interface) fit within the scheme of the P-stack. The current definition of the message sets requires only the use of the Common Management Information Service Element (CMISE) protocol. This may change as the definition of these standards progresses.

Managed Objects and Attributes

A managed object is a representation (view) of a data processing, data communications, or telecommunications resource. As such, a managed object may be manipulated through the use of an OSI Management protocol (e.g. the OSI P-stack). A managed object may represent a physical item of equipment, software component, some abstract collection of information, or a part or combination of these items. Every managed object has a set of management information (attributes) associated with it, and a managed object identifier that is used to identify and access the managed information through OSI management services.

A managed object class is a generic classification shared by a set of similar managed objects that have similar properties and are used for like purposes. A managed object may be identified as a distinguishable unit, and managed object classes are defined because it is convenient to manipulate those sets of management information as individual, nameable units. Managed objects of the same class share the same set of attributes, although the values of that information may differ from one instance of the class to another.

Something that has been defined as a managed object may consist of many physical (and logical) components. However, the attributes of the object represent the only knowledge that OSI management services have about this object.

For purposes of network management and control, managed objects may be created and deleted. Creation of a managed object is achieved by placing its identifier and a set of management information appropriate to its class in the Management Information Base (MIB) so that it may be accessed through the OSI management services. Deletion is achieved by removing this information from the MIB. The class of a managed object defines the set of attributes, and the value constraints for that information, possessed by every instance of that class and thus defines the management information that must be supplied when a managed object instance is created. The creation and deletion of a managed object ordinarily has some effect upon the operation of the open system where the managed object resides.

Management-Specific Application Service Elements

Management-specific service elements are the protocols necessary to provide the functionality of the five management areas. Although the ISO/IEC is currently working in these areas, the standards are being pushed by the OSI/NMF, which has published their recommended Application Service Elements (ASEs) [3]. The ASEs are provided through the use of CMIS services.

Although many of these ASEs are currently being defined, the following three are offered as examples (later, a networking scenario is presented illustrating the use of these ASEs):

- Change Attribute Service
- Report Attribute Change Service
- Create Object Service

Figure 8 illustrates the message flows (at a high level) associated with these services.

The Create Object Service

This service is provided via the CMIS M-CREATE service. It allows one CMISE user (the invoker) to request another CMISE user (the performer) to create a new managed object of a specified class. The new managed object is created with attributes as specified by its class definition. The attribute values can be specified explicitly by the invoker, derived from the attribute values of a reference object of the same class specified by the invoker, or set according to the default values specified in the definition of the managed object class. The identification of the instance of the new managed object can be either explicitly supplied by the invoker or automatically generated by the performer. Table I is a list of the parameters required by this service.[2]

Table I. Create Object Service Parameters

Parameter Name	Request/ Indication	Response/ Confirmation
Invoke Identifier	M	M (=)
Managed Object Class	M	C
Managed Object Instance	U	C
Access Control	U	–
Reference Object Instance	U	–
Attribute List	U	C
Current Time	–	U
Errors	–	C

The Change Attribute Service

This service uses the CMIS M-SET service to allow a CMISE invoker to request a CMISE performer to change the attribute value(s) of one or more managed objects in the provider's MIB. The parameters associated with this service are illustrated in Table II.

The Report Attribute Change Service

This service allows one CMISE user to announce to another CMISE user that the values of one or more attributes of one or more managed objects in the announcer's MIB have changed.

[2]In the tables in this article, "M" indicates a mandatory parameter, "=" indicates that the value of the parameter is equal to the value of the parameter in the column to the left, "U" indicates that the use of the parameter is a service-user option, "–" means that the parameter is not present in the interaction, and "C" means the parameter is conditionally present (conditions are defined by the text that describes the parameter).

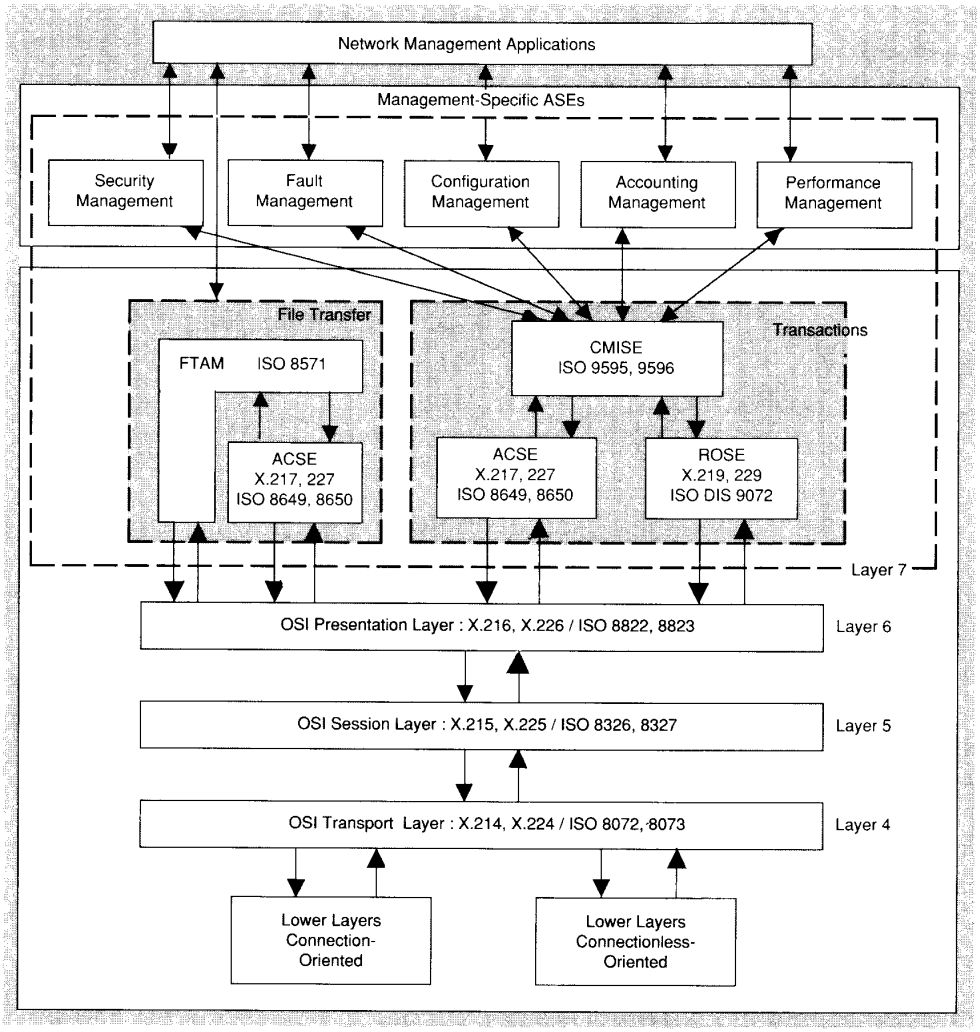

Fig. 7. OSI/NMF recommended OSI protocol stack.

This service is provided via the CMIS M-EVENT-REPORT service, and the parameters are listed in Table III.

Examples

Corporate Banking Voice Network

The Network Configuration

This example is predicated on a hypothetical configuration of The "Bigger Bank's" (BB) corporate voice network. This configuration is illustrated in Figure 9.

In addition to its headquarters, Bigger Bank has several regional offices and several branch offices. The branch offices (only one of which is shown in Figure 9 for simplicity) are connected to the headquarters by Centrex lines provided by the Regoinal Bell Operating Company (RBOC). The Centrex services are provided by a Class 5 Central Office (CO) switch, which in turn is served by an OSS supporting an OSI P+M interface. The regional offices are large enough to cost-justify PBXs that are directly connected to the Large Private PBX/Tandem switch in the headquarters building with private tie-lines.

The private NMS in this network provides many NM services, including station administration functions (moves, adds, and changes) for all sets, including those connected to the corporate network through an RBOC CO. This is possible

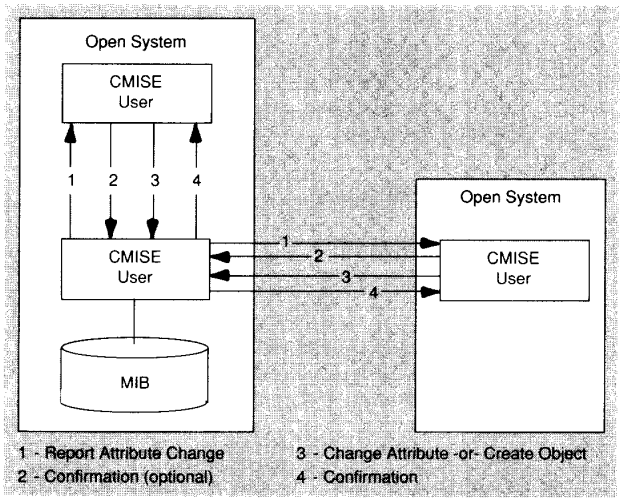

Fig. 8. Application service message flows.

through the OSI P+M interface using CMISE-based system management functions. When a station administration func-

Table II. Change Attribute(s) Service Partameters

Parameter Name	Request/ Indication	Response/ Confirmation
Invoke Identifier	M	M (=)
Linked Identifier	–	C
Mode	M	–
Base Object Class	M	–
Base Object Instance	M	–
Scope	U	–
Filter	U	–
Access Control	U	–
Synchronization	U	–
Managed Object Class	–	C
Managed Object Instance	–	C
Attribute List	M	C
Current Time	–	U
Errors	–	C

tion is carried out on a set served by Centrex, the RBOC equipment must be utilized and the RBOC must be able to accurately charge Bigger Bank for the use of their resources. This may be done through the use of Accounting Management services.

Accounting Management Actions

Consider the following situation. The Bigger Bank system administrator is carrying out a series of changes on the Centrex stations in a specific branch office. Each change represents a billable transaction from the viewpoint of the RBOC, even though the RBOC personnel are not required to physically take part in the transaction. Figure 10 illustrates the message flows associated with a single change of a Centrex user profile.

The following list is an explanation of the message transfers illustrated in Figure 10:

- *Message 1*: Change Attribute service (using CMIS M-SET). The network system administrator (using the private NMS) requests a change to a Centrex user's profile. This is carried out by the Configuration Control object within the private NMS requesting the Configuration Control object within the OSS to change the specified attribute of the Station object within the OSS management information base.
- *Message 2*: Change Attribute service confirmation. This message is a preliminary confirmation by the Configuration Control object within the OSS notifying the Configuration Control object within the NMS that the change attribute request has been received (but not necessarily executed).
- *Message 3*: Change Attribute service (using CMIS M-SET). The Configuration Control object within the OSS (acting as the M-SET service invoker) changes the attribute of the specified Station object.
- *Message 4*: Change Attribute service confirmation. The Station object notifies the Configuration Control object that the attribute has been changed.
- *Message 5*: Report Attribute Change service (using CMIS M-EVENT-REPORT). The Station object reports the change of the attribute to the Account Management Control object within the OSS. The attribute change represents a utilization of logical resources. The reason that the Station object "knows" to report the change activity to the Account Management Control object is that the Account Management Control object (at some earlier time) instructed it to do so by using the Initiate Resource Utilization Reporting service.
- *Message 6*: Report Attribute Change service confirmation.
- *Message 7*: Create Object service (using CMIS M-CREATE). The Account Management Control object creates a record of the resource utilization within the Account Log.
- *Message 8*: Create Object service confirmation.
- *Message 9*: Report Attribute Change service (using CMIS M-EVENT-REPORT). The OSS's Configuration Control object reports the change of the Station object's attribute to the Configuration Control object within the private NMS.
- *Message 10*: Change Attribute service (using CMIS M-SET). The Configuration Control object within the private NMS changes the attribute of its corresponding instance of the Station object within its own MIB.
- *Message 11*: Change Attribute service confirmation.

Table III. Report Attribute Change Service Parameters

Parameter Name	Request/ Indication	Response/ Confirmation
Invoke Identifier	M	M (=)
Mode	M	–
Managed Object Class	M	U
Managed Object Instance	M	U
Event Type	M	C (=)
Event Time	U	–
Event Argument	M	–
Current Time	–	U
Event Result	–	C
Errors	–	C

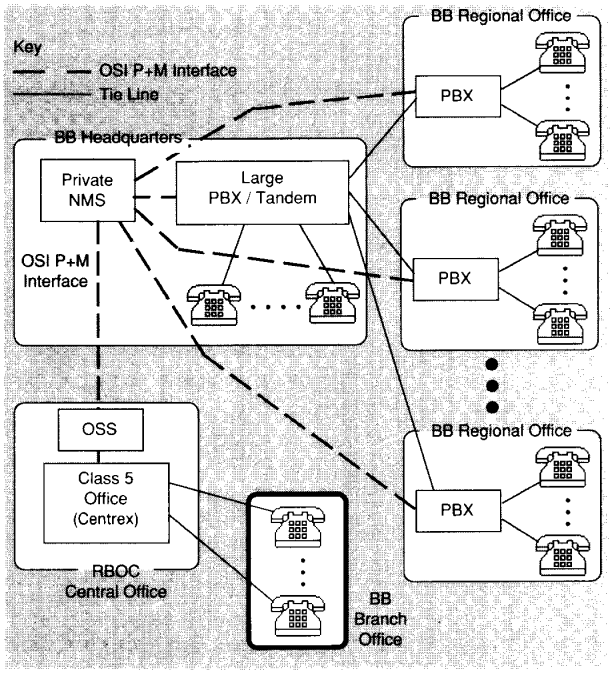

Fig. 9. Bigger Bank corporate voice network.

OSI-Based Configuration Management Applied to Automatic Call Distribution

Existing Automatic Call Distribution Systems

Automatic Call Distribution (ACD) is a major software component of many advanced PBXs. ACD provides equal distribution of calls to predesignated answering positions. If all positions are busy, calls are queued in order of their arrival, taking into account the call's priority. Figure 11 is an illustration of this system.

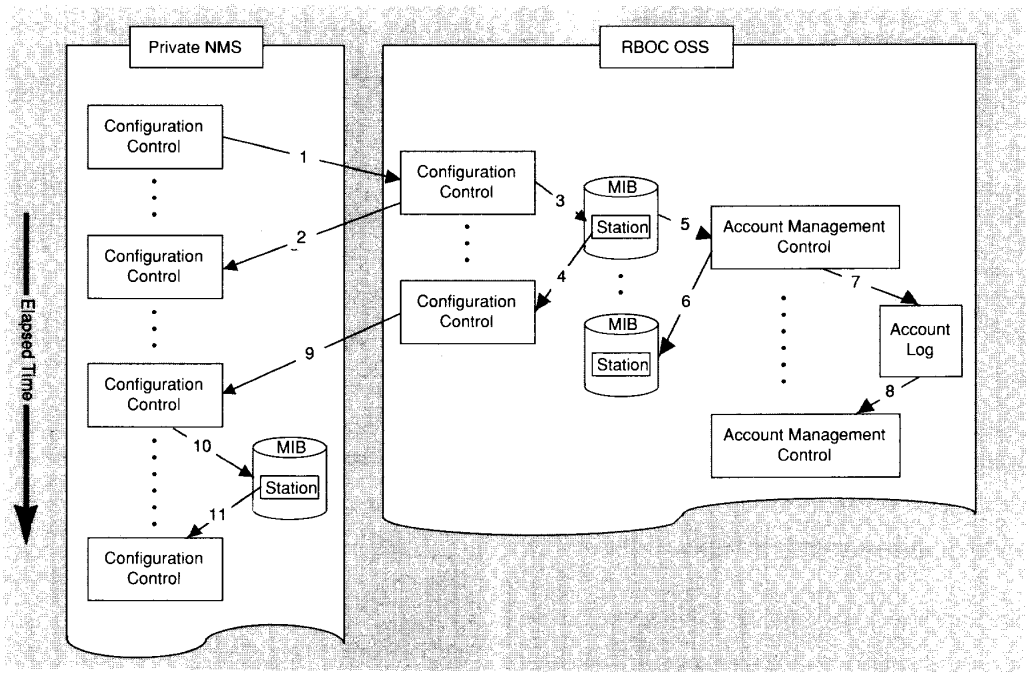

Fig. 10. Billed Centrex change operation message flow.

Attendant consoles are always in one of four states: Idle, Busy, Not Ready, or Deactivated. These four states may be mapped into to the following states, defined in the OSI/NMF State Model [3]:

- Idle: Enabled/Unlocked
- Busy: Busy/Unlocked
- Not Ready: Enabled/Locked
- Deactivated: Disabled/Locked (or Disabled/Unlocked)

The state transition diagram is illustrated in Figure 12.

Typical ACD systems do not keep track of individual attendant positions. Rather, they track the membership of three queues: the Idle Queue, the Busy Queue, and the Not Ready Queue. When a call comes in, it is terminated on the first member of the Idle Queue and this station is moved to the Busy Queue. When an attendant finishes serving a customer or is disconnected from the customer, this station is moved from the Busy Queue to the Idle Queue. An administrative action is required to move a station to or from the Not Ready Queue.

Since queues are prone to corruption, ACD systems must perform audits to check for irregularities in these queues and take corrective action to repair corrupted queues. ACD log subsystems provide information gained as a result of these audits and, in some cases, can make repairs.

The OSI Solution Applied to Automatic Call Distribution

The real-time performance of ACD systems would benefit if the operational mode of querying the queues and flushing the ones that are found to be in error was replaced with a less severe management method such as the one described in this section.

To start, the queues and the attendant positions should be defined as objects that comprise the ACD system (also an object). These objects are related directly (and indirectly) through group relationships illustrated in Figure 13.

These objects (and their respective relationships) would be contained in an MIB controlled by the ACD management system. The existing software that pairs incoming customers with

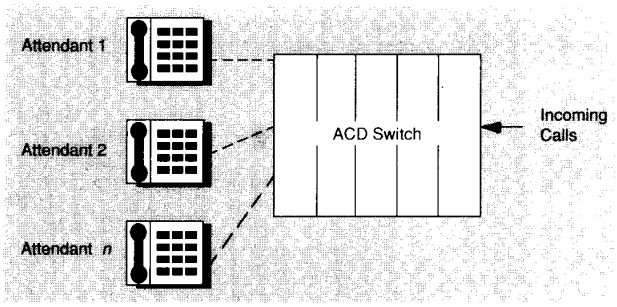

Fig. 11. Automatic call distribution.

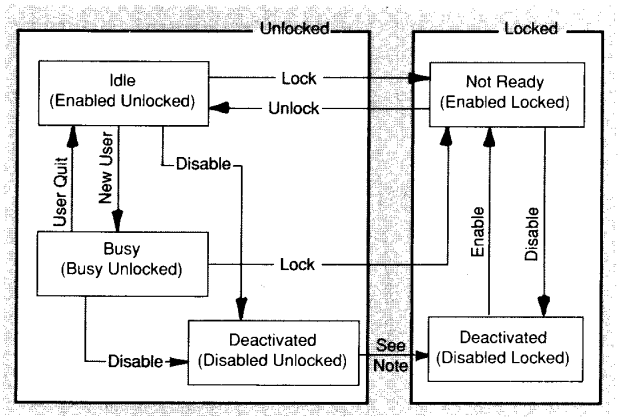

Note: After disabling an attendant console (moving it to the "Deactivated" state from either "Idle" or "Busy") it is necessary to move it to the "Not Ready" queue before it may be moved to any other queue. This implies that any station taken to the "Deactivated" (Disabled Unlocked) state by disabling it, will automatically transition to the "Deactivated" (Disabled Locked) state.

Fig. 12. ACD attendant position state transitions.

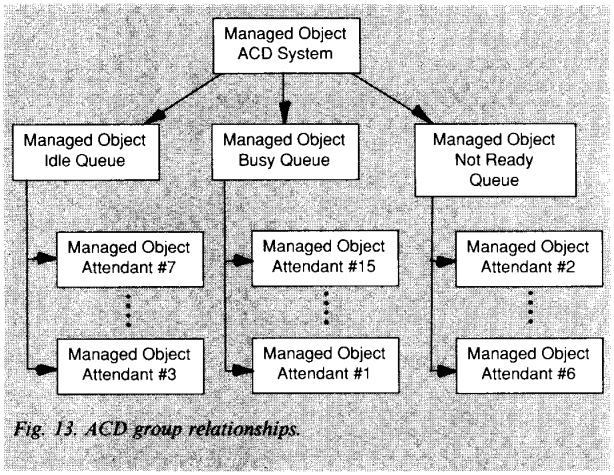

Fig. 13. ACD group relationships.

attendant positions in the idle queue could continue to operate as it does today. However, the queues (as represented by objects in an MIB) could be managed by OSI-based services without the disruptive practice of auditing and flushing corrupt queues.

Consider the following routine activity within an ACD system:

- Attendant Position #5 is terminated upon by an incoming customer because it is the first attendant position in the Idle Queue. Attendant Position #5 is then moved from the Idle Queue to the Busy Queue.

Figure 14 illustrates the initial state of the managed objects and their attributes that are involved in this activity, and Figure 15 illustrates the final state of the same objects after the attendant position has changed its queue membership. This change could be accomplished by a single CMISE operation (i.e., the Change Attributes Service) between the ACD Management Algorithm and the MIB.

Note that Figures 14 and 15 do not show the managed objects corresponding to the Not Ready Queue or any of the attendant positions (other than Attendant Position #5). Note also that the attribute lists shown are not necessarily exhaustive.

Conclusions

As the implementation of OSI-based systems progresses, achieving widespread interoperability will become technically feasible. Service providers will be able to provide tariffed NM services and redeploy the human resources presently required for the tedious task of NM data integration into more productive and lucrative positions. Corporations will at last be able to efficiently manage their networks through automated integration of NM data and centralized control.

There are still several hurdles to clear before OSI-based systems become the norm. Network modelling (creating common objects and messages) is still in its infancy. Even as generic network models are standardized, totally new objects will have to be defined to support new services and new devices. A common human-machine interface does not exist for the portrayal of NM data. In the absence of standards, it may prove difficult for an NM system to display a piece of information in anything other than an *ad hoc* manner. Existing *de facto* standards (e.g., SNA, DECNET, etc.), which are closed and proprietary, are firmly embedded in most networks all over the world. The conversion of these networks to OSI will be a tremendous under-

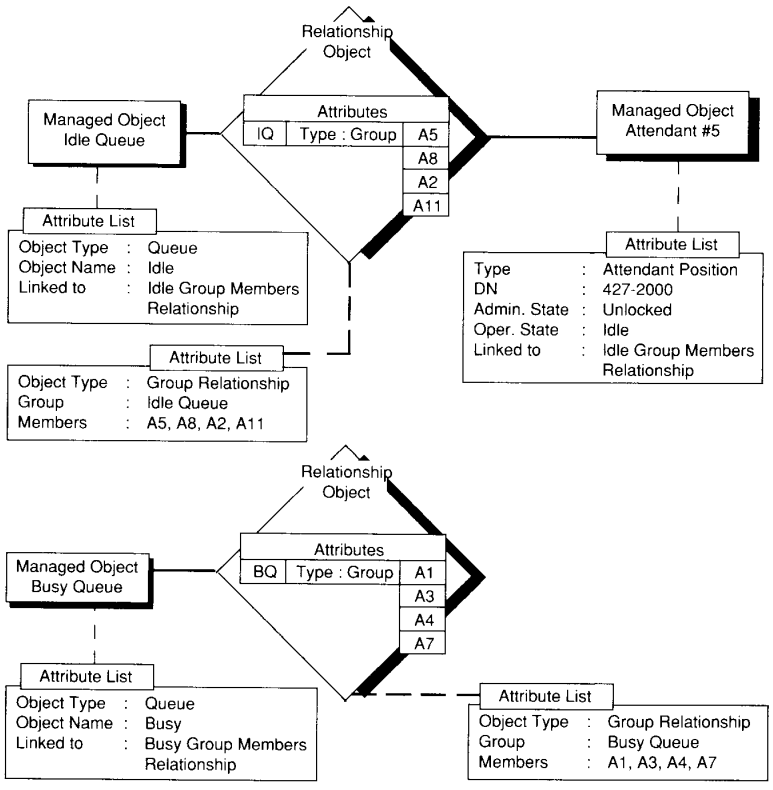

Fig. 14. ACD managed objects—initial state.

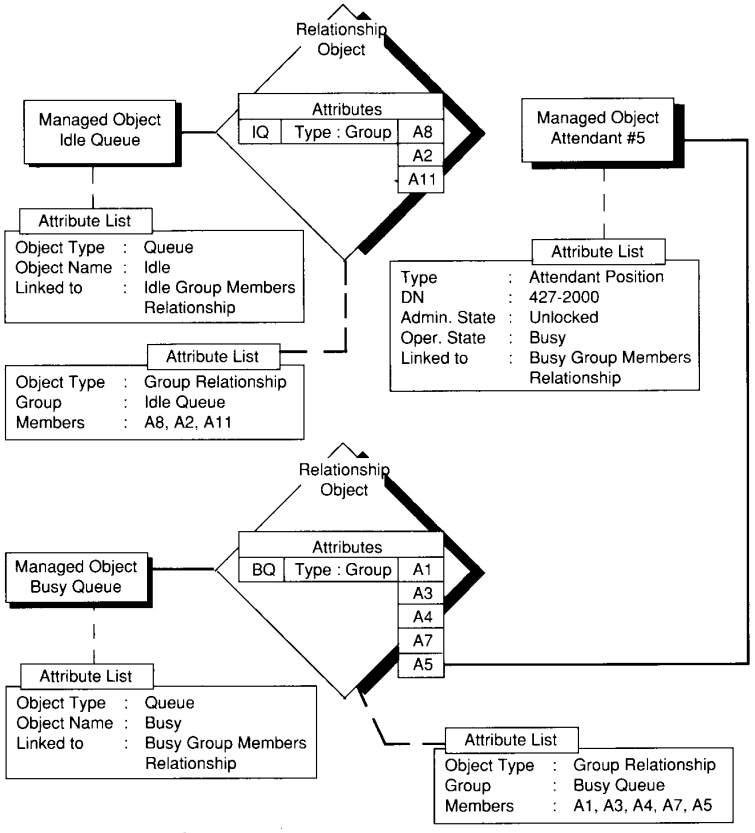

Fig. 15. ACD managed objects—final state.

taking and not without a lot of pain and anguish for users and service providers alike.

Finally, where the domain of the standards ends must be well understood. In Figure 7, the topmost element consists of the NM applications. Everything below the NM applications exists for the benefit of the applications. The applications themselves will never be standardized because the potential power of efficient value-added network management applications represents tremendous revenue opportunities for those who embrace the task of migration to OSI.

References

[1] *OSI/NM Forum Protocol Specification*, Issue 1, Jan. 1989.

[2] ISO/IEC JTC1/SC 21 N 2058, *Information Processing Systems—Open Systems Interconnection—Management Information Service Definition—Part 1: Overview*, Dec. 1987.

[3] *OSI/NM Forum Application Services*, Issue 1, June 1989.

Biography

Ian Sugarbroad holds Master of Science and Master of Business Administration degrees from the University of Western Ontario, Canada, and a Bachelor of Science degree from the University of London, England.

He is Vice President of Network Systems for the Data Communications and Networks group of Northern Telecom Inc., located in Richardson, Texas. He is responsible for product line management and technology development for large-scale corporate telecommunications networks. He also represents Northern Telecom on the board of the OSI Network Management Forum, where he is currently serving as Secretary.

Previously, as General Manager for Broadband Integrated Services Digital Network (BISDN), he managed Northern Telecom's initial ISDN application trials, including the first ever in the public network, launched in Phoenix, Arizona in the fall of 1986. A former research physicist and financial analyst, he has held several marketing, administrative, manufacturing, and technology positions since joining Northern Telecom in 1976, including product line management, strategic marketing, and international marketing positions for for transmission products.