NIST Special Publication 500-175

Computer Systems Technology

U.S. DEPARTMENT OF COMMERCE
National Institute of Standards and Technology

NIST

# Management of Networks Based on Open Systems Interconnection (OSI) Standards: Functional Requirements and Analysis

Robert Aronoff
Michael Chernick
Karen Hsing
Kevin Mills
Daniel Stokesberry

NIST Special Publication 500-175

# Management of Networks Based on Open Systems Interconnection (OSI) Standards: Functional Requirements and Analysis

Robert Aronoff
Michael Chernick
Karen Hsing
Kevin Mills
Daniel Stokesberry

National Computer Systems Laboratory
National Institute of Standards and Technology
Gaithersburg, MD 20899

November 1989


U.S. DEPARTMENT OF COMMERCE
**Robert A. Mosbacher, Secretary**
NATIONAL INSTITUTE OF STANDARDS
AND TECHNOLOGY
Raymond G. Kammer, Acting Director

NIST

## Reports on Computer Systems Technology

The National Institute of Standards and Technology (NIST) (formerly the National Bureau of Standards) has a unique responsibility for computer systems technology within the Federal government. NIST's National Computer Systems Laboratory (NCSL) develops standards and guidelines, provides technical assistance, and conducts research for computers and related telecommunications systems to achieve more effective utilization of Federal information technology resources. NCSL's responsibilities include development of technical, management, physical, and administrative standards and guidelines for the cost-effective security and privacy of sensitive unclassified information processed in Federal computers. NCSL assists agencies in developing security plans and in improving computer security awareness training. This Special Publication 500 series reports NCSL research and guidelines to Federal agencies as well as to organizations in industry, government, and academia.

**Library of Congress Catalog Card Number: 89-600778**
**National Institute of Standards and Technology Special Publication 500-175**
**Natl. Inst. Stand. Technol. Spec. Publ. 500-175, 130 pages (Nov. 1989)**
**CODEN: NSPUE2**

**U.S. GOVERNMENT PRINTING OFFICE**
**WASHINGTON: 1989**

For sale by the Superintendent of Documents, U.S. Government Printing Office, Washington, DC 20402

ABSTRACT

To provide for management of future interoperable multi-vendor networks, the ISO and other international organizations are currently developing management standards for communications networks based on the Open Systems Interconnection (OSI) Reference Model. This report examines current and proposed network management systems to determine both user and functional requirements for network management. It then compares the derived functional requirements to the emerging standards being developed by the ISO and others to determine where and how the requirements are being met by these emerging standards. In those cases where requirements are not being met, these deficiencies are noted.

The examination of requirements is generally restricted to those that are necessary for interoperability. These are organized and examined in six broad areas: Architecture, and the management functional areas of configuration management, fault management, security management, performance management and accounting management. This report also contains a discussion of requirements that transcend these areas and a discussion of future requirements beyond the scope of current standardization. Such requirements, while not necessary for interoperability, are useful nonetheless. Examples include automated network management and a standard operator interface. Finally this report also contains a discussion of applying OSI management to the emerging Integrated Services Digital Network (ISDN) technology.

Keywords:
Automated Network Management; ISDN Management; Network Management; OSI Functional Requirements; OSI Management

Note: Drafts of this report were previously released under the title "Network Management Functional Requirements."

Acknowledgments

The authors would like to express their appreciation to the following people who have helped produce this report by their expertise and insightful comments: Dr. Paul J. Brusil of The MITRE Corporation, Bedford, MA; Eugene Berger of The MITRE Corporation, Houston, TX; Dr. Dennis Branstad of the National Institute of Standards and Technology; Wayne McCoy of the National Institute of Standards and Technology; Richard Colella of the National Institute of Standards and Technology; and the many others who have contributed to this report and earlier drafts.

TABLE OF CONTENTS

# 1. EXECUTIVE SUMMARY

This executive summary gives a management level overview of the background, purpose and content of this report. It also outlines the major issues and conclusions within the report. The network management functional requirements have been analyzed by members of the staff of the National Institute of Standards and Technology (NIST), formerly the National Bureau of Standards, as one element of a program of work that is jointly sponsored by several government agencies.

NIST and its sponsors have identified the present absence of network management products and services as a serious shortcoming to the construction and operation of large-scale, integrated, multi-vendor networks. This program of work seeks to expedite the availability of commercial network management products through the active participation in the continuing progression of international standards and the specification of implementation agreements based upon those standards.

This report builds upon previous work at NIST - it is an update and expansion of an earlier draft on the same subject produced by NIST as the second of two related reports on network management standardization for the United States Air Force Mission Effective Information Transmission System (MEITS) program. The first report [NBS87] documented the state of affairs with respect to network management standardization, pointing out several areas of potential problems where the NIST and the MITRE Corporation, a research partner of the NIST in this area, could focus their efforts. The second report concentrated on the user and functional requirements that implementations of the network management standards must satisfy. An important aspect of the report was the identification of significant issues in the development of international standards for management of communications systems. The initial draft of this report was issued in October 1987 (under the title "Network Management Functional Requirements") and distributed to approximately 60 experts who attended a workshop in Bedford, MA on October 28-30, 1987. The purpose of the workshop was to review the contents of the report to find areas of disagreement, missing issues, and other opinions. The draft was also widely distributed for comment throughout the network management community. The present report incorporates the output of the workshop discussions and subsequent comments received.

This report also includes several major additions to the previous draft in areas that have become increasingly important since its release. This includes a discussion on management for Integrated Services Digital Network (ISDN), a new technology rapidly gaining user interest. Other major areas are now discussed in greater detail: specifically, the areas of security management, accounting management, and automated network management systems (the application of expert systems to network management problems).

## 1.1. Background

The focal point for international standardization of communications systems is the Open Systems Interconnection (OSI) reference model and the

related standards developed within the International Organization for Standardization (ISO). OSI standardization has reached the point where implementation agreements are in place for all seven layers of the Open Systems Interconnection (OSI) reference model. Products are on the market, or soon will be, supporting electronic mail and file transfer applications over a variety of local and wide-area networks. The number of applications supported will soon increase to include virtual terminal, directory services, and transaction processing. The deployment of commercially available, nonproprietary, interoperable, multi-vendor data communication products is about to commence; however, to truly realize multi-vendor networks of significant scale, interoperable network management must be achieved.

OSI management functionality supports the location and correction of faults, the establishment and adjustment of configurations, the measurement and tuning of performance, the control of security, and the collection and reporting of billing and accounting information. Such functionality is needed in end systems (hosts), intermediate systems (routers), and other network elements (e.g., bridges, switches, modems, and multiplexors).

The OSI management standards are in the middle stages of their development, but they are beginning to progress rapidly. Availability of a complete set of definitive management standards cannot be expected before 1992. An additional 1 to 2 years will probably be required before implementations based on these standards become commercially available. Two of the standards pertinent to the structure of management information and five of the standards pertinent to the system management functions reached DP status in ISO in December 1988. The management framework standard and the standards for the Application Layer Common Management Information Protocol (CMIP) and Services (CMIS) are now at the DIS level. Standards for the specification of managed objects are now in the early stages of development in ISO, ANSI, CCITT, and IEEE.

The U.S. Government cannot afford to wait until post-1994, the earliest that full OSI management standards, implementor agreements, and products are expected to be available. Five years without any OSI management will permit continued expansion of vendor market share through means other than cost and performance, i.e., through proprietary, noninteroperable network management. Accordingly, the U.S. Government requires initial network management specifications that provide a useful subset of the full OSI management functions. It is desirable to specify the initial subset in such a way that it is easy to add other capabilities to reach the full set of management functions.

Recognizing such a need, the NIST has initiated a program to further the development of OSI systems management standardization. The immediate goal is publication of the interim network management specification as a Federal Information Processing Standards (FIPS) by 1990. The desire is to base the interim network management FIPS upon implementor agreements developed by consensus within a public forum composed of both vendors and users of network management products. The long-term goal is to continue to pursue the development of full network management standards leading to vendor products. Once implementor agreements are available for full network management, the

FIPS can be revised to reflect the new capabilities. This FIPS will be a suitable reference for GOSIP, the Government Open Systems Interconnection Profile.

The NIST network management program includes three major activities: development of the implementation agreements, active participation in the basic network management standards process, and research that supports these activities through development of prototype implementations of network management systems.

The focal point of the activity to develop suitable implementation agreements is the NIST OSI Implementors Workshop (OIW). Approved international standards (ISs) for OSI do not lead directly to interoperable implementations in multi-vendor products. The typical IS contains a number of incompatible subsets and options that hinder interoperability. To achieve interoperable commercial products, the NIST established, in 1983, an open forum where implementors and users of OSI products meet to reach specific agreements concerning the protocols, subsets, and options to be implemented. The output of these workshops is a documented set of agreements that point the way to implement interoperable OSI products. Several groups have adopted the workshop output as the basis for functional profiles, including General Motors for Manufacturing Automation Protocol (MAP), Boeing Computer Services for the Technical and Office Protocols (TOP), and the U.S. Government for GOSIP. In addition, the Corporation for Open Systems (COS) uses the workshop output as the basis for conformance testing profiles.

The OIW is organized as a set of special interest groups (SIGs) addressing such subjects as the lower and upper OSI layers, electronic mail, file transfer, virtual terminal, security, office document interchange, and directory services. The NIST has successfully supported the establishment of a SIG on the topic of network management and NIST network management experts have provided leadership in the activities of this SIG, where the required OSI network management implementor agreements are now being developed.

The standards participation activity has taken NIST staff members into a number of fora including national and international OSI management standards committees (ANSI ASC X3T5.4 and ISO/IEC JTC 1/SC 21/WG 4), OSI layer management committees (ANSI ASC X3S3.3 and ISO/IEC JTC 1/SC 6), the MAP network management task force, the IETF on network management, COS Network Management Subcommittee (NMSC), the International Federation for Information Processing (IFIP) network management working group (IFIP WG 6.6), and the Institute of Electrical and Electronics Engineers (IEEE) local area network management subcommittee (P802.1). The sheer number of groups working on network management issues illustrates both the importance and complexity of the field. Table 1 presents a list of standards organizations working on OSI NM.

The third area of NIST network management activity concerns laboratory research and development of prototype implementations of management systems. The proposed standards are not based upon any existing network management technology. In addition, the proposed standards surround and cut across all

**Table 1: Network Management Standards Activities**

| Management Element | Standards[1] Group | Work Items | Status[2] | Estimated Completion Date |
|---|---|---|---|---|
| Architecture | ISO SC21/WG4 | OSI Management Architecture | IS | Complete |
| | IEEE 802.1 | LAN layer-management Architecture | WD | Undecided |
| | CCITT SG VII | Telephony network management Architecture | Work starting | 1990 |
| Management Communication Services and Protocols | ISO SC21/WG4 & CCITT SG VII & IAB NetMan | Common Management Information Services (CMIS) and Common Management Information Protocol (CMIP) | DIS<br>Work starting<br>RFC | 1989-1990<br>1990<br>Complete |
| | IEEE 802.1 | LAN Layer-management protocol | WD | Undecided |
| | IAB | Simple Network Management Protocol, a transition protocol for managing the internet before OSI's CMIP/ CMIS are deployed | RFC | Complete |

---

1 ISO: International Standardization Organization; SC: Subcommittee; WG: Working Group; CCITT: International Consultative Committee on Telegraphy and Telephony; SG: Study Group; IAB: Internet Activity Board; ANSI: American National Standards Institute; ASC: Accredited Standards Committee; MIB: Management Information Base.

2 WD: Working Draft; DP: Draft Proposal; DIS: Draft International Standard; IS: International Standard; RFC: Request For Comment (IAB's equivalent of a standard); ISO standards work proceeds from WD through DP to DIS to IS.

Table 1: Network Management Standards Activities (cont.)

| Management Element | Standards Group | Work items | Status | Estimated Completion Date |
|---|---|---|---|---|
| System Mgmt Functions | ISO SC21/WG4& ANSI T1M1.5& CCITT SG VII | Configuration Management & Fault Management, | WD | Undecided |
| | | Performance, Accounting and Security Management, | WD | 1991-1993 |
| | | Common Functions such as state management, error reporting used in systems management | DP | 1991 |
| Managed Objects | ISO SC21/WG4& ANSI X3T5.4 | Defining structures, formats and guidelines for managed object definitions (structure of management information) | DP | 1991 |
| | ISO SC21/WG4 | Defining parameters to be managed for systems (WG4: systems identification and serial numbers, for example) | Ranges from work starting to DP | Undecided |
| | ISO SC21/WG5 | Defining parameters to be managed for upper-layer protocols. For example, which system is to initiate sending | Ranges from work starting to DP | Undecided |
| | ISO SC6/WG2& ISO SC21/WG4 | Defining parameters to be managed for lower-layer protocols. For example, timers specifying re-transmission timeouts and timers registering number or packets sent | WD | 1991 |
| | IEEE 802.2- | Defining parameters to be | Ranges | Undecided |

Table 1: Network Management Standards Activities (cont.)

| Management Element | Standards Group | Work items | Status | Estimated Completion Date |
|---|---|---|---|---|
| Managed Objects (cont.) | 802.10 | managed for lower-layer protocols for LANs and metropolitan area NWs includes security | from beginning effort to DIS | |
| | ANSI ASC X3T9.5 | Defining parameters to be managed for high-speed fiber-optic LANs | Work starting | Undecided |
| | ANSI ASC T1M1.5 | Defining parameters to be managed for telecommunication devices such as multiplexers | WD | Undecided |
| | CCITT various SGs | Defining parameters to be used in communications such as those for ISDN | Work starting | Undecided |
| | IAB MIB WG | Defining parameters to be managed for the Internet's Transmission Control Protocol/Internet Protocol | RFC | Version 1 Complete |

seven layers of the OSI reference model. The developers of network management services and protocols need the feedback that only laboratory experience can provide. Thus, the NIST has outlined a general program of laboratory research and development intended to address the needs of its sponsors, to provide insight and feedback to standards setters, and to point the way for transition between emerging Department of Defense (DoD) network management protocols and the developing OSI management protocols.

### 1.2. Purpose

The purpose of the study, documented by this report, is three-fold: 1) to determine the functional requirements that OSI systems management standards makers must satisfy, 2) to evaluate the progress toward meeting the requirements, and 3) to examine what position the U.S. Government should take within GOSIP regarding network management standards for OSI systems. The report is intended to provide standards setters with a set of achievable goals, to indicate areas where the standards process is lagging, and to provide a perspective on what portion of network management functionality is covered by the emerging standards.

Within the body of the present report there are many user requirements for network management that are outside the scope of the standards. In each instance these extra-standardization requirements are identified as such. The most important effect of such requirements on standards making is that emerging standards must not preclude the possibility of satisfying these user requirements. These factors make the process of establishing network management standards very complicated indeed. We must continuously ask the questions illustrated in figure 1: "Do the pieces fit?" and "What pieces are missing?".

### 1.3. The Approach

The approach used during the study is illustrated in figure 2. User requirements were identified from several sources, including vendor implementations, user operated networks, and the MAP 3.0 network management specifications. User requirements tend to be fairly general and mission-oriented. To provide a basis for detailed comparison with existing and emerging standards, functional requirements are established from an analysis of the user requirements. Where user requirements are outside the areas appropriate for international standardization, functional requirements are not derived.

The functional requirements and retained user requirements are applied to evaluate the emerging and existing OSI management standards in an iterative manner. The results of the evaluation identify requirements that are satisfied and those that are not. In addition, major areas of concern are identified.

The document consists of five major sections, following the introduction. Section 3 documents network management requirements as gathered from users and

# DO THE PIECES FIT?

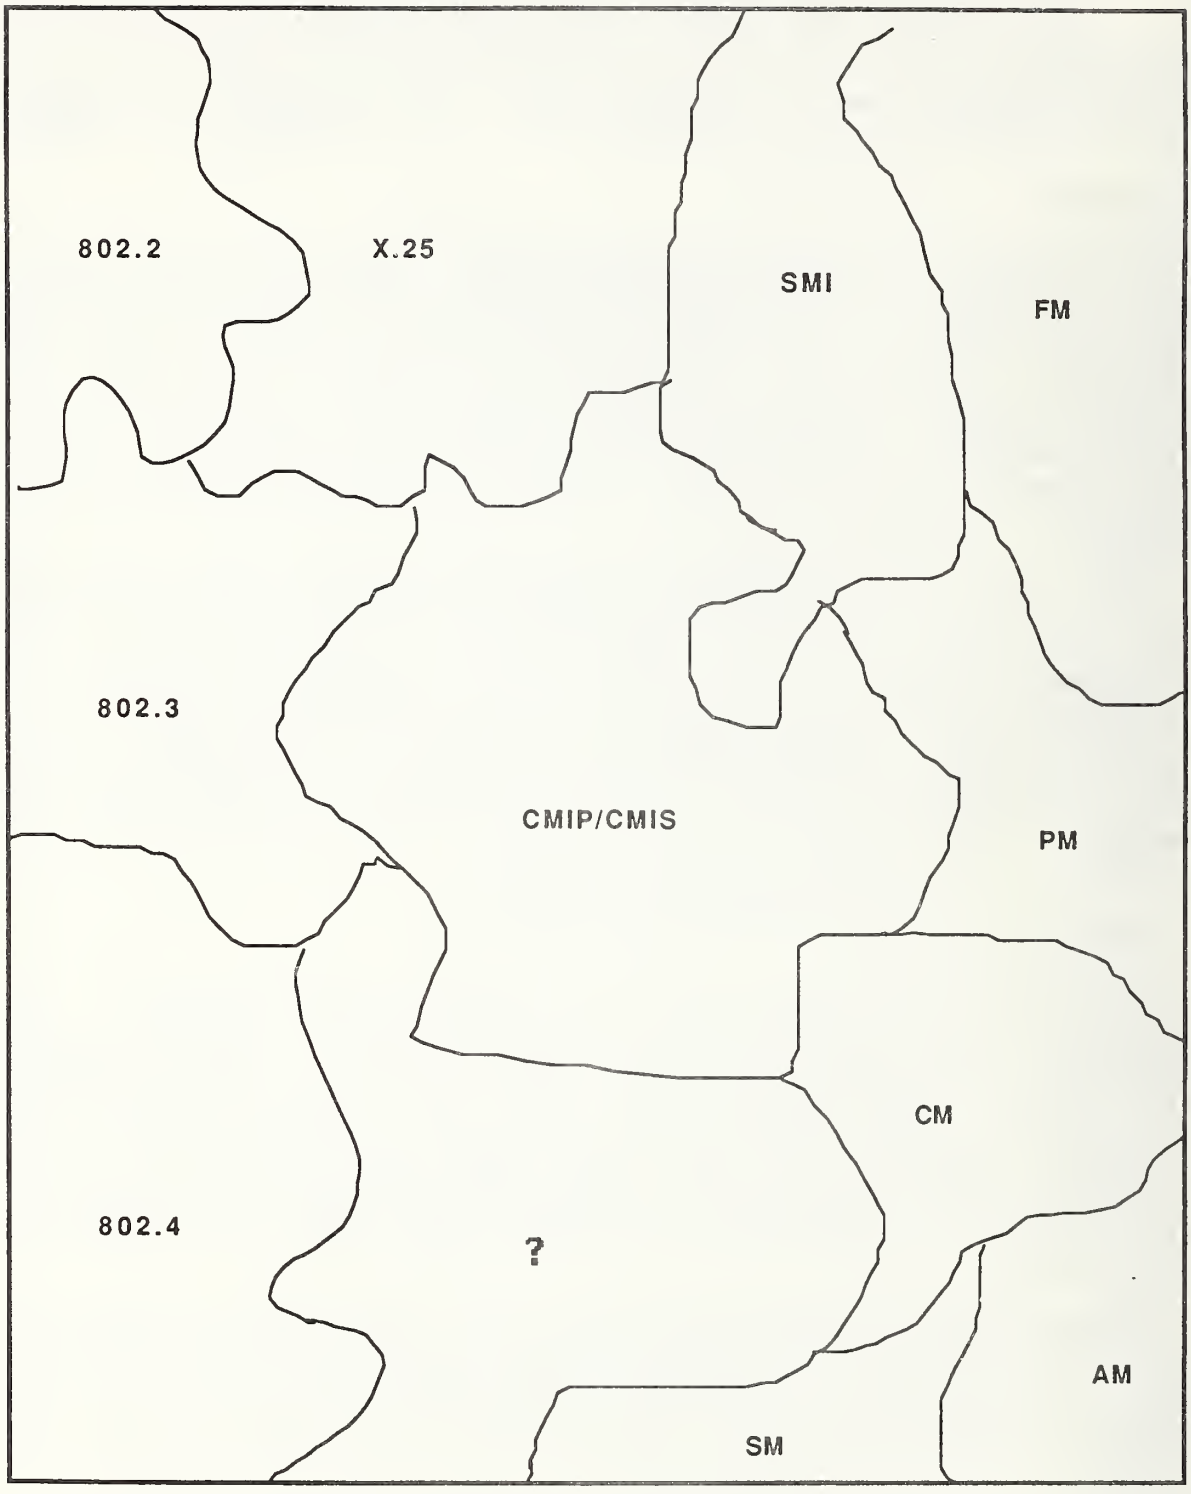


FIGURE 1

## OSI Management Standards Development Process

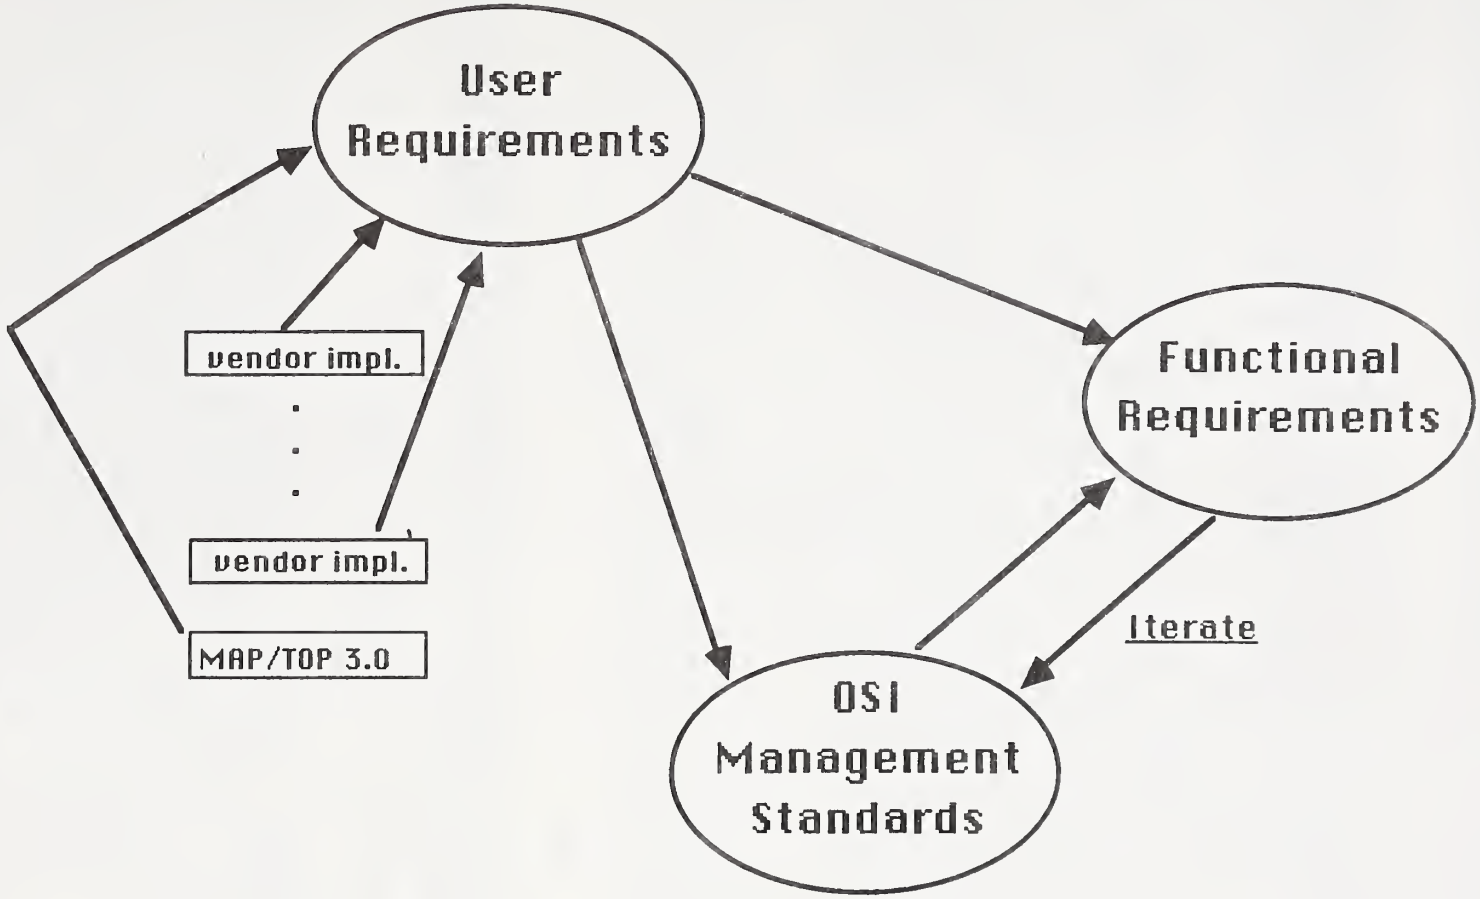


**FIGURE 2**

as results of the functional requirements analysis. This section is most important because the set of requirements can be used by any organization to independently evaluate the progress of network management standards or vendor network management products.

Section 4 presents the results of an evaluation of the extent to which existing and emerging network management standards are meeting user needs. During the course of the evaluation over a dozen major issues were uncovered, reflecting, possibly serious, deficiencies in the existing OSI systems management standards work.

Section 5 examines 20 areas of special concern or where additional research and standardization work would be useful. This section discusses issues that do not easily fit into the traditional functional categories of management. These issues and their implications are reviewed below.

Section 6 discusses the need, in the long run, to develop automated management systems that incorporate the principles of artificial intelligence to provide the services required by the operators and users of complex networks in the future. Concepts of expert systems are introduced and their usefulness for the purposes of management are examined. This discussion, of course, extends well beyond the scope of standardization activities.

Section 7 focuses on the requirements for a particularly interesting new concept, the Integrated Services Digital Network that is currently receiving so much attention by many users of network services. This section analyzes appropriate methods to identify the specific requirements needed to manage systems that include ISDN services and components.

## 1.4. Major Issues

A review of sections 4 and 5 reveals four significant sets of issues with respect to network management requirements and standardization. These sets of issues are summarized below.

### 1.4.1. Technological Assumptions and Efficiency

The first issue concerns the technological assumptions around which the network management standards are being built, and the result these assumptions will have upon the efficiency of implementations resulting from the standards. All the current definitions of OSI system management protocols reside at layer seven of the OSI reference model and use Common Management Information Service (CMIS) and Protocol (CMIP) to transfer management information between open systems. CMIS is connection-oriented, requiring the establishment of an Application association, a Session connection, a Transport connection, and, depending upon the underlying communications technology, Network and Link Layer connections. The network resources are modeled, from a management viewpoint, as managed objects with attributes. Management operations to acquire information and control resources are defined for these attributes. The information to be transferred is encoded in an abstract syntax notation (ASN.1) and decoded at the receiving end. Although not specifically covered

by the standards, the transfer of management information may well be in-band, i.e., the information may be transferred over the communications systems being managed.

These decisions are being taken despite the fact that many vendor solutions and user environments assume the existence of local area networks with datagram services, including broadcast and multicast. It is also common to have vendor management implementations sit atop a Transport Layer service and use a fixed binary encoding for data. Many vendor solutions use in-band signalling on local networks and move to out-of-band signalling upon entering a wide-area network environment, sometimes through separate virtual circuits and sometimes through physically separate resources. Connection-oriented versus connectionless issues are defined and discussed in section 5.4.

Whether the emerging network standards will permit users to achieve effective network management with efficient services is an open question that is addressed in section 5.3. What will be the overhead required by OSI management, reflected in extra network traffic and extra computing power dedicated to network management? What will be the response time to management situations permitted by a seven-layer connection-oriented implementation of network standards? What functions will be lost or made inefficient without broadcast and multicast services? What is the relationship between layer management and systems management? On what basis should the choice between the two be made -- should it be based on efficiency? The layer management issue is raised in section 4.1.

The need for flexibility often results in a sacrifice in efficiency. Thus, section 5.14 discusses the need for systems that not only are efficient and effective but also are extensible to meet future technological advances.

### 1.4.2. Multiple Manager Considerations

The second significant issue reported by this study is the lack of sufficient standards consideration being given to the need for multiple managers. Multiple managers exist in almost every real network management solution implemented today. Several reasons for requiring multiple managers are pointed out in section 5.5. First, large networks are often composed of a number of distinct subnetworks, each with its own management system. Thus, to manage a sizeable network and to allow for expansion of initial networks, most network management systems permit the existence of multiple managers, each responsible for some management domain. Often, the activities of the multiple managers are coordinated by higher levels of managers providing a network management hierarchy. The Systems Management Overview introduces necessary concepts such as management domains, but it offers little guidance on the establishment and coordination of the activities of the management systems among the domains. An important consideration is to determine which groups should develop the appropriate managed object definitions for domains.

Another scenario seen often today involves redundant managers within the same domain. When a network requires extreme reliability, as is often the case in commercial or military networks, the network management system must provide performance, configuration, security, and fault control to the level

expected. As a result, the reliability requirements for the network management system are usually higher than for the network as a whole. This often leads to redundant network managers each checking the other in a hot stand-by mode. Even for less stringent requirements a redundant network manager may wait in cold stand-by to detect the failure of the main manager so that network operations can continue after a minimal interruption. To date, there has been no definitive work on establishing standards that coordinate redundant managers.

A less often imagined requirement for multiple managers exists when the network management functions are partitioned such that individual managers are dedicated to specific areas, such as security, accounting, performance, or configuration control. This approach is normally used only in large networks or to meet specific needs such as those imposed by security requirements.

Despite the many examples of requirements for manager-to-manager operations, the OSI systems management standards activities are primarily aimed at only manager-to-agent operations (the agent is the system being managed). The scope of the standards work has recently been expanded to accommodate those manager-to-manager operations that can be modeled by manager-to-agent operations on managed objects that represent a view of another manager's information. Further work to refine and expand this concept is needed.

### 1.4.3. Management of Other Than OSI End Systems

The third significant issue identified during this study is the very real need to manage non-OSI end systems (OSI end systems support all seven layers of OSI protocols). Real networks consist of such mundane, but essential, devices as communications links, modems, routers, bridges, repeaters, and switches. None of these real devices are OSI end systems, yet they must all be managed, as discussed in sections 5.6, 5.11 and 5.17. Proprietary solutions for management of these devices already exist and it is therefore unlikely that proposed integrated management solutions will be acceptable unless they include the same functionality. Three examples of this issue are described below.

Within OSI end systems many implementation-dependent resources such as buffers exist. How these resources are used can have significant effects on performance, security, fault detection, and configuration control. These non-OSI resources must be managed, preferably in logical manner, consistent with the management of the OSI resources.

During a period of migration to OSI, or facing the prospect of maintaining existing systems mixed with OSI systems, network management gateways will probably be required. This will extend the need to manage resources that are equivalent to OSI resources, yet are not OSI resources. Thus, the emerging standards must provide for interoperable means of extending management through the specification of the appropriate managed objects. Further, such extensibility reinforces the need for supporting the issues raised in section 5.14.

The emphasis within the OSI systems standards bodies has been on management of OSI end systems. This is a most severe shortcoming because users will certainly need management of other than OSI end systems. As a simple example, the OSI standards permit the existence of a three-layer network routing device. Such a device must be managed. The emerging OSI network management standards have not fully addressed how to handle such devices. This uncertainty will mean increased cost to users as implementors struggle to find the best approach.

### 1.4.4. Lagging Standardization Process

In an earlier report [NBS87], the NIST described the organizations involved in making network management standards and predicted a timeline for the development of international standards for OSI systems management. The report predicted that the full range of management standards now planned would not be complete before 1992. Some think that this prediction is overly optimistic and see 1994 as a more realistic date. For example, the first OSI layer standards (for Transport and Network) to include OSI managed objects are not expected to reach DP status until 1990. There is currently no schedule for the management of upper layer services and protocols. Products that implement the standards can be expected 2 years after a standard is complete, thus the first full OSI systems management products can be expected sometime between 1994 and 1996.

This situation exists against a backdrop of OSI product availability for basic data communications beginning in 1986 and continuing through 1990. It appears that, at best, the OSI management products will be available 4 years later than needed. For example, significant work has only recently started in standards for performance and accounting management. The advancement of security management is tied to the pace of activities in subcommittees SC 6 and SC 21 of ISO/IEC JTC 1. A number of issues, as outlined above, have not even been addressed.

Further complicating progress is the organization of responsibilities established for the various standards activities. The establishment of a management framework as a revision to the basic OSI reference model placed responsibility for OSI systems management into ISO/IEC JTC 1/SC 21/WG 4. This responsibility includes the services and protocols required to support the framework, where the management services are divided into five distinct functional areas: configuration, fault, security, performance, and accounting management. The standards describing the systems management functions do not match the functional areas on a one-to-one basis. Standards being developed to meet the needs of one of the five areas will be available to satisfy the needs of the other areas as well. Further, the objects to be managed must be defined by the individual layer standards groups and this adds to the potential for confusion. While this organization of tasks permits the maximum parallel development of standards, resulting iterations due to lack of coordination or the time devoted to coordination will further slow the emergence of standards.

This situation creates a void that could continue to lock users into single vendor solutions even though OSI protocols are adopted by the users.

Lack of interoperable network management on a multi-vendor basis creates a problem for data communications users. Several approaches can be adopted by users. For example, General Motors has defined the minimal acceptable network management requirements for the Manufacturing Automation Protocol (MAP). This allows MAP users to specify OSI products that include interoperable network management in the areas of performance, configuration, and fault, primarily for the lower four layers. To date, however, few products that meet the MAP specification have been produced and the network management draft standards referenced in MAP have been superseded by newer versions.

Another possible solution is to operate nonproprietary, interoperable protocols without appreciable network management. This solution has been used throughout many organizations that have adopted the DoD protocols (e.g., TCP/IP). This is also the solution embodied within the first and second versions of GOSIP.

Another likely solution is to specify publicly documented network management protocols supported by a proprietary vendor within a framework of OSI protocols. For example, an organization could require implementations of OSI protocols to support a particular vendor's proprietary network management protocols. The practical effect of this solution is to limit the number of vendors meeting the requirement, maybe even to one. This approach seems unlikely to have the desired effect.

In the past 2 years, several organizations that develop implementation agreements that build upon the base standards have come into prominence. The most important ones are the Network Management SIG (NMSIG) of the NIST OSI Implementors' Workshop and the OSI/NM Forum, a corporation that is a consortium of over 60 vendors. Other groups have also initiated significant activities. These groups, including COS, Standards Promotion and Application Group (SPAG), Communications Network for Manufacturing Applications (CNMA), the IETF, and the Open Software Foundation (OSF), seek to add value to the emerging standards to define systems which meet the near term needs for network management. An important consideration is to devise a migration path to the long term solutions that the full set of standards will permit, i.e., a plan that considers upward compatibility.

Other areas that deserve more attention in the standards arena include software distribution (e.g., downline load) and initialization and shutdown of systems. While these areas do not need to be addressed for the first phase of standardization, they must be addressed before management standards can be considered to be complete.

In summary, even under the most optimistic views, the pace of network management standards development has created a 4 year void, causing a problem for users while permitting vendors additional time to expand market share and to lock users into proprietary network solutions. Such a situation demands cooperative actions by both users and vendors. The best hope lies in the work of the implementation agreements groups and, of these groups, the NMSIG exhibits the best characteristics of pursuit of OSI management principles in an open public forum.

## 1.5. Conclusions

This study covers an intricate subject cutting across protocols, layers, devices, and organizations. The focus, and thus, the conclusions relate specifically to user needs and to OSI standards supporting network management. The direction of the OSI standards process is fundamentally correct. A management framework is required and specific management functions are needed for the purposes of configuration, performance, fault, accounting, and security management. The current plan of work for OSI standards meets these needs.

Questions of whether the framework will meet every need in an efficient way exist. Is the seven layer stack the best means of transferring management information? Is it even an effective means given the inefficiency that may result? Can the standards ignore the broadcast and multicast capabilities of local area networks? Will the framework accommodate out-of-band signalling without sacrificing interoperability?

The standards process must evolve to accommodate multiple managers by providing definitions of managed objects suitable for manager-to-manager operations. In addition, the standards must also accommodate management of other than OSI end systems. Resolving these issues will add time to the development of OSI management standards, already too late to satisfy many users.

The major conclusion of this study is that OSI users must make decisions concerning the specification of network management services within the emerging functional profiles. The MAP NM specification has already embodied the conclusion of General Motors -- they cannot wait until 1994. The IETF, a standards making body for the U.S. DoD, has examined the progress of the OSI management standards and concluded that there is an interim need for non-OSI management protocols as well as the OSI management protocols. Meanwhile, the NMSIG is actively pursuing its goal of developing implementation agreements based on OSI standards. The NMSIG plans to publish its first set of stable implementation agreements as soon as a sufficient set of OSI standards reach technical stability. Present ISO schedules predict technical stability of this set by April 1990.

The clear choice for FIPS is to build on the NMSIG implementation agreements. This course is wise because a significant set of users and vendors are actively participating in their development and some set of vendors will implement them. Products will be developed. In addition, drafts of the NMSIG Phase One agreements have been provided as contributions to national and international standards organizations so that the emerging international standards will meet the needs identified by the NMSIG. The Phase One specification is limited to configuration and fault management and requires the use of non-OSI managed object definitions, since a sufficient set of OSI managed objects will not be available for several years. Further, it does not include any mechanisms to provide for security of the management system. The NIST recommends that this deficiency be corrected to meet the needs of Federal Government users.

Outside the realm of standardization many questions remain to be investigated. Should individual organizations develop man-machine interface standards and, if so, what makes an appropriate man-machine interface? Can artificial intelligence techniques be applied to network management? If so, how? Who defines and enforces the appropriate policies that guide the operation of the management systems? How can management gateways be implemented? Can management data be collected for use in off-line planning of network capacity maintenance? If so, what data is needed and should a standard format be used to permit development of competing software packages to analyze the data? How will the emerging technology of ISDN be managed? Can OSI systems management be easily applied to ISDN?

Network management has not been a science in the past. Network managers have practiced an art based upon rules-of-thumb and accumulated experience. The recent explosion of local area networks coupled to wide-area networks and populated with personal computers and professional workstations will require new artists who understand the use of much more sophisticated tools. Standardization within such an infant field is an ambitious undertaking, and yet it must be done. Static, slow changing networks such as the traditional voice telephone network do not accommodate the dynamic, fast changing requirements of data communications within modern business organizations; therefore, the well-tried telephone network management schemes are not likely to meet the variety of user requirements identified in this study. Network management standardization is and will be a stimulating and significant area of development of the systems of the future.

# 2. INTRODUCTION

## 2.1. Background

### 2.1.1. The Growth of Large, Complex Networks

Computer utilization strategies continue to undergo significant evolutionary change. Former strategies, promoting the use of increasingly larger centralized computer facilities to accommodate growing user needs, are giving way to current trends to decentralize and distribute functionality to many smaller computing units. This difference in approach is prompted, in large part, by the changing economics within the computer industry. It is less and less expensive to have more and more memory, storage capability and processing power housed in ever smaller packages. A natural consequence of this new strategy has been an ever-growing need to move data between these units through data communications networks.

The recent past has seen the proliferation of these data communications networks in all areas of endeavor (e.g., for electronic mail and data communications, and in offices and factories). These networks are designed both in accordance with standardized concepts (e.g., the International Standards Organization's (ISO) Open Systems Interconnection (OSI) reference model, and the Department of Defense's (DoD) Transmission Control Protocol/Internet Protocols (TCP/IP)), and in accordance with proprietary architectures (including, for example, IBM's System Network Architecture (SNA) and the Xerox Network System (XNS)). As these networks have proven their utility, they have been expanded, increased in complexity, and have been interconnected with other networks. Each subnetwork within these larger interconnected networks may contain large numbers of nodes that serve various functions such as terminal servers, file servers, print servers, process controllers and communications modems. Within these subnetworks, different physical transmission media (e.g., telephone lines, satellite links, broadband cable, and baseband cable) as well as different media access strategies (e.g., token-passing bus, Carrier Sense Multiple Access with Collision Detection (CSMA/CD), and token-passing ring), can be used. Moreover, various communications protocols, singly or in combination, can be employed on any one of these subnetworks. These individual subnetworks can be interconnected through bridges and gateways which handle routing and protocol mapping to provide the potential for worldwide accessibility.

### 2.1.2. The Need for Network Management

It is increasingly apparent that management (primarily monitoring and control) of these networks and network systems is absolutely essential. Systems of such complexity cannot be maintained without some oversight capability to assure continued reliable operation at an expected level of service. For example, when components malfunction or are withdrawn from service for replacement, repair, or periodic maintenance, or when nodes are added to or removed from networks, either temporarily or permanently, these networks need to be appropriately reconfigured. Performance degradation, also, must be diagnosed and remedied if users are to be provided a reliable

quality of service. And, value-added network service providers need accounting capabilities to support billing for services.

### 2.1.3. Standardization: The Path to Interoperability

Networks, both by design and by necessity, are used and will continue to be used to interconnect diverse systems, devices, and subnetworks providing a wide range of functionality. With such heterogeneity, it is highly unlikely that a network's components will be restricted to any one vendor's products.

As with other areas of standards development, such interoperability cannot be achieved without commonly accepted standards. The ISO OSI reference model includes provision for network management capability, but aggressive pursuit of management standards has been lacking until recently. Now, however, the standards community, vendors, and users are exhibiting heightened interest in network management. Network management tools and capabilities are finally being recognized as essential network operating components and, therefore, are achieving commercial acceptance to the extent of becoming advertising and selling criteria promoted by vendors. Users, as a result, are becoming increasingly active in interest groups that represent their needs and concerns with regard to the emerging standards and products. Users, vendors, and the standards community are united in pursuit of this common goal of establishing useful and cohesive standards which are sufficiently comprehensive to allow (and perhaps even point the way for) vendors to implement network management (NM) suitable for OSI systems.

The most fundamental concepts of OSI management are described in the Management Framework document (MF) [FRMWK] to the OSI Reference Model [OSIREF]. The MF describes, in the form of a general conceptual overview, the set of problems to be addressed by OSI Management. It partitions the functions of OSI management into five distinct categories - Configuration Management (CM) [CONFIG], Fault Management (FM) [FMWD], Security Management (SM) [SECURE], Performance Management (PM) [PMWD], and Accounting Management (AM) [AMWD]. These five categories are referred to as OSI Systems Management Functional Areas. Other important documents are being advanced which detail, with greater specificity, the sets of functions needed to accomplish OSI management. (See the NIST project's phase one document [NBS87] for specific insight into the nature of, and relationship among, these documents. However, the reader is cautioned that some portions of NBS87 are no longer correct because of recent restructuring of standards and documents by ISO.)

## 2.2. Approach

Progression of the work toward appropriate network management solutions for open systems depends upon a proper and full understanding of the problem. In order to aid in this process, the National Institute of Standards and Technology (NIST), has undertaken a three-phase project: 1) to study the emerging NM standards; 2) to formulate the set of requirements for NM and compare these with the facilities and solutions offered in the emerging management standards; and 3) to establish a laboratory for the investigation of NM issues with the first laboratory task being an implementation of the

common management information protocol (CMIP) to demonstrate the feasibility of OSI network management.

Phase one of this NIST project surveyed network management (NM) standards activity and the emerging network management standards [NBS87]. The efforts of ISO and other standards making bodies, such as ANSI and the IEEE, were reviewed in their roles to develop OSI management standards. These efforts, having been underway for many years, have now (Summer of 1989) progressed beyond the point of broad concepts to the phase of detailing the mechanisms and information structures needed to implement interoperable management for OSI systems.

### 2.2.1. Identification of Issues and Requirements

The phase two study, of which this paper is a major element, has as its specific goal the formulation of the set of functional requirements for the management of OSI-based networks. OSI-based networks are formed by the interconnection of computing systems through the use of OSI services and protocols (i.e., the well-known seven-layer stack). The management of these services and protocols of any computing system on an OSI-based network is to be accomplished through the use of OSI management standards now being developed primarily by Working Group 4 (WG4) of ISO/IEC JTC 1/SC 21. These emerging standards are directly concerned with the management of the communications aspects of OSI systems and not directly concerned with so-called "network management."

Thus, a distinction must be made between "network management" and "management of OSI-based networks." Network management, as commonly used in the telecommunications industry, is concerned with the management of the elements that interconnect computing systems. These elements include, for example, switches, multiplexors, modems, and circuits. By contrast, OSI management standards, as currently being developed, do not specifically address "network management," but are primarily concerned with the communications aspects of full seven-layer OSI systems. Therefore, it is important to note that OSI management is not the same as network management. OSI management is not oriented toward the management of circuits, a concern of traditional network management. Furthermore, OSI-based networks primarily, although not exclusively, employ packet switching.

While the above distinction must be made, nevertheless, it is believed that OSI management can be applied to the management of telecommunications networks beyond the current focus of OSI management standardization. For example, within the United States there is currently an effort by Accredited Standards Committee (ASC) T1M1 to apply OSI management to telephony elements. The extension of OSI management to more general network management and suggestions for accomplishing this extension is a concern of this report. Such extension efforts should lead to efficient, consistent, integrated network management in the future, although, as discussed in this report, there are technical problems to be addressed.

Now that the above distinctions have been illuminated, it is useful to employ a term that encompasses both the concepts of OSI management and

traditional network management. Therefore, for the purpose of this report (unless otherwise noted), the term "network management" will refer to the concepts of management of OSI-based computing systems (i.e., "OSI management" as being developed by ISO) as well as the more traditional management of interconnecting communication elements such as circuits, switches, and multiplexors.

A further distinction is made in this paper between user requirements and functional requirements. User requirements specify the services that the manager of a network needs and/or expects in order to do his job. Functional requirements are specifications of the elements needed to provide these end services. In this sense, functional requirements are derived from user requirements and are to be used in design and implementation specifications. Some user requirements may be specified in sufficiently elemental detail that they serve not only as user requirements, but also as functional requirements. Most user requirements, however, are of such a broad general nature that they require subsequent analytical breakdown to discover the functional elements needed to achieve their ends.

In an attempt to assure that as many user requirements as possible are met by the emerging standards and that no reasonable ones are precluded, this study examines implementations of current vendor-supplied proprietary network management systems, contributions from various user groups involved with network management, and the most current technical literature on the subject. The requirements which initially motivated the development of a standard solution must also be revisited to determine, first, if standard mechanisms have been devised which adequately address these requirements and, second, if there are additional requirements which must be considered in light of increased experience, knowledge, and understanding of the problems and issues of network management.

The standards address the issues and requirements that relate to interoperability. There are also many issues and requirements related to networks and network management where interoperability is not involved. For example, network management includes many distinct tasks, such as planning, installation, and operation. Some of these tasks may be accomplished in a manual or off-line manner. Network planning may be manual or facilitated by automated methods. Installation is primarily a manual task. In addition, some operational functional requirements have little or no effect on interoperability. For example, it is generally agreed that it is a network management requirement that the Man-Machine Interface (see fig. 3) be consistent and easy to use. While a consistent interface is certainly desirable, the management system supplied by one vendor may be totally, functionally interoperable with another vendor's system to be managed, and yet the interfaces to the human operator may be proprietary.

### 2.2.2. Analysis of the Scope and Completeness of the Standards

The only standards required for interoperability are those concerned with the information in the electronic messages passed between the systems from diverse vendors, and since the purpose of this survey is to detect problems in the set of emerging standards for network management, it is those functional

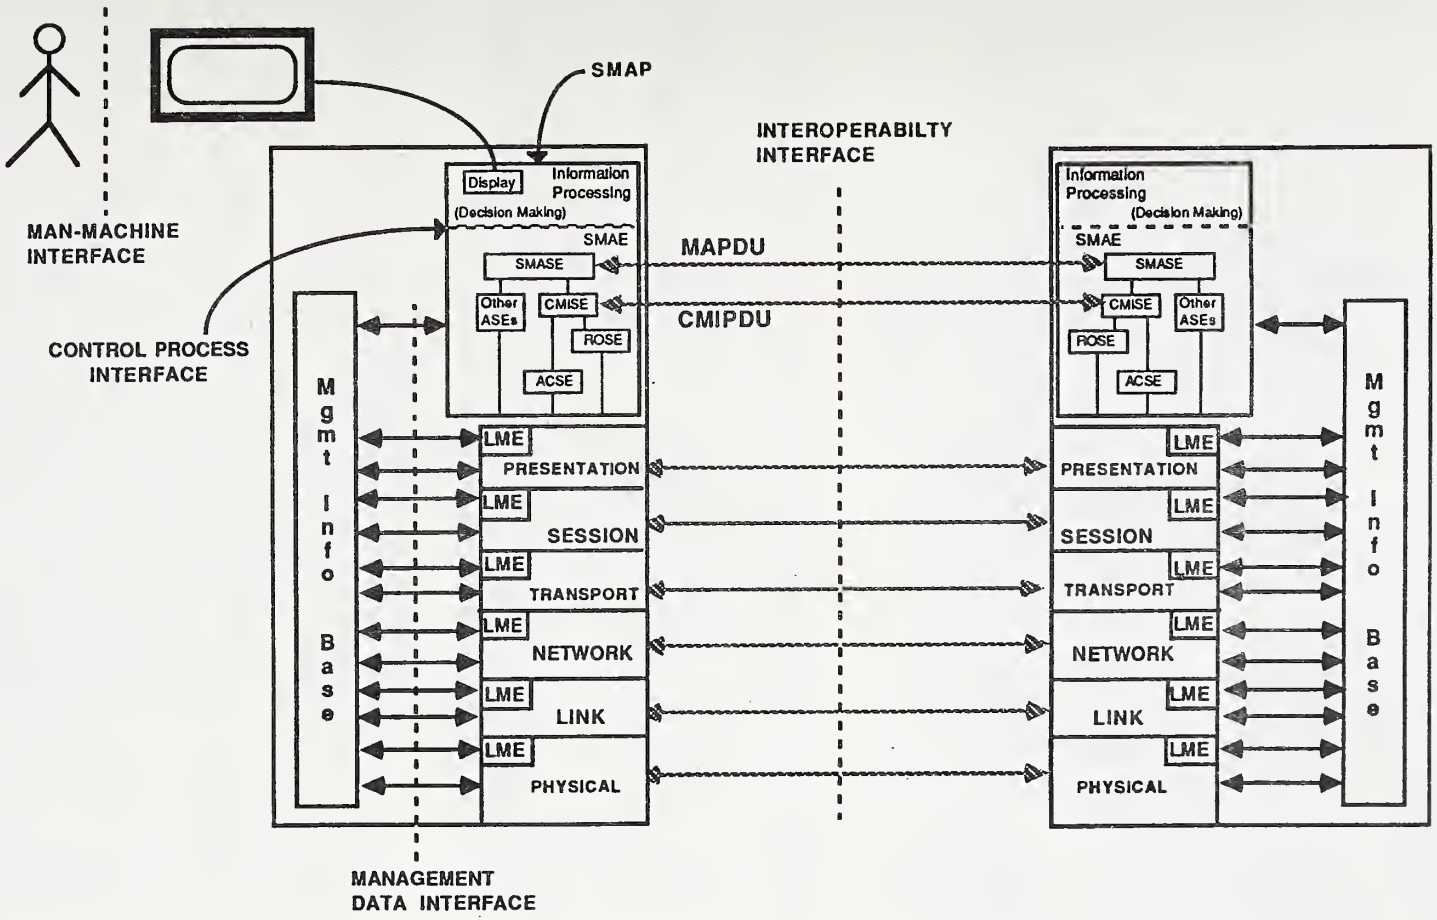


# SIMPLIFIED OSI MANAGEMENT MODEL
# FIGURE 3

requirements for OSI network management related to and required for network operation which are of primary concern within this survey. Thus, while it is recognized that there are other important functional requirements, this survey focuses on those requirements that directly affect interoperability and the emerging set of OSI management standards.

Developing this set of requirements is important for, at least, two reasons. First, such a set of requirements will provide guidance for NM implementors and users by outlining the full range of NM functionality to be considered by implementors and users of NM systems. Second, it will provide the NM standards bodies with a set of functional requirements so that user needs can be fully incorporated in the emerging standards. By comparing requirements with the facilities and solutions offered in the emerging management standards, this study addresses the question whether the emerging standards contain errors, deficiencies or omissions which may, at some point, require the standards to be retrofitted or revised. The information derived from this study will be distributed to the appropriate standards committees to promote the development of interoperable, comprehensive network management products capable of meeting real management needs. (See the Appendix for a detailed outline of phase two of this NIST project.)

## 2.3. Organization of the Paper

The following presents an organizational overview of the major sections of this paper, briefly describing the elements discussed in each section.

Section 1 -- The Executive Summary gives an overview of this study, discusses policy issues, highlights major conclusions, and suggests future directions.

Section 2 -- The Introduction describes the background and motivation for the study, presents the goals of the study, and describes activities undertaken to achieve a comprehensive analysis of network management requirements.

Section 3 -- The Network Management Functional Requirements section presents collated user requirements for network management based upon analysis of actual networks and relevant information from other sources. The user requirements are organized for presentation according to categories of functionality initially based upon ISO functional categories (e.g., CM, PM), but extended to include additional functionality where appropriate.

Section 4 -- The Analysis section offers a functional comparison among standards, implementations, and requirements gleaned from other sources. This comparison relates management functions implemented in actual systems, or specified as necessary by users, with the functions provided for in the emerging standards.

Section 5 -- The section on Additional Issues suggests areas of future interest in network research and standardization, discusses issues related to determining the appropriate scope of Management for OSI-based networks, and discusses issues that transcend the traditional ISO functional categories of CM, PM, FM, SM, and AM.

Section 6 -- The section on Automated Network Management Systems discusses the need to apply "expert systems" technology (a branch of Artificial Intelligence) to network management problems. Concepts of expert systems are introduced. This section provides insights into how automated management may improve operations of networking, in general, and network control centers, in particular. The reader is cautioned that automated management systems are beyond the scope of standardization.

Section 7 -- The Integrated Services Digital Network (ISDN) section focuses on concerns in identifying ISDN NM requirements. ISDN NM standards and the NM activities of the North American ISDN Users' Forum (NIU) are discussed. Mechanisms to identify ISDN NM requirements are suggested.

## 3. NETWORK MANAGEMENT FUNCTIONAL REQUIREMENTS

A primary goal of network management (NM) users is to have a network management system (NMS) that allows the control and monitoring of a network composed of products from different vendors. An important prerequisite for such a management system is the development of interoperable NM products. However, interoperable NM products cannot be assured unless they are produced in accordance with an agreed upon set of standards.

Many NM standardization activities are currently underway to develop such a set of standards. As with any large undertaking of such complexity and broad scope, the authors of this study feel it is important to reexamine NM functional requirements, at this stage in the development of these standards, to insure that the emerging standards, in fact, meet user requirements.

This study has gathered and examined NM user requirements from a variety of sources. Based on an analysis of user requirements, we have prepared a set of functional requirements designating the functional elements which must be incorporated in NM in order to meet users' needs. The primary sources for the user requirements include IFIP WG 6.6 user requirements [IFIPUSR], WAN and LAN vendor NM products, MAP 3.0 Network Management requirements, private communication with a network services provider, and comments received from an earlier draft of this study.

NM user requirements may be viewed from two different perspectives; that of the end user of a network (i.e., not a network administrator), and that of the operator, administrator or manager of a network. End users expect reliable network services with consistently good performance. They may want to retrieve accounting and performance information and be notified of any configuration changes that affect them. The network operator, administrator, and manager require sophisticated network management tools that are both necessary and sufficient to perform all types of functions to support the network services expected by end users. In order to prepare a set of NM functional requirements that encompass both network end user and network manager's needs, requirements from both communities are considered. The term, "NMS users," used in this section indicate both the end users and the network administrators including operators and managers. For functions that are usually only required by the network administrators, the term "network administrator, NM manager or operator" is used.

The functional level requirements are grouped into seven categories that are similar to the functional areas defined within the ISO Management Framework [FRMWK]. The categories are architecture, configuration, fault, security, performance, accounting, and others. We have chosen to organize the functional requirements in this manner for ease of comparison with OSI NM standards, and subsequent analysis. Each functional requirement section follows a user requirements section that describes and justifies its needs. Some user requirements may lead to more than one functional requirement and, conversely, one functional requirement may satisfy all or part of more than one user requirement. Therefore, the correspondence between user requirements and functional requirements is not necessarily one-to-one.

## 3.1. Architecture

Due to the complexity and broad scope of areas and functions that are involved with a network management system (NMS), it is a logical approach to begin with a model for the management system when developing standards. The model should describe the architecture (structure) of the system, identify the purpose of each component within the system, and identify the relationships among all the components within the system so that it can provide a formal basis for developing the elements of the network management system in a systematic manner.

The NMS provides its users with the capability to control and monitor the communications in a system interconnection (networking) environment. Since the communications aspect of these interconnected systems already has a standard model (i.e., the OSI reference model), it is only natural that a system to manage implementations of this communications model employ an architecture that is compatible with the standard communications model. In this sense, the NMS model needs to be designed around the OSI reference model.

Most networks of any considerable size have a network control center where the network management functions are coordinated and controlled. The staff of a network control center typically includes: 1) technicians who are charged with resolving problems as they occur, 2) operators (administrators) who are responsible for management of the physical resources of the network, 3) consultants who run the user help desk that serves as the single point of contact for end users with problems or questions, and 4) the managers who are responsible for monitoring and controlling the operational network.

The network control center usually oversees configuration, fault, and security management functions, while most accounting management functions are performed in the billing/accounting department and most performance management functions are handled by the performance analysis and/or capacity planning organization. In some organizations the security management functions are administered by security group within the organization. Organizations such as network design/architecture, network engineering and development, software control and strategic planning, use diverse management information and functions for performing some of their jobs. Network management information is therefore widely needed and used. An effective network management system is one that enhances the productivity, effectiveness, and responsiveness of the network control center staff and, ultimately, the end users of the communications system.

The architectural requirements specify network management functions that are essential building elements for a NMS from an architectural point of view. They include the model, the directory services which are required by all the NMS components, the management protocols and services for transferring the network management information, layer management, and specific resources to be managed and for controlling them.

### 3.1.1. User Requirements

Most users of a network management system require a single integrated system that allows them to remotely monitor and control the operation of the network. To provide this integrated network management system, a network management model has to be in place to give an overview of the system, to indicate how the management information can be transferred and collected, and to identify the functions that need to be standardized for providing an interoperable and integrated network management system.

Many users of a NMS require network-wide distributed NM control for large networks. Such a NMS may contain many, possibly loosely coupled, network managers rather than a single centralized network manager. However, these distributed managers may themselves be centrally managed. This concept of distributed managers implies the need for manager-to-manager protocols and services. According to the IFIP user requirements document [IFIPUSR], the user requirements that support the distributed NM control are the following:

a) Distributed NM control better reflects the structure of a large network, with each network manager supporting some segment of a large network. However, it may be more logical to manage some segments centrally.

b) In instances where there is a strong geographic locality of reference, network management information can be put logically closer to the NM user, reducing communication time and cost.

c) Network-wide management dispersion reduces the likelihood that a disaster will affect more than one part of the network.

d) When one NM control node goes down for maintenance or updates, the NM functions within its domain can be redirected to another node.

e) Each network manager in a domain tailors its user interface to its own needs (e.g., German, English, Japanese). The only point common to all the network managers is that they all support the same standard when accessing each others' network management information.

A NMS needs to be capable of performing the required network management functions on any component of a multi-vendor network. Therefore, the network management system should support collection and distribution of management information from heterogeneous components.

NMS users need the flexibility to select any cost-justified technological alternative that best meets their business requirements. NMS users should not be unduly restricted (i.e., "locked-in") by the existing NMS.

NMS users expect the network management system to maximize the availability of network component resources, whether those resources are communications links, processing nodes, software or data. The network

management system should be robust enough to continue network operation in spite of most common hardware and/or software failure.

### 3.1.2. Functional Requirements

#### 3.1.2.1. Model

The model of a network management system (NMS) describes the essential NMS components and the functions of each component. The model needs to clearly define the organization of the network management system including the relationships, interactions and interdependencies among its components. In addition, the model should allow for:

a) hierarchical and distributed control of access to and manipulation of network management information and network resources;

b) centralized management of distributed network managers;

c) the flexibility to accommodate and support new technologies so that there is no need to build a new NMS in order to manage a network with new technology;

d) additional proprietary network management solutions;

e) additional proprietary network management security control;

f) future network management system expansion; For example, as new types of managed objects are identified, the NMS developed based on network management standards should be able to perform actions on these new objects with minimal detrimental effects by using the standard methods to specify and add new managed objects to the data bases or directories.

g) redundant managers or out-of-band signaling to ensure maximum availability of the network management system;

h) message distribution in a hierarchical distributed network management system. This will be required for broadcasting status, caution, warning, bulletins and management directives (e.g., start all performance functions, stop all activity to prepare for network shutdown).

#### 3.1.2.2. Services and Protocols

The network management system requires manager-to-agent management information exchange services and a supporting protocol. These services and protocol are necessary (but not sufficient) for any network management system.

Manager-to-manager management information exchange services and protocol are required to support multiple managers within a network management system. Manager-to-agent protocols may be sufficient to support peer manager-to-

manager operations. However, various types of manager-to-manager protocols may be required when hierarchical management structure is considered.

The services provided by these protocols must be sufficient to perform all the functions specified in the following five sections (i.e., configuration, fault, security, performance and accounting management.). The form of management information exchange must be able to support all the required services across various levels of hierarchical management and control of distributed management subsystems.

#### 3.1.2.3. Resource Identification

The OSI and related non-OSI resources to be managed must be identified. Each resource can be viewed as a managed object with attributes. The definition includes the managed object, its attributes, the set of operations that can be performed on the attributes, and the semantics of these operations. These resources consist of layer, system, and network component level entities. A buffer is an example of a non-OSI resource while a Transport connection or a Transport Layer retransmission count are examples of OSI resources. The resources associated with manager-to-manager functions also need to be identified. These include all types of data base logs such as security events, failure events, and configuration changes.

#### 3.1.2.4. Information Structure

The network management system needs a common structure for heterogeneous management information to facilitate exchanging information across various vendors' products. In other words, the difference in data formats, source, structure and semantics should be transparent to the NMS users. A standard method to describe the common structure for heterogeneous management information is therefore required.

#### 3.1.2.5. Layer Management

Those resources identified as necessary and important for management need to be collected, stored and made available to remote management systems by local management systems. Each local layer entity must have sufficient functionality to support the local management system.

#### 3.1.2.6. The Directory

To provide quick access to the desired management information, directories are needed to identify applications, users and resources in a network. Directories also should contain or at least be capable of determining the routing information for interconnected networks, perhaps by pointing to the appropriate OSI directory services. The Directory standard should specify:

a) where directory information is stored,
b) how information in the directory is created and updated,
c) how access to the information is controlled,
d) the structure of the information stored in the directory, and

e) how the information is used by local and remote management systems.

The directories require accurate and organized information from the network components. Before building useful directories, the four functional requirements described in sections 3.1.2.2 - 3.1.2.5 (i.e., the management protocols and services, the layer management, the common management information structure, and the identification and definition of network resources) must be in place.

### 3.1.2.7. Network Management Communications Overhead and Performance

Most users expect the NMS to impose minimal overhead on network operations, and dictate that the NMS functions should not interfere with routine, ongoing workloads. Many users require that the NMS perform to a user specified level and out-of-band signaling that transfers management commands and information to remote devices without affecting normal data channels is one way to meet these requirements.

### 3.1.2.8. Support for Efficient Information Transfer

Network management traffic can often be categorized into a large number of small exchanges or a few high volume exchanges. Information exchange may have other diverse data volume and data transfer frequencies. For example, software distribution and "up-line" retrieval of statistics blocks are examples of high volume transfers which require file transfer capability. Transaction processing, order entry, and DBMS updates may require low volume, reliable transaction services. Standard ways of transferring various amounts of management information from one node to another need to be provided by the network management system. The information transfer can be done either directly through a management protocol or through the use of other services such as FTAM or through transaction processing protocols.

### 3.1.2.9. Standardization of Terminology

To provide NMS users with an integrated view of network management capability, terminologies used across multi-vendor NM products should be defined in the standards. Examples of such terms are the definition of faults, the definition of security terms, the definition of performance measures, the definition of terms used in the model and the services and protocols, and the definition of configuration states and the relationships among configurable network components.

## 3.2. Configuration Management

Modern data communication systems are composed of individual components and logical subsystems (e.g., queue managers in an operating system) that can be configured to perform many different applications. The same device, for example, can be configured to act either as a gateway or as an end system node, or both. Once the manager decides how he intends to use the device, he

can choose to establish values for the appropriate set of attributes associated with this device. The device can be considered a system resource, or managed object where "A managed object is the OSI Management view of a system resource that is subject to management, such as a layer entity, a connection or an item of physical communications equipment." [SMO] Furthermore, "Attributes are properties of managed objects. An attribute has an associated value, which may have a simple or complex structure." [SMO] The value of an attribute "... may determine or reflect the behaviour of the managed object." [MIM]

Configuration management (CM) is that aspect of network management which embodies the functionality to, among other things, assign that set of attributes to the device. Configuration management is concerned with initializing a network and gracefully shutting down part or all of the network. It is also concerned with maintaining, adding, and updating: 1) the relationships among components, 2) the status of the relationships among the components and 3) the status of the components themselves during network operation. By its nature, CM interacts with other aspects of network management, to a greater degree than other functional areas, to provide important monitoring and reconfiguration services.

During initialization, configuration management identifies and specifies the characteristics of the network components and resources (managed objects) which will constitute the network. The managed objects include both high level composite objects (e.g., an end system or gateway), and lower level atomic objects (e.g., a Transport Layer retransmission timer). The configuration manager provides the capability to set attribute values individually or collectively to predefined default values. This process causes these managed objects to commence operation in the proper states, possess the proper attribute values, and form the desired relationships with other network components.

While the network is in operation, configuration management functions monitor the network components and may reconfigure managed objects when desired or necessary. In this regard, configuration management functions may be allied with the functions of other management areas and used to support their operations. For example, if the performance management developed Workload Monitoring Function determines that network performance is degrading (e.g., increased response times are causing excess retransmissions), or the fault management developed Error Reporting and Information Retrieval Function detects a malfunctioning component, the services of a configuration management function (e.g., State Management Function) can be enlisted to modify the appropriate managed objects to remedy these situations. The actions taken in these cases might include increasing the appropriate retransmission timeout periods and reconfiguring the network to work around the malfunctioning component until it is repaired.

The OSI Management Framework document (ISO 7498-4), defines CM as follows:

> "Configuration identifies, exercises control over, collects data from and provides data to open systems for the purpose of preparing for,

initialising, starting, providing for the continuous operation of, and terminating interconnection services. Configuration management includes functions to:

a) set the parameters that control the routine operation of the open system;
b) associate names with managed objects and sets of managed objects;
c) initialise and close down managed objects;
d) collect information on demand about the current condition of the open system;
e) obtain announcements of significant changes in the condition of the open system;
f) change the configuration of the open system." [FRMWK]

The following two sections present a more in-depth view of Configuration Management, presenting first the users' view of this aspect of network management followed by the functional elements which are needed to provide such capabilities.

### 3.2.1. User Requirements

Startup and shutdown operations on a network are the specific responsibilities of configuration management. It is often desirable, and even necessary, for these operations on certain components to be performed unattended on distributed systems (e.g., starting or shutting down a remote line multiplexor).

Network operators or administrators need the capability to identify the components that comprise the network and to define the desired connectivity of these components. Users who regularly configure a network with the same or similar set of resource attribute values need ways to define and modify default attributes and to load these predefined sets of attribute values into the specified network components. This avoids specifying the same resource attributes and values every time.

Network managers or operators need the capability to change the connectivity of network components when users request such changes or when reconfiguration is mandated by performance, fault, or security requirements. Reconfiguration of a network is often desired in response to performance evaluation or in support of network upgrade, fault recovery, or security checks.

Network users often need to, or want to, be informed of the status of network resources and components. Therefore, when changes in configuration occur, users should be notified of these changes. Configuration reports can be generated either on some routine periodic basis or in response to a request for such a report. Before reconfiguration, the operator or the manager often wants to inquire about the status of resources and their attributes.

Company administrators and network managers usually want only authorized users and operators to manage and control network operations such as software distribution and updating.

Network capacity planning, network performance and usage trend analysis, and the management of the inventory of information system components (including software, hardware and microcode) are also network managers' NM requirements. These requirements, however, are beyond the scope of this study because they are not needed for interoperability. (See sec. 5.7 for further discussion.)

### 3.2.2. Functional Requirements

#### 3.2.2.1. Defining Resources and Attributes

Mechanisms are needed to allow the NMS users to specify resources and the attributes associated with a resource. Attributes can be, for example, name, address, identification number, states, operational characteristics, software version number, and release level. Network resources include network physical resources (e.g., modems, the communications media, or computers), and network logical resources (e.g., timers, counters, virtual circuits, and connections).

The NMS users should be allowed to specify the range and type of values to which the specified resource attribute can be set. The range can be a list of all possible states, or the allowed upper and lower limits for parameters and attributes. The type of value allowable for an attribute can also be specified.

#### 3.2.2.2. Setting and Modifying Attribute Values

Mechanisms are needed to allow the NMS users to set and modify values of resource attributes (e.g., activate and deactivate ports, set and trace a retransmission timer value, and monitor and adjust buffer allocation).

The NMS users require mechanisms to load predefined default attribute values such as default states, values and operational characteristics of resources on a system-wide, individual node, or individual layer basis.

The NMS must allow users to set clocks for network components.

#### 3.2.2.3. Defining and Modifying Relationships

The NMS users must have the ability to specify relationships among network resources. A relationship usually describes an association, connection or condition that exists between network resources or network components. These relationships can take the form of a topology, a hierarchy, a physical or logical connection or a management domain. What is meant here by a management domain is a set of resources that share a set of common attributes or a set of common resources that share the same management authority.

Mechanisms are needed to allow the NMS users to add, delete, and modify the relationships among network resources. The NMS must also allow its users to expand the network or change existing relationships among resources without

taking all or part of the network down (i.e., the relationships may be modified on-line during network operation).

#### 3.2.2.4. Examining Attribute Values and Relationships

Mechanisms are needed to allow the NMS users to examine resources. This requires the ability to locally or remotely examine the attributes associated with the resources and the current values of these attributes.

Mechanisms are needed to allow the NMS users to locally or remotely examine the existing relationships among network resources.

To provide the above two capabilities, the NMS must be able to keep track of configuration changes from which the existing network resources and attributes, their status, and relationships can be determined.

#### 3.2.2.5. Distributing Software Throughout the Network

The ability to distribute software throughout the network is essential. This requires facilities to permit software loading requests, to transmit the specified versions of software, to notify the NMS user at the completion of software loading, and to update the configuration tracking systems.

The NMS user needs mechanisms (e.g., downline loading capability) to examine, update and manage different versions of software and routing information. For example, users can specify the loading of different versions of software or routing tables based on a specified condition, such as error rate.

#### 3.2.2.6. Initializing and Terminating Network Operations

The NMS must provide mechanisms to allow its users to initialize and close down network, or subnetwork, operation. Initialization involves, among other things, verifying that all settable resource attributes and relationships have been properly set, notifying users of any resource attribute or relationship still needing to be set, and validating the users' initialization command. For termination, mechanisms are needed to allow the NMS users to request retrieval of specified statistics blocks or status information before the termination procedures have completed.

Mechanisms are needed to allow the NMS users to remotely reinitialize (reboot) a system.

#### 3.2.2.7. Verifying NMS Users' Authorization

The most privileged NMS users have the ability to specify: 1) the hierarchy of authorization for performing various configuration functions and 2) the methods used for assigning and validating various levels of authorization.

Mechanisms are required to allow only authorized NMS users to perform various configuration functions. This is related to security issues, but it

is required by configuration management to ensure that only authorized personnel can gain access to or change network configuration information as well as start or stop a network's operation.

#### 3.2.2.8. Reporting Configuration Status

Notification of configuration changes in resources and in relationships among resources must be available to the NMS users. In order to accomplish this, managing systems (i.e., those that manage other systems) must be able to inform agent systems (i.e., those that are managed) under what conditions, and where, configuration change notification is to be sent by the agent system.

Mechanisms are needed to allow the NMS users to request and obtain configuration reports. The configuration reports focus on such things as network connectivity, network topology and node resources, attributes, and values. Furthermore, these configuration reports may display routine snapshots of the network configuration and the status of the components (e.g., a NMS can display a snapshot of network topology every 5 minutes alternating with performance snapshots).

The NMS users must have the ability to broadcast or multicast configuration news (information about network configuration) to other network managers. Such news can include notification of when certain components or facilities will become available or unavailable.

### 3.3. Fault Management

To maintain proper operation of a complex system such as a computer network, care must be taken that the system as a whole, and each essential component, individually, is in proper working order. Accomplishing this requires the ability to take corrective action and make repairs when needed. Where down-time cannot be tolerated, it is essential to anticipate problems so that preventive maintenance procedures can be invoked to avoid the actual occurrence of problems or faults.

Fault Management (FM) is that aspect of Network Management (NM) which attends to these concerns. It represents a logical division of labor within network management activities. FM seeks to maintain system operation as close as possible to a fault free level. The OSI Management Framework document [FRMWK] defines FM as follows:

> "Fault management encompasses fault detection, isolation and correction of abnormal operation of the OSI Environment. Faults cause open systems to fail to meet their operational objectives and they may be persistent or transient. Faults manifest themselves as particular events (e.g., errors) in the operation of an open system. Error event detection provides a capability to recognize faults. Fault management includes functions to:
>
> a) maintain and examine error logs;
> b) accept and act upon error detection notifications;

c) trace and identify faults;
d) carry out sequences of diagnostic tests;
e) correct faults."

Central to the definition of fault management is the fundamental concept of a "fault." Faults are to be distinguished from errors. A fault is an abnormal condition which requires management attention (or action) to repair. A fault is usually indicated by failure to operate correctly, or by excessive errors. Certain errors, (e.g., CRC errors on communication lines), may occur occasionally and are not normally considered to be faults. Ordinarily, these errors are handled by an (N)-layer entity as part of (N)-layer operation.

The fault administrator (whether it be a person, a dedicated "expert" computer process, a particular facet of a general management process, or some combination of these) can be envisioned as an overseer of network activities who is continually seeking the answers to the following three questions: 1) Is there (or, perhaps, will there be) a fault? 2) If there is a fault, where is it and what are the offending components? and finally 3) If there is a fault and its cause has been identified, how can it be repaired? Thus, the fault administrator is charged with probing and/or monitoring the network for the purpose of: 1) detecting the existence or imminent occurrence of faults, 2) diagnosing the cause of the fault, and 3) setting in motion corrective measures to either repair or bypass the offending component in order to return the system to the highest level of operation as quickly as possible.

Detection of faulty behavior can be achieved in various ways. One way is to test components directly with loopback tests. Another approach is to send threshold information to subordinate fault management agents on the appropriate network nodes. When a threshold is exceeded, the agent notifies the manager that such an event has occurred. This remote servicing of the manager's request is possible because of the common understanding and identification of the resources of concern, and because both the manager and agent are speaking the same language (protocol). Faults may also be detected, in some cases, by functions designed for other aspects of NM (e.g., Performance or Configuration Management) which inform the FM administrator that there is a problem requiring attention.

These thresholds and critical events are used to indicate the actual existence of a fault, or the imminent occurrence of a fault which may still be avoidable. In addition, since not all faults can be detected on-line, out-of-band signaling or other off-line techniques, such as a remote user telephoning the network administrator to report his system malfunction, are still reasonable techniques for detecting network faults and notifying FM to take appropriate action.

The fault administrator can establish a set of a priori assumptions and decisions about which components in the system are subject to faults and, therefore, are of concern to him. He may use certain preprogrammed defaults in this process. He also determines the error levels and other parameters which serve as thresholds to signal fault conditions for each of these resources. In addition, he has the ability, possibly based upon his analysis of actual network operational data, to add new assumptions, change detection

criteria, and even add components to his list of resources potentially subject to faults. The FM administrator is supported by the FM system in all these activities.

Once having detected a fault, the fault administrator must next ascertain the cause of that fault. Possibly enough information has been gleaned from the detection phase to diagnose the cause. If not, however, additional probing will be required to identify the cause of the fault and, if appropriate, the location of the offending component.

Finally, the FM system can attempt to correct the fault using the observational and analytical data derived from the detection and diagnosis stages. Often, FM may enlist the aid of functions developed and designed for other aspects of NM to effect the correction of the fault (e.g., The State Management Function developed for Configuration Management may be used to reconfigure the network around a faulty component). Sometimes problems can be corrected by on-line measures such as remotely rebooting a system. At other times, either the problem cannot be corrected on-line and service personnel must be dispatched for the repair, or the repair cannot be done quickly enough and the component must be bypassed by either on-line or off-line measures.

The following functional requirements address, in greater detail than this brief introduction, the issues involved in providing these types of fault management capabilities to the user.

### 3.3.1. User Requirements

Users expect fast and reliable problem resolution. While they expect very high quality network services delivered on a consistent basis, most end users understand that even the most advanced technology will at times suffer failure. Most end users, therefore, will tolerate an "occasional" outage. When these infrequent outages do occur, however, the end user, generally, expects the problem to be corrected immediately. To provide this level of fault resolution requires very reliable error detection and diagnostic management functions. Redundant paths between major nodes have been implemented in some networks to enhance the possibility of "fault-tolerant" operation and some users even demand redundant fault management to increase network reliability. Most users desire immediate notification when outages occur.

Users expect to be kept informed of the network status, including both scheduled and unscheduled maintenance. In the event of failure, users expect to be notified of the approximate time that the service will be resumed. Users expect reassurance of correct network operation through mechanisms that run confidence tests or analyze dumps, logs, event reports or statistics blocks.

When a fault occurs, a partial or complete resolution can be implemented (i.e., the correction may be temporary or permanent in nature). In a simple case, such as a printer failure, the system may bypass the reported offender by directing printer output to an alternate printer.

After correcting a fault and restoring the system to its operational state, the fault management system must ensure that the problem is truly resolved and that no new problems are introduced. This requirement is called "problem tracking and control."

If the fault management system restores system operation by substituting redundant components, the fault management system should ensure that any failed components be identified so that repairs can be made as quickly as possible to maintain redundancy.

Moreover, FM must provide status reporting of failed items. This should be done in a coordinated manner with CM developed functions since CM also will require status reporting of subsequently repaired resources. The correct current status of resources (operating or redundant) should be available to the administrator at all times. Fault recovery actions such as scheduling repair actions and determining the format of trouble tickets are user requirements not critical for interoperability and, therefore, beyond the scope of our study.

Failure of fault NM functions should not affect regular network operation.

### 3.3.2. Functional Requirements

#### 3.3.2.1. Detecting and Reporting Faults

The NMS must provide mechanisms to allow its users to log events and errors. This includes the facilities to allow users to create and change logging filter settings (logging criteria), to initialize and stop the event/error logging on a routine or a demand basis, to specify what information is to be logged for what duration while the logging filters are satisfied and to specify where the logged information is to be sent. It may be sent either to a local information store or be sent remotely.

The NMS must be able to monitor the specified events or errors. This includes the facilities to allow the NMS users to specify the events or errors to be monitored, to specify the starting and stopping time for monitoring, to specify how frequently the monitored events or errors are to be polled and recorded, and to specify the threshold level or the count when a notification of abnormality should be given. For example, an event can be activities associated with: 1) a specific timer or counter, 2) a group of counters, 3) other parameters, or 4) a virtual circuit.

The ability to anticipate faults as a result of analyzing errors and/or events is needed. (The errors and events may be either monitored or logged.) For example, when a particular line shows an unusually high error rate then the situation must be investigated and appropriate action taken, if necessary, without causing problems for the network users. The method of analyzing monitored and/or logged information does not need to be standardized. However, analysis of errors and faults is a very important and necessary function.

Mechanisms are needed to generate event reports and to send them to the NMS users when a fault has occurred or when it is anticipated (e.g., a threshold level is reached).

A NMS must allow its users to broadcast or multicast notification of faults to network managers or to user specified network entities. This can be done either by a multicast/broadcast facility within the NM facilities, or, alternatively, by coordinated use of a separate facility to provide this messaging service.

#### 3.3.2.2. Diagnosis of Faults

The NMS must allow its users to activate predefined diagnostic and testing procedures for the purpose of determining or verifying where the faults are and for the purpose of testing network components before they are put in use in a network. The diagnostic and testing procedures may be executed either on-line or off-line. At first, the NMS users have to be able to define diagnostic and testing procedures or to select test procedures from a set of already defined procedures. Then the users of NMS need the mechanisms to collect test data and to have access to other network information for analyzing test data in order to isolate faults and to identify the possible causes of it.

The NMS users, including NM expert systems, must be able to request fault related data, such as dumps, statistics blocks and status information for the purpose of diagnosing faults. The NMS must have ways of providing this requested information. In cases where expert systems are used in diagnosing and/or correcting faults, NM users need the ability to inquire as to the status of the diagnosis and/or correction.

NMS users need to analyze event reports of faults, error conditions, and other information they requested such as dumps and the results of testing and diagnosis. Although the method of analysis does not necessarily need to be standardized, NMS users do need standard mechanisms to exchange information for use in proprietary fault diagnostic and analysis programs. And they need standard definitions of the measures and metrics upon which the analysis is based. Fault testing and isolation may need to be standardized along with the measures and derived metrics from the analysis so that all parties involved have a common understanding of the nature and meaning of events.

#### 3.3.2.3. Correction of Faults

The NMS must allow its users to change or reset resource attribute values, take components or lines down, or put components or lines back in service. As corrective actions in response to a fault, NMS users need the ability to request reconfiguration of all or part of the network. (Functions derived for Configuration Management may be invoked to satisfy this requirement).

NMS users require mechanisms to track network operations following fault correction attempts to ensure that the faulty situations are, in fact, corrected. With respect to this capability, NMS users must be able to specify what parameters are to be tracked and how long the tracking is to be maintained.

#### 3.3.2.4. Robust Fault Management

The NMS users require sufficient robustness from fault management, perhaps in the form of redundant fault managers, to ensure that fault management functionality will be available even in the event of a major FM component failure or during maintenance.

### 3.4. Security Management

Security itself is concerned with more than secure communications. Applications using stand-alone and networked computer resources, and those requiring varying degrees of security control or protection, span all sectors of society. These sectors include, but are not restricted to, banking, business, insurance, credit bureau services, legal, national security, and military. Vast computer networks, for example, have evolved to handle transaction processing for banks and retail establishments. Smart cards, on-line sales-inventory control systems, computerized buying services, electronic mail, and electronic office memoranda are just some of the commercial applications which can no longer rely on such traditional security methods as physical control of access to facilities and paper audit trails in order to insure integrity of data and desired levels of privacy. The specific security requirements of government, security agencies, and the military also present significant needs for security which are often described as more stringent and somewhat different from commercial security needs.

With respect to communications interests, networks have the disadvantage of being highly distributed and of affording relatively easy access to network facilities and resources. An interesting problem which arises in this regard for the military, for example, relates to problems of secrecy on LANs. Since LANs operate in a broadcast mode requiring the destination station to recognize messages addressed to it, it is necessary for every station on that LAN to be capable of recognizing or deciphering at least the destination address field of the message to determine if the message is for it. Therefore, in this regard, LANs are susceptible to traffic analysis.

Security, generally, provides for the confidentiality and privacy, integrity, and appropriate availability of data and data processing capabilities (often referred to, appropriately enough, as "CIA," the principles of security). Security "refers to a complex of ... procedural, logical and physical measures aimed at prevention, detection and correction of certain kinds of misuse ... together with the tools to install, operate and maintain these measures ... [Security refers to the] characteristics of data processing systems that give resistance to attack and misuse, intentional or otherwise." [ECMASEC] Security elements perform such functions as that of

limiting access to particular users and/or applications, and limiting access to and corruption of stored and transmitted data.

In order to provide this protective, secure environment, thereby supporting the "CIA" security principles, supportive security services are required. The following is a list of supportive services (usually referred to mnemonically as the "3A's" and the "5S's").

The 3A's are:

Authentication,
Authorization,
Audit/accountability,

and the 5S's are:

Secret,
Sealed,
Sequenced,
Signed,
Stamped.

Security between peers, as with any other communication related activity, requires certain understandings and agreements as to services, mechanisms, and information interpretation. Therefore, in order to perform the security function, a particular security policy must be agreed upon between interacting open systems. Agreement upon the security policy is essential because of the potential for different open systems to adopt policies which will not interoperate. A security policy generally specifies:

-- "how data transmissions between open systems will be protected from unauthorized reception or corruption;

-- how access to resources in one open system will be granted to entities on another open system;

-- how the identity of entities wishing to intercommunicate will be determined with certainty;

-- how and to whom significant events relating to security will be reported; and

-- how audit trail information will be collected, and how and to whom that information will be reported." [SECURE]

Additionally, the security policy should specify how communication services between open systems are reliably provided and maintained. Moreover, the concept of "security domain" depends upon this security policy concept in that a "security domain" comprises a "bounded group of security objects and security subjects to which a single security policy, executed by a single security administration," applies. [ECMASEC]

After a security policy has been selected to govern part or all of a particular communication session between open systems, the security policy is realized by properly configuring the security services, mechanisms and security related information which will control security during that session. These security service applications require functionality both to provide their particular security services (e.g., authentication), and to manage those security services (e.g., key or credential distribution). Security management (SM), then, is needed in this regard, to provide the means by which the security services, mechanisms and security related information are managed.

The issue of importance here is: What is the scope of security management? Security management is involved with such activities as generating, distributing, and storing encryption keys. Password and other authorization or access control lists must be maintained and distributed. Moreover, security management is concerned with monitoring and controlling access to computer networks, or access to all or part of the management information obtained from the network nodes. Logs are an important tool for security management and, therefore, security management is very much involved with the collection, storage, and examination of audit records and security logs, as well as with the enabling and disabling of these logging facilities. Security management oversees the facilities needed to secure communication on a network as well as those needed to secure the management operations themselves.

The OSI Management Framework document [FRMWK], defines SM as follows.

"The purpose of security management is to support the application of security policies by means of functions which include:

a) the creation, deletion and control of security services and mechanisms;
b) the distribution of security-relevant information;
c) the reporting of security-relevant events."

More specifically, SM encompasses three categories of management: system security management, security service management, and security mechanism management [SECURE]. That is, security management deals with the management of "security aspects of the overall OSI environment[,] ... particular security services, such as peer entity authentication and access control," and the security support mechanisms used to provide the security services [SECURE].

Security management can be further characterized as "that aspect of systems management which defines management information exchanges for performing the task of administering systems security." [SECURE] Security management primarily differs from other systems management functional areas only in the class of objects it manages, not in the operations which it uses to manage them. A small sample list of the types of objects appropriate for security objects includes: keys, authentication information, access right information, and operating parameters of security services and mechanisms [SECURE].

Management activities required to support security functions are generally of three types (administration, detection, and recovery). The first

of these types, security "administration," refers to both the gathering (reading) of system security management information and the addition, modification, or deletion (writing) of this system security management information.

The second type of management activity, security event "detection," deals with the auditing of system security operations. Auditing is normally considered to have four components: audit trail content specification, audit trail analysis, audit reporting, and audit trail archiving.

The final type of management activity, security "recovery," is concerned with recovering from an actual, or suspected security attack. This can entail such corrective measures as altering security procedures or modifying security information at appropriate nodes in the system.

The following two sections present a more in-depth view of security management, presenting first the users' view of this aspect of network management followed by the functional elements which are needed to provide such capabilities.

### 3.4.1. User Requirements

In the limited space available here to discuss security management requirements, it would be impossible, and, in fact, undesirable, to enumerate the scores of individual user requirements that have been documented in the references of this report as well as in numerous other documents and implementations not referenced here. However, it is appropriate to attempt to characterize the general nature of these user requirements so as to give guidance in generating and evaluating relevant functional requirements.

Since the function of security is to protect the integrity and confidentiality of "security objects" (i.e., entities "in a passive role to which a security policy applies" [ECMASEC] -- e.g., programs and data), it is essential that security management assure that the services providing this protection are fully functional and have all the support that they need to operate. In fact, security management must be very robust and, to that end, requires a high degree of fault tolerance, for example, to maintain its own "CIA."

Security management provides mechanisms for the protection of network resources and user information. Network security functions should be available for authorized users only. Appropriate validation procedures should be provided to ensure this obligation. In view of this, SM must provide support for security functions by assisting in transferring, monitoring, and controlling security related data; by providing for the recording of attacks and attempted attacks on systems as well as the capability of archiving and retrieving this information; and by providing the capability for notification of security related activity of interest.

The several concepts that follow, often referred to as constituting the major OSI security goals, comprise the set of security facilities that security management is intended to keep intact and properly functioning.

1. The need to prove the identity of security subjects (i.e., those attempting to access some security object). This is generally referred to as Authentication.

2. The need to verify authorization for access to some security object. This activity is generally termed Access Control.

3. The need to prevent disclosure of information. This is generally called Confidentiality.

4. The need to detect various activities such as modification, loss, insertion, replay, or reflection of information. This is referred to as Integrity.

5. And finally, it is often necessary to provide for third party registration of activity in order to be able to be certain that some activity has indeed occurred. Non-repudiation is the term applied to this activity.

Security of the physical location of the network control center in terms of access, data storage, fire, flood and power supply disruption, while all legitimate user concerns, are generally considered to be beyond the scope of OSI security and therefore are also beyond the scope of OSI security management as investigated by this study.

### 3.4.2. Functional Requirements

An analysis of security management yields the following categorization of management activity and requirements into three general areas. These are:

1. **The Ability to Control Access to Resources**

   The security manager must be able to grant or restrict access to the entire network or selected critical parts of the network. The following capabilities can be used to enable the security manager to fulfill this requirement:

   -Authorization control,
   -Authentication control,
   -Control access to security codes,
   -Control access to source routing and route recording,
   -Control access to directories and information bases,
   -Control of updates to directories (including addition, deletion, and modification of directory entries),
   -Control of the distribution of directory information and routing tables,
   -Control of the setting of threshold levels and accounting tables,
   -Prioritized access to requested network resources,

-Maintenance of general network user profiles, and usage profiles for specific resources, for the purpose of controlling access to security resources.

2. **The Ability to Archive and Retrieve Security Information**

Security management requires the ability to gather appropriate information, store the information and access that information for analysis and control purposes. This requirement entails the following capabilities:

-Event logging,
-Monitoring security audit trails,
-Monitoring usage and the users of security related resources,
-Reporting security violation,
-Receiving notification of security violation,
-Maintaining and examining security logs,
-Maintaining redundant or backup copies for all or part of the security related files.
-Maintaining general network user profiles, and usage profiles for specific resources, to enable reference for conformance to designated security profiles.

3. **The Ability to Manage and Control the Encryption Process**

The security manager must both be able to encrypt its communications, when desired, and facilitate the encryption process, in general. This requires the following functionality:

-Encryption (e.g., encryption algorithm selection),
-Key management.

## 3.5. Performance Management

Modern data communications network systems are composed of multiple complex components which must intercommunicate and share data and resources. In some cases, it is critical to the effectiveness of an application that the communication over the network be within certain performance limits. On the other hand, while it is usually desirable to perform at the highest level, performance characteristics may not always be critical, and, at times, less than optimal performance levels can be tolerated.

Performance management (PM) is that aspect of Network Management (NM) which attends to these concerns. It represents a logical division of labor within network management activities. The OSI Management Framework [FRMWK], defines PM as follows:

"Performance management enables the behavior of resources in the OSIE and the effectiveness of communication activities to be evaluated. Performance management includes functions to:

a) gather statistical information;
b) maintain and examine logs of system state histories;
c) determine system performance under natural and artificial conditions;
d) alter system modes of operation for the purpose of conducting performance management activities."

Whereas fault management is concerned with whether all or part of the network is working, performance management is concerned with how well the network or its parts are working. Performance management deals with the quality and effectiveness of network communications. It involves the processes of quantifying, measuring, and reporting error levels, the responsiveness, availability, and utilization of individual network components and the network as a whole. An airline reservations system, which normally provides reasonably fast response times to queries, will be performing at an unacceptable level if response times triple or quadruple, angering both ticket agents and customers. Are delays within acceptable limits for the transfer of data from one station to another? Are response times within reason for the virtual terminal or database application, or are these times so long that the productivity in the office environment is absolutely effected?

At least conceptually, performance management of computer networks includes two broad functional categories -- monitoring and tuning. Monitoring is the performance management function which tracks activities on the network. The tuning function enables performance management to make adjustments to improve network performance. Performance management enlists these mechanisms to provide an awareness of the degree to which the network is fulfilling the service expectations of the users and the degree to which the overall resources of the network are being used. What is the level of bandwidth utilization? Is there excessive traffic? Has throughput been reduced to unacceptable levels? Are there bottlenecks? Are response times increasing? Are performance levels within the limits that the user expects or was promised? These are just some of the issues of concern to a performance administrator.

To deal with these concerns, performance management must initially focus on some basic (default) set of resources to be monitored in order to assess performance levels. This includes associating appropriate metrics and their values with relevant network resources as indicators of different levels of performance. For example, how many retransmissions on a Transport connection should be considered a performance problem requiring attention? Performance management, therefore, must monitor many resources which provide information in determining the performance level of networks. By collecting this information, analyzing it, and then using the resultant analysis as feedback to the set of threshold values of the metrics, the performance administrator can become more and more adept at recognizing situations indicative of present or impending performance degradation.

It should be noted that the definition of PM quoted above from the framework document does not include performance control (needed for tuning) in PM. However, for discussion here it seems appropriate to follow through with the insight that once a problem has been recognized, it is appropriate to try to remedy the situation, regardless of who or what performs that function.

The first stage of PM, monitoring, entailing observation and probing of designated resources, serves to gather information which is stored for later analysis. Particularly in large network environments, the performance administrators (e.g., a human administrator of the performance management software), use this information to analyze the network operation, either manually or through automated methods, and determine areas of performance degradation. His analysis also yields appropriate threshold values which are indicative of various levels of performance. When problems are recognized, a managing system might employ functions developed for fault and/or configuration management to diagnose and rectify the situation. This is necessary to avoid duplication of functionality.

Systems management developed functions which can set and configure managed objects, can assist performance management developed functions to control a degrading performance situation. Often this is accomplished by setting the necessary parameters for management of the traffic on the network. Sometimes, however, additional network resources must be allocated or procured to solve the problem (e.g., more resources may need to be procured, as indicated by capacity planning analysis). If the problem is severe and unexpected, one natural place to turn for assistance is to functions developed for fault management. In fact, it is a natural symbiotic relationship for the monitoring and analytical functions developed for PM reasons to provide input to functions developed for fault management when a situation becomes or is approaching a pathological state.

In order to evaluate whether acceptable performance levels are being maintained, it is necessary to have a well-defined expectation of what they should be. Likewise, it is essential to be able to determine what it is that the performance parameters are indicating.

The following requirements address, in greater detail than this brief introduction, the issues involved in providing these types of performance management capabilities to the user.

### 3.5.1. User Requirements

The behavior of network resources and the effectiveness of interconnection activities need to be evaluated. Before using a network, potential users often want to know information such as the average and worst case network response times for their applications (e.g., VT or FTAM), variability in response, and the reliability and availability of network services. The NMS must also be able to monitor the usage of various resources and be able to focus monitoring efforts upon those resources that are most important.

Network managers need performance statistics to help them plan, manage, and maintain large networks. Performance statistics can be used to recognize potential bottlenecks before they cause problems for the end users. Appropriate corrective action can then be taken. This action can take the form of changing communications traffic routes to balance or redistribute traffic load during times of peak usage or when a bottleneck is identified by a quickly growing load in one area. Over the longer term, capacity planning based on performance information can indicate the proper decisions to make, for example, with regard to expansion of the number of communications lines in that area.

Performance tuning involves first recognizing and diagnosing the existence of performance deficiencies and the load associated with these performance deficiencies, then identifying where and how performance tuning should be done. The use of various performance testing algorithms can help diagnose problem areas and determine where and how to tune the network. Once modifications have been made, it is essential to track the results of these performance tuning efforts to assure the effectiveness of these interventions.

End users expect network resources to be managed in such a way as to consistently afford their applications minimal response time and minimal delays. The end user community will not accept excessive or unanticipated variations in network performance. As performance degrades, business processes within the organization may falter. Therefore, when performance degrades because of traffic levels, high priority messages required for profits or safety should not be affected.

The users want to know that the network has adequate capacity to handle their loads under normal and adverse (e.g., heavily loaded) conditions. Also, as the network grows in size and usage, they are concerned that performance levels can be maintained. They also want to be assured that the network is reliable (i.e., as immune as possible to component failures and transmission losses).

Even though managers of large networks often use analytical and/or simulation models to predict network performance and to help plan for network expansion, actual network traffic and performance measurement is generally still needed to verify modeling results and to provide real network activity profiles as input to these models. Performance statistics should allow analysts to trace the reasons for performance results and anticipate performance changes. In addition, these statistics should allow administrators to anticipate when capability margins are close to being exceeded.

### 3.5.2. Functional Requirements

#### 3.5.2.1. The Ability to Monitor Performance

The NMS needs to be able to monitor performance relevant events, measures and resources. This includes the facilities to allow NMS users to select the events, resources, or measures to be monitored, to specify the starting and stopping times for monitoring, to specify how frequently the monitored events,

measures or resources are to be polled and recorded, to specify other related performance information to be collected during each polling, and to specify the threshold level when a notification of performance abnormality or degradation should be given. The NMS must ensure that statistical measures are based upon a minimally accepted number of samples. When time stamps associated with the reported information or activities or events are involved, either adjustment to the same time instance is required or synchronized clocks among devices are needed.

The collected performance information needs to be parsed and reduced for performance analysis. The NMS users need to be able to select the performance analysis criteria and algorithms for various performance measures.

#### 3.5.2.2. The Ability to Tune and Control Performance

When an indication of a performance abnormality or degradation is reported, the NMS user needs mechanisms to execute predefined performance tests, and to collect test results for the purpose of diagnosing network performance anomalies and determining the appropriate performance tuning strategy. The same set of functions can be used to evaluate the results (i.e., the effectiveness) of performance tuning.

NMS users need to be able to change (potentially non-OSI) resource allocation (e.g., buffer allocation method, flow control allocation level), to modify resource (managed object) attributes and to set managed object attribute values in order to provide better performance or resolve performance problems that cause bottlenecks or prevent the flow of high priority data. Some of these capabilities are the same as those required for configuration management.

#### 3.5.2.3. The Ability to Evaluate Performance Tuning

The NMS users need the ability to keep track of the performance tuning results in terms of user specified measures or criteria. The set of facilities required to satisfy this evaluation (tracking) function may duplicate some of those functions that are required for performance monitoring and tuning.

#### 3.5.2.4. The Ability to Report on Performance Monitoring, Tuning, and Tracking

Notification of abnormal performance changes must be able to be spontaneously generated and sent to the NMS users who have previously requested such notification. Recognition of such abnormal performance changes may occur, for example, when performance measure thresholds are exceeded.

The network managers need to initially select a predefined domain (a set of systems to be managed) or subdomain as a base for generating performance trending reports. In addition, NMS users need to be able to create, delete and modify domains or subdomains within a network in order to view the performance characteristics of a portion of a network that is of particular interest to them.

Mechanisms are needed for NMS users to request and obtain performance reports based on user specified criteria (e.g., throughput of a specified domain in the past 24 hours, the past week, month or year). The performance trending reports that represent current (real-time) network performance or network performance histories can be saved for a specified period of time to be used as benchmarks. These reports can relate to an individual domain, subdomain or the network as a whole. Accumulation of daily performance reports into weekly, monthly and yearly reports should be possible.

The NMS should have the ability to display routine snapshots of network performance in terms of those user specified measures (e.g., network utilizations versus load, traffic or load distribution among subnetworks, ratio of overhead packets to data packets, detail of traffic profile and peak hour rates, and average network response time). The NMS should have the ability to compute and display statistics of standard metrics such as average, median, maximum, minimum, ratios and standard deviation.

#### 3.5.2.5. The Ability to Test Capacity and Special Conditions

In order to assure that capacity margins for network components are sufficient, it may be necessary to run tests to determine the effects of additional network loading under natural or artificial conditions. By imposing test loads on the network, managers can perhaps more accurately predict when additional equipment must be brought on-line than they could using analytic or simulation models. Tests could also be designed to establish the effects of equipment failures to determine the reliability of the networks. It might be possible to run tests during normal operations to determine incremental effects, but often tests should be run during late evening or early morning off-peak hours under artificial conditions so as not to disturb real user traffic or to obtain test results under more controllable conditions that do not include "random" background traffic. Such tests can also be used to validate simulation or analytic models.

### 3.6. Accounting Management

"Accounting management enables charges to be established for the use of resources in the OSIE, and for costs to be identified for the use of those resources. Accounting management includes functions to:

a) inform users of costs incurred or resources consumed;
b) enable accounting limits to be set and tariff schedules to be associated with the use of resources;
c) enable costs to be combined where multiple resources are invoked to achieve a given communication objective." [FRMWK]

Accounting management provides information about the use of resources for cost analysis, tracing network usage, and user billing. The results of statistical analysis of accounting information will help plan network expansion or the types of network services and may further indicate the trend of the development of new network technology. Information gathered for

accounting may also prove useful in the development of "expert systems" for automated network management. (See sec. 6.)

### 3.6.1. User Requirements

When a network or a system of interconnected networks are used by more than one organization or cost center within an organization, there is a need to apportion costs for network services in proportion to the amount of resources consumed by each chargeable user. Therefore, information on usage of resources needs to be recorded, collected and archived to provide the necessary information for proper distribution of resource costs.

NMS end users and administrators need to be able to specify the kinds of accounting information to be recorded at various nodes, the desired interval between sending the recorded information to higher level management nodes, and the algorithms to be used in calculating and reporting the accounting information. Accounting reports should be generated in user specified form and sent to user designated output devices.

In order to limit access to accounting information, the NMS must provide the capability to verify user's authorization to access and manipulate that information. As with other areas of network management, accounting management should have minimal effect on network performance.

The ability to adjust billing rates and accounting factors are legitimate administrative requirements, but they are outside of the scope of this study.

### 3.6.2. Functional Requirements

#### 3.6.2.1. The Ability to Record and Generate Accounting Information

The NMS must allow its users to specify the accounting information to be collected as well as the duration of the collection period, or the criteria to be used to determine the duration of the collection period. The definition of accounting information and units must be defined by related standards making groups. The format and options associated with this accounting information also need to be defined by NM standards. Examples of accounting information include network user network connection time, quantity of data transmitted, and class of service provided for each user or group of users.

Mechanisms are needed for the NMS, through layer management entities or the layer entities, to record and/or collect user distinguishable accounting information, and to generate accounting messages. The generated accounting messages are to be forwarded to the default or NMS user specified files or nodes.

#### 3.6.2.2. The Ability to Specify Accounting Information to Be Collected

The NMS must allow its users to specify what accounting information is to be collected (e.g., connection time or transmission time) and how it is stored (e.g., selecting predefined accounting units or even calculation algorithms for statistics about resource usage and individual or group user charges).

Daily raw accounting information or calculated statistics are saved in data bases for use in periodic weekly, monthly or quarterly accounting reports.

### 3.6.2.3. The Ability to Control the Storage of and Access to Accounting Information

The NMS must provide standard procedures to retrieve and store accounting information and standard ways to name archived accounting information files. Accounting information can be stored on disk, or output to the printer or screen.

Access to accounting information is limited to authorized personnel only. If current access authorization algorithms are not adequate, new algorithms must be developed. To make use of these algorithms, mechanisms must be provided to allow authorized NMS users to select or change authorization algorithms.

### 3.6.2.4. The Ability to Report Accounting Information

The NMS must be capable of reporting degree of resource usage and resource usage charges at an NMS user specified level. In other words, the NMS user can specify what information is to be reported, in what form to report it, and to which network user's accounting profile or to which output device to report it.

Mechanisms are needed to allow NMS users to create and transmit selectively or broadcast, as appropriate, accounting news such as network resource billing rate changes or accounting limit changes.

### 3.6.2.5. The Ability to Set and Modify Accounting Limits

The network manager must be able to read, set and change accounting limits for various groups of users. For example, to balance the usage of network transmission capacity among users, a user may be limited in his access to the network to various levels at different times depending on the overall network load or accounting (finance) policy changes.

Mechanisms are needed to allow the network manager to change the priorities assigned to the network users for access to network resources, including, for example, the priorities in using various classes of network services at various layers.

### 3.6.2.6. The Ability to Define Accounting Metrics

In order to obtain and use the accounting information, the standard metrics and the definition of accounting information unit need to be established. For example, the type of accountable units needs to be defined such as the call (connection) duration or the number of bits, characters, blocks or files transmitted. A standard method for defining these accounting metrics will allow expansion or changes of accounting metrics.

### 3.7. Other Requirements

The aspects of network management which are not critical for interoperability and therefore do not require standardization lie mainly in the areas of man-machine interface, analysis, and the management of large amounts of management data. These areas are important, however, since they may be keys that determine the popularity and usefulness of a network management system. They may be major factors in discriminating among implementations of NM standards.

The man-machine interface should offer the NMS user a standard management application interface. The interface should enable the NMS user to quickly and easily comprehend the network management system's capabilities, to use the NMS efficiently, and to allow flexibility in performing the desired operations.

The inclusion of "help text" to explain the use of and purpose of the network management commands can be quite useful. Easy and efficient input to the NMS can be provided by menu-driven management commands, programmable function keys, and mechanisms to permit users to build command files. The output generated may include 1) color graphics (e.g., different colors can be used to indicate the severity of faults or the degree of traffic load distributed over a single network or interconnected networks), 2) choices in formatting and presenting displays and reports, and 3) generation of real-time and/or historical displays or reports concerning areas of particular interest determined by user specified criteria.

Analysis of collected management information is very important for successful network management. Efficient analysis tools help, for example, to diagnose and isolate faults. The analysis results can be used to predict network performance, to forecast network expansion needs, to plan network upgrades, to balance network load among network users or groups of network users, to minimize the network management cost, and to optimize overall network performance.

Because of the size and complexity of data communications networks, expert systems may be required for the management of network management information. (See sec. 6.) Successful network management presupposes that the NMS can process information at a rate faster than the information is generated.

# 4. ANALYSIS OF FUNCTIONAL REQUIREMENTS

This section contains an analysis of the network management functional requirements identified in section 3 and their relationships to the emerging OSI standards. Since the OSI standards are not yet complete, the latest available ISO documents were used as a basis of comparison. These are primarily those available as a result of the December 1988 ISO/IEC JTC 1/SC 21/WG 4 meeting in Sydney, Australia and include working drafts, Draft Proposed standards (DPs), Draft International Standards (DISs), International Standards (ISs), and reasonably stable editing drafts. This section is organized similarly to section 3, i.e., by Functional Areas (FAs), such as Performance Management or Configuration Management. Within each of these areas, an attempt is made to identify problems, omissions and errors in the emerging standards with respect to the (previously identified) functional requirements.

An earlier draft of this report, including a preliminary analysis, was made available for comment by network management experts in October 1987. Further refinement of that initial analysis was performed as a result of discussions held at a NBS/MITRE sponsored Workshop on Network Management Functional Requirements held October 28-30, 1987 in Bedford, MA. This analysis reflects the results of that earlier draft and workshop.

There are some requirements, concerns and issues that transcend functional areas. These are discussed in section 5 of this report.

## 4.1. Analysis of Architectural Requirements

The primary source of architectural information on OSI management is the OSI Management Framework [FRMWK]. This document defines and outlines some major components of OSI management such as systems management, layer management, the Management Information Base (MIB), and the OSI management functional areas (e.g., Fault, Accounting, and Performance Management). This document has reached International Standard (IS) status in 1989 as a result of an editing meeting in October 1988 as well as comments made at the December 1988 SC 21/WG 4 meeting in Sydney.

Another important source of architectural information is the Systems Management Overview [SMO]. The SMO further expands on concepts introduced in the Management Framework. The SMO introduces concepts on management processes acting in either a managing or an agent role. It discusses functional aspects of systems management, management domains, Application Layer concepts relating to systems management, and systems management standards. This document reached DP status as a result of decisions made at the December 1988 Sydney WG4 meeting.

### 4.1.1. Model

The model of a network management system (NMS) needs to clearly describe its components and their organization and relationships. The Management

Framework (MF) includes a somewhat superficial description of its components. While it introduces important concepts, it lacks detail. More details relating to systems management are supplied by the Systems Management Overview. The SMO describes important components and concepts for management, but, quite correctly, omits discussions of layer management issues.

Perhaps a layer management overview document is needed to present the principles and theory of layer management as opposed to systems management. Such a document could serve as a guideline to layer management standards developers and implementors. Layer management might be used when systems management is nonfunctional. Such a situation might occur, for example, if a Transport Layer parameter were incorrectly set such that Application associations were inhibited. Link Layer management might be used in such admittedly unusual circumstances to "reset" Transport without dispatching service personnel. It should be cautioned, however, that in general, layer management should not duplicate functionality that the full seven layer stack provides.

While layer management and systems management are distinct, they are interrelated. Systems management is concerned with managing resources, many of which are contained within layers. The SMO does little to clarify issues relating to interactions between layer management and systems management.

Taken together the MF and the SMO outline the basic components of management and provide details on system management but do not provide in sufficient detail how, even conceptually, these components interact with layer management entities. The relationships, interactions, and interdependencies among NMS components need additional refinement. This need has been recognized by several organizations such as the COS Network Management Subcommittee (NMSC) and suggested in similar work by the IEEE 802.1 Systems Management Project.

In section 3.1.2.1, eight additional requirements for the OSI architectural model are identified. These include a need for hierarchical and distributed control of access to and manipulation of network management information and network resources. These forms of control are allowed by the OSI management architecture but additional work is needed on multiple manager considerations. (See sec. 5.5.)

The second additional requirement identified is a need for centralized management of distributed network managers. This is allowed by the OSI management architecture but additional work is needed. (See sec. 5.5.).

The third requirement from section 3.1.2.1 is the need for architectural flexibility to accommodate new technologies such that a new NMS is not needed with the introduction of each new technology. Although it is impossible to predict with certainty that all new technologies can be accommodated, it appears that the proposed OSI management architecture is sufficiently flexible to accommodate many technological changes. The OSI management model uses the services of the seven-layer OSI model, so to a great extent the management model is as adaptable to changing technology as is the underlying seven-layer OSI model.

A fourth requirement identified is the need to accommodate additional proprietary network management solutions. This need appears to be met by the OSI management model. (See sec. 5.14 on "Extensibility.")

A related additional requirement is the need to accommodate proprietary (non-standard or ad hoc) network management security control mechanisms. While it appears that this requirement can be met, additional analysis by security experts is required to determine this with certainty.

The sixth additional requirement identified is the need to allow for future network management expansion, primarily the ability to manage new object types. This requirement is met by the OSI model, since the proposed method for identifying managed object uses ASN.1, a rich and flexible data description tool. (See also sec. 5.14 on "Extensibility.") While the details of naming and identifying managed objects is not yet complete, the ISO/IEC JTC 1/SC 21/WG 4 December 1988 meeting in Sydney produced a Management Information Model [MIM] that clarifies many issues. This document was updated and recommended to advance to DP status at the SC 21/WG 4 interim meeting held in April 1989.

The seventh additional requirement for the architectural model is the need to support the use of redundant managers and out-of-band signaling to ensure NMS availability. The OSI management model allows for redundant managers. However, additional work is needed to specify how these redundant managers are to work in concert. (See sec. 5.5 and 5.16.) The use of out-of-band signaling is not precluded by the management model, but so far there has been little or no development work to support this concept.

There is currently no specific support for message distribution in a hierarchical distributed NMS. Current CMIS is not multicast and is uses connection-oriented underlying services. (See sec. 5.4.) However, there is nothing in the current OSI management architecture to exclude future extensions for broadcast/multicast services. [Note: It is generally assumed that connection-oriented Transport will be used to support CMIP, but, in the future, it might be possible to use Transport's connectionless datagram service.] For statistics involving non-real-time scheduling (i.e., not urgent), OSI-based X.400 mail systems could be used.

### 4.1.2. Services and Protocols

The requirement, stated in section 3.1.2.2, for a manager-to-agent information exchange service and protocol may be met by the Common Management Information Service and Protocol (CMIS/P) currently under development by ISO/IEC JTC 1/SC 21/WG 4. There are some questions that remain to be answered with regard to CMIS/P's lack of broadcast/multicast (multipeer) capabilities and the related issue of connection-oriented versus connectionless services. (See sec. 5.4.) These questions suggest a serious deficiency in meeting the information exchange and protocol requirements of OSI management. Some OSI experts and groups have proposed using X.400 mail and FTAM to support management activities.

Another requirement stated in section 3.1.2.2 is for manager-to-manager information exchange service and protocol. This capability is essentially provided by the concepts of managing and agent processes, as well as management domains as discussed in the Systems Management Overview [SMO]. Much more development work is required by the ISO to support multiple managers in a NMS. (See sec. 5.5.)

### 4.1.3. Resource Identification

There is a requirement that both OSI and non-OSI resources to be managed must be identified. This identification process for OSI resources is being met by the various layer groups within the ISO and other standards making organizations (e.g., IEEE 802). A versatile method for naming and identifying resources, i.e., managed objects, is proposed in the Management Information Model [MIM].

With regard to non-OSI resources, there is no recognized authority to perform the identification process. Therefore, it is probable that each vendor will extend resource identification to those proprietary extensions within his implementation as he sees fit. The use of such extensions may lead to significant loss of interoperability within OSI management products. (See sec. 5.14, "Extensibility.")

To help assure that similar managed objects are treated in a similar, consistent manner, the ISO is developing a specification for the Structure of Management Information (SMI). SMI includes four parts, each a separate document. While these are by no means complete, they greatly contribute to the understanding of how managed objects are to be remotely accessed and managed. The four parts are:

Part 1: Management Information Model [MIM] This document contains two important parts. The first describes an information model which includes Object-Oriented Design Principles and Managed Object Classes and Inheritance Relationships. The second part is concerned with Principles of Naming including Containment Hierarchy and Name Structure. Taken as a whole this document, while not yet complete, does much to ensure that managed objects are treated in a consistent, versatile, logical manner from implementation to implementation and to ensure that they can be named and located efficiently and consistently.

Part 2: Definitions of Support Objects [DSO] This document provides the definitions of object classes which may be useful (or required) to support management. The four object classes currently defined (more may added later) include: "Top" (the root of the object class tree), "Discriminator" (used as a template for event management), "Event Forwarding Discriminator" (used for establishing criteria for generating event reports), and "Service Access Discriminator" (used by receivers of reports to define criteria needed to select incoming reports for processing). The ASN.1 definitions are given in an annex.

Part 3: Definition of Management Attributes [DMA] This document provides definitions of attributes that may be useful for management. Probably the more interesting of the 31 definitions given are: Count, Gauge, Threshold, Counter-Threshold, Gauge-Threshold, and Tide-mark. The ASN.1 definitions are included in an annex.

Part 4: Guidelines to the Definition of Managed Objects [GDMO] This document provides useful information to layer standards groups for defining managed objects needed to manage their layers. Primarily, important concepts from the other three parts of SMI are introduced. In addition, useful discussions on event definitions, action definitions and conformance issues are included. There is also an annex on Notational Tools for Managed Object Definitions, including ASN.1 as appropriate.

### 4.1.4. Information Structure

The requirement is stated in section 3.1.2.4 for an information structure or common view of managed objects in a heterogenous management environment. In such an environment, there may be differences in data formats, structure, and semantics. This requirement appears to be met by the proposed OSI management standards through the use of tools such as ASN.1 and specifications of the Structure of Management Information (SMI).

### 4.1.5. Layer Management

The requirement for layer management is met by the emerging OSI architecture and related standards, although the effort is not moving along as quickly as desirable, especially at the upper layers. There are also concerns about the lack of specificity in the interrelationships between systems management and layer management, as discussed in section 4.1.1. The development of the Guidelines to the Definitions of Managed Objects [GDMO] should accelerate the work of layer management.

### 4.1.6. The Directory

Section 3.1.2.6 states the architectural requirements for Directory in OSI management. It is beyond the scope of this report to discuss whether or not the current Directory proposal can meet these requirements since such a discussion would require analysis by Directory experts.

### 4.1.7. Network Management Communications Overhead and Performance

Section 3.1.2.7 states the requirement that NMS users expect minimal impact of the NMS system itself on network performance. The use of out-of-band signaling to meet this requirement is suggested there. While its use is not precluded by the management architecture, there has been little or no development work so far to support this concept. (See sec. 5.16.)

### 4.1.8. Support for Efficient Information Transfer

There is a requirement that software distribution and dumping of statistics blocks employ a file transfer capability. Such capability is not precluded by the OSI Management Framework (MF). However, current efforts on information exchange services and protocols have focused on the use of the Common Information Management Service and Protocol (CMIS/P) for all management information exchange purposes. (See sec. 4.1.2.)

When a file transfer capability is needed, the ISO File Transfer, Access and Management (FTAM) services and protocol can be used. While FTAM appears to offer more than sufficient functionality to meet the needs of OSI management bulk data transfer, issues relating to upper layer architecture must be clarified to enable the use of FTAM and CMIP over a single management association. Multiplexing FTAM and CMIP over a single association, while not always necessary, is certainly a desirable capability and may be a virtual requirement for future management systems. The OSI/NM Forum, a consortium of over 40 major vendors of network management products and services, has adopted FTAM as part of its Interoperable Interface Protocols. Therefore, the eventual use of FTAM in OSI management is virtually certain.

### 4.1.9. Standardization of Terminology

An architectural requirement noted in section 3.1.2.9 is the need for standardization of the terminology used across multi-vendor network management products. This requirement calls for common definitions of terms used for configuration states and relationships, for security, for the model, for performance, and more. It is not clear that this requirement is met by current ISO efforts. (Note: There is an effort by the ISO to define terms used by OSI, but there is no known effort to define terms specifically used within OSI management.)

## 4.2. Analysis of Configuration Management Requirements

Configuration management (CM) is one of the five management functional areas defined in the Management Framework [FRMWK] undergoing standardization, and has been identified by the COS Network Management Subcommittee and others as the most important of these areas for several reasons. One is that the functions it developed are used to set and determine the states of network components -- a fundamental requirement for management. Another is that a network cannot be managed without CM, since it provides key services such as bringing a network into operation, determining its status, and bringing it, or parts of it, down.

CM (along with Fault Management) is a management functional areas scheduled for IS status by April 1991 (relatively early completion) by the ISO/IEC JTC 1/SC 21 subcommittee. Three functions developed by CM have already reached Draft Proposed (DP) standard status, more than any other management functional area. The three functions are: Object Management, State Management, and Relationship Management.

In section 3, a set of functional requirements was identified as necessary for CM. In this section, we will analyze those requirements to determine whether or not the emerging Configuration Management Functions can meet them.

Before proceeding, it is important to note some change of terminology from that used in section 3, where, for example, the functional requirements for CM were discussed using terms such as "resources" and "attributes" to describe objects subject to management and their associated attributes. In this section, the term "attribute" is retained, but "resource" is replaced by the term "managed object" to keep in concert with the latest ISO working papers, which recognize that OSI resources are a special case of potential objects subject to management. The ISO terminology is used here to facilitate understanding of the reader familiar with the ISO working papers.

It is also worth noting that ISO/IEC JTC 1/SC 21/WG 4 at its recent meeting in December 1988 in Sydney, made the decision to divorce functions from the management functional areas that developed them. For example, Configuration Management developed the Object Management Function [OMF]. The document describing this function ([OMF]) is now no longer part of the Configuration Management Working Document. The Object Management Function is available for systems management use by any authorized user. There is no artificial restriction that it must be used for configuration purposes. For example, it could and probably will be used to set threshold values for alarms most directly related to fault or perhaps performance management purposes.

#### 4.2.1. The Specification of Resource Attributes

The ability to specify attributes associated with network resources essentially requires a Network Management System to be able to define new managed object types. While there are no plans to allow such definitions on-line, implementors are free to use ASN.1 to describe new proprietary managed object classes using the Guidelines for the Definitions of Managed Objects [GDMO]. Permitted attribute values can be included as part of the definition. Permitted state values can also be included when a managed object class is defined. Furthermore, a technique known as "Polymorphism," discussed in the Management Information Model [MIM], can be used to extend object classes to support new types of equipment and technology without making older management systems obsolete.

#### 4.2.2. Setting and Modifying Attribute Values

The requirement to remotely set and modify attribute values is primarily met by services defined in the Object Management Function [OMF]. The function does not specifically provide the capability to set system clocks, but this can be provided if a system clock is defined as an attribute to be managed. Such a definition should be provided by SC 21/WG 4.

#### 4.2.3. Defining and Modification of Relationships

The requirement to remotely define and modify relationships among network

resources is met by services defined in the Relationship Management Function document [RMF].

### 4.2.4. Examination of Attribute Values and Relationships

The ability to remotely examine attribute values is provided by the services of the Object Management Function [OMF]. The ability to remotely examine relationships among network resources is provided by the Relationship Management Function [RMF].

### 4.2.5. Distribution of Software Throughout the Network

The requirement for software distribution capability is recognized as an outstanding issue in the Configuration Management Working Document [CONFIG]. It is further discussed in SC21 N3310 "Software Management Function" [SFMF]. However, this work is still formative. SC21 recognizes the need for software management, but up to now little work has been done on this issue. It clearly has lower priority than other functions, but has not been overlooked, simply deferred. It is apparent that much more work is required before this function can be considered to be stable. This distribution function, common in vendor-proprietary network management systems, is a serious omission from needed OSI management capabilities.

The requirement for distribution of routing information is a matter of concern to the Network Layer (Layer 3), beyond the scope of this study. However, there may be accounting, performance, or security concerns that require management intervention into routing decisions that cannot be solved by mathematical algorithms alone. This management intervention may require additional interfaces to the Network Layer that have not yet been defined. Further study of this problem should be considered.

### 4.2.6. Initialization and Termination of Network Operations

The basic requirement to provide these capabilities appears to be partially satisfied by the State Management Function [SMF]. States and transitions are defined in the SMF to allow for the initialization and termination of managed object operations. However, initialization and shutdown of entire managed objects is still in the list of CM Outstanding Issues [CONFIG]. Further work in this area is clearly needed. The requirement to allow for the remote reboot of a system is also in the CM Outstanding Issues list. OSI Systems Management does not specifically address initialization and termination of network operations but only considers these operations on open systems. The initialization and termination of network operations must be accomplished through the use of yet to be developed functions for systems initialization and termination.

### 4.2.7. Verification of NMS Users' Authorization

Verification of each user's authorization for performing configuration management functions is primarily provided as part of the normal security mechanisms available to all OSI resources and can be managed through the functions provided by the Security Management Functional Area. However, it is

important that the proper access controls be assigned to all critical system resources to ensure that the security mechanisms perform their intended functions. This assignment may, perhaps, be best discussed in a "guide to implementors and/or managers" document. Such a document (or documents) should be produced to guide OSI system managers as how to best protect their resources from intrusion. Discussions on security related issues should be included in the Guidelines to the Definitions of Managed Objects [GDMO]. Security is not currently considered in this document.

### 4.2.8. Reporting on Configuration Status

The ability to report on configuration status, as discussed in section 3.2, includes several requirements. The first consists primarily of the ability of agent systems to report changes to their managers as these changes occur, i.e., asynchronous event reporting. For the most part this requirement is met by a combination of concepts described in the State Management Function [SMF] and the Attribute Change Event Report Service of the Object Management Function [OMF].

The second requirement discussed in section 3.2.2.8 is essentially met by the ability of manager systems to poll their agents in order to determine each agent's configuration status. This ability is provided by a combination of services included in the Object Management Function.

One requirement discussed in section 3.2 that is not met by the current management functions, or other OSI services, is the ability to broadcast or multicast messages about configuration changes. (See sec. 5.4 for a discussion of connection-oriented versus connectionless services.) There is a clear need for broadcast/multicast services that is not being met by the emerging OSI standards.

## 4.3. Analysis of Fault Management Requirements

Fault management (FM) is one of the five Management Functional Areas defined by the ISO that is undergoing standardization. FM, progressing second only to Configuration Management, has developed one function, entitled "Error Reporting and Information Retrieval" [ERIRF] that has reached Draft Proposed (DP) standard status.

In section 3 a set of functional requirements was identified as necessary for FM. In this section, we will analyze these requirements to determine whether or not the emerging functions and other documents developed by FM can meet them.

The primary source document used to perform this analysis was the emerging FM Management Functional Area Fault Management Working Document [FMWD]. This document, a reorganized update to older Fault Management documents, was produced according to editor's instructions developed at the December 1988 meeting of SC 21/WG 4 in Sydney. It reflects new document structure in that each function developed by one of the Management Functional Areas is to be produced as a stand-alone document. (Therefore, for example,

this latest FM Working Document does not include Error Reporting and Information Retrieval.)

The FM Working Document does include background discussions and overviews of FM concepts. It also includes discussions of functions that have not yet reached DP status, including the Confidence and Diagnostic Testing Function. Annexes to the document include discussions of FM concepts and requirements, a model for the Confidence and Diagnostic Testing Function, and FM outstanding issues.

The Error Reporting and Information Retrieval Function [ERIRF] provides two services. These are (not surprisingly) the Error Reporting Service and the Information Retrieval Service. The former provides a service in which error reports are generated by a management system when some preselected event occurs (e.g., power low, error count exceeded). The latter allows a management process to employ a CMIS M-GET to retrieve pertinent error information.

### 4.3.1. Detecting and Reporting Faults

The ability to detect and report faults includes several requirements. The first of these is the ability to log events and errors, i.e., the ability to locally record event reports within the managed system that generated the event (i.e., "logging") and then subsequently retrieve these event reports remotely (i.e., "log retrieval"). Such capability is to be provided by the Log Control Function [LCF], which is still in the early stages of development. Eventually it is expected to meet the requirements for remote log access and maintenance.

The second requirement is the ability to monitor events and errors. This requirement appears to be met by a combination of services including some from the Management Services Control Function [MSC] to select criteria for enabling reporting of events, and some from the Error Reporting Service of the Error Reporting and Information Retrieval Function to transfer the reports. The information to be recorded with each event is specific to the managed object generating the event and must be specified by layer groups, i.e., the groups responsible for defining OSI layer standards, (or, in the case of nonstandard extensions, implementors) who define the managed object class. Advice for defining event reports is specified in the Guidelines for the Definitions of Managed Objects [GDMO].

The third requirement is the ability to anticipate faults. This appears to be partially addressed by services offered by the Error Reporting and Information Function and by the Managed Service Control Function, which when applied to thresholds and gauges, can be used to support this requirement.

The fourth and last requirement is the ability to broadcast or multicast notification of faults to network managers and user specified nodes. This requirement is <u>not</u> met by current nor actively proposed FM services. Issues involving multiple managers and broadcast/multicast services are discussed in section 5, "Additional Issues," of this report. (See sec. 5.4 and 5.5.)

### 4.3.2. Diagnosis of Faults

The ability to diagnose faults includes several requirements. The first is the ability to activate diagnostic and testing procedures. This requirement will likely be met by the still formative Confidence and Testing Function, which is discussed in the Fault Management Working Document [FMWD]. This function will provide not only the ability to activate testing, but also to test, report results, terminate, and report failure. It will provide the following services:

(a) Connectivity Test,
(b) Data Integrity Test,
(c) Protocol Integrity Test,
(d) Data Saturation Test,
(e) Connection Saturation Test,
(f) Response Time Test,
(g) Imaging Loopback Test,
(h) Function Test,
(i) Diagnostic Test.

The second requirement is the ability to request dumps, statistic blocks and status information. It appears that this requirement will be met by the test reporting aspect of the various FM Confidence and Diagnostic Testing Functions test services, although further refinement is needed, both in specific managed objects upon which tests are connected and especially with regard to the ability to produce dumps. (This is an area related to software distribution, discussed in the Configuration Management analysis and recognized by SC 21/WG 4 as needing additional refinement.)

The third and last requirement is the ability to analyze the results of diagnosis and testing. This requirement is an area that is beyond the scope of current standardization efforts, i.e., each vendor/implementor may apply his own analysis methods.

### 4.3.3. Correction of Faults

Fault correction includes two requirements. The first is the ability to manipulate managed object attribute values and states. This requirement can be effectively satisfied through the services of the Object Management Function [OMF].

The second requirement is the ability to track corrections to fault conditions. This requirement will apparently be partially met through the use of the services of the FM Confidence and Diagnostic Testing Function, in addition to which vendor-proprietary, implementation-dependant, methods of cataloging (maintaining records of) faults are needed. Logging, to be defined by the Log Control Function, may also be useful.

### 4.3.4. Robust Fault Management

This requirement for a redundant set of fault management facilities (stated in sec. 3.3.2.4) is not met by current nor actively proposed by current FM proposals. Concerns that peripherally touch upon redundancy are discussed in the Configuration Management Working Draft [CONFIG], where Annex A cites a requirement to activate a "standby configuration". In addition "backup status" is discussed in the FM developed Error Reporting Service [ERIRF]. However, the current ISO standards do not appear to discuss the overall philosophy of dealing with system robustness issues. Concerns involving multiple managers and network management robustness are discussed in section 5, "Additional Issues," of this report. (See sec. 5.5 and 5.16.)

## 4.4. Analysis of Security Management Requirements

Security management (SM) is a difficult area for the management community to deal with since SM mechanics are specialized and expertise about security tends to be secretive. Furthermore, SM mechanisms exist throughout several layers and through most vendor-specific (implementation-dependent) aspects of an implementation. For these and other reasons, related standards for security architecture and mechanisms are still very much in the formative stages.

Nevertheless, security management is of considerable importance since many applications are quite dependent upon the security mechanisms to assure reliable application operation. This is evident in the banking community, for example, with regard to the need for key management. This function is concerned with the generation and distribution of keys used for encipherment and message authentication. Responsibility for the development of standards relating to key management is not yet clear. It is probably within the scope of SC 21/WG 4, or possibly SC 27 or TC 68.

The primary documents addressing the OSI Security Management standardization concerns are the Security Architecture document [SECARCH] and the Security Management working document [SECURE]. While the Security Architecture document is currently at the DIS level and provides a reasonably complete picture of security architecture, the Security Management document is still a working document (currently the Fifth Draft), requiring a great deal more work to reach completion. However, although this working document is currently rather sketchy, it nevertheless does appear to address most of the SM requirements mentioned previously in section 3, and show potential for addressing those requirements not yet explicitly handled.

The Annex to this report includes an additional presentation of work done by the European Computer Manufacturers Association (ECMA) with regard to Security and Security Management. Although this group is not sponsored by ISO, it has, nevertheless, produced an important reference document, [ECMASEC], which is used by ISO developers and, in fact, offers a considerable amount of detail and analysis of Security and Security Management issues beyond the current ISO work. For this reason, a discussion of information from this ECMA document has been included as an Annex to this report in order

to provide supplementary material to this Section and greater definition to the specification of requirements for Security Management.

The following list presents the four functions for Security Management defined by the Security Management document [SECURE].

1. The Security Audit Trail Function "provides for the collection and review of security events for the purpose of monitoring the operation of that portion of the security policy implemented in the open systems."

   This function enables the SM user to "maintain a record of security-related events which occur in the SM domain." It further enables the user to "review and analyze these events in order to detect security breaches, malfunctions and effectiveness of the security services and mechanisms which are implemented pursuant to the security policy."

2. The Security Alarm Function "provides event reporting for detecting security attacks and malfunctions."

   This function assists the SM user in knowing "when the security services and mechanisms are malfunctioning, and when system security has been attacked or breached." This function provides current monitoring and notification capabilities, as opposed to the historical record-keeping and record reviewing capabilities of the previous function.

3. The Security Object and Attribute Management Function "provides the means to manage security objects."

   This function enables the SM user to "control security-related objects and attributes ... [by such means as] creating and deleting objects, changing their attributes, changing their state, and affecting the relationships between security-related objects."

   It is interesting to note that this function and the previously described function are just some of the examples of functions which have not been definitively categorized as belonging to security rather than security management, or vice versa. The still fuzzy distinction made between security and security management is very much a factor in this regard. For example, it is clearly within the purview of SM to define when security services and mechanisms are malfunctioning. However, it is not quite so evident who bears the responsibility (i.e., security or security management) for defining what constitutes a security attack or a security malfunction. Thus, jurisdiction for these determinations is not, at present, unambiguously assignable to the standards groups responsible for security, or to those responsible for security management.

4. The Security Alarm and Audit Trail Management Function "provides the ability to establish and configure security audit trails and alarm reporting relationships."

   This function permits the SM user to "control the operation of the audit trail and security alarm functions, selecting which events are to be reported, to whom and under what circumstances."

All four of these SM functions are intended to be implemented either by a direct use of basic CMIS services, or by indirect use of these CMIS services as filtered through higher order functions such as the Object Management Function, the State Management Function, the Relationship Management Function, the Management Service Control Function, or the Log Control Function. Such an implementation strategy, however, can mean some potential problems for the architecture of a trusted system. For example, it has been suggested that the services used to implement security management must, for the most part, be included in a trusted kernel, otherwise, the "security" provided is not reliable. However, the larger the security kernel is, the greater the problem of evaluating its trustedness becomes.

### 4.4.1. The Ability to Control Access to Resources

The ability to control access to resources is primarily an administrative concern of security management, (as contrasted with a detection or recovery concern), which involves permitting or disallowing access to security related parts of the network (i.e., security objects). The Security Management document, [SECURE], specifies that security objects are merely objects with associated attributes to be managed in the same way as other objects and attributes are managed. These security objects, as indicated in section 3.4.2, are all in the nature of data objects containing support information used in determining access rights to these or other objects. This information may take the form, for example, of passwords, routing tables, security codes and other information appropriate for "security credentials."

Since the Security Object and Attribute Management Function, defined above, is intended to manage security objects, it naturally is this function which provides the capability to control access to security resources. This function provides this capability by enabling the user of SM to create and delete objects, change their attributes or state, or affect the relationships between security objects. Access control to security resources might be particularly served by this function, for example, in regard to "updating access control lists, setting initial key values and seeds, and creating capabilities credentials" (i.e., who is permitted to do what with this object).

This function, moreover, is specified to perform its services by invoking the services of the Object, Relationship, or State Management functions. And, in true recursive fashion, these functions, in turn, rely on the basic CMIS services to ultimately provide the desired functionality.

Thus, although the Security Management document [SECURE] is still somewhat sparse, the basic mechanism appears to be presented therein to provide the ability to control access to resources, an activity which primarily entails maintenance and control of objects and their attributes. Further clarification and specification of the services, however, will be needed to provide sufficient guidance to the implementor.

### 4.4.2. The Ability to Archive and Retrieve Security Information

In addition to being able to control access to resources, it is necessary to keep account of activity, or attempted activity, with these security objects in order to detect and recover from attempted, or successful, security attacks. This section deals with these issues.

The ability to archive and retrieve security information involves providing the capabilities to create and delete security logs or audit trails, read from and write to these logs, start and suspend these logging or auditing activities, and monitor these audit trails or security logs to identify security violation activity and provide reporting and notification of these violations or attempted violations. It also involves the ability to directly monitor and report on security activity with, for example, the intent of reporting or notifying the appropriate SM user when suspect security activities are recognized. This last capability refers to current monitoring and notification rather than the archiving of this information for later evaluation. These activities address themselves more to the detection concerns of SM than to the administrative concerns discussed in the section above.

The Security Management document [SECURE] defines functions capable of providing the required capabilities to support the requirement to archive and retrieve security information. These functions include three of the four mentioned above (i.e., the Security Audit Trail Function, the Security Alarm Function, and the Security Alarm and Audit Trail Management Function). These functions enable the SM user to monitor security events and to store and retrieve historical information regarding these events.

These functions accomplish their desired results by invoking assistance from the Object Management Function, the State Management Function, the Relationship Management Function, the Management Service Control Function, and the Log Control Function. Furthermore, these latter functions rely on the basic CMIS services (such as CMIS M-EVENT-REPORT, M-GET, M-SET, M-CREATE, and M-DELETE) to ultimately provide the desired functionality.

As in the discussion in the previous section, although the Security Management document [SECURE] is still being formed, the basic mechanisms have been presented which can provide the ability to archive and retrieve security information, an activity which primarily entails monitoring security objects, setting thresholds, reporting events, and storing security information for later evaluation. Further clarification and specification of the services, however, will be needed to provide sufficient guidance to the implementor.

### 4.4.3. The Ability to Manage and Control the Encryption Process

The ability to encrypt and control the encryption process is another of the primarily administrative concerns of security management, which involves supporting the encryption activities of SM. The support for this capability offered in the Security Management document [SECURE] lies in the use of the Security Object and Attribute Management Function. This function enables the SM user to manage the security objects which can include, for example, encryption algorithm designations or encryption keys used in the encryption process. Distribution of these security objects (e.g., key distribution) is an important application for this function.

As we have seen above, this function is specified to perform its services by invoking the services of the Object, Relationship, or State Management functions, which, in turn, rely on the basic CMIS services to ultimately provide the desired functionality.

Furthermore, although as mentioned above, the Security Management document [SECURE] is still somewhat incomplete, the basic mechanisms exist to provide the ability to encrypt and control the encryption process, an activity which primarily entails distribution, maintenance and control of objects and their attributes, such as encryption keys. Further clarification and specification of the services, however, will ultimately be needed to provide sufficient guidance to the implementor.

## 4.5. Analysis of Performance Management Requirements

Performance management (PM) is an area that had, until recently, received very little attention from the standards community. However many of its requirements, as stated in section 3.5, are beginning to be addressed by ISO/IEC JTC 1/SC 21/WG 4. The main document produced so far by ISO regarding performance management is ISO/IEC JTC 1/SC 21 N 3313, the Performance Management Working Document [PMWD]. Perhaps the most important part of that document is the proposed Workload Monitoring Function. Other important parts include:

- monitoring and tuning models,
- a list of potential functions to be developed by PM,
- a discussion of management relevant to PM,
- performance management requirements,
- performance management outstanding issues, and
- proposed metrics.

The PMWD does not and should not include definitions of performance-related resources (managed objects). These definitions, which belong in the Definitions of Support Objects [DSO] or in appropriate layer management standards, should include information and actions needed to support performance management. This may include units of measure and format of data to be recorded. Such definitions must be developed by the layer groups defining management for layer resources. For each non-OSI resource for which

performance management is desired, some group yet to be determined must define the related information and actions to be support by the resource.

### 4.5.1. The Ability to Monitor Performance

The ability to monitor performance will be provided primarily by the PM Workload Monitoring Function (and possibly other monitoring functions - see below) for the generation of event reports. This function allows a management system to monitor workload (i.e., resource utilization) and generate event reports as utilization approaches capacity. It also includes a model for overload conditions (i.e., service request denial due to lack of capacity). Although the Workload Monitoring Function is not yet complete, it is expected to reach DP at the November 1989 meeting of ISO/IEC JTC 1/SC 21.

The ability to support performance monitoring on a polling (non-event report) basis is provided by the Object Management Function [OMF], which allows attribute values to be retrieved.

There are no current plans to develop standards that enforce statistical measures to be based upon a minimally accepted number of samples. However, implementors or operators can determine the proper number of samples either on a theoretical basis or by experience. There is no requirement that the sampling parameters (e.g., time interval between samples) be standardized. Well-tuned performance managers will use the minimal number of samples to infer agent system performance.

The PM Working Draft includes a discussion of several proposed monitoring functions which are still extremely formative. As they are further developed, they could be applied to a variety of performance monitoring problems. However, it is still early to predict exactly what services they will provide. These proposed functions are:

- o Throughput Monitoring Function,
- o Response Time Monitoring Function, and
- o Queue Monitoring Function.

While the emerging performance management standards are still in a formative stage, it appears that they will meet the requirements stated in section 3.5.2.1. These include, for example, the ability to select the events, resources, or measures to be monitored, the ability to specify measures or resources to be polled and recorded, and the ability to specify the threshold level used to trigger the notification of a performance abnormality.

### 4.5.2. The Ability to Tune and Control Performance

The ability to execute performance tests and to collect the results from those tests will be provided by the Confidence and Diagnostic Testing Function which is under development and included in the Fault Management Working Document [FMWD]. The analysis of test results is beyond the scope of standardization and will be implementation specific. However, a model for

performance monitoring and tuning, still under development, is included in the PM Working Document [PMWD]. The model outlines an iterative method for performance monitoring and tuning and suggests a possible decomposition of functionality for performance monitoring and tuning. The PM Working Document also cites a proposed Performance Tuning Function, but the description is not clear as to what services the function would provide. It is likely that tuning will always involve a combination of standardized services and non-standardized analysis.

The PM Working Document also includes proposed Statistical Analysis Functions, which will allow standard metrics to be computed on agent systems, and then sent to manager systems. However, this work and related work on metrics is still formative.

The ability to change resource allocations and to set or modify resource (managed object) attribute values is provided by the Object Management Function [OMF], developed for configuration management.

### 4.5.3. The Ability to Evaluate Performance Tuning

The ability to evaluate performance tuning (a requirement stated in sec. 3.5.2.3) is provided by the use of the Object Management Function [OMF] as well as by non-standardized analysis methods, possibly supported by the proposed Statistical Analysis Functions and metrics included in the PM Working Document.

### 4.5.4. The Ability to Report on Performance Monitoring, Tuning, and Tracking

Notification of abnormal performance changes (a capability whose requirement was stated in sec. 3.5.2.4) can be provided through the use of the Performance Workload Monitoring Function and possibly by the other monitoring functions included in the PM Working Document. (These are cited above in the introduction to sec. 4.5.)

The concepts of management domains is discussed in the Systems Management Overview [SMO]. However, the specification of such domains is not complete. Methods for creating, deleting, and modifying domains or subdomains within a network still must be developed. Domain specification and maintenance are problems clearly not unique to performance management and must be solved by a domain model that includes consideration of requirements of all management functional areas, i.e., security, configuration, fault, accounting, as well as performance.

Generation of performance reports based on user specified criteria can be accomplished through use of appropriate parameters applied to the functions under development by performance management. The storage and manipulation of trending reports for benchmarks is beyond the scope of standardization. This requirement could be satisfied by incorporating a data base management system (DBMS) into OSI management. Such incorporation in the near future is unlikely. It might be possible to define a standard application program interface (API) for management, to allow independent DBMS vendors to "hook"

into management. (See sec. 5.19 for a discussion on application program interfaces.)

To provide snapshots of network performance in terms of user specified measures, a set of potential measures must be assembled, so that implementors can develop software to support them. The task of assembling such potential measures probably belongs to layer groups who must specify managed objects with the appropriate capabilities to produce such measures. It is not clear that such an effort is currently being pursued.

Statistics, which can be applied to the measures, will include average, maximum, minimum, and others. (The list will expand as the standards are developed.) In addition, the concepts of rates, usually event occurrences divided by time units, should be supported, but is now only formative in the PM Working Document [PMWD]. Support for statistics, also known as metrics, are under development within SC 21/WG 4. Discussions of such concepts are included in the PMWD. Related concepts such as composite or summary information about managed objects is also included in the PM Working Document. While the development of metrics within the PM standards effort is still formative, the effort appears to be accelerating. It is likely that within a few years sufficient progress will be made to meet users' requirements in this area.

### 4.5.5. The Ability to Test Capacity and Special Conditions

The ability to test resource capacity will be provided by services included in the Confidence and Diagnostic Testing Function (CDTF) currently under development by SC 21/WG 4 and discussed in the Fault Management Working Document [FMWD]. (The services under development for testing are listed in sec. 4.3.2.)

Additional work will be required to define PM methods for establishing modes for network quiescence to allow testing under controlled conditions. It must be possible to establish a network "testing mode" that would alter normal network procedures and perhaps priorities to ensure accurate test results. This and other testing matters should be undertaken in concert with Fault Management's development effort on the CDTF. It is important that PM's requirements for testing be included in the Confidence and Diagnostic Function.

## 4.6. Analysis of Accounting Management Requirements

Accounting management (AM) is an area that had, until recently, received very little attention from the standards community. However many of its requirements, as stated in section 3.6 and in section 5.9 under "Tariff Line Management," are beginning to be addressed by ISO/IEC JTC 1/SC 21/WG 4. The main document produced so far by ISO regarding accounting management is ISO/IEC JTC 1/SC 21 N 3314, the Accounting Management Working Document [AMWD]. Important parts of that document include:

- o scope and model of accounting activity,
- o an OSI management function for communication instance accounting,

- o an OSI management function for accessing accounting logs, and
- o a list of accounting management outstanding issues.

### 4.6.1. The Ability to Record and Generate Accounting Information

The Accounting Management Working Document [AMWD] describes some types of information that may be gathered about the usage of resources in the OSI environment. Information that may be collected includes, for example, duration of communications resource usage, number of service data units used, quality of service provided, and reason for communication termination. The definition is too heavily linked to accounting for communications resources. While communication resources are a great concern of AM, there may be other resources within an open system that must be accounted for as well. For example, the use of computing resources, or special equipment such as high speed laser printers, might be subject to accounting, although they may not always be directly related to the support of OSI communications.

The AMWD does not and should not include definitions of accountable resources (managed objects). These definitions, which belong in the Definitions of Support Objects [DSO] or in appropriate layer management standards, should include information and actions needed to support accounting. This may include units of measure and format of data to be recorded. Such definitions must be developed by the layer groups defining management for layer resources. For each non-OSI resource for which accounting is desired, some group yet to be determined must define the related accounting information and actions to be support by the resource.

The Communications Instance Account (CIA) function under development and described in the AM Working Document includes the ability to account for each instance of communications. The CIA does not allow for specification of accounting data collection periods nor criteria to support such specification. However, such capability could easily be included in each implementation without requiring additional standardization.

### 4.6.2. The Ability to Specify Accounting Information to be Collected

The information to be collected is dependant on what the definition of the managed object permits, those options supported by implementors in the case of optional attributes, and by the attributes that the operator chooses to be collected. The management standards make no statement as to the format in which information is to be stored on open systems. The linkage of data base management systems (DBMSs) to accounting management systems is not now subject to standardization and is thus dependant on the implementors' willingness to provide such a linkage. If a standard application program interface were to be specified for the management system (see sec. 5.19), users or third party vendors could provide such a linkage. However it is provided, linkage between systems management and a DBMS would go a long way towards meeting the requirements for storage of AM records to provide periodic weekly, monthly, or quarterly accounting reports.

### 4.6.3. The Ability to Control Storage of and Access to Accounting Information

The Communications Instance Accounting function under development will allow for standard procedures to retrieve accounting information. Disposition of the information after retrieval is not subject to standardization.

Security (e.g., confidentially) controls to be applied to such information must be considered by security management in general and the definers of the resource to which accounting is to be applied, e.g., the layer group defining the attributes associated with each managed object.

### 4.6.4. The Ability to Report Accounting Information

The emerging standards for OSI systems management do not specify in what format information is to be presented to users, operators, or administrators. Each implementation is free to present the information in whatever format the implementor desires. However, to support interoperability, the emerging standards will specify representations of management information that is to be exchanged among open systems.

There is currently no support in OSI systems management for the broadcasting or multicasting of news such as network resource billing rate changes or accounting limits. (See sec. 5.4.)

### 4.6.5. The Ability to Set and Modify Accounting Limits

While the definition is still incomplete, the Communications Instance Accounting function will support the ability to read, set, and change accounting limits for communications resources. (This is a requirement stated in sec. 3.6.2.5.)

Mechanisms for changing priorities assigned to network users for access to resources can be accomplished through the use of the Object Management Function [OMF], developed by configuration management.

### 4.6.6. The Ability to Define Accounting Metrics

Standard accounting metrics will be defined by the layer groups who define managed objects which represent accountable resources. Such definition is supported by related concepts defined in the Accounting Management Working Document [AMWD].

## 4.7. Analysis of Other Requirements

Section 3.7 cited several functional requirements for network management that are not needed for interoperability and are therefore beyond both the scope of current standardization efforts and the scope of this paper. (See sec. 5.18 on "Scope of Standardization.")

However, it should be noted here that several of these cited requirements are discussed within this paper. Additional comments on the requirements for

a man-machine interface are discussed in section 5.12. Except in section 3.7, the requirement for "help text" is not discussed elsewhere in this paper but is part of the man-machine interface issue. Data analysis requirements are discussed in section 5.7, "Management Assistance in Off-line Tasks." Finally, the last requirement cited in section 3.7, the potential need for "expert systems" in network management, is discussed in section 5.1 and in section 6.

## 5. ADDITIONAL ISSUES

This section contains discussions of various issues that have been identified as important network management functional requirements, but which either 1) transcend the categories of section 4, or 2) require additional discussion because of their importance or in order to determine if they are worthy of further consideration. In addition, issues are included here that should be considered for future standardization after the current set of management standards stabilizes (e.g., expert systems in network management).

### 5.1. Application of Expert Systems to Network Management

The use of expert systems, an artificial intelligence technique, to aid in the management of networks appears to have great potential. In situations where complex decisions must be made quickly and reliably, expert systems seem particularly useful. Systems that require extremely high reliability (e.g., space stations) or that cannot depend on highly trained operators (e.g., some commercial and military networks) are also candidates for the use of expert systems.

While there may be enormous potential for the application of expert systems to network management, there have been few, if any, detailed studies upon which to base any development work at this time. Some vendors have recently announced and/or delivered products that include expert systems technology. The application of such techniques will require considerable experience with real networks and protocols to refine useful rules, parameters, and heuristics for management. Thus, while it may be premature to apply expert systems to all network management problems, it is time to consider in earnest how such techniques might be enhanced by studying of real networks to refine baseline rules, parameters, and other related information.

For further discussion on the application of expert systems to network management, see section 6, "Automated Network Management Systems."

### 5.2. Management Information Base (MIB) Design

As specified in the OSI Management Framework (ISO 7498-4), the Management Information Base (MIB) is "that information within an open system which may be transferred or affected through the use of OSI Management protocols." The Management Framework further describes the MIB as the set of managed objects within an open system. Its design is considered to be an implementation issue not subject to standardization. However, MIB organization is crucial to the operation of an efficient, robust OSI systems management implementation, and its design is extremely important. For example, efficiency considerations suggest that the MIB be organized such that the most often needed information be most rapidly accessible. Although, since the MIB is a conceptual schema, what appears to be most rapidly accessible may require considerable overhead to access. The important point is that MIB design and implementation be efficient.

Robustness considerations dictate that default values of critical communications parameters (e.g., transport inactivity timer) be maintained on each managed system in case communication with a manager system is lost. Such critical defaults must also be available on the managed systems before communication with managers is established.

The MIB should, theoretically, give a homogeneous view of an otherwise heterogeneous data base. It is important that all implementations share a common understanding of MIB contents including, for example, access characteristics and operations on managed objects. All managers should "see" an agent's MIB in a clear, consistent manner. This relates to SMI issues, including, for example, the Management Information Model [MIM], and the Definitions of Management Attributes [DMA]. These help to provide common definitions of object types with standardized semantics.

It would be useful to allow generalized data base operations to be permitted on elements of management information thus providing, for example, selective retrieval. Such operations would prove useful in a hierarchical management environment where managing systems could selectively retrieve information from managed systems. A DBMS could be used to support such activities and serve to archive structured information relating to network performance or faults, for example. Smaller managed systems need not support DBMS activities. Requirements for a DBMS are currently beyond the scope of OSI standardization, but should be given future consideration.

### 5.3. Network Management Efficiency Issues

A prime requirement of a successful network management system is that it should not seriously degrade the network operation it is intended to maintain and that its decisions be made and its control actions taken quickly before network conditions significantly change (otherwise, instability may arise in the measurement/control-action feedback loop). The network management system must provide needed services while consuming as few network resources as possible. This implies that both network bandwidth and computing resources be conserved. For example, network management personnel would probably be upset to learn that traffic in support of management consumes over 30 percent of available bandwidth or computing cycles. Yet, if not properly designed and maintained, such an arbitrarily chosen figure could easily result or be exceeded.

One large commercial network operator has stated that management operations consume about 2 percent of the available network bandwidth. On this network, accounting data are separately transmitted and use about another 2 percent of bandwidth. A designer for a large consulting firm stated a design goal that network management use between 1 and 5 percent of bandwidth. A figure of 5 percent has been suggested by others as the absolute maximum allowable bandwidth consumption by management operations. Therefore, it can be concluded that relatively low use of network bandwidth for management operations is a realistic goal in OSI-based systems, although there may be specialized networks with special requirements (e.g., extremely high reliability) where additional management overhead can be tolerated.

Network managers and operators may want to isolate and measure the effect of network management operations on the network. Then, perhaps, rational decisions can be made as to when there is too much network management. The results of such measurements can be used in conjunction with other measurements, such as network utilization profiles, to enable optimal scheduling of certain resource-consuming management operations at times of low network utilization, minimizing management activity at times of peak network use.

To help in assuring that network management operations do not introduce significant detrimental effects on other network operations, tools for measuring these effects should be provided. Without proper tools, these effects might not easily be detected.

In the interim, as standards are being developed, simulation modeling can be used to help determine the effect of network management on efficiency. Modeling should be accurate enough to provide a reasonable sense of the impact of network management on efficiency without the expense of first building a network, and then determining this impact. Furthermore, simulation is a useful way for network standards designers and network implementors to determine optimal values for network parameters under controlled conditions that are difficult to duplicate on real networks.

One technique to reduce management consumption of network bandwidth is to employ a hierarchy of managers as discussed in section 5.5 below, entitled "Multiple Manager Considerations." (This technique is sometimes used by vendors in proprietary network management schemes.) Using local managers to manage small LAN segments which then convey summary reports to more global managers reduces the amount of overall traffic on a network, since the most commonly sent management messages are confined to those network segments where they are needed. A global manager could likewise control local systems by passing parameter setting commands to local managers to be forwarded to one or more local systems.

A specific example of an efficiency related issue arises when considering the difficulties of designing network management standards where the data must be represented in an unambiguous, machine-independent manner. Abstract Syntax Notation One (ASN.1) has been developed for this and other purposes. While data encoded using ASN.1 are machine independent, they are usually not as compact as machine-dependent representations, thus requiring extra bandwidth to be sent from one OSI system to another. Moreover, each item of data must be encoded on one side of a management association and decoded on the other. This requires machine cycles that might otherwise be applied to solving user problems. Network management implementors must develop fast, efficient methods for encoding and decoding ASN.1 data. (In proprietary network management schemes where communicating systems are of the same vendor design, the problem of machine-independent data representation usually does not exist -- data are transferred using the machines' internal representations.)

The current OSI systems management architecture uses Common Management Information Services and Protocol (CMIS/CMIP) as its primary method for

transferring management information between open systems. A major concern about using CMIS/CMIP is its efficiency in terms of encoded representations. To access remote data, for example, general managed object identifiers (both class and instance) must be provided. This generality may require the use of very long strings to access even small counters (e.g., 16 bit counters). If proper management dictates interrogating many counters frequently, much of a network's bandwidth (and possibly implementation-dependent buffer capability) may be consumed by these large object identifiers. Techniques could be developed for establishing contexts of objects such that full object identifiers need not be conveyed on each CMIS/CMIP operation.

One method that could be employed is to exploit the scoping and filtering rules of CMIS. By proper design and implementation of managed object classes and instances, "base" managed objects, selected as the subject of managed operations, could effectively be used to imply the selection of large sets of related objects in a single CMIP service data unit. Thus the impact on network bandwidth, and probably other OSI system resources as well, would be significantly reduced.

A possible technique, suggested in the Management Framework, but rarely discussed, that might reduce management resource usage, is to allow Layer Management Entities to convey management information directly between peers without employing the Application layer for transfer of information. This approach, while not nearly as general as CMIS/CMIP, might be employed on local network segments to reduce overhead. Standards developers, however, are, in general, very much opposed to such an approach since such techniques inhibit full management interoperability among open systems (e.g., Link layer management protocols can not be operated across Network layer routers) and might require layer management to replicate some higher layer functionality. While this is admittedly true, the use of such techniques may need to be considered not only for efficiency but also, for example to enhance management robustness. For example, if upper layer communications were inhibited, systems management would probably cease to function, but management at a lower layer might be able to resolve the problem without sending out personnel to examine and service the equipment. (See sec. 5.16.)

## 5.4. Connection Mode and Peer Mode Issues

In many communications systems, data transfer is either connection-oriented or connectionless. There are advantages and disadvantages of each connection mode, depending on the situation. The issue of a connection-oriented versus connectionless approach to network management is an important and controversial one. (It should be pointed out that the connection mode for management communications discussed here is at the Application layer. Perhaps, "association-oriented" might be a preferred term. Although some underlying layers may, and often do, use a different approach, the issue under consideration here is the orientation of the management protocol used in the Application layer.) Because of the history of their development, the OSI services and protocols tend to favor connection-oriented solutions with single peers at each end of the connection. For many problems, the connection-

oriented solution, especially on wide area networks (WANs) is the obvious choice.

The situation on local area networks (LANs) is often quite different, however. LANs tend to use a connectionless (often multi-peer, broadcast or multicast) mode of communication, and LAN vendors often use a multi-peer mode of communication for network management purposes. Unlike WAN technology, LAN communication costs are low and bandwidths high. For these and other reasons, the underlying services (at the Physical and Link layers) are most often connectionless and multi-peer on LANs.

The service and protocol developed by ISO for OSI management, CMIS/CMIP, is single-peer, connection-oriented. (Developers of CMIS/CMIP have shown little interest in either multi-peer or connectionless communications.) In order to transfer management information between open systems using CMIS/CMIP, single-peer connections (associations) must be made. If only occasional pieces of information need to be transferred, maintaining such connections could be wasteful of resources. In the case of multi-peer broadcast mode LANs, the natural topology suggests that connectionless services and protocols are more efficient. In addition, LANs have such extremely low error rates that the error checking capabilities of connection-oriented transport protocols (used by CMIP), can reasonably be dispensed with for non-critical reporting operations. (Note: Although there is a connectless mode transport protocol, current CMIS/CMIP does not have a mechanism to employ it.)

Furthermore, on LANs it is often the case that management information is to be broadcast to all stations on the network for such conditions as announcements of network outages, reconfiguration, or requests for management services. The use of a single-peer connection-oriented management protocol is questionable in these situations and the use of multi-peer mode layer management services and protocols might be efficient in the LAN environment.

There is an effort within ISO/IEC JTC 1/SC 21/WG 6 to provide for connectionless services for upper layers, including Association Control Service Element (ACSE), Session, and Presentation considerations. However, this work, so far, has not embraced multi-peer mode transmission needed for management. Working papers on multi-peer transmission, rumored to be available, are not yet sufficiently mature to be used as a basis for Application layer systems management.

However, to facilitate multi-peer transmission, it might be possible for management to utilize network configuration information available at the Network layer for systems management. As part of Network layer End System/Intermediate System (ES/IS) operations, a data base of information about the configuration of open systems (end systems primarily and sometimes, intermediate systems) is maintained. This information is normally obtained using the multi-peer broadcast/multicast protocols available at lower layers. Whether system management can utilize this information for domain establishment and management purposes needs further consideration. Also, to be considered is whether additional standards are required to support this concept of Network/Application management information sharing or whether it is simply an implementation issue.

## 5.5. Multiple Manager Considerations

In the simplest case where network management protocols and standards need to be applied, there is a single, manager system managing one or several agent systems. For relatively small networks, a single manager may oversee perhaps up to several hundred agent systems. Beyond that, it is somewhat infeasible to have a single manager system. (Simulation modeling or actual implementation experience will indicate more accurately what is a realistic bound.)

Even without management considerations, large networks are often decomposed into smaller networks for routing and bandwidth considerations. Large networks may be composed of segments using different underlying technologies that are interconnected by gateways. Contention based (e.g., CSMA/CD), token bus, and token ring LAN's may be interconnected. Multiple LAN's in diverse locations may be interconnected through X.25 gateways. It would be difficult to conceive of a single network manager controlling such a large network by direct X.25 connections to each system. Cost considerations alone would make such a solution infeasible. It is much more plausible to require that each local LAN or LAN segment contain a manager which conveys summary information to some central manager using X.25 connections.

Therefore, it is inevitable that real OSI-based networks contain multiple managers in some form. Multiple managers allow for a local manager on a "subnetwork" to control the systems on that subnet and relay information to a higher level manager, forming a hierarchy of manager systems. (This is, in fact, currently done on many vendor-proprietary management architectures.) Such hierarchical management reduces the path length that management information must travel and thus minimizes the reduction of available network bandwidth while decreasing the probability that an agent system will be isolated as a result of a perhaps distant network failure. It also provides for off-loading from a single manager system that simply may not have the computing power to manage more than a few hundred systems directly.

Size, however, does not provide the only reason for considering multiple managers. Even in relatively small, local LAN segments, multiple managers may be employed as a strategy to enhance the robustness of the system. Redundant managers can provide backup in case the primary manager fails or is undergoing maintenance. This can be an important consideration in a real time or near real time LAN.

Multiple manager concepts may make the development of some management systems less complex, because a single manager system may not include the software (and sometimes, hardware) required to handle all possible variations of equipment in a network. A manager of a CSMA/CD LAN need not implement functions special to a token ring network and vice versa.

Finally, yet another reason for multiple managers is to provide specialized managers which are concerned with only certain objectives. For

example, in a highly secure network, a security manager may be desirable. Such a manager might control, for example, passwords, access control lists, and association and connection establishments. This "security" manager would be dedicated to security issues and would not be primarily concerned with configuration, fault, and performance considerations.

For these reasons, it is important that future ISO work include developing standards that allow for the interconnection of multiple managers on a single or interconnected set of networks. The current ISO Systems Management Overview [SMO] includes the concept of management domains and allows for multiple "managing processes" and "agent processes" on a single network, or even on a single open system. This concept could be incorporated into a management model that supports multiple managers on a single OSI-based network.

For example, further refinement is needed on managed objects to represent the management view of a domain. Mechanisms are needed to establish domains in a uniform, well-understood, standard manner.

To support communications among multiple managers, special manager-to-manager protocols may be required, although it is plausible that CMIS/CMIP could provide the required functionality. Further study in this area is needed. Questions include: Do CMIS and CMIP provide sufficient functionality to provide data transfer services for multiple manager operations? And, which standards group(s) should be developing models and managed objects to help coordinate management domains?

### 5.6. Physical Device Management

The management of physical devices (e.g., switches and modems) used for communications must be addressed by OSI management. Although their proper operation is essential to maintaining communications, these devices usually cannot be addressed as OSI systems by management protocols.

The most obvious approach to managing them is to consider them to be resources of the open system to which they are attached. Then they can be addressed as resources (managed objects) as part of that system. This approach seems quite feasible in most cases.

Problems arise, however, when the physical devices are not directly attached to an open system but, rather, to an active network component that is less than a full seven layer system, such as a bridge, router, or gateway. Since these so called "thin stack" components need special consideration in these respects as well, they are discussed in more detail in section 5.11

### 5.7. Management Assistance in Off-line Tasks

Many of the tasks associated with the management of computer networks are primarily manual or off-line in nature. (By off-line, we mean those tasks do not constitute a part of normal, real time network operations; rather, they

may be performed in a batch environment.) These tasks include such things as planning or installation. The emerging management standards address the operational phase of network management. This phase must be automated, since monitoring and control of networks must be done in real time.

The data gathered as part of network operations may be applied to other tasks as well, however. For example, performance related data can be analyzed off-line to suggest possible new network configurations. Trend analysis can be applied to operational data for the purpose of capacity planning. The data can also be used as input to simulation models to gain further understanding of operational networks and project the effects specific changes may have on them.

In order for any of these "off-line" tasks to be accomplished in a straightforward manner, the operational data must be accessible on-line to reduce the time and effort required to make these data available for analysis. This serves to reduce any possible errors due to manual processing (e.g., operator retyping of data for analysis), thereby making the whole process more reliable. Provision for "capture" of on-line operational data should be included in network management implementations. While it would be most useful for the format of these data to be standardized so that any number of analysis packages could be used, such standardization is not an initial requirement, although the format of any such data must be well documented. In any case, data format for management data is beyond the scope of current standardization efforts.

### 5.8. Multilayer Considerations

Responsibility for the development of an architecture for the management of OSI systems has rested with ASC X3T5.4, here in the U.S., and with ISO/IEC JTC 1/SC 21/WG 4, internationally. The results have included the OSI Management Framework (ISO 7498-4), OSI Basic Reference Model Part 4, Management Information Service and Protocol Draft International Standards, Functional Areas documents, and Draft Proposed International Standards for management functions as well as for the Structure of Management Information. While most of these documents have not yet been completed, they all pertain to the management of OSI resources contained in and defined by the various layers.

Although overall architectural development for OSI management standards has been primarily centralized, the nature of management standards development requires that the standards making bodies for each layer define the resources (i.e., the managed object classes) to be managed within that layer, since those standards making bodies are most familiar with the layer. Not only must the resources to be managed be defined, but also the set of allowed operations on those layer resources must be defined. For example, ASC X3S3.3 in the U.S. and ISO/IEC JTC 1/SC 6/WG 2 internationally are the standards making groups responsible for the Network and Transport layers. These groups are developing papers describing resources to be managed within these layers and the operations that can be performed upon them.

Therefore, while overall architectural issues have been considered by X3T5.4 and layer management issues have been, or soon will be, considered by the various layer groups, it is not clear who has responsibility for coordination among layers to control any possible detrimental effects that management of one layer may have on other layers. The net effect of changing parameters for the optimization of one layer may, in some cases, cause overall negative effects.

For example, consider a situation in which, in an attempt to increase throughput at the Transport Layer, the maximum TPDU size is increased from 1024 octets to 2048 for connections to be maintained over a large concatenated network where the maximum packet size that can be transmitted end-to-end is 1500 octets. The effect of the increase in TPDU size is to cause packet fragmentation and reassembly at the lower layers, which, in turn leads to a decrease in throughput. In addition, the extra packets introduce increased processing overhead at each node in the Transport connection path.

Another related potential problem is synchronization of changes across layers. For example, consider the adjustment of maximum PDU sizes. In the event that the above limit of 1500 octets is not a physical constraint but rather the maximum NPDU size, and management recognizes that larger PDU sizes at both Network and Transport can be used, it is important that the changes to each layer be synchronized, otherwise unnecessary fragmentation or other inefficiencies may result.

Also related is fault determination and recovery across layers. If a component at one layer fails, it may trigger fault notifications at other layers, or possibly, multiple notifications in a single layer. For example, if a modem fails, a Physical layer alarm might be generated. At the same time, multiple Transport layer alarms might also be generated for each transport connection using the modem. There is a single root problem of the fault. However, multiple alarms could cause multiple recovery actions to be attempted, all essentially useless unless the root problem is corrected or alternate Physical layer routing can be provided. Multiple recovery actions can lead to "thrashing" as each layer tries to utilize new resources to resolve the problem.

## 5.9. Tariff Line Management

Tariff line management is primarily a decision-making process that requires detailed cost information. As such, tariff line management is outside the scope of OSI management standardization. However, it is partially dependent upon cost information supplied by OSI accounting management.

Since the task of specifying standards for accounting management has only recently begun in earnest by ISO, it is important that tariff line management be considered in specifying standards for accounting management. For example, it should be possible, directly or indirectly through the use of AM services, for an individual user of network communication facilities to determine the cost (in dollars or other national currency) incurred for the use of those facilities. It should not be necessary to obtain such information in a manual

or off-line manner. Furthermore, it should be possible to use AM in any billing services to provide sufficient information about communications usage charges to reconcile the billed costs with projected tariff costs.

(Note: Since different organizations, including both OSI service users and providers (some of whom are potentially common carriers), may want to implement different policies for providing charging information, not all aspects of accounting management need be standardized. However, the standards, and especially models for accounting, should be sufficiently flexible to allow for some OSI service providers to accommodate on-line charge determination to users while others may not implement such capability.)

### 5.10. The Directory (Formerly Directory Services)

It is the purpose of this study to discuss functional requirements for network management. The scope and maturity of documents available leaves it unclear as to the exact nature of the relationship between The Directory and network management. The Directory can be treated as another Application layer entity (or set of entities) within the OSI model. From this perspective, there does not appear to be any unique requirements by The Directory of network management. The Directory benefits indirectly from the specific management services, just as any Application layer entity might. There may be unique requirements imposed on the Directory by Security Management, but these are issues for future study.

### 5.11. Management of Bridges, Routers, and Gateways

Management of devices (systems) that contain some, but not all, of the seven-layer OSI functionality is a well-known problem within the OSI management standards development community. The Management Framework (ISO 7498-4) [FRMWK] is concerned primarily with "Systems Management", the management of full seven-layer systems using Application layer protocols. The use of (N)-layer management is generally discouraged, when systems management can be used. However, the framework recognizes that such functionality does not always exist. It cites the problems involved in the management of a "Relay" (Network layer router) and "Broken" Open Systems, where there exists only a "Minimum Communication Capability". Since these devices (referred to as "thin stack" machines) are less than full seven-layer systems, they can have no management application processes through which to communicate. Furthermore, there is a question as to how to address these devices, since they are not OSI end systems.

In the past, solutions to these problems have been slowed not so much because technical solutions are that difficult, but because devising solutions in conformance with the letter and spirit of previous drafts of the Management Framework (and the rest of the seven-layer model) were not obvious. By discussing the existence of less than seven-layer "OSI" systems, the most recent drafts (and final International Standard text for 7498-4) take a more realistic view of the problem of managing these systems than did previous drafts which ignored the problem.

We can anticipate that in the future, new technology and "smart" implementations will include more complete OSI functionality at little to no extra cost. However, there will still be a need to manage older installed equipment that may have many more years of useful service life.

Developers of non-OSI-based networks (e.g., proprietary and DoD TCP/IP) have developed solutions to handle these less than seven-layer devices. They were able to do so since they employed management models that more adequately recognized and dealt with less than full seven-layer functionality. Within the OSI community, both the IEEE and the MAP/TOP Network Management group have developed alternatives to and extensions to the OSI management model that facilitate management of devices such as bridges, gateways, and routers. The standards development community within ISO should attempt to integrate these and other concepts for (N)-layer management into OSI Systems Management, now that the need for the management of less than full seven-layer systems has been formally recognized.

### 5.12. Man-Machine (User) Interface Considerations

There are many man-machine (user) interface requirements to be considered for network management. For example, operators of network management control centers usually require consistent, menu driven, human engineered interfaces. Full examination of requirements for man-machine interfaces (MMI) is an extensive task beyond the scope of this document for two reasons.

The first is that the human user interface to OSI management systems is not currently subject to standardization. The ISO currently is developing standards that allow for interoperability among open systems and standards for a user (human) interface are not needed to achieve that goal.

The second reason is that examination of user interface considerations require skills different from that of the authors of this report. Many of these considerations are based on psychological criteria which we are ill prepared to discuss. Rather than attempt a superficial treatment of these issues, and since it is already beyond the scope of OSI standardization, we prefer to leave it to others to examine these requirements.

However, before totally dismissing the issue, we note that in our reading and discussions with users and potential users of OSI systems, there appears to be a large demand for a common human (user) interface to network management systems. This interface must be vendor independent. For example, command syntax must not vary. If faults are indicated by blue flashing squares on the operator console supplied by vendor A, then they must also be represented that way on consoles supplied by vendor B.

Users require a standard interface for several reasons. They do not wish to retrain operators each time a new network management system is delivered. They want the system to be reliable and not confuse the operator. And, they want the ability to allow operator training to be standardized. Since the needs of each user organization may vary with respect to man-machine interface

requirements, we understand MMI standardization to be a user organization concern rather than an international standards issue.

### 5.13. Templates of Norms/Baseline Values

A very common technique used in network management is to compare "normal" values, or ranges of values with current values. Significant differences may indicate a problem or possibly an area of particular interest. Many vendors exploit this "difference from the norm" technique by providing templates of norms or baseline values to be used as comparison filters. This method is used so often that it has become a defacto requirement for network management.

Initial baseline values can be established by off-line management methods. Such techniques may include analytic methods, simulation models, or previous experience with the networks employing similar communications technology.

Once the initial baseline values (or ranges of values) are established, they are stored as a template for later use in comparison with current values. Often these initial values are inappropriate, causing too many significant events or problem areas to be reported. Changes to network topologies, usage patterns, and new applications will generally cause the initial baseline values to be inappropriate after an extended time period.

This type of situation requires that baseline values be updated. Network management, therefore, must provide methods for maintenance of and adjustments to these value templates. Such methods may include techniques borrowed from data base management technology. For efficiency and robustness, the templates must usually be stored locally on each system to be managed, thereby creating what is, in essence, a distributed data base management problem.

Templates of values appear most often at the man-machine and process control interfaces (See fig. 3). The process control interface occurs only on manager (not agent) systems, and is conceptually between a process that "talks" to the operator (the "MGMT PROCESS" in fig. 3) and the process that communicates with peers on agent systems. Since these interfaces are not necessarily standardized, it is not necessary for interoperability that the use of templates be reflected in standards. However, efficiency concerns dictate that templates be considered in Management Information Base (MIB) design since unnecessary mapping of values (and their interrelationships) back and forth across open systems is wasteful of both communications bandwidth and machine cycles.

Baseline values are potentially a very useful source of information for capacity planning. As these values are adjusted, long term trends can be observed more easily than by monitoring current values of managed objects. With even simple trend analysis and a knowledge of the underlying capacity of network resources, the point at which demand will exceed capacity can often be easily determined.

Another possible use of the baseline values is for future implementations of expert systems to network management. The baseline values stored in the network management templates may serve as initial values for such systems.

## 5.14. Extensibility

Computer and communications technology is constantly changing. New developments, such as high density RAM memories, small high capacity Winchester disks, and fiber optics, are changing the methods by which we solve problems. Yesterday's optimal solution is not always viable and is generally not optimal in today's world.

Yet for all these changes, we often use less than optimal solutions, and for a very practical reason. We have software developed for architectures introduced 5, 10, or 20 years ago. Even though newer architectures offer vast improvements in performance, the cost of conversion is too high. We become locked into outdated solutions. Sometimes, even vendors would be happy to move onto more advanced architectures, but they too must support old solutions since their customers are reluctant to convert.

Therefore, it is important when introducing new software systems to allow flexibility. Methods introduced today may be around for decades. Upward compatibility with as yet undeveloped technologies must be anticipated. Future technologies must be accommodated within the constraints of today's emerging standards. In the jargon of the computer industry, "the hooks must be provided" for future extensions.

Today's emerging OSI management standards have been designed with flexibility in mind. Resource identifiers are generalized. ASN.1 encoding rules allow for an almost limitless possibility for future data structures. There are few fixed field sizes to limit the usefulness of the management protocols.

Over a year ago changes to the ISO working papers reflected the fact that the ISO has recognized that OSI management techniques can be (and probably must be) applied to non-OSI resources. These resources may be, in fact, not related to OSI communications at all, but simply remotely managed objects.

OSI management standards must be extendable so as to be able to manage objects for which the standards were not originally designed. They should be able to perform actions on objects (or, perhaps more appropriately direct remote systems to perform actions on objects) that were not foreseen when the standards were originally produced. They should work efficiently, in terms of bandwidth and computing overhead, over underlying network technologies not yet conceived.

However, flexibility can have a price. Efficiency may have been sacrificed for the sake of providing generality in the emerging management standards. Each time a datum is to be transferred from one open system to another, it must be ASN.1 encoded, transmitted, and decoded. This requires

extra processing time and communications bandwidth compared with sending it unencoded.

Even though they appear to provide sufficient flexibility, it is, of course, impossible to predict if the emerging standards can meet these constraints for the next 10 to 20 years. However, unless industry wants to redesign and retrofit standards, as well as implementations, it is imperative that the standards being developed today allow for extensibility.

Perhaps the area in which extensibility is most important is that involved with defining resources (managed objects) that are not standardized by any recognized standards organization, but whose management is necessary for proper network operation. Management of such resources is of almost immediate importance, even without the introduction of new technologies or methods.

Consider, for example, the management of buffer space within an OSI implementation. Buffers are an implementation-dependent resource. Yet improper allocation of buffers may adversely affect performance, even to the point of causing network faults. To perform buffer management with today's emerging standards, one must resort to protocol extensions. Buffer resources must be addressable and allowed actions on them must be defined. Since it is not currently subject to standardization, buffer management will likely be defined, at least initially, by implementors. Therefore some form of implementors' agreements (or at least public disclosure) will be necessary to achieve fully functional multivendor interoperability.

In defining extensions, it is important to define a minimal useful subset. Duplication of functionality should be avoided. For example, there should be one correct method for addressing the newer resources (managed objects) to avoid duplication of development efforts initially, and simplify software maintenance later. However, because of the ways extensions are developed, avoiding duplicate function is not always possible.

Three alternatives in which extensions can be added have been suggested:

(a) Ad-hoc vendor-proprietary:
    Desirable features:
    - full level of functionality can be obtained;
    - allows testing and rapid development;
    - later, can be published.
    Undesirable feature:
    - does not provide for interoperability among different vendors.

(b) Extensions developed by recognized standards making bodies:
    Desirable features:
    - avoids duplication of effort;
    - insures generality across vendors.
    Undesirable features:
    - tends to be suboptimal for any given implementation-dependent resource;

- may be slow in coming to market as standardization process is slower than proprietary solutions;
- violates spirit of implementation-dependent resource standardization.

(c) Privately developed (i.e., as in alternative (a) above), and then made public:

Desirable features:
- allows for rapid proprietary implementation and testing;
- should lead to interoperability (eventually);

Undesirable features:
- may evolve into a "defacto" standard which is less than optimal for many implementations (i.e., an inefficient solution to the management of such managed objects);
- may need to assign new "public" object identifier to old "private" object.

## 5.15. Scalability

Scalability is the ability to manage various size networks with a common set of techniques in an efficient manner. Techniques that may work well on a small centralized network of 50 to 250 nodes may not perform well, or indeed not at all, on networks of thousands or tens of thousands of nodes, and vice versa.

In developing OSI management standards, it is important to consider methods that will work correctly and efficiently on all sizes and configurations of networks. It is not obvious that the emerging OSI management standards can address all of the problems involved in managing very large networks. Perhaps simulation modeling can be applied to aid in discovering potential problems and verifying the applicability of the currently emerging standards to large networks.

Although there is nothing to indicate that these emerging standards are not capable of managing very large networks, it is clear that additional standards or additional managed objects for domain management must be considered. For example, methods that allow multiple manager systems are needed for efficiency and robustness. It would be inefficient, for example, for a wide area network manager to maintain connections with multiple managed systems on a single local area network. Rather a single local manager might maintain multiple connections to each local system and report to a senior network manager over a single connection. (See sec. 5.5.)

Multiple manager systems are common in today's large commercial networks. Methods are needed for combining such systems through hierarchical domains of responsibility. Some of these concepts are introduced in the Systems Management Overview [SMO], but additional definitions to provide interoperability among multiple managers, standardization of the large configuration data bases used to maintain information on network components and topology may need to be considered. (See sec. 5.2.)

## 5.16. Robustness

Robustness, the ability to continue functioning under adverse conditions, must be provided by network management products and services. The emerging network management standards must provide, whenever possible, features that will make management resistant to faults. For example, there must be provision to support multiple redundant managers by developing standard ways for them to interact for the purpose of providing redundancy and/or fault tolerance. Furthermore, management information may need to be sent along special high priority connections. (A technique sometimes called "out-of-band signaling.") Management protocols (e.g., CMIP) must recover in the face of communications errors.

It is impossible for robustness to be provided by the emerging network management standards alone. Much of an implementation's ability to exhibit robustness is based on implementation-dependent features and much is based on standards developed for the various layers. However, the overall network management standards and architecture must allow implementors, to the greatest extent possible, to include features for enhancing robustness, in a standard, interoperable manner.

Features identified as important for robustness include:

- Multiple redundant managers,
- Distributed as well as centralized managers,
- Out-of-Band or high priority signaling (Manager-to-Agent), (Note: This is **not** a CMIS service.)
- Throttling of error messages to prevent a single fault from causing the network to be flooded with high priority alarm messages, (Note: It appears that this can be provided by event reporting discriminators which limit alarm messages as described in the Management Service Control Function [MSC]).
- Network manager to function even if its resources have been exhausted (e.g., no more space on disk for logging error reports),
- Failure of manager (or loss of communication to manager) to not cause agent systems (i.e., managed systems) to fail. (This is primarily an implementation design matter.), and
- Normal communications operation does not require systems management.

It is not clear that many of these features are provided by the emerging OSI standards. Further study and resolution of issues relating to robustness are needed.

## 5.17. Merging Existing Networks

The topology of a computer network often reflects the management structure of the organization using that network. Consequently, when real world management structures undergo changes through reorganization or mergers, these changes must often be accompanied by similar changes in the associated computer networks. The formidable task of merging two or more already existing networks is left to the network manager.

One of the significant considerations in merging networks is the issue of naming and addressing. Problems can arise both when there are, and there are not, conflicts in the naming conventions used. In the first instance, there may be many conflicts in the conventions used for naming resources. These conflicts must be resolved. Tools must be provided for updating configuration tables and other resources that control network configuration, thus enabling orderly transitions. On the other hand, problems may still emerge where there is no conflict between conventions used for resource naming. If the networks being merged follow the same conventions, there may be a "name space collision" of resource names. Clearly, tools to provide an orderly transition are required in this case, as well.

The merger of OSI and non-OSI networks brings to light even more complex issues. Unfortunately, these issues are too complex to be discussed in detail in this paper. Since there is no standard non-OSI network, a solution that may work for connecting one type of proprietary network, may not be similarly successful in connecting a different proprietary network architecture to an OSI-based network. In the near term it is probable that vendors themselves will develop gateways to OSI networks and that such gateways will be used for transition or migration.

At the NIST, we have developed Application layer gateways for TCP/IP (DoD) architectures to be connected to OSI-based networks. These gateways are for file transfer (FTP/FTAM) and for mail (SMTP/X.400). Perhaps these gateways may serve as a model for how to merge dissimilar networks.

Using emerging, proposed OSI management standards, the OSI/Network Management Forum is developing methods for managing merged proprietary network technologies. These methods should also be useful for transitioning from existing networks and technologies to future OSI-based networks. However the Forum's methods involve the use of managed objects that precede normal standardization processes and therefore these managed objects themselves will need to be extended into future OSI-based standards.

A similar problem exists within the NIST OSI Workshop NMSIG Implementation Agreements. As the standards evolve and the various phases of the IAs are developed, it is essential that the extensibility is maintained.

### 5.18. Scope of Standardization

An important issue that the reader must understand deals with the question: what is the appropriate subject of international standardization? To fully appreciate this issue, figure 3, a simplified view of a model of OSI Systems Management, should be discussed. In figure 3 an example of a "managing system" is shown on the left and a "managed (or "agent") system" is shown on the right. The typical OSI end system, or host computer, consists of seven layers that must be managed (Physical through Application). At the Application layer, management must be provided for a number of Application Service Elements (ASEs) which provide services such as electronic mail, file transfer, remote terminal, The Directory, and systems management (SMAE). Each

layer is supported by a layer management entity (LME) responsible for providing layer specific management. The LMEs may provide management services through (N)-Layer management protocols (not illustrated in fig. 3). Also they may assist systems management indirectly through implementation-specific means or they may exchange information through the Management Information Base (MIB), discussed below.

A Systems Management Application Process (SMAP) conceptually exists to provide OSI systems management. The SMAP, as shown in figure 3, consists of two portions, Information Processing and Systems Management Application Entity (SMAE). The Information Processing portion provides management decision-making logic, and in the case of managing systems, operator displays and other operator services.

The SMAE provides OSI functions to support management. These functions include those developed for Configuration, Performance, Fault, Security, and Accounting. Examples of such functions include the Object Management Function and the Error Reporting and Information Retrieval Function. The functions are implemented in the Systems Management Application Service Element (SMASE). The SMASE exchanges Management Application Protocol Data Units (MAPDUs) with its peer element to effect systems management. To perform its duties, the SMASE may call upon the services of the Common Management Information Service Element (CMISE), which in turn invokes the services of the Remote Operations Service Element. The CMISE employs the Association Control Service Element to establish associations for the purpose of exchanging Common Management Information Protocol Data Units (CMIPDUs) with its peer on another management system. The SMASE may employ other ASEs such as File Transfer, Access and Management (FTAM) to accomplish its functions.

The SMASE portion of the SMAP provides functions that are currently subject to OSI standardization. The other portion illustrated, Information Processing, is not currently subject to standardization. The conceptual interface between the two portions is represented in figure 3 as the control process interface. The decision-making process may be implementation-specific and vendor-proprietary. The implementor may choose his own or others' proprietary algorithms for effecting management. He may choose to employ artificial intelligence automation techniques such as expert systems to aid in decision-making or possibly refer decisions to an operator (via a display or other operator interface) for resolution. Having a human operator creates another interface, the man-machine interface (MMI).

Elements of the SMAP, layer entities, and LMEs cooperate exchanging information through the local management information base (MIB), often called "a conceptual repository of management information." Access to the MIB is through a conceptual "management data interface." This interface, provides a "local view" of the MIB's structure and contents, is not specified by emerging ISO management standards. However, a view of a remote MIB, including the structure, contents, and access of such a remote MIB, are currently undergoing standards development within ISO. Such standardization results in a common understanding of activities within distributed OSI systems, without which meaningful management would be difficult, if not impossible.

Currently, only the management information that crosses the interoperability interface, as shown in figure 3, is subject to appropriate international standardization. Although the standardization of the control process and management data interfaces would aid software portability, these interfaces are not currently being considered for standardization because they do not affect interoperability of open systems. The standardization of the MMI, which leads to commonality and consistency for network management operators is, perhaps, an area of administrative concern because portability of network management operations personnel saves considerable money in reduced training needs and increased operational effectiveness. However, the standardization of the MMI on an international scale is not appropriate at this time because vendors can use enhanced MMI as a product differentiation feature and because each user organization is likely to have markedly differing MMI needs. In the future as implementations of integrated network management systems based on international standards proliferate, such standardization may be appropriate.

Within the body of the present report many user requirements for network management exist which have no direct effect upon the process of developing international standards. In each instance these extra-standardization requirements (e.g., man-machine interface) have been identified as such. The most important effect of such requirements on standards making is that emerging standards must not preclude the possibility of satisfying user requirements not themselves subject to standardization. These factors make the process of establishing network management standards very complicated indeed.

### 5.19. Standard Application Program Interface

The current OSI standardization for management does not include a standard for an Application Program Interface (API). An API is a method for specifying the interface that an application program uses to interact with some service. For example, in graphics there is a standard interface, known as GKS, for specifying exactly how objects are to be drawn on graphics devices. There are similar standards under development for operating systems calls for a portable operating system called POSIX.

In the world of OSI standardization, the interfaces between layers has always been treated as an abstraction. Indeed, conformance has usually involved specification of protocol data units to be exchanged between open systems, with only abstract upper and lower layer boundaries specified (for which there is no conformance).

Recently there have been reports of several vendors agreeing on an API standard for X.400 message handling systems. There has been a recent call within the Internet Engineering Task Force (IETF) for consideration of an API for management. Perhaps the time is right for consideration of the development of an API standard for OSI management.

If we consider the Simplified OSI Management Model depicted in figure 3, we can see an interface labeled "Control Process Interface" between the

Information Processing portion of the Systems Management Application Process and the Systems Management Application Entity portion. If this interface were to be standardized, there might be several benefits to OSI management users.

First such a standard API would allow third party vendors to develop portable software for special graphics interfaces. Without a standard control process API, a third party vendor would be forced to sign agreements with OSI management vendors for rights to examine each vendor's management source code and then develop custom software for that vendor's implementation. Relicensing and redesigning would be necessary to port to another vendor's product. With a standard API, a third party implementation might require nothing more than recompilation.

Another benefit of a standard API is a third party data base management system. Network management users (administrators, operators, and planners) want the ability to store data for long periods of time and produce trend reports, as well as other analysis based on OSI management data. Without a standard API, these users might be forced to use an OSI management vendor's DBMS, or to seek a third party vendor who has developed special DBMS software. This is similar to the special graphic software discussion above.

A third area that might benefit from a standard API is automated network management, e.g., the use of expert systems. Again the benefits to the user are the same as those used in the discussion of graphics and DBMSs.

Yet another point to be made in favor of standardizing an API for management is that it would allow each user to customize his own reports in any manner he saw fit. There may be no need to employ a third vendor for customers who have the programming expertise to develop their own solutions. And these customers would be reasonably assured that a "custom" solution would port to another vendor's implementation with no major code redesign.

A final point in favor of a standard API for management is conformance testing. A standard API would make available a common interface for all vendors management systems. A conformance test system could be easily recompiled to test each vendor's management system, with little need for major test system recoding.

There are a few reasons why such standardization should not be considered. First it might be difficult to specify an API. What programming language or languages should be supported? There is also the matter of efficiency. Would a standard API force implementations to be less efficient than otherwise? Would such a standard stifle vendor creativity?

While the answers to these and other questions that must be posed are not clear, there are some extremely attractive potential benefits for OSI management users in the standardization of a control process management API.

## 5.20. Taxonomy of Managed Objects

One of the great problems in developing a Management Information Library (MIL) of Managed Object (MO) classes is determining some scheme for categorizing (classifying) the MOs. This is not a simple task for several reasons. One is that it is difficult to anticipate future enhancements or changes in technology that will cause one MO class to suddenly merge with or acquire attributes of another class. Another is that different MO designers often have different "views" of the same or similar situations. For example, circuit vendors (e.g., common carriers) may want to define a "circuit" MO with associated termination characteristics. Modem vendors may have a view that prefers to define a "modem" class MO with associated circuits.

There has been little background study in how to organize a taxonomy of managed objects, but initial attempts by various groups to begin the task have shown that it is difficult to obtain consensus, primarily of the problem of differing "views" of the world by different MIL designers. Without proper taxonomy design now future MIL development could be constrained or be unnecessarily difficult.

# 6. AUTOMATED NETWORK MANAGEMENT SYSTEMS

## 6.1. Need for Automated Assistance to Network Management

Network management is one of the largest unresolved problems facing the data communications community today. International standards for network management are now in development; a minimal set of international standards to provide initial functionality is not expected for at least 2 years. Implementations of these standards may lag the standards development by as much as 2 more years.

It is already clear, however, that the emerging standards will not solve the entire network management problem; they only define the tools necessary to implement interoperable management systems. Effective network management requires personnel as well as tools. Personnel are required to determine how to configure, control and monitor network systems. Decision-making for network management requires highly skilled, trained personnel. Control processes needed to oversee increasingly complex networks must be capable of handling a wide variety of configurations, functions and parameters. Consequently it will be difficult for even highly trained and skilled network management personnel to manage networks in a consistent and knowledgeable manner to achieve effective real-time network control.

Automated network management, employing "expert systems" (defined below), will become essential components of network management to assist network managers in keeping interconnected heterogeneous network systems under control and in providing the desired network services in an efficient, secure and cost-effective manner. In fact, without automated network control, the usage and application of all types of communications network services may be limited by the availability of qualified network operators. (Note that automated network management is not currently being considered for standardization since it is not required to achieve interoperability.)

A rather infeasible alternative to automated network management might involve, for example, the hiring of operators who: 1) have Ph.D.s in Data Communications, 2) will work for $12,000 a year, and 3) will work third shift. A more reasonable approach, however, is to employ expert systems to support network management operators.

The design and implementation of expert network management systems is presently a subject of increasing interest among users, vendors and researchers. Although no products have yet been announced, some prototype expert network management systems have been developed. The automation of network control systems has been identified as an important future need by network management experts at various meetings and workshops. NM experts at the May 1989 IFIP NM Symposium in Boston, MA indicated that the development of an intelligent network management system is the next essential step in the development of future management systems.

## 6.2. Introduction to Expert Systems

Knowledge-based expert systems (to be referred to here, for convenience, simply as expert systems) are intelligent computer programs that use knowledge and inference procedures to solve problems that ordinarily require significant human expertise and intelligence. The knowledge of an expert system consists of facts and heuristics. The "facts" constitute a body of information that is widely shared, publicly available, and generally agreed upon by experts in a field. The "heuristics" are mostly private, little-discussed rules of good judgment (rules of plausible reasoning, rules of good guessing) that characterize expert-level decision-making in the field [FEIG82]. Such a collection of knowledge (i.e., including both facts and heuristics) is commonly referred to as a knowledge base. The performance level of an expert system is primarily a function of the size and quality of the knowledge base that it possesses.

Inference procedures are the control structures that organize and control the steps taken to solve problems. The inference procedures are often referred to, in an expert system, as an inference engine or rule interpreter. An inference engine executes the knowledge (e.g., the rules, in a rule-based system) stored in the knowledge base to provide problem solutions to the user. The knowledge necessary to perform at such a level, plus the inference procedures used, can be thought of as a model of the expertise of the best practitioners of the field.

Although the knowledge base is generally unique to a particular domain, the inference engine may be common to a number of domains that have similar characteristics. A number of inference engines have already been developed for various types of expert systems (e.g., for decision-tree, rule-based, and frame-based expert systems).

The main task in developing network management expert systems is to derive a knowledge base for the various network management functions from human "expert" network managers. In order to solve these types of problems currently addressed only by human experts (e.g., in engineering, medicine, and computer configuration) machine problem-solvers need to "know" what the human problem-solvers know about that subject. Therefore, a human "domain expert" usually collaborates to help develop the knowledge base. Throughout the past 2 decades, AI researchers have been learning to appreciate the great value of domain-specific knowledge as a basis for solving significant problems [FRED84]. Experts are often those who know more facts and heuristics about a domain than nonexperts [GEVART85].

The theory and practice used to construct a knowledge base form a new discipline called "knowledge engineering." Knowledge engineering attempts to provide a mechanism for people to capture, store, distribute, and apply knowledge systematically.

Expert systems are usually used as "assistance to" rather than "replacements for" human experts. In view of this, then, with regard to network management, the objective of such a system is to assist the network operator or manager by providing a set of recommended actions as problem

solutions. For example, when a fault occurs, instead of just reporting an alarm or providing huge amounts of data, an automated network control system might present the operator with a list of possible actions along with the fault indication.

The applicability of expert systems is virtually limitless. Such systems can be used to diagnose, monitor, analyze, interpret, consult, plan, design, explain, learn, and conceptualize. Current expert system applications include medical diagnosis, equipment repair, computer configuration, speech and image recognition, financial planning, military intelligence and planning, and VLSI design.

The great promise of expert systems lies in their ability to provide effective problem-solving assistance by processing "knowledge." Conventional data processing programs, on the other hand, are locked into well-defined and relatively inflexible algorithms needing complete and precise data; and these algorithms are limited to using only the body of data prescribed for them. To date, no great difference exists between the hardware used for conventional data processing and that used in knowledge processing; rather, the major differences between these two approaches are in the software. Conventional software represents and manipulates "data", whereas knowledge-based systems represent and manipulate "knowledge." Conventional computers do not "know" how to interrelate data to produce knowledge because no one has been able to tell them how to do it, until now. Conventional software can use only algorithms, while knowledge-based software can use heuristics in addition to algorithms. This is of particular value when the knowledge cannot be made precise and certain or when it is not economical to be precise. For the first time ever, computer-provided solutions may say, "I may be wrong, but my best guess is... ". The use of repetitive processes in conventional software contrasts with the use of inferential processes in knowledge-based software. That is, knowledge-based software can take "facts," and relate them in logical ways, thereby producing new "facts." "Artificial Intelligence technologies (in particular, expert systems) are regarded as the second giant step toward human's full realization of the power of computers." [ARTI86]

### 6.3. Approaches to Developing Automated Network Management

The following diagram depicts a model of an ideal expert system. Major components of an expert system are: 1) a knowledge base of domain facts and heuristics associated with the problem, 2) a set of inference procedures for utilizing the knowledge base in the solution of the problem, and 3) a working memory for keeping track of the problem status, the input data for the particular problem, and the relevant history of problem solving activities. In a more sophisticated system, an explanation module can be included, allowing the user to challenge and examine the reasoning process underlying the system's answers.

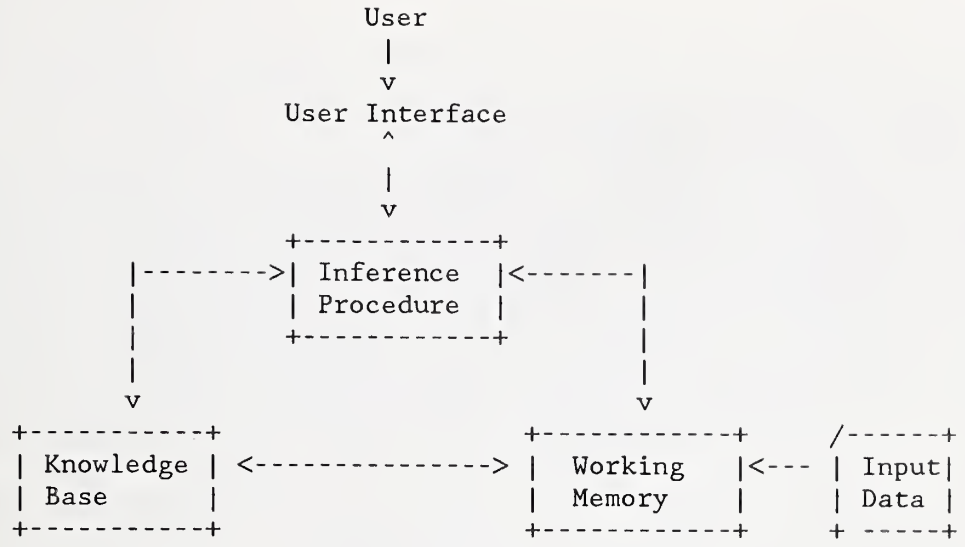


Network management provides a wealth of opportunities for the application of artificial intelligence techniques which may ultimately have far-reaching effects on many, if not all, areas of network management. To realize the goals of truly automated network management systems, however, research is needed:

1) to develop effective man-machine interfaces for monitoring and controlling network activities. Currently, menu-driven, multi-window screen displays, and graphics interfaces are often employed.

2) to evaluate the efficiency of various knowledge-based inference rules and search strategies for network management application.

3) to select and acquire the appropriate tools (shells) that have implemented the desired inference procedures for use in the implementation of automated network management systems.

4) to engineer the network management knowledge into a set of rules/ knowledge base. Knowledge representing network management expertise must first be solicited from network management experts, operators, administrators and managers/analysts, and then organized and integrated into inference engine executable rules.

5) to select and design network management processes that consist of those management functions/activities which can demonstrate the usefulness of automation (e.g., performance analysis and fault isolation and recovery), and

6) to integrate the above components into an expert network management system that allows automated NM control and monitoring.

## 6.4. Impact of Automated Network Management

Expert systems for network management will help to solve the network management automation problem faced by our nation's industry and government. Managers of several large operational networks have indicated that the future of network management is to automate NM's decision-making processes. AI techniques will help network management perform required control and monitoring functions with regard to, for example, dynamic selection and adjustment of parameters for flow control, traffic routing, and quality of service. Such automated assistance can also help reduce extensive operator training requirements. Other NM functions such as routine network testing and diagnosis, software distribution, statistics gathering, network capacity planning, dynamic tuning of network parameters, load balancing, analysis of failures, and repair scheduling and inventory control may all benefit from the help of expert systems. Without the help of expert systems, it may become economically infeasible, or even impossible to perform part or all of these tasks.

Research into the use of expert systems in network management will not only help accelerate the attainment of automated network management systems in particular, but will also advance the development of AI technology in general. Professor R. Goodman of Caltech., at the May 1989 IFIP NM symposium indicated that the work done to identify NM requirements for automated decision-making had had the related effect of significantly accelerating the development of several areas in AI as well. Examples include that 1) the research in parallel inferencing techniques has been stressed due to the requirement of real-time network control; and 2) the development of machine learning has more rapidly advanced due to the requirements of automated network topology identification.

Today most of the basic research in AI is being conducted primarily in universities and some nonprofit laboratories. In the United States, DARPA, NIH, and NSF have traditionally provided most of the funding for this research. Both Japan and the UK recently have undertaken significant new research programs in AI with government support [HELLI86]. The three major subfields of AI today include knowledge engineering, natural language processing, and vision and robotics [FRED84]. Substantial commercial interest has developed in all three subfields. In addition, all three subfields today exhibit commercial applications of the technology and commercial tools to support additional applications [MERITT86]. We believe that expert systems applications in NM will be a prominent area of AI applicability.

## 7. ISDN NETWORK MANAGEMENT

Integrated Services Digital Network (ISDN) may be defined as an end-to-end digital network that provides customer services using existing subscriber loops. The CCITT-issued 1988 Blue Book(s) describe the ISDN concepts and standards. The development of ISDN is a still-evolving international movement which pledges to provide digital transmission to customer premises, and to allow integrated access to voice and data services including possibly video and other types of services by using circuit and/or packet switching services. The main feature of ISDN is the provision of a wide range of service capabilities with the promise of great potential in future applications. ISDN is intended to provide easier access both to knowledge and to distributed processing. ISDN is also intended to provide new and better customer service applications. It is expected to stimulate the development of new third party enterprises and to provide other advantages not yet imagined.

Fundamentally, ISDN offers its end users an interface into intelligent networks. End-user access to intelligent networks through ISDN may occur over a basic-rate interface (BRI) or a primary-rate interface (PRI). The BRI is composed of two B channels plus a D channel (denoted 2B+D) where each B channel transfers user information such as data, voice and video at a rate of 64 Kb/sec., and the D channel transfers signalling and control information, and sometimes user data, at a rate of 16 Kb/sec. The PRI may be composed of either 23 B channels plus a D channel (23B+D) for the 1.544 Mb/sec line rate commonly used in North American telephony systems, or 30 B channels plus a D channel (30B+D) for the 2.048 Mb/sec line rate commonly used elsewhere.

A key issue which still needs to be addressed by the providers of ISDN services and equipment is that of network management. These providers (and potential providers) must consider user needs. For example, one large potential ISDN customer has indicated that end users must have access to maintenance functions in terminal equipment and terminal adapters (used for interfacing non-ISDN terminals and ISDN). Users will not tolerate proprietary network management systems for ISDN, since ISDN is so heavily multi-vendor oriented [WEIS89]. The trend has shown that ISDN end users strongly desire access to network management functions and information. In this regard, the message-based call control over the D channel in ISDN may offer great potential for user management of network capabilities inherent in the signalling functions.

The goal of ISDN network management is similar to the goal of OSI management -- i.e., to have a single system that manages all network components, such as terminal adapters, switches, work stations, digital phones, and network services that include, for example, telephone connections and supplementary services. Network control and management will be vital to the operation of integrated digital networks. Usage and application of ISDN services will be affected by the availability and effectiveness of the network's management system.

Adoption of ISDN necessitates new equipment, including, for example, new telephones, terminals, and central office line modules. As standards for new equipment are developed, NM requirements for this equipment must be

considered. Furthermore, as many additional types of ISDN services are developed and provided to users, the associated functions and information necessary for management of these new services should likewise be identified and defined.

Because ISDN is still evolving, its management requirements are not yet well-defined. A method is needed to help in the process of identifying these requirements. This section of this report proposes such a method. Section 7.1 describes the areas where the ISDN NM requirements should be investigated. Section 7.2 describes the generic elements of a network management system regardless of the network environment. ISDN NM standards and the NM activities of the North American ISDN Users' Forum (NIU) are discussed in section 7.3. Finally, section 7.4 describes outstanding issues in ISDN network management.

## 7.1. How to Identify ISDN Network Management Requirements

ISDN service providers and carriers need management tools to ensure that network performance meets users' criteria. The network control center needs various tools and mechanisms to operate, maintain and monitor the entire network. It is anticipated that many users of ISDN services may develop applications for providing specialized services to their customers. Such users become "value-added" vendors. These "value-added" vendors will demand the quality of service they pay for so that they can, in turn, provide the level of performance their customers expect. Applications for the commercial use of ISDN are being considered by various potential user and implementor groups. These applications, not surprisingly, are expected to have their own network management requirements.

In order to sort out all these concerns and requirements, we have divided the NM requirements of ISDN into three layers starting with the requirements of ISDN end users and proceeding through the requirements associated with ISDN services and network operation and then to the requirements of managing the core of the network, the communications protocols. The outer layer consists of various end users of ISDN services. These users normally want to have control over their terminal equipment as well as to have access to certain network statistics. The middle layer comprises two parts: 1) ISDN services applications which may have both common and unique NM requirements depending upon the application, and 2) network planning/design, installation, administration, operation, and maintenance. The inner layer consists of the protocols that specify hand-shake rules between communication devices and consequently make information transfer possible between computers/terminals through communication media and interconnected communications equipment. The performance and reliability of these communications protocols can be managed through options and settable parameters. Each of these areas will be examined in the following four subsections as a first step toward analyzing ISDN NM requirements.

### 7.1.1. End Users NM Requirements

To end users, ISDN represents a set of standard interfaces for connecting various terminal types to a network facility that should provide lower cost circuits than available today. Customer premise equipment (CPE) will include adapters for connecting non-ISDN terminals to the interface, and network termination units for connecting ISDN terminals and/or terminal adapters to the network ISDN switches.

ISDN subscribers will want to have access to maintenance and configuration functions in the terminal equipment and terminal adapters. In other words, from the CPE point of view, management information needs to be kept and exchanged among all the CPE components. However, it is still not at all clear whether end users can or will be able to access information in network termination units. Some end users of ISDN services (e.g., third party vendors) may want to have certain control over the level of quality of services that they provide to their customers, or they may simply want to be able to specify and/or change the required level of services.

Recommendations for NM requirements need to be solicited from all classes of users of integrated network management systems which can then be offered as input to the relevant standards groups. One of the important outcomes of reaching a good definition of the overall ISDN NMS concept will be to clarify what maintenance and control information, alarms, and actions are to be network controlled and what operation and management information/notification and logs are accessible by the end users and various other types of NM users. Some important questions in this area which still need attention include whether or not the end user may initiate diagnostic tests for his or her own customer premises equipment, whether the end user may manage and/or load a profile which determines the characteristics of his or her terminal equipment, and whether network generated alarms may be received by end users, and if so, which alarms need to be sent to the end users.

### 7.1.2. Application-Oriented NM Requirements

ISDN (including the standards, the manufacturers of ISDN equipment and the customers for this technology) is almost in place. One element still missing is the identification of useful commercial applications (e.g., scenarios). The North American ISDN Users' Forum (NIU), which meets quarterly to address ISDN users' concerns and requirements, has received approximately 70 ISDN applications for analysis as of its March 1989 meeting. Some of these applications are network management oriented ones, such as the multi-media services application profile. This scenario delineates a unique NM problem in the ISDN environment related to the provision of a variety of types of services, including, for example, voice, data, imagery, and video. Since each service type has different performance parameters, monitoring requirements, and connection and operation limitations, management of the ISDN portion of the network must be developed with these service types in mind.

Although most of the applications are not NM-oriented, nevertheless, in most cases there will still be some NM requirements associated with these ISDN

applications. For example, one possible ISDN application is to allow one end user/customer to talk to another end user/customer through one channel and, at the same time, to allow data files to be exchanged between the two users through another channel. If the channel used for transferring files goes down during the transfer and the resume functions fail, what management operations and functions would be required to recover from this fault situation by the end users and/or by the switch operators? Should the end-user send notification to the operator or administrator for fault recovery actions or simply let the system (protocols) detect the fault and provide the necessary fault management actions? Is the switch operator allowed to divert the rest of the file transfer to the other available channel (e.g., the D channel) if the fault functions available to the operator do not resolve the problem quickly? While many NM requirements will probably be common among applications, it is equally likely there will be unique management requirements for certain applications.

Another example of specific NM requirements for ISDN applications concerns the management of ISDN terminal profiles. Consider the situation in which a customer service agent transfers a customer's call to another customer service agent. In such a situation, should access to management information about this customer be transferred to the new agent? The management information in question may include the customer's terminal profile which defines the functionality of the terminal that the customer is using. When the new agent has a choice to supply more than one type of service in response to the customer's request, that new agent wants to know what types of services this customer's terminal supports. Access to the information contained in the user's terminal profile may be restricted to network management personnel such as the administrator, who may be solely responsible for updating customers terminal profiles. However, some customers may require access rights to update their terminal profile(s) and to change their terminal configurations themselves.

NM functions and requirements also affect the types of source and destination points of the management operations. They include user-to-user management operation, user-to-local-management-facility management operation, user-to-central-management-facility management operation, local-management-facility-to-local-management-facility-management operation, and local-management-facility-to-central-management-facility operation. Requirements pertaining to all of these interfaces and operations must be investigated. The development of models for each of these operations is an important step in identifying the NM requirements for these operations/services and for required control of access to them.

NM functions and information required by those applications should be identified and provided to the standards-making groups. As each application is submitted and analyzed by the NIU, whatever management requirements are decided upon will be incorporated into each application's document and then the NM requirements from all these applications will be formally presented. To further this process, network service providers, including third party service providers, should determine their own NM requirements as they develop services and then submit these management requirements to the standards-making groups.

Since user requirements for ISDN applications are still being identified and analyzed, it is clear that NM application requirements cannot yet be completely identified.

### 7.1.3. Network Installation, Administration, Operation and Maintenance Related NM Requirements

From the network perspective, ISDN comprises equipment and carrier plants that use standard interfaces to and from ISDN switches at central and toll offices. Thus, it is understandable that a major focus of NM requirements for carriers is to ensure that equipment is indeed compatible and in proper working order.

Operators at the switches must make sure that the network is functioning well. Their responsibilities may include running routine diagnostic tests of network components as well as identifying and repairing faults, possibly including the replacement of defective equipment. Such actions often will be most effectively performed remotely by means of network management tools which retrieve information and send control and testing commands to managed components. NM administrators of clusters of PBXs or switches may, moreover, want to remotely access maintenance and configuration functions in the switches. If so, they will likely need sophisticated tools in order to control network resources so as to provide the required level of service. NM requirements vary considerably depending, in part, upon the operators'/administrators' responsibility and interest. Moreover, these responsibilities and interests may span many interfaces and domains.

In some organizations, a network manager controls configuration changes in his domain, monitors traffic, balances (changes) the routing of traffic, and authorizes the use of functions that are network controlled. Network performance analysts, capacity planning personnel, tariff and accounting specialists, network security personnel, and market analysts all need to access management information to fulfill their respective job functions.

Many network management functions are required for routine operation and maintenance of the network. For examples, NM functions are needed to handle such problems as breaks in the physical connection, downed subscriber loops, or routine and before-use confidence tests of equipment. Management functions are also required to analyze results of these tests. Although the analysis methods do not need to be standardized, existence of these methods is nevertheless required in order to determine whether the equipment under test is operational. Since operating companies usually have many requirements in this regard, these companies should provide a prime source for definition of these requirements to the standards groups.

Network management requirements are not static. They change from time to time, from configuration to configuration, from network to network, and from user to user. Consequently, the identification of network management functions is a dynamic and evolving process. Moreover, management information may reside in various devices supplied by different vendors. Access to

management information may require crossing many domains that are controlled and billed by different organizations. This most likely will create complicated problems, particularly in the areas of security and accounting management.

### 7.1.4. Protocol-Based NM Requirements

Network management functions can be categorized according to the types of problems they address (e.g., configuration, fault, performance, accounting, and security). Within each of these categories, some of the NM functions are protocol related, while others are not. For example, the management of flow control window size is a protocol related management function, while the recording of the number of times that a link goes down per unit time period is not a protocol related management function. Nonprotocol related parameters are those parameters which do not affect protocol behavior and performance. End users' NM requirements usually are mapped into nonprotocol related NM information and operations. For example, end users often want to know what network service charges have occurred thus far in a month. Requirements such as automatic reconfiguration and reconnection, and adaptive routing, may involve both protocol related and nonprotocol related information and operations.

Each party responsible for the standards development of the communications protocols knows the details of these protocols the best. Therefore, when protocols are developed, the experts should identify the operations, facilities, parameters, and objects for the purposes of management, control, and maintenance of these protocols. Guidelines for performance tuning and fault recovery should also be considered to facilitate the management of later implementations and, at the same time, assist in evaluating the viability and performance of the protocols being standardized.

## 7.2. Elements of an ISDN NM System

In order to provide management functions such as fault, configuration, performance, accounting, and security management, ISDN, like other network technologies, requires common management services, protocols, parameter descriptions, and managed objects. In addition, it is important to develop a model for ISDN network management as a base for the development of other management elements. However, many ISDN management standards developers are just beginning to deal with these concepts.

NM users want to see a consolidated set of standards that allow the development of integrated management systems that monitor and control both ISDN and OSI networks. Since ISDN is still evolving and its potential remains to be explored, it is not certain that the OSI NM architecture fits the ISDN environment. As the ISDN standards mature, the feasibility of using OSI NM to meet ISDN management requirements will be an appropriate area of investigation.

The OSI management standards groups already have expended considerable effort into identifying what must be standardized to provide interoperable NMS on multi-vendor networks. The fruits of this labor can and should be used by the ISDN community as a basis for development of network management for ISDN.

### 7.3. ISDN NM Standards and Implementors' Agreements on Standards

To achieve the goal of having a single integrated network management system across multi-vendors' communication and/or management products employing complex integrated technology, it is necessary to develop general management principles and a common network management architecture. The required capabilities and interfaces can then be identified with respect to the architecture. NM standards specify all those elements necessary to achieve interoperability among multi-vendor products. Those areas of network management that do not affect interoperability are not currently being considered for standardization. Examples of such areas include the user interface to the network management system and methods for the analysis of management information.

A number of committees within CCITT and ANSI are developing management standards. Within CCITT, management responsibilities are divided among Study Groups dedicated to one or more specific management areas. The following is a list of management related CCITT Recommendations (i.e., standards) :

CCITT - I.601 - General maintenance principles of ISDN subscriber access and subscriber installation

I.602 - Application of maintenance principles to ISDN subscriber installation

I.603 - Application of maintenance principles to ISDN basic rate accesses

I.604 - Application of maintenance principles to ISDN primary rate accesses

I.605 - Application of maintenance principles to ISDN static multiplexed basic accesses

Q.940 - ISDN user-network interface protocol for management general aspects

In ANSI, the T1 subcommittee delegates responsibility to T1M1 for management (which it refers to under the title: Operations, Administration, Maintenance and Provisioning). However T1S1 bears responsibility for the function-independent aspect of the management protocol. [Note: Clarification of the division of responsibility in developing NM standards for ISDN in ANSI can be found in the letter entitled "Liaison to T1E1 and T1S1 on ISDN User/Network Management and Maintenance" (reference number: T1E1.4/89-017) from the chair of T1M1 in January 1989.]

In addition to the ISDN NM standards being developed, implementation agreements will be needed once these standards are in place. In anticipation of this need, the network management group of the ISDN Implementors' Workshop

(IIW) of North American ISDN Users' Forum (NIU) is identifying the ISDN NM requirements that will eventually influence the development of the implementation agreements (IAs) on ISDN NM standards. To avoid redundant work in developing NM standards and IAs, the NMSIG of the OSI Implementors' Workshop and the NIU NM Working Group have met to establish a liaison mechanism. The coordinated effort between these two groups entails first investigating the similarities of the NM model in each environment with respect to NM requirements and then identifying the fruitful areas of collaboration in the work of the two groups. The current collaboration effort includes the NMSIG reviewing ISDN NM applications, and then forwarding draft agreements of the managed object templates to the ISDN group for comment regarding completeness of the templates for ISDN NM applications.

## 7.4. Issues

Because of the wide variety of domain types, technologies, and requirements involved in the management of ISDN, an initial practical approach to identifying NM requirements is to partition the set of ISDN management problems so as to be able to focus on each individual part. Section 7.1 discussed one way of accomplishing this by partitioning the ISDN NM requirements into four parts. Within each part, prioritization of work items is necessary. Identification of detailed requirements (e.g., objects and actions) of essential functions such as loopback tests for fault recovery will have higher priority, than, for example, identification of those functions which are important but not essential, such as the performance measurement of channel throughput under various loads. The development of an ISDN NM model may be necessary to help clarify the detailed NM functions and facilities, and the relationships among them.

An important NM issue is the transfer of management information across domains which may represent different ownership and control principles. The security and accounting concerns in this regard will need special study to fully understand the detailed requirements. Furthermore, it will be necessary to determine whether the common management protocol and services specified in Q.940 are sufficient to meet all the NM requirements identified.

Secured access to management functions is receiving considerable attention by the standards development groups. Since even authorized peer-entities are capable of making errors, the more serious of which can jeopardize the whole network and consequently affect other network users, operating companies will find it essential to seriously consider all aspects of user access to certain management operations and parameters. However, a trend is developing whereby ISDN users and customers expect the NM systems to provide many more tools and much more management information. Using on-line terminals, line printers and other visual equipments, users will want to be able to control and monitor their own private portions of networks via the direct customer interface. This clearly raises the issue of how to draw a boundary line between network controlled and customer/end-user controlled management operations and parameters.

In response to users' demand for end-to-end multivendor network management, vendors have begun to release proprietary implementation information to enable communication between their products and products of their competitors. They have also developed an interest in determining end user NM requirements. Some end user NM requirements may be satisfied by some of the ISDN supplementary services. (The ISDN standards define two types of services: basic and supplementary. The supplementary services can be implemented in conjunction with basic services to increase the number of features and functions available to the end user. Call waiting, call transfer and call pick-up are examples of supplementary services). In some cases, this raises questions as to who should maintain common information needed both by the service providers to implement the desired services, and by the end users/customers for network management purpose, and how access to such information should be controlled. Examples of situations of these cases include:

> The "Selective Call Forwarding" service, a refinement of the "Call Forwarding" supplementary service, allows the user to forward only those calls that originate from a predesignated set of network numbers. In view of the fact that many users want to be able to change this list of predesignated calling numbers on the fly, the issue arises as to whether this is a network management service reserved for ISDN service providers or a generic service available to end users?
>
> The "Charging Advice" service, listed as a subordinate service of the "Message Detail" supplementary service, provides the charged user with periodic real-time call detail elements and/or cumulative charges during a call. The information can be presented verbally or by an alphanumeric display. If, however, the NM user wants to have the "charging advice" information of all calls sent to one specific address to be stored for later analysis, how does the NM system interact with the supplementary services to share access to this information?

Other NM issues that need consideration are listed below:

- Since many ISDN applications are expected to require some management information, the best time to start identifying NM requirements for these applications is when the applications are first developed.

- The experience gained from the development of OSI NM standards suggests that ISDN NM requirements should be identified and shared with the standards-making groups in the early stages of standards development.

- Automated network management and control (see sec. 6 of this report for a discussion of this issue) is an increasingly important issue in ISDN due to the complexity of management information and the integration of varied technologies.

- Currently, there is one numbering plan for public switched telephone networks and one numbering plan for public switched packet data networks for addressing. Eventually the ISDN, telephone, and data networks will operate under a uniform numbering plan. Although the evolution of telephone and data numbering plans to support ISDN is a separate issue, it must be resolved before integrated management of ISDN, telephone and data networks can occur.

- Use of the D-channel signalling that is separated from the users' B-channel traffic will add both flexibility and complexity to network management as consequence of the interaction between the two flows and their associated resources.

## 7.5. Summary

The realization of ISDN will entail piecemeal replacement of communications equipment, which will be required to simultaneously interoperate fully with existing public switched telephone networks, and existing local loops. One can expect more and more network management requirements to surface as ISDN evolves and matures.

## 8. APPENDIX: NIST Phase Two Project Goals and Methodology

The following paragraphs provide an overview of a five step methodology for accomplishing the NIST phase two project goals (as described in sec. 2.2). These steps do not relate solely to this paper. Rather, they outline the activities surrounding the development of this paper, including initial research, analysis of findings, production of initial output, review and updating of output, and finally, application of results.

### 8.1. Investigate Functional Requirements

In order to gain as comprehensive a view as possible of network management requirements in terms of what functions, resources and operations are mandatory and desirable, this report documents both user requirements and functional requirements for network management. Vendor (proprietary) products were surveyed and analyzed to determine what facilities existing networks provide for network management. Other sources of information on network management requirements, such as the MAP/TOP NM 3.0 Specification, were analyzed as well. In addition, the authors have participated in the efforts of IFIP WG 6.6, which has a continuing study of NM user requirements. Finally, functional requirements are also pursued through participation in the Network Management Special Interest Group (NMSIG) of the NIST/OSI Implementor's Workshop.

### 8.2. Examine Scope of Standards

Prior to evaluating whether the functional requirements are successfully addressed by the solutions proposed in the developing standards, one must examine the scope of emerging management standards for OSI-based networks [NBS87] such as the management framework (MF), the common management information protocol (CMIP), the common management information service (CMIS), and the systems management overview (SMO). Recognizing that standards are never perfect and that changing technology, as well as users' changing perspectives concerning their needs, may alter functional requirements in the long run, we have chosen to concentrate our efforts on the current set of standards due to become International Standards (ISs) within the next 2 to 4 years.

### 8.3. Identify Incompatibilities Between Standards and Requirements

At this point, with information available about both the requirements and the related standards, it is appropriate to determine the functional comprehensiveness of the standards by comparing these emerging standards, from ISO and other standards making bodies, with the desired set of requirements. This comparison will demonstrate where the standards are fulfilling needs and where needs are still unmet. The existence of errors, omissions, or deficiencies in the emerging standards will be documented.

The results of this comparison focus on two main areas of interest. First, attention is concentrated on areas of incompatibility between standards and requirements which lead to situations in which requirements cannot be satisfied because those mechanisms needed to provide the functionality are precluded by the nature of the standard. Second, attention is then focused on those requirements not addressed by the emerging standards.

### 8.4. Solicit Additional Inputs

The very important next step was to solicit additional opinions and critical input on this study. A preliminary draft of the paper was produced and distributed for public review and comment. A workshop was held in late October 1987 to discuss issues related to requirements for network management. The preliminary draft of this paper was a major input to this workshop with the intent of focusing the discussion. The output from the workshop and from other public comment has been incorporated in this the final version of the paper.

### 8.5. Participate in Standards Formations

Once the study has been finished and the strengths and weaknesses of the emerging standards have been identified, the results will be made available to the ANSI working groups here in the United States, and to ISO groups internationally, in order to help assure that user needs are addressed by the work on NM. A common goal of those involved in standards development is that the emerging standards be versatile and extensible in order to meet future demand and technological changes. No reasonable management approach should be precluded by these standards.

## 9. REFERENCES [1]


[ALS] "Information Processing Systems - Open Systems Interconnection - Application Layer Structure (ALS)", ISO/IEC DIS 9545, 15 September 1988.

[AMWD] "Information Processing - Open System Interconnection- Management Information Service Definition - Accounting Management Working Document", ISO/IEC JTC 1/SC 21 N 3314, December 1988.

[ARP] "Information Processing Systems - Open Systems Interconnection - Common Management Information Protocol: Add/Remove Protocol", ISO/IEC 9596/PDAD 2, (ISO/IEC JTC 1/SC 21 N 3305), January 1989.

[ARS] "Information Processing Systems - Open Systems Interconnection - Common Management Information Service: Add/Remove Service", ISO/IEC 9595/PDAD 2, (ISO/IEC JTC 1/SC 21 N 3306), January 1989.

[ARTI86] "Artificial Intelligence Letter", January 1986.

[CANGETP] "Information Processing Systems - Open Systems Interconnection - Common Management Information Protocol: CancelGet Protocol", ISO/IEC 9596/PDAD 1, (ISO/IEC JTC 1/SC 21 N 3304), January 1989.

[CANGETS] "Information Processing Systems - Open Systems Interconnection - Common Management Information Service: CancelGet Service", ISO/IEC 9595/PDAD 1, (ISO/IEC JTC 1/SC 21 N 3303), January 1989.

[CMIPSPEC] "Information Processing Systems - Open Systems Interconnection - Management Information Protocol Specification - Part 2: Common Management Information Protocol", ISO/IEC DIS 9596-2, 22 December 1988.

[CMISDEF] "Information Processing Systems - Open Systems Interconnection - Management Information Service Definition- Part 2: Common Management Information Service Definition", ISO/IEC DIS 9595-2, 22 December 1988.

[CONFIG] "Information Processing - Open Systems Interconnection- Working Draft of the Configuration Management Overview", ISO/IEC JTC1/SC21 N 3311, 16 January 1989.


---

[1] It should be noted that because standards are not yet stable but are continuing to evolve it is possible that newer versions of the documents exist.

[DMA] "Information Processing Systems - Open Systems Interconnection - Structure of Management Information - Part 3: Definitions of Management Attributes", ISO/IEC DP10165-3, (ISO/IEC JTC1/SC21 N 3302), January 1989.

[DSO] "Information Processing Systems - Open Systems Interconnection - Structure of Management Information - Part 2: Definitions of Support Objects", ISO/IEC DP 10165-2, (ISO/IEC JTC1/SC21 N 3301), January 1989.

[ECMASEC] "Security In Open Systems - A Security Framework", ECMA TR/46, July 1988.

[ERIRF] "Information Processing - Open Systems Interconnection-Systems Management - Part 4: Error Reporting and Information Retrieval Function", ISO/IEC DP10164-4, (ISO/IEC JTC1/SC21 N 3298), 31 January 1989.

[FEIG82] Feigenbaum, "Knowledge Engineering for the 1980s", Department of Computer Science, Stanford University, Stanford, California, 1982.

[FMWD] "Information Processing Systems - Open Systems Interconnection - Systems Management - Fault Management Working Document", ISO/IEC JTC1/SC21 N 3312, January 1989.

[FRED84] Frederick, Hayes-Roth, "The Knowledge-Based Expert System: A Tutorial", IEEE Computer, September 1984.

[FRMWK] "Information Processing Systems - Open Systems Interconnection - Basic Reference Model - Part 4: Management Framework", ISO/IEC 7498-4 : 1989(E), 1989.

[GEVART85] Gevarter, William B., "Intelligent Machines", Prentice-Hall, 1985.

[GDMO] "Information Processing Systems - Open Systems Interconnection - Structure of Management Information - Part 4: Guidelines for the Definition of Managed Objects", ISO/IEC DP10165-4, (ISO/IEC JTC1 SC21 N 3437), 10 February 1989.

[HELLI86] Helliwell, John, "Guru: Brave New Expert System?", PC Magazine, May 27, 1986.

[IFIPUSR] "Network Management - User Requirements Analysis", by Rajan Rathnasabapathy, NT/BNR Richardson, IFIP WG 6.6, July 22-24, 1987, Revised - August 1987.

[LCF] "First Working Draft for Systems Management: Log Control Function", ISO/IEC JTC 1/SC 21 N 3309, January 1989.

[MERITT86] Meritt, Bonnie, "Trends of Expert Systems Building Tools: Personal Consultant Plus", Texas Instruments Engineering Journal, January/February 1986.

[MIM] "Working Draft for Structure of Management Information - Part 1: Management Information Model", ISO/IEC JTC1 SC21 N 3324, January 1989.

[MSC] "Information Processing Systems - Open Systems Interconnection - Systems Management - Management Service Control", ISO/IEC DP10164-5, (ISO/IEC JTC1/SC21 N 3299), January 1989.

[NBS87] Chernick, C. Michael; Mills, Kevin; Aronoff, Robert; Strauch, John W. "A Survey of OSI Network Management Standards Activities", National Bureau of Standards, (U.S.) NBSIR 87-3593; 1987 July. 58 p.

[OMF] "Information Processing - Open Systems Interconnection-Systems Management - Part 1: Object Management Function", ISO/IEC DP10164-1, (ISO/IEC JTC1/SC21 N 3295), 31 January 1989.

[PMWD] "Information Processing - Open System Interconnection-Management Information Service Definition - Part 6: Performance Management Working Document (Third Draft)", ISO/IEC JTC 1/SC 21 N 3313, January 18, 1989.

[RMF] "Information Processing - Open Systems Interconnection-Systems Management - Part 3: Relationship Management Function", ISO/IEC DP10164-3, ( ISO/IEC JTC1/SC21 N 3297), 31 January 1989.

[SECARCH] "Information Processing Systems - Open Systems Interconnection - Basic Reference Model - Part 2: Security Architecture", ISO/IEC DIS 7498-2, 18 June 1987.

[SECURE] "Fifth Draft for OSI Security Management Working Document", ISO/IEC JTC1/SC21 N 3315, December 1988.

[SFMF] "OSI Software Management Functions", ISO/IEC JTC 1/SC 21 N 3310, January 1989.

[SMF] "Information Processing - Open Systems Interconnection-Systems Management - Part 2: State Management Function", ISO/IEC DP10164-2, (ISO/IEC JTC1/SC21 N 3296), 31 January 1989.

[SMI] "Information Processing Systems - Open System Interconnection-Management Information Services - Structure of Management Information", ISO/IEC JTC 1/SC 21 N 2684, April 1988.

[SMO] "Information Processing Systems - Open Systems Interconnection - Systems Management: Overview", ISO/IEC DP10040, (ISO/IEC JTC 1/SC 21 N 3294), January 1989.

[WEIS89] Weissberger, Alan J., "ISDN Maintenance and Management: How Can We Make It a Reality?", Tutorial Presentation at NIU ISDN Users' Workshop, Network Management Working Group, 19 January 1989.

## 10. ANNEX: A Further View of Security Facilities and Their Management Requirements

The security community generally refers to the "top ten" security functions when discussing the functionality of a security system. Listed below are these 10 facilities with brief descriptions of their functionalities. The section following then specifies the management functionality required to support these facilities.

### 10.1. Security Facilities as Specified by ECMA

The "top ten" security functions, termed facilities by ECMA (as referenced in [ECMASEC]) are:

1. **Subject Sponsor**

   "The Subject Sponsor is the intermediary between the security subject [i.e., an entity in an active role to which a security policy applies] and the other security facilities. It is the only entity in a distributed system that is aware of all of a subject's current and sometimes concurrent activities while accessing protected objects or applications. The subject involved may be either a human end user or an accessing application."

2. **Authentication Facility**

   This facility receives and checks authentication information and reports approval or disapproval to other security facilities based on the in force security policy.

3. **Association Management Facility**

   This facility provides security for associations:

   - it uses the authorization facility to authorize communication between two entities,
   - it uses the authentication facility or other means to assure the identity of communicating entities.
   - it controls security of the communication service over which the association exists, (e.g., by controlling routing of the association over only trusted paths).

4. **Security State Facility**

   This facility maintains security state information pertaining to authentication records of associated communicating subjects and objects, including security

attributes of the associations. This state information is collected by each facility reporting its activity to this facility.

5. **Security Attribute Management Facility**

"This facility provides for the creation, distribution, revocation, archiving and destruction of security attributes of subjects and objects within a given security domain:

- subject-related access privilege attributes for known subjects, which may be human subjects or services and applications in an active role.

- object-related access control attributes for protected objects, which may be services or applications,

- object-related access control attributes for objects belonging to a particular service or application."

6. **Authorization Facility**

This facility approves or disapproves requests for access to security objects by security subjects. Decisions can be based on different criteria such as access control lists.

7. **Inter-domain Facility**

"This facility maps one domain's interpretation of security attributes (subject identity, object identity, authentication and authorization data) into another domain's interpretation. It helps Association Management form associations between entities in different domains."

8. **Security Audit Facility**

This facility collects, records, and analyzes event information from other security facilities. Reports are made to the security administrator or security recovery facility when appropriate.

9. **Security Recovery Facility**

This facility responds, based on the recovery policy of the security administrator, to information from the security audit facility or other designated input. Actions may include locking out some security subject or security object considered to be under a security attack. For example, a terminal or a user from which too many

incorrect passwords are received may be denied any access pending some management action.

10. **Cryptographic Support Facility**

This facility provides cryptographic services to other security facilities. These services include encryption and decryption of data, data integrity check computation, data origin authentication, nonrepudiation of origin, nonrepudiation of receipt, and key management (including, for example, key distribution and translation).

## 10.2. Security Facility Management Requirements as Specified by ECMA

The functional requirements to follow are intended to identify what is needed to assure the performance of the "top ten" security facilities listed above. This list includes the specific aspects of concern for security management for each of these facilities [ECMASEC].

1. **Subject Sponsor -- facility management requirements**

   - setting or changing of policy parameters (e.g., timer values, actions to be taken)
   - set, change, or delete rules for auditable events

2. **Authentication Facility -- facility management requirements**

   - create (enter) credentials for a particular ID
   - delete credentials for a particular ID
   - change credentials for a particular ID
   - set, change, or delete credentials change date
   - suspend credentials for a particular ID
   - set, change, or delete rules for auditable events

3. **Association Management Facility -- facility management requirements**

   - set or change rules for selecting Security Service parameter values
   - set, change, or delete rules for auditable events

4. **Security State Facility -- facility management requirements**

   - no management requirements are specified for this facility because the security state which is stored

by this facility is provided (written) by other security facilities (e.g., the Association Management or Authorization Facility) and is retrieved (read) by other security facilities according to their needs.

5. **Security Attribute Management Facility -- facility management requirements**

   - set, change, or delete privilege attributes
   - set, change, or delete control attributes
   - set, change, or delete rules for auditable events

6. **Authorization Facility -- facility management requirements**

   - definition of authorization rule sets
   - activation/de-activation of rule sets
   - set, change, or delete rules for auditable events

7. **Inter-domain Facility -- facility management requirements**

   - set, change, or delete privilege attribute translation rule
   - set, change, or delete control attribute translation rule
   - set, change, or delete rules for auditable events

8. **Security Audit Facility -- facility management requirements**

   - set, change, or delete the interchange format for audit information
   - set, change, or delete rules for audit information analysis
   - set, change, or delete specifications of alarm generating events (This would include applying these operations to event forwarding discriminators, for example.)

9. **Security Recovery Facility -- facility management requirements**

   - set, change, or delete rules for taking recovery action

- set, change, or delete parameters for corrective actions

- set, change, or delete rules for auditable events

10. **Cryptographic Support Facility -- facility management requirements**

    - secure installation of algorithms and master keys

    - deletion of inactive keys

    - deletion of active keys in recovery situations

    - set, change, or delete rules for auditable events

NIST-114A
(REV. 3-89)

U.S. DEPARTMENT OF COMMERCE
NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY

# BIBLIOGRAPHIC DATA SHEET

1. PUBLICATION OR REPORT NUMBER
NIST/SP-500/175

2. PERFORMING ORGANIZATION REPORT NUMBER

3. PUBLICATION DATE
November 1989

4. TITLE AND SUBTITLE

MANAGEMENT OF NETWORKS BASED ON OPEN SYSTEMS INTERCONNECTION (OSI) STANDARDS: FUNCTIONAL REQUIREMENTS AND ANALYSIS

5. AUTHOR(S)

ROBERT ARONOFF, MICHAEL CHERNICK, KAREN HSING, KEVIN MILLS, DANIEL STOKESBERRY

6. PERFORMING ORGANIZATION (IF JOINT OR OTHER THAN NIST, SEE INSTRUCTIONS)

U.S. DEPARTMENT OF COMMERCE
NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
GAITHERSBURG, MD 20899

7. CONTRACT/GRANT NUMBER

8. TYPE OF REPORT AND PERIOD COVERED
FINAL

9. SPONSORING ORGANIZATION NAME AND COMPLETE ADDRESS (STREET, CITY, STATE, ZIP)

Same as item #6

10. SUPPLEMENTARY NOTES

Library of Congress Catalog Card Number: 89-600778

[ ] DOCUMENT DESCRIBES A COMPUTER PROGRAM; SF-185, FIPS SOFTWARE SUMMARY, IS ATTACHED.

11. ABSTRACT (A 200-WORD OR LESS FACTUAL SUMMARY OF MOST SIGNIFICANT INFORMATION. IF DOCUMENT INCLUDES A SIGNIFICANT BIBLIOGRAPHY OR LITERATURE SURVEY, MENTION IT HERE.)

To provide for management of future interoperable multi-vendor networks, the ISO and other international organizations are currently developing management standards for communications networks based on the Open Systems Interconnection (OSI) Reference Model. This report examines current and proposed network management systems to determine both user and functional requirements for network management. It then compares the derived functional requirements to the emerging standards being developed by the ISO and others to determine where and how the requirements are being met by these emerging standards. In those cases where requirements are not being met, these deficiencies are noted.

The examination of requirements is generally restricted to those that are necessary for interoperability. These are organized and examined in six broad areas: Architecture, and the management functional areas of configuration management, fault management, security management, performance management and accounting management. This report also contains a discussion of requirements that transcend these areas and a discussion of future requirements beyond the scope of current standardization. Such requirements, while not necessary for interoperability, are useful nonetheless. Examples include automated network management and a standard operator interface. Finally this report also contains a discussion of applying OSI management to the emerging Integrated Services Digital Network (ISDN) technology.

Note: Drafts of this report were previously released under the title "Network Management Functional Requirements".

12. KEY WORDS (6 TO 12 ENTRIES; ALPHABETICAL ORDER; CAPITALIZE ONLY PROPER NAMES; AND SEPARATE KEY WORDS BY SEMICOLONS)

Automated Network Management; ISDN Management; Network Management; OSI Functional Requirements; OSI Management

13. AVAILABILITY

[X] UNLIMITED

[ ] FOR OFFICIAL DISTRIBUTION. DO NOT RELEASE TO NATIONAL TECHNICAL INFORMATION SERVICE (NTIS).

[X] ORDER FROM SUPERINTENDENT OF DOCUMENTS, U.S. GOVERNMENT PRINTING OFFICE, WASHINGTON, DC 20402.

[X] ORDER FROM NATIONAL TECHNICAL INFORMATION SERVICE (NTIS), SPRINGFIELD, VA 22161.

14. NUMBER OF PRINTED PAGES
130

15. PRICE

ELECTRONIC FORM

☆ U.S. GOVERNMENT PRINTING OFFICE 1989— 261 - 913/05117

## ANNOUNCEMENT OF NEW PUBLICATIONS ON COMPUTER SYSTEMS TECHNOLOGY

Superintendent of Documents
Government Printing Office
Washington, DC 20402

Dear Sir:

Please add my name to the announcement list of new publications to be issued in the series: National Institute of Standards and Technology Special Publication 500-.

Name _____

Company _____

Address _____

City _____ State _____ Zip Code _____

**(Notification key N-503)**

# NIST Technical Publications

## *Periodical*

**Journal of Research of the National Institute of Standards and Technology**—Reports NIST research and development in those disciplines of the physical and engineering sciences in which the Institute is active. These include physics, chemistry, engineering, mathematics, and computer sciences. Papers cover a broad range of subjects, with major emphasis on measurement methodology and the basic technology underlying standardization. Also included from time to time are survey articles on topics closely related to the Institute's technical and scientific programs. Issued six times a year.

## *Nonperiodicals*

**Monographs**—Major contributions to the technical literature on various subjects related to the Institute's scientific and technical activities.

**Handbooks**—Recommended codes of engineering and industrial practice (including safety codes) developed in cooperation with interested industries, professional organizations, and regulatory bodies.

**Special Publications**—Include proceedings of conferences sponsored by NIST, NIST annual reports, and other special publications appropriate to this grouping such as wall charts, pocket cards, and bibliographies.

**Applied Mathematics Series**—Mathematical tables, manuals, and studies of special interest to physicists, engineers, chemists, biologists, mathematicians, computer programmers, and others engaged in scientific and technical work.

**National Standard Reference Data Series**—Provides quantitative data on the physical and chemical properties of materials, compiled from the world's literature and critically evaluated. Developed under a worldwide program coordinated by NIST under the authority of the National Standard Data Act (Public Law 90-396). NOTE: The Journal of Physical and Chemical Reference Data (JPCRD) is published quarterly for NIST by the American Chemical Society (ACS) and the American Institute of Physics (AIP). Subscriptions, reprints, and supplements are available from ACS, 1155 Sixteenth St., NW., Washington, DC 20056.

**Building Science Series**—Disseminates technical information developed at the Institute on building materials, components, systems, and whole structures. The series presents research results, test methods, and performance criteria related to the structural and environmental functions and the durability and safety characteristics of building elements and systems.

**Technical Notes**—Studies or reports which are complete in themselves but restrictive in their treatment of a subject. Analogous to monographs but not so comprehensive in scope or definitive in treatment of the subject area. Often serve as a vehicle for final reports of work performed at NIST under the sponsorship of other government agencies.

**Voluntary Product Standards**—Developed under procedures published by the Department of Commerce in Part 10, Title 15, of the Code of Federal Regulations. The standards establish nationally recognized requirements for products, and provide all concerned interests with a basis for common understanding of the characteristics of the products. NIST administers this program as a supplement to the activities of the private sector standardizing organizations.

**Consumer Information Series**—Practical information, based on NIST research and experience, covering areas of interest to the consumer. Easily understandable language and illustrations provide useful background knowledge for shopping in today's technological marketplace.

*Order the* **above** *NIST publications from: Superintendent of Documents, Government Printing Office, Washington, DC 20402.*

*Order the* **following** *NIST publications—FIPS and NISTIRs—from the National Technical Information Service, Springfield, VA 22161.*

**Federal Information Processing Standards Publications (FIPS PUB)**—Publications in this series collectively constitute the Federal Information Processing Standards Register. The Register serves as the official source of information in the Federal Government regarding standards issued by NIST pursuant to the Federal Property and Administrative Services Act of 1949 as amended, Public Law 89-306 (79 Stat. 1127), and as implemented by Executive Order 11717 (38 FR 12315, dated May 11, 1973) and Part 6 of Title 15 CFR (Code of Federal Regulations).

**NIST Interagency Reports (NISTIR)**—A special series of interim or final reports on work performed by NIST for outside sponsors (both government and non-government). In general, initial distribution is handled by the sponsor; public distribution is by the National Technical Information Service, Springfield, VA 22161, in paper copy or microfiche form.

**U.S. Department of Commerce**
National Institute of Standards and Technology
(formerly National Bureau of Standards)
Gaithersburg, MD 20899