NISTIR 5309

A Context Analysis of the Network Management Domain

Conducted as Part of the Domain Analysis Case Study

Christopher Dabrowski
Susan B. Katz

U.S. DEPARTMENT OF COMMERCE
Technology Administration
National Institute of Standards
and Technology
Computer Systems Laboratory
Gaithersburg, MD 20899

NIST

NISTIR 5309

A Context Analysis of the Network Management Domain

Conducted as Part of the Domain Analysis Case Study

Christopher Dabrowski
Susan B. Katz

U.S. DEPARTMENT OF COMMERCE
Technology Administration
National Institute of Standards
and Technology
Computer Systems Laboratory
Gaithersburg, MD 20899

December 1993

U.S. DEPARTMENT OF COMMERCE
Ronald H. Brown, Secretary

TECHNOLOGY ADMINISTRATION
Mary L. Good, Under Secretary for Technology

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
Arati Prabhakar, Director

ACKNOWLEDGMENTS

We would like to acknowledge the valuable contributions of GTE Government Systems Division in Chantilly, Virginia, including Richard Kramer, Richard Law, Cathy Lytle, Michael Moore, John Rosner, Ted Ruegsegger, and Linda Sveinsson. These domain experts provided the bulk of the expertise on network management systems that made this report possible. Captain Bob Schlicher of the Air Force BMC^3 Program Office ESC/XRS (SMC/MGB) (Ballistic Missile Defense Organization (BMDO)) together with George McWilliams and Martin Cover of the Colsa Corporation contributed valuable information on network management systems. Nikhil Parekh of the Network Management Group at I-NET, Inc., in Bethesda, Maryland also contributed valuable advice. We would also like to acknowledge Sholom Cohen, Patrick Donohoe, John Leary, Spencer Peterson, and Jay Stanley of the Software Engineering Institute at Carnegie Mellon University for their help in the use of the Feature Oriented Domain Analysis (FODA) Method. Hasan Sayani of ASTEC, Inc., contributed valuable advice on several occasions. Within NIST, we would like to acknowledge the efforts of Wo Chang, Rob Densock, Elizabeth Fong, Karen Hsing, Arnold Johnson, Margaret Law, Fran Nielson, James Watkins, and others who contributed expertise or who reviewed early copies of this report.

PREFACE

The Computer Systems Laboratory (CSL) within the National Institute of Standards and Technology (NIST) has a mission under Public Law 89-306 (Brooks Act) to promote the "economic and efficient purchase, lease, maintenance, operation, and utilization of automatic data processing equipment by federal departments and agencies." When a potentially valuable technique first appears, CSL may be involved in research and evaluation. Later on, standardization of the results of such research, in cooperation with voluntary industry standards bodies, may best serve federal interests. Finally, CSL helps federal agencies make practical use of existing standards and technology through consulting services and the development of supporting guidelines and software.

Certain commercial software products and companies are identified in this report for purposes of specific illustration. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the products identified are necessarily the best available for the purpose.

TABLE OF CONTENTS

LIST OF FIGURES

EXECUTIVE SUMMARY

A key to increasing software producibility in the development of large, reliable software applications is the systematic reuse of existing software products. Domain analysis is a pivotal technique for developing reusable products that can be used to engineer software systems. The Domain Analysis Case Study was created to investigate domain analysis methods. This report describes the application of the first phase of a domain analysis effort to the domain of network management systems--software systems that manage communications networks. The first phase of the domain analysis process is called the context analysis phase. The executive summary describes the key concepts of domain analysis, provides an overview of the content of the report, and describes future work in the Case Study.

The Problem: Inefficiency in Software Development

Large-scale software development is an expensive and time-consuming process. In the coming decades, large organizations within government and industry will be engaged in developing millions of lines of code to provide needed software systems. When developed according to software engineering standards, these large systems will also have volumes of documentation. This documentation will include operational concept documents, system requirements specifications, software requirements specifications, interface requirements specifications, software design documents, and others.

One of the problems inherent in current software engineering practice stems from the focus on individual system development. Within government and industry, many software systems performing the same or closely-related functions are developed independently of each other. Duplication occurs in the development of both entire software applications as well as key subsystems that support many larger software systems. Furthermore, the developed systems are often unreliable, inefficient, expensive to maintain, and are not portable.

A Solution: Systematic Reuse of Software Products

The inefficiency of redundant development can be resolved by leveraging knowledge gained in one effort to benefit subsequent efforts. The process of leveraging knowledge to develop multiple systems is called software reuse. The benefits of software reuse are improved software quality, reduction of development time, better system integration, increased software reliability, and increased ease of system maintenance. Reuse should also reduce development costs; however, there is not yet enough experience with the application of reusable products to guarantee cost savings for each system developed.

To realize software reuse will require the investigation of new techniques and supporting technologies. One of these techniques is domain analysis.

An Enabling Process for Software Reuse: Domain Analysis

Domain analysis is a systematic approach to developing reusable products that will enable large-scale software reuse. Domain analysis can be viewed as systems analysis performed for a class of systems having similar requirements and capabilities. The class of systems is known as a domain.

Domain analysis uses many of the same analysis and representation techniques used in traditional software engineering. However, in domain analysis, these traditional techniques are used to model commonalities and differences among a class of systems in a domain, rather than modeling individual systems.

The products of domain analysis are reusable models of requirements and designs for the systems within a domain. The requirements model is called the domain model, and the design model is termed the domain architecture. The domain model describes requirements that are common to all systems within the domain. The domain model also distinguishes requirements that may vary among systems or that are unique to particular systems. The domain architecture provides a generic design for systems in the domain that can be adapted to provide designs for individual systems.

The domain model and domain architecture can be used to systematically guide the development of individual systems that fall within a domain. The domain model is used to derive individual system requirements. The domain architecture is used to derive system designs. The domain model and domain architecture also provide a basis for developing a taxonomy of software components in a domain. This taxonomy can be used in a reuse library system to help locate components that can be potentially reused both during system development and system maintenance.

The use of the products of domain analysis to specify individual systems that fall within a domain has the following benefits:

- Improving quality of system specifications by helping to assure that they are correct and decreasing chances of omissions and weaknesses.
- Reducing development time by making it unnecessary to completely redevelop system specifications.
- Facilitating integration of systems through use of standard descriptions of external system interfaces.
- Increasing software reliability by aiding validation and verification of developed systems. (This is accomplished by providing another level of traceability known as the domain level.)

- Supporting system maintenance through use of standardized models of system requirements and designs to guide system modification.

Domain analysis is useful principally to systems analysts, software engineers, and system maintainers. However, the identification and description of software domains is also of benefit to planners and managers for obtaining a better understanding of an organization's current and potential software resources.

The Domain Analysis Case Study

In recent years, several domain analysis methods have been proposed. However, these methods have not been sufficiently tested for widescale use. Both the products and processes of domain analysis need to be studied and, where appropriate, improved. The methods then need to be transitioned from the research environment into practical use. These are the tasks addressed by the Domain Analysis Case Study.

The Domain Analysis Case Study is currently investigating the use of a particular domain analysis method. This method, called the Feature Oriented Domain Analysis (FODA) method (described in more detail in sec. 1.5 of the report), was developed by the Software Engineering Institute at Carnegie Mellon University. The FODA method consists of three phases:

- The context analysis phase, which provides a formal description of the scope and boundary of the domain to be analyzed.

- The domain modeling phase, in which the domain model is developed.

- The architecture modeling phase, in which the domain architecture is developed.

The context analysis sets the stage for the subsequent phases of the FODA process.

The First Step: Context Analysis

The success of domain analysis, in part, hinges upon being able to define software domains, to describe their boundaries, and to define their interfaces. In order to do this, the context of a domain must be considered. Context refers to the environment in which a class of software systems operates. The context analysis identifies the external entities in a domain's context with which systems in the domain interact. External entities may include other software systems, data stores, equipment, and human beings. Analysis of the interactions between a domain and its external entities helps determine the scope and boundary of a domain and provides a basis for specifying external interfaces for systems in the domain.

Another important goal of the context analysis is to describe variability in key characteristics of external entities that affect systems in the domain. In particular, it is important to describe how variability in external entities:

(1) Causes variation in the interactions between the domain and its external elements, and

(2) Influences variation in requirements for systems within the domain.

Understanding how variability in the context affects the domain provides additional information that can later be used for specifying system interfaces. Understanding the affect of variability in context also provides a basis for identifying both commonalities and variabilities in a domain's systems. This information, in turn, is a key to describing reusable specifications for a class of systems.

The context of a domain is represented using context diagrams. Context diagrams show interactions and data flows between a domain and its external entities. The context diagrams are accompanied by text that documents how variability in the context affects both variability in data flows and variability in requirements for systems within the domain.

The Network Management Domain

The domain to which the context analysis has been applied is the network management domain. A network management system monitors and controls a communications network to ensure its continuous operation, efficiency, and integrity. This domain was chosen because network management systems provide essential services to communications networks. As such, this domain should provide broad opportunities for applying domain analysis to develop reusable components.

The results of the context analysis are described in sections 2 through 6 of the report. Section 2 provides a description of communications networks. The functions of network management are described in section 3 of the report. These functions are organized into five functional areas. For the purposes of the context analysis, these functional areas may be regarded as subdomains of the network management domain. The five subdomains are:

- **Fault Management**

  The detection, reporting, diagnosis, correction, and prevention of faults that occur in a communications network.

- **Performance Management**

  The control of the quality, effectiveness, and efficiency of network communications.

- **Configuration Management**

  The tracking and control of network resources and their current and potential connections.

- **Accounting Management**

  The collecting and storing of information on the utilization of network resources for the purpose of generating bills.

- **Security Management**

  The monitoring and controlling of access to network resources.

Section 3 also provides an overview of the organization of network management software systems. A discussion is also provided of how these software systems implement network management functions. Then, using the descriptions of the communications network and network management as a basis, section 4 of the report formally defines the context of the network management domain according to the FODA method. A context diagram is provided that describes interactions between network management and its external entities.

Alarm Surveillance Systems

Within fault management, the report focuses on the critical subdomain of alarm surveillance. Alarm surveillance systems monitor the communications network and report faults that affect the operation of the network. Alarm surveillance systems will be the subject of the domain modeling phase of the FODA method. GTE Government Systems, Inc., of Chantilly, Virginia is providing expertise in alarm surveillance systems necessary to develop the domain model.

In section 5 of the report, the alarm surveillance subdomain and its context is defined in detail. Separate context diagrams are provided for the fault management and alarm surveillance subdomains. Section 5 goes on to describe the major external entities in the context of alarm surveillance systems. The most important external entity is the communications network itself. A discussion is provided of the variability in important characteristics of communications networks. The discussion focuses on how this variability affects alarm surveillance systems. Among the characteristics discussed are the size of the network (number of devices), the geographic scope of the network, the configuration of the connections between the devices (network structure), the types of devices, and the use of protocols or protocol suites for communications between devices. For each identified variability, a description is provided of how the variability affects data flows across domain boundaries and how it affects requirements for alarm surveillance systems.

In section 6, the report provides an assessment of the potential for developing reusable components in the alarm surveillance subdomain. This assessment is a critical aspect of the context analysis phase since it helps determine the feasibility of developing reusable components in the subsequent domain modeling effort. The assessment also provides an estimate of the economic value of developing components within the alarm surveillance subdomain and the potential payoff for future software engineering efforts.

Section 7 provides a glossary of terms used in the report. This glossary will be further evolved during the subsequent domain modeling stage. Section 8 lists the references.

Future Work

This is the first report of the Domain Analysis Case Study. Subsequent reports will describe the domain modeling and architecture modeling phases. The products of the Case Study will be tested in actual system development. The results of the domain modeling phase, the architecture modeling phase, and product testing will be provided in a final report on the Case Study. The final report will make recommendations on using domain analysis to support software reuse. The final report will also recommend specific characteristics a domain analysis method should have to be useful for real-world software development.

1. INTRODUCTION

A key to increasing software producibility in the development of large, reliable software applications is the systematic reuse of existing software products. Domain analysis is a pivotal technique for developing reusable products that can be used to engineer software systems. The Domain Analysis Case Study was created to investigate domain analysis methods. This report is a product of the Domain Analysis Case Study. This report describes the application of the first phase of a domain analysis effort to the domain of network management systems--software systems that manage communications networks. The first phase of the domain analysis process is called the context analysis phase.

1.1 Background

The Software Producibility Manufacturing Operations Development and Integration Laboratory (Software Producibility MODIL) was established at the National Institute of Standards and Technology (NIST) in 1992. The Software Producibility MODIL was one of four MODILs instituted at national laboratories by the Ballistic Missile Defense Organization (BMDO).[1] The purpose of the MODILs was to investigate new technology that could be transitioned from the research community to the BMDO. The objective of the Software Producibility MODIL was to investigate new software technology that could be used to improve software development within BMDO.

The initial focus of the Software Producibility MODIL was software reuse. Domain analysis is a process that is widely regarded as important in enabling software reuse. The application of domain analysis results in the development of reusable software products that can be used to engineer new systems. The Domain Analysis Case Study was created by the Software Producibility MODIL to investigate domain analysis methods.

The first phase of the Domain Analysis Case Study was completed by the Software Producibility MODIL under the sponsorship of BMDO. Because of the potential application of the results of this work to many different domains, the Domain Analysis Case Study has been continued and expanded by NIST.

1.2 The Purpose of This Report

The Domain Analysis Case Study is currently investigating the use of a particular domain analysis method. This method, called the Feature Oriented Domain Analysis (FODA) method (described in more detail in sec. 1.5), was developed by the Software Engineering Institute at Carnegie Mellon University. The purpose of this report is to

---

[1]The BMDO was formerly known as the Strategic Defense Initiative Organization (SDIO).

describe the results of the application of the first phase of the FODA method--the context analysis phase. The context analysis defines the scope and boundary of a domain within which reusable software products will be developed. The domain to which the context analysis has been applied is the network management domain. This domain is concerned with software systems for management of communications networks.

This report will focus on a portion of the network management domain, called the target subdomain, that has been selected for more detailed analysis. The target subdomain is alarm surveillance. Alarm surveillance systems detect and report problems affecting the operation of the communications network. In the application of subsequent phases of the FODA method, comprehensive models of alarm surveillance systems will be produced. These models will serve as reusable components that can be used to develop requirements and designs for individual alarm surveillance systems. The results of the application of subsequent phases of FODA will be described in the final report for the Domain Analysis Case Study. The final report will recommend the characteristics a domain analysis method should have to be used for real-world software development.

1.3 The Audience for This Report

The audience for this report includes, but is not limited to:

- BMDO and other government personnel and contractors who are interested in domain analysis methods and are, in particular, interested in applying the FODA method.
- Experts in network management and developers of network management systems within government and industry who may wish to use the products of the Case Study.
- Purveyors of domain analysis methods who wish to gain further understanding of the application of domain analysis methods.

This report will be available to all interested parties. The authors welcome comments and cooperative follow-on research efforts.

1.4 The Domain Analysis Case Study

The Domain Analysis Case Study is examining the potential use of domain analysis methods. The Case Study will determine the usefulness of domain analysis methods in developing reusable components.

1.4.1 Domain Analysis Methods

A domain can be defined as a class of systems with similar requirements and capabilities. In this report, the domain of interest is network management; the class of systems is network management software systems.

Domain analysis can be thought of as systems analysis applied across the systems within a domain. It is a process by which information used in developing software systems within a domain is identified, captured, and organized so that it can be reused to engineer new systems [PRIE90]. The product of a domain analysis effort is a set of reusable components that can be adapted to develop application systems within a specific domain. Reusable components, as defined in [DoD92] and [KATZ93], include requirements, designs, and code modules.

In the last several years, research in methods for domain analysis has resulted in the emergence of several methods in addition to FODA, including [DISA93], [HOLI91], [PRIE91], [SPC92], and [STARS93]. These methods are currently being applied to a variety of domains. An overview of the state of the art in domain analysis methods is provided in [DABR93] and [WART92].

1.4.2 The Goals of the Domain Analysis Case Study

The Case Study will do the following:

(1) Test the viability of domain analysis as a process to aid software reuse,

(2) Identify the processes and products necessary to meet the goals of domain analysis in developing reusable components, and

(3) Identify the characteristics that determine a method's usability, i.e., evaluate the ease with which developers may use the method to create the products of domain analysis.

The strategy of the Case Study is based on direct "hands on" experience in the use of the FODA method on a practical problem. During the Case Study, key aspects of the use of the method will be documented. This includes documenting the use of selected FODA procedures to develop domain analysis products. It also includes documenting the use of domain analysis products to engineer applications within the domain.

1.4.3 The Selection of the FODA Method for the Case Study

The FODA method was selected for use in the Domain Analysis Case Study because of its potential for practical use in developing reusable components. A pivotal aspect of FODA is the use of feature analysis. Features are user-visible characteristics of systems within a domain (described more fully in sec. 1.5.2). The resulting feature model can be utilized to assist application system developers in using reusable components to create new systems.

In addition, the FODA method prescribes processes and products that are similar to those in other domain analysis methods. The selection of FODA allows the investigation of aspects of a domain analysis method that are shared by other methods. The results of the Case Study will be provided to the purveyors of the FODA method to allow them to further develop the method.

1.4.4 The Selection of the Network Management Domain

The network management domain was selected in the Case Study to meet the current and future needs within BMDO. Network management systems will be required to support BMDO software systems using communications network services. Network management is an example of a horizontal domain--a domain that provides information or services to more than one other domain. It is expected that several BMDO program elements will undertake the development of network management systems in the near future.

For the purposes of the Case Study, it is necessary to reduce the size of the domain analysis effort to an appropriate level. This is accomplished by selecting a smaller area within the network management domain for more detailed examination. As indicated in section 1.2, this area--termed the target subdomain--is alarm surveillance. Alarm surveillance systems, described in detail in section 5, perform the most critical functions of network management [ADAMS91]. The alarm surveillance target subdomain will be analyzed in detail in the subsequent phases of the FODA process. The existence of a core set of features common among systems in the target subdomain, coupled with variations between systems within the subdomain, provides ample opportunity for investigating the use of domain analysis methods.

The definitions used in describing the domain are based on the OMNIPoint specifications produced by the Network Management Forum [NMF92a], [NMF92b], [NMF92c], OSI International Standards Profiles, as well as other pertinent literature. GTE Government Systems is providing the domain expertise necessary to support the domain analysis effort.

1.4.5 Future Products of the Case Study

The information produced during the course of the Case Study will be used to recommend requirements for domain analysis methods. This will lead to the evolution of domain analysis processes that enable the development and application of more effective reusable components. These recommendations will be presented in the Case Study's final report. In addition, the products of this particular domain analysis effort will be made available for review by interested parties. These products are described in the next section.

1.5 Overview of the FODA Process

The FODA method and its current applications are described in [KANG90], [PETE91], and [COHEN92]. At present, the FODA method consists of three phases. The first phase--called the context analysis phase--is the subject of this report. The second and third phases--the domain modeling and architecture modeling phases--will be briefly described to provide the reader with background.

1.5.1 The Context Analysis Phase

The context analysis phase establishes the scope and boundary of a domain of interest. The context analysis specifies what major functions and capabilities are within the domain and what functions and capabilities are excluded from the domain.

In the FODA approach, context describes the circumstances, situation, or environment in which a particular domain exists. The context analysis identifies important elements that are external to the domain. It defines the relationships and interactions that exist between the domain and these external elements. An important part of a context analysis is to understand how the function and behavior of external entities may vary. It is important to understand how this variation (1) affects interactions between the domain and its external elements, and (2) in a broad sense, influences requirements for systems within a domain. Understanding variation in context provides a basis for detailed analysis of variability among systems in the domain during the subsequent domain modeling and architecture modeling phases.

The context of a domain is represented in a context model. Two principal aspects of the context model are structure diagrams and context diagrams.

- A structure diagram is an informal block diagram in which a domain of interest is placed relative to higher, lower, and peer-level domains. Higher-level domains are domains to which the domain of interest provides services necessary to carry out the higher-level domain's functions. Lower-level domains provide services to the domain of interest. All other related domains are considered to be peer domains. A structure

diagram may consist of a series of layers, each of which contains domains that provide and receive similar types of services.

- A context diagram shows data flows between the domain of interest and the external entities with which it communicates. The context diagram should be accompanied by documentation on how variability in the context affects variability in the data flows between the domain and its external entities.

Another product of the context analysis is the Domain Dictionary. The definitions in this dictionary provide a common understanding of important domain terms and concepts that are used in the context model.

The principal product of the context analysis--the context model--should be considered a reusable component that can be used by application developers to identify whether their application falls within the scope of the network management domain and the alarm surveillance target subdomain. The context model should also be used to help the application developer understand the relationship between an alarm surveillance application and systems in external domains. This is necessary to ensure proper use of the domain model and architecture and prevent wasted effort. To aid the application developer, this report provides instructions on how the context model is to be used.

1.5.2 The Domain Modeling Phase

The domain modeling phase is directed towards describing the capabilities of, and requirements for, systems within a domain of interest whose scope was defined in the context analysis phase. The product of this phase is a domain model consisting of the following three components:

- The entity-relationship model describes the major domain entities, their internal structure, and relationships between entities in the domain of interest. The entity-relationship model represents the essential domain knowledge that is necessary to create the feature model and functional model described below.

- The feature model represents the end user's view of the capabilities of systems in the domain and describes visible aspects of systems that directly affect end users. It contains features that are mandatory for all systems in the domain as well as features that are optional or that represent alternative choices. The feature model can be used by application developers to define capabilities that their applications should have. The feature model parameterizes other models produced by the domain analysis.

- The functional model describes the functions, data flows, and control flows of systems *within the domain of interest*. (This is in contrast with the data flows *between the domain of interest and external entities in the context model*, discussed in

sec. 1.5.1.) The functional model identifies those capabilities and requirements that are common among the systems and describes the parameters by which systems can vary.

During domain modeling, the Domain Dictionary that was begun during the context analysis phase is extended and refined. The domain model is a reusable component that can be used to determine requirements for applications within the domain.

1.5.3 The Architecture Modeling Phase

During the architecture modeling phase, a domain architecture is developed. The domain architecture describes a high-level design or "blueprint" for implementing software applications in a domain. The domain architecture is also a reusable component. The entire architecture, or selected parts of it, can be adapted to develop designs for specific applications within a domain. To accomplish this, the domain architecture may contain information on how the requirements described in the domain model are to be implemented as software systems. The domain architecture describes how requirements in the domain model are mapped to specific design components within the architecture. Use of a domain architecture eliminates the need for having to develop designs for individual software applications from scratch.

1.6 How to Read This Report

The organization of the report is as follows:

- Section 2 provides an overview of communications networks. This discussion provides the important background necessary to understand the context of network management.

- Section 3 defines network management, describes its principal functions, and discusses network management software systems. External functions and activities that use services provided by network management are identified. A FODA structure diagram for network management is provided.

- Section 4 presents a more formal description of the context of network management according to the FODA method. A context diagram for network management is provided.

- Section 5 defines the target subdomain--alarm surveillance--and describes its context. A context diagram is also provided for the alarm surveillance subdomain. The section discusses variability in the context of the alarm surveillance systems and its impact on the target subdomain.

- Section 6 discusses the future development of reusable components in the target subdomain. The section discusses important issues that need to be addressed in preparing for domain modeling.

- Section 7 contains the domain dictionary.

- Section 8 provides the references.

This report may be read in different ways by different parts of the intended audience. Persons unfamiliar with communications networks should read sections 2 and 3. Developers of network management systems may wish to focus on sections 3 through 6. Persons with a general interest in domain analysis and purveyors of methods should probably read the entire report.

Underlined words in the text of this report are found in the domain dictionary in section 7. These words are underlined on their first occurrence and appear without underlining afterwards.

2. COMMUNICATIONS NETWORKS

This section provides background on communications networks, which are the entities managed by network management systems. The purpose of the overview is to provide an understanding of the environment in which network management systems operate. This section lays a foundation for the definition of network management in section 3 and for the formal description of its context in section 4. It also provides background necessary to understand the context of the target subdomain, described in section 5. Readers familiar with communications networks may wish to skip this section and proceed to section 3.

As the usage of computer equipment has increased, many organizations find themselves with many diverse types of equipment. This equipment may have been purchased to perform various types of tasks and may come from different manufacturers. Furthermore, there is often a need to integrate these resources. This integration may serve to coordinate or pool the capabilities of the machines or to allow communication between physically remote pieces of equipment.

The motivations for connecting equipment via a network to permit data transmission include:

- **Remote Access by Terminals.**

  Both data and computer programs may be stored on, and used from, a remote data store or computer.

- **Computer Resource Sharing.**

  Specialized resources, such as printers or platforms providing greater processing power, can be shared. Distributed processing then becomes possible.

- **Computer Resource Backup.**

  Access to multiple computer resources allows alternative ways of providing the required functions. Redundant links, redundant data stores, and redundant processing are thus feasible for applications that require a high probability of success and availability for operations.

2.1 Definition of a Network

A network is a set of devices that are connected by one or more forms of transmission media. The devices may include computers, terminals, printers, robots, data stores, and other computer or communications resources. A transmission medium is any mechanism that supports propagation of digital or analog signals. Examples of transmission

media are cables such as leased lines from common commercial carriers, fiber optic cables, satellite channels, etc. Physical transmission lines are also sometimes referred to as circuits, channels, or trunks [TANEN88].

2.1.1 Size and Scope of Networks

In its simplest form, a network can consist of a group of devices attached to a single cable or joined to each other by separate links. Geographically, a network can be located at a single physical site (such as a building) or span a considerable area. A group of networks at distant points can be linked together into a larger network by cable or satellite links. Three classes of networks based in part on geographic extent are recognized:

- Local Area Networks (LANs) are networks that are usually limited to a single building or closely-spaced group of buildings.

- Metropolitan Area Networks (MANs) usually denote networks located in a single city. They may be composed of LANs or even smaller MANs within the same city.

- Wide Area Networks (WANs) may cover several sites that are geographically distant, such as sites located in different cities or even different continents.

A complex network can consist of WANs which span continents or geographic regions within continents and connect smaller, more localized MANs or LANs. Within BMDO, LANs are used to provide intrabase communications. Inter Site Networks (ISNs), similar to WANs, are generally used to provide communications between different BMDO bases. In this report, network size is defined as the total number of devices that must be managed within the network. The geographic area encompassed by the network will be called its scope.

2.1.2 Structure of Networks

In this report, the structure and geographic layout of an individual network, or a group of networks that form a larger network, will be referred to as a network topology. Topology influences the organization of a network management system across a network, as will be described in section 3. According to [TANEN88], there are two general types of designs for networks:

(1) Point-to-point networks, and

(2) Broadcast networks.

A point-to-point network contains many cables or leased lines, each one connecting a pair of devices. Devices may communicate directly with other devices to which they are connected

or indirectly through intermediary devices. When a message (often called a packet) is sent across the network, it is received at each intermediary device in its entirety. It is then stored at that device until the required output line is free, and then it is forwarded. Thus a point-to-point network is also referred to as a store-and-forward or a packet-switched network.

In the case of broadcast channels, a single communication channel is shared by all the devices on the network. Packets sent out by any device are received by all others. The address field within the packet designates the intended recipient so the message is ignored by other devices. Broadcast systems generally also allow packets to be addressed to all destinations by using a specifically designated code in the address field. Some broadcast systems also support multicasting: transmission to a subset of devices on the network. Most LANs, and a small number of WANs, are broadcast networks.

Broadcast networks may have several different topology types such as ring, bus, and satellite. Point-to-point networks include star, connected ring, or irregular topologies. Figure 2.1 ([STAL88] and [TANEN88]) depicts four of the topologies listed above.

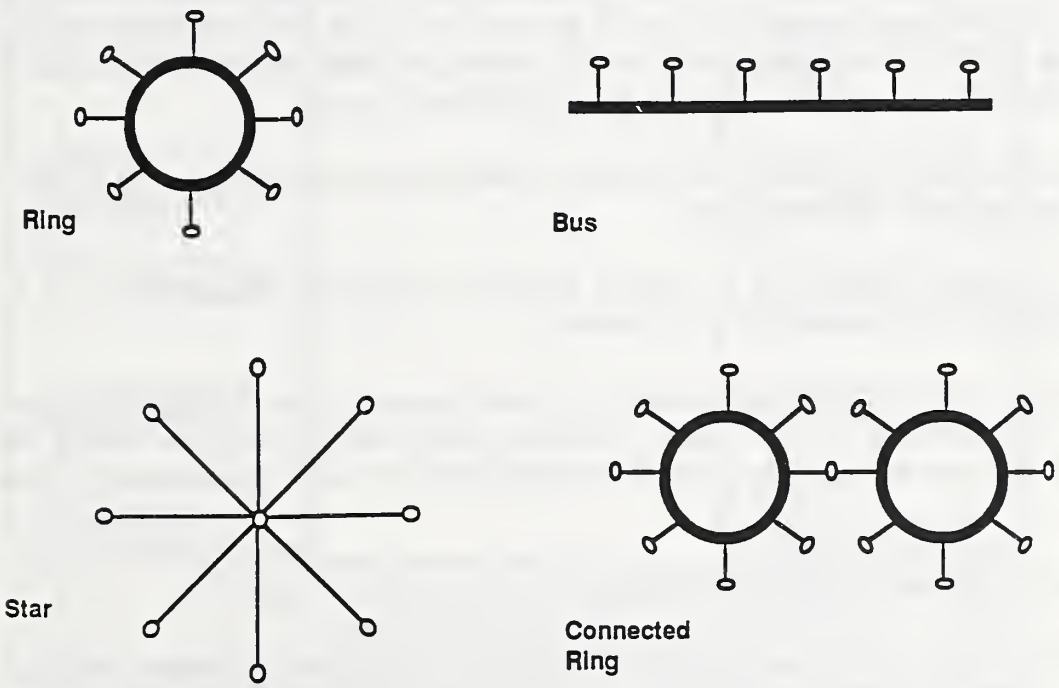

Figure 2.1: Various Network Topologies

2.1.3 Types of Network Devices

A network device is any device that can send, receive, or transmit digital signals across a network. In general, three classes of devices can be distinguished:

- o End-System Devices are platforms containing hardware and possibly software systems. These devices may contain resources accessed via the network, such as printers. Or they may access the resources of the network, as is the case for terminals. In addition to computers, terminals, and printers described above, end-system devices may also be radar systems or platforms that host weapons systems software.

- o Intermediary Devices are devices that influence the transmission path in some way. Examples of intermediary devices, from the simplest to the most sophisticated, are bridges, routers, and gateways. Bridges are devices that store and forward complete packets from one physical address on one network to another on a different network. Bridges are designed to connect only two networks (often LANs) that use the same protocol. (See sec. 2.2.1 for a discussion of network protocols.) Routers are devices that make decisions about which of several paths network traffic will follow. Gateways are a special kind of sophisticated router. Gateways are dedicated computers that connect two or more networks and make routing decisions regarding the traffic between those networks. In contrast to bridges and routers however, gateways may connect networks that use different protocols.

- o Link Devices are pieces of transmission media (defined at the beginning of sec. 2.1) that propagate digital signals.

Another important characteristic of network devices is bandwidth, defined as the rate at which information is transmitted in bits/second.

Within each of the general classes of network devices, there are many categories of devices. The existence of these classes provides a partial basis for variation in both the context of network management and in the requirements for network management systems.

2.2 Network Communications Services

The purpose of a network is to facilitate the transmission of information among remote points by providing a set of services for this purpose. Network communications services are the capabilities provided by software systems and devices that enable transmission from one point to another point within a network. These services transform a message on an end-system device into digital signals, determine a route to the desired destination point or points, and transmit the signals across the transmission medium to another device or group of devices. At the destination point, signals are transformed into a

form that is readable to the receiving device. Network communications services may be directed at transmission of three basic types: data, voice, and image. At present, within BMDO, the transmission of data and voice is of great concern.

2.2.1 The Use of Protocols

Network communications services involve the use of protocols. Protocols are sets of rules that specify when messages are to be passed, how they are to be routed, their exact formats, and what is to be done in the event of various types of errors. To reduce the complexity of providing communications services, most networks organize their protocols in a series of layers, each one built upon its predecessor [COMER91]. Each layer uses the services of the layer below it and in turn provides a value-added service to the layer above it. Thus each layer is shielded from having to know the details of how those layers with which it interfaces perform their functions.

The number of layers and the function of each layer varies from one implementation to another; however, the general concept remains the same. Figure 2.2 depicts the layering of protocol services.

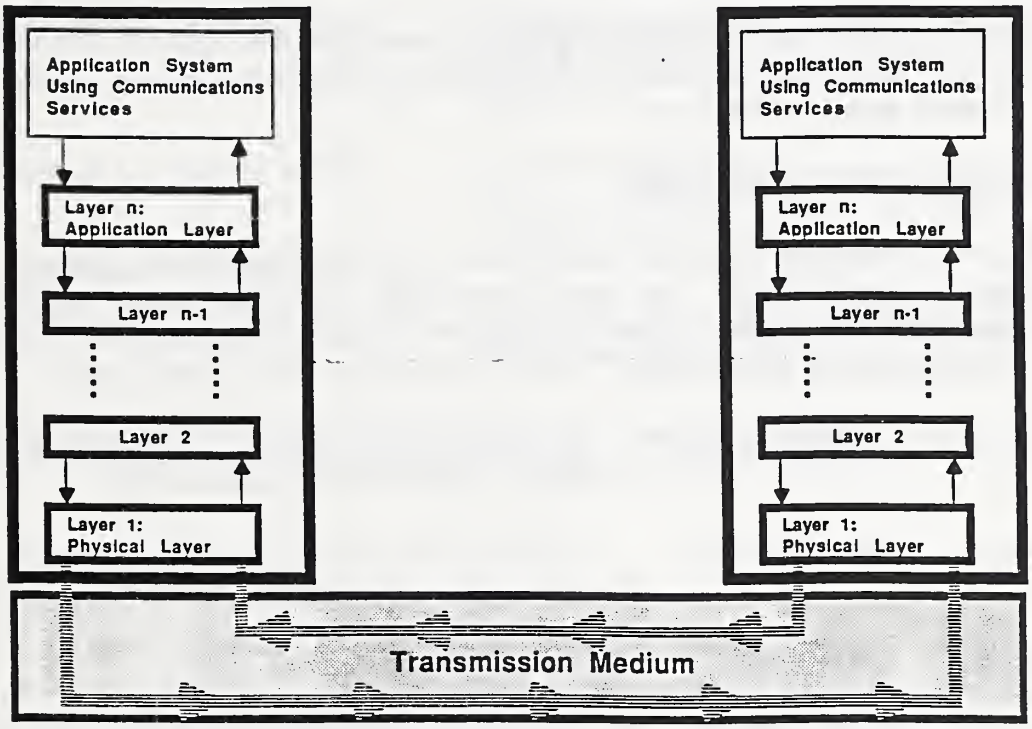

Figure 2.2: Layered Communications Services

The top layer is the Applications Layer that provides services directly to application systems using communications services. The bottom layer is the Physical Layer that transmits digital signals across the communications channel. The intermediary layers perform such functions as transforming messages into forms suitable for transmission, determining addresses of destinations, determining routes across the network, establishing and releasing connections, and performing reliability and error checking functions.

In figure 2.2, each layer on one side carries on a virtual conversation with the corresponding layer on the other side. The protocol(s) defined for each layer are used to carry out the conversation. The virtual communication occurs between peer layers on remote devices. However, the actual flow of data and control occurs between the layer interfaces on each device. In figure 2.2, when an application needs to transmit data, that data is passed from the application layer downward. Each layer passes data and control to the layer immediately below it. Finally, Layer 1 sends the message to the communications medium, where the actual transmission to the remote device occurs. When the message is received at the remote end of the transmission medium, it is passed from Layer 1 upwards through the protocol layers, to the application that will be using the information.

2.2.2 Standards for Network Protocols

Suites of standard network protocols have been developed that organize protocol communications by specifying the role of each layer and the interfaces between layers. Two prominent standard protocols suites are:

- **The Internet Suite of Protocols.**

  This suite of protocols is usually referred to as TCP/IP (Transmission Control Protocol/Internet Protocol). The TCP/IP suite grew out of early research by the Advanced Research Projects Agency (ARPA) in DoD and has since spread to many areas of government and industry. The TCP/IP suite consists of four layers:

Layer 4:	**APPLICATION**	Provides services directly to users such as file-transfer protocols and electronic mail.
Layer 3:	**TRANSPORT**	Provides end-to-end communication between applications and verifies correct packet arrival.
Layer 2:	**INTERNET**	Handles packet routing and integrity.
Layer 1:	**NETWORK INTERFACE**	Provides an interface to network hardware or device drivers. Also called the Data Link Layer.

o **The Open System Interconnection (OSI) Suite.**

The OSI suite is the product of over a decade of work in the international standards arena. OSI was developed under the auspices of the International Organization for Standardization/International Electrotechnical Committee (ISO/IEC). The OSI suite consists of seven layers described in [ISO/IEC 10040] and [STAL93]:

Layer 7:	**APPLICATION**	Provides services directly to users such as file-transfer protocols.
Layer 6:	**PRESENTATION**	Defines and transforms the format of data to make it useful to the receiving application.
Layer 5:	**SESSION**	Establishes, manages, and terminates connections between applications, and provides checkpoint recovery and security mechanisms.
Layer 4:	**TRANSPORT**	Ensures error-free, in-sequence exchange of data between end points.
Layer 3:	**NETWORK**	Provides routing (path control) services to establish connections across communications networks.
Layer 2:	**DATA LINK**	Provides reliable transfer of data across physical links, error and flow control, link level encryption and decryption, and synchronization.
Layer 1:	**PHYSICAL**	Provides for the transmission of unstructured bit streams over the communications channel.

Both protocol suites also provide special protocols for conveying network management information. Network management protocols will be described in section 3.

2.2.3 Proprietary Protocols

In addition to standard protocols, there exist a number of proprietary or vendor-specific protocols, created by private companies for the network devices they manufacture. Many networks contain large collections of such devices. Because these individual collections of devices use proprietary protocols, they must often be controlled by management tools, called proprietary element management systems, supplied by vendors for

this purpose. These systems serve as interfaces to proprietary portions of the network.[2] The consequences of having existing proprietary element management systems, which often predate newer parts of a network, will be discussed later in this report.

2.3 Definition of the Managed Network

In this report, a particular network, the set of network communications services provided on the network, and any other services or resources that are involved in the operation of a network will be referred to as the managed network. A managed network will be managed by a specific network management system. The set of individual elements of the network that are subject to network management will be considered network resources. End users are either (1) persons who directly or indirectly use network communications services that can be accessed through an end-system device, (2) pieces of equipment, such as satellites, that utilize communications services, or (3) software systems, such as engagement planning systems, that reside on an end-system device and use such services. The managed network is the most important aspect in the context of network management.

---

[2]An element management system may be either proprietary or standards-based. Standards-based element management systems manage devices that are based on either OSI or TCP/IP.

3. MANAGEMENT OF COMMUNICATIONS NETWORKS

In this section, the major functional areas of network management are defined. The role of network management software systems in managing communications networks is described. Network management is distinguished from other functions necessary to support the operation of a communications network. A FODA structure diagram summarizes the relationship between network management, the managed network (defined in sec. 2.3), and the other functions needed to support the network. In section 4, the information in this section and in the preceding section will provide the basis for defining the context of network management.

3.1 Definition of Network Management

Network management is the discipline that describes how the managed network is monitored and controlled to ensure its continual operation, efficiency, and integrity. Network management both provides services to, and utilizes the services of, the managed network. Network management provides the services needed to monitor and control communications activity on the managed network. Network management serves the network by helping to keep it running. A network management system also uses the physical network connections and network protocols in order to transmit messages. (This will be described further in sec. 3.2.) The set of services provided by network management has been defined by OSI [ISO/IEC 7498-4]. These services are organized into five management functional areas (MFAs), described below. For the purposes of the context analysis, these MFAs may be regarded as subdomains of network management.

3.1.1 Fault Management

Fault Management is the detection, reporting, diagnosis, correction, and prevention of faults and fault conditions. A fault is a malfunction or abnormal pattern of behavior that is causing or will cause, an outage, an error, or degradation of communications services. A fault may be either a physical malfunction in a piece of equipment or a control malfunction in a software or firmware system. A fault condition is a set of circumstances under which a fault is likely to occur.

A particularly important function of fault management is to monitor the network to detect faults and fault conditions. Monitoring may be accomplished in different ways. Individual devices on the network may be periodically queried--or polled--to determine their status. Devices may emit alarms--messages reporting that a fault has been detected. Traps are messages emitted by devices indicating that, based on current behavior patterns and indications, a fault is likely to occur. A combination of polling, alarm generation, and trapping may also be used. Monitoring is an important function of alarm surveillance, the target subdomain described in section 5.

An outage, an error, or a rapid performance degradation may be caused by single or multiple faults (or transmission media degradations). On the other hand, a single fault, such as a broken cable that links a group of devices in a point-to-point network, may manifest itself as multiple device outages. These outages may possibly be noted by different users at different times. A fault may be isolated by narrowing down the location of the fault to a particular device, or devices, in which the fault occurred. Once isolated, the fault may be corrected by replacing the defective network resource, by reconfiguring the network to bypass the defective network resource, or by repairing the resource. If the fault or fault condition is noted early by the network administrator or network management system, it may be resolved or circumvented before it actually becomes a problem for the end user.

3.1.2 Performance Management

Performance management is concerned with the quality, effectiveness, and efficiency of network communications. It is primarily concerned with measuring and quantifying the responsiveness, availability, and utilization of individual network resources and of the network as a whole. Performance management activities include monitoring performance parameters that describe the efficiency of the network, such as response time and traffic throughput. They also include reporting the results of monitoring activity, adjusting and tuning network resources to improve efficiency, and evaluating the results of tuning actions. Examples of specific performance management activities are:

- Near real-time reporting of the utilization, response time, or throughput for individual network resources, such as links or routers.

- Long-term collection of statistical data on the utilization, response time, or throughput for all the network resources on a particular link or set of links. Such data may be used for trend analysis and capacity planning.

- Assisting in dynamic network optimization (by balancing traffic loads across all network resources on alternative communications paths that link points in a network).

3.1.3 Configuration Management

Configuration management is the tracking and control of network resources and their current and potential connections. It includes creating and maintaining an accurate inventory of:

- All resources associated with the network.

- Each network resource's operating characteristics, described by values of specific internal settings or variables.

- Each network resource's logical and physical connections to other resources.

- The network topology or any portion of it.

Configuration management controls both the information about the network and the physical configuration of the network itself. Information about the network includes the network's current configuration, allowable configurations, the history of past configurations, and the status of managed devices (including transmission media). This information may be maintained in a configuration database. The physical configuration, called the network configuration, consists--at any point in time--of the particular devices on the network, the operating characteristics of these devices, and the logical and physical connections between devices. Configuration management provides a means to change the network configuration.

Configuration management supports the other functions of network management. Configuration management provides information about network resources in support of the fault management, performance management, and accounting management functions. By controlling the network configuration, configuration management can be used to perform corrective measures such as rerouting traffic to bypass a faulty link. Similarly, a network configuration can be altered to balance traffic loads and optimize the performance of the network.

3.1.4 Accounting Management

Accounting management is the collecting and storing of information on the utilization of network resources. This information is used to generate bills for end users or end-user groups. Accounting management tracks network resource utilization on a per-user and/or per-group basis and enforces account limits. Accounting management may also provide information such as reports on utilization for different categories of end users. Such reports can be used to perform trend analysis and to predict utilization patterns and costs.

3.1.5 Security Management

Security management is the process of monitoring and controlling access to network resources. This includes granting authorization for use of, and controlling ongoing access to, network resources. In a network with stringent security requirements, encryption procedures may be used. The associated encryption key distribution and accounting functions are part of security management.

Security management procedures include monitoring usage of secure network resources, recording information about usage of secure resources, as well as detecting and reporting attempted or successful violations. Security management procedures may involve distribution of security-relevant information such as security violation reports. Security

reports may be used to track long-term security problems and to assist in trend analysis of security violations. Security management also includes the maintenance of the physical integrity of the network.

3.2 Network Management Systems

For the purposes of this report, a network management system is a software system that performs the functions of network management for a communications network. It should be noted that not all the functions described in section 3.1 are automated in software systems. For instance, within fault management, network monitoring functions are generally performed by network management systems. However, many aspects of isolating and correcting faults are built into the hardware (such as specialized build and test equipment) or performed manually by human beings.

3.2.1 How the Functions are Executed: Managers and Agents

In a TCP/IP or OSI network as well as in many networks that use proprietary protocols, a network management system can be characterized at an abstract level as having 2 parts:

(1) Managers, software systems that control and coordinate the functions of network management for a managed network, and

(2) Agents, software systems that reside on managed devices and are responsible for network management functions on those devices.

Managers reside on network management stations. A network management station contains, in addition to the manager, the workstations, displays, data processing software, and other equipment necessary to assist the network administrator in managing the network. In a larger network, the management station may be located at a network operations center or network control center. The network operations center is an installation that contains the personnel, equipment, and other network resources needed to operate and maintain the portion of the network within its jurisdiction.

Agent systems monitor the device's internal status. To do this, an agent system may receive information from background processes. For example, a background process running on the device may periodically check internal parameters that describe a device's status. Such a process may compare current parameter values with predetermined threshold values and provide the results to the agent system. If a threshold is exceeded, a fault may have occurred. This information is communicated by the agent to the manager in one of two ways. In a TCP/IP network, manager systems poll agents that respond with information about possible faults. In an OSI network, agents may independently generate alarms that are

sent to the manager. Agent systems also carry out other management operations on the device. For instance, an agent may implement changes to the configuration of a device specified remotely by a manager. In a TCP/IP or OSI network, every managed device has an agent. This may not necessarily be the case for devices that support proprietary protocols.

3.2.2 The Management Information Base

The Management Information Base (MIB) [ROSE90a] is a standardized description of the information that must be maintained by managed devices. A MIB is implemented on a device as a collection of variables that can be used to describe a device's internal state. An agent system may use a MIB to record information about changes in a device's state. MIB variables are also used to control devices (as will be described in the next section).

MIBs are used in TCP/IP and OSI networks. In recent years, a wide variety of MIB definitions have been created, partly to manage the different kinds of vendor-specific devices. Many of these MIB definitions have been standardized. The general framework within which MIBs can be defined is specified by the Structure of Management Information [ROSE90b], a topic beyond the scope of this report. Understanding variation in the type of information that must be maintained in MIBs for different types of devices will be a significant aspect of the domain modeling phase.

3.2.3 Communication of Management Information

The communication of information needed to manage the network, both from manager to agent and agent to manager, is accomplished through use of network management protocols. These protocols specify a set of low-level operations, referred to in this report as management operations. Management operations are used in two ways: (1) to ascertain a managed device's status by retrieving values of MIB variables, and (2) to control a device by altering MIB variable values. For instance, a manager may initiate a series of management operations in order to poll a set of devices. In response, the agent systems on the devices retrieve values of MIB variables and send poll responses back to the manager. Alarms sent by agents may also contain values of MIB variables that describe a device's status.

To control a device or to change its configuration, the manager initiates operations that change the values of specific MIB variables. These changes are carried out on the device by its agent system. The changes to variable values have the effect of triggering internal processes that alter the devices's operating status--such as causing it to reboot. In addition to management operations, a protocol may also specify the type, format, and information content of alarms and other notifications sent by agents to managers.

Conceptually, management protocols reside in the application layer and use the communications services provided by the lower layers (as described in sec. 2.2). Two prominent management protocol standards are:

- **The Simple Network Management Protocol (SNMP)**

  SNMP is used with the TCP/IP suite of protocols. At present, the protocol, described in [CASE90] and [CASE93], specifies a set of management operations for retrieving and altering information in MIBs, authorization procedures for accessing MIB tables, and mappings to lower TCP/IP layers. An expanded version of SNMP is being created, called SNMPv2, which will provide additional services.

- **Common Management Information Service (CMIS) and Common Management Information Protocol (CMIP).**

  Within OSI, CMIS [ISO/IEC 9595] defines a set of network management services. This includes, but is not limited to, management operations as defined above, called management-operations services. It also includes notifications that can be sent by agents, referred to as management-notification services. CMIP supports the services defined in CMIS by specifying lower-level data units for transmitting CMIS operations and notifications [ISO/IEC 9596-1].

In addition to standard protocols, vendors of network devices have also created proprietary management protocols for use by their devices. Although there is a trend toward use of standard protocols, many proprietary systems still exist. In practice, a network management system may have to manage devices that use standard management protocols as well as devices and proprietary element management systems (described in sec. 2.2.3) that use proprietary management protocols.

3.3 Architectures for Network Management Systems

In general, within a TCP/IP or OSI network, each manager controls a number of managed devices containing agent applications. For larger networks, factors such as size, scope, topology, available bandwidth (the amount of data that can be transmitted over a digital channel in bits/second), and administrative considerations may require use of many managers. These managers may control different portions of the network and operate at different network management stations or different network operations centers. The network management architecture describes the organization of the management of a network including the distribution of responsibility and control among different managers. The network management architecture can be thought of as the part of the overall network configuration (discussed in sec. 3.1.3). As such, the management architecture describes (1) connections between manager systems and parts of the managed network, and (2) connections between different manager systems.

3.3.1 Basic Types of Architectures

Three basic types of network management architectures may be described:

- o Centralized, with one manager software system at a single location that controls the entire network.

- o Distributed, with different managers having responsibility for different parts of the network (these managers are in a "peer to peer" relationship with respect to each other and coordinate their actions).

- o Distributed hierarchical, with multiple levels of managers in which higher-level managers control lower-level managers.

The network topology (described in sec. 2.1.2) in part determines the management architecture (and is discussed in more detail in sec. 5).

3.3.2 Factors Influencing Network Management Architectures

In addition to such factors as size, scope, topology, and administrative considerations, the design of network management systems may have to take into account a number of other considerations. Three of these are described below.

- o **Interfacing to Element Management Systems.**

  Proprietary element management systems, introduced in section 2.2.3, are management tools supplied by vendors for managing collections of devices that use the same proprietary protocol. A network management architecture for a large network may include several proprietary element management systems. A manager must be able to communicate with an element management system either using a standard protocol or a third-party translation package. The manager must be aware of what kind of information is provided by the element management system about the devices under its control. The manager must be aware of what management operations can be performed on the devices controlled by the element management system. Providing "seamless" interfaces to multiple element management systems is an important consideration in designing a network management architecture.

- o **Requirements for Survivability.**

  During a conflict, an important requirement of DoD network management systems is that they be able to continue to operate under adverse circumstances. To fulfill this requirement may necessitate a significant degree of flexibility in a network management architecture. During a crisis or war, it may be necessary to alter the

architecture and the existing relationships between network management systems. It may also be necessary for particular network management systems to assume different responsibilities and change the way they perform management functions.

o **Requirement for Manager-to-Manager Communications.**

The existence of network management architectures with multiple management stations requires using conventions for communicating among managers. "Manager-to-manager" communications is an important new area of network management.

3.4 Functions Related to Network Management

In addition to the network management functions described in section 3.1, there are other management and management-related activities that must be performed to support a communications network and ensure that it operates smoothly. These include provisioning end users with communications services, help desk functions, and business-related functions. They also include the planning, design, and engineering activities necessary to evolve the network over time. These functions are not part of network management domain and are not performed by network management software systems.[3] However, carrying out these functions may require either use of information provided by a network management system or use of the network management system services. Because of their interactions with network management, the four functions described below form part of the context of the network management domain and need to be accounted for in this report.

o **Provisioning**

Provisioning refers to the activities necessary to provide new communications services to end users. Functions of provisioning include:

- processing of requests for communications services,
- assignment of network resources,
- equipment installation and configuration, and
- verifying that the original request has been satisfied.

Provisioning uses functions provided by configuration management (described in sec. 3.1.3). Configuration management functions are used to set the operating characteristics and to define the connectivity of the equipment that is being installed.

---

[3]It should be noted that automation of the help desk functions could result in their becoming part of the network management system.

o **Providing End Users With Help Desk Facilities**

This includes the activities associated with assisting end users in their use of communications services. A help desk receives user inquiries, provides assistance, disseminates information about the network, and receives problem reports. When possible, reported problems may be resolved on the spot. Otherwise they are forwarded for handling to the network operations center (though not necessarily the network management system itself). To support its function, the help desk may receive (1) reports on fault occurrences provided by fault management, and (2) reports on utilization of network resources provided by performance management.

o **Performing Business Administration Functions for the Network**

Network business administration includes:

- program and budget management,
- negotiating and monitoring procurement of contractor services and equipment,
- maintaining directory services to identify end users and network resources, and
- charging and billing end users of network services.

Network business administration uses information on utilization of network resources provided by accounting management for charging and billing.

o **Systems Planning, Design, and Engineering**

Planning and design is concerned with the evolution of the network to support future end-user requirements [NMF92d] and [MOLL92]. This includes requirements analysis of future network services needs, trend analysis, capacity planning, and design and modeling of networks. Systems engineering refers to the activities of implementing the network design developed during systems planning and design. Systems planning may use long-term utilization and trend analysis information provided by performance management. Engineering may require the services of configuration management to configure newly installed equipment.

These functions are often performed offline by personnel located at a network operations center. Because of their interactions with network management, the functions described above are shown in the structure diagram that appears in section 3.5. They are also represented as external domains and entities in the context diagram presented in section 4. In some views, the functions described in this section, together with network management, are said to constitute a larger functional area called enterprise network management [MOLL92].[4]

---

[4]Systems management is another term sometimes used for enterprise network management. However systems management can also mean many other things in a computer environment,

3.5 A Structure Diagram For Network Management

The structure diagram shown in figure 3.1 organizes the functional areas associated with the communications network into conceptual levels. Each level provides services that support the level above it. The use of conceptual levels is analogous to the layered protocols described in section 2.2 *but focuses on relationships between domains rather than protocol layers*. Three levels are shown:

- Level 1, the lowest level, contains the network devices and transmission medium.

- Level 2 contains:

  (1) the layered communications protocols, such as the OSI and TCP/IP, and

  (2) operating software system services.

  The layered communications protocols at level 2 utilize the hardware systems at level 1 to convey digital signals. As described in section 2.2.1, it is possible to divide the communications protocols into separate layers, in which each layer provides services that support the layer above it. However in the structure diagram, only the application layer is explicitly depicted. The application layer contains the network management protocols that provide an interface between network management (at level 3) and the lower communications protocol layers. Operating system services are used by end-user application systems at level 3 (as well as network management systems). Operating system software runs on hardware platforms included in level 1.

- Level 3 shows the network management domain and various end-user application systems, such as command and control software, that reside at the same level conceptually. These systems utilize the communications services and operating system services in level 2. It should be noted that a network management system also provides services to the managed network, i.e., the management services described in section 3.1. (An arrow in the downward direction denotes this relationship.)

- Above level 3 reside the provisioning, help desk, network business administration, systems planning, systems design, and systems engineering functions. These non-computerized activities use services provided by network management (described in sec. 3.4). They also provide necessary support services to the managed network (also described in sec. 3.4).

The managed network, defined in section 2.4, consists of the network devices and transmission medium in level 1 and the layered communications services at level 2.

---

such as management of individual servers or processors.

In figure 3.1, subdomains are represented as blocks that appear within other blocks. The structure diagram shows the five system functional areas, described in section 3.1, as subdomains of network management. (In sec. 5, fault management will be further decomposed to identify the target subdomain.) Within the level 1 and 2 blocks are shown the major components of the managed network and the major types of network devices.

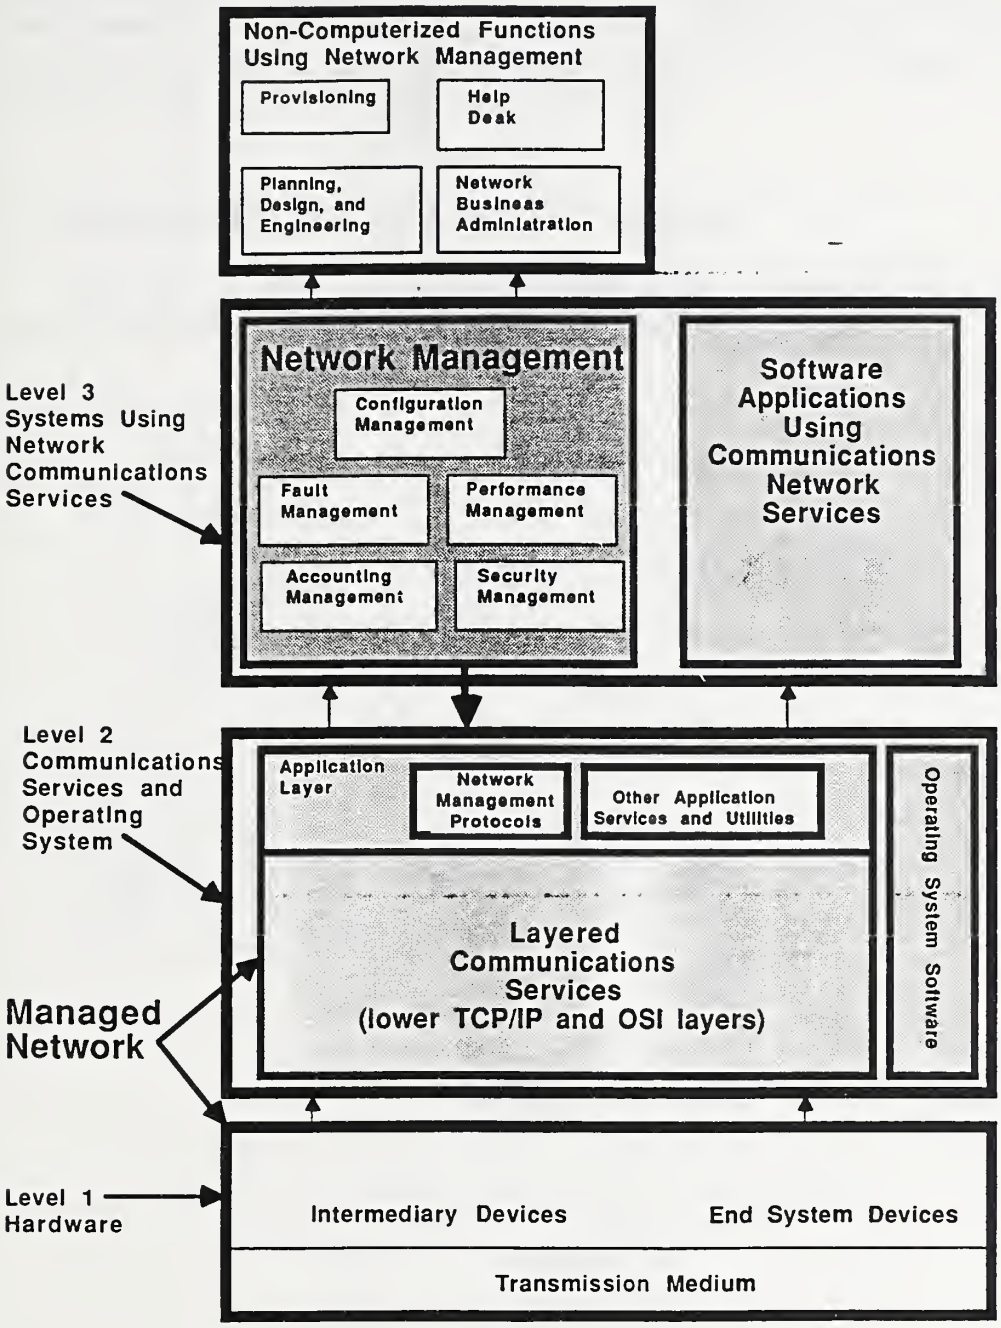

Figure 3.1: The Structure Diagram for Network Management

4. THE CONTEXT OF NETWORK MANAGEMENT

Using the discussion of the communications network and network management presented in the preceding sections, the context of the network management domain is defined according to the FODA process. A context diagram is presented that shows data flows between a network management system and the external entities in its context. Guidance is provided to application developers on how to use the context diagram to determine if their application falls within the domain.

4.1 The Context Diagram For Network Management

The context diagram for network management is presented in figure 4.1. The data flows shown in this diagram are identified and described in section 4.2.

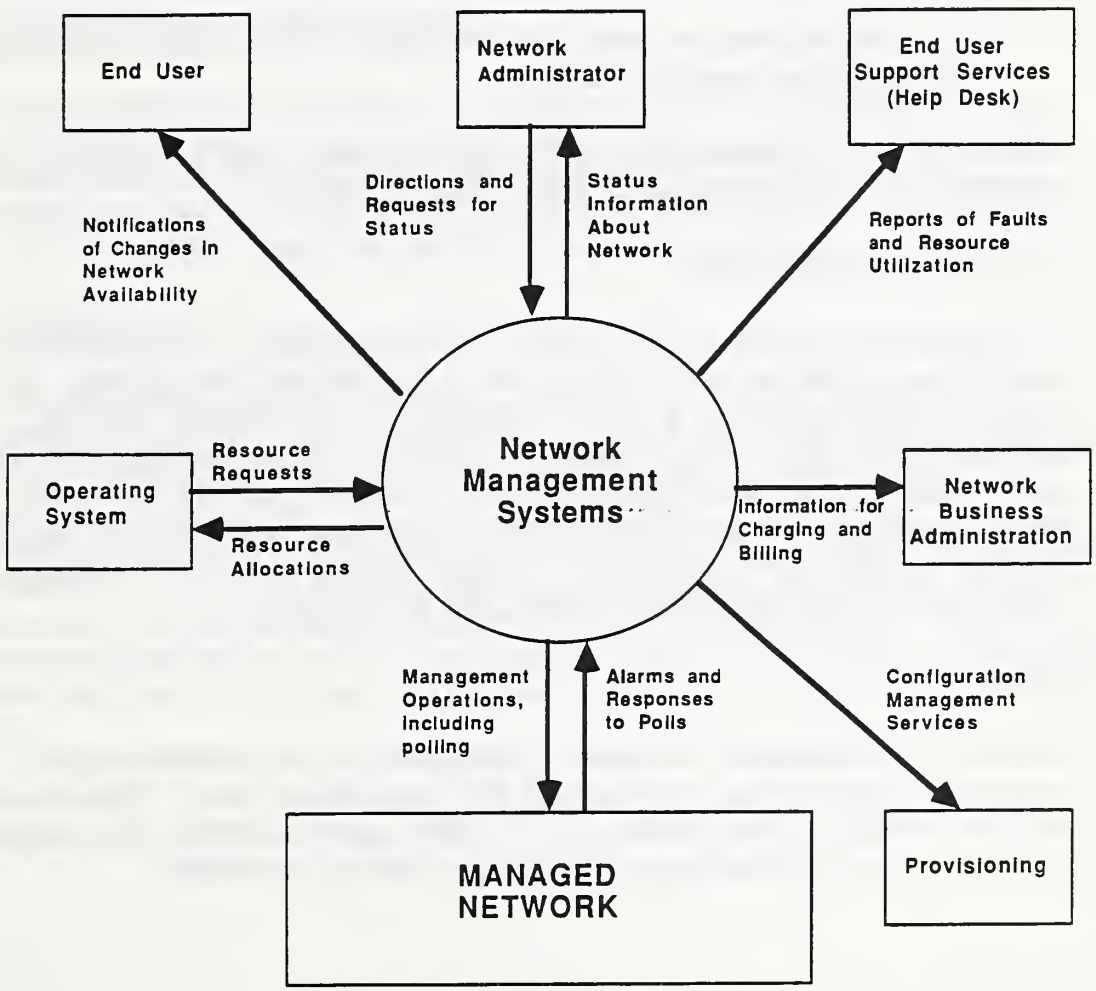

Figure 4.1: The Context Diagram for Network Management

4.2 Description of External Entities and Data Flows

Data flows between network management and each of the external entities in its context are described. (To match the external entities in the context diagram to the entity descriptions in the previous section, the reader should start with the Managed Network entity and proceed in a clockwise direction.)

- **The Managed Network.**

  The manager part of a network management system performs management operations on the managed network (defined in sec. 2.3) and receives responses to its actions. (Manager software systems were described in sec. 3.2.1.) For networks that support TCP/IP or OSI protocols, data flows between manager and the agent depict:

  (1) Management operations sent from manager to agent, and

  (2) Agent responses to management operations, alarms, and other notifications transmitted from agent to manager.

  Internal interactions between agent systems and the devices they control are described in section 3.2. These interactions are not shown in figure 4.1.

- **The Host Operating System.**

  The network management system interacts with host operating systems running on the network management stations (defined in sec. 3.2.1) and managed devices. The network management system software requests computing resources from the operating system and receives allocations of resources. These interactions are carried out by means of interprocess communications on manager and agent platforms. (It should be noted that, in some cases, secure network management systems include operating systems. Also, in the future, it is possible that distributed operating systems and network management systems may be combined.)

- **End users.**

  Under some circumstances, the network management system provides selected information about the state of the network directly to the end users. The interaction could be carried out by an application layer utility such as electronic mail, supported by the lower-level communications services provided by the network.

- **External network administrators.**

  These persons (or automated systems) request and receive information about the state of the network from the manager part of the network management system. They may also provide directions to the network management system and initiate management operations. These interactions may be carried out by means of an operator interface with graphics user interface capabilities.

- **Help Desk.**

  A help desk located at a network operations center, described in section 3.4, may send reports to the network administrator describing perceived problems on the network. The context diagram shows the flow of data from network management to the help desk. The data describes the utilization of network resources and fault occurrences.

- **Network Business Administration.**

  The network business administration function, described in section 3.4, may be carried out by support personnel located at the network operations center. The context diagram shows the flow of data from network management (accounting management) to network business administration. The data describes the utilization of network resource by end users and is used for billing purposes.

- **Provisioning.**

  The provisioning function, described in section 3.4, is carried out by support personnel located at the network operations center. The context diagram shows use of configuration management services to modify the network configuration.

In the case of data flows to the end user, end-user support services, network administration, and provisioning, it should be noted that the actual transfer of information may be carried out by another party, such as an automated system or the network administrator. Data flows between network management and systems planning, design, and engineering are omitted for the sake of simplicity. All necessary data stores are considered to reside within individual subdomains.

4.3 Guidelines for Using the Context Diagram

The context diagram can be used by the application developer to assess whether or not the application being developed falls within the network management domain. The following guidelines can be used:

(1) The functions to be performed by the application to be developed should fall within those listed in section 3.1.

(2) If there is a discrepancy between the scope of the application and the scope of the context model, then one or both of them may need to be reevaluated and modified. In some cases, a description of a network management system to be developed may include the provisioning function and end-user support services (help desk facilities). These two functions are often co-located with a network management station at a network operations center. Though they may interact with a network management system, they do not constitute part of the network management domain as described in section 3.1.

(3) A context diagram can be created for the network management system application being developed.[5] Similar to the context diagram for a domain, the application's context diagram shows data flows between the application system and its external entities. The data flows to and from the application should be a subset of the data flows to and from the network management domain. (See secs. 4.1 and 4.2.) That is, all the data flows associated with the application and its context should appear in the context diagram for the domain (fig. 4.1). The description of the data flows identified during the context analysis phase should be expected to evolve during the domain modeling phase. Therefore, the context diagram should be flexible enough to accommodate change.

To be considered within the scope of the network management domain, the application to be developed should satisfy guidelines (1) and (3). If the application does not fully satisfy guideline (2), further analysis by the application developers may be needed.

---

[5]See [YOUR89] for a description of how to do this.

5. THE CONTEXT OF THE TARGET SUBDOMAIN

In section 3.1, fault management was identified as one of the five subdomains of network management. Within fault management, a specific functional area--alarm surveillance--has been selected as the target subdomain for the Domain Analysis Case Study. Alarm surveillance systems will be the subject of detailed domain modeling during the execution of the next phase of the FODA domain analysis method.

In this section, the functions of fault management are first described in more detail. The context of fault management is defined and a context diagram is provided. Since alarm surveillance falls within fault management, this description provides a basis for understanding the place of the target subdomain within network management. The target subdomain is then itself described and its context is defined. A context diagram for the target subdomain is presented. This is followed by a description of variabilities in the context of alarm surveillance systems. The impact of these variabilities on data flows across domain boundaries and on requirements for systems within the domain is discussed. Finally, a description is provided of how the context model for the target subdomain is to be used by potential application developers.

5.1 Fault Management

Figure 5.1 shows the relationship between network management, fault management, and the functional areas within fault management. For the purposes of this report, these functional areas--alarm surveillance, trouble tracking, fault diagnosis, and fault correction--will be considered subdomains of fault management.

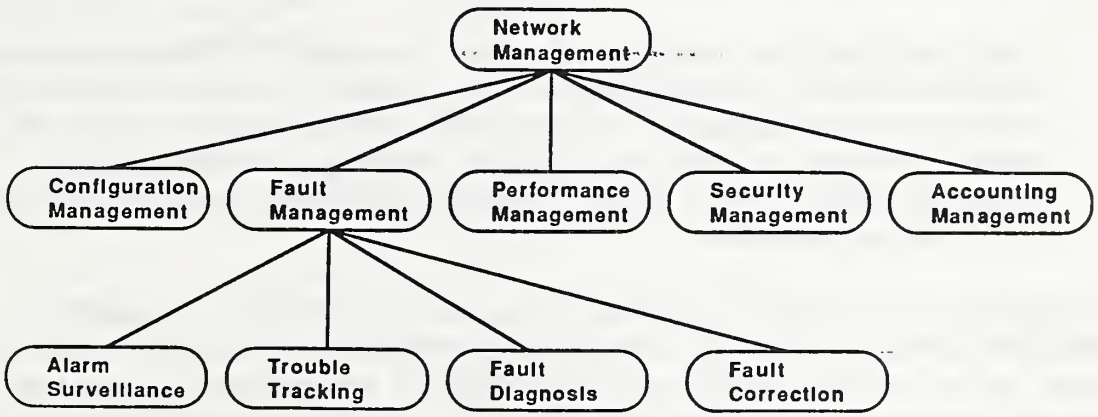

Figure 5.1: Subdomains of Network Management

5.1.1 The Subdomains of Fault Management

The identification of the subdomains of fault management in this section is based on the OMNIPoint specifications ([NMF92a], [NMF92b] and [NMF92c]) and is generalized for management of non-OSI networks.

- Alarm Surveillance is concerned with: (1) the monitoring, detection, and reporting of faults and fault-related events or conditions that occur within a network; (2) the logging of this information for future use in fault detection and other network management activities; and (3) the analysis and control of alarms, notifications, and other information about faults to ensure that network management resources are directed toward reported faults that affect the operation of the network.

- Trouble Tracking is the process of recording, updating, maintaining, and distributing information about the progress of a reported fault. An important service of trouble tracking systems is to provide information that assists support personnel in resolving complex faults in a timely manner. The information provided by trouble tracking may include step-by-step procedures that describe how to diagnose and correct specific types of faults. Trouble tracking may also include procedures for assessing the impact of a fault upon end users of the network.

- Fault Diagnosis is the process by which (1) the location of a detected fault is isolated to a specific network element, such as a particular device or link in which the fault occurred, and (2) if necessary, additional information about the fault and its circumstances is obtained that allows prescription of corrective actions. Fault diagnosis includes the definition and use of diagnostic and testing procedures tailored to particular kinds of faults. Diagnostic and testing procedures are often performed off-line by support personnel.

- Fault Correction refers to the action(s) taken in response to a fault that are necessary to restore network communications services. Typically, corrective actions may include reconfiguring a portion of the network, replacing faulty resources, or repairing devices. In some cases, it may be necessary to re-route traffic to temporarily bypass a faulty device. Correction procedures are often performed off-line by support personnel.

The functions of the subdomains of fault management are interrelated. The alarm surveillance function monitors the network. In a simplified case, information about a detected fault is routed by alarm surveillance to network support personnel who use diagnostic procedures to isolate the fault. The faulty component is then bypassed if necessary and taken off-line for repair. For more complex problems, faults that cannot be resolved immediately are forwarded to the trouble tracking function. Trouble tracking procedures provide guidance on how the fault is to be resolved and may specify the circumstances under which particular diagnostic and correction procedures should be used.

5.1.2 The Context Diagram For Fault Management

The context diagram for fault management is presented in figure 5.2. This context diagram shows the flow of information between the fault management and the external entities that make up its context. The data flows between fault management and the managed network, the network administrator, and end-user support services are the same as those shown in the context of network management in figure 4.1. Two external entities are added to the context of fault management that do not appear in figure 4.1: configuration management and performance management. In figure 4.1, these external entities are not shown because they are within network management. However, in the context of alarm surveillance, configuration management provides information to alarm surveillance about network resources that need to be monitored. Performance management provides long-term statistical information on usage.

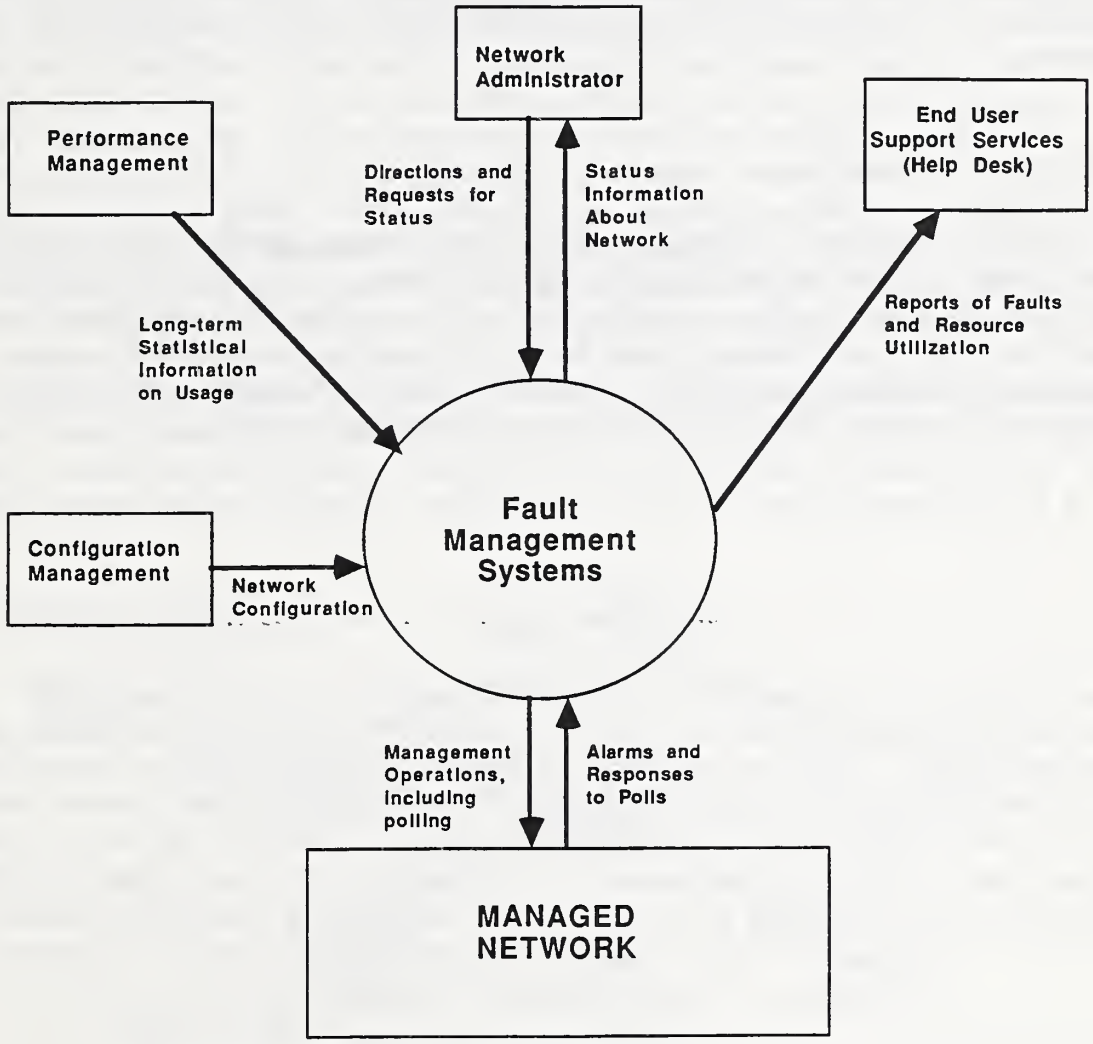

Figure 5.2: The Context Diagram for Fault Management

In the case of data flows between fault management and configuration management, performance management, and end-user support services, it should be noted that the actual transfer of information may be carried out by another means, such as an automated system or the network administrator.

The context diagram does not show interactions between fault management and:

- Network Business Administration.
- Provisioning.
- End-User Support Services.
- Operating System.

These external entities most often interact with other subdomains of network management (as discussed in sec. 3.4) and may interact with fault management only in exceptional circumstances.

5.2 The Target Subdomain

The alarm surveillance subdomain consists of the following functions, shown in figure 5.3: fault detection, alarm analysis (consisting of alarm filtering, alarm correlation, fault prediction) and log control. Alarm surveillance systems are software systems that perform these functions. Alarm survelllance systems may be implemented as part of the manager software system. (Managers were defined in sec. 3.2.) Specific functions such as monitoring or filtering and correlation may also be implemented as separate systems.

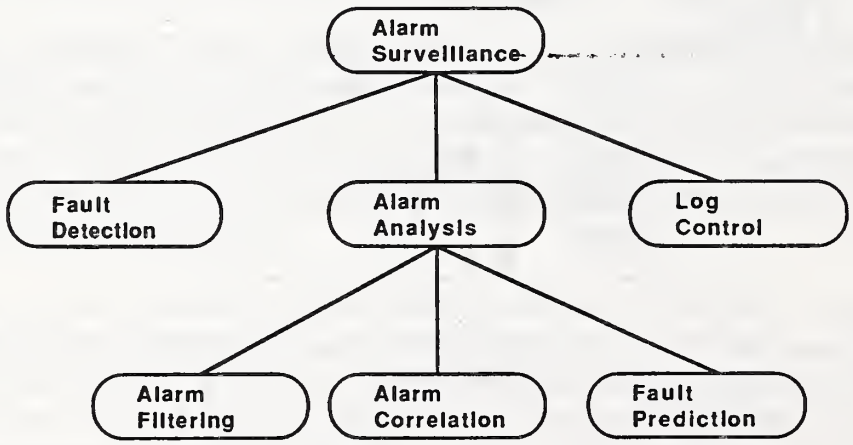

Figure 5.3: The Functions of Alarm Surveillance

5.2.1 Fault Detection

Fault detection is the process of discovering and reporting active faults, potential faults, and fault conditions. As discussed in section 3.1, faults and fault conditions may be detected by monitoring the managed network and its environment. The monitoring strategy determines whether the network is monitored through polling, receipt of alarms and traps, or a combination of these techniques.

The choice of monitoring strategy is driven by the characteristics of the managed network. Defining a monitoring strategy involves taking into consideration such factors as the number of devices to be managed, the location of the devices, the relative priority of different devices for receiving management services, the types of faults that are likely to occur, the available bandwidth on the network, and whether the management protocols used by the devices support polling, traps, or alarm generation. Tradeoffs between use of polling versus alarm generation are also a factor. Frequent polling of large numbers of devices consumes more network resources than alarm generation. However, polling allows a manager (defined in sec. 3.2.1) to exercise more control over monitoring activity. By polling, a manager system is not solely dependent on agent systems to send notifications describing a device's status. In the event of the complete device failure, the lack of response to a poll still provides information to the manager. The effect of variability in network characteristics on choosing a monitoring strategy will be discussed further in section 5.4. Detailed description of monitoring strategies is reserved for the domain modeling phase.

Alarms and other reports describing faults may contain information about the network device or other network resource in which the fault originated, the type of the fault, its severity, and its suspected cause. In addition to automated detection, faults may also be detected by persons and reported to the help desk or point-of-contact facility (sec. 3.4).

5.2.2 Alarm Analysis

Alarm analysis consists of alarm filtering, alarm correlation, and fault prediction. Alarm filtering is the process of controlling the flow of information about faults and fault-related events, including alarms, poll responses, and other notifications. Alarm filtering involves selecting which messages to forward and determining the destination to which the messages should be forwarded. Filtering also includes the ability to suspend and resume the reporting of alarms and other notifications. The process of selecting which alarms and notifications to forward is accomplished through use of a set of selection criteria, referred to as discrimination criteria. These criteria may be:

- stored in tables,
- expressed as activity or severity thresholds,
- expressed as rules composed of Boolean combinations of factors, or
- represented using methods provided by expert system technology.

Alarm filtering is intended to ensure that (1) network management resources are directed only to the reported faults whose resolution will be necessary to effectively manage the network, and (2) network performance is not degraded by transmitting redundant or useless information.

Alarm correlation is the process of comparing information in alarm messages to determine if they have the same or related causes. The information provided by alarm correlation can be used to diagnose faults and to support the alarm filtering function. Correlation constitutes a spectrum of diverse activities, ranging from simple short-term association of related alarms to trend analysis. For instance, alarms generated within a short time span by different devices may be correlated to reveal a single underlying problem. Long-term correlation of alarms coming from a single device (or perhaps a group of associated devices) can be used for trend analysis. Fault prediction is the process of using information about alarms and other events stored in a log to predict faults or identify developing fault conditions.

Filtering, correlation, and fault prediction functions can be combined in systems that provide enhanced analysis capabilities. In addition to alarms, poll responses, and other notifications, these systems may receive and analyze additional information such as historical information on faults stored in logs (described in sec. 5.2.3). Such systems may filter out irrelevant information, correlate and cluster information pertaining to a single problem or potential problem, and forward the identified problem to an appropriate location for handling. Some systems are capable of performing a simple corrective action, such as a device "reset" operation, in response to certain kinds of faults. Often fault analysis functions can be implemented as expert systems [CRONK88], [RABIE88], and [SUTT88].

5.2.3 Log Control

Log control is the process of selecting and storing alarm messages and other related information using a permanent storage facility such as a database. The functions of log control are described in [ISO/IEC 10164-6].

5.2.4 Functions Outside the Target Subdomain

The functions outside the target subdomain are the other functions of fault management: trouble tracking, fault diagnosis, and fault correction. Other subdomains of network management that are at the same level as fault management are also external to the target subdomain. These include performance management, configuration management, security management, and accounting management. Subdomains of fault management and network management that form part of the context of alarm surveillance are added to the context diagram in the next section.

5.3 The Context Diagram for Alarm Surveillance

The context diagram for alarm surveillance is presented in figure 5.3. This context diagram shows the flow of information between alarm surveillance and the external entities that make up its context. The data flows to and from the managed network, the network administrator, and end-user support services are the same as those shown in the context diagrams for network management (fig. 4.1) and fault management (fig. 5.2). The data flows to and from performance management and configuration management are the same as those shown in the context diagram for fault management (fig. 5.2).

An additional external entity--trouble tracking--appears in the context of alarm surveillance. Trouble tracking (described in sec. 5.1.1) is a subdomain of fault management. In figure 5.2, this entity is not shown because it is internal to fault management. Alarm surveillance forwards information about faults that cannot be resolved immediately to trouble tracking.

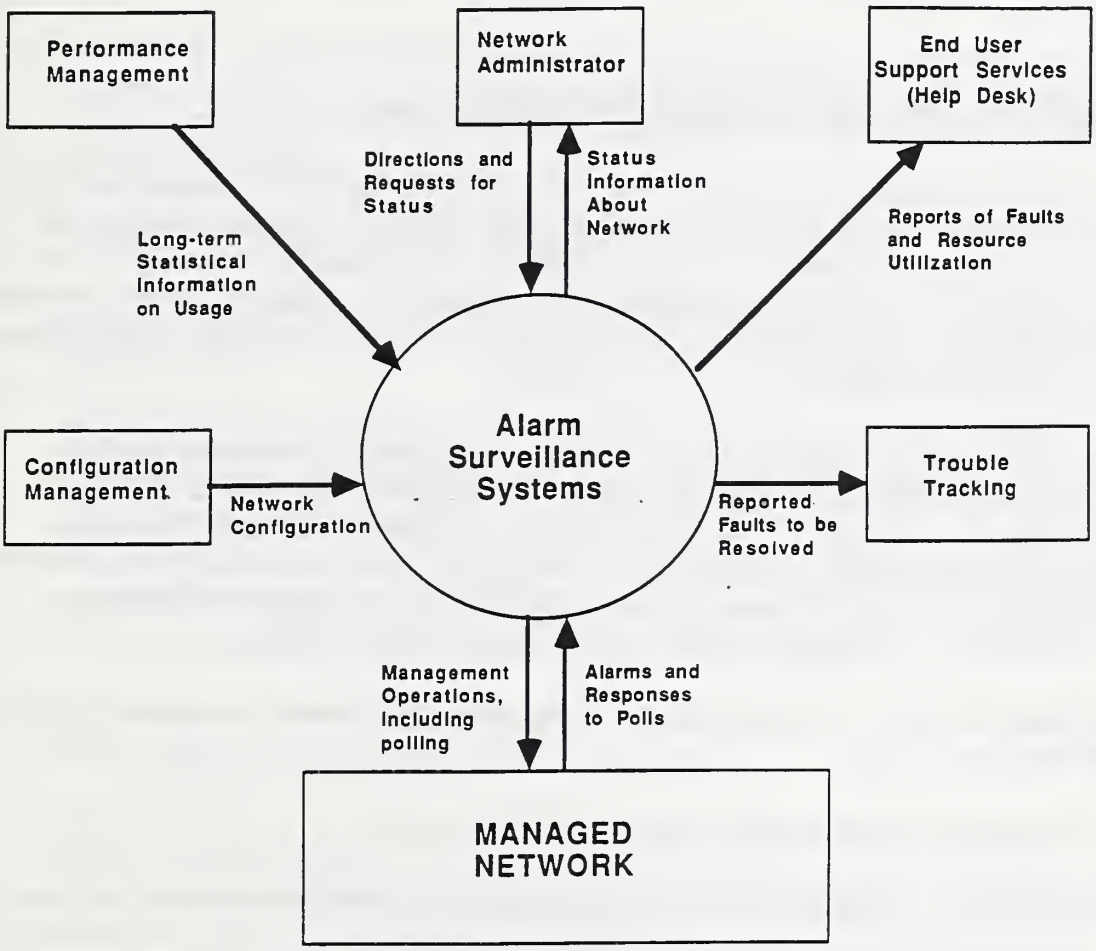

Figure 5.4: The Context Diagram for Alarm Surveillance

In the case of data flows to and from configuration management, performance management, end-user support services, and trouble tracking, it should be noted that the actual transfer of information may be carried out by another party, such as an automated system or the network administrator.

As in the context diagram for fault management (fig. 5.2), the context diagram for alarm surveillance does not show interactions with:

- Network Business Administration.
- Provisioning.
- End-User Support Services.
- Operating System.

These external entities most often interact with other subdomains of network management (as discussed in sec. 3.4) and may interact with alarm surveillance only in exceptional circumstances.

5.4 Variation in the Context of Alarm Surveillance Systems

This section identifies variability in functions and behavior of the managed network domain and its impact on the data flows described in section 5.3 and on the requirements for systems within the target subdomain. The importance of understanding variability in the context of a domain was discussed in section 1.5.1. During the subsequent domain modeling phase, the variabilities identified below will provide the basis for identifying context features that form an important component of the feature model.

In a broad sense, requirements for network management systems are based on differences in the kind of devices to be managed, their function, their number, and their location relative to the network management station. Operational requirements for the communications network that affect network management are also important factors. It is the variability in these factors that forms the basis for variation in the systems within the domain. Accordingly, variability in context is organized into four areas:

(1) Physical variations in the structure of the network including its size, scope, and topology (sec. 5.4.1),

(2) Variability in devices to be managed (sec. 5.4.2),

(3) Variability in operational requirements for network management (sec. 5.4.3), and

(4) Variability in the administrative and political divisions of the managed network (sec. 5.4.4).

For each area of variability in the managed network domain described below, the following information is provided:

(1) An identification and description of the variability,

(2) How this variability affects the flow of data across domain boundaries, and

(3) In some cases, how this variability affects the operating features of, and requirements for, alarm surveillance systems.

5.4.1 Physical Variations.

This includes variation in the size, bandwidth, scope, and topology of the network.

a) Network size.

Variation in context: Network size is determined by the number of devices that must be managed. (See sec. 2.1.1.) Network size can affect the frequency of faults occurring on the network. (Type and frequency of faults occurring in individual devices is discussed in sec. 5.4.2.)

Resulting variation in systems in the domain:

(1) Large networks may be partitioned into smaller segments that are monitored by different managers. This influences the choice of network management architecture.

(2) Network size and available bandwidth (defined in sec. 2.1.1 and identified as an area of variability below) influence the choice of monitoring strategy. The number of devices to be managed in the network partially determines whether polling, alarm generation, or a mixed strategy should be used. Smaller bandwidths have the affect of limiting the amount of management information that can be transmitted. In general, frequent polling results in greater consumption of network resources used to provide communications services than does alarm generation.

(3) The choice of network management architecture (defined in sec. 3.3) and the choice of monitoring strategy may be closely related.

(4) Use of out-of-band channels may be desirable. Out-of-band channels are alternative communications links that can be used instead of channels normally used for transmission. Out-of-band channels may include satellite links and underground fiber optics systems. In addition, since network management transmissions use smaller bandwidths than regular traffic, out-of-band channels can employ telephone lines (such as the Government Emergency Telephone System--GETS).

b) Bandwidth

Variation in context: Bandwidth (defined in sec. 2.1.2) is the rate at which information can be transmitted in bits/second. A network may contain different bandwidths along different links.

Resulting variation in systems in the domain: Bandwidth together with network size influences the choice of monitoring strategy and choice of network management architecture, as described in section 5.4.1a.

c) The scope of the network.

Variation in context: The scope of the network refers to the geographic area encompassed by the network. LANs, MANs, and WANs, described in section 2.1.1, are distinct categories of networks having different scopes. The managed network may include any or all of these three categories in any number.

Resulting variation in systems in the domain: The scope of the network influences the choice of monitoring strategy. For instance, in a satellite network, the decision on whether to use polling or alarm generation may in part, be based on the time delay due to transmission distances. (In cable links, time delay is not as great a factor.)

d) The topology of the network.

Variation in context: Network topology refers to the structure, interconnectivity, and geographic layout of a group of networks that form a large network. Different types of topologies were described in section 2.1.2. A network may include any or all of these topologies. Topology may have specific characteristics which, when combined with other contextual variabilities, may impact requirements for network management. For instance, a particular network may have an irregular topology consisting of distinct segments that are geographically distant. Because of topology and distance, the connections between the component segments may be difficult to reinforce with redundant communications links. In the event of a military engagement, such a network may be easily disrupted, leaving the separate segments isolated from each other.

Resulting variation in systems in the domain: Topology may influence the network management architecture. Parts of the network that would be left isolated in the event of a failure of a critical link may require their own, independent, management systems to ensure rapid response to faults.

5.4.2 Variation in Managed Devices

Variation in devices on the network consists of: variation in functions performed by particular devices; variation in the type and frequency of faults that may occur on particular devices; variation in the information about the device that should be maintained in order to manage it; and variation in the use of different management protocols, both in individual devices and proprietary element management systems.

a) Function of devices.

Variation in context: Major classes of managed devices, including end-system devices, intermediary devices, and link devices were discussed in section 2.1.3. Devices performing functions that are essential to critical missions may have higher priority in receiving management services. It may be essential to quickly detect faults and fault conditions in high-priority devices and then rapidly substitute redundant devices for faulty ones or initiate other corrective procedures. For instance, the maintenance of certain link devices that connect critical portions of the network may have priority over less critical end-system devices. The number of devices that are judged to provide mission-critical services within a network may vary. In some cases, the devices providing mission-critical services are viewed by network administrators as the set of relevant devices on the network. In such cases, these devices are the only devices that are monitored, thus effectively reducing network size (discussed in sec. 5.4.1). Variability in device function and priority is therefore closely related to variability in network size.

Resulting variation in system requirements: The existence of mission-critical devices and priority levels influences the choice of monitoring strategy. Frequent polling of high-priority devices permits closer monitoring of devices, thus helping to ensure rapid response to faults and fault conditions. Monitoring solely through use of alarm generation may be reserved for low-priority devices. In the case of complete device failure, the agent system on the device may not be able to generate an alarm. The number of high-priority devices affects the choice of network management architecture as well as monitoring strategy. A greater number of managers may be required to monitor networks with many high-priority devices. To provide closer monitoring of high-priority devices, polling may be chosen over alarm generation.

b) Type and frequency of faults.

Variation in context: Different types of devices can produce different types and frequencies of faults. Basic types of faults defined for OSI systems are described in [ISO/IEC 10164-4].

Resulting variation in data flows: The variability will directly influence the content of the data flow coming from the managed network to the alarm surveillance system.

Resulting variation in systems in domain: The type and quantity of faults to be detected will (1) affect the choice of monitoring strategy, and (2) influence requirements for alarm analysis functions.

c) **Information about the status of devices.**

Variation in context: Variation exists in the internal characteristics of a device that needs to be examined in order to assess and control its state. Manufacturers of devices that respond to SNMP or CMIP protocols may define different MIBs to represent characteristics necessary to manage different types of devices (discussed in sec. 3.2). In addition some devices, often using proprietary protocols, may be designed or configured to automatically emit notifications that contain a wide range of information describing the status of a device.

Resulting variation in data flows:

(1) With respect to devices that support MIBs, the data flow will vary in terms of:

a) the MIB variables that need to be examined to determine the status of particular devices, and

b) the variables that must have their values changed to alter the device's operating status.

(2) The type and frequency of notifications emitted by devices (using both standard and proprietary protocols) influences requirements for the alarm analysis function. In some cases, a large amount of irrelevant information may need to be filtered out.

d) **Network management protocol used by devices.**

Variation in context: As described in section 3.2.3, different devices may respond to different management protocols (either standards-based or proprietary). Element management systems control collections of devices that use the same protocol (either standards-based or proprietary).

Resulting variation in data flows: The data flows from the managed network to alarm surveillance may vary in the management protocol used by a specific device. Management protocols may vary in the content, number, and type of the messages passed, and in the procedures for sending messages.

Resulting variation in systems in domain: The choice of monitoring strategy will be affected by whether the device (or element management system) uses a management protocol that supports polling or alarm generation.

5.4.3 Operational Requirements for the Communications Network.

This includes operational requirements for both the managed network and for network management systems.

a) Survivability and operating modes.

Variation in context: A DoD communications network may have to respond to changes in operating conditions. This includes shifts from normal peacetime operations to a crisis state, wartime state, temporary disabled conditions (not necessarily arising from military conflicts), or even reconstruction. Reconstruction is a state in which a damaged network is reconstituted. The critical function of network management must be able to continue in the event of war, crisis, the advent of a disabled condition, or during network reconstruction.

Resulting variation in systems in the domain:

(1) To achieve survivability, the network may have to assume different architectures under different conditions. For instance, in a war or crisis state, the intermediate-level manager in a distributed hierarchical network management architecture may become inoperable. (Managers were defined in sec. 3.2.1.) As a result, the architecture could change from distributed hierarchical to distributed "peer to peer."

(2) The functions of alarm surveillance may change. For instance, different operating modes may change the monitoring strategy of a network management system.

(3) The communication of management operations may have to be switched from normal channels to out-of-band channels (described in sec. 5.4.1).

b) Staffing requirements and the need for automation.

Variation in context: Different networks have different requirements for containing the cost of staffing network management stations and network operations centers. Around-the-clock operations require more personnel than single shifts. (It has been estimated that it takes 5 people to staff a single around-the-clock position versus 1.2 people for a single shift [GTE93].) In addition, the number of positions on a shift is affected by such factors as the number of faults that are expected to occur on the network during a given time period.

Resulting variation in systems in the domain: A requirement to lower staff levels may in turn require greater sophistication in the capabilities provided by network management software. For instance, it may be necessary to use expert systems to perform analysis activities that would otherwise be done by humans. (See sec. 5.2.2.)

5.4.4 Administrative and Political Considerations.

Administrative and political considerations center on questions of who has jurisdiction over a managed network.

Variation in context: The network may be divided along administrative lines. Different devices, or different parts of the network, may belong to different organizations or may be subject to the jurisdiction of different organizational components.

Resulting variation in systems in the domain: Administrative boundaries and organizational policy may, in part, determine the network management architecture (discussed in sec. 3.3). For instance, the division of world-wide DoD operations into "combat theaters" affects the composition of the network management architecture.

5.5 Using the Context Diagram for the Target Subdomain

The context model can be used by the application developer to determine whether or not the network management system to be developed contains subsystems that fall within the target subdomain.

The functions to be performed by subsystems within the application should fall within those listed in section 5.2. Not all of the functions listed in the target subdomain need to be included in the application. For instance, the application may not require certain functions such as log control or fault prediction. However, most network management systems will require fault detection and reporting capabilities.

The domain model that will be developed in the next phase of the FODA process will be relevant to those subsystems within the application that fall within the scope of the target subdomain.

6. PREPARING FOR THE DOMAIN MODELING PHASE

This section discusses issues that need to be addressed in preparation for the domain modeling phase. It provides an assessment of the potential for developing reusable components in the target subdomain. Topics that will need to be investigated during the domain modeling phase are identified. Sources of information that can be used to support the domain analysis effort are also identified. The procedure for developing the domain model is described.

6.1 Evaluating Reuse in the Target Subdomain

Before deciding to do domain analysis, it is important to evaluate the extent to which the target subdomain can be exploited to develop reusable components. Within the target subdomain of alarm surveillance, the evaluation includes the following four factors:

(1) The criticality of the alarm surveillance subdomain to the organization,

(2) The number of different alarm surveillance systems that are likely to be developed within the organization,

(3) The extent of variability in the context of the alarm surveillance systems to be developed, and

(4) The maturity and stability of knowledge about alarm surveillance.

Factors (1), (2), and (3) help determine the potential payoff in developing reusable components. Using domain analysis to develop reusable components is most appropriate for critical domains in which it is anticipated that a number of systems will be developed. The context of systems within the domain should vary significantly. (If there is little or no variation, one system can be developed and fielded in multiple locations.) Factor (4) is essential in assessing the quality of the reusable components to be developed. It is hard to develop good reusable components in domains in which knowledge will change extensively.

6.1.1 Assessing the Criticality of Functions in the Target Subdomain

Automated fault detection is a basic function necessary to manage large, complex communications networks. This function is perhaps the most essential of network management functions [ADAMS91]. Understanding the application of fault detection in different contexts will be an important factor during the domain modeling phase. Alarm analysis, while less necessary for smaller networks, takes on added importance in larger networks. Log control is necessary to support monitoring and detection as well as to support alarm analysis.

6.1.2 Assessing Future Development in the Domain

Assessing the development of future systems in a domain depends, in part, on both the number of systems to be developed and the variability in the context of the domain. It is expected that many organizations within government and industry will undertake the development of alarm surveillance systems in the coming years. Variability in the context of alarm surveillance is significant. This topic was discussed in detail in section 5.4.

6.1.3 Assessing the Maturity and Stability of Knowledge in the Target Subdomain

The maturity of the alarm surveillance subdomain is indicated by the overall quality of the knowledge of domain experts and the quality of knowledge contained in other domain resources. This includes:

- The extent and amount of knowledge about building alarm surveillance systems.
- The likelihood of extensive changes to the knowledge in the near-term, i.e., the permanence or stability of this knowledge. (If the knowledge were to undergo radical changes in the near term, the domain model would become obsolete.)
- Whether or not knowledge about building alarm surveillance systems has been exercised extensively and successfully in problem solving.

With respect to the functions in the target subdomain, knowledge about fault detection appears to be the most extensive. Many software systems have been developed that perform fault detection. Standardized profiles have also been created that describe important parts of this function. (See sec. 6.4 below.) Alarm analysis is also a mature area for which standardized profiles have been created. As indicated in section 5.2, a number of expert systems have been developed that perform fault detection and alarm analysis. Fault prediction is a much newer, possibly less well understood area.

6.2 Major Areas Within the Target Subdomain to be Explored

Section 1.5.2 described the three major components of the FODA domain model: the entity-relationship model, the feature model, and the functional model. This section discusses specific topics and issues that will need to be addressed to develop the three components of the domain model for alarm surveillance.

6.2.1 Fault Detection

In the area of fault detection, the following information will need to be included in the domain model.

- The feature model will describe the capabilities of fault detection systems as seen by network administrators, system developers, and end users. The feature model will include alternative monitoring strategies that can be chosen such as polling, alarm generation, and combinations of polling and alarm generation. An understanding is needed of why different strategies are chosen and precisely how variation in context affects the choice.

- The entity-relationship model will describe the major data structures needed to support the alarm surveillance function.

- The functional model will capture the functions, flow of data, and constraints associated with particular monitoring strategies.

In addition, the analysis process will uncover other information that will need to be included in the domain model.

6.2.2 Alarm Analysis

In the area of alarm filtering, the following information will be included in the domain model.

- The feature model will describe the capabilities of alarm analysis systems that directly affect network administrators. An understanding is needed of the effect of variation in context on the choice of specific filtering and correlation techniques (described in sec. 5.2.2).

- The entity-relationship model will identify the major data structures needed to represent discrimination criteria and support filtering and correlation functions.

- The functional model will describe the functions, flow of data, and constraints associated with:

  - filtering using Boolean combinations or production rules;
  - possibly, filtering using more sophisticated automated reasoning techniques that will require the capabilities provided by expert systems [CRONK88], [HONG91], [RABIE88], [RICH91], and [SUTT88];
  - possibly, processes for automatically learning discrimination criteria that may later be implemented using neural networks approaches [HERT91] and

[WASS89], or other machine learning techniques [CARB90] and [SHAV90]; and

- using different methods for correlating alarms and notifications.

Variability described in the functional model and the entity-relationship model should be linked to alternative and optional features in the feature model.

o Basic requirements should be provided for alarm correlation functions.

Because of the extensive diversity of alarm correlation functions, the domain model may be limited to specifying what correlation functions are to be performed.

6.2.3 Other Areas

It is anticipated that less information about log control and fault prediction will be required. In the case of log control, basic data structures should be included in the entity-relationship model. A high-level functional model should depict the flow of information to and from the log control function. The functional model of fault prediction may be limited by the extent of knowledge about this function.

6.3 Sources of Domain Knowledge

This section lists sources of domain knowledge that will be used in this study. Several categories of sources are relevant. These include:

o Human experts who specialize in development of alarm surveillance systems.

o System requirements and design documents for alarm surveillance systems.

o Research and commercial literature that describes the functions of alarm surveillance.

o Reports by private companies that develop DoD network management systems.

o Information obtained from vendors of commercial network management tools, including manuals and other documentation.

o Public domain network management software and accompanying documentation.

Experts from GTE Government Systems, Inc., located in Chantilly, Virginia, are providing domain expertise for the project.

6.4 Standards For Network Management and Communications Services

The domain model will have to take into consideration those standards for network management and network communications that are most relevant to the target subdomain. Three classes of standards are most important: reference model standards that define network management concepts and operations, network management protocols, and basic communications protocols.

In the category of reference model standards, the Open Systems Interconnection (OSI) family of international standardized profiles (ISPs) provides basic definitions for network management functions. In addition to the definitions provided in this report for the five major functional areas (defined in sec. 3), ISPs have been published that describe aspects of fault detection (called Alarm Reporting) [ISO/IEC 10164-4], alarm filtering [ISO/IEC 10164-5], and log control [ISO/IEC 10164-6]. The OMNIPoint Specifications and Technical Reports [NMF92a], [NMF92b], and [NMF92c] created by the Network Management Forum, describe a set of guidelines for the specification and implementation of selected network management applications. These guidelines are based on the use of standardized descriptions of management information, base standards, ISPs, and conformance requirements. All of the standards listed above will be part of the foundation of the domain model.

In the category of network management protocols, the role of the SNMP and CMIP protocols was described in section 3. The current state of these standards and their current strengths and limitations in performing management operations must be taken into account in the domain model. The Management Information Base (MIB), also discussed in section 3, contains a description of low level data structures that may influence the domain model. In addition, the Structure of Management Information (SMI) provides a meta-model of how information is represented within the MIB.

The last category, the basic standards for carrying out communications services, include the TCP/IP and OSI suites of protocols as well as a large number of other standards that address specific areas of network communications. These standards will also have an impact on the domain model.

6.5 Overview of the Domain Modeling Process

The domain model for alarm surveillance will be developed iteratively. Each of the functional areas of alarm surveillance described in section 5.2 will be modeled. Guidelines for modeling individual functional areas were discussed in section 6.2. Initially, a high-level entity-relationship model, feature model, and functional model will be developed for the entire domain. The models for the major functional areas--fault detection, alarm filtering, and log control--will be all be extended to greater levels of detail. Thus, each of the major functional areas will have its own entity-relationship model, feature model, and functional model.

The entity-relationship model, feature model, and functional model are interrelated. Changes to one model may affect the evolution of the other models. In this way the development of each model can be expected to "drive" the development of the others. The process of developing increasingly detailed models for the functional areas will continue until all are specific enough to permit the identification of requirements for applications. At that point, the domain model will be regarded as complete. The development of the domain architecture will commence after the initial modeling of all major functions in the target subdomain.

7. DOMAIN DICTIONARY

A dictionary of major terms and concepts is presented. This dictionary is preliminary. It will be expanded and refined during the subsequent domain modeling phase. Within individual definitions, terms in bold print are defined as separate entries. Sources for definitions are cited in brackets. The use of the notation "+ed" indicates that the original definition has been edited. The notation "(NIST)" indicates that definition was developed by the Domain Analysis Case Study team.

Accounting Management:

The collecting and storing of information about the utilization of **network resources** for purposes of generating bills for **end users** and/or end-user groups. [NMF92c] +ed

Agent:

A software system residing on a **managed device** that is responsible for carrying out **network management** functions on the device, including the execution of **management operations** initiated by the **manager**. [FIPS179], (NIST)

Alarm:

A particular type of **notification** that conveys information about a detected **fault**. An alarm may include information about the type of fault that has been detected, information about the circumstances surrounding the fault, and an estimate of the severity of the fault. [ISO/IEC 10164-4] +ed

Alarm Analysis:

The process of **alarm filtering**, **alarm correlation**, and **fault prediction**. See Alarm Filtering, Alarm Correlation, and Fault Prediction. [GTE93]

Alarm Correlation:

The process of comparing information in different **alarms** to determine if they arose from the same, or related, **faults** or **fault conditions**. [GTE93], [NMF92b] +ed

Alarm Filtering:

The process of controlling the flow of **alarms** and other **notifications** describing events that have occurred to ensure that (1) the resources of **network management** are directed to the **faults** whose resolution is necessary to effectively manage the **network**, and (2) network performance is not degraded through transmission of redundant or useless information. Filtering includes using **discrimination criteria** to select alarms to be forwarded and to select alarms to be discarded. Filtering also includes determining the destination to which alarms should be forwarded. [GTE93], [ISO/IEC 10164-5] +ed

Alarm Management:
See Alarm Surveillance.

Alarm Reporting:
An OSI term that refers to the communication of information about a possible detected **fault**. This information generally includes the identification of the **network device** or **network resource** in which the fault was detected, the type of the fault, its severity, and its probable cause. The fault is reported in a **notification** called an **alarm**. See Alarm and Fault Detection. [ISO/IEC 10164-4] +ed

Alarm Surveillance:
The set of functions that enable (1) the monitoring of the **communications network** to detect **faults** and fault-related events or conditions; (2) the logging of this information for future use in **fault detection** and other **network management** activities; and (3) the analysis and control of **alarms**, **notifications**, and other information about faults to ensure that the resources of network management are directed toward faults that affect the operation of the communications network. Analysis of alarms consists of **alarm filtering**, **alarm correlation**, and **fault prediction**. [NMF92b], [GTE93] +ed

Alert:
See Alarm.

Bandwidth:
The rate at which information can be transmitted in bits/second. (NIST)

Centralized Network Management Architecture:
A **network management architecture** in which one **manager** at a single location controls the entire **communications network**. See Network Management Architecture, Distributed Network Management Architecture, and Distributed Hierarchical Network Management Architecture. [GTE93] +ed

CMIP:
See Common Management Information Protocol.

CMIS:
See Common Management Information Service.

Common Management Information Protocol (CMIP):
A **network management protocol** that supports the basic services defined in **Common Management Information Service (CMIS)** by specifying lower-level data units for transmitting CMIS operations and **notifications**. CMIP is described in [ISO/IEC 9596-1]. See Common Management Information Service (CMIS).

Common Management Information Service (CMIS):

A **network management protocol** that defines a set of basic network management services used by **network management systems** in the **Open Systems Interconnection** suite of protocols. CMIS includes but is not limited to **management operations**, called management-operations services, that are initiated by **managers**. It also includes **notifications** that can be sent by **agents**, referred to as management-notification services. CMIS is described in [ISO/IEC 9595].

Communications Network:

A system consisting of a **network** and a set of **communications services** that enable transmissions (voice, data, or other forms) to take place between persons, software applications, or other equipment. The users of the communications network are called **end users**. (NIST)

Communications Services:

The set of capabilities provided by software systems and **network devices** that enable transmissions from one point to another in the **network**. (NIST)

Configuration Database:

A database describing **network configurations**. The database may include a history of changes made to the network configuration. (NIST)

Configuration Management:

As used in **network management**, the tracking and control of the **network resources** and their current and potential connections. Configuration management includes creating and maintaining an accurate inventory of:

- the resources associated with the **communications network**,
- each network resource's operating characteristics--described by values of specific internal settings or variables,
- each network resource's logical and physical connections to other resources, and
- the **network topology** or any portion of it.

Configuration management controls the **network configuration**. It provides the means to change the operating characteristics of individual network resources, the logical and physical connections of resources, and the network topology. [NMF92e] +ed

Discrimination Criteria:

In **alarm filtering**, the criteria used to select whether or not to forward or discard an **alarm** or **trap**. See Alarm Filtering. [ISO/IEC 10164-5] +ed

Distributed Hierarchical Network Management Architecture:

A **network management architecture** with multiple levels of **manager** software systems in which (1) different managers have responsibility for different parts of the **communications network**, and (2) higher-level managers control lower-level managers. See Network Management Architecture, Centralized Network Management Architecture, and Distributed Network Management Architecture. [GTE93] +ed

Distributed Network Management Architecture:

A **network management architecture** in which different **manager** software systems have responsibility for different parts of the **communications network**. In a distributed architecture, managers are said to be in a "peer to peer" relationship with respect to each other and do not control each other. See Network Management Architecture, Centralized Network Management Architecture, and Distributed Hierarchical Network Management Architecture. [GTE93] +ed

End-System Device:

A platform containing hardware and possibly software systems that utilize the **communications services** provided on the **communications network**. (NIST)

End User:

The user of the services provided by a **communications network**. A user may be a person or a software application. (NIST)

Enterprise Network Management:

The complete set of activities required to establish, maintain, operate, and administer a **communications network**, its business functions, and its use by the **end-user** community. Enterprise network management includes **network management** as well as **provisioning**, providing help desk facilities, **network business administration**, and systems planning, design, and engineering. [MOLL92] +ed

Error:

The deviation of a system from normal operation that may have been caused by a **fault**. [ISO/IEC 10164-4]

Fault:

A physical malfunction or abnormal pattern of behavior that is causing or will cause, an **outage**, **error**, or degradation of **communications services** on a **communications network**. [ISO/IEC 10164-4], [GTE93] +ed

Fault Condition:

A set of circumstances associated with a **network resource** or group of resources that is likely to lead to a **fault**. [GTE93] +ed

Fault Correction:

The corrective action(s) taken in response to a **fault** necessary to restore **communications services**. Corrective actions may include temporarily reconfiguring a portion of the **communications network** or repair of **network devices**. [NMF92c], [NMF92d] +ed

Fault Detection:

The process of discovering and reporting active **faults**, potential faults, and fault-related conditions. See Alarm Reporting. [NMF92d]

Fault Diagnosis:

The process by which (1) a detected **fault** is isolated or narrowed down to a specific **communications network** element in which the fault occurred, and (2) additional information about the circumstances of the fault is obtained that will be needed for **fault correction**. See Fault Correction. [NMF92c] +ed

Fault Management:

The detection, reporting, diagnosis, correction, and prevention of **faults** and **fault conditions**. Fault management includes **alarm surveillance**, **trouble tracking**, **fault diagnosis**, and **fault correction**. [NMF92c] +ed

Fault Prediction:

The process of using current and historical information about **alarms** and other events to predict **faults** or to identify developing **fault conditions**. [GTE93] +ed

Fault Tracking:

See Trouble Tracking.

Intermediary Device:

A **network device** that influences the transmission path in some way. Typically, intermediary devices are network devices that connect different subcomponents of the **network**. Examples are bridges, routers, and gateways. (NIST)

Internet Suite of Protocols:

A suite of **protocols** that grew out of early research by the Advanced Research Projects Agency (ARPA) in DoD and has since spread to many areas on government and industry. The Internet suite is usually referred to as **TCP/IP (Transmission Control Protocol/Internet Protocol)**. (NIST)

LAN:

See Local Area Network.

Link Device:

A **network device** that consists of, or includes, a **transmission medium** that propagates digital signals. (NIST)

Local Area Network (LAN):

A **communications network** that is usually limited to a single building or closely-spaced group of buildings. (NIST)

MAN:

See Metropolitan Area Network.

Managed Device:

A **network device** that is monitored and controlled by a **network management system**. Generally, a managed device contains an **agent** system that responds to a **network management protocol** and can execute **management operations** on the device. See Agent. (NIST)

Managed Network:

The set of all **network resources**, including **network devices** and **communications services**, that are subject to **network management** and that are controlled by a particular **network management system**. (A **communications network** may also contain network resources that are not subject to network management. These are not part of the managed network.) See Communications Network. See Managed Device. (NIST)

Management:

Generically, the monitoring and control of a set of resources, activities, or events. See Network Management. (NIST)

Management Application:

See Manager.

Management Information Base (MIB):

A standardized description of the information that must be maintained by **managed devices** responding to **network management protocols**. A MIB is implemented on a device as a collection of variables that can be used to describe and control the device's state. MIBs are described in [ROSE90a]. See Management Operation. [ROSE90a] +ed

Management Operation:

A low-level operation specified in a **network management protocol** that a **manager** can perform on a **network resource** in a **managed network**. A management operation may be a "GET" operation that obtains information from a network resource's **Management Information Base**. This information can be used to

interpret the state of a **managed device** or other network resource. A management operation may also be a "SET" operation that alters variable settings in order to control the operating status or configuration of a managed device or other network resource. (NIST)

Management Station:
See Network Management Station.

Manager:
A major part of a **network management system**--a manager is a software system that initiates actions for monitoring and controlling a set of **network devices** or other **network resources**. The actions of a manager include but are not limited to initiating **management operations** that are carried out by **agent** systems residing on network devices. Such devices are referred to as **managed devices**. Managers also respond to **notifications** sent by agent systems describing the status of a managed device. See Agent, Managed Device, Managed Network, Management Operation, and Network Management System. [FIPS179] +ed

Metropolitan Area Network (MAN):
A **communications network** located in a single city or metropolitan area. A MAN may be composed of LANs or even other MANs. (NIST)

MIB:
See Management Information Base.

Monitoring Strategy:
The procedure for monitoring a **communications network** to detect **faults** and **fault conditions** and to determine the status of the network's **managed devices** and other **network resources**. A monitoring strategy may involve a procedure for periodically **polling** individual **managed devices** to determine their status, receiving **alarms** automatically transmitted by **agent** systems residing on the devices, or a combination of these methods. See Alarm and Polling. (NIST)

Network:
A set of devices such as computers, terminals, and printers that are physically connected by a **transmission medium** so that they can communicate with each other. These devices are called **network devices**. See Communications Network, Managed Network, and Network Device. (NIST)

Network Business Administration:
The performance of business related functions for a **network** associated with budgeting, staffing, training, equipment procurement, and charging and billing **end users** of network services. [MOLL92]

Network Configuration:

A specific set of **network resources** that form a **communications network** at any given point in time, the operating characteristics of these network resources, and the physical and logical connections that have been defined between them. (NIST)

Network Control:

The initialization and shut down of **network resources**. (NIST)

Network Control Center:

See Network Operations Center.

Network Device:

A device that is part of and can send or receive electronic transmissions across a **communications network**. Network devices include: **end-system devices** such as computers, terminals, or printers; **intermediary devices** such as bridges and routers that connect different parts of the communications network; and **link devices** or transmission media. [ROSE91], [TANEN88] +ed

Network Directory Services:

The capabilities of storing and retrieving information about the identities and addresses of **end users** and other **network resources** that can be accessed through the **communications services** provided by the **network**. [MOLL92] +ed

Network Management:

The discipline that describes how to monitor and control the **managed network** to ensure its operation and integrity and to ensure that **communications services** are provided in an efficient manner. As described in [ISO/IEC 7498-1], network management consists of **fault management**, **configuration management**, **performance management**, **security management**, and **accounting management**. (NIST)

Network Management Architecture:

The distribution of responsibility for **management** of different parts of the **communications network** among different **manager** software systems. The network management architecture describes the organization of the management of a network. The three types of network management architectures are the **centralized network management architecture**, the **distributed network management architecture**, and the **distributed hierarchical network management architecture**. (NIST)

Network Management Protocol:

A **protocol** whose purpose is to convey information pertaining to the **management** of the **communications network**, including **management operations** from **managers** as well as responses to **polling** operations, **notifications**, and **alarms** from **agents**. [ROSE91] +ed

Network Management Station:

A computing platform on which the **manager** system runs. In addition to the manager, the station contains the workstations, displays, automatic data processing hardware and software, and ancillary equipment and facilities used to provide the ability to manage the **communications network**. (NIST)

Network Management System:

A software system that performs functions of **network management** for a **communications network**. This system may include both **manager** and **agent** systems. Since network management is not fully automated, a network management system performs a subset of network management functions as defined in [ISO/IEC 7498-1]. (NIST)

Network Operations Center:

An installation or site that contains the personnel, equipment, and other **network resources** needed to operate, control, and maintain the portion of the **communications network** within its jurisdiction. (NIST)

Network Resource:

An element of a **communications network** to be managed, including a **network device** or a **communications service**. [NMF92a] +ed

Network Size:

The total number of **network devices** that must be managed within the **network** and all its subcomponents. (NIST)

Network Topology:

The term has two meanings: (1) the structure, interconnectivity, and geographic layout of a group of **networks** forming a larger network, and (2) the structure and layout of an individual network within a confined location or across a geographic area. (NIST)

Notification:

A message emitted by an **agent**. A notification may describe an event that has occurred within the **managed device**. [ISO/IEC 10040], [ISO/IEC 10164-4] +ed

Open System Interconnection (OSI) Suite of Protocols:

A suite of **protocols** developed under the auspices of the International Organization for Standardization/International Electrotechnical Committee (ISO/IEC). (NIST)

Outage:

The period of time for which a **communications service** is unavailable [ISO/IEC 10164-4].

Performance Management:

The process of monitoring and controlling a **communications network** to ensure that it operates efficiently. This function includes (1) the collection and evaluation of data that measures the efficiency of the communications network in meeting its operational objectives; and (2) the controlled change of any factors to improve that efficiency. [NMF92c] +ed

Polling:

The process of sending messages to individual **managed devices** to determine their operational status. (NIST)

Problem Report:

A problem reported by an **end user** that may or may not be a **fault** [NMF92d].

Proprietary Element Management System:

A software system intended to monitor and control a collection of **network devices** that respond to a **proprietary protocol**. Proprietary element management systems may be developed and supplied by commercial vendors. See Standards-Based Element Management System. [GTE93] +ed

Proprietary Protocol:

A **protocol, network management protocol,** or suite of protocols developed by a private company to manage **network resources** manufactured by that company. See Proprietary Element Management System. (NIST)

Protocol:

A formal description of message formats and rules that must be followed to exchange messages. Protocols can describe high-level exchanges between application programs (e.g., the way in which two programs transfer a file across an internet). Protocols may also describe low-level details of machine-to-machine interfaces (e.g., the order in which the bits from a byte are sent across a wire). Most protocols include both intuitive descriptions of expected interactions as well as more formal specifications based on finite state machine models. [COMER92] +ed

Provisioning:

The process of installation and customer assignment of communications equipment. Typically, provisioning is an off-line process and may take days to complete [NMF92d].

Resource:

See Network Resource.

Scope:

The geographic area that a **network** spans. (NIST)

Security Management:

The process of monitoring and controlling access to **network resources**. This includes monitoring usage of network resources, recording information about usage of resources, detecting attempted or successful violations, and reporting such violations. [ISO/IEC 7498-4], [NMF92c] +ed

Service Provisioning:

See Provisioning.

Simple Network Management Protocol (SNMP):

A **network management protocol** used with the **TCP/IP** suite of protocols. SNMP specifies a set of management operations for retrieving and altering information in a **Management Information Base** (MIB), authorization procedures for accessing MIB tables, and mappings to lower TCP/IP layers. An expanded version of SNMP is being created, called SNMPv2, that will provide additional services. SNMP is described in [CASE90] and [ROSE91].

Size:

See Network Size.

SNMP:

See Simple Network Management Protocol.

Standards-Based Element Management System:

A software system that monitors and controls a collection of **network devices** that respond to either the **Open System Interconnection (OSI) Suite of Protocols** or **TCP/IP (Transmission Control Protocol/Internet Protocol)**. Standards-based element management systems may be developed and supplied by commercial vendors. See Proprietary Element Management System. [GTE93] +ed

Systems Management:

See Enterprise Network Management.

TCP/IP (Transmission Control Protocol/Internet Protocol):

See Internet Suite of Protocols.

Topology:

See Network Topology.

Transmission Medium:

A mechanism that supports propagation of digital signals. Examples of a transmission medium are cables such as leased lines from common commercial carriers, fiber optic cables, and satellite channels. [TANEN88]

Trap:

A message indicating that a **fault condition** may exist or that a **fault** is likely to occur. See Alarm. [GTE93]

Trouble Tracking:

The process of recording, updating, maintaining, and forwarding information about the progress of a reported **fault** to support personnel or automated systems in order to ensure the fault is diagnosed and corrected in a timely manner. (NIST)

WAN:

See Wide Area Network.

Wide Area Networks (WAN):

A **communications network** that covers several sites that are geographically distant. A WAN may span different cities or even different continents. (NIST)

8. REFERENCES

[ADAMS91] Adams, E., "Global Commonality in User Requirements," in Integrated Network Management, II: Proceedings of the IFIP TC6/WG6.6 Second International Symposium on Integrated Network Management, I. Krishnan and W. Zimmer (eds.) held Washington, DC, April 1-5, 1991, pp. 171-181.

[CARB90] Carbonell, J. (ed.), Machine Learning: Paradigms and Methods, Elseiver Science Publishers B. V., Amsterdam, The Netherlands, 1990.

[CASE90] Case, J., et al., A Simple Network Management Protocol, Request for Comments 1157, DDN Network Information Center, SRI International, May 1990.

[CASE93] Case, J. and M. Rose, "The Simple Network Management Protocol (SNMP) for Internet Network Management," Tutorial presented at INTEROP 93, Washington, DC, March 1993.

[COHEN92] Cohen S., J. Stanley, S. Peterson, and R. Krut, Application of Feature-Oriented Domain Analysis to the Army Movement Control Domain, CMU/SEI-91-TR-28, Software Engineering Institute, Carnegie-Mellon University, Pittsburgh, PA, 1992.

[COMER91] Comer, Douglas E., Internetworking with TCP/IP, Volume I: Principles, Protocols, and Architecture, Second Edition, 1991, Prentice-Hall, Inc., Englewood Cliffs, NJ 07632.

[CRONK88] Cronk, R., P. Callahan, and L. Bernstein. "Rule-Based Expert Systems for Network Management and Operations: An Introduction," IEEE Network, Volume 2, Number 5, September 1988, pp. 7-21.

[DABR93] Dabrowski, C. and T. Kirkendall, Preliminary Report on Domain Analysis Methods, Produced for the Software Producibility MODIL, National Institute of Standards and Technology, Gaithersburg, MD, February 1993.

[DISA93] Domain Analysis and Design Process, Version 1, The Defense Information Systems Agency, Center for Information Management, Software Reuse Program, Document No. 1222-04-210/30.1, March 30, 1993.

[DoD92] DoD Software Reuse Initiative Vision and Strategy, Department of Defense, Document No. 1222-04-210/40, 1st edition, July 1992.

[FIPS179] Government Network Management Profile (GNMP), Federal Information Processing Standards Publication 179, National Institute of Standards and Technology, December 1992.

[GTE93] Conversations With Domain Experts from GTE Government Systems, April-June, 1993.

[HERT91] Hertz, J., S. Krogh, and R. Palmer, Theory of Neural Computation, Addison-Wesley Publishers, Redwood City, CA, 1991.

[HOLI91] Holibaugh, R., "Joint Integrated Avionics Working Group (JIAWG) Object-Oriented Domain Analysis Method (JODA)--DRAFT DOCUMENT," Revision Number 3.1, Software Engineering Institute, Carnegie-Mellon University, Pittsburgh, PA, January 1991.

[HONG91] Hong, P. and P. Sen, "Incorporating Non-deterministic Reasoning in Managing Heterogeneous Network Faults," in Integrated Network Management, II: Proceedings of the IFIP TC6/WG6.6 Second International Symposium on Integrated Network Management, I. Krishnan and W. Zimmer (eds.) held Washington, DC, April 1-5, 1991, pp. 481-491.

[ISO/IEC 7498-4] International Organization for Standardization, "Information Processing Systems - Open Systems Interconnection - Basic Reference Model - Part 4: Management Framework," 1989.

[ISO/IEC 9595] International Organization for Standardization, "Information Technology - Open Systems Interconnection - Common management information service definition [for CCITT Applications]," ISO JTC1/SC21, 1990.

[ISO/IEC 9596-1] International Organization for Standardization, "Information Technology - Open Systems Interconnection - Common Management Information Protocol - Part 1: Specification," ISO JTC1/SC21, 1990.

[ISO/IEC 10040] International Organization for Standardization, "Information Technology - Open Systems Interconnection - Systems Overview," ISO JTC1/SC21, 1990.

[ISO/IEC 10164-4] International Organization for Standardization, "Information Technology - Open Systems Interconnection - Systems Management - Part 4: Alarm Management," ISO JTC1/SC21, 1991.

[ISO/IEC 10164-5] International Organization for Standardization, "Information Technology - Open Systems Interconnection - Systems Management - Part 5: Event Report Management Function," ISO JTC1/SC21, 1991.

[ISO/IEC 10164-6] International Organization for Standardization, "Information Technology - Open Systems Interconnection - Systems Management - Part 6: Log Control Function," ISO JTC1/SC21, 1991.

[KANG90] Kang, K., S. Cohen, J. Hess, W. Novak, and S. Peterson, Feature-Oriented Domain Analysis (FODA) Feasibility Study, CMU/SEI-90-TR-21, Software Engineering Institute, Carnegie-Mellon University, Pittsburgh, PA, 1990.

[KATZ93] Katz, S., C. Dabrowski, and M. Law, Glossary of Reuse Terms, Draft Produced for the Software Producibility MODIL, National Institute of Standards and Technology, Gaithersburg, MD, July 1993,

[MOLL92] Moller, R., J. Rutter, and B. Zielinski, " Network and Systems Management Functional Definition," Report Number MTR92B0000121, Mitre Corporation, August 1992.

[NMF92a] Network Management Forum, "Reconfigurable Circuit Service: Configuration Management Ensemble," OMNIPoint Specifications and Technical Reports, Book I and II, August 1992.

[NMF92b] Network Management Forum, "Reconfigurable Circuit Service: Alarm Surveillance Ensemble," OMNIPoint Specifications and Technical Reports, Book I and II, August 1992.

[NMF92c] Network Management Forum, "Application Services: Path Tracing Function," OMNIPoint Specifications and Technical Reports, Book I and II, August 1992.

[NMF92d] Network Management Forum, Statement of User Requirements for Management of Networked Information Systems, October 1992.

[NMF92e] Network Management Forum, "Reconfigurable Circuit Service: Configuration Management Ensemble," OMNIPoint Specifications and Technical Reports, Book I and II, August 1992.

[PETE91] Peterson S., and S. Cohen, A Context Analysis of the Movement Control Domain for the Army Tactical Command and Control System (ATCCS), CMU/SEI-91-SR-3, Software Engineering Institute, Carnegie-Mellon University, Pittsburgh, PA, 1991.

[PRIE90] Prieto-Diaz, Ruben, "Domain Analysis: An Introduction," ACM SIGSOFT Software Engineering Notes, Vol. 15, No. 2, April 1990, pp. 47-54.

[PRIE91] Prieto-Diaz, R., The Reuse Library Process Model, IS-40.2 03041-002, STARS Reuse Library Program, New York, March 1991.

[RABIE88] Rabie, S., A. Rau-Chaplin, and T. Shihahara. "DAD: A Real-Time Expert System for Monitoring of Data Packet Networks," IEEE Network, Volume 2, Number 5, September 1988, pp. 29-34.

[RICH91] Rich, E. and K. Knight, Artificial Intelligence (second edition), McGraw-Hill Publishers, 1991.

[ROSE90a] Rose, M. (ed.), Management Information Base Network Management of TCP/IP based Internets: MIB-II, Request for Comments 1158, DDN Network Information Center, SRI International, May 1990.

[ROSE90b] Rose, M. and K. McCloghrie, Structure and Identification of Management Information for TCP/IP based Internets, Request for Comments 1155, DDN Network Information Center, SRI International, May 1990.

[ROSE91] Rose, M. The Simple Book: An Introduction to Management of TCP/IP-based Internets, Prentice-Hall, Inc., Englewood Cliffs NJ, 1991.

[SHAV90] Shavlik, J. and T. Dietterich (eds.), Readings in Machine Learning, Morgan Kaufmann Publishers, Inc., San Mateo, CA, 1990.

[SPC92] Domain Engineering Guidebook, SPC-92019-CMC, version 01.00.03, Software Productivity Consortium, Herndon, VA, December 1992.

[STAL88] Stallings, W., Data and Computer Communications, (second edition), Macmillan Publishing Company, New York, NY, 1988.

[STAL93] Stallings, W., Networking Standards: A Guide to OSI, ISDN, LAN, and MAN Standards, Addison-Wesley, Reading, MA, 1993.

[STARS93] Organizational Domain Modeling, Volume I - Conceptual Foundations, Process And Workproduct Description, Informal Technical Report for the Software Technology for Adaptable, Reliable Systems (STARS), Report Number STARS-UC-05156/024/00, July 31, 1993.

[SUTT88] Sutter, M. and P. Zeldin, "Designing Expert Systems for Real-Time Diagnosis of Self-Correcting Networks," IEEE Network Magazine, Volume 2, Number 5, September 1988, pp. 43-51.

[TANEN88] Tanenbaum, A., Computer Networks, Prentice-Hall, Englewood Cliffs, NJ, 1988.

[WART92] Wartik, S. and R. Prieto-Diaz, "Criteria for Comparing Reuse-Oriented Domain Analysis Approaches," International Journal of Software Engineering and Knowledge Engineering, Volume 2, Number 3, September 1992, pp. 403-431.

[WASS89] Wasserman, P., Neural Computing: Theory and Practice, Van Nostrand Reinhold Publishers, New York, NY, 1989.

[YOUR89] Yourdon, E., Modern Structured Analysis, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1989.