METRIC

MIL-HDBK-818-1

31 October 1992

MILITARY HANDBOOK

SURVIVABLE ADAPTABLE FIBER OPTIC EMBEDDED NETWORK (SAFENET) NETWORK DEVELOPMENT GUIDANCE

AMSC N/A

AREA MCCR

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.

FOREWORD

1. This military handbook is approved for use by all Departments and Agencies of the Department of Defense.

2. Beneficial comments (recommendations, additions, deletions) and any pertinent data which may be of use in improving this document should be addressed to: Space and Naval Warfare Systems Command, Code 231-2B2, Washington, D.C., 20363-5100, by using the self-addressed Standardization Document Improvement Proposal (DD Form 1426) appearing at the end of this document or by letter.

3. This document provides guidance for the development of fiber optic local area networks which are intended for use in support of mission critical computer resources. It is a companion document to MIL-STD-2204, the SAFENET standard.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1. SAFENET OVERVIEW

1.1 Background. SAFENET is a network standard defined in MIL-STD-2204. It defines standards for a set of components which can be used in subsystems of mission critical computer resources. Historically, Navy computers have usually been interconnected with point-to-point interfaces. However, the distributed architectures of modern Navy combat systems require a greater degree of system integration than these point-to-point interfaces can support. For instance, the degree of connectivity of some systems is currently limited by the number of I/O ports in the computers. The SAFENET standard was developed to solve these connectivity and system integration problems by providing these computers with the capability to communicate with multiple application programs and devices over a single I/O port through the use of a computer network.

SAFENET provides standards for an integrated set of components for use in communication subsystems which use a local area network (LAN) for intercomputer and computer to peripheral data transfer. It is based on existing and proposed commercial network standards to permit the use of commercial components where applicable. The basic element of SAFENET is a dual-ring token passing LAN which interconnects the attached computers and peripherals.

1.2 Network features. In addition to reducing the number of I/O ports required to achieve interconnectivity, SAFENET provides several other features:

a. Flexible implementation - Except for the physical medium specification, the SAFENET standard does not restrict the hardware aspects of an implementation in any way. All requirements (other than for physical medium) are in terms of logical interfaces and protocol definitions, which restrict software functionality rather than the actual hardware. This means that a SAFENET interface can be implemented either as an embedded interface within a host computer, or as a stand-alone interface unit to which devices can be connected.

b. Survivability - SAFENET contributes to system survivability by providing automatic fault isolation and reconfiguration as a function of its dual-ring LAN. Additional survivability may be required in some applications, and SAFENET accommodates this possibility by defining the physical topology components from which

a network may be built, rather than defining the topology itself.

c. Fiber optics - SAFENET employs a fiber optic transmission medium. This medium has many advantages over copper cable including low weight, high bandwidth, and insensitivity to electromagnetic interference (EMI).

d. Multiple vendors - SAFENET is based on commercial standards and does not require any proprietary technology. For these reasons it should be possible to obtain SAFENET compliant networks and components from multiple vendors.

e. Coexistence with alternate I/Os - Use of SAFENET on one or more of a computer's I/O ports should not interfere with that computer's ability to communicate point-to-point over other I/O ports.

1.3 SAFENET profile. SAFENET employs a layered protocol architecture which is based on the ISO Open Systems Interconnection (OSI) reference model (ISO 7498) for computer networks. Within this layered architecture SAFENET specifies one or more protocols at each layer. The complete set of SAFENET specified protocols is known as the SAFENET profile.

SAFENET provides some implementation options in the form of different profiles. For instance, the OSI profile is available for full OSI compliant networking, while the lightweight profile can be used to support real-time data transfer. A SAFENET station may implement either of these profiles alone, or both of them (the combined profile). Note that both the OSI and lightweight profiles use the same token ring LAN for actual data transfer. The concept of alternate OSI and lightweight profiles is illustrated in Figure 1.

SAFENET also provides flexibility in the implementation of its physical topology. Rather than defining a limited number of specific topologies, components are defined from which to build a network. Thus the network may be designed for optimum use in any specific application.

1.4 Profile selection. The following paragraphs provide guidance on the selection of SAFENET profiles. It is intended to help system designers determine the appropriate profile for the requirements of their network.

1.4.1 Features common to all profiles. There are several features which are common to all SAFENET profiles. Foremost

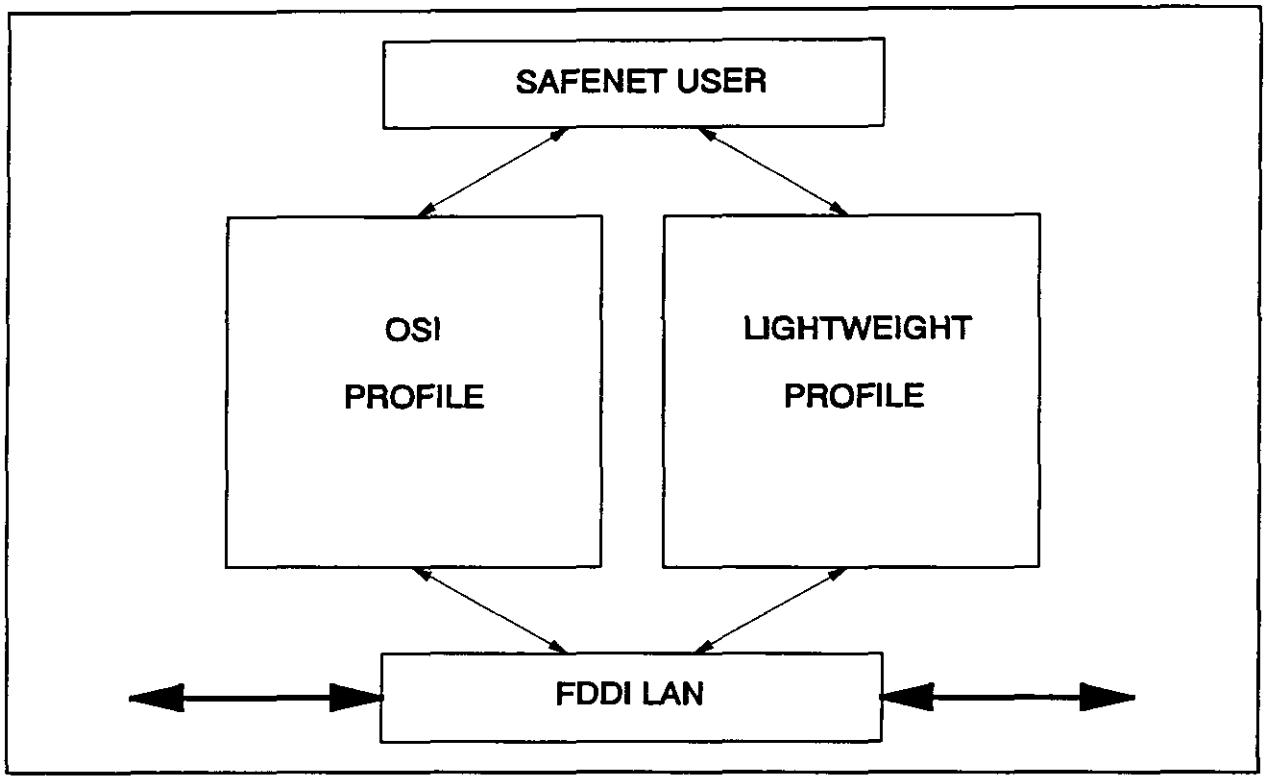

Figure 1. Alternate Profiles

among these is the local area network and the logical link control protocol. Because all SAFENET stations share common protocols through the physical and data link layers, it is possible to attach stations which implement different profiles to the same LAN. Of course, stations using different profiles may not be able to interoperate, but they will be able to share the use of the LAN.

Another important feature which is common to all profiles is the SAFENET Time service. This service provides the capability for all the stations on the LAN segment to synchronize their clocks.

1.4.2 OSI profile. The OSI profile is intended for situations where either the interoperability of independently developed SAFENET nodes is a driving consideration, or the functionality of File Transfer, Access, and Management (FTAM) is required, or the system is of such complexity that network management is required. While the OSI profile provides these capabilities, it does so at the expense of increased data

transfer latency and an inability to send the same data to multiple users simultaneously (multicast).

Within the OSI profile, the FTAM protocol provides robust file transfer and management services. Additionally, the Private Communication (Private Comm) interface provides a user with the capability to directly establish connections with remote users for the purpose of passing messages. The services of the SAFENET OSI profile can be used to support additional OSI-defined application layer services such as Message Handling System (MHS) (X.400), which provides an e-mail capability. Some of these services may require extensions to the underlying presentation and session services. This is permissible as long as such implementations are interoperable with other implementations of SAFENET at the intersection of their capabilities. Finally, the network management services provide a user with the capability to manage the network resources.

1.4.3 Lightweight profile. The lightweight profile is intended for situations where either data transfer latency is critical, or multicast data transfer is required, or system specific support services need to be implemented in place of ISO standard session, presentation, and application layer services. While the lightweight profile provides these capabilities, it does so at the loss of the ISO standard protocols and network management mechanisms. Furthermore, since the lightweight profile permits system specific implementation, interoperability is limited to those stations which have implemented identical lightweight profiles.

Within the lightweight profile, the lightweight application service definition provides the user with streamlined access to the transport layer services, while permitting the use of system specific support services at the application layer. Additionally, Xpress Transfer Protocol (XTP) and the ISO Connectionless (CL) transport protocol are available for efficient, low latency data transfer, along with multicast data transfer capabilities.

1.4.4 Combined profile. The combined profile is intended for situations where the capabilities of both the OSI and lightweight profiles are required. In the combined profile, all the services and capabilities described above are available to the user. In addition, this profile provides management support for the protocols which comprise the lightweight profile. The robust functionality of the combined profile comes at the expense of the complexity of the system and the cost of its development.

2. SAFENET USER SERVICES

2.1 Introduction. The purpose of this section is to provide an overview of the application, presentation, and session layers of SAFENET, which are collectively called the SAFENET user services.

2.2 User services overview.

2.2.1 Assumptions and guidelines. The general assumptions made and guidelines followed in specifying the SAFENET user services and application interfaces are described below.

First, the physical implementation of the SAFENET environment has no effect upon the SAFENET environment application interface. This does not mean that the physical implementation does not affect the SAFENET environment user, but only that it does not affect the application interface to the SAFENET environment subsystem. For example, the actual physical implementation might affect the overall performance of the application program involved and therefore might have to be taken into account during program analysis and design.

Second, the physical location of the information objects remote to a SAFENET environment user is transparent to that user. The global name or logical name is all that the user needs to know. This information object might be resident in a peripheral device, a remote computer, or within the local computer itself.

Third, the user services perform all functions requiring communication with remote nodes through the transfer services. This means that the only interface between the SAFENET environment users and the lower layers of the OSI reference model is through the transfer services interface.

Fourth, all SAFENET environment users will be provided access to the capabilities of SAFENET either through the MAP application interface or through an application interface specific to the lightweight application services.

2.2.2 Conformance, interoperability and performance. The issues of conformance, interoperability, and performance need to be considered at several levels in a system design. The following points are given as examples.

First, interoperability at the application layer interface using a standard application layer protocol does not guarantee that two processes in different application programs can operate

together. Information of the type found in Interface Requirements Specifications (IRSs) and Interface Design Documents (IDDs) are also required.

Second, the use of an application layer service such as FTAM or Directory Services may considerably ease the problem of making two application programs operate together over a SAFENET environment.

Finally, performance needs to be considered from several points of view, including the amount of capability or functionality offered by a service and the resources (e.g., time, memory, transmission bandwidth) required to carry out the functionality.

It is the intent of SAFENET to offer the system designer a number of options so that an appropriate balance among the often conflicting system requirements for performance and interoperability can be achieved. However, doing this in the face of the incomplete state of OSI standards and the Government Open Systems Interconnection Profile (GOSIP) upon which this work is based means that certain key components cannot be specified. Therefore, interoperability of all SAFENET functions is not guaranteed simply by conformance to current SAFENET definitions.

Two examples follow. First, the area of network management services is not adequately defined. Second, time efficient user services, tailored to current Navy practice in inter-process message passing, do not exist in the OSI protocol set. Therefore, the concept of lightweight user services is necessary to permit the local definition of the needed services and protocols. This represents a direct tradeoff of performance for interoperability.

2.2.2.1 SAFENET environment structure. In discussing the issues of conformance, interoperability, and performance in SAFENET a brief discussion of the structure of a real open system as it might appear in Navy applications using SAFENET is required. The real open systems of primary concern to SAFENET will generally be found within the confines of a single ship or submarine. They will also be found in land-based test or development sites for the above.

Open systems for Navy applications will be designed under the guidance of MIL-STD-490A and DOD-STD-2167A. These standards call out a series of documents which together establish the functionality and performance requirements of the system or subsystem being designed and implemented, how its components interact with each other, and how the system and its components

interact with other systems or subsystems with which it may be joined to carry out higher level functions. These documents include the System/Segment Specification (SSS), Software Requirements Specification (SRS), Software Top Level Design Document (STLDD), Interface Requirements Specification (IRS), and Interface Design Document (IDD). At the level of the application-program in an open system these establish conformance, interoperability, and performance requirements. They indicate the type, direction, and meaning of information flow among system components. This type of design information is a necessity whether the data transmission media is shared memory, a computer backplane bus, or a LAN segment using SAFENET components.

Issues of performance, interoperability, and conformance at the system and application program level are the responsibility of the system designer. SAFENET assists the system designer by supplying data transmission components of known performance characteristics and interoperability upon which the system can be built.

An open system using SAFENET components is considered to be built of nodes. Depending on the complexity of the system the nodes may all be on one LAN segment or may be on a number of LAN segments interconnected by routers. In addition to using SAFENET, the nodes may be connected through the use of other communication media. In addition to being connected to a SAFENET LAN segment a node could be connected to a non-SAFENET LAN segment so that it could process information from both sources. It is also possible that a single SAFENET LAN segment or a set of interconnected LAN segments could be shared by nodes belonging to several different independently developed systems. These systems could be completely separate, except for sharing the network, or they could be cooperating in performing some higher level system function.

2.2.2.2 Conformance. The term conformance is used in two different senses. In the first sense it indicates whether a service, function, or component is required in a node so that it may be called a SAFENET node. In the second sense it means that an item, if implemented, must be in accordance with its specification. In certain cases, to permit maximum tailoring to specific user needs or because of the lack of adequate definition in existing standards, SAFENET provides only a functional description. In such cases, interoperability is a local matter and is solely the concern of the system designer and implementer.

2.2.2.3 Interoperability. Typical Navy mission critical computer resource systems are composed of networks, or

hierarchies of independently developed systems, built by different vendors under different constraints. Since options are available within the SAFENET standard there is a need to ensure interoperability among component systems. A major impetus behind standardization comes from the desire to have different vendor's equipment communicate with each other.

As stated earlier, SAFENET constitutes a communication architecture consisting of a set of standardized options which will be integrated into Navy systems. Interoperability can therefore be achieved only when different nodes demonstrate the following characteristics. First, the nodes share the same protocol stack. Second, the data flows established are configured with the same parameters. Third, policies administered throughout the network are implemented consistently.

At a different level, the phrase "SAFENET interoperability" must be qualified by the system environment for which the component is intended. Components designed for different system architectures that include the same protocols may be capable of exchanging data but are not necessarily interoperable. Possible design tradeoffs involve issues such as performance, policy (including management policy), parameterization, operating system, and implementation architecture (embedded or non-embedded).

2.2.2.3.1 Application process interoperability. After the need for and meaning of information interchange has been established by an appropriate design document, interoperability among different application processes, including those in application programs, on a SAFENET network is ensured by adherence to SAFENET protocols. Nodes which do not provide the MAP 3.0 application interface, along with any necessary user defined interface, must then provide a defined interface to the lightweight application services. This then becomes the minimum level of functionality required for interconnection between application programs on SAFENET. It must be noted that this level of functionality will not provide interoperability without a further definition of at least directory service functions. In the current version of SAFENET this definition is a local matter.

2.2.2.3.2 LAN segment interoperability. Physical and logical interoperability among SAFENET LAN segments and between SAFENET and non-SAFENET LAN segments can be accomplished with the use of routers. For effective communication, the application processes must use common protocols and be in agreement as to the meaning of the communication. Interoperability between application processes on different SAFENET LAN segments will be facilitated by common use of SAFENET protocols. Interoperability

between SAFENET and non-SAFENET application processes may also be facilitated through use of SAFENET protocols by both parties. Any other scheme for achieving interoperability is beyond the scope of SAFENET and will not be treated further in this document.

2.2.2.4 Coexistence of OSI and lightweight protocols. Both OSI and lightweight protocols can coexist in the same transfer services partition with the understanding that end-user components which use either protocol will be able to communicate only with end-user components using the same protocol. Access to the OSI connection-oriented transport protocol can only be gained through the MAP application interface. Access to the lightweight transport protocols, XTP and OSI connectionless transport, can only be gained through the lightweight application services, which may be implemented to provide only those services required by a specific application. In many cases, it is expected that this application specific set of services will effectively allow direct access to the desired lightweight transport service.

2.2.3 User services functions. The SAFENET user services are the vehicle through which the designer of a subsystem or real open system using the SAFENET environment as a component provides communication capabilities to an application program. The user services, using the available lower layer functions of the SAFENET environment, provide the following capabilities within a real open system: information transfer, LAN node management, and system coordination.

It should be noted that in a specific system design additional functions will usually be required and some of the functions listed may not be needed or may only be required in a very limited form. Non-SAFENET components may also participate in providing the above functions. Thus in order to balance timing, cost, reliability, damage resistance, and other needed system properties a number of SAFENET and non-SAFENET communication options are needed. These permit the system designer the capability of selecting the most suitable solution to the requirements of the system being designed.

The remaining paragraphs in this section provide a brief description of the features of each of these functions. Note that this functional view of an implementation of SAFENET has a many-to-many mapping to the components as they are implemented in accordance with the OSI reference model.

2.2.3.1 Information transfer. The information transfer function is concerned with the transfer of information among SAFENET environment users. This information may consist only of

data, or it may indicate the need for, or cause, an action or change of state of the user. Data which needs to be exchanged or shared may range from a simple message or set of variables to complex file structures. Timing requirements for the accomplishment of the desired service can range from milliseconds to minutes and can be either real time or non-real time. A specific system may have a wide range of requirements. For these reasons several different data transfer services are offered by SAFENET.

2.2.3.2 LAN node management. The LAN node management function is concerned with the control of the total set of functions in a LAN node. This includes subfunctions such as enabling or disabling the LAN node in the SAFENET environment, reconfiguration within the LAN node in the SAFENET environment operations, local and network error monitoring and display, error threshold determination and assignment, and processing of local SAFENET environment status information.

2.2.3.3 System coordination. The system coordination function is concerned with the operation and synchronization of the real open system as a whole. This includes subfunctions such as software reconfiguration within the system, downloading and initialization of software across the SAFENET environment, communication coordination with peripheral devices connected directly to the SAFENET environment, initial start-up procedures of the system including its SAFENET environment, time synchronization of users within the network, and global control of the hardware, firmware and software contained within the system. It is also concerned with such issues as logical-name registration and resolution to physical N-address, priority definition and assignment, and service type resolution.

2.2.4 Application program computational model. It is expected that in an open system of any complexity a variety of computational models or designs will be required in the application program to carry out the needed functions. If SAFENET user services are used at each node in such a system then the various implementations of the user services will have to interface not only with each other through their transfer services but to the application programs with their varying computational models.

The following has been assumed concerning the application program computational model in selecting the SAFENET user services standards.

First, the SAFENET user services will be interfaced to application programs which pre-date the definition of SAFENET to

support the integration of existing combat system elements with new elements. It will be, however, most frequently interfaced with application-programs written specifically with SAFENET in mind.

Second, the SAFENET user services will be used in a multi-tasking environment that supports inter-process communication and coordination.

Third, the SAFENET user services are targeted for use with application programs written in programming languages that can associate multiple data structures with the same data space. The language must also allow the users to define data types and complex data structures. Examples of such languages are C and Ada.

Finally, the design of the user services at any given time cannot address all possible application program requirements. Therefore it must be extensible to meet new requirements.

2.3 Components of user services.

2.3.1 Application layer. The ISO application layer provides those services which are directly visible to the end user. When the user is a process or program an Application Program Interface (API) or set of calling sequences is needed to provide access to the services offered by application layer services. Since suitable standards were not available at the time of publication of this document SAFENET does not require standard API's for the application layer services it requires except for the Private Communication interface derived from the MAP 3.0 specification. This is supplied to provide a simple message passing capability in the OSI profile. The following sections provide a brief description of the SAFENET defined services in the application layer.

2.3.1.1 Association Control Service Element. The Association Control Service Element (ACSE) is a component of the OSI profile which provides services for the establishment and termination of application associations. ACSE is intended to provide a standard service for applications to communicate common parameters, such as titles, N-addresses and application context, during the request for an association.

2.3.1.2 Directory services. In the OSI profile, the National Institute of Standards and Technology Stable Implementation Agreements directory system provides a repository for information about network objects. A typical directory user would supply an application-entity title and get back the

presentation address of the entity. The typical directory consists of a Directory Information Base (DIB), a directory service called the Directory System Agent (DSA), and several Directory User Agents (DUAs).

The lightweight profile must also be supported with some form of directory services. How this capability is provided is a local matter. This could be implemented within lightweight user services as a fixed table provided as a part of the system design, or as a dynamic system similar to the directory system discussed above.

2.3.1.3 Private Communications. In the OSI profile, the Private Comm Application Interface is designed to support communications between two application processes which need to exchange messages in a manner which requires the full functionality of the OSI profile. It allows application processes virtually direct access to the ACSE services for the establishment of associations (connections), and to the presentation services for the exchange of data.

2.3.1.4 File Transfer, Access and Management. In the OSI profile, the FTAM protocol provides a set of services for transferring information conveniently between application processes and filestores. It supports several file types including sequential, random access, and single key indexed sequential files. FTAM provides the ability to: work with binary or text files, create and delete files, transfer entire files, read and change file attributes, erase file contents, and work with specific records within a file.

2.3.1.5 Common Management Information Service Element. The Common Management Information Service Element (CMISE) provides a common message framework for management procedures supplying both data-manipulation and notification/operation-oriented services.

2.3.1.6 System Management Application Service Element. The System Management Application Service Element (SMASE) specifies the management functions which are supported in a system node, and defines the semantics and abstract syntaxes of information transferred as relevant to OSI management.

2.3.1.7 Remote Operations Service Element. The Remote Operations Service Element (ROSE) provides a simple, uniform service for remotely invoking operations and then receiving correlated replies to those operations.

2.3.2 Presentation and session layer. In the OSI profile the presentation and session layers must adhere to OSI standards. In the lightweight profile these layers are null.

2.3.3 Transfer services interface. The interface to SAFENET transfer services is based upon the NATO Network Independent Interface (NIIF) whenever it is required to be defined.

3. NETWORK MANAGEMENT

3.1 Introduction. The purpose of this section is to provide an overview of SAFENET network management.

3.2 Network management overview. The SAFENET architecture includes a network management structure which supports the development of network performance measurement, fault tolerance, dynamic configuration control, and network security. This structure is defined in ISO/IEC 7498-4 Management Framework and ISO/IEC 10040 System Management Overview. While aspects of network management as defined by ISO standards are still evolving, enough is stable to provide a sound and useful basis for managing complex networks. The following subsections provide an overview and general requirements derived from the current NIST OIW Stable Implementation Agreements and from the underlying ISO standards (ISO 7498, ISO/IEC 7498-4, ISO/IEC 10040, ISO/IEC 10165-1).

3.2.1 Structure. Network management can be viewed as the set of operational and administrative mechanisms necessary to:

a. bring up, enroll, and/or alter network resources,
b. keep network resources operational,
c. fine tune these resources and/or plan for their expansion,
d. manage the accounting of their use, and
e. manage their protection from unauthorized use/tampering.

The structure provided in the current ISO standards and in the NIST OIW Implementation Agreements is designed to support all the above functions. The mechanisms included in the SAFENET profile will primarily support a subset of the capabilities in the first three areas.

A complete set of management capabilities in a network system will require ISO system management functions, ISO layer management functions, and local or system node management functions. This last is outside the scope of the ISO standards.

System management in the ISO standards is based on the following concepts:

f. managed objects,
g. layer management,
h. managing applications,
i. management agent applications,

j. capability in each instantiation of a communication service (or equipment) to carry out management actions or provide management information,
k. protocols for communication between managing applications and agent applications,
l. interfaces between the agent application and layer management entities in the same system node.

3.2.1.1 Managed objects. A managed resource is any resource such as a layer entity, a connection, or an item of physical communication equipment within the managed system that is subject to management. A managed object is the abstracted view of such a resource that represents its properties as seen by, and for the purposes of, management. The ISO standards include a subset (ISO/IEC 10165-1, 2, 4) which describe how objects are to be defined. It is the responsibility of the group defining a communication service or device to provide the definition of the managed object classes required by that entity. The definition of a managed object class includes its attributes (for example the parameters which can be read or set), control actions which can be taken on the real entity it represents, and notifications of important events which can be transmitted to a managing application. Thus the definition of the managed object classes relevant to an entity define the management operations which can be used with respect to that entity and the notifications which the entity will make available to a managing process.

The set of managed object class definitions representing the managed resources in a real open system is referred to as its Management Information Base (MIB). This is a conceptual grouping of all the information, actions, and notifications. It is not intended to imply any particular way of storing information. In particular it does not imply a central storage of all information.

The NIST OIW includes in its implementation agreements a non-normative appendix which provides the definitions of certain managed object classes. This listing is referred to as the Management Information Library (MIL). Object definitions to be used in SAFENET are provided in MIL-STD-2204.

3.2.1.2 Layer management. The ISO architecture recognizes that some management functions will need to be carried out primarily within specific layers of a profile. The specification of such functions is included in the standards for the layer in question. Any managed object definitions required should be in accordance with the rules for defining managed objects as defined in MIL-STD-2204. In SAFENET, LAN station management (SMT) is an

important example of this. This is defined in FDDI SMT. Among other things SMT controls the initialization and reconfiguration processes for the LAN. While some aspects of this can be controlled by system management much is carried out without direct reliance on system management.

3.2.1.3 Managing application. System network management will be carried out by some combination of people and software implementing system policy. The software implementing or assisting people in implementing management policy may be centralized, federated, or distributed. The ISO standards allow for all cases. The managing applications which make up this software, and the definition of policies to be implemented by such software, are outside the scope of ISO communication standards. The ISO standards define the capabilities available to carry out policy and provide the communication protocols needed. There may be a number of system nodes which contain managing application processes. Each one may, for example, be concerned with a different aspect of system management (e.g., security, fault recovery) or be concerned with managing a subset of the nodes in the system.

3.2.1.4 Management agent process. An agent application is an application making use of system management services, which for a particular exchange of systems management information performs management operations on managed objects and emits notifications of events on behalf of managed objects. This definition and the standards allow for the case in which an application acts as both manager and agent. The same application process may take on the managing role and agent role in different instances of management communication.

In the OSI profile each SAFENET LAN node should have a management agent process. This agent process should carry out all supported management interactions with the managed objects in the LAN node of which it is a part. The management interactions (services) supported are defined in MIL-STD-2204. Figure 2 shows the relationships between a process acting in the manager role and one acting in the agent role. Layer management or local management functions is not shown in this figure.

3.2.1.5 Layer management entities. Any instantiation of a communication service should provide the capability to carry out management actions and provide the management information specified in the definition of the managed object classes relevant to it. The instantiation of these capabilities is referred to as the layer management entity. The SMT of FDDI is an example of a layer management entity. The layer management entity does not have to be a separately defined and identified

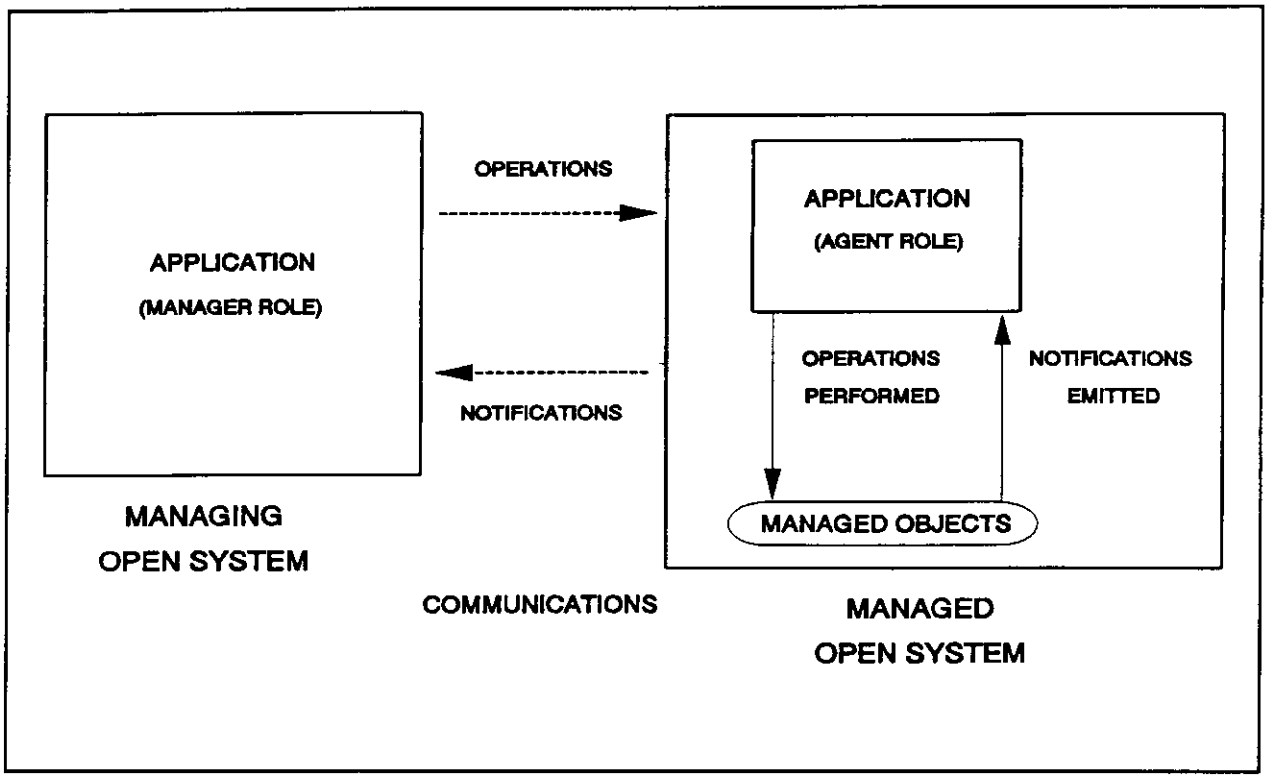

Figure 2. Systems Management Interactions

part of an implementation.

The managed object class definition provides the semantics and syntax of the information, and describes actions which the communication object may be required to carry out and the events for which it will generate notifications. The specification of the layer management entity interface to the local management agent application is a local matter. If both are software entities, then the interface may be the subject of other standards.

3.2.1.6 System management communication protocols. The ISO standards provide the communication protocols and service definitions for interchanges between applications acting in the managing and agent roles. That portion of the application process which implements these standards and is logically a part of the application layer is called the System Management Application Entity (SMAE). Figure 3 shows application layer services that are included in the SMAE. The SMAE includes the management services and protocols defined by the ISO standards which are used by an application process acting in the manager

role, agent role, or both. The ISO standards do not define the application interface between the SMAE or its components and the remaining functions of the management application process. The SMAE is logically comprised of Common Management Information Service Element (CMISE), System Management Application Service Element (SMASE), Remote Operation Service Element (ROSE), and Association Control Service Element (ACSE). These last two application service elements are also intended to be used by other application layer capabilities (for example FTAM and directory services). In addition to the services included in the SMAE a system management application process may make use of other services such as FTAM. An instantiation of the SMAE may also contain locally defined capabilities such as the mechanism for gaining access to local managed objects. The ISO standards do not specify whether or how managing processes in different system nodes will communicate with each other when multiple managers exist in a system. Within the scope of SAFENET, FTAM or the Private Communications application interface could be used for this purpose. Capabilities defined in other ISO application layer standards such as Transaction Processing (ISO/IEC 10026 parts 1-4) could also be used but are outside the scope of SAFENET. The components of the SMAE are defined in MIL-STD-2204.

3.2.1.7 Local interfaces. The means by which an application acting in either the managing role or the agent role carries out management operations or receives notifications within the system node in which it is resident is a local matter. The exception to this is when the Network Independent Interface (NIIF) is required, in which case management information should be passed between lower layer entities and the application by means of the NIIF management operations.

3.3 Limitations. A conforming implementation of SAFENET should provide mechanisms for carrying out network management. Some capabilities, such as providing selected fault detection and recovery functions, are inherent in the layer standards themselves and are available whenever the layer standard is implemented. For example, a conforming FDDI implementation automatically reconfigures to recover from certain breaks in the cable plant. As outlined in the following paragraph, the complete provision of network management in a system requires much additional capability beyond that supplied if only the functions and mechanisms specified in the SAFENET profile are implemented.

System designers and implementors using components based on SAFENET need to establish their requirements for the way in which network components are interconnected, and for capabilities such as system initialization, dynamic configuration control, fault

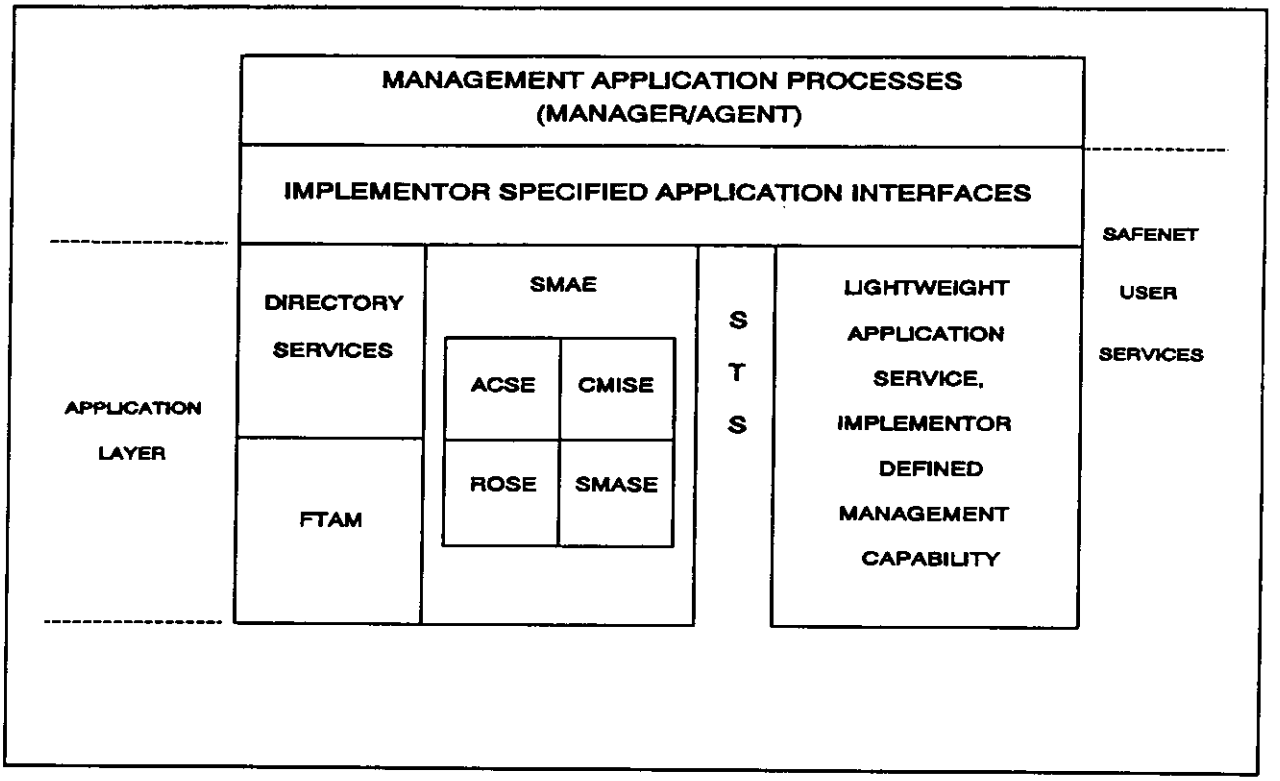

Figure 3. Management and the Application Layer

detection and recovery, survivability, safety, security, and time of day distribution which are supported by network management. The selection of the components to be connected to the network, the network topology, and the management policies to be implemented in managing processes will determine the extent to which system requirements can be achieved by a LAN based on the SAFENET profile. The selection process, management policies, and the functions implemented in managing processes are outside the scope of this handbook.

4. SAFENET LIGHTWEIGHT APPLICATION SERVICE DEFINITION

4.1 Introduction. The purpose of this section is to provide guidance on the implementation of the SAFENET Lightweight Application (SLA) service definition.

4.2 Service definitions. The SLA services provide service definitions in accordance with ISO/IEC TR 8509. Elements of a service definition are described below.

4.2.1 Service primitives. A service primitive is an abstract, implementation independent representation of an interaction between a service user and the service provider. A primitive has associated with it a direction, which may be from a service user to the service provider or from the service provider to a service user, and one or more parameters. The four types of primitives used in SAFENET service definitions are listed below:

a. Request: An interaction in which a service user requests some action of the service provider. The direction associated with this primitive type is from service user to service provider.

b. Indicate: An interaction in which the service provider indicates to a service user that it has performed some action. The action performed may have been initiated either by the service provider itself or by a service user at a remote service access point. The direction associated with this primitive type is from service provider to service user.

c. Response: An interaction in which a service user reports to the service provider that it has completed some action in response to an indicate primitive. The direction associated with this primitive type is from service user to service provider.

d. Confirm: An interaction in which the service provider reports to a service user the completion of some action previously initiated by a request primitive. The direction associated with this primitive type is from service provider to service user.

4.2.2 Service parameters. Each primitive of a particular service has associated with it a set of service parameters, and each of these parameters has a particular applicability. Parameter definitions are provided along with the specific

service definitions (see 4.3.1.1). The parameter applicabilities used in SAFENET are described in the following paragraph.

4.2.3 Parameter applicabilities. Associated with each parameter used in the service primitives is an applicability which describes that parameter's status with respect to the primitive in which it is named. The applicabilities relevant to the primitives of SAFENET service definitions are listed below:

a. M (indicates the parameter is mandatory)
b. U (indicates the parameter is user optional)
c. = (indicates that the value supplied in an indication or confirm primitive is always identical to that supplied in the previous request or response primitive issued at the peer service access point)
d. NA (indicates the parameter is not applicable)

4.2.4 Time-sequence diagrams. Time-sequence diagrams are used in conjunction with service definitions to illustrate the interactions of service primitives over time. The time-sequence diagrams which apply to SAFENET services are shown in Figure 4.

4.3 Lightweight application services. The SLA service definition describes services which can be provided by the lightweight profile to a SAFENET user. These services are divided into two categories: data transfer services and directory services. The lightweight data transfer services are services which are essentially exported from the transport layer to the application layer (see 4.3.2). The lightweight directory services include registration and deletion of both individual and group logical names (see 4.3.3). In any application in which a process should be coded independent of its own network address or of the network addresses of processes with which it must communicate, a directory service capability is required. Most, if not all, fault tolerant designs have this as a requirement. Thus it is expected that implementations of the SAFENET lightweight suite will implement both a data transfer service and a matching directory service.

4.3.1 Lightweight service parameters. The parameters associated with the lightweight application service primitives and the selections associated with the QOS parameter are defined below. See 4.2 for additional information on primitives, parameter applicabilities, and time-sequence diagrams.

4.3.1.1 Lightweight service parameter definitions. The definitions of the parameters used in the lightweight application service definition are given in the following paragraphs.

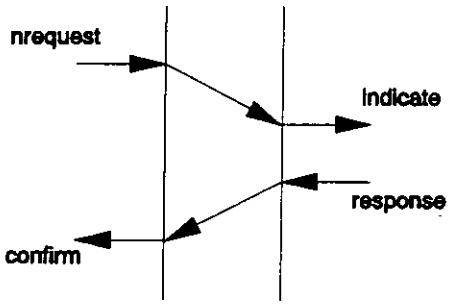

Figure 3(a)

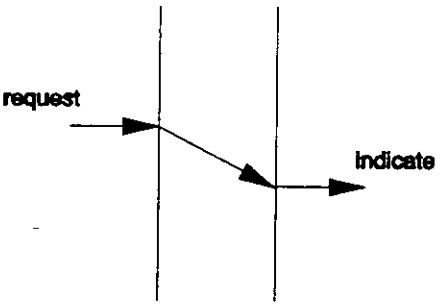

Figure 3(b)

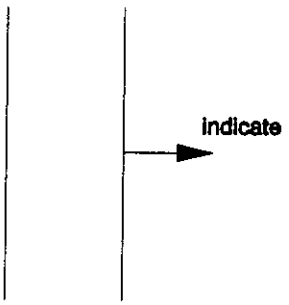

Figure 3(c)

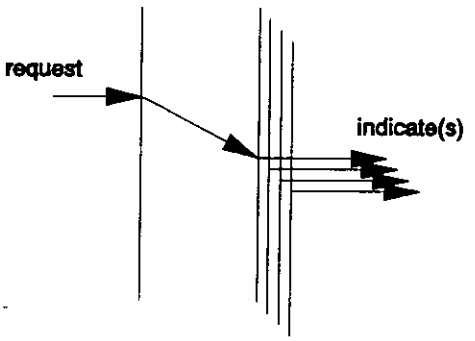

Figure 3(d)

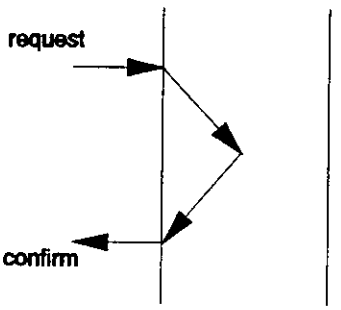

Figure 3(e)

Figure 4. Time Sequence Diagrams for SAFENET Services

4.3.1.1.1 Connect-ID. This value is used to identify uniquely the local end of a connection.

4.3.1.1.2 Disconnect-Mode. This specifies to the service provider the mode of disconnect desired. The two possible modes are graceful and immediate. The graceful mode specifies that the user expects any data pending on the connection to be delivered before the connection is terminated, and the immediate mode specifies that the connection should be terminated without attempting to deliver any additional data.

4.3.1.1.3 Dst-Logical-Name. This logical name is used to identify the destination entity in a peer to peer service request.

4.3.1.1.4 Group-Name. This logical name is used to specify a multicast group.

4.3.1.1.5 Logical-Name. This name uniquely identifies an application entity and frees the application user from the burden of having to specify complete network addresses.

4.3.1.1.6 Membership-List. This list of application entity logical names is used as a mechanism for restricting membership in multicast groups.

4.3.1.1.7 Message. This is a single unit of user data. No restrictions on the size or content of a message are defined here nor are there intended to be any implications on transmission characteristics such as reliability.

4.3.1.1.8 QOS. This set of selections is used to specify to the service provider the quality of service expected by the user of a given service primitive. The QOS selections may be either absolute requirements or indications to the service provider of the relative importance of the request.

4.3.1.1.9 Reason. This parameter is used to identify the nature of a failed request or aborted connection.

4.3.1.1.10 Request-ID. This parameter is used to identify the local end of a unitdata transfer request. It is used to identify uniquely the recipient of an asynchronous notification.

4.3.1.1.11 Request-Message. This is a message associated with a transaction request primitive.

4.3.1.1.12 Response-Message. This is a message associated with a transaction response primitive.

4.3.1.1.13 Result. This parameter is used to report the success or failure of a service request.

4.3.1.1.14 Src-Logical-Name. This parameter is used to specify the logical name of the service user invoking a primitive.

4.3.1.1.15 Status. This parameter is used to indicate delivery status information to lightweight service users.

4.3.1.1.16 Transaction-ID. This parameter is used to identify uniquely the local end of a transaction. This parameter is used by the service provider to associate a request and response(s).

4.3.1.1.17 T-Service. This parameter is used in the unitdata primitives to specify the use of the ISO connectionless transport protocol instead of the default XTP.

4.3.1.1.18 XTP-Control. The following parameters specify options or mechanisms offered by XTP that could be useful to an application. In all cases there must be default modes or values, should an application not specify them, so that the protocol will operate properly.

4.3.1.1.18.1 Checksum. This parameter enables/disables the data checksum within an XTP or CLTP packet. Because of the processor time required to calculate the data checksum, an application may choose to disable the checksum function. Note that this parameter is applicable to both XTP and CLTP.

4.3.1.1.18.2 Flow-Control. This parameter determines whether flow control is used by XTP to limit the amount of data that can be transmitted before it is acknowledged.

4.3.1.1.18.3 Aggressive-Acknowledgement. This parameter allows the receiver of data to respond with an acknowledgement of a missing packet as soon as it receives an out of order packet, rather than waiting for an acknowledgement request from the sender.

4.3.1.1.18.4 Reservation. This parameters mandates the use of flow control based on available application buffer space instead of XTP buffer space.

4.3.1.1.18.5 Special-Data. This parameter specifies an eight octet value, supplied by the application, to be transferred to the receiver(s) in the btag field of an XTP packet, and delivered to the receiving application(s).

4.3.1.1.18.6 Acknowledgement-Policy. This parameter specifies the means by which data transmission will be acknowledged. This parameter is divided into two parts:

a) the granularity of acknowledgements - octets, packets, messages, or once, and
b) the number of granularity units to be transmitted before an acknowledgement is requested.

4.3.1.1.18.7 Rate-Control. This parameter allows the receiver to specify a transmission data rate to the sender. This can be used to prevent overflow of a slow receiver because of a faster transmitter.

4.3.1.2 Quality of service selections. The quality of service (QOS) parameter is defined to be a set of values from which the lightweight application service user should select to specify to the service provider certain performance and reliability characteristics. Each selection may be further qualified to indicate to the service provider whether the selection imposes an absolute requirement on transmission characteristics or is simply guidance to the service provider. The set of available selections depends upon the service being requested. The QOS selections available to the user and the services to which they apply are shown in Table II. These QOS selections are defined below.

Table II. Quality of Service Selections

	QOS Selections			
Service	SAC	TD	MXL	PR
SLA-CONNECT	X	-	X	X
SLA-DISCONNECT	-	-	-	X
SLA-STREAM-DATA	-	-	-	X
SLA-UNITDATA	X	X	-	X
SLA-MCAST-UNITDATA	X	X	-	X
SLA-MCAST-STREAM	X	X	-	X
SLA-TRANSACTION	X	-	X	X
SLA-NOTIFY	-	-	-	X

4.3.1.2.1 Selectable acknowledgment control (SAC). This selection provides a mechanism for a user to specify the desired level of acknowledgment control. The two primary settings are acknowledged transfer and unacknowledged transfer. Acknowledged transfer here indicates acknowledgement by the peer service

provider. An implementor may provide for intermediate levels of acknowledgement in the case of multicast transfers. That is, the multicast may be considered a success if some subset of the total potential set of transfers is successful.

4.3.1.2.2 Transit delay (TD). This selection specifies a maximum time for delivery of a single message to a remote user.

4.3.1.2.3 Maximum latency (MXL). This selection specifies a maximum time between a request and its corresponding confirm.

4.3.1.2.4 Priority (PR). This selection specifies the priority of the primitive relative to other transfer requests.

4.3.2 Lightweight data transfer services. The lightweight data transfer services are those services which apply directly to the transfer of messages between two or more lightweight application service users. These services are all supported by XTP. Only the SLA-UNITDATA and SLA-MCAST-UNITDATA services are supported by the ISO connectionless transport protocol.

4.3.2.1 SLA-CONNECT service. The SLA-CONNECT service allows a service user to establish a connection between itself and another service user for the purpose of exchanging data. Upon successful completion of the connection establishment, a connection ID should be supplied to the user by the service provider to identify the local end of the connection. As a means of identifying the connection in the event that a disconnect needs to be initiated by the service provider before the connection is established, the service user should propose a connection ID. Multiple connections may exist simultaneously between the same peer service users. Once a connection is made, the peer service users may exchange messages over the connection with the assurance that all messages should be delivered correctly and in the order in which they were submitted. The service user has the option of sending data along with the connection request. If this option is used, the service provider should not wait for a successful connection to be established but instead should attempt to deliver the user's data along with the connect indicate. The QOS selections which apply to this primitive are selectable acknowledgment control, maximum latency, and priority.

4.3.2.1.1 SLA-CONNECT primitives. The service primitives used to describe the SLA-CONNECT service should be request, indicate, response, and confirm. Their interactions are shown in the time-sequence diagram of Figure 4(a) (see 4.2.4).

4.3.2.1.2 SLA-CONNECT parameters. The SLA-CONNECT service parameters and their applicabilities should be as listed below:

Parameter Name	Request	Indicate	Response	Confirm
Src-Logical-Name	M	=	M	=
Dst-Logical-Name	M	=	M	=
Connect-ID	M	M	=	M
QOS	U	=	=	=
Message	U	=	NA	NA
Checksum	U	NA	NA	NA
Flow-Control	U	NA	NA	NA
Aggressive Acknowledgement	U	NA	NA	NA
Reservation	U	=	NA	NA
Special-Data	U	=	NA	NA
Acknowledgement -Policy	U	NA	NA	NA
Rate-Control	U	NA	NA	NA

4.3.2.2 SLA-DISCONNECT service. The SLA-DISCONNECT service allows either an application service user to terminate an existing connection between itself and another service user (called a normal close), or the service provider to terminate an existing connection between peer service users (called an abnormal close). The reason parameter should indicate at a minimum whether the connection is being closed normally or abnormally. A normal close should occur when the service user determines that the connection is no longer needed, while an abnormal close should be initiated by the service provider in response to an error condition or the inability to maintain the required level of service. The only QOS selection which applies to this service is priority.

4.3.2.2.1 SLA-DISCONNECT primitives. The service primitives used to describe the SLA-DISCONNECT service should be request and indicate. Their interactions are shown in the time-sequence diagrams of Figure 4(b) for a normal close, and Figure 4(c) for an abnormal close (see 4.2.4).

4.3.2.2.2 SLA-DISCONNECT parameters. The SLA-DISCONNECT service parameters and their applicabilities should be as listed below:

Parameter Name	Request	Indicate
Connect-ID	M	M
Disconnect-Mode	U	=
Reason	NA	M

4.3.2.3 SLA-STREAM-DATA service. The SLA-STREAM-DATA service allows peer service users to exchange data over previously established connections. The service provider guarantees that the user's data is delivered to the receiving entity in the same order in which it was submitted and that the data is delivered reliably within the quality of service selections set during connection establishment. A disconnect initiated by the service provider should inform the user when the reliable delivery service cannot be provided or when any required QOS constraints can no longer be met. Priority is the only QOS selection which may be set on individual stream-data requests. Specifying this selection should not affect the ordering of data transferred over the connection; however, the current priority of the connection may be varied over time using this option.

4.3.2.3.1 SLA-STREAM-DATA primitives. The service primitives used to describe the SLA-STREAM-DATA service should be request and indicate. Their interactions are shown in the time-sequence diagram of Figure 4(b) (see 4.2.4).

4.3.2.3.2 SLA-STREAM-DATA parameters. The SLA-STREAM-DATA service parameters and their applicabilities should be as listed below:

Parameter Name	Request	Indicate
Connect-ID	M	M
QOS	U	=
Message	M	=
Special-Data	U	=
Rate-Control	U	NA

4.3.2.4 SLA-UNITDATA service. The SLA-UNIT-DATA service allows peer service users to exchange messages without the overhead of maintaining a connection. This service should not provide for the ordered delivery of messages transferred between the same two service users. The QOS selections which apply to the unitdata service are selectable acknowledgment control, transit delay, and priority. Failure to meet these requirements should be reported to the user via the notify service.

While XTP should be the default provider of this service, the user may alternately select the ISO connectionless transport protocol as the service provider. If ISO connectionless transport is selected the user should ensure that the message being transferred will fit within the constraints of the data fields of the underlying data transfer protocols.

4.3.2.4.1 SLA-UNITDATA primitives. The service primitives used to describe the SLA-UNITDATA service should be request and

indicate. Their interactions are shown in the time-sequence diagram of Figure 4(b) (see 4.2.4).

4.3.2.4.2 SLA-UNITDATA parameters. The SLA-UNITDATA service parameters and their applicabilities should be as listed below:

Parameter Name	Request	Indicate
Src-Logical-Name	M	=
Dst-Logical-Name	M	=
QOS	U	=
Request-ID	M	M
T-Service	U	=
Message	M	=
Checksum	U	NA
Special-Data	U	=

4.3.2.5 SLA-MCAST-UNITDATA service. The SLA-MCAST-UNITDATA service allows a service user to exchange messages with members of a multicast group. The service provider should attempt to deliver the message to each active member of the multicast group. The QOS parameters which apply to this service are selectable acknowledgment control, transit delay, and priority. Failure to meet these requirements should be reported to the user via the notify service.

While XTP should be the default provider of this service, the user may alternately select the ISO connectionless transport protocol as the service provider. If ISO connectionless transport is selected the user should ensure that the message being transferred will fit within the constraints of the data fields of the underlying data transfer protocols.

4.3.2.5.1 SLA-MCAST-UNITDATA primitives. The service primitives used to describe the SLA-MCAST-UNITDATA service should be a request followed by zero or more indicates. Their interactions are shown in the time-sequence diagram of Figure 4(d) (see 4.2.4).

4.3.2.5.2 SLA-MCAST-UNITDATA parameters. The SLA-MCAST-UNITDATA service parameters and their applicabilities should be as listed below:

Parameter Name	Request	Indicate
Src-Logical-Name	M	=
Group-Name	M	=
QOS	U	=
Request-ID	M	M
T-Service	U	=
Message	M	=
Checksum	U	NA
Special-Data	U	=

4.3.2.6 SLA-MCAST-STREAM service. The SLA-MCAST-STREAM service allows a user to establish a connection with group members and send data over the connection. The service provider guarantees that the user's data is delivered to the group in the same order in which it was submitted. The reliability is based on the XTP multicast reliability paradigm and, with selectable acknowledgement control set to acknowledged, will guarantee delivery to at least one of the group members with a high degree of confidence of delivery to other active group members. The QOS selections which apply to this service are selectable acknowledgement control, maximum latency, and priority.

4.3.2.6.1 SLA-MCAST-STREAM primitives. The service primitives used to describe the SLA-MCAST-STREAM shall be a request followed by one or more indicates. Their interactions are shown in the time-sequence diagram of Figure 4(d) (see 4.2.4).

4.3.2.6.2 SLA-MCAST-STREAM parameters. The SLA-MCAST-STREAM service parameters and their applicabilities shall be as listed below:

Parameter Name	Request	Indicate
Src-Logical-Name	M	=
Group-Name	M	=
QOS	U	=
Connect-ID	M	M
Message	M	=
Checksum	U	NA
Flow-Control	U	NA
Aggressive Acknowledgement	U	NA
Reservation	U	=
Special-Data	U	=
Acknowledgement-Policy	U	NA
Rate-Control	U	NA

4.3.2.7 SLA-TRANSACTION service. The SLA-TRANSACTION service allows a service user to initiate a request/response sequence with a peer service user. The response may optionally

contain a user message. The service provider should be responsible for associating the request and response and ensuring the completion of the service within the QOS selections specified. The orderly delivery of requests with identical quality of service from a source service user to a destination service user should be guaranteed. The parameter transaction-ID is similar to a connection-ID in that it is a local identifier and is used to match the transaction response to the transaction request. The notify service may be used to inform the user of an inability to complete successfully the transaction service within any required QOS constraints. The QOS selections which apply to the transaction service are selectable acknowledgment control, maximum latency, and priority.

4.3.2.7.1 SLA-TRANSACTION primitives. The service primitives used to describe the SLA-TRANSACTION service should be request, indicate, response, and confirm. Their interactions are shown in the time-sequence diagram of Figure 4(a) (see 4.2.4).

4.3.2.7.2 SLA-TRANSACTION parameters. The SLA-TRANSACTION service parameters and their applicabilities should be as listed below:

Parameter Name	Request	Indicate	Response	Confirm
Src-Logical-Name	M	=	M	=
Dst-Logical-Name	M	=	M	=
QOS	U	=	=	=
Request-Message	M	=	NA	NA
Transaction-ID	M	M	=	M
Request-ID	M	M	=	=
Response-Message	NA	NA	U	=
Checksum	U	NA	NA	NA
Special-Data	U	=	NA	NA

4.3.2.8 SLA-NOTIFY service. The SLA-NOTIFY service is used to provide service users with information concerning datagram and transaction reliability. This service should minimally advise the user of failures by the transaction or unitdata services. In addition, in the case of multicast services, the notify primitive should be used to provide the user with information on partial delivery status. The only QOS selection applicable to this service is priority.

4.3.2.8.1 SLA-NOTIFY primitives. The service primitive used to describe the SLA-NOTIFY service should be indicate. Its time-sequence diagram is shown in Figure 4(c) (see 4.2.4).

4.3.2.8.2 SLA-NOTIFY parameters. The SLA-NOTIFY service parameters and their applicabilities should be as listed below:

Parameter Name	Indicate
Request-ID	M
Status	M
Reason	M

4.3.3 Lightweight directory services. The lightweight directory services describe functions which are not provided by the transport layer protocols of the lightweight profile, but which are necessary for use of the lightweight data transfer services. The directory services provide two types of functionality. The first is logical name to address mapping which frees the user from the burden of having to know and specify fully qualified addresses for peer to peer process communication. The second function is to provide the capability for minimal group management. If a system uses only the combined profile, OSI directory services may be used to provide this functionality. In all other cases lightweight directory services and OSI directory services should be independent.

The use of multicast transport services will require the implementation of all the directory services listed below. If only point to point services are used then only the SLA-REGISTER and SLA-CANCEL services are required.

This section lists the services to be supplied, but does not provide the protocols which will be needed in many cases to implement these services when only the lightweight suite is implemented. An explicit protocol and explicit implementation of the services are not required if static tables are used to provide mapping. In general, static tables supply very limited fault tolerance and will usually not be satisfactory in mission critical applications where fault tolerance is important.

4.3.3.1 SLA-REGISTER service. The SLA-REGISTER service allows an application process to identify itself to the network. Associated with each logical name in the local entity's database should be the names of any multicast groups to which it belongs. An application process may register more than one logical name, but the logical names for each registration should be unique. Once an application process has identified itself to the network, it may begin exchanging messages with other application processes via its logical name.

4.3.3.1.1 SLA-REGISTER primitives. The service primitives used to describe the SLA-REGISTER service should be request and confirm. Their interactions are shown in the time-sequence diagram of Figure 4(e) (see 4.2.4).

4.3.3.1.2 SLA-REGISTER parameters. The SLA-REGISTER service parameters and their applicabilities should be as listed below:

Parameter Name	Request	Confirm
Logical-Name	M	=
Result	NA	M

4.3.3.2 SLA-CANCEL service. The SLA-CANCEL service allows an application process to remove its logical names from the network. Once an application process has removed its name or names completely from the network, it should not be allowed to exchange data with other application processes. All open connections involving the withdrawing application process should be closed as a result of the completion of this service.

4.3.3.2.1 SLA-CANCEL primitives. The primitives used to describe the SLA-CANCEL service should be request and confirm. Their interactions are shown in the time-sequence diagram of Figure 4(e) (see 4.2.4).

4.3.3.2.2 SLA-CANCEL parameters. The SLA-CANCEL service parameters and their applicabilities should be as listed below:

Parameter Name	Request	Confirm
Logical-Name	M	=
Result	NA	M

4.3.3.3 SLA-OPEN-GROUP service. The SLA-OPEN-GROUP service allows an application process to establish a new multicast group. This service should identify the new group name to the network and establish the requesting application process as the first member of the newly opened multicast group. An application process may send messages to any group, but it should only receive messages addressed to one of the groups to which it belongs. The application process which opens the multicast group should have the ability to restrict membership in that group by including a group membership list which should be consulted when a join request is issued. If a non-null membership list is associated with the group, an application should not be permitted to join unless its logical name appears in the membership list.

4.3.3.3.1 SLA-OPEN-GROUP primitives. The service primitives used to describe the SLA-OPEN-GROUP service should be request and confirm. Their interactions are shown in the time-sequence diagram of Figure 4(e) (see 4.2.4).

4.3.3.3.2 SLA-OPEN-GROUP parameters. The SLA-OPEN-GROUP service parameters and their applicabilities should be as listed below:

Parameter Name	Request	Confirm
Logical-Name	M	=
Membership-List	U	NA
Group-Name	M	=
Result	NA	M

4.3.3.4 SLA-CLOSE-GROUP service. The SLA-CLOSE-GROUP service provides the user with the capability to delete a group from the network.

4.3.3.4.1 SLA-CLOSE-GROUP primitives. The service primitives used to describe the SLA-CLOSE-GROUP service should be request and confirm. Their interactions are shown in the time-sequence diagram of Figure 4(e) (see 4.2.4).

4.3.3.4.2 SLA-CLOSE-GROUP parameters. The SLA-CLOSE-GROUP service parameters and their applicabilities should be as listed below:

Parameter Name	Request	Confirm
Logical-Name	M	=
Group-Name	M	=
Result	NA	M

4.3.3.5 SLA-JOIN-GROUP service. The SLA-JOIN-GROUP service allows an application process to add itself to a multicast group. This service should establish the requesting application process as an active member of the desired multicast group. An application process may send messages to any group, but it should only receive messages addressed to one of the groups to which it belongs. If the group is a restricted group, the logical name of the application process issuing the join request should have to appear in the membership list in order to gain admittance to the group.

4.3.3.5.1 SLA-JOIN-GROUP primitives. The service primitives used to describe the SLA-JOIN-GROUP service should be request and confirm. Their interactions are shown in the time-sequence diagram of Figure 4(e) (see 4.2.4).

4.3.3.5.2 SLA-JOIN-GROUP parameters. The SLA-JOIN-GROUP service parameters and their applicabilities should be as listed below:

Parameter Name	Request	Confirm
Logical-Name	M	=
Group-Name	M	=
Result	NA	M

4.3.3.6 SLA-LEAVE-GROUP service. The SLA-LEAVE-GROUP service allows an application process to leave any multicast group to which it belongs. This service should remove the requesting application process from the active membership list of the specified multicast group.

4.3.3.6.1 SLA-LEAVE-GROUP primitives. The service primitives used to describe the SLA-LEAVE-GROUP service should be request and confirm. Their interactions are shown in the time-sequence diagram of Figure 4(e) (see 4.2.4).

4.3.3.6.2 SLA-LEAVE-GROUP parameters. The SLA-LEAVE-GROUP service parameters and their applicabilities should be as listed below:

Parameter Name	Request	Confirm
Logical-Name	M	=
Group-Name	M	=
Result	NA	M

5. NATO NETWORK INDEPENDENT INTERFACE

5.1 Introduction. The NATO Network Independent Interface (NIIF) defines the complete interface between a user system implementing OSI layers 5-7 and a Data Transfer System (DTS) implementing OSI layers 1-4, as shown in Figure 5. Note that a user system (as defined here) corresponds to SAFENET user services, and a DTS corresponds to SAFENET transfer services plus LAN services. In this section, the terms user system and data transfer system will be used to conform with the NATO documents.

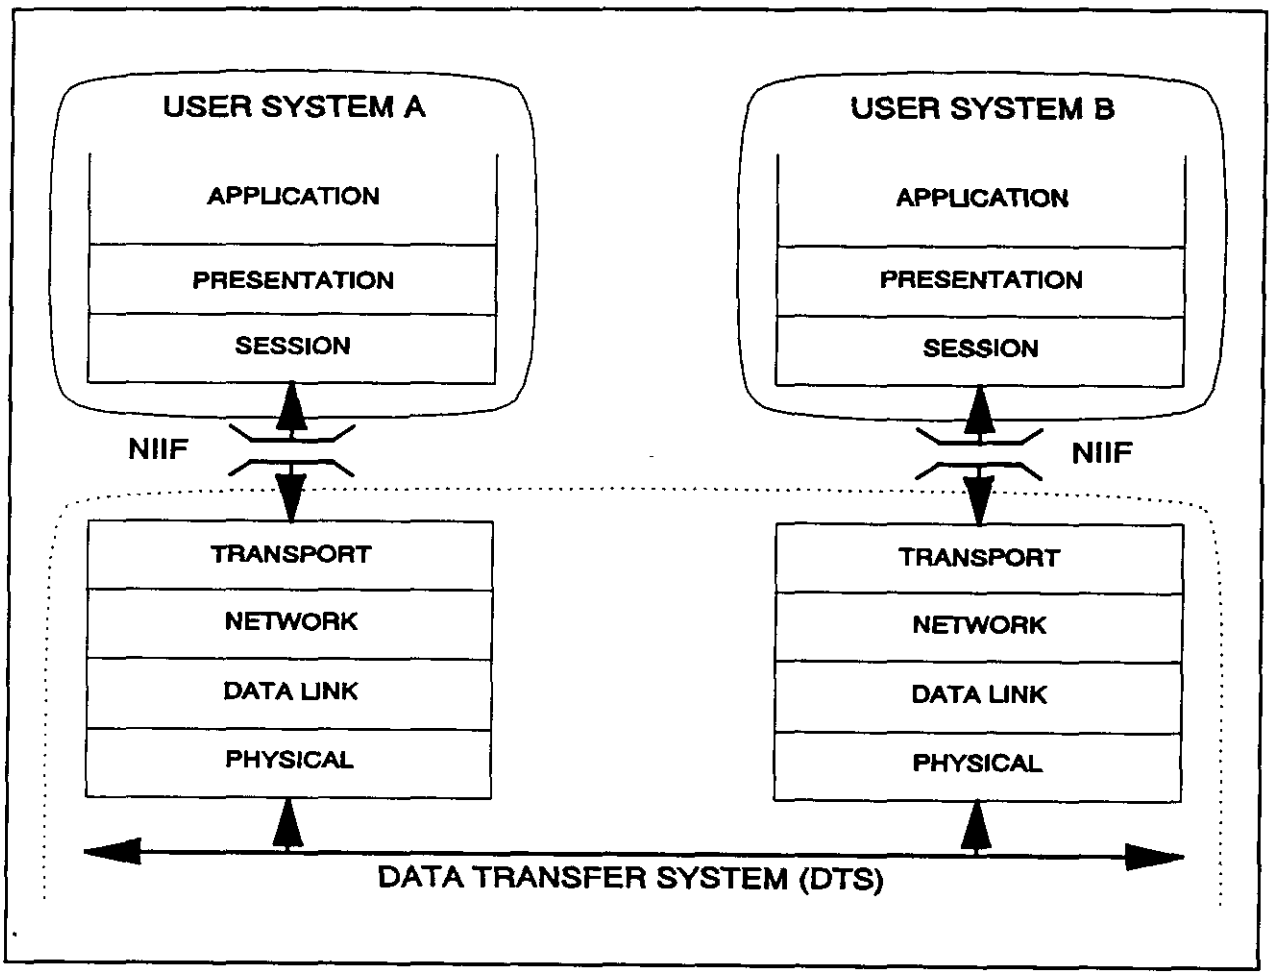

Figure 5. Location of NIIF in the OSI Reference Model

Access to the data transfer functions of the DTS through the NIIF allows user systems to transfer information to other user systems. The NIIF promotes interoperability of separately

developed user system and DTS equipment. Use of the NIIF is required in SAFENET whenever these portions of the SAFENET architecture are implemented independently.

The concept of the NIIF is compatible with the layered structure of the OSI reference model and its associated peer to peer protocol standards. The OSI protocol standards apply only to communication between cooperating peer entities. They do not identify, define, or otherwise refer to explicit interfaces between adjacent layers. These interfaces are considered to be a local matter to be solved in the specific implementations. Thus, the NIIF constitutes an explicit definition of a vertical interface which can be applied to specific implementations.

5.2 NIIF services. The link between the NIIF and the OSI reference model is through the associated NATO Naval Transport Service Definition. This document is part of the family of documents that define the NIIF. It identifies the services provided by the transport layer to the session layer. A service is actually the method by which the user system gains access to a particular protocol entity within the DTS. This transport service definition allows the user system to access the full range of services available from the data transfer system. The NIIF services are: connectionless mode data transfer service, connection mode data transfer service, real-time data transfer service, and management service.

5.2.1 Connectionless mode data transfer service. The connectionless mode data transfer service provides the user system with the ability to transmit a single unit of data without the requirement of establishing a connection. For the NIIF, this service is defined in the Proposed Draft Standard for NATO Naval Intraship Transport Service Definition - Addendum 1 to Annex A: Connectionless Mode Data Transmission. This addendum defines data transfer services based on the Connectionless Mode Service Addendum to ISO 8072.

In this mode messages will be individually handled, and sequential delivery is not assured. The user system is not assured of notification of the loss or duplication of a message. All information necessary for delivery of the data unit is present at the time of invocation of the service. No connection establishment is required in order to exchange data, and no relationship exists between successive uses of the service.

5.2.1.1 T-UNIT-DATA primitive. The T-UNIT-DATA primitive is the only primitive defined for the connectionless mode service. The T-UNIT-DATA primitive is used to transfer a single Transport Service Data Unit (TSDU) from the source Transport

Service Access Point (TSAP) to the destination TSAP in a single invocation of the service. Use of the T-UNIT-DATA primitive is diagrammed in Figure 6.

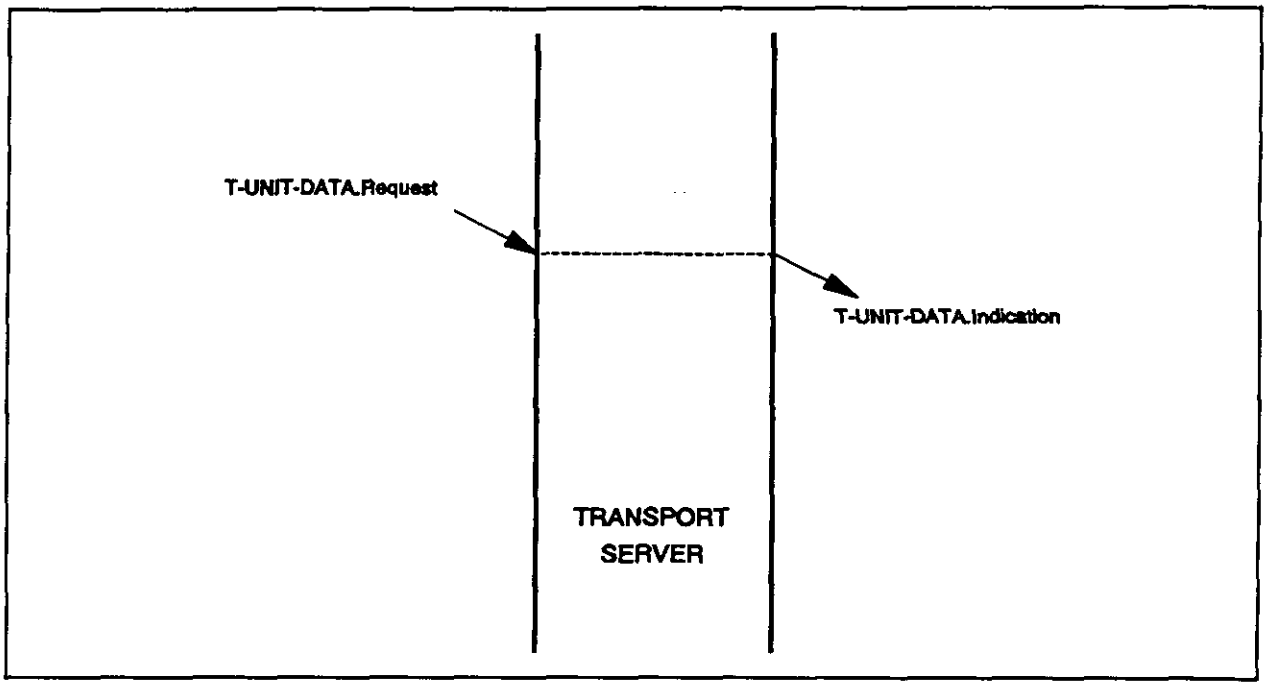

Figure 6. Connectionless Mode Data Transfer Primitive

5.2.2 Connection mode data transfer service. The connection mode data transfer service provides the user system with reliable, sequential data transfer within the limits of the quality of service (QOS) parameters established for the connection. For the NIIF, this service is defined in the Proposed Draft Standard for NATO Naval Intraship Transport Service Definition - Annex A: Connection Mode Data Transmission. This annex defines data transfer services based on the ISO 8072 Connection Mode Transport Service.

The NIIF provides for fixed and selectable quality of service parameters, but no negotiation. The service depends on the implementation to be aware of its specific limitations with respect to the available services. In other words, the available QOS should be known for the specified system or a range of values from which to choose should be available for the user system at the invocation of the service.

The connection mode service allows the user system to establish connections with other user systems within a virtual

address environment. Virtual addresses are independent of the underlying physical address structure and are identified to the data transfer system during the connection establishment phase.

Primitives for the connection mode service fall into three classes: connection establishment, data transfer, and connection release. Use of the connection mode primitives is diagrammed in Figure 7.

5.2.2.1 T-CONNECT primitive. The T-CONNECT primitive provides a mechanism for establishing connections between two (or more) user systems. This service definition allows for selection of a specific set of QOS parameters. The other parameters are fixed by the system design.

5.2.2.2 Data transfer primitives. Once a connection has been established the T-DATA and T-EXP-DATA primitives may be used for the exchange of TSDUs.

5.2.2.3 T-DISCONNECT primitive. At the completion of the data transfer process, or in response to faults, the T-DISCONNECT primitive may be used to release a connection.

5.2.3 Management service. The management service provides the user system the means with which to manage the resources of the DTS. For the NIIF, this service is defined in the Proposed Draft Standard for NATO Naval Intraship Transport Service Definition - Annex B: Management Service Definition. These management services are defined in the context of current OSI and IEEE network management standards.

The management service provides access to the Layer Management Entities (LMEs) of each local protocol entity within the data transfer system, as well as those of remote protocol entities. Since this service provides for layer management within the local protocol entities, it is not strictly a peer-to-peer communication. The management service primitives consist of the following groups: parameter control, action control, and event reporting. Use of management primitives is diagrammed in Figure 8.

5.2.3.1 Parameter control primitives. Parameters that control specific characteristics and attributes of the layer service providers within the DTS can be set, conditionally set, and examined using this group of primitives. The specific parameter encodings are defined in the resource specification. All parameter control Invoke primitives are processed by the appropriate LME, and a Reply primitive is returned.

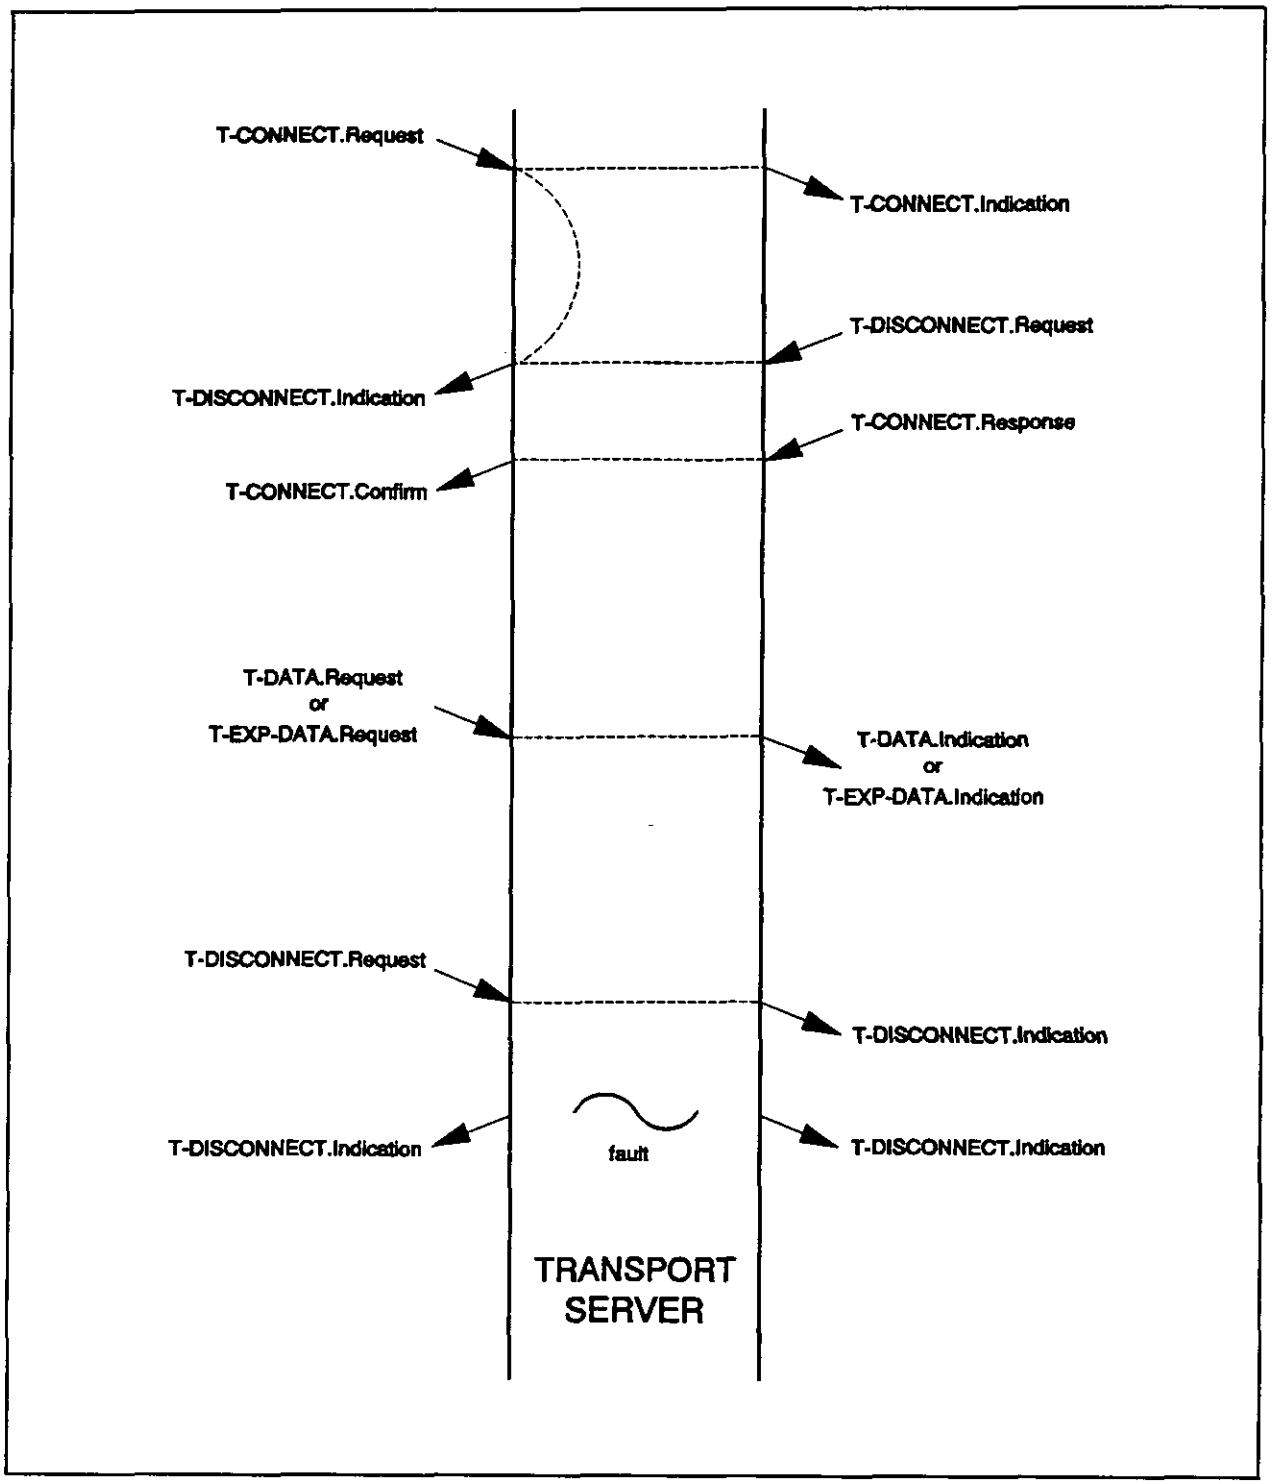

Figure 7. Connection Mode Data Transfer Primitives

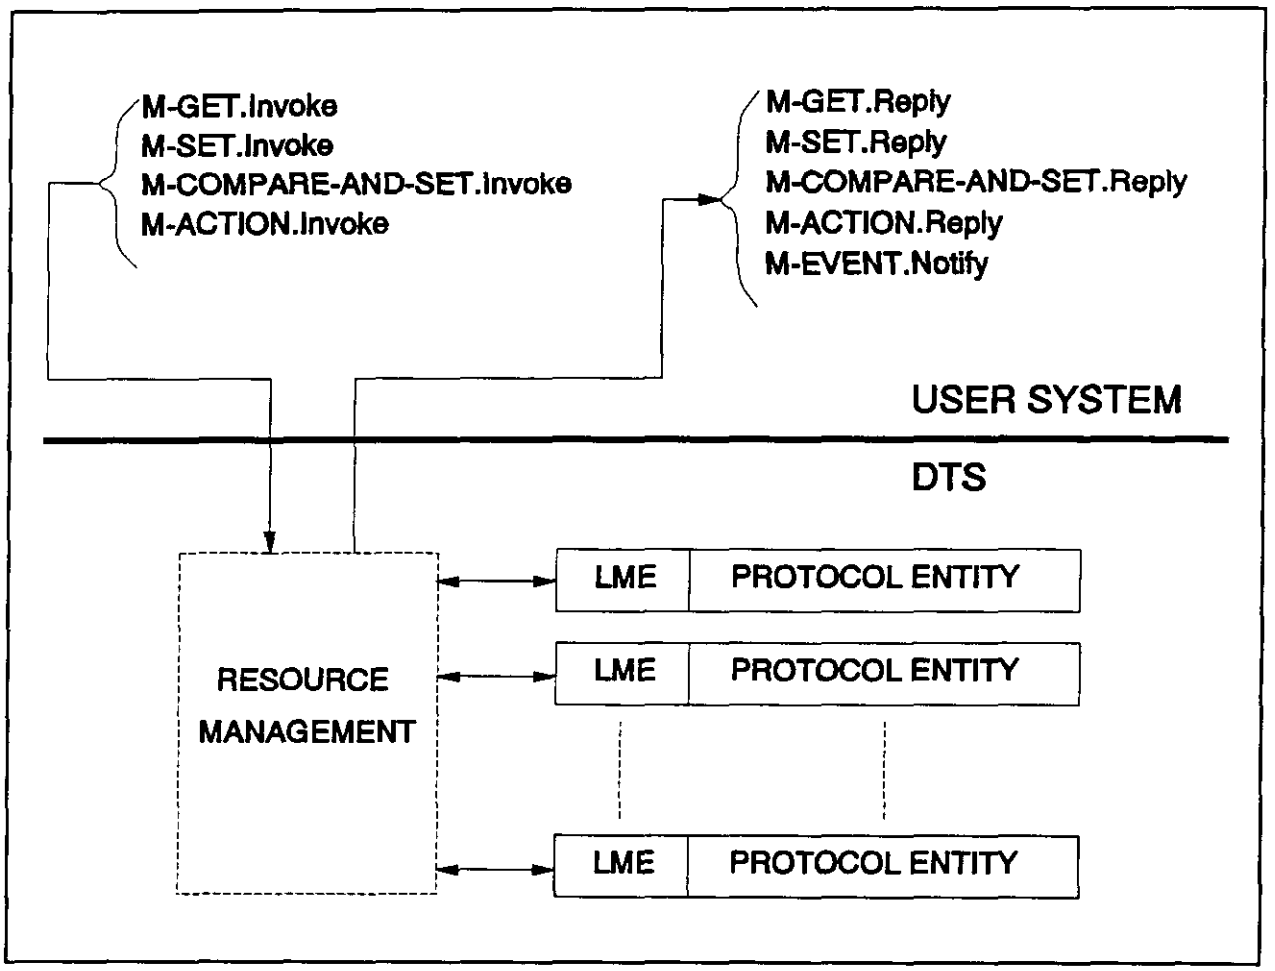

Figure 8. NIIF Management Primitives

5.2.3.1.1 M-GET primitive. This primitive allows the user system to obtain the value of a specific resource parameter.

5.2.3.1.2 M-SET primitive. This primitive allows the user system to perform a set operation on a specific layer parameter.

5.2.3.1.3 M-COMPARE-AND-SET primitive. This primitive allows the user system to compare the value of a parameter to a specified value and, if they are equal, to perform a set operation on a specific resource parameter.

5.2.3.2 Action control primitive. The M-ACTION primitive allows the user system to invoke a specified state transition or sequence of events within the specified resource.

5.2.3.3 Event reporting primitive. The M-EVENT primitive allows the DTS to report significant events occurring within a specified resource to the user system.

5.3 NIIF specifications. The NIIF is intended to support both external and embedded network interfaces. It also capitalizes on readily available physical interface standards (e.g., VME bus, STANAG 4153). The NIIF specifications are divided into the following two groups: interface data unit specifications, and local profile specifications.

5.3.1 Interface data unit specifications. The NIIF data transfer and management primitives are passed between the user system and the DTS in the form of Interface Data Units (IDUs). These IDUs explicitly define the format and content of each primitive. Within an IDU, parameters have specified field positions and field sizes. Usually, the range and meaning of parameter values will also be specified. An IDU may also contain a Protocol Data Unit (PDU). In a specific network implementation which employs the NIIF, the implementer must define values and ranges for the QOS parameters and management parameters.

5.3.2 Local profile specifications. The transfer of IDUs across the NIIF between the user system and the DTS requires each side to provide functions equivalent to the lower two layers of the OSI reference model. The physical transfer mechanism is selected for a specific implementation through a local profile.

Figure 9 illustrates the NIIF local profile model. The exchange of IDUs between a user system and a DTS may be based upon any one of a number of profiles, with each profile supporting a different physical interface. Each side (user system or DTS) of an NIIF local profile is partitioned into three segments: the logical segment, convergence segment, and physical segment.

5.3.2.1 Logical segment. The logical segment provides standard procedures for the control of the exchange of data across the interface that are independent of the physical medium being used. This segment is common to all interface profiles. This is analogous to the logical link control sublayer of the OSI data link layer (layer 2).

5.3.2.2 Convergence segment. The convergence segment implements the functions needed to interface a particular physical interface with the logical segment. A different convergence segment is required for each type of physical interface. This is analogous to the media access control sublayer of the OSI data link layer.

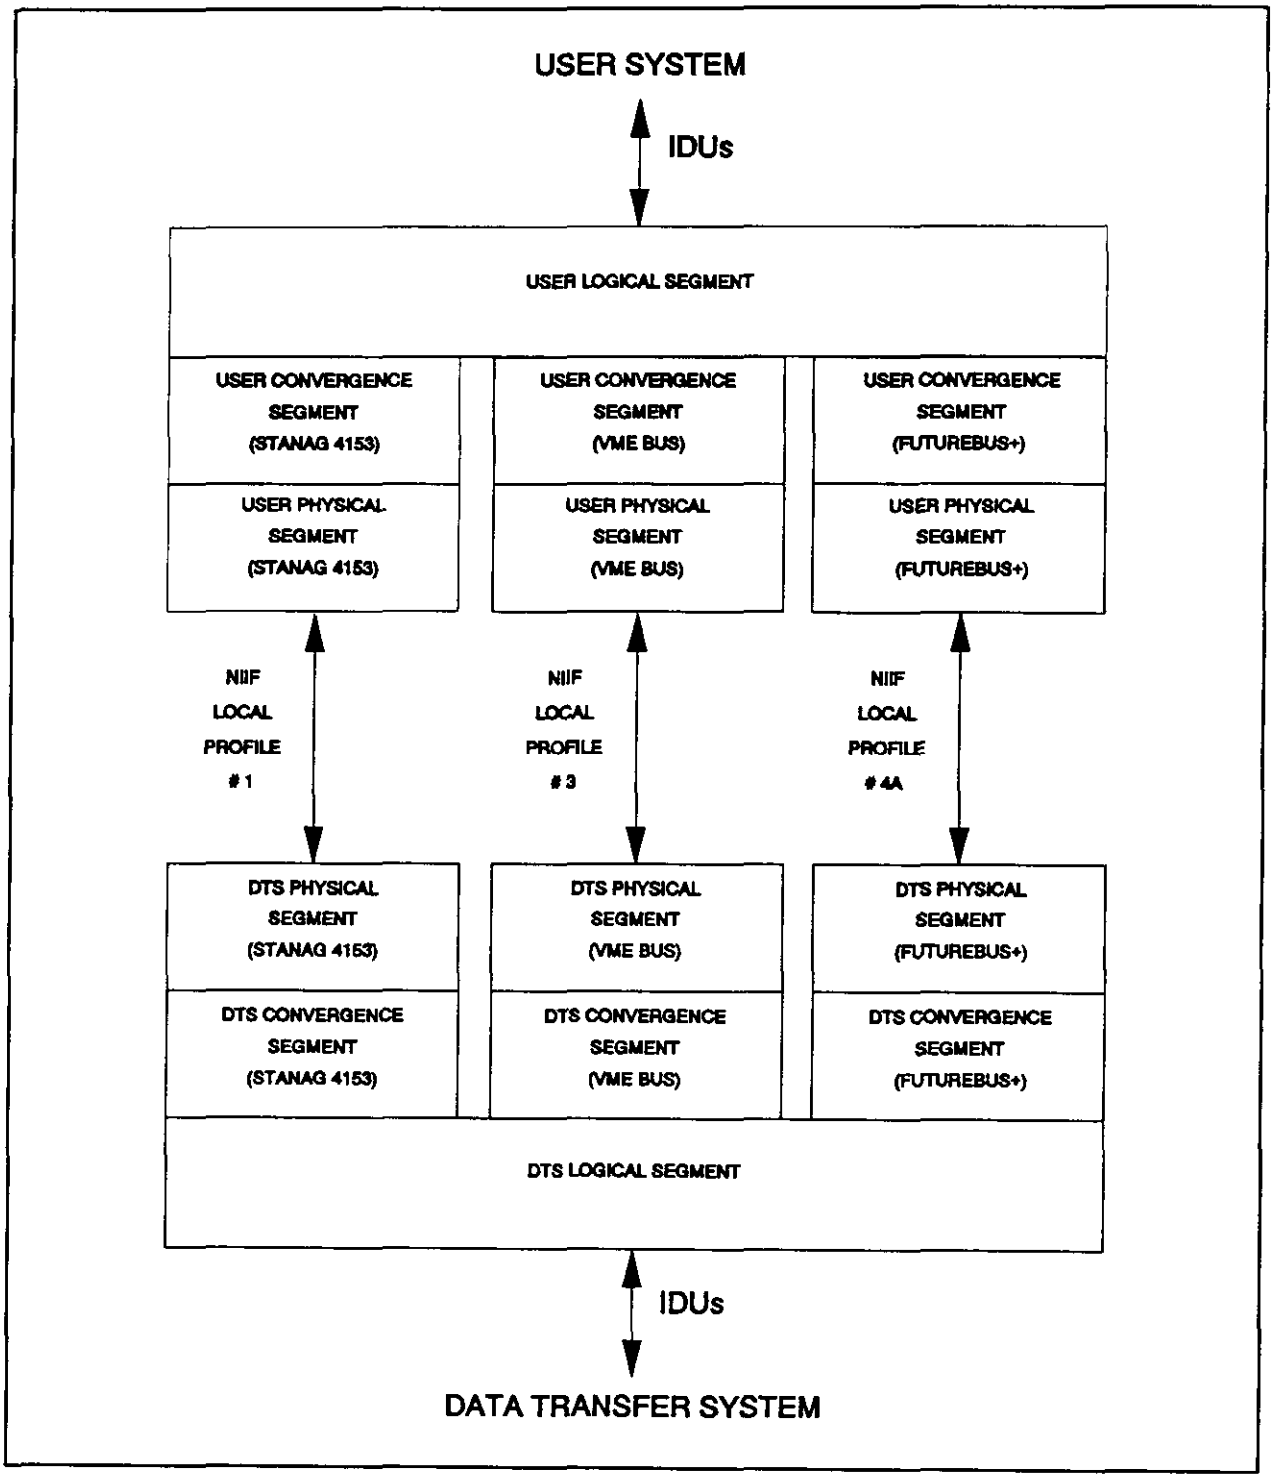

Figure 9. NIIF Local Profile Model

5.3.2.3 Physical segment. The physical segment provides the specific electrical or optical interface being used to transfer data. The physical segment is expected to be a readily available interface specification which can be used to attach the user system and the DTS. It may be a standard type of input/output (such as RS-232) or a standard backplane (such as VME). This is analogous to the OSI physical layer (layer 1).

5.3.2.4 Physical interface selection. The NIIF neither stipulates, nor limits, the choice for the local physical interface. However, the NIIF has specified certain physical interfaces for which it has defined an appropriate convergence segment. Selection of a physical interface which is not specified by NIIF will require that the implementer define an appropriate convergence segment.

6. XTP Implementation Notes

6.1 Introduction. XTP provides SAFENET users with fast, efficient service for those systems which may need this type of functionality. XTP implementors have documented certain behaviors of the XTP protocol which may not be immediately obvious. In some cases, unexpected results occur. These anomalies are not protocol problems, but are conditions which should be understood by XTP users and implementors. Some of these behaviors are presented in this section.

6.2 Context Identification. The source NSAP address and key are used by the destination host to identify a context. Suppose multiple network interfaces are available at a given node and each interface uses a different NSAP. If an interface other than the one on which a connection is established is used (i.e., due to FDDI reconfiguration, load sharing across the interfaces, etc.), context identification fails. Data may be delivered incorrectly or duplicate data may be delivered.

A number of implementation solutions exist. System designers must select a solution which meets their unique requirements.

6.3 Retry Frenzy. It is possible for an excessive number of retry packets to be generated if reliable transfer of a large message is attempted to an off-line node. Suppose for example that Node A is trying to establish an implicit connection with Node B to send a large block of data reliably. Node A sends a FIRST packet followed by a number of DATA packets. If Node B is powered down, the packets are lost. When WTIMER expires in Node A, the FIRST and DATA packets are sent again. If retry_count is set sufficiently high, this process could unnecessarily use bandwidth.

To reduce the amount of traffic generated in cases like this, a number of implementation specific solutions exist. One could, for example, set retry_count to some low number. While this reduces the amount of bandwidth used by the process, it may not be suitable for all applications. An alternative approach would be to limit the amount of traffic sent by lengthening the time between FIRST packet retransmissions. An exponentially increasing timer could be introduced.

6.4 Context Termination. If a system designer decides to use a series of reliable datagrams, there could be a context resource problem if a large number of high rate messages are sent

to the same destination. Since a context must be maintained until CTIMER expires after a connection is closed, a large number of closed contexts may have to be maintained in a node. This problem does not occur if a three-way handshake is used.

6.5 Multicast Reliability. A potential problem exists with multicast return key mapping. Received return keys, that is those that are "primed", are assumed to be unique since they are derived from the unique non-primed keys previously generated at that node. Multicast violates this principle, however, in that packets with return keys can be broadcast to the entire group, not just the originating node.

Suppose Node A and Node B are both multicast senders. Node C is a multicast receiver. Node A sources group 1, or G1, and Node B sources group 2, or G2. Node A, Node B, and Node C are members of both G1 and G2. Suppose Node A and Node B both select the same key, say K, to identify the group. When Node C sends a CNTL packet with key=K' to G1 or G2, the recipients can not tell if it is for G1 or G2.

SAFENET solves this problem by sending the multicast group originator's MAC address in the XROUTE and XKEY fields of CNTL packets. The MAC address is used to identify connections using the same return key and group address.

7. SAFENET TRANSFER SERVICES

7.1 Introduction. The purpose of this section is to provide an overview of the transport, network, and logical link control layers of SAFENET, which are collectively called the SAFENET transfer services. This section primarily describes the OSI protocols specified for use in SAFENET.

7.2 Transport layer services.

7.2.1 OSI transport protocols. In SAFENET the transport layer includes two OSI protocols: connection-oriented OSI Transport Class 4 (TP4) defined by ISO 8073 and connectionless transport defined by ISO 8602. TP4 provides the necessary connection-oriented services for the SAFENET OSI profile, which supports applications requiring reliable, acknowledged information transfer. Connectionless transport provides a basic datagram service for the SAFENET lightweight profile, which supports applications where messages fit into one network packet and the reliability of the underlying network service is adequate.

The OSI model is often divided into the lower layer protocols and the upper layer protocols. The transport layer is considered the highest of the lower layer protocols. The transport layer provides an end-to-end protocol, i.e., a protocol between the source and the ultimate destination. ISO 7498 states, "The transport-service provides transparent transfer of data between session-entities and relieves them from any concern with the detailed way in which reliable and cost effective transfer of the data is achieved." The higher layer protocols can assume a communications system which is reliable enough to meet their needs.

7.2.1.1 Transport Class 4 protocol. The functions performed by the OSI TP4 protocol which are applicable to the SAFENET environment are establishment and release of transport connections, error detection and recovery, flow control, segmentation and reassembly of messages, in-order delivery of messages, expedited data service, provision of user requested Quality of Service (QOS), and an optional checksum.

One TP4 entity communicates with another via Transport Protocol Data Units (TPDUs). The connection is established by the exchange of a Connection Request (CR) and a Connection Confirm (CC) TPDU. The error detection and recovery and flow control functions are accomplished using AK (acknowledge) TPDUs transmitted by the receiving station. Each DT (data) TPDU has a

field containing a sequence number. Every AK TPDU contains the value of the next sequence number expected by the receiver, thereby acknowledging the receipt of all DT TPDUs with lower sequence numbers. The AK also contains a credit field. The transmitting station is allowed to send DT TPDUs with sequence number K where:

$$\text{AK seq no} \leq K < (\text{AK seq no} + \text{AK credit})$$

This range of sequence numbers is known as the window. There are transmission timers on the transmitting station which are restarted when AKs are received. If an AK is not received before the retransmission timer expires, then unacknowledged DT TPDUs are retransmitted.

The user may wish to send Transport Service Data Units (TSDUs) which are too large to send in a single Network Service Data Unit (NSDU). The maximum size of the NSDU is dependent on the type of network. The NSDU plus lower layer headers must fit into one network packet. One of the services offered by TP4 is to segment TSDUs which are too large to fit into a single NSDU into multiple TPDUs which do fit into single NSDUs. At the destination the TPDUs are reassembled into a TSDU and delivered to the user.

A transport connection is normally disconnected by exchange of a Disconnect Request (DR) TPDU and a Disconnection Confirm (DC) TPDU. The transport level may also issue indications to both users that the connection has been disconnected. This is a non-graceful close, i.e., all pending TPDUs are discarded.

The user may request that an Expedited Data (ED) TSDU containing 1 to 16 bytes of user data be sent. ED TPDUs are acknowledged by Expedited Acknowledge (EA) TPDUs. Only one ED TPDU may be outstanding at a time on a connection. The rule for transmission of an ED TSDU is that it may not be delivered to the destination user after DT or ED TSDUs which were submitted subsequently to the transport layer on that transport connection. So the ED TSDU must be served at least in FIFO order. The transport layer expedited data service was not intended to be a general purpose service for sending high priority messages; rather, it was intended to support synchronization features in the session layer.

One of the interesting features of TP4 is how much is left to the implementor. For example, the amount of credit transmitted in AK TPDU can be based on the amount of buffer space the receiver has but it may also be based on other considerations. The destination system may send any credit it

pleases; the amount is not specified in the protocol. Another example is retransmission timers. It is possible to have a timer for each TPDU or one timer per connection. When the timer expires (assuming one timer per connection), there are different possible strategies. One is to reset the timer and start sending from the first unacknowledged DT TPDU and continue until all unacknowledged DT TPDUs have been retransmitted or an AK is received. Another strategy is to reset the timer, send only the first unacknowledged DT TPDU, and wait for an AK. Each DT TPDU does not have to be acknowledged by a separate AK TPDU. An AK TPDU can acknowledge more than one DT TPDU. When to send AKs is up to the implementor. Fewer AKs can result in higher throughput because of less AK processing. However, fewer AKs can also mean that an error takes longer to detect. All these choices have advantages and disadvantages depending on the environment. It should be noted that SAFENET permits the use of any TP4 acknowledgement strategy which conforms to the standard.

7.2.1.1.1 TP4 quality of service parameters. Ideally, a transport protocol implementation should be able to provide the requested Quality of Service (QOS) to the user, monitor whether the QOS goals are being met, take corrective action if QOS goals are not met, and inform the user if QOS cannot be met. Unfortunately, this is a research topic and no current implementation of TP4 does these things. QOS is an area where the OSI standards are likely to change. The following QOS parameters are specified in the transport service standard ISO 8072:

a. Transport Connection (TC) establishment delay
b. TC establishment failure probability
c. Throughput
d. Transit delay
e. Residual error rate (corruption, duplication, loss)
f. Resilience of the TC (probability of transport service initiated disconnect)
g. Transfer failure probability (probability that QOS goals are not met during an observation period)
h. TC release delay
i. TC release failure probability
j. Priority (order in which connections shall have their QOS degraded if necessary to meet the QOS requirements of higher priority connections. The use of a higher priority does not imply that the connection will be served before a connection of lower priority.)
k. TC protection (extent to which a transport service provider attempts to prevent unauthorized monitoring or manipulation of user data)

It should be noted that in ISO 8072, QOS is negotiated among the two transport service users and the transport service provider.

7.2.1.2 Connectionless transport protocol. The OSI connectionless transport protocol is specified in ISO 8602 and the service is specified by ISO 8072/ADD1. The connectionless transport protocol is very simple and offers an unreliable datagram service. Each TSDU must fit into an NSDU since connectionless transport does not provide segmentation and reassembly.

The user may request QOS parameters for transit delay, protection, residual error probability, and priority. In the connectionless transport case, a TSDU of a higher priority is processed by the transport service provider before one of a lower priority. If the network service provider is unable to provide a quality of service to meet the residual error rate specified by the transport service user, the connectionless transport protocol will insert a checksum field that must be checked at the destination.

7.3 Network layer services.

7.3.1 OSI network protocol. The full connectionless network layer protocol defined by ISO 8473 is the network layer protocol required by OSI TP4. If connectionless transport is used in a SAFENET implementation, it is required that all multicast transmissions be over a single local area network. Unicast connectionless transport is not restricted to a single LAN segment.

7.3.2 Connectionless network services. The connectionless network layer provides an unreliable datagram service to deliver NSDUs to Network Service Access Point (NSAP) addresses which may be either on the same subnetwork or on a different subnetwork. The network layer user may also specify the following QOS parameters for each NSDU: transit delay, protection (security), cost, residual error probability, and priority. In contrast to the priority parameter specified in the transport QOS, use of priority in the network service may imply that a NSDU with higher priority is serviced before an NSDU of lower priority.

The primary functions performed by the connectionless network layer are forwarding Network Protocol Data Units (NPDUs) according to the routes determined, NPDU lifetime control, and the segmentation and reassembly of NPDUs which are too large to fit into the subnetwork service data units. Additional functions which may be optionally provided are security, complete or

partial source routing, complete or partial route recording, priority, error reporting, congestion notification, QOS maintenance, and NPDU header padding. The capability to add a checksum for the NPDU header must be present. It is only used, however, when selected by the originating network entity.

7.3.2.1 Network layer routing. The primary service of the network layer is to route a PDU to its destination, even if intervening subnetworks are involved. Two basic strategies can be employed: adaptive routing and source routing.

Adaptive routing is the default strategy for OSI connectionless network service. It requires that the sender supply only the address of the destination. When the NPDU arrives at the router, that router makes the decision on where to send it next. This continues until the PDU arrives at the destination.

Source routing requires that the sender supply a list of intervening routers that must be visited on the way to the destination. Partial source routing occurs when this list of intervening routers is incomplete. When partial source routing is used every router on the list must be visited, and additional routers may also be visited.

7.3.2.2 Segmentation and reassembly. At the source network entity or at any intermediate network entity, the size of the NPDU may be larger than the maximum service data unit size supported by the underlying subnetwork service being used to transmit the NPDU. In this case, the NPDU is segmented into smaller NPDUs by the network protocol (unless the non-segmenting subset is being used). Segmentation could take place several times. Generally, reassembly of the smaller NPDUs into the original NPDU takes place at the final destination, but reassembly can take place at some intermediate point. Each NPDU has a data unit identifier which is not reused within the lifetime of the NPDU. The data unit identifier plus the source and destination NSAP addresses identify the individual segments of an original NPDU. The individual segments are also numbered in order. When the first segment of an NPDU shows up at a reassembly point, a timer is set. If the timer expires before all the segments are received, all received segments of the NPDU are discarded on the assumption that some segments have been lost.

7.3.2.3 NPDU lifetime control. Every NPDU is initialized with a counter. Each time that an NPDU is processed by an intermediate system, the counter is decremented by at least one. When the counter is decremented to zero the NPDU is removed by

the next intermediate system. This keeps NPDUs from circulating in the network indefinitely.

7.3.2.4 Congestion control. When congestion is encountered by an NPDU at an intermediate system a bit in the QOS maintenance field may be set which reports this congestion to the destination station. It is not specified in the standard how to determine whether congestion has occurred, although some guidance is given. This congestion information could be used to make routing decisions. However, in SAFENET use of the QOS maintenance field for routing decisions is specifically forbidden.

7.4 Logical link control services.

7.4.1 Logical Link Control protocol. The SAFENET Logical Link Control (LLC) sublayer consists of the type 1 connectionless LLC protocol specified in the IEEE 802.2 standard. The main function provided by LLC is protocol stack selection through assignment of Link Service Access Points (LSAPs) to protocol stacks. The OSI connectionless network protocol uses LLC type 1 services.

LSAP FE (hex) is used to identify the OSI Connectionless Network Protocol Entity and the Associated OSI routing protocols specified by SAFENET. LSAP AA (hex) is used to identify the direct use of LLC services by XTP.

8. SAFENET CONFIGURATION OPTIONS

8.1 Introduction. SAFENET provides a number of options which allow the system designer to meet the specific requirements for a particular system. Options are available which allow the issues of survivability, availability, cost, weight, and physical size to be traded off for a particular system. These options are provided to make SAFENET useful in environments which, for example, require very little survivability and also in environments which require the highest level of survivability. This section is intended to familiarize the system designer with the options that are available for consideration.

8.2 Criteria for Selection. SAFENET is intended to provide the tools necessary to build a survivable system for mission critical applications. Because some SAFENET applications may not fall into this category, it is not desirable to limit them only to mission critical applications. To keep land based and non-mission critical systems simple and inexpensive, SAFENET provides for options to meet the budget and needs of a particular application.

A system designer should determine the environment of a particular system in order to select proper options and tools from the SAFENET standard. The following items should be taken into consideration when determining the best components for a system:

Availability - what is required in order for the system to remain operational,
Level of criticality - how critical is it that the system remain operational,
Expense - how much will components and cabling cost,
Required robustness in connection - what type of station connections are available,
Required robustness in reconfiguration - can stations be connected after primary paths are broken,
Ability for future expansion - will the network grow over time,
Area of installation - what type of barriers or equipment are near the cables,
Ease of maintenance - what is involved in maintaining the cable plant and stations.

8.3 Fiber. There are many important considerations when designing a local area network. Among them are the interconnection of fibers, cable topology, connection of stations, use of trunk coupling units, and optical power budget.

8.3.1 Fiber splices vs. connectors. Unlike copper cabling, which has traditionally been used in networking environments, fiber optic cables are not as easily maintained or modified. When a fiber is cut, either for installation or as a result of damage, either the fiber can be spliced together or a connector can be used. The manipulation of fibers may require chemicals or other materials which may be prohibited under certain circumstances. The type of splice or connector used also affects the optical power budget.

8.3.1.1 Splices. Splices connect two fibers by heat welding, chemical fusing, or mechanical pressure. These methods produce the best type of fiber connection and result in low loss. However, splicing requires both an experienced technician and accurate equipment. Additionally, each type of splice has its own characteristics for loss. Although splicing results in low loss, some loss is encountered and the number of splices on a single cable is limited. Splices have the limitation that once made, they cannot be removed easily without breaking the fiber another time. For long runs of fiber, splices should be used to reduce the overall loss of the connection. Splicing is the recommended method for connection in survivable networks.

8.3.1.2 Connectors. Connectors are mechanical coupling devices which connect two fibers together. When used, the network can be maintained through simple replacement of short cables. Since connectors are mechanical devices which hold two fiber cables together, they are easy to open up quickly to allow for data monitoring. This is useful for monitoring stations and diagnosing problems. Additionally, once a connector is installed on a fiber, the fiber can be relocated to a different area and the interconnection of fibers can be accomplished more easily in the field than is possible with splices. Connectors are best suited to benign environments where devices are relocated frequently, and networks that can generally handle short interruptions.

8.3.2 Connections at stations. Stations can be connected to the data path using either an ST connector, a Media Interface Connector (MIC), or a Fiber Optic Interface Connector (FOIC).

8.3.2.1 ST. The ST connector is the easiest connector to attach to a fiber and is the least expensive. This connector is frequently used in benign environments because it is unkeyed (making the fiber usable for either transmit or receive) and self contained. It is often used for network monitoring taps or for attaching to a single fiber for repairs when splicing is not used. The ST connector has the advantage that when multi-fiber

cables are used, the fibers can be used interchangeably for either transmit or receive.

8.3.2.2 Media Interface Connector (MIC). The MIC connector is the ANSI X3T9.5 approved connector for all commercial hardware. This is the most popular type of connector. The connector contains both the transmitting and receiving fiber in one shell. The MIC receptacle is polarized to prevent improper attachment of input and output fibers. The MIC receptacle also provides a physical key to prevent improper MIC plug attachment, and to control the types of FDDI ports (such as A, B, S, or M) which can be attached with the fiber. Additionally, these connectors offer a quick connect/disconnect which makes them very useful for equipment or networks which are frequently relocated. Land based operations which do not require environmental protection benefit from these connectors because they contain both fibers (transmit and receive), are keyed so that the user cannot connect the wrong station to the network, and are quick to connect and disconnect.

8.3.2.3 FOIC. The Fiber Optic Interface Connector is a circular, heavy duty, multi-fiber connector specified in MIL-STD-2204. This connector provides for backup fibers which can be used in the event of damage to the cable. This is the connector required by SAFENET for severe environments in which survivability is paramount. Due to the MIL qualification, this connector is the most expensive of the three.

8.3.3 Trunk coupling units. A trunk coupling unit (TCU) is a device which contains fiber optic switches which either pass data around the station (when the station is disabled or deinserted from the network) or allow the data to pass to the station (when the station is enabled and inserted onto the network). The TCU is needed for survivable interconnection of concentrators or stations on the trunk ring of a network, or concentrators connected using the concentrator tree topology, because they allow for devices to be powered off without breaking the token path for an extended period. Due to the nature of the operation of a fiber optic switch, a TCU presents a considerable optical power loss in the data path and if several serial TCU's are used some compensation will be required, either with repeaters or extended power levels. A trade-off is required relative to the use of TCU's between survivability requirements, cost of TCUs vs. repeaters, and the use and cost of extended power levels. TCUs also require an additional electrical cable interface to provide power to the switches to control station insertion.

The internal functionality of concentrators allow them to control the attachment of stations which are powered down, or operating in error, without the use of TCUs. Additionally, when the trunk ring is composed of three or less stations, the use of a TCU is unnecessary. The interconnection of three or less devices by a trunk ring does not present the isolation problems which exist in larger systems.

8.3.4 Power levels. SAFENET provides for three different power levels which allow complete interconnection of components. These power levels range from standard FDDI to increased power at the transmitter and increased sensitivity at the receiver. The various power levels provide for increased survivability in case more than one node needs to be bypassed.

8.3.4.1 Standard FDDI. Standard FDDI power levels are those specified by FDDI PMD (ISO 9314-3). These are considered to be a commercial standard and are required for the minimum FDDI transmitting power and minimum receiver sensitivity to guarantee a power budget of 11 dBm. Since this power level is considered the FDDI standard, it is obviously the most common power level available in commercial equipment.

8.3.4.2 Increased receiver sensitivity. To allow for at least one station to be bypassed between active stations and still meet the military requirements on optical power tolerances, the sensitivity of the receiver may be increased. The benefit of using this increased sensitivity is realized in systems where repeaters are extensively used, or when it is unlikely that more than one station will be down between active stations. Increasing the receiver sensitivity effectively increases the dynamic range by 4 dBm.

8.3.4.3 Increased transmitter power/receiver sensitivity. Increasing the transmitter power and the receiver sensitivity allows SAFENET networks to bypass more stations without loss of the network. For mission critical stations connected to the trunk ring, this is the recommended SAFENET configuration. This increases the dynamic range by 8 dBm over that of FDDI. See MIL-STD-2204 for a complete definition of the three power levels described above.

8.4 Cable topology. Cables typically contain more than one fiber. The system designer must choose the proper number of fibers per cable to minimize the cost of installing the cable plant, maintaining the fibers and connectivity, and allowing for expansion.

8.4.1 Dual fiber pair cables. Under the basic FDDI reconfiguration policies, the fibers of the network are normally run in dual fiber pairs (i.e., FDDI linked pairs). In this case the primary and secondary data path fibers are located in the same cable. Since the FDDI reconfiguration mechanism known as wrap bypasses both fibers when one of the fibers is broken, the undamaged fiber is unused when this occurs. Consequently, there is no inherent advantage in separating the individual fibers. Using dual fiber cables makes overall cable plant management easier and reduces the number of actual cables to be pulled for a network. This is the most common cable plant for commercial networks.

8.4.2 Multi-fiber cables. Multi-fiber cables normally contain multiples of two fibers and can save on the overall cost of pulling the cable plant through an installation. The multi-fiber cables can be used to provide backup fibers or for increasing the number of FDDI linked pairs in a single cable. For harsh environments, multi-fiber cables are primarily used to allow for standby replacement fibers, but require human intervention to switch to another fiber. In benign environments, multi-fiber cables can be used to connect a number of stations in a close location to the network without having to pull fibers for each station.

8.4.3 Diverse routing vs. single cable topologies. Implementors may choose to install SAFENET with the optical fibers either in a single cable, or in two cables routed separately to the network stations. Factors which may affect this decision are the location where the network is to be installed, cost of components, and pre-existing constraints on the physical layout, which may occur on a ship.

One option is to enclose the FDDI dual rings together in the same cable. This will save on component cost, since less fiber optic cable is necessary. In some cases it may be the only viable option, for example if the number of available cable ways on a ship is severely limited. There are survivability implications to consider for this case. Having only one cable reduces the total number of places where the network can sustain damage, but if damage does occur it is quite likely that both rings will be severed.

The other option is to have separately routed dual rings. Since there is roughly twice as much cable in this configuration, this increases the exposure to probable damage. This also increases the component cost. If the implementor so chooses, the FDDI global hold policy may be employed in this configuration.

This survivability feature provides benefit only if the rings are cabled separately.

8.5 Network devices and connection options. Network stations and network devices such as repeaters are connected to a SAFENET network through various techniques. This section covers the types of devices and connections within a system.

8.5.1 Repeaters. Repeaters can be used to boost the optical power to allow for more stations to be bypassed at any time. Repeaters are recommended at any point in the network where available power to a downstream station is a concern.

Repeaters can be built using a variety of methods, and they should permit Line States to be retransmitted without filtering (i.e., disabling the repeat filter function from the PHY). The clock frequency and symbol jitter should meet FDDI PHY requirements.

8.5.2 Concentrators. Concentrators allow single attachment devices to connect to the FDDI data path while still maintaining the fundamental FDDI linked pairs. Concentrators in SAFENET contain a minimum of one MAC. The internal functionality of concentrators allow them to control the attachment of stations which are powered down, or operating in error, without the use of TCUs.

8.5.2.1 Single MAC. Single MAC concentrators cannot perform the more advanced functions such as multiple local path management. Single MAC concentrators are the least expensive concentrators allowed in a SAFENET system and provide minimal functionality.

8.5.2.2 Dual MAC. A concentrator with two or more MACs provides functionality beyond that of a single MAC concentrator. It has the additional capability of using one of the MACs for diagnostics, and for management of additional local data paths. One advantage of having multiple MACs within a concentrator is that the network can be separated into subnetworks using a local path such that the MAC can test equipment prior to inserting it onto the operational ring.

8.5.3 Station attachment. Stations may be either single attached or dual attached to the FDDI data path. Further, dual attached stations may be connected in such topologies as the trunk ring or the concentrator tree, or may be dual-homed.

8.5.3.1 Single attachment. Single attachment stations connect to the network through a concentrator. They are

particularly useful for duplicated resources where loss of one station is tolerable.

8.5.3.2 Dual attachment. A dual attachment station can be connected as either a peer or a slave connection. The robust types of connections which can be made with a dual attachment station make them essential for a survivable system.

8.5.3.2.1 Trunk ring. When a station is attached to the trunk ring, it is connected to both the primary and secondary rings and participates in trunk ring reconfiguration. A station which attaches to the trunk ring must also provide an electrical connection for controlling a TCU if it is to be used for a tactical mission critical system.

8.5.3.2.2 Dual homing. A station which is attached to two concentrators is dual homed and provides a standby port for reconfiguration. A dual homed station does not require an electrical interface because a TCU is not required for attachment, though TCU's may be used if desired in this case.

8.5.3.2.3 Concentrator tree. Concentrators having both ports connected to one other concentrator form the basis for the concentrator tree topology. This topology uses TCUs to prevent segmentation when one of the concentrators is powered down. Stations may be dual homed or single attached to the concentrators in the tree, or the concentrators may act as stations independently. The concentrator tree uses the characteristics of dual-homing to achieve a survivable topology.

8.5.4 Multiple subnetworks. The use of multiple subnetworks is the connection of multiple independent FDDI networks through the use of routers. Multiple subnetworks provide the advantage that each ring may be independently designed without concern about the impact on other components or systems in the network. Each subnetwork contains its own token and all data is maintained on the individual rings independently of data on other subnetworks. Data is passed to other networks only when the destination address is not on the local network.

9. SAFENET OPTICAL POWER BUDGET

9.1 Introduction. System optical power budget calculations for SAFENET are presented in this section for two purposes. The first is to demonstrate that the network will operate as required using the signal and component values specified in SAFENET. The second is to provide the system designer with some insight into the problem of determining the placement of repeaters in the network. In the calculations shown, the physical media consists of worst case cable lengths. Tuned splices are used for connections to TCUs.

9.2 Power budget calculations.

9.2.1 Operating conditions. A correct implementation of SAFENET will provide for a substantial unallocated power margin in the power budget. This power margin (M) allows for losses due to factors such as component ageing, environmental conditions, power supply variations, and repairs made to the physical medium. This power margin does not include losses due to temperature variation since each component is required to remain within specification over its operating temperature range. For a SAFENET network operating under normal conditions the power margin (M) is as follows:

$$M = 6 \text{ Db}$$

This margin may not be adequate for a SAFENET network operating in an extremely adverse environment. The effect of a such an environment on the power margin is discussed in 9.2.3.

Any multi-mode optical link is subject to higher order mode loss (HOL). This term accounts for the loss of power due to the dissipation of higher order optical modes as an equilibrium mode distribution is approached in the optical link. For SAFENET, the higher order mode loss will be assumed to have the following maximum value:

$$HOL = 1.0 \text{ dB}$$

In SAFENET the loss due to signal dispersion (i.e., the dispersion power penalty) is not taken into account. It is assumed that the station receiver sensitivity is appropriately derated to account for this loss.

All other values used in the following calculations will be taken directly from the physical layer signalling requirements and the physical medium requirements specified in the SAFENET

standard, MIL-STD-2204. These values will be cited in 9.2.2 below.

9.2.2 Worst case system power budget. The worst case system power budget is significant because it shows whether or not a SAFENET implementation using the specified signals and components will work under normal operating conditions. The functional requirement is that the optical link between the transmitting and receiving stations be able to operate with at least four intervening bypassed stations.

The first step is to compute the power budget (P) available to the optical link. The equation is as follows:

$$P = [\ Po(min) - Pi(min)\ ] - M \qquad \text{(Eq. 1)}$$

Po(min) is the minimum required output signal power
Pi(min) is the minimum required input signal power
M is the unallocated power margin

Po(min) = -16 dBm (see MIL-STD-2204)
Pi(min) = -35 dBm (see MIL-STD-2204)
M = 6 Db

So: $P = [-16 - (-35)] - 6$; $P = 13$ dB

The next step is to compute the link budget to ensure that the link losses do not exceed the available power budget (P). However, optical links in SAFENET are of variable length, depending upon the presence of intervening bypassed stations in the link. For this reason, the optical link will be considered in two parts: the minimum link (ML), and the variable link (VL).

The minimum link is diagrammed in the Figure 10. As shown, the minimum link in SAFENET is a link between physically adjacent stations on a ring. The optical signal on this link will pass through, in order, the following components:

-- an FOIC and an optical interface cable
-- 2 splices and a TCU (transmit path)
-- a trunk cable
-- 2 splices and a TCU (receive path)
-- an optical interface cable and an FOIC

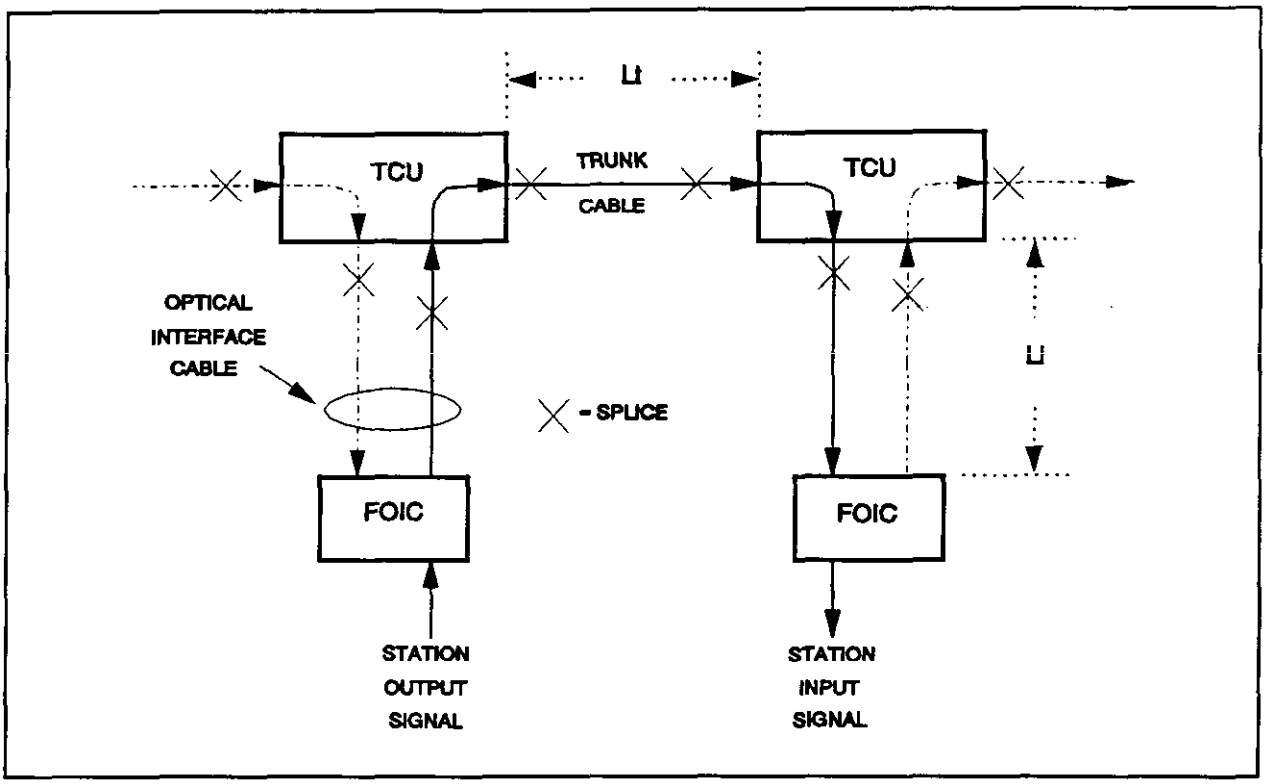

Figure 10. Diagram of the Minimum Link

Therefore, the minimum link equation is as follows:

$$ML = (Lt+2Li)*Cf + 2Sw + 4Sp + 2Co + HOL \qquad \text{(Eq. 2)}$$

Lt is the length of a trunk cable run
Li is the length of an optical interface cable run
Cf is the length-dependent loss of the fiber optic cable
Sw is the optical power loss of a bypass switch (TCU) along the bypass, transmit, and receive paths
Sp is the optical power loss of an splice
Co is the optical power loss of an FOIC
HOL is the higher order mode loss

Lt = 0.2 Km	(max)	(see MIL-STD-2204)
Li = 0.1 Km	(max)	(see MIL-STD-2204)
Cf = 2.0 dB/Km	(max)	(see MIL-STD-2204)
Sw = 0.8 dB	(max)	(see MIL-STD-2204)
Sp = 0.2 dB	(max)	(see MIL-STD-2204)
Co = 1.0 dB	(max)	(see MIL-STD-2204)
HOL = 1.0 dB		

So for the worst case:

ML = (0.4)*2.0 + 2(0.8) + 4(0.2) + 2(1.0) + 1.0 ; ML = 6.2 dB

The variable link is diagrammed in Figure 11. As shown, the variable link is that series of optical components which will be inserted into the optical link when a station is bypassed at its TCU. The variable link consists of a trunk cable run, a TCU (bypass path), and two splices. Therefore, the variable link equation is as follows:

VL = Lt*Cf + Sw + 2Sp (Eq. 3)

All terms and values are as defined above. So for the worst case:

VL = (0.2)*2.0 + 0.8 + 2(0.2) ; VL = 1.6 dB

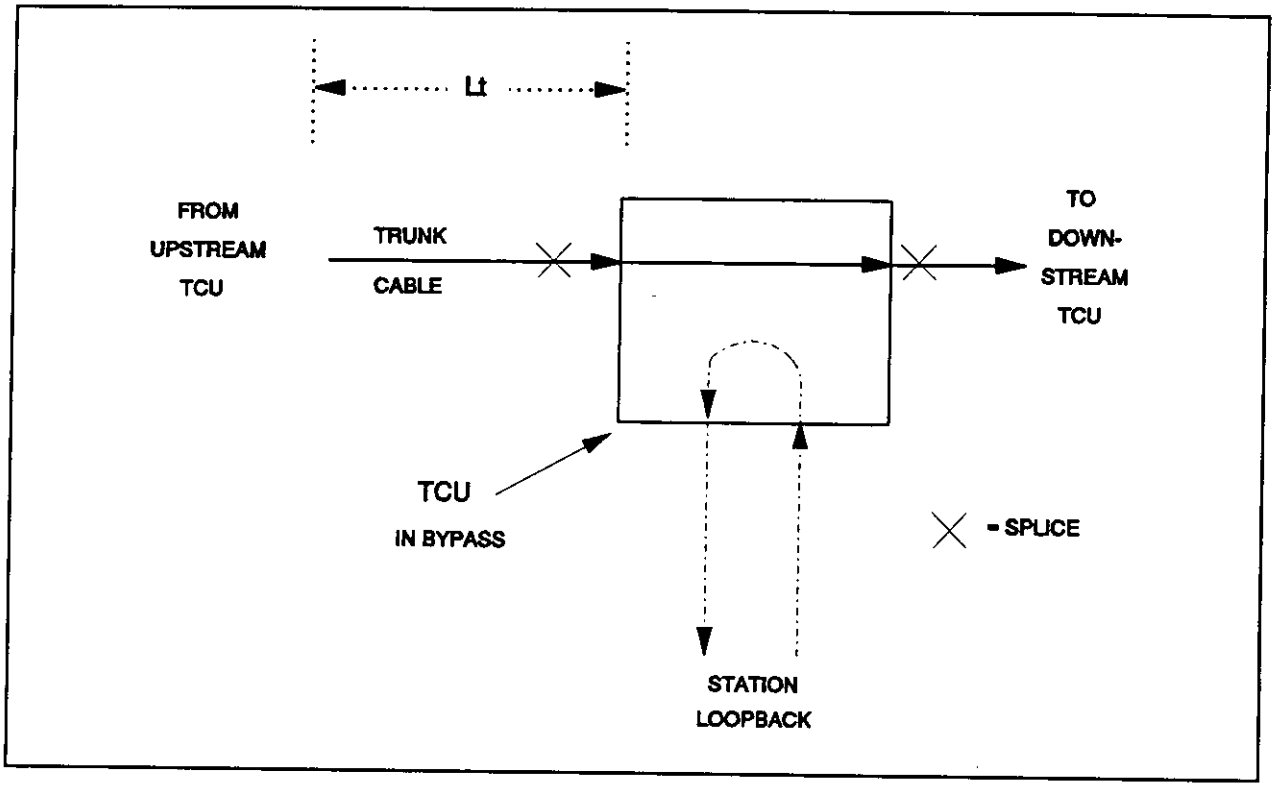

Figure 11. Diagram of the Variable Link

An expression for a total link (L) can be written in terms of the minimum link and the variable link. The equation for the total link is as follows:

$$L = ML + N*VL \qquad \text{(Eq. 4)}$$

N is the number of variable links in the total link, i.e., it is the number of bypassed stations intervening between the transmitting and the receiving station.

For a SAFENET network to operate in the worst case, it should be able to maintain an optical link when N = 4. That is, the following condition must be met:

$$P \geq L \ ;$$

which gives:

$$P \geq ML + 4*VL \qquad \text{(Eq. 5)}$$

Using the calculations given above:

$$L = 6.2 + 4*(1.6) \ ; \qquad \underline{L = 12.6 \text{ dB}}$$

Since P = 13 dB, P is greater than L, and the condition cited in Equation 5 has been met. Therefore, even with worst case component losses a SAFENET network should be able to operate with four consecutive bypassed stations.

A different way to show this would be to solve for the maximum number of bypassed stations which could be allowed in an operating optical link. This will occur when P = L, which leads to the following equation:

$$P = ML + N*VL \ ;$$

which gives:

$$N = (P-ML)/VL \qquad \text{(Eq. 6)}$$

For the worst case values this gives:

$$N = (13 - 6.2)/(1.6) \ ; \qquad \underline{N = 4}$$

As shown before, even in the worst case a correct implementation of SAFENET should be able to operate in the presence of four consecutive bypassed stations.

9.2.3 Adverse environmental conditions. The system power budget under extremely adverse environmental conditions is similar to the worst case power budget calculated in 9.2.2 for normal operating conditions. The primary difference is that a power margin of 6 dB may not be sufficient to include all the

losses which could be caused by these environmental conditions. Conditions which may be of interest to a system implementer include shock, vibration, humidity, sand and dust, compression, mechanical stress, and nuclear radiation.

To account for the effect of these conditions, a loss margin value must be assigned to each optical component in the system. This component margin represents the additional loss which a component may be expected to contribute to the link budget as a result of the simultaneous occurrence of all the above adverse environmental conditions. Under these conditions, a components loss will equal its normal loss plus its loss margin.

An assumed value will be assigned to each component for the purpose of illustrating the effect on the system power budget. The loss margins given below are not intended to include the effects of radiation. These component loss margins are defined as follows:

Mcf is the length-dependent loss margin assumed for fiber optic cable
Msw is the loss margin assumed for a bypass switch (TCU) along the bypass, transmit, and receive paths
Msp is the loss margin assumed for an splice
Mco is the loss margin assumed for an FOIC

Mcf = 1.3 dB/km
Msw = 0.7 dB
Msp = 0.1 dB
Mco = 0.25 dB

Using these component margins, a composite loss margin can be determined for a minimum link by modifying Equation 2 as follows:

$$M(ml) = (Lt+2Li)*Mcf + 2Msw + 4Msp + 2Mco \qquad \text{(Eq. 7)}$$

Using the loss margins listed above:

$$M(ml) = (.4)*1.3 + 2*(.7) + 4*(.1) + 2*(.25)$$

So: $\underline{M(ml) = 2.82 \text{ dB}}$

Similarly, a composite loss margin can be determined for a variable link by modifying Equation 3 as follows:

$$M(vl) = Lt*Mcf + Msw + 2Msp \qquad \text{(Eq. 8)}$$

So: $M(vl) = (.2)*1.3 + 0.7 + 2*(.1)$; $\underline{M(vl) = 1.16 \text{ dB}}$

To finish computing the system power budget for adverse environmental conditions, new values for available power budget (P), minimum link (ML), and variable link (VL) need to be determined. These will be called Pa, MLa, and VLa to represent the adverse condition. Pa is determined by using Equation 1.

Pa = [Po(min) - Pi(min)] - Ma

Ma is the unallocated margin in adverse environmental conditions. Since component losses are now being considered separately, Ma only needs to account for active device aging and any gradual degradations which occur in the link. The following value is assigned:

Ma = 3 dB

Thus: Pa = [-16 - (-35)] - 3; Pa = 16 dB

MLa is simply computed by adding the previously determined value of ML to the minimum link margin calculated above.

MLa = ML + M(ml) (Eq. 9)

MLa = 6.2 + 2.82; MLa = 9.02 dB

VLa is computed in similar fashion.

VLa = VL + M(vl) (Eq. 10)

VLa = 1.6 + 1.16; VLa = 2.76 dB

Using these values, the number of stations which can be bypassed in the worst case under these adverse environmental conditions (Na) can be determined by using Equation 6.

Na = (Pa - MLa)/VLa

So: Na = (16 - 9.02)/(2.76); Na = 2

Therefore, this calculation demonstrates that given the component loss margins assumed above a SAFENET network should be able to operate an optical link with two intervening bypassed stations, even under extremely adverse environmental conditions. Two points need to be emphasized regarding this calculation.

First, a calculation like this can only be meaningful if the component loss margins used represent the known, measured characteristics of the actual components under the adverse

conditions. Currently, this type of information is not available for most components and must be estimated, modelled, or assumed.

Second, these loss margins represent the optical response of the network components to the simultaneous occurrence of several unrelated adverse environmental conditions. The system designer should consider carefully whether designing to meet this set of conditions is necessary to ensure the reliable operation of the network.

9.3 Repeater placement. It was shown in 9.2.2 that under normal operating conditions even a worst case implementation of SAFENET will be able to maintain an optical link across four bypassed stations. The system designer can use this information to determine the interval at which repeaters should be inserted in the network. For the case given above, a repeater should be attached at every fifth TCU to ensure network continuity.

However, not all SAFENET installations will need to be designed for worst case link losses. For example, suppose it is known that a specific SAFENET installation is going to be limited to the following cable run lengths:

Lt = 0.03 Km (max)
Li = 0.02 Km (max)

Using this information, the maximum number of bypassed stations (N) can be recomputed. For the following link calculations all component loss values will be held to the previously cited maximums. The new minimum link value can be calculated from Equation 2, resulting in:

ML = 5.54 dB

The new variable link value can be calculated from Equation 3, resulting in:

VL = 1.26 dB

To find the new value for N, Equation 6 is used as follows:

N = (13 - 5.54)/(1.26) ; N = 5

These calculations show that a SAFENET network with these maximum run lengths should be able to operate in the presence of five consecutive bypassed stations. Thus, the system designer would only have to use repeaters at every sixth TCU, rather than every fifth TCU. Of course, similar computations of N can be

performed for other conditions, such as using components whose optical losses are known to be less than the specified maximums.

10. SAFENET RECONFIGURATION AND SURVIVABILITY

10.1 Introduction. This section is provided to illustrate some of the different methods of using the FDDI reconfiguration policy and the components defined by SAFENET to produce survivable systems. It is by no means an exhaustive list and does not intend to preclude legal topologies that may be desired by some system designers. This section does not discuss all legal topologies that may be of use in systems which do not require survivability beyond what can be obtained with standard FDDI components.

MIL-STD-2204 provides implementation flexibility by defining allowable components rather than specific topologies. Recognizing that different systems may have unique sets of requirements, the system designer is allowed this flexibility so as not to preclude implementations that may meet survivability requirements of a particular system.

10.2 Background. In SAFENET, the FDDI LAN standard includes mechanisms which support operating the token ring network as a pair of counter-rotating rings. This dual-ring architecture provides the capability for the network to reconfigure automatically in the presence of faults, such as a break in the trunk cable. This automatic reconfiguration capability is inherent in FDDI, where upon a failure or casualty to either ring the stations on either side of the break isolate it from the network by connecting the two rings together. This is referred to as ring wrap. This is also an important and inherent survivability feature of SAFENET.

The SAFENET topology provides additional survivability by requiring each station that attaches to the trunk ring to connect by way of a trunk coupling unit (TCU), as shown in Figure 12. In the event of a station failure the TCU can be used to isolate the failed station from the ring using a procedure known as station bypass. A SAFENET station is capable of transmitting data to a receiving station past multiple bypassed stations, as described in section 9.

The SAFENET topology provides additional survivability by permitting the connection of stations and concentrators as defined in the FDDI Station Management (SMT) standard and shown in Figure 13. When a device is dual homed to concentrators in an FDDI network, in the event of a failure such that the concentrator attached to the B port becomes isolated from the trunk ring, the device will switch to receive data from the A

port. Essentially, this method provides simple redundancy by providing a backup link to the data path.

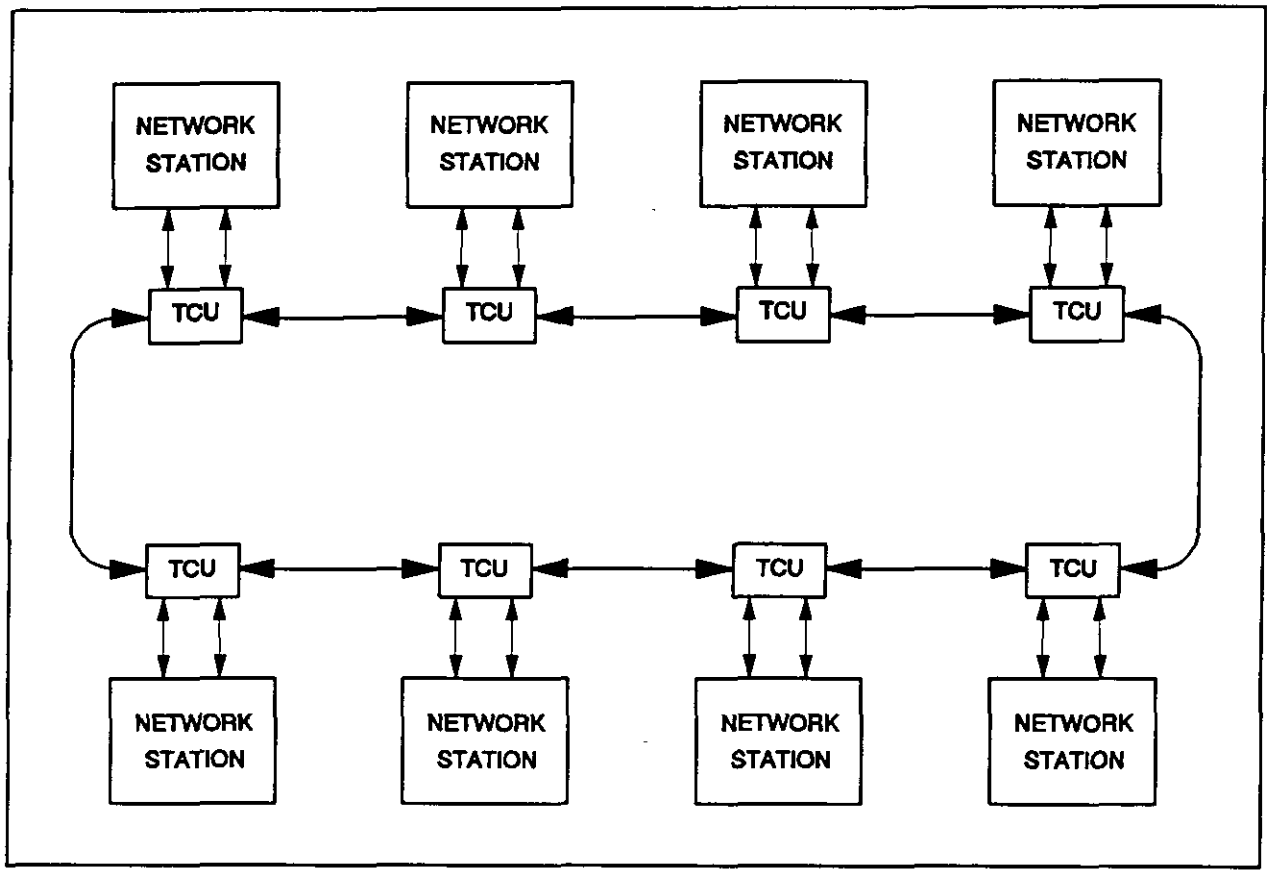

Figure 12. Basic Network Topology

The SAFENET topology also provides survivability by permitting the key network components to be located apart from each other. The trunk coupling units can be located away from their attached stations and concentrators can be located with attached stations but away from the trunk ring. While the network components can be located apart from each other, it is recommended that the trunk ring fibers be cabled together (unless the FDDI Global Hold policy is implemented) so as to provide fewer places for the network to absorb damage. In addition, SAFENET allows the implementation of multiple subnetworks for added survivability. All of these features allow the network to absorb some damage without losing its ability to operate.

10.3 Survivability techniques. A SAFENET network consists of the components defined in MIL-STD-2204. These components may

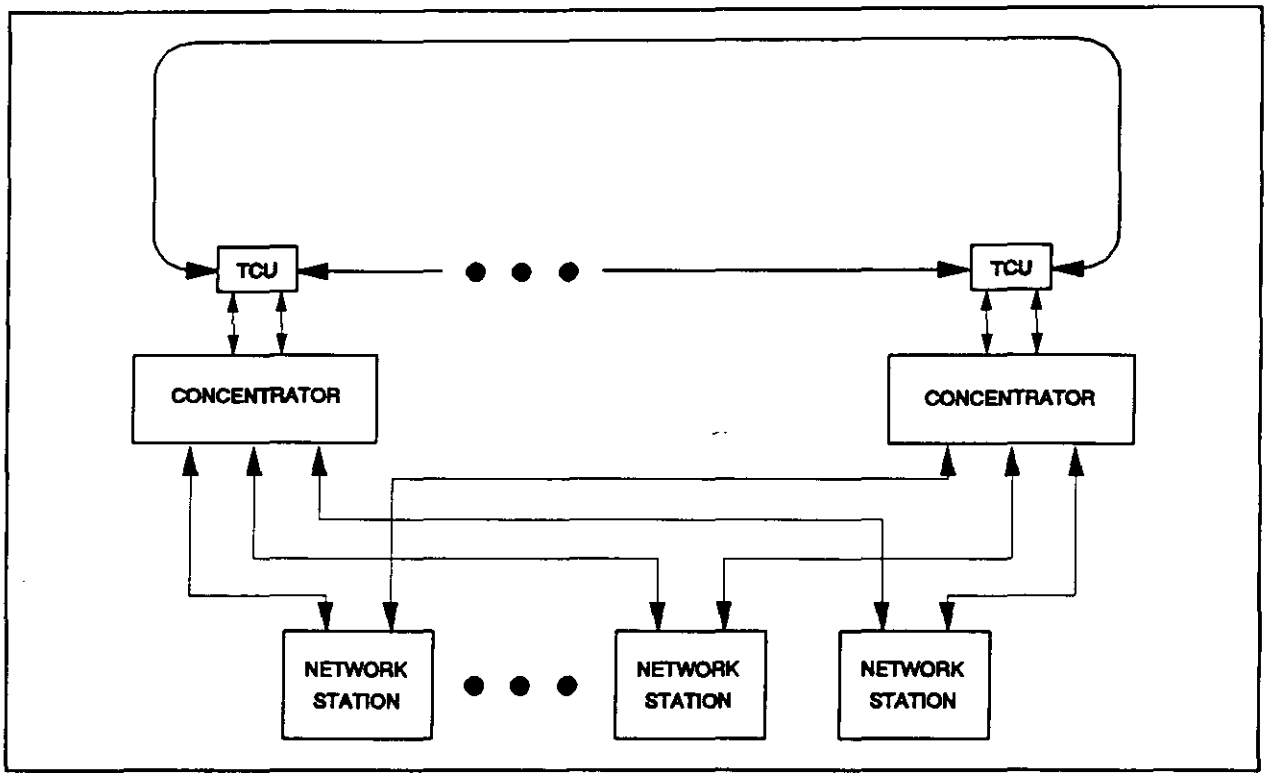

Figure 13. Basic Dual Homing Topology

be implemented in various configurations, some of which may be more survivable than others. Examples of some survivable topologies are given below. This is not intended to provide an exhaustive listing of network topologies and should be considered only as a sampling of the techniques available. Other implementation techniques are possible and are not intended to be precluded by SAFENET.

10.3.1 Dual homed topology. An example dual homed physical topology is shown in Figure 14. As shown, this topology is based on the use of an FDDI dual-ring, with the primary and secondary rings of this dual-ring cabled together to form a single FDDI dual-ring trunk. Dual homed DASs, concentrators, and SASs can be used in the dual homed topology.

10.3.1.1 Dual homing reconfiguration. The dual homing reconfiguration method employs the standard FDDI connection rules to enable network stations to switch their communications from one attachment on the FDDI dual-ring (via a concentrator) to a second such attachment in the case of a hard fault in the first attachment. Implementation of the dual homing reconfiguration

method does not require any particular protocol extensions or mechanisms beyond those provided by the FDDI standards. A description of the dual homing mechanism is provided in the following paragraphs.

10.3.1.2 Dual homing mechanism. The dual homing mechanism is based on the connection rules established in the FDDI standard. A SAFENET defined dual homed dual attachment station will include two FDDI defined ports: an A Port and a B Port. Each of these ports is attached to an M Port on two different concentrators. Thus, two FDDI connections are formed: an A-to-M connection and a B-to-M connection. FDDI defines these connections to be redundant tree connections, and further states that the B-to-M connection will have precedence over the A-to-M connection in a single MAC node.

The result of these connection rules is that a properly connected dual homed network station will be actively inserted on the FDDI LAN through the concentrator to which it has attached its B Port through a tree connection. Furthermore, it will have established a back-up connection to the FDDI LAN through the concentrator to which it has attached its A Port. With these connections in place, the following fault responses will be performed automatically in accordance with the FDDI protocols:

a. In the event of a fault on the dual-ring trunk, the stations on both sides of the fault will perform wrap reconfiguration.
b. In the event of a fault to a dual-homed DAS's A Port tree connection or to the concentrator attached to the dual homed DAS's A Port, no reconfiguration will occur as these provide the back-up tree connection.
c. In the event of a fault to a dual-homed DAS's B Port tree connection or to the concentrator attached to the dual homed DAS's B Port, the station's communications will switch to the back-up tree connection (supporting A Port).

10.3.1.3 Dual homing components. The components of the dual homed topology are concentrators and attached network stations. See MIL-STD-2204 for complete definitions of these components and methods of interconnecting them.

10.3.1.3.1 Dual homed dual attachment stations (DASs). Dual homed DASs used in the dual homing topology are attached to two separate concentrators through tree connections. Each dual homed DAS uses two tree connections, one for each concentrator connection. The concentrators to which the dual homed DAS is

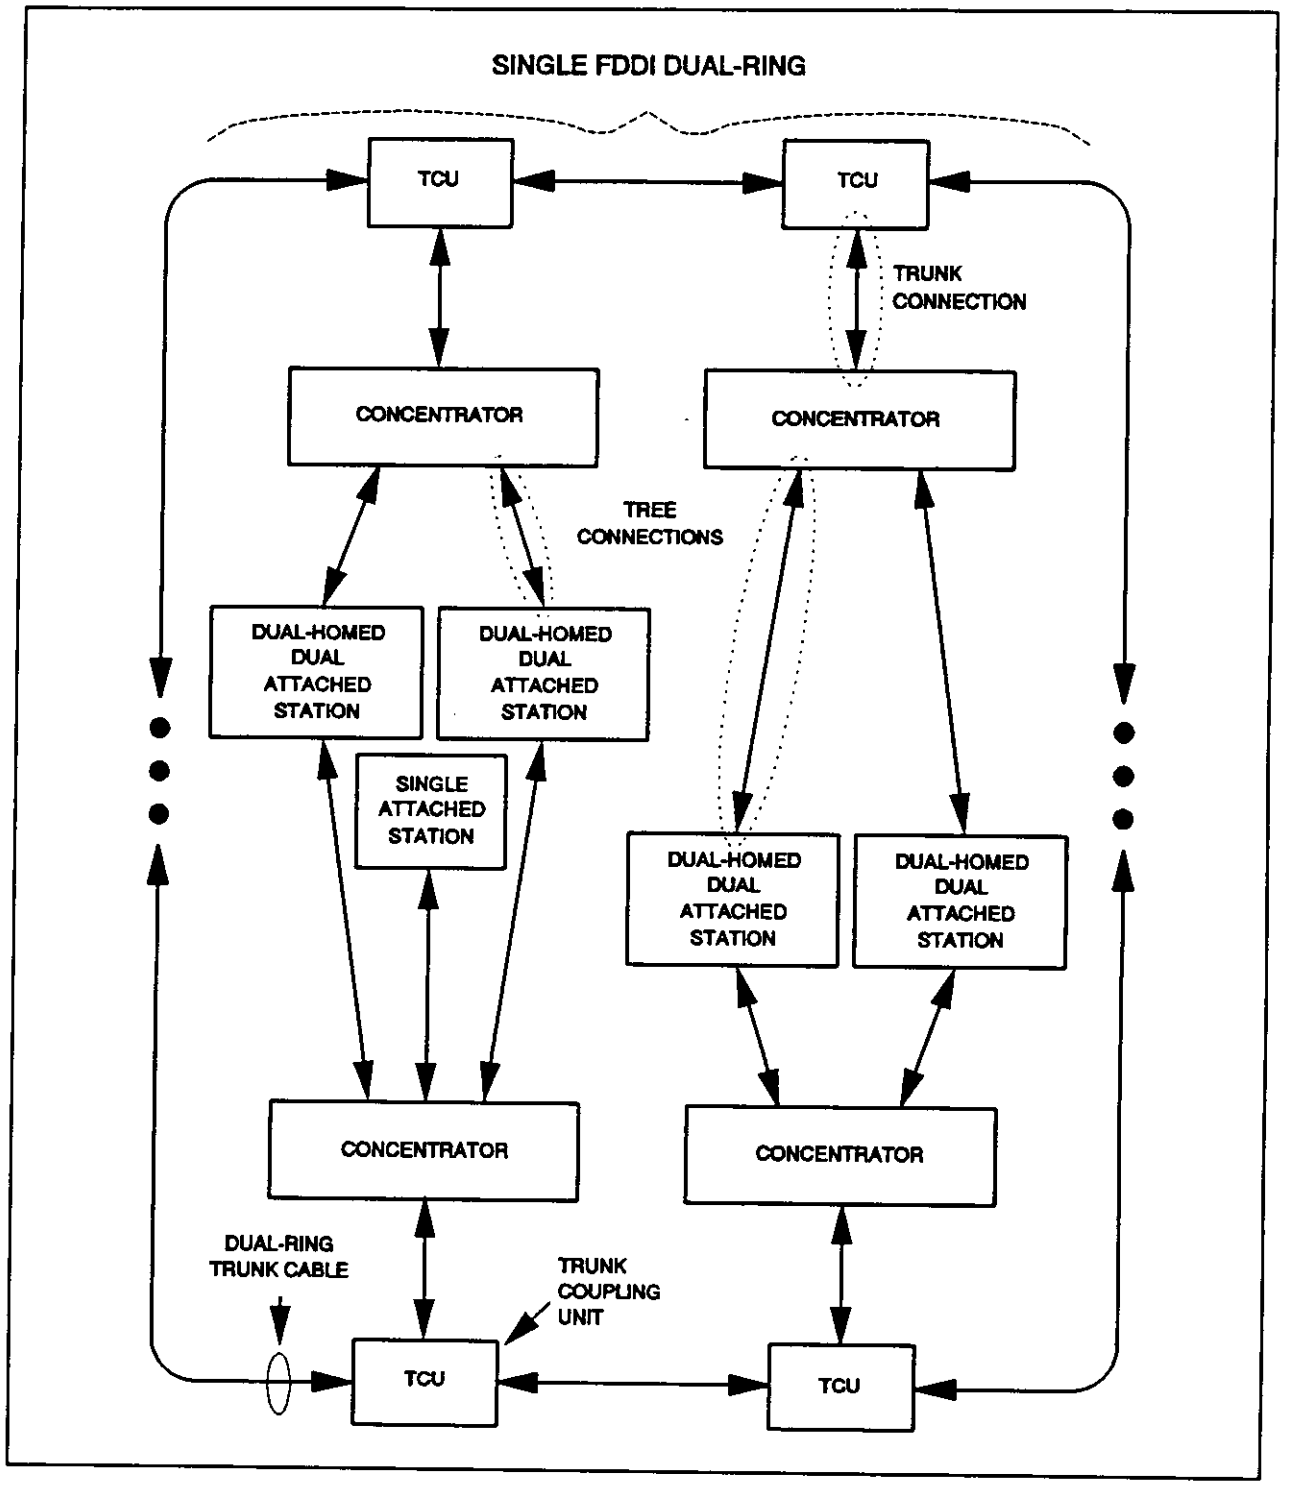

Figure 14. Dual Homed Topology

attached should be located apart from each other to avoid simultaneous damage.

10.3.1.3.2 Concentrators. Each concentrator used in the dual homed topology is attached to the dual-ring, through a trunk connection, by a trunk coupling unit (TCU), or dual homed to concentrators through tree connections. Thus each concentrator uses one trunk connection and one or more tree connections, or only tree connections, depending on the number of attached stations being supported and its position in the topology.

10.3.1.3.3 Single attachment stations (SASs). SAFENET allows the use of SASs. If the function of a station is not mission-critical or is duplicated elsewhere in the network it may be desirable to connect it as a SAS. Connection of a SAS is to a concentrator through a tree connection. Figure 14 illustrates the connection of SASs in this reconfiguration approach.

It is important to note that a SAS, though allowed in the dual homed topology, is not itself a dual homed station. The concentrator to which the SAS is attached, the SAS itself, and the corresponding tree connection are single points of failure, and damage to any one will result in loss of communication to the station.

10.3.2 Multiple subnetwork topology. An example multiple subnetwork physical topology is shown in Figure 15 for a two subnetwork implementation. As shown, this topology is based on the use of two or more FDDI dual-rings, with the primary and secondary rings of each dual-ring cabled together to form a set of two or more FDDI dual-ring trunks. The multiple subnetwork topology may use any combination of the SAFENET network attachment types. The attachment technique used for the independent networks may be the same (as shown in Figure 15), or they may use different techniques. It is acceptable to connect some stations in a multiple subnetwork topology to only one of the subnetworks.

10.3.2.1 Multiple subnetwork reconfiguration. The multiple subnetwork reconfiguration method employs the dynamic routing capabilities of the network layer protocols defined in SAFENET. This enables network stations to route their communications dynamically among two or more independent FDDI LANs (subnetworks) to which they are attached. Implementation of the multiple subnetwork reconfiguration method does not require any particular protocol extensions or mechanisms beyond those provided by the network layer protocols required by SAFENET. A description of the multiple subnetwork mechanism is provided in the following paragraphs.

10.3.2.2 Multiple subnetwork mechanism. The multiple subnetwork mechanism is based on the requirement that the network stations implement more than one complete FDDI station. Within a given multiple dual attachment network station or multiple single attachment network station there will reside at least two complete and independent DAS/SAS entities. Each of these DAS/SAS entities will include its own SMT entity, and it will be associated with its own LLC protocol entity. Every DAS/SAS entity within the station will be associated with the same set of network layer protocols through their associated LLC entities.

This arrangement provides the network layer of a multiple DAS or multiple SAS with a direct interface to two or more FDDI subnetworks. The network layer protocols will therefore be able to select, through their dynamic routing capabilities, which of these subnetworks to use for each instance of data transfer. The result is that in the event of a failure of a particular subnetwork, a multiple DAS or multiple SAS will already have at least one other subnetwork available for its communications. Note that multiple SASs connect via concentrators.

10.3.2.3 Multiple subnetwork components. The components of the multiple subnetwork topology are multiple network stations and concentrators. See MIL-STD-2204 for complete definitions of these components and methods of interconnecting them.

10.3.2.3.1 Multiple network stations. When multiple network stations are used in this topology, each network station is attached to a different FDDI network. Multiple DAS's may be attached to a different network through trunk connections or by a pair of concentrators via tree connections. Multiple SAS's may be connected by concentrators via tree connections.

10.3.2.3.2 Concentrators. When concentrators are used in the multiple subnetwork topology to support the use of multiple DAS's or SASs, they may be attached to a dual-ring, through a trunk connection, by a TCU. Thus each concentrator may use one trunk connection and one or more tree connections, depending on the number of multiple DAS's or SASs being supported.

10.3.2.4 Multiple subnetwork routing. This information is not intended to be the complete definition of SAFENET routing, but is provided here for convenience. Other network designs which do not require the updating of routing tables, or which use other routing schemes, are not intended to be precluded. Complete information on routing is found in MIL-STD-2204.

Access to a subnetwork is gained via the SAFENET Network layer protocol (CLNP) through a subnetwork point of attachment

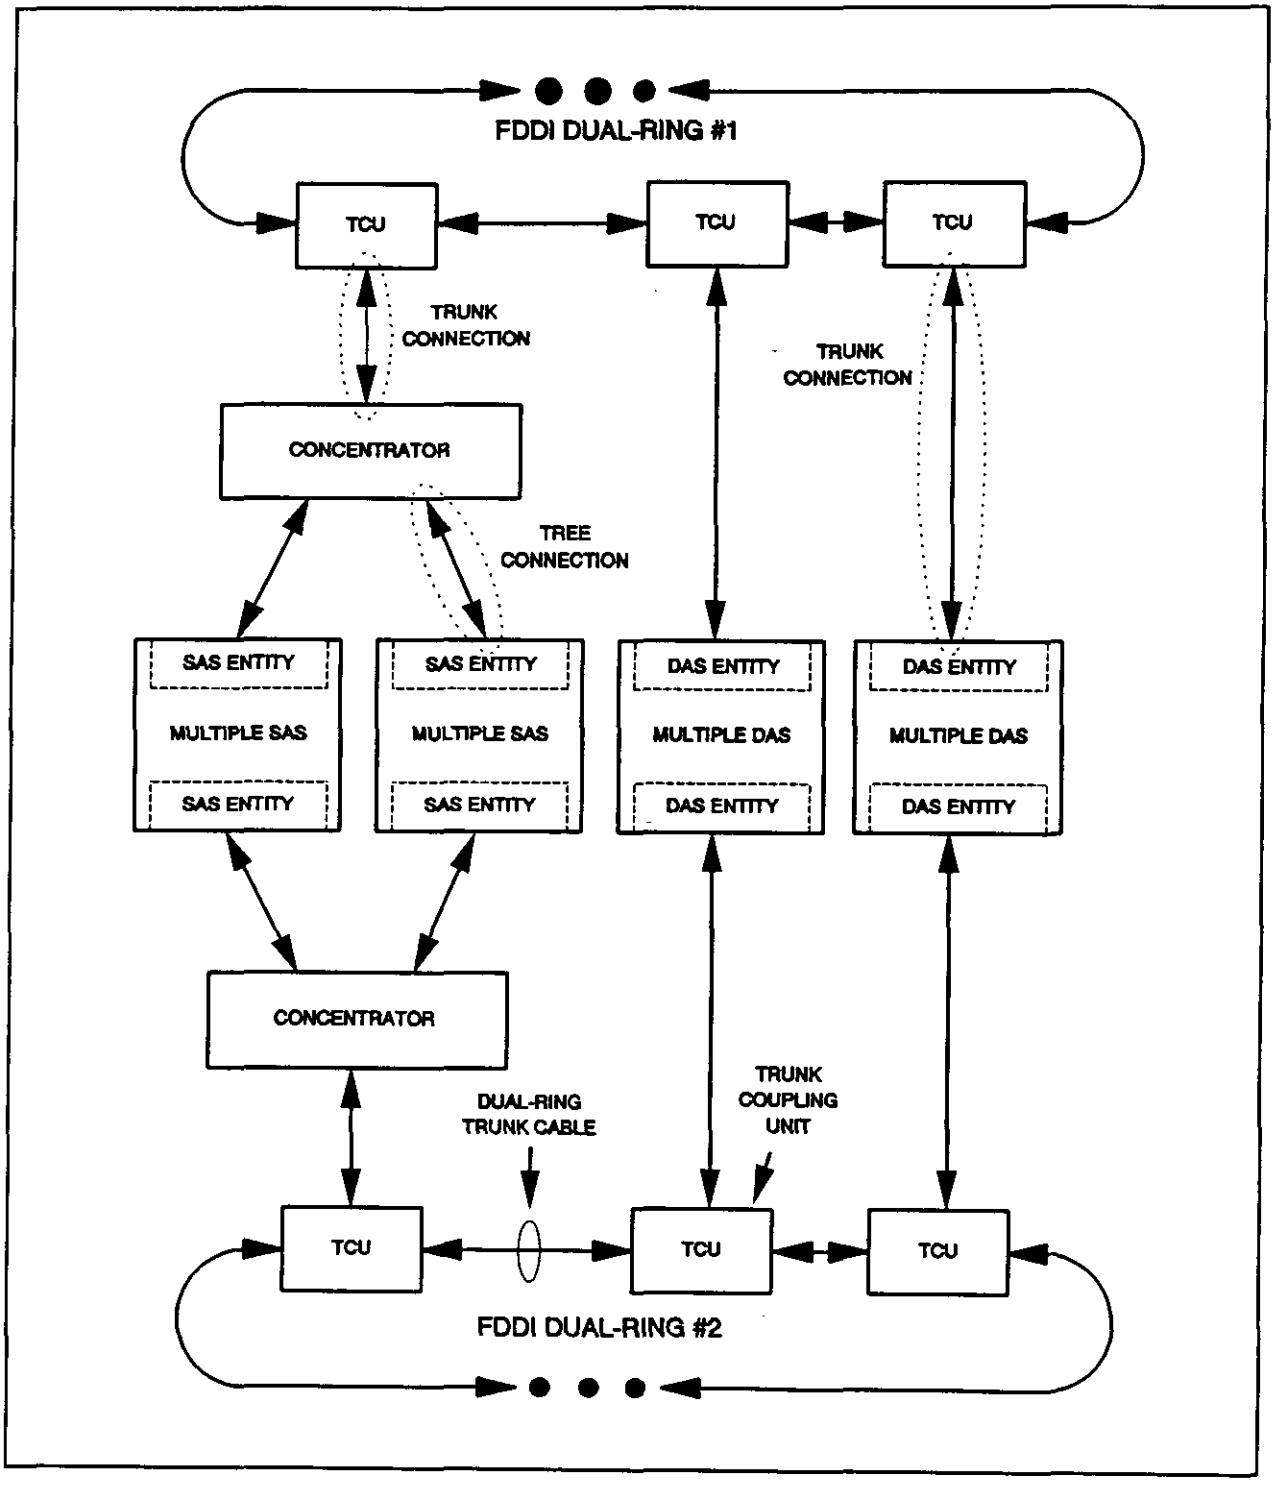

Figure 15. Multiple Subnetwork Topology

(SNPA). Any one subnetwork is totally independent of any other subnetwork and thus does not share any Data Link or Physical components.

Stations attached to multiple subnetworks can be either End Systems or Intermediate Systems. While the multiple subnetworks provide the redundant data paths to support survivability and availability, a means of dynamically changing the network layer routing tables is needed. In SAFENET the ES/IS routing exchange protocol is required and meets the end systems' routing needs. In order to reroute data units dynamically when the underlying FDDI rings undergo segmentation, a number of ISs are required and the IS/IS routing exchange protocol is needed to provide dynamic route determination. This type of reconfiguration capability requires that multiple interconnected ISs survive on all of the isolated segments. System implementors need to design the number and placement of ISs to meet the survivability requirements of the system being developed.

Stations which use multiple subnetwork routing do not provide standardized methods for users to select which subnetwork a particular communication will take. QOS capabilities (if implemented) may affect such a selection. It is permissible for a local means to be supplied by SAFENET product suppliers to influence the selection of specific subnetworks for specific data transfers.

10.3.3 Concentrator tree topology. An example concentrator tree topology is shown in Figure 16 for a four concentrator implementation. As shown, this topology uses the mechanisms of dual-homing for dual attachment stations. Although the devices in Figure 16 are shown as concentrators, they actually may contain the full SAFENET protocol stack and function as stations on the network.

10.3.3.1 Concentrator tree reconfiguration. The concentrator tree topology provides reconfiguration by using the dual homing policy to provide for a segment hop (switching to a backup fiber segment) when faults occur. This is achieved due to the FDDI SMT rules, which state that a single MAC station favors the B port for data transfer and uses the A port only in the event of loss of activity on the B port. Effectively, the concentrators withhold the A port as long as the B port is active.

In figure 16, in the event that a break occurs on one of the cables on the B side of the concentrator, the concentrator which falls directly after the break will switch to the A port. By the nature of the operation of a concentrator, all M ports are placed

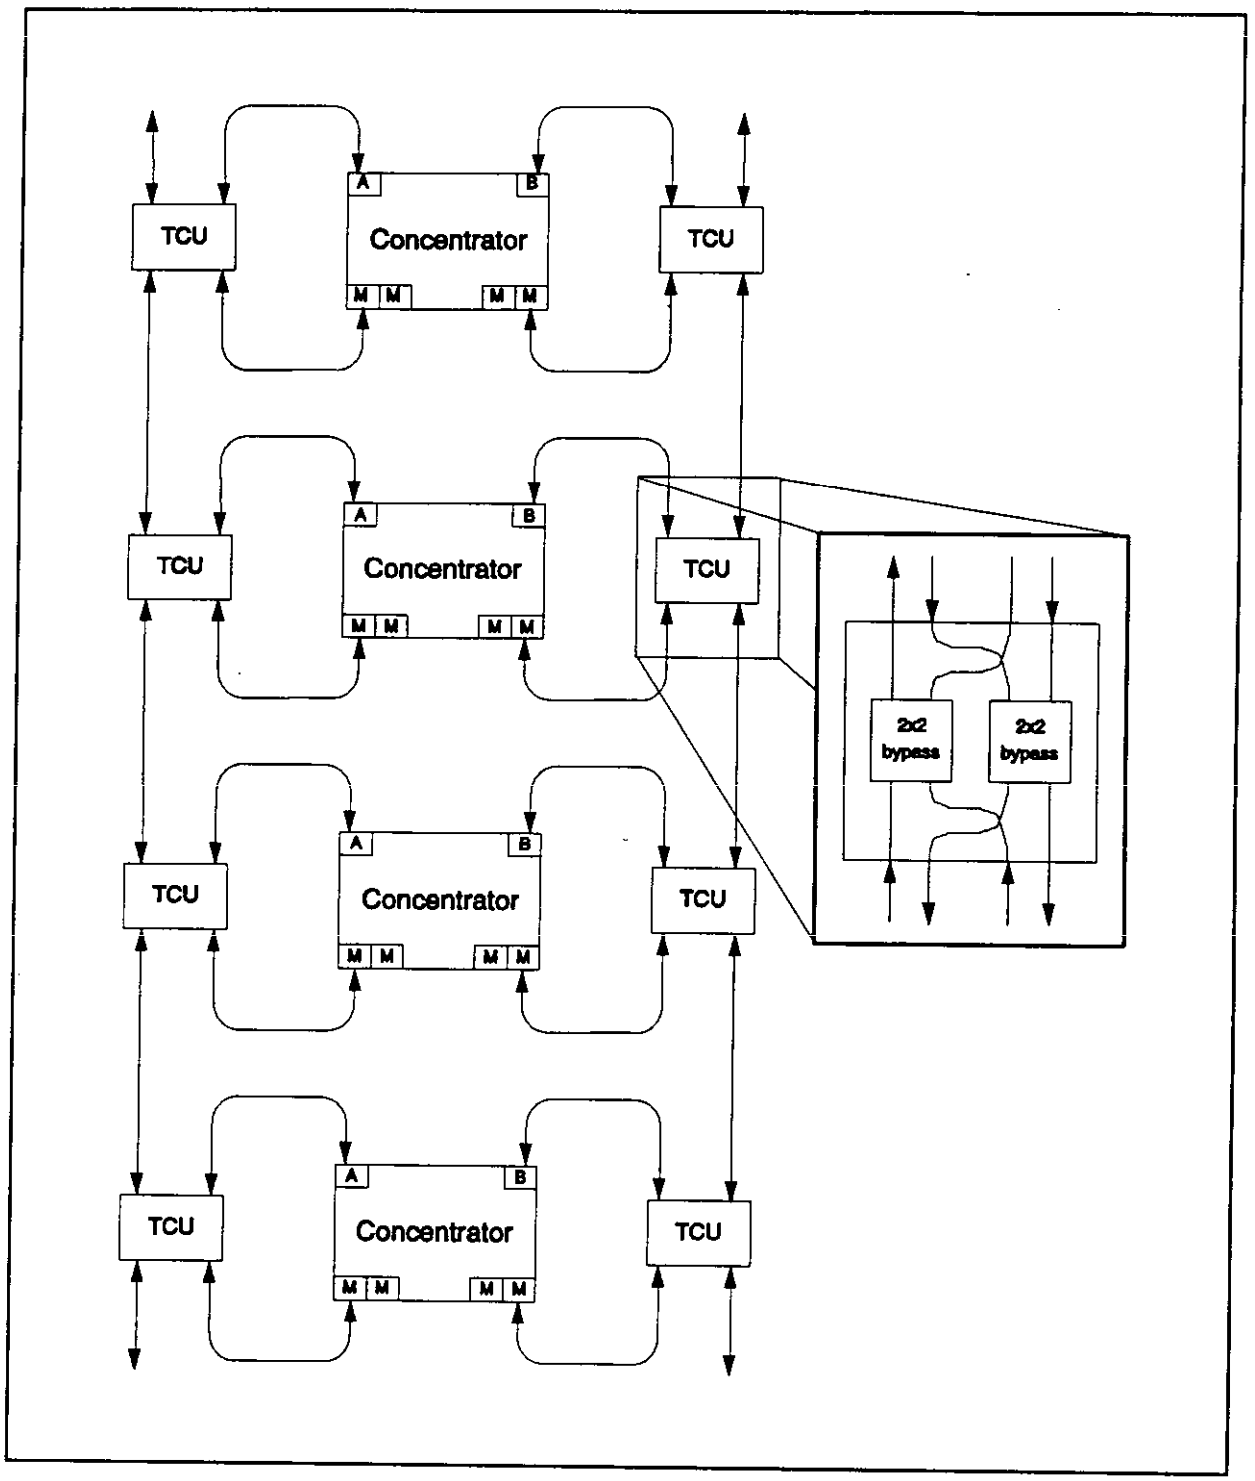

Figure 16. Concentrator Tree Topology

in the active data path and the A port immediately becomes the active input of the concentrator. The other concentrator attachments in the tree remain active through the B ports. The M ports of the affected concentrator are switched to the active A port and data transfer resumes. Additionally, the M port above the break is removed from the data path due to the break and the data path is not broken by the M port.

10.3.3.3 Concentrator tree components. The components of the concentrator tree topology are concentrators and TCUs. See MIL-STD-2204 for definition of a concentrator and methods of interconnection.

The concentrators are connected through TCUs (as shown in figure 16). Tree connections are used to connect the ports of the concentrator to the TCU and an electrical interface cable is used to control the actions of the TCU. Note that the TCU is the exact same device as used on the trunk ring, but the tree connection to the TCU is different than that of a standard station connection.

11. SAFENET SECURITY GUIDANCE

11.1 Introduction. This section provides guidance for specifying and developing implementations of SAFENET profiles which provide added support in achieving data security. An approach is provided for identifying the necessary security services and allocating the services to the SAFENET protocols and management elements. Implementation alternatives are also provided.

11.2 Applicable Documents. The following is a partial listing of applicable documents that may be of interest to those implementing a system with security needs.

11.2.1 Government Documents.

11.2.1.1 Specifications, Standards, and Handbooks.

DOD 5200.28-STD	Department of Defense Trusted Computer System Evaluation Criteria. December 1985.
CSC-STD-002-85	Department of Defense Password Management Guideline. 12 April 1985.
CSC-STD-003-85	Computer Security Requirements - Guidance for Applying the Department of Defense Trusted Computer System Evaluation Criteria in Specific Environments. 25 June 1985.
CSC-STD-004-85	Technical Rational Behind CSC-STD-003-85 Computer Security Requirements - Guidance for Applying the Department of Defense Trusted Computer System Evaluation Criteria in Specific Environments. 25 June 1985.

(Applications for copies should be addressed to the INFOSEC Awareness Division, ATTN:IAOC, Ft. George G. Meade, MD 20755-6000.)

DOD 5200.28	Security Requirements for Automated Information Systems (AISs). March 1988.

MIL-STD-2204	Survivable Adaptable Fiber Optic Embedded Network.

(Unless otherwise indicated, copies of federal and military specifications, standards, and handbooks are available from the Standardization Documents Order Desk, Bldg. 4D, 700 Robbins Ave., Philadelphia, PA 19111.)

11.2.1.2 Other Government Documents, Drawings, and Publications.

NATIONAL COMPUTER SECURITY CENTER

NCSC-TG-001	A Guide to Understanding Audit in Trusted Systems. 1 June 1988.
NCSC-TG-002	Trusted Product Evaluations, A Guide for Vendors. 22 June 1990.
NCSC-TG-003	A Guide to Understanding Discretionary Access Control in Trusted Systems. 30 September 1987.
NCSC-TG-004	Glossary of Computer Security Terms. 21 October 1988.
NCSC-TG-005	Trusted Network Interpretation of The Trusted Computer System Evaluation Criteria. 31 July 1987.
NCSC-TG-006	A Guide to Understanding Configuration Management in Trusted Systems. 28 March 1988.
NCSC-TG-007	A Guide to Understanding Design Documentation in Trusted Systems. 2 October 1988.
NCSC-TG-008	A Guide to Understanding Trusted Distribution in Trusted Systems. 15 December 1988.
NCSC-TG-009	Computer Security Subsystem Interpretation of the Trusted Computer System Evaluation Criteria. 16 September 1988.

NCSC-TG-011 Trusted Network Interpretation Environments Guideline -- Guidance For Applying the Trusted Network Interpretation. 1 August 1990.

NCSC-TG-013 Rating Maintenance Phase Program Document. 23 June 1989.

NCSC-TG-014 Guidelines for Formal Verification Systems. 1 April 1989.

NCSC-TG-015 A Guide to Understanding Trusted Facility Management. 18 October 1989.

NCSC-TG-017 A Guide to Understanding Identification and Authentication in Trusted Systems. September 1991.

NCSC-TG-019 Trusted Product Evaluation Questionnaire. 16 October 1989.

NCSC-TG-025 A Guide to Understanding Data Remanence in Automated Information Systems. September 1991.

NCSC-TG-026 A Guide to Writing the Security Features User's Guide for Trusted Systems. September 1991.

C Technical Report 79-91 Integrity in Automated Information Systems. September 1991.

(Applications for copies should be addressed to the INFOSEC Awareness Division, ATTN:IAOC, Ft. George G. Meade, MD 20755-6000.)

DEPARTMENT OF THE NAVY

OPNAVINST 5239.1A Department of the Navy Automated Data Processing Security Program. August 1982.

(Applications for copies should be addressed to the National Technical Information Service, 5285 Port Royal Rd., Springfield, VA 22161.)

SECNAVINST 5239.2 — Department of the Navy Automated Information System (AIS) Security Program (Draft IRM/C1). 15 November 1989.

(Applications for copies should be addressed to the Navy Printing Office, Standardization Documents Order Desk, Bldg. 4D, 700 Robbins Ave., Philadelphia, PA 19111.)

Computer Security Guidebook for Mission-Critical Computer Resources Managed Under the Research, Development, and Acquisition Process. 28 October 1990.

(Applications for copies should be addressed to the Space and Naval Warfare Systems Command, Warfare Systems Engineering Policy Division (SPAWAR 2241), Washington, DC 20363-5100.)

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY

SDN.301 Rev 1.5 — Secure Data Network System - Security Protocol 3. May 15, 1989.

SDN.401 Rev 1.3 — Secure Data Network System - Security Protocol 4. May 2, 1989.

SDN.801 Rev 1.3 — Secure Data Network System - Access Control Concept. July 26, 1989.

SDN.802 Rev 1.0 — Secure Data Network System - Access Control Specification. July 25, 1989.

SDN.802/1 — Secure Data Network System - Access Control Specification - Access Control Information Specification Addendum 1. July 25, 1989.

SDN.902 Rev 3.2 — Secure Data Network System - Key Management Protocol - Definition of Services Provided by the Key Management Application Services Element. August 1, 1989.

SDN.903 Rev 3.2 — Secure Data Network System - Key Management Protocol - Specification of the Protocol for Services Provided by the Key Management

Application Services Element. August 1, 1989.

(Applications for copies should be addressed to the National Technical Information Service, 5285 Port Royal Rd., Springfield, VA 22161.)

GOSIP 2.0 - U. S. Government Open Systems Interconnection Profile (GOSIP), Federal Information Processing Standards Publication (FIPS PUB) 146-1, 3 Apr 1991

(Applications for copies should be addressed to the U. S. Department of Commerce, National Technical Information Service (NTIS), 5285 Port Royal Road, Springfield, VA 22161)

11.2.2 Non-Government Publications.

INTERNATIONAL ORGANIZATION FOR STANDARDIZATION (ISO)

ISO 7498 Information Processing Systems - Open Systems Interconnection - Basic Reference Model. October 15, 1984.

ISO 7498-2 Information Processing Systems - OSI Reference Model - Part 2: Security Architecture. July 19, 1988.

ISO 9594-7 Information Technology - Open Systems Interconnection - The Directory - Part 7: Authentication Framework. 1990 (e).

ISO DIS 10164-7 Information Technology - Open Systems Interconnection - Systems Management - Part 7: Security Alarm Reporting. July 1991.

ISO DIS 10164-8 Information Technology - Open Systems Interconnection - Systems Management - Part 8: Security Audit Trail Function. June 1991.

ISO/IEC JTC1/SC21N6693 Information Technology - Open Systems Interconnection - Security Frameworks in Open Systems - Part 1: Overview. November 1991.

ISO CD 10181-2	Information Technology - Open Systems Interconnection - Security Frameworks in Open Systems - Part 2: Authentication. May 13, 1991.
ISO CD 10181-3	Information Technology - Open Systems Interconnection - Security Frameworks in Open Systems - Part 3: Access Control. June 24, 1991.
ISO/IEC JTC1/SC21N6692	Information Technology - Open Systems Interconnection - Security Frameworks in Open Systems - Part 4: Non-Repudiation. November 1991.
ISO/IEC JTC1/SC21N6690	Information Technology - Open Systems Interconnection - Security Frameworks in Open Systems - Part 5: Integrity. November 1991.
ISO/IEC JTC1/SC21N6691	Information Technology - Open Systems Interconnection - Security Frameworks in Open Systems - Part 6: Confidentiality. November 1991.
ISO/IEC JTC1/SC21N6169	Information Technology - Open Systems Interconnection - Security Frameworks in Open Systems - Part 7: Audit. May 1991.
ISO DIS 10736	Information Technology - Telecommunications and information exchange between systems - Open Systems Interconnection - Transport Layer Security Protocol. October 11, 1991.
ISO DIS 10736 PDAM 1	Information Technology - Telecommunications and information exchange between systems - Open Systems Interconnection - Transport Layer Security Protocol - Amendment 1: Security Association Establishment. December 1991.
ISO CD 10745	Information Technology - Open Systems Interconnection - Upper

	Layers Security Model. June 28, 1991.
ISO CD 11577	Information Technology - Telecommunications and Information Exchange Between Systems - Open Systems Interconnection - Network Layer Security Protocol. November 13, 1991.
ISO/IEC JTC1/SC21N6587	Second Working Draft - Security Exchange Overview and Specification Framework (GULS Part 1). December 5, 1991.
ISO/IEC JTC1/SC21N6588	Second Working Draft - Security Exchange Service Element Service Definition (GULS Part 2). December 5, 1991.
ISO/IEC JTC1/SC21N6589	Second Working Draft - Security Exchange Service Element Protocol Specification (GULS Part 3). December 6, 1991.
ISO/IEC JTC1/SC21N6590	Second Working Draft - Security Transfer Syntax, Encoding Procedures, and Syntax Notation (GULS Part 5). December 5, 1991.

(Applications for copies of International Organization for Standardization (ISO) documents should be addressed to the American National Standards Institute, 1430 Broadway, New York, NY 10018)

AMERICAN NATIONAL STANDARDS INSTITUTE

X3T4/92-007	Nonrepudiation using Symmetric Encipherment Algorithms. December 13, 1991.

(Applications for copies of American National Standards Institute (ANSI) documents should be addressed to the American National Standards Institute, 1430 Broadway, New York, NY 10018)

INSTITUTE OF ELECTRICAL AND ELECTRONIC ENGINEERS

P802.10B/D6	IEEE Standard for Interoperable Local Area Network Security Part B

- Secure Data Exchange. November 16, 1990.

P802.10C — IEEE Standard for Interoperable Local Area Network Security Part C - Key Management Proposal. May 21, 1991.

(Applications for copies of Institute of Electrical and Electronic Engineers (IEEE) documents should be addressed to the Institute of Electrical and Electronic Engineers, 445 Hoes lane, Piscataway, NJ 08854-4150)

11.3 Acronyms. The following acronyms are used in this section of the handbook.

AIS	Automated Information System
ANSI	American National Standards Institute
CD	Committee Draft
CLNP	Connectionless Network Protocol
COMSEC	Communications Security
DIS	Draft International Standard
DoD	Department of Defense
DOS	Denial of Service
ES	End System
FDDI	Fiber Optic Distributed Data Interface
GOSIP	Government OSI Profile
GULS	Generic Upper Layer Security
IBAC	Identity Based Access Control
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IS	Intermediate System
ISO	International Organization for Standardization
JTC1	Joint Technical Committee 1
KMP	Key Management Protocol
LAN	Local Area Network
NLSP	Network Layer Security Protocol
NSA	National Security Agency
OSI	Open Systems Interconnection
PDU	Protocol Data Unit
QOS	Quality of Service
RBAC	Rule Based Access Control
SAFENET	Survivable, Adaptable, Fiber Optic, Embedded Network
SDNS	Secure Data Network System
SDU	Secure Data Unit
SILS	Standard for Interoperable LAN Security

SP3	Security Protocol 3
SP4	Security Protocol 4
TLSP	Transport Layer Security Protocol
TNI	Trusted Network Interpretation
TNIEG	Trusted Network Interpretation Environments Guidelines
TP4	Transport Protocol Class 4
WD	Working Draft
XTP	Xpress Transfer Protocol

11.4 Background. The Survivable Adaptable Fiber Optic Embedded Network (SAFENET) set of communication services offers limited support in meeting the security needs of a system that must process information subject to Department of Defense (DoD) regulations concerning data security. Capabilities of published national and international standards for communication services applicable to the provision of security services are currently limited. A secure implementation of a system that uses SAFENET will require security issues to be addressed at the SAFENET interfaces and at the interfacing elements (e.g., operating system, database management system, network management).

As with the other sections of the handbook, this section is informative only. To ensure a secure SAFENET implementation, specific language detailing the security requirements should be included in procurement specifications. This section provides guidance for specifying the necessary security requirements (i.e., security services) for SAFENET implementations.

In addition to providing technical support in meeting the security needs of a system that uses the SAFENET protocols, a network may require certification/accreditation by appropriate authority before being deployed. Depending on the security policy in effect for a particular SAFENET installation and the ownership of the network and connecting systems, this may be done as part of the security certification/accreditation of the system or it may be done separately.

11.5 Security design process. This section describes a five step process that leads to specification of all necessary security services in a system security design while maximizing the use of standards and products. The five steps - policy development, risk assessment, security service selection/allocation, security standards and products selection, additional security mechanism selection (to fill any residual security holes) - are described in detail below.

11.5.1. Policy development. The first step in the SAFENET security design process is the development of a system security

policy based on the operational concept for the system, high level threat knowledge, and guiding DoD security policy. The system security policy should address types of data to be processed by the network, levels of each type of information and how it should be protected, and network connectivity. System security policy must address:

a. information and equipment protection policy (e.g. confidentiality, integrity)
b. human and equipment/process access control
c. human authentication
d. system availability
e. assurance levels
f. system interface policy.

System security policy must also identify the system users, their clearance levels and the type of access each type of user will be allowed. Navy security policy is defined in OPNAVINST 5239.1A and SECNAVINST 5239.2.

11.5.2 Risk assessment. The risk assessment process spans the entire system life cycle. Risk assessments are conducted to determine the degree of risk associated with threats and vulnerabilities in the system during the specification, design, implementation, and installation/operation phases. The following items should be considered when conducting a risk assessment.

11.5.2.1 Threats. A threat is any circumstance or event with the potential to cause harm to a system in the form of destruction, disclosure, or modification of information or equipment, or denial of service. A description of some possible threats to systems that will be SAFENET-compliant is provided in paragraph 11.4.1.

11.5.2.2 Existing security countermeasures. A security countermeasure is anything that reduces the vulnerability of a system to known threats. When designing a system that will include SAFENET, it is important to identify any known countermeasures inherent in the system concept or operating environment. Countermeasures may be provided through physical mechanisms, host computer systems, or particular products to be used in the system.

11.5.2.3 Vulnerabilities and vulnerability points. A vulnerability is a weakness in the system that could be exploited by a threat. A vulnerability point is the system component where a vulnerability exists. Each vulnerability and vulnerability point within the system should be identified and documented.

11.5.2.4 Threat scenarios. A threat scenario is a postulated series of events which collectively result in subversion of the system security. Applicable threat scenarios should be defined to illustrate the perpetration of successions of threats at multiple vulnerability points that result in unauthorized disclosure, modification or destruction of information or equipment, or denial of service.

11.5.2.5 Evaluation of risk. Risk is the potential for successful exploitation of a vulnerability, given the identified and proposed countermeasures. The damage that may be associated with system risks include unauthorized disclosure of information, unauthorized modification or destruction of information or equipment, and denial of system services. The potential risk of each vulnerability being exploited should be determined from an analysis of existing countermeasures identified and their effectiveness in countering the vulnerability. The degree of risk (i.e., a combination of the sensitivity of the vulnerable information, equipment, or communication service and the probability of exploitation) associated with a particular threat/vulnerability scenario should be divided into three categories:

a. Low. The threat scenario is considered to be very unlikely to occur, to have a low impact if it does occur, or to be completely controlled by existing countermeasures.

b. Medium. The threat scenario is considered to have a moderate likelihood of occurrence, to have a moderate impact if it does occur, or to be partially controlled by existing countermeasures.

c. High. The threat scenario is considered to have a high likelihood of occurrence, to have a significant impact if it does occur, or not to be controlled by existing countermeasures.

11.5.3 Security service selection/allocation. Once the risk assessment has been completed, the appropriate security services may be chosen and allocated to system elements. The services should be chosen and allocated to reduce the risk of system vulnerabilities, to supplement existing countermeasures, and to optimize the performance and cost objectives for the system. A description of security services is provided in paragraph 11.4.2.

11.5.4 Security standards/products selection. Following identification of desired security services and their allocation

to system elements, appropriate standards or products must be chosen to provide these security services. SAFENET encourages use of products that conform to national or international standards rather than proprietary products. A discussion of available security standards is provided in paragraph 11.5.

11.5.5 Additional security mechanism selection. In some cases, the existing security standards and products will not be sufficient to provide all required security services. In these cases, it becomes necessary to define system-specific security mechanisms to "fill the holes". Any system-specific security mechanisms should be presented to the appropriate standards organization to encourage standardization of the mechanism.

11.6 Threats and services. The following paragraphs identify the threats that may exist in a SAFENET system. Following the discussion of threats is a discussion of security services in paragraph 11.4.2. The security services are then matched with the threats in paragraph 11.4.3.

11.6.1 Threats. The following paragraphs outline threats that should be considered when conducting a risk assessment.

11.6.1.1 System access. Threats to system access include attacks on the system from an intruder with access to the application software or operating system. This assumes an intruder has access to some form of user-machine interface in the system.

11.6.1.1.1 Browsing. Browsing includes attempts by a user or intruder to gain unauthorized access to network information.

11.6.1.1.2 Misuse. Misuse is the use of processing or communication services for other than official or authorized purposes. Misuse includes both inadvertent and intentional execution of malicious functions, performance of undesirable functions, and errors of commission, omission, and oversight.

11.6.1.1.3 Penetration. System penetration includes attacks by unauthorized persons attempting to gain system access by defeating the system security perimeter. Penetration could lead to the realization of other threats (e.g., browsing, misuse).

11.6.1.2 Communication access. The communication access threat includes passive and active attacks on information transferred over communication channels. The definition of this threat assumes an intruder has access to communication media or components.

11.6.1.2.1 Eavesdropping. Eavesdropping is the passive surveillance of communication channels to gain illicit access to transmitted information. It is perpetrated through physical, electrical, and radio frequency taps into the communication channel.

11.6.1.2.2 Traffic analysis. Traffic analysis is the passive collection, analysis, and interpretation of communication patterns to infer operational and logistics-related information. It is perpetrated through eavesdropping.

11.6.1.2.3 Emanations attacks. Emanations attacks are perpetrated through passive collection, analysis, and interpretation of electromagnetic emissions from electrical and electronic equipment.

11.6.1.2.4 Spurious communication. Spurious communication includes intruder attempts to establish communication connections or associations using a false identity through replay of legitimate initiation sequences or communication messages.

11.6.1.2.5 Message stream modification. Message stream modification includes attempts to modify, delete, or insert information into a message stream while it is in transit over a communication channel.

11.6.1.2.6 Denial of service attacks. Denial of Service (DOS) attacks prevent or delay the performance of legitimate or critical communication services and includes the flooding of a communication path with unauthorized messages.

11.6.1.3 Tampering. Tampering includes attacks on physical and logical components in the network including substituting rogue components in place of legitimate components, physical theft of components, damage to components, and electrical probing within the system. This assumes that an intruder has physical or electrical access to the system components and their internal structures.

11.6.1.3.1 Damage. Damage attacks against system components destroy them or render them unable to perform their intended mission.

11.6.1.3.2 Substitution. Substitution is the introduction of unauthorized and potentially malicious components to intercept communications, generate incorrect or misleading information, masquerade as legitimate components, or perform other undesirable functions.

11.6.1.3.3 Theft. Theft is the unauthorized removal of system components.

11.6.1.3.4 Probing. Probing attacks on internal system components either obtain sensitive information contained within the components or explore the details of sensitive technology.

11.6.2 Security services. Security services are applied to a system to prevent or reduce a system's vulnerability to threats and the resulting damage to the system or its information. Security services are categorized into access control, accountability, confidentiality, data integrity, and service availability.

11.6.2.1 Access control. Access control is the prevention of unauthorized use of a resource, including the prevention of use of a resource in an unauthorized manner.

11.6.2.1.1 Identity based access control. Identity Based Access Control (IBAC) limits access of specific users to system resources based on the user's identity and need-to-know.

11.6.2.1.2 Rule based access control. Rule Based Access Control (RBAC) limits the access of users to system resources based on the user possessing specific privileges (e.g. user clearance) that comply with the corresponding system access rules.

11.6.2.1.3 Labeling. Labeling is a service that supports access control. Labels are markings bound to information, system resources, or users that name or designate its security attributes.

11.6.2.2 Accountability. Accountability is the service that enables security-relevant activities on a system to be traced to individuals who may then be held responsible for their actions.

11.6.2.2.1 Authentication. This process establishes the validity of a claimed identity. The are two parts to authentication are described below.

11.6.2.2.1.1 Peer-entity authentication. This corroborates that a peer entity in an association is the one claimed. This service, when provided by the (N)-layer, provides corroboration to the (N+1)-entity that the peer entity is the claimed (N+1)-entity.

11.6.2.2.1.2 Data origin authentication. This corroborates that information transmitted over the network has been transmitted by the claimed source.

11.6.2.2.2 Audit. This service maintains system security-relevant activity information in audit logs and provides outputs, both in report and alarm formats, depending on the criticality of the information.

11.6.2.3 Confidentiality. The confidentiality service ensures information is not made available or disclosed to unauthorized individuals, entities, or processes.

11.6.2.3.1 Connection-oriented confidentiality. This service protects all (N)-service data units from unauthorized disclosure during communication between (N+1) peer entities, where a security association is established for the transfer of data and for the application of confidentiality service between the entities themselves, and between each entity and the physical layer.

11.6.2.3.2 Connectionless confidentiality. This service protects (N)-service data units from unauthorized disclosure during transmission from one (N+1)-entity to one or more (N+1)-entities, where each entity has an association with the physical layer, and no association is established for the transmission of data or for the application of the confidentiality service between the layer peer-entities themselves.

11.6.2.3.4 Traffic flow confidentiality. This service protects information which might be derived from the observation of traffic flows by hiding information or providing misleading information.

11.6.2.4 Data integrity. Data integrity is the property that verifies data has not been altered or destroyed in an unauthorized, undetected manner.

11.6.2.4.1 Connection-oriented integrity. This service provides for the integrity of all (N)-service data on a security association and detects any modification, insertion, deletion, or replay of any data within an entire Secure Data Unit (SDU) sequence.

11.6.2.4.2 Connectionless integrity. This service provides for the integrity of a single SDU. Normally, connectionless integrity determines whether or not the received SDU has been modified.

11.6.2.4.3 Non-repudiation. This process provides unforgeable proof of information receipt or origin.

11.6.2.4.3.1 Non-repudiation of origin. This process provides unforgeable proof of information origin to either the recipient or a third party.

11.6.2.4.3.2 Non-repudiation of receipt. This process provides unforgeable proof of information receipt to either the originator or a third party.

11.6.2.5 Service availability. Service availability is the ability of the system to detect, recover from, and resist DOS conditions.

11.6.2.5.1 Protocol based denial of service protection. Protocol based DOS protection is the ability of communication protocols to identify or support the identification of a DOS condition.

11.6.2.5.2 Continuity of operations. Continuity of operations is the ability of a system to recover from DOS conditions, and to continue operations following the detection of a DOS condition.

11.6.2.5.3 Physical security. Physical security, while beyond the scope of the SAFENET standard, can be used to provide protection against a number of threats identified in paragraph 11.4.1. Physical security includes such things as guards, combination locks, protected spaces, and TEMPEST protection.

11.6.3 Relating services to threats. Table I illustrates the correlation between threats and the security services usually applied to reduce the vulnerability of a system to these threats.

11.7 Standards. The ISO security architecture framework does not specify security protocols for the implementation of security services. However, for open systems to make use of security services, the format of security information in Protocol Data Units (PDUs) of existing open system standard protocols, as well as any security protocols, will have to be standardized. The following sections provide information on existing and proposed standards for security services and protocols.

The following sections include a number of references to the Secure Data Network System (SDNS) and two protocols from ISO, the Transport Layer Security Protocol (TLSP) and the Network Layer Security Protocol (NLSP). The SDNS project was initiated by the National Security Agency (NSA) to investigate methods of

Table I. Correlation Between Threats and Security Services

Threat	Services to Mitigate Threat
System Access	
System Browsing	IBAC, RBAC, Labeling
System Misuse	Audit, IBAC, RBAC, Labeling
System Penetration	Authentication
Communication Access	
Eavesdropping	Connection-Orientated Confidentiality, Connectionless Confidentiality
Traffic Analysis	Traffic Flow Confidentiality
Emanations	Physical Security (TEMPEST)
Spurious Communications	Authentication, Connection-Oriented Integrity, Connectionless Integrity
Message Stream Modification	Non-Repudiation, Connection-Oriented Integrity
Denial of Service Attack	Continuity of Operations, Protocol Based DOS Protection
Tampering	
Damage	Physical Security
Substitution	Authentication, Physical Security
Theft	Physical Security
Probing	IBAC, RBAC, Authentication

implementing security in distributed computer networks for U.S. Government and private sector users. However, some of the SDNS specifications are not complete nor totally consistent with a number of other security projects in the national and international standards arena. As a result, the SDNS specifications are subject to modification for various reasons as they progress through the standards process. For example, ISO NLSP is derived from SDNS Security Protocol - Layer 3 (SP3), but there are differences between the two protocols. The Government Open Systems Interconnection Profile (GOSIP) intends to conform

to ISO Open Systems Interconnection (OSI) protocols to provide interoperability in a heterogeneous environment. Therefore, it is very likely that ISO instead of SDNS security protocols will be incorporated in later versions of GOSIP.

International Standards may take a number of years to go from a new proposal to a final standard in the international community. Standards in the international community require a minimum of four years to reach the international standard stage. There is no maximum amount of time. Draft standards are normally considered stable when they have reached the Draft International Standard (DIS) stage, though this may vary depending on the document and the support for it within the international community. The information in this section is intended to provide information on work that is completed as well as work that is in progress. Thus, Working Drafts (WD) and Committee Drafts (CD) have been included in the discussion. It should be noted that WDs are at least three years and CDs are at least two years from being an international standard.

11.7.1 Access control. Access control is an important consideration in any system. It is necessary to guarantee that only authorized individuals gain access to classified or sensitive information. There is one proposed standard for the use of access control in open systems, the Access Control Framework. This document discusses both IBAC and RBAC.

The SDNS has published a set of access control documents, covering both IBAC and RBAC in the SDNS environment. The following paragraphs discuss these protocols.

Generic Upper Layers Security (GULS) is an evolving security standard that deals with a generic application layer security exchange. GULS could be used as a mechanism for exchange of access control information and labels between two application layer entities.

11.7.1.1 Identity based access control. Most standards of IBAC deal with access to files within a computer system. IBAC could also be used to provide users with the ability to refuse connections with other users or to refuse incoming messages. These capabilities would reside at the application layer. Implementation should be done with care so as to allow important or emergency messages to override the IBAC function.

IBAC may also be provided at the transport layer. IBAC at the transport layer is based on host identities and can be used to limit network connectivity.

11.7.1.2 Rule based access control. SAFENET implementations should allow users access to classified or sensitive information based on clearance, authorization, and need-to-know attributes. In a network environment these concepts need to be applied not only to the local system but also when access is requested to other users, hosts, and channels.

RBAC for user-to-user access is an application layer service. This service is important when a user operating at a higher security level communicates with a user operating at a lower security level, whether or not information at the higher level is potentially transferred.

RBAC for host access and RBAC for channel access can be transport or network layer services. These services provide checks to make sure that a host is authorized to receive the level of information being sent or is authorized to send information to a destination host, and make sure a channel can protect the level of information being transmitted. The following standards can be used, in conjunction with key management and distribution, to provide this service.

11.7.1.2.1 Transport layer security protocol. The TLSP is a layer 4 protocol. By using the security label field of the TLSP protocol, TLSP will verify that the level parameter falls within the set of acceptable security levels for the channel. If not, the data is discarded.

11.7.1.2.2 Network layer security protocol. The NLSP is a layer 3 protocol. By using the security label field of the NLSP protocol, NLSP will verify that the level parameter falls within the set of acceptable security levels for the channel. If not, the data is discarded.

11.7.1.2.3 Security protocol 4. The SDNS Security Protocol 4 (SP4) is a layer 4 protocol. By using the security label field of the SP4 protocol, SP4 will verify that the level parameter falls within the set of acceptable security levels for the channel. If not, the data is discarded.

11.7.1.2.4 Security protocol 3. The SDNS SP3 is a layer 3 protocol. By using the security label field of the SP3 protocol, SP3 will verify that the level parameter falls within the set of acceptable security levels for the channel. If not, the data is discarded.

11.7.1.3 Labeling. Labeling is required to support both IBAC and RBAC. In the SAFENET context, users, hosts, channels, data packets, and files may require labels. Version 2.0 of GOSIP

includes a section on security labeling. This section deals primarily with labeling at the network layer (specifically, through extensions to the Connectionless Network Protocol (CLNP)). If used in conjunction with SAFENET, it would be possible to use these labels at the application, transport and network layers to indicate the security level of all users, hosts, channels, data packets and files. Future versions of GOSIP are expected to expand the security labeling structure.

11.7.2 Accountability. The following paragraphs discuss protocols which provide various accountability services.

11.7.2.1 Authentication. Authentication Framework is the one proposed standard for the use of authentication in open systems. This document discusses all variations of authentication in a network or host system.

The Directory includes an authentication service. This service can be used to authenticate hosts and users on a network. It requires that a repository (either central or distributed) for directory information be maintained. The central repository would include an authentication certificate for each user or host that could be verified during the authentication process. The directory is already part of the SAFENET standard.

As noted in the access control discussion, GULS is an evolving security standard that identifies a generic application layer security exchange. GULS could be used as a mechanism for exchange of authentication certificates between two application layer entities.

11.7.2.2 Audit. Audit Framework is the one proposed standard for the use of audit in open systems. This document discusses the audit mechanism and the type of events that should be audited in an open system.

The Audit Trail Function provides a description of audit event logs and formats to be used in creating an audit trail. This document does not discuss the uses of audit or specific auditable events.

Once the audit trail has been captured, it needs to be protected and analyzed. The audit trail is usually protected from modification or even examination by any system user. A separate System Auditor is usually the only user allowed access to the file. The analysis of the file can be long and difficult. When large systems are involved an automated analysis tool is the only way to provide complete analysis of the file in a realistic time period. The file can get very large even when care is exercised in selecting events to be audited. Any system should

contain the capability to select audited events. The system should allow events to be chosen for individuals, hosts, or the entire system.

11.7.3 Confidentiality. Confidentiality may be provided through physical or encryption oriented methods. Physical methods include protection of the communication media from detection or the disabling of the communication media when a compromise is detected. Specific methods for implementing physical protections are beyond the scope of this document. The following paragraphs will discuss standards associated with encryption oriented methods for providing confidentiality.

The Confidentiality Framework is currently a WD. It provides an overview of confidentiality services and a discussion of the characteristics of encryption algorithms. It does not discuss specific encryption algorithms.

11.7.3.1 Data confidentiality. The following sections discuss the characteristics of existing protocol standards in this area.

11.7.3.1.1 Transport layer security protocol. TLSP is a layer 4 protocol. It is specified as an addition to existing layer 4 protocols (such as Transport Protocol Class 4 (TP4)). It is intended to provide algorithm independent connection oriented and connectionless confidentiality. TLSP can not be used with the Xpress Transfer Protocol (XTP) in the lightweight profile.

11.7.3.1.2 Network layer security protocol. NLSP is a layer 3 protocol. It can be implemented at either the top (above CLNP) or bottom (below CLNP) of the network layer. Since SAFENET requires CLNP, it is necessary for NLSP to be used in the connectionless mode. If NLSP is implemented at the bottom of the network layer, data must be decrypted when in transit through routers since the routing would be done above NLSP. NLSP is also algorithm independent.

11.7.3.1.3 Security protocol - layer 4. SP4 is a layer 4 protocol. It is similar to TLSP and provides both connection oriented and connectionless confidentiality. SP4 is algorithm independent. SP4 can not be used with XTP in the lightweight profile.

11.7.3.1.4 Security protocol - layer 3. SP3 is a layer 3 protocol. SP3 is implemented at the top of the network layer. SP3 is algorithm independent and provides connectionless confidentiality.

11.7.3.1.5 Standard for interoperable local area network security. The Standard for Interoperable Local Area Network (LAN) Security (SILS) is a layer 2 security protocol. It is designed to provide encryption over a LAN and provide connectionless confidentiality. If used in SAFENET, SILS could be implemented as part of either the OSI or lightweight profiles, but all traffic would have to be decrypted as it passed from one SAFENET LAN segment to another.

11.7.3.2 Traffic flow confidentiality. Traffic flow confidentiality has an inherent weakness on any LAN. Since all traffic is essentially broadcast traffic, every station has the ability to monitor traffic flow. Any attempt to prevent this would require extensive modification to the LAN protocol concept.

If traffic flow confidentiality is required on traffic between LAN segments, it is possible to disguise the size of the traffic, but not the fact that traffic is occurring. NLSP provides a traffic flow function that pads out each message to a predetermined length.

If a SAFENET system is to be implemented as a series of interconnected LANs, it is possible to use selective routing to bypass one LAN in favor of a more secure route. These types of implementations could use the existing routing protocols, End System (ES)/Intermediate System(IS) and IS/IS, to do this. Both protocols would be required to make routing decisions based on security parameters. SAFENET does not specify routing based on security.

11.7.4 Integrity. The following paragraphs will discuss standards associated with maintaining data integrity.

11.7.4.1 Data integrity. Data integrity deals specifically with the identification of errors or intentional modifications to the data as it passes across the communication system. While most communication protocols include mechanisms to check for errors, these mechanisms are well known and may be circumvented when an intentional modification is made.

The Integrity Framework is currently a WD. It provides an overview of integrity services. The following sections discuss the characteristics of existing standards in this area.

11.7.4.1.1 Transport layer security protocol. As noted in paragraph 11.5.3.1.1, TLSP is a layer 4 protocol. As a function of the cryptographic algorithm, it can provide connection oriented and connectionless integrity detection. TLSP can not be used with XTP in the lightweight profile.

11.7.4.1.2 Network layer security protocol. As noted in paragraph 11.5.3.1.2, NLSP is a layer 3 protocol. As a function of the cryptographic algorithm, it can provide connectionless integrity detection. If connection oriented integrity is desired, a sequence number can be added to the integrity check via a Quality of Service (QOS) parameter.

11.7.4.1.3 Security protocol - layer 4. As noted in paragraph 11.5.3.1.3, SP4 is a layer 4 protocol. Through its cryptographic algorithm, SP4 can provide connection oriented or connectionless integrity detection. SP4 can not be used with XTP in the lightweight profile.

11.7.4.1.4 Security protocol - layer 3. As noted in paragraph 11.5.3.1.4, SP3 is a layer 3 protocol. Through its cryptographic algorithm, SP3 can provide connectionless integrity detection.

11.7.4.1.5 Standard for interoperable LAN security. As noted in paragraph 11.5.3.1.5, SILS is a layer 2 security protocol. Through its cryptographic algorithm, SILS can provide connectionless integrity detection. If used in SAFENET, SILS could be implemented as part of either the OSI or lightweight profiles.

11.7.4.2 Non-repudiation. Guidance for non-repudiation is provided in the Non-Repudiation Framework. This document is a WD. The framework provides information on non-repudiation mechanisms and issues surrounding the implementation of a non-repudiation service. There are currently no standards for non-repudiation mechanisms. Work is ongoing in ISO and ANSI on standards for non-repudiation mechanisms, but only one working draft is currently available. This document is called Nonrepudiation Using Symmetric Encipherment Algorithms.

11.7.5 Service availability. The following paragraphs will discuss standards associated with service availability.

11.7.5.1 Protocol based denial of service protection. There are no standards that discuss protocol based DOS protection. However, existing communication protocols, such as TP4 and Fiber Distributed Data Interface (FDDI), provide active acknowledgements of receipt of data units. Any type of active acknowledgement or negative acknowledgement provides a measure of protection. The key is to identify when a host, link, or router goes down so that the network manager can be notified and alternative means of communication implemented as quickly as possible to minimize service loss or information loss.

11.7.5.2 Continuity of Operations. As with protocol based DOS protection, there are no standards for continuity of operations. Existing communication protocols can be used to provide some measure of protection. The protection provided is limited by the architecture of the system. If there are no alternate routes, operations can not be continued if the primary link goes down. FDDI provides a measure of protection (through its reconfiguration capabilities) as do routing protocols such as IS/IS and ES/IS (through the use of alternate routes). System architecture should be defined to use SAFENET reconfiguration options in a manner suitable to provide the required continuity of operations.

11.8 Architectural considerations. The interconnection of a distributed set of nodes on a network plays an important role in aggregate system security. Appropriate levels of information classification, user clearances, and system assurances must be considered in defining permissible network connections. Management of network information may require communication between groups of systems that might otherwise not be required to communicate. Distributed directory and network management databases may contain sensitive information that must be adequately protected. The distribution and management of key material and access control certificates must be accomplished through secure communication services. Security management policies must be established to assist the network security manager in properly controlling information flow through multiple SAFENET interconnections.

The information contained in this section represents only a subset of the architectural security issues that apply to SAFENET implementations. Each system will require security engineering to ensure that the architecture and subsequent design and implementation is tailored to address the specific threat environment.

11.8.1 Connectivity. When integrating systems on a network, each communicating system must provide adequate assurance to comply with the aggregate risk index as defined by the Trusted Network Interpretation Environmental Guidelines (TNIEG). The permissible connections should be defined in the network security policy to limit host and user accesses to only those required to support the intended system mission.

The cascade problem, as defined in the Trusted Network Interpretation (TNI), should be mitigated through strict network host access controls that permit only those hosts processing information at the proper levels and assurances to communicate. The problem exists when a penetrator can take advantage of

network connections to compromise information across a range of security levels that is greater than the accreditation range of any of the component systems he must defeat to do so. Figure 17 shows how information may be transferred from System A to System B. The Top Secret information in System A may be compromised to Confidential if a penetrator is able to defeat the protections of the two systems. Figure 18 shows an example of the nesting condition. The nesting condition is satisfied if the accreditation ranges of every node pair within a network are either disjoint (have no common elements) or nested (one is a subset of the other). The nesting condition is one method of protecting against the cascade problem.

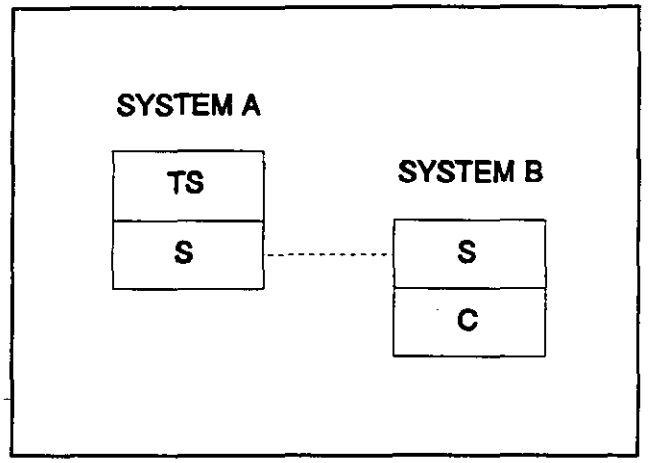

Figure 17. The Cascade Problem

Host access controls should be implemented according to a clearly defined network security policy. The network security administrator should be responsible for executing this policy through careful control of the distribution of key material, host authentication certificates and RBAC. Controlled access to end system services by users is an important matter that must be handled outside the implementation of SAFENET by the host operating systems and applications software.

11.8.2 Management. While the SAFENET handbook service definitions provide a basis for constructing secure networks, additional controls must be properly implemented and managed. The administration of system security will include the distribution of security information, management of an audit repository, and implementation of access controls to network security information. The functions provided by the security

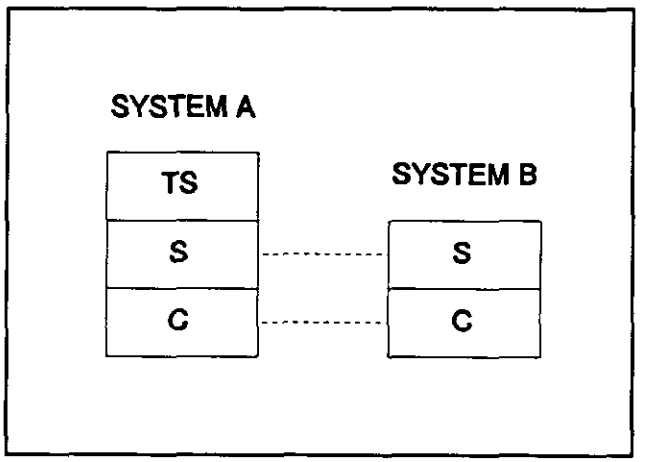

Figure 18. The Nesting Problem

manager must be distinctly separated from those provided to support the remainder of the network management functions.

11.8.2.1 Distribution of security information. Privately held information will be necessary to control access to sensitive and critical information and processes distributed across the network. Security control information may include user certificates for authentication, host certificates for host access controls, audit information, and access control lists. To ensure resources protected by this information remain adequately protected, the distribution, management, and storage of security-relevant information must be adequately protected through access controls, confidentiality, and integrity services.

The distribution of security information may take one of three forms:

a. Physical distribution, where security information is directly loaded into each host participating in secure processing and communications

b. Distributed information transfer, where security information is transferred to all hosts by some central security manager or managers

c. Distributed through a combination of physical distribution for initial "boot" information, followed by network distribution of remaining information.

Each of these approaches has positive and negative impacts on system security and management. Physical distribution of all information may involve a significant amount of physical distribution of information, presenting a large administrative burden to the security manager in tracking the physical location of the media containing the security information. On-line distributed transfers, while minimizing administrative burdens, may increase the risk of information disclosure or modification while security information is transferred over the network. As a compromise, a combination of distribution methods may best minimize administrative burden and system security risk. The combination may additionally provide the necessary protection to security information transferred over the network. The specific techniques chosen for a given network implementation must be based on a trade-off of network complexity, physical proximity of hosts, cost, performance, and sensitivity and criticality of information and processes.

In general, security control information should be treated at the level of sensitivity and criticality of the information

and processes that it protects. The following paragraphs describe the security control information that should be managed and protected.

11.8.2.1.1 Authorized host connections. Authorized host connection information may take the form of ES/ES permissible interconnect tables, additional host access control lists, and cryptographic key linkages. Since this information provides initial access to an end system, it should be protected to the level of the most sensitive information or processes residing on the host system.

11.8.2.1.2 Subject certificates. The authentication framework discusses subject certificates to authenticate subjects residing on a network, which are subsequently used by access control mechanisms as unique, unforgeable identifiers. The subject certificates may optionally be contained within a distributed directory service database, if the access controls within the database are managed to permit only the security manager and the authentication process to gain direct access to information contained within the certificate. Individual subjects should require only indirect access to this information to prompt the process to perform some form of authentication. This access may take the form of a direct reference by name, provided the operating system provides local authentication of the requesting subject to prevent local (within the same host) misuse of the certificate.

11.8.2.1.3 Host certificates. Host certificates, when not under the control of approved Communications Security (COMSEC) mechanisms, should be controlled to limit access to only a trusted implementation of the network layer software and hardware. When the host certificates take the form of Type I COMSEC keys, the certificates should be managed following the measures described in the key management discussion (see paragraph 11.6.2.4). Since the certificates provide a control point for all information entering a given host, physical distribution methods should be used to load the certificates into each host. After the initial physical distribution, on-line secure network-based transfers may be used to distribute additional host certificates.

11.8.2.2 Audit capture. Implementation of mechanisms to capture audit information will require the definition of the amount, type, and storage location of security-relevant information to be captured. In general, all audit information must initially be captured in all network components participating in the SAFENET LAN segment. While this is the most expedient method of developing and certifying a collection

mechanism, analysis of the information to identify and control potential security problems may require some form of centralized collection and processing of information. To ease the burden on security personnel in processing audit information, an application may be necessary to collect audit information using network management services and to generate audit information. Distributed security management applications may be necessary to identify critical events and transmit alarm messages to a centralized security manager.

11.8.2.3 Manager connectivity. When a network based on SAFENET comprises multiple, operationally independent collections of interconnecting hosts, the addition of dedicated hosts to provide network and security management services poses a potential risk to system security. The alternatives in connecting network and security management systems include:

a. Separate network manager hosts for each group of communicating systems

b. A central repository that only collects information and is incapable of transmitting (another system would be responsible for issuing audit information requests to be directed to the recipient)

c. A highly trusted central network management host.

The choice of alternatives will depend on the architecture of the overall system. Access to network manager systems should be limited to those individuals cleared to the level of the most classified or sensitive information contained within the systems. Since this information may contain records of user level connections and information transfers, the network/security management personnel should be authorized to obtain access to the information prior to accessing the network management facility.

11.8.2.4 Key management. If cryptographic mechanisms are used in the system, some form of key management will be required. The key management system functions include key generation, key distribution, key transfer, key update, rollover, destruction, periodic modification of, and accounting for all keys used by the encryption services provided within the system. The following sections discuss the three existing key management standards.

11.8.2.4.1 Secure data network system key management protocol. The SDNS Key Management Protocol (KMP) was designed to provide key management services for all SDNS protocols (SP3, SP4). It is an application layer protocol and can be used to

provide services to support cryptographic mechanisms (for data confidentiality and integrity) at any layer of the OSI model.

11.8.2.4.2 Standard for interoperable LAN security. SILS also provides a key management standard. It was originally designed to support the cryptographic mechanisms of SILS but can be used with cryptographic mechanisms at any layer. It is also an application layer protocol.

11.8.2.4.3 Transport layer security protocol addendum. ISO has been working on a security association establishment protocol for use with TLSP. This protocol is specific to TLSP and operates at the transport layer.

11.8.3 Assurance. Assurance is the method of determining the trust that can be placed in a processing device and software. The assurance of a device or system depends on the design, analysis, testing, and performance of the installed devices. Assurance is not determined by the security services that are implemented within a system. The final arbitrator of the assurance of a system is the Certification Authority or the Designated Approving Authority.

The level of assurance required by a system depends on the type and classification level of data that it will carry and the clearance level and type of users that will be allowed access to the system. The TNIEG can be used to determine the required level of assurance for a given application.

12. SAFENET TIME SERVICE

12.1 Introduction. The purpose of this section is to provide a discussion on some of the critical aspects for system engineers working with the SAFENET Time Service (STS). There are a number of factors of the STS that require the attention of the system engineer. The performance and stability of this service is heavily influenced by decisions made at the time of system design and deployment. Additional guidance on the implementation of NTP can be found in RFC 1305.

For the purposes of this discussion, the following definitions are provided. Precision (referred to as clock granularity in MIL-STD-2204) is the resolution of the value of GlobalTime provided by the network clock. Stability is a measure of how well a network clock maintains a constant frequency. Accuracy is how closely GlobalTime compares with the desired reference time for the system. The concept of reference time is discussed further in 12.3. A host is considered the network clock in the local SAFENET node; a peer is the network clock in the remote SAFENET node with which the NTP protocol machine in the local node is exchanging time information. A synchronization subnet is the logical tree structure organized for the purposes of synchronizing all the nodes in a SAFENET system. Epoch is the date chosen for the beginning of the SAFENET timescale (January 1, 1970).

12.2 The SAFENET timescale. MIL-STD-2204 defines a specific timescale for the STS. This timescale is based on the definition of GlobalTime. The two factors that distinguish this timescale are the origin and the manner in which leap seconds are handled. The SAFENET timescale has an origin of January 1, 1970. It is initialized to zero at 0 hours, 0 minutes, and 0.0 seconds on that day, and counts consecutive seconds continuously from that point. This approach results in a continuous timescale that covers approximately 136 years. This timescale is defined for use within the boundaries of the STS. Conversions to and from other time formats are outside the scope of SAFENET. However, the system engineer must be aware of Coordinated Universal Time (UTC) and the difference between GlobalTime and UTC in order to use external references and design conversion routines effectively.

UTC is the conventional civil timescale. UTC is generated by national standards bodies using very precise atomic oscillators. As a result of slight variations in the speed of rotation of the earth, the timescale generated by these atomic oscillators must be adjusted. This adjustment is referred to as a leap second. When the error approaches 0.7 seconds, a leap

second is added to or deleted from the UTC timescale on the last day of either June or December. This event is determined in advance and notification is distributed by various standards organizations including the U.S. Naval Observatory.

To avoid one second discontinuities in clocks in SAFENET nodes, the STS chose not to incorporate leap seconds. The STS maintains a continuous timescale. When a leap second is added to or deleted from UTC, this information is distributed to all instances of the STS using the NTP protocol; however, this information does not have any impact on the value of GlobalTime in the network clock. Leap second information is maintained by the time manager for use by routines that perform conversions to other time formats and from external references. Leap seconds are an issue only at the boundaries of the STS.

12.3 Reference time. Reference time is an abstraction used here to help the reader visualize the functioning of the STS. The clock synchronization service in essence synchronizes the network clocks in a SAFENET network to a reference time. Reference time can be either a time standard or an ensemble timescale. A time standard can be an external source of time information such as the Global Positioning System (GPS) or a calibrated atomic clock. Alternatively, a time standard can be a particular network clock designated as the reference for that system. Reference time could also be an ensemble timescale. An ensemble timescale is one that is derived from a number of sources using a predefined mechanism (election, arbitrary selection, etc.). There are a number of ways to define an ensemble timescale. In any case, it is the responsibility of the system engineer to evaluate the requirements of the system and choose an approach for providing reference time to that system. The provision of a form of reference time to a system is outside the scope of SAFENET.

12.4 System configuration. Proper configuration of the synchronization subnet is critical for acceptable performance of the STS. There are two aspects of configuration that must be addressed by the system engineer. For each SAFENET node or host in the SAFENET system, the set of peers with which the particular host will exchange information must be defined. In addition, the type of service employed between each host/peer pair must be specified. Based on this configuration information and the current conditions in the SAFENET network, a synchronization subnet is dynamically formed. This synchronization subnet reconfigures within the limits established by the configuration chosen. Finally, the behavior and performance of the synchronization subnet is affected by a number of parameter choices within the NTP protocol.

12.4.1 Peer selection. The first aspect of configuration is the choice of peers with which a particular host (SAFENET node) will communicate for the purpose of exchanging time information. In the current NTP specification there is no means for automatically discovering time servers in the network. A manager needs to specify manually the set of peers with which each host will exchange NTP PDUs. This specification can be done via a configuration file or through manipulation of managed objects. Considerations to be made when choosing the set of peers include the distance between the two nodes, the expected quality of the peer's network clock, the accuracy requirements for the host's network clock, and the expected synchronization load at the peer. Each host should maintain associations with at least three peers for better results when selecting a synchronization source and during various failure scenarios.

12.4.2 Modes of service. The second aspect of configuration is choosing the type of service used between each host/peer pair. There are three service classes available in the NTP protocol. These service classes are a procedure call service, a symmetric service, and a multicast service. These service classes are implemented using a selection of five modes of service (server, client, symmetric active, symmetric passive, and broadcast). Each network clock can operate in multiple modes of service based on its configuration and that of the various peers with which it is communicating. From that set of peers, a host will select a synchronization source using the peer selection algorithms. This selection will be based on the rules associated with the various modes of service. In addition, only choices of modes of service that form an identified class of service are allowable within NTP.

The procedure call class of service defines the interaction between nodes acting in the server and client modes of service. In this case, the server is willing to provide synchronization to the client but not to accept it. The client, on the other hand, is willing to accept synchronization from the server but not provide it. The client makes a request to the server and the server responds. In this class of service, the accuracy benefits of the two way exchange of time are achieved; however, the configuration is more static because only the server can be chosen as a synchronization source.

The symmetric class of service is a more generalized situation. In this class of service, a host, operating in either symmetric active or symmetric passive modes, is willing to provide synchronization to or accept synchronization from a peer, also operating in either symmetric active or symmetric passive modes. This relationship also benefits from the two way exchange

of time. In addition, a "dynamically reconfigurable, hierarchically distributed" configuration can be achieved because individual nodes are willing to both provide synchronization and receive it.

The final class of service provided by NTP is multicast. In this situation, a server is willing to provide synchronization to clients but not to accept NTP PDUs from that client. The client uses an assumed delay instead of calculating it based on the two-way exchange of time. This mode of service may be less accurate; however, it is simpler and may be suitable for some systems or for a subset of the machines in a system.

12.4.3 Hierarchial structure. The NTP system forms a hierarchial synchronization subnet following the conventions of the telecommunications industry. Using the peer selection algorithms, each host selects a single peer as its synchronization source. Based on these choices, a synchronization subnet is formed. The servers closest to external sources or to reference time are assigned a low number, called stratum, indicating their proximity to the root of the tree and indirectly their accuracy. As one progresses outward toward the leaves in the tree, each successive server encountered has a stratum number that is the stratum value of its own synchronization source, incremented by one. As a result, the stratum of a particular network clock dynamically reflects how far that node is from the root of the synchronization subnet. The system engineer can manipulate the structure of the synchronization subnet by the choice of peers and modes of service. The network clocks close to the root of the synchronization subnet should be some of the better quality clocks in the system in order to promote good accuracy and stability.

12.4.4 Parameter values. There are a number of parameters defined in the NTP specification. The values of these parameters have been explicitly engineered for the Internet environment. MIL-STD-2204 requires that the STS be able to modify these parameters. For the best results, at least a few of these parameters will need to be adjusted for the SAFENET environment. The system engineer will need to spend some time characterizing the requirements of the specific system and engineering the parameter values appropriately. It is noted here that some current implementations of NTP treat these parameters as compile time constants. It is the intention of the STS that these be modifiable during execution. Some of the parameters that may require special engineering for a SAFENET system are briefly discussed below. This is only overview material. Additional guidance is provided in RFC 1305.

12.4.4.1 Protocol parameters. There are a number of protocol parameters that have an impact on the performance of the STS. The minimum and maximum polling intervals affect how quickly the STS can initialize and react to changing conditions in the synchronization subnet. Another consideration when modifying this parameter is the network traffic that will result when smaller intervals are used. In addition to the polling intervals, the maximum distance parameter is used to control how many exchanges between the host and peer occur before the clock considers itself initialized. As a result, this parameter is also used to help specify the initialization time. If the value is set to the maximum dispersion value, initialization can occur after a single exchange of timestamps. Each time this value is halved, the number of timestamp exchanges required is increased by one. The next parameter, the maximum clock age, identifies the amount of time after a server has become disconnected from a synchronization source that it considers itself to have correct time. This also affects how long a client will continue to accept time from a server when that server has lost its synchronization source. The maximum skew parameter specifies the maximum frequency variation permitted by a client. It is specified as the maximum offset error that can occur over the interval specified by the maximum clock age. The minimum and maximum dispersion parameters affect a number of things, including the number of timestamps that must be exchanged before initialization or tracking and the importance of various samples in the peer selection process.

There are also a number of parameters that affect the configuration of the synchronization subnet. The maximum stratum parameter limits the depth of the synchronization subnet. For well managed systems, it may be acceptable to have a more shallow synchronization subnet than is currently used in the Internet. The minimum and maximum select clock parameters control the number of peers that a host will consider when choosing a synchronization source. Appropriate use of these parameters allows the system to use enough peers to maintain a robust system without using too large a number and consuming an excessive amount of resources.

12.4.4.2 Clock parameters. The local clock model used in NTP has a number of relevant parameters. These need to be examined when building a Network Clock implementation for SAFENET. Of particular interest are the adjustment interval and the maximum aperture. The adjustment interval specifies the amount of time between each clock adjustment. The maximum aperture specifies the interval over which an offset can be gradually applied to the Network Clock. If the offset value

determined is greater than the maximum aperture, a step adjustment may be required.

12.5 Performance parameters. The following performance parameters are recognized as being vital to the operation of the clock synchronization service. While no performance requirements are imposed by the STS, these parameters need to be specified in a system design in order to determine the correct operational parameters of the time synchronization protocol.

12.5.1 Clock synchronization accuracy. The required clock synchronization accuracy needs to be specified. This specification is made by identifying the following condition. The network clock in a LAN node shall be maintained such that the difference between any two network clocks sharing a peer path does not exceed a given value (i.e., 0.5 binary milliseconds). Also, the difference between any two clocks in a SAFENET system (including those residing on separate interconnected LAN segments) shall not exceed twice the given value (i.e., 1.0 binary millisecond) relative to a common frame of reference.

12.5.2 Clock initialization and fault detection times. The clock initialization time and the fault detection time parameters need to be specified in order to determine appropriate polling rates for the time synchronization protocol in a specific network. These two parameters are related but distinct. Clock initialization time refers to the time that it takes for a network clock to be synchronized to within the clock synchronization accuracy upon initialization or reset. Fault detection time refers to the time it takes to recognize a faulty clock in an operational network. The first of these parameters impacts the polling interval during initialization, and the second impacts the polling interval during normal operation of the network. Also impacting these parameters is the configuration of the system. Care must be taken when configuring the system in order to obtain reasonable initialization times across multiple LAN segments.

12.5.3 Use of an external time reference. The process of synchronizing a network clock in a SAFENET network to an external time reference is a local matter. While SAFENET establishes no performance requirements for this function when it is present, the system engineer needs to choose external time references that meet the synchronization requirements of the system. In addition, the STS will need the specifications of any external references chosen in order for the protocol to select reliable peers. This would include such things as the frequency tolerance and the synchronization distance.

12.6 Network clock implementation. A Network Clock is the combination of the hardware and software used to provide a local source of time. The implementation of a Network Clock is a local matter. However, some guidance on this topic is provided here. Additional guidance is provided in RFC 1305.

The two basic characteristics of a clock are phase and frequency. In this context, to synchronize a clock is to synchronize both its phase and its frequency. As stated in MIL-STD-2204, the network clock must be implemented in such a way that both the phase and frequency corrections provided by the clock synchronization service have a positive impact on the performance of the network clock. The phase of an individual network clock can be synchronized in a relatively short period of time; however, synchronizing the frequency of a clock requires significantly longer periods of time. The system engineer needs to make design trade-offs between quality of synchronization achievable, the time required to achieve that required level of synchronization, the expense of any specialized components used, and the network resources used in the process.

Once the system engineer has implemented the network clock, certain performance characteristics of that network clock need to be identified and specified for the accurate performance of the STS. In particular, the system engineer needs to be able to bound the frequency offset of the network clock implementations in the absence of any outside synchronization. This is in order to maintain correct error bounds and provide accuracy guarantees. In addition, such parameters as stability, wander, jitter, and availability are used to characterize a particular network clock.

12.7 Management of the STS. The management aspects of the STS are an important part of this service. MIL-STD-2204 specifies managed objects for use by the time manager. All remaining specification is left to the system engineer. A number of functions that should be supported are presented below. In addition, a few critical issues are discussed. Additional specification is needed for a completely interoperable implementation of the time manager.

12.7.1 Time management functions. A number of functions are necessary for an implementation of a time manager. Some of these functions are specified below. The method of implementing these functions is not specified. Additional functions may be required for a complete realization of the STS. These additional functions are also not specified.

12.7.1.1 Time synchronization protocol configuration. The time manager needs to provide mechanisms for the dynamic

manipulation of peers, polling intervals, and modes of operation. The time manager needs to have the capability, within limitations specified at system design time, to change these selections dynamically during operation of the system.

12.7.1.2 Time synchronization protocol parameters. The time manager needs to provide access to the basic parameters of the time synchronization protocol. For efficient operation, all instantiations of the protocol needs to use the same parameter values. The time manager needs to have mechanisms to set these parameters dynamically and distribute them to all instantiations of the protocol.

12.7.1.3 Leap second indication. The time manager needs to provide a mechanism for a network manager to set an indication that a leap second is about to be added to or deleted from the UTC timescale. The time synchronization protocol (NTP) then handles distribution of this indication to all clients. The time manager also needs to keep track of the total number of leap seconds that have accumulated in the UTC timescale. This is to facilitate the conversion to and from other time formats.

12.7.1.4 STS initialization. The time manager needs to provide support for initialization of the STS. In particular, the time manager needs to be able to support various values of protocol parameters in order to meet the requirements for time initialization. The sequence of events for this process is yet to be defined.

12.7.2 Critical issues. There are a number of critical issues involved with the time manager. A few of these are briefly discussed here. First, there is no interface specified between NTP and the time manager. Some guidance is provided in RFC 1305; however, this is not enough to ensure independent interoperable implementations of the time manager. In addition, there is no standard method to stop and restart an STS entity. In order to support dynamic reconfiguration of modes of service and parameter values in an interoperable manner, this must be specified. These are currently unspecified in MIL-STD-2204. The system engineer needs to be aware of these issues.

Custodians:
Army - ER
Navy - EC
Air Force - 10

Review Activities:
Army - TM, SC, IE, ET, AC
Navy - CG, TD
Air Force - 02, 17, 21, 23, 26

Preparing Activity:
Navy - EC
(Project - MCCR 0037)

STANDARDIZATION DOCUMENT IMPROVEMENT PROPOSAL

INSTRUCTIONS

1. The preparing activity must complete blocks 1, 2, 3, and 8. In block 1, both the document number and revision letter should be given.
2. The submitter of this form must complete blocks 4, 5, 6, and 7.
3. The preparing activity must provide a reply within 30 days from receipt of the form.

NOTE: This form may not be used to request copies of documents, nor to request waivers, or clarification of requirements on current contracts. Comments submitted on this form do not constitute or imply authorization to waive any portion of the referenced document(s) or to amend contractual requirements.

I RECOMMEND A CHANGE:	1. DOCUMENT NUMBER	2. DOCUMENT DATE (YYMMDD)
	MIL-HDBK-818-1	92/10/31

3. DOCUMENT TITLE
Survivable Adaptable Fiber Optic Embedded Network (SAFENET) Network Development Guidance

4. NATURE OF CHANGE *(Identify paragraph number and include proposed rewrite, if possible. Attach extra sheets as needed.)*

5. REASON FOR RECOMMENDATION

6. SUBMITTER

a. NAME (Last, First, Middle Initial)	b. ORGANIZATION	
c. ADDRESS (Include Zip Code)	d. TELEPHONE (Include Area Code) (1) Commercial (2) AUTOVON (If applicable)	7. DATE SUBMITTED (YYMMDD)

8. PREPARING ACTIVITY

a. NAME	b. TELEPHONE (Include Area Code) (1) Commercial	(2) AUTOVON
Commander Space & Naval Warfare Systems Command Code 231-2B2 (Attn: Gregg Sawyer)	(703) 602-3966	332-3966
c. ADDRESS (Include Zip Code) Washington, D.C. 20363-5100	IF YOU DO NOT RECEIVE A REPLY WITHIN 45 DAYS, CONTACT: Defense Quality and Standardization Office 5203 Leesburg Pike, Suite 1403, Falls Church, VA 22041-3466 Telephone (703) 756-2340 AUTOVON 289-2340	