STARS

University of Central Florida
STARS

Institute for Simulation and Training

Digital Collections

1-1-1995

Communication Architecture For Distributed Interactive Simulation (CADIS): Military Standard (draft)

University of Central Florida Institute for Simulation and Training

Find similar works at: https://stars.library.ucf.edu/istlibrary
University of Central Florida Libraries http://library.ucf.edu

This Research Report is brought to you for free and open access by the Digital Collections at STARS. It has been accepted for inclusion in Institute for Simulation and Training by an authorized administrator of STARS. For more information, please contact STARS@ucf.edu.

Recommended Citation

University of Central Florida Institute for Simulation and Training, "Communication Architecture For Distributed Interactive Simulation (CADIS): Military Standard (draft)" (1995). *Institute for Simulation and Training*. 38.
https://stars.library.ucf.edu/istlibrary/38

University of Central Florida
STARS
Showcase of Text, Archives, Research & Scholarship

INSTITUTE FOR SIMULATION AND TRAINING

[redacted]

MAY 15, 1992

MILITARY STANDARD (DRAFT)

COMMUNICATION ARCHITECTURE FOR DISTRIBUTED INTERACTIVE SIMULATION (CADIS)

IST-TR-92-18

[redacted]

IST

INSTITUTE FOR SIMULATION AND TRAINING

Contract N61339-91-C-0091
PM TRADE
DARPA

May 15, 1992

Military Standard

Draft

Communication Architecture for Distributed Interactive Simulation (CADIS)

Institute for Simulation and Training
12424 Research Parkway, Suite 300
Orlando FL 32826

University of Central Florida
Division of Sponsored Research

IST-TR-92-18

PRELIMINARY DRAFT

COMMUNICATION ARCHITECTURE

FOR

DISTRIBUTED INTERACTIVE SIMULATION

[CADIS]

PRELIMINARY DRAFT

FORWARD

1. This draft military standard has been prepared by the Institute for Simulation and Training for PM TRADE and DARPA based on currently available technical information but it has not been approved for promulgation. It is subject to modification. However, pending its promulgation as a coordinated military standard, it may be used.

2. Beneficial comments (recommendations, additions, deletions) and any pertinent data which may be of use in improving this document should be addressed to: Dr. Bruce McDonald, Institute for Simulation and Training, 12424 Research Parkway, Suite 300, Orlando, FL 32826 by using the self-addressed Standardization Document Improvement Proposal Form appearing at the end of this document or by letter.

3. This standard contains requirements for the communication architecture to be used to support distributive interactive simulation applications. This document contains the standards that specify the communication architecture, the recommended practices for implementing the communication architecture, and the rational to support these standards and practices. This document is not intended to be either a communications or a DIS tutorial. This document is not intended to be a guide to good system design. Information in these areas will only be presented to illustrate a point being made.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Figures

1. SCOPE

1.1 Scope. This standard establishes the requirements for the communication architecture to be use in a distributed interactive simulation application. This document contains both the standards and the recommended practices for implementing the communication architecture and the rationale behind them.

1.2 Application. This document has three main purposes. The first is to provide government agencies that are procuring DIS applications with the information necessary to write specifications. As such the document establishes a series of standards for network services, protocols, network performance, security, and network management. When invoked in a specification or statement of work, these requirements will apply to the communication architecture supporting simulation devices intended for participation in a Distributed Interactive Simulation (DIS). The contractor is responsible for invoking all the applicable requirements of this Military Standard on any and all subcontractors he may employ. The document also contains a set of recommended practices to provide guidance in those areas where standards have not been set or are not appropriate.

The second purpose of this document is to provide system designers with the information necessary to develop key areas of the system. In addition to the standards and recommended practices, the document contains extensive rationale supporting the choice of key items that have become part of the standard. This rational is intended to give the system designer a better understanding of why some choices were made and what impact deviation from them might have on the system being designed.

A third purpose is to provide the characteristics of the Wide Area Network (WAN) communications service that will be required when interconnecting DIS applications at different locations. This information, together with the development plans and rollout schedules for the various DoD systems that will use DIS, will be used by the Defense Information Systems Agency (DISA) to plan evolving DoD common-user wide area networks (the Defense Information Systems Network or DISN) with the functional, performance and security characteristics to support these systems. (Rationale: DMSO needs to provide this information to DISA for the whole spectrum of modeling and simulation, and this document provides an excellent source of the communications services characteristics for the DIS community.)

2. APPLICABLE DOCUMENTS

2.1 Government Documents.

2.1.1 Specifications, standards, and handbooks. The following specifications, standards, and handbooks form a part of this document to the extent specified herein. Unless otherwise specified, the issues of these documents are those listed in the issue of the Department of Defense Index of Specifications and Standards (DODISS) and supplement thereto, cited in the solicitation.

SPECIFICATIONS

MILITARY

MIL-STD-1777 12 August 1983	Military Standard, Internet Protocol
MIL-STD-1778 12 August 1983	Military Standard, Transmission Control Protocol
MIL-STD-1780 10 May 1984	Military Standard, File Transfer Protocol
MIL-STD-1781 10 May 1984	Military Standard, Simple Mail Transfer Protocol

(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.)

2.1.2 Other Government documents, drawings, and publications. The following other Government documents, drawings, and publications form a part of this document to the extent specified herein. Unless otherwise specified, the issues are those cited in the solicitation.

FIPS PUB 146-1 April 1991	U.S. Government Open Systems Interconnection Profile (GOSIP)Version 2.0 draft.

(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.)

2.2 Non-Government Publications. The following documents form part of this document to the extent specified herein. Unless otherwise specified, the issues of the documents which are

DoD adopted are those listed in the issue of the DODISS cited in the solicitation. Unless otherwise specified, the issues of documents not listed in the DODISS are the issues of the documents cited in the solicitation.

INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS (IEEE)

IEEE P1278[1]	IEEE Standard for Information Technology, Distributed Simulation Applications, and Data Entity Interchange Formats
IEEE 802-1990	IEEE Standards for Local and Metropolitan Area Networks - Overview and Architecture

(Applications for copies should be addressed to the Institute of Electrical and Electronics Engineers, 445 Hoes Lane, Piscataway, NJ 08854-4150.)

INTERNATIONAL ORGANIZATION FOR STANDARDIZATION (ISO)

ISO 7498-1984	Information Processing Systems - (CCITT X.200) Open Systems Interconnection - Basic Reference Model.
ISO 7498 AD 1-1987	Information Processing Systems - Open Systems Interconnection - Basic Reference Model - Addendum 1: Connectionless Mode Transmission.
ISO 7498-2-1989	Information Processing Systems - Open Systems Interconnection - Basic Reference Model - Part 2: Security Architecture.
ISO 7498-3-1989	Information Processing System - Open Systems Interconnection - Basic Reference Model - Part 3: Naming and Addressing.
ISO DIS 7498-4	Information Processing Systems - Open Systems Interconnection - Basic Reference Model - Part 4: Management Framework, 1989.

---

[1] This document was previously the draft military standard Protocol Data Units for Entity Information and Entity Interaction in a Distributed Interactive Simulation.

ISO 7776	X.25 LAPB Compatible DTE Data Link Procedures
ISO 8072-1986 (CCITT X214)	Information Processing Systems - Open Systems Interconnection - Transport Service Definition.
ISO 8073-1986 (CCITT X.224)	Information Processing Systems - Open Systems Interconnection, Connection Oriented Transport Protocol Specification.
ISO 8208	X.25 Packet Level Protocol for Data Terminal Equipment
ISO 8348	Network Service Definition
ISO 8473-1988	Information Processing Systems - Open Systems Interconnection - Data Communications Protocol for Providing the Connectionless Mode Network Service.
ISO 8878	X.25 to Provide OSI Connection-Mode Network Service
DIS 8880	Protocol Combinations to Provide and Support the OSI Network Service
ISO 9542-1988	Information Processing Systems - Open Systems Interconnection - Telecommunications and Information Exchange Between Systems - End System to Intermediate System Routing Exchange Protocol for Use in Conjunction with the Protocol for the Provision of the Connectionless-Mode Network Service.
ISO DIS 9545 (CCITT X.207)	Information Processing Systems - Open Systems Interconnection - Extended Application Layer Structure (ALS)(1988).
ISO DIS 9595	Common Management Information Service Definition
ISO DIS 9596	Information Processing Systems - Open Systems Interconnection - Management Information Protocol

	Specification - Common Management Information Protocol (CMIP) (1988).
DIS 10030	End System to Intermediate System Routing Exchange Protocol
DP 10040	Systems Management Overview
DP 10164	OSI System Management
DP 10165	Structure of Management Information
DP 10589	Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol

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

(Non-Government standards and other publications are normally available from the organizations that prepare or distribute the documents. These documents also may be available in or through libraries or other informational services.)

2.3 Order of precedence. In the event of a conflict between the text of this document and the reference cited herein, the text of this document takes precedence. Nothing in this document, however, supersedes applicable laws and regulations unless a specific exemption has been obtained.

3. DEFINITIONS

3.1 Acronyms used in this standard. The acronyms used in this standard are defined as follows:

ANSI	-	American National Standards Institute
BER	-	Bit Error Rate
CCITT	-	Consultative Committee for International Telegraphy and Telephone
CLNP	-	Connectionless Network Protocol
CMIP	-	Common Management Information Protocol
CMIS	-	Common Management Information Services
COTS	-	Commercial Off The Shelf
DIS	-	Distributed Interactive Simulation: The present name of the standard for interactive simulations.
ES	-	End System: an OSI DTE which is the source and/or sink of communicated process information.
ES-IS	-	End System to Intermediate System: Usually associated with routing.
FDDI	-	Fiber Data Distribution Interface
FTAM	-	File Transfer Access and Management
GOSIP	-	Government Open Systems Interconnection Profile
HW	-	Hardware
IEEE	-	Institute of Electrical and Electronics Engineers
IDRP	-	Inter Domain Routing Protocol
IP	-	Internet Protocol
IS	-	Intermediate System: an OSI DTE which interfaces, routes, maintains MTA information, and dynamically stores naming and addressing files for the ESs.
IS-IS	-	Intermediate System to Intermediate System: Usually associated with routing.
ISDN	-	Integrated Services Digital Network
ISO	-	International Organization for Standardization
LAN	-	Local Area Network
LEM	-	Local Exercise Manager
MAN	-	Metropolitan Area Network
MHS	-	Message Handling System
NIST	-	National Institute for Standards and Technology
OSI	-	Open Systems Interconnection
PC	-	Personal Computer
PDU	-	Protocol Data Unit
PM TRADE		U.S. Army Project Manager Training Devices
SIMNET	-	Simulation Network: An R&D effort which demonstrated the ability of simulators to interact dynamically over a LAN.
STRICOM		U.S. Army Simulation, Training, and Instrumentation Command (formally PMTRADE)
SW	-	Software
TCP	-	Transmission Control Protocol
TP	-	Transport Protocol
UCF/IST		University of Central Florida, Institute for Simulation and Training

UDP - User Datagram Protocol
WAN - Wide Area Network
XTP - eXpress Transfer Protocol
X.25 - A link and network layer set of protocols internationally accepted for DTE DCE interchange.

3.2 Application interface. The programming access mechanism to the communication resources of a network.

3.3 Application layer (layer 1). The layer of the ISO reference model which provides the means for user application processes to access and use the network's communications resources.

3.4 Broadcast mode (BC). A transmission mode in which a single message is sent to all network destinations, i.e. one-to-all. Broadcast is a special case of multicast.

3.5 Connectionless (CL). A mode of information transfer between peer entities in which each data transfer is independent of and not coordinated with previous or subsequent transfers and in which no state information is maintained.

3.6 Connection-oriented (CO). A mode of information transfer between peer entities in which a logical association is established prior to the exchange of data and which is maintained for the lifetime of the exchange process.

3.7 Datagram. A unit of data that is transferred as a single, no-sequenced, unacknowledged unit.

3.8 Data link layer (layer 2). The layer of the ISO reference model which provides the functional and procedural means to transfer data between stations, and to detect and correct errors that may occur in the physical layer.

3.9 Distributed Interactive Simulation (DIS). An exercise involving the interconnection of a number of simulation and/or real devices in which the simulated entities are able to interact within a computer generated environment. The simulation or real devices may be present in one location or be distributed geographically.

3.10 Emitter. A device that is able to discharge detectable electromagnetic or acoustic energy.

3.11 Entity. An element of a simulated world (such as a vehicle) that is generated and controlled by one or more host computers. An entity may also be an element of the simulated world, such as cultural features including building and bridges, that may be subject to changes in appearance as a result of the simulation exercise.

3.12 Exercise. See Simulation Exercise.

3.13 Host or Host computer. A computer attached to a network.

3.14 Interoperability. The capability, promoted but not guaranteed by joint conformance with a given set of standards, that enables heterogeneous equipment, generally built by various vendors, to work together in a network environment.

3.15 Local Area Network (LAN). A communications network designed for a moderate size geographic area and characterized by moderate to high data transmission rates, low delay, and low bit error rates.

3.16 Long-Haul network. See Wide Area Network.

3.17 Multicast mode (MC). A transmission mode in which a single message is sent to multiple network destinations, i.e. one-to-many.

3.18 Network layer (layer 3). The layer of the ISO reference model which performs those routing and relaying services necessary to support data transmission over interconnected LAN segments.

3.19 Network management. The collection of administrative structures, policies and procedures which collectively provide for the management of the organization and operation of the network as a whole.

3.20 Node. A general term denoting either a switching element in a network or a host computer attached to a network.

3.21 Non-Real time service. Any protocol function which does not require real time service. (see Real Time Service.)

3.22 ISO Reference Model (ISORM). A model that organizes the data communication concept into seven layers and defines the services that each layer provides.

3.23 Physical layer (layer 1). The layer of the ISO reference model which provides the mechanical, electrical, functional, and procedural characteristics access to the transmission medium.

3.24 Presentation layer (layer 6). The layer of the ISO reference model which frees the application processes from concern with differences in data representation.

3.25 Protocol. A set of rules and formats (semantic and syntactic) which determines the communication behavior of (N)-entities in the performance of (N)-functions.

3.26 Protocol Data Unit (PDU). A unit of data specified in an (N)-protocol and consisting of (N)-protocol-information and (N)-user-data. The term is used in this standard to refer to data that is passed on a network between application processes.

3.27 Protocol suite. A defined set of protocols within the communication architecture profile which constitutes a permitted implementation.

3.28 Real time service. A service which satisfies timing constraints imposed by the service user. The timing constraints are user specific and should be such that the user will not be adversely affected by delays within the constraints.[2] (DIS requires 5% of all data be processed within 100ms and 95% be completed within 300ms[3], therefore the DIS real time threshold is 100ms.)[4]

3.29 Reliable service. A communication service in which the number and type of errors that the user finds in the data is acceptable for the application. Reliable communication may require specific mechanisms in order to achieve the user's requirements: error detection and notification, such as bit errors based on a too high BER as defined by the user or error detection and correction from PDU errors, such as bit errors, duplicated PDUs, missing PDUs, or out-of-sequence PDUs.

3.30 Session layer (layer 5). The layer of the ISO reference model which provides the mechanisms for organizing and structuring the interaction between two entities.

3.31 Transport layer (layer 4). The layer of the ISO reference model which accomplishes the transparent transfer of data over the established link, providing an end-to-end service with high data integrity.

3.32 Wide Area Network (WAN). A communications network designed for large geographic areas. Sometimes call Long-Haul Network.

---

[2] Some data communications, e.g. voice, may require compensation to meet the timing constraint.

[3] These numbers are based on limited experience and are provided only as an experimental baseline.

[4] The amount of delay acceptable in a given application depends on the nature and intended use of the application. For some applications the acceptable delay may be less than 100ms or greater than 300ms.

3.33 Unicast mode (UC). A transmission mode in which a single message is sent to a single network destination, i.e. one-to-one.

3.34 Unreliable service. A communication service in which transmitted data is not acknowledged. Such data typically arrives in order, complete, and without errors. However, if an error occurs, nothing is done to correct it (e.g., there is no retransmission).

4. GENERAL REQUIREMENTS

4.1 Communication Architecture Overview

4.1.1 Introduction. The purpose of the communication subsystem for DIS is to provide an appropriate interconnected environment for effective integration of locally and globally distributed simulation entities. There are many diverse aspects of this integration, ranging from the nature of the entities represented within the common simulated environment, to the common communication interface used for receiving packets of information from other simulators. This standard is concerned only with the necessary communication system standards which must be accepted and adopted for supporting the integrated framework.

The Protocol Data Units (PDUs) defined in the DIS Standard are the "lingua franca" by which any two simulators or simulation sites can communicate. This includes simulators of different and unrelated design and architecture. No restriction is placed on what the participating simulator or site is, only on the way it communicates with the outside world.

Where the DIS PDUs define the information passed between simulators and simulation sites, this standard will define how those simulators, simulation sites, and other DIS entities can be connected in a modular fashion to facilitate the communication at the local and global levels. This will be done through the required use of communications standards which promote interoperability, such as the International Organization for Standardization (ISO) Open Systems Interconnection (OSI) reference model and the Government OSI Profile (GOSIP).

4.1.2 Approach to Communication Architecture. This standard specifies an evolutionary approach to be taken for implementation of a communication architecture for DIS. This approach, defined in terms of phases, is needed due to the fact that certain requirements for DIS cannot be met with the current OSI/GOSIP standardized communication network protocols. As protocols are developed which meet these requirements and are adopted by standards committees, they will be incorporated into this standard.

The communications architecture for DIS employs a layered model which is based on the seven layer OSI Reference Model (OSIRM) (see ISO 7498). The standards define the communication functions of the network by dividing them into a hierarchical set of layers. Each layer performs an integral subset of special functions required to communicate with another layer of similar type. There are seven layers in the OSIRM: Application, Presentation, Session, Transport, Network, Data Link, and Physical (Layers 7-1, respectively).

The functions provided by each layer, as presented by Tannenbaum[5], are summarized below:

LAYER	FUNCTION
1-Physical	Concerns the transmission of an unstructured bit stream over the physical medium; deals with the mechanical, electrical, functional, and procedural characteristics to access the physical medium.
2-Data Link	Provides for the reliable transfer of information across the physical link; sends blocks of data (frames) with the necessary synchronization, error control, and flow control.
3-Network	Provides upper layers with independence from the data transmission and switching technologies used to connect systems; responsible for establishing, maintaining, and terminating connections.
4-Transport	Provides reliable, transparent transfer of data endpoints; provides end-to-end error recovery and flow control.
5-Session	Provides the control structure for communication between applications; establishes, manages, and terminates connections (sessions) between cooperating applications.
6-Presentation	Provides independence to the application processes from differences in data representation (syntax).
7-Application	Provides access to the OSI environment for users and also provides distributed information services.

4.1.2.1 Approach to OSI Compliance. The strategy for OSI compliance is based on a phased, evolutionary approach. The first step to this evolution is the recommendation of an interim architecture, based on available network products and services, which is capable of supporting current exercises and communications experiments. The interim architecture will then transition to OSI/GOSIP standards over a period of years, as multicast communications protocol standards are adopted to support DIS.

4.1.2.2 Rationale: Benefits of DIS Compliance to the OSI/GOSIP Architecture. DIS compliance with the OSI/GOSIP architecture provides the following benefits: reduced cost,

---

[5] Tannenbaum, Computer Networks. Prentice Hall: 1988.

increased interoperability, and increased application-level functionality. Efforts to ensure conformance to OSI/GOSIP standards and ensure interoperability between products of different vendors means that computer networking can be done as an integration of multi-vendor, commercial-off-the-shelf (COTS) components. Easy access to vendor interoperable COTS OSI/GOSIP products gives wider availability to networking capabilities at a reduced cost.

Not only will OSI/GOSIP standards provide interoperability between products, but international interoperability will also be increased. The OSI standards are international in scope and will be used by North Atlantic Treaty Organization (NATO) allies, among others. Using OSI standards opens the possibility that interoperation with our NATO allies will be accomplished within the framework of international standards.

4.1.3 Key Assumptions For This Standard. This document makes a number of assumptions about underlying requirements of the DIS application and how they will be applied. These issues and assumptions must be understood if the application of this document is to be successful.

4.1.3.1 Long Haul Connection. Simulators at different sites will have to be connected via a Wide Area Network (WAN). This document defines the functional and performance characteristics which must be satisfied by the communications service, including the WAN. It is the goal of this communications architecture that the WAN be based on standards such as frame relay, Switched Multimegabit Data Service (SMDS), Broadband ISDN, and SONET. The provision of the WAN will depend on the evolution of these high speed communications services in the marketplace and the particular organization using the DIS applications.

Wide area networks today do not in general support multicasting. If two or three sites using DIS are to participate in a demonstration or exercise, they could be interconnected by point-to-point circuits or a network with sufficient capacity to support repeated transmission to each site. This, however, would become uneconomic for a larger number of sites.

For DoD systems that will be using DIS, the Defense Modeling and Simulation Office (DMSO), in collaboration with the Defense Advanced Research Projects Agency (DARPA) and the Defense Information Systems Agency (DISA), is establishing a testbed Defense Simulation Internet (DSI) to provide an interim WAN that supports multicasting and satisfies the latency requirements of DIS (see section 4.2.2.2). The testbed DSI is an extension of DARPA's Terrestrial Wideband Network (TWBNet) and will be available as a pilot operational system to support DIS and other applications until replaced by a permanent fully operational

system. This target system will be part of the Defense Information System Network (DISN) and will employ commercial standards-based services such as those listed in the following section.

The nature and development of WANs for DIS application is taking two distinct paths. The first is the establishment of a permanent infrastructure that will connect all DIS simulator sites. Although physically one large network, it will support multiple exercises via the creation of an individual logical network for each exercise. This approach is called the Defense Simulation Internet. The second approach is the establishment of Ad Hoc WANs as necessary to support exercises and tests. The primary mechanism for this is the bandwidth-on-demands services starting to be offered by the major communications suppliers (e.g. AT&T, MCI, Sprint). The concept is that a network connecting any set of simulator sites can be created quickly and efficiently from commercial services without the cost of maintaining a permanent infrastructure. The Advanced Distributed Simulator Technology (ADST) program is exploring this approach. This document does not assume either of these approaches and will support both of them.

4.1.3.2 Video Conferencing & Other Tools. A number of DIS documents, including the concept of operations, have identified a video conferencing requirement. This is to support exercise planning, briefing, and debriefing, but specific requirements (e.g. number of sites, functionality) have not been identified. The communications industry is creating new ways to achieve such video conferencing, but mature products are not yet available. Video conferencing is very demanding of network capabilities and will have a major impact on any DIS network design. Because the requirements for video conferencing are not clearly identified and because industry offerings are not stable, video conferencing is not addressed in this document. This requirement will however be addressed in future versions as the requirements and available services become better understood.

4.1.3.3 Multiple Exercises. DIS has the ability to accommodate multiple exercises over the network by assigning each exercise a different exercise ID. Those entities participating in the exercise will contain the assigned exercise ID in their PDUs. When the exercises involved require any level of security, then security mechanisms must be in place for the level of security needed. This standard will not address or support multiple exercises that have different levels of security.

4.2 Service Requirements. This section describes the services required to be provided by the communication architecture for DIS applications. These services are divided into two categories: communication service requirements and performance requirements. The service requirements for the

communication subsystem of the DIS standard are based on experience with state-of-the-art distributed simulation activities as well as projections based on anticipated use and evolution of the technology base.

4.2.1 Communication Service Requirements. Distributed simulation environment support requires various types of communication. The communication requirements encompass control and data. Data communications, including voice, may be with or without real time requirements and will likely be augmented to include such things as video and other forms of pictorial information. Upon the introduction of each of these forms of traffic, they shall share communications facilities instead of having disjoint facilities for each.

A summary of the communication service requirements is shown in Table I.

TABLE I. DIS Communication Service Requirements

Unicast
Multicast
Broadcast
Real Time Operating Speeds
Non-Real Time
Small Packets
Bulk Transfer
Reliable
Unreliable
Low Interpacket Dispersion for Voice/Video
Multicast Implementation
Multicast Management
Authentication/Access Control
Non-Blocking Interface
Flow Control
Low Latency Packet Delivery
Security
Flexible Entity Naming & Addressing
High Throughput

4.2.1.1 Service Requirements of PDUs. Each DIS PDU requires certain services to make its communication practical. These services are grouped into broad classes of operation for DIS.

4.2.1.1.1 Communication Classes Based on Requirements. This section establishes DIS communication classes based on the application service characteristics for both the required and recommended interim DIS PDUs. Each DIS PDU requires certain service characteristics to make its communication practical.

These characteristics are grouped into broad classes of operation for DIS.

4.2.1.1.1.1 Application Requirements. The DIS application (PDUs) has been characterized using the following subset of communication service requirements: unicast, multicast, broadcast, reliable, unreliable, real time, non-real time, packet size, and bulk transfer. The application service characteristics are used to define a service model necessary to support DIS communication. The service model developed from the PDU characterization shall be used to develop the interface to the application and lower layers.

4.2.1.1.1.2 DIS PDU Service Characterization. DIS functional requirements are to provide: Entity Information, Entity Interaction, DIS Management, and Environment Information. Within each functional category, PDUs have been defined or recommended to satisfy specific requirements. The October 1991 version of the DIS standard defines ten required PDUs and six recommended interim PDUs.[6] The application services for required and recommended DIS PDUs are defined in Tables II and III, respectively.

Although packet size and bulk transfer are included as application requirements in 4.2.1.1.1.1, it is not presented in the summary tables for the following reason. Inter-entity communication in a distributed interactive simulation environment consists largely of packets sent between two or more of the simulation participants. These packets are usually small, <250 bytes, and constitute the majority of PDU traffic. All PDUs listed in Table II and III fall into the "small packet" characterization. There are situations which mandate non-real time, point-to-point, reliable bulk transfer, however. Such situations arise when moving large items such as database files or video images. The bulk transfers fall into the Network and/or Simulation Management functions, but there are currently no PDUs which reflect this type of interaction. Consequently, bulk transfer is considered a special case.

---

[6] The October version of the DIS standard specifies three recommended PDUs for Update Threshold Control. As of this writing, those PDUs have been removed from the standard and, therefore, will not be included in this characterization.

TABLE II. Required DIS PDU Communication Services

	Reliable	Best Effort	BC	MC	UC	Real Time
Entity State		*		*		*
Fire	*desired*	*		*		*
Detonation	*desired*	*		*		*
Service Request		*			*	(few seconds)
Resupply Offer		*			*	(few seconds)
Resupply Received		*			*	(few seconds)
Resupply Cancel		*			*	(few seconds)
Repair Complete		*			*	(few seconds)
Repair Response		*			*	(few seconds)
Collision	*				*	*

TABLE III. Recommended DIS PDU Communication Services

	Reliable	Best Effort	BC	MC	UC	Real Time
Emitter	*desired*	*		*		*
Radar	*desired*	*		*		*
Activate Request		*			*	
Activate Response		*			*	
Deactivate Request		*			*	
Deactivate Response		*			*	

DIS Management will require additional capability beyond the Activate and Deactivate PDUs. Although these capabilities have not yet specified, Table IV projects additional application requirements for these areas.

TABLE IV. DIS Functional Requirements Communication Services

	Reliable	Best Effort	BC	MC	UC	Real Time
Network Management		*			*	
Simulation Management	*			*desired*	*	

4.2.1.1.1.2.1 Entity Information. The Entity State PDU (ESPDU) constitutes the bulk of network traffic for a simulation exercise. Currently, the appearance updates represented by the ESPDU are of most interest to exercise participants within a limited radius of the initiating entity. Any exercise participant which is not in the area of interest, but receives the ESPDU, will have to filter out this unwanted information. Therefore, Entity State has a strong requirement for multiple multicast interactions. Multicast interactions deliver identical packets to multiple recipients as part of a single sender operation. A multicast data transfer provides co-located entity groups the capability of communicating state information based on locale in the simulated exercise.

In addition to their multicast requirements, ESPDUs must be delivered in real time but do not need to be transmitted reliably. Dead Reckoning (DR) algorithms are used to predict the entity's position over time in order to preserve network bandwidth by reducing the frequency at which state information is required. Reliability need only be a best effort. If an ESPDU is lost, the DR models used to reduce network traffic may also be able to account for the lost packet.

4.2.1.1.1.2.2 Entity Interaction. Entity Interaction PDUs have varied characteristics. Within the Weapons Fire category, the Fire PDU (FPDU) and the Detonation PDU (DPDU) have the same service characterization. Similar to the ESPDU, both the FPDU and the DPDU have a strong multicast requirement. This requirement allows only those entities within the area of interest to receive information about weapons firing and detonation.

These PDUs also have a real time requirement, and should be more reliable than ESPDUs. Whereas ESPDUs can rely on DR to extrapolate position after packet loss, FPDUs and DPDUs are not as robust. When a weapon impacts, it is crucial that everyone in the multicast group receive that information so "killed" targets do not continue to play in the exercise. A high degree of reliability is desired for the FPDUs and DPDUs, however current multicast protocols do not provide this service. Therefore, FPDUs and DPDUs must use a best effort real time multicast service.

The Logistics Support PDUs (i.e., Service Request, Resupply Offer, Resupply Received, Resupply Cancel, Repair Complete, and Repair Response) represent activities which, although long in duration, do not require real time service. The resupply and repair interactions require a simple reliable transaction (request/reply) paradigm. This reliability is built into the application by pairing the acknowledgement (or reply) PDU with the request (e.g., Service Request and Resupply Offer PDUs). The Logistics Support PDUs do not require multicast, because only the entities involved in the service are interested. Therefore, the Logistics Support PDUs are characterized as requiring an unreliable unicast service.

The last required category of PDUs in Entity Interaction is Collisions. Collision PDUs require a real time, unicast service. Again, only the entities involved in the collision will be interested in this information. Changes in entity appearance resulting from the collision will be communicated using ESPDUs.

The only category of PDUs not required for Entity Interaction is Electronic Interaction. Electronic Interaction currently consists of two recommended PDUs, Emitter and Radar. Both PDUs desire a reliable real time multicast transmission but, as stated before, this is not available. Therefore, these PDUs are characterized as requiring best effort real time multicast.

4.2.1.1.1.2.3 DIS Management. There are no PDUs specified for Network Management. Network management will be handled by a standard network management protocol (e.g., Simple Network Management Protocol or Common Management Information Protocol) and will not require DIS PDUs to accomplish the management of the physical network. Network management is accomplished with an unreliable unicast service.

The Simulation Management category of PDUs is responsible for the activation and deactivation of simulation players. The request to activate or deactivate entities in a simulation exercise (i.e., Activate and Deactivate Request and Response PDUs) requires a simple reliable transaction (request/reply) paradigm. The reliability is built into the application by pairing the acknowledgement (or reply) PDU with the request

(e.g., Activate Request and Activate Response PDUs). This service is characterized as non-real time unicast. Other possible functions of Simulation Management include management and control messages spanning multiple exercises. This type of service desires a reliable multicast transmission, however reliable multicast is not currently available. Therefore, this type of service is characterized as reliable unicast. In addition to the packet form of interaction, there are situations which mandate non-real time, point-to-point, reliable bulk transfer. Such situations arise when moving large items such as databases or video images. Standard file transfer protocols such as File Transfer Protocol (FTP) or File Transfer Access and Management (FTAM) will be used.

There are no PDUs required or recommended for Performance Measures. If PDUs are developed for this functional area, the required services will fall into one of the established service classes.

4.2.1.1.1.2.4 Environment Information. There are no PDUs required or recommended for Environment Information. If PDUs are developed for this functional area, the required services will fall into one of the established service classes.

4.2.1.2 Communication Classes. From the previously stated rationale, three service models emerge as characterizing the DIS application.

CLASS 1 **Unreliable Multicast**
A mode of operation where the multicast service provider uses no added mechanisms for reliability except those inherent in the underlying service.

CLASS 2 **Unreliable Unicast**
A mode of operation where the unicast service provider uses no added mechanisms for reliability except those inherent in the underlying service.

CLASS 3 **Reliable Unicast**
A mode of operation where the unicast service provider uses whatever mechanisms are available to ensure the data is delivered in sequence with no duplicates and no errors.

The service model is shown in Table V.

TABLE V. DIS Application Service Model

CLASS 1 Unreliable Multicast	**CLASS 2** Unreliable Unicast	**CLASS 3** Reliable Unicast
Entity State	Service Request	Collision
Fire	Resupply Offer	Simulation Management
Detonation	Resupply Received	
Emitter	Resupply Cancel	
Radar	Repair Complete	
	Repair Response	
	Network Management	
	Activate Request	
	Activate Response	
	Deactivate Request	
	Deactivate Response	

4.2.2 Performance Requirements

4.2.2.1 Network Bandwidth Requirements. Network bandwidth requirements are subject to "best guess" estimation procedures due to the combination of man-in-the-loop and non-deterministic simulated adversaries in DIS[7]. The traffic requirements result from the frequency at which the PDUs are generated; the size of

[7] The terms bandwidth and traffic shall be used interchangeably, with both referring to the total number of bits per second which the network must carry.

the individual PDUs remains relatively stable. See Appendix B for a detailed explanation of bandwidth estimation procedures.

4.2.2.2 Latency Requirements. The following standards shall be required for communications latency values in the DIS environment:

100 milliseconds	Total latency permitted between the output of a PDU at the application level of a simulator and input of that PDU at the application level of any other simulator in that exercise when that exercise contains simulated units whose interactions may be tightly coupled.
300 milliseconds	Total latency permitted between the output of a PDU at the application level of a simulator and input of that PDU at the application level of any other simulator in that exercise when that exercise contains only simulated units whose interactions are not tightly coupled.
50 milliseconds	Maximum dispersion of arrival times of the voice PDU at the application level of the device converting digital voice to analog.

4.2.2.2.1 Rationale. Some interactions between simulated entities are very tightly coupled in time. That is, the action of an individual controlling one of the entities may be a reaction to the activity of another. How tightly these interactions are coupled in time depends on the performance of the unit being controlled. High performance units, that is those units that react quickly to a human controllers input, tend to be very tightly coupled. An example of this is one simulated fighter aircraft flying in close formation with another. Units that respond to control inputs less quickly, such as ships, are only loosely coupled.

The issue of communications latency is directly related to how tightly a simulated entity is coupled to the entity to which it is reacting. The more tightly coupled two simulated entities are, the less latency is permitted in the communications that carry the state data of each to the other. Allowable latency under different circumstances is the subject of considerable debate. Little research of the quality that can serve as the basis of standards for latency has been done. The best information available is from the flight simulator industry, which for many years has been struggling with a related issue called transport delay. Flight simulator experience provides the following:

1. Humans cannot distinguish differences in time that are less than 100 milliseconds. This is due to physiological factors of the human body. This effectively provides a floor latency/transport delay value. That is, with a human in the control loop, there is no benefit to be gained from latency/transport delay less than 100 milliseconds.

2. In situations where latency/transport delay reaches 300 milliseconds, pilots start compensating for the lag in response. The result is a phenomenon known as Pilot Induced Oscillation (PIO). Such PIO can range from a minor annoyance to total loss of control.

The flight simulation community has also experimented with schemes to compensate for transport delay by predicting the behavior of the device being controlled. This approach showed promise, but the main emphasis in dealing with transport delay has been in reducing the delay by faster processing and better communications within the simulator. The DIS community has also begun to explore prediction of position as a means to compensate for latency in tightly coupled interaction. Northrop has done the most work in this area. Studies reported to the DIS community[8] suggest that sophisticated prediction algorithms can compensate for up to 750 milliseconds of latency in the interaction of high performance aircraft carrying out radical maneuvers.

The position of simulated vehicles is not the only consideration in dealing with latency. DIS networks will also carry voice in the simulation of tactical radio nets. A speaker's voice will be converted from analog from to a digital data stream that will be treated as just another series of PDUs. At the listener's position these will be converted back into analog form and will be output to speakers and/or headphones. Latency in such voice communications carries its own considerations.

In the case of an overseas phone call that was routed via a geosynchronous satellite, latency of a half second or more is inherent in such communication. In normal conversation this is annoying but the speakers can generally adjust to it without difficulty. However, in the heat of a simulated operation such delays would render a simulated radio net unusable and would not be acceptable. Also, there is no prediction mechanism that can compensate for delays in voice traffic.

---

[8] Position paper "Techniques for Extrapolation, Delay Compensation, and Smoothing with Preliminary Results and an Evaluation Tool," S. Goel, K. Morris, IST-CR-91-13, Summary Report: The Fifth Workshop on Standards for the Interoperability of Defense Simulations.

The dispersion of the arrival times of voice PDUs is also important. In the process of converting analog voice to a digital data stream, the analog signal is sampled at regular intervals and each sample is converted to a digital message. For the reconstruction of the voice back to analog these messages should ideally arrive at the same regular interval. However, due to a variety of factors, there will some dispersion of arrival times. If the dispersion is too great, voice quality will suffer and may be unintelligible. The mechanism of converting voice from digital to analog form can handle some dispersion in arrival times. It is also possible to deliberately hold incoming voice PDUs in an accumulating FIFO buffer and then meter them to the voice reconstruction mechanism at a same rate at which the voice was sampled. This technique would eliminate the effects of delay dispersion, but would do so at the cost of additional overall latency.

4.2.2.2.2 Allocation of Latency Values. In designing systems that meet the total latency standards defined above, it is important to allocate these latencies in a reasonable manner. For example, if one designs a simulator with a latency of 45 milliseconds between the application layer and the media (layer 1) of the LAN to which it is attached, it will still meet the standard of 100 milliseconds for total latency with similar simulators on the same LAN. However, if it becomes part of an exercise that includes simulators from other geographic sites, the total latency will likely exceed 100 milliseconds due to the latency consumed by the WAN connecting the sites.

For this reason, the following standard allocation of latency shall be:

10 milliseconds	Maximum latency between the application and physical layers of any DIS simulator.
80 milliseconds	Maximum latency occupied by the media, bridges, routers, gateways, encryption/decryption devices, long-haul media, and other components of a network that connects simulators whose interactions may be tightly coupled.
280 milliseconds	Maximum latency occupied by the local media, long-haul media, bridges, routers, gateways, encryption/decryption devices, and other components of a network that connects simulators whose interactions are not tightly coupled.

Figure 1 summarizes the latency standards.

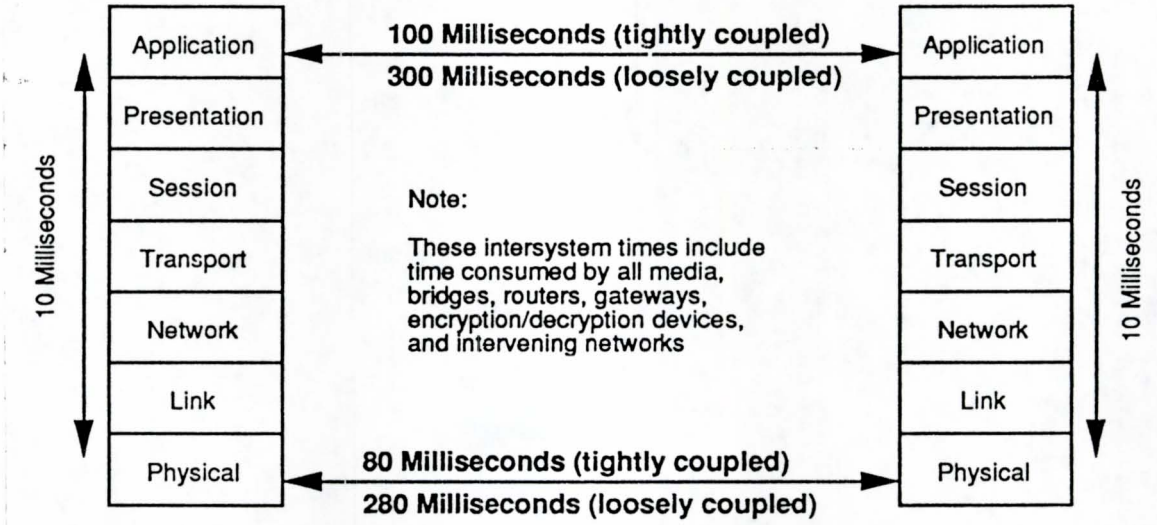

Figure 1. Standard Latency Values

4.2.2.3 Error Detection/Correction. Section 4.2.1 of this document, which defines the required services of the communications architecture, identifies PDUs which must be delivered reliably. This is a specific term that means that each of those PDUs must be delivered to its destination in the order in which it was sent and without error. Implied in this definition is that the receipt of each PDU must be acknowledged and retransmitted if necessary. Such acknowledgement and retransmission will be handled by the error detection/correction mechanism of the of the protocols used. That is, there is no action required at the application level other than to indicate that a particular PDU is to be sent reliably. The receiving application can assume that all PDUs sent reliably are in order and intact.

PDUs not requiring reliable delivery will be given best effort delivery. These PDUs make up the bulk of network traffic and include those PDUs that are multicast to all simulators in a DIS exercise. Acknowledgement and retransmission, associated with reliable delivery, is not feasible due to the additional latency and network bandwidth that would be required. Instead the application, via its dead reckoning mechanism, is tolerant of occasionally missed packets. There is also the possibility that a PDU with corrupted data may be received. The processing of such corrupted data may create unacceptable behavior in the receiving simulator. To prevent this the DIS communications architecture will include in its best effort delivery a checksum mechanism in its transport level protocol. This checksum will include the entire PDU. If a checksum error is detected in a

received PDU, the PDU will be discarded by the communications software. That is, it will not be made visible to the application.

4.3. The Communication Architecture Protocol Suites for DIS. The DIS communication architecture will be divided into three phases. The following diagram shows the three phases of protocol suites that shall be used for the communication architecture:

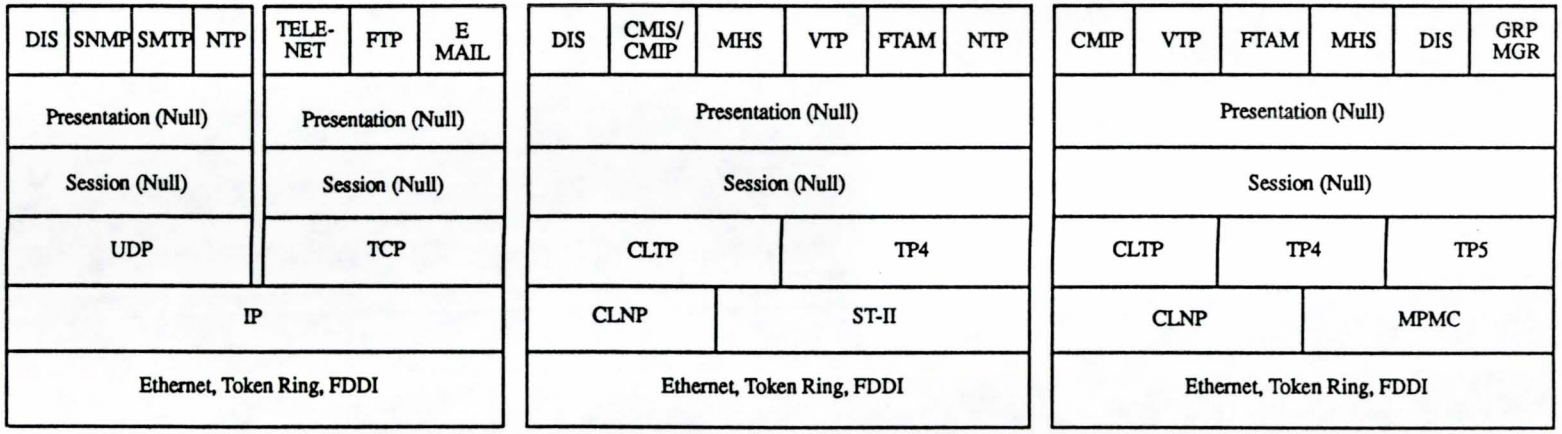

Figure 2: DIS Three Phase Communication Architecture

0330-2540.1

Phase 0 is a proof-of-concept for DIS communication applications. The Phase 0 communication architecture is composed of the following Internet standards:

DIS	- Distributed Interactive Simulation PDUs
SNMP	- Simple Network Management Protocol (RFC 1157)
SMTP	- Simple Mail Transfer Protocol (RFC 821)
Telnet	- (Terminal Protocol) (RFC 854)
FTP	- File Transfer Protocol (RFC 959)
NTP	- Network Time Protocol (RFC 1119)
UDP	- User Datagram Protocol (RFC 768)
TCP	- Transmission Control Protocol (RFC 793)
IP	- Internet Protocol (RFC 791)
ICMP	- Internet Control Message Protocol (RFC 792)
ARP	- Address Resolution Protocol (RFC 826)
RARP	- A Reverse Address resolution Protocol (RFC 903)
Open	- STream-II (RFC 1190) or XTP (by PEI)

Phase 1 is proof-of-concept of the DIS OSI communication infrastructure and hybrid implementation of the multicast protocol. The Phase 1 communication architecture is composed of the following ISO and Internet standards:

DIS	- Distributed Interactive Simulation PDUs
CMIP	- Common Management Information Protocol (ISO 9596)
MHS	- Message Handling System (CCITT X.400)
VTP	- Virtual Terminal Protocol (ISO 9041)
FTAM	- File Transfer Access and Management (ISO 8571)
NTP	- modified Network Time Protocol (RFC 1119)
CLTP	- ConnectionLess Transport Protocol (ISO 8602)
TP4	- Transport Protocol Class 4 (ISO 8073)
CLNP	- ConnectionLess Network Protocol (ISO 8473)
Open	- STream-II Protocol (RFC 1190) or XTP (by PEI)

Phase 2 is an enhanced OSI architecture based upon lessons learned in Phase 1, added functionality, and final versions of OSI/GOSIP multicast protocols. The Phase 2 communication architecture will be composed of the following ISO standards:

DIS	- Distributed Interactive Simulation PDUs
CMIP	- Common Management Information Protocol (ISO 9596) with possible extensions for multicast group management
MHS	- Message Handling System (CCITT X.400)
VTP	- Virtual Terminal Protocol (ISO 9041)
FTAM	- File Transfer Access and Management (ISO 8571)
OSITP	- OSI Time Protocol (undefined)
CLTP	- ConnectionLess Transport Protocol (ISO 8602)
TP4	- Transport Protocol Class 4 (ISO 8073)

TP5 - Transport Protocol Class 5 to provide a reliable multicast service (undefined)
CLNP - ConnectionLess Network Protocol (ISO 8473)
MPMC - Multipeer/Multicast Protocol to provide bandwidth reservation (undefined)

Note: All three phases of the communication architecture shall require a group manager function to specify the group membership management, group initiation, and group communication termination.

4.4 Interoperability Requirements. Interoperability can virtually always be achieved when the same engineer (or group of engineers) undertakes the communications problem at all sites. Interoperability that the DIS initiative is attempting requires that interactive operations be achieved by separate engineers at each site, consulting a standard document, and without communicating with each other. Such interoperability requires that the interface of each simulator to the network be specified down to the hardware plug. Simulators that comply, can be plugged in and will interoperate.

Strict interoperability requires that the standard take care of all technical aspects of linking together parts of the network. Only administrative details are left to be negotiated by the participants and the network. A good illustration of strict interoperability is the current telephone service. Telephone sets, modems, and fax machines may be produced by anyone. So long as they comply with certain specifications and FCC rules, they can be plugged in and will interoperate. The user must negotiate with the network for access to a socket and network address (telephone number). There may exist private telephone nets that are not open to all. So long as the open standard is complied with, access to the private net is denied by choice, not by problems of compatibility.

Much progress has been made over the past decade on standardizing approaches to interconnecting computer systems. Three aspects of the simulation application distinguish DIS from the more general computer/communication interconnection. These are: 1) real time delivery requirements for interactive, man-in-the loop behavior 2) multicast delivery options for convenient updating of shared data items and 3) military security considerations. The approach to interoperability specified in this standard is to adopt the more general communication framework, augmenting it only as necessary to meet the specific additional requirements of distributed interactive simulation.

Any approach taken toward communication interoperability must apply as well to as wide a variety of existing simulators as possible, preferably all. This interoperability integration must

be possible with minimal disruption of existing simulators, even at the expense of optimality and efficiency. To accomplish this for the widest class of existing simulators, including those already interconnected as well as those running stand alone, only the minimum properties should be standardized to accomplish the integration. This allows as many pre-existing configurations as possible to remain compliant with the minimum change, as well as accommodating the maximum flexibility for future innovation with minimum disruption to working systems.

4.4.1 Examples Employing the Interoperability Architecture. There are many approaches to integrating a simulator into an integrated DIS exercise which fit within the framework outlined above. From an architectural point of view, the following list enumerates a variety of possible simulator organizations, all of which are appropriate for meeting DIS interoperability requirements:

- a simulator and its DIS communication interface can coexist on a single host computer.

- a single host can run multiple simulations, using the same or different DIS host identities for these entities.

- a dedicated processor front end can be used for implementing the communication interoperability (as well as other DIS) requirements for one or more back end simulators. This approach is sometimes referred to as an "application gateway". One of the primary advantages to this approach is minimal interference with currently operational simulators. The interconnection of the application gateway with the simulator is not subject to the DIS standardization effort. A reasonable implementation of such a component might be useable by various classes of simulator.

- a simulation implementation can span multiple computers, either as part of a multiprocessor system, locally distributed, or even with geographically distributed components. With such arrangements, from the vantage point of the network, a single component is designated as representing the simulation in its entirety. Any information distribution among the components is entirely the responsibility of the simulator.

Figure 3 illustrates the variety of interconnection modes which can be supported by the communication architecture approach indicated in this report.

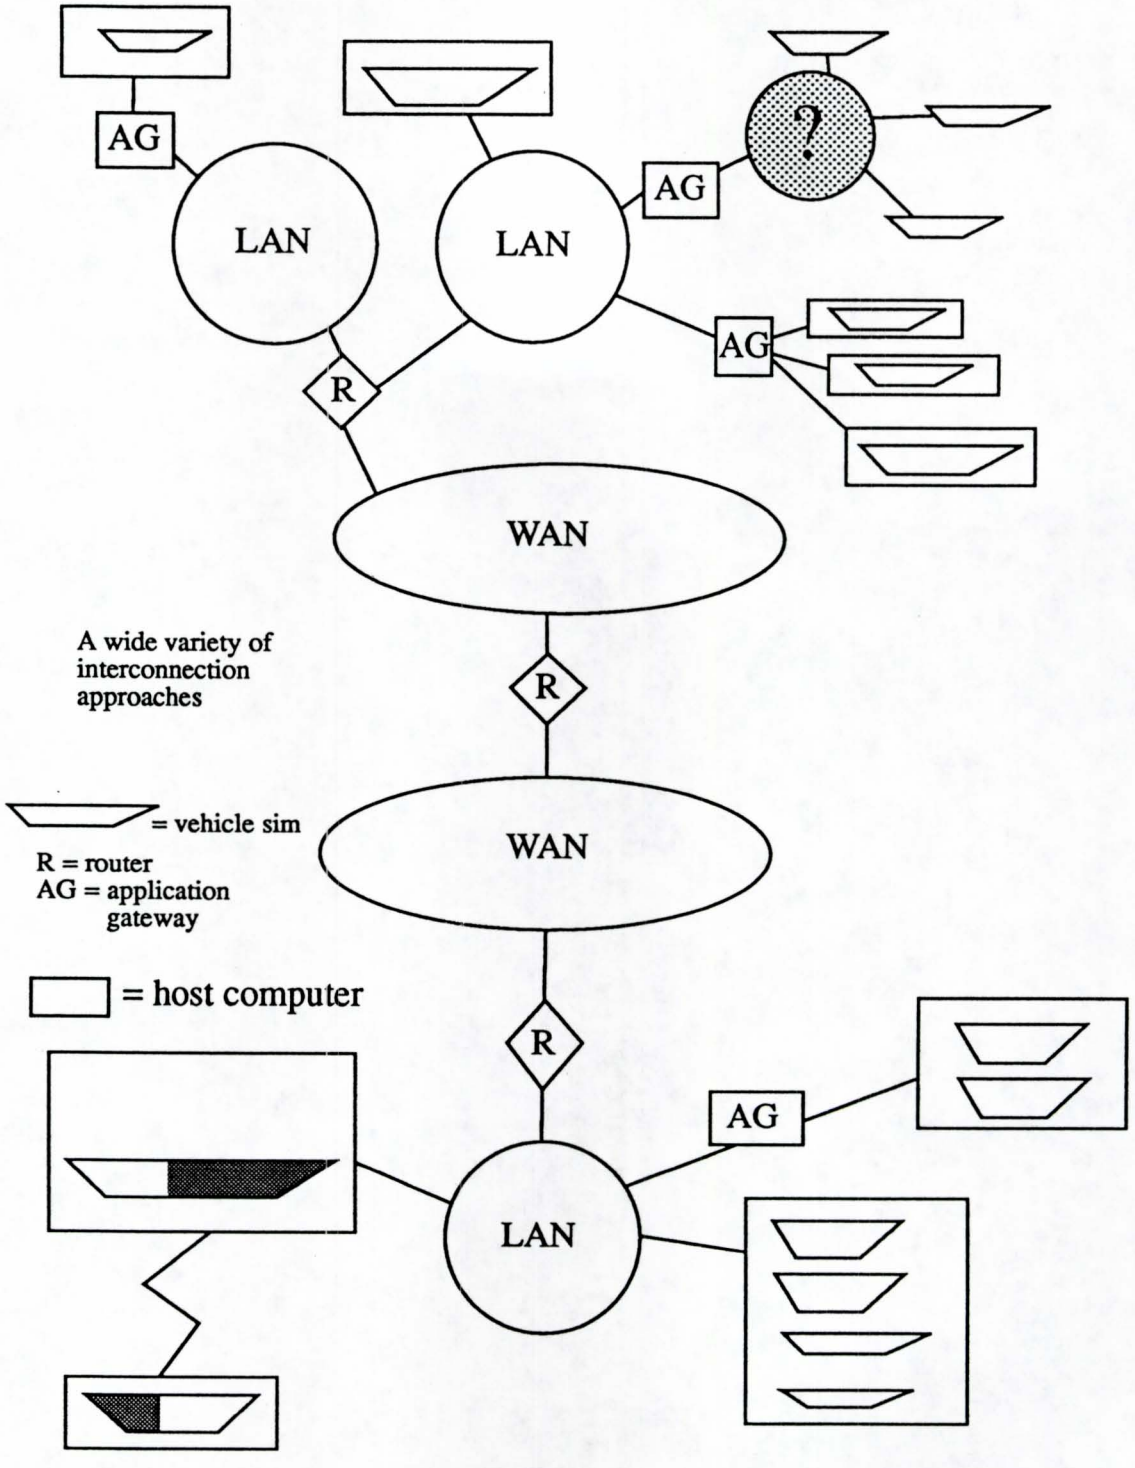

Figure 3. Variety of Interconnection Modes

4.5 Communication Management Requirements

4.5.1 Network Management. The approach to network management is usually dependent upon the type of network employed. Thus, there is a generally recognized and sanctioned way to manage an OSI based network, in the form of OSI network management protocols and service definitions. Similarly, an Internet based network is typically managed by Internet network management protocols. In analyzing the requirements for network management posed by distributed interactive simulation, the conclusion is that the requirements are no different than for a non-simulation network. Hence, the most prudent course of action would be to adopt the network management approach that comes with the protocol suite selected for handling interoperability.

4.5.1.1 Basic Functions. Exercise communications management is a set of facilities to monitor and control the networks that join simulators and other DIS components at a site and sites with each other. By monitoring we mean the ability to determine the status of a network component. By control we mean the ability to set parameters of a network component. The monitoring and control of network components is often referred to as "network management".

DIS requirements for network management are essentially the same as for any other distributed application. One can think of an exercise as having two phases, initialization and operation. During the initialization phase, one would use network management monitoring facilities to check the status of lines, host interfaces, routers, and other network components required for the exercise. Control functions would be used to boot devices with the appropriate parameters, enable interfaces, and so on. The exact set of functions used would depend on the equipment being used, the extent to which its configuration can be changed, and the nature of the network or networks involved.

During operation of an exercise, network management functions would be used to detect and troubleshoot problems. Monitoring functions are used to detect apparent connectivity or equipment failures. Once a problem is detected, operators select appropriate monitoring functions to retrieve parameter values or other information needed to determine the exact cause. Finally, operators can use control functions to reboot equipment, activate alternate interfaces, or take other corrective action. As is the case for initialization, the exact functions used would depend on the nature of the problem, the equipment, and of the networks involved.

It should be noted that some facets of network operation are not typically automated or performed remotely. For example, a network operator might command the use of a dial-up line, but the use of leased lines must typically be arranged for in advance.

Also, while a network operator might command the use of back-up equipment when primary equipment fails, it is sometimes necessary for a technician to remove and replace failed components.

4.5.1.2 Potential Implementation Mechanisms. DIS shall use standard network management protocols to manage the communications infrastructure. Simple Network Management Protocol (SNMP) is a network management protocol frequently used in conjunction with the Transmission Control Protocol (TCP)/Internet Protocol (IP) stack. Common Management Information Services (CMIS)/Common Management Information Protocol (CMIP) is used in an Open Systems Interconnection (OSI) environment. The choice of network management protocol would depend on the other protocols (TCP/IP or OSI) being used in the network.

4.6 Network Security Requirements. Security pertains to the protection of data that is transferred on the network. It may not be necessary for all applications of DIS to require security services. The need for security measures within a given DIS network depends primarily on the requirements of its users. Any implementor of DIS should give careful consideration to the security requirements of the application. For those applications which require some level of security protection, guidance will be provided in future versions of this standard. See Appendix C for a discussion of network security for DIS.

5. DETAILED REQUIREMENTS

5.1 DIS Communication Architecture. Section 4.3 defines the protocol suites that shall be used in the phased approach for the DIS Communications Architecture. In order to guarantee the smooth transition from one phase to the next, this section will address the detail requirements and the migration of each phase.

5.1.1 Phase Zero - Initial Internet Architecture. The Phase 0 architecture is based on Internet network products and communications service which are available today and can be used to support current exercises and early implementations of DIS applications. At each site, there is an Ethernet (IEEE 802.3) Local Area Network (LAN) with a local broadcast capability. For testing, demonstrations, and exercises involving multiple sites, the LANs will be interconnected using a wide area network that can provide the required communications services at those locations. The Phase 0 architecture is shown in Figure 4.

Each simulator has IP with both UDP and TCP. UDP is used to support not only the simulation application, but also the NTP and SNMP. TCP is used to support FTP, Telnet, and SMTP. Telnet is needed for remote debugging and for configuration management chores. FTP is needed for retrieval of databases. E-mail is not really needed on the simulators, but may be convenient to have on some simulation machines.

As part of the standard operation of IP-over-Ethernet, the mapping between IP-addresses (single and multicast) and the corresponding local (Ethernet) addresses, is handled by ARP and by ICMP.

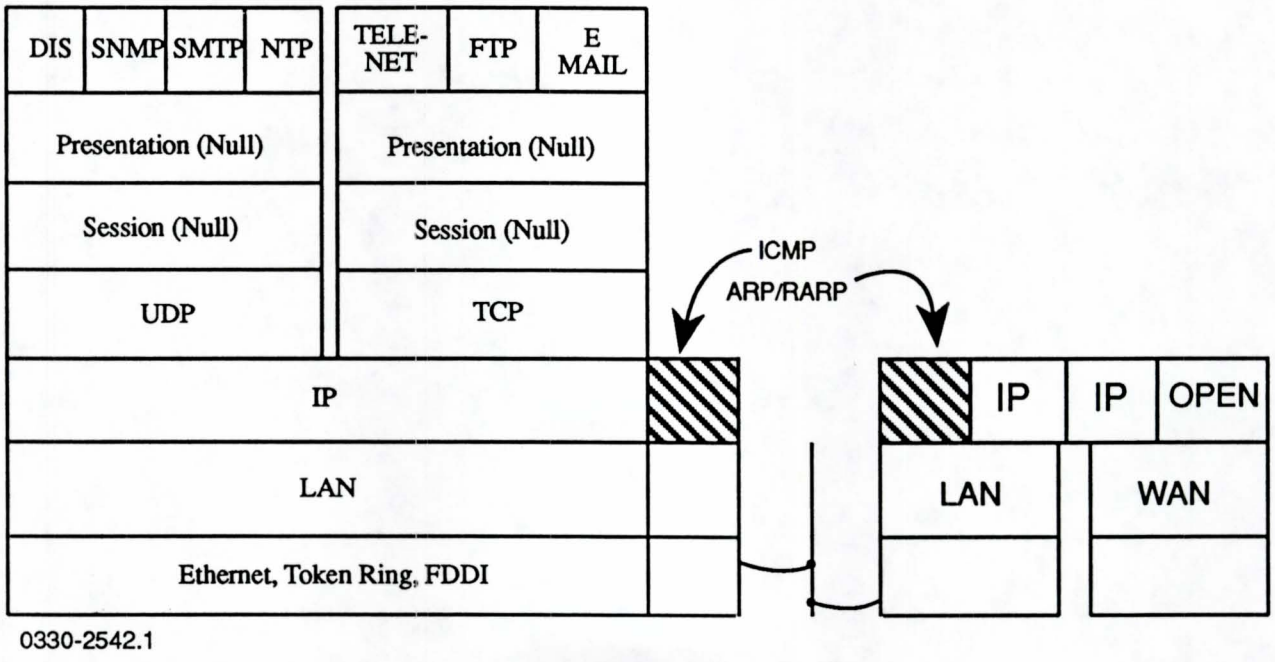

Figure 4. Phase 0: Internet Architecture

5.1.2 Phase One - Hybrid Internet/OSI Architecture. The Phase 1 hybrid architecture is based on OSI network products which are available today. Interim wide area network facilities are used as in Phase 0. The Phase 1 architecture is shown in Figure 5. Phase 1 represents a possible interim transition step for DIS applications that start with Phase 0 and are migrating to Phase 2. In addition, systems that are under development using DIS may start with Phase 1, depending on the timing of their program schedule.

membership management, group initiation, and group communication termination.

The following items are considered developmental based on product availability: multicast (ST-II), group management, and CLTP.

5.1.2.1 Migration Path to Phase One. The transition from Phase 0 to Phase 1 will require protocols at all levels to change. The protocol migration is shown in Table VI.

TABLE VI. Transition From Phase 0 to Phase 1

Phase 0 Internet Standards		Phase 1 OSI/Hybrid Standards
SNMP	-->	CMIP
NTP	-->	modified NTP
TELNET	-->	VTP
FTP	-->	FTAM
SMTP	-->	MHS
TCP	-->	TP4
UDP	-->	CLTP
IP	-->	CLNP
Open	-->	ST-II
		XTP

5.1.2.2 Migration Process to Phase One. Two types of milestones can be used to determine when the transition from Phase 0 to Phase 1 should occur. The first set of milestones are the "maturity criteria" which includes: protocol maturity, product availability, product maturity, product cost, and implementations. The second set of milestones, "risk criteria", include: required development and development cost.

5.1.3 Phase Two - Full OSI/GOSIP Architecture. The proposed Phase 2 architecture incorporates the future OSI multicast protocols into the GOSIP compliant network. The Phase 2 architecture is shown in Figure 6.

TABLE VII. Transition From Phase 1 to Phase 2

Phase 1 OSI/Hybrid Standards		Phase 2 OSI Standards
CMIS/CMIP	-->	CMIP w/multicast extensions
NTP	-->	OSI Time Protocol
VTP		
FTAM		
MHS		
TP4		
	-->	TP5
CLTP		
CLNP		
ST-II or XTP	-->	MPMC

5.1.3.2 Migration Process to Phase Two. The transition from Phase 1 to Phase 2 should be based on the multicast protocols being standardized by ISO. After adoption, the functionality of the new protocols should be demonstrated and tested in prototype implementations. Once testing is completed and the protocols have been validated, the new architecture should be assessed by the maturity criteria established for the Phase 0 to Phase 1 transition.

6. NOTES

6.1 Role of the Communication Architecture. The OSI Reference Model is probably the most widely referenced communication architecture model, and we adopt its use here. Under this model, the communication interconnection problem is broken down into seven layers, each with specific responsibility in carrying out part of the overall communication integration. The development of this reference model was in large measure motivated by and patterned after the success of the DARPA Internet program, which was the pioneer of the general machine interconnection technology base. Along with the development of the reference model, OSI has developed a series of protocols which in some cases mirror comparable entities in the Internet, and in other cases extend and formalize concepts only primitively developed by the Internet program. Currently, there are two dominant suites of protocols (Internet and OSI) which fit within the Reference Model communication architecture and are instantiations of a solution to the general communication interoperability problem. They differ in their details, in their maturity, in their number of options, in their flexibility, in their performance, in their number of currently available commercial products, in their number of fielded systems, and in their organizational support, among other factors.

Within level 3 of this reference model there is functionality which is key to a generalized interconnection model. This so called "internet level" provides for packets of information to be transparently delivered from system to system across almost arbitrary interconnections of local and wide area networks. By adopting the low cost conventions of providing for remote delivery even when delivery is actually local, and through the provision of gateway processors linking the local and wide area networks, a single approach (from the application perspective) can handle both the local and global cases, as well as transparently handle any needed change from one to the other. Under this approach, any reasonable selection for the layers below will be perfectly acceptable and work. These decisions can be handled locally on a case by case basis or by policy over some administrative domain if deemed appropriate. Since building to the level three interface admits a mixing and matching approach to all of the levels below without sacrificing interoperability. Levels above do need to be matched. However, in our immediate case, handling interoperability for these functional elements has already been subsumed into the current DIS PDU standard. This approach ensures the maximum interoperability with the minimum of specification and new development.

6.2 Generalized Functional Architecture. The Communications community thinks in terms of a vertical layering of communications functions. The accepted nomenclature (adopted

by the International Standards Organization) refers to seven layers. Table VIII identifies the levels and illustrates their meaning in the context of the networking of simulators.

TABLE VIII: Seven Layer OSI model applied to Simulation

Number	Name	Content
7	Application	Kind of data exchanged (position, orientation,...) Dead reckoning rules. Rules on determining hit or miss and damage.
6	Presentation	of position (local vs geocentric coordinates) of orientation (Euler angles, Quaternions, SPV) units: English, metric, degrees, BAMs... encoding: integer vs float, big vs little endian.
5	Session	Procedure for starting and ending an exercise. Rules for joining and leaving an exercise. Freeze.
4	Transport	Addressing from end user to end user. Assuring communications reliability, if required.
3	Network	Addressing information from node to node.
2	Link	Framing of information on a physical link. Flags, zero bit insertion. Conflict resolution.
1	Physical	Wire, optical fiber, radio transmission. Voltage levels, impedance values, clock rates.

The DIS PDU document addresses levels 5 through 7. It does so without separating the levels. Levels 4 and below are defined in the remainder of this section.

There are a variety of existing protocols and interfaces which populate the functional areas for levels 1-4. The two most prominent suites of protocols which are collectively put forth as solutions to the interoperability problem are the DoD (Internet) suite and the OSI (GOSIP) suite. At this stage of evolution, the two are conceptually similar, but vary considerably in the details and in maturity. Both suites emphasize the network transparency from level 3 and above, as discussed previously. This means that one simulator is completely isolated from the selections made at levels 1 and 2 for every other simulator or

collection of simulators, by adopting one of the "internetwork" layer standards as the base level for interoperability. This provides the freedom to delegate to local decision making the protocols used for the lower levels (assuming the selections conform with overall, real time performance objectives). The current real work of this document focuses essentially on levels 3 and 4. A plan which starts from the more mature Internet suite and evolves as appropriate over time toward the GOSIP suite is the most prudent path at this time.

6.3 OSI Compatibility. The OSIRM was developed in 1977 by the International Organization for Standardization in response to the need to interconnect heterogeneous computers. OSI defines a framework for the interaction of users and applications in a distributed environment.

The Government Open Systems Interconnection Profile is the U.S. Government program for adoption of OSI across all Federal agencies. The purpose of GOSIP is to provide: networking connectivity, through GOSIP network architecture; interoperability, through standard "profiles" of OSI protocols; and competition, through focus on small number of subnetwork technologies and interoperable applications.

6.3.1 Desired Extensions & Additions to OSI. The DIS multicast requirement is not presently found in OSI, however, work is underway to develop these standards. Currently, there are six American National Standards Institute (ANSI) working groups participating in the development of multicast standards: X3T5.1 (OSI Architecture), X3T5.4 (OSI Management), X3T5.5 (OSI Upper Layers), X3T5.7 (OSI Security), X3S3.3 (Network and Transport Layers), and X3S3.7 (Public Data Networks). The goal of the Multipeer/Multicast (MPMC) effort is to develop a complete set of standards which will provide DIS with a full range of multicast functions and capabilities.

To include multipeer/multicast in the OSIRM, are required the following extensions and additions to current ISO standards:

OSI Reference Model, including Part 1: Multipeer Addendum to the Basic Reference Model, Part 2: Security, Part 3: Naming and Addressing, and Part 4: Management Framework;

Application Layer, including the Application Layer Structure and Extended Application Layer Structure;

Transport Layer, including Connectionless Transport Protocol and Connection Oriented Transport Protocol;

Network Layer, including Connectionless Network Protocol, X.25 Packet Level Protocol for Data Terminal Equipment, X.25 to Provide OSI Connection Mode Network Service;

Routing, including End System to Intermediate System Protocol (ES-IS), Intermediate System to Intermediate System Protocol (IS-IS), and Intra-Domain Routing Protocol (IDRP);

Data Link, including X.25 LAPB Compatible DTE Data Link Procedures; and

Network Management, including Common Management Information Service (CMIS), Common Management Information Protocol (CMIP), Systems Management Overview, OSI System Management, and Structure of Management Information

Other extensions to the OSI architecture include a time protocol. The Navy's SAFENET project is currently modifying the Internet Network Time Protocol (NTP), designed for Transmission Control Protocol/Internet Protocol (TCP/IP) networks, to run on an OSI stack. This work is being forwarded to the appropriate ANSI working groups.

APPENDICES

APPENDIX A

REPRESENTATIVE PROFILES

10 Communication Architecture Profile For Phase 0

This communication architecture uses the DoD family of protocols is based on IP, with TCP (MIL-STD-1778) for reliable communication and UDP for real-time communication.

10.1 Exercise Management. In each simulation site there are several simulators and a Local Exercise Manager (LEM), all interconnected by a LAN, to which we refer as "Ethernet", even though it may be FDDI or other LANs.

The LEM is a software module, which does not need dedicated hardware, and may be implemented on any of the simulators, for example. The LEM is in communication with other LEMs, and in particular with the Global Exercise Manager (GEM) for the purpose of coordinating the entire exercise.

After the LEMs agree about the parameters of an exercise they communicate them to all the participating simulators, using a "session-level" type communication. This setup includes the identifications of all the simulators involved in the exercise, their roles, the exact presentation schemes to be used, the exact geographic database to be used, the maximal bandwidth that each simulator is allowed to load on the network, and the IP-multicast-addresses (IPMCA) assigned to the entire exercise.

10.2 Communication Setup. The setup communication, between the LEM and the simulators is conducted by using Telnet over TCP (over IP, over Ethernet). The setup process may use both manual and automatic procedures. As a part of the general setup, database files (e.g., geographic) are loaded, by using FTP (MIL-STD-1780), from designated directories. FTP also operates over TCP (over IP, over the Ethernet). The real-time communication (e.g., of PDUs) is carried by UDP. These packets are encapsulated inside Ethernet packets. The entire configuration is managed (and verified) by using SNMP in the simulators. This allows a remote network management process to check the status of each simulator.

In each case, of the real-time simulation messages are broadcast to all the participants in the exercise, locally over the Ethernet, and remotely over WANs.

It is expected that future simulators will require time synchronization. This may be achieved by using the Internet time synchronization protocol (a.k.a. the Network Time Protocol, NTP), over UDP, on the Ethernet. The time protocol is defined in RFC1119 and RFC1129.

The real-time communication for the support of distributed interactive simulation requires that a given bandwidth is delivered without exceeding a given delay. In practice this

required both broadcast and bandwidth performance guarantees. Phase 0 only has a requirement for broadcast, so multicast is not an issue in this profile. Since these issues (bandwidth+delay and broadcast) are at the network level (level-3 of the ISORM) it is possible to address them at the gateway between the LAN and the WAN. If the WAN provide these services there is no need for this gateway to be involved. However, in the most general case this gateway should handle them. In cases that the Commercial Off The Shelf (COTS) gateways and WANs do not provide this functionality it can be achieved it by adding a front-end to the gateway on a general purpose computer, preferably with two Ethernet interfaces to allow inserting this front-end "in series" with the gateway.

To guarantee interoperability each simulator should comply with the Host-Requirement, as specified in "Requirements for Internet hosts - communication layers" (RFC1122) and in "Requirements for Internet hosts - application and support" (RFC1123). This would guarantee the "invisible support" as required for interoperability (including ARP, re-direct, etc.). A good source of information is "Perspective on the Host Requirements RFCs" (RFC1127).

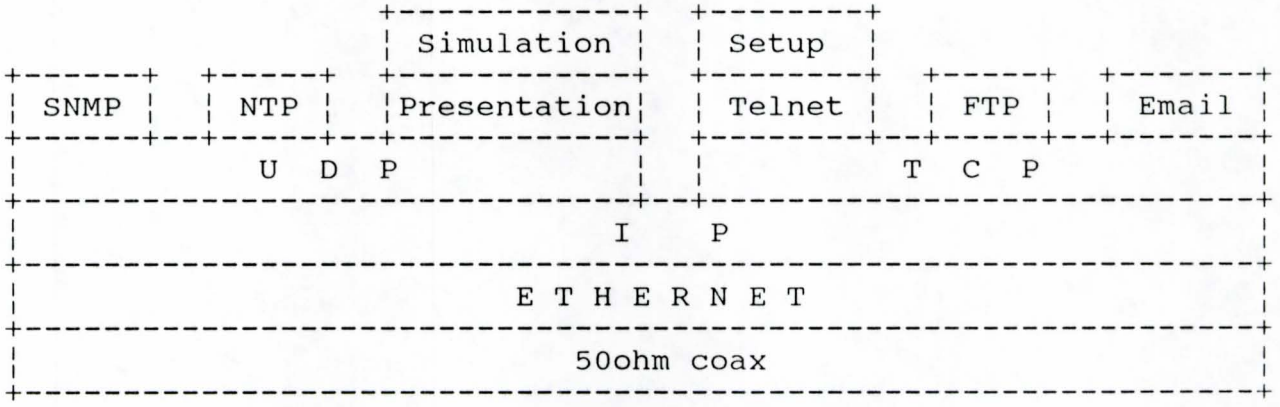

Figure 7. The Protocol Structure in the Simulators

Notes:

* The ISORM level of the Ethernet is 2, of IP is 3, and of TCP and UDP is 4. The ISO level of the simulation is 7, and its presentation level is 6.

* The simulation session level, 5, does not show explicitly.

* Each of Telnet, FTP, and E-mail span levels 5 through 7.

20 Security Architecture Profile for Phase 0.

The Phase 0 Distributed Interactive Simulation scenario consists of multiple IP-based simulators, in each of several sites, participating in a real-time exercise. Each site uses a LAN like an Ethernet or FDDI which could be interconnected securely via a black WAN to the other sites. Some of the simulators, at some sites, will participate in a given exercise, at a system-high level, using a red LAN, i.e., all the subscribers to that LAN will have access to all the traffic associated with the exercise. The system-high operation may interfere with the ability of a site to participate in several independently secure exercises simultaneously.

At each site there will be a Local Exercise Manager (LEM). The LEM is a software process that will participate in the set up of that exercise and would know which other sites participate. It would also distribute specific parameter values for each exercise (such as the bandwidth allocated to the site, update frequency, the choice of coordinate system, and the version of the geographic database in use). It is anticipated that some manual set up will be required initially for each exercise, either of a new type, or with a new set of participants.

The LEM will use TCP to communicate with each of the simulators, with LEMs at other sites, and with the GEM, the Global Exercise Manager. The individual simulators will use TCP (with FTP) for reliable loading of critical files, such as programs and geographic databases. The real-time communication among the simulators during the exercise will use UDP.

For security considerations the LEM is divided into separate black and red components.

The WANs may, or may not, support bandwidth reservation and multicast, such as the TWBnet. For generality the assumption is made that the WANs support neither. The example addressed here represents the worst case scenario. Two-by-two approaches have to be considered:

(H) System-high operation
(C) Controlled access to exercises

and

(2) Network security at Level-2
(3) Network security at Level-3

For the total of four approaches:

(H.2): System-high operation based on network security at Level-2
(H.3): System-high operation based on network security at Level-3

(C.2): Controlled access based on network security at Level-2

(C.3): Controlled access based on network security at Level-3

We assume that initially (H.X), system-high operation at Level-2 or at Level-3, would be used with encryption devices per site, regardless of the number of simulators there. This is done because of the understanding that eventually the system will migrate to (C.X), controlled local access to exercises, with individual encryption devices for all the simulators participating in the exercise, and for the red LEM.

The migration from (H.X) to (C.X) will require some development costs, with no significant implications for the architecture, the software, or the hardware. This will require the addition of another network encryption device and a red LEM per local physically separated community of interest (i.e., using a separate encryption key when leaving their system high LAN).

For maximal flexibility there may be a dedicated encryption device for each host. However, the budget may not be able to support a large number of such security devices since the cost of NSA-approved encryption devices does not follow the trend of consumer electronics (commercial computers, included) and does not decrease dramatically annually.

In principle, approaches (H.2) and (H.3) are similar, even though they differ in many details. For brevity we describe here only (H.2). This is not a recommendation to prefer (H.2) over (H.3), (C.2), or (C.3).

The actual choice among (H.2) and (H.3) and among the various devices available on the market (through the Commercial COMSEC Endorsement Program (CCEP), of National Security Agency (NSA)) could be made only after objective engineering tradeoffs are taken into account. Considerations of interest include:

* Performance (both in bps and pps)
* Keying
* Security management
* Multicast (and crypto synchronization),
* Real-time behavior (and packet loss)
* Security doctrine (modes/policy)
* Configuration
* Error response
* Error characteristics
* Scalability
* Network management
* Interaction with other protocols and services (e.g., Redirect, ARP, and SQ)
* cost.

20.1 Approach-(H.2): System-High at Level-2. Figure 8 shows the configuration. The local system-high LAN supports the simulators. The encryption is performed at the Ethernet level, by Xerox's XEU.

The LEM job is divided among two units, the Black-LEM (B-LEM) and the Red-LEM (R-LEM). The configuration information that is entered into the B-LEM is then distributed over airgaps to the R-LEM, to the Ethernet bridge, and to the KML (for the XEU). It is also distributed to the Reno (a.k.a. XET, the front end for the gateway, see later) over a black-LAN. The R-LEM is on the red system-high LAN. The B-LEM cannot be on the red LAN because it has to provide information to the black Reno. Neither part of the LEM requires a dedicated hardware unit. Both are software modules that can be run at a user level, the B-LEM on the black side (e.g., in a Reno) and the R-LEM on the red side (e.g., on any simulator). R-LEMs communicate with other R-LEMs, only in a secure mode over the WAN once the appropriate keying arrangements are made. B-LEMs can always talk with each other over the WAN in an unsecured mode.

Each exercise has its own IP MultiCast Address (IPMCA) (and Ethernet MultiCast Address (EMCA)) used for all the communication with the other simulators, at the other sites. The transmission scenario, after the initial set up is as follows. Please consult the diagram of (H.2), below.

20.1.1 Example Scenarios. Each simulator prepares data for transmission, in red UDP/IP packets, addressed to the IPMCA assigned to the exercise. These red packets leave the simulator into the red LAN using the EMCA[9]. (The IPMCA was provided earlier by the R-LEM to the simulators (using TCP).)

An Ethernet-bridge, B, on the LAN transfers only the red packets with the EMCA to the XEU. The EMCA was provided earlier to the bridge by the B-LEM[10].

The XEU verifies that the originating host (identified by its individual Ethernet address) is indeed authorized to send red packets to this EMCA. The EMCA and the list of the Ethernet addresses of all the authorized simulators should have been provided earlier by the B-LEM via the Key Manager/ Loader (KLM) to the XEU. The XEU generates the black version of these red

---

[9] The EMCA is derived from the IPMCA according to the standard IP operation procedures as defined in RFC1112, "Host extensions for IP Multicasting", by S.E. Deering, Aug-01-1989.

[10] the exact way for doing that depends on the particular bridge in use.

packets, by encrypting them in their entirety, by adding its own headers (for crypto synchronization, etc.), and also by adding black Ethernet headers with the original EMCA in the black.

The KML is physically located at one site (per exercise). It generates physical keys that have to be distributed to the sites. New keys have to be distributed to all sites from the KML at least once a year unless security compromises occur. Whenever a site adds devices with new Ethernet addresses (e.g., new simulators) it would require a new key to be generated for that site only for the exercises in which these new devices are to participate.

The packets are then given by the XEU, over the black Ethernet, to the Reno, that operates totally in the black[11].

The Reno recognizes the exercise packets by their EMCA and encapsulates them in IP (or ST) packets, as required by the gateway to the WAN. These packets are IP-addressed (or ST-addressed) to Renos at the other participating sites. Being a general purpose computer the Reno can easily be programmed to set the priority field, or to open a connection with bandwidth reservation, as required. The black Reno packets are sent over the black LAN to the black gateway for transmission over black WANs. That gateway should be a COTS unit, optimized for the WAN in use.

It is possible for a sending Reno to strip off the black Ethernet headers for the transmission over the WAN, to be re-inserted by the receiving Reno. This saves some WAN bandwidth at a cost of additional processing (a typical engineering tradeoff). If an IP-WAN does not support IP-multicast, the Reno can replicate the black packets to achieve the desired multicast. The IP-addresses to which the packets should be forwarded were provided earlier by the B-LEM to the Reno, over the black LAN.

All the knowledge about the WAN is in the Reno, because the WAN connects (practically) only the Renos. Hence, issues such as the choice of IP vs. ST for the long haul are separated by the Renos from the rest of the system that uses IP for end-to-end communication. Note that because of the security there is no

---

[11] The Reno (a.k.a. XET is a software package developed by Xerox, that can run on any general purpose computer (e.g. any 386 system or a SPARC). The function of the Reno is to serve as a front-end for the gateway, by encapsulating the black Ethernet packets (produced by the XEU) inside black packets suitable for handling by the WAN, and by performing the inverse task at the receiver end. If needed, the Reno may perform the gateway front-end tasks described in the Architecture-Profile section above.

direct interaction between the end-to-end IP and the IP that is used over the WAN. These two IPs do not have to share even their address space.

Therefore, changing the WAN (say from a general IP-network, to a TWBnet, to its future successors, and then to a gigabit speed IP-network) could be handled by changes only to the Reno.

As advances in technology (and especially in DoD procurement) make better networks available, only the gateway selection and some Reno software may have to be modified, isolating the local site, the simulators, the LEMs, and the encryption gear from the need to adapt to the upgrading of the WANs.

The reception scenario, after the initial set up is as follows. Please consult the diagram of (H.2), (Figure 8) below.

The black WAN deliver black packets to the Reno over the black Ethernet. The Reno recognizes them, by their IPMCAs (or by their ST addresses) as belonging to a particular exercise, de-encapsulates them from the IP (or ST) packets, and gets the black Ethernet packets with the appropriate headers as expected by the XEU. These packets are delivered over the black Ethernet to the XEU, that decrypts them and recovers the original red Ethernet packets, with the original EMCA.

The XEU transmits these red packets on the red system-high Ethernet, through the bridge, and makes them available to all the local simulators, where the Ethernet headers are discarded, and the original red UDP/IP packets are received and processed.

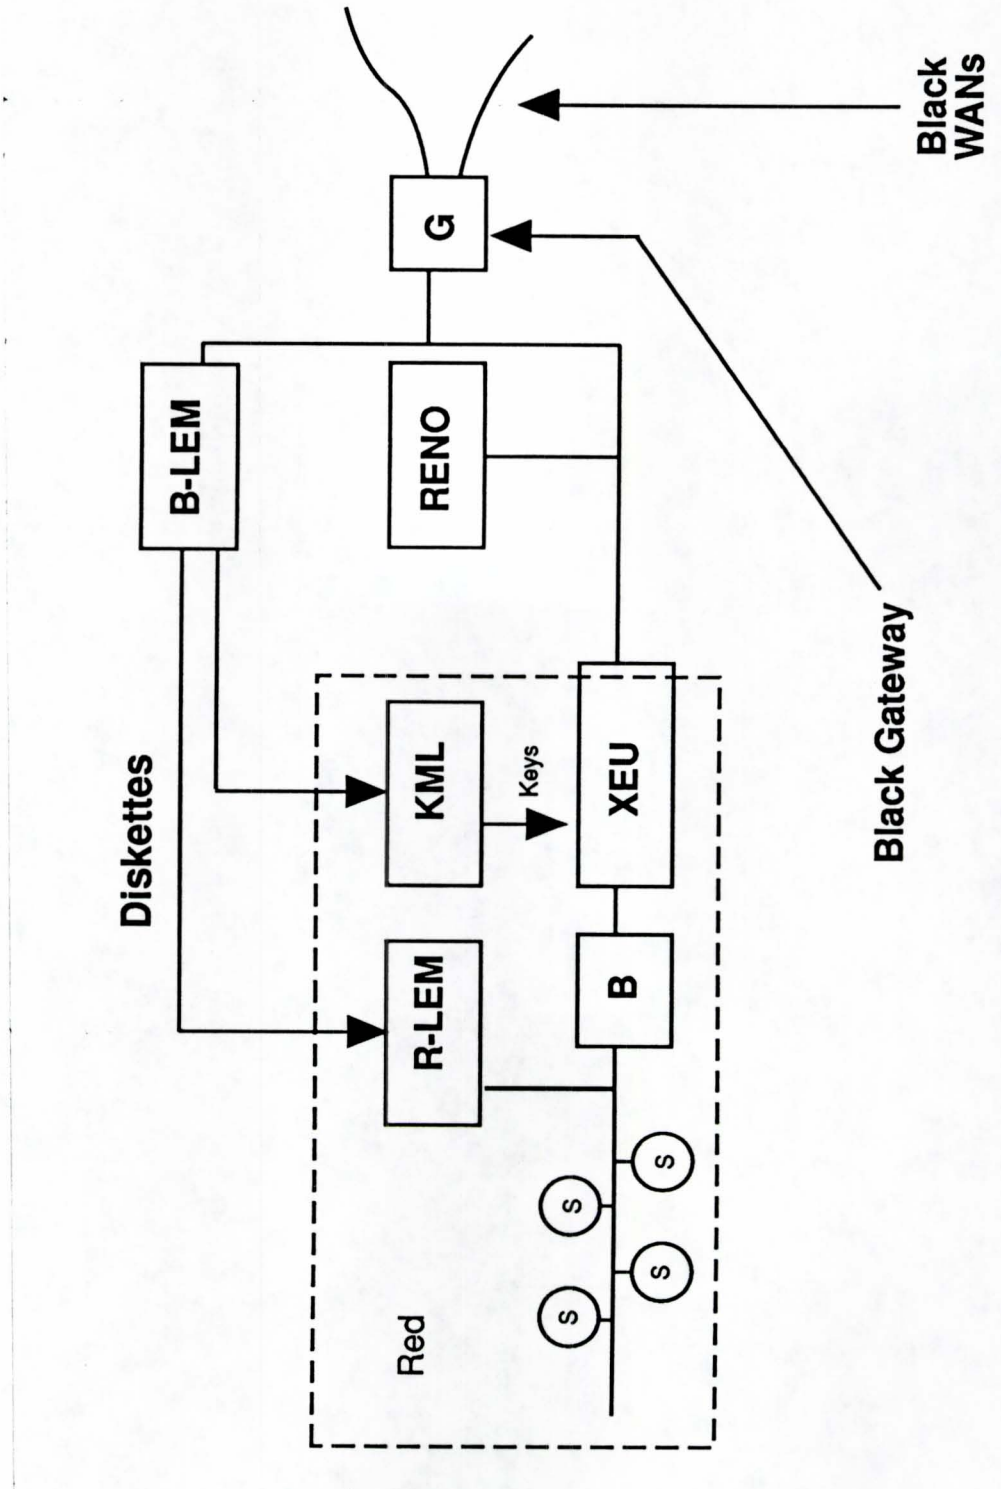

Figure 8. Diagram of (H.2): Net-Security at Level-2, System-High

Notes for the (H.2) diagram:

- Diskettes are used to carry information from the B-LEM to the R-LEM and to the KML, if it's at that site. The setting of the bridge depends on the particular bridge in use.

- The KML exists only in a few locations. Only one KML is in charge of any exercise. It generates the keys for all the participants in that exercise, and from it they have to be securely distributed to the various sites.

- If needed, for performance, the Reno may use two Ethernet interfaces to allow splitting the black Ethernet, with one segment between the XEU and the Reno, and another for the WAN gateway(s), the B-LEM, and the Reno. This makes the LAN, the bridge, the XEU, the Reno, and the Gateway to be "in-series", as shown in the following diagram.

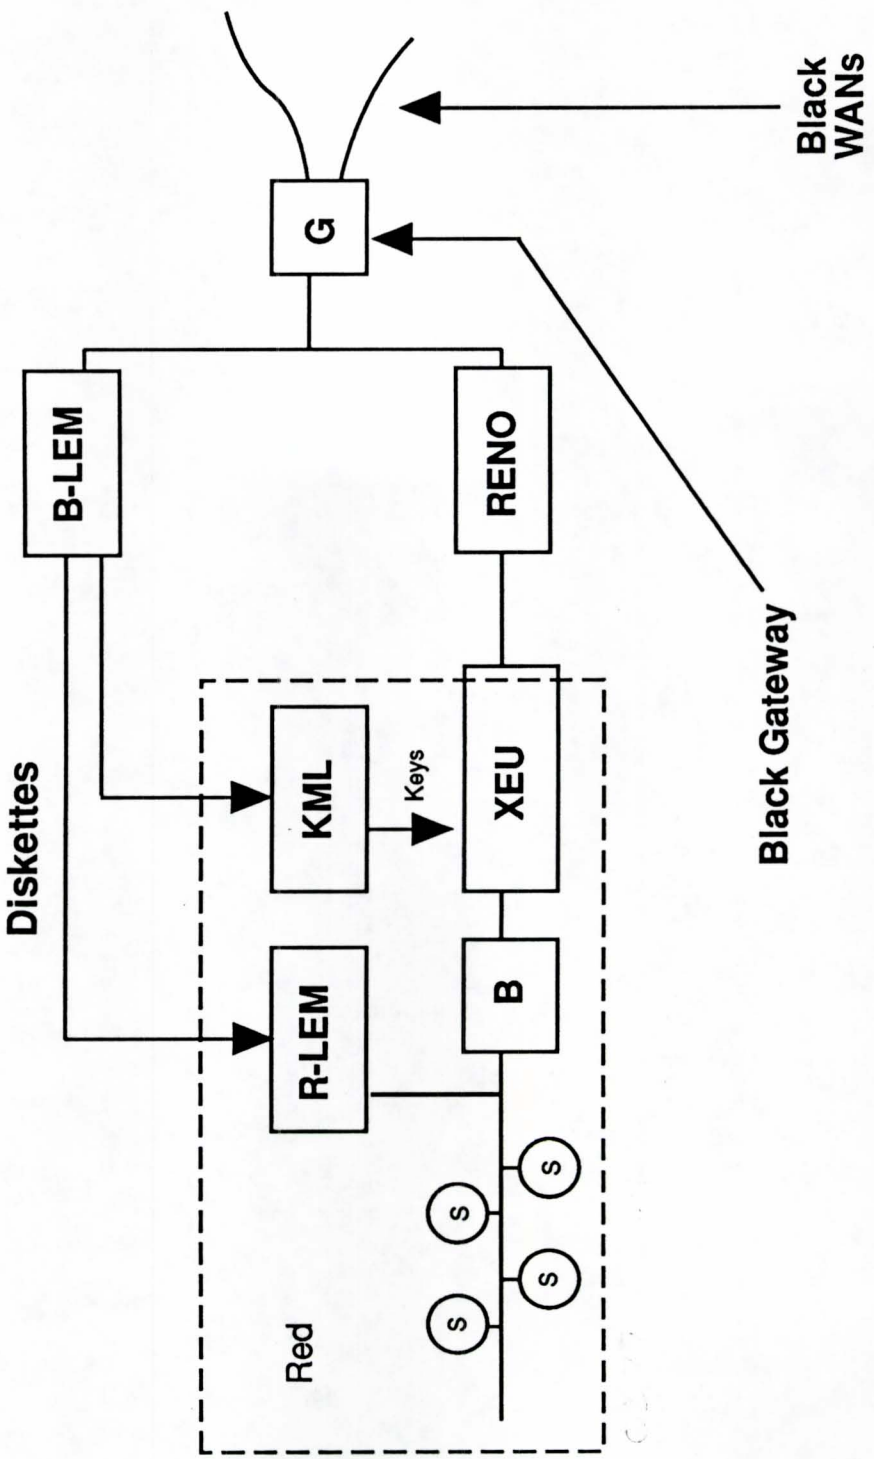

Figure 9. (H.2): Net-Security at Level-2, System-High with Reno in-series

APPENDIX B

BANDWIDTH ESTIMATION PROCEDURES

10 Factors which influence DIS bandwidth. There are a number of factors which have a major influence on DIS bandwidth. At the very highest level, they include:

Total number of entities
Mixture of entity types.
Type of exercise or scenario
Choice of dead reckoning algorithm (and positional/angular thresholds)
Security requirements

For the current set of approved DIS PDUs, the majority of network traffic will to Entity State PDUs (ESPDUs). ESPDUs are required to be sent at some minimum rate (e.g. every 5 seconds) by every entity and may be sent much more frequently depending on entity dynamics. The start-up of a session will also see high traffic but that is deterministic. The PDUs used to initialize an exercise or entity (such as the recommended Activate PDUs) represent a significant amount of data to be sent via the net, but they can be transmitted at a controlled rate. In the near term, the inclusion of Emitter PDUs may add a significant traffic load to the network, depending on the degree of electronic warfare (EW) present in a given exercise. Similarly the future inclusion of simulated tactical communication links (both voice and data) will undoubtedly have a substantial impact on bandwidth.

In addition to the above there are also additional bandwidth requirements due to communications "overhead". A given PDU of "n" bits in length requires the addition of both headers and trailers in order to satisfy routing and data integrity requirements. The proposed UDP/IP protocols add 32 bytes (8 for UDP and 24 for IP). The underlying media adds further overhead, such as FDDI's 20 to 28 bytes of preamble, header and trailer information. A method to reduce this load is to concatenate PDUs at the application layer such that the overhead bits are applied to groups of PDUs rather than to every PDU. This approach, however, imposes an additional computational load on each host. This trade-off of processing load vs network traffic requires further study before serious recommendations can be made.

Another source of "overhead" traffic are security measures. The degree of overhead depends on at what layer (of the OSI seven layer stack) the security measures are implemented.

10.1 Estimating Exercise Bandwidth Requirements. In general, there is no single set of formulae for accurately estimating the bandwidth requirements of any given DIS exercise since, by nature, they have a combination of man-in-the-loop and non-deterministic simulated adversaries. As such, each entity in a given exercise generates network traffic at a varying rate. The rate varies depending on the particular involvement of that

entity with others. For example any vehicle that is in transit to or from its assigned duty area will exhibit very predictable dynamics and therefore generate low network traffic. Conversely, an entity entering into conflict or close cooperation with another will typically generate a high level of traffic. In both cases the traffic is a result of the frequency at which the PDUs are generated, while the size of the individual PDUs remain relatively stable. Estimating sizes of PDUs for selected entity types is a comparatively straightforward process while estimating the frequency at which they are generated is fairly complex and more subjective.

As stated earlier the Entity State PDU will be the main source of network traffic. There are currently nine other PDU types required by the DIS standard, with several others recommended. Of the nine required, six are related to logistics (e.g. repair and resupply) and are expected to occur so infrequently as to have little or no effect on network bandwidth requirements. Another, the Collision PDU, also falls into this category. The remaining two are the Fire PDU (FPDU) and Detonation PDU (DPDU), and conceivably can occur frequently enough at certain stages of battle to be considered in bandwidth calculations. In addition, the Emitter PDU (EPDU), one of the emerging recommended messages, is likely to be a major contributor in the near future. These four PDU types have the following formula for determining their sizes (in bits):

PDU	FORMULA	REMARKS	
ESPDU	1152+128A where	A=	# of articulated part records
FPDU	704		
DPDU	800+128H	H=	# of articulated parts hit
EPDU	192+E(160+B(304+96T))	E=	# of emitters
		B=	# of beams per emitter
		T=	# of targets per beam

Given the above, it is possible to estimate the PDU sizes for classes of entity types. For example, for a given type of tank the minimum number of articulated part records may be 5 (azimuth and azimuth rate for turret, elevation of the barrel, and up/down position for two hatches) and the number of emitters 1 (laser range finder). For a fighter aircraft the number of articulated parts could easily be 20 (8 weapon stations, 2 drop tank stations, 6 vertical control surfaces, 2 horizontal control surfaces, landing gear, and speed brake) with 3 emitters (radar, jammer, and laser designator). Similar assumptions can be made regarding surface ships. The following table presents estimates of PDU sizing for these three classes of entities (without any overhead bits).

TABLE IX. PDU Sizing Estimates

ENTITY CLASS	A	H	E	B	T	ESPDU	FPDU	DPDU	EPDU
TANK	5	1	1	1	1	1792	704	928	752
AIRCRAFT	20	2	3	1	2	3712	704	1056	2160
SURFACE SHIP	50	5	10	1	5	7552	704	1440	9632

The next step in estimating the bandwidth requirements of a given exercise is to approximate the rates at which each entity class will issue each of the above PDU types. Since this rate can vary a great deal within a given exercise, one method of estimation is to give values representing some average low and high rates. The final step is to determine the number of each major entity type which will participate in the exercise. Given all of these factors, the determination of a range of probable network traffic can be easily calculated. Figure 10 presents an example of such an analysis for three different types of exercises. The examples include tactical voice and data links as sources of network traffic (65 Kbs for each voice channel and actual values for Link-4A, Link-11, and Link-16). Figure 11 presents the results of the same analysis in graphical format.

SAMPLE PDU SIZING

	A	*H*	*E*	*B*	*T*	*ESPDU*	*FPDU*	*DPDU*	*EPDU*
TANK	5	1	1	1	1	2220	1132	1356	1180
AIRCRAFT	20	2	3	1	2	4140	1132	1484	2588
SURFACE SHIP	50	5	10	1	5	7552	1132	1868	10060
						OVERHEAD BITS/PDU=			428

SAMPLE RATES PER ENTITY TYPE PER PDU TYPE

	LOW RATE (HZ)				*HIGH RATE (HZ)*			
	ESPDU	*FPDU*	*DPDU*	*EPDU*	*ESPDU*	*FPDU*	*DPDU*	*EPDU*
TANK	0.2	0	0	0.2	2	0.1	0.1	1
AIRCRAFT	0.2	0	0	0.2	8	0.1	0.1	4
SURFACE SHIP	0.2	0	0	0.2	1	0.1	0.1	2

SAMPLE EXERCISE TRAFFIC ESTIMATES

% ENTITIES AT HIGH RATE	*0%*	*20%*	*40%*	*60%*	*80%*	*100%*
% ENTITIES AT LOW RATE	*100%*	*80%*	*60%*	*40%*	*20%*	*0%*
SAMPLE EXERCISE #1						
600 TANKS	408,000	1,030,656	1,653,312	2,275,968	2,898,624	3,521,280
100 AIRCRAFT	134,560	982,320	1,830,080	2,677,840	3,525,600	4,373,360
10 TACTICAL VOICE LINKS	650,000	650,000	650,000	650,000	650,000	650,000
TOTAL TRAFFIC	*1,192,560*	*2,662,976*	*4,133,392*	*5,603,808*	*7,074,224*	*8,544,640*
SAMPLE EXERCISE #2						
24 SHIPS	84,538	201,896	319,254	436,612	553,970	671,328
50 AIRCRAFT	67,280	491,160	915,040	1,338,920	1,762,800	2,186,680
3 TACTICAL DATA LINKS	237,500	237,500	237,500	237,500	237,500	237,500
TOTAL TRAFFIC	*389,318*	*930,556*	*1,471,794*	*2,013,032*	*2,554,270*	*3,095,508*
SAMPLE EXERCISE #3						
200 TANKS	136,000	343,552	551,104	758,656	966,208	1,173,760
100 AIRCRAFT	134,560	982,320	1,830,080	2,677,840	3,525,600	4,373,360
5 TACTICAL VOICE LINKS	325,000	325,000	325,000	325,000	325,000	325,000
3 TACTICAL DATA LINKS	237,500	237,500	237,500	237,500	237,500	237,500
TOTAL TRAFFIC	*833,060*	*1,888,372*	*2,943,684*	*3,998,996*	*5,054,308*	*6,109,620*

Figure 10. Sample Network Traffic Analysis

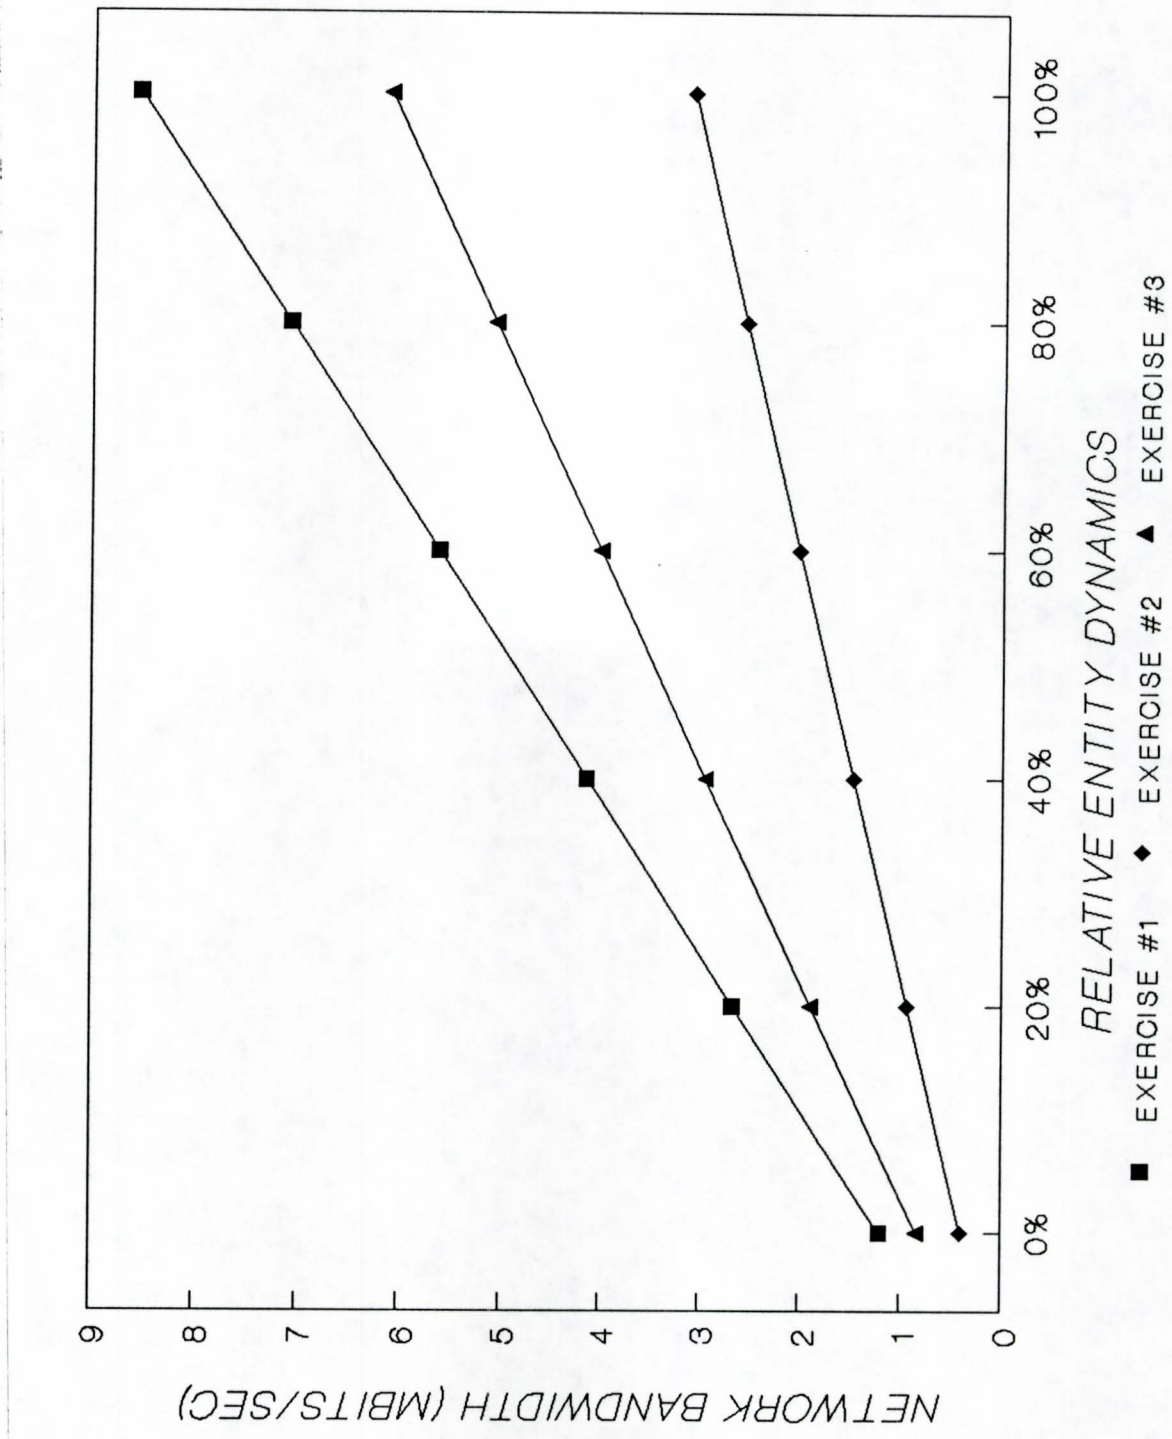

Figure 11. Sample Exercise Bandwidth

As shown in the figures, the network traffic can vary as much as ten to one depending on the relative number of entities which are in a high dynamic environment. The low end of the charts are certainly the minimum bandwidth requirements since they are based on all entities in a quiescent mode (i.e. ESPDUs only once every 5 seconds). The high ends of the charts are more subjective since it makes assumptions as to the maximum rates each entity type will exhibit, but in any case are not probable since they represent all entities simultaneously engaged in heavy combat. Given those assumptions, such charts may be used as a guide to sizing a network for any type of exercise.

Some final points to be made about the above discussion:

- The sample bandwidth values shown are only for illustration, and should not be used in formal specifications.

- The Emitter PDU used here is in accordance with the latest format proposed by the Emissions Subgroup, not the format shown in the existing version of the DIS specification. This latest version results in less overall network traffic since it is only issued on change of the emitter data (the older version had to be issued at least as often as ESPDUs).

- The analysis does not account for the transitory existence of entities in the form of guided weapons released by various types of weapon systems. These will add still further traffic and will most likely be present during the same period of time where high vehicle dynamics are also occurring - during engagement of groups of opposing forces.

- No data compression is assumed. For reduction of PDU traffic it is not considered viable at this time due to the large computational load it would place upon each entity host computer. It should be seriously considered for tactical voice links, however, since the task is simplified by the fact that the computer does not need to know what is actually in a voice message; the compression and decompression can then be done by hardware, external to the computer system. The signal can be compressed by hardware at the source, sent over the network in its compressed form, and fed directly to decompression hardware at the listener. A variety of commercial devices currently exist to support this, some offering time stamping of the audio stream for synchronization. Standards are emerging with the growth of multimedia computing technology, and could be considered for use in the DIS application.

10.2 Estimating Traffic in terms of PDUs and Packets per Second. Once it has been established that the underlying media is capable of handling network traffic (i.e. from the bits per second standpoint), the next figure of merit to analyze is that of the number of messages to be handled in a given unit of time. This factor provides a relative figure-of-merit for the type of processing power necessary for a given set of communications protocols.

Figure 12 presents another look at the sample exercise data presented earlier. Here, in addition to the total traffic for each exercise in bits per second there are two additional fields showing the number of PDUs per second as well as packets per second. The following assumptions were made in developing these estimates:

1. Packet length is the standard 1500 byte Ethernet datagram size.

2. PDUs can be concatenated such that each packet contains several PDUs. The PDU sizes here are taken to be without overhead bits; a single set of overhead bits is applied to the entire packet.

3. The "host" composing the packets always waits until the 1500 byte limit is filled. In actual practice the efficiency factor will probably be lower (to avoid excessive latency), resulting in an actual packet rate that falls somewhere between the two values (PDUs/sec and packets/sec) shown.

4. Voice packets are produced at 32 Hz, Link-11 and Link-4A at 4 Hz, and Link-16 (JTIDS) at 16 Hz.

SAMPLE PDU SIZING

	A	H	E	B	T	ESPDU	FPDU	DPDU	EPDU
TANK	5	1	1	1	1	2220	1132	1356	1180
AIRCRAFT	20	2	3	1	2	4140	1132	1484	2588
SURFACE SHIP	50	5	10	1	5	7552	1132	1868	10060
						OVERHEAD BITS/PDU=			428

SAMPLE RATES PER ENTITY TYPE PER PDU TYPE

	LOW RATE				HIGH RATE			
	ESPDU	FPDU	DPDU	EPDU	ESPDU	FPDU	DPDU	EPDU
TANK	0.2	0	0	0.2	2	0.1	0.1	1
AIRCRAFT	0.2	0	0	0.2	8	0.1	0.1	4
SURFACE SHIP	0.2	0	0	0.2	1	0.1	0.1	2

SAMPLE EXERCISE TRAFFIC ESTIMATES

% ENTITIES AT HIGH RATE	0%	20%	40%	60%	80%	100%
% ENTITIES AT LOW RATE	100%	80%	60%	40%	20%	0%
SAMPLE EXERCISE #1						
600 TANKS	408,000	1,030,656	1,653,312	2,275,968	2,898,624	3,521,280
100 AIRCRAFT	134,560	982,320	1,830,080	2,677,840	3,525,600	4,373,360
10 TACTICAL VOICE LINKS	650,000	650,000	650,000	650,000	650,000	650,000
TOTAL TRAFFIC	*1,192,560*	*2,662,976*	*4,133,392*	*5,603,808*	*7,074,224*	*8,544,640*
PDUs/SEC	*600*	*1172*	*1744*	*2316*	*2888*	*3460*
PACKETS/SEC	*81*	*187*	*293*	*399*	*505*	*610*
SAMPLE EXERCISE #2						
24 SHIPS	84,538	201,896	319,254	436,612	553,970	671,328
50 AIRCRAFT	67,280	491,160	915,040	1,338,920	1,762,800	2,186,680
3 TACTICAL DATA LINKS	237,500	237,500	237,500	237,500	237,500	237,500
TOTAL TRAFFIC	*389,318*	*930,556*	*1,471,794*	*2,013,032*	*2,554,270*	*3,095,508*
PDUs/SEC	*54*	*185*	*316*	*448*	*579*	*711*
PACKETS/SEC	*32*	*74*	*115*	*157*	*199*	*241*
SAMPLE EXERCISE #3						
200 TANKS	136,000	343,552	551,104	758,656	966,208	1,173,760
100 AIRCRAFT	134,560	982,320	1,830,080	2,677,840	3,525,600	4,373,360
5 TACTICAL VOICE LINKS	325,000	325,000	325,000	325,000	325,000	325,000
3 TACTICAL DATA LINKS	237,500	237,500	237,500	237,500	237,500	237,500
TOTAL TRAFFIC	*833,060*	*1,888,372*	*2,943,684*	*3,998,996*	*5,054,308*	*6,109,620*
PDUs/SEC	*304*	*652*	*1000*	*1348*	*1696*	*2044*
PACKETS/SEC	*61*	*139*	*217*	*296*	*374*	*452*

Figure 12. Analysis with Bits, PDUs, and Packets/Sec

APPENDIX C

NETWORK SECURITY FOR DIS

10. Introduction. The goal of this section is to identify a number of security requirements made evident by the broad outlines of DIS, and by common understanding regarding the environment in which DIS will perform. The section will also give a thumbnail sketch of the world of information security, so that our problems can be seen in a larger perspective. This section is not intended to be a comprehensive analysis of DIS security requirements. Such requirements will be a complex function of the system itself as it evolves, and of the needs of its primary intended users.

20. A Security Vocabulary. Arguably the most important task in defining a security specification for DIS is the acceptance of a common vocabulary for discussing security issues. The most widely accepted system so far developed is the DoD Trusted Computer Systems Evaluation Criteria (TCSEC), and its follow on "interpretations" for Networking, Secure Database Standards, and Integrity Criteria. TCSEC is also known as the Orange Book, while the Trusted Network Criteria (TNI) is known as the Red Book. These works have both popularized and made explicit such terms as "Security Policy", "Multilevel Security", "Discretionary Access Control", "Trusted Path", etc., as well as the familiar rating categories C1, C2, B1, B2, B3, A1.

The National Computer Security Center (NCSC) actually undertakes to evaluate and certify production systems according to these criteria. The certification process is long, expensive, and uncertain, with relatively few corporations submitting their products (much less completing the process).

For many environments, including military ones, uncertified systems can provide usable and reliable security features, and frequently comprise the only option. Certainly it is true that a systems designer's choices are extremely limited if he chooses to utilize only NCSC-certified products. The bright side of the picture is that choices are becoming far more numerous.

There are many good reasons for using the DoD security vocabulary. For one thing, it is fairly explicit, and addresses, in one form or another, virtually every conceivable aspect of computer security. It is not necessary to commit to the evaluation categories, or to specific formulas of risk assessment, to benefit from the vocabulary, concepts, and methodologies which have been developed. In addition, the primary clients for DIS will, at least initially, come from the DoD, and the classified nature of information exchanged on distributed simulation nets makes the DoD approach appropriate.

An interesting method for specifying security criteria can be found in the TNI Red Book (Part II), which defines Function to be the security problem being addressed, Mechanism as the means to address the problem, and Assurance as the anticipated

effectiveness of the security mechanism. A method such as this can be useful (irrespective of DoD, TCSEC, et. al) in resolving the frequently ambiguous concepts of computer and networking security. For example, a password mechanism might be used to implement the function of user identification and authentication; the assurance provided by the mechanism would result from specifics of the implementation.

30. DIS Security Requirements. A comprehensive list of DIS security requirements is not available, nor is there one in preparation. Yet certain specific security needs are already discernible.

30.1. Encryption

30.1.1. Confidentiality Requirement. It is known that messages exchanged during a military simulation will contain sensitive data regarding weapons systems characteristics and warfare tactics. A DIS exercise may also be the rehearsal of an operational mission, and as such the data exchanged will be extremely sensitive. Clearly such information must be protected from eavesdropping by simulation non-participants, much in the same manner as telemetry data is protected. Eavesdropping can occur via wiretapping, which is monitoring by entities not legitimately connected to the net, or by users who are legitimately connected but are accessing message data not intended or them.

The primary mechanism for thwarting eavesdroppers is encryption of messages on the network. The architectural level at which encryption/decryption occurs is significant: encryption at the link level (L2) is more efficient, while encryption at the session layer or higher (L5+) allows users to be differentiated by different encryption keys, and protects messages for a greater part of their passage through the operating system of the host. Encryption is used for other tasks as well, in particular the authentication of user identities (see below).

Strength and throughput are two main parameters for evaluating an encryption algorithm. Very secure algorithms exist, but generally compromise the transmission rate on a network, i.e., there is a trade-off between assurance and performance. DIS requirements in this area appear to be of the high-assurance/ high-performance variety.

30.1.2. Evaluating Encryption Mechanisms. Assigning and distributing keys on a dynamic or per-session basis can be a major difficulty. Fortunately, the area of Key Management Systems (KMS) is one in which many standards and commercial products exist. There are also a number of very strong encryption algorithms for which key distribution is not required. These are called public key (PK) encryption algorithms; each user has an encryption key which can be made public. Unfortunately,

PK algorithms are amongst the slowest in terms of throughput. For this reason, they are used to distribute small but critical amounts of data, for example, the encryption keys to a faster algorithm, or authentication data known as digital signatures.

One of the strongest high-throughput algorithms is known as DES, for Data Encryption Standard. It was developed at IBM, and has been widely implemented in commercially available security products, including the Verdix secure Ethernet interface card. It is usually implemented as an LSI chip, with throughput rates approaching 10Mb. In spite of its strength, the NSA has not approved it. Nonetheless, a number of products utilizing DES have been certified by the NCSC.

A public-key algorithm, known as "El Gamal", has been approved by the NSA, and shows strong promise of being incorporated into emerging OSI security standards. It is a close cousin of the earlier Rivas-Shamir-Edelman (RSA) algorithm which has been commercially integrated into various products for some time. Both algorithms rely on the difficulty of factoring large primes, and are of approximately equal strength.

30.1.3. DIS Encryption. Nodes on a DIS network will not transmit a great deal of data, but will receive data from all the other nodes in the simulation. Thus a fast algorithm is required, if only on the decoding end. Algorithms such as DES are suitable for Ethernet-range speeds; a notable example of this being the Verdix VSLAN-100, which provides MLS-type security features (see below). For FDDI, encryption could nullify the speed advantages of the 100Mbs technology.

FDDI fiber optics are relatively safe from wiretapping, so conceivably the need for encryption on a FDDI ring is reduced. It is clear, however, that the encryption problem for FDDI networks will eventually have to be solved. For one thing, eavesdropping by legitimately connected FDDI hosts must be prevented; likewise there is an eventual need for session level isolation and access control in multi-user application gateways. In the short term, only single simulations will run on the DIS net, and nonparticipants can be physically excluded; thus link level encryption for FDDI can await the emergence of a suitably fast technology.

30.2. Multilevel Security

30.2.1. Access Control Requirement. We know that DIS will eventually allow multiple simulations to run simultaneously on a single network. Enforcing the separation of simulations becomes a security issue when differing classification levels coexist, as, for example, when a highly classified weapons development simulation is run together with a simulated battle scenario, presumably at a lower classification level. Participants in one simulation are likely not cleared to view the data being

exchanged in the other. It is also conceivable, that multiple levels of data classification could be required within a single scenario, as when, for example, a highly classified test weapon is introduced into a "normal" battle scene.

Moreover, DIS security issues go far beyond the protection of run-time simulation messages. Computers that participate in more than one level of exercise will be required to store and to internally manipulate data of varying classifications, and to insure that only users with proper clearance can access classified data. This raises issues of Multilevel Security at both the Operating System and Database levels. The following sections discuss these issues in more detail.

30.2.2. MLS Mechanisms. Multilevel (or label-based) security is the primary requirement of the "B" security levels defined in the DoD TCSEC. These levels are B1, B2, B3, in ascending order of strength, based upon the degree of assurance provided by the enforcement mechanism. The mechanisms are called Mandatory Access Controls (MAC) because data transfer is based solely upon the classification and clearance labels involved, and not upon identity-based privileges.

In the first example above, where two simulations coexisted with different levels of classification, access control could be enforced using an encryption mechanism, provided it differentiated individual simulation users. In general, though, the problems discussed above require maintenance and isolation of security levels within the individual systems, and a compatible network transport service to ensure consistency of classification between communicating nodes.

Multilevel Security is implemented by defining a class of protected data objects, and attaching security labels or "classifications" to them. Autonomous entities (users and processes) are known as subjects; these also receive labels, or "clearances", by which their access to the protected objects is regulated The set of subjects and objects, together with the rules for access, is formally known as the Security Policy. The enforcement mechanism for the policy is frequently referred to as the Reference Monitor.

The Bell-Lapadula Model is a Security Policy associated with the B2 security level. It contains the familiar "read down, write up" provisions which guarantee that classified information cannot flow to a lower-classified entity within a system. The Bell-LaPadula model is the most widely accepted and implemented access control model currently in existence.

The Red Book interprets these concepts in the context of interconnected systems, or networks. Especially important are the issues of compatibility between security labels when connecting MLS systems, and of authenticating user/process

identities and privileges across the network. For these and other reasons, network security products are considered much more difficult to evaluate and certify than stand-alone systems, and there are significantly fewer certified security products in this area.

30.2.3. DIS Multilevel Requirements. As we have seen, the need for Multilevel security in DIS is closely tied to existence of multiple concurrent simulations, or multiple levels of simulation hosted on single machine. If each simulation node is assigned a single security level, and is only allowed to participate in exercises at that level, a network level MLS system, such as that provided by the Verdix VSLAN Ethernet interface, should suffice (or encryption can be used to separate security levels). If a system must handle multiple levels of classified data concurrently, an MLS Operating System will be required. Storage of multiple levels of classified DIS information within the system will require an MLS database facility. Product availability in these and other areas is discussed in the next section.

Prototype DIS implementations will most likely not support multiple simulations, pending widespread availability of a standardized multicast mechanism. Hence the need for MLS operating systems and databases for the simulation hosts is not immediate (which is fortunate, as these products are currently slow and difficult to integrate).

30.3. Identity and Authentication. In a distributed interactive simulation, it is important to guarantee that participants are, in fact, who they say they are; this is known as the Authentication problem. Identification of entities can occur at varying levels of granularity: the level of host on a network, the level of human users on the network, or the identifications of individual processes. In the initial DIS environment, simulation hosts will participate in only one simulation at a time; it seems reasonable, therefore, to initially propose a per-node granularity of authentication.

We have observed that encryption is frequently used to provide authentication. An example of this is the Digital Signature mechanism mentioned above. Digital Signature is a mechanism that can be provided by a reversible Public Key encryption algorithm, that is, one for which each user's public key and private key exactly reverse the operation of each other. In such a case, a user X can send a message encrypted with his private key; anyone can decrypt it with X's public key to confirm that only X could have sent it.

Authentication is used in secure networks as a basis for automatically assigning clearances on each node of the network, for example when a user executes a remote login. Frequently a more elaborate form of digital signature is used, called a

certificate, which in addition to authenticating identity, contains additional network-wide security information as well.

Generally speaking, a network authentication mechanism of this type will interface with a Key Management unit that generates and maintains cryptographic keys at the necessary degree of granularity. The specification and distributed operation of such a mechanism is quite difficult, for example, the interaction for crypto-based authentication mechanisms of SP3 (the OSI/SDNS Layer 3 security protocol) and the distributed Key Management Base have yet to be completely specified.

30.4. Integrity. In DIS, as in most environments, there is the need to insure that data is not corrupted, either deliberately or by accident. This issue of Integrity is an important security problem, and applies to message data, stored information, and dynamically manipulated information within an operating system. Again, cryptography plays an important role in data integrity verification (for example by checksums), and many network authentication services also support point-to-point integrity policies. More involved approaches formulate integrity policies for system data, with models and enforcement mechanisms similar to the Bell-LaPadula policy used for B2-level access control. In some respects an integrity policy is the reverse of Bell-Lapadula, as the "write-up" operation permitted by the latter actually allows one to corrupt protected data at a classification level higher than oneself.

30.5. Audit. A critical facility for all secure systems, including networks, is the audit facility. The audit facility maintains logs of security-relevant events in tamperproof, restricted access locations; typical examples of logged events include attempted logins and access to critical data. A fair number of commercial audit products exist.

Audit trails can be maintained on individual systems, but a network audit facility is also desirable in DIS. Coordinating a distributed audit facility can be a problem, and might require utilities like NFS and yellow pages (secure versions of which are currently under development). The main problem with audit is storing and analyzing the enormous amount of data that can be generated. The primary approach to this problem is to specify a limited set of audit events; this greatly reduces the data volume. Many audit systems will have built in "triggers", or thresholds, that expand the level of audit detail in areas where certain conditions have been exceeded. Likewise, there are processing tools to make the analysis of audit data easier, should that prove necessary.

30.6. Physical Security. Physical security consists of functions that can be performed by "physical" mechanisms, i.e., those that are not part of the computer operating system. A list of examples might include:

1) Protected cables, locked rooms, security guards, removable media

2) Computer locks, disk drive locks, hardening against radiation leakage & radiation damage

Physical security can come from unexpected sources, for example fiber-optic networks, which almost as a bonus provide several measures of physical security. Fiber optic networks 1) are difficult or impossible to tap undetected, 2) immune to EMR damage, and 3) do not leak EMR that can be monitored.

In general, however, the methods of physical security lie outside our scope of interest.

40. Security Products. A list of certified network products can be found in the Information Systems Security Products and Services Catalogue, published by the NSA.

50. Conclusions. The security situation for DIS is complicated by the desire for standards and interoperability, as well as by a real dearth of available products. Implementors of critical features, such as networking, operating systems, and database security will have to confront major systems integration and standards-conformance problems. At the same time, the classified environments in which DIS must operate will make adherence to formal standards of evaluation and certification more critical than in commercial environments.

DIS is not the only standards effort hampered by the complexities of computer and network security. A long term DIS strategy for implementing security features must be developed. An evolutionary approach which utilizes available technologies and anticipates emerging ones (just as in the case of network architecture) should serve this effort well. The situation is one in which standards are evolving from implementations, not vice-versa.