


DOD PROTOCOL REFERENCE MODEL(U) SYSTEM DEVELOPMENT CORP
SANTA MONICA CA 30 SEP 82 SDC-TM-7172/201/01
DCA100-82-C-0036

1 / 1

F/G 17/2

NL

END
DATE
FILMED
4 83
DTIC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

SDC

System Development Corporation

2500 Colorado Avenue, Santa Monica, CA 90406 Telephone (213) 820-4111

TM

a working paper

This document was produced by
System Development Corporation in performance of Contract DCA100-
82-C-0036

series base no. vol. issue
7172/291/01

author
Sytek Staff

technical *R/LM*
Richard L. Mandell

release *GACole*
Gerald A. Cole

for
Charles A. Savant

date
9/30/82

DTIC
SELECTED
APR 8 1983
H

DCEC PROTOCOL STANDARDIZATION PROGRAM

DoD PROTOCOL REFERENCE MODEL

SEPTEMBER 1982

This document was produced under subcontract to SDC by
Gregory B. Ennis, David J. Kaufman, and Kenneth J. Biba
of Sytek, Inc.

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

83 04 07 038

DA 126556

DTIC FILE COPY

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 7172/201/01	2. GOVT ACCESSION NO. A126 550	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) DoD Protocol Reference Model		5. TYPE OF REPORT & PERIOD COVERED interim technical report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Sytek Staff		8. CONTRACT OR GRANT NUMBER(s) DCA100-82-C-0036
9. PERFORMING ORGANIZATION NAME AND ADDRESS System Development Corporation 2500 Colorado Ave. Santa Monica, CA 90406		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS P.E. 33126K Task 1053.558
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Communications Engineering Center Switched Networks Engineering Directorate 1860 Wiehle Ave., Reston, VA 22090		12. REPORT DATE 30 Sep 82
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) N/A		13. NUMBER OF PAGES 68
		15. SECURITY CLASS. (of this report) Unclassified
16. DISTRIBUTION STATEMENT (of this Report) Approved for Public release; distribution unlimited		15a. DECLASSIFICATION DOWNGRADING SCHEDULE N/A
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) N/A		
18. SUPPLEMENTARY NOTES This document represents results of interim studies which are continuing at the DCEC of DCA.		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Protocols, Data Communications, Data Networks, Protocol Standardization, Protocol Reference Model		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This document is proposed as a baseline Reference Model serving the development of standard protocols for the Department of Defense. As such, it attempts to describe the design principles which are implicit in the protocols developed under the ARPANET and internet programs; it also attempts to prescribe principles for the development of future protocols under the ongoing DoD Protocol Standardization Program managed by the Defense Communications Agency.		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

CONTENTS

1. INTRODUCTION.....	1
2. DESCRIPTION OF APPROACH.....	4
2.1 DoD Networking Requirements.....	4
2.2 The DoD Program of Protocol Standardization.....	6
2.3 Relationship with the ISO Model.....	7
3. BASIC CONCEPTS OF THE MODEL.....	12
3.1 Networks, Hosts, and Processes.....	12
3.2 The Internet.....	13
3.3 Entities and Protocols.....	14
3.4 Services and Service-Access-Points.....	15
3.5 Connections and Connectionless Service.....	16
3.6 Layers.....	17
4. CURRENT PROTOCOLS AND THE BASIC MODEL CONCEPTS.....	19
4.1 Layering.....	19
4.2 Service-Access-Points.....	20
4.3 Connections and Connection-Endpoint-Identifiers.....	20
5. THE PRESENTATION LAYER.....	22
5.1 Differences between DoD and ISO Presentation Layers.....	22
5.2 THE DoD PRESENTATION LAYER.....	25
5.3 Current DoD Presentation Protocols.....	27
6. THE SESSION LAYER.....	28
6.1 Differences between DoD and ISO Session Layer.....	28
6.2 THE DoD SESSION LAYER.....	31
6.3 Current DoD Session Protocols.....	36
7. THE TRANSPORT LAYER.....	37
7.1 Differences between DoD and ISO Transport Layer.....	37
7.2 THE DoD TRANSPORT LAYER.....	38
7.3 Current DoD Transport Protocols.....	45
8. THE INTERNET LAYER.....	46
8.1 Differences between DoD and ISO Internet Layer.....	46
8.2 THE DoD INTERNET LAYER.....	47
8.3 Current DoD Internet Protocols.....	54
9. THE NETWORK LAYER.....	55
9.1 Differences between DoD and ISO Network Layer.....	55
9.2 THE DoD NETWORK LAYER.....	56
9.3 Current DoD Network Layer Protocols.....	62
10. THE DATA LINK LAYER.....	63
10.1 Differences between the DoD and ISO Data Link Layers.....	63
10.2 THE DoD DATA LINK LAYER.....	64
10.3 Current DoD Data Link Protocols.....	67

11. REFERENCES.....	68
---------------------	----

1. INTRODUCTION

As in any engineering discipline, the development of computer communications systems is facilitated by the use of design principles. These principles give guidance on "good engineering practice" within the discipline of computer network design.

Fundamental to the design of any computer communication system are the rules by which information is exchanged. Such rules are embodied in the protocols which must be followed if communication across a given network is to succeed. The discipline of protocol design has its own set of guiding principles for "good engineering practice"; these principles have been found by experience to yield better network designs if followed. For example, experience has taught that if the set of protocols used within a network have a certain hierarchical structure, then the overall network will be easier to design, easier to modify, and easier to understand. Such principles have often gone unstated, but their existence and their utilization have allowed network design to move beyond the black-art stage.

A Protocol Reference Model is a document which presents principles of protocol design. Such a document may be "descriptive" in that it describes the principles which were indeed followed within a particular design effort; it may be "prescriptive" in that it prescribes rules-of-thumb for future protocol designers.

A Reference Model document serving a particular network design effort will in fact be both descriptive and prescriptive. The development of a set of protocols is an evolutionary process, and successful protocols have a long life-cycle. Consequently, the design of new protocols must always take into consideration the existing environment of older protocols, and should be cognizant of the principles that went into their design. This is the benefit of a descriptive Reference Model as new protocols are developed. The prescriptive aspects of a Reference Model can then evolve, drawing upon the experience and insights gained as new protocols are introduced into the system.

This document is proposed as a baseline Reference Model serving the development of standard protocols for the Department of Defense. As such, it attempts to describe the design principles which are implicit in the protocols developed under the ARPANET and Internet programs; it also attempts to prescribe principles for the development of future protocols under the ongoing DoD Protocol Standardization Program managed by the Defense Communications Agency.

The most well-known protocol Reference Model is the Reference Model for Open Systems Interconnection, developed by ISO [1]. This ISO document is the Reference Model being used within the international effort which is heading towards a set of standard protocols for commercial systems. As such, it presents the design principles which are being applied as these new protocols are defined. Since the ISO committees are still in the early stages of this protocol design effort, the ISO Model is at present mostly prescriptive; we can anticipate changes in the ISO document as the new protocols are implemented and tested.

The ISO Reference Model is commonly referred to as "the" Reference Model. However, it can no longer be viewed as the embodiment of universal principles appropriate for all network design efforts. The ISO Model originates from a particular protocol development effort - admittedly one with wide scope. The often-misunderstood goal of the developers is a set of protocols which can be implemented by computer system vendors to allow communication across commercially-provided networks. As a potential user of commercial systems and commercial networks, DoD is clearly interested in the ISO effort. However, DoD cannot use the ISO Model as a description of design principles appropriate for the development of standard protocols within the military environment.

The reasons for the divergence of the DoD and ISO Reference Models fall into two categories:

1. DoD-specific communications requirements (such as security and survivability requirements) have a major impact on the shape of the most appropriate protocol architecture. These concerns have not been uppermost in the minds of the ISO developers, and predictably are not reflected in the ISO Model.
2. The fundamental principles embodied in the ISO Model are felt to place far too many restrictions on the designer of DoD-standard protocols.

These differences will be explored in detail throughout this document.

One of the primary roles of a Reference Model document is as an aid to the understanding of an ongoing network development effort. The present document should communicate the basic principles of the "DoD approach to protocol design" not just to those actively involved in the DoD Standardization Program, but also to those on the outside of this effort. The near-ubiquity of the ISO document and the language it has popularized means that this goal can best be achieved by using elements of the ISO language when possible. It must, however, be recognized that the implications of some important concepts (such as the concept of a protocol residing in a given layer) differ greatly between the two documents. Of course, both documents to some extent share common concepts; much of the original ARPANET work has found its way into the ISO effort. Thus there is a certain commonality of language which should help those familiar with the ISO model in their reading of the present document.

The document presented here is proposed as a guide to the basic direction of DoD protocol specification efforts. Section 2 describes the basic approach taken in the development of the model. The role of the model in the overall DoD protocol standardization effort is described, and the relationship between the DoD and ISO models is discussed. Section 3 presents the basic concepts used in the model. A discussion of how mechanisms within current DoD protocols relate to these basic concepts is presented in Section 4.

Beginning with Section 5, the layers of the DoD Reference Model are presented. Each layer is first discussed informally, outlining the deficiencies of the corresponding ISO layer from the DoD perspective. A "formal" description of the layer is then given, in the style of the ISO document. This description should make explicit the differences between the DoD model and the ISO

30 September 1982

-3-

System Development Corporation
TM-7172/201/01

Reference Model. Finally, various existing and planned DoD protocols are discussed as they relate to the given layer.

2. DESCRIPTION OF APPROACH

The next sections discuss three sets of issues which determine the shape of the DoD Protocol Reference Model. First we identify certain DoD-specific networking requirements, paying particular attention to how such requirements may affect a Reference Model. Secondly, the role of such a model within the overall DoD program of protocol standardization is discussed - the view of the model's developers on this subject obviously affects the proper interpretation of the model's contents. Finally, the similarities and differences between the DoD and ISO Models are briefly described.

2.1 DoD Networking Requirements

To a large extent, DoD's needs are similar to those of many other users of computer networks. However, DoD's requirements in some specific areas demand the development of a DoD-specific protocol architecture. The development of protocols for DoD networks must take into consideration the following DoD concerns:

- anticipated DoD network applications
- internetworking with present and planned DoD (and non-DoD) systems
- security requirements
- robustness and other DoD quality-of-service issues
- support of realtime and tactical communications
- phased evolution from existing DoD systems and protocols

The following paragraphs briefly describe how each of these concerns has affected the proposed DoD Reference Model.

2.1.1 Anticipated Applications

The DoD protocol architecture must support a wide variety of anticipated applications, including advanced services such as computer-based messaging, multi-media teleconferencing, and distributed database access. Voice applications (real-time and store/forward messaging) are also anticipated to play a major role in future DoD networking.

The development of other network architectures often proceeds from an assumption that such high-level services are beyond the purview of the basic set of standard protocols. Within DoD, however, such applications can be explicitly incorporated into the architecture, and indeed must be if true DoD-wide communication is to be possible. Consequently, the DoD Reference Model should point towards the future development of standard protocols for such sophisticated applications.

2.1.2 Internetworking

The DoD architecture model must facilitate the interconnection of networks which use vastly different internal routing schemes, including broadcast-based local area networks, long-haul DoD networks such as DDN, demand-access satellite networks, circuit-switched tactical networks, and X.25 public networks. It must be pointed out that the administrative differences among such networks are often a greater impediment to intercommunication than the technological differences.

Experience has shown that the datagram approach typified by the DARPA Internet Protocol (IP [2]) is capable of providing the functionality necessary for such internetworking. Such an approach forms the basis of the DoD architecture's internetting scheme.

However, some applications require a level of performance which datagram internetting may not be capable of providing. Consequently, the DoD architecture must allow for alternative internet routing mechanisms which can support, for example, delay-sensitive applications.

2.1.3 Security

DoD's security requirements must be fully integrated into the DoD Reference Model. The architecture includes concepts which support DoD's security needs, particularly in the areas of data protection, access control, authentication, and accountability. The present Reference Model document includes explicit architectural mechanisms to support the necessary exchanges between "users" and appropriate access controllers. However, a full description of the security architecture is not included here.

2.1.4 Robustness and Quality-of-Service Requirements

DoD applications will have a wide variety of delay and reliability requirements. The model's approach to such quality-of-service requirements must be integrated coherently throughout the entire architecture. Such issues are present in all network architectures, but are particularly important to DoD for crisis management as well as for peacetime support of special DoD applications. Flexibility in the offered quality-of-service is important if voice and other "non-data" applications are to be supported.

Reliable data transmission is important for many applications. But for DoD, reliable management of distributed applications is equally important. Such requirements are generally labeled as "survivability" concerns. The architecture can support DoD's survivability requirements by providing facilities for maintaining control of a distributed application in the presence of node and link failure.

2.1.5 Support of Realtime and Tactical Communications

In addition to their support of standard file transfer, messaging, and terminal access functions, DoD networks must often support realtime traffic. This includes such applications as radar tracking and device control. Although

these applications typically require specialized transmission technologies, it is critical that their required protocols be part of the integrated DoD architecture. This is to allow communication among separately designed tactical systems as well as between tactical and other DoD networks.

A key problem in the tactical communications arena is the impact of very low bandwidth channels on the design of appropriate protocols. The Reference Model must provide guidance to the designers of such systems - allowing use of standard protocols whenever possible, and promoting the coherent inclusion of tactical systems within the overall framework of DoD communications.

2.1.6 Evolution From Existing DoD Protocols

Finally, the development of the DoD architecture must proceed with an awareness that DoD has a substantial investment in current systems and protocols. The Reference Model must describe the design principles behind the existing DoD protocols, particularly IP [2] and TCP [3]. A major purpose of the Reference Model is to guide the evolution of DoD networks from existing designs to a more complete system capable of meeting DoD's future networking needs.

2.2 The DoD Program of Protocol Standardization

A strong program of DoD protocol standardization is required to effectively manage the proliferation of separate DoD communication systems. One important piece of such a program is the development of a DoD Protocol Reference Model.

The Model will serve to illuminate the structure of the set of protocols. The desirability for a new protocol can emerge from new user needs or from new technology. Through a comparison of the new protocol's informal description against the Model document, proper placement of the protocol within the architecture can be determined. Such an analysis can also serve to refine ideas on the anticipated protocol's service and mechanisms.

These informal ideas are made rigorous through the development of a set of specifications which together define the protocol. A service specification states precisely what services the protocol will provide to its users. A mechanism specification describes precisely how the protocol is to provide its service (adding functionality to the services offered by the lower level protocol's). An interface specification defines the manner in which the protocol's services may be invoked. Together, these specification documents place the protocol precisely within the Reference Model, describing the interaction between the new protocol and other existing or planned protocols.

During the detailed development of the specifications, the Model document guides the protocol designer by ensuring that the overall network architecture is considered. The designer can ensure that the new protocol will provide useful services without redundancy. Implementation, validation, and installation of protocol modules can then proceed without explicit reference to the Reference Model.

2.3 Relationship with the ISO Model

The DoD Reference Model is in some respects similar to the ISO document; in other respects there are significant differences.

2.3.1 Reasons for Similarities

The DoD Reference Model must describe the basic concepts of network communications as dictated by DoD requirements. The primary "audience" of this description consists of the designers and users of networks, protocols, and systems for DoD applications.

However, another important audience consists of those outside of the DoD community. It is often desirable for DoD to use commercially available networks and networking products if possible, so long as such use satisfies fundamental DoD requirements. Consequently, the communication of the DoD approach to networking to the developers of commercial standards and products is an important DoD objective.

For this reason, the present DoD Reference Model shares much of the language with the ISO Reference Model document. The use of some common language facilitates comparisons between the ISO and DoD models, and it is hoped that it will allow continued interaction between the commercial and DoD protocol standardization efforts.

DoD use of commercial network systems will also be facilitated if there is a high degree of correlation between the services DoD requires of a set of protocols, and the services provided by commercial systems. This is commonly stated as a desire for "functional equivalence" between different protocols, and one of the goals of a Reference Model is to identify the most appropriate classification of protocols according to the type of services they provide. Although DoD requirements for specific protocol services are often different from those within the commercial world, the basic method of classification (the layers) may share a common framework. If the DoD and ISO models have such a common framework, the ideal of functional equivalence may be less unattainable as protocol standards based on the two models are developed.

Thus there are three basic reasons why the DoD Reference Model document is similar to that of ISO:

1. The development of the ISO model drew upon many of the concepts which originated in ARPA-sponsored research, and consequently the two models share some common history.
2. Through the use of some common language and concepts it is hoped that the communication of DoD networking approaches to those familiar with the ISO effort is facilitated.
3. By agreeing in many instances on the appropriate division of functionality among layers, the ease with which DoD may use commercial networking systems may be enhanced.

The development of the DoD Reference Model document has been guided by a number of considerations. For the above reasons, it is in many respects similar to the ISO document. However, other important considerations have precluded the wholesale adoption of the ISO document for the DoD.

2.3.2 Reasons for Differences


The omission of some DoD-specific requirements from the ISO Model has already been briefly discussed, and will be touched upon at various points further in the document. Here we wish to make explicit the fundamental architectural differences between the two models.

There are five fundamental differences:

1. the meaning of "layer",
2. the manner in which one protocol may use the services of another,
3. the importance of internetworking,
4. the utility of connectionless services, and
5. the approach to management functions.

The concept of "layer" is fundamental to the ISO Model. The concept of "protocol hierarchy" is fundamental to the DoD Model. The distinction between these two concepts has a major impact on the interpretation of the documents and on the actual design of corresponding protocols and implementations.

The ISO "layer" concept seems to have originated in the following way: it has been found that well-designed computer networks have their protocols organized hierarchically. By this we mean that one protocol interpreter provides a communications service to its users by adding value to the services provided by one or more other protocols; and "good engineering practice" dictates that at no point should one protocol be implicitly using its own services, i.e. the protocols form a hierarchy. Thus we find protocol architectures often depicted via graphs of the following form:



The next step in the development of the "layer" concept is the recognition that in many respects, protocols in the same row of such a picture seem to have certain features in common. This yields the naming of the rows, and the attempt to describe in an abstract fashion what features are held in common by all protocols within a given row. Thus we get "Layers".

There is no argument with the basic thrust of this exercise. Indeed, protocols within a row do have some features in common, the most important of which is the fact that they all use the services of protocols further down. Unfortunately, the fact that they use similar services does not imply that they provide similar services. Consequently, any attempt to rigidly specify the services which protocols within a given layer provide will likely exclude certain protocols which in fact use the same lower-level services.

This is the origin of a major disagreement between the ISO and DoD Reference Models. Within the ISO Model, the concept of layer has become of paramount importance, overshadowing the more fundamental notion of protocol hierarchy. This has yielded several principles implicit in the ISO Model, such as the requirement that the users of an (N)-protocol need be (N+1)-entities. This precludes allowing Protocol "C" from using the services of Protocol "H" - a restriction which does not arise from the "arrange protocols hierarchically" maxim, but only from the "place protocols in layers" maxim.

Many of the arguments within the ISO effort surrounding "minimal layers", "sublayers", and "which layer does this belong to" can be traced to the fundamental overemphasis on layers within the model. The concept of layers does in fact have utility as an explanatory aid - it is indeed true that protocols at a given vertical position in a hierarchy do seem to have some features in common. However, the layers should not have a "life of their own", dictating protocol designs.

Good protocol engineering does not require that protocols fit into layers; it is only required that protocols be arranged hierarchically. This latter principle is the basis of the DoD Reference Model. Nevertheless, the DoD Reference Model is described as consisting of layers; this is an explanatory technique rather than a basic principle of protocol design.

The second fundamental difference is related to the first. Although it is possible to interpret the ISO Model differently, the following is a common understanding of the document:

- (N)-entities must exchange data using services provided by (N-1)-entities.
- (N-1) entities provide their service by exchanging data-units which contain (N-1) control information and data from the (N)-entities.
- (N-1) entities must be involved in every data transfer between the (N)-entities.
- (N) control information is passed to the remote side as (N-1) data.

We want to make it clear that within the DoD architecture, communication need not be so restricted. In particular, the DoD Reference Model in no way precludes the provision of services to higher level users via the following techniques:

- "Escape characters" which allow the placement of control information within a data stream (e.g. Telnet [4]). It is unnatural to describe such protocols using the concept of "data units" containing both control and higher-level data.
- "Control connection and data connection", in which higher-level data and control information may never share a "data unit". (Similar to the File Transfer Protocol [5]).
- "Using (N-1) control to signal (N) control" - example: closing a TCP connection implicitly terminates a Telnet connection, without requiring that Telnet control information be passed as TCP data to indicate that the Telnet connection is to be closed.
- "Getting out of the way" - a protocol interpreter may be involved at the start of a sequence of data transfers (e.g. a name service protocol), but need not "touch" the data during the remaining transfers.

In fairness it should be pointed out that the ISO Model may not in fact preclude any of the above ways in which a protocol provides services to users. However, the language of the ISO Model seems to be patterned on a paradigm of "encapsulation" as the primary means by which higher level control information is handled by the lower level protocol. The more fundamental notion is this: a protocol provides a set of services to its users, and the users may make any use they wish of these services.

The third fundamental difference between the ISO and DoD Models is the relative importance given to the problems of internetworking. This is most easily seen by considering that the DoD Reference Model has explicitly identified an "Internet Layer".

It is clear that the ISO Model did not originally place the proper emphasis on the problem of multiple network routing. Although many efforts based on the ISO work are now taking the Network Layer to include an Internet Sublayer, this seems to confuse rather than clarify the issue (since it raises questions as to the meaning of sublayers). Since in practice the implementation of the internet protocols will likely be provided by a different vendor/developer from the implementation of the subnetwork protocols, it seems that this concept deserves its own layer.

It is also difficult to correctly present the DoD approach to internetworking within the confines of the ISO Model due to the ISO Model's lack of emphasis on "connectionless" services. This is the fourth of the fundamental differences between the two models.

The primary use of connectionless service within the DoD architecture is for internetworking - IP being the paradigm. However, connectionless services are of

great utility in other contexts as well. For example, interactions with a Name Server may involve the use of a protocol similar to the "User Datagram Protocol" (UDP [6]), which allows exchanges of data between processes without the establishment of a connection. Similarly, many local networks operate on a connectionless basis. The DoD Reference Model explicitly identifies connectionless services as equal in importance to connection-oriented service throughout the architecture.

The final fundamental difference between the ISO and the DoD Models is the manner in which various management-related functions are treated. Here we refer to functions such as the naming of resources, the control of access to resources, and the accounting for resource and network usage. The ISO Model has difficulty treating these, due in part to its architectural restrictions. For example, many of these functions are most naturally handled via connectionless services; they typically involve entities which are only involved at the initiation or termination of a sequence of data transfers; and many communicants may be involved besides the "users". It is clear that protocols must be defined: how are such protocols to be placed within the rigid layers of the ISO Model?

Given the DoD Model's less restrictive approach to "layers", it is possible to describe a uniform approach to many of these system-management protocols. The DoD Model identifies such protocols as typical of "session layer" services, which involve the coordinated action of multiple entities (e.g. access controllers or name servers) on behalf of users and administrators, helping to establish and manage user communications. Such protocols are placed within the session layer by virtue of their use of "transport layer" services (e.g. UDP) and their common features.

It is important to recognize that this does not imply that all data transfers between application programs require handling by some session layer protocol; such an implication would exist with the ISO session layer. There is in fact little in common between the ISO and DoD session layer descriptions.

Of course, there will be many protocols defined which sit at other locations within the protocol architecture which may play a role in the management of the network and of network resources. It is not implied that all such protocols are most appropriately defined as "session" protocols. However, it does appear that many similar protocols will exist which can be described as performing session-layer functions within the DoD Model.

Another reason why the ISO Model cannot deal properly with such services is the difficulty of developing true international standards for e.g. access control. In contrast, the Department of Defense can develop such standard protocols, and indeed must if communication across access control domains is to be a part of the internetworking services provided by TCP and IP. The placement of these services within the DoD Reference Model thus points towards future areas for protocol standardization.

These are the major differences between the DoD and the ISO Reference Models; additional differences will arise throughout the actual text of the Model.

3. BASIC CONCEPTS OF THE MODEL

The DoD Reference Model is based on a set of fundamental concepts. "Fundamental" concepts in any subject are difficult if not impossible to rigorously define, and the language of computer networking and protocols has its share of hazy terms. Mutual understanding of such a language is typically reached through its usage rather than through mastering of definitions. In this section we describe the manner in which the language of these fundamental concepts is used within the DoD Reference Model document, and why these concepts have shaped the model into its present form.

3.1 Networks, Hosts, and Processes

Although much is made of the "merging" of communications and computing, it is still clear that the distinction exists. Computers play an ever-increasing role in the transmission of information, and computations often involve communications - but nonetheless it is generally possible to identify a system's primary role as one of computation or one of communications. The term used to describe machines whose primary role is one of computation is host, while a set of equipment which acts in a coordinated fashion to allow communications between hosts forms a network.

Hosts can often support multiple simultaneous computational activities, which are often called processes. Each process "resides" in a particular host. It is communication between processes which must in fact be supported by the network.

These three concepts have yielded a fundamental principle in the DoD Reference Model: the transfer of information to a process can be accomplished by first getting it to the host in which the process resides, and then getting it to the process within the host. These two "demultiplexing" operations can be handled independently. Consequently, a network need only be concerned with routing information between hosts - so long as the hosts agree among themselves on how information is to be directed between processes.

As will become clear, this principle is one of the primary influences on the basic shape of the DoD Reference Model. A network conforming to the Reference Model must be able to distinguish among all the hosts which are attached to it, and to route information accordingly. Such host identification is typically implemented as an "addressing scheme" for the network, with a specific format for host addresses and a careful management of address assignments to specific hosts. Using these addresses, the network may be directed to transfer information to a specified host.

It should be pointed out that once a network develops an addressing scheme, one quickly achieves a de facto definition of "host" from that network's point of view: hosts are precisely identified with addresses. It is in this way that one finds a single machine can be two hosts (it has two addresses), or that a machine which is seemingly attached to a network may in fact not be a host from the network's viewpoint (it has no address). Such anomalies often lead to confusing discussions about hosts.

In much the same manner, with the introduction of host addresses one achieves a de facto definition of "network" from a host's point of view: a network consists of the communications system which allows data exchanges between hosts sharing a particular addressing format, with the administration of address assignments under common management. In this way we see that a network is not to be identified with a set of communications equipment drawing from a common transmission technology, but rather is to be identified with a centrally administered address and routing scheme. Thus the public phone system is a single network, even though a wide variety of transmission technologies are used (satellites, microwave links, and digital exchanges), whereas a broadband cable system with each channel separately administered may be considered multiple networks.

3.2 The Internet

At the heart of the DoD communication problem is the realization that no one network technology is adequate for all of the needs of DoD users. Networks can be built from a variety of technologies, including circuit-switched and packet-switched land lines, satellites, coaxial cable local networks, and packet radio. Each technology has certain advantages with certain types of traffic and applications. Consequently the proliferation of different network technologies will continue within the DoD environment.

In addition, DoD consists of a wide variety of separate agencies and services, each with their own mission. This has yielded a multitude of separately administered networks, each based on the technologies which are most appropriate for its particular applications. DoD also wishes to make use of public networks if possible; such a network forms an additional example of a separately administered network within the DoD environment.

Although many DoD applications do not require communication beyond hosts on a single network, it is imperative that the profusion of separate networks not preclude communication between hosts on distinct networks. Thus the DoD Reference Model takes as one of its fundamental requirements the ability to communicate across multiple networks.

The approach taken to internetwork communication is similar to that taken to communication within a single network: hosts must be addressed using a common format, with assignment of addresses to hosts centrally administered. This is accomplished by assigning an identifying number to each network which is to participate, so that hosts may be uniquely identified via a network identifier and a host address on that network. The collection of networks which have been assigned such an identifying network number are collectively referred to as the internet.

Network's do not lose their individual identity by participating in the internet. The network's addressing scheme for intranetwork routing needn't be dictated by the internet's administration. However, if a host on the network is to participate in internet communications, it must be capable of translating between internet and intranet address formats.

Internet routing is the responsibility of internet gateways. A gateway is a host on multiple networks - i.e. it is assigned a network address by the administrations of those networks and is capable of exchanging data with other hosts on those networks. In addition, a gateway must implement the routing procedures defined by the internet administration.

The simplest internet routing mechanism assumes the least about the routing mechanisms of the individual networks. This "least mechanism" internet scheme assumes only that each network in the internet is capable of delivering a data unit from one host to another with some high probability of success. The network service assumed here is often referred to as "datagram service", and the style of internetting it leads to is called "datagram internetting". Gateways in such a scheme are similar to packet switches: they receive a datagram from one network, and determine from the internet address of its intended destination how it is to be forwarded. Each datagram is handled independently, and if a datagram is lost, duplicated, or damaged within a network, it will not be recognized by the internet.

Datagram internetting is the basic approach taken by the DoD Reference Model. It is embodied in the DoD standard "Internet Protocol" (IP), which defines a datagram format and the procedures for handling datagrams within both hosts and gateways. However, other internetting schemes are possible, and are also consistent with the Reference Model. These alternatives may provide enhanced services to the hosts (for example to provide minimization of delay variances between speech packets traversing the internet), and may take advantage of network services beyond the minimal datagram service.

3.3 Entities and Protocols

A protocol is a set of rules describing the manner in which an existing communication service can be used to achieve information transfers. Within computer communication networks, the interpreters of a protocol are typically implemented within software. One such interpreter may communicate with another according to the rules of the protocol, using a "lower" communication service to exchange data and control information.

Often a protocol is defined which takes an existing communication service and "adds value" to it, providing an enhanced communication service to its users. In this way the concept of "protocol hierarchy" arises.

Protocol interpreters - and the users of protocols - are implemented in different fashions within different hardware/software environments. In many instances one finds a protocol interpreter implemented as a "process", in other cases it may be defined as a "procedure"; often protocol interpreters are integrated into an operating system kernel or are made to appear as device drivers. Protocol interpreters may be implemented as a collection of processes, or using a combination of various implementation techniques. In all cases, however, it appears to the remote protocol interpreters across the network as though there were some active being - an "entity" - who responds correctly to the rules of the protocol and is the user of the underlying communication resource.

In many respects, the term "entity" is a synonym for "protocol interpreter". However, "entity" has the following connotations which are not explicit in the term "protocol interpreter":

1. An entity may in fact be capable of interpreting multiple protocols.
2. Although entities must interpret the same protocol in order to communicate, they need not be the same in other respects.

The second point is illustrated in so-called "asymmetric" protocols, in which the two entities are playing different roles yet communicating via a common protocol. Name service protocols are an example; one can also point to the distinction between a "Host IP Entity" and a "Gateway IP Entity". The specification of an entity must not be confused with the specification of a protocol.

3.4 Services and Service-Access-Points

A protocol is often defined to provide an enhancement of an existing communications service for the benefit of some "users". For example, a given communications service may not have an acceptable error rate, and a protocol can be defined to decrease the perceived error rate (e.g. through retransmissions). The entities implementing such a protocol are providing a service to their users, and an important part of the specification of such a protocol is a specification of these services.

Typically, a protocol is defined not just to provide services to a single user, but to provide the same service to multiple users. This requires that the user entities and the protocol entity agree on a method by which user entities can access a complete set of services independent from other users' access to services. This idea is expressed within the Reference Model using the term service-access-point. A service-access-point is the abstraction used throughout the Model for the means by which a user entity may access the services of a given protocol entity.

A protocol typically includes a defined means of addressing a service-access-point. In this way, user entities may be identified within the protocol by their address, i.e. by the address of the service-access-point to which they are attached. It should be pointed out that although a specific user entity may be attached to multiple service-access-points, this need not be known to the protocol entity, who treats all service-access-point independently.

It is definitely not true that every entity provides services to a set of users. For example, a "routing" entity (such as a Gateway IP Entity) need not have a "service specification" describing what services it offers at its service-access-points. The "services" provided by such an entity are on behalf of remote entities, rather than on behalf of a local "higher" user entity.

3.5 Connections and Connectionless Service

Many protocols make use of the concept of a "connection". Connection-oriented communications between entities typically require the establishment of state information within the entities which persists (appropriately updated) throughout the duration of the data transfers. Such state information is present, for example, if reliable data transfers are required, and the paradigm of a connection-oriented service is the "reliably sequenced and flow controlled virtual connection". It should be pointed out, however, that connection-oriented protocols are also of utility for other applications which may have specific delay requirements that require reservation of resources throughout the network.

The initialization of such state information often requires the exchange of control information prior to the transfer of any user data. A similar exchange is often required to terminate the connection. This is often expressed by describing the connection as passing through an establishment phase, a data transfer phase, and a termination phase. Although data may indeed be transferrable in all phases, such connection-oriented protocols nonetheless have these distinct phases.

"Connectionless" transfers, in contrast, do not require such phases. Individual data transfers are treated independently, with no state information maintained by the communicants between separate data transfers. Connectionless service is often of great utility for applications which cannot tolerate the overhead of connections, and which do not need the enhanced services provided by a connection-oriented protocol.

Of particular importance within the DoD Reference Model is the use of connectionless transfers to support internetworking. This approach arises from several considerations, including:

1. The variety of approaches taken within separate networks for connection-oriented services precludes the easy concatenation of distinct connections across multiple networks.
2. Since the concatenation of reliable connections does not necessarily yield a reliable connection, "end-to-end" acknowledgments and retransmissions are required to achieve reliability. This removes the necessity of reliable connections across individual networks.
3. All networks can support the independent transfer of separate packets from an entry point to an exit point, with some probability that packets may be lost, duplicated, or otherwise mishandled.
4. Some internetwork applications do not require reliable connections.

These considerations argue for the "datagram" approach to internetworking taken by the DoD Reference Model.

Connection-oriented protocols typically allow users to establish multiple simultaneous connections at a given service-access-point. Such connections

are identified via a "connection-endpoint-identifier", which is a locally significant identifier used between the user entity and the connection-oriented protocol entity to name the connection.

3.6 Layers

The DoD Reference Model draws upon the concept of "layers". The layer concept is an explanatory technique allowing for multiple protocols of a similar nature to be described.

Distinct protocols may have some features in common. For example, all protocols using the services of a given protocol will have at least that in common: they require similar services. Other protocols may provide similar services to their users, such as "process-level addressability", or "virtualization of data formats". The attempt to describe the similarities within such sets of protocols can be a useful exercise.

The only restriction placed on protocol architectures by good engineering practice is that they be hierarchically arranged. Protocols which occur in a similar vertical position within a hierarchy are often found to have significant features in common beyond their (perhaps coincidental) architectural placement. By placing protocols in layers, it is hoped that some general common features of a set of protocols may be easily described.

One of the major purposes of the DoD Reference Model is as an explanatory document covering DoD standard protocols. As an agent of exposition, the layer concept promises to aid the reader in his understanding, and for this reason the DoD Reference Model is organized by layers. However, it must be remembered that the placement of a protocol within a specific protocol architecture depends upon the specific services that it requires and the specific services that it provides. Such considerations may make it difficult to properly place a protocol within a specific layer without misleading the reader.

Given the descriptive nature of the Reference Model layers, the following principles apply:

1. Each layer may consist of multiple protocols.
2. In general, protocols within a layer have features in common which can be described by describing the layer.
3. In general, a protocol in a layer uses the services of protocols in the next lower layer, and provides services to the protocols of the next higher layer.

However, the Reference Model in no way prohibits the direct use of any protocol's services by any higher layer entity. In this respect, the layer concept is seen not to be a fundamental principle of the DoD Reference Model. The fundamental principle is the arrangement of protocols into a hierarchy; the use of layers allows for general features of the hierarchy to be described.

30 September 1982

-18-

System Development Corporation
TM-7172/201/01

The layers of the the DoD Reference Model include the following:

- Presentation Layer, containing protocols which perform virtualization of data formats for specific types of applications.
- Session Layer, containing protocols which help to coordinate the establishment of user communications through the mediation of specialized entities.
- Transport Layer, containing protocols which provide for process-to-process communication across one or more networks.
- Internet Layer, containing protocols which perform routing between networks.
- Network Layer, containing network-specific protocols which allow for data transfers across a single network.
- Link Layer, containing protocols which manage the transfer of data across a single physical channel.

The descriptions of these layers form the major sections of the DoD Reference Model.

4. CURRENT PROTOCOLS AND THE BASIC MODEL CONCEPTS

One of the primary goals of the reference model development effort is to maintain architectural compatibility with the DoD internetting architecture (TCP/IP). In this section, we examine how closely the fundamental mechanisms within TCP/IP fit the basic concepts described in the previous section.

4.1 Layering

The current DoD protocols are commonly placed within a four layered hierarchy:

application	(FTP, Telnet, Name Server Protocol ...)

transport	(TCP, UDP)

internet	(IP)

subnetwork	(1822, X.25, ...)

Each layer has its "entities" (the various protocol interpreters), and entities within a layer may provide services to entities within higher layers.

The layering of the DoD Reference Model is very similar, with two exceptions:

1. The protocols within the topmost layer within the above diagram have only one thing in common: they use the services of TCP and UDP. The DoD Reference Model refines this layer by recognizing that some of the protocols provide "virtualization of data format" services (and are placed in the Presentation Layer), while others involve interactions with specialized entities (such as name servers or access controllers) which help to manage the data transfers required by the Presentation Layer or other entities. These other protocols are placed within the Session Layer for expository purposes.
2. The above picture does not refine the manner in which subnetwork services are accessed into a layered structure. The DoD Reference Model includes a separate Link Layer, recognizing that such link protocols do form an important class of protocols, providing services which are often quite distinct from network protocols. Standardization of such protocols will continue to be an important DoD concern, and consequently the DoD Reference Model explicitly identifies a Link Layer.

4.2 Service-Access-Points

In the DoD Reference Model, an (N)-entity accesses the services of an (N-1)-protocol at a "service-access-point". It is the service-access-points of a protocol which are assigned "addresses", allowing those (N)-entities currently using an (N-1)-service-access-point to be referenced via the access-point's address.

Similar concepts are found within TCP and IP; however, they don't quite match the exact concept of a service-access-point and the related addressing architecture of the Reference Model.

4.2.1 TCP Service-Access-Points

The TCP "port" seems to naturally correspond to the service-access-point concept. Remote entities can be reached by referring to a port to which they are currently attached. Multiple connections can be terminated at a single port. The binding between a higher-level entity and a TCP port need not be permanent. All these are properties of service-access-points as well.

However, TCP ports play an additional role which distinguishes them from service-access-points. A TCP connection is determined by the pair of ports at the two ends. It is not possible to have multiple simultaneous TCP-connections between the same pair of ports. Thus ports are used to identify connections in addition to their role as access-points for the TCP service.

This aspect of TCP ports is similar to the concept of a "connection-endpoint-identifier" within the Reference Model, which is discussed below. Thus the correspondence between TCP-port and service-access-point is not exact.

4.2.2 IP Service-Access-Points

IP includes a Protocol Identifier mechanism for distinguishing among multiple higher-level entities. This shares many properties with the concept of a service-access-point.

The primary difference between the "Protocol Identifier" and a true service access point is the implicit correspondence between the Protocol Identifiers at the two ends of a transfer. Although IP implementations could perhaps allow the protocol entity with identifier 5 to exchange datagrams with a protocol entity with identifier 7, such is not ordinarily the case.

In contrast, service-access-point addresses need not have any relationship at the two ends of a data transfer.

4.3 Connections and Connection-Endpoint-Identifiers

The ISO Reference Model is heavily connection-oriented; the DoD Reference Model incorporates connectionless transfers as well. Nonetheless, connections still play an important role in the DoD Reference Model, and the key architectural concept of "connection-endpoint-identifier" must be compared with similar notions in TCP.

30 September 1982

-21-

System Development Corporation

TM-7172/201/01

In the DoD Reference Model, there may be any number of simultaneous connections between two service-access-points. These connections are distinguished by "connection-endpoint-identifiers" which have purely local significance; i.e. they can be different at the two ends of the connection.

TCP connections are between port pairs; only one connection may exist between a port pair. Thus connections are identified at both ends by the port pair. This differs from the connection-endpoint-identifier concept in two ways:

1. Port pairs do not identify connections in a manner which is of only local significance.
2. Multiple simultaneous connections are not possible between the same port pair.

The first item is relatively insignificant, since additional local identifiers can always be introduced. The second item, however, is a major architectural distinction between TCP and the DoD Reference Model.

5. THE PRESENTATION LAYER

5.1 Differences between DoD and ISO Presentation Layers

The Presentation Layer allows user-entities to communicate without the necessity of a common syntax. By providing the necessary syntax transformations (while preserving meaning), protocols within the Presentation Layer avoid the $n \times m$ problem which would result if user-entities needed to translate the syntaxes of every other user-entity.

The primary model for Presentation Layer operation (at least implicitly) involves the concepts of "transfer syntax" and "local syntax". In this model, a presentation-entity exchanges data with its local user-entity using a local syntax which they agree upon. For transfers between presentation-entities, data is represented in a transfer syntax which may or may not be different from the original local syntax. The receiving presentation-entity passes the data to its user-entity using a third syntax, local to that interface.

A more general model involves the notion of "network virtual resource". Here the user-entities are provided with a local representation of a shared resource (e.g. a terminal, or a file) which they can manipulate using locally-defined access methods. The presentation-entities map the local representation's data structures into a "standard" (or "virtual") representation of the resource. Modifications made by an user-entity to its local representation of the resource are communicated between presentation-entities as changes to the virtual resource. Upon detection of a change in the virtual resource, presentation-entities make appropriate modifications to its user-entity's local representation.

Actually, the virtual resource only exists as "images" of its current state as perceived by the different presentation-entities. Since updates to the virtual resource's state cannot be instantaneously communicated between presentation-entities, some temporary inconsistencies may exist among the different views. Presentation Layer functions must ensure that such inconsistencies do not lead to deadlocks or permanent data misrepresentations.

With the virtual resource approach, communication between user-entities occurs through changes in the state of a data structure which represents a shared resource. User-entities induce and detect these changes by accessing their local representation of the shared resource. There need be no explicit "connection" between user-entities over which data is to be "transferred". Of course, connections of some sort might exist at lower layers to facilitate the communication of resource state changes between presentation-entities, but such connections can be hidden from the user-entities. The advantage of such an approach is that multiple user-entities can be accessing the same shared resource.

The ISO Presentation Layer includes some of the concepts of network virtual resources. In particular, the concept of "presentation-image" is introduced as the presentation-entity's data structure representing its local view of the shared resource's state. However, there is no analogous concept introduced giving the user-entity's local data structure - i.e. there is no "local

representation" in the sense used above. In fact, communication between user-entities is couched in "connection-oriented" terms. Data is transferred over presentation-connections, with "code and character conversions" performed in the Presentation Layer.

It is not clear in the ISO Model how the presentation-entities are to use session-connections to provide the presentation-service. It is stated that "there is a one-to-one correspondence between presentation-service-access-point-address and session-service-access-point-address. There is no multiplexing in the Presentation Layer". However, one can easily invent situations in which multiple session-connections are required to support a single data structure for user-entity access (for example, one session-connection for real-time images with a second for the cursor). The ISO description does not make explicit whether such an application would require two separate presentation-connections (which would then be synchronized within the Application Layer). Such a problem does not arise with the modifications proposed for the DoD Presentation Layer.

The DoD Presentation Layer explicitly incorporates the concept of user-entity communication via manipulation of a shared resource. Multiple images of the resource's state are accessed by the entities involved (local representations accessed by user-entities, which are mapped into the virtual representation by the presentation-entities). The Presentation Layer is responsible for ensuring consistency among the various representations. This is accomplished via communication between the presentation-entities.

Data transfers between presentation-entities will in general use the services of protocols within the Transport Layer. Session Layer services (e.g. access control, name service) may also be invoked to support the management of the user communications.

This approach has the following advantages for DoD:

1. It is more consistent than the ISO Presentation Layer, in which user-entities access presentation-connections (no explicit data structure) which are mapped to presentation-images (a data structure) within the Presentation Layer.
2. Information of vastly different formats (e.g. the image and cursor data described above) can easily be accommodated within the "local representation" data structure to be accessed by an user-entity. Such a situation is not easily managed within the ISO presentation-connection approach. Since "multimedia" services are anticipated to be a major DoD requirement, it is important that they be handled cleanly.
3. The shared resource can be accessed by more than two user-entities. For example, a virtual file service may be distributed over multiple nodes. The distributed nature of such a service could be hidden from the user user-entity. Such a presentation-service would make effective use of the multi-entity "session" construct in the DoD Session Layer.

4. Teleconferencing (another important future DoD service) requires a sophisticated Presentation Layer. Multiple user-entities must communicate, either through a conference controller or "directly". Syntax negotiations and transformations are complex due to the multiplicity of participants. Data from multiple sources must be time-multiplexed to the receiving user-entity. It is anticipated that sophisticated Session Layer protocols will be required to support such teleconferences, in conjunction with Presentation Layer protocols which can provide for multi-entity access to a single shared resource.

Finally, the issue of quality of service is absent from the ISO Presentation Layer. In the DoD model this is explicitly addressed. An user-entity can negotiate, for example, delay and reliability parameters which affect the manner in which the presentation-entities communicate resource state information. For the most part, these parameters will be merely passed down to the lower protocols uninterpreted by the presentation-entities. However, they may affect internal Presentation layer functions such as data compression.

In summary, the differences between the DoD and ISO Presentation Layers are:

1. User-entities may be provided a data structure which represents the state of a shared resource in a syntax appropriate to the local system.
2. User-entities communicate via manipulations of their images of the shared resource. The Presentation Layer protocols are responsible for ensuring the consistency of the different images.
3. Multiple user-entities can share a resource. Such entities can even exist in different systems - in this sense the resource can be distributed.
4. The data structure provided by a presentation-entity for an user-entity's local image can represent multimedia resources.
5. Quality of service parameters can be negotiated to determine the way in which presentation-entities exchange the information necessary to maintain consistency across the multiple images.

The DoD Presentation Layer is viewed as the proper location of a wide variety of sophisticated services required for future DoD communications: virtual terminal services, virtual filestores, distributed database access, teleconferencing, and advanced messaging support. Whereas the ISO Presentation Layer is capable of eliminating problems of syntax distinctions at the ends of a point-to-point connection, the DoD Presentation Layer can handle syntax problems of multi-entity applications - even hiding the distributed nature of a shared resource. The possibility of multi-entity coordination provided by the DoD Session Layer allows for these functions to be placed coherently within the DoD Presentation Layer.

5.2 THE DoD PRESENTATION LAYER

5.2.1 Purpose

The purpose of the protocols within the Presentation Layer is to represent information to communicating user-entities in a way that preserves meaning while resolving syntax differences. These protocols allow user-entities to manipulate remote resources using locally defined access methods. In this way user-entities do not require knowledge of how the resource may be represented to other user-entities.

In the DoD architecture, the syntaxes used by application processes that wish to communicate may be very similar or quite dissimilar. When they are similar, the data transformation and formatting functions may not be needed at all; however, when they are dissimilar, protocols within the Presentation Layer provide the means to converse and decide where needed syntax conversions will take place.

User-entities communicate by manipulating a shared resource. Such manipulation is achieved through local manipulation of the resource's local representation as provided by the local presentation-entity. It is the responsibility of the Presentation Layer protocols to ensure that the resource's representation to other user-entities properly reflect such manipulations.

5.2.2 Services Provided to Users

Protocols within the Presentation Layer may provide the following services:

- a) selection (through negotiation) of syntax for local representation of the resource;
- b) manipulation primitives allowing modification of local representation;
- c) communication of local changes to those remote entities sharing access to the resource;
- d) access primitives allowing those changes to the resource made by remote entities to be visible to the local user-entity (i.e. remote changes are reflected in the local representation);
- e) selection (through negotiation) of quality of service to be provided, i.e. delay, reliability, and security requirements for the data transfers necessary to communicate resource modifications among participating presentation-entities.

5.2.3 Functions within the Presentation Layer

When one user-entity modifies a resource, the Presentation Layer protocol must ensure that these modifications are visible to other user-entities which access that resource. This involves the transfer of data between the originating user-entity and its local presentation-entity, between presentation-entities, and between a second presentation-entity and its user-entity during

the resource access. There are several syntactic versions of the data being transferred: the syntax used by the user-entity of the originator of the data, the syntax used by the user-entity accessing the data, and the syntax used to transfer the data between presentation-entities ("transfer syntax"). It is clearly possible that any two or all three of these syntaxes may be identical. Protocols within the Presentation Layer contain the functions necessary to transform between the transfer syntax and each of the other two syntaxes as required.

There is not a single predetermined transfer syntax for all systems conforming to the DoD architecture. The transfer syntax to be used on a presentation-connection is negotiated between the correspondent presentation-entities. Thus, a presentation-entity must know the syntax of its local system and the agreed transfer syntax. Only the transfer syntax needs to be referred to in the Presentation Layer protocols.

A user-entity's local representation of a resource is provided by its local presentation-entity. Manipulations of the local resource representation by the user-entity are made visible to the remote user-entities via corresponding changes in the remote resource representation. Manipulations of the resource by remote user-entities will be visible to the presentation-entity via changes in the local representation.

Protocols within this layer perform the following functions to help accomplish the above services:

- a) establishment of communications using the services of Transport Layer protocols, which may also involve invocation of Session Layer services
- b) negotiation of the transfer syntax;
- c) data transformation and formatting;

Syntax negotiations consist of the dialog between the presentation-entities on behalf of the user-entities to determine the form that data will have while in the DoD internetworking environment. The negotiations will determine what conversions are needed (if any) and where they will be performed.

5.3 Current DoD Presentation Protocols

The proposed DoD Presentation Layer is intended to provide applications access to a "virtual resource". This service includes not only the code and format transformations provided by ISO presentation-protocols, but also can make a "distributed resource" (e.g. a distributed file system) appear (if desired) as a single data structure.

No existing DoD protocol is currently structured in this fashion to allow uniform access to a distributed resource. However, Telnet and FTP provide example implementations of the virtual resource concept. The proposed DoD architecture should allow the introduction of more sophisticated resource management techniques required for multi-media connections (e.g. simultaneous image and data), teleconferencing, and distributed database access. Such applications may require multi-entity coordination protocols within the DoD Session Layer.

The model followed within the DoD Presentation Layer involves "images" of a data structure which represent a local entity's current view of that structure's state. It is the responsibility of the Presentation Layer to ensure that the multiple simultaneous (distributed) images are updated appropriately to prevent serious long-term inconsistencies. For example, a presentation protocol may be introduced which "presents" to its attached user-entity a representation of a graphic terminal screen with cursor information. This local representation would be in fact a data structure updated appropriately by the presentation-entity in response state-information transmitted by a correspondent presentation-entity. The transfers of such state-information between presentation-entities may involve multiple connections provided by the lower level protocols, (for example, separate connections for image and cursor data). These multiple lower-level connections would not be necessarily visible to the user-entities, who are merely accessing a certain data structure provided by the presentation-entities.

A wide variety of possible presentation-protocols are thus envisioned for the DoD Presentation Layer, supporting the "virtual resource" requirements of DoD applications. The model described above is consistent both with "point-to-point" presentation services required for e.g. virtual terminal support, as well as more complex distributed resources which are anticipated for the future DoD environment.

6. THE SESSION LAYER

6.1 Differences between DoD and ISO Session Layer

The intent of the Session Layer is to help communicating user-entities manage their data transfers. The ISO Session Layer provides services allowing two presentation-entities to establish a session-connection between themselves. A session-connection is mapped onto a transport-connection, but the state of a session-connection can be maintained even in the event of transport-connection failure.

The services provided through ISO session-connections additionally include:

- quarantine service, allowing a sender to request that certain data not be made available to the receiver until explicitly released by the sender,
- interaction management, allowing the users of a session-connection to exchange the "turn",
- expedited data transfer, and
- synchronization services, allowing presentation-entities to "mark" specific synchronization points and reset the session-connection.

Although these services are useful, the Session Layer's enhancement of basic Transport Layer services are so minor that it is hard to justify their incorporation in a separate layer. If one informally views the concept of a "session" as a well-defined period of time during which application-entities coordinate their actions (through communication) to perform a specific job, it is clear that the ISO Session Layer is deficient in several respects.

First, the session support services (e.g. quarantine and interaction management) seem to be directed towards record-oriented transaction exchanges. Voice and other real-time stream applications are not adequately supported by the ISO Session Layer. This is particularly apparent when one examines the issue of quality of service negotiation between the ISO Session Layer and its users: there is no discussion of this in the ISO document. Flexible support for a wide variety of application traffic is particularly important for DoD.

Secondly, session-connections will only support the coordinated activity of two presentation-entities. The coordination of multiple application entities communicating among themselves is left entirely up to the higher layers. Since the ISO document nowhere explicitly addresses this issue, we must assume that this is intended to be the user's responsibility.

The DoD Reference Model takes a completely different approach to the Session Layer. There seems to be a class of protocols which involve the interaction between "users" and one or more "specialized" entities, particularly prior to the establishment of user-user communication. For example:

- A user wishes to establish communication with a "named" resource, which requires interaction with a name server to determine the appropriate

location information.

- The network administration wishes to restrict user access to certain resources, and accomplishes this by forcing every user connection establishment to be mediated by an access controller.
- A low-priority user is dominating a particular resource, and an agent of the network administration must preempt this user's communication on behalf of a user with a critical current need.
- Upon termination of a user connection, usage statistics must be communicated to a network accounting agent.

The protocols required for such interactions seem to require Transport Layer services - e.g. UDP or perhaps TCP. In addition, a variety of protocols "related" to those in the above examples seem to follow the same pattern: for example, protocols to register name/address bindings at a name server, or protocols to distribute keys which enforce access control decisions.

It is clear that DoD will develop standard protocols for such functions. It is possible to conceive of future protocols which similarly involve interactions with multiple specialized entities, supporting sophisticated multi-user presentation-level protocols for teleconferencing, distributed file systems, and other future DoD applications. The reliable management of critical distributed applications in the presence of node or link failures will require coordinating the activity of redundant entities. Thus the important DoD concept of survivability requires that mechanisms for distributed application control exist in the architecture. Such protocols are considered to be session protocols within the DoD Reference Model. This is a very different Session Layer from that within the ISO Model.

The ISO Model does not explicitly address how such multi-entity/multi-connection applications might be supported within the architecture. Thus ISO's is a "point-to-point" architecture, giving the rules governing communication between two application-entities. The manner in which these two entities may be participating in a coordinated action with other entities is irrelevant to the ISO model. Management of such multi-entity functions is placed entirely within the user domain.

Although the above applications are quite different in their specific requirements, the manner in which the necessary user-user communications might be established, monitored, and terminated shows some similarities. Typically at the time of session establishment, some interaction is required between the user entities and some specialized control entities (e.g. name servers, access controllers, or key generators). Other entities may play an additional controlling or monitoring role during the session, requiring perhaps the ability to query the status of the participating entities. Finally, session termination may again require interaction with designated specialized entities.

Some of these functions are similar to those provided by the ISO Session Layer for point-to-point session-connections, e.g. interaction management and maintenance of session-connection state in the face of transport-connection

abortions. Use of multiple transport-connections between two session-entities has already been incorporated within the ISO Model.

The proposed DoD Architecture differs from the ISO Session Layer in its concept of "session". A session is established among user-entities. Session establishment may involve interaction with specialized session-entities. Once established, a session may in fact be identical to a single transport-connection - that is, the session-entities which participated in the establishment of the session may play no role once the session is established. This is in strong contrast to the basic principle in the ISO Model that every data transfer must be "touched" by a session-entity.

Some specialized session entities may be capable of forcibly terminating a user session, e.g. for security or precedence reasons. Thus the session concept fits in with fundamental DoD requirements.

A variety of session protocols are anticipated, supporting different applications and providing different management services. Some session protocols (e.g. access control protocols) may be virtually invisible to the users. Others (such as those providing name service) may be seen by the users as providing substantial additional functionality over the "raw" transport layer services.

In summary, there are three ways in which the DoD Session Layer differs from the ISO Session Layer:

1. Quality of service parameters are explicitly described which are negotiable across a session-service-access-point, ensuring that real-time and other classes of applications will be adequately supported by the Session Layer protocols.
2. Distributed applications are directly supported through the "session" concept, which allows for the existence of specialized administrative entities which participate in the establishment, monitoring, and termination of user communications.
3. The DoD Session Layer incorporates some of the access control, accounting, and authentication functions necessary for management of the inter-network system.

This approach to the Session Layer seems to be particularly relevant to anticipated DoD applications. The placement of some access control and name service functions within the layer points towards future DoD standards activity. In addition, the enhanced robustness and control provided by anticipated session layer protocols serve some of DoD's most fundamental requirements.

6.2 THE DoD SESSION LAYER

6.2.1 Purpose

The purpose of the Session Layer is to provide the means necessary for cooperating user-entities to organize and synchronize their data exchanges. To do this, protocols within the Session Layer provide services allowing the establishment of a session among user-entities, and to support their orderly data exchange interactions. Proper management of session establishment, session data transfers, and session termination may involve the coordinated action of various session-entities, both on behalf of the participating user-entities and on behalf of system and network administrations.

To implement the transfer of data between the user-entities, a session uses the services of the Transport Layer. These services may be provided by a connection-oriented or by a connectionless transport protocol. Even when the underlying transport service used for data transfers is connectionless, it may be desirable to invoke session-establishment and management services for the purposes of name resolution or access control. Transport Layer services are used to support the exchanges required for proper management of the session establishment and termination. These need not be provided by the same transport protocol which is used for transferring user data.

Sessions among user entities are created when requested by a user-entity at a session-service-access-point. During the lifetime of a session, session services are used by the user-entities to regulate their communication, ensuring orderly message exchange across the session. The session exists until released by the user-entities, unless a privileged session-entity participating in the overall management of the session requests the session's termination. Such may occur in situations requiring forced session termination either for precedence or security considerations. While the session exists, session services maintain the state of the session even over data loss by the Transport Layer.

A session may be initiated by a user-entity, and may be joined by other user-entities. Session establishment may involve the participation of several specialized session-entities in order to provide services such as:

- a. name to session-address mapping
- b. determination of predefined session attributes
- c. session authorization and access control
- d. session accounting
- e. coordination of multi-user sessions.

Such specialized session-entities may act on behalf of system and network administrations.

User-entities may establish sessions by specifying a the intended session's attributes. Such attributes may include transport-addresses of intended participants, access control requirements, specification of redundant entities, and required transport-connections. Modification rights to some session-type attributes may be restricted to certain authorized system or network administration entities.

6.2.2 Services Provided

6.2.2.1 Attribute Definition

Protocols within the Session Layer may provide a session attribute definition service, which enables user-entities to define the attributes of a session type. These attributes determine actions to be taken by session-entities in the management of a session of the given type. Such attributes may include the following:

- a. access restrictions
- b. necessary user-entity session participation (e.g. for monitoring or accounting purposes)
- c. transport services required for support of the session

The entity attribute definition service enables user-entities to make their own attributes known to the session-entities. Such attributes may include the following:

- a. user-titles (allowing the Session Layer to determine "name/address" bindings)
- b. availability for remotely initiated sessions
- c. redundancy of role with other user-entities (e.g. back-ups) within particular session types.

Declaration of session type attributes and entity attributes may be independent of any particular session establishment. Some user-entities may be privileged to manipulate attributes of other user-entities.

6.2.2.2 Session Establishment

Session establishment services enable user entities to establish a session among themselves.

The session establishment service allows the user-entities cooperatively to determine the values of session attributes at the time the session is established. The session establishment request may also reference a predefined set of session attributes (identified by a session-type-identifier).

The session establishment service provides to the user-entities a session-identifier which uniquely specifies the session within the environment of the

participating session-entities. This identifier may be used by the user-entities to refer to the session during the lifetime of the session.

6.2.2.3 Session Termination

The session-termination service allows the user-entities to release a session in an orderly way without any loss of data. In addition, a session may be terminated by a specialized session-entity participating in the overall management of the session. Such a specialized session-entity may be acting on behalf of system or network administrations. Participating session-entities may be informed of the reason for session termination.

6.2.2.4 Session Abort

The session abort service informs the user-entities of session aborts. A session may be aborted due to unrecoverable errors at the Transport Layer, or upon demand by a specialized session-entity participating in the overall management of the session. Such a specialized session-entity may be acting on behalf of system or network administrations.

6.2.2.5 Quality of service negotiation

Prior to the establishment of each session, the correspondent user-entities and their associated session-entities must agree on the quality of service to be provided over each session. This is achieved through the negotiation of session quality of service parameters.

The following parameters affecting the session establishment, data transfer, and termination phases can be negotiated:

Establishment Phase

- establishment delay
- security classification
- precedence

Data Transfer Phase

- bit reliability
- delivery reliability
- sequence reliability
- absolute delay
- delay variance
- "reliability versus delay"

Termination Phase

- reliability of termination
- termination delay

Since the Session Layer is not primarily responsible for maintaining the end-to-end quality of service for data transfers, these quality of service parameters for the most part are passed down to the Transport Layer for interpretation.

Protocols that provide reliable stream, real-time stream, and transaction classes of service are anticipated to be important service offerings of the Transport Layer. Consequently these same services are available to the users of the Session Layer for their data transfers. A session protocol may, however, provide some enhancement even during the data transfer phase. For example, "real-time" session-layer services would provide some assurances of delay and delay variance minimization, corresponding to the transport service which is used. If serious problems within the Transport Layer or below lead to the disconnection of a transport-connection supporting a real-time session, the session protocol could take steps to immediately establish a new transport-connection.

6.2.2.6 Normal data exchange

The normal data exchange service allows a sending user-entity to transfer a unit of data to a receiving user-entity.

6.2.2.7 Exception reporting

The exception reporting service permits the user-entities to be notified of unanticipated situations not covered by other services, such as unrecoverable session malfunctions.

Session attributes as determined during the session's establishment may require that the session-entities take special action upon occurrence of specified exception conditions such as failure of a session or of an entity. Such actions may include notification of designated user-entities and establishment of new transport-connections (e.g. to a designated "back-up" entity).

6.2.3 Functions within the Session Layer

The functions within the Session Layer are those which must be performed by a session-entities in order to provide the services required of a specific session protocol.

6.2.3.1 Session establishment and control functions

A protocol within the Session Layer may include functions which allow user-entities to define and modify session and entity attributes. These attributes could be used to determine how sessions of a particular type are to be managed.

Maintenance of this attribute information may be the responsibility of specialized session-entities, (e.g. name servers or access controllers). Such specialized session-entities may be involved during the establishment of a session. Other specialized session-entities may be required to participate in the management of a session during its lifetime. Transport services will be used for information exchanges between these session-entities.

6.2.3.2 Normal data exchange

Session protocols do not substantially add value to the normal data exchange services provided by the transport protocols. In fact, once the appropriate transport service has been decided upon to support a given session, it may be that session entities no longer play a role, having existed primarily to help establish the required transport service on behalf of the users. With such a session service, the users would use transport services directly upon establishment of the session.

6.2.3.3 Session recovery

In the event of reported failure of an underlying transport-connection, a session protocol may contain the necessary functions to regain a transport-connection to support the session, which continues to exist. The session-entities involved would notify the user-entities via the exception reporting service that service is interrupted and would restore the service only as directed by the user-entities. This permits the user-entities to resynchronize and continue from an agreed state.

Alternatively, the session-entities may take restorative action without intervention from the user-entities. In this case the user-entities would be notified that such an event has occurred.

Restorative actions may include establishment of transport-connections to designated entities. Session attributes (as determined at the time of session-establishment) can indicate what actions should be taken by the session-entities in response to specified exception conditions.

6.2.3.4 Session Termination

A protocol within the Session Layer contains the necessary functions to release the session in an orderly way, without loss of data, upon request by the user-entities. The Session Layer also contains the necessary functions to abort the session with the possible loss of data.

6.3 Current DoD Session Protocols

At present, there are no DoD standard protocols which explicitly conform to the proposed DoD Session Layer. A variety of access control and authentication protocols have been developed; the incompatibilities among these protocols points to the necessity for future standardization activity in this area.

The Internet Name Server Protocol [7] is in many respects an example of a session protocol as described in the DoD Model, particularly in its extended form (allowing TCP and UDP ports to be assigned to names). Companion protocols (yet undefined) allowing the registration of names within the name database would also be examples of session protocols.

Teleconferencing, multi-media services, distributed database access, and critical data collection processes may all require the coherent management of multiple entities, perhaps through multiple connections established in concert with appropriate session protocols. These DoD applications are not well enough defined at present to allow for the identification of session services which might be desired.

7. THE TRANSPORT LAYER

7.1 Differences between DoD and ISO Transport Layer

The Transport Layer incorporates end-to-end functions which enhance the data transfer services provided by the Internet Layer.

Relatively minor differences exist between the DoD and ISO Transport Layer. One difference is the enhanced emphasis on connectionless as well as connection-oriented services. For example, UDP is a perfectly legitimate transport protocol.

Another difference lies in the wording of the quality of service references. The DoD Transport Layer must offer a flexible range of services, supporting bulk transfer applications such as file transfers, delay-sensitive applications (e.g. voice), and transactions. Although the ISO model intends to support such a variety of services, the Transport Layer is heavily connection-oriented, with an inadequate discussion of quality of service negotiable parameters (for example, delay variation is not included).

The following paragraphs briefly describe the differences between the DoD and ISO Transport Layers.

References to the lower layer service reflect the DoD Internet Layer's connectionless orientation.

ISO's discussion of quality of service parameters has been replaced. Specific parameters are identified, which are closely matched with the parameters negotiable across the other layer interfaces (e.g. between the transport and network entities). These parameters are intended to be sufficient to allow identification of a broad set of service classes which are required to support DoD applications. These include delay-sensitive and reliable transaction services.

Graceful closes have been explicitly introduced as an important service which can be requested for a transport connection. According to the negotiated termination quality of service, data delivery during the termination phase may or may not be guaranteed - the termination is "graceful" if delivery is guaranteed. This can be viewed as part of the general quality of service structure.

7.2 THE DoD TRANSPORT LAYER

7.2.1 Purpose

Protocols within the transport layer provide transparent transfer of data between transport addresses. Transport Layer protocols relieve the transport users from any concern with the detailed way in which cost effective transfer of data meeting the users' service requirements is achieved.

The transport users are identified to the Transport Layer by transport-service-access-point addresses ("transport-addresses"); the data transfer service is provided to the addressable entities without regard to their location. Transport-addresses are drawn from a large enough address space to ensure that many users within a single host may be distinguished. Transport protocols are commonly referred to as "process-to-process" protocols, as opposed to internet or network level protocols which will typically allow addressing of individual hosts but not of large numbers of users within hosts.

Different protocols within the Transport Layer provide different services. For example, these different services may include reliably sequenced transfers over transport-connections, unacknowledged datagram transfers, and delay-minimized stream transfers for "real-time" applications. Each protocol may also allow additional quality of service parameters (specific to that protocol's service) to be negotiated between the user-entities and the transport-entities.

The transport-service is provided by the Transport Layer performing all necessary functions in conjunction with the utilization of the most appropriate underlying facilities and quality of service available from the protocols within the Internet Layer.

The Transport Layer is required to optimise the use of the available communications resources to provide the performance and level of security required by each communicating transport user at minimum cost. This optimisation will be achieved within the constraints imposed by considering the global demands of all concurrent transport users and the overall limit of resources available to the Transport Layer.

Since the internet service provides transport of data units from one transport-entity to another, including the case of using multiple networks, and relieves the Transport Layer of any concern with switching, routing, and relaying, all protocols defined in the Transport Layer will have end-to-end significance, where the ends are defined as the correspondent transport-entities which may reside in hosts attached to different networks.

Transport functions resident in the Transport Layer allow the Internet Layer to use more than one communication resource (e.g. the transfer of data units by the Internet Layer may involve the use of a public packet switched network, used in tandem with a circuit switched network).

The transport functions invoked in a transport protocol to provide a requested service quality may depend on the quality of the internet-service.

7.2.2 Services Provided

7.2.2.1 General

A protocol within the Transport Layer uniquely identifies each of its users by its transport-address. Users of a transport protocol are provided with the means to request the transfer of data between two transport-addresses with specified delay, reliability, security, and other requirements. The user of the Transport Layer must determine which transport protocol is best able to provide the desired service.

Some transport protocols may require the establishment of a transport-connection to provide their service. Such connections have a connection-establishment phase, data transfer phase, and connection-termination phase. Quality of service parameters for the connection-establishment phase, data transfer phase, and the connection-termination phase may be specified by the user.

Other transport protocols may not require the establishment of a connection. Such connectionless services could be provided with a variety of reliability and delay quality of service parameter values.

Connection-oriented transport protocols may allow more than one transport-connection to be established between the same pair of transport-addresses; the means by which the user can distinguish between the transport-connection-end-points will be provided by the Transport Layer, in terms of "transport-connection-end-point-identifiers".

The existence and performance of each transport-connection is independent of all other such connections, except for the limitations imposed by finite resources available to the transport protocol providing the service.

7.2.2.2 Establishment services

If a transport protocol provides connection-oriented service, then the following services are provided by the transport protocol at the transport-service-access-point:

a) Transport-connection establishment

Transport-connections are dynamically established to a peer transport-address. The quality of service of the transport-connection may be negotiated among the users and the transport-service via various parameter combinations such as throughput, transit delay, connection set-up delay and various guaranteed values of parameters affecting the connection establishment phase, data transfer phase, and connection termination phase. Such quality of service parameters could include:

Connection Establishment Phase

- reliability of connection establishment

- connection establishment delay
- connection security classification
- connection precedence/priority

Data Transfer Phase

- bit reliability
- delivery reliability
- sequence reliability
- absolute delay
- delay variance
- "reliability versus delay"

Connection Termination Phase

- reliability of connection termination
- connection termination delay

A connection-oriented transport protocol will allow the establishment of transport-connection only when both peer user-entities of the given transport-connection agree on the quality of service selected.

7.2.2.3 Data transfer services

This service provides data transfer in accordance with the agreed upon quality of service. When this quality of service cannot be maintained and all possible recovery attempts have failed, then the transport-connection is terminated and the transport users are notified.

- a) Transport-service-data-unit transfer provides the means by which arbitrarily selected transport-service-data-units are delimited and transparently transferred in sequence from one sending transport-service-access-point over a transport-connection. This service is subject to flow control.
- b) Expedited-transport-service-data-unit transfer provides an additional means of information exchange on a transport-connection. They are subject to their own set of transport-service and flow control characteristics. The maximum size of expedited-transport-service-data-units is limited.

7.2.2.4 Termination services

This service provides the means by which either session-entity can terminate the connection and have the correspondent session-entity informed of the termination.

According to the negotiated termination quality of service, data delivery during the termination phase may or may not be guaranteed.

7.2.3 Functions within the layer

7.2.3.1 General overview

The functions performed by a protocol within the Transport Layer may include:

1. mapping transport-addresses onto network-addresses;
2. transfer of user data between transport-addresses;
3. establishment and termination of transport-connections;
4. end-to-end sequence control on individual connections;
5. end-to-end error detection and any necessary monitoring of the quality of service;
6. end-to-end error recovery;
7. end-to-end flow control on individual connections;
8. supervisory functions
9. expedited transport-service-data-unit transfer.

Different transport protocols will include different functions depending on the actual transport service which they are to supply. For example, a connectionless transport protocol will not include the connection-oriented functions.

The following sections describe the types of functions which may be defined in a transport protocol. Only the first is required for a connectionless transport protocol.

7.2.3.2 Addressing

When a transport-service user requests that data be transferred from one transport-service-access-point to another (which may require the establishment of a transport-connection), the transport entity needs to determine the internet-address identifying the transport-entity which serves that correspondent user-entity, i.e. which maintains that correspondent transport-address.

Because the transport-entities support services on an end-to-end basis by means of end-to-end functions, no intermediate transport-entity is involved as a relay between the end transport-entities. Therefore the internet-addresses on which the transport entity maps transport-addresses are those identifying the end transport-entities.

One transport-entity may serve more than one user-entity. Therefore several transport-addresses may be associated with one network-address within the same transport-entity. Transport protocols include sufficient range within their space of transport-addresses to be able to identify many users; transport-addresses are typically used to identify processes within a host.

Corresponding mapping or switching functions must then be performed within transport-entities to provide these facilities.

7.2.3.3 Connection Multiplexing

Transport connections may be implemented using the basic connectionless internet service, or may be mapped onto internet-connections. In order to optimize the use of internet-connections, the mapping need not be on a one-to-one basis. A cost effectiveness analysis needs therefore to be made in each particular implementation, to determine whether connection multiplexing needs to be performed or not.

7.2.3.4 Phases of Connection Operation

For connection-oriented transport service, the phases of operation within a transport protocol are as follows:

- a) establishment phase;
- b) data transfer phase;
- c) termination phase.

The transfer from one phase of operation to the other will be specified in detail within the protocol for the Transport Layer.

7.2.3.5 Establishment phase

The goal of the establishment phase is to establish a transport-connection between the two transport users. The functions of the transport protocol during this phase must match the requested class of services with the services provided by the Internet Layer, as follows:

- a) select the internet-service which best matches the requirements, taking into account charges for various services; it may, however, be that a particular transport protocol can only use a single internet-service, in which case no selection will be possible
- b) if necessary, decide whether to multiplex transport-connections onto a single network-connection;

- c) establish optimum transport-protocol-data-unit size;
- d) select the functions that will be operational upon entering the data transfer phase;
- e) map transport-addresses onto internet-addresses;
- f) provide a means to distinguish between different transport connections between the same pair of transport-service-access-points (connection identification function);
- g) transportation of user data.

7.2.3.6 Data transfer phase

The purpose of the data transfer phase is to transport transport-service-data-units between the two transport-service-users connected by the transport-connection. This purpose is achieved by means of transmission of transport-protocol-data-units and by the following functions, each of these being used or not used in accordance with the result of the class of service selection performed in the connection establishment phase.

- a) blocking is a function used to collect several transport-service-data-units into a single transport-protocol-data-unit; the destination transport-entity separates the blocked transport-protocol-data-units;
- b) concatenation is a function used to collect several transport-protocol-data-units into a single network-service-data-unit; the destination transport-entity separates the concatenated transport-protocol-data-units;
- c) segmenting is a function used to split a single transport-service-data-unit into multiple transport-protocol-data-units; the destination transport-entity reassembles the segmented transport-protocol-data-units;
- d) multiplexing is a function used to share a single internet-connection used between two or more transport-connections, or to split a single transport-connection onto multiple internet-connections;
- e) flow control is a function used to regulate the flow of transport-protocol-data-units between two transport-entities on one transport-connection;
- f) error detection is a function used to detect the loss, corruption, duplication, misordering, or misdelivery of transport-protocol-data-units;
- g) error recovery is a function used to recover from detected and signalled errors;
- h) expedited data is a function used to bypass the flow control of normal transport-protocol-data-unit; expedited transport-protocol-data-unit flow is controlled by its own flow control;

30 September 1982

-44-

System Development Corporation
TM-7172/201/01

- i) transport-service-data-unit delimiting is a function used to determine the beginning and ending of a transport-service-data-unit;
- j) transport-connection identification is a function to uniquely identify a transport-connection between the pair of transport-entities supporting the connection.

7.2.3.7 Termination phase

The purpose of the termination phase is to terminate the transport-connection and may include the following functions:

- a) notification of reason for termination;
- b) identification of the transport-connection terminated;
- c) possible additional information.

30 September 1982

-45-

System Development Corporation
TM-7172/201/01

7.3 Current DoD Transport Protocols

At present, there are two well-defined DoD Transport Layer protocols: TCP and UDP.

The service offered by TCP connections is directly reflected in the description of "transport-connections". Modifications to the ISO Transport Layer were made to take into account certain services offered by TCP but not included within OSI - for example, graceful closes. UDP fits within the Model as a connectionless transport protocol.

It is clear that additional transport-level services will be required by DoD applications. Real-time connection service (offering "guarantees" on delay and delay dispersion) could also be incorporated within the Transport Layer. Some of the functions performed by NVP [10] may be properly included within a more general-purpose real-time transport protocol.

30 September 1982

-46-

System Development Corporation
TM-7172/201/01

8. THE INTERNET LAYER

8.1 Differences between DoD and ISO Internet Layer

The Internet Layer relieves the higher layers from any concern with how data is routed, allowing the Transport Layer protocols to be "end-to-end".

There are many different approaches to internet and intranet routing. One approach - typified by DARPA IP and the Xerox Pup [10] architecture - is to transfer universal internet datagrams across the various subnetworks by using each subnetwork's individual internal routing scheme between gateways. Thus internet routing is datagram-based, but each of the subnetworks can have connectionless or connection-oriented intranetwork routing independent of the internetting architecture.

Alternatively, internetwork routing can be based on connections. Such an approach is typified by X.75, in which an internet connection is formed by concatenating several intranetwork X.25 connections between gateways.

It seems clear that any general internetting strategy must be capable of working with a wide variety of separate intranet routing mechanisms. Networks offering connection-oriented network service will certainly exist (e.g. X.25 public nets). But datagram networks will also be very prevalent, and in fact will probably come to predominate as local area networks (which are for the most part datagram nets) proliferate.

There is no ISO Internet Layer. The internetting function was intended to be provided through network-connections traversing "tandem subnetworks". This is similar to the connection-concatenation approach of X.75. For the reasons stated in Section 3 of this Reference Model document, DoD requires a different approach.

The basic approach to internetting taken within the DoD Internet Layer is connectionless - the paradigm being IP. However, connection-oriented internet protocols are not precluded by the model; these may be required for the provision of internet services for delay-sensitive applications such as voice.

8.2 THE DoD INTERNET LAYER

8.2.1 Purpose

The Internet Layer manages transfers of data units on behalf of users who may not be attached to a common network. Users of the internet-service are relieved from all concerns regarding the topology of the collection of networks and gateways which together form the Internet.

The Internet Layer includes separate protocols which may provide different types of service. The basic service is a connectionless internet service in which internet-service-data-units are independently transferred between internet-service-access-points. Other internet protocols may provide an internet service which requires the establishment of internet-connections. Such protocols may, for example, provide a service which minimizes delay variations between separate data unit transfers.

The Internet Layer provides to its user entities independence from routing and switching considerations associated with the transfer of internet-service-data-units across multiple networks. It makes invisible to transport-entities how the Internet Layer uses the services of individual networks to provide the internet service.

In other words, the users of the Internet Layer may be end-system oriented. For example, protocols within the Transport Layer operate only between end-systems. Any relay functions of protocols used to support the internet service between the end-systems are transparent to the users of the Internet Layer.

8.2.2 Services Provided

8.2.2.1 General

The basic service of the Internet Layer is to provide the transparent transfer of all data submitted by the users. This service allows the structure and detailed content of submitted data to be determined exclusively by layers above the Internet Layer.

All services are provided at a known cost.

The Internet Layer contains functions necessary to provide the Transport Layer with a firm Internet/Transport boundary which is independent of the underlying networks in all things other than quality of service. Thus the Internet Layer is assumed to contain functions necessary to mask the differences in the characteristics of different transmission and network technologies into a consistent internet service.

When user entities request service from a connection-oriented protocol within the Internet Layer, the service provided at each end of an internet connection shall be the same even in the case of a internet-connection spanning several networks where each of the networks offers dissimilar services.

The Internet Layer includes protocols providing a variety of types of service. Three basic types of service have been identified which require separate internet protocols:

1. Connectionless internet service
2. Reliable sequenced internet service
3. Real-time internet service

Reliable sequenced service and real-time service may require the establishment of internet-connections between the internet-service-access-points.

Users of a particular internet-protocol's service may negotiate additional quality of service parameters to be used in the internet transfers on their behalf.

In the case of connection service, the quality of service parameters are established for each internet-connection. While this quality of service may vary from one internet-connection to another it will be agreed for a given internet-connection and the same at both internet-connection endpoints.

8.2.2.2 Internet-Addresses

Users of the internet service may be identified by means of internet-addresses (i.e. internet-service-access-point addresses). Internet-addresses are provided by the Internet Layer and can be used by user-entities to uniquely identify other user-entities, i.e. internet-addresses are the means by which user-entities can communicate using the internet-service.

This may be independent of the addressing needed by the underlying networks.

8.2.2.3 Internet-Service-Data-Unit Transfer

The Internet Layer provides for the exchange of internet-service-data-units between internet-addresses. These units have a distinct beginning and end and the integrity of the unit content is maintained by the Internet Layer. The internet-service-data-units are transferred transparently between user-entities.

The service offered by a connectionless protocol within the Internet Layer provides no guarantees of internet-service-data-unit delivery to the intended internet-address, nor is the received sequence of internet-service-data-units guaranteed to be in the same order as the transmitted sequence. In addition, internet-service-data-units may be duplicated within the internet.

8.2.2.4 Internet-Connections

If an enhanced level of internet-service (beyond the basic connectionless service) is desired, it may be necessary to establish a internet-connection between the internet-service-access-points, using the services of a connection-oriented internet protocol. Different such connection-oriented

Internet protocols may exist to provide different internet services.

A internet-connection is the means of transferring data between user-entities when reliably sequenced transfers or real-time service is desired. Protocols within the Internet Layer provide the means to establish, maintain and terminate internet-connections.

More than one internet-connection may exist between the same pair of internet-addresses.

8.2.2.5 Internet-Connection-Endpoint-Identifiers

Connection-oriented internet protocols provide to their users an internet-connection-endpoint-identifier which identifies the internet-connection-endpoint uniquely with the associated internet-address.

8.2.2.6 Quality of Service Parameters

Internet Layer protocols offer a variety of classes of service the user entities. User-entities attached to a internet-service-access-point can further negotiate additional quality of service parameters with the associated internet-entity, or request to override the default parameters associated with that class of service.

Data transfers provided by a connectionless internet protocol may also be governed by user-selectable quality of service parameters on each internet-service-access-unit. These parameters may include:

- level of bit reliability
- level of delivery reliability
- delay
- security level

Such parameters would be used by the Internet Layer entities during the routing process (to determine, for example, which of several possible networks to choose).

Internet-service-access-units are delivered to the destination user-entity intact (though with possible bit corruptions commensurate with the desired level of bit reliability). Any fragmentation of the internet-service-access-unit which may have occurred within the Internet Layer is hidden from the user-entities.

Similarly, internet-connections are governed by quality of service parameters negotiated between the user-entities and the internet-entities. The Internet Layer establishes and maintains a selected quality of service for the duration of the connection.

The quality of service parameters include:

- a) Residual errors which may arise from alteration, loss, duplication, disordering, misdelivery of internet-service-data-units, or others;
- b) service availability which is the probability that a requested internet-connection can be established.
- c) reliability, which is the mean time between failures and mean time to repair an established internet-connection;
- d) throughput, which is the information transfer capacity;
- e) transit delay, which includes variations on the transit delay;
- f) delay for internet-connection establishment.

Different internet protocols will offer different standard combinations of these quality of service parameters to provide, for example, reliable sequenced network service or real-time network service.

8.2.2.7 Error Notification

Unrecoverable errors detected by the Internet Layer will be reported to the user-entities.

Error notification may or may not lead to the termination of a internet-connection, according to the specification of a particular internet service.

8.2.2.8 Sequencing

The Internet Layer may provide the service of sequenced delivery of internet-service-data-units over a given internet-connection when requested by the user-entities.

8.2.2.9 Flow Control

A user entity which is receiving at one end of a internet-connection can cause the internet-service to stop transferring internet-service-data-units over the internet-service-access-point. This flow control condition may or may not be propagated to the other end of the internet-connection and thus be reflected to the transmitting user-entity, according to the specification of a particular internet service.

8.2.2.10 Congestion Control

The Internet Layer provides the service of congestion control, which ensures that total network resources are maintained in a non-saturated state. This internet service attempts to avoid serious service degradations caused by extreme quantities of offered traffic. The invocation of congestion control within the Internet Layer may result in the internet-service stopping the transfer of internet-service-data-units over the interface between the

Transport and Internet Layers.

8.2.2.11 Termination Services

A user of a connection-oriented internet service may request termination of a internet-connection. The internet-service may not guarantee the delivery of data preceding the termination request and still in transit, according to the specification of the specific internet-service. The internet-connection will be terminated regardless of the action taken by the correspondent transport entity.

8.2.3 Functions within the Internet Layer

8.2.3.1 General

The functions to be provided within the Internet Layer are outlined below.

8.2.3.2 Addressing

Each internet-entity is uniquely identified within the scope of a single network by its network-address, i.e. the address of the network-service-access-point to which the internet-entity is attached. Such a network-address will be in the particular format defined for the particular network.

To uniquely identify an internet-entity within the scope of the entire internet it is necessary to define network-identifiers. A network-identifier uniquely names a particular network within the internet. The internet is in fact the collection of networks which have been granted a network-identifier by the internet administration. Each internet-entity can then be uniquely identified within the internet via the combination of a network-identifier and a network-address (in the format of the particular network).

A universal format for network-addresses within the internet can be defined. If a particular network uses a network-address format which is not compatible with the universal format, then the internet-entities attached to that network must then be responsible for mapping between network-addresses in the universal format and those in the network's own particular format.

8.2.3.3 Routing

Routing is performed to select the appropriate route between internet-addresses, and for transferring internet-protocol-data-units between them.

Routing functions may involve intermediate nodes, acting as relays between end internet-entities. Such intermediate nodes are called internet-gateways. Internet-entities are responsible for determining if the transfer of internet-protocol-data-units requires that they be routed through an internet-gateway, and for choosing which internet-gateway. This involves determining the network-address of the appropriate internet-gateway, and then using the service of the network layer to effect the transfer.

8.2.3.4 Fragmentation and Reassembly

Internet-entities attached to a given network are responsible for ensuring that internet-protocol-data-units submitted to the network for transmission are not larger than the maximum size network-service-data-unit for that network. Within an internet-gateway, this may involve fragmentation of a received internet-protocol-data-unit into smaller internet-protocol-data-units prior to routing through another network. Internet-entities are responsible for ensuring that such fragmentation is transparent to the users of the internet service.

8.2.3.5 Internet-connections

This function includes mechanisms for providing internet-connections between transport-entities.

Since the basic service of the Internet Layer is provided by a connectionless protocol, the use of an internet-connection will only be required if a quality of service not obtainable from the connectionless service is required. For example, to minimize delay variations it may be necessary to route internet-protocol-data-units through a fixed sequence of gateways which have reserved resources for the internet-connection. Such functions would be the responsibility of the appropriate connection-oriented internet protocol.

Internet-connections may require the use of network-connections.

8.2.3.6 Error detection

Error detection functions are used to check that the quality of service provided at a internet-service-access point is maintained. Error detection in the Internet Layer uses error notification from the Network Layer. Additional error detection capabilities might be necessary to provide the required quality of service.

8.2.3.7 Error recovery

This function includes mechanisms for recovering detected errors. Depending on the quality of the network service provided, these functions may vary.

8.2.3.8 Sequencing

This function includes mechanisms for providing the service of sequenced delivery of internet-service-data-units over a given internet-connection when requested by transport-entities.

8.2.3.9 Flow control

This function includes mechanisms for providing the service of flow control of internet-service-data-units over a given internet-connection when requested by transport-entities.

30 September 1982

-53-

System Development Corporation
TM-7172/201/01

8.2.3.10 Congestion Control

This function includes mechanisms for protecting the internet system from service degradations due to extreme quantities of offered traffic.

8.3 Current DoD Internet Protocols

The primary protocol within the Internet Layer is IP. This is a connection-less protocol which provides a fragmentation and reassembly service as described in the Model. Most DoD applications will use IP for internet transfers.

IP includes within its mechanisms a means by which error reports and other control messages may be exchanged. These mechanisms use the services of IP, and are specified separately from IP as the "Internet Control Message Protocol". However, ICMP is not a "user" of IP in the sense that TCP is a user, rather it is in fact an integral part of IP itself.

Another protocol associated with the Internet system is the "Gateway Gateway Protocol" (GGP). Although in many respects this protocol may be viewed as a Transport Protocol, we prefer to describe it as belonging to the Internet Layer. The reason for this classification is that GGP plays an important role in the overall provision of internet service to users. It certainly does not fit the "process addressability, end-to-end service" model of a transport protocol. GGP is a good example of why the DoD Reference Model takes a much less restrictive view of the layer concept than does the ISO Model; protocols such as GGP are very difficult to fit within a strict layer.

One final DoD internet protocol which should be mentioned is the "Stream Protocol" (ST [11]). This protocol provides internet data transfers for delay-sensitive applications. Resources within gateways are reserved on a "per-connection" basis to ensure that delay variances are minimized between separate transfers. This experimental protocol is intended to support applications such as packet voice, and is a good example of a connection-oriented internet protocol.

9. THE NETWORK LAYER

9.1 Differences between DoD and ISO Network Layer

The Network Layer describes the services which are typically provided by a single network within the Internet.

The ISO Network Layer is heavily oriented towards the concept of the "network-connection". The use of the phrase "tandem subnetworks" is typical of ISO's connection-orientation even for internet routing. It seems the model for Network Layer services assumed by ISO is the connection-oriented X.25 interface offered by public packet-switched networks.

The DoD Reference Model takes a more general view of the Network Layer by recognizing connectionless Network services. As local networks proliferate within the DoD environment, it is likely that networks offering connectionless services will in fact predominate. Since DoD's internetting approach assumes only connectionless service from the constituent networks, and since some DoD higher level protocols are themselves connectionless, the provision of connectionless network service is to be encouraged within the DoD Model.

Thus the basic differences between the DoD and ISO Network Layers are as follows:

1. The DoD Network Layer allows (indeed encourages) connectionless service, whereas the ISO layer is very connection-oriented.
2. All internetting aspects of the DoD Model are placed within the DoD Internet Layer; this issue is confused within the ISO Model, with some saying it is present in the current Network Layer, others saying it requires a separation of the Network Layer into "sublayers".

9.2 THE DoD NETWORK LAYER

9.2.1 Purpose

The Network Layer provides the functional means to exchange network-service-data-units between internet-entities. Such exchanges use the services of an individual network within the internet. Thus network-layer protocols are not responsible for routing of data units between different networks within the internet.

Internet entities may be able to use various different network protocols to transfer data across a given network. One network protocol may provide a connectionless network service in which network-service-data-units are independently transferred between network-service-access-points. A different protocol may provide for the establishment of network-connections, which may be used by the internet-entities.

A specific network may not offer multiple types of service to its attached internet-entities. For example, a network may only provide a connectionless service, or may only provide a connection-oriented service. The fact that not all networks offer a connection-oriented service is one reason why the connectionless internet-service is taken to be the basic internet-service, since connection-oriented services can easily support connectionless services, but not vice versa.

All networks in the internet must provide a connectionless network service to the Internet Layer. It is not required that all networks in the internet provide connection-oriented network service.

The Network Layer provides to the internet-entities independence from routing and switching considerations associated with the transfer of network-service-data-units across a single network. It makes invisible to internet-entities how the Network Layer uses underlying resources such as data-link-connections to provide the network-service. Routing among multiple networks is the responsibility of the Internet Layer.

9.2.2 Services provided to the Internet Layer

9.2.2.1 General

The basic service of the Network Layer is to provide the transparent transfer of all data submitted by the Internet Layer. This service allows the structure and detailed content of submitted data to be determined exclusively by layers above the Network Layer.

All services are provided to the Internet Layer at a known cost.

Three basic types of service provided by protocols within the Network Layer have been identified:

1. Connectionless network service

2. Reliable sequenced network service
3. Real-time network service

These different services may be provided by different protocols.

All networks are required to provide a connectionless network service. Reliable sequenced service and real-time service may require the establishment of network-connections between the network-service-access-points.

In the case of connection service, the quality of service parameters are established for each network-connection. While this quality of service may vary from one network-connection to another it will be agreed for a given network-connection and the same at both network-connection endpoints.

9.2.2.2 Network-Addresses

Internet-entities are known to the Network Layer by means of network-addresses (i.e. network-service-access-point addresses). Network-addresses are provided by the Network Layer and can be used by internet-entities to uniquely identify other internet-entities, i.e. network-addresses are the means by which internet-entities can communicate using the network-service. The Network Layer uniquely identifies each of end systems (represented by internet-entities) by their network-addresses.

9.2.2.3 Network-Service-Data-Unit Transfer

The Network Layer provides for the exchange of network-service-data-units between network-addresses. These units have a distinct beginning and end and the integrity of the unit content is maintained by the Network Layer.

The network-service-data-units are transferred transparently between internet-entities.

9.2.2.4 Network-Connections

If an enhanced level of network-service is desired, it may be necessary to establish a network-connection between the network-service-access-points. Such a service may be provided by a connection-oriented protocol within the Network Layer. For example, a network-connection may be the means of transferring data between internet-entities when reliably sequenced transfers or real-time service is desired. A connection-oriented protocol within the Network Layer provides the means to establish, maintain and terminate network-connections.

More than one network-connection may exist between the same pair of network-addresses. Only specific network-service-access-points offer connection services.

9.2.2.5 Network-Connection-Endpoint-Identifiers

At those network-service-access-points which offer network-connection services, the Network Layer provides to the internet-entity a network-connection-endpoint-identifier which identifies the network-connection-endpoint uniquely with the associated network-address.

9.2.2.6 Quality of Service Parameters

Protocols within the Network Layer may offer a variety of classes of network service to the internet-entities. Internet-entities attached to a network-service-access-point can further negotiate additional quality of service parameters with the associated network-entity, or request to override the default parameters associated with that class of service.

A connectionless protocol within the Network Layer may provide user-selectable quality of service parameters on each network-service-access-unit. These parameters may include:

- level of bit reliability
- level of delivery reliability
- delay
- security level

Such parameters would be used by the Network Layer during the routing process (to determine, for example, which of several possible data links to choose).

Network-service-access-units are delivered to the destination internet-entity intact (though with possible bit corruptions commensurate with the desired level of bit reliability). Any fragmentation of the network-service-access-unit which may have occurred within the Network Layer is hidden from the internet-entities.

If a network provides a connection-oriented service, each network-connection is governed by quality of service parameters negotiated between the internet-entities and the network-entities. The Network Layer protocol establishes and maintains a selected quality of service for the duration of the connection.

The quality of service parameters include:

- a) Residual errors which may arise from alteration, loss, duplication, disordering, misdelivery of network-service-data-units, or others;
- b) service availability which is the probability that a requested network-connection can be established.
- c) reliability, which is the mean time between failures and mean time to repair an established network-connection;

- d) throughput, which is the information transfer capacity;
- e) transit delay, which includes variations on the transit delay;
- f) delay for network-connection establishment.

9.2.2.7 Error Notification

Uncoverable errors detected by the Network Layer protocols may be reported to the internet-entities.

Error notification may or may not lead to the termination of a network-connection, according to the specification of a particular network service.

9.2.2.8 Sequencing

The Network Layer may provide the service of sequenced delivery of network-service-data-units over a given network-connection when requested by the internet-entities.

9.2.2.9 Flow Control

A internet-entity which is receiving at one end of a network-connection can cause the network-service to stop transferring network-service-data-units over the service interface between the Internet and Network Layers. This flow control condition may or may not be propagated to the other end of the network-connection and thus be reflected to the transmitting internet-entity, according to the specification of a particular network service.

9.2.2.10 Congestion Control

The Network Layer provides the service of congestion control, which ensures that total network resources are maintained in a non-saturated state. This network service attempts to avoid serious service degradations caused by extreme quantities of offered traffic. The invocation of congestion control within the Network Layer may result in the network-service stopping the transfer of network-service-data-units over the interface between the Internet and Network Layers.

9.2.2.11 Termination Services

A internet-entity may request termination of a network-connection. The network-service may not guarantee the delivery of data preceding the termination request and still in transit, according to the specification of the specific network-service. The network-connection will be terminated regardless of the action taken by the correspondent internet entity.

9.2.3 Functions within the Network Layer

9.2.3.1 General

The functions which may be provided a Network Layer protocol are outlined. Only the first is a function which is required for connectionless network protocols.

9.2.3.2 Routing and switching

Routing and switching are performed for selecting the appropriate route between network-addresses, and for transferring network-service-data-units between them.

Routing and switching functions may involve intermediate nodes, acting as relays between end network-entities.

In a data network, an intermediate node is a point where one or more network-entities may interconnect data circuits at the Physical Layer and data links at the Data Link Layer.

The service offered to internet-entities over a network-service-access point allows transfer of network-service-data-units to other network-service-access-points. Such transfers may involve switching entities within the network.

9.2.3.3 Network-Connection Establishment

A connection-oriented protocol within the Network Layer includes mechanisms for establishing network-connections between internet-entities. This may include negotiation of quality of service parameters which are to govern the established network connection.

9.2.3.4 Multiplexing

In order to optimize the use of data-link-connections, the Network Layer may multiplex network-connections onto data-link-connections.

9.2.3.5 Error detection

Error detection functions are used to check that the quality of service provided at a network-service-access point is maintained. Error detection in the Network Layer uses error notification from the Data Link Layer. Additional error detection capabilities might be necessary to provide the required quality of service.

9.2.3.6 Error recovery

This function includes mechanisms for recovering detected errors. Depending on the quality of the network service provided, these functions may vary.

30 September 1982

-61-

System Development Corporation
TM-7172/201/01

9.2.3.7 Sequencing

This function includes mechanisms for providing the service of sequenced delivery of network-service-data-units over a given network-connection when requested by internet-entities.

9.2.3.8 Flow control

This function includes mechanisms for providing the service of flow control of network-service-data-units over a given network-connection when requested by internet-entities.

9.3 Current DoD Network Layer Protocols

The large number of different DoD networks, and their wide variation in design and purpose, necessitates a DoD Network Layer capable of supporting multiple intra-network routing schemes. In this section we discuss present and proposed DoD intranetwork protocols as they relate to the DoD Network Layer.

The primary DoD internetting strategy is embodied in the Internet Protocol. IP is a datagram internetting scheme whereby internet datagrams are routed among hosts and gateways using the local subnetwork's internal routing mechanisms. The requirements of IP thus clearly yield the model for the "basic" connectionless service as described in the DoD Network Layer. IP sits in the Internet Layer, and would call on the services of the protocols within the Network Layer for routing within a subnetwork. All protocols within the internet Layer's use the services of the protocols within the Network Layer for intranetwork routing.

Examples of such Network Layer protocols include:

- "Host/IMP" with 1822
- X.25
- various local area network broadcast routing schemes
- connection-establishment control mechanisms within circuit-switched networks

Note that the DoD model's Network Layer does not preclude any combination of Internet protocol and Network protocol for a specific data transfer. For example, an IP module could use an X.25 virtual circuit to exchange datagrams with another IP module. However, some combinations are certainly more useful than others, and a user's quality of service requests to the Network Layer will affect the choice of network mechanisms.

10. THE DATA LINK LAYER

10.1 Differences between the DoD and ISO Data Link Layers

The ISO Data Link Layer allows network-entities to communicate via "data-link-connections", which are built using physical circuits provided by the Physical Layer. The primary purpose of the Data Link Layer is to detect and possibly recover from errors introduced at the Physical Layer. Depending on the requested quality of service, the data-link-connection may additionally perform sequence and flow control functions.

The ISO Data Link Layer does not take into consideration the type of service offered in many contemporary systems built using a broadcast medium (e.g. coaxial cable or radio). In these systems, there is little notion of "connection" at the link level. Instead, "frames" are transferred between data-link-entities in a connectionless (and unreliable) fashion. Error detection is generally performed at the link level, but often there is no attempt at sequence control or error recovery - such mechanisms are left up to the higher layers.

The DoD Data Link Layer differs from the ISO Data Link Layer primarily in the addition of connectionless transport. Quality of service parameters can be requested for such transfers, with the recognition that some quality of service requests can only be satisfied through the establishment of a data-link-connection.

10.2 THE DoD DATA LINK LAYER

10.2.1 Purpose

The Data Link Layer provides functional and procedural means to exchange data-link-service-data-units among user-entities. Such exchanges may involve the establishment of a data-link-connection between the user-entities. The data-link service uses one or several physical-connections.

The objective of this layer is to detect and possibly correct errors which may occur in the Physical Layer. This may be achieved using either connection-oriented or connectionless protocols within the Data Link Layer.

10.2.2 Services Provided

10.2.2.1 Data transfer

The Data Link Layer provides a service which allows a user-entity to transfer data-link-service-data-units to another user-entity. Transfers are governed by selectable quality of service parameters including bit reliability, delivery reliability, sequence reliability, delay and delay variation.

The basic mode of data transfer is connectionless. However, the provision of certain combinations of quality of service parameters may require the establishment of a data-link connection.

If connectionless transfers are provided by a data link protocol, the ability to transfer a single data unit to multiple user-entities may also be supported.

The size of the data-link-service-data-units may be limited by the relationship between the physical-connection error rate and the Data Link Layer error detection capability.

10.2.2.2 Data-link-connection

The Data Link Layer may provide one or more data-link-connections between two user-entities. A data-link-connection is always activated and deactivated dynamically. A data-link-connection is always built over one or several pairs of physical-connection-endpoints.

10.2.2.3 Data-link-connection-endpoint-identifiers

The Data Link Layer provides data-link-connection-endpoint-identifiers that can be used by a user-entity to identify another user-entity; for example, when data-link-connections are built upon multipoint physical connections.

10.2.2.4 Sequencing

When required, the sequence integrity of data-link-service-data-units will be maintained.

10.2.2.5 Error notification

Notification is provided to the user-entity when any unrecoverable error is detected by the Data Link Layer.

10.2.2.6 Flow control

Each user-entity can dynamically control (up to the agreed maximum) the rate at which it receives data-link-service-data-units from a data-link-connection. This control may be reflected in the rate at which the Data Link Layer will accept data-link-service-data-units at the correspondent data-link-connection-endpoint.

10.2.2.7 Quality of service parameters

Quality of service parameters may be optionally selectable. For connection-oriented service, the Data Link Layer establishes and maintains a selected quality of service for the duration of the data-link-connection. The quality of service parameters include:

- a) Mean time between detected but unrecoverable errors
- b) Residual error rate, where errors may arise from: alteration, loss, duplication, disordering, misdelivery of data-link-service-data-units, and other causes.
- c) Service availability
- d) Transit delay
- e) Throughput, which is the information transfer capacity.

Different data-link-service-access-points may provide different classes of data-link-service, where a service class is a predefined range of values for specific quality of service parameters.

10.2.3 Functions within the layer

10.2.3.1 Connectionless Data Transfers

This function provides for the transfer of data units between user-entities without requiring the establishment of a data link connection.

This function may include the capability of transferring a data unit to multiple user-entities via a single transmission.

10.2.3.2 Data-link-connection activation and deactivation

This function activates and deactivates data-link-connections on existing activated physical-connections. When a physical-connection has multiple end-points (e.g. multipoint connection), a specific function is needed within the Data Link Layer to identify the data-link-connections using such a physical-

connection.

10.2.3.3 Data-link-service-data-units mapping

This function maps data-link-service-data-units into data-link-protocol-data-units on a one to one basis.

10.2.3.4 Data-link-connection downward-multiplexing

This function performs downward-multiplexing of one data-link-connection onto several physical-connections.

10.2.3.5 Delimiting and Synchronization

These functions allow the recognition of the sequence of bits transmitted over the physical-connection as a data-link-protocol-data-unit.

10.2.3.6 Sequence control

This function maintains the sequential order of data-link-service-data-units across the data-link-connection, if such service is requested by the user-entities.

10.2.3.7 Error detection

This function detects transmission, format and operational errors occurring either on the physical-connection, or as a result of a malfunction of the correspondent data-link-entity.

10.2.3.8 Error recovery

This function attempts to recover detected transmission, format and operational errors and notifies the user-entities of those which are unrecoverable.

10.2.3.9 Flow control

This function permits provision of the flow control service.

10.2.3.10 Identification and parameter exchange

This function performs data-link-entity identification and parameter exchange.

10.2.3.11 Conveying data circuit control to the Network Layer

The Data Link layer conveys to the network layer the capability to request assembly of data circuits within the Physical Layer (i.e. the capability of performing control of circuit switching).

10.2.3.12 Data Link Layer Management

The Data Link layer protocols deal with some management activities of the layer (such as activation and error control).

10.3 Current DoD Data Link Protocols

The ISO Data Link Layer is designed to incorporate many current Data Link protocols, such as HDLC and ADCCP. These protocols typically have various "modes of use" which are compatible with the ISO data-link-service description.

As an extension of the ISO Data Link Layer, the DoD model will similarly cover many of the existing data link protocols. The link protocols which were used as the model during the development of the ISO architecture are all connection-oriented - i.e. there is a specific link establishment phase, followed by a data transfer phase, followed by a link disconnection phase. Data transfers are often flow controlled, acknowledged and retransmitted, allowing the data link protocol to provide a reliable stream service to its users. The description of these services within the ISO Data Link Layer allow HDLC, ADCCP, Bisync, and other standard link protocols to be used within ISO Open Systems.

However, the proliferation of broadcast-media local area networks (e.g. cablebus systems such as Ethernet) requires that enhancements be made to the ISO Data Link Layer to cover "connectionless" datagram-oriented link protocols. These link protocols (with no flow control or retransmission mechanisms) can be incorporated directly into the DoD model. Many of these connectionless protocols include "broadcast" and "multicast" services as described in the model.

It is anticipated that future DoD standardization efforts will be directed at these local network link-level protocols.

11. REFERENCES

1. Data Processing Open Systems Interconnection - Basic Reference Model, Draft Proposal 7498, ISO/TC97/SC 16 N 719, August 1981.
2. Postel, J. Internet Protocol Specification, Internet Protocol Transition Workbook, Network Information Center, March 1982 (RFC 791, September 1981)
3. Postel, J. Transmission Control Protocol, Internet Protocol Transition Workbook, Network Information Center, March 1982 (RFC 793, September 1981)
4. Postel, J. Telnet Protocol Specification, Internet Protocol Transition Workbook, Network Information Center, March 1982 (RFC 764, June 1980)
5. File Transfer Protocol Specification, Internet Protocol Transition Workbook, Network Information Center, March 1982 (RFC 765, June 1980)
6. User Datagram Protocol Specification, Internet Protocol Transition Workbook, Network Information Center, March 1982 (RFC 768, August 1980)
7. Name Server Protocol Specification, Internet Protocol Transition Workbook, Network Information Center, March 1982 (IEN 116, August 1979)
8. Forsdick, H. and McKenzie, A., AUTODIN FTP, Summary, IEN 101, May 1979.
9. Higginson, P. and Bennett, C., NIFTP: Summary and Assessment, IEN 99, May 1979.
10. Neigus, N., File Transfer Protocol for the ARPA Network, Arpanet Protocol Handbook, DCA 1978
11. D. Boggs et al. Pup: An Internetwork Architecture, Xerox PARC CSL-79-10, July 1979
12. J. Forgie, ST - A Proposed Internet Stream Protocol, IEN 119, September 1979
13. D. Cohen, Specifications for the Network Voice Protocol, ARPANet RFC 741, January 1976
14. M. Padlipsky, A Perspective on the ARPANET Reference Model, August 1982.

DATE
FILMED
— 8