


NAT'L INST. OF STAND & TECH R.I.C.



A11104 907204

NIST
PUBLICATIONS

NISTIR 5793

Data Communications Strategy

**Jerry Mulvenna
Tim Boland**

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

QC
100
.U56
NO. 5793
1993

NIST

Data Communications Strategy

**Jerry Mulvenna
Tim Boland**

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

January 1996



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

TECHNOLOGY ADMINISTRATION
Mary L. Good, Under Secretary for Technology

NATIONAL INSTITUTE OF STANDARDS
AND TECHNOLOGY
Arati Prabhakar, Director

DATA COMMUNICATIONS STRATEGY

Jerry Mulvenna
Tim Boland

January 7, 1996

This document was written under funding from
the Internal Revenue Service

CONTENTS

| | |
|--|----|
| 1. Introduction..... | 1 |
| 1,1 Background..... | 1 |
| 1.2 Purpose..... | 1 |
| 1.3 Scope..... | 1 |
| 1.4 Organization of the Document..... | 2 |
| 1.5 Acknowledgements..... | 2 |
| 2. Computer Networking Services and Architecture..... | 3 |
| 2.1 Internet Protocol Suite (IPS)..... | 3 |
| 2.1.1 General Process..... | 3 |
| 2.1.1.1 History..... | 3 |
| 2.1.1.2 Specific Process..... | 4 |
| 2.1.1.3 Internet Perspective..... | 5 |
| 2.1.2 Architecture..... | 5 |
| 2.1.3 Services..... | 6 |
| 2.1.3.1 Application Layer Services..... | 6 |
| 2.1.3.1.1 Remote Login..... | 6 |
| 2.1.3.1.2 Mail..... | 7 |
| 2.1.3.1.3 File Access and Transfer..... | 8 |
| 2.1.3.1.4 Directory Service..... | 9 |
| 2.1.3.1.5 Information Retrieval..... | 10 |
| 2.1.3.2 Transport/Network Layer Services..... | 11 |
| 2.1.3.3 Network Management..... | 11 |
| 2.1.3.4 Security..... | 13 |
| 2.1.3.5 Application Program Interfaces..... | 13 |
| 2.1.3.6 Conformance/Interoperability Testing..... | 13 |
| 2.2 Open Systems Interconnection (OSI) Protocol Suite..... | 13 |
| 2.2.1 OSI Background..... | 13 |
| 2.2.2 Architecture..... | 15 |
| 2.2.2.1 General Description..... | 15 |
| 2.2.2.2 OSI-IPS Architectural Comparison..... | 16 |
| 2.2.3 Services..... | 17 |
| 2.2.3.1 Application Layer Services..... | 17 |
| 2.2.3.1.1 Mail..... | 17 |
| 2.2.3.1.2 File Access..... | 19 |
| 2.2.3.1.3 Terminal Emulation..... | 20 |
| 2.2.3.1.4 Transaction Processing..... | 20 |
| 2.2.3.1.5 Remote Database Access..... | 21 |
| 2.2.3.1.6 Directory Service..... | 22 |
| 2.2.3.1.7 Additional OSI Applications..... | 23 |
| 2.2.3.2 Network Management..... | 24 |
| 2.2.3.3 Security..... | 25 |
| 2.2.3.4 Application Program Interfaces..... | 26 |
| 2.3 Coexistence and Convergence..... | 27 |
| 2.4 Consortia..... | 28 |
| 3. Strategy..... | 30 |
| 3.1 General..... | 30 |
| 3.2 Testing..... | 33 |
| 3.3 Security..... | 34 |
| 3.4 Network Management..... | 36 |

| | | |
|----------|---|----|
| 3.5 | Electronic Mail and Directory Services..... | 36 |
| 3.6 | Electronic Data Interchange..... | 40 |
| 3.7 | Message Attachments..... | 42 |
| 3.8 | File Access..... | 43 |
| 3.9 | Information Retrieval (IR)..... | 43 |
| 3.10 | Remote Database Access (RDA)..... | 43 |
| 3.11 | Transaction Processing..... | 44 |
| 3.12 | Network Infrastructure..... | 44 |
| 3.13 | APIs..... | 45 |
| 3.13.1 | Electronic Mail API..... | 45 |
| 3.13.1.1 | Client APIs..... | 45 |
| 3.13.1.2 | Server Provider Interface..... | 45 |
| 3.13.1.3 | Specific Implementation..... | 45 |
| 3.13.1.4 | Mail API Summary..... | 46 |
| 3.13.2 | Transport API..... | 46 |
| 3.13.3 | Other APIs..... | 46 |
| 3.14 | Collaborative Computing..... | 46 |
| 3.15 | Gateways..... | 47 |
| 3.16 | Consortia..... | 48 |
| 4. | Summary..... | 49 |
| | Appendix A: Testing..... | 50 |
| | Appendix B: List of Acronyms..... | 51 |
| | References..... | 53 |

Disclaimer

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

1. Introduction

1.1 Background

Federal government agencies are no longer required to acquire Open Systems Interconnection (OSI) products in support of their networking requirements, because it was recognized that there was no single solution for meeting these requirements. Technology is advancing at a rapid rate, and cost is an important factor in any procurement decision. Users must balance the services provided with the expected cost and be prepared to make tradeoffs. Today, there are many disparate information systems, providing a variety of network services. Interoperability and interworkability among these systems is problematic at best. Accomplishment of user missions is limited by capabilities of existing equipment. To a large extent, users are still locked into proprietary solutions, and this situation is not entirely unpleasing to their vendors. Administratively, organizations are continuing to decentralize, so the need for networking services is increasing.

1.2 Purpose

The absence of a single solution for providing global data communications interoperability gives users more flexibility in making procurement decisions. It also gives them more responsibility for making the right choice. In order to make more informed decisions, users need to be more knowledgeable about the current state of computer networking technology and the planned advances in that technology. Users must also be aware of the vendors' plans for marketing that technology. Only by integrating all of this information with their own requirements will users be able to make intelligent procurement decisions.

This document will survey the services that are currently provided by non-proprietary communications technology and that are expected to be provided in the foreseeable future. When the services provided by the different protocols are similar, the differences will be highlighted to assist users in making the right choice. There will also be an assessment of the degree to which the major protocols or protocol suites have gained or are expected to gain marketplace acceptance. It will be impossible to accurately predict all of the circumstances that can influence a future procurement so whenever possible, the contingency factors that will affect future events will be specified. In summary, it is the intent of this document to arm the reader with the best information available to assist in the procurement of non-proprietary communications products and services.

1.3 Scope

Although the mandate to acquire OSI technology to provide interoperability between heterogeneous computer systems has been removed, the National Institute of Standards and Technology (NIST) continues to believe that solutions to computer networking

requirements that do not lock users into the services of a single vendor are in the long-range interests of most users. Accordingly, this report will focus on providing non-proprietary solutions to those requirements.

1.4 Organization of the Document

This section (section 1) provides an introduction to the rest of the document. Section 2 contains information about the most widely implemented non-proprietary data communications protocols and the services that they provide. Section 3 contains alternative strategies for users to consider, given the information in section 2.

1.5 Acknowledgements

The National Institute of Standards and Technology (NIST) wishes to acknowledge the assistance of Daniel Blum (Rapport Communications), Jack Finley (General Services Administration), and Phill Gross (MCI), who provided information used in the preparation of this report.

2.0 Computer Networking Services and Architectures

2.1 Internet Protocol Suite (IPS)

2.1.1 General Process

2.1.1.1 History

In the early 1970's, the Defense Advanced Research Projects Agency (DARPA) of the Department of Defense (DOD) instituted research in order to develop a mechanism for preserving critical information in case of foreign attack. The research was considered experimental, and involved the ability to transfer information easily and quickly from one computer to another. Initially several machines were connected at several academic communities (Massachusetts Institute of Technology (MIT), Stanford, Utah, University of California at Berkeley); thus the Internet was born. In 1977, early versions of the Transmission Control Protocol (TCP) and Internet Protocol (IP) began to be developed on the experimental Internet. TCP and IP collaborated to route and reliably transfer data packets between the source and destination end systems. Before TCP and IP were developed, ad hoc solutions were used. A more mature and bigger Internet was created in 1980 when DARPA began to connect computers on many of its research networks using the new TCP/IP. In 1981 the Internet effort had grown to the point that the DARPA Program Manager formed an advisory group, called the Internet Advisory Control Board (IACB) to advise DARPA in managing the operation of the Internet. Initially consisting of eight members, it is essentially the same management structure that is in place today. DARPA funded the inclusion of TCP/IP into UNIX 4.1 Berkeley Software Distribution (BSD) in 1981 and the distribution of source code by Sun Microsystems in 1982. Informal testing groups resolved any operational problems. The transition to Internet technology accelerated in January of 1983 when DARPA mandated that all computers connected to its research networks use TCP/IP to provide a networking infrastructure.

Another key factor in the rapid expansion of the Internet was the involvement of the National Science Foundation (NSF), which funded regional networks in 1985, and funded a national backbone to link those regional networks in 1987, using TCP/IP. The National Science Foundation Network (NSFNET) in 1985 used TCP/IP over T-1 links; this represented an important Government commitment. By 1988 hardware/software platforms were available which incorporated TCP/IP.

In 1986, the Internet Architecture Board (IAB) created the Internet Engineering Task Force (IETF) and the Internet Research Task Force (IRTF). The former was chartered to provide near-term solutions to Internet operational problems and to develop near-term enhancements for the Internet. The latter group was asked to pursue topics of long-term interest that carry some technical risk. The IETF benefited from regional network and NSF interest in their activities.

Attendance at the first IETF meetings numbered a few individuals. Today attendance is over 700. Currently the Internet is estimated to include about 40,000 networks and 20 million users, and the number of networks and users connected is growing every year. In sum, several factors were key to the explosive growth of the Internet Protocol Suite: (1) DARPA funded the academic community to carry out much of the original research. This developed expertise and stimulated interest in these protocols among the creative members of this community, which, in turn, stimulated additional research and development work not funded by DARPA, and led to expanded usage of the protocols in the academic world., (2) DARPA funded the inclusion of the Internet Protocol Suite in UNIX 4.1 BSD, (3) Sun Microsystems included the Internet Protocol Suite source code in its Sun Operating System, and (4) NSFNET came into existence.

2.1.1.2 Specific Process

The process by which Internet Protocol Suite (IPS) standards are developed has evolved over time as the scope and importance of the IPS, the Internet and related commercial markets have grown. In 1992 the IETF process for approving standards became more formal; IAB members must now be nominated and approved according to established procedures. IPS standards development is primarily the responsibility of the IETF. The IETF consists of numerous working groups (WGs) that are formed to address specific, narrowly focused development or operational issues relating to the Internet.

The primary goal of most IETF working groups is the technical development of IPS standards. Individuals interested in the content and progression of IPS standards attend the three or four IETF meetings that are held annually and participate in the deliberations of the working groups. All documents being developed by IETF working groups are required to be available on-line.

While the IETF working groups develop the technical specifications, the progression of these documents through the standards process is controlled by an oversight committee called the Internet Engineering Steering Group (IESG). For the purposes of such oversight, IETF working groups are organized into specific subject areas. Each area has one, or more, director(s) who together with the IETF chairperson comprise the IESG. The IESG makes all decisions regarding the progression of specifications through the IPS standards process.

The IETF forms one half of the Internet Society's (ISOC) standards and research infrastructure. The ISOC is a professional society that is concerned with the growth and evolution of the worldwide Internet, with the way the Internet can be used, and with the social, political and technical issues that arise as a result. Previously IETF attendance was mostly by the research and academic community, but now attendees with commercial interests predominate. Non-US participation in IETF meetings is growing as a percentage of total attendees, but non-US participants are still a minority of

all participants. The Internet Research Task Force (IRTF) forms the other half of the infrastructure and is tasked with long-term research issues related to the evolution of the Internet. The IAB, which reports to the ISOC board of trustees, is tasked with the broad oversight and coordination of all IETF and IRTF activities.

The primary means of Internet standardization is through requests for comments (RFCs). RFCs began as an informal means of documenting technical information for the original ARPANET. Today, the RFC mechanism is part of a formal Internet standards process.

2.1.1.3 Internet Perspective

There are still several issues and problems regarding the Internet. There is no central Internet "authority". Although agreements exist between service providers, it is still difficult for the user to determine whom to contact when something goes wrong. The same informality that has allowed the explosive growth of the Internet has made it more difficult to determine responsibility for its operation.

The Internet has developed in a non-commercial environment into the leading major international non-regulated network of value added services. Furthermore, the Internet has promoted the implementation of de facto standards intended for international recognition. An agreement was concluded in June 1994 between the International Organization for Standardization (ISO) (see sec. 2.2.1) and the IETF that will lead to some mutual recognition of Internet and ISO protocols and the organization of joint activities.

The Internet is evolving towards providing commercial services for the general public. Participation by the ARPA, Defense Communications Agency (DCA), NSF, National Aeronautics and Space Administration (NASA) and other government agencies has been an important factor in the development of the Internet, but the involvement of these government agencies is declining. The original ARPANET plan did not include application services such as mail, lists, or news. People who wanted these services developed them. For example, the User's Network (USENET) has developed into a worldwide distributed conferencing system.

2.1.2 Architecture

The IPS architecture consists of four layers: (1) the Network Interface (hardware) Layer, (2) the Internet Layer, (3) the Transport Layer, and (4) the Application Layer. Each layer provides service to the layer above. The function of the Network Interface Layer is the transfer of data frames between directly connected machines, to support the internetworking capability of the Internet layer. Sequencing and error detection may be provided. These functions can be provided by various subnetwork technologies. These subnetwork technologies are independent of the IPS architecture (as well as other architectures, for example,

OSI). The IPS must be mapped onto specific subnetwork services. Examples of subnetwork technologies are: local area networks (LANs), point-to-point links, X.25, switched multimegabit data service (SMDS), frame relay, integrated services digital network (ISDN), and asynchronous transfer mode (ATM)/synchronous optical network transmission (SONET).

The function of the Internet Layer is to provide internetworking among machines not directly connected. The internetworking protocol for the Internet Layer is the Internet Protocol (IP); the function of this protocol is to route data between source and destination end systems on the same network or different networks, through a series of intermediate systems if necessary. Routing protocols may be intra-domain or inter-domain; domains may be organized along administrative or technical lines. A network address is used to identify the destination system.

The function of the Transport Layer is to provide end-to-end transfer of IP datagrams between source and destination machines. The protocols that provide this capability are the Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP). TCP is a connection-oriented reliable transport service, and UDP is an unreliable connectionless transport service. TCP offers a stream service to applications. Flow control, congestion control, error detection, framing, and demultiplexing services are provided by the Transport Layer.

The IPS applications are remote login (TELNET), file transfer protocol (FTP), and the Simple Mail Transfer Protocol (SMTP). These applications present a data stream to the transport service, and provide a customized interface to the applications user. All functions at this layer are intrinsic to particular applications. Services offered to the user are specific to individual protocols.

2.1.3 Services

This section describes the current services provided by the Internet Protocol Suite (IPS) and anticipates those services which are likely to be provided in the foreseeable future. To the extent possible, there will be an assessment of marketplace acceptance of implementations providing these services.

2.1.3.1 Application Layer Services

2.1.3.1.1 Remote Login

The IPS virtual terminal protocol, TELNET, was designed with scroll mode terminals in mind, although some page mode terminal properties can be negotiated. Scroll mode terminals do not have any local editing capability. The TELNET protocol is concerned with setting up and manipulating two simplex data streams, one in each direction.

TELNET allows a user to connect and run applications on a remote

computer. To initiate a TELNET session, a user first identifies the remote host with an Internet address. The remote system then prompts for a username and a password to ensure access only by authorized users. Once connected, a user can run applications on the remote host just as if connected to a local host. Once the TELNET session begins, applications run on the processor of the remote computer, not on the local system that initiated the session.

TELNET offers three basic services. First, it defines a network virtual terminal that provides an interface to remote systems around which client programs are built. Second, it includes a mechanism that allows the client and server to negotiate options and it provides a set of standard options (e.g., an option controls whether data passed across the connection is binary or ASCII text). Finally, TELNET treats both ends of the connection symmetrically. So, instead of forcing one end to connect to a user's terminal, TELNET allows either end of the connection to be a program. No additional TELNET services are currently planned, except, possibly, TELNET encryption.

Most vendors who market the UNIX operating system include the Internet TELNET as part of the operating system. For this reason, the Internet TELNET application is widely available. Government agencies and universities rely on the Internet and TELNET for remote login. Most database services and electronic bulletin boards use IPS TELNET as one of the primary access methods.

2.1.3.1.2 Mail

The IPS electronic mail application is known as the Simple Mail Transfer Protocol (SMTP). SMTP is implemented in accordance with RFC-822 and RFC-821. RFC-822 defines a commonly implemented message format used in Internet mail, USENET, BITNET, and other messaging systems. RFC-821 defines the protocol necessary to convey that format between systems.

Communication between a client and server consists of readable text. Although SMTP rigidly defines the command format, a transcript of interactions between a client and server is easily read by humans. Once communication has been established, the originator's system can transmit one or more mail messages, terminate the connection, or request the server to turn roles of sender and receiver around so that messages can flow in the opposite direction. The recipient's system must acknowledge each message. It can also abort the entire connection or abort the current message transfer.

The SMTP protocol focuses specifically on how the underlying mail delivery system passes messages across a link from one machine to another. It does not specify how the mail system accepts mail from a user or how the user interface presents the user with incoming mail. Also, SMTP does not specify how mail is stored or how frequently the mail system attempts to send messages. There is an

extended SMTP (RFCs 1425-28) (not widely implemented), that allows a sending Message Transfer Agent (MTA) to determine the capabilities of a prospective receiver. Once it is clear that the receiver is capable, the sender can use additional extensions to signal such parameters as the length and type of message to be sent.

RFC-822 was initially written to convey ASCII text. The Multi-media Internet Mail Extensions (MIME) standard was developed to expand the message format to include non-ASCII text and multi-media attachments. MIME provides three capabilities: (1) it standardizes several encoding techniques that can be used when needed to render binary data into a character form, (2) it provides a tagging scheme to label each attachment as to its information type, and (3) it provides for marker strings that allow MIME-enabled User Agents (UAs) to determine the beginning and end of each attachment.

Most vendors marketing the UNIX operating system include SMTP as part of the operating system. For this reason, SMTP is widely available from most major vendors and many small software houses. Public domain versions of MIME are also available. Basic Internet messaging has been implemented and deployed more widely than any other messaging technology, standardized or proprietary.

PEM (Privacy Enhanced Mail) does for the security of Internet mail what MIME does for its media capabilities. Like MIME, PEM builds on the RFC 822 message format to allow users to encrypt and digitally sign messages so as to guard their privacy, authenticate their origin, and prevent their modification. PEM has not been that successful; however, because (among other reasons) it is not geared for a multilevel security environment.

2.1.3.1.3 File Access and Transfer

The IPS File Transfer Protocol (FTP) is a simple user-oriented protocol designed to transfer files between remote hosts. FTP uses the IPS TELNET protocol to set up a connection between the FTP user and server. The services provided by FTP are login (USER,PASS,ACCT), close (QUIT), transfer of entire files (GET,PUT), and a file directory capability. A limited number of file types (e.g., ascii, binary, image) can be transferred.

FTP allows authorized users to log into a remote system, identify themselves, list remote directories, copy files to or from the remote machine, and execute a few simple commands remotely (e.g., to obtain help with the remote machine file name syntax). FTP understands a few basic file formats and can convert among popular representations (e.g., between EBCDIC and ASCII character sets). Many FTP implementations provide statistics on transfer rates and diagnostic aids such as packet tracing. FTP allows a user to access multiple systems in a single "session". Commercial implementations of FTP are widely available for most common computing environments. Public-domain implementations of FTP are readily available for MS-DOS and UNIX environments.

The Network File System (NFS), developed by Sun Microsystems, offers the ability to share file systems among computers on a TCP/IP network. FTP transfers individual files from one computer to another. NFS takes this a step further, making all or part of the file system of a remote computer seem like it is resident on a local computer.

An NFS server makes part of its file structure available to be shared by other computers on the network. The NFS server must include security features to ensure that only authorized computers gain access to its file systems. NFS uses the User Datagram Protocol as the underlying transport mechanism. An NFS client can mount the file system of a remote host on a TCP/IP network.

Microcomputer-based TCP/IP products often do not include NFS, or offer it as an added-cost option. Some may offer an NFS client, but not include NFS server capabilities.

2.1.3.1.4 Directory Service

In the traditional IPS environment, directory services are usually provided by one of three basic alternatives: (a) the WHOIS service currently provided under the InterNIC Directory and Database Services contract; (b) the "Host Table" service currently provided under the InterNIC contract; and (c) the "Domain Name System" (DNS).

The WHOIS database system accepts simple query requests and performs a broad or constrained search of the WHOIS databases for the requested information. Because WHOIS is a centralized database service, it does not scale easily to serve the needs of a global user community. Problems dealing with the growth in size and use of the database have already led to breaking the database into several physical pieces that are logically integrated using ad hoc methods.

The Host Table service, like WHOIS, is a centralized query/response system hosted on an InterNIC server. In this context, access is by means of a local application program which establishes a TCP connection to port 101 at the service host. Using the Hostname Server Protocol, the NIC Host Table service accepts simple query requests, performs a search of the Host Table for related entries, and then returns the related records. The problems associated with Host Table service mirror those of WHOIS; the service is not scalable to the needs of a global network infrastructure.

Unlike Host Table and WHOIS, the DNS is a distributed database system. The contents of the database can be viewed from a network global perspective as being allocated among many servers. A typical DNS Resolver offers the following three functions: (1) translation of a host name to the corresponding host address, (2) translation of a host address to the corresponding host name, and (3) general lookup of host information.

The DNS translates an IP address in its registered canonical form to the dotted decimal notation; for example, "boland@snad.ncsl.nist.gov" may be translated to "129.6.55.1". People remember and manipulate names much better than numbers, but low-level TCP/IP programs deal only with numerical addresses. It is common practice to assign each computer on a network a name that consists of a unique machine name prepended to the domains associated with the network. TCP/IP networks almost always include a "nameserver" computer that performs DNS. It constantly listens for requests for its services.

The DNS is a highly specialized database system. Because it is specialized to handle a small number of record types and has limited search capability, it is able to provide high performance. However, specialization also limits its applicability. (In general, DNS functionality is relatively difficult to extend when compared to the flexibility inherent in the OSI Directory architecture.) A digital signature service is being planned for DNS in the future.

Products needed to access and use the WHOIS and Host Table services are widely available. Likewise, DNS server and resolver products are widely available; a public domain implementation is also available. The WHOIS service is widely used by the IPS community. The Host Table service, despite its problems, continues to be used by legacy applications that have not been converted to use DNS. The DNS is widely deployed throughout the IPS community and is used extensively by some implementations of Internet electronic mail. The DNS distributed database contains information to locate the appropriate server that could find the IP addresses for any system connected to it.

The OSI Directory Service is being seriously considered as a solution due to limitations in the current alternatives. For example, the combined limitations of the traditional directory services for Internet Protocol Suite users have led to extensive pilot programs throughout the world that are experimenting with the use of the OSI Directory Service to provide a more flexible and scalable directory.

2.1.3.1.5 Information Retrieval

The IPS supports user-defined protocols to access information available via the Internet; there are no direct OSI counterparts to these protocols. The Internet has an enormous volume of information available (such as anonymous FTP archives and databases). There are a number of resource discovery tools being designed and tested on the Internet, including Archie, Gopher, Wide Area Information Service (WAIS), and Mosaic. The placement of source code at low cost into creative hands has led to the development of widely-used functionality totally independent of the Internet standards process.

The primary purpose of Archie is to provide a convenient way for

users to browse the anonymous FTP archives available throughout the Internet. The shortcomings of Archie include replication of information and nonoptimal query keys. The original implementation was done at McGill University in Montreal, Canada.

Gopher was initially intended as a campus-wide information system for the University of Minnesota. Thus, the emphasis is on providing a convenient structure for information of the scope and size needed by a campus community. Since not all of the information needed by a user community is typically stored on on-campus systems, the design of Gopher gives the user convenient access to the information without the user having to know precisely where it resides. The potential drawback to Gopher is a lack of scalability. To ameliorate this problem, a specialized index-server has been developed for Gopher called Veronica. Veronica collects lists of Gopher items from Gopher servers and allows users to perform keyword searches on the lists.

The WAIS is a full-text search and retrieval system. It allows users to search and retrieve information from indexed databases. The information can include text, formatted documents, pictures, spreadsheets, graphics, sound, and video.

The World Wide Web (WWW) is a collection of information sources, interlinked using a hypertext markup language (HTML). While Gopher tends to be text-based and hierarchical, the Web relies more on hypertext and multimedia presentations. Web servers implement a protocol called Hypertext Transfer Protocol (HTTP) and use documents stored in HTML format. One can traverse the WWW by following the links established between documents. Mosaic and Netscape are user interfaces for providing such traversal.

2.1.3.2 Transport/Network Layer Services

The IPS transport service is provided by the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP). The service provided is transmission of IP packets between end systems on the same network or different subnetworks. TCP provides a reliable end-to-end packet delivery mechanism, and UDP provides an unreliable packet delivery mechanism. The IPS network layer service is the Internet Protocol (IP).

2.1.3.3 Network Management

The network management application services in the IPS are provided by the Simple Network Management Protocol (SNMP); SNMP comprises three RFCs (RFC 1155 - Structure and Identification of Management Information for TCP/IP Networks, RFC 1157 - Simple Network Management Protocol, and RFC 1213 - Management Information Base for Network Management of TCP/IP-based Internets).

The IPS Structure of Management Information (SMI) standard provides a methodology for abstractly defining managed objects, which in this model are variables that represent atomic data elements.

These managed objects are collected into a database, called the Management Information Base (MIB), which defines the set of variables that SNMP servers maintain as well as the semantics of each variable. MIB variables record the status of each connected network, traffic statistics, counts of errors encountered, and the current contents of internal data structures such as the machine's IP routing table. Each managed object contains a name, a type syntax, access information, and status. The Internet SMI allows MIB variables to be collected into lists and tables.

SNMP helps network managers locate and correct problems in a TCP/IP Internet. Managers invoke an SNMP client on their local computer and use the client to contact one or more remote SNMP servers. SNMP uses a fetch-store paradigm in which each server maintains a set of conceptual variables that include simple statistics, such as a count of packets received. SNMP messages either specify that the server should fetch values from variables or store values in variables, and the server translates the requests to equivalent operations on local data structures. SNMP defines both the syntax and meaning of the messages that clients and servers exchange.

An example of SNMP operation might involve storing the time a system has been operational. Many systems simply record the time at which the system started, and compute the time that the system has been operating by subtracting the startup time from the current time. Thus, SNMP software can simulate a MIB "variable" that contains the time since last startup. It performs the computation whenever a request arrives to read a value from the MIB variable, and the remote site remains unaware of the computation.

The theory of operation of IPS network management is less complex than that of OSI network management (see sec. 2.2.3.2). For example, data relating to managed objects must be requested by the manager - data is not automatically provided. IPS network management uses polling from the manager to the agent to gather information, whereas OSI network management uses a notification mechanism from the agent to the manager. However, by establishing an appropriate time interval between information transfer, IPS network management can simulate OSI network management.

Currently, SNMP Version 1 (v1) seems to be the protocol of choice for managing the resources associated with TCP/IP networks. SNMPv1 has been widely implemented by many vendors; the implementation cost has been relatively low due to the simple design of the protocol and the MIB. However, SNMPv1 may not provide sufficient security, does not support bulk data transfer, is not robust, and does not provide sufficient scalability (limited number of agent devices per management system). Recently, a new version of SNMP, SNMPv2, is making its way through the Internet standards process. This new version promises enhanced management capabilities, including Secure SNMP and support of bulk data transfer. SNMP v2 is not likely to be widely implemented or deployed in the next several years because of its cost and transition issues from SNMPv1. SNMPv2 may provide more security than is needed by most

users.

2.1.3.4 Security

Network security services for the IPS include authentication, access control, encryption, and data integrity. Security services are provided at the Internet and Application Layers. A protocol is being defined to provide security services for IP; this protocol is called IPSEC (IP Security Protocol); functionally it is the same as the OSI Network Layer Security Protocol (see sec. 2.2.3.3) but employs different mechanisms. IPSEC provides all of the security services defined above; it is to be used if a common level of security is desired among all applications. Also, it only provides security between adjacent systems; it does not provide true end-to-end (or application) security.

2.1.3.5 Application Program Interfaces

There are no application program interfaces that are currently widely used by the IPS.

2.1.3.6 Conformance/Interoperability Testing

Testing on the Internet is ad-hoc and informal. There are no specific codified interoperability criteria, so the level of interworking is not documented anywhere. Sometimes a "TCP/IP bakeoff" (interoperability testing tournament) occurs, where implementations are tested against each other, using informal accumulations of tests. The usual result of such a bakeoff is a media event, announcing the interoperable participants. There is no publicly available formal record of the level of testing conducted.

2.2 Open Systems Interconnection (OSI) Protocol Suite

2.2.1 OSI Background

The objective of the early developers of the Open Systems Interconnection (OSI) standards was to provide an internationally recognized method for providing interoperability among heterogeneous computer systems. When the decision was made to develop these standards in the late 1970's, the first step was to develop an OSI Reference Model. The Reference Model partitioned the communication functions into seven layers. Each layer provided a service to the layer above while shielding the upper layer from the details of how the service was actually implemented. The Reference Model allowed the concurrent development of standards at the different layers; for example, standards for Message Handling Systems at the Application Layer (layer 7) could be and were developed concurrently with the Transport and Network Layer standards at layers 4 and 3 respectively.

OSI standards were developed in Joint Technical Committee 1 (JTC1) of the International Organization for Standardization (ISO)/

International Electrotechnical Commission (IEC), as well as by the International Telecommunications Union - Telecommunications Standardization Sector (ITU-T) (formerly CCITT). The American National Standards Institute (ANSI) is the ISO/IEC member representing the United States. The outputs of the OSI ISO/IEC standardization process are classified as either ISO/IEC international standards or technical reports.

The ITU-T is an international organization whose aim is the production of standards for international interworking of telephone and other telecommunications systems. The ITU-T has adopted a four-year work cycle. ITU-T Recommendations are voted on and, if approved, published at the end of this four-year cycle. The State Department is the US member to ITU-T. There is currently an Internet Society (ISOC)-ITU reciprocity agreement.

The ITU-T differs from the ISO/IEC in that the former has full members that are usually the telecommunication carriers for public data networks of each nation, whereas ISO/IEC membership is open only to national standards organizations. The ISO/IEC may cooperate with the ITU-T in the development of joint standards (such as Directory Services).

The need to achieve consensus has prevented the international standards organizations from restricting options (e.g., there are five allowable classes of Transport). Low-level decisions, such as allowable values of specific parameters (e.g., maximum length of message) are not made. Consequently, in addition to base standard specifications, the OSI standards process contains an additional level of standardization known as functional standards. Functional standards, also known as profiles, contain selections of specific OSI standards along with additional refinements (e.g., selection of protocol options, specification of parameter values) appropriate for a particular user community. Open, public regional workshops around the world develop functional standards, and harmonize these standards where differences occur.

The international OSI standards process is slow; it may take five years or more from the time a proposed work item is submitted to the time it is standardized. There is a seemingly endless round of ballots and ballot resolution meetings. Regular standards meetings are held relatively infrequently and the pace of change is painfully slow. The ballot procedures themselves are unnecessarily complex and slow progress. The additional time required by regional workshops to develop functional profiles also delays the implementation of products based on these standards.

The need to gain concurrence caused OSI standards to be too complex. Workshops had to determine what subset of services to implement, and there were no concurrent implementations to test workability of standards. The complexity of OSI standards increased the cost of the standards to develop.

Information relating to the OSI standards in progress was not

easily obtainable. On-line distribution was not available, and only a few centers were authorized to sell paper copies of the standards. The documents were not generally available in electronic form.

The ISO/IEC JTC1 has a scope of standards development similar to that of the IETF. IETF working groups are similar to specific technical projects within ISO/IEC working groups; however, in comparing the review and approval process, there are several differences. The ISO/IEC/ANSI process employs a consistent mechanism of explicit voting by member organizations. The IETF process employs oversight committees to make all final standards decisions.

Other differences are cultural. In order to have a "reality check" on the standards being developed, the IETF develops implementations of the standards. The lessons learned in implementing the standards causes vendors to make modifications at an early date. The failure of the OSI community to develop implementations in parallel with standards has led to incompletely specified standards despite their complexity. The IETF interests are overwhelmingly dominated by North American concerns. A focused market means that in the IETF, because participation is predominately North American, fewer interests have to be satisfied in producing the standards. The OSI community is more diverse than that of the IETF. The lack of common interests means that there may be less cooperation in the OSI world.

Both OSI and IETF participants are recognizing both the shortcomings of their respective processes and the benefits of collaboration, when appropriate. For example, the IETF is seeking to get more international participation in its activities. The IETF is also seeking to formalize its process to a greater extent, and is working with the OSI community to determine any basis for future collaboration. OSI organizations are now making their documentation more publicly available, and are also seeking a basis for future collaboration.

2.2.2 Architecture

2.2.2.1 General Description

The OSI layers are as follows: (1) Physical Layer, (2) Data Link Layer, (3) Network layer, (4) Transport Layer, (5) Session Layer, (6) Presentation Layer, and (7) Application Layer. Each layer provides services to the layer immediately above. Although the generic service descriptions are specified, implementation considerations are not.

The Physical Layer is concerned with transmitting raw bits over a communication channel; that is, when one system sends a 1 bit, it is received by the adjacent system as a 1 bit, not a 0 bit. The main task of the Data Link Layer is to take a raw transmission facility and transform it into a line that appears free of

transmission errors to the Network Layer. It accomplishes this task by having the sender break the input data up into data frames, transmit the frames sequentially, and process the acknowledgement frames sent back by the receiver.

Layer (3) is the Network Layer, which routes data between end systems on the same or different networks. OSI includes both a connection-oriented network service (CONS) and connectionless network protocol (CLNP). "Connection-oriented" involves setting up a predefined link before user data is sent; a link is in effect for the life of the connection (analogy: telephone conversation). Connectionless packets may travel different paths to the destination (analogy: postal letter).

The Transport Layer provides for the end-to-end transfer of data packets between end systems. There is both a connection-oriented service and a connectionless service. There are five classes of connection-oriented transport service that assume varying levels of Network Layer services.

The Session Layer (layer 5) allows cooperating application entities to organize and synchronize conversation and to manage data exchange. Session services include dialogue control and token management.

The Presentation Layer (layer 6) specifies or, optionally, negotiates the way information is represented for exchange by application entities. The Presentation Layer is concerned only with the syntax of the transferred data.

The Application Layer (layer 7) allows for protocols and services required by particular user-designed application processes. Functions satisfying particular user requirements and application service elements that can be used by more than one application are contained in this layer.

2.2.2.2 OSI-IPS Architectural Comparison

Although the OSI and IPS Transport and Network layers map one for one, the functions of OSI layers 5, 6 and 7 are bundled into the IPS Application Layer. At the Transport Layer the IPS has a single protocol for providing all transport services; this protocol corresponds closely to OSI Transport Class 4 (TP4). At the network layer the IPS provides a single internetworking protocol (IP); this protocol provides nearly identical functionality to OSI CLNP, although different routing protocols and address structures are used in support of each. These similarities mean that OSI applications can, with minor modification, use the TCP/IP infrastructure, and IPS applications can, with minor modification, use an OSI infrastructure.

IPS and OSI use different forms of network addresses; these address forms have implications in the behavior of OSI and IPS routing protocols. The traditional IPS address is composed of four

numerics, separated by delimiters (dots); the first portion of the address (one to three numerics) is a network identifier, and the second portion (one to three numerics) is a host identifier. The number of numerics to allocate to the network identifier portion and number to allocate to the host identifier portion depends upon the IP address's class, but the total of the two numbers must add up to four. There is a shortage of Internet address space because of the inefficient assignment of IP addresses in the past coupled with the increased demand for those addresses due to the many new systems being attached to the Internet. A new version of IP, IPv6, that is currently under development, will provide a larger address space, and routing aggregation techniques (such as Classless Interdomain Routing (CIDR)), as well as more efficient assignment of addresses. OSI address space depletion has never occurred because there is more address space, the method of address assignment makes more efficient use of this space, and fewer OSI addresses have been assigned.

2.2.3 Services

This section describes the current services provided by the OSI protocol suite and anticipates those services which are likely to be provided in the foreseeable future. To the extent possible, there will be an assessment of marketplace acceptance of implementations providing current services.

2.2.3.1 Application Layer Services

2.2.3.1.1 Mail

The OSI electronic mail service is specified in the ITU-T X.400 Series of Recommendations for Message Handling Systems [REF 1]. The original recommendations were first issued by the ITU-T, then known as the CCITT, in 1984. A major revision to those recommendations was issued in 1988. The functional model developed in the 1984 Recommendations consisted of two major components - The Message Transfer System (MTS) and cooperating User Agents (UAs). The MTS is composed of a series of Message Transfer Agents (MTAs) that are responsible for relaying the message from the originator's UA to the recipients UA. The message may be routed through intermediate MTAs that can perform Application Layer routing based on address information contained in the message. The originator's MTA, the recipients' MTA and the intermediate MTAs may be managed by different organizations or administrations. An administration is either a country's Postal Telephone and Telegraph (PTT) service (in most countries, there is only one PTT) or, in the United States, a common carrier recognized by the ITU-T, such as Sprint or AT&T.

The MTS is responsible for accepting the message from the originator's User Agent and delivering the message to the User Agents of all recipients addressed in the message. The originator's User Agent supplies to the MTS the message content,

the address(es) of the message recipients, and the MTS services that are being requested. In addition to message submission and delivery interaction, User Agents have many functions that are outside the realm of standardization. The originator's User Agent assists in the creation and editing of a message; the recipients' User Agent(s) can use certain fields in the message to control the way that the message is presented to the recipient. Only the message submission and delivery interaction with the local MTA needs to be standardized. The functional model allows for different types of User Agents to communicate with each other. As long as the recipient's User Agent can interpret the data sent by the originator's User Agent, meaningful communication is possible. The Message Transfer System does not examine the message content unless a content conversion service is requested. Although many types of User Agents can potentially use the Message Transfer Service, a common use of the MTS is to send a message from an originator to one or more recipients. For that reason, the ITU-T standardized an Interpersonal Messaging User Agent in the 1984 X.400 Series of Recommendations. The 1988 X.400 Series of Recommendations [REF 2] expanded the model that was developed for the 1984 X.400 Recommendations. The 1988 Recommendations recognized that some User Agents, may not have all the functionality and storage capacity that was assumed by the 1984 Recommendations. For example, most LAN E-mail systems operate according to a client-server paradigm in which the client, a simple User Agent, sends messages to and fetches messages from an intermediary system on the LAN which functions as a Message Store and provides the interface to the outside world. The ITU-T, recognizing this, standardized a protocol for a Message Store to communicate with a remote Message Transfer Agent and also a protocol for a simple User Agent to communicate with a remote Message Store. In practice; however, the protocol used between the User Agent and Message Store is almost always proprietary.

Additional services were also incorporated into the 1988 X.400 Recommendations. Security services were added which enable the various components of the Message Handling System to verify the origin of messages and the integrity of the content as well as prevent unauthorized disclosure of the message content. These services can be provided within the Message Transfer System or by the User Agent and/or Message Store in such a way as to be transparent to the Message Transfer System. The 1988 X.400 Recommendations provide for the capability of interacting with an implementation of the X.500 Directory Service, including the ability to expand a Distribution List contained in a message enroute to its destinations. The 1988 X.400 Recommendations provide for the ability to interact with a Physical Delivery Access Unit which would serve as a bridge between X.400 and other delivery mechanisms, enabling X.400 messages to be routed to a remote system and allowing final delivery to be made; for example, by the postal service. X.400 implementations have been upgraded to provide an interface with X.500 Directory Service implementations. The other services included in the 1988 X.400 Recommendations have, for the most part, not been implemented and will not be implemented unless

there is sufficient user demand.

There was a growing awareness in the messaging community that the X.400 Message Transfer Service could be used to provide the computer-to-computer transfer of order forms, invoices, authorizations for payment and other business transactions. This transfer of structured business data is called Electronic Data Interchange (EDI). Even though many of the most commonly used data forms are standardized by the ANSI Accredited Standards Committee X12 and UN/EDIFACT committees, EDI messaging is structured so that any EDI format, including privately defined data forms, can be conveyed by the X.400 Message Transfer Service.

A user that wishes to use the X.400 Message Transfer Service to send an EDI transaction has different requirements from a user that wants to send an interpersonal message. For example, when an EDI transaction is sent authorizing payment of funds, it is important that the recipient know that the message was sent by the originator and that the amount of funds authorized to be deducted from an account has not been tampered with. It is also important that the originator know that the recipient has received the message intact. In 1991, the ITU-T standardized an EDI User Agent [REF 3] that provides the services required by EDI messaging users. The security services specified in the 1988 X.400 Recommendations were critical for meeting these requirements.

2.2.3.1.2 File Access

The OSI file transfer application is known as File Transfer, Access, and Management (FTAM) [REF 4]. FTAM allows the user to transfer a complete file or a specified portion of a file (down to record-level access). The file formats allowed are numerous, from a simple text file to indexed sequential files; the file characters may be ASCII text, integer, floating point, or Kanji characters (among other choices). An example is: read the 233rd record of a large file directly, update it, and write it back directly. Users can remotely modify filename, file size allowed, and other attributes. Users can specify access control on particular records of files, to prevent others from accessing those records. Users can easily create or delete files remotely; a number of types of file transfer are possible. Users can recover from errors in sending a long file without having to retransmit the entire file. FTP can only transmit complete files, with a limited number of formats. With FTP, however, users can remotely create or delete directories.

Most current FTAM products only implement complete file transfer (the same functionality provided by FTP) and a limited file management capability. Some products support FTAM gateways to FTP and other file transfer applications. There are not many FTAM products in existence, because users have not demanded the additional features that FTAM provides (e.g., record-level access) that are not provided by FTP. No additional features are anticipated for the FTAM standard.

Application gateways have been developed to allow the FTAM and FTP protocols to interoperate; these gateways have not been widely available. RFC 1415 provides a technical specification of a FTAM-FTP gateway. The FTAM service available through such an application gateway is a subset of the full FTAM functionality.

2.2.3.1.3 Terminal Emulation

The OSI virtual terminal (VT) application [REF 5] provides mechanisms to effectively insulate application processes from the specific characteristics of the terminals with which they communicate. This allows terminals to access applications running on a variety of systems, and vice versa, regardless of the supplier of the terminal or host system. The ultimate goal of a virtual terminal (VT) application is that terminals, regardless of model or design, should be able to access application programs resident on either a local or remote system. For example, a VT100 terminal can login to two different computers simultaneously that are not VT-100-compatible. Similarly, any application program should be able to communicate with any terminal, regardless of its model or design and regardless of whether it is resident on the same or a different system. The VT standard does this by defining a series of generic terminal classes, or screen representations, to which real terminals map their information.

Several VT profiles have been defined, supporting several screen formats. One such profile (TELNET) supports a simple line at a time or character at a time dialogue. The transparent profile supports the exchange of uninterpreted sequences of characters. This provides the ability to control terminals directly through the use of embedded control characters and escape sequences. The forms profile supports forms-based applications that provide local entry and validation of data by the terminal system.

Deployment of OSI virtual terminal products is not widespread today, because the protocol is complex and users have not required the additional functionality provided by the VT application. No further extensions are planned to the VT standard. Although a VT/TELNET gateway is technically feasible, it has not been in demand. Although the VT TELNET profile is similar in functionality to IPS TELNET, usage of IPS TELNET vastly predominates.

2.2.3.1.4 Transaction Processing

The OSI Transaction Processing (TP) application [REF 6] allows users to perform multiple operations as an atomic, single, committed action. If any one of the operations fails, then the entire transaction fails. Users can specify the degree of control and notification they have over this transaction. For example, making a plane reservation, a rental car reservation, and a hotel reservation for a business trip would result in high overhead and disjoint control if each action were taken separately. It is more efficient to do all of these in a single committed operation,

since if a user can not get a plane reservation, the car reservation is probably irrelevant. The TP application allows the user to find out whether all of the actions can be completed successfully, and then, if they can, direct that they be carried out as a single action. TP applies among distributed systems, and ensures: 1) atomicity (the total work is performed or nothing is done), 2) consistency, 3) isolation (while the work is being performed inconsistent data is not available to other transactions), and 4) durability (the work is fault-tolerant). This last point is especially important in the context of database management. It means that enough information will be retained so that in the event of a system failure the information on the database can be reconstructed. Some potential uses of TP are in banking transactions, supply and accounting systems, and network management.

2.2.3.1.5 Remote Database Access

The Remote Database Access (RDA) standard [REF 7] provides protocols for establishing a remote connection between a database client and a database server. The RDA standard addresses distributed database processing in this client/server environment. RDA specifies a two-way transfer syntax and, when combined with a database specialization, semantics for database operations.

The RDA standard is specified in two parts. Part 1 defines the RDA Generic Model, Service, and Protocol. Part 2 defines the RDA Specification for Database Language SQL [REF 8]. RDA conformance can only be expressed in conjunction with a specific database language. The RDA SQL Specialization allows the connection of RDA clients with RDA servers conforming to database language SQL. Both the client and the server must conform to the RDA SQL Specialization protocol; however, only the server need provide an SQL conformant client database management system. The client can be an application that simply sends SQL statements to the server. SQL is thus far the only specialization developed to complement the RDA Generic Model, Service, and Protocol.

Users can read database entries and perform updates to databases. Users can cancel update operations without penalty and receive instant status information on their update operation. Users can combine database entries to satisfy specific queries, and can execute their requests immediately or save them for execution at a later time. It is possible to have a one-phase commit (which allows updates at one remote site per operation), or two-phase commit (which allows updates at multiple remote sites in the same transaction). RDA can be used to support access to SQL systems, and Call Level Interface (CLI) [REF 8] functionality could be used as a standard interface to various databases.

The SQL Access Group (now part of X/Open (see sec. 2.4)) has divided RDA into subsets in order to specify basic interoperability. Various groups are prototyping RDA over TCP/IP using RFC1006 (see sec. 2.3). A Transport-Independent RDA

specification (a separate specification) is under consideration by various groups for seamless operation over a variety of transport stacks.

Current RDA implementations use one-phase commit; implementations that incorporate two-phase commit depend upon the popularity of TP and may not be available in the near term. There is no equivalent application in the IPS. Some extensions are being considered to the RDA standard that would support a wider range of SQL-type operations.

2.2.3.1.6 Directory Service

The OSI Directory (joint ISO-ITU-T activity) [REF 9] is a distributed database, capable of storing information about people and objects in various servers distributed across a network. It is these servers, acting in concert, which provide the potentially global access to information made possible by Directory technology.

Distributing information in this manner has advantages over the conventional method of centralizing information storage. The information is kept "close" to those people or processes which are most likely to make most frequent use of it and are most likely to be responsible for keeping it up-to-date - this is likely to reduce access time and network costs, and increase the likelihood of the accuracy of the stored information. Since the information is distributed across several servers, the impact of a given server becoming inactive, for whatever reason, is only to make unavailable the information for which that server is responsible, rather than bringing down the entire database, as would be the case if a centralized server were to go down. The Directory has the capacity to grow indefinitely in size and storage capacity through the simple addition of new nodes. Such growth might be achievable but would be less practical with a centralized system. The Directory offers the opportunity to unify information resources across the globe, as opposed to the insularity which tends to occur when organizations rely on proprietary, centralized databases.

The physical location of the accessed information is transparent to the user. The Directory possesses the necessary knowledge to locate requested information, regardless of where the information might be on the network. All information received will usually appear as if it came directly from the local server. From the user's point of view, the Directory behaves exactly as if the user was accessing a centralized server.

The Directory user accesses the Directory via a client process known as a Directory User Agent, or DUA. The DUA interfaces with the Directory using a standard protocol between itself and one of the Directory servers, termed Directory System Agents, or DSAs. Usually the DSA contacted would be the one closest, in terms of connection cost or organizational affiliation, to the Directory User Agent.

The DSAs making up the Directory also communicate via a standard protocol which embodies a set of operations which may be performed by the Directory (e.g., retrieving information, adding information, etc.). Each DSA knows how to contact one or more additional DSAs (at least one). This is the mechanism through which a Directory request can be propagated through the system: if a particular DSA is unable to satisfy a request, the request is forwarded to another DSA which is more likely to have the necessary information, and so on.

The entry is the fundamental unit of information in the Directory. All Directory entries are made up of a collection of attributes, each of which describes a particular quality or aspect of the object represented by the entry. For example, the entry for "Tim Boland" may have attributes of type "telephone number", "street address", and "email address". Entries are grouped into classes on the basis of shared properties.

The Directory Schema constitutes the framework within which Directory information is stored. It consists of a set of rules and definitions which define the naming of entries, the content of attributes and entries, the structure of the Directory as a whole, and the hierarchical relationships between entries (represented in a structure called the Directory Information Tree, or DIT). Each entry in the Directory is identified by at least one unique name, called the entry's Distinguished Name.

The Directory offers the following services to its users: (1) read (retrieve information contained in an entry), (2) list (used primarily to browse the DIT), (3) compare (compare a user-supplied value against one in the directory), (4) search (look for entries which match specified criteria), (5) add entry (add a new entry to the Directory), (6) remove entry (delete an entry from the Directory), (7) modify the contents of a directory entry, and (8) modify Distinguished Name (change the Distinguished Name of an entry). These services may be presented in a variety of ways, depending on the directory user agent product in use. Attached to a DUA may be a user interface, an X.400 application, or a database application, among other choices.

Additional services provided by the Directory are security and replication. The security services include: (1) authentication of Directory users to establish their identity, and (2) access control procedures to prevent unauthorized access to Directory information. Replication is the means by which information held by a DSA may be copied to one or more other DSAs in order to increase efficiency and decrease access time.

2.2.3.1.7 Additional OSI Applications

The Information Retrieval (IR) application [REF 10] supports the open interconnection of information clients with information servers by specifying an OSI application layer protocol for intersystem search and retrieval of information. The IR

application provides retrieval (but not update) of information and the IR protocol specifies basic information retrieval operations, a common syntax for queries and the means to express their semantics, and the means to allow the partner systems to share an understanding of the information retrieved.

The Manufacturing Messaging Specification (MMS) standard [REF 11] provides for client/server message based communications between programmable devices in a computer controlled environment. MMS defines messages useful for information interchange. It does not define a complete set of services for remote device programming.

2.2.3.2 Network Management

Network management allows users to configure network resources, detect and correct faults, account for network use, monitor and adjust performance, and manage security mechanisms. A user can also determine what network resources need to be managed. Network management is provided by a family of standards [REF 12] (formally known as OSI Systems Management) covering the areas of management communications, management information, and systems management functions and services. The management communications protocol is the Common Management Information Protocol (CMIP). To provide interoperability among network management systems, each system must have a common "view" of management information. System resources to be managed are represented by managed objects using an object oriented model. OSI Systems Management may be applied to system resources (e.g., buffers, disks), not just to networks.

An example of OSI network management is event notification of excessive CPU utilization in two interconnected systems to a remote manager. The manager may then take action to offload some of the work.

For OSI network management, work is progressing on additional managed object definitions and management profiles. Standards work is also continuing in the areas of performance, security, accounting, configuration, and fault management.

The OSI network management protocols have been relatively slow to be deployed, with only a few companies having products available. Since the release of OmniPoint (Open Management Interoperability Points: TM - Network Management Forum), more vendors have implemented OSI network management, particularly for larger, more complex (enterprise) networks in the worldwide telephony industry.

A user should employ IPS network management (SNMPv1, see sec. 2.1.3.3) when the devices to be managed are simple and fewer in number (for example, those that appear on a LAN), because many relatively inexpensive products are available. For more complex devices (or for a larger number of different devices), such as those that appear on an enterprise network, OSI network management may be a better approach, if vendors support OSI management in their manageable products. SNMPv1 is a "lightweight" alternative

to CMIP.

The IPS and OSI network management protocols can and have been deployed over alternative transports. For example, there are several mixed protocol stack approaches defined, including CMIP for the Internet, CMIP over logical link control (LLC), and SNMP over OSI. In addition, there are a set of specifications known as ISO Internet Management Coexistence (IIMC) that defines methods for the management of SNMP resources using OSI managers and vice versa. There has been limited deployment of implementations based on the IIMC specifications.

2.2.3.3 Security

The security services provided in the OSI model are authentication, access control, integrity, confidentiality, and non-repudiation. Data confidentiality services protect against unauthorized disclosure. Protection of medical records to insure a patient's privacy is an example of the need for confidentiality. Data integrity services protect against unauthorized modification, insertion and deletion. Electronic funds transfer between banks requires protection against modification of the information. Authentication services verify the identity of communicating peer entities and the source of data. Owners of bank accounts require assurance that money will be withdrawn only by them. Access control services allow only authorized communications and access to system resources. Only financial officers are authorized access to a company's financial plans. Non-repudiation, with proof of origin, provides to the recipient proof of the origin of data and protects against any attempt by the originator to falsely deny sending the data. Non-repudiation, with proof of delivery, provides to the sender proof of the delivery of data. The non-repudiation service can be used to prove to a judge that a person received or sent a message (e.g., a purchase order).

The OSI architecture defines security protocols at the Transport Layer, Network Layer, and the Application Layer. NLSP (Network Layer Security Protocol) provides security services at the Network Layer, and TLSP (Transport Layer Security Protocol) provides security services at the Transport Layer. Different application layer protocols may integrate security services. For example, for Message Handling Systems, these services are accomplished by having a special module between the User Agent and the Message Transfer Agent which uses public and private keys to provide authentication and integrity (on a message-by-message basis).

As an illustration, authentication/integrity are provided in a message as follows: (1) the originator computes a checksum of the text, then the originator encrypts the checksum using the originator's private key, (2) when the recipient receives the text, it decrypts the checksum using the originator's public key, and then, independently, computes the checksum and compares the two values, and (3) if the values are equal, the message has not been changed and the message has been sent by the originator. If a

confidentiality service is required, the outgoing message is encrypted using a specified key. Each new message can be encrypted with a different key, and the key that is used is encrypted using each recipient's public key. When the recipient receives the message, the recipient's private key is used to decrypt the key that was used to encrypt the message and then that key is used to perform the message decryption.

At the Network Layer, the security services are provided between adjacent intermediate systems. At the Transport Layer, the security services are provided between the communicating end systems. The security services provided at the Transport and Network Layers apply to all applications; the security services provided at the Application Layer can be on a per-application basis. Thus far, implementation of OSI security services has been at the Application Layer. There has been, thus far, little or no user interest in having these services provided at the Network or Transport Layers.

2.2.3.4 Application Program Interfaces

An application programming interface (API) is a conceptual boundary between an application process or program and a service provider. Through this boundary, information is transferred in a structured manner. APIs are programmable interfaces which link application software on a variety of application platforms to a collection of hardware and software components.

The advantages of API use are modularity and portability. Modularized code is written in localized routines or macros so that service provider software may be separated from other applications code. This means that changes in the former may occur without affecting the latter, and vice versa. Portable code can be moved, or ported, to a different environment, and it should operate in the same fashion as before.

Standardized OSI APIs are being developed, at the Transport, Session, Presentation, and Application Layers. These APIs describe methods for accessing the services of these layers in a uniform way. The Institute of Electrical and Electronic Engineers (IEEE) has standardized APIs for Message Handling Systems, Directory Services, FTAM, Generic Application Layer, and Presentation Layer.

Other de facto APIs are being developed, such as the Microsoft Messaging API (MAPI). MAPI is a high-level API involving common messaging functions which provides ease of use and simplicity. Support of MAPI by OSI messaging is anticipated. MAPI is largely independent of the underlying message system used. Simple MAPI (the common messaging functions themselves) is available now, and extended MAPI (with extra features such as file folder support) should be available in the near future. Simple MAPI may be considered a subset of extended MAPI for the purposes of interoperability.

2.3 Coexistence and Convergence

Coexistence and interoperability of different protocol suites (for example, OSI and IPS) has emerged as a strong requirement. Although protocol convergence is an ideal target, this may still be a long-term goal, in particular for application protocols, where gateway operation may be the only way of interoperability for a long time. However, for some applications, common databases are being considered, e.g., for systems management.

A hybrid stack is a selection of protocols in each functional layer such that the protocols selected come from different protocol suites. The hybrid stack that is most likely to be deployed is layers 5 through 7 from the OSI protocol suite, operating with layers 1 through 4 from the IPS. Hybrid stacks are used to take advantage of existing or enhanced capabilities, and allow a user to custom-design an environment. The powerful capabilities of OSI applications can be used with the existing IPS infrastructure to capture the advantages of both. A hybrid stack can interoperate only with another identical hybrid stack.

The most popular solution for allowing OSI applications to run over TCP/IP is specified in RFC1006 [REF 13]. RFC1006 allows a TCP interface to appear as an OSI Transport service interface. RFC1006 allows OSI application protocols to be operated over an IP network, by specifying the provision of the OSI Transport Service over the Internet TCP. RFC1006 does not provide a way for OSI applications on an OSI network to interwork with OSI applications on Internet networks.

XTI (X/Open Transport Interface) has been developed by the X/Open consortium to provide applications independence from any particular transport provider. That does not mean that it automatically makes the application entirely independent of the class of service provided, i.e., connectionless and connection-oriented with or without orderly release. However, a truly transport-independent application can be written using XTI if no assumptions are being made by the application about the nature of the underlying transport service, i.e., as long as it does not expect that all concepts used in one protocol exist or have the same semantics in other ones. Existing applications have to be modified to use XTI. Industry-consensus specifications of how to use XTI over OSI, TCP/IP, RFC1006, and Systems Network Architecture (SNA) full duplex have been developed and published.

The X/Open Multi-Protocol Transport Networking (XMPTN) specification is designed for interworking transport users (applications) that have the same requirements for transport services, or, in other words, "match" each other. This means that the investment in current applications is preserved.

A dual protocol host (or dual stack) has the complete OSI and Internet protocol suites available as part of its networking capabilities. A user of such a host would have the option of

invoking the IPS application protocols or the analogous OSI application protocols. A dual protocol host can be used directly by users with accounts on it to communicate to any OSI or IPS destination. It can also be used as a staging point for manual interoperation between a host that has only IPS protocols and a host that has only OSI protocols.

2.4 Consortia

As an alternative to the OSI and IPS approaches, vendors and users have consolidated operations in consortia. A consortium is an organization whose aim is to develop or implement open systems solutions. There are usually extensive requirements for full membership, and members usually pay membership fees. There may be differing levels of membership, and consortia are usually non-accredited. Specifications are produced which reflect member consensus. The processes for formation, document approval and operation are specific to each consortium. There are two types of consortia: those that create specifications (like X/Open), and those which develop working code (like the Open Software Foundation (OSF)). Consortia may anticipate future standards development and implementation, or define alternative solutions. Their work may be submitted to accredited national standards bodies. Consortia are usually vendor dominated. Members of consortia may be free to develop innovative solutions free of other constraints.

An example of a consortium-produced open systems solution is the Open Software Foundation's Distributed Computing Environment (DCE). DCE is a set of services and tools that support the creation, use, and maintenance of distributed applications in a heterogeneous computing environment. The DCE components fall into two categories: tools for developing distributed applications, and services for running distributed applications. The tools, such as DCE Remote Procedure Call (RPC) and DCE Threads, assist in the development of an application. The services, such as the DCE directory service, security service, and distributed time service, provide the support required in a distributed system that is analogous to the support that an operating system provides in a centralized system. DCE's set of services is integrated and comprehensive. DCE claims to provide interoperability, portability and data sharing across heterogeneous platforms. The DCE is independent of the underlying networking infrastructure; DCE applications can use the Network and Transport Layers of either the OSI or Internet protocol suite.

As another example of a consortium produced open systems alternative, the Common Object Request Broker Architecture (CORBA) was developed by the Object Management Group (OMG). CORBA defines an interface to an Object Request Broker (ORB), which provides the mechanisms that allow objects to make requests on other objects and receive responses to those requests. By building objects atop an ORB and by relying on the ORB for communications between those objects, an application developer does not have to worry about where other objects are and how requests made on those objects

actually get to them. An ORB has software to insulate end users from this knowledge. CORBA also defines a protocol allowing interoperability among different vendor's object request brokers. CORBA is an approach to distributed-object computing. Vendors offer products that support the CORBA technology, although there are interoperability problems caused by incomplete specifications. CORBA applications could, in theory, use either IPS or OSI communications protocols.

3.0 Strategy

3.1 General

The two principal non-proprietary data communications protocols are the Open Systems Interconnection (OSI) protocol and the Internet Protocol Suite (IPS). Any strategy for using non-proprietary data communications technology must consider to what extent each of these protocol suites satisfies the user's requirements. Many users will find that neither protocol suite, by itself, is sufficient and that both protocol suites or hybrids of both protocol suites are needed. In addition, certain consortia developed applications may provide services that are critical to a user's mission. The international standards process that produced the OSI protocols has not led to the widespread availability of computer networking products at prices that most users found attractive. Products result from an interaction between vendors and users in which vendors try to understand, even anticipate, user needs and then develop products in response to those needs. The vendors' goal is always to increase their share of the market in which they compete and to maximize profits. Users vote on the vendors' strategy with their pocketbooks. If the user reaction is favorable, vendors will produce more of the same product which they can now make and market at a lower price, further stimulating demand. If the user reaction is unfavorable, vendors will rethink their marketing strategy and make the necessary adjustments.

There was interaction between vendors and users in the development of the OSI standards. The international organizations that produced the OSI standards and the workshops that produced the functional profiles based on the standards welcomed users. Although the process was dominated by vendors, a significant number of users, mostly from the technical staff of large corporations, regularly participated. These same organizations, however, did not rush forward to purchase products after they had been implemented. In fact, one result of user participation in the standards process was that the users themselves had varying requirements and represented additional interests, each of which had to be accommodated in developing the standards. This increased the complexity of the standards, increased the cost of implementations based on the standards and, consequently, decreased user interest in purchasing the resulting products.

For any commodity, the final user's decision to buy or not to buy is always based on the services provided, the cost, and the available alternatives. The technical staff of the vendor and user organizations that developed the OSI standards did not control the purse strings of their organizations and dialogue about the relationship between services and cost did not occur. In fact, the complexity of some applications, resulting from the tendency of the standards organizations to please everybody, was seen as a drawback, not as an enhancement, and not desirable at any price. The end result is that the OSI protocols are likely to be evaluated by users on a case-by-case basis. Some of the OSI protocols (e.g.,

the Virtual Terminal Protocol) were dead on arrival or, to be more precise, they never arrived. They were implemented by only a few vendors and the specifications are still gathering dust in the files of the large vendors. Use of other OSI protocols (e.g., the File Transfer Access and Management (FTAM) Protocol) will likely be confined to certain niches of the market. Still other OSI protocols (e.g., X.400, X.500 (Directory Service)) should enjoy increasingly greater user acceptance because they provide services that users value and the ability to procure a hybrid stack implementation allows users to procure these services without committing to acquisition and use of the full seven layers of the OSI protocol suite.

It was previously stated that the available alternatives are a major factor in any procurement decision. Because of the complexity, delays and costs associated with the OSI protocols, the Internet Protocol Suite, a protocol suite not originally designed with the business user in mind, became an increasingly viable alternative for procurers of computer networking services. Although use of the Internet Protocol Suite was fueled by the explosive growth of the Internet, there was no constraining linkage and use of the IPS spread rapidly to non-Internet networks. Although the IPS was originally designed for research and development users in the Department of Defense, service upgrades are being made to enable use in a wider business environment. The frequent service upgrades to both the IPS and OSI protocol suites complicates the procurement decision process and requires a continuing assessment of the available technology.

This section will examine some of the technical and marketplace factors that should influence a user's procurement decision. Some of the guidance will be general in nature; other guidance will be specific to the services being procured. When more than one viable alternative exists, users must thoroughly understand their own requirements as well as the services provided by the alternative protocols. Only then will they be able to perform an effective matching of requirements and available technology that is essential when making good procurement decisions. All of the following principles are not specific to the procurement of computer networking products and services. Even though some would consider these principles an application of common sense, they are discussed because they are frequently violated in the procurement process.

1. **Determine the community of users with whom you currently interoperate and those with whom you intend to interoperate in the future. Develop interoperability agreements with that community of users.**

The original vision of the international standards community was that their efforts would result in one set of internationally recognized data communication protocols which would be implemented worldwide. The protocols which they standardized would allow, for example, a Federal agency user in the United States to send mail to a French corporation and to receive mail from a university in Japan

with the assumption that all users would have the same networking infrastructure and the same application layer services. There would be a globally distributed Directory Service supporting the transfer of electronic mail. There would also be the capability of retrieving files, accessing remote data bases, and performing transaction processing using the same internationally recognized set of protocols. A standardized set of applications would satisfy most of the requirements of business users; they would only need to develop the special purpose applications of interest to their own community and even these would interoperate over the standardized networking infrastructure. That vision, however, has been overtaken by reality. The data communications world will consist of multiple non-proprietary and proprietary alternatives for data transfer. Any user choice can lock out future communication with other users. It is critical that representatives of all organizations or groups that need to interoperate participate in selecting the protocols that will permit interoperation. This may require inter-agency groups to be formed or even groups that include non-government organizations that want to have on-line access to government data. By "circling the wagons" around the community of users with common interests, and making decisions that are agreed by and apply to that community, procurement authorities will insure that they will not be adversely impacted by developments over which they have no control.

- 2. Determine the computer-networking services which are essential to the conduct of business in your organization. Determine those services which, though desirable, are not essential to the business enterprise.**

The services that are procured determine the cost. Separating the essential from the merely desirable allows the procurement authorities some flexibility when negotiating with the vendors. In making this distinction, it should not be assumed that services that are required by some personnel need to be provided to all personnel.

- 3. All policy decisions should take into account the Information Technology budget. Standards groups within an organization should not set policy requiring the procurement of computer networking services if there is no money to pay for them.**

This rule is often violated by Federal agencies. The standards group has limited interaction with the budget group and the budget group frequently cannot anticipate the availability of funds beyond the current fiscal year. Even when there are budgetary uncertainties, the standards group should still issue computer networking policy which sets the direction for the organization. This policy should have, in advance, the endorsement and support of everyone affected by the policy, particularly top management.

- 4. Be knowledgeable about the networking services provided by current products and about vendor plans for augmenting those services in the future.**

The removal of the requirement to acquire OSI products allows Federal agencies more options when procuring computer networking products but it also requires them to be more informed about the functional capabilities, performance, cost and availability of each option. Users need to regularly consult the trade press and their vendors to anticipate advances in computer networking technology so that they do not procure functionality that will soon be obsolete. It is important that the interaction between users and vendors be a dialogue; the issuance of user requirement documents or profiles will do nothing to influence vendor actions unless these documents are backed up with a significant amount of procurement funds.

5. Develop a realistic workable transition strategy that is appropriate for all users affected by the procurement.

User resistance to the introduction to new technology can be expected unless the users understand how the new technology will improve efficiency and make their job easier. If the transition involves some temporary hardships, users must be able to see how these inconveniences are a necessary part of an overall rational transition strategy. The transition strategy that is adopted should be evaluated at regular intervals to determine whether the original expectations are being met and, if not, to make the necessary adjustments.

3.2 Testing

Users should consider carefully the type and extent of testing that will be required and avoid unnecessary duplication. There are four categories of testing that should be considered when procuring computer networking products: conformance testing, interoperability testing, performance testing and functional testing. The purpose of each type of testing is discussed in Appendix A. Appendix A also discusses current testing procedures and specifies the organizations that are responsible for evaluating and registering test results.

Since there is a cost associated with all new testing that is mandated in connection with a procurement, both users and vendors should have an interest in keeping new testing to a minimum and making use of any information about the technology to be procured that is publicly available, even if it is not formally recorded.

Some of the factors that should be considered in making the testing decision are:

1. Level of maturity of the technology being procured.

The fact that the computer networking protocols in the product to be procured have been operating successfully in a real-world environment over an extended period of time should give most users an increased level of confidence in that product even if the protocols have not undergone formal conformance or interoperability testing. Testing which focuses on areas of specific concern to the user may be all that is warranted.

2. Magnitude of the Procurement

In a small procurement, additional mandated testing can add significant overhead to the procurement cost. In a large procurement, the testing may add only a small increment to the procurement costs.

3. Number of Current Validated Products

If multiple products satisfying a user's requirements already exist on a register, the user should procure one of those products instead of requiring additional testing by a vendor whose products have not yet been formally certified.

4. Availability of Tests and Test Systems

Users should not mandate testing if there are no test systems or only test systems of dubious quality available to perform the testing. Users should have a thorough understanding of the tests and systems to be used before they require testing in a procurement.

5. Criticality of the Services Being Acquired

Users should develop a worst-case scenario for the impact on their organization if an error or errors should occur in the computer networking protocols that are being procured. If the result is only a minor disruption in the activities of the organization, users may determine that an increased confidence level is not worth the cost of additional testing. If, however, a single error has the potential of causing a severe security breach or a large financial loss, then the additional cost of testing may be a very wise investment.

3.3 Security

Each organization should determine its requirement for the following security services:

- (1) Data Integrity services protect communicated data from unauthorized modification, insertion, or deletion.
- (2) Authentication services identify and authenticate the data originator to all recipients. Authentication services are necessary for implementing access control.
- (3) Access control services allow only authorized communications and access to system resources.
- (4) Non-repudiation protects against any attempt by the originator of the communicated data to later falsely deny sending the data.
- (5) Data confidentiality services protect against unauthorized disclosure.

Some organizations may find that their authentication requirements can be met by unprotected or protected passwords. Other

organizations will find that they have a requirement for strong authentication services. Strong authentication services are based on the use of digital signatures which rely on the use of public key cryptographic techniques. These techniques also provide assurances of data integrity and protect against originator repudiation as a by-product to the strong authentication service. The National Institute of Standards and Technology issued Federal Information Processing Standard (FIPS) 186, Digital Signature Standard [REF 14], which became effective on December 1, 1994 as the government standard for use by all Federal agencies which require a public key cryptographic signature system for unclassified sensitive information.

Public key cryptography makes use of two paired keys: a public key and a private key. The public key can be known by anyone. The private key is secret and its use is controlled by its owner. Encryption performed using a private key can be decrypted using the corresponding public key; encryptions performed using a public key can be decrypted using the corresponding private key. In order for public key cryptography to work effectively, there must be a Certification Authority which is responsible for the public/private key pairs and for securely transferring the private key to the associated user. Discussions are underway within the Federal government relating to how this service will be provided and who should provide it. It is expected that it will be several years before a Certification Authority service will be implemented and available to all Federal agency users.

The Digital Signature Standard does not provide a confidentiality service. If confidentiality is required, the data originator could first encrypt the message using the Data Encryption Standard (DES) before signing it, using the algorithm specified in the Digital Signature Standard.

There is another technique, however, of providing a confidentiality service in such a way that it does not jeopardize effective law enforcement, public safety, and national security. Furthermore, the infrastructure required to implement this technique is currently operational. The technique is based on a special tamper resistant hardware encryption device (Capstone Chip) which implements a strong encryption algorithm (SKIPJACK) and a Key Escrow System which gives the government access to a device unique key that unlocks all communications encrypted by the chip pursuant to a lawful court order.

This method of providing confidentiality services has been endorsed by the Department of Defense (DOD) and will be implemented in the recently awarded Defense Messaging System. The SKIPJACK algorithm will be implemented on a FORTEZZA Personal Computer card. Several hundred thousand DOD users will receive FORTEZZA cards to encrypt and authenticate messages in the Defense Messaging System. The FORTEZZA card, which is inserted into a PC card reader, contains specific cryptographic algorithms, unique encryption and decryption keys and other information that can be used to encrypt messages

sent to other FORTEZZA card users and decrypt and authenticate messages received from other FORTEZZA card users. This system requires that users procure a FORTEZZA card and that their PCs be equipped with a card reader. The cost of the FORTEZZA cards is declining (currently about \$70) but the requirement for a PC card reader may be a significant barrier to widespread deployment outside of DOD. Another alternative that is being investigated and tested in a pilot program is to put the Data Encryption Standard (DES) [REF 15] and DSS algorithms and required key material on a 3.5 inch floppy drive. Since no modification to existing PCs would be required, this method should be particularly attractive to agencies like the IRS who have a need to securely send messages to and receive messages from private citizens.

3.4 Network Management

There are three main protocols for consideration: (1) CMIP (see sec. 2.2.3.2), (2) SNMPv1 (see sec. 2.1.3.3), and (3) DCE management (see sec. 2.4). It is likely that SNMPv1 will be the network management protocol implemented on most LANs; there are many SNMPv1 products available, but SNMPv1 is lacking in security functionality. CMIP may be more appropriate for larger, more complex enterprise environments, but there are not many products available. Similarly, there are also not many DCE management products available.

3.5 Electronic Mail and Directory Services

In July 1993, the Office of Management and Budget (OMB) chartered an Electronic Mail Task Force (EMTF) to examine the state of electronic messaging among Federal agencies and to make recommendations on what should be done to provide interoperable business-quality electronic mail throughout all Federal government agencies.

In April, 1994, the EMTF issued its final report which defined business-quality E-mail as "a service that appears to the user to be a single, unified electronic postal system that offers robust and trustworthy capabilities with legally-sufficient controls for moving all forms of electronic information among employees at all levels of government, and with the public we serve; and, like the nation's telephone network, is affordable, ubiquitous, efficient, accessible, easy-to-use, reliable, cost-effective and supported by an effective directory service." The recommendations of the EMTF included the following: (1) require Government-wide E-mail connectivity, (2) establish a Government-wide E-mail Directory, (3) establish an E-mail Program Office.

The EMTF report recommended that the DOD Defense Messaging System (DMS) operational characteristics specifications be an important source of information to be used when defining the operational characteristics of business-quality Government-wide E-mail. The Defense Message System is an X.400-based system that includes service elements added to the commercial Interpersonal Messaging

User Agent (IPM-UA) to meet the Department of Defense (DOD's) business-class messaging requirements. The contract award to implement the system was made in the spring of 1995.

Security services, including authentication, integrity, and confidentiality, are provided by algorithms and keys that are stored on a FORTEZZA card which is input to a Personal Computer Memory Card Interface Adapter (PCMCIA). The FORTEZZA PCMCIA card interacts with an Application Programming Interface (API), which expects as input the requested security services and the algorithms and keys to be used to provide those services.

DMS Directory Services will use Directory Service Agents and Directory User Agents that conform to the ITU-T X.500 Series of Recommendations.

In March 1995 DOD issued electronic messaging policy and implementation guidance that stated that the DMS will meet all DOD messaging requirements and will be used to provide a single, seamless, end-to-end global electronic mail service within DOD. The policy statement further directed that DOD should migrate to DMS-compliant messaging as rapidly as possible and imposed a moratorium on the acquisition of non-compliant electronic mail systems unless a transition path to full compliance could be documented.

The EMTF recognized however that Federal agencies outside of DOD may have different requirements for business-class messaging than the Department of Defense. The EMTF recommended that an Electronic Mail Program Management Office (PMO) be established to coordinate, shape and implement policies that would provide electronic interoperability among Federal agencies and between each Federal agency and the outside community that it serves. The Office of Management and Budget (OMB) chartered and funded a group within the General Services Administration (GSA) to perform this task.

Business quality messaging, as currently envisioned by the GSA E-mail PMO, includes the ability to send and receive messages containing both text and binary data using either the non-proprietary X.400 or Simple Mail Transfer Protocol (SMTP) applications. SMTP must include MIME functionality in order to process binary data. Directory services will be provided by implementations based on the ITU-T X.500 Series of Recommendations. Application layer gateways will be needed if there is a requirement to interconnect X.400 and SMTP mail systems.

The GSA E-mail PMO has developed a two-year plan to provide business quality messaging to Federal agencies by 1997. Government-wide business quality messaging however will not happen without the active participation of experts from all Federal agencies. The GSA E-mail PMO will act as a coordinator and catalyst to bring Federal agency experts together to develop the detailed functional documents that specify Federal agency electronic mail and directory service requirements and the strategy

required to implement those requirements. The GSA E-mail PMO will also be responsible for removing all barriers that would inhibit the implementation of that strategy.

Given the aforementioned Federal government activity, it is recommended that each Federal agency take the following steps:

(1) **Define what business quality messaging means for your agency**
Some of the questions that need to be answered include: What services are essential? What services, although not essential, are desirable? Are the same services required for all intra-agency and inter-agency communication or do different intra-agency groups have differing requirements? To what extent are the requirements of the agency affected by the decisions made by users outside the agency - decisions over which the agency may not have full control? What are the agency's security and performance requirements? What compromises will have to be made because procurement funds are limited?

(2) **Work with other Federal agencies to develop a common specification for business quality messaging**
It is in the interest of each Federal agency to align their requirements for business quality messaging with those of other agencies. The E-mail PMO has established the means by which the necessary inter-agency interaction can occur. The E-mail PMO has set up two on-line forums to facilitate communication among Federal agencies for the purpose of achieving a consensus on required E-mail services. Federal agency personnel can subscribe to one or both on-line forums by sending an e-mail message to listproc@etc.fed.gov with
SUBSCRIBE email-1 NAME OF SUBSCRIBER
and/or
SUBSCRIBE X.500-1 NAME OF SUBSCRIBER
in the body of the message. To get a list of additional on-line groups that deal with related issues, put LIST in the body of the message. In addition, a volunteer working group called the Electronic Messaging and Directory Working Group (EMADWG) has been formed to address a large range of messaging issues at the working level. For further information about E-mail PMO sponsored activities, contact Jack Finley at Jack.Finley@gsa.gov.

(3) **Work with other Federal agencies to develop a directory service schema to support business quality messaging**
The E-mail PMO is working with independent contractors to develop a directory services schema that can be used by all Federal agencies. A directory schema consists of rules which define the naming of entries, the content of entries and attributes, the structure of the directory as a whole and the hierarchical relationships between entries. An entry is the universal unit of information storage in the directory. An entry consists of one or more attributes which can be mandatory or optional for that entry. Each entry belongs to an object class which governs the attributes that the entry must or may

contain and position that the entry may take in the Directory Information Base (DIB). The DIB is organized in a tree-shaped hierarchy known as the Directory Information Tree (DIT) in which each entry has exactly one superior entry but may have many subordinate entries. The Directory Service Schema will specify a set of structure rules which define the hierarchical relationships that will exist between entries of different object classes.

The Directory Services schema that is being developed by the E-mail PMO will be available in draft form for comment and approval by Federal agencies. Federal agencies should ensure that all requirements that they have in common with other agencies are incorporated into the final version. The previously mentioned X.500 on-line forum and the EMADWG will play an important role in the review of the draft document. Agencies should understand that the schema that is being produced will allow them to add object classes and attributes of particular interest to their agency.

(4) Determine how the directory service will be administered within your agency

The DIT can be broken up into subtrees. The E-mail PMO will administer the portion of the DIT dealing with the U.S. Government. The X.500 Directory is modeled as a distributed service; it can consist of a number of Directory System Agents (DSAs), each of which holds a portion of the overall Directory Information Base. The E-mail PMO, under its authority to administer the portion of the DIT under Country=US, Organization=U.S. Government, will assign Organizational Unit names to requesting government organizations and implement a knowledge reference DSA which will contain the address of all DSAs that implement the Directory Service at the Organizational Unit level of the DIT.

Each organization that requests an organizational unit name from the E-mail PMO must decide how to structure and administer the subtree in the DIT that contains the assigned Organizational Unit as its root entry. The requesting organization must know the extent of its administration authority. For example, does the Department of Treasury have the authority to administer the name space for the Internal Revenue Service? If so, then the Internal Revenue Service becomes a subordinate Organizational Unit under the Department of Treasury. Does the Internal Revenue Service have the authority to request and administer its own Organizational Unit subtree? There is no correct answer but policy must be determined before a request for the assignment of an Organizational Unit name is made.

Once the extent of the administration authority of an organization is determined, the next issue to be determined is how the administrative authority is to be partitioned within the organization. Each agency must determine how many DSAs

will provide users with transparent access to the data stored in their portion of the DIT and what part of the DIT, or what naming context, is the responsibility of each DSA.

In summary, Federal agencies should collaborate to develop a common specification for business quality messaging. The E-mail PMO has set up mechanisms to facilitate collaboration. Although the EMTF has stated that the DMS operational characteristics should be an important input to this process, agencies are not precluded from determining that the SMTP application meets the business needs of their agency and implementing an SMTP-X.400 gateway to reach X.400 users; users of X.400 systems will also require a similar gateway to reach SMTP users. The X.500 Directory Services application, however, appears to be the only viable means of providing government-wide directory services. Users should understand that use of X.500 Directory Services does not preclude running this application over the TCP/IP transport and network protocols (in fact, this is the most likely scenario). In addition, implementing the X.500 Directory Service does not imply use of the X.400 mail service. The X.500 Directory allows an attribute type hierarchy to be set up that allows X.400 addresses, SMTP addresses, and mail addresses used by proprietary systems to be stored and accessed.

The OSI (X.500) Directory Service has a good chance of gaining marketplace acceptance, because it offers more functionality than its IPS counterparts, and because users seem to want and require that additional functionality. There are no current ISO/IEC or ITU-T standardization activities relating to major functional extensions for OSI Directory Services.

3.6 Electronic Data Interchange

Electronic Data Interchange (EDI) is the computer-to-computer transfer of structured business data, including Request for Quotations (RFQs), purchase orders and insurance claim forms.

Early electronic data interchanges were based on proprietary formats agreed between two trading partners. In the late 1970's, the American National Standards Institute (ANSI) chartered the Accredited Standards Committee (ASC) X12 to create a set of standard data formats to facilitate the electronic exchange of business information. Each year, the ASC X12 Secretariat publishes the entire set of the latest X12 standards and the new draft standards.

Independently, the United Nations Economic Commission for Europe created another family of standards known as Electronic Data Interchange for Administration, Commerce and Transport (EDIFACT). Although X12 is more mature and provides functions not currently present in EDIFACT, many of these functions are currently under development in the EDIFACT standards community. The differences in the X12 and EDIFACT syntaxes; however, make interoperation impractical. Although the ASC X12 committee has agreed to develop standards based both on the X12 and EDIFACT syntaxes, many X12

users view switching standards as a cost with little financial return and have resisted a move to the EDIFACT standards.

Thus, development of X12 standards will continue and a vote to reexamine this decision will be taken every three years - the next vote will occur in 1998. The delay in sole support of the EDIFACT syntax increases the likelihood that a decision to support only the EDIFACT standards will never occur. As more and more X12 transaction sets are standardized, users will have a greater dependence on the X12 standards and a transition to the EDIFACT formats will become even more difficult than it is now. Thus, although the EDIFACT standards may be used by organizations with international electronic data interchange requirements, inertial forces are likely to lead to continued use of the X12 standards for national electronic data interchange in the foreseeable future.

The software component that governs the conversion of application data to and from EDI interchanges is called an EDI translator. For outbound transactions, an application writes the transaction data to a file. The translator formats the data according to the appropriate EDI syntax rules and produces an EDI file that is ready to be communicated to the trading partner. The process occurs in reverse for inbound transactions.

The communication service is not part of the translation process. The EDI standards do not specify how EDI documents are to be transmitted to a trading partner. Most EDI trading partners currently use the services of a third party Value Added Network (VAN). EDI VANs provide a communications network to connect trading partners, regardless of individual hardware platforms or communications protocols. Each partner connects to the VAN and the VAN manages the connections to all of the trading partners. VANs also serve as a document clearinghouse, providing a mailbox service to store received interchanges until a user is ready to download them. VANs may provide other services such as security, directory, or translation.

In order to provide an additional method for transmitting an EDI document, the ITU-T (then the CCITT) issued Recommendations F.435 [REF 16] and X.435 [REF 3] to specify the services and protocol for an EDI User Agent (EDI-UA) which will use the services of the Message Transfer Agent [REF 2], specified in the 1988 Series of Recommendations for Message Handling Systems to transfer an EDI document that is output from an EDI translator as a body part in the content of an X.400 message. The EDI document can be an X12 transaction set, an EDIFACT message or even a privately defined format. Although the format differs, an X12 transaction set and an EDIFACT message are functionally similar, i.e., both can be used to represent a single Request for Quotation (RFQ) or a single purchase order. The vendors that market X.400 compliant mail systems generally do not market EDI translators; the EDI translator has to be procured separately. Users will have to know which EDI translators have a compatible interface with the X.435 EDI User Agent that they are procuring.

Each organization will have to make a decision about the communications service that it will use to transfer EDI data. There are strong obstacles that have to be overcome in order for X.435 to become the preferred means of transferring EDI documents. Organizations are not likely to purchase X.400 systems primarily for use in EDI; a full-scale commitment to use X.400 services will probably be needed in order to justify the procurement costs. Even in those organizations that choose to procure X.400 systems, there may be strong resistance to changing the current modus operandi if it appears to be working. For this reason, organizations should not procure X.435 User Agents until a comprehensive EDI policy is agreed both by the organization and its suppliers. For guidelines for the evaluation of EDI products, consult [REF 17].

3.7 Message Attachments

A variety of word processing applications are in widespread use. Frequently, there is a requirement to send the output of one of these applications as an attachment to an electronic mail message. Some word processing applications, including WordPerfect and Microsoft Word, output files in binary format, other formats such as LATEX and Rich Text Format are in ASCII. If an attachment is in binary format and all mail systems through which the message passes cannot process binary data, then the attachment has to be converted to seven bit ASCII before transmission and converted back to binary format upon receipt before further processing can be performed on the message. In UNIX systems, the programs UUENCODE and UUDECODE are commonly used to perform the required encoding and decoding functions. In addition, if a large amount of data is being sent, it may be desirable to compress the data before transmission. If a binary to ASCII conversion is necessary, the compression will occur before the conversion to ASCII and the decompression will occur following the ASCII to binary conversion upon receipt. The compression/decompression and format conversion operations are not performed by the mail system; they have to be performed by the originator prior to transmission and by the recipient after reception. The family of programs that perform the compression and decompression, as well as the encoding and decoding programs, have to be compatible.

If the recipient's electronic mail end system does not contain a word processing application that is capable of processing the attachment, then the message has to be downloaded to a system containing such an application (using; for example, the File Transfer Protocol (FTP)). Any required decoding or decompression in order to create a file that can be input to the word processing application can be done either by the electronic mail end system or by the system containing the word processing application after the file is transferred.

New versions of some word processing applications will recognize and convert formats other than their own. For example, Word Perfect 6.0 will recognize a file in Microsoft Word format and convert it to Word Perfect 6.0 format. It can then be read or

printed in the same manner as any other Word Perfect 6.0 file.

Another complicating factor in the forwarding of attachments can occur when the message is processed by a gateway that performs the conversions required to permit two different mail system to interoperate. The gateway must be able to determine if the attachment is in ASCII or binary format and, if the attachment is in binary format, further determine that the mail system to which the message is being forwarded can process binary data. If that mail system cannot process binary data, the gateway must either call a program such as UUENCODE to convert the attachment to ASCII text or refuse to forward the message. Recent versions of the X.400 application recognize and process binary data; SMTP must include MIME functionality in order to process binary data. The bottom line is that as long as heterogeneous document types exist, users that mail documents as attachments to electronic mail messages must be knowledgeable about the ability of the interconnected mail systems to handle binary data and also be knowledgeable about the document processing applications that are available to all recipients.

3.8 File Access

FTAM products are not expected to be plentiful, because of their expense and the existence of readily available alternatives, such as the File Transfer Protocol (FTP). Only a small fraction of the functionality specified in the FTAM standard has been implemented in FTAM products. Thus, in terms of functionality available in products, FTP has equivalent functionality in terms of file transfer, and in addition has file directory access and manipulation functionality that FTAM products do not possess.

3.9 Information Retrieval (IR)

IR products are expected to become numerous in the next few years, in part because of the growing popularity of free text search and retrieval techniques. Existing products contain many options relating to retrieval efficiencies and particular algorithms to be used. Agencies need to weigh carefully the various options and features of the available product offerings, in order to make an informed selection. It should be noted that in many cases, software may need to be added to existing IR capability to provide a particular optimization or a customized user interface.

3.10 Remote Database Access (RDA)

The RDA protocol establishes a remote connection between a database client and a database server and allows the client to access remote distributed databases. RDA allows access to the database systems built by different vendors using the SQL database access language. RDA does not preclude the use of other database access languages; however, only the SQL specialization has been standardized and implemented.

Distributed Query Processors (DQP) can easily be built which parse the query submitted by a client and use RDA to access the appropriate database(s) using a directory stored in the DQP. If the query requires that more than one remote database system be accessed, the returned data can be stored in the DQP until all data is received, then the data can be merged and sent to the client.

RDA eliminates the dependency of users on a single database system. For that reason, even though RDA products are beginning to appear in the market, some of the large database systems vendors are not marketing RDA technology and probably will not unless there is user demand. There are some overhead and service limitations associated with the use of the RDA protocol. Single-vendor database systems are generally more efficient, but, for many users, this fact is offset by the ability to access heterogeneous database systems.

In summary, products exist but may not be aggressively marketed by vendors, who would prefer to sell their proprietary systems. Until RDA vendors can fully establish themselves in the marketplace, users will have to determine if the services RDA provides are worth minor service and performance degradations.

3.11 Transaction Processing

On-line Transaction Processing (OLTP), developed by X/Open, will be popular for transaction processing applications. It has many of the features of full OSI TP, and may be considered a "subset" of OSI TP. Products are available, supported by many vendors. There are few full OSI TP products available, and few are anticipated. Users should consider available offerings (such as OLTP) for their transaction processing functionality, and weigh the benefits and features of any options or enhancements provided by vendors. No further extensions are planned to the TP standard. There is no application corresponding to OSI TP in the IPS.

3.12 Network Infrastructure

Some network technologies available to users are Asynchronous Transfer Mode (ATM), frame relay, Switched Multimegabit Data Service (SMDS), narrowband and broadband Integrated Services Digital Network (ISDN), Distributed Queue Dual Bus (DQDB), Fiber Distributed Data Interface (FDDI), high-speed Ethernet, switched Ethernet, and wireless LANs. Many of these technologies are available in products today, and many more products will be available in the next few years. Thus there is a large array of choices available to the user. It is likely that there will be a multiplicity of network technologies into the foreseeable future. In addition, more and more agencies are replacing traditional media with high-speed fiber.

It should be noted that applications described in this paper can employ any of these network technologies (since the transport functionality "shields" the user from changes in underlying technology). These network technologies can be used in combination

for maximum effectiveness. For example, an approach gaining in popularity is to use switched Ethernet over an ATM backbone. For more information on ATM, consult [REF 18]. The use of high-speed fiber in place of traditional media may cause changes in existing Transport and Network layer specifications, including IPng.

3.13 APIs

3.13.1 Electronic Mail API

There are several different types of electronic mail application program interfaces (APIs). Each type will be described listing the advantages and disadvantages and example implementations.

3.13.1.1 Client APIs

The client service calls in this type of API are not tied to a particular electronic mail implementation. They are generic (i.e., in the form ReadMail or SendMail). The client service calls are mapped into service calls to another client API for a specific electronic mail implementation. The mapping will vary with each new implementation. Because the service calls are generic, the application cannot make full use of the services provided by the underlying implementation. The Common Mail Call (CMC) API developed by the X.400 API association is an example of this type of API.

3.13.1.2 Server Provider Interface

This type of API specifies a server provider interface which facilitates the mapping of the generic service calls into the services provided by different electronic mail implementations. The API can be architected so that both client and server make service calls to an operating system. The Microsoft Application Programming Interface (MAPI), for example, provides an interface between the Common Mail Call API and different electronic mail implementations through service calls to the Microsoft operating system. Alternatively, the API can be implemented outside of the operating system with a plug-interface specified for both client and server. This type of API minimizes the software that has to be written by both client and server.

3.13.1.3 Specific Implementation

This type of API allows the client to make maximum usage of the services provided by the underlying messaging system. The service calls are not generic but are tailored to the services of the mail system. The disadvantage of this type of API is that it binds the client to the services of one mail system; changing the mail system may require significant modifications to the client application. X/Open, the X.400 API Association and the IEEE developed an API which is linked to the services provided by the X.400 Message Transfer System (MTS). This API makes X.400 MTS services available to portable user agents and to other mail systems that wish to use

X.400 as a gateway.

3.13.1.4 Mail API Summary

In choosing an electronic mail API, an organization must consider whether there is a strong commitment to one electronic mail system by that organization or whether a more flexible solution is required. Other factors to be considered include availability and cost. Currently, the Microsoft Application Programming Interface is the most widely implemented electronic mail API.

3.13.2 Transport API

RFC1006 allows OSI application protocols to be operated over an IP network, by providing the OSI Transport Class 0 Service over the Internet TCP. RFC1006 products are growing in number.

XTI (X/Open Transport Interface) provides independence from any particular transport provider. A truly transport-independent application can be written using XTI if no assumptions are being made by the application about the nature of the underlying transport service. XTI functionality, along with Sockets functionality, is included in an IEEE-published Protocol Independent Interface (PII) specification.

3.13.3 Other APIs

Several APIs have been developed by the IEEE, including: (1) Association Control Service Element (ACSE)/Presentation Layer API, (2) FTAM API; and (3) Directory Services API. None of these is likely to gain wide acceptance in products over the next several years, for several reasons, including the close connections with the actual protocols with which they interface. APIs which provide a higher degree of independence are likely to gain more acceptance by the user community.

3.14 Collaborative Computing

Collaborative computing may be defined as the ability of two or more users in different locations to interact in production as though they were in the same room (an example is videoconferencing). Concurrent engineering refers to the development of the distributed infrastructure necessary for collaborative computing and other applications. Groupware may be defined as the software (or middleware) necessary for groups to operate effectively in a production-oriented environment.

Collaborative computing applications may use the protocols discussed in this paper to provide the necessary functionality for productive interaction, or these applications may be built directly over existing infrastructures. The nature of the existing infrastructure may determine the types of functionality that can be supported in collaborative computing applications. The amount and diversity of product offerings is expected to expand greatly over

the next several years, with the tendency towards more sophisticated products and functionality (mirroring the development of more sophisticated network technologies).

Standards bodies (i.e., ITU-T (formerly CCITT)) and consortia (i.e., Object Management Group) are developing collaborative computing specifications. In addition, vendors are developing their own approaches; these approaches may represent "bundled" functionality, or stand-alone functionality, and vendors are also joining forces to combine functionality into larger packages. NIST is setting up a laboratory to demonstrate and test collaborative computing applications. The most advanced products offer a complete set of tools in one product. Examples of such tools include: (1) searching on electronic mail, (2) maintaining "chat" sessions, (3) searching, managing and sharing documents, and (4) sophisticated communication between servers. Growth areas include the development of a multimedia capability, improved security and the integration of existing infrastructures. The separate development of WWW servers and Internet newsgroups may impact the collaborative computing market in the years to come.

3.15 Gateways

Gateways for applications other than messaging applications are likely to be unwieldy and cumbersome. The messaging application is most adaptable to being gatewayed because it, alone, of all the standardized applications, is modeled as a store and forward application. The implementation of an application gateway for other applications will transform those applications into store and forward applications. Since they were not so designed, the end result is likely to be a noticeable performance degradation.

Mail gateways are application layer gateways that can vary in design. Mail gateways can be tailored for two specific mail applications. This maximizes the services that can be provided, at the expense of flexibility. Another solution is to have a common mail gateway to serve all applications. The X.400 protocol is most suitable for this purpose because it provides a superset of the services of most mail applications. But any gateway can only provide the service subset that the two gatewayed mail protocols and the gatewaying protocol have in common. A user served by mail application A that is gatewayed to mail applications B, C, and D will most likely be provided with services that will differ depending upon the application that generated the message. This can become more annoying and exasperating to users as more and more mail applications are gatewayed and may lead to decreased usage of electronic mail for communication.

The bottom line is that mail gateways may have an important role to play in a transition to a single mail system or in allowing interoperability with a community of users that will continue to be served by another mail application. Users, however, should not look upon them as an interoperability panacea. They have significant drawbacks that should be understood before they are

widely implemented. The fewer gateways that are implemented, the better.

3.16 Consortia

It is considered that some consortia have formed for the wrong reasons (that is, they are reactive, not proactive). In addition, consortia are usually closed, so there has long been a resistance to adopting their standards (possible antitrust concerns). However, the specifications may be closely tied to operational environments, and may be based on implementation experience. So, for these reasons, plus the fact that comprehensive solutions may be included, there is strong consideration towards adopting these consortium-based models by some end users.

It is still unclear which consortia-based networking products, if any, will gain widespread acceptance. Vendors may be using consortia to promote their own proprietary solutions, and users are likely to evaluate implementations based on consortia-developed standards on a case-by-case basis.

As an example, more detail needs to be added, and more implementation experience gained, before CORBA gains widespread acceptance. In particular, it is unclear whether CORBA is merely a linkage connecting proprietary ORBs or is truly an open systems solution. A viable CORBA-to-DCE link should greatly enhance CORBA's acceptance. The CORBA RPC should be considered as an alternative to DCE RPC and OSI in terms of establishing communications.

DCE products have a presence in the marketplace currently, but the complexity of DCE poses concerns for future growth. Applications which have important security requirements may continue to take advantage of DCE capability. User requirements are being gathered for the next generation of DCE functionality, called "DCE 2000".

4.0 Summary

The absence of a single solution for providing worldwide computer interoperability places more responsibility on procurement authorities to make the right choice. In order to make intelligent procurement decisions, procurement authorities will have to be informed, both about the technical features and marketplace acceptance of these products, to a degree that was not previously anticipated. This Data Communications Strategy document provides procurement, managerial and technical personnel with some of the factors that need to be considered when making these decisions.

APPENDIX A. Testing

There are four categories of testing that should be considered when procuring computer networking products: conformance testing, interoperability testing, performance testing and functional testing. Conformance testing verifies that an implementation acts in accordance with a particular specification. Interoperability testing duplicates the "real-life" environment in which an implementation will be used. Performance testing measures whether an implementation satisfies the performance criteria of the user. Functional testing determines the extent to which an implementation meets user functional requirements.

If appropriate, users may require vendors to demonstrate conformance and interoperability by consulting NISTIR 4594, "GOSIP Conformance and Interoperation Testing and Registration" [REF 19]. A register of OSI products is maintained for the convenience of those agencies that wish to acquire products based on OSI standards.

The Joint Interoperability Test Center (JITC) at Fort Huachuca maintains an online database facility that provides information for the following list of registers: (1) OSI Abstract Test Suite (ATS) (a test suite that is independent of any specific implementation), (2) Assessed Means of Testing (MOT) (specific implementation of one or more ATS), (3) National Voluntary Laboratory Program (NVLAP) Accredited Test Laboratories (laboratories found competent to perform specified testing operations), (4) Conformance Tested OSI Products (a list of products that have gone through the Conformance Testing process), (5) Interoperability Test Suites (ITS) for OSI Products (public test scripts specified in terms of abstract services required for interoperability), (6) Reference Entities for Means of Testing Assessment(s) (publicly available implementations that can be used as a standard in evaluating other implementations or test systems), and (7) Interoperability Test and Registration Services (organizations which provide interoperability test results). These files are available for downloading and may be accessed via anonymous ftp over the Internet from IP address - 138.27.7.2. For any questions or comments dealing with this database please contact: Joint Interoperability Test Center - TCBB, Fort Huachuca, AZ 85613-7020, or email: C3A-TCB@huachuca-EMH2.army.mil. A Uniform Resource Locator (URL) is available as follows: <ftp://138.27.7.2/public/GOSIP.html>. This URL contains information on the above-mentioned files, plus additional information on related topics.

Interoperability test results have been available from OSINET (a consortium of U.S. vendors and users), however, OSINET ceased to exist as of December 31, 1995, and functions were transferred to OSIONE (a parent organization to OSINET, and an organization with worldwide membership). The secretariat for OSIONE is in Brazil; for more information, contact toledo@brisa.org.br.

APPENDIX B. List of Acronyms

| | |
|---------|--|
| ACSE | Association Control Service Element |
| ANSI | American National Standards Institute |
| API | Application Programming Interface |
| ARPA | Advanced Research Projects Agency |
| ASCII | American Standard Code for Information Interchange |
| ATM | Asynchronous Transfer Mode |
| CCITT | Consultative Committee for International Telegraphy and Telephony |
| CIDR | Classless Interdomain Routing |
| CLI | Call Level Interface |
| CLNP | Connectionless Network Protocol |
| CMC | Common Mail Call |
| CMIP | Common Management Information Protocol |
| CONS | Connection-oriented Network Service |
| CORBA | Common Object Request Broker Architecture |
| DARPA | Defense Advanced Research Projects Agency |
| DCA | Defense Communications Agency |
| DCE | Distributed Computing Environment |
| DIT | Directory Information Tree |
| DNS | Domain Name Server |
| DOD | Department of Defense |
| DQDB | Distributed Queue Dual Bus |
| DQP | Distributed Query Processors |
| DUA | Directory User Agent |
| DSA | Directory System Agent |
| EBCDIC | Extended Binary Coded Decimal Interchange Code |
| EDI | Electronic Data Interchange |
| EDIFACT | Electronic Data Interchange for Administration, Commerce and Transport |
| EDI-UA | Electronic Data Interchange-User Agent |
| EMTF | Electronic Mail Task Force |
| FTAM | File Transfer, Access and Management |
| FTP | File Transfer Protocol |
| GOSIP | Government Open Systems Interconnection Profile |
| GSA | General Services Administration |
| HTML | Hypertext Markup Language |
| HTTP | Hypertext Transfer Protocol |
| IAB | Internet Architecture Board |
| IACB | Internet Advisory Control Board |
| IEC | International Electrotechnical Commission |
| IEEE | Institute of Electrical and Electronics Engineers |
| IESG | Internet Engineering Steering Group |
| IETF | Internet Engineering Task Force |
| IIMC | ISO Internet Management Coexistence |
| IP | Internet Protocol |
| IPM-UA | Interpersonal Messaging User Agent |
| IPS | Internet Protocol Suite |
| IR | Information Retrieval |
| IRTF | Internet Research Task Force |
| ISDN | Integrated Services Digital Network |
| ISOC | Internet Society |
| ITS | Interoperability Test Suite |

ITU-T International Telecommunications Union-Telecommunications
 JITC Joint Interoperability Test Center
 JTC1 Joint Technical Committee 1
 LAN Local Area Network
 LLC Logical Link Control
 MAPI Microsoft Application Programming Interface
 MIB Management Information Base
 MIME Multi-media Mail Extensions
 MIT Massachusetts Institute of Technology
 MMS Manufacturing Messaging Specification
 MOT Means of Testing
 MTA Message Transfer Agent
 MTS Message Transfer System
 NASA National Aeronautics and Space Administration
 NIST National Institute of Standards and Technology
 NSF National Science Foundation
 NSFNET National Science Foundation Network
 OLTP On-line Transaction Processing
 OMB Office of Management and Budget
 OMG Object Management Group
 ORB Object Request Broker
 OSI Open Systems Interconnection
 PCMCIA Personal Computer Memory Card Interface Adapter
 PEM Privacy Enhanced Mail
 PII Protocol Independent Interface
 PMO Project Management Office
 RDA Remote Database Access
 RFC Request for Comments
 RFQ Request for Quotations
 RPC Remote Procedure Call
 SMI Structure of Management Information
 SMDS Switched Multimegabit Data Service
 SNMP Simple Network Management Protocol
 SMTP Simple Mail Transfer Protocol
 TCP Transmission Control Protocol
 TP Transaction Processing
 UA User Agent
 UDP User Datagram Protocol
 USENET User's Network
 VAN Value Added Network
 VT Virtual Terminal
 WAIS Wide Area Information Server
 WWW World Wide Web
 XTI X/Open Transport Interface

REFERENCES

1. CCITT Recommendation X.400 series (Red Book, 1984), Message Handling Systems
2. CCITT Recommendation X.400 series (Blue Book, 1988), Message Handling Systems
3. CCITT Recommendation X.435 - Message Handling Systems: EDI Messaging System (1991)
4. Information Processing Systems - Open Systems Interconnection - File Transfer Access and Management ISO 8571
5. Information Processing Systems - Open Systems Interconnection - Virtual Terminal Basic Class ISO 9041
6. Information Processing Systems - Open Systems Interconnection - Transaction Processing ISO 10026
7. Information Processing Systems - Open Systems Interconnection - Remote Database Access ISO 9579
8. ISO/IEC 9075 - Database Language SQL
9. CCITT Recommendation X.500 series The Directory, 1992
10. Information and Documentation-Search and Retrieve Application Protocol Specification for Open Systems Interconnection, ISO 10163
11. Industrial Automation Systems - Manufacturing Message Specification ISO 9506
12. Information Technology - Open Systems Interconnection - Systems Management ISO 10164
13. RFC 1006, 1987. Rose, M. "ISO Transport Services on Top of the TCP." Internet Request for Comments (May)
14. Federal Information Processing Standard 186, Digital Signature Standard, 1994, available from NTIS
15. Federal Information Processing Standard 46-2, Data Encryption Standard (DES), 1993
16. CCITT Recommendation F.435 - Message Handling Systems: EDI Messaging Service (1991)
17. Guidelines for the Evaluation of Electronic Data Interchange Products, NIST CSL/NSA Technical Report, August 1995
18. Asynchronous Transfer Mode Procurement and Usage Guide, NISTIR

4435, October 1994

19. GOSIP Conformance and Interoperation Testing and Registration, March 1991, Version 1.0, NISTIR 4594

