X/Open Guide

Guide to IPS - OSI Coexistence and Migration

X/Open Company, Ltd.

X/Open Guide

Guide to IPS - OSI Coexistence and Migration

ISBN: 1 872630 22 7
X/Open Document Number: XO/GUIDE/91/040

Set in Palatino by X/Open Company Ltd., U.K.
Printed by Maple Press, U.K.
Published by X/Open Company Ltd., U.K.

Any comments relating to the material contained in this document may be submitted to X/Open at:

X/Open Company Limited
Apex Plaza
Forbury Road
Reading
Berkshire, RG1 1AX
United Kingdom

or by Electronic Mail to:

XoSpecs@xopen.co.uk

Contents

Preface

This Document

This document is a Guide (see above). It is a comprehensive review of the issues of coexistence and migration between the IPS and OSI protocol suites. It is intended for:

- network component developers
- applications developers
- systems integrators
- end user representatives
- procurement individuals and organisations
- migrators and other network specialists

A broad overview of the issues is presented to aid planners and network designers in choosing the best migration and coexistence strategies for a particular situation. In addition, the more detailed information in separate sections of this guide is designed to aid software developers in creating the necessary tools, applications, and network software.

An overview of the OSI and IPS protocol suites is provided so that differences can be highlighted and solutions offered.

This guide is organised so that it can be read selectively, depending upon the needs of the reader. All readers should review the first three chapters and **Section 4.1**, **Overview**, for relevant information, because these chapters contain requisite background information for the remaining chapters. **Section 1.3**, **Scope and Audience**, explains which parts are relevant to a particlar audience.

This guide is organised as follows:

- **Chapter 1** is an introduction.
- **Chapter 2**, **Problem Statement**, outlines the problems of coexistence and migration by giving examples of several typical situations. These examples set the stage for the solutions presented in the remaining chapters.
- **Chapter 3**, **Migration Endpoints**, is an overview of the IPS and OSI protocols. This section provides the necessary background in the protocols to understand the coexistence and migration techniques that are presented in the next chapter. The IPS and OSI protocols are first described, then the important similarities and differences are examined.
- **Chapter 4**, **Techniques**, presents several coexistence and migration solutions. The first part of this chapter provides a high-level overview of the techniques to give a general understanding of how they work and what problems they are designed to solve. The remainder of the chapter describes each of these techniques in greater detail.

- **Chapter 5**, **Policies**, specifies some general policies for solving coexistence and migration problems. A general rationale for these policies is given.
- **Chapter 6**, **Tools**, describes the tools that are required to implement the policies described in the previous chapter. The particular characteristics of each tool are presented, and the basic requirements for their operation are specified. This information is useful for creating these tools, as well as for examining the suitability of available tools.
- **Chapter 7**, **Application to Scenarios**, describes the application of coexistence and migration techniques to the example situations described in **Chapter 2**, **Problem Statement**. For each scenario, migration and coexistence strategies are suggested, and the possible problems that accompany these strategies are examined.
- **Appendix A** is a glossary containing definitions of the terms and abbreviations used in this guide.
- There is an index at the end.

Trademarks

DECnet® is a registered trademark of Digital Equipment Corporation.

Ethernet® is a registered trademark of Xerox Corporation.

IBM® is a registered trademark of International Business Machines Corporation.

SAA™ is a trademark of International Business Machines.

SDLC is a product of International Business Machines.

SNA is a product of International Business Machines.

NetView® is a registered trademark of International Business Machines.

Token Ring is a product of International Business Machines.

UNIX® is a registered trademark of UNIX System Laboratories in the U.S.A. and other countries.

VAX® is a registered trademark of Digital Equipment Corporation.

VMS® is a registered trademark of Digital Equipment Corporation.

X/Open™ and the ''X''™ device are trademarks of X/Open Company Limited in the U.K. and other countries.

Referenced Documents

The following documents are references in this specification. Because of the large number of references, this section is divided into:

- ISO specifications
- Requests for Comments (RFCs)
- General references that are neither of these
- Other sources of information used in the compilation of this guide.

International Standards

The following standards, published by the International Organization for Standardization and International Electrotechnical Committee, are referenced:

ISO 4903
15-pole DTE / DCE Interface Connector and Contact Number Assignments.
International Standard 4903. 1989.

ISO 7498-1
Information Procession Systems — Open Systems Interconnection — Basic Reference Model.
International Standard 7498-1. 1984.

ISO 7498-2
Information Procession Systems — Open Systems Interconnection — Basic Reference Model — Part 2: Security Architecture.
International Standard 7498-1. 1984.

ISO 7776
High-Level Data Link Control Procedures — X.25 LAPB-Compatible DTE Data Link Layer Procedures.
International Standard 7776. 1986.

ISO 8072
Information Processing Systems — Open Systems Interconnection — Transport Service Definition.
International Standard 8072. 1986.

ISO 8073
Information Processing Systems — Open Systems Interconnection — Connection-oriented Transport Protocol Specification.
International Standard 8073. 1986.

ISO 8208
Information Processing Systems — Data Communications — X.25 Packet Level Protocol for Data Terminal Equipment.
International Standard 8208. 1987.

ISO 8326
Information Processing Systems — Open Systems Interconnection — Basic Connection Oriented Session Service Definition.
International Standard 8326. 1987.

ISO 8327
Information Processing Systems — Open Systems Interconnection — Basic Connection Oriented Session Protocol Specification.
International Standard 8327. 1987.

ISO 8348 ADD 2 Information Processing Systems — Data Communications — Network Serviced Definition — Addendum 2: Network Layer Addressing.
Addendum 2 to International Standard 8348. 1988.

ISO 8473 Information Processing Systems — Data Communications — Protocol for providing the Connectionless-mode Network Service.
International Standard 8437. 1988.

ISO 8571-x Information Processing Systems — Open Systems Interconnection — File Transfer, Access, and Management (FTAM) — Part x: ...
International Standard 8571. 1988.

ISO 8602 Information Processing Systems — Open Systems Interconnection — Protocol for Providing the Connectionless-Mode Transport Service.
International Standard 8602. 1989.

ISO 8649 Information Processing Systems — Open Systems Interconnection — Service Definition for the Association Control Service Element.
International Standard 8649. 1988.

ISO 8649 ADD 2 Information Processing Systems — Open Systems Interconnection — Service Definition for the Association Control Service Element — Addendum 2: Connectionless mode ACSE Service.
Final Text of Draft Addendum 2 to International Standard 8649. June 1990.

ISO 8650 Information Processing Systems — Open Systems Interconnection — Protocol Specification for the Association Control Service Element.
International Standard 8650. 1988.

ISO 8802-2 Information Processing Systems — Local Area Networks — Part 2: Logical Link Control.
International Standard 8802-2. May 1989.

ISO 8802-3 Information Processing Systems — Local Area Networks — Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA / CD).
International Standard 8802-3. 1990.

ISO 8802-4 Information Processing Systems — Local Area Networks — Part 4: Token-Passing Bus Access Method and Physical Layer Specifications.
International Standard 8802-4. 1990.

ISO 8802-5 Information Processing Systems — Local Area Networks — Part 5: Token Ring Access Method and Physical Layer Specifications.
Draft International Standard 8802-5. July 1990.

ISO 8802-7 Information Processing Systems — Local Area Networks — Part 7: Slotted Ring Access Method and Physical Layer Specifications.
Final text of Draft International Standard 8802-7. July 1989.

ISO 8822 Information Processing Systems — Open Systems Interconnection — Connection-Oriented Presentation Service Definition.
International Standard 8822. 1988.

ISO 8823 — Information Processing Systems — Open Systems Interconnection — Connection-Oriented Presentation Protocol Specification.
International Standard 8823. 1988.

ISO 8824 — Information Processing Systems — Open Systems Interconnection — Specification of Abstract Syntax Notation One (ASN.1).
International Standard 8824. 1990.

ISO 8825 — Information Processing Systems — Open Systems Interconnection — Specification of Basic Encoding Rules for Abstract Syntax Notation One (ASN.1).
International Standard 8825. 1990.

ISO 8878 — Information Processing Systems — Data Communications — Use of X.25 to Provide the OSI Connection-mode Network Service.
International Standard 8878. 1987.

ISO 8881 — Information Processing Systems — Data Communications — Use of the X.25 Packet Level Protocol in Local Area Networks.
International Standard 8881. 1889.

ISO 9040 — Information Processing Systems — Open Systems Interconnection — Virtual Terminal Service: Basic Class.
International Standard 9040. 1990

ISO 9041 — Information Processing Systems — Open Systems Interconnection — Virtual Terminal Protocol: Basic Class.
International Standard 9041. 1990

ISO 9066-x — Information Processing Systems — Text Communication — Reliable Transfer — Part x: ...
International Standard 9066-x. 1989.

ISO 9072-x — Information Processing Systems — Text Communication — Remote Operations — Part x: ...
International Standard 9072-x. 1989.

ISO 9545 — Information Processing Systems — Open Systems Interconnection — Application Layer Structure.
International Standard 9545. 1989.

ISO 9594-x — Information Processing Systems — Open Systems Interconnection — The Directory — Part x: ...
International Standard 9594-x. 1991.

ISO 9595 — Information Processing Systems — Open Systems Interconnection — Common Management Information Service Definition (X.710).
International Standard 9595. 1991.

ISO 9596 — Information Processing Systems — Open Systems Interconnection — Common Management Information Protocol Specification (X.711).
International Standard 9595. 1990.

ISO 9804 — Information Processing Systems — Open Systems Interconnection — Service Definition for the Commitment, Concurrency and Recovery

Service Element.
International Standard 9804. 1990.

ISO 9805 — Information Processing Systems — Open Systems Interconnection — Protocol Specification for the Commitment, Concurrency and Recovery Service Element.
International Standard 9805. 1990.

ISO 10021-x — Information Processing Systems — MOTIS — Part x: ...
International Standard 10021-x. 1990.

ISO 10040 — Information Processing Systems — Open Systems Interconnection — Systems Management Overview (X.701).
Draft International Standard 10040. September 1990.

ISO 10164-x — Information Processing Systems — Open Systems Interconnection — Systems Management — Part x: ...
Draft International Standard 10164-x. September 1990.

ISO 10165-x — Information Processing Systems — Open Systems Interconnection — Structure of Management Information — Part x: ...
Draft International Standard 10165-x. September 1990.

ISO TR 9577 — Information Technology — Protocol identification in the network layer.
ISO Technical Report 9577. 1990.

Requests for Comments

The following RFCs are referenced:

RFC-768 — Request for Comments 768, User Datagram Protocol.
Postel, Jon B. August 1980.

RFC-783 — Request for Comments 783, TFTP Protocol (Revision 2).
Sollins, K.R. June 1981.

RFC-791 — Request for Comments 791, Internet Protocol.
Postel, Jon B. September 1981.

RFC-792 — Request for Comments 792, Internet Control Message Protocol.
Postel, Jon B. September 1981.

RFC-793 — Request for Comments 793, Transmission Control Protocol.
Postel, Jon B. September 1981.

RFC-821 — Request for Comments 821, Simple Mail Transfer Protocol.
Postel, Jon B. August 1982.

RFC-822 — Request for Comments 822,
Standard for the Format of ARPA Internet Text Messages. Crocker, David H. August 1982.

RFC-826 — Request for Comments 826, An Ethernet Address Resolution Protocol.
Plumer, D. November 1982.

RFC-854 — Request for Comments 854, TELNET Protocol Specification.
Postel, Jon B. May 1983.

RFC-877 Request for Comments 877, Standard for the Transmission of IP Datagrams Over Public Data Networks.
Malis, A.G. December 1983.

RFC-894 Request for Comments 894, A Standard for the Transmission of IP Datagrams over Ethernet Networks.
Hornig, C. April 1984.

RFC-920 Request for Comments 920, Domain Requirements.
Postel, Jon B.; Reynolds, J.K. October 1984.

RFC-959 Request for Comments 959, File Transfer Protocol.
Postel, Jon B. October 1985.

RFC-987 Request for Comments 987, Mapping Between X.400 and RFC822.
Kille, Stephen E. June 1986.

RFC-1006 Request for Comments 1006, ISO Transport Services on top of the TCP.
Rose, Marshall T. and Cass, Dwight E. May 1987.

RFC-1026 Request for Comments 1026, Addendum to RFC 987 (Mapping Between X.400 and RFC-822).
Kille, S. September 1987.

RFC-1034 Request for Comments 1034, Domain names - concepts and facilities.
Mockapetris, P.V. November 1987.

RFC-1035 Request for Comments 1035, Domain names - implementation and specification.
Mockapetris, P.V. November 1987.

RFC-1042 Request for Comments 1042, A Standard for the Transmission of IP Datagrams over IEEE 802 Networks.
Postel, Jon B. and Reynolds, J. February 1988.

RFC-1069 Request for Comments 1069, Guidelines for the use of Internet-IP addresses in the ISO Connectionless-Mode Network Protocol.
Callon, Ross and Braun, Hans-Werner. February 1989.

RFC-1070 Request for Comments 1070, Use of the DARPA / NSF Internet as a Subnet for Experimentation with the OSI Network Layer.
Hagans, Robert A.; Hall, Nancy E.; and Rose, Marshall T. February 1989.

RFC-1086 Request for Comments 1086, ISO-TPO bridge between TCP and X.25.
Onions, Julian P. and Rose, Marshall T. December 1988.

RFC-1095 Request for Comments 1095, Common Management Information Services and Protocol over TCP / IP (CMOT).
Warrier, Unni and Besaw, Larry. April 1989.

RFC-1098 Request For Comments 1098, A Simple Network Management Protocol.
Case, Jeffery D.; Fedor, Mark S.; Schoffstall, Martin L.; and Davin, James R. April 1989.

RFC-1122
Request for Comments 1122, Requirements for Internet Hosts - Communication Layers
Braden, R. October 1989.

RFC-1123
Request for Comments 1123, Requirements for Internet Hosts - Application and Support.
Braden, R. October 1989.

RFC-1155
Request for Comments 1155, Structure and identification of management information for TCP/IP-based internets.
Rose, M.T.; McCloghrie, K. May 1990.

RFC-1156
Request for Comments 1156, Management Information Base for network management of TCP/IP-based internets.
Rose, M.T.; McCloghrie, K. May 1990.

RFC-1157
Request for Comments 1157, Simple Network Management Protocol (SNMP).
Case, J.D.; Fedor, M.; Schoffstall, M.L.; Davin, C. May 1990.

RFC-1188
Request for Comments 1188, Proposed standard for the transmission of IP datagrams over FDDI neworks./fP
Katz, D. October 1990.

RFC-1189
Request for Comments 1189, Common Management Information Services and Protocols for the Internet (CMOT) and (CMIP).
Warrier, U.S.; Besaw, L.; LaBarre, L.; Handspicker, B.D. October 1990.

General References

The following general publications are referenced.

AT&T
Network Programmer's Guide - UNIX System V/386.
AT&T. 1988.

CCTA
U.K. Government OSI Profile.
Version 3.0. Central Computer and Telecommunications Agency (CCTA). January 1988.

EC
Council Decision of 22 Decmeber 1986 on standardisation in the field of information technology and telecommunications.
European Community. EC Directive 87/95/EEC, Official Journal of the European Communies, L36/31.

Groenbaek
Conversion between the TCP and ISO Transport Protocols as a Method of Achieving Interoperability Between Data Communications Systems.
Groenbaek, I. IEEE Journal on Selected Areas in Communications, SAC-4, March 1986, pp. 288-96.

Tanenbaum
Computer Networks.
2nd ed. Tanenbaum, A. Prentice-Hall. 1988.

NIST
The Government Open Systems Interconnection Profile.
FIPS PUB 146. U.S. Department of Commerce, National Institute of Standards & Technology (NIST). August 1988.

NTIS	Protocol Interoperability between DDN and ISO Protocols. U.S. Department of Commerce, National Technical Information Service (NTIS). August 1988.
XGIPS	Guide to the Internet Protocol Suite X/Open Company Ltd. 1991.
XOSIP	Comparison Study of OSI Profiles. X/Open Company Ltd. 1990.

General Sources of Information

The following publications were used in compiling the information presented in this guide, but are not explicitly referenced within it:

Bucciarelli et al.
Connectionless Services in the OSI Reference Model.
Bucciarelli, P. and Caneschi, F. Proceedings of the ICCC, 1984, pp. 564-569.

Burg et al.
Of Local Networks, Protocols, and the OSI Reference Model.
Burg, F. M.; Chen, C. T.; and Folts, H. C. Data Communication, November 1984, pp. 129-150.

Caneschi et al.
Standardizing the Presentation Layer: Why and What.
Proceedings of the Seventh International Conference on Distributed Computer Caneschi, F. and Merilli. Systems, IEEE, 1987, pp. 35-39.

Caneschi Hints for the Interpretation of the ISO Session Layer.
Caneschi, Fausto. Computer Communication Review, 16, August 1986, pp. 34-72.

Case et al. Network Management and the Design of SNMP.
Case, Jeffrey D.; Davin, James R.; Fedor, Mark S.; and Schoffstall, Martin L. ConneXions, 3, March 1989, pp 22-6.

Chong Software Development and Implementation of NBS Class-4 Transport Protocol.
Chong, H. Y. Computer Networks and ISDN Systems, 11, May 1986, pp. 353-65.

Cohan et al.
The ISO Reference Model and Other Protocol Architectures.
Cohan, Danny and Postel, John B. Proceedings of the IFIP Congress, Paris, France, 1983.

Cole et al. OSI Transport Protocol—User Experience.
Cole, Robert and Lloyd, Peter. Open Systems 86, Online Publications, 1986, pp. 33-43.

Comer Internetworking with TCP/IP: Principles, Protocols, and Architecture.
Comer, D. Prentice-Hall. 1988.

Conard Services and Protocols of the Data Link Layer.
Conard, J. W. Proceedings of the IEEE, 71, December 1983, pp. 1378-83.

Day et al. The OSI Reference Model.
Day, John D. and Zimmermann, Hubert. Proceedings of the IEEE, 71, December

1983, pp. 1334-40.

Emmons et al.
OSI Session Layer: Services and Protocols.
Emmons, W. F. and Chandler, A. S. Proceedings of the IEEE, 71, December 1983, pp. 1397-400.

Frankel et al.
An Overview of the Army /DARPA Distributed Communications and Processing Experiment.
Frankel, Michael S.; Graff, Charles J.; Dworkin, Larry U.; Klein, Theodore J.; and desJardins, Richard L. IEEE Journal on Selected Areas in Communications, SAC-4, March 1986, pp. 207-15.

Henken Mapping of X.400 and RFC822 Addresses.
Henken, G. Computer Networks and ISDN Systems, 13, 1987, pp. 161-64.

Henshall et al.
OSI Explained. End to End Computer Communication Standards.
Henshall, J. and Shaw, A. Ellis Horwood. 1988.

Klerer The OSI Management Architecture: an Overview.
Klerer, S. M. IEEE Network Magazine, 2, March 1988, pp. 20-9.

Knightson Standards for Open Systems Interconnection.
Knightson, K. G. Mc-Graw Hill. 1988.

Leiner et al.
The DARPA Internet Protocol Suite.
Leiner, Barry M.; Cole, Robert; Postel, John; and Mills, David. IEEE Communications Magazine, 23 March 1985, pp. 29-34.

Lew at al. Getting there from here: Mapping from TCP/IP to OSI.
Lew, H. Kim; and Jung, Cyndi. Data Communication, August 1988, pp. 161-175.

Linnington
The Virtual Filestore Concept.
Linnington, P. F. Computer Networks, 8, 1984, pp 13-16.

McClelland
Services and Protocols of the Physical Layer.
McClelland, F. M. Proceedings of the IEEE, 71, December 1983, pp. 1372-77.

McCoy Request for Comments 1008, Implementation Guide for the ISO Transport Protocol.
RFC-1008. McCoy, Wayne. June 1987.

McLeod-Reisig et al.
ISO Virtual Terminal Protocol and its Relationship to Mil-Std TELNET.
McLeod-Reisig, S. E. and Huber, K. Proceedings of the Computer Networking Symposium, IEEE, 1986, pp. 110-119.

MITRE The Department of Defense Open Systems Interconnection (OSI) Implementation Strategy.
The MITRE Corporation. May 1988.

Rose Request for Comments 1085, ISO Presentation Services on top of TCP/IP-based Internets. RFC-1085. Rose, Marshall T. December 1988.

Rose Transition and Coexistence Strategies for TCP/IP to OSI.
Rose, M. T. IEEE Journal on Selected Areas in Communications, 8, January 1990, pp. 57-66.

Rose The ISO Development Environment: User's Manual. 5.0 edition. Rose, Marshall T. The Wollongong Group. March 1989.

Rose The Open Book: A Practical Perspective on OSI. Rose, Marshall T. Prentice-Hall. 1990.

Rose et al OSI Transport Services on top of the TCP.
Rose, Marshall T. and Cass, Dwight E. Computer Networks and ISDN Systems, 12, 1986.

Stallings Data and Computer Communications.
2nd ed. Stallings, W. Macmillan. 1988.

Chapter 1

Introduction

The International Organization for Standardization (ISO) has been working on a suite of networking protocols since the early 1970s. The resulting set of protocols is referred to as the Open Systems Interconnection (OSI) suite of protocols. These protocols are defined in a set of published standards that have evolved and been refined over the years through the cooperation of standards organisations from each participating country. Since the 1988 publication of the OSI suite of protocols, many government and private organisations have started requiring all newly installed networks to use the OSI protocols. Because of the widespread acceptance of this networking standard, many people believe that OSI systems will eventually dominate, and perhaps even eliminate, other networks. Regardless, it is clear that there will be a lengthy period of migration during which many new OSI networks will be installed and many older networks will be converted to OSI. During this period of migration, OSI networks must coexist with and, as much as possible, work harmoniously with other networks.

This guide deals with the specific issues of coexistence and migration between networks based upon the Internet Protocol Suite (IPS) and networks based upon the OSI suite. IPS networks (also referred to as TCP/IP networks) currently enjoy wide use throughout the world, with over 200,000 hosts and over a million users. IPS networks have given reliable service for many useful applications over a period of many years. Nevertheless, it is expected that many, if not most, of these networks will eventually be replaced by OSI networks. In addition, IPS networks that remain in service need to coexist with the emerging ISO networks.

This guide is designed to ease the transition between IPS and OSI networks in a number of ways:

- by defining problems of coexistence between IPS and OSI
- by defining problems in migration from IPS to OSI
- by describing techniques and tools which enable coexistence between the two protocol suites and facilitate migration
- by providing guidelines for the development of policies to solve these problems and achieve successful coexistence and migration

This guide is designed for a wide audience and contains broad overviews of the issues and their solutions, as well as detailed explanations for implementors and software developers. Readers should consult **Section 1.3**. **Scope and Audience**, for advice on how to approach this guide.

The choice of appropriate solutions often involves trade-offs among many different factors. In addition to the technical issues described in the guide, there are often considerations of cost of installation, and cost to the user community as a result of disruptive changes in the computer system. This guide attempts to present all of the major issues that must be considered before implementing a coexistence and migration plan.

The scenarios described in **Chapter 2**, **Problem Statement**, reflect the widespread presence of proprietary protocols in real networks. This guide addresses only IPS migration to OSI. It does not specifically address proprietary networking standards, although the techniques described in **Chapter 4**, **Techniques** may be equally applicable. Users should look to the suppliers of proprietary protocols for the provision of appropriate tools for migration and coexistence.

1.1 MOTIVATION FOR MIGRATION TO OSI

There is a range of possible reasons for an organisation's decision to migrate its networks to use OSI protocols. They can be categorised as either Standardisation Motivations or Cost-Benefit Motivations. The former seek to achieve interoperability among heterogeneous systems and protect an organisation's investment in technology, the latter seek to gain some commercial or operational benefit from the migration. They are dealt with individually below, along with a discussion of reasons which currently inhibit organisations from migrating to OSI protocols.

1.1.1 Standardisation Motivations

Facilitating widespread interoperability of computer systems and enabling the procurement of vendor-neutral networking solutions were major original motivations for the development of the OSI protocol suite. Since then, various organisations have mandated the adoption of these standards.

Government Standards

ISO is composed of standards organisations from the member countries, so it is not surprising that in many countries, public organisations are required to use OSI protocols. Standards established by these organisations often carry the force of law in government purchases, and many private organisations routinely follow these standards. The U.S. Government has published a Federal˜Information˜Processing˜Standard (FIPS 146) that now applies to all its computer equipment purchases, requiring all computer networking products to comply with OSI where feasible. In addition, the European˜Community has published a directive (87/95/EEC) which relates to open government procurement in general and as such applies to procurement of computer systems and networks. (In this context open procurement means specifying the function required rather than the form of the solution when requesting tenders)

FIPS 146 includes a standard known as the Government OSI Profile (GOSIP). Similar standards have been established, or are in the process of being established, in several European countries. A European GOSIP is under development. Such profiles seek to ensure that all networking equipment interoperates by specifying that all interconnected networks use a common set of features. The referenced **Comparison Study of OSI Profils** (XOSIP) document describes and compares the existing profile standards.

Industry Standardisation

Companies purchasing and installing computer networks are concerned about protecting their investments. Many companies are expected to require the use of the OSI protocol suite, as defined by an interconnection profile such as GOSIP. By choosing OSI products, the company guarantees that equipment and software is available from a variety of sources. By using a profile, such as one of the GOSIPs, and specifying conformance and interoperability testing requirements, a user can ensure that the selected components work together successfully. Because of the broad-based support for this protocol suite, support and upgrades are likely to be available for many years.

Many vendors incorporate OSI as their strategic method for interoperability between heterogeneous systems. They are also integrating OSI protocols into their network architectures. A large proportion of vendor development effort is currently being directed at implementing the OSI protocols and applications which utilise them.

1.1.2 Cost-Benefit Motivation

For many organisations, especially commercial companies, migration can proceed only if supported by a clear cost-benefit justification. Currently, there is little real experience of implementing OSI networks on which to base a business case, however there are a number of factors which must be considered. These are discussed individually here.

Functionally Rich Applications

Many people may choose to install OSI networks because of the rich set of application services that are currently available. One of the most prominent of these is the Message Handling System (MHS). This set of services for electronic mail allows messages to be sent composed of multiple body parts of different types. For example, a message can contain ASCII text, digital voice, facsimile, or any number of other defined types. This differs from most current electronic mail systems, which allow only a single message type to be sent. MHS users can also send text in a variety of character sets for different languages. This multi-lingual feature alone may be sufficient incentive to cause many people to install OSI-based networks.

In addition to MHS, other applications such as Directory Services (DS), File Transfer, Access and Management (FTAM), and Virtual Terminal (VT) services have been defined. Standards in other areas are also underway. For example, a Transaction Processing standard is being defined that can serve as the basis for other applications, such as airline reservations. Typically, these applications offer important features that make OSI networks very attractive.

Existing and Developing Infrastructure

Most European WANs are already OSI-based and commercial OSI-based networks are available in the U.S. The presence of these networks encourages the use of the complete OSI stack, because fully homogeneous networks are easier to create and maintain.

Education, Research, and Training

Recent university graduates frequently have extensive training in the OSI protocol suites. This increasing pool of OSI talent makes it easier to develop and maintain OSI networks and applications.

Network Coexistence

Often, an organisation may have many different networks serving different groups of computers. For example, SNA networks are widely employed among large main-frame computers, while IPS networks are widely used among workstations and by scientific and educational users. These networks, and many others, frequently coexist in the same organisation. As a result, many organisations have electronic mail systems that are not conveniently interconnected, files and databases that are not easily shared, and networks that are independently maintained and administered. Furthermore, for interconnected networks such as electronic mail, gateway hardware and software must be installed and maintained.

One of the primary reasons for standardisation in government and industry is to reduce the proliferation of incompatible networks. In the previous example, where IPS and SNA networks coexist, the only way to operate under one homogeneous network may be for both networks to convert to OSI. Other solutions requiring extra gateway and networking software may be unworkable or unavailable. While network conversions are often costly, the long-term benefits, derived from increased connectivity and decreased administrative and maintenance costs, will cause many organisations to install OSI networks. In many cases, the original wiring and networking hardware can be retained, but most of the networking software must be replaced.

OSI networks can also be used to interconnect other networks, either by providing application gateways or as a direct network connection. These mechanisms are discussed in **Chapter 4**, **Techniques**.

1.1.3 Factors Hindering Migration

Whilst there is a growing recognition that adoption of OSI protocols is both desirable and inevitable, there are still a number of factors which are inhibiting organisations from proceeding. Users and system vendors can cite commercial and technical reasons for not considering migrating networks to the OSI protocol suite. The networking scenarios described in **Chapter 2**, **Problem Statement**, highlight a number of such concerns voiced by user organisations. Some of these are valid, however many are founded on out-of-date information or are the result of natural suspicion of technology that it is in some respects ‘‘leading edge’’. In addition, **Chapter 3**, **Migration Endpoints**, and **Chapter 7**, **Application to Scenarios**, highlight the lack of standards and products in some areas.

It is clear that migration of networks to OSI will only become commonplace when users are confident that OSI networks can deliver the key functionality that they require, with appropriate levels of performance and reliability. The inhibiting factors must be addressed by technical developments; by the existence of real commercial OSI networks which demonstrate the practical possibilities; and by the availability of independent performance and interoperability assessments. This section discusses some of the major inhibiting factors in more detail.

Standards and Product Availability

For many users, OSI protocol standards are still regarded as being incomplete or too volatile to be implemented safely. This perspective is a result perhaps of publicity during the 1980s which suggested that OSI networks were ‘‘just around the corner’’. The situation now is quite different, with stable international standards in place for the OSI transport service and for most of the primary applications. **Chapter 3**, **Migration Endpoints**, discusses the status of the ISO standards for each layer of the OSI reference model. There are some important exceptions; the virtual terminal standards have been finalised but have failed to gain widespread acceptance by suppliers, and the standards to support FTAM-based transparent file access are missing.

There is a similar problem of outdated perspective concerning the availability of OSI-based products and the commitment of suppliers to this market. It is clear that a major proportion of vendors' resources is being directed at producing and marketing OSI networking products. Buyers' guides such as that published by the UK government procurement organisation demonstrate availability of products from a wide range of suppliers, with firm commitments to expand their product ranges to encompass the latest standards available.

Interoperability Problems

Users are aware of the potential incompatibilities implied by the many options available at various levels in the OSI protocol stack. In particular, companies with bases in both North America and in Europe are concerned about the incompatible standards adopted in those regions at the network and transport layers. The former concern is being addressed actively by the development of various profiles and by initiatives aimed at branding implementations which conform to certain profiles or have undergone interoperability testing. The latter is more problematic and requires a clear statement about how the two transport stacks are to interoperate.

Performance

The OSI protocol suite is perceived as being ‘‘heavy-weight’’ in comparison with existing inter-networking protocols such as IPS. Such comparisons are not always valid, for example, the two transport stacks can be meaningfully compared, however above that, OSI standardises a number of functions which are left up to individual applications in the case of IPS. In addition, the OSI protocol suite is more ambitious in some of its objectives which tends to make the protocols more complex and increases processing and communications overheads. (For example, OSI's global addressing scheme leads to a much increased address length, something which IPS is just beginning to address.)

Aspects of performance which receive particular attention are usage of memory, processor and communications bandwidth. It is beyond the scope of this guide to present a comparison of the performance of the two protocol suites. So far, most comparisons have either been theoretical or have been based upon prototype implementations, neither of which is a good basis for making a comparison. With the increasing availability of OSI products, it is only now becoming possible to perform meaningful practical comparisons. Current activity in the OSI market is concentrating on making products available, increasing competition is likely to focus development efforts on performance aspects in the near future.

Lack of Experience

Lack of experience is perhaps the most inhibiting user concern, and must be dealt with even when users are confident that a migration to OSI protocols is both feasible and desirable. Apart from the obvious skills required to install and administer and operate an OSI network there is the understandable suspicion of an unknown technology. This suspicion covers all of the concerns discussed above, plus, in addition, concerns about hidden costs associated with the installation of OSI.

It is, perhaps, in this area that this guide is most able to advance the cause of OSI, by demonstrating the range of techniques available to an organisation for introducing OSI protocols into its current networks, and by providing practical advice on how to develop policies for successful coexistence and migration. The most effective way of overcoming such concerns is by providing practical techniques for the incremental introduction of OSI protocols, allowing the organisation to develop skills and gain experience without compromising operationally critical systems or applications.

1.2 MIGRATION AND COEXISTENCE

Migration and coexistence are separate but interrelated issues. At any particular installation where IPS networks are being replaced by OSI networks, there is usually a period of migration. The IPS software is not removed and the equivalent OSI software installed in one go, instead the IPS and OSI networks and applications coexist throughout the migration process. For example, the following is just one of many possible scenarios of how migration from IPS to OSI might occur:

1. A standard IPS network exists at a location to support a networked CAD/CAM application and Internet-style corporate mail.
2. An OSI network is installed to support a new CAD/CAM application. However, the Internet-style mail is retained for compatibility with the rest of the corporate mail system. During this stage of the migration there is a period of coexistence. Somehow, users on the OSI network must be able to exchange mail with people on the IPS network. Problems such as this must be handled during coexistence.
3. Eventually, the rest of the corporation converts to a standard OSI mail system (MHS) so that few or no IPS applications or network software remain.

In practice, networks are much more complex than this. They usually support many applications and are often interconnected with many other networks, frequently extending beyond the local installation. The highly interconnected nature of most networks virtually guarantees that there is a period of coexistence during any migration from IPS to OSI. Usually, there are a large number of coexistence problems with an even larger number of possible solutions. Consequently, at some installations, the migration plan may consist of several stages and may employ a variety of coexistence techniques.

1.3 SCOPE AND AUDIENCE

For detailed descriptions of the OSI and IPS protocol suites, consult the relevant specifications. The companion document to this guide, **X/Open Guide to the Internet Protocol Suite** (XGIPS) contains a complete description of current common practice for IPS implementations. The appropriate **Requests for Comments** (RFCs) listed in **Section 3.3**, **IPS - The Internet Protocol Suite**, contain the authoritative definitions of IPS networks. For the OSI suite, the OSI documents listed in **Section 3.2**, **OSI - The Open Systems Interconnection Standard**, contain the authoritative definitions. Current common practice for OSI protocols is defined by the various functional standards, for example, ISPs, GOSIPs, and others. For a definition of functional standards, see the referenced XOSIP document.

The following list describes the targetted audiences and identifies the relevant parts of this guide for each audience.

Network component developers:
Some coexistence solutions require specially designed network components. In particular, the application gateway, transport relay, and network service tunnel techniques require software that allows interconnection of applications on different networks. Users rarely interact directly with these network components and, if everything works correctly, may even be unaware of their existence. The detailed explanations of these components in **Chapter 4**, **Techniques**, describe how they work and outline the issues that must be addressed in their design.

Applications developers:
Some of the migration and coexistence solutions require the creation of new applications. The dual stack and common API solutions may use a universal application, which allows the same interface to be presented to the user, regardless of the underlying network. The hybrid network solution uses an application that runs over a network of a different type. **Chapter 4**, **Techniques**, has detailed explanations of how these applications must be written to operate with the two types of network interfaces. Differences between the two networks are highlighted so that applications developers can make appropriate decisions in dealing with these differences.

Systems integrators:
Systems integrators are the people who put together a network. This includes installing and configuring hardware and software to make a network function. Networks are notoriously difficult to set up correctly and keep running. Any given network may have several different kinds of hardware and software attached to it, and each endpoint may require separate administration. These difficulties are compounded when IPS and OSI networks coexist. In addition to administering two separate networks, systems integrators must install and maintain additional software and hardware to allow these networks to coexist. **Chapter 5**, **Policies**, **Chapter 6**, **Tools**, and **Chapter 7**, **Application to Scenarios**, are directly relevant to the needs of the systems integrator. **Chapter 6**, **Tools**, is especially relevant, because it describes the tools that must be used for integrating IPS and OSI networks.

User representatives:
It is generally the users that experience the greatest disruption during periods of migration or coexistence. For this reason, user representatives should take an active

part in all network planning. **Chapter 2**, **ProblemStatement**, gives an overview of some typical migration and coexistence problems for a useful perspective on the potential problems. The applications sections of **Chapter 3**, **Migration Endpoints**, should be at least scanned to learn about the features of each relevant network application. The overview section of **Chapter 4**, **Techniques**, should be reviewed to understand the possible solutions. Finally, **Chapter 5**, **Policies**, and **Chapter 7**, **Application to Scenarios**, give some general policies and show how they can be applied. The relative merits of each approach are given so that informed decisions can be made to minimise the effect on users.

Procurement individuals and organisations:

In some government and private organisations, OSI compliance is often required for new network purchases. In addition, procurement individuals often must coordinate purchasing policies across several departments. **Chapter 5**, **Policies**, and **Chapter 7**, **Application to Scenarios**, contain information on networking policies and their implementation.

Migrators and other network specialists:

The problems of migration and coexistence between IPS and OSI systems are sufficiently complex that it is often necessary to consult specialists in migration and network integration. These specialists, if they do not already have a thorough background in IPS and OSI networks, should read the appropriate sections of the first three chapters. The remaining chapters of the book contain detailed information essential to anyone needing a full understanding of all aspects of coexistence and migration.

1.4 CONVENTIONS AND TERMINOLOGY

Throughout this guide the term OSI refers to all layers of the OSI protocol suite as they are defined in **Chapter 3**, **Migration Endpoints**. The term \s-1PS refers to all parts of the TCP/IP protocol suite and the associated applications defined in the appropriate RFCs listed in **Chapter 2**, **Problem Statement**, and in the XGIPS document (X/Open, 1990). Wherever possible, this guide derives its terminology from these primary specification documents. Definitions for these terms are included in **Appendix A**, **Glossary**.

Chapter 3, **Migration Endpoints**, describes the OSI and IPS protocols to aid in the understanding of coexistence and migration strategies presented in **Chapter 4**, **Techniques**. This guide is not intended as a complete definition of these protocols; for that, the relevant specifications should be consulted.

Chapter 2

Problem Statement

This chapter describes five scenarios where IPS protocols are used and where the potential exists for migrating network applications and the underlying network technology to one based on OSI protocols. The scenarios describe the problems to which the migration and coexistence techniques, tools and policies (described later in this guide) can be applied. This provides a basis for **Chapter 7**, **Application to Scenarios**, which demonstrates their application and discusses problems that may be encountered in applying them in real situations.

The type of problems that may be encountered in situations where IPS-OSI coexistence or migration is being planned may be divided into three categories:

Provision of the required functions:

When migrating a network to the OSI protocols, all *key* user and operational functions must be preserved. In addition, it must be possible to manage the migration from IPS to OSI such that functions are maintained during the period of coexistence. In this context, *functions* includes routing, security and network management.

In a real migration, it is almost certain that some functions are either lost or changed significantly. It is the task of a migration policy to identify where this is acceptable within the context of a specific organisation and network.

Provision of the required performance:

The replacement of IPS by OSI and, more importantly, the coexistence of OSI and IPS, must not compromise network performance. In this context, performance includes response time, bandwidth, availability and resilience.

Strategic Issues:

The migration strategy is made more difficult if it addresses only technological issues. OSI must be *sold* to management, sometimes in situations where the communications infrastructure is business driven and the technical issues of OSI in relation to IPS are not fully understood.

Requirements and concerns that come under the first two categories are easier for users to express objectively. The third class of problems is more subjective but can have a very significant effect on the speed with which IPS networking facilities are migrated to OSI.

2.1 DEVELOPMENT OF THE SCENARIOS

The scenarios were developed using information collected from a number of users of UNIX systems and IPS protocols. Some general underlying information was collected by means of a questionnaire; then some of the completed questionnaires were followed up with telephone and face to face interviews.

In selecting the companies the aim was to achieve a good geographic and industry coverage. Companies selected were, in the main, situated in Europe and the U.S.A., with a small Japanese sample. The industry coverage included oil, finance and automobile, with some general commercial and public sector organisations.

A total of 54 companies were selected and contacted by phone. Of these, 42 were identified as suitable and willing to receive the questionnaire. Of the companies that did receive the questionnaire, 24 completed and returned it. The quality and completeness of the replies was variable, but in general provided adequate information to build up the scenarios and identify those organisations willing to provide further information.

Face to face interviews were carried out with six of the organisations. Their selection was based upon geographical location and usage of IPS within the organisation. Of the six, two were based in the U.K. and four in the U.S.A.

Each scenario describes a particular class of problems, hence each scenario is a composite, based primarily on one organisation but borrowing ideas from other similar organisations. The composite nature of the scenarios also helps to preserve the confidentiality promised to the organisations that contributed to the study. For the same reason some fairly arbitrary changes have been made to the information provided by companies where these do not change the underlying networking requirements. In addition to the purely technical information described in each scenario, additional information is given (such as Information Technology (IT) strategy, or IT trends in the industry) in order to present a framework within which an organisation's networking requirements can be viewed.

2.2 OVERVIEW

The five scenarios described here cover a representative range of IPS environments:

Scenario 1 **Basic IPS Networking**

This scenario describes a very basic IPS environment. Consequently, it must be addressed by any IPS-OSI migration and coexistence strategy. This scenario is fundamental to the whole migration process - the communications facilities described in the completed questionnaires were in many instances either similar to the facilities described in this scenario or were more complex versions of them. Thus, the techniques and policies and tools described in this guide must address the requirements of this scenario as a minimum.

Scenario 2 **IPS as Common Interworking Protocols**

This scenario describes an environment where IPS is used as a means of communicating between different systems. The completed questionnaires showed that IPS is often seen as both vendor independent and likely to be supported by many different vendors. Consequently it is seen as a good candidate for a *common language*, often in situations in which UNIX is not used at all. This role of IPS is likely to be taken over by OSI.

Scenario 3 **Engineering Workstation Environment**

This scenario describes an environment where the improvement in cost and performance of computing systems has led to the widespread introduction of workstations and the requirement to interconnect them. This type of requirement is typical of an engineering environment and appeared in many of the questionnaires. It is characterised by a distributed processing environment and small domains with local administrative control. As the cost of workstations has decreased, they are appearing in many new situations replacing tasks previously done on powerful minicomputers; where these systems are networked, IPS protocols are generally used.

In this kind of environment users often have a big say in the type of facilities they use and there is currently little outside pressure for migration to OSI. Some positive advantage to the end users themselves must be demonstrated if migration is to occur.

Scenario 4 **Summary Operational Requirement**

This scenario takes the form of a summary of an Operational Requirement (OR) for the expansion of an organisation's networking infrastructure, from a small number of UNIX systems with minimal IPS networking, to a much larger network based on OSI protocols. This scenario is primarily intended to cover various network performance issues. This type of information was not requested in the questionnaire, since it cannot be easily provided by users. However, the targets specified in the scenario are typical of ORs issued by public or large private sector organisations. The format of this scenario is fundamentally different from the others. The reason for this is to provide a natural situation in which to express network parameters, such as transmission delay and availability, and show how they are affected by the tools and techniques discussed in this guide.

Scenario 5 **User Concerns**

This scenario describes a hierarchically structured network with heavy interoperation requirements between systems. The main problems arise from managing the migration and demonstrating the benefits of OSI. This scenario has some technical issues, but is also an amalgamation of the subjective worries and prejudices expressed by a number of people in the different organisations contacted.

2.3 SCENARIO 1 - BASIC IPS NETWORKING

2.3.1 Introduction

This scenario is intended to present the fundamental networking facilities one would find in most IPS based networks. Before more difficult problems can be solved, **Chapter 7**, **Application to Scenarios**, must demonstrate that the tools and techniques discussed in this guide are able to:

- duplicate, using OSI, the functions currently in use
- manage the migration with little impact on users

IPS is confined to a number of Ethernet segments. The IPS network is large, but it has a simple topology and traffic levels are comparatively low. Connections between Ethernet segments are exclusively via MAC-level bridges.

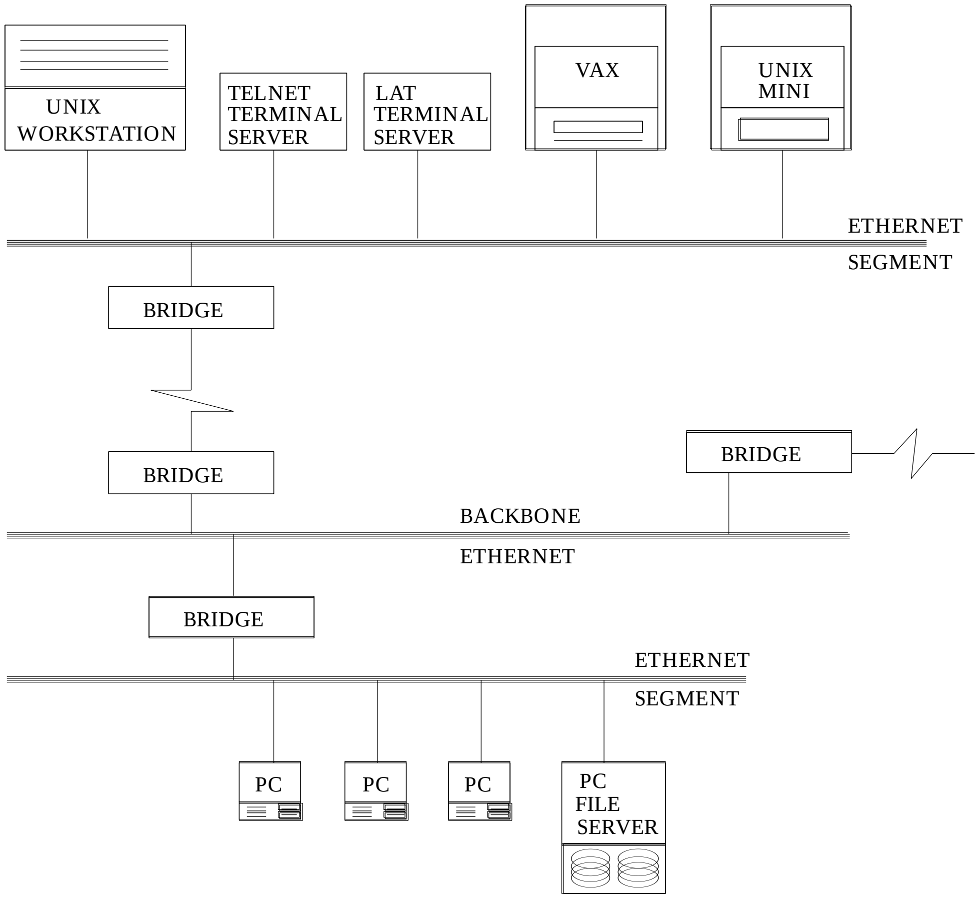

Number of Ethernet Segments	= 100
Total Number of Devices	= 2500
Devices per Segment	= 20-50

Scenario 1 Basic IPS Network

2.3.2 Type of Organisation

This scenario describes the communications requirements of a U.S.A.-based multinational company involved in the manufacture of a range of chemical and pharmaceutical products. The company uses a range of computer equipment at a number of sites. It has offices, manufacturing plants, distribution warehouses and so on, at a number of sites both in and outside the U.S.A. This scenario concentrates on the requirements arising at one of the major sites of the company where heavy usage is made of UNIX-based systems.

For most of its life the company has had a centralised policy. About five years ago, though, top management decided to break the company into separate business administrative units, each managed by a vice president with profit and loss responsibility. This has produced a complex hierarchical structure. The company is divided into a small number of Business Groups, each in turn divided into a number of Business Units. A number of corporate functions such as Corporate Planning, Marketing and Central Sales are still held centrally and support all the Business Units.

The Business Groups are fairly autonomous and vary in the type of activity they are involved in (ranging from heavy chemical engineering to consumer products).

2.3.3 Background

In recent years technology has become increasingly more important. The corporate IT needs have traditionally been met primarily by IBM and, to a lesser degree, Digital. This means that the communications protocols most widely used are SNA and DECnet. At the company's headquarters there is a large IBM data centre, with 60-70 SNA links to about 200 locations throughout the world.

UNIX has developed more recently in the Scientific and Engineering areas. The improvement in price and performance of computers and their ease of use have reached the point where most applications can be executed on workstations. Many workstations are specifically designed to accomplish functions with integrated hardware, software, and communications. The evolution of user needs and local network technology means that most operational functions are accomplished in LAN environments.

The UNIX systems are networked via Ethernet and it is here that extensive use of IPS protocols is made. However, SNA communications exceeds everything else in terms of volume in all types of traffic (electronic mail, file transfer, terminal access). An effect of this is that corporate strategy and communications infrastructure investment tends to address more fully the needs of the IBM-based systems. A major issue is compatibility with SNA and, to a lesser degree, DECnet.

There are two sites that make extensive use of IPS protocols. Both are located in the same area and are linked with fast MAC-level bridges. The topology of the two sites is very similar and consists of a backbone Ethernet connecting a number of Ethernet spurs. The types of applications that run at these locations are:

CAD
modelling
Simulation
some Office Automation applications (database, spreadsheet)
a small amount of software development.

At the largest of these sites there are about 400 IT staff, of which 90% are using proprietary systems skills.

IT Strategy

A Corporate Information Systems Unit is responsible for the development and management of the corporate IT infrastructure, and such corporate activities as the establishment of corporate standards, new product evaluation, participation in external standards activities, supplier relationships and high-level consulting.

On a day-to-day basis, responsibility is greatly devolved, with the various business groups and units being responsible for all of their IT requirements. Business managers are responsible for selecting the mix of resources that yields the maximum value for their operations. The result of this is that the company's IT needs are very much business-driven. Software packages are used to support virtually all generic business functions. Custom applications are used only when they offer a clear competitive advantage to business.

There is strong corporate guidance expressed in an Infrastructure Strategy. This is divided into four IT Systems Architectures, which in turn are divided into a number of component strategies:

Infrastructure Strategy	
Systems Architectures	Component Strategies
Applications	Electronic Mail FileTransfer Terminal Access EDI
Telecomunications	TCP/IP FDDI Ethernet Token Ring ISDN PABX
Computer Systems	
Databases	

The strategy is expressed in terms of

Corporate Standards	:	these must be adhered to
Recommended Practices	:	these are essentially good technical advice
Technical Briefings	:	these are for information only.

The size and autonomy of the various groups and units means that before a component strategy can be accepted there is a hierarchy of managers that have to ''buy in'' to the strategy. This is particularly so for something new, such as a migration to OSI. Further reasons for having to promote OSI are:

- in most areas the current IT practitioners do not really understand OSI
- the lack of awareness of the range of OSI products available
- the perceived lack of immediate and clearly identifiable benefits from OSI
- the varying levels of support of OSI by computer vendors

Within the organisation, OSI is seen as having an important future potential, but currently there is only minor involvement with OSI. There has been a pilot project running high-level OSI protocols over TCP (this was undertaken mainly because of the interest of the individual driving the project).

2.3.4 Topology

The LAN topology at the major site using UNIX systems consists of about 100 Ethernet segments (located in a number of buildings), connected via MAC-level Bridges to a backbone segment. The total number of devices supported is about 2500. All connections between LANs are high speed links (both fibre-optic and T1 lines are used).

These LAN segments carry the following protocols:
IPS
DECnet / LAT
PC Protocols.

Although different segments belong to and support different units, in general the various protocols listed above share cabling and bridges. There are links to other sites for IPS and DECnet.

The SNA links are carried on a different cabling system, therefore IPS can be considered in isolation from SNA.

There has been some interest in routers, but there has been no decision to start using these in preference to bridges.

2.3.5 Network Configuration

The statistics at the main UNIX site are as follows:

Number of segments	=	100	
Number of backbones	=	1	
Total Devices	=	2500	
PCs	=	200	
Minis	=	400	
LAT Terminal Servers	=	500	
Telnet Terminal Servers	=	200	
Print Servers	=	50	
Workstations	=	1000	
Devices per segment	=	20	-500
External links	=	1	Internet (for R&D)
	=	10	9.2 kbps to sites outside U.S.A.
	=	6	56 kbps to sites in U.S.A.

2.3.6 Traffic

Data and applications are distributed, and the business requires a lot of interaction between units. This gives rise to many-to-many file transfer and electronic mail traffic. For a given user most terminal access is to a specific host.

There is no time-critical traffic over the backbone. Real-time, process control traffic is limited and is carried on a few local segments only.

FTP is used to transfer files primarily between workstations and, to a lesser degree, UNIX minicomputers. Most file transfer traffic (80-90%) is local to individual segments. Typical file sizes over the network are 10-100 Kbytes, with very few exceeding 20 Mbytes.

Electronic mail is not used extensively, hence traffic is small.

There are a number of Telnet (and LAT) terminal servers, which require good response times, consistent with dumb terminal usage.

The peak load on a segment is typically 5-15% (however, some segments can reach higher loads). The load on the backbone Ethernet is also in the same range, rarely exceeding 15%.

2.3.7 Network Management

At the corporate level, network management is driven by the requirements of the SNA network, where Netview is used extensively.

There is a small management centre at each of the two main UNIX sites, which looks after the Ethernet backbones and all the external connections. They are also responsible for addressing and distributing routing tables for all interconnected segments and the distribution of a directory accessible via electronic mail. No high-level protocols are used for controlling the Ethernet segments, but a lot of interest is developing in network management using SNMP.

The control of individual segments is left to the units owning the segments and is generally done on an ad-hoc basis. This inevitably results in a varied level of network management that depends on the competence and enthusiasm of the individuals concerned.

2.3.8 Interfaces

Normally applications do not have direct access to network facilities. Both file transfer and mail facilities are accessed either via scripts or are user driven.

It is expected that more OSI content will be introduced into the Digital systems with new releases of communications software. In anticipation of the arrival of this there has been a move on the Digital systems to develop subroutines that hide all details of the Network Interface from application developers.

2.3.9 Security

There are no special security requirements, with the possible exception of some encryption in the Research and Development (R & D) labs.

2.3.10 Summary

The main issues covered in this scenario are:

- IPS is used as a LAN protocol suite.
- The applications to be migrated are file transfer, electronic mail, and remote terminal access.
- Traffic volumes are fairly low. LAN peak loading rarely exceeds 15%. Electronic mail accounts for very little traffic.
- There are no time-critical applications.
- Most file transfer and mail connections are many-to-many.
- Network management is *ad-hoc*, but it is increasing in importance.
- There are no special security requirements.
- Although there are corporate IT standards, the use of OSI must be sold to those most affected by it. They are unlikely to migrate from IPS to OSI if there are any substantial disadvantages (even if these are short term and are outweighed by long term benefits).
- IPS is used in a small part of the organisation, consequently it is not seen as the primary architecture by higher management.
- Those involved in running the network, although aware of OSI, have little experience of it.

2.4 SCENARIO 2 - IPS AS COMMON INTERWORKING PROTOCOLS

2.4.1 Introduction

This scenario involves the potential replacement of IPS with OSI as a means of interconnecting dissimilar systems. It describes a very limited use of IPS in an environment dominated by proprietary LAN and WAN techniques. This probably represents a large class of IPS user organisations.

Because of the support of IPS by many different vendors, IPS is sometimes used to connect dissimilar systems, even in situations where there is very little use of the UNIX operating system. This role is likely to be taken over by OSI. A strategic switch from IPS to OSI for interconnecting dissimilar systems in a non-UNIX environment may present opportunities to sell into new organisations for companies that can help users solve the IPS-OSI migration problems.

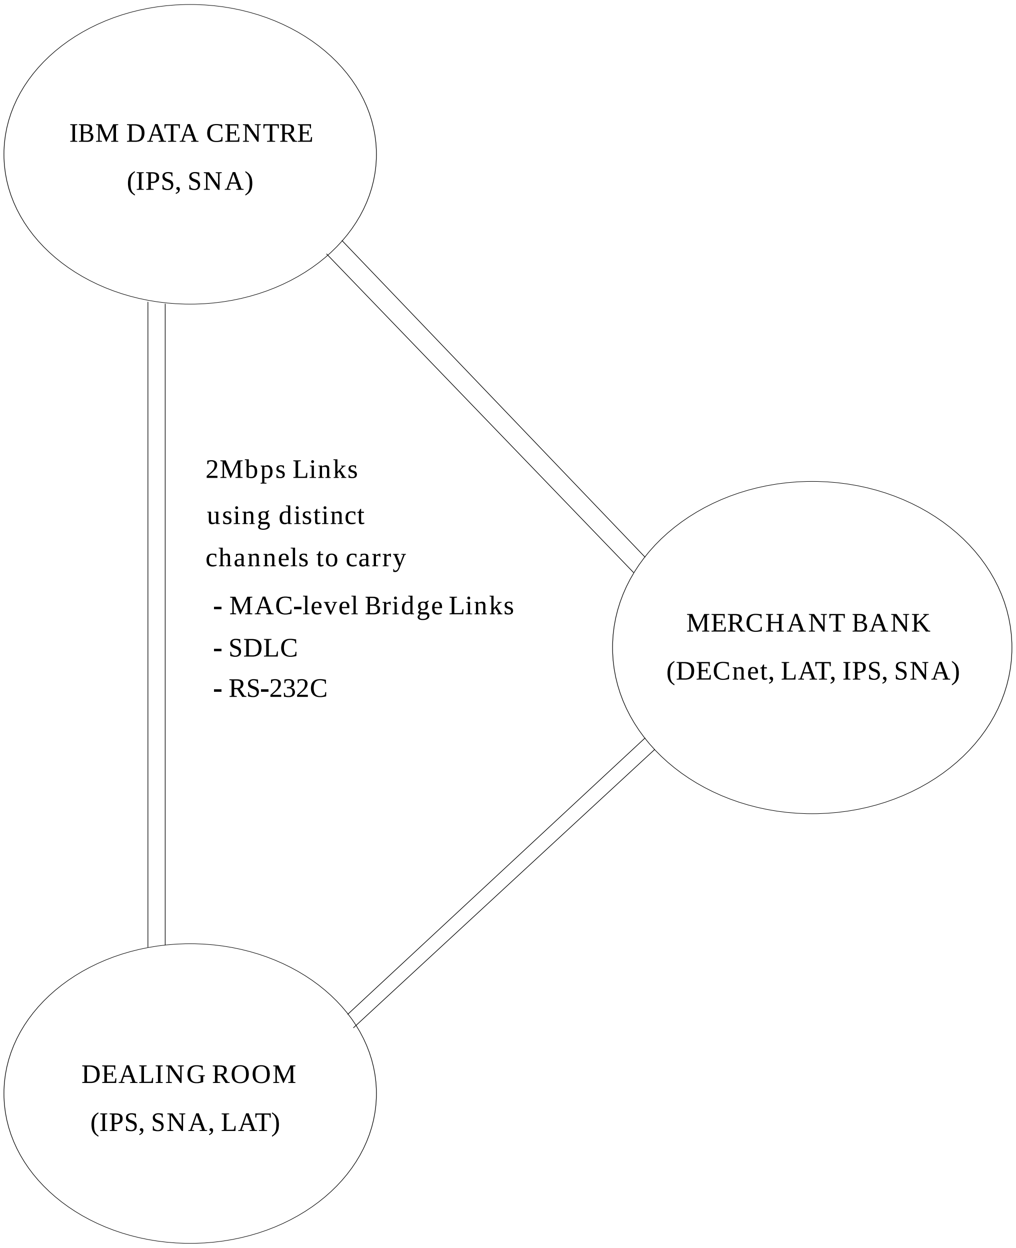

Scenario 2 IPS as Common Interworking Protocols

2.4.2 Type of Organisation

The company is a major merchant bank involved in a wide range of financial services. This includes merchant banking, trading in securities and equities, foreign exchange, investment management and all normal dealing services.

The company is based in central London, with its main offices in the City of London. However, the company also has large offices in two other main locations in the South-east of England; one of these locations houses the company's IBM data centre.

Over the years, international banking has grown to depend more and more on IT, and the company has developed and uses an extensive range of computer systems to support its day to day business. The provision of effective and reliable communications facilities is a key component of the supporting infrastructure.

2.4.3 Background

There are a number of discrete business communities and each has developed its own IT solutions. This has led to a variety of hardware systems to support their diverse activities:

- securities back office, investment management, accounting
- office automation
- merchant banking activities
- securities trading
- traded options
- equity research

The equipment used to support the different business functions is also used for electronic mail, consequently there has been a proliferation of different mail systems. The various computer systems are fairly autonomous, with the exception of the need for file transfer between IBM and Digital systems.

The IBM user community is by far the largest; therefore, it has a strong influence on the IT infrastructure. It recognises that it has integration problems and is looking for solutions. The IBM community is showing a lot of interest in SAA, DISOSS, and Netview as a means of achieving better integration.

2.4.4 IT Strategy

There is no overall IT strategy and the procurement of IT equipment and communications has been piecemeal, driven by the business requirements of individual departments with little regard to standards and future interoperation.

The provision of communications facilities is driven by short term business needs to keep in step with the facilities used by competitors. Hence, long term strategic planning is often sacrificed for short term expediency. The formulation of a networking strategy does not have strong corporate backing, being driven mostly by technical personnel. It is very difficult to put forward solutions that have long term, technical advantages because more expedient solutions tend to have a much better pay-off in the short term and hence a much stronger support from users.

There is little awareness of OSI at the higher management level. There is some interest and recognition of the benefits of OSI (especially X.400), but this is mostly at the technical level. X.400 is considered the most immediately useful OSI application, being seen as the most likely way to have interoperation between the mail systems of the various user communities. There is little interest in FTAM since FTP satisfies all current file transfer requirements.

The company's IT department has an unofficial policy to decoupled application programs from the details of the underlying communications facilities. A project has been put forward to produce a simplified, generic networking interface. However, upper management has not endorsed it and there has been little enthusiasm from users, consequently the project has made little progress.

The IPS protocols are currently seen as the *de facto* standards for interoperation of dissimilar systems. In particular, FTP is used to support file transfer between different systems, especially between IBM and Digital.

2.4.5 Topology

The three major sites support different systems. The city office supports the dealing room and has special communications requirements, with specialist financial information feeds. There is a possibility that TCP/IP may be used to distribute some of these feeds. In addition, there are SNA/SDLC connections to the IBM Data Centre. There are also a number of terminals connected to LAT terminal servers (all LAT terminal servers are connected to Ethernet).

The site operating the IBM data centre is supported by an SNA network. It supports the company's securities settlements, accounting and investment management and has SNA/SDLC links to the city office. The IBM data centre operates Netview, but does not make extensive use of it. However, Netview is seen as a long term strategic management tool.

The third site runs a number of Digital VAXs and IBM 6150s which support the company's merchant banking activities.

In addition, there are a number of other systems at all three sites, most importantly Wang supporting the corporate office automation system.

The three sites have Ethernet LANs based on Ungermann-Bass NetOne equipment. They are interconnected with 2 Mbps lines which are subdivided into 64 kbps and 9.6 kbps channels that support:

- MAC-level Ethernet bridge connections, carrying IPS, LAT and DECnet
- SDLC, carrying SNA
- Asynchronous terminal traffic (RS-232-C)

There are also some 9.6 kbps links to sites outside the UK.

The various divisions of the Merchant Bank have a degree of autonomy and each has independently selected its own hardware vendor. This has led to a variety of proprietary communications protocols. IPS serves a very important function in linking together, via file transfer, the IBM and Digital systems. In addition, the City Office is looking at IPS to support some specialist applications; these may require direct access to TCP/IP, Telnet and RPC.

2.4.6 Traffic

IPS traffic is confined to file transfer between Digital VAXs and IBM mainframes. This is an off-line activity that occurs overnight (between 6pm and 8am). Despite the lack of enthusiasm from users and upper management for decoupling applications from networking facilities, the IT department has guidelines that require applications to have no knowledge of the underlying communications facilities or the topology of the network. These guidelines have been followed in the case of file transfer.

File transfer is driven by requests from the receiving computer, that is there are no unsolicited transfers. The process is outside the control of applications programs, being driven entirely via scripts. Updating scripts to drive a different underlying file transfer process is seen as a fairly trivial task.

2.4.7 Security

Security is very important. An industry-specific encryption algorithm is used and two levels of passwords are used on the Digital VAX machines. For file transfers a special directory is set up, with the remote machine having only read access to files in this directory.

2.4.8 Summary

The main issues covered in this scenario are:

- IPS is used in links between computer systems.
- The main application to be migrated is file transfer.
- Although there are at present no X.400 applications, there is interest in X.400 as a means of linking together different mail systems.
- Traffic volumes are high but are on dedicated circuits; file transfers generally occur overnight.
- IPS Network management is not an issue.
- There are stringent security requirements.
- Users are not interested in the underlying technology infrastructure. Therefore, selling OSI is difficult in the areas where user requirements are currently met with IPS.
- Reliability is very important.
- The migration must not bring the network down for any significant period of time.
- IPS usage is functionally limited but its presence has stimulated an interest in open systems and networks.

2.5 SCENARIO 3 - ENGINEERING WORKSTATION ENVIRONMENT

2.5.1 Introduction

This is a scenario involving workstations running the UNIX operating system and interconnected via IPS. This is very typical of engineering and scientific environments. In such an environment, high traffic rates are typical. There is a requirement to connect to non-UNIX systems (IBM) used in other parts of the company.

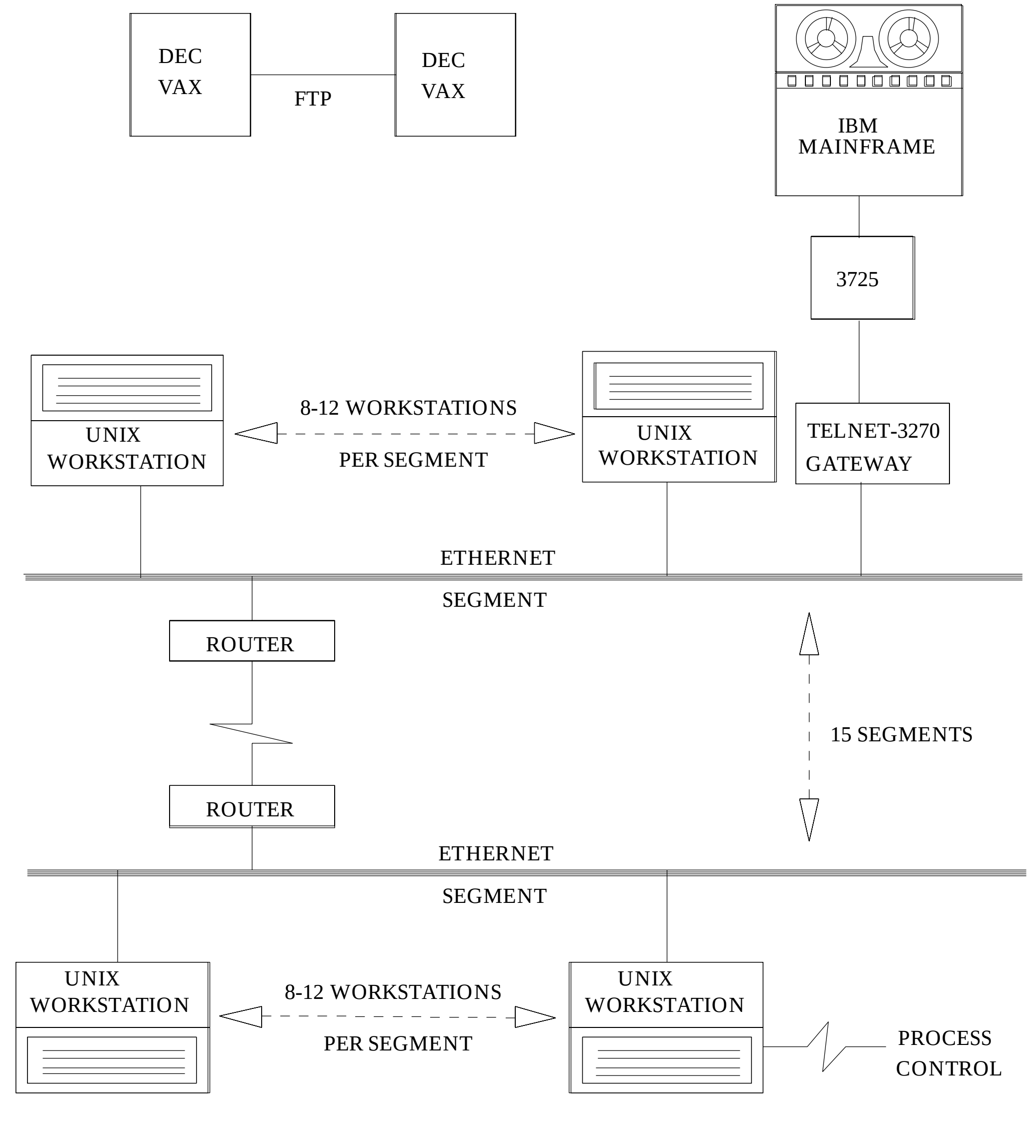

Scenario 3 Engineering Workstation Environment

2.5.2 Type of Organisation

This is an organisation involved in Oil Exploration. Business factors have encouraged the oil producers to maximise their profitability by taking responsibility for all aspects of the oil business from the exploration for new oil, through extraction and refining to eventual sale.

The oil industry is divided broadly into three operational areas: exploration, refining and sales. UNIX systems are used most heavily in the exploration area, which requires the use of very high-powered computer resources. Before deciding to drill for oil in a particular area, an oil company invests a great deal in geological research to determine whether there is a reasonable chance of finding oil. The large volume of seismic data generated by this research provides the input data for seismic modelling software running on high powered supercomputers.

The organisation has about 30,000 staff world-wide, with most located in a small number of offices in the U.S.A. Most of the data processing of the organisation is supported on IBM mainframes.

The exploration department is responsible for the provision of its own facilities at its site. However, it must access corporate SNA services via gateways.

2.5.3 Background

Until recently most data processing has been based on IBM mainframes. However, there has been a decision to adopt the UNIX operating system as a common platform for data processing operations in scientific and engineering areas. Consequently, over the past three years there has been a significant introduction of UNIX workstations into the organisation.

The procurement of IT equipment and communications has been driven by vendors. The exploration department has been fairly independent of the rest of the organisation, being driven by the requirements of the department.

Individual user requirements are met with workstations. All the workstations are networked using IPS protocols over Ethernet. In addition to these, there is a much smaller number of departmental minicomputers (Digital VAXs). The choice of IPS protocols was made because they are non-proprietory and the only real option offered by the workstation vendors selected.

There has been little recognition of OSI and there is no real drive to OSI nor pressure for the adoption of OSI protocols. There have been some experiments with X.400 but there are no plans to incorporate this into the existing electronic mail network.

X.25 is used for wide area connectivity and this is seen by users as a step towards OSI. There are plans to use X.25 to link some of the remote locations.

The generic requirements of the exploration department make heavy use of file transfer. The department is trying to meet these requirements with off-the-shelf products.

2.5.4 Topology

The organisation has an extensive corporate network, based on SNA. This connects 5 major sites using high speed (T1 or 512 kbps) lines. There are also a number of remote locations that are spurred from the nearest major site. Over the past year some of these have been upgraded to support X.25 links. In some cases IPS is used over X.25. Traffic over the X.25 links is invariably file transfer or alarms to minicomputers with X.25 links (there are no X.25 gateways on the LAN).

The exploration department is based in a number of offices located at a single site. These are supported by 15 Ethernet segments, each connecting 8-12 workstations. Until recently these Ethernet segments were connected with MAC-level bridges, however, these are being replaced with routers. Important considerations that have led to the choice of routers over bridges are:

- ability to filter by address and by protocol
- much better management potential (SNMP management capability is a requirement for all new routers).

Most segments have more than one link, however, some of the segments do not use dynamic routing. Addresses are distributed on a weekly basis. The network topology is not hidden from users, who are able to specify IP addresses rather than hosts names.

On one of the segments there is a connection to the department's IBM mainframe. This in turn has connections to the corporate SNA network and to a supercomputer.

Most PCs are standalone; none are connected to the Ethernet LANs though there is some localised use of Token Ring (again not connected to the Ethernet LAN).

The other minicomputers operated by the department are not connected to any network, but some have direct links for file transfer.

2.5.5 Network Configuration

The Exploration Department has the following systems:

UNIX Workstations	150
VAX/VMS	15
UNIX minicomputers	15
PCs	50

The number of workstations is likely to increase disproportionately over the next few years. Although the choice of workstation vendor is under the control of end users one workstation vendor has been predominant in the past. Recently a shift to a different workstation vendor has been observed.

The number of PCs is also likely to increase significantly over the next few years, however, these are used mostly in office functions, either stand-alone or networked in small groups. Currently there are no PC connections to the LAN supporting the workstations, although this need may arise in the future.

2.5.6 Traffic

The three main IPS application protocols - that is Telnet, SMTP and FTP - are all used to varying degrees.

Telnet is used between workstations, but in this role its use is limited and not very important. It is much more important in accessing IBM mainframe applications, where a Telnet to 3270 gateway is used. This is capable of concurrently supporting 16 LU sessions.

The corporate IBM mainframes support a company wide electronic mail system based on PROFS. This provides the only corporate mail system. There is a requirement for electronic mail between workstations and between workstations and the corporate IBM systems. The former is met using SMTP, while the latter is supported via an SMTP-PROFS gateway. Peak traffic levels are fairly low.

By far the heaviest usage of network bandwidth arises from file transfer. File transfer connectivity using the IPS FTP protocol is as follows:

- workstation - workstation
- workstation - IBM mainframes
- Digital - IBM
- Digital - Digital (where these are not connected via DECnet)
- UNIX Minicomputers - IBM

The use of FTP between such a variety of vendor equipment has led to a number of problems arising from incompatibilities between the FTP implementations and data representations of the different vendors.

Traffic between workstations is of the order of:

2-40 Mbytes, several times a day per workstation
100-150 Mbytes, typically, once a day per workstation.

FTP dependencies are not built into applications. However, the network topology is not hidden from users. Addresses are distributed on a weekly basis.

The typical traffic load per segment is 5-15%, with peak loads reaching as high as 80%.

There is no time-critical traffic on the LAN. Some of the workstations are involved in process control and in this role they filter host-bound reports and alarms to be logged and made available to other workstations on the LAN.

2.5.7 Network Management

Management of the LANs has been fairly *ad-hoc*. The management of individual segments is the responsibility of local personnel. However, as the number of networked systems is increasing, network management is becoming increasingly more important and a central management team is being established. They are looking at the network management offerings of the hardware vendors to help in this area. Currently, they do not use SNMP, but this may become important in the future.

2.5.8 Performance

The majority of applications are not time critical. The main performance requirement of the network is to handle the large traffic volumes being generated by file transfer between the workstations, and between the workstations and the IBM mainframes. To control this problem, routers are being used to isolate traffic in segments.

2.5.9 Availability and Resilience

Network availability and resilience has not yet been an issue. The absence of 'servers' on the LAN has meant that single points of failure are limited to low level network devices.

As the networking facilities become more sophisticated then this aspect of networking may be reconsidered.

2.5.10 Summary

The main issues covered in this scenario are:

- IPS is used as a LAN protocol.
- The choice of IPS was based primarily on availability and its open nature.
- Routers are used in preference to bridges.
- The main IPS protocols to be migrated are FTP and Telnet.
- Traffic volumes are high.
- IPS Network Management is not yet an issue.
- There are stringent security requirements.
- Reliability is very important.
- The migration must not adversely affect network connectivity, availability or bandwidth.

2.6 SCENARIO 4 - SUMMARY OPERATIONAL REQUIREMENT

2.6.1 Introduction

This scenario summarises a hypothetical OR for the expansion of the networking facilities of a large organisation. It describes an environment where topology, traffic, bandwidth, availability and resilience must all be considered.

The aim of this scenario is to demonstrate that an existing UNIX-based network can be extended and migrated to OSI and at the same time achieve good network performance. Performance figures can be calculated for an existing network, but these are not necessarily good targets for an OSI migration strategy. More realistic targets are those commonly expressed in ORs issued by large public or private sector organisations.

If the tools and techniques discussed in this guide result in additional devices being placed in the path between devices then this can have a detrimental affect on network performance. This scenario addresses the effects of additional devices introduced into the network (because of IPS-OSI coexistence or OSI migration) on such network parameters as:

- response times
- bandwidth
- reliability
- resilience

The scenario assumes a UK-based organisation in order to require the use of the Connection-oriented network protocol and Transport Class 2. It also specifies that some existing equipment uses the Connectionless Network Protocol in order to investigate how this can interoperate with the Connection-oriented network protocol.

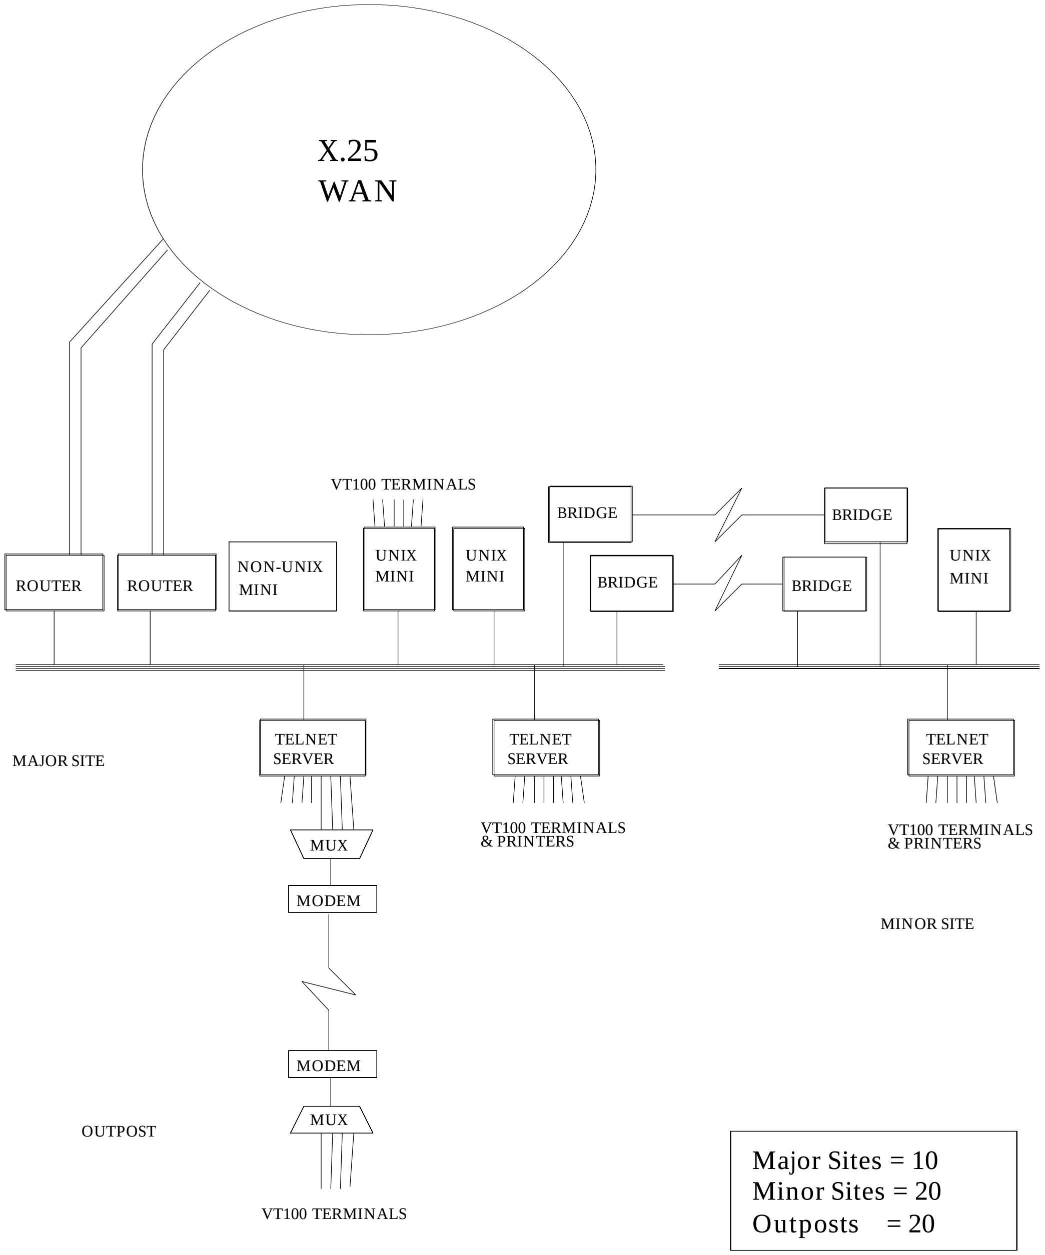

Scenario 4 Summary Operational Requirement

2.6.2 Type of Organisation

This is a UK government organisation which is in the process of updating its IT infrastructure and has issued an OR that states what is required from its networking infrastructure. Being a government organisation, it is under very strong pressure to conform, as far as possible, to the UK GOSIP specification.

The organisation has about 10,000 staff in 50 sites located throughout the UK. The 50 sites are divided into 10 major sites 20 minor sites and 20 outposts. One of the major sites is the organisation's headquarters.

The organisation is responsible for the provision of its own facilities and supporting services necessary to achieve its objectives. However, for communications between major sites, it is required to use an X.25 network run by a separate organisation.

2.6.3 Background

IT has made great inroads into all aspects of the organisation. In recognition of this an IT Strategy has been formulated. Until recently there has been no overall IT Strategy and the procurement of IT and communications equipment has been piecemeal, driven by the requirements of individual departments with little regard to standards and future interoperation. Most systems have been mainframe based, with a small number of minicomputer-based systems being introduced recently, mostly UNIX-based. Also the number of microcomputers has greatly increased.

The organisation's infrastructure has evolved to be very centralised, with most data processing needs being met by mainframes at the HQ site. A large number of terminals operating in block-mode are connected to the mainframes. Most are situated at the HQ site, with a small number situated at some of the major sites and connected to the mainframe via direct lines. In addition, at two major sites there are six non-UNIX minicomputers that are currently connected to the mainframe site via point-to-point X.25 lines.

2.6.4 IT Strategy

Networking facilities have been very limited, with terminals connected via direct lines to the mainframes. IPS networking using Ethernet exists at two major sites; all other minicomputer hosts are stand-alone with terminals directly connected to them.

The potential for gaining significant operational advantages from IT has been recognised and the organisation has commissioned a consultancy firm to develop an IT strategy. This included a communications infrastructure component. One of the fundamental precepts of the IT Strategy is a move to a systems architecture employing a high degree of distribution both of data and of processing. This means a significant increase in mini and microcomputers situated at all major and minor sites. Only a limited amount of mainframe processing remains.

All new systems are supported on mini and microcomputer, with mainframe based operations gradually being moved to minicomputer based systems. The IT strategy requires the greatest growth in minicomputer systems to occur over the next five years.

The applications architecture laid out in the IT strategy calls for significant file transfer and terminal interaction, mostly with local, but also with remote minicomputers. The strategic implementation provides the potential for word processing facilities on local processors and electronic mail over both the local and wide areas. The IT Strategy dictates that industry standards are used where applicable:

- UK GOSIP and POSIX for minicomputers and
- UK GOSIP and MS-DOS for microcomputers.

A standard RDBMS (with distributed capability) and 4GL are required.

2.6.5 General Objectives for Upgrading the Network

The overall objective for data communications is to install an infrastructure to embrace all known and forecast communications requirements for the next five years. The objectives are:

1. To provide local network facilities that:
   - support all present UNIX-based systems;
   - support the current levels of traffic; and
   - be capable of gradual expansion and managed reconfiguration to meet future user requirements.
2. To provide local networking facilities which, together with the wide area facilities being provided outside this project, meet the required network performance targets.
3. To provide a network that can support the organisation's requirements for multi-vendor connectivity and operability in order to conform to EEC Directive EC 87/95. For levels 1-4 UK GOSIP-T profile (equivalent to ISO International Profiles TC51 and TC 111x) must be supported.

The network has no security requirements, except for normal password facilities.

2.6.6 Topology

The network has a 3-tier hierarchy, as follows:

- Each major site has one or more links to the X.25 WAN (determined by traffic/availability/protocol requirements).
- Each minor site has one or more connections to a designated major site.
- Each outpost is directly connected to its nearest major or minor site (traffic considerations warrant just one 9.6 kbps line).

Minicomputers are situated at major and minor sites. Outposts only have terminal links to a minicomputer at a major or minor site.

One of the major sites is the organisation's headquarters. This is where the mainframes are located.

A key question for this topology is where to locate gateways in a hierarchical and geographically distributed network, and the effects of these on networking parameters.

2.6.7 General Requirements

The network must support all existing and all new UNIX minicomputers. All terminals must be linked via terminal servers, connected to ISO 8802/3-conformant LANs. The capability is required to allow users to select facilities on the network via menu(s). It must be possible to present different menus to different users or groups of users.

The organisation would like to be presented with different options for supporting the non-UNIX minicomputers. These have (the currently used) X.25 capability, but can also be easily upgraded to attach to a LAN with the following protocol profile:

- ISO 8802/3,
- Null Internet,
- Transport Class 4

A particular point of interest is that, in a LAN environment, these systems support only the Null Internet protocol at level 3.

The organisation would ideally want all devices to be handled by the same networking facilities, provided that this did not lead to significant extra cost or technical problems.

There should be a two-tier network management facilities structure. There should be centralised systems management facilities at the headquarters major site for management control and configuration of the local network facilities at all sites working in conjunction with local management facilities. These centralised facilities are called the Central Management Facility (CMF).

At each major site there is a Local Management Facility (LMF). This must be able to manage all communications facilities and systems at the major site and all linked minor sites. Each LMF must be able to function independently of all other LMFs or the CMF.

Network Management has become a very important requirement and users are realising that network management must not be confined to managing the physical links. The perception amongst many users is that ISO is a long way from producing stable network management standards.

The wide area element of the network is to be provided by another project. This network provides a packet switched interface based on the 1984 CCITT Recommendation X.25; this will be migrated to the 1988 version of X.25 when this becomes available from the WAN vendor.

2.6.8 Network Configuration

Currently there are:

- 10 UNIX-based minicomputers
- 6 Non-UNIX minicomputers
- 500 asynchronous, character mode terminals
- 20 Printers
- 100 PCs
- 2 IPS LANs (using Ethernet) each located at a major site

There are, in addition, other devices (minicomputers and terminal clusters) that are outside the scope of the new network.

Over the next five years the number of systems is projected to grow to the following:

Overall

- 50 UNIX-based minicomputers
- 6 Non-UNIX minicomputers
- 2200 asynchronous, character mode terminals
- 300 Printers
- 400 PCs
- Ethernet LANs to support the devices listed above.

Per Site

Device Type	Major	Minor	Outposts
UNIX minicomputers	2	1	0
UNIX terminals	100	50	8
Printers	20	9	1
PCs	20	10	0
Ethernet LANs	<----- as required ---->		

Two sites have three non-UNIX minicomputers each (as described previously).

2.6.9 Traffic

The IT Strategy fundamentally changes the pattern of traffic between sites. Currently, a large amount of traffic takes the form of transactions being passed between major sites and the HQ where the mainframes (and major databases) are held.

Under the Strategy each major site has its own database(s) and the majority of transactions are local. However, collation of management information, distribution of software upgrades, and the need to access existing data held on the central mainframes means a continuing requirement for traffic to and from the HQ.

Provision of more computing capacity to the users is realistically expected to generate demand for facilities which in turn brings about an increase in data communications traffic levels. An annual growth of 20% has been applied to the WAN traffic figures to reflect this increase.

LAN Traffic

The following assumptions should be used to calculate the traffic generated in the local areas:

All VDUs are in operation at the same time in echoplex mode.

Each VDU operator works at 20 words a minute (6 chars per word).

Print traffic is on average 250 bytes per second.

At least one file transfer is in progress at any one time. Average file size 25 Kbytes. The range is 100 bytes to 100 Kbytes, with few files exceeding 10 Mbytes.

Electronic mail traffic can be assumed to be negligible.

WAN Traffic

The figures below give the peak aggregate traffic that cross the X.25 WAN in bytes per second for all traffic type (terminal, file transfer and electronic mail) in five years time.

Major Site	Outgoing Traffic	Incoming Traffic
1	8000	5500
2	25000	10000
3	4500	22000
4	5000	23000
5	25000	6000
6	52000	22000
7	4600	19000
8	500	600
9	2200	4000
10	5000	17000

2.6.10 Performance

Different performance levels for network delay (response times) are required for WAN and LAN communications:

- For communications involving the WAN a response time of under 3 seconds is required for completion of a simple enquiry type transaction, given a 1 second turnaround time in the host.
- Over the WAN, there is a limited amount of echoplexing where the required mean response time for transmission and acknowledgement of a single character should be as short as practical and should be under 1 second under normal load.
- Within each local area (there are 10 local areas, each includes a major site and its connected minor sites and outposts), response times must be consistent with the need to provide echoplexing communications between asynchronous terminals and hosts to support such applications as word-processing. The response times must be such that they are not noticeable to a typist.
- There is also file transfer traffic for local and remote destinations and the network must support the traffic volumes listed elsewhere in this scenario.

2.6.11 Availability

The organisation requires that any user device has a minimum of 99.3% network availability to any other device on the network. The availability of the X.25 network must be included in all availability calculations.

2.6.12 Resilience

Resilience is an important factor of the network and is governed by the number of devices dependent on any network component. It is a requirement that single failures should not isolate more than 16 devices. For sites with four or fewer terminals alternative links to their parent sites need not be provided.

Access to major sites must have at least dual routes controlled by network personnel, and traffic from one being diverted automatically to another in the event of failure. For minor sites, manual intervention is allowed (however, a single failure must not isolate a minor site).

2.6.13 Inter-device Connectivity and Protocol Support

The network and all new devices on the network must comply with OSI protocols. Within the OSI network there should be no theoretical limitations hindering any-to-any connectivity.

The network must provide connectivity between any pair of minicomputers. This should also extend between (current) minicomputers using IPS and (new) minicomputers using OSI. The network must support IPS protocols on an interim basis for any systems that may require it. The degree of interoperation between IPS- and OSI-based UNIX hosts must be stated.

All current PCs are connected as VT100 terminals to minicomputers. However, in the future it is expected that there will be a significant use of networked PCs, some requiring terminal and file transfer access to minicomputers.

Any dumb (VT100) terminal (or intelligent device able to emulate VT100) must be able to communicate with any UNIX host both inter and intra site.

All new dumb screens are VT100-compatible and are connected to terminal servers (or similar). Where existing dumb terminals are VT100 compatible they should be migrated to use terminal servers unless such terminals are in the same or adjacent rooms as the minicomputer to which they are connected (in which case no benefit is seen in changing their mode of connection).

The network must provide connection for asynchronous printers. For printers serving only one host, the network sets up a default connection. However, some printers are shared between systems, therefore the network must support printer sharing, printer control (printer busy, paper out) and printer access control.

UNIX minicomputers should be connected to the network by means of ISO 8802/3 adaptors. If reverse terminal servers are used, these should have hunt group capability.

2.6.14 WAN Characteristics

The service performance levels guaranteed by the WAN provider are as follows:

- The probability of a packet, or its data contents, forming part of an established call being correctly delivered is greater than 99.99%.
- The probability of a corrupt data packet being detected is greater than 99.94%
- The mean transit delay of a packet across the network, under busy traffic load must not exceed 400 milliseconds; 99% of the traffic having a transit delay of one second or less.
- The WAN must provide a continuous service with an availability of 99.90% for a single port on a single switch or 99.9999% for a single port on two switches.

2.6.15 Summary

The main issues covered in this scenario are:

- Support on a common infrastructure both for connectionless-mode and connection-mode network protocols.
- Interworking between IPS and OSI UNIX systems.
- Effect of migration on
  - — response times
  - — bandwidth
  - — reliability
  - — resilience.

2.7 SCENARIO 5 - USER CONCERNS

2.7.1 Introduction

The aim of this scenario is to illustrate the type of non-technical problems that migration and coexistence come up against. This scenario lists a number of subjective worries and prejudices expressed by a number of people in the organisation contacted. Objections to OSI listed below are all reports of views expressed in interviews or questionnaires, but they come from several organisations.

This scenario is perhaps the most important because it shows that in addition to delivering key user and operational functions, an OSI solution must also address the user's technical concerns (and possibly prejudices) if it is to gain acceptance.

This scenario describes a hierarchically structured network in a scientific environment with heavy interoperation requirements between systems. It is a high profile, public sector system and the use of reliable, proven technology is essential. The role of the system is essentially a planning one and involves a succession of long term projects, often lasting two to three years. The system was established some years ago and has met one of its primary objectives of reducing the planning effort by about 60%.

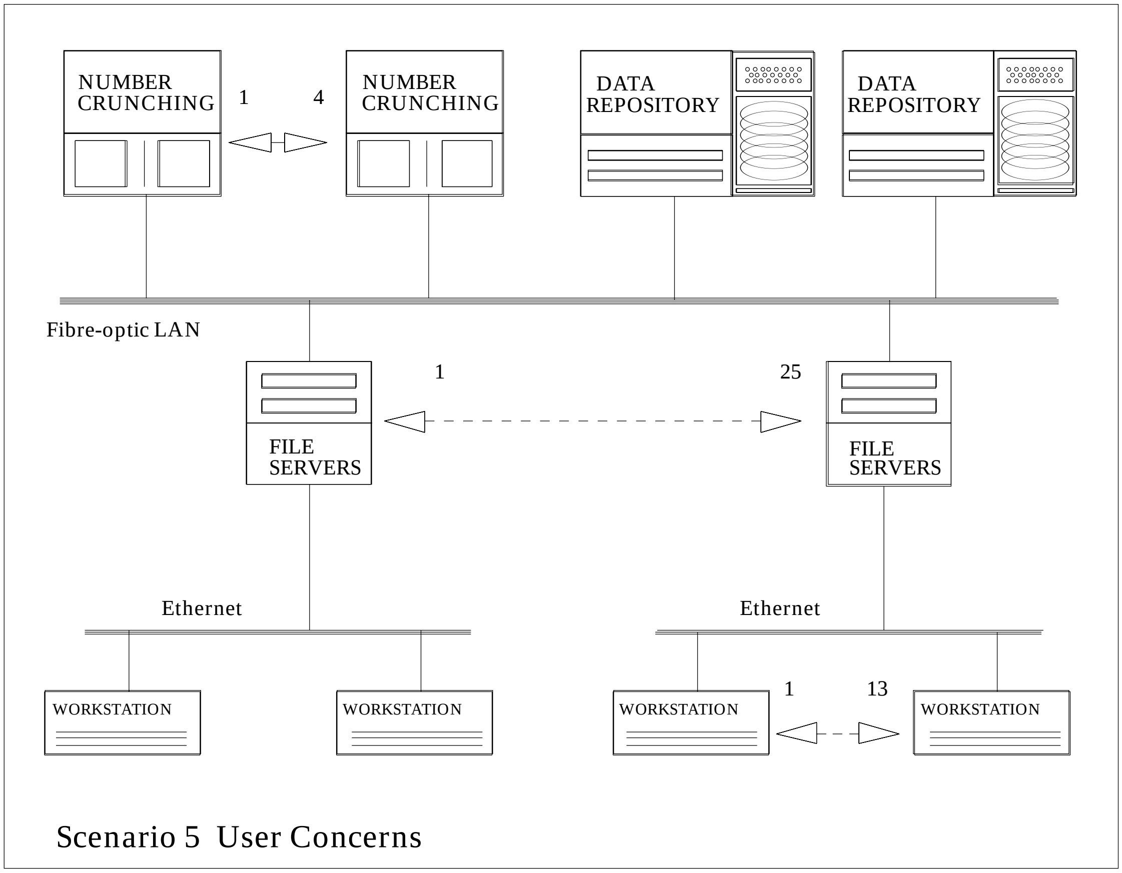

Scenario 5 User Concerns

2.7.2 Topology and Devices

This organisation uses high-powered, networked workstations, running normal applications such as database, graphics and spreadsheets, as well as some special custom applications. These are served by a number of file servers, which in turn provide access to powerful machines supporting massive databases. These file servers also provide access to more powerful machines that provide additional processing for computation-intensive applications.

The system relies upon good, reliable, LAN communications. Because of the type of hardware being used the only real choice of networking technology available at the time the network was being set up was IPS. The LAN architecture selected consists of a fibre-optic LAN connecting the database and number crunching machines, with normal copper Ethernet cable connecting all the other machines. Applications and data are interchangeable and shareable between workstations. The hardware for this system can be logically viewed in three tiers:

Tier 1

Two very large data storage and archival systems. They are UNIX-based, multi-processors with good I/O subsystems and large amounts of off-line storage.

Four powerful minicomputers, providing additional computational power to the workstations;

The above are interconnected by fibre-optic LAN.

Tier 2

25 fileservers; these are powerful UNIX-based multi-processors. They are all connected to the fibre-optic LAN and each supports an Ethernet LAN to which a number of workstations are connected.

Tier 3

25 LANs and attached workstations; each LAN is connected to one of the fileservers and supports 7-20 workstations, with an average of 13.

2.7.3 Connectivity

Required communications are between workstations on the same and other Ethernets and between workstations and LAN servers.

The type of links required are primarily file transfer and, to a lesser degree, remote login (mostly to run remote shell). Typical file transfers are:

- Small (10 Kbytes - 100 Kbytes), several times a days
- Large (10 Mbytes - 100 Mbytes), once per day and once per week.

Recently there has been interest in X-Windows and there are tentative plans to introduce X-Windows on workstations and run some X-Windows aware programs on some of the server machines.

2.7.4 Choice Of Network Architecture

The IPS protocol suite was chosen for the network because of the need for reliable, proven, off-the-shelf products. OSI was considered as an option but there were a number of reasons why it was not chosen:

- No external communications requirement was anticipated, so OSI's global connectivity was of no advantage.
- Standard IPS protocols could satisfy all project communications requirements; the system did not require OSI's additional functions.
- There was (and is) no installed OSI user base, hence, no real experience of using OSI.
- In this field there is an overwhelming penetration of IPS; the project did not wish to use unproven technology.
- OSI has a number of optional features that can cause incompatibilities. Very few OSI conformance testing facilities existed when the project was established and the current conformance testing facilities are viewed with scepticism.
- The user perceives that OSI has adopted two mutually incompatible sets of options at levels 3 and 4 (the Connection-oriented Network Protocol combined with Transport Class 2, and the Connectionless Network Protocol combined with Transport Class 4). This nullifies OSI's advantage if the need to communicate with other scientific institutions arises, as these are likely to be European as well as U.S.A.-based.
- Lack of experience with OSI products raised the possibility of indirect and hidden costs. OSI may be offered as an added feature at extra cost.
- There was a perception that neither business partners nor customers nor suppliers were migrating to OSI.
- OSI was not (and is not) universally supported, hence, an OSI-only policy could have compromised choice of vendors. Vendor commitment to OSI was (and is) uncertain.
- In the past vendors have often made unrealistic promises.
- There have been reports of performance degradation in the use of OSI protocols, which were perceived to be heavy-weight by comparison with IPS.
- OSI network management standards had not reached a stable stage.

The recent move to use OSI in all public sector computer systems has raised the issue of whether to migrate the communications architecture of the system to one based exclusively on OSI protocols. A migration plan must overcome both the technical problems involved in the migration, but also, must overcome the considerable hostility that exists towards a move from IPS to OSI.

2.7.5 Summary

There are two sets of issues in this scenario:

- how to migrate a hierarchically structured network
- how to overcome the objections to OSI and demonstrate that the use of OSI has a net advantage.

Chapter 3

Migration Endpoints

This chapter provides some basic background information about IPS and OSI networks, and serves as an introduction to the migration and coexistence techniques discussed in the next chapter. **Section 3.1**, **The OSI Basic Reference Model**, introduces the model used to discuss networks. The organisation of the first three sections of this chapter is based on this model. **Section 3.2**, **OSI - The Open Systems Interconnection Networking Standard**, describes the OSI protocol stack. **Section 3.3**, **IPS - The Internet Protocol Suite**, describes IPS networks. A complete description of current common practice for this type of network is available in the referenced **XGIPS** document. (For more detailed information on either of these types of network, consult the appropriate specifications.) Following this background information, **Section 3.4**, **Functional Comparison of IPS and OSI**, then compares the two networking standards, summarising their similarities and differences and highlighting the issues that must be addressed in the coexistence and migration solutions. Finally, **Section 3.6**, **Additional Applications**, discusses additional application protocols which are not part of the official IPS protocol stack but which it is necessary to consider when planning a migration or coexistence strategy.

3.1 THE OSI BASIC REFERENCE MODEL

The ISO Basic Reference Model for Open Systems Interconnection (ISO 7498) models communication between computer systems. It is generally referred to as the OSI Reference Model. Although IPS networks actually predate this model, it fits the layering of the IPS protocol stack rather well. As a result, the OSI Reference Model provides a useful structure for comparing IPS and OSI networks.

In the OSI Reference Model, communications-related processing is organised into seven layers, each layer having well-defined characteristics.

The principles that were followed by ISO in defining the layers can be summarised as follows.

a. The number of layers should not be larger than necessary.

b. Interactions between layers should be kept simple.

c. Functions that are essentially different in kind should be in separate layers.

d. Conversely, functions that are similar to each other should be in the same layer.

e. Past experience should be used as a guide.

f. Relationships between layers should not be affected by changes within individual layers; such changes may be required to take advantage of advances in technology.

g. There should be layer boundaries corresponding to interfaces for which standards would be useful.

h. Different levels of data abstraction should be handled in different layers.

i. Changes within one layer should not necessitate changes within other layers.

j. Each layer should have boundaries only with the layer immediately above it and the layer immediately below it.

The resulting model is illustrated in the diagram below.

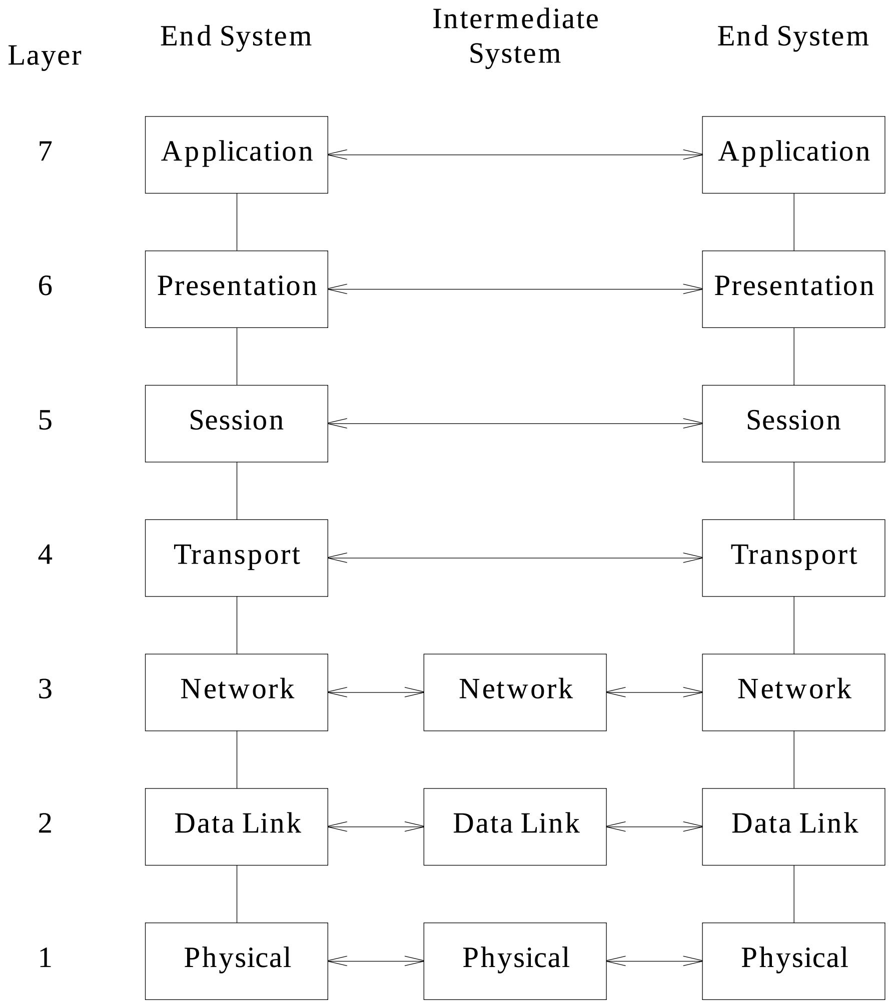

Computer systems containing application processes that communicate with application processes in other computer systems are modelled as *end systems*. Within each end system, information for transmission to another end system is passed by the application process to the communications environment at the top layer of the model: the application layer. It is then passed down through successively lower layers until it reaches the physical layer (which is the lowest layer). At the physical layer, the information is given a physical form (for example, as a sequence of voltage levels) for transmission to another system. In the receiving end system, information received at the physical layer is passed up through successive layers until it reaches the application layer, at which it passes from the communications environment to the communicating application process.

As well as providing for transfer of information between end systems that are connected to the same local or wide area network, the model also provides for information interchange between end systems that are connected to different networks, possibly of different types. In this case, the information passes through a series of other systems, which are referred to as *intermediate systems*. These relay the information across a series of networks, which are referred to as *subnetworks*. Within an intermediate system, incoming information is passed up only as far as the network layer (layer 3) and is then passed back down to the physical layer for transmission to another intermediate system or to the destination end system. Thus the upper four layers (which provide end-to-end functions) are not present in intermediate systems.

Conceptually, each layer in an end system communicates only with the same layer in other end systems, using the services of the layers below. The mechanism used to achieve this is as follows.

In the sending system, each layer encapsulates data by adding information in the form of a header or, sometimes, a trailer, before passing it to the layer immediately below. Data from upper layers is not examined or changed. In the receiving system, each layer strips and processes the information added by the same layer in the sending system and then passes the remaining information to the layer immediately above.

The only exception to this occurs at the presentation layer, which may compress, encrypt, or otherwise transform information from the application layer. Here, the semantics of the information is preserved but its syntax may be altered.

The main functions of each layer are as follows:

Physical layer

The physical layer is the lowest layer in the model. It provides for physical connections and bit transmission between systems. Specifications at this layer cover aspects such as voltage control, pin assignment, bit timing and line discipline.

Data link layer

The data link layer provides for logical connections and transfer of data units between systems that are directly connected by physical connections. Its functions include: data link connection management, framing, error detection and correction, and sequencing and flow control.

Network layer

The network layer provides for logical connections and transfer of data units between systems that are part of the same network but are not necessarily connected to the same subnetwork (the connection between them may involve a chain of intermediate systems). Its functions include: routing and relaying, connection management, multiplexing, sequence control, error detection and recovery, flow control, segmentation and blocking, and congestion control.

Transport layer

The transport layer provides transfer of data across a network between processes within systems. It is concerned with the reliability and cost of the data transfer. Its functions include: connection management, sequence control, error detection and recovery, flow control, segmenting, blocking and concatenation, and multiplexing.

Session layer

The session layer provides the means for communicating processes to organise and synchronise their dialogue and to manage their data exchange. Its functions include: connection management, flow control, dialogue discipline negotiation, and error recovery.

Presentation layer

The presentation layer provides for the representation of information that is transmitted between communicating systems or is referred to in communications between systems. Its functions include: data compression, encryption, and canonical data representation.

Application layer

The Application Layer is the highest layer in the model. It provides the means for processes in communicating systems to access the communications environment. Its functions include specific functions such as messaging and file transfer plus general supporting functions such as identification and authentication of communications partners and control of associations with them.

3.2 OSI - THE OPEN SYSTEMS INTERCONNECTION STANDARD

3.2.1 Background

Development of the OSI suite of standards was instigated by the International Organisation for Standardisation (ISO), which is an organisation of national standards bodies.

ISO collaborates with the International Electrotechnical Commission (IEC). Together, the two bodies provide the framework for world standardisation. In the field of Information Technology, they have formed a Joint Technical Committee (ISO/IEC JTC1). This committee is now the ultimate authority for work on OSI.

ISO/IEC is responsible for the overall framework of OSI and has developed many of the OSI standards. Some OSI standards have however been created by other standards organisations or have been created jointly by other standards organisations and ISO.

One such body is the International Telegraph and Telephone Consultative Committee (CCITT), an international organisation whose main members are national postal, telephone and telegraph authorities (PTTs) and other public telecommunications service operators. It makes recommendations which effectively constitute international telecommunications standards. Some of the OSI standards are derived from CCITT recommendations and some others are the results of collaboration between ISO and CCITT and are published by both organisations.

Another organisation whose standards have been adopted by ISO is the Institute of Electrical and Electronic Engineers (IEEE). This is a United States professional body, membership of which is however open to people from other countries.

At each layer except the application layer, there are ISO standards that define the services provided by that layer to the layer above. The application layer is the highest layer and does not provide a service to any other layer; at this layer there is an ISO standard that defines the relationship between the communications environment and the applications that use it.

At each layer except the physical layer, a number of communications protocols are specified by ISO standards. These protocols are used by communicating entities that are at the same layer of the model but in different systems. At the physical layer, there are standards that specify connectors, electrical characteristics and other aspects of a physical connection.

The service and protocol standards describe the layers in a manner that is independent of any programming language or operating system conventions. The method of implementation of the protocols at each layer is not constrained, nor is it required that an implementation recognise the layer boundaries in its design.

The OSI reference model allows for two possible modes of communication: connection mode and connectionless mode (the original version of ISO 7498 covers connection mode only; connectionless mode is covered in ISO 7498 Addendum 1.) In connection mode, two communicating entities establish a connection, transfer data over it and then release it. In connectionless mode, one entity can send a piece of information to another but there is no concept of a connection between them; each unit of data is transferred in a single, self-contained operation.

The choice of connection-mode or connectionless-mode transmission has no implications at the physical layer but each of the next five higher layers (data link, network, transport, session and presentation) may provide either a connection-mode or a connectionless-mode service. These layers and the application layer may use either connection-mode or connectionless-mode protocols as appropriate.

Interoperability is one of the primary goals of the standards and perhaps one of the best tests of a given implementation. The OSI standards are designed to be widely applicable and are defined with many optional features which will actually inhibit interoperability unless agreements are reached on which options are implemented. To overcome this potential problem, *profiles* have been defined to support particular applications in particular types of network. A number of such profiles have been or are in the process of being adopted as *International Standardised Profiles* (ISPs) by ISO / IEC JTC1.

The OSI suite of standards is being produced to reduce the proliferation of proprietary networks and to promote connectivity between systems by making them open. Work on these standards began in the 1970s, but only by around 1988 were they considered sufficiently complete that it was practical to create actual working networks based upon them.

The process of creating these standards is lengthy and complex, because the cooperation of many organisations is required and the acceptance procedure is protracted. To begin establishing a standard, a working group creates an initial version of the standard, often based on an existing standard submitted by a member organisation. The document then goes through the formal stages established by ISO where it is revised and refined until it is ready to be submitted as an *International Standard*, subject to a voting procedure. Although many of the OSI standards are full International Standards and are stable, some of those referred to in this section have not yet reached International Standard status, so some modification can be expected.

The following sections describe each layer in more detail. These descriptions provide a background for the later explanations of coexistence and migration between OSI and IPS. Because of this, particular emphasis is placed on the transport and application layers, where coexistence and migration problems are most likely to be encountered.

3.2.2 OSI Physical Layer

The service provided by the physical layer is defined in ISO 10022. There are a number of standards specifying how the service is implemented over different physical media. The most commonly encountered ones are as follows.

Table **3-1.** OSI Physical Layer Standards

Standard	**Description**
CCITT V.24	DB-25 connector, EIA-232-D
CCITT V.35	DB-25 connector, high speed interface, EIA-422-D
ISO 4903	Point-to-point connection, 8-pin connector, can be circuit switched. Equivalent to CCITT X.21
ISO 8802-3	CSMA / CD, derived from Ethernet, uses baseband coaxial cable of various sizes plus (optionally) optical fibre and runs at speeds from 1 to 10 Mbps.
ISO 8802-4	Token Bus, broad-band coaxial cable; 1, 5, or 10 Mbps, most commonly encountered in factory automation (MAP) networks.
ISO 8802-5	Token Ring, shielded twisted pair; 1 or 4 Mbps.

The ISO 8802 standards are derived from, and, in the cases of the standards referred to above, are almost identical to, the IEEE 802 series standards. These standards define three layers: the physical layer, the medium access layer (MAC), and the logical link layer. The logical link layer is described in the ISO 8802-2 specification and clearly corresponds to a part of the data link layer of the OSI model. The MAC layer has some characteristics of both the physical and data link layers of the OSI model, while the ISO 8802 physical layer corresponds to a part of the OSI physical layer. The three ISO 8802 specifications listed in the table above describe the physical layer and the MAC layer for three different types of network.

Ethernet is the LAN standard defined by DEC, Intel and Xerox (the DIX consortium) on which IEEE 802 and ISO 8802 are based. It is still in common use in networks throughout the world. Ethernet and ISO 8802-3 are identical at the physical layer but have slightly different packet headers. This has important implications for some coexistence strategies. These issues are discussed in greater detail in **Section 3.4**, **Functional Comparison of IPS and OSI**, where these network standards are compared.

3.2.3 OSI Data Link Layer

The OSI data link service is defined in ISO DIS 8886. A number of different protocol standards are defined at this layer to cater for differences of physical media and type of subnetwork. The most commonly encountered ones are:

Table **3-2.** OSI Data Link Layer Protocol Standards

Standard	**Description**
ISO 7776	HDLC LAP-B compatible data link layer procedures - derived from CCITT X.25
ISO 8802-2	Logical Link Control (LLC) used above the IS0-8802-3, ISO 8802-4, and ISO 8802-5 protocols

The ISO 7776 protocol is mostly used on point-to-point links in Wide Area Networks (WANs), over V.24, V.35 or X.21. It supports a connection-mode service.

The ISO 8802-2 specification defines the (in IEEE 802 terms) Logical Link Control (LLC) layer (often referred to as the LLC sublayer of the OSI data link layer). It specifies a common set of protocols that is used in Local Area Networks (LANs) conforming to ISO 8802-3, -4 and -5. It supports both connection-mode and connectionless-mode services.

The ISO 8802-3, -4, and -5 specifications are also relevant at this layer because they define some data link layer services such as error detection and flow control as part of their (in IEEE 802 terms) MAC layer specifications. This set of services is often referred to as the MAC sublayer of the OSI data link layer.

3.2.4 OSI Network Layer

The OSI network layer service is defined in ISO 8348 and in ISO 8348 Addenda 1-4. At this layer, different protocols are specified for connection-mode and connectionless-mode operation.

Table **3-3.** OSI Network Layer Protocol Standards

Standard	Description
ISO 8208	Connection-mode Network Service (CONS) Packet Level Protocol. Derived from CCITT X.25
ISO 8878	Use of X.25 to provide the connection-mode network service
ISO 8881	Use of X.25 over local area networks
ISO 8473	Protocol for Providing the Connectionless-mode Network Service (CLNS).

The network service provides information transfer across a network which may consist of a number of subnetworks, possibly of different types. For example, an item of information could pass from the sending end system across an ISO 8802-3 LAN subnetwork to an intermediate system, then across an X.25 WAN subnetwork to a second intermediate system and finally across an ISO 8802-5 LAN subnetwork to the receiving end system.

Both a connection-mode service (the Connection Oriented Network Service - CONS) and a connectionless-mode service (the Connectionless Network Service - CLNS) are defined by OSI standards. This distinction has particular importance for some coexistence and migration solutions because IPS services at this layer are connectionless-mode and, while the OSI connectionless-mode service is similar in many respects to the IPS service, some OSI products support only the connection-mode service.

In addition to providing integration between different types of subnetwork, the network layer includes the function of ensuring that data is routed properly. While passing from one end system to another, data may be routed through several intermediate systems. In the sending end system, and in each intermediate system, the outgoing route is selected on the basis of:

- the network address contained in the data, where transmission is in connectionless-mode or where the data contains a connection-mode connection set up request

- the connection identifier contained in the data, where transmission is in connection-mode and the data does not contain a connection set up request
- routing information held in the end system or intermediate system.

The routing information in each system may be there as a result of administration procedures or may have been obtained through communication with other systems using routing protocols. The OSI standards include specifications of two routing protocols, for use in conjunction with provision of the connectionless-mode network service:

- the ES-IS protocol (defined in ISO 9542), which enables end systems and intermediate systems on the same subnetwork to exchange information about each other
- the IS-IS protocol (defined in ISO DIS 10589), which enables intermediate systems to exchange information about each other.

OSI network addressing is defined in ISO 8348 Addendum 2. Network addresses are partitioned into two parts: the Initial Domain Part (IDP) and the Domain Specific Part (DSP). The IDP contains two pieces of information: the Authority and Format Identifier (AFI) and the Initial Domain Identifier (IDI). The AFI indicates the addressing authority and the format for the IDP. ISO assigns AFI values. Each AFI value is associated with a particular addressing authority and implies a specific format and length for the IDI. The IDI then defines the agency responsible for the format and value of the DSP.

<table>
<tr><td colspan="2">IDP</td><td rowspan="2">DSP</td></tr>
<tr><td>AFI</td><td>IDI</td></tr>
</table>

As an example, ISO has assigned to US GOSIP the IDI 0005 under AFI 47 (decimal). AFI 47 specifies that the IDI is interpreted as four decimal digits, and that the DSP is represented in binary form. In the US GOSIP specification, the first octet of the DSP is the format for the remainder of the DSP. Currently the only DFI assigned is 80 which identifies a DSP of seven fields (including the DFI itself):

Table **3-4.** Example US GOSIP DSP (DFI=80)

Field Contents	**Width**
Domain Format Identifier (80)	1 octet
Adminstrative Authority	3 octets
Reserved	2 octets
Routing Domain	2 octets
Area	2 octets
System	6 octets
Network Selector	1 octet

3.2.5 OSI Transport Layer

The OSI transport service is defined in ISO 8072 and ISO 8072 Addendum 1. Only two OSI protocols are specified at this layer.

Table **3-5.** OSI Transport Layer Protocol Standards

Standard	Description
ISO 8073	Connection-oriented transport protocol specification
ISO 8602	Connectionless transport protocol specification

Transport Layer Services and Protocols

The transport layer optimises the use of the available network service to provide transparent transfer of data between end systems.

The ISO transport service includes both a connection-mode service (the Connection Oriented Transport Service - COTS - defined in the body of ISO 8072) and a connectionless-mode service (the Connectionless Transport Service - CLTS - defined in ISO 8072 Addendum 1). Currently, the connectionless-mode transport service is rarely encountered.

A COTS connection has three phases:

Connection Establishment
The transport service user initiates this phase by issuing a transport connection request. This contains the calling and called transport addresses and specifies the quality of service (QOS) requested and whether expedited data should be enabled or not. The transport service establishes the connection and responds to the transport service user, indicating the QOS that applies and whether expedited data is available. (The transport service user may choose to disconnect the call if the QOS is unacceptable or if the requested expedited data service is not available.)

Data Transfer
After the connection has been established, data can be transferred between the two endpoints. Separate transfer requests are used for regular data and for expedited data (if available). Flow control may be invoked to limit the data rate for regular data.

Connection Release
The connection may be terminated by the transport service user in either of the connected end systems or by the transport service itself.

The COTS is provided using the Connection Oriented Transport Protocol. This protocol can operate over any of three types of network layer service:

Type A
Acceptable rate of errors signalled by the network service (for example by disconnect or reset) and acceptable rate of errors that are not signalled (so that, for practical purposes, received data can be assumed not to have been corrupted).

Type B

Unacceptable rate of errors signalled by the network service but acceptable rate of errors that are not signalled (errors are too frequent but any received data can be assumed not to have been corrupted).

Type C

Unacceptable rate of unsignalled errors (errors are not detected and data may be corrupted).

Note that, for Type A and Type B networks, the network layer detects lost, corrupted, or re-ordered data, indicates the error to the transport layer and does not pass any of this data to the transport layer. For Type C connections, the network layer does not indicate errors and passes all data, correct or incorrect, to the transport layer. If reliable data transfer is required over a Type C network, it must be provided in the transport layer.

The connection-mode transport service may be used over either the connection-mode or the connectionless-mode network service. The connection-mode network service may be of types A, B or C. The connectionless-mode network service has similar characteristics to a Type C network.

Five classes of the Connection Oriented Transport Protocol are defined:

Class 0 - Simple

This class may break up and reassemble TPDUs if they are too large for the underlying network layer. Expedited data is not available. If the network layer indicates a disconnect or reset, a transport layer disconnect is initiated. This class is intended to operate over a Type A network service.

Class 1 - Basic Error Recovery

This class provides recovery from error conditions by keeping copies of the TPDU as they are sent out and by waiting for acknowledgements from the other end. If a network disconnect occurs, the connection is re-established, and the unacknowledged TPDUs are resent. Unacknowledged TPDUs are also resent if the network indicates a reset. Expedited data is not available. This class is designed to operate over a Type B network service.

Class 2 - Multiplexing

This class provides multiple logical connections between two systems. Multiplexing saves the overhead and expense of setting up a new network connection when a connection already exists between two systems. Class 2 can also provide flow control for each transport connection. Because no error recovery is provided, this class may only operate over Type A network services.

Class 3 - Error Recovery and Multiplexing

The features of both Class 1 and Class 2 are combined in this class. The Class 1 features make it suitable for operation over both Type A and Type B network services.

Class 4 - Error Detection and Recovery

This is the only transport class suitable for operation over a Type C network service (and hence the only class that is used over the connectionless-mode network service). It checks for the integrity of the received data with a checksum and has a timeout mechanism for acknowledgements. In the case of checksum errors or timeouts, the missing or erroneous data is simply retransmitted. Also, incorrectly ordered data is correctly reassembled. Expedited data is available. As it is possible to operate over the Connection Oriented Network Service, the class 4 transport protocol also supports multiplexing.

These classes are designed to support the different types of network layer service available and to allow the upper OSI layers to choose the quality of service that is needed. The following table shows the transport classes that apply to each network type:

Table **3-6.** OSI Network Characteristics and Transport Classes

Network Characteristics			
Undetected Errors	**Reported Errors**	**Multiplexing**	
Acceptable	Acceptable	N	0,1
Acceptable	Acceptable	Y	2,3
Acceptable	Unacceptable	N	1
Acceptable	Unacceptable	Y	3
Unacceptable	Unacceptable	Y	4
Unacceptable	Unacceptable	N	4

The CLTS does not have any connection-establishment or connection-release phases. It simply provides for transfer of units of data, which carry both the sender's and the receiver's addresses.

The CLTS may be used over either the connection-mode or the connectionless-mode network service.

Transport Layer Application Programming Interface (API)

Although the Basic Reference Model does not recognise application access to the services of the transport layer, such APIs do exist. One such interface is the X/Open Transport Interface (XTI), which provides access to the transport services of a number of protocol suites. The transport layer is the lowest layer to which application access is commonly provided.

3.2.6 OSI Session Layer

The OSI session service is defined in ISO 8326 and in ISO 8326 Addenda 1-3. Only two OSI protocols are defined at this layer.

Table **3-7.** OSI Session Layer Protocol Standards

Standard	**Description**
ISO 8327	Connection oriented session protocol specification
ISO 9548	Connectionless session protocol specification

The OSI session service includes both a connection-mode service (the Connection Oriented Session Service defined in the body of ISO 8326) and a connectionless-mode service (the Connectionless Session Service defined in ISO 8326 Addendum 3). Currently, use of the connectionless-mode service is rarely encountered.

The Connection Oriented Session Service is provided by using the Connection Oriented Transport Service. The basic connection and transfer functions of the Connection Oriented Session Service map onto the equivalent ones in the COTS:

Connection Establishment

The transport layer parameters are present and additional parameters are included to support session layer services such as Token Management and Synchronisation. A single transport connection may be serially reused by a number of session connections (perhaps to avoid the overhead of transport connection set-up).

Data Transfer

In addition to the regular data and expedited data that the transport layer offers, the session layer adds *Typed Data* and *Capability Data*. Typed data is used in conjunction with half-duplex data services to send data when the sender does not have the data token. Capability Data is used in conjunction with activity management services to exchange data outside the context of an activity.

Connection Release

At the session layer, connection release is *orderly*, that is it prevents the loss of data. This enhances the connection release service provided by the transport layer which can lose data during termination of a connection. Orderly release becomes an important issue when dealing with some coexistence and migration solutions, because IPS does not have an explicit session layer and provides this service through the TCP protocol which, in other respects, can be classified as belonging to the OSI transport layer. In addition to the orderly, synchronised release, two abnormal termination primitives are available that may result in data loss. One of these is initiated by the session layer for unrecoverable errors, and the other can be initiated by the user of the session layer.

In addition, there are a number of other groups of functions which may be negotiated for use during the data transfer phase:

Token Management

Token-management options may be invoked when the connection is established. There are four different tokens, each controlling the exchange of a different kind of information. For example, the *data token* allows for half-duplex data communication between session layer entities. This is a useful feature for applications that must send and receive data in a controlled and coordinated fashion. The other tokens are used for negotiating release of the connection and synchronising the exchange of information. Applications that use tokens are in one of two states: they either have the token, in which case they may send information, or they do not have the token and are waiting for information from the other end. Primitives are provided for requesting and exchanging these tokens.

Synchronisation

Synchronisation, also called dialogue control, allows sessions to be interrupted and restarted at a known point. As data is sent, checkpoint numbers are included. The receiving end may then return acknowledgements as these checkpoints are received. If a connection is interrupted, a new connection may be established, and the data transfer can continue after the last confirmed checkpoint. This is an optional feature, established at connect time, which gives a level of protection from network failures and loss of data at the application.

Activity Management

Activity management allows data exchange on a connection to be controlled. Data exchanges are contained within *activities*, each of which has specific start and end indications, which is helpful in some kinds of transaction processing. Activities can be nested, and an activity can be temporarily interrupted to allow a higher-priority activity.

Exception Reporting

If abnormal conditions require the attention of the session layer user, the condition is reported instead of the session being terminated. In the case of an exception report, the user at the other end must either correct the condition or terminate the session.

3.2.7 OSI Presentation Layer

The OSI presentation service is defined in ISO 8822 and in ISO 8822 Addenda 1-3. Two protocols are specified at this layer and there are also standards specifying how information is represented.

Table **3-8.** OSI Presentation Layer Standards

Standard	**Description**
ISO 8823	Connection oriented presentation protocol specification
ISO 9576	Connectionless presentation protocol specification
ISO 8824	Specification of Abstract Syntax Notation 1 (ASN.1).
ISO 8825	Basic encoding rules for Abstract Syntax Notation 1 (ASN.1).

Two modes of OSI presentation service are defined: a connection-mode service (the Connection Oriented Presentation Service defined in the body of ISO 8822) and a connectionless-mode service (defined in Addendum 1 of ISO 8822). The connectionless-mode service is currently rarely encountered in use.

The services offered by the presentation layer are intended to preserve the meaning of the data. This is the only layer in which the data itself may be directly modified. Currently, there are no ISO defined data compression and encryption methods, so the applications must define them, if required.

The ISO 8824 and ISO 8825 standards define a portable method for representing various kinds of information. Applications may exchange data that contain many different types of information. For example, a mail item might consist of the addresses of both sender and recipient, the current time and date, the priority, a subject field, and the message itself. Other mail items might contain a different assortment of data types. ASN.1 defines the way in which each piece of data is encoded so that it can be correctly interpreted regardless of the machine type.

The formal ASN.1 definition is quite complex; simply put, it defines a formal grammar for describing data of any type. This grammar is used to create standardised descriptions of data. Associated with it is a set of Basic Encoding Rules (BER) that define how data described by ASN.1 can be represented as strings of bits. Using this representation, applications running on different systems, possibly with different information architectures, can communicate. A total of 27 universal data types are defined, of which integer and date are examples. The syntax is quite flexible, because it allows new data types to be defined and even allows the syntax to be redefined through a macro facility. Compilers are available for translating ASN.1 notation.

3.2.8 OSI Application Layer

The internal structure of the applications layer and the manner in which data passes between an application and the communications environment is defined in ISO 9545.

Application layer functions are partitioned into application service elements (ASEs). Each ASE performs a particular service (which is described in a service definition standard) using a particular protocol (described in a protocol specification standard). Each application may incorporate one or more ASEs to perform its communications functions (the ASEs use lower layer functions, accessed via the presentation service, as appropriate.) An application may also use the presentation service directly.

An ASE may be generic, that is applicable to a range of applications (for example ACSE, the Association Control Service Element, which is used by a number of applications, such as Message Handling and File Transfer), or application-specific (for example CMISE, the Common Management Information Service Element).

Generic ASEs

In addition to the services provided by the lower layers, many applications share a need for other services that are not available in these layers. To meet these additional needs, ISO has defined ASEs that provide functions that any application can use. Some of the most important ones are described here.

Table **3-9.** Generic OSI Application Service Element Standards

Standard	Description
ISO 8649	ACSE service definition
ISO 8650	ACSE protocol specification
ISO 9804	CCR service definition
ISO 9805	CCR protocol specification
ISO 9072	ROSE - model, service definition and protocol specification
ISO 9066	RTSE - model, service definition and protocol specification

Association Control Service Element (ACSE)

A cooperative relationship between two applications level entities is referred to as an *association*. The ACSE provides functions that enable applications entities to set up and manage associations with other application entities.

Commitment, Concurrency, and Recovery (CCR) Service Element

Many applications with coordinated data bases (for example, banking) require simultaneous actions at multiple points without error. The transactions must either succeed at all points or fail at all points. The CCR service element provides services which can be used to implement several styles of synchronisation. These include checkpointing, rollback. and two-phase commit.

Reliable Transfer Service Element (RTSE)

This service element is designed to facilitate the reliable transfer of potentially large blocks of data. RTSE also provides confirmation of the data transfer.

Remote Operations Service Element (ROSE)

ROSE allows one application entity to invoke another remotely to perform a function.

Specific Applications

The following applications are described in subsequent sections of this chapter:

- Message Handling (MHS)
- Directory Services (DS)
- File Transfer, Access, and Management (FTAM)
- Virtual Terminal (VT)
- Common Management (CMIS / CMIP)

The standards for these service elements are stable and some implementations are already available.

There are a number of other important applications, not described here, that are currently in the process of being defined by various ISO working groups. Electronic Data Interchange (EDI) provides a standard method for exchanging business documents such as purchase orders, invoices, customs declarations, and so on. Transaction Processing (TP) provides services for applications that require coordinated manipulation of data. Remote Data Base Access (RDA) enables an application executing on one system to access a data base resident on another. These applications, and many others, are expected to enjoy widespread use on OSI networks.

3.2.9 Electronic Mail - OSI Message Handling Service (MHS)

ISO uses the term MOTIS to refer to this service, but the more popular term, MHS, is used here. The original MHS standards were created by CCITT as the X.400 series of recommendations. As a consequence, the term X.400 Mail is frequently used to refer to this application. The specifications in the most recent version of the X.400 series recommendations (1988) are essentially identical to those in the MOTIS standards issued by ISO as the various parts of ISO 10021.

MHS Structure

The MHS standards divide the MHS functions into three major parts: the *User Agent* (UA), the *Message Transfer Agent* (MTA), and the *Message Store*. These provide the Inter-Personal Messaging Service (IPMS), the Message Transfer Service (MTS) and the Message Store Service (MSS) respectively. These parts and their relationships are illustrated in the following figure:

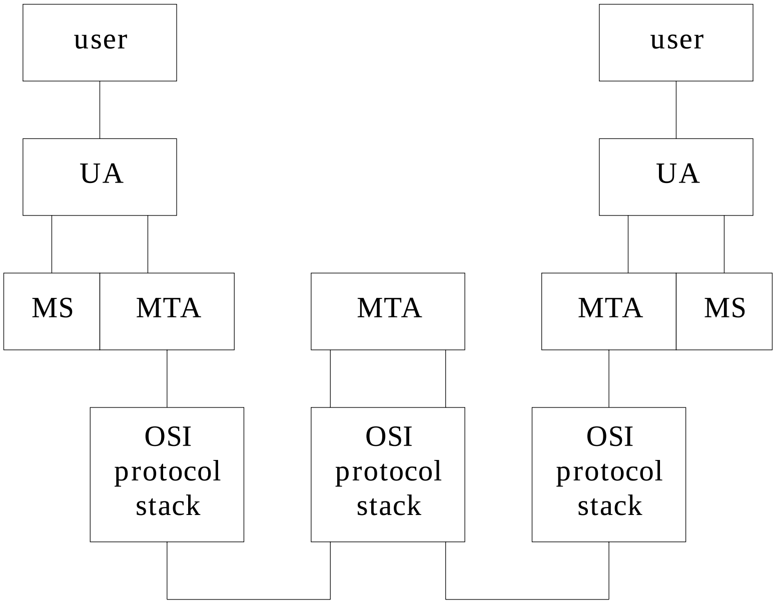

The UA provides the interface between the Message Handling Service and a user of that service. It is responsible for adding an *envelope* to the message *contents*, and passing it on to an MTA. The MTA is responsible for forwarding the message. It examines the address or addresses on the envelope, replaces the envelope with a new one and either delivers the message or passes it on to another MTA. The message may pass through several MTAs before it reaches its destination. The destination MTA can then pass the message directly to the receiving user's UA or to the MS, from which the UA can retrieve it subsequently.

Initially, the only type of UA that was defined was the Inter-Personal Messaging User Agent (IPMUA) which supported an electronic mail service for use by people. Subsequently there has been work on an EDI UA to support electronic exchange of information by computer systems.

The *envelope* consists of fields that are defined in the MHS specification. The *recipient address* is used by each MTA to determine the route. The *originator address* is used to return undeliverable items or for delivery confirmation. Other fields include those used to control the priority, date and time of delivery.

The *content type* field of the envelope indicates the type of information contained in the message. The contents can be any kind of data, including strictly binary data like encrypted databases and executables. Encrypted contents are usually indicated by the appropriate field in the envelope but this is not required; all that is required is that the *content type* be understood by the receiving UA.

Other *extended* fields may also be defined by agreement between the MTAs. If one of these *extended* fields is encountered that is unknown to the MTA, a flag in the field is examined to see if it is critical to mail transfer. If it is, the mail is returned; otherwise, the message is simply passed on. The number of defined fields is quite large, although there are only a few required fields.

For messages between people sent by IPMUAs, the contents have two parts: the header and the body. Like the envelope, the header consists of a set of defined fields. Again, a large number of fields are defined, although a typical message only uses a few of them.

These fields are only examined by the receiving UA. They may contain information that duplicates information in the envelope.

Addressing

Addresses in MHS are referred to as *originator/recipient addresses* or simply as *O/R addresses*. These addresses contain enough information so that each MTA can correctly forward the item until it reaches the destination UA. A user may also create a mail item using an *O/R name*. An O/R name must contain enough information so that a directory search results in an O/R address. An O/R name may contain a complete O/R address, in which case no directory look-up is required. (Directory services are described in the next section.) By definition, all O/R addresses are also O/R names but, conversely, not all O/R names are O/R addresses, because O/R names may be incomplete.

Each O/R address consists of a set of attributes. The following table shows the possible variants of the form of an O/R address and the attributes that may be present in each (a Postal O/R Address variant is also defined in CCITT Recommendation X.402 but is not shown in the table).

Mnemonic O/R Address

Country name

Administration domain name

Private domain name

Personal name

Organisation name

Organisation unit names

Domain-defined attributes

Numeric O/R Address

Country name

Administration domain name

Numeric user identifier

Private domain name

Domain-defined attributes

Terminal O/R Address

Country name

Administration domain name

Network address (for example: as defined by CCITT Recommendation X.121)

Private domain name

domain-defined attributes

Message Contents

One of the most attractive features of MHS is its ability to carry different types of messages in the body parts of the contents; for example, Telex, Videotex, digitised voice, facsimile, and several other formats are allowed, in addition to standard IA5 text. If the UA cannot handle the particular contents of a message, the MTA may convert the data before handing it to the UA. Before the conversion can take place, the appropriate fields on the envelope must indicate the data types contained in the message and the types of conversion, if any, that are permitted. For example, if one of the bodies in a mail item contains Videotex data, and the UA understands only ASCII data, the envelope fields must indicate that the contents contain Videotex data and that approximate conversion is allowed, because Videotex data may not be fully translatable to ASCII.

3.2.10 OSI Directory Service (DS)

The Directory Service and the protocols that it uses are specified in the various parts of ISO 9594, which are aligned with the CCITT X.500 series recommendations.

The DS provides a generalised mechanism for determining addresses, given one or more attributes. It was originally designed for use by the MHS application but can also be used by other applications such as network management.

The DS manages a directory that is arranged hierarchically. For example, the root might contain country names and successive levels might be based upon a variety of organisational and geographic considerations.

The DS includes two sets of functions: those performed by a Directory User Agent (DUA), a conceptual entity that acts on behalf of a user of the directory, and those performed by a Directory System Agent (DSA).

Directory User Agents (DUAs) can query the directory by specifying a set of attributes that cause the Directory System Agent (DSA) to descend through the directory hierarchy to get an entry or entries. The DSA can return an error indicating that no entry has been found or that access permission has not been granted. The DSA can also tell the DUA to query another DSA or the original DSA may do the query itself. For example, the DUA could make a query using the following attributes:

```
country=Holland, city=Apeldoorn, name=Meindert Hobbema
```

In this example, the DSA would search the root directory for the country, Holland, and then the city directory for the city, Apeldoorn, and then the city directory for the name, Meindert Hobbema. The DSA would then return the requested values (for example; telephone number, network address, and so on). Several conditions might prevent a successful search for this entry: the user might not have access to the country directory, one of the subdirectories might not be available, or the return values could be disallowed because there are too many matches.

Another query might look like this:

```
country=Holland, org=Philips, org unit=R&D, name=Meindert Hobbema
```

Here, the search is based upon organisational, rather than geographical, features.

Although information in the directories is arranged hierarchically, from the DUA's standpoint, the search is simply based on the set of attributes. In the above example, the city attribute of Apeldoorn could have been added to the list of attributes. The DSA must determine how the search is to proceed, based on the attributes it is given.

The DS can be implemented in a variety of ways. On receiving a query, a DSA examines its own data to see if it can answer the query. If it cannot, it may query another DSA for the information or it may tell the DUA where the query should be made. Different DSAs might contain replicated data. Any of the entries can be modified at any time, and the querying arrangements between the DSAs can change as well. As a result, it is possible for values returned to the DUA to vary from one query to the next.

The DS supports aliases. For example, the name *payroll* could be an entry in the directory that refers to a particular mail address. Users of an MTA connected to this DSA could simply mail items to this O/R name and let the DSA find the correct O/R address. This feature spares users from supplying the full address and keeping track of address changes. An alias can also refer to a list of entries, so mailing lists are supported as well.

3.2.11 Remote File Access - OSI File Transfer, Access, and Management (FTAM)

The OSI File Transfer, Access and Management service and associated protocols are defined and specified in the various parts of ISO 8571.

The central concept behind FTAM is the virtual filestore. The virtual filestore is a set of files, each with contents and a set of attributes. The possible attributes are of two different types: *file attributes* and *activity attributes*. The file attributes describe the general characteristics of the file and are identical for all users of the file. These attributes include the user names and dates for the most recent access or modifications to the file, the file name, access permissions, and so on.

Activity attributes describe attributes that may vary from one user to the next. The attributes include the name and status of the current users, as well as access-control information to allow file locking for concurrent file access. The extensive list of file attributes is designed to include the features found on file systems in a wide variety of computer operating systems. The local implementation must deal appropriately with attributes that are not available in the local system.

One commonly available file-system feature that is not reflected in FTAM is the file directory. Currently, FTAM cannot directly reflect the hierarchical structure of the UNIX file system. An OSI working group is developing a file directory facility that may eventually be added to the FTAM standard.

FTAM offers a series of primitives for interacting with the virtual filestore. These primitives are arranged into *regimes*. The regimes are nested so that a regime must be entered before the associated primitives are available to the application. Notice that each regime contains a primitive for leaving that regime and another primitive for entering the next lower regime. Some of the primitives operate on File Access Data Units (FADUs), structured units defined for the particular file.

Association Regime

The association regime is entered whenever an association is established between the requesting application and the application providing the virtual filestore. Password protection is invoked when the association is established, and the set of services that is available is negotiated at this time. A new file can be created in this regime.

File Selection Regime

The file selection regime is entered when a file is selected for possible use or a new file is created. File attributes may be selected or modified, but the file open regime must be entered before the file can be modified.

File Open Regime

The file open regime is entered in preparation for data transfer or deletion of FADUs. The file pointer may be changed and FADUs may be erased. The activity attributes are established for the file at this time.

File Transfer Regime

The transfer regime is entered whenever data is read or written. When the transfer is completed, the file transfer regime is exited. Checkpointing is available to allow a restart of data transmission in case of failures at the lower layers.

3.2.12 Terminal Service - OSI Virtual Terminal (VT)

Table **3-10.** OSI Virtual Terminal Standards

Standard	Description
ISO 9040	Virtual terminal service definition - basic class
ISO 9041	Virtual terminal protocol specification - basic class

The VT service enables a user whose terminal is connected to one system to interact with an application executing on another system. ISO 9040 defines the Basic Class VT service and ISO 9041 specifies the protocol used to provide it. The Basic Class VT supports character-oriented devices. Both standard asynchronous ASCII terminals and IBM 3270-type terminals are supported. Bit-mapped terminals are not supported.

A Virtual Terminal Protocol Machine (VTPM) operates in each end system. The VTPM in the application end system receives terminal commands from the application and passes keyboard input to the application. The VTPM in the terminal end system receives keystrokes from the keyboard and passes terminal commands to the terminal. The state of the VT is kept in the Conceptual Communications Area (CCA) in each VTPM. The two VTPMs keep the state of the CCAs updated by communicating with each other using the VT protocol. The following diagram shows the relationship between the components.

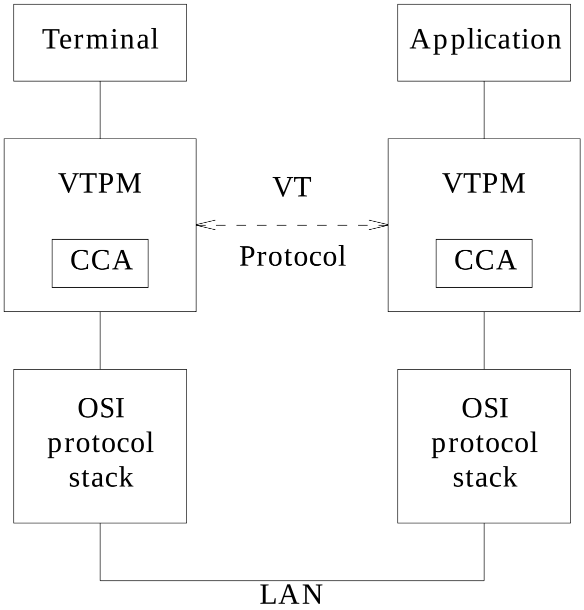

The description of a virtual terminal is contained in the Virtual Terminal Environment (VTE). The VTE contains a large set of parameters that describe the terminal characteristics. These parameters contain information about such things as keyboard characteristics, screen dimensions, colours, and fonts. Once these terminal characteristics are negotiated (usually during association establishment) the CCAs can be updated by exchanging VT protocol messages.

3.2.13 OSI Systems Management

OSI systems management defines a model for the management of networked computer systems and provides an application service element and definition of a Management Information Base for use in building distributed network management applications.

The core standards are now at full IS status and agreements are in place which allow implementation of the management services and of OSI-based management applications. There is significant implementation activity in this area and products based upon draft standards and earlier agreements may be available. However, the final standards are too recent for a significant number of installed systems to exist.

Table **3-11.** OSI Systems Management Standards

Standard	**Description**
ISO 9595	Common Management Information Service
ISO 9596	Common Management Information Protocol
ISO DIS 10040	Systems Management Overview
ISO DIS 10164	Systems management functions (such as object management, alarm reporting and event reporting)
ISO DIS 10165	Structure of management information - model, definition and guidelines for definition.

An overview of the management model and its relationship to the specifications which define OSI systems management is given in ISO DIS 10040. The following figure, based upon figures in the ISO specification, gives a much simplified view; its elements are summarised below.

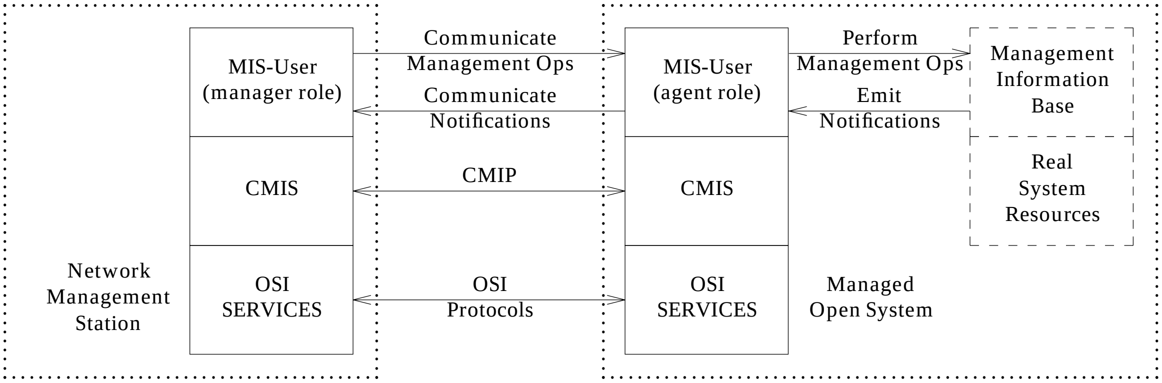

Management Information Base

A managed system is described by a set of *managed objects* which are conceptual representations of the system's actual resources (such as a communications device or protocol layer implementation). Information about these objects is contained in a conceptual repository of information known as the *Management Information Base.* The definition of a managed object includes its attributes, the operations that may be performed on it (including the conditions required for them to succeed), and the events that it can generate. For example, an object representing a communications controller might have attributes representing its operational status, its physical network address, traffic counters and so on.

There are two classes of managed object: objects specific to an individual layer, which are termed *(N)-layer managed objects*, and objects relevant to more than one layer (perhaps the whole system) or specific to a systems management function, which are termed *system managed objects.*

The various parts of ISO 10165 introduce the model for describing system resources, define generally useful managed object classes and lay down guidelines for the definition of objects.

Management Roles and Operations

The model terms an application performing systems management an *MIS-User* and defines two management roles: *manager*, requesting management operations and receiving notifications; and *agent*, performing management operations and forwarding notifications. When two management applications establish an association, the endpoints negotiate which roles each endpoint is to perform. It is possible for the endpoints to assume both manager and agent roles; in this case an endpoint selects an appropriate role for each interaction it initiates.

A *management operation* is one of the following:

- a request to get or set the value of an attribute (such as getting the value of a traffic counter for a communications controller or resetting the counter to zero)
- a request to perform an action on an object (for example, performing a self test or setting loop-back operation on a controller)
- a request to create or delete an object (for example, configuring an additional controller into a managed system)

It is the task of the agent MIS-User to interpret these logical operations in an appropriate way for the real resources of the managed system.

Management operations may use scoping and filtering to perform an operation on more than one object. Thus scoping might be used to get all elements of a routing table. Alternatively, a filter could be added to select only entries within a certain cost range.

A *notification* is an event, generated by a real resource of a system. It may be used to indicate a change of state (such as ‘‘interface out of service’’), or the creation or deletion of a managed object. An agent may forward a notification to one or more managers or store it locally. Again the agent must map the events generated by an actual resource in terms of the logical notifications defined.

Both operations and notifications can be defined as confirmed or un-confirmed or both, in which case the initiating application defines whether a confirmation is required.

OSI Communications Aspects

OSI systems management provides an application service element (the Common Management Information Service Element - CMISE) for applications which wish to perform management activities. This ASE provides the primitives which are used to implement the operations and notifications described above. In order to allow for restricted implementations, the protocol supports negotiation of the set of functions to be supported by the endpoints of an association.

ISO 9595 defines the services provided for systems management applications. ISO 9596 defines how these services are implemented using the services of ACSE, ROSE and the presentation layer.

Functional aspects

The Systems Management specifications contained in the various parts of ISO 10164 define management capabilities that address particular requirements, adding value beyond that provided by the basis management communication services.

Organisational Aspects

The management framework supports the concept of management domains, allowing the network to be partitioned in different ways for different functional purposes (for example, distributing responsibility for managing security whilst centralising fault and configuration management).

The model allows an OSI Systems Management application to manage systems and resources which are not part of an open system, or which are part of a separate network of open systems. This is supported by the concept of a *cascaded operation* where a managed object represents a resource outside the managed system. The result of applying an operation to such an object is an interaction between the managed system and the remote resource, possibly using further OSI Systems management operations or some other management protocol. This might be used to manage networks based upon other protocol suites or to manage devices (such as MAC-level bridges) which do not support the full OSI stack.

Future Developments

One shortcoming of OSI systems management is that its requirement for a full OSI stack for communication between manager and agent is inappropriate for management of simple network devices such as LAN bridges. Current activity is aimed at developing methods of providing CMIS directly over the data link layer. This is known as ‘‘CMOL’’.

3.2.14 OSI Security

Table **3-12.** OSI Security Standards

Standard	Description
ISO 7498-2	Security architecture
ISO 8649 AMD 1	ACSE service definition - Authentication during association establishment.
ISO 8650 AMD 1	ACSE protocol specification - Authentication during association establishment.

The security architecture addendum to the OSI Basic Reference Model defines a framework for addressing security within the OSI protocols. It does little more than define the types of security mechanism supported by OSI protocols (access control, authentication, identification, encryption, and so on) and identify the layers within the model where these security mechanisms can reside.

Two ISO steering committees are involved in further work.

SC21, working group 1, is addressing general architectural aspects of the OSI protocols. As part of their task, they are developing a group of security frameworks for open systems, each addressing a specific security mechanism identified in the Security Architecture. Whilst some of these frameworks are at an advanced stage, work on developing specifications for the services and protocols themselves (under the control of other SC21 working groups) is less advanced. An amendment has been agreed to ACSE, specifying services and protocols for authentication during association establishment. Work is in progress on a number of aspects of security management within the OSI Systems Management standards.

The ISO steering committee SC27 is addressing Common Security Techniques for IT Applications. Working groups within this committee are developing specifications for techniques (such as encipherment using public and private keys) that are generally applicable. Some of these techniques may be applicable to the security mechanisms being developed by SC21.

In conclusion, there is a significant amount of work being done on security mechanisms for use in OSI protocols, however, there are as yet few emerging standards for their practical application.

3.3 IPS - THE INTERNET PROTOCOL SUITE

3.3.1 Background

In the late 1960s, the Defense Advanced Research Projects Agency (DARPA), part of the U.S. Department of Defense (DoD), started supporting research on computer networks. The resulting network, ARPANET, was first implemented in 1969 and is still used within the U.S. military and many participating universities and private companies. The use of this networking standard increased rapidly, eventually including a variety of large and small networks throughout the world. This large collection of interconnected networks, which includes the original ARPANET, is commonly known as the *Internet*.

The set of protocols that these networks use is referred to in this document as the Internet Protocol Suite (IPS). IPS standards are established by the Internet Activities Board (IAB), which also provides a forum for other networking issues. These standards cover several areas, including:

- User applications.
- Communications protocols.
- Addressing and routing.

IPS standards are defined in Requests for Comments (RFC). An RFC is a document, created by any member of the Internet community, that is given official status (an RFC number) by the IAB. Only some of these RFCs become standards, by an official vote of the IAB. An Internet Standard can be supplanted by a newer RFC, at which point it becomes an Historical Standard. RFCs may also be submitted that are only experimental and not intended for standardisation. In contrast to the ISO standards, the Internet standards can evolve fairly quickly, in some cases taking less than a year to go through the standardisation process.

An important characteristic of the Internet protocols is that they usually reflect actual implementation experience. That is, a standard is usually considered only after a number of people have had extensive experience with it and have demonstrated its viability. This emphasis encourages the development of standards that are straightforward and easy to test.

The US DoD has endorsed the IPS standards and has created comparable Military Standards as well as conformance suites for these standards. There are some differences between the Military Standards and the Internet Standards.

The TCP and IP protocols are used in many networks. These networks are often simply referred to as TCP/IP networks. Implementations of these protocols vary, so it is important to be aware of common practice when developing software for these networks. Different implementations support differing sets of features, and interpretations of many of the features varies. The referenced **XGIPS** document describes the features that should be supported in any implementation of this protocol suite.

3.3.2 IPS Physical Layer

A number of different physical layers are used in IPS networks although none of these are specified by RFCs. All of the physical layers described for the OSI networking standard can also be used in IPS networks. IEEE 802.3 and Ethernet CSMA/CD networks are probably the most common, but X.21/X.25, and IEEE 802.5 are also sometimes used.

3.3.3 IPS Data Link Layer

The data link layer in IPS is primarily concerned with mapping the network layer on top of the physical layer. There are RFCs for mapping Ethernet, IEEE 802, X.25, and FDDI LANs.

Table **3-13.** IPS Data Link Layer RFCs

RFC	Description
RFC-894	IP mapping to Ethernet
RFC-877	IP mapping to X.25 public data network
RFC-1042	IP mapping to IEEE 802.2
RFC-1188	IP mapping to FDDI
RFCF-826	Address resolution protocol.

RFC-826 specifies the Ethernet Address Resolution Protocol (ARP) for mapping between network layer addresses and Ethernet addresses. The ARP is used to provide address mapping for Ethernet data encapsulation and IEEE 802.2 data encapsulation.

3.3.4 IPS Network Layer

Internet Protocol (IP)

Table **3-14.** IPS Network Layer RFCs

RFC	Description
RFC-791	Internet Protocol (IP)
RFC-792	Internet Control Message Protocol (ICMP)

The Internet Protocol provides a connectionless and potentially unreliable service over packet-oriented network systems. The IP is normally used under the TCP layer, which provides a reliable connection-oriented service for applications. The IP provides two basic functions:

Addressing and Routing

Addresses are always 32 bits in length, with four possible formats. Class A, B, and C formats are intended for normal point-to-point data exchange while the usage of Class D addresses are the subject of continuing research.

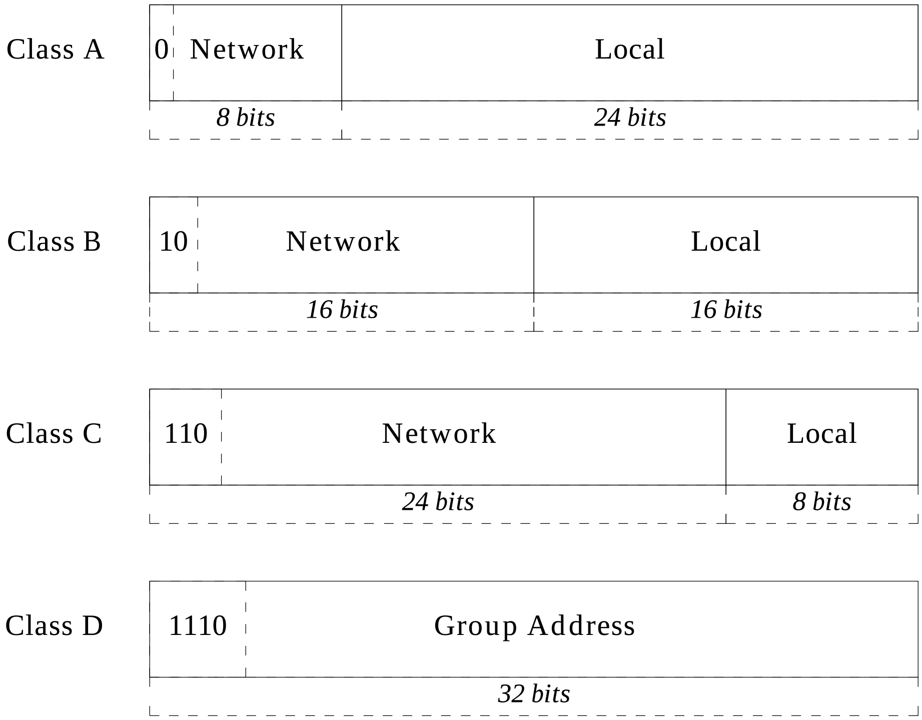

This scheme can accommodate a wide range of possible network configurations. Network addresses are normally written with the four octets expressed as decimal numbers separated by periods (for example 128.212.32.10). Conversion from host names to network addresses is done by a separate IP module. Each IP network node maintains a routing table which is used to forward packets.

Fragmentation and reassembly

IP datagrams can be as long as 65536 bytes, which may exceed the maximum packet size of the underlying network. If necessary, the IP layer can break long packets into smaller packets for transmission and reassemble them at the other end.

The Internet Protocol treats each datagram separately; each datagram is routed using the address contained in it. Because routing tables can change dynamically, successive datagrams to the same destination can take different routes. The upper layer (usually TCP) must guarantee that the datagrams are reassembled in the correct order.

Each datagram is composed of a header and the data. The data is simply passed intact between the communicating upper layers.

Internet Control Message (ICMP)

In addition to these datagrams, the network layer can send ICMPs. These messages, defined in RFC-792, are used to report problems on the network, diagnose the network, and deal with network addressing. TCP and UDP (described in the next section) can send ICMP messages as well.

3.3.5 IPS Transport Layer

Table **3-15.** IPS Transport Layer RFCs

RFC	Description
RFC-793	Transmission Control Protocol (TCP)
RFC-768	User Datagram Protocol (UDP)

Transmission Control Protocol (TCP)

The Transmission Control Protocol provides a reliable, connection-oriented service between applications. It is normally used over the IP layer. The TCP layer provides the following basic functions:

Basic Data Transfer
The TCP treats the data as a continuous stream of octets between the endpoints. This stream is broken into segments for transmission across the network. A *push* function is available to force the immediate transmission of any pending data.

Reliability
The TCP can recover from lost, erroneous, duplicated, or out-of-order data. This allows it to be used over an unreliable network. Each data segment is assigned a sequence number and protected with a checksum. The sequence number allows the receiving end to restore out-of-order data and to discard duplicated data. The receiving end also discards erroneous data with incorrect checksums. Each data segment has a sequence number that is individually acknowledged by the receiving end. A failure to acknowledge a segment means that the data was either lost during transmission or discarded at the other end because of a bad checksum. Segments that are not acknowledged by the timeout period are retransmitted.

Flow Control
Every TCP segment contains a window indication that tells the transmitting end how much data it may send past the current acknowledged segment. This allows the receiving end to control dynamically the amount of data being sent.

Multiplexing
The TCP provides for many simultaneous connections between any given pair of endpoints. A number of addresses (also called ports) can exist with a given host. These addresses are then associated with endpoints established in the applications. Each endpoint can support multiple connections. The TCP layer manages these multiple connections so that every data segment is delivered to its correct endpoint.

Connections
Every connection has three phases: establishment, data transfer, and release. The TCP must manage these phases over potentially unreliable networks. Reliable connection establishment is guaranteed by a three-way handshake. Reliable data transfer is guaranteed by the methods described above. Release occurs via a three-way handshake, which guarantees that all data has been transferred before the connection is dropped.

Precedence and Security

TCP users may optionally specify precedence and security parameters. Default values are used if these parameters are not supplied.

Each segment is composed of a header and the data. The data is received from the application and is passed intact to the communicating application.

User Data Protocol (UDP)

In addition to these TCP segments, the application can also send UDP messages. These messages, defined in RFC-768, allow datagrams to be sent across the network without first establishing a connection. They contain the necessary addressing information and have no confirmation or guaranteed delivery.

3.3.6 IPS Session Layer

There is no explicit session layer defined for IPS networks.

3.3.7 IPS Presentation Layer

There is no explicit presentation layer defined for IPS networks. Individual applications must implement any required presentation layer functions directly.

3.3.8 IPS Application Layer

The Internet Protocol Suite does not have an explicit application layer, instead a series of standard applications are defined. The following sections discuss those which have an equivalent in the OSI protocol suite.

Application Services

Unlike OSI, IPS does not define any generic application services. such services are left to individual applications to implement as appropriate.

3.3.9 Electronic Mail - IPS Internet Mail

Internet Mail is a widespread and popular IPS application. It provides electronic mail service to virtually every other Internet user. In addition, there are gateways from the Internet Mail system to other systems, such as BITNET and USENET, that allow messages to be passed across a wide variety of networks and systems. Within many organisations, Internet Mail has largely replaced paper mail.

Table **3-16.** IPS Electronic Mail RFCs

RFC	Description
RFC-821	Simple Mail Transfer Protocol (SMTP)
RFC-822	Format of text messages
RFC-920	Domain requirements

An Internet Mail message consists of a series of header lines and a blank line followed by the body of the message. The format of the message and its headers is defined in RFC-822. Each header starts with a reserved header-field name terminated by a colon, followed by the header value. The field names are restricted to printable ASCII characters (33-126), and the values may be any ASCII character except CR and LF. Each header line starts on a new line, although, for display purposes only, long lines may be wrapped to the next line with leading white space. The headers are divided into the following types:

Trace
These fields contain information that allows the mail path to be traced. This information can be used to evaluate mail routes and to return items to the originator.

Originator
These fields identify the sender, the sender's address and the return address.

Receiver
These fields contain the recipients addresses. Carbon copy and blind carbon copies are supported.

Reference
These fields contain identification for the mail item itself and identification of related messages.

Other
Subject and comment fields are available as well as a field to indicate encryption.

In addition, user-defined fields are available and may be used for any purpose. User-defined fields always start with *X-*. For example, this extension is used by gateways to the OSI Mail Handling System (MHS), to pass header fields that have no comparable RFC-822 field.

RFC-822 also defines name-domain addressing. Here are two examples:

einstein@physics.princeton.edu (Albert Einstein)
daley@chicago.illinois.us

Upper case and lower case are treated as equivalent. Everything inside parentheses is treated as a comment. In the example above, the first address is for a person named *einstein* whose top-level domain is *edu*, with *princeton* as a subdomain and *physics* as a further subdomain.

The policy for establishing and administering these domains is described in RFC-920, which also defines a restricted number of top-level domains. These include *COM* for commercial organisations, *EDU* for educational institutions, *GOV* for government, *MIL* for military, and *ORG* for other large organisations. In addition, the ISO standard two- or three-letter country codes can be used for top-level domains. Each top-level domain and many subdomains have a central administrative authority that registers the names used in that domain. Note that the address does not specify how the message is to be routed; that is determined by the mail system.

RFC-822 specifies 7-bit data for the body of the message. There is no provision for indicating the content type of the message so recipients may not be able to interpret non-ASCII messages. In fact, some mail nodes strip the high-order bit as the message passes through the system. This means that some mail text with 8-bit data, such as ISO 8859 encoded text, may be corrupted as it passes through the network. In addition, some mail nodes may even modify the body of the message by wrapping lines longer than 72 characters.

The transfer of mail from one host to another occurs via the SMTP protocol defined in RFC-821. The SMTP protocol is used after two hosts have established a connection for the transfer of mail. SMTP is simply a series of commands that allow message to be passed from one host to the next. The receiving host may deliver the item locally or pass it to another host. Each host may add *Received:* lines to the header. These lines can be examined at any point to evaluate the routing methods at each of the nodes. Each node may also manipulate the address to conform to the addressing conventions of the receiving host. In SMTP, this address is passed separately from the rest of the message. The SMTP node does not modify the *To:* field in the message.

3.3.10 Directory Service - IPS Domain Name Service

Table **3-17.** Internet Standards

RFC	Description
RFC-1034	Domain names - concepts and facilities
RFC-1035	Domain names - implementation and specification

The IPS DNS provides a distributed, hierarchical name look-up service. Within an IPS network, the DNS is the only global mechanism used. The primary purpose of the DNS is to provide host domain name to internet address translation and the reverse mappings. While the DNS does have support for finding other information, the other services, such as mail box look-up and the list of protocols supported, are not widely used. There may also be local options for non-DNS name look-up such as /etc/hosts and Yellow Pages.

3.3.11 Remote File Access - IPS File Transfer Protocol (FTP)

Table **3-18.** IPS File Transfer RFCs

RFC	Description
RFC-959	File Transfer Protocol (FTP)
RFC-783	Trivial File Transfer Protocol (TFTP)

The FTP allows users to interact with file systems on remote machines. Its primary purpose is to allow transfer of files between hosts, although other operations are available. Files may be transferred between the local host and a remote host, or between two remote hosts. The FTP uses the Telnet protocol (described in the next section) to pass commands to the hosts. Any file transfers then take place over separate network connections that are established for that purpose. There are many FTP commands; not all are widely implemented. They are described below in groups:

Access Control
These commands verify the user, account and password, and allow the user to change directories and file-system structures.

Transfer Parameters
These commands allow display and manipulation of the file-transfer parameters. The parameters indicate the type and structure of the file being transferred and control the transfer mode.

FTP Service

This series of commands allows direct manipulation of the files. This includes starting and terminating file transfers; creating and removing files and directories; and displaying directories, status, and help.

The FTP is designed to support file transfer across a variety of systems, where each system may have different file structures and different local representations of data. After the parameters are set correctly, appropriate conversions take place as the file is being transferred.

Every command generates at least one reply, which is prefixed by a number. This number informs the local user or program about the state of the server. The state diagrams for the commands and the resulting replies are given in RFC-959.

3.3.12 Terminal Service - IPS Telnet

Table **3-19.** IPS Terminal Service RFCs

RFC	Description
RFC-854	Telnet protocol
RFC-1123	Requirements for support of Telnet options

Telnet provides a general, terminal-oriented communications facility between two endpoints. The Telnet connection simply passes a stream of 8-bit bytes in either direction. Commands may be inserted into this byte stream by the Telnet process at one end and interpreted by the Telnet process at the other end. The display device (referred to as the NVT printer in RFC-854) is assumed to be able to display the 95 printable ASCII characters. In addition, the display device is expected to interpret eight different control characters. All other control characters may be ignored.

The Telnet server is expected to be able to respond to a variety of commands. These commands are always preceded by the IAC character (255) which is the only reserved character in the data stream. A character with a value of 255 may be passed intact from end to end by inserting two IAC characters in a row. Some Telnet commands control display functions, such as erase line and erase character, while others negotiate terminal options. The currently recommended options are specified in a series of RFCs. A complete list of these RFCs is included in RFC-1123, and current practice is described in the referenced **XGIPS** document. These options control many Telnet features, including terminal type, binary data, echoing, and window size.

3.3.13 IPS Systems Management

Within the Internet Protocol Suite, systems management protocols and products are currently the subject of a considerable amount of attention. This has been prompted by the need to manage the rapidly expanding Internet and also by the need to plan for the migration of significant networks, such as the Internet, to OSI protocols and for coexistence between IPS and OSI networks. The IAB has adopted a strategy which sanctions two management protocols: SNMP, based upon an existing protocol used for management of IPS networks; and CMOT, providing OSI systems management using the transport service provided by TCP/IP. To allow both approaches to be developed while limiting the effect on the software and systems being managed, the definition of the IPS resources to be managed is defined separately, and both protocols must use this definition. This strategy has been overtaken by events; commercial implementors have embraced the SNMP protocol enthusiastically and largely ignored the rival CMOT protocol.

Table **3-20.** IPS Systems Management Standards

RFC	Description
RFC-1155	Structure and identification of management information objects
RFC-1156	Management information base for TCP / IP
RFC-1157	Simple Network Management Protocol (SNMP)
RFC-1189	CMIP over TCP / IP (CMOT)

Management Information Base (MIB)

Both IPS management protocols model the resources of the managed system in a Management Information Base (MIB). Management operations on these conceptual resources are mapped locally into appropriate interactions with the real system's resources. The MIB is defined by RFC-1155 and RFC-1156, which describe a highly restricted syntax for object definition and define the actual objects required to manage the protocols which make up the transport and lower layers of the IPS.

Simple Network Management Protocol (SNMP)

Defined in RFC-1157, the SNMP architecture models a network as a set of *network elements* managed by a set of *network management stations*. The following figure gives a simplified view of the model.

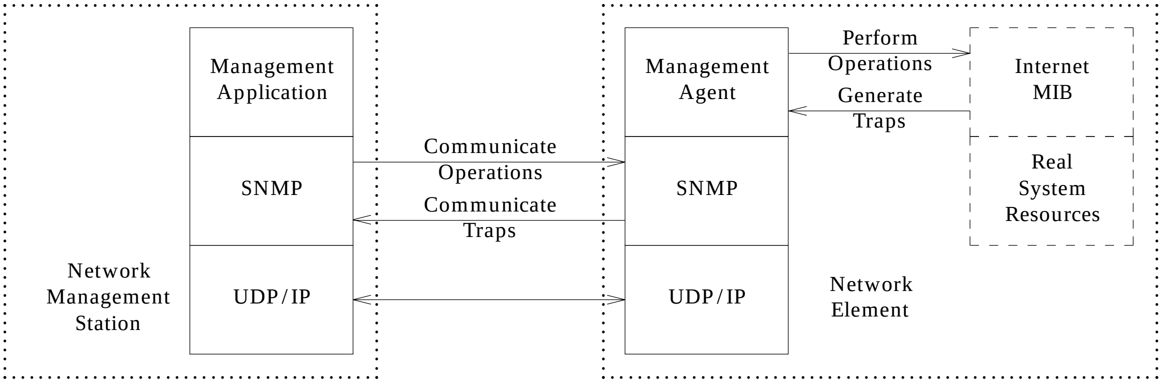

The protocol has three architectural goals, designed to minimise the impact of management software on systems which may have little spare processing or memory capacity:

- A strictly limited set of management functions is defined, a network management station may either *set* or *get* one or more attributes at a network element. The *get next* operation allows the elements of a table to be accessed. In addition a network element may send a *trap* to one or more network management stations, giving enough information to allow the manager to poll the network element for more details.

- The syntax of the management information, defined by the SMI, is a highly restricted subset of the ASN.1 abstract syntax, selected for its simple processing requirements on a wide range of common processing architectures.
- SNMP requires only a datagram transport service, in this case provided by UDP. This limits the complexity of additional protocol layers which may have to be introduced into a low-level device to enable network management.

SNMP supports access control, based upon a *community string* carried in each message, restricting how a system's objects are accessed. SNMP also supports the concept of a *proxy agent*, allowing a managed system to forward management requests to a system which uses a different management protocol or is part of a separate management domain.

Common Management Information Services and Protocols for the Internet (CMOT)

Defined in RFC 1189, CMOT consists of a set of agreements, some presented explicitly and some pointed to in other documents (such as the OIW Stable Agreements, published by NIST), by which the ISO systems management protocol may be run in an IPS environment. The protocol is run on top of a Lightweight Presentation Protocol (LPP) defining two levels of service; one using the TCP transport protocol and the other using UDP. The following figure shows how CMIS / CMIP fits into the TCP / IP environment.

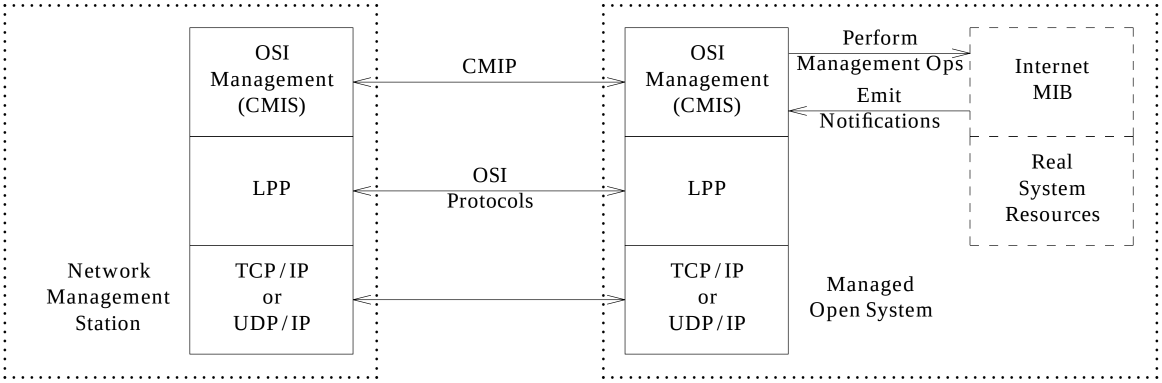

It is likely that the CMOT protocol will disappear completely, being replaced by standard CMIS / CMIP running over a hybrid stack which provides a a full implementation of the OSI upper layers with a convergence protocol providing TP0 transport layer services on top of TCP.

3.3.14 IPS Security

The original design of the IPS protocols did not address security issues. Since then, both the IETF and the DoD have done work on IPS security, defining a security option at the IP layer. This option can be used to control the routing of a single datagram to ensure that it only passes over links with an appropriate level of protection. The definition of the format and use of this option is in the hands of the DoD. There are no current RFCs on the subject, although there is an industry group which is working with the IETF to try and define its use in implementing secure workstations.

Other security mechanisms must operate at the physical layer (encrypting IP packets for secure transmission) or within the applications themselves (perhaps using some of the emerging standards for public and private key encryption). There are no IPS-specific standards for such mechanisms.

3.4 FUNCTIONAL COMPARISON OF IPS AND OSI

The IPS and OSI protocols have very different development histories and some differing underlying requirements. Nevertheless, there are many basic similarities between the two systems.

The similarities between these two systems form the basis for most of the coexistence and migration solutions described in the next chapter. The most important similarities are in the following areas:

- Both IPS and OSI networks can share the same physical links. Ethernet, IEEE 802.3, and ISO 8802-3 can run on the same LAN without interference, and X.25 can be used at the lower layers in both LANs and WANs.
- At the transport layer, the OSI Class 4 service is similar to the service provided by TCP in IPS. The interface to these services is referred to as the *Transport Service Interface.*
- Many OSI and IPS applications are comparable to each other. For example, MHS and IPS mail (RFC-822 and RFC-1098) provide similar services and also structure their data in similar ways.

The following figure shows an approximate mapping of the OSI protocols onto those of IPS, using the the OSI Basic Reference Model as a guide.

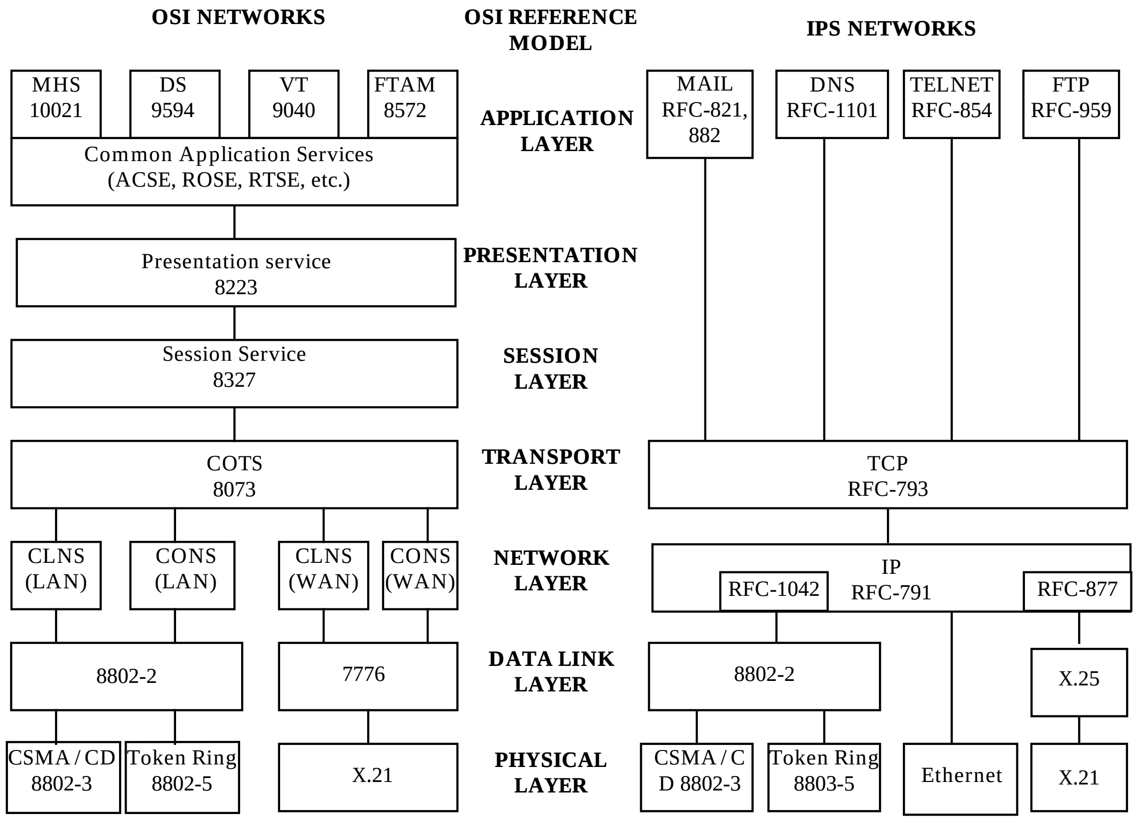

Figure **3-1.** Comparison of OSI and IPS Layers

3.4.1 Ethernet, IEEE 802.3, and ISO 8802-3

IPS networks commonly run on Ethernet or IEEE 802.3 LANs, while OSI requires conformance to ISO 8802-3. All of these LANs use the same physical layer, which permits building wiring and network interface cards to be retained when migrating from IPS to OSI.

Also, each of these standards specifies a unique data-packet format that allows them to coexist on the same LAN. The following figure illustrates the LAN packet level differences between these protocols:

Ethernet (IPS)	Destination Address	Source Address	Type (>1536)
	6 bytes	*6 bytes*	*2 bytes*

IEEE 802.3 (IPS)	Destination Address	Source Address	Length (0 - 1500)	LLC Header (DSAP = 0xAA)
	2 or 6 bytes	*2 or 6 bytes*	*2 bytes*	*3 bytes*

SNAP Header	PIF (OUI = 0x0D 0x00 0x00, LOCAL = Ether-Type)
	5 bytes

ISO-8802-3 (OSI)	Destination Address	Source Address	Length (0 - 1500)	LLC Header (DSAP = 0xFE)
	2 or 6 bytes	*2 or 6 bytes*	*2 bytes*	*3 bytes*

IPS packets running over Ethernet can be distinguished from OSI packets by the 2-byte value following the source address. For IPS packets, this value indicates the type of data that follows and always has assigned values greater than 1536 (0x600). For OSI packets, this value indicates the length of the data portion of the packet and must have a value of less than 1500.

IPS packets running over IEEE 802.3 can be distinguished from OSI packets by the DSAP value in the LLC header. This value is 0xFE for OSI indicating the presence of a network layer packet conforming to the ISO technical recommendation on protocol identification in the network layer - ISO/IEC TR 9577. For IPS the value is 0xAA, identifying the presence of the Sub Network Access Protocol (SNAP) defined by IEEE 802.1b. A SNAP packet consists of a Protocol Identification Field header (PIF) plus the the packet data. The first part of this header is the Organizationally Unique Identifier (OUI) used to identify a controlling organisation (in this case the IAB). The format of the second part of the PIF is defined by the organisation in question, the IAB have specified that it should contain the *Ether-Type* field normally found in an Ethernet packet. Thus, by using SNAP, both IP and ARP packets (and any other packets usually carried in Ethernet packets) can be carried in an 8802-2 LLC packet.

Thus as long as the data link layer driver can interpret all the above formats, it can route the packets to the appropriate protocol stack.

3.4.2 TCP and OSI Transport Class 4

Both TCP and OSI Transport Class 4 are designed to provide a reliable, connection-mode service over a potentially unreliable network layer. Much of the design of the ISO Transport Class 4 protocol (TP4) is based upon TCP. Both protocols have similar methods for creating, using, and releasing a connection and have similar header fields. However, TCP and OSI TP4 also have many differences. These differences have important implications when considering some coexistence and migration solutions. Significant mismatches between the two systems makes interoperability difficult or even impossible. The following list of differences is based upon analyses by Tannenbaum (1988), Groenbaek (1986), and a U.S. Department of Commerce publication (NTIS, 1988).

TPDU Types

TCP has only one type of Transport Protocol Data Unit (TPDU), while TP4 has nine. As a result, TCP TPDUs are longer, because each TPDU must include all fields whether they are needed or not. The minimum size of a TCP header is 20 bytes, in contrast to a minimum of five in TP4.

Connection Collision

If two TP4 endpoints try to set up a connection with each other at the same time, two independent connections results. Under these conditions, TCP either establishes a single connection or rejects both connections, depending upon the implementation. This is because the TCP uses the Transport Service Access Point address (TSAP) to identify the connection, while TP4 uses unique connection identifiers.

Addressing

TCP has a fixed, 4-octet addressing scheme with a 2-octet port number, while TP4 has an unrestricted address format, although a limit of 32 octets is usually imposed.

Expedited Data

Under TP4, expedited data is sent using a separate TPDU. The remote endpoint is informed that expedited data is available and may retrieve it at any point. TCP supports a similar notion but urgent data is a single byte inserted into the normal data stream rather than carried in a separate data packet through the network. The remote endpoint is informed that urgent data is available but it must read all queued normal data to receive it. For TP4, the expedited data service is negotiated when the connection is established whereas in TCP urgent data is always available.

Quality of Service

TCP has a single-octet field that specifies the quality of service from a set of well-defined service parameters. On the other hand, TP4 has a number of fields for specifying a much wider selection of service parameters that are negotiated when the connection is established. The TCP parameters are a subset of the TP4 parameters.

User Data

Unlike TP4, TCP has no provision for transferring user data while the connection is being established.

Data Stream

Both TP4 and TCP transfer data as a stream of octets. TP4 transfers these as a series of ordered messages. In some implementations of TCP, the sender can create segment-like boundaries in the data by using the *push* function. This is similar to the concatenation function that is available in the OSI session layer.

Flow Control

TCP provides explicit flow control via the *window size* parameter in the TPDU. With TP4, explicit flow control can be negotiated when the connection is set up, or the implicit flow-control method of the network layer can be used.

Sequence Numbers

TP4 provides protection against window size changes arriving out of order, by attaching sequence numbers to each message. TCP sequence numbers are on a per-octet basis. Data is reassembled and acknowledged only when it is contiguous.

Connection Release

TCP provides an orderly release of the data connection that prevents the loss of any data. In the OSI model, orderly release is taken care of by the session layer. As a result, any *close* request issued to TP4 may result in closing the connection before the data transfer is complete.

3.4.3 Internet Mail and MHS

Internet Mail and MHS organise the parts of the mail item somewhat differently. MHS places information that is relevant to the User Agent in the header fields of the message content and places information that is relevant to the Message Transfer Agent in the fields of the message envelope. Some information, such as originator and recipient addresses, may appear in both places. With Internet Mail, all of this information is placed in a single header. The following figure compares the formats of MHS and Internet mail items:

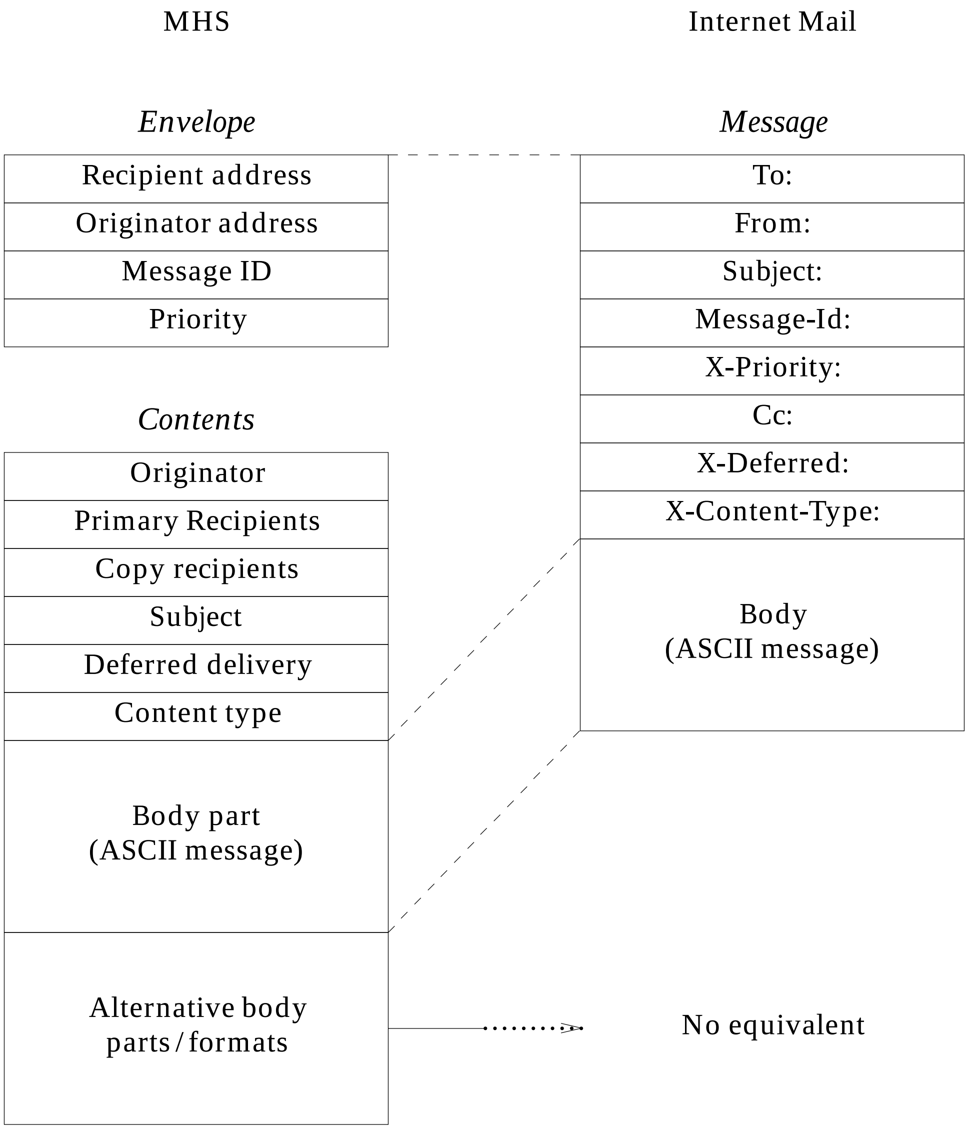

In many cases, there is a direct mapping between the Internet header fields and the MHS header fields. All but one of the Internet header fields (*Keywords*) have comparable MHS fields, but many MHS fields do not have comparable Internet fields. When MHS messages must be converted to the Internet format, the additional information can be carried in Internet Mail through the extended headers (headers that start with *X-*) or by inserting appropriate text into the body of the message.

Internet Mail uses the name@domain addressing format; these addresses are specified using a subset of the ASCII character set. MHS uses O/R addressing. Although the MHS and Internet addresses appear quite different, they are similar enough in concept that it is often possible to create a reasonable mapping between the two (RFC-987). An example mapping might be as follows:

IPS	**MHS**
darwin@biology.cambridge.ed.uk	Country: England Administrative domain name: Education Private domain name: Cambridge Organisation name: Biology Personal name: Darwin

An Internet Mail item can contain only one body, which is usually assumed to contain standard, human-readable, 7-bit text. RFC-822 defines the body to be made up of lines of ASCII text. In practice, other character sets are often used, although some Internet mail systems may strip the high-order bit. In contrast to Internet Mail, MHS allows an unlimited number of bodies in a mail item. Furthermore, each of these bodies can be encoded in a variety of ways; they can contain different character sets, as well as binary data such as voice or facsimile.

3.4.4 FTP and FTAM

FTP and FTAM both provide access to files on a remote system. Comparisons of the two systems fall into two areas: file formats and file access.

FTAM defines a general, flexible, and complex means of describing a file. FTP, on the other hand, defines only a few file types that are known to be in common use. As a result, all FTP file types have a comparable representation in FTAM, but the reverse is not true.

FTAM defines a set of services available for file access. FTAM does not define a user interface but provides command primitives from which an application can be built. State diagrams and tables define relationships between the command primitives. Primitives are available for the transfer of data, as well as the creation, modification, and deletion of files. Normally, a single user command results in the execution of several primitives.

FTP defines the commands available to the user. Commands are defined for the transfer of data, as well as the creation and deletion of files.

Some features on one system have no counterpart on the other system. For example, FTAM includes primitives for viewing and manipulating a large set of file attributes, while FTP has no comparable ability. Until recently, FTAM had no facilities for changing, creating or deleting directories. However, the US GOSIP has defined the NBS-9 document type that can be used to represent a UNIX directory. ISO has included an identical document type in the FTAM service definition.

3.4.5 Telnet and VT

Telnet and VT both provide remote login and basic terminal capabilities, but otherwise, they are very different. Telnet defines a simple printer device with basic carriage motion and line-feed capabilities that all implementations must support. It then allows more complex functions to be negotiated if available. Telnet provides some basic terminal functions and some options that can be negotiated. In practice, few implementations provide very many of the possible options. The referenced **XGIPS** document, provides a discussion of current common practice. VT, on the other hand, acts like a high-function video terminal that is mapped to the local terminal type. It has a rich set of functions that allow screen-based applications to interface with a wide variety of terminal types, one of which is that of the Telnet terminal.

3.4.6 Systems Management

Both SNMP and CMIS/CMIP management protocols operate from a broadly similar management framework. Both model the system being managed as a collection of managed objects, requiring the management application running in a managed system to map between logical operations on these objects and interactions with the system's actual resources. The two systems have significantly different ways of naming the objects and their attributes. OSI management's scope goes beyond the definition of a management protocol and management information, defining operations specific to a range of actual systems management tasks which are regarded as application matters by SNMP and are not standardised.

Both protocols use a tightly specified Structure of Management Information (SMI) to define the syntax for describing these objects. The two protocols, however, have very different design philosophies. The SNMP aims to reduce the load on the systems managed, using a minimal protocol and a highly restricted syntax for defining managed objects. The OSI protocol, on the other hand, is much more comprehensive, using a slightly more complex protocol and allowing the full generality of the ASN.1 language for the definition of managed objects in a MIB with a more complex structure. It is possible to represent SNMP-defined MIBs fully using the OSI syntax but not vice-versa.

SNMP is designed to run over an unreliable transport service such as a UDP, although it is not restricted to do so. CMIS/CMIP uses the services of the ACSE and ROSE service elements in a connection-mode environment.

The demands of the connection-mode protocol and the arbitrarily complex syntax of managed objects suggests that the OSI protocol imposes a higher load on the managed system. In addition, SNMP has a polling philosophy for event handling - unsolicited events carry the minimum of information and the management station must poll the appropriate agent for details of the event. CMIS/CMIP notifications, on the other hand, are self-contained, carrying complete information about the event being reported.

3.5 ADDITIONAL APPLICATIONS

This section covers a number of networking applications which do not form part of the IPS and OSI protocol suites but which are widely implemented and used in current IPS networks and therefore must be considered when migration and coexistence issues are being discussed.

3.5.1 X Window System

The X Window System is is used by many workstations to support windowing of applications on the workstation screen. The X Window System was originally developed as a research project at the Massachusetts Institute of Technology and is now controlled by the X-Consortium, an organisation which develops new versions of the protocol and distributes sample implementations.

The X Window System architecture views the workstation screen as a shared resource and defines an *X-server* which controls its displays (usually one) and input devices (usually a keyboard and a pointer device). An application process, known as an *X-client*, may make a connection to an X-server and open windows on the screen in which to interact with the user. The server allows the user to select any of the windows on the screen and to interact with the application running in it. In practice, the appearance and disposition of windows on the screen and the way in which the user interacts with them are controlled by a special application known as a window manager, which allows he use rot start and stop applications and to ontrol how windows are arranged on the screen.

Servers and clients may be distributed among the hosts on a network, a client addresses a server using a *server name* of the format *host:server.screen*. The protocol allows X-clients to send messages to each other and to be informed of significant events occurring in the server. X Windows requires a connection-oriented transport service and is widely implemented on top of TCP.

Version 11 of the protocol is in the process of being standardised by ANSI. The standard will contain a mapping for the X protocol over an OSI network. This mapping, known as *Xosi*, uses the services of ACSE to establish connections and the presentation layer to transfer X requests between servers and clients. It is expected that this specification will be progressed to the status of draft ANSI proposal and then submitted to ISO as a proposed standard for window management.

In order to produce an early agreement which can form the basis for implementation, a proposal which is aligned with the ANSI work is being progressed within EWOS. The EWOS standard includes agreements on how a client should include an X-Server name in the association request to facilitate an application gateway between the IPS and OSI protocol suites.

These specifications are not yet stable. Fundamental issues (such as how X requests are encoded in presentation PDUs) have yet to be resolved and there is currently no agreement on how the directory service is to be used by a client to resolve a server identifier into a specific presentation address.

There is some disagreement about the applicability of these standards to high throughput, low delay local area networks where the majority of current implementations run. It has been proposed that X could utilise a direct interface to the OSI transport service, this however is opposed as it does not conform to the OSI reference model. An alternative which is gaining support is the definition of a set of agreements which allow a light-weight (meaning low processing-overhead, high throughput) implementation of the upper layers. This would be achieved by restricting an implementation to use fixed headers for the presentation and session layer PDUs; the headers can be preformatted and reused during a connection. The EWOS standard specifies appropriately restricted use of ACSE, presentation and session layer services, and includes an annex offering guidance on how to implement light-weight application services.

It is likely in the near future that agreements and standards will be in place allowing interoperable X Window System products. In addition, agreements should allow the specification of application-layer gateways between servers and clients running over different protocol suites.

3.5.2 Transparent File Access

Transparent file access (TFA) refers to the ability to access file systems located at a remote system as though they were stored on a locally mounted disk. An application uses the normal system calls (such as *read*( ) and *write*( )) for accessing and managing files, the operating system arranges for the requests to be applied to files resident on a remote system.

TFA in an IPS Environment

For the IPS protocol suite TFA is a relatively mature technology with a number of proprietary products in active use for sharing files amongst UNIX hosts and making files available to PCs. In such products, *servers* export file systems, which *clients* can then mount as if they were on local devices. Requests to read a remotely mounted file are translated into request messages which are transmitted to the appropriate server. The server accesses the required file and sends a response message containing the required data.

There is considerable competition between these competing products, based upon such aspects as how much of the standard system call semantics are delivered for remote files, the range of file types supported remotely and access control. One key aspect is statelessness versus statefulness,

A stateful implementation maintains state in the file system server concerning which resources its clients have mounted, and file access state such as reference counts and locks. A stateful implementation is likely to support most of the UNIX file system and system call semantics. A stateful implementation may use a connection-mode protocol (in this case TCP), using the reliable delivery mechanism to help client and server monitor each other's health, however, this is not required. As the name suggests, in a *stateless* implementation the server is not aware of which clients have its systems mounted and does not maintain per-file state information. A stateless implementation is likely to be based upon a connectionless service (in this case UDP), providing its own mechanisms for recovery when it runs over a less reliable network service. A stateless implementation is likely to provide a lower level of support for UNIX file and system call semantics but provides much more resilience for the clients against server failure. A client may pend while a server is re-booted and resume file access as soon as the server host has restarted. All the popular TFA mechanisms are implemented using remote procedure calls.

TFA in an OSI Environment

No OSI applications currently exist to provide TFA in the OSI environment although at least one IPS-based networked file system can run over an OSI connection-mode transport service using the UNIX System V TLI interface. A standard for providing TFA functions using OSI application services is essential to allow homogeneous system groups to share files easily in OSI-based networks. A few points relevant to providing TFA in the OSI environment are discussed here. The services of FTAM provide most of the functions required to support transparent file access. The FTAM service primitives allow an application to access and manage files resident on a remote system. A remote file's data may be accessed at any level which is appropriate to the file's internal structure. FTAM's virtual filestore is designed specifically to hide the differences in the way data storage and access are arranged in a homogeneous collection of systems.

One area, important to TFA but not currently addressed by FTAM is mapping of file names and filestore directory structure between homogeneous systems. The FTAM virtual filestore has no concept of filestore directory structure, despite the prevalence of structures such as hierarchical directories in real filestores. In addition, the FTAM filename file attribute is defined as a simple, uninterpreted string. As discussed in **Section 3.2**.12, **Remote File Access - OSI File Transfer, Access and Management (FTAM)**, work is in progress to surmount these shortcomings. The IEEE P1003.8 Networking committee has considered what subset of its P1003.1 operating system interface may be supported for networked file systems. As part of this work they have considered the characteristics of an FTAM-based file system. The subcommittee's brief does not extend to considering a protocol for implementing FTAM-based TFA.

Thus most of the infrastructure is in place for the support of OSI-based TFA; what is needed is a standard for the application itself.

Comparison of IPS and OSI TFA

The lack of standards for providing TFA in an OSI environment make a comparison impossible. It is however worth considering how suitable the two protocol stacks are for supporting TFA. Given the importance of file sharing in workstation environments, such standards are essential if an OSI-based network is to provide the kind of TFA service which LAN users are used to and which is required in many current departmental networks.

It may be possible to use existing IPS protocols over OSI transport services. This violates the OSI Reference Model and must be considered as a migration technique rather than an OSI application. IPS-based protocols which operate over TCP may operate in an OSI environment using the connection-mode transport protocol.

Implementing TFA over the connection-oriented FTAM service is possible but requires agreements on protocols and data representations. Even with these requirements satisfied, FTAM-based TFA is likely to provide services more suited to limited file sharing, probably involving wide-area networks with low-throughput and relatively high error rates. FTAM requires a separate association (and hence presentation connection) per open file. Thus as currently used there would be a high overhead per accessed file. It is possible that an FTAM profile utilising one of the multiplexing transport classes (2 or 3) might be used to reduce the overheads by opening a single connection to each host with which files are shared.

3.5.3 Berkeley ‘r’ utilities

Most UNIX and derivative systems which participate in IPS-based local area networks support a set of programs known as the *‘r’ utilities.* These utilities (*rlogin* - remote login, *rcp* - remote copy, and *rsh* - remote shell) were developed as part of the Berkeley 4.xBSD UNIX derivative and are designed to emulate UNIX command style and syntax in a networked environment. They provide a number of useful features not present in the equivalent standard IPS utilities, in particular, password security may be bypassed among groups of trusted hosts by the use of *host equivalence* feature.

The functions and protocols implemented by these utilities are defined only by the manual pages and source code of the original implementations; no formal specifications exist. There is no standard activity aimed at providing these utilities in an OSI environment although equivalent functions may be available using applications based upon the standard protocols such as FTAM and VT.

The utilities are widely used in development environments; it is not clear how widespread their use is in commercial networks and therefore how significant they are to the subject addressed by this guide. Consequently, this guide limits itself to highlighting the existence of the utilities and the need to consider how equivalent functions are to be provided as part of a coexistence or migration strategy.

Chapter 4

Techniques

This section describes the available techniques for coexistence and migration between IPS and OSI. **Section 4.1**, **Overview**, gives a high-level description of the available techniques. This overview is intended for people who need a general familiarity with these techniques. Writers of network and application software should consult the remaining portions of this section that contain more detailed descriptions.

For simplicity, a single physical LAN is shown in most of the diagrams in this chapter. It is important to realise that, despite the single physical connection, the two protocol stacks are not compatible; OSI protocol stacks can only interoperate with other OSI protocol stacks, and IPS protocol stacks can only interoperate with other IPS protocol stacks. Because IPS and OSI use different, non-overlapping conventions at the data link layer, the two logical networks can operate completely independently of one another on the same physical network. Because OSI and IPS protocol stacks can usually share the same physical network, most of the coexistence and migration strategies discussed in this section can occur with little or no change of hardware.

The diagrams are simplified in other ways as well. The actual network may be very complex, with many sub-networks and interconnecting bridges and gateways. For the purposes of discussion in this chapter, a single LAN is shown in most of the diagrams.

4.1 OVERVIEW

This section gives a broad architectural overview of the various techniques. It serves both as a summary for those who only need a general understanding of the techniques and as a useful introduction for others needing more detailed information. For some readers, this section should suffice to identify the relevant techniques for their situation. Those readers can then consult the appropriate detailed explanations in later sections.

For every coexistence or migration problem, there may be more than one available technique. The choice of a technique depends on many factors including:

- The need to support specific applications.
- The need to connect to existing networks and applications.
- The availability of coexistence and migration solutions.

Each of the following techniques is designed to solve particular coexistence problems where portions of OSI and IPS networks must interoperate. Some of these solutions may add new capabilities but may have some limitations as well. Some of the techniques rely on modifications to end systems and affect every user in the network. Others provide gateway services to whole communities of users and avoid, at least in the short term, modifying the software running in the end systems. The advantages and disadvantages of each solution are presented in the detailed sections that follow.

4.1.1 Dual Stack

A dual stack exists whenever OSI and IPS networks coexist on the same system. Users of an end system which supports dual stacks may choose OSI or IPS applications as needed. This is a basic technique that supports a number of other techniques, such as switchable applications and application gateways. For example, a user would choose an OSI mail application to exchange mail with users on other OSI systems, and an IPS mail application to exchange mail with users on IPS networks. In this way, a single user could send mail to users whose systems only have a single OSI or IPS protocol stack. Note that users on single stack OSI systems cannot exchange mail with users on single stack IPS systems; that would require the application gateway described in the next section.

4.1.2 Application Gateway

Application gateways allow IPS applications to interact with comparable OSI applications. An application gateway is often used when OSI and IPS networks coexist, it runs on a system which supports dual protocol stacks. **Figure 4-1**, **Application Gateway**, illustrates this technique.

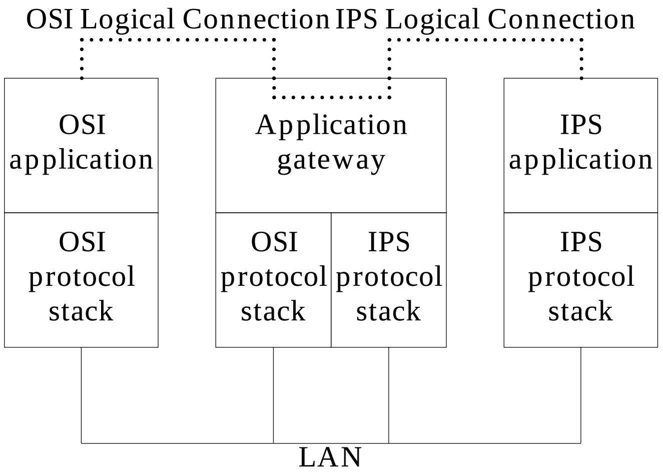

Figure **4-1.** Application Gateway

For example, users on an OSI network may need to exchange electronic mail with users on an IPS network. If MHS is used on the OSI network, and Internet Mail is used on the IPS network, a specialised gateway can be installed that provides a mail connection between these two groups of users. The gateway has separate logical connections to the two networks and understands the mail protocol for both networks. The gateway also performs any necessary transformation on the mail item as it passes from one mail system to the next.

This technique is most effective in cases where the services of the OSI and IPS applications are closely matched or where the OSI services are a proper superset of the IPS services. Where this is not the case, the services offered by the gateway can be restricted and transparency can be compromised.

4.1.3 Universal Applications

A universal application is one that can run in multiple networking environments. Two classes of universal applications are discussed here:

Switchable Applications

The figure below shows a switchable application, which operates in a dual-stack environment. It is a single application which selects the appropriate protocol stack at when a connection is created, depending on the remote endpoint being accessed. In some cases, the application can be designed so that the existence of separate networks is transparent to the user. **Figure 4-2**, **Switchable Application**, illustrates this technique.

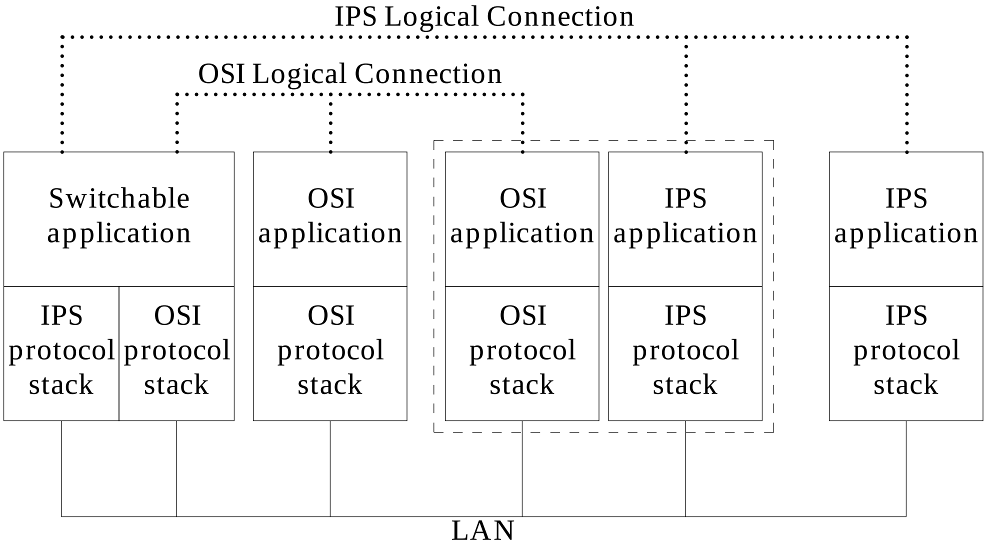

Figure **4-2.** Switchable Application

This technique is particularly applicable to client-server applications, where either client or server endpoints can be made switchable. A client example might be a switchable FTP utility which can connect both to IPS-based FTP servers and OSI-based FTAM servers (using X/Open's BSFT profile for implementing the FTP utility over FTAM services). A server example might be a LAN Manager server which accepts simultaneous connections from clients using IPS and OSI protocols as well as the common LAN server protocols.

Portable Applications

One way of easing the transition between IPS and OSI networks is to provide portable applications that are network-transparent. This means that the application presents essentially the same interface to the user, regardless of which type of network is being used. Unlike the switchable application described previously, this type of application deals with only one network on a given system. Portable applications are particularly useful as a part of some migration strategies, because users face minimal disruption when moving from one network to another. **Figure 4-3**, **Portable Application**, illustrates this technique.

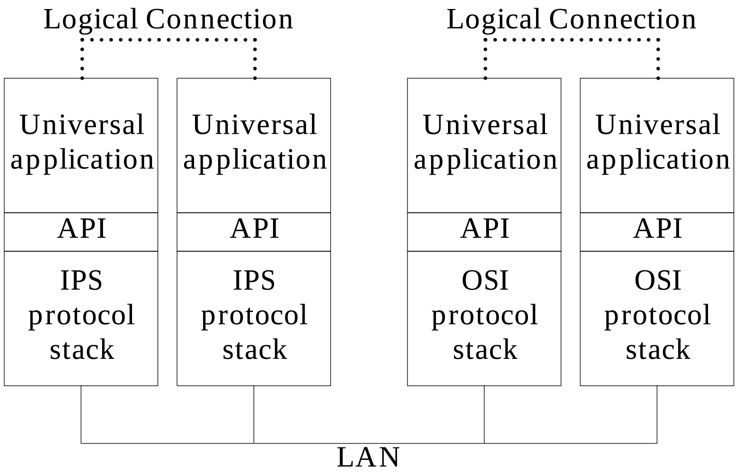

Figure **4-3.** Portable Application

4.1.4 Common API

A common application programming interface (API) is one way to simplify the creation of universal applications. It allows the application developer to write the application to a common, network independent interface, such as XTI, that can be used in both OSI and IPS environments (as long as the application does not make use of protocol-dependent features).

4.1.5 Hybrid Stacks

Hybrid stacks exist whenever an OSI or IPS application runs over a network of the other type. This approach may be desirable to keep the user interface constant even though the underlying network is changing. Hybrid stacks are particularly useful whenever a transport relay is used to connect IPS and OSI networks. (Transport relays are described in a later section.) Because IPS and OSI protocol stacks differ, hybrid applications require an additional software layer. This layer provides the appropriate interface between the application and the protocol stack. RFC-1006 describes a mapping that can be used to allow OSI applications to run over IPS networks. **Figure 4-4**, **Hybrid Stack - OSI Services over IPS**, and **Figure 4-5**, **Hybrid Stack - IPS Services over OSI**, illustrate two of the possible configurations.

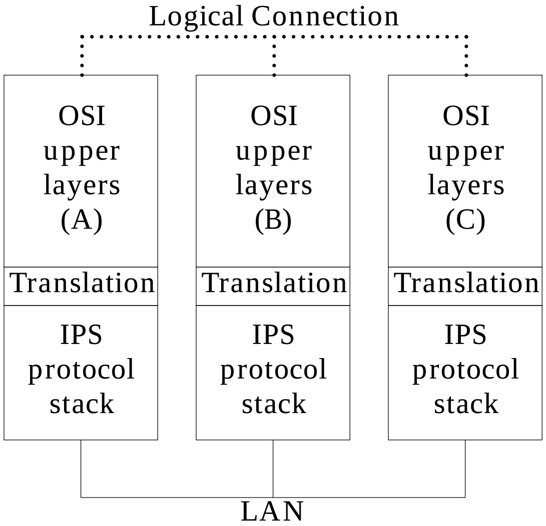

Figure **4-4.** Hybrid Stack - OSI Services over IPS

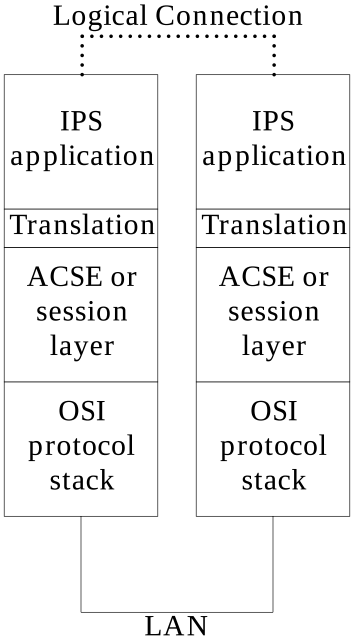

Figure **4-5.** Hybrid Stack - IPS Services over OSI

4.1.6 Network Service Tunnel

Network service tunnels connect two compatible networks by using an intermediate network of a different type. This technique can be used to connect end systems which use one protocol suite in an organisation whose infrastructure is based upon a different protocol suite. For example OSI-based end systems can be interconnected across an organisation's existing IPS network.

The example shown below gives one possible example, where two OSI networks are connected through an IPS network. By using the same scheme, IPS networks could also be connected through an OSI network. Note that the intermediate network only serves as a bridge; no applications services are available on the intermediate network. **Figure 4-6**, **Network Service Tunnel**, illustrates this technique.

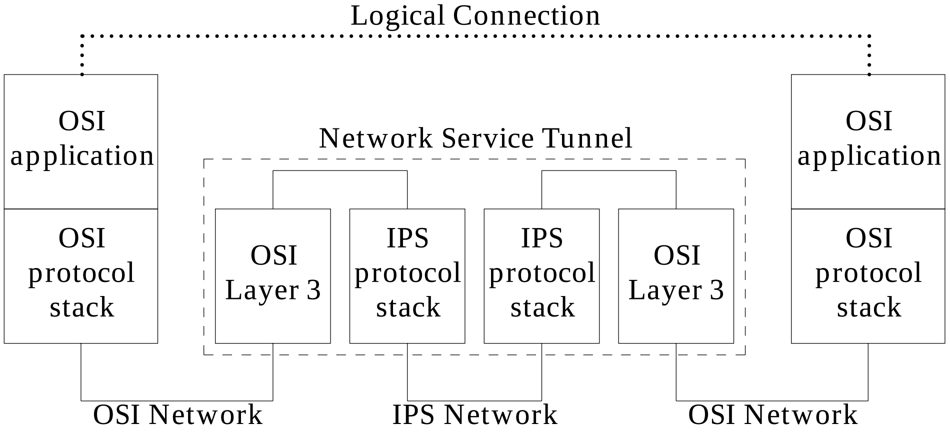

Figure **4-6.** Network Service Tunnel

4.1.7 Transport Relay

The transport relay allows compatible applications to communicate with one another when one of the applications is running on a hybrid network. **Figure 4-7**, **Transport Relay**, illustrates one example of this technique, with an OSI application running on top of an IPS network. For the applications on this hybrid network to communicate with compatible applications on an OSI network, a transport relay is required between the two networks.

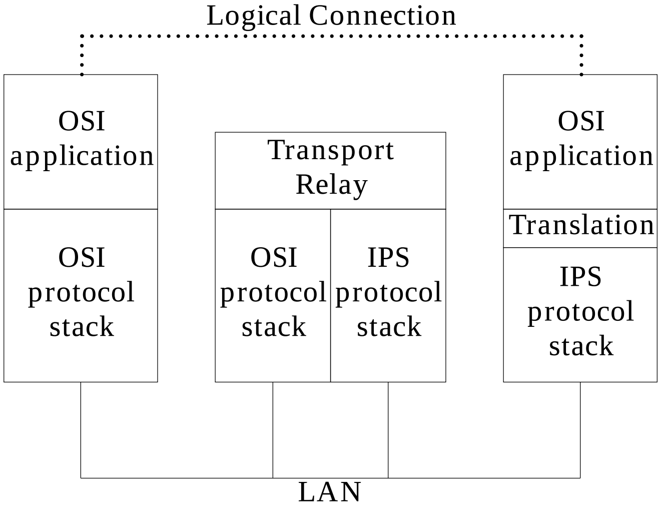

Figure **4-7.** Transport Relay

Note that a comparable situation exists (not shown), in which IPS applications running on an IPS network communicate with other IPS applications running on a hybrid OSI network.

4.1.8 Summary of Overview

The following table summarises the techniques according to the kinds of problems that they are designed to solve. For detailed explanations of the merits of each technique, consult the relevant sections in the remaining parts of this chapter.

Table **4-1.** Coexistence and Migration Solutions

Problem	**Solution**	**Comments**
Connect OSI and IPS applications.	Application gateway	A single gateway can often serve many networks, although each type of application needs a separate gateway process.
Provide a common user interface to both OSI and IPS services.	Portable Application	A single application can interface with OSI services on an OSI network, or with IPS services on an IPS network.
	Switchable application plus a dual stack	An application can interface to either an OSI or IPS network depending on the type of application at the other end.
Run OSI applications over IPS networks or IPS applications over OSI networks.	Hybrid Stack	A special application is needed that can only communicate with similar applications on the same network.
	Transport Relay	This allows applications running on a hybrid stack to communicate with compatible applications on a standard network.
Connect OSI networks over an intermediate IPS network or connect IPS networks over an intermediate OSI network.	Network Service Tunnel	Applications on remote networks are connected using an incompatible network as a bridge. No application services are available on the intermediate network.

The following sections describe each of these techniques in greater detail.

4.2 DUAL STACK

The most obvious method of moving users from an IPS environment to an OSI environment is to use a dual stack (IPS and OSI). Both IPS and OSI protocols are co-resident within the same system. This approach allows a user to choose which protocol stack to use and works well if all systems on the network can be modified to support both protocol stacks. Communication with single stack systems can be difficult, because the user must select the correct application depending on the type of network at the other end.

Applications available to the user can be IPS only, OSI only, or universal (use either protocol). Ideally, most applications would be universal, so that a user is not required to select a different version of the application depending on the network to which the connection is to be made.

Figure 4-8, **Dual Protocol Stacks**, illustrates how the protocol stacks may coexist in a host.

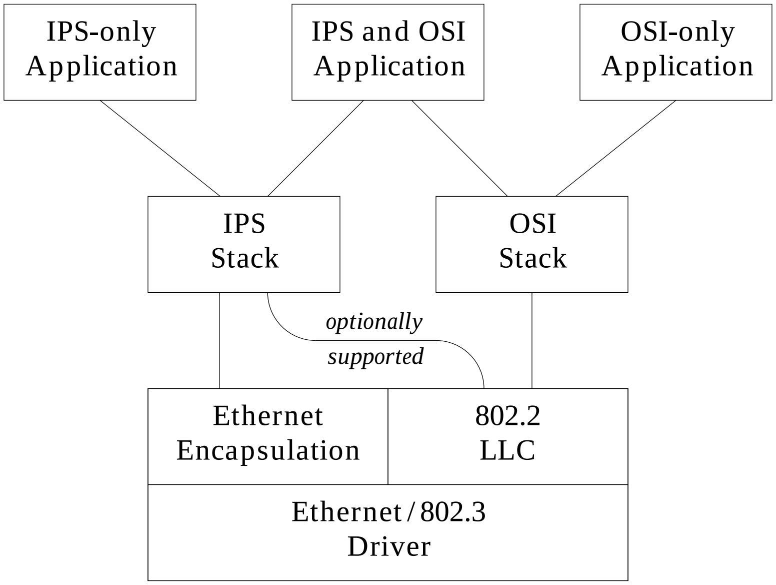

Figure **4-8.** Dual-protocol Stacks

4.2.1 Coexistence in the Same Host

Both the IPS and OSI protocol stacks can coexist on the same host at the same time and use the same network interface. For IEEE 802 LANs that implement the IEEE 802.2 Logical Link Control (LLC) Class 1 protocol, the IPS and OSI protocol stacks use different LSAP values. As a result, both protocols can use the same LAN interface without conflict.

For the special case of Ethernet / 802.3 networks, IPS systems typically use the LAN as an Ethernet, while OSI uses it as an 802.3 LAN. Although this may initially appear to be in conflict, in practice it is not a problem. Briefly, Ethernet uses the two octets following the six octets of destination address and the six octets of source address as a *type* field. 802.3 uses the same two octets as a *length* field. All of the XEROX or IEEE **registered** Ethernet types are values that are greater than the legal length values for 802.3 frames. This allows the device driver to determine the LLC protocol type by examining this 2-octet field.

While some IPS implementations support the use of 802.2 LLC Class 1, they are not commonly used. Interoperability with existing IPS systems on an Ethernet / 802.3 LAN requires this dual mode of operation. There is no equivalent need for dual interfaces at the data link layer with the other 802 LANs, because both IPS and OSI use the 802.2 LLC.

In order to facilitate the implementation of dual stacks, a common interface should be defined for LAN and other network interface device drivers.

4.2.2 Comments

Dual protocol stacks are one of the mechanisms that allow for the coexistence of IPS and OSI networks within the same host and, more importantly, within the same organisation. A Dual-Protocol Stack does not itself assist application or user portability, or interoperability between IPS and OSI systems, but it does provide a necessary platform upon which to build more complete solutions. A dual-stack system is a prerequisite for OSI and IPS interoperability for several of the other techniques discussed in this document, including application gateways and transport relays.

The presence of dual stacks implies that there are effectively two separate networks coexisting on the same physical network and on the dual-stack hosts. The two networks have their own management and configuration requirements, which results in greater systems management overhead.

Pros:

- The dual-stack approach provides all the services of both the IPS and OSI networks to users of the dual-stack system.
- Dual stacks are a necessary component of a number of other solutions discussed in this document. A complete migration strategy requires at least one dual-stack system on the network.

Cons:

- It is expensive, and not always possible, to implement dual stacks on all hosts on the network. Some systems may not have OSI implementations available, and others may have memory or other constraints that prevent multiple protocols from being resident concurrently. This is particularly true of systems which implement protocols on an front-end processor board.
- Where both dual- and single-stack systems are used, the user may be confused because some services are either not accessible on the single stack host or are accessible only via a relay or gateway host.

4.3 APPLICATION GATEWAY

An application gateway allows users to continue using their familiar applications. If an application user on one network needs to access the corresponding service on the other network, an application gateway on a dual-stack host between the two may act as an intermediary. Application gateways are of most use in environments where it is not practical or possible to convert all hosts to OSI or to support a dual-stack implementation on most hosts.

The operation of application gateways may be categorised as either store-and-forward or end-to-end. The store-and-forward technique is widely used in applications such as electronic mail, where the gateway can receive the entire message before converting and then forwarding it to another mail system. The end-to-end technique is more appropriate when there is interaction between the ends of the applications, such as during a remote login session.

With an application-level gateway, an IPS application interoperates with an OSI application service, and *vice versa*, at the application layer protocol level. That is, an IPS application on an IPS-only system can use the equivalent OSI service on an OSI-only system by connecting to an application-level gateway on a system with both IPS and OSI implementations.

4.3.1 Store-and-forward Gateways

A store-and-forward gateway first performs a whole transaction with one system and then performs an equivalent transaction on the other. For example, electronic-messaging gateways typically accept a connection from a remote system and then, if possible, accept a message intended for a system on a different network. After the message is entirely received, it is translated into the appropriate form for the other network and submitted to that network's messaging system. The store-and-forward version of the application gateway technique is illustrated in **Figure 4-9**, **Store-and-forward Gateway**.

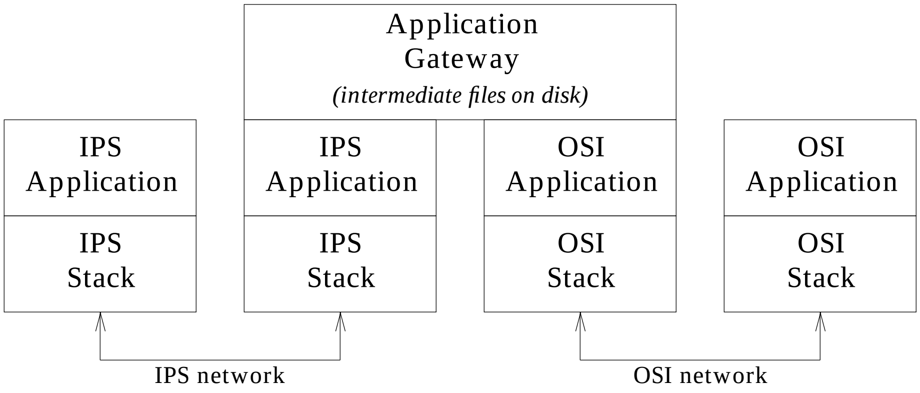

Figure **4-9.** Store-and-forward Gateway

An example of this type of application gateway is an MHS / RFC-822 mail gateway. The translations are carried out in accordance with RFC-987 and further clarified in RFC-1026. The gateway is not completely transparent to the users. For a user on the RFC-822 side, a message received from an MHS user typically has additional message header lines and some noticeable differences in the address format. An MHS user sees some RFC-822 headers moved into the body of the message to prevent the loss of information. The result is that the differences are visible to the users.

The connections between the gateway and the two hosts may either be consecutive or concurrent. If the connection is consecutive, a connection is set up between the host and gateway, the transfer takes place, and the connection is closed before the connection is set up with the next host and the next transfer takes place. If the connection is concurrent, the connections between the gateway and both hosts are set up before any transfer takes place. In both cases, the data is stored temporarily on the gateway before being transferred to the destination system.

When a message is sent from an RFC-822 service to MHS system with consecutive connections, the following occurs:

- The RFC-822 sending system establishes a connection to the RFC-822 service of the gateway.
- Using the standard protocol for the IPS mail system, the sending system sends the message.
- The gateway saves the message and acknowledges that it has received it.
- The sender disconnects from the gateway and is no longer involved in delivering the message.
- The gateway system converts the message into the MHS format and uses the MHS Message Transfer Agent protocol to put the message into the MHS Message Transfer Service.
- After the MHS MTS accepts the message, the gateway itself is no longer involved.
- If the MTS rejects the message, the gateway creates an error response message and sends it to the RFC-822 user as a mail message.

Table **4-2.** RFC-822 to MHS Gateway with Consecutive Connections

RFC-822 Mailer		**RFC-822/MHS Gateway**		**MHS MTA**
Connect	→	Accept Connection		
Submit Message	→	temporary storage		
	←	Acknowledge Message		
Close		convert message		
		Connect	→	Accept Connection
		Submit Message to MTA	→	
		Close		Deliver Message

When a message is sent from an RFC-822 service to MHS system with concurrent connections, the following occurs:

- The RFC-822 sending system establishes a connection to the RFC-822 service of the gateway.

- The gateway system establishes a connection with the destination MHS Message Transfer Service after the address has been determined.
- Using the standard protocol for the IPS mail system, the sending system sends the message.
- The gateway saves the message.
- The gateway system converts the message into the MHS format and uses the MHS Message Transfer Agent protocol to put the message into the MHS Message Transfer Service.
- After the MHS MTS accepts the message, the gateway confirms the transfer with the sending RFC-822 system.
- If the MTS rejects the message, the gateway reports the error to the sending RFC-822 system.
- The connections between the gateway and the two hosts are closed.

Table **4-3.** RFC-822 to MHS Gateway with Concurrent Connections

RFC-822 Mailer		**RFC-822/MHS Gateway**		**MHS MTA**
Connect	→	Accept Connection		
		Connect	→	Accept Connection
Submit Message	→	temporary storage		
		convert message		
		Submit Message to MTA	→	
	←	Acknowledge Message	←	Acknowledge Message
Close		Close		Close
				Deliver Message

Similar techniques may apply to other applications. This technique could be used to implement an FTP-FTAM gateway, but it would have some performance problems. For example, if a local FTP user attempts to copy a multimegabyte file from the remote system to the local system, the gateway must have sufficient storage to hold the entire file. The time for the entire transaction would be the total time for the two transfers, since data transfers would not be going on in parallel.

The temporary storage and the additional transfer time are not usually problems for electronic messaging, since messages are typically fairly small and immediate response is not required. The store-and-forward technique is not applicable to a remote login session.

4.3.2 End-to-end Gateways

An end-to-end gateway operates quite differently from a store-and-forward gateway. Rather than performing a complete transaction in one system and then translating and performing the equivalent on the other system, an end-to-end gateway translates and forwards primitives as they are received. The end-to-end version of this technique is illustrated in **Figure 4-10**, **End-to-end Gateway**.

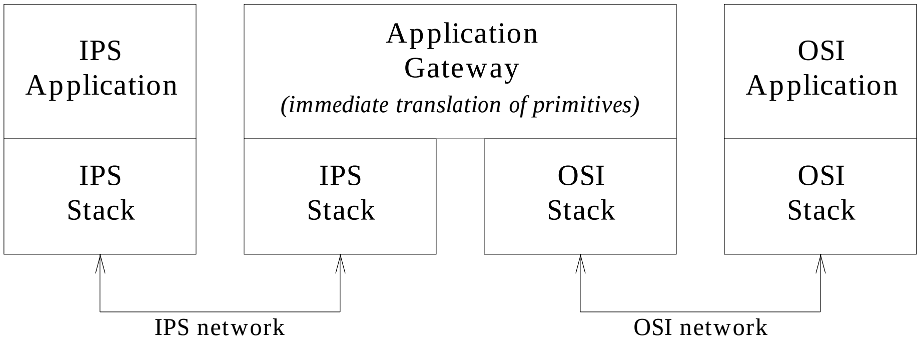

Figure **4-10.** End-to-end Gateway

An example of this type of application gateway is an FTP-FTAM gateway. The following events occur when an FTP user on an IPS system attempts to transfer a file to an OSI system that is running FTAM:

- The FTP user on the IPS system connects to the FTP-FTAM gateway.
- The gateway establishes a connection to the remote FTAM server and begins translating FTP primitives to the equivalent FTAM primitives. The FTP *USER* and *PASS* commands are translated into the FTAM F-INITIALIZE primitive.
- After the FTAM session has been fully established, the gateway indicates success or failure of user authentication by sending the appropriate response to the FTP user.

If the user then wants to send a file to the remote system, an FTP *STOR* command is sent to the gateway, which translates the command into the appropriate sequence of primitives (FTAM F-CREATE, F-OPEN, and so on) and sends them to the FTAM server. If the FTAM server accepts the request, the gateway indicates the status and begins the data transfer.

This is a simplified scenario; the details of the actual translation of primitives can be quite complex. If an application-protocol primitive does not directly map into primitives of the other application protocol, the gateway must attempt to emulate the primitive or must reject the request.

Table **4-4.** FTP to FTAM Gateway

FTP		**FTP-FTAM Gateway**		**FTAM**
Connect	→	Accept Connection		
USER ftamuser@ftamsite	→	parse username		
		Connect	→	Accept Connection
	←	user logged in	←	user logged in
. . .		. . .		. . .
STOR filename	→	F-SELECT	→	open target file
open file		F-OPEN		
		F-WRITE		
send data	→	F-DATA	→	write data
close file	→	F-DATA-END	→	close file
		F-TRANSFER-END		
		F-CLOSE		
		F-DESELECT		

The end-to-end technique is the method of choice for a remote login gateway. A Telnet - VTP gateway translates primitives of one protocol to the corresponding primitives of the other protocol as they are received.

Performance degradation is particularly noticeable using this type of application gateway; a remote login user is more aware of gateway performance than users of some other applications, because each key-stroke may generate a response. In addition, for some protocols, delays introduced by the intermediate node may cause timeouts to occur at the end-nodes, requiring some kind of subterfuge by the gateway to avoid this.

In the United States, the National Institute of Science and Technology (NIST) has recommended application-level gateways as part of the preferred migration path from IPS-based networks to OSI. NIST is also defining the protocol service mappings necessary for some applications gateways. The MHS / SMTP specification, in Estelle, is available.

4.3.3 Specifying the Remote System to the Gateway

Once an application gateway is available, users need to know how to go through it to access a remote system. Each gateway has its own mechanism but all must observe one rule: use of the gateway must not require changes to any system or software other than the gateway itself.

A common technique for providing gateway access is to extend the user name. For example, an IPS FTP user could directly connect to the FTP-FTAM gateway using the following login sequence:

```
ftp ftp-ftam-gateway
Name: ftamuser@ftamhost
Password:
```

In the example above, the address of the remote FTAM server is specified by appending **@ftamhost** to the name of the user on the FTAM host; files on the FTP-FTAM gateway are accessed by not specifying an FTAM address. An equivalent sequence is used by an FTAM user to access a remote FTP server.

In the case of a Telnet-VTP gateway, the gateway can prompt for the remote host and allow the native login mechanisms to work normally after the session is established.

4.3.4 Comments

Application gateways enable users familiar with the applications of one type of network may use the services on the other type of network without difficulty. This increases the number of services available to users, and increases the level of interoperability between otherwise dissimilar networks.

Pros:

- Users can take their time to learn the new OSI applications, because the IPS applications are still available.
- As systems are converted from IPS- to OSI-based protocols, full connectivity is maintained without expending significant effort on the IPS side.
- Users on either type of system can access the other type of system with no new applications needed on either system.

Cons:

- Overall performance may be a problem with some application protocols, especially those that need end-to-end behaviour.
- It is difficult, and in some cases impossible, to implement all the functions of a given application across an application gateway, and some information may be lost.
- Much work is involved in developing the application gateways. The problem of unwanted protocol time-outs may increase gateway complexity. Fortunately, only a few gateways are needed on each network.
- Users have little incentive to learn the new OSI applications, therefore prolonging the migration process.

4.3.5 Gateway Server

An alternative use for application gateways is to provide indirect access to OSI services for users on an IPS network. This differs from the above types of gateway because there is no mapping performed between the services of the OSI and IPS applications. In fact this technique applies in cases where there is no equivalent IPS application. Instead an OSI application is implemented at the gateway server and remote clients access the application's services using a simple remote access protocol over the IPS network. This is similar to the way in which PC network clients, which cannot access the IPS protocols directly, access IPS services which reside on a LAN server.

For example, an application resident on an IPS-based workstation may wish to access the OSI Directory Services. The application's directory access requests are converted into remote procedure calls to a real Directory User Agent which resides on gateway server.

The remote access protocol may be application-specific, passing high-level requests to the gateway server. A single request such as ‘‘Lookup name’’ is passed directly to the gateway server where it results in a potentially complex series of actions by the OSI application. Alternatively, an API to the OSI services can be implemented remotely, allowing the application full access to the API’s functions. The latter case allows the application to use a standard interface which aids portability, but each call to an API function can result in an RPC request to the gateway server.

Pros:

- Provides access to OSI services from workstations which either have not yet been migrated to use the OSI protocols, or are unable to support them at all.

Cons:

- As for the other application gateway, response times may suffer as a result of the additional link in the service access path.

4.4 UNIVERSAL APPLICATIONS

As discussed in **Section 4.1**, **Overview**, there are two classes of universal application, each of which is addressed below. In addition, the applicability of Common APIs to these applications is discussed.

4.4.1 Switchable Applications

A switchable application runs in a dual-stack system, and is able to interface to either protocol stack, depending on the network connectivity of the endpoint to which a connection is requested. The technique can be applied to both the client and server ends of an application. A server application is designed to service requests from either protocol stack, allowing single stack clients from either network to make requests. For a client, each workstation must have a dual stack in order to access the required services on either network.

A client application must have a mechanism for determining to which network the user intends to connect. If the dual nature is to be transparent to the user, the differences in the underlying protocols must be masked by the application, this requires careful selection of the services used and may also require significant programming effort within the application itself. An automatic method of determining which protocol stack to use is necessary. The techniques for achieving transparency have not yet been fully developed but access to a common host and service directory is essential in order to avoid the overhead of multiple directory requests when obtaining the network address of the required connection endpoint.

An application that is written to a common application program interface (API) for both protocols is more likely to achieve the desired transparency. It is possible to develop switchable applications that use multiple APIs, each with its own directory service, but this approach would be more complex.

Pros:

- A user on a dual-stack workstation has a single interface for accessing services on either network.
- A single server can support clients on either network.
- Allows access to services on either network without the performance penalties associated with the use of application gateways
- A client application may be able to determine which stack to use for a connection without the user being aware of the choice.

Cons:

- The user of a dual-stack system may have to use different name syntaxes to identify the required service depending on which network it resides.
- For switchable client applications, additional stack must be installed in all workstations that support the application.
- The performance penalties associated with dual-stack operation may reduce the performance of the workstation or server.

- A client system may experience additional delays associated with accessing two directory services.

4.4.2 Portable Applications

A portable application is a means of providing continuity of function and interface to users whilst the underlying networking software is being migrated from one protocol suite to another. The comments given above about network transparency for switchable applications apply equally well to these applications, except that there is no need to select the stack to be used for a connection.

For portable applications, a common API can be used to increase application portability and ensure transparency of the underlying protocol suite.

Pros:

- Allows the underlying network to be migrated to OSI with minimal disruption to the user's environment.

Cons:

- May restrict access to the additional functions provided by the equivalent OSI applications.
- Significant effort is required to port an IPS application to the OSI environment.

4.5 COMMON API

Using an application program interface (API) common to both the IPS and OSI protocol suites can simplify porting of applications between the two protocol suites or the creation of a switchable application. It also maximises the reusability of network applications. A common API can exist at the transport layer or at higher-layer interfaces such as remote procedure calls (RPC) or application-specific interfaces. Care must be taken in choosing an appropriate form of API and in selecting the network services used if the goal is to use the same application over multiple protocol stacks.

Transport API

The XTI, TLI, and BSD sockets interfaces are examples of a transport-level API. As long as applications developers are careful to use a minimum of features provided by the underlying transport protocol, any transport-level API allows an application to be reasonably portable between IPS and OSI protocol stacks.

In addition to having a common set of primitives for accessing the transport protocol, it is essential to have a common directory service for performing address and service look-up. This is especially important if the same application is expected to work on multiple protocols concurrently. The current specifications of XTI and TLI do not contain such a directory service. The version of TLI in the System V Release 4 UNIX system does include a directory service mechanism. The BSD socket interface has host and service look-up mechanisms, but their use has been fairly specific to the IPS addressing format.

The lack of a standard API for directory service causes developers to implement their own. For example, the System V Release 3 versions of RFS and UUCP each have their own mechanisms for address look-up in different locations and have different formats for the information.

It is also advisable to avoid the use of protocol-specific options. Frequently, options of one protocol suite are not applicable to the other protocol suite, or the formats and values of the options may differ.

Pros:

- It is possible to develop applications that operate over any transport protocol or can be ported without difficulty for use over multiple transport protocols.

Cons:

- Transport protocols may differ in their services. In addition, protocols whose services map similarly to the API primitives may differ semantically. Developers must use care to prevent seemingly identical applications from behaving differently when run on different protocol stacks.

- The task of porting all of an organisation's own applications to run over the new protocol stack is a major undertaking.

Remote Procedure Call API

A common API to a mechanism for a remote procedure call (RPC) may be preferable for distributed applications that are portable across multiple network protocol stacks. An RPC mechanism can provide the application writer with a uniform and well-defined set of services across multiple protocol suites. There are a number of candidates for a standard RPC mechanism.

Pros:

- The specific features of the underlying transport protocols are hidden from the developer. A consistent API is presented that should have the same semantics regardless of the underlying transport protocol.

Cons:

- Choice of a specific RPC mechanism implies use of a corresponding RPC protocol and services. There are currently no standard RPC mechanisms.
- Not all applications can be implemented with a single RPC mechanism.
- RPCs are not useful for implementing the traditional IPS protocols such as VT and FTP.

ACSE/Presentation API

A possible approach, which has been utilised by proprietary systems but which has yet to be taken up by industry or standards bodies, is to access the connection management and data transfer services of the ACSE and Presentation Layer via a transport-level interface such as the XTI API. This would provide IPS applications with an OSI-based service which has fewer differences from the IPS transport service and to which applications can be ported with fewer changes than required by a direct interface to the primitives of these OSI services.

Application-specific API

Application-specific APIs are being developed for the OSI application protocols. Both proprietary and industry-standard APIs for FTAM, MHS (X.400), and DS (X.500), APIs either have been developed or are near completion. Application-specific APIs for other protocols are expected to follow. Much of this work is being driven by industry-sponsored bodies such as X/Open.

While application-specific APIs are not particularly applicable to supporting IPS applications in an OSI environment, however, it is possible that such an API could be used to support an OSI application using both IPS and OSI application services (for example, a CMISE API could support OSI systems management over both CMIP and SNMP protocols).

The next section discusses an alternative method of supporting applications over alternative protocol stacks.

4.6 HYBRID STACK

The hybrid protocol stack approach uses parts of one protocol stack on top of the other protocol stack. Basically, the hybrid stack approach provides for using IPS-based applications over an OSI protocol stack or for using OSI-based applications over an IPS protocol stack. Both types of hybrid stacks are covered in more detail below.

4.6.1 OSI Services over IPS

Many people in the Internet and academic arenas are implementing OSI applications over IPS. This technique may have application in commercial environments too, allowing an organisation to gain experience with the new OSI applications and benefit from their richer functions whilst leaving the underlying network infrastructure un-disturbed. This can ease the eventual migration to OSI-based networks, because the replacement of the lower network layers is largely imperceptible to the user.

The ISO Development Environment (ISODE) is one implementation of this technique. It currently provides implementations of FTAM, Virtual Terminal, and Directory Services as applications over IPS. An MHS implementation is expected in a future release.

There are no commercial implementations of this technique at present, current implementations, such as ISODE, have the status of demonstration and research tools and are unlikely to be of product quality (in terms of performance, reliability and so on).

An alternative, although less general, approach is to port each application directly to the IPS environment and not spend the time to make the IPS environment look like the OSI environment. This approach requires more work and is not a general solution.

4.6.2 Implementation of OSI Services over IPS

This approach provides an implementation of ISO Transport Class 0 (TP0) on top of a TCP virtual circuit and then implements the upper layers of the ISO protocol stack above the TP service. A proposed standard method of implementing TP0 services has been published in RFC-1006.

This technique uses TCP as the underlying reliable network layer protocol in place of the X.25 service normally used as the network service for TP0. The TP0 implementation typically issues network layer primitives, which are mapped onto the underlying TCP. For example, if TP0 wants to initiate a network layer connect request, a TCP active open is completed and a TCP connect request is initiated. The completion of the TCP connect is translated into a network connection confirmation. Similarly, other TP0 requests to the network layer are converted to appropriate TCP service requests. All TP0 TPDUs are sent as TCP data. The receiving TP0 is responsible for extracting complete TP0 TPDUs from the TCP data stream.

OSI Services above TP0

Supporting TP0 is not sufficient for supporting OSI applications, it is a transport platform upon which the OSI upper layers can be implemented. Assuming that the TP0-over-TCP implementation provides the same transport-service interface to the system that an ISO TP4 implementation does, implementation of the OSI upper layers should be nearly, if not exactly, identical in both cases. OSI applications that were intended to run over either TP0 or TP4 should require no changes to run over the TP0-over-TCP protocol stack. Thus, the OSI applications are portable across both IPS and OSI systems.

By providing OSI services in this manner, users of the IPS system have an OSI development environment that allows them either to port existing applications to OSI, or to develop new applications for an OSI network prior to moving to an OSI-only system. **Figure 4-11**, **Forms of OSI Transport Service**, illustrates how the RFS-1006 convergence protocol is used to make the TCP transport service support OSI upper-layer protocols.

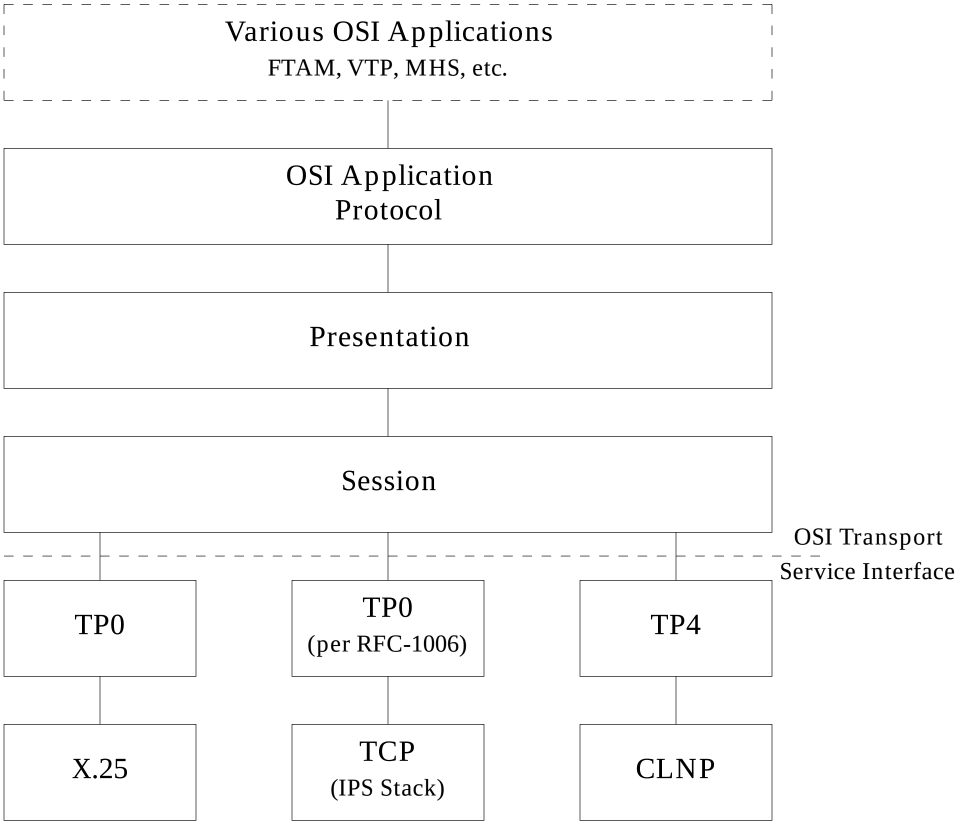

Figure **4-11.** Forms of OSI Transport Service

4.6.3 Comments

The OSI-over-IPS hybrid stack brings with it an entire OSI environment. While users on an IPS system may have to learn new applications, current OSI users may move to an IPS system with few difficulties. After the IPS users become familiar with the OSI applications, they can move between the two networks quite easily. New IPS applications developed with the TP0 interface can be ported to a pure OSI environment with little or no difficulty.

Pros:

- TP0 over TCP provides an OSI environment within the IPS environment, facilitates development of new OSI applications, and does not affect the existing IPS network.
- IPS users do not need a full OSI stack in order to become familiar with the OSI applications. The replacement of the IPS stack by an OSI stack is relatively easy.

Cons:

- Implementation of TP0 involves some work on the IPS system and may require kernel changes.
- It does not provide a way for an OSI application user to interoperate with standard IPS applications.
- The unpredictable behaviour of the urgent (expedited) data may cause confusion. The expedited data may appear at different times in the data stream,

4.6.4 IPS Applications over OSI

Modification of IPS applications to run over an OSI protocol stack is a low-effort approach to moving users from an IPS platform to an OSI platform, because IPS applications already exist and are in wide use. This approach may be useful in a dual-stack environment.

The developer may choose to implement the application directly over the OSI transport, over the OSI session, or over an OSI application protocol such as ACSE. The X/Open Interworking Group has formally recommended implementing applications using ACSE. The session layer has the necessary services, but ISO does not specify it as an application interface. Furthermore, IPS does not have an explicit session layer so it is not a natural boundary for IPS applications. On the other hand, there are a number of products currently available which use an interface to OSI transport for their network services.

Since IPS users are already familiar with the IPS applications, they do not need to learn a new set of applications when an OSI system is installed. They may learn the new OSI applications at their leisure.

4.6.5 Implementing Directly on OSI Transport

This is the most straightforward approach to providing the same services over OSI that are currently available over IPS. With some effort, a single application can be created that works on either protocol stack. TCP services map fairly well onto TP4. Normal TCP data maps to normal TP4 data. The TP4 Transport Service Data Unit (TSDU) boundaries would be ignored in an IPS-originated application. TCP urgent data is roughly equivalent to TP4 expedited data. However, in order to implement IPS applications that make use of TCP urgent data, developers must take care to minimise the problems.

While there is a close functional mapping between TP4 and TCP, some differences can lead to slightly different service interfaces between the two protocol stacks. These differences are:

TP4 Feature	TCP Feature
packet-oriented	stream-oriented
expedited data	urgent data
abortive close	graceful close
network-specific addressing	single format

Pros:

- Little work is required to develop the IPS applications directly over an OSI stack.
- The same application can be made to work with either or both protocol stacks in a user-transparent way, thereby insulating the users from changes in the underlying network.
- Users on either type of network can access services on the other type of network if a transport level bridge exists.

Cons:

- This approach might lengthen the period of migration to OSI, because IPS applications are likely to remain on the network for a period of time.
- It does little to provide a way for an IPS user to interoperate with a standard OSI application.
- The unpredictable behaviour of the urgent (expedited) data may cause confusion. The expedited data may appear at different times in the data stream, resulting in different behaviour in the two networks.
- The abortive nature of the TP4 close could cause lost data. The FTP protocol of IPS normally closes the network data stream to indicate end of file. If the close causes data to be discarded, the remote system does not receive all the data and is not notified of the error.
- Good network-transparent applications are difficult to develop.

TCP Urgent and OSI Expedited Data

IPS applications that use TCP urgent data assume that the indication of Urgent data is in the correct place in the data stream. If an OSI protocol stack is used to replace TCP, the application may get the indication in the wrong place and not flush the data stream correctly. This may not be a significant problem on a reliable, high-speed local-area network, but could be a problem on a slower or less reliable network where retransmissions are frequent.

The mapping of TCP urgent data to OSI expedited data differs in three important ways. The major difference is that TCP urgent data is always transmitted in line (in-band) with normal data, while OSI expedited data is handled out-of-band with normal data. One effect of this is that OSI expedited data may arrive at the application in advance of where it would appear in the stream of data if it were TCP urgent data.

The second difference is that multiple TCP urgent pointers are potentially merged such that only the last urgent pointer is seen by the application. However, each OSI expedited data message is received independently by the application.

Finally, the two protocols vary in their handling of asynchronous notification of the arrival of urgent data. The TCP specification indicates that some form of out-of-band notification should be made to the TCP user when urgent data arrives in the TCP receive queue but that the actual urgent data should stay in the data buffer where it was received. In contrast, the OSI protocols deliver expedited data to the user before normal data that may have been sent prior to the expedited data.

To illustrate, assume that an application has done several send operations, and then an event occurs that should cause the in-transit data to be flushed. In the TCP environment, the application protocol typically indicates the desire to flush all data up to a certain point by sending a TCP urgent indication. The receiving application protocol probably specifies that it wants asynchronous notification of the receipt of urgent data, and reads and discards data past the point where the urgent data had been inserted.

TCP Graceful Close and OSI Abortive Close

The difference between the TCP graceful close and the OSI abortive close is significant. A graceful close means that the receiver waits for data to drain before actually closing the connection, thereby avoiding the loss of any data. The OSI abortive close discards any unsent data and forces the connection closed.

IPS applications using TCP frequently assume that the close is graceful. An example is the FTP protocol. When using the stream mode of data transfer, FTP indicates end-of-file by closing the TCP virtual circuit. If FTP were used directly over TP4, the close might cause some data to be discarded, but the user would see no indication of the data loss.

4.6.6 Implementing with ACSE and Presentation Services

The Association Control Service Elements (ACSE) protocol provides a fairly simple association and control protocol that establishes connections (associations) and provides a common application interface to presentation and session services. It is possible to use these OSI services to implement IPS applications. X/Open is currently producing an API to access the primitives of these services.

This approach has been proposed to ANSI X3 for implementing the X Window System over OSI. This approach seems to be favoured over implementing X directly over the transport layer.

The ACSE primitives that must be implemented are A-ASSOCIATE and A-RELEASE, for establishing and closing the association. By using ACSE, the close operation of A-RELEASE is effectively a graceful close, because the A-RELEASE is received by the remote ACSE and acknowledged before a T-DISCONNECT is performed. Data is transferred with the presentation layer P-DATA primitive.

The presentation layer has the P-EXPEDITED-DATA primitive, which can be used to send urgent data by IPS applications. The difference between TCP urgent data and OSI expedited data (discussed in the previous section) arises here as well.

The parameters associated with the ACSE primitives are described in RFC-1026. Underlying the ACSE services are the presentation and session layer services. The P-DATA primitive need support only the *user data* parameter.

This approach can eliminate some but not all of the problems associated with implementing IPS applications over OSI, providing a more reliable mechanism for implementing IPS applications over OSI, and has little effect on users. The task of porting an IPS application to use a direct interface to the ACSE and Presentation primitives is more difficult that using the OSI transport mechanisms directly, however it may be possible to use a higher-level, simpler, API (for example accessing the connection management and data transfer functions of ACSE and the Presentation Layer via the XTI API).

Pros:

- This approach makes the IPS application fit the seven-layer model, with all layers conforming to the OSI protocols.
- Using ACSE eliminates data loss with the TP4 abortive close.

Cons:

- Implementing an application using ACSE primitives directly requires much more work than implementing directly on the OSI transport layer. The ACSE services are significantly different from those provided by the transport-level interfaces, so it may require considerable modification of the application code to adapt to the ACSE interface.
- There is a potential performance penalty as a result of the additional processing that occurs in the ACSE and session layer.

4.7 NETWORK SERVICE TUNNEL

A network service tunnel carries network-level information of one protocol family across a dissimilar intermediate network. In IPS and OSI networks, network service tunnels carry IPS IP datagrams from one IPS network to another over an intermediate OSI network, or carry OSI CLNP datagrams over an IPS network. The network protocol of the intermediate network is treated as a data-link protocol rather than as a network protocol.

The major advantage of this mechanism is that a network can use existing networks and routers that do not implement all of the protocols. The major disadvantage is that there are significant difficulties associated with managing medium or large networks based upon network service tunnels. This is due to the requirement to maintain two separate networks and to maintain the routing table used by the guest protocol to navigate its packets across the host network.

4.7.1 OSI CLNP over IPS

RFC-1070 describes a method of implementing an IPS network service tunnel between two OSI networks. RFC-1069 describes the address mapping for this service. Because of problems encountered in address handling and routing propagation using the methods described in these RFCs, an Internet Engineering Task Force (IETF) has been created by the IAB to solve these problems. The IETF is planning to issue new RFCs for an OSI-over-IPS network service tunnel.

The CLNP PDUs are entirely encapsulated within an IPS IP datagram and transmitted over (tunnelled through) the IPS network. The PDU size should be chosen by CLNP to avoid fragmentation by the IPS IP. The receiving IPS IP must be prepared to reassemble the PDU if fragmentation occurs. When the IP datagram is received at the destination, the CLNP PDU is extracted from the IPS IP datagram and placed on the CLNP receive queue.

The intermediate IPS network is entirely transparent to the OSI user and to the OSI protocol stacks (other than the OSI routers) that are interfaced to the IPS networks. The IPS network may actually have somewhat better routing performance due to the maturity of IPS routers. **Figure 4-12**, **Network Service Tunnel - OSI Over IPS**, illustrates OSI protocols crossing an IPS network.

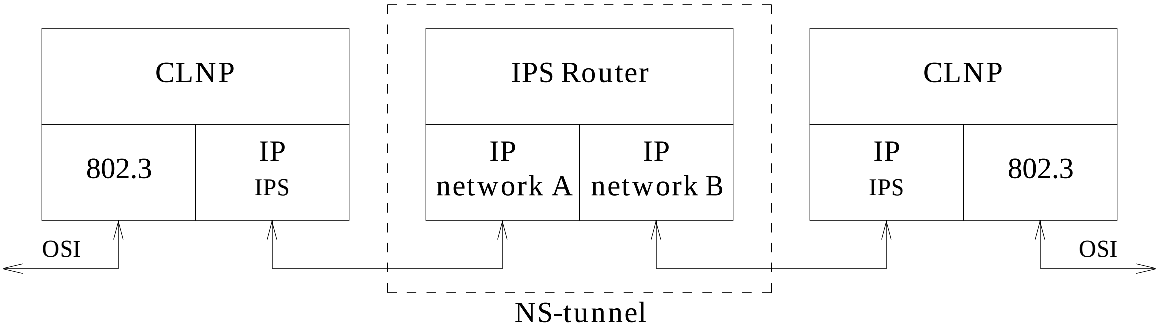

Figure **4-12.** Network Service Tunnel - OSI Over IPS

Pros:

- Allows the use of existing IPS networks as subnetworks for OSI networks.
- Implementation is fairly simple.
- The tunnel is transparent to the OSI application user.

Cons:

- How much utility a network service tunnel has is unclear. Many of the commercial IPS IP routers support, or are planning to support, the OSI protocols.

4.7.2 IPS IP over OSI

An approach similar to that used in OSI CLNP over IPS networks could be used to send IPS IP packets over an OSI network. In the U.S., where IPS networks currently outnumber OSI networks, there is little need for this type of mechanism.

However, in Europe there are many IPS LANs with public X.25 networks available for LAN-LAN interconnection. In this environment, network service tunnels are a practical method for connecting IPS LANs over large distances. The ISO Technical Report TR9577 defines how IP packets may be embedded in X.25 packets in such a way that the receiving endpoint can determine by which protocol suite the packet is to be processed.

4.8 TRANSPORT RELAY

A transport relay maps transport service primitives of one transport protocol to the transport service primitives of another transport protocol. The purpose of a transport relay is to connect the services of one network to those on another dissimilar network. In the IPS-to-OSI case, a transport relay allows an end-to-end connection from an IPS-only system to an OSI-only system.

To understand how transport relays work, assume that a user of an OSI application on an OSI-only system wants to access the same OSI services running on an IPS system. There is a common service interface at the session layer. A transport relay establishes a connection between the two sets of transport protocols on a dual-stack host and then transfers the session primitives between the two transports. This provides end-to-end connectivity between applications on different networks.

4.8.1 Bridging Function

Transport relays can be used between TP0 transports on different underlying networks, between TP0 and TP4 transports, and between TP4 and TCP. Because both transport protocols are connection-oriented, the transport relay must have an active mechanism for establishing a connection on the second network. After the connection is established, the transport relay can then perform the relaying function of the transport services.

4.8.2 TCP-TP4 Relay

Since TCP and TP4 are functionally similar, a transport relay, or bridge, can provide end-to-end connectivity between the IPS and OSI networks. IPS applications running over TP4 can then have full connectivity across the combined IPS and OSI networks.

Because of semantic differences in the two protocols, it is not feasible to implement a full TP4-to-TCP gateway. The two protocols can be connected through a relay of the associated transport service primitives. That is, when a TCP connection to the OSI network is attempted, the relay notices the incoming TCP *T-CONNECT.indication* as a request to make *T-CONNECT.request* to a TP4 system. A fully functional relay maps each incoming *T-PRIMITIVE.indication* to an outgoing *T-PRIMITIVE.request* on the other network. **Figure 4-13**, **Transport Layer Relay (TCP-TP4)**, illustrates this technique.

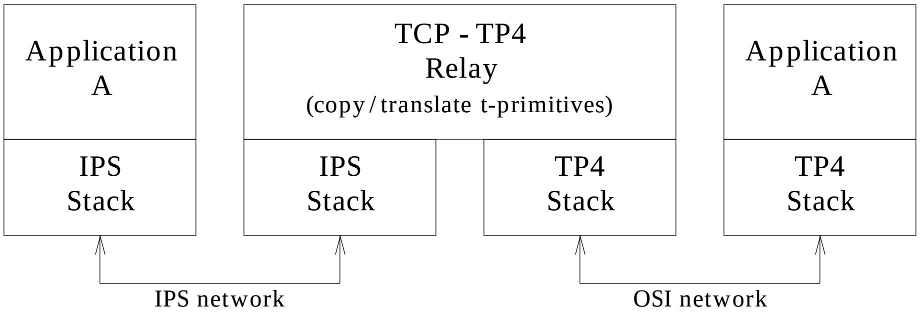

Figure **4-13.** Transport Layer Relay (TCP - TP4)

Given the differences in addressing schemes, it can be difficult to determine the connection point within the other network. This problem appears when an IPS application running on an IPS-only host attempts a connection to an IPS application server running on an OSI-only host. A TCP application tries to connect to a remote address that has a 32-bit internet address plus a 16-bit TCP port number. An OSI system needs additional addressing information which is variable in length. Without modifying the IPS-only TCP software, there is no way to carry the additional OSI addressing information.

In the transport relay, the TCP/IP protocols would be implemented in such a way as to accept all known internet addresses for the OSI hosts. One possible solution is to use an internet network or subnetwork for those internet addresses. The relay system must look like an entire OSI subnetwork. The same situation exists in the TP4 to TCP direction; the TP4 side of the relay must accept the incoming *T-CONNECT.indication* for hosts on the IPS side of the network. Some form of mapping of TCP port to TSAP must also be performed.

Connections are finally established when the target system returns a connection-confirmed indication. After an end-to-end connection is established, the IPS applications can communicate with each other subject to the limitations described earlier in **Section 3.4**.2, **TCP and OSI Transport Class 4**. In particular, the problems of TCP urgent data and OSI expedited data, and the TCP graceful close and OSI abortive close still exist.

Table **4-5.** Transport Relay

Transport A		**Transport Relay**		**Transport B**
T-CONNECT.request	→	T-CONNECT.indication	→	T-CONNECT.request
T-CONNECT.confirm	←	T-CONNECT.confirm (associate A and B connections)	←	T-CONNECT.confirm
T-DATA.request	→	T-DATA.indication	→	T-DATA.request
T-DATA.indication	←	T-DATA.indication	←	T-DATA.request
. . .		. . .		. . .
T-DISCONNECT.request	→	T-DISCONNECT.indication	→	T-DISCONNECT.request
T-DISCONNECT.confirm	←	T-DISCONNECT.confirm	←	T-DISCONNECT.confirm

4.8.3 Comments

Implementing a bridge between TP4 and TCP increases the interoperability of the two types of networks. In addition it allows users to continue using IPS applications in the OSI environment when the two networks must coexist. It is useful only when IPS applications must remain unchanged in the OSI network environment.

Pros:

- The technique provides connectivity between IPS and OSI networks for IPS-only applications. Without this technique, users of IPS applications are limited to the services on their own network, independent of whether the network is IPS or OSI based.

Cons:

- The mapping between TCP and TP4 can lead to anomalous behaviour under some circumstances. A user has no indication when errors occur.
- The relay must maintain state information on established connections. If either side of a connection fails, the cause may not be obvious. If an intermediate relay goes down, it may be difficult to re-establish a connection. Any failure of the relay or endpoint connection must be reported to the user.

4.8.4 TP0/TCP-TP0/X.25 Relay

When the relay connects a TP0/TCP stack to a TP0/X.25 stack, the service primitives are simply transferred between the networks. Because the two networks use the same transport services, there are no problems resulting from an incomplete mapping of the services from one network to the other as there are with a TCP-TP4 relay.

ISODE contains an ISO-TP0 bridge between TCP and X.25. This bridge is also defined in RFC-1086. It includes a simple registration protocol to determine the address mapping and makes the IPS host appear as part of the X.25 address space. **Figure 4-14**, **Transport Layer Relay (TP0/TCP-TP0/X.25)**, illustrates this technique.

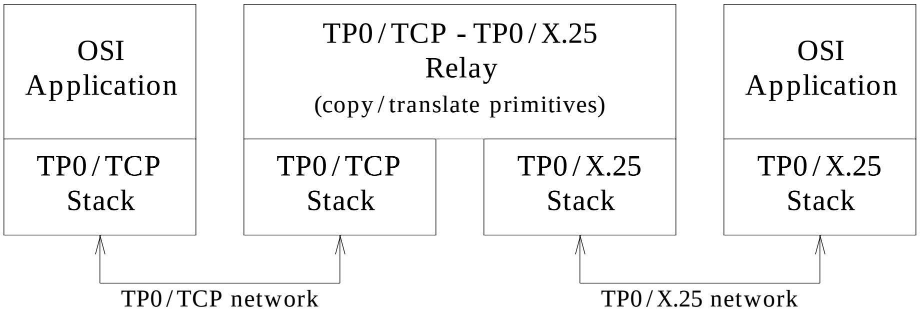

Figure **4-14.** Transport Layer Relay (TP0/TCP-TP0/X.25)

As in the TCP/TP4 case, connection establishment is the biggest problem. When an IPS host tries to connect to an X.25 host, the IPS host first connects to the TP0 relay at TCP port 146. Once connected, the IPS host sends a 1-octet function value, indicating that an active connection is requested, and then sends the X.121 address of the X.25 host. If the IPS host wants to allow incoming connections from an X.25 host, the 1-octet function value is set to indicate ‘‘listen’’, the IPS hosts X.121 address is sent, and the TP0 relay listens for X.25 connections for that address. In this case, the X.121 address must be a subaddress of the TP0 relay.

4.8.5 Comments

A bridge or relay between the two TP0 networks increases the interoperability of the IPS and OSI networks. OSI applications running on an IPS network can interoperate with OSI applications in an OSI X.25 environment.

Pros:

- The technique provides connectivity for OSI applications between the IPS and OSI networks.

Cons:

- The relay must maintain the state of the connection. If either side of a connection fails, the cause may not be obvious. It is difficult to re-establish a connection if an intermediate relay goes down.

4.8.6 TP0-TP4 Relay

A TP0-TP4 relay is similar to the TP0-TP0 relay. A connection-establishment protocol must be developed. After a connection is established, the transport service primitives are copied between the two active connections as with the other two forms of transport relay.

The TP0-to-TP4 relay is useful in a pure OSI network. TP0 over X.25 is in common use in some parts of the world and is the specified transport mechanism in the 1984 X.400 specification. TP4 is being mandated in other areas, particularly the US GOSIP requirements. Because both networks exist, and may need to interconnect, such a relay may be necessary for those networks where TP4 cannot be negotiated. **Figure 4-15**, **Transport Layer Relay (TP0-TP4)**, illustrates this technique.

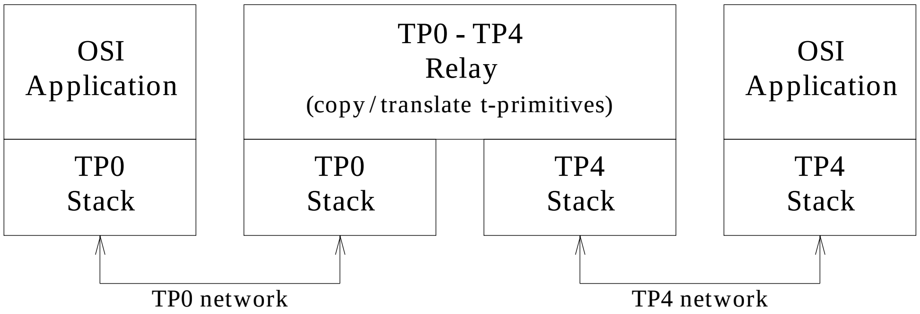

Figure **4-15.** Transport Layer Relay (TP0 - TP4)

4.8.7 Comments

This particular approach is very similar to the TP0/TCP-to-TP0/X.25 relay in terms of its functions and interoperability. It has the added benefit that it could be extended to bridge between TP0/X.25 and a TP4 network as well.

Pros:

- A working solution provides interoperability with a wider set of networks, including X.25-based TP0 networks with a GOSIP TP4 network.
- An OSI application user in the IPS environment has full and fairly transparent access to services provided in a true OSI environment.

Cons:

- The relay must maintain the state of the connection. If either side of a connection fails, the cause may not be obvious. It is difficult to re-establish a connection if an intermediate relay goes down.

4.8.8 Addressing

Network addresses are usually obtained through a directory look-up and are, therefore, transparent to the user. When a transport relay is used to connect two OSI applications, the conventions established in RFC-1069 must be followed.

The U.S. Department of Defense, a large IPS user, has specified an address format that includes the IPS address in the DSP part of the OSI network address. In this format, the DSP is partitioned into five fields: Version, Area Number (AN), Autonomous System Number (AS), IP Address, and IP Protocol ID (PID). **Figure 4-16**, **IPS Address Included in ISO Network Address (U.S. DoD)**, illustrates how a OSI network address of this type might appear.

	IDP		DSP				
US DOD	AFI	IDI	version	AN	AS	IP address	PID
	47	0006	01	0000	0207	80D420FE	06

Figure **4-16.** IPS Address Included in ISO Network Address (U.S. DoD)

A group of other IPS users has created a separate method for including IPS addresses in an OSI network address. This scheme uses an AFI already assigned to TELEX, and a reserved IDI value that indicates the format of the DSP. **Figure 4-17**, **IPS Address Included in ISO Network Address (Internet)**, illustrates how a OSI network address of this type might appear. Other organisations may create new address formats as their requirements dictate.

	IDP		DSP		
OSI over IPS	AFI	IDI	subnet	IP address	port
	54	00728722	03	80D420FE	00102

Figure **4-17.** IPS Address Included in ISO Network Address (Internet)

4.9 SUMMARY

Each of the methods covered in this survey has some merit and each provides a unique solution to the problems of coexistence and migration. None of the individual methods solve all of the problems associated with migrating users and services from an IPS platform to an OSI platform.

There must be a small number of dual-stack systems on the network since dual stacks are required to provide interoperability between the IPS and OSI networks. If all systems could be made dual-stack at the same time, the migration effort would be simpler. However, it is rarely possible to modify every system on a network.

Application gateways provide a solution for users of systems where dual stacks cannot be installed. Providing application gateways enables the IPS user who is forced to stay on an IPS-only system to use the existing IPS applications and access services on the OSI network. There may be some performance penalties associated with the gateways, but this may be more acceptable than a lack of access to the services on a pure OSI system.

Hybrid stacks can be used to provide either an OSI-like environment on an otherwise IPS system or an IPS-like environment on an OSI system. The OSI-over-IPS environment can be used as a development platform for familiarising users with OSI applications or implementing new OSI applications prior to the introduction of a real OSI network. This approach can smooth the transition to an OSI network.

The other direction, IPS applications over OSI may provide familiar applications to users who must move from IPS to OSI. However, care must be taken to avoid the problems that result from the mismatch in transport protocol semantics. Also, retaining IPS applications may unnecessarily prolong the migration to OSI and, in addition, requires extra effort to support the extra IPS applications.

Network service tunnels provide a mechanism for connecting two compatible networks via an intermediate network of another type. This is useful where direct connections or connections through a compatible network are unavailable.

Transport level bridges are another mechanism for providing access to services between IPS and OSI networks. Their primary purpose is to allow hybrid stack applications to access the associated service on the other network. For example, a bridge allows an FTAM user on an IPS system to access an FTAM server on an OSI system.

A real migration strategy consists of some combination of the methods described. The exact combination depends upon a number of factors including: environment that currently exists, the desired OSI features, and the availability of coexistence techniques. **Chapter 2**, **Problem Statement**, of this guide describes some typical examples of network installations and particular migration and coexistence needs. The remaining chapters of this guide give some recommendations for applying the techniques described in this chapter to solve these particular migration and coexistence needs.

Chapter 5

Policies

This chapter provides a guide for those involved in determining the policy of an enterprise with respect to the transition from IPS networks to OSI-based networks and the coexistence of both networking technologies in the same enterprise. Policies need to be defined whether the enterprise chooses to migrate or coexist. A policy is a set of decisions determining:

- whom to migrate
- what to migrate
- in which order to migrate
- when migration occurs
- the extent of interoperation between IPS and OSI during migration

The motives for migration of an organisation's network to OSI might differ. For government agencies and some enterprises, the appeal lies in achieving a high degree of vendor independence through having their networks follow a non-proprietary architecture based on a set of international standards. Government organisations are also bound by directives from national and international bodies regarding their procurement practices. For most of the commercial enterprises, the migration probably occurs due to business pressures because the emergence of implementations of OSI services allows them to do business in new cost-effective ways.

For most organisations, the balance of the cost of adopting OSI protocols against the business and efficiency benefits to be gained is of paramount importance. Such organisations must strive to formulate a policy which maximises the benefits and minimises the costs and inconveniencies.

A policy should be based on a strategic vision of the future network and its functions as experienced by users in an enterprise-wide context. Since every policy must be specific to the enterprise in question, this guide restricts itself to listing the factors to be considered when formulating a policy.

5.1 COEXISTENCE AND MIGRATION

There are three options to be considered when discussing coexistence and migration: *pure coexistence, coexistence with transparent interoperability* and *total migration*. Each can be thought of as points on the graph below.

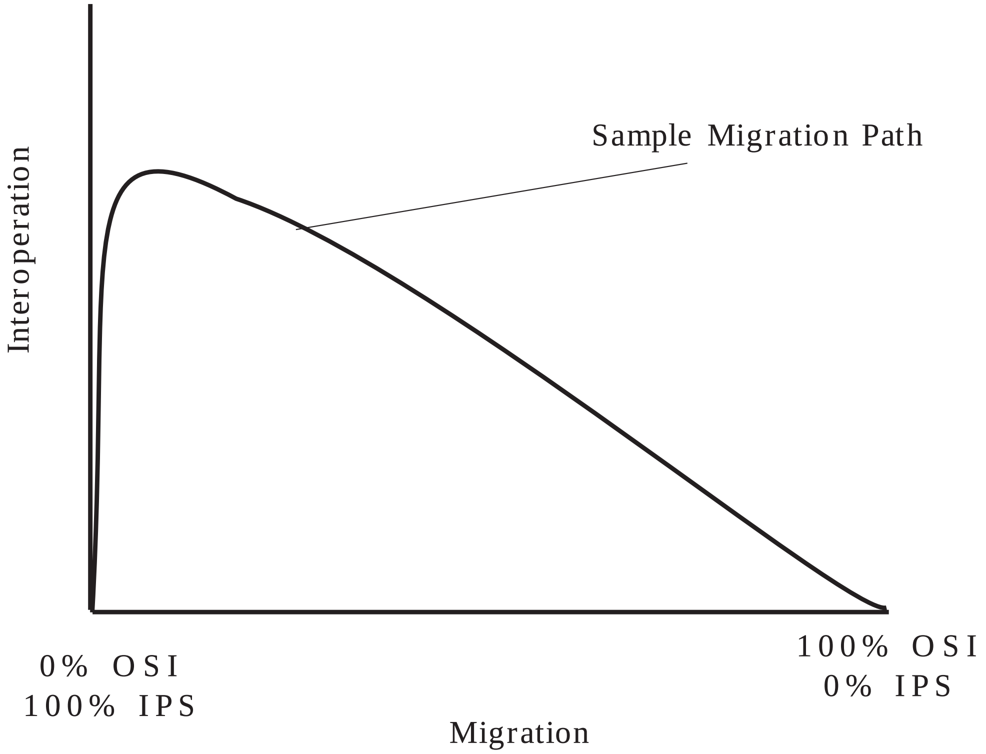

Figure **5-1.** Migration / Coexistence Model

There are two types of coexistence to be considered.

- Pure Coexistence, where two protocol suites exist but do not interact in any way.
- Coexistence with transparent interoperability, where a user using one protocol suite can interact with a user using the other protocol suite without knowing there are more than one.

Even pure coexistence requires management to ensure that the two protocol suites do not interact destructively. Coexistence with interoperablity exhibits many of the problems that are faced when migrating from one protocol suite to another but involves the conscious decision to stop at a given point.

Many of the issues that arise when determining a policy for migration occur on the path to total migration and often relate to interoperability. Each enterprise must choose the method and the phasing of the migration and decide on the amount of interoperation that is needed during the migration.

5.2 FACTORS DETERMINING POLICY

There are many factors to be considered when determining the policy of migration and coexistence. This section considers each of them in turn.

5.2.1 Cost and Benefits

For most organisations, especially those in the commercial sector, the balance of the cost of implementing OSI protocols against the commercial benefits that can be obtained is the most important factor in determining policy. Thus the policy developed must demonstrate clear commercial advantages such as efficiency of operation, and justify the relative costs of the policy options chosen and rejected.

The costs to be assessed include:

- Capital costs of new hardware and software.
- Costs associated with operating the network, both during migration and afterwards.
- Skills which must be acquired (by training or recruitment) to implement the policy.

5.2.2 Users

In order for policy to be determined, an inventory of what users are doing at the moment and what they are likely to do in the future is required. All interactions in the system must be identified, including external contacts with other organisations' networks. This profile can then be used to make predictions about future needs and take decisions about the phasing of migration. Migration has several consequences for users:

- Retraining;
- New or changed documentation, and
- A possible reduction in the levels of service.

Retraining

Whatever the plan for migration, users need retraining at some point in time since in most cases both the interface to a service and the services it provides change. The policy determines the sophistication of the user interface and the skill levels required for different categories of users. In the case of simple migration the user may need to be trained on the new user interface for each application used. If interoperability is required during migration, it may be necessary to provide a common interface to corresponding IPS and OSI applications. This provides some flexibility to the migrator in that applications can be replaced below the user interface without the users knowing. This approach does have implications. For example it is possible to provide an interface for file transfer that is consistent with that of FTAM, but the user is unable to access some features because the underlying application is an IPS one and does not support it. Such interfaces must be highly user-friendly to cope with complexities such as changes in addressing and to satisfy the differing requirements of those users who are used to IPS or OSI services.

Documentation

Migration to OSI applications requires the provision of new user documentation and possibly new operating procedures. More importantly, special migration documentation has to be written for the transitional period. The documentation of how, what, when and the facilities provided during migration are crucial to the success of the migration process.

Service Levels

Service levels perceived by the users of the network may be affected in a number of ways:

- Throughput or response times (such as character echo) may be reduced by the presence of additional gateways in the path to a service.
- Performance of workstations and servers may be increased where multiple protocol stacks are supported.
- However the load placed on the network infrastructure may increase response times and reduce throughput.
- Maintenance response times may be increased and availability reduced due to the load on the department's supporting the services. If this is the case it is vital that the services are prioritized to give the users the support required for the most important applications.

5.2.3 Operational Consequences

During the period of migration the choices that can be made are constrained by the availabilty of operational staff that are skilled in both protocol suites and understand the issues when interoperating and running them concurrently. The departments responsible for providing services to the organisation's users experience an increased work load as they have to support multiple applications. This may require more staff and new skills. The operational policy must define:

- The service levels that users can expect
- The priority of the business with respect to applications (for example: is mail more important than file transfer?)
- Staffing levels and the skills required

5.2.4 Applications

As stated earlier, many of the problems occur on the path to migration, where some form of interoperability is required. This is especially true of applications.

The major policy decision to be made regarding applications is whether to preserve the currently available user interface and functions in the new environment, or to migrate the users to applications which are native to the OSI environment. The former case has the advantage of reducing the disruption caused to users as a result of the migration, although there are aspects, such as changes to address syntax and gateway navigation, which are difficult or impossible to hide from the user. The main disadvantage of preserving the user interface is that the user is unable to take advantage of the advanced functions that may be available unless significant effort is put into enhancing the current interface. A possible policy is to allow users to retain the security of the existing interface whilst providing them with information about the advantages of the advanced utility, supported by training programs, to encourage migration.

Three OSI applications that are sufficiently mature to be candidates for migration are discussed individually below:

X.400-based Electronic Mail

The OSI application that is most likely to be adopted early is electronic mail as it provides extra functions over and above that provided by the equivalent IPS application.

There are several aspects that need to be considered when determining the policy of migration. IPS mail and OSI mail are substantially different in the functions provided. These differences affect all levels of the network structure. The policy must determine:

- Whether to preserve the user interface
- The functions required
- The connectivity required

For example, OSI mail has the capability to deliver mail that consists of many different types of information, such as image, voice, binary programs, electronic business documents, and so on. IPS, on the other hand, predominantly uses ASCII text transfer. If only the user interface is preserved, connectivity is increased but the users are unable to take advantage of the extra functions.

OSI mail relies upon the OSI Directory Services application to find out where the mail should be delivered, IPS mail uses the domain name service. These services are substantially different and require a translation to be made when an OSI mail user tries to send mail to an IPS mail user. The policy adopted must state where this translation takes place. This may affect the names that are used at the user interface. For example, if a decision is made to migrate an IPS user agent to run over X.400 protocols, a further decision must be made to preserve the IPS naming that the users know or to use fully fledged Originator-Recipient names.

File Transfer

There is no standard mapping defined between IPS and OSI file transfer services as there is for X.400 based mail and IPS mail, however, there are a products on the market which implement proprietary mappings. The policy adopted depends on the needs of the users and the cost of adoption of this policy; the options are:

- Define organisation's own proprietary mapping and implement some form of translation software
- Migrate the file transfer service for all users to FTAM
- Use commercially available software which defines its own mapping
- Coexist indefinitely

As with all applications a decision needs to be made whether to migrate user interfaces or not. For example one could use the FTP command set with an underlying FTAM service. In this case the issue of names at the user interface discussed above also applies.

Virtual Terminal

The complex nature of the OSI VT standard has meant that terminal server and host-based products have been slow to appear. With the emergence of profiles for supporting Telnet access and a simple forms mode of operation, a number of suppliers are now beginning to announce plans for VT-based products. Until such products are more widely available, policies for terminal access must concentrate on the coexistence of current terminal access mechanisms within the migrated network.

Application Coexistence

When considering providing transparent user access to network applications it is necessary to explore whether or not there is a common subset of functions that exists between the two applications. New applications should be written to an appropriate standard API in order to ensure wide portability and possible independence from underlying network technologies. An enterprise's policy must include guidance on the use of APIs for the commissioning of new applications. Such a policy might require an application to use the highest level API which is available on the required platforms and which provides the services required.

5.2.5 Infrastructure

Some but not all network components and devices are affected by the migration to OSI. Again the issues arise when the two protocol suites have to interoperate. There are two aspects of the infrastructure that require consideration when determining a policy for an enterprise. They are the physical network components and the management of those components.

Network Components

As an example, the table below is a selection of Ethernet LAN components, with suggestions about the impact on those components with respect to coexistence, migration and interoperation.

- Cabling
  - — **Coexistence:** Current cabling supports multiple protocol suites.
  - — **Interoperability:** Not an issue.
  - — **Migration:** Not an issue.
- Media Access Units
  - — **Coexistence:** Base hardware can cope with either OSI or IPS frames. Requires a driver that can pass on frames to different modules of upper layer software
  - — **Interoperability:** Not an issue.
  - — **Migration:** Not an issue.
- Bridges

  - **Coexistence:** MAC-level bridges operate at the frame level, have no concept of upper layers of a protocol and can be made to work transparently with both protocol suites.
  - **Interoperability:** Not an issue.
  - **Migration:** Not an issue. See the explanation under coexistence.

- Routers
  - **Coexistence:** A router which could process both protocol suites is possible; the underlying hardware of most current routers could do the job. Many of the existing routers load their software from the network or other media and would be re-configurable.
  - **Interoperability:** A router configured for one protocol suite could not interoperate with a router configured for another. But it is possible to build a router which supports both protocol suites and can route packets from either suite across common links.
  - **Migration:** As mentioned previously, many of the routers which exist today have loadable software and could be re-configured for OSI operation.
- Terminal Servers
  - **Coexistence:** It is likely that terminal servers will become available which support both IPS and OSI protocols. Such servers have heavy memory and processor requirements in order to support the multiple stacks simultaneously.
  - **Interoperability:** A terminal server configured for one protocol suite cannot interoperate with a terminal server configured for another.
  - **Migration:** Terminal servers may be upgradable by changing the embedded software.

It should be noted that migration of network management protocols and tools affects most of the above networking devices, as they must be manageable in the new environment. This may require protocol support, gateways and converted object definitions.

LAN Gateways

When planning to use gateways in a LAN environment the effect on traffic must be considered. Two aspects should be noted:

Traffic doubling:
The relay function of the gateway causes the LAN traffic to be doubled on a LAN segment when the gateway operates over a dual stack on a single adaptor. A policy must consider the effect on overall traffic levels.

Gateway locality:
Unless gateways are available on the local segment, intra-segment traffic may have to cross a bridge to find the relay and then cross back again (this might be acceptable for mail gateways which operate in a store-and-forward mode but not for file-transfer gateways). A policy must define the criteria for deciding the number and siting of gateways.

Management

The management of a network during the migration period or when coexistence is being considered, requires careful thought and planning. Products are now becoming available that are able to manage both protocol suites by the use of common APIs, gateways and the standardisation of the objects being managed. Unless the policy mandates the adoption of such a product the organisation must maintain expertise and management systems for two distinct networks. It must be emphasised that extra cost is incurred even if the two protocol suites do not interoperate. There is the additional burden when the major applications also interoperate, as they require maintenance if they are to be made to work consistently and reliably.

5.2.6 Security

For some organisations, such as government institutions or banking organisations, the issue of network security adds a further dimension to the problems associated with formulating a migration policy. When choosing coexistence techniques, security considerations may make most of the techniques discussed in the guide unacceptable without significant modifications.

For example, the use of network tunnelling techniques may subvert network layer security because the protocols used to transport the guest packets across the intermediate network might route data without regard to its security requirements. Application gateways are another potential area of risk unless data to be protected is encrypted above the application layer.

5.2.7 External Influences.

There may be influences which are external to the enterprise and which have a major impact on the policy of an organisation that is migrating to OSI networks. It is important not to ignore these sources of requirements when gathering information to allow policy decisions to be made.

Government Directives

For some organisations, such as government departments and public bodies, procurement of computer systems is bound by directives which have the force of law. As discussed in **Section 1.1.1**, **Standardisation Motivations**, two such directives are FIPS-146 from the U.S. Department of Commerce, and 87/95/EEC from the European Community, both of which relate to aspects of procurement, and which have led to the development of government profiles (GOSIPs) for use in the tendering process.

Business Partners

Changes in the way that an organisation's business partners, clients or suppliers operate may force a change. For example, using Electronic Data Interchange (EDI) over networks as a method of exchanging stock information, orders and invoices. Links with business partners may also force the pace of adoption of OSI mail and directory services, even though the majority of the enterprise's users may only need to send text around.

5.2.8 Timing

In determining the order in which network components, users and applications are migrated, the policy maker is influenced by a number of factors.

Development Lead Times

The migration policy adopted may require significant development effort to port the organisation's own applications to the OSI environment. The availability of resources for such activities must be taken into account.

Commercial Considerations

The timing of migration steps must avoid jeopardising business-critical functions such as reporting cycles.

Product Availability

Clearly a policy must take account of the availability of products and the functions provided when scheduling the steps in a migration plan.

Incremental Implementation

The coexistence techniques presented in this guide allow an organisation to migrate its networks and systems in an incremental fashion to allow experience in operation and using the new applications and infrastructure to be gained progressively. Such policies may protract the migration timescales but have the advantage of building confidence in manageable steps.

Chapter 6

Tools

The purpose of this chapter is to identify the tools which can be used to achieve IPS - OSI migration and coexistence. **Chapter 4**, **Techniques**, presents a discussion of the techniques that can be used to enable the coexistence and interoperability of IPS and OSI protocol suites and to support the migration from one to the other. Building on that, this chapter describes the specific software tools required to implement the coexistence and migration techniques proposed. Implementation of a particular technique may require a combination of tools; in addition a particular tool may be applicable to a number of techniques. There are additional tools which may be used to support coexistence and migration in a general way as well as providing specific support for individual techniques (for example the integrated directory service).

For each tool included in this chapter a number of aspects are identified:

- The techniques that it implements and the environment in which it is applied - this serves as a cross-reference to **Chapter 4**, **Techniques**.
- The major requirements that it must meet - this provides a guide, for software implementors and buyers, to the high-level requirements that a software component must satisfy. For clarity, the requirements are grouped to distinguish functional requirements, operational requirements, and standards requirements.

Only tools that are applicable in real life scenarios are listed. For example, **Chapter 4**, **Techniques**, discusses end-to-end application gateways as a technique for applications in general. However, this chapter does not list an end-to-end electronic mail gateway because it is generally accepted that the store-and-forward gateway technique is more appropriate to this application. This assertion is based upon two observations: the architecture for both IPS and X.400 mail protocols support a store-and-forward mode of operation; existing gateways between the Internet and other mail networks (including X.400) are store-and-forward gateways. Nor is a virtual terminal gateway listed, in this case because the OSI VT protocol is regarded as the least viable of the OSI protocols and is not sufficiently widely implemented to figure in real migration scenarios. This assertion is based on three observations: there appear to be very few VT protocol implementations on the market; interoperability demonstrations concentrate on the X.400 and FTAM applications; few of the X/Open shareholder members appear to regard the VT protocol as a viable basis for implementation of products.

It must be emphasised that some of the tools listed here have not yet been implemented by open systems suppliers. Some may never be implemented, in the case where the requirements can be met using alternative tools, or where there is not perceived to be a big enough market. Also there is no requirement for suppliers to to recognise the component boundaries implied by this chapter when implementing products.

The subsequent sections of this chapter describe individual tools. Each section includes the following headings:

- **Description** - a short, high-level description of the tool.
- **Applicability** - a cross reference to **Chapter 4**, **Techniques**.
- **Environment** - a cross reference to other tools which this tool requires or supports.
- **Functional Requirements** - the high-level user requirements that the tool must satisfy. In this context, the term user may refer either to a human or to other software components that use the services of this tool.
- **Operational Requirements** - high-level operational and administrative requirements that the tool must satisfy. This subsection highlights operational requirements specific to the tool rather than those generally understood to apply to similar network components.
- **Standards requirements** - this subsection highlights any agreements for use of protocol options or services required by a tool in addition to the currently agreed standards, profiles and implementors' agreements.

6.1 SWITCHABLE APPLICATION

6.1.1 Description

A switchable application is one which can run on top of both IPS and OSI stacks, the specific stack to be used being selected at run time depending on the end system with which the communication is to be established. It provides a user with a single consistent user interface for file transfer or electronic mail over either stack. It may be based upon an existing IPS application, for example an implementation of the FTP utility which has been modified to perform both FTP- and FTAM-based file transfer.

6.1.2 Applicability

Such an application supports the dual-stack technique. It is required in each dual-stack end system where the application is to operate over both stacks.

6.1.3 Environment

This class of application runs on top of a dual stack. For IPS applications, a transport level Common API can be used to support the application over both stacks, alternatively the application can be modified to use an application-layer interface to the OSI stack. An OSI application requires a hybrid stack in the IPS environment as the transport service is not sufficient to support it. The automatic stack selection, described below in the **Functional Requirements** section for this tool, may require a directory service which can support the decisions that the application is required to make.

6.1.4 Functional Requirements

The principal requirement for a switchable application concerns the method adopted for selecting which stack to use to implement a request. This may be manual, the user must identify which stack to use, or automatic, in which case the application itself decides. In the latter case information from a directory service, if available, might be used to make the decision. There are a number of criteria that a policy for automatic selection of protocol stack may be based upon, such as performance, administrative or security requirements.

When an end system is also a dual-stack system the application must have a policy for choosing which stack to use. A blanket policy may be adopted (for example: always use OSI if possible), alternatively the choice may be based upon the type of application or some attribute of the end system (such as the subnetwork that it is attached to). The security requirements of the application might also be considered in this case.

6.1.5 Operational Requirements

There are no specific operational requirements applicable to this tool.

6.1.6 Standards Requirements

When an existing application is ported to run over an alternative protocol stack it is necessary to decide the way it is to be implemented using the services in the environment to which it is ported. The following aspects must be considered:

- At which layer is the application to interface to the protocol stack? As discussed in **Chapter 4**, **Techniques**, an IPS application may interface to the OSI stack at the transport layer; alternatively the OSI common application services or an application-specific service such as FTAM may be used. At the transport layer, there is also the choice of a streams or socket interface to be considered. For many applications there is little or no choice in the matter, the issue being decided simply by the availability of APIs across the required range of platforms.
- What services are to be used? This selection of services defines a profile for the protocol stack required to support the application.
- How is the application to use the available primitives? For example, if the Telnet protocol were to be implemented using OSI's ACSE and Presentation services, which PDUs would be used and how would the Telnet protocol be carried in them?

Where an application is distributed, such aspects must be standardised to ensure that the clients and servers interoperate successfully. An example of such a standard is X/Open's profile for simple FTAM-based file transfer, **X/Open Preliminary Specification - Byte Stream File Transfer (BSFT)**. This profile defines how the functions of the familiar Berkeley FTP file transfer utility may be implemented using the services of OSI's FTAM. It consists of:

- A suggested manual page for the utility's user interface
- A specification of the services required from the local FTAM *initiator*
- A specification of the services required from the remote FTAM *responder*
- A mapping of the FTAM virtual file store onto the local file system

X/Open is well placed to promote the adoption of standards in this area.

In order to enable an application to interface to protocol stacks from diverse suppliers, the API to the relevant layer services must be standardised. X/Open is currently addressing this issue in cooperation with formal standards organisations, by developing APIs to a number of OSI services and application-specific APIs.

6.2 PORTABLE APPLICATION

6.2.1 Description

A portable application is similar to the switchable variety discussed above, In this case the application is ported to run in the new environment so that protocol stacks can be switched without affecting the interface with which the user is familiar.

6.2.2 Applicability

This class of application supports the portable application technique, where the migration policy requires individual end systems to be switched from one protocol stack to the other rather than supporting dual-stack operation. A ported application is required for each end system which is to be migrated from one protocol stack to another.

6.2.3 Environment

The environmental requirements of this type of application are similar to those of a switchable application except that only a single stack is required at any one time.

6.2.4 Functional Requirements

The primary requirement is that the same user interface be available whichever stack is in use. Thus there is no user requirement for a single software product which is compiled to run on either an IPS or an OSI stack, although the benefits to a software vendor are obvious. However, a single software product, written to an appropriate common API, can have a number of desirable characteristics even for the user, such as helping to minimise the number of user-discernible differences in the functions of the application in the two environments, and the operational benefits of using a single product.

6.2.5 Operational Requirements

There are no specific operational requirements applicable to this tool.

6.2.6 Standards Requirements

The standards requirements of this type of application are similar to those discussed in **Section 6.1**, **Switchable Application**.

6.3 COMMON API

6.3.1 Description

A common API may be used to support a switchable or portable application. Suppliers of such applications can write their software to an API which provides an interface to both protocol stacks. A number of levels of API have been discussed in **Chapter 4**, **Techniques**, including transport level, RPC and application-specific APIs.

Common API is included here as a tool because it may be necessary to purchase an API (such as the XTI transport layer API) as an additional layered product to support a product where the API is not bundled in with the operating system or protocol stack licence.

6.3.2 Applicability

Common APIs are used to support the switchable application, portable application and dual-stack techniques.

6.3.3 Environment

Not applicable.

6.3.4 Functional Requirements

There are no specific functional requirements applicable to this tool.

6.3.5 Operational Requirements

There are no specific operational requirements applicable to this tool.

6.3.6 Standards Requirements

There are no specific standards requirements applicable to this tool.

6.4 COEXISTENT PROTOCOL STACK

6.4.1 Description

In order to configure more than one protocol stack in a system, the protocol stacks must coexist without interference. Many UNIX-based systems have a native implementation of the IPS protocol stack. For dual protocol stack operation it is necessary to install an OSI protocol stack alongside the native IPS software, probably sharing a network interface card, without interference between the two protocol suites.

6.4.2 Applicability

Coexisting stacks support dual stacks, application gateways and transport gateways. Depending on the coexistence and migration policies adopted, dual stacks may be installed in end systems that are to access both networks, or just in those systems that provide a gateway between the two networks.

6.4.3 Environment

Not applicable.

6.4.4 Functional Requirements

For a host to support dual protocol stacks the network software must coexist at each level where the two protocol stacks touch (for example at the points where the protocol stacks share some resource). Two specific requirements are:

API Level

Where the two stacks are to be accessible via a common API, that API must support both stacks simultaneously and provide mechanisms to allow an application to select which stack to use. This join is currently possible only at the transport layer where APIs such as XTI and TLI currently provide support for both IPS and OSI protocol suites.

It is only possible to introduce protocol stack implementation from a third-party where the API is designed to accommodate multiple stacks and the mechanisms for doing so have been published.

Data Link Layer

Where a policy requires both protocol stacks to share the same network interface the data link layer must provide a mechanism to multiplex the two network-layer protocols.

Again, the interface to the data link layer must have been designed to support multiple protocol stacks. The specification for this interface must be in the public domain to allow implementors to use it for their products. An example of such an interface specification is the proprietary NDIS standard for sharing LAN adaptors under the DOS operating system.

6.4.5 Operational Requirements

It is essential for such protocol stacks to be network-manageable. It is highly desirable that both protocol stacks can be managed from a single network-management protocol, however, before this is possible there must be agreement on the objects to be managed for each protocol stack and their representation in the relevant object syntaxes.

6.4.6 Standards Requirements

It would greatly enhance the development of portable, coexistent protocol stacks if the interfaces mentioned above were standardised. Currently, only proprietary standards exist.

6.5 OSI UPPER LAYERS WITH TCP INTERFACE

6.5.1 Description

In order to run OSI applications over IPS protocols, it is necessary to provide the OSI application services in that environment. An implementation of the OSI upper layers (session, presentation and application) is required which uses the RFC-1006 convergence protocol to interface to the services provided by the TCP protocol.

There is some debate about whether the hybrid stack technique is a viable mechanism for commercial systems or merely a research tool, however at least two major systems vendors have announced plans to supply products which utilise this technique.

6.5.2 Applicability

This tool implements the hybrid stack technique and supports the switchable application and portable application techniques where an OSI application is to run over an IPS transport service.

6.5.3 Environment

Runs over a TCP-based transport service.

6.5.4 Functional Requirements

The principle functional requirement is to support the appropriate service profiles for the applications to be run. These are defined for each application.

The upper layer implementation must interface to each host's transport interface to access the services of the TCP protocol. There are a number of possible interfaces; X/Open's XTI interface is a widely accepted standard.

6.5.5 Standards Requirements

In order to support applications from diverse suppliers, the API to the application layer services must be standardised. X/Open is currently addressing this issue in cooperation with other standards organisations, by developing a number of application-independent and application-specific APIs.

As discussed in **Section 4.8.2**, **Addressing**, the RFC-1006 implementation must support the RFC-1069 method of representing IP addresses as ISO 8348/DAD2 network addresses.

6.6 TRANSPORT RELAY

6.6.1 Description

A transport relay sets up and maintains connections between application endpoints running on different protocol stacks. It interfaces to the protocol stacks at the transport layer, receiving primitive indications from one stack and translating them into primitive requests on the other.

Three types of relay are discussed in **Chapter 4**, **Techniques**:

- TCP-TP4
- TP0 / TCP-TP0 / X.25
- TP2 / TP0-TP4

6.6.2 Applicability

Supports the transport relay technique. A transport relay allows an OSI application running on top of a hybrid stack to interoperate with an end system which supports a full OSI stack.

6.6.3 Environment

A transport relay runs on top of a dual stack. In the case of TP0 over TCP, the RFC-1006 convergence protocol is required to enhance the service to make the TCP service look like TP0.

6.6.4 Functional Requirements

It is desirable to have a single relay product which can perform all of the above relaying functions (plus TP0 / TCP-TP4) depending upon its hardware and software configuration.

There are a number of addressing requirements for such a tool:

- The calling system must be able to identify that a relay is required to access the called system and must be able to select an appropriate relay for the connection.

  An appropriate relay may be selected by use of static configuration tables or may be supported by a directory service. In either case the chosen method has significance for availability, resilience and administrative requirements of the relay tool.

- The relay system must map the address of the called system on the local network to its address on the remote network. Again this may be accomplished by use of static configuration tables or using a directory service.

In addition a transport relay implements a sensible strategy for relaying error conditions from one connection on a link to the other in such a way that the two connections remain synchronised and the users of the connection get meaningful error indications.

6.6.5 Operational Requirements

The operational requirements for a transport relay (such as availability, resilience, scalability, performance, maintenance) are similar to those for conventional gateways and routers. In particular, it is important that the performance of the relay is appropriate to throughput required.

For high throughput applications such relays must be implemented as embedded systems, rather like current gateway and router products. However, where the transport relay is only required for a short migration period, products which can run on general-purpose processors are more appropriate.

6.6.6 Standards Requirements

It may be necessary to standardise the mechanism for communicating to the transport relay, the ultimate destination for the connection so that the relay has sufficient addressing information to make the onward connection.

6.7 MULTI-PROTOCOL ROUTER

6.7.1 Description

A multi-protocol router is used to bridge Local Area Network traffic across intervening WANs. The WAN may be a mesh of point-to-point links or an X.25 network (public or private). A router participates in the network layer and routing protocols on the local LAN, receiving packets destined for remote LANs and transferring them to the appropriate router across the WAN. Thus two FTAM-based applications stacks may interoperate across an intervening IPS network.

In the case of an X.25 WAN, the routers set up virtual circuits amongst themselves; these virtual circuits are then used as a data link layer for the exchange of IPS IP or OSI CLNS packets. In the case of a point-to-point WAN, the routers implement a routing protocol amongst themselves to route packets to the destination LAN.

6.7.2 Applicability

Multi-protocol routers are used to implement the Network Service Tunnel technique. In particular, such routers are used to provide coexistence where no interoperability is required, that is where LAN hosts do not need to interoperate with WAN hosts.

6.7.3 Environment

Such relays are generally implemented as embedded systems due to the high performance and diverse interface support requirements.

6.7.4 Functional Requirements

There are two principle requirements for packet routing:

- **Local routing protocols** - The router must participate in the local routing protocols to gather information about the destinations that can be reached from the local LAN and to advertise destinations that can be reached via its own connections.
- **Global routing** - The routers must exchange routing information amongst themselves to distribute information about hosts which can be reached on their local networks. This may be done by forwarding the local routing protocols or by implementing a separate global routing protocol, capable of distributing routing information for diverse routing domains.

6.7.5 Operational Requirements

The general operational requirements for these routers are the same as those discussed above for transport relays.

In the case of an X.25 WAN, there is an administrative requirement to control the virtual circuits across the intermediate network in order to balance the cost of maintaining an idle link against the cost and delay involved in re-opening a closed link.

6.7.6 Standards Requirements

In general, multi-protocol router products from different vendors do not interoperate. In order to create 'open' multi-protocol networks, agreements are needed on the way in which the various network layer protocols are to be encapsulated for transfer across the intervening WAN. In addition, if the local routing protocols are not to be propagated across the WAN, agreement is required on the global routing protocol to be used instead. It seems likely that a forum of router suppliers and users would be the appropriate organisation to develop such agreements, rather like the current OSI Network Management Forum.

6.8 ELECTRONIC MAIL GATEWAY

6.8.1 Description

A store-and-forward gateway is required to support application-level interoperability between IPS RFC-822 and OSI MHS mail applications, providing a relay function between corresponding applications in each protocol suite.

6.8.2 Applicability

This tool implements the application gateway technique.

6.8.3 Environment

The gateway operates in a dual-stack environment accessing the IPS stack at the transport interface and the OSI stack at the application layer.

6.8.4 Functional Requirements

As described in **Chapter 4**, **Techniques**, there are two modes of operation for a store-and-forward gateway: *consecutive* or *concurrent*. The choice depends upon balancing immediacy of response against duration of connection. Concurrent mode gives immediate notice of whether the destination system has accepted the message; consecutive mode closes the connection as soon as the gateway confirms reception of the message. Current IPS mail applications tend to work in an off-line mode: the message composing utility (user agent) passes a completed message to a background mailer (message transfer agent) for delivery. Thus consecutive mode provides an acceptable service. However, as the choice of mode may depend upon the role of the gateway, it should be configurable to operate in either mode.

RFC-1026 specifies the operation of an MHS/RFC-822 gateway, defining the subset of functions from each mail protocol that is supported and defining the mapping between them.

As with the transport relay, a mechanism is required to allow an application to identify that a mail recipient is in another addressing domain and to select an appropriate gateway to which to route the message. Current solutions require the cooperation of the mail sender. It may be possible to use the services of the DNS and OSI directory service to perform the relay transparently.

6.8.5 Operational Requirements

There are no specific operational requirements applicable to this tool.

6.8.6 Standards Requirements

There are no specific standards requirements applicable to this tool.

6.9 FILE TRANSFER GATEWAY

6.9.1 Description

An application gateway is required to support interoperability between the IPS FTP and OSI FTAM-based file transfer applications. This may either be an end-to-end or store-and-forward gateway, the choice depends on individual requirements. The gateway provides an application-level relay function between corresponding applications in each protocol suite.

6.9.2 Applicability

This tool implements the file transfer gateway technique.

6.9.3 Environment

The gateway operates in a dual-stack environment.

6.9.4 Functional Requirements

As with the transport relay, a mechanism is required to allow an application to identify that the destination host is in another addressing domain and to select an appropriate gateway to which to route the message. In the absence of any transparent mechanism for achieving this, the file transfer user must be aware of the name of the gateway and must be able to specify the target host-name in the remote protocol domain.

6.9.5 Operational Requirements

There are no specific operational requirements applicable to this tool.

6.9.6 Standards Requirements

Unlike the RFC-1026 standard for electronic mail, there is currently no agreed mapping of FTP primitives onto equivalent FTAM primitives. The existence of such a standard is a prerequisite for the production of an open FTP/FTAM gateway. The IAB are the most appropriate organisation to undertake this work.

6.10 INTEGRATED DIRECTORY SERVICES

6.10.1 Description

In a network which includes both IPS and OSI protocol suites, the administrative effort involved in operating two separate directory services may be significant, particularly when migration is in progress and the address and name spaces are changing to reflect the new configuration of the network. In addition, a number of the tools discussed in this chapter require support from a directory service to provide mapping between protocol domains.

A directory service which supports the OSI and IPS name spaces and also provides information useful to the operation of transport relays and application gateways may satisfy these requirements.

6.10.2 Applicability

Such a service is useful where both IPS and OSI address spaces coexist in a network. It may support transport relay and application gateway techniques.

6.10.3 Environment

Depending on the method adopted for accessing the service, this tool may require the support of a dual stack or hybrid stack.

6.10.4 Functional Requirements

The possibilities for such tools are only just being explored so it is too early to discuss specific requirements. However, the following requirements are worth noting:

- Integrate the DNS and Directory Services server and database maintenance functions.
- Provide transport relay or application gateway information to enable a calling endpoint to identify a suitable intermediate system for the connection.
- Provide support to a transport relay or application gateway to map destinations in one network to the required destination in another.
- Provide support to dual-stack systems to assist in the selection of the appropriate stack for each connection.

6.10.5 Operational Requirements

There are no specific operational requirements applicable to this tool.

6.10.6 Standards Requirements

The IAB is currently addressing the issue of supporting DNS-style name services within an OSI Directory Service. X/Open and the X.400 APIA Association have published an API to the OSI directory service; this API may need extending with additional operations and objects to support DNS.

In addition, the other functional requirements listed above for this tool may require standardisation of additional objects and attributes in the directory.

6.11 INTEGRATED NETWORK MANAGEMENT

Where both IPS and OSI protocols exist in a single network, the problem of managing the network's resources becomes increasingly complex. The problem is compounded by the variety of methods that suppliers have adopted for managing network components, with the two open protocol standards both being capable of managing each others objects. It is essential that network management products recognise the long-term coexistence of these alternative, competing standards, and address the requirement for management in a multi-protocol network.

A number of products are becoming available which are capable of supporting multiple network management standards, including proprietary systems, from a single integrated management application. During the migration of the network from one protocol suite to another, when the configuration of the network can be highly dynamic, the support of an integrated management tool may be crucial to controlling the network's resources and managing the migration process smoothly.

6.11.1 Applicability

Such a tool has application to any policy which supports the coexistence of two protocol suites in a complex network.

6.11.2 Environment

Such a tool operates in a dual-stack environment to access the two networks. Alternatively, it may access proxy hosts to relay requests to nodes not contactable directly.

6.11.3 Functional Requirements

An integrated management tool must satisfy the following requirements:

- It must support the major, open management standards (SNMP, CMIS / CMIP, CMOT). In addition, emerging standards for managing devices which may not run a full OSI stack should be supported.
- It must provide a single model of the network being managed, converting between the objects in the model and those used by the underlying protocols.
- Where a network node is accessible from either protocol suite, it should be possible to manage all the node's resources using either protocol suite.
- It should support a range of options for protocol conversion. Allowing the conversion to take place on the management host, on the system being managed, or in an intermediate system (*proxy management*), depending on the performance and scalability requirements of the devices and management stations involved.
- Scalability, performance and administrative requirements make the support of distributed management essential.
- The need to shield the operator from knowledge about the underlying addressing of individual network elements make the support of an integrated directory service an essential requirement.

- When this application resides on a multi-stack host, it is particularly important that selection of the appropriate protocol stack for accessing a particular resource is automatic because of the potentially large number of resources which must be accessed.

6.11.4 Operational Requirements

There are no operational requirements specific to this tool.

6.11.5 Standards Requirements

As mentioned above, standards for management of simple network devices such as LAN bridges are being developed. In addition, the object definitions for manageable resources are also being developed and extended by organisations such as the IAB and OSI / NM Forum. These organisations are also working on the definitions required to support management of one protocol suite's resources using the management protocols of the other.

Chapter 7

Application to Scenarios

Chapter 2, **Problem Statement**, presents a number of scenarios where IPS protocols play a significant role in an organisation's networking strategy. The purpose of these real-life examples of IPS usage is to provide a context in which the techniques and policies presented in this guide can be considered, allowing their applicability and comprehensiveness to be assessed.

It is the objective of this chapter to complete the guide by showing how these techniques, policies and tools might be applied to the scenarios. To meet this objective, example coexistence and migration policies are presented here for each scenario. Each policy describes how OSI protocols might be introduced into the network; how existing IPS-based components and applications can coexist with the new protocols; and how some or all of the IPS infrastructure and applications might be migrated to OSI-based solutions. The policies presented are only examples, in most cases there are a number of possible solutions, each with its advantages and shortcomings.

The coexistence and migration of proprietary protocols is outside the scope of this guide. However, due to the widespread presence of proprietary protocols in the networks described, their significance cannot be ignored. Consequently, the policies make general statements about how the introduction of OSI protocols affects these protocols and how coexistence and migration might be achieved.

These policies have been developed by considering each scenario from the perspective of a consultant preparing proposals for a an organisation that is considering incorporating OSI protocols into its IT strategy. To provide a useful summary of the various aspects of each proposed policy or policy component, a migration checklist is included with the following sections:

Technology Technological implications of implementing the policy. These may include:

- New hardware and software components which must be purchased to implement the policy.
- New functions that are gained by implementing the policy.
- Old functions lost and additional restrictions introduced as a result of the policy.

Performance In general, this section lists performance considerations rather than actual figures. In the context of coexistence and migration, performance is considered relative to the performance of the system prior to the introduction of the OSI components and any coexistence tools. Performance characteristics may include the following aspects of a network.

- Response times
- Bandwidth
- Processing overhead of protocol software
- Memory overhead of protocol software
- Availability and resilience.

Costs Cost implications of implementing the policy. Those associated with the new hardware and software components and skills requirements are not listed here. These additional costs may include:

- The cost of developing, porting or modifying the client's application software to run in the new environment.
- Costs associated with extra bandwidth requirements.
- Costs associated with reduced system capacity.

Skills Additional skills required by users and operational staff as a result of implementing the policy. This may require training of users and operational staff, or may involve the acquisition of new staff with the required skill.

The checklist should include a time plan, to consider such aspects as availability of technology, availability of capital and other business aspects which affect when and in what order things are done. However, as such considerations are highly specific to particular organisations they are not considered further here unless there are specific sequence requirements which result from the policies proposed.

It should also be noted, as discussed in **Chapter 5**, **Policies**, that documentation is crucial to the successful implementation of any policy. Clearly, detailed documentation of the policy itself is required, setting out its objectives and the strategy for achieving them. In addition, user and operational documentation is required, both for the final system and for the intermediate states that are encountered during the migration.

As discussed in **Chapter 1**, **Introduction** and **Chapter 5**, **Policies**, the motivation of commercial organisations when considering adoption of new technologies, is that it allows them to conduct their business in new cost-effective ways. Whilst the policies presented in this chapter identify the potential benefits associated with adoption of OSI protocols, it is not practical to include detailed cost-benefit analyses. Consequently, for the purposes of this guide, it is necessary to assume that this basic motivation has been established prior to these policies being produced.

The subsequent subsections of this chapter discuss each of the scenarios introduced in **Chapter 2**, **Problem Statement**, in turn. Finally, a summary discusses how well the technical and policy sections of the guide have met the challenge of producing solutions to the scenarios.

7.1 SCENARIO 1 - BASIC IPS NETWORKING

This scenario addresses the basic functions that are found in most IPS networks. It is a LAN-based scenario in which a large number of UNIX workstations exchange files and, to a small extent, electronic mail. In addition asynchronous terminals access UNIX and other hosts by means of Telnet- and LAT-based terminal servers. The description of the scenario indicates that apart from the technical considerations, a migration policy must consider the business structure of this company. The hierarchical structure of the company and the devolved manner in which IT decisions are made by individual business units means that a simple corporate edict setting a policy for migration to OSI is unlikely to succeed on its own. Consequently, a corporate policy for the introduction of OSI is proposed which gives commercial incentives for business units to adopt the new technology and provides corporate support services to help them do so. The policy is presented as a Corporate IT policy plus a number of component IT policies which support it:

Corporate IT Policy -
This policy sets the corporate objectives for adoption of OSI-based products and lays down a strategy for achieving those objectives.

Appropriation Policy -
This policy lays down rules for the purchas of major new systems by the business units.

Workstation Migration Policy -
This policy describes how existing groups of IPS-based workstations may be migrated to use OSI protocols.

Workstation Coexistence Policy -
This policy describes how existing IPS-based systems interoperate with OSI-based systems.

Asynchronous Terminal Policy -
This policy considers the large population of asynchronous terminals attached to terminal servers.

Business Unit Infrastructure Policy -
This policy covers local management and administration strategy.

Corporate Infrastructure Policy -
This policy covers the corporate provision of OSI-based support services and gateways to promote the adoption of OSI technology and provide interoperability between OSI systems and the corporate data processing systems.

The proposed corporate and component policies are described individually below. For each component policy the migration checklist is used to summarise the key implications of the policy.

7.1.1 Corporate IT Policy

The corporate policy for adoption of OSI-based technologies has a number of objectives:

- Improve communications among business units and between them and the corporate IT users using electronic mail.
- Achieve wider availability of corporate and business unit data using file transfer.
- Facilitate the adoption of new technologies that enhance business performance (EDI, for example).
- Keep abreast of emerging technology in order to benefit from suppliers' investment in this area (it is perceived that future development of networking products will have an OSI bias).

To achieve these objectives a strategy has been developed which gives corporate support to OSI systems, provides perceptible benefits for units which adopt OSI systems, and places the cost and administrative burden of coexistence on those non-OSI users who need to interact with the OSI-based systems:

- Any major investment in new systems must include OSI protocol support capability (current or announced) in its communications requirements. This is justified by the need to ensure the future compatibility of major acquisitions.
- The corporate IT group will implement a company-wide initiative to promote distributed network management and administration systems. As part of this initiative, each business unit must develop a strategy for supporting local management of its network components and administration of its name and address spaces.
- The corporate IT group will implement a company-wide initiative to promote electronic mail connectivity amongst the company's business units and between them and the proprietary corporate IT systems. The intention is to provide an X.400-based mail backbone to improve intra-company communication.
- The corporate IT group will implement a company-wide initiative to promote file transfer connectivity between OSI-based systems and the proprietary corporate IT systems. The intention is to provide an incentive for business units to migrate to OSI protocols to gain access to corporate data using file transfer.
- The corporate IT group will provide central OSI services to support the corporate initiatives. These central services include application gateways into the proprietary corporate systems for electronic mail and file transfer, and provision of top-level directory services and systems management functions.
- Gateways between OSI- and IPS-based systems (other that electronic mail) must be provided and operated by the business unit wishing to make the access. The intention is to provide an incentive for migration of a business unit's systems to OSI protocols, to reduce administrative burden and improve access to data on OSI-based systems.

These strategies are implemented by a number of component polices, each addressing a related group of users or services. Each of these component policies is described below in more detail.

7.1.2 Workstation Migration Policy

This policy applies where a business unit wishes to migrate workstations and servers currently using IPS protocols to the OSI protocol suite.

The workstations are migrated in work groups, switching the stacks from IPS to OSI in all workstations in a group in one go. This means that in general, a workstation only supports a single protocol stack at any point during the migration. It also means that application gateways must be provided during the migration period to maintain connectivity with non-migrated work groups.

File Transfer Gateways

File transfer access to non-OSI file systems is by means of file transfer gateways operated by the business unit that benefits from the access. During migration of a number of associated work groups, temporary gateway provision may be required to maintain access between migrated and non-migrated groups. Once the migration is complete, continued gateway provision is only required if access is required to systems which are not to be migrated.

An exception to the single stack-per-system rule may be made where heavy file transfer traffic with non-OSI systems by a particular host would otherwise distort the requirements for gateway capacity. In such a case the workstation or server is provided with dual-stack capability to allow direct access to endpoints on both OSI and IPS networks.

File Transfer Application Migration

File transfer application support is migrated from the FTP utility to an equivalent application based upon X/Open's BSFT specification. This allows the services provided to the user (either interactive or command script) to remain virtually unchanged when a workstation migrates to OSI and FTAM.

Where users on dual-stack systems have interactive access to file transfer applications (as opposed to file transfer under the control of command scripts) the application must be switchable. This allows the user to use a single utility, regardless of whether the remote file system resides on an OSI- or IPS-based system. The application should be capable of selecting the appropriate protocol stack automatically (the user does not have to decide which is the correct stack to use for a request).

Where the file transfers are controlled by scripts, a switchable application is not essential. Each script can be customised to use the appropriate transfer utility. Alternatively the scripts can be designed to use a look-up table. When the network topology changes (as a result of file systems migrating to OSI), the scripts or look-up tables must also be modified.

Electronic Mail Application Migration

Application migration occurs in two stages to minimise the disruption to users. In the first stage, the existing mail agent is ported to the OSI environment. This provides the basic level of service that users rely upon and ensures that the user's filing systems, filtering, mail lists and other automatic mail applications should work in the new environment without too many changes other than to addressing information. The more sophisticated features of the OSI mail protocols are not accessible. Porting an IPS mail utility to X.400 is a significant development and may be in the control of the suppliers of the current mail utility.

The second stage migrates the users to a new X.400-based user agent. If the current usage of electronic mail is fairly unsophisticated, the cost of doing this is restricted to training the users to use the new mailer. It has the benefit of providing easy access to the additional functions of the OSI mail protocol.

Continued access to electronic mail users on IPS-based systems and to users on corporate IT systems is provided by the corporate X.400 gateways.

Other Considerations

In order to ease the management problems associated with reconfiguring systems and administering two related and dynamic name and address spaces during the migration period, an early implementation of integrated management and administration tools would be of significant benefit. To support the distributed management and administration strategy, new systems must be compatible with OSI network mangement and directory services.

Migration Checklist

Technology The following hardware and software components are required to support this policy:

- **OSI protocol stack** - An OSI stack is required for each workstation and each server. In general, workstations and servers only support a single protocol stack, thus there is no direct requirement for this stack to coexist with any other.
- **X.400 electronic mail application** - An X.400 User Agent, either with a local or server-based Message Transfer Agent, is required per workstation.
- **Portable IPS electronic mail application** - To support the easy transfer of users to OSI-based systems, the user agent used over IPS should be provided on top of the OSI protocol stack.
- **File transfer application** - A BSFT-based file transfer utility to support the current style of interactive and script-driven file transfer activities is required per workstation. Depending on the availability of standard FTAM APIs and software bundling policies, this application may be supplied by the OSI stack supplier or by a third-party software supplier.
- **File transfer gateways** - Where required for application access, one or more store-and-forward or end-to-end gateways to support access to files on IPS-based hosts. Such gateways may require dedicated processors to support the required transfer levels. Alternatively, the gateway software may reside on a dual-stack host which has sufficient spare storage and processing capacity.

- **Coexistent OSI Protocol Stack** - Where heavy file transfer traffic with IPS-based systems is required by a particular workstation or server, it can be provided with dual-stack capability to allow direct access to endpoints on both OSI and IPS networks.
- **Switchable Application** - Where users have interactive access to file transfer on a dual-stack system, a switchable application is required to utilise the appropriate protocol stack depending on the endpoint to be accessed. The application must support automatic selection of the appropriate stack given a simple host name for the remote file system.

In general, no application functions are lost as a result of this policy, however, where requirements for file transfer with IPS-only hosts do not justify the provision of a gateway, some access may be lost. Some gains in function may result from more sophisticated mail applications. In addition, file transfer access to the corporate hosts is possible.

Performance File transfer gateway throughput and resilience is the most significant performance aspect of this policy. Sufficient processing power and disk storage must be supplied to support the expected file transfer loads. In addition, it may be necessary to duplicate gateways to provide the appropriate level of resilience and availability.

It should be noted that the traffic doubling effect of relaying file data across a gateway must be considered when calculating LAN capacities. Siting of gateways should be considered to ensure that file traffic does not have to cross between LAN segments unnecessarily.

Cost The costs associated with this policy reflect the list of hardware and software components given above.

Skills Requirements for user training are equivalent to those required when any new system is installed. Where existing support staff are to take over responsibility for the administration and operation of the new system, there are significant training requirements to acquire the new skills needed.

7.1.3 Appropriation Policy

This policy applies to business units (and to the corporate IT facility) which are planning significant purchases of new systems. The appropriation policy requires that all such new systems either support the OSI protocol suite or have a definite upgrade path to use OSI in the near future.

Workstations and servers must support an OSI protocol stack and appropriate electronic mail and file transfer applications. The considerations relating to protocol stacks, provision of file transfer gateways, and selection of OSI-based applications, as discussed above in **Section 7.1.1**, **Workstation Migration Policy**, are applicable to this policy too. The users involved may not currently be using IPS-based applications and therefore compatibility with previous usage may not be an issue.

Migration Checklist

The check list for this policy is essentially the same as that given for the workstation migration policy given above.

In addition it should be noted that IPS software is often bundled into the price of the operating systems licence for many UNIX and derivative systems, especially those based upon the BSD variant. A decision to adopt OSI technology is likely to imply extra costs for purchasing an additional, unbundled, protocol stack and associated OSI applications. For some systems the OSI implementation may only be available from a third-party supplier.

7.1.4 Workstation Coexistence Policy

The coexistence policy addresses the requirements of workstations and servers that continue to support only IPS protocols.

Where there is a requirement to access files stored on OSI-based systems, individual business units must provide file transfer gateways. Access to the company-wide X.400 backbone is provided by the corporate X.400 gateways and should not require additional software or hardware within individual business units.

Migration Checklist

Technology New hardware and software requirements are as follows:

- **File Transfer Gateway** - Where work groups within a business unit require access to files on OSI-based hosts, one or more store-and-forward or end-to-end gateways are required to support access to files on OSI-based hosts.
- Such gateways may require dedicated processors to support the required transfer levels.
- Alternatively, the gateway software may reside on a dual-stack host which has sufficient spare storage and processing capacity.

Where file transfer gateways are not provided by a business unit, applications may lose access to files stored on IPS systems which subsequently migrate to OSI protocols.

Performance Where access to data is essential to important business functions, it is important that an appropriate level of throughput and resilience is provided. This may imply duplication of gateway capacity.

Where file transfer gateway traffic is high, there may be additional LAN capacity requirements due to the traffic doubling effect of the relay function.

Costs The equipment costs associated with this policy reflect the list of new hardware and software requirements given above.

Skills Additional operational and administrative skills are required to manage the file transfer gateways where provided. In addition, users of dual-stack systems may require training to use a new utility for file access or to understand host names in the new protocol domain.

7.1.5 Asynchronous Terminal Policy

This policy considers the large population of asynchronous terminals connected to terminal servers running Telnet and LAT, and accessing multi-user hosts. In the short term, migration of the terminal servers to OSI is not considered viable or desirable. There are a number of reasons for this.

- There is a fairly large investment in the current terminal server hardware. It may be possible to upgrade the current servers to support OSI, however this is unlikely to be viable without a memory upgrade to accommodate the extra protocol requirements.
- There are few host-based implementations of the VT protocol available.
- For character-based applications there are few clear benefits in adopting VT technology, there is however a possible performance penalty.

A consequence of this policy is that the UNIX and non-UNIX hosts on the network must continue to support the non-OSI virtual terminal protocols (Telnet and LAT).

Migration occurs in the medium to long term as a result of the following factors:

- Availability of OSI-based terminal servers and host VT implementations - few suppliers currently offer this option
- The need to buy new servers and to replace old ones
- A requirement to connect directly to OSI-only hosts which may be added to the network
- A desire to phase out IPS-based systems to simplify network management
- A need to remove the interconnectivity problems associated with supporting both Telnet- and LAT-based terminal servers and hosts

As a result of this non-migration, file transfer access may be required to OSI-based systems. See **Section 7.1.2**, **Appropriation Policy**, for discussion of how this access is achieved.

The migration checklist given for the workstation coexistence policy applies equally to this policy, but in the context of multi-user hosts rather than workstations,

7.1.6 Business Unit Infrastructure Policy

To support the distributed management and administration strategy, each business unit must develop a strategy for the management of its systems and administration of its namespace. As IPS-based systems are likely to persist for some time, the strategy must take account of the need to integrate management and administration tools for both protocol suites. The strategy must be compatible with the corporate provision of top-level support for these functions. Depending on cost considerations and the size of the configuration, a business unit may defer actual implementation of this strategy or contract the task to some other part of the organisation.

7.1.7 Corporate Infrastructure Policy

The corporate IT policy provides services to support the installation of OSI-based systems, and encourage migration of existing systems to the OSI protocol suite. There are three main parts to the policy:

- The provision of an X.400-based mail backbone to interconnect all corporate and business unit systems. The backbone consists of X.400 gateways to IPS and the proprietary corporate IT systems.
- The provision of file transfer gateways to the proprietary corporate IT systems.
- Top-level provision of distributed network management services to support local management of network resources by the individual business units. The central management tool controls the corporate-supplied parts of the network and allows the corporate IT function and individual business units to delegate control of certain resources.
- Top-level provision of distributed directory services to support local administration of namespaces and addressspaces by the individual business units. This allows inter-unit traffic to access address and O/R Name information for systems in other business units.

The currently installed LAN technology should support the OSI protocols without alteration.

Migration Checklist

Technology The following hardware and software is required to support the infrastructure policy:

- **Electronic mail gateway** - provision of store-and-forward application layer gateways to support electronic mail connectivity company-wide. This includes gateways from the OSI-domain to both IPS-based and proprietary corporate systems.
- **Proprietary file transfer gateways** - provision of e ither store-and-forward or end-to-end application layer gateways to support file transfer connectivity between the OSI domain and the proprietary corporate systems.
- **Directory services** - provision of top-level directory services. Allowing administration of the namespace for corporate hardware and services and supporting devolved administration of each business unit's own namespace.
- **Systems management** - provision of top-level systems management services. Allowing management of corporate hardware and services and supporting devolved management of each business unit's own resources.

Performance As centrally supplied services, prediction of the load to be supported is difficult, especially as none of the services involved are available in the current network.

Resilience and availability require some level of duplication of the gateways and services. Depending on the availability required, backup services may be manually or automatically switched. Initially, none of the services are mission-critical and a lower level of resilience can be tolerated, however the situation must be monitored to ensure that the quality of service is upgraded to match the level of resilience placed on it.

Costs

The costs implied by this policy reflect the list of software and hardware given above in the Technology section.

In addition there is a significant cost associated with configuring and operating the X.400 backbone and the other OSI support services.

Skills

The administrative and operational activities implied by the services being provided suggest that significant additional skills are required to implement these policies. Not just to implement the central services but also to support the business units within the distributed management and directory services applications.

In addition the business units require support in the form of technical and operational assistance, and user training to take advantage of the company-wide electronic mail connectivity.

7.2 SCENARIO 2 - IPS AS A COMMON INTERWORKING PROTOCOL

This scenario addresses the use of IPS as a common protocol for exchanging files between systems that support proprietary communications protocols, in this case IBM and DEC systems.

Current file transfer is controlled by scripts which collect files from a remote system during quiet periods on the network. Strict security is enforced at the system-level by allowing remote access only to specified spool directories and at the application layer by encrypting files that are placed there for collection. Format conversion of the file data transferred between hosts is performed by the hosts themselves, based upon knowledge of the differences in representation between the architectures. This is a manageable problem where only two hosts are involved, the number of potential transformation algorithms required increases significantly where more host types are involved in transfers.

The point of this scenario is to demonstrate that a function currently commonly performed by IPS is a good candidate for OSI products. As such, the proposed solution does not concern itself with the issues of migration, rather it proposes a policy for implementing the current system using OSI protocols. Two policies are proposed to provide the above functions using OSI protocols:

Basic Policy -
This is a simple replacement of the current FTP-based file transfer by the FTAM equivalent. The only benefit of this policy is a potential increase in the range of proprietary systems which can participate in such file transfers, as a result of wide-spread support of the OSI protocols on previously closed systems. Given the introduction of TCP/IP implementations into the proprietary mainframe arena in recent years, it will be some time (if ever in the case of some older operating systems) before OSI becomes more ubiquitous than its IPS predecessor.

Advanced Policy -
This policy aims to take advantage of some of the more advanced features of FTAM and the other OSI protocols to provide a more elegant solution to the user's needs. This policy uses high- and low-level APIs to the services of the FTAM protocol, consequently its implementation requires detailed knowledge of the FTAM protocol services and the FTAM Virtual File System. It should be noted that current implementations of FTAM and the upper layer protocols may not provide the advanced features required to implement this policy.

The two policies proposed are discussed below in more detail.

7.2.1 Basic Policy

The basic policy, implementing the current script-based file transfer using OSI protocols, requires an OSI protocol stack in each host. This stack may be required to coexist with a native protocol stack of the host system. In addition, to maintain compatibility with the current scripts which control file transfers, a BSFT-compatible file transfer utility is supporting an FTAM-based file transfer utility such as BSFT.

Migration Checklist

Technology The following software and hardware components are required to implement the basic policy:

- **Coexistent OSI protocol stack** - Depending on the OSI product strategies of the various host system suppliers, this stack may be supplied by them or by a third-party network software supplier. The stack should be able to share the LAN adaptors with the other protocol stacks accessing the LAN. However where this cannot be achieved, which may be the case if the OSI stack is supplied by a third party, it may be necessary to use a second adaptor to support the file transfer traffic.
- **File Transfer utility** - A BSFT-compatible file transfer utility which supports the current script-based FTP usage. Depending on the availability of standard FTAM APIs and software bundling policies, this application may be supplied by the OSI stack supplier or by a third-party software supplier.

The OSI solution should work with the same LAN technology currently supporting the IPS-based solution.

Using the BSFT-compatible utility, the functions delivered are virtually identical to that provided by the FTP-based system. In the short term, the interoperability of the FTAM solution is likely to be less than that of the existing solution. In the longer term, it is possible that wider interoperability may be achieved.

Performance None.

Costs The costs associated with this policy are those associated with the software and hardware listed above, plus the cost of rewriting the file transfer scripts to use the new file transfer utility. This should be trivial if BSFT-conformant utilities are used.

Skills A small amount of systems programming skills are required to modify the file transfer scripts to use the new file transfer utility and to use OSI names or addresses. In addition it is necessary to configure and operate an OSI network with a small number of nodes; this can be done using static configuration tables and manual management of the nodes involved.

7.2.2 Advanced Policy

There are three specific areas where advanced FTAM features may benefit this organisation, given an appropriate level of implementation of optional features in the participating OSI protocol stacks:

- **File-level Security** - The FTAM protocol supports an access control list file attribute for individual files, allowing specific actions on the file to be controlled by user name, application entity name and password. An FTAM responder implementation that supports this attribute can be used to enforce stricter access security than is currently possible using the FTP protocol.

- **Encryption** - Among the functions of the presentation layer, as defined by the OSI reference model, are encryption and decryption of the data being transferred. As discussed in **Chapter 2**, **Migration Endpoints**, there is currently no ISO standard method of encryption for use at this layer. However, given suitable agreements for encryption methods to be used, applications can use presentation layer encryption to protect the data being transferred automatically instead of having to encrypt and decrypt as an off-line activity.

- **Contents Transformation** - An important objective of the FTAM protocol is to shield applications from differences between filestores by implementing a virtual filestore and by defining a range of document types to represent common file structures. Applied to the scenario under consideration, this allows the format conversion, currently performed by applications which wish to share files, to be left to the file transfer applications. This may be achieved by using one of the record-based document types (NBS-6, NBS-7, NBS-8 or NBS-11) defined by the NIST OSIWG, or by using facilities to define and transfer application-specific document types. Both methods require a low-level interface to the services of FTAM. It should be noted that the former method relies upon document types that are currently not widely implemented (if at all), while the latter requires a sophisticated FTAM implementation and an ASN.1 compiler.

This policy has similar requirements to the basic policy for a coexistent OSI protocol stack, but in this case, instead of an FTAM-based file transfer utility, an FTAM implementation which provides a low-level interface to the FTAM services is required.

Migration Checklist

Technology New software and hardware required to support this policy are as follows:

- **Coexistent protocol stack** - This requirement is the same as for the basic policy.

- **FTAM implementation and API** - This must provide access to the FTAM and upper-layer features required to implement the customised file transfer applications proposed:

- **ASN.1 compiler** - For the record-based and application-specific document types it is necessary to produce ASN.1 definitions of the records and document types required for use by the application or the FTAM implementation.

This policy offers all the features of the FTP-based solution apart from its script-based simplicity. Additional functions delivered are integrated file access and encryption security features, and integrated file format conversion.

Performance The file transfer performance of this solution is likely to be slower due to the additional transfer-time functions implemented (encryption and format conversion). However, this is more than compensated for by the simplifications achieved by making file transfer a single stage operation.

Costs Apart from the costs associated with appropriating the software and hardware listed above, there is significant cost associated with the development of the FTAM-based application support software required to take advantage of the advanced FTAM features.

Skills A high level of system-programming and FTAM skills is required by this solution.

7.3 SCENARIO 3 - ENGINEERING WORKSTATION ENVIRONMENT

This scenario considers the exploration department of an oil company. To some extent it mirrors both the previous scenarios in its use of FTP to achieve file transfers between workstations and from workstations to non-UNIX minicomputers and mainframes. In this case however there is widespread use of non-networked systems and point-to-point connections to support specific file transfer requirements. This fragmentation of networking solutions has the effect of restricting opportunities to share data, processing and communications resources, and fragments the task of administration and operational control of the resources.

In contrast to the first scenario, there is no corporate IT policy in this organisation. There are however, centralised data-processing facilities based on IBM mainframes and accessed over an SNA network. In order to reflect this lack of centrally-defined IT policy and to contrast with the policies presented for scenario 1, a departmental policy is proposed here.

Rather than make specific proposals for migration of existing systems to OSI protocols, the policy concentrates on the provision of an infrastructure in which such a migration can occur at some point in the future, when business or technical consideration make it desirable. Given the points of similarity with the two preceding scenarios, it is assumed that the policies presented for their migration are applicable to this scenario also. This is particularly true of the policy for file transfer amongst heterogeneous systems presented for scenario 2, given the range of operating systems and computer architectures involved in this scenario.

The proposed policy for this scenario is presented as a departmental policy with a number of supporting component policies:

Departmental Policy -
This defines the objectives of the policy and introduces the overall strategy for achieving those objectives.

Departmental Connectivity Policy -
This policy proposes a strategy for expending network connectivity to all the departmental systems, where possible consolidating X.25 and point-to-point network connections into a departmental network.

PC Connectivity Policy -
This policy proposes installation of work group LANs with connectivity to the departmental network.

Infrastructure Policy -
This policy proposes a departmental administration and network management infra-structure to support current operations and facilitate migration to OSI at some point in the future.

7.3.1 Departmental Policy

The objectives of the departmental policy are as follows:

- Achieve global connectivity of all departmental hosts from PC to mainframe to facilitate sharing of data, CPU and networking resources.

- Position department to take advantage of advances in networking technology and to facilitate the development of a common networking environment for all applications.
- Integrate and concentrate administrative and operational activities to reduce resources required and increase effectiveness.

The strategy adopted in order to achieve these objectives does not in itself involve the installation of OSI-based networking components or the migration of current systems from IPS to OSI. Instead the strategy adopted is founded on the installation of multi-protocol router products which can tie the current combination of LAN, X.25 and point-to-point protocols into a single network, using the existing protocol suites. Similarly, a multi-protocol network management system is specified, providing integrated management based on both proprietary and open management protocols.

This strategy allows the current mix of protocols to continue to be supported in a single integrated network, while allowing the future introduction of OSI-based transport, application and management protocols without needing to upgrade the infrastructure in significant ways.

For network administration the principal tool required is a multi-protocol directory service, capable of administering the IPS, proprietary and OSI namespaces and providing support for the tools used during coexistence and migration. Until such an integrated solution is available, it is proposed to implement separate IPS and proprietary administration systems.

As discussed above this departmental policy is supported by three component policies which address Departmental Connectivity, PC Connectivity, and Infrastructure.

7.3.2 Departmental Connectivity Policy

This policy addresses the global connectivity objective of the departmental policy by building a single integrated network from the current mix of Ethernet LAN, X.25 and point-to-point links. To achieve this it is proposed that the current LAN routers are replaced with a small number of expandable high-performance multi-protocol routers. These routers form a resilient hub for the interconnection of all the current LANs and point-to-point links. The LANs are linked to the centrally located routers by remote repeaters using fibre or high-speed twisted-pair connections. All point-to-point or X.25 links currently connecting remote sites to local systems are re-terminated at router interfaces. (Exceptions to this are the SNA link accessed using the Telnet/3270 gateway and any links which use ad-hoc protocols or proprietary protocols that the routers cannot handle.)

The routers support a range of network-layer and routing protocols to handle both IPS and proprietary protocols. In addition, support of OSI internet and routing protocols is required, either now or in the future. Value added functions such as X.25 gateway and packet assembley/disassembley (PAD) may also be required to support some of the current connections.

Where possible, minicomputers which are currently without network connections are connected into a convenient LAN by installing an adaptor card and associated software, or are wired directly to the router using either X.25 or a point-to-point protocol.

All hardware and software installed must be compatible with the tools described in **Section 7.3.4**, **Infrastructure Policy**.

Managing the migration from the existing mesh topology is a sensitive matter if network availability is to be maintained, especially if existing plant and LAN components are to be re-used in new roles. It is necessary to migrate the current LANs in an incremental fashion, running the new routers in parallel with the existing ones to allow changes to be tested and reversed if problems occur.

Migration of the X.25 and point-to-point links is just as sensitive but less problematical as links can be reterminated and the appropriate routes installed on an individual basis. Unless there are other routing considerations to be taken into account, this process can happen independently from the LAN migration.

The following subsections use the migration checklist to discuss specific aspects of the interconnection policy.

Migration Checklist

Technology The following specific hardware and software components are required to implement the interconnection policy:

- **Multi-protocol routers** - Two or more high-performance routers are required to provide a resilient multi-protocol backbone network. The router selected should support the full range of devices to be connected, allowing expansion through additional interface cards. The range of internetwork and routing protocols supported is also an important consideration. Additional features such as bridging capability, PAD and X.25 Gateway capabilities, and SNA frame relay may be considered.

  It may be possible to use the currently installed routers in this role, however the hub topology of the new network differs significantly from the current mesh topology so they may not meet the requirements stated above.

- **Remote repeaters** - A pair is required for each LAN segment not co-located with the router to which it is to be connected. The halves of a remote repeater are connected using up to 1 kilometer of optical fibre. Where resilience is of key importance, a connection to more than one router must be provided.

  Again, it may be possible to use existing routers or bridges as a substitute for the remote repeaters. This has the advantage of allowing copper links as well as optical fibre, the choice depending upon plant availability, environmental considerations, distance and throughput.

- **Communications adaptors and software** - Each of the currently non-networked minicomputers requires either a LAN adaptor or point-to-point communications card and associated software to connect it into the new topology.

Performance The key elements in this policy are the router hubs. They must provide sufficient performance to support the required level of inter-LAN and remote traffic. By providing multiple routers at the hub and linking each LAN to more than one router, resilience is provided. Connection of remote network links to the router increases the resilience of the network by allowing them to be accessed by any host on the network, thus transfer and monitoring applications are no longer tied to a single host.

Costs The costs associated with this policy are primarily associated with the list of software and hardware given above. In addition there may be costs associated with changing current applications to work within the new topology. For example an application that collects events from a remote site may have to interface to a gateway implemented by the central router instead of accessing a direct X.25 interface.

Skills It is likely that the skills required to configure and operate the new routers are essentially the same as for the current mesh topology, thus no significant new skills are required for this policy.

7.3.3 PC Connectivity Policy

Networking of the departmental PCs is considered separately from the general connectivity policy because there is currently little or no networking of PCs in the department being considered.

The policy starts by connecting PCs into work groups based on Ethernet or Token Ring LANs with file and printer servers. The technology involved in such configurations is well understood and is not considered in detail here.

These LANs are then connected into the departmental network by supporting TCP/IP and in future OSI stacks on the LAN servers. File transfer, electronic mail and other network applications are then supported by *gateway server* applications which use an appropriate LAN transport protocol to link a PC-resident API or application stub with a server-resident OSI application. This avoids the need to support a TCP/IP or OSI stack on every PC in a network which requires one software licence per PC and consumes significant amounts of PC resources. Instead only the gateway servers require licences for the protocol software and resources to run it.

The migration checklist for this policy considers only the aspects associated with the interconnection of the PC LANS into the departmental network discussed above.

Migration Checklist

Technology The following hardware and software components are required to implement the PC interconnectivity policy:

- **Server protocol stack** - Each server host that is to provide gateway services between a LAN and the departmental network requires a protocol stack to support the required applications. In the short term this is likely to be an IPS stack, in the future an OSI protocol stack is required, possibly coexistent with the IPS stack.

- **Network connection** - The physical connection from the PC-based LAN into the departmental network may use a repeater connection into the hub routers. Alternatively, it may be a connection into a local Ethernet LAN using a bridge. The former method of connection requires a router interface for each LAN to be connected. Where a high level of resilience is required, these connections may be duplicated.

- **Client/server applications** - In order to access network resources, the PCs and servers implement split applications where the user interface is resident on the PCs and the rest of the application, including the network interface, is resident on a server host which supports the appropriate protocol stack. Initially the applications required are IPS-based file transfer, remote login and electronic mail. In the future OSI-based applications will be required. Additional applications might include remote database access applications.

As there is currently little PC networking there are only gains in function as a result of this policy.

Performance In this policy, the server hosts are key to the performance. By acting as gateways to OSI, the servers must have sufficient processing power, memory and storage capacity to support the PCs in a work group. The MAC-layer bridges connecting the work group LAN to the departmental LAN must be duplicated if resilience is required.

Costs Apart from those associated with the additional components listed above, there are costs associated with the development of applications to take advantage of the new connectivity. However, as these costs are directly linked to additional functions rather than to migration of existing applications, they need not be justified in this checklist.

Skills Excluding the skills required to set up the general PC network, additional skills are required to install and manage the gateway protocol stacks and to implement the applications required.

7.3.4 Infrastructure Policy

The infrastructure policy specifies network mangement and directory services required to configure and control the newly integrated departmental network. The intention is to install tools that can manage and support the current IPS and proprietary protocols and applications, and are capable of managing OSI-based applications and protocols as they are introduced into the network in the future.

Migration Checklist

Technology The following hardware and software components are required to implement this policy:

- **Multi-protocol network management** - Management stations are now available which support both IPS and OSI style network management plus a number of proprietary management protocols. There are a number of aspects which must be considered when selecting a management tool:

  — Protocol conversion - There are a number of methods used to integrate the underlying protocols, ranging from full protocol conversion on the management station, to doing the work in the component being managed.

  — Protocol stack support - It is not appropriate to support a full protocol stack on some network devices (such as bridges). There are standards emerging for the management of such components using only data link layer protocols.

  — Corporate management policies - Future requirements to integrate with corporate network management standards may also require consideration.

- **Directory services** - It is proposed that initially, due to a lack of integrated directory products, administration and service of the name and address spaces of each supported protocol are addressed separately.

  In the future, migration to an integrated solution should be considered to support the dynamic situation associated with migration from one protocol suite to another, and to provide support for the gateways, relays and multiple stack hosts that are used during such a migration.

Performance The selected tools must have sufficient capacity to support the event logging and directory look-up loads expected in the departmental network.

Costs None other than those associated with the component list above.

Skills The installation and operation of management and administration tools implies significant new skill requirements.

7.4 SCENARIO 4 - SUMMARY OPERATIONAL REQUIREMENT

Scenario 4 involves a government organisation which is installing a large OSI-based network based upon distributed site LANs tied together with a private X.25 WAN. The only element of coexistence and migration in the scenario involves integration of two existing LANs which connect a small number of IPS-based UNIX systems.

Presented as an operational requirement, the scenario requires calculations of the theoretical availability and capacity of the network proposed. Whilst the capacity of the proposed network is obviously of prime importance in a real situation, it is not directly relevant to this guide and is not calculated here. Resilience and availability are calculated to show the effect of the extra components required by the proposed migration policy (all calculations assume an MTTR of eight hours). The scenario requires conformance to the UK Government's GOSIP profile. Version 3.1 of the profile, dated 31/1/90, has been used in these policies.

As with previous scenarios, the proposed solution is presented as a number of policies, each addressing a particular aspect of the proposed system. In this case the proposed policies are:

Connectivity Policy
This policy addresses the LAN/WAN infrastructure and the protocol stacks required to support the proposed network. This includes the provision of the required performance, availability and resilience targets.

Application Policy
This policy addresses the initial applications to be run on the networks hosts.

IPS Host Policy
Coexistence and migration of the current IPS-based hosts and LANs to OSI protocols are addressed here. This includes the effect on the performance, availability and resilience of the systems during the migration.

Management Policy
Provision of central and local network management and administration systems.

7.4.1 Connectivity Policy

This policy reflects the LAN/WAN topology presented in the statement summary operational requirements given in the scenario. The diagram presented in **Chapter 2**, **Problem Statement**, should be consulted.

OSI LAN Profile

The OSI-based hosts and terminal servers at each major and minor site are grouped on a single 8802/3 LAN. The existing IPS-based LANs are considered below under **IPS Host Policy**. The specified profile for the LAN (UK GOSIP-T) allows either a connection-oriented or connectionless data link layer. Given the limited availability of connection-oriented products, the following protocols are specified at layers 1-4:

Layer	Standard	Comment
Physical layer	8802/3	Ethernet LAN.
Data link layer	8802/2 (LLC Type 1)	Connection-less data link layer protocol.
Network layer	8473	Connection-less network service (CLNS).
Transport layer	8072 (TP4)	Connection-oriented transport service (COTS)

The variant of 8802/3 to be used (for example 10BASE5, 10BASET, and so on) does not affect the protocol layers above. Consequently, the physical aspects of the LAN technology to be used are not discussed here as not enough is known about the geography of the sites.

Terminal and Printer Connectivity

New terminals are connected to the major and minor LANs using terminal servers. Additional servers can be added as the terminal population expands.

As discussed in the policies proposed for scenario 1, currently there are few host-based implementations of the OSI Virtual Terminal Protocol. In this scenario, however there is no large installed base of IPS-based terminal servers to consider. The solution depends on whether an interoperable combination of terminal server and host-based OSI protocol stacks can be obtained. If not, then IPS-based terminal servers must be provided initially. This has a number of operational implications:

- Hosts must support dual protocol stacks.
- Both IPS and OSI domains must be managed and administered.
- In the future the terminal servers and hosts access software must be migrated to OSI.

In order to meet the resilience requirements, 16-port servers are specified to limit the number of terminals affected by a single terminal server failure. For the purposes of the availability calculations, the MTBF figures for these devices is assumed to be of the order of 50000 hours. This gives a component availability (termed Ats) of 0.999840.

Remote Terminal and Printer Connectivity

Outpost sites are connected to a convenient major site using leased lines. Up to 16 devices (terminals and printers) are connected using a statistically multiplexed 9.6 kbps line. At the major site the terminals are connected to the LAN as discussed above for local terminals. The limited number of devices involved mean that resilience requirements can be met without providing backup for the leased line and modems. Availability has not been calculated for these devices.

UNIX Hosts

Major sites are supported by two UNIX hosts connected to the LAN, minor sites are supported by a single host. Resilience requirements are met because any failure isolates only the host in question. For the purposes of availability calculation, the hosts are assumed to have MTBF figures of the order of 20000 hours, giving an availability for a UNIX host (termed Aux) of 0.99960.

WAN Connectivity

Major sites are connected over a private X.25 network - currently X.25 (1984) but to be upgraded to X.25 (1988). Each LAN is connected to the WAN using dual routers for resilience. The routers use the *Network Tunnel* technique to transfer the connectionless network layer packets across the WAN in a single routing hop. ES-IS protocol is used to route data on the local LANs (in this context *local LAN* includes all minor site LANs bridged into the major site LAN).

If the terminal servers installed to support terminal-host access are IPS-based, the routers must also transport IP packets across the WAN. This requires a multiprotocol router, which also routes IP packets and supports an IP routing protocol such as RIP on the local LANs.

The resilience requirements are met by duplicating the routers, ensuring that they are connected to separate switches on the WAN. For the purposes of the availability requirements, the routers are assumed to have an MTBF of 50000 hours, giving a component availability (termed Art) of 0.999840. The WAN availability is quoted as 0.999999 for the two port configuration (termed Ax25), this is assumed to include the availability of the connections between the routers and the appropriate switch nodes. The availability of two routers in parallel (termed Art2) is:

$$\mathrm{Art2}=1-((1-\mathrm{Art})*(1-\mathrm{Art}))\approx 1$$

Consequently the availability of the LAN-LAN connections across the X.25 WAN (termed Awan) is effectively 1.

Major/Minor Site Connectivity

The minor sites are each connected to a major sites by duplicated MAC-Layer bridges. These bridges use the IEEE 8802/1 spanning tree algorithm to ensure that only one bridge is in operation at any time.

The duplication of the bridges provides the required resilience. The bridges are assumed to have an MTBF of 50000 hours, giving a component availability (termed Abg) of 0.999840. In addition, the leased line connecting a pair of remote bridges (termed All) is assumed to have an availability of roughly 0.9990. These are combined to give an availability for a pair of remote bridges of:

$$\mathrm{Abg}*\mathrm{All}*\mathrm{Abg}=0.99868$$

Two of these in parallel give a combined availability (termed Amm) of effectively 1.

Migration Checklist

The migration checklist for the connectivity policy is given here:

Technology The software and hardware required to support this policy is as follows:

- **Terminal Servers** - 16-port terminal servers supporting asynchronous terminals and printers. Depending on the decision concerning the use of the OSI VTP protocol these must either support Telnet over the IPS protocols or the appropriate OSI protocol stack. If IPS-based servers are chosen, they must be upgradable (software and any required memory expansion) to a suitable OSI stack at some point in the future.
- **Host Protocol Stack** - Each UNIX host attached to the new LANs must support an OSI stack; the profile for the stack is listed above. Where IPS-based terminal servers have been chosen, or where required to support the Migration Policy for existing UNIX hosts, this stack must coexist with an IPS protocol stack.
- **Multi-Protocol Routers** - Each major site LAN requires a pair of routers to route OSI CNLP and IP packets across the X.25-based WAN. The router must support both ES-IS and RIP routing protocols on the local LAN. As the protocols are being tunnelled through the WAN, the routing protocols used between the routers is a private matter as long as all routers are provided by a single supplier.
- **Remote Bridges** - For each minor site to major site link two pairs of remote MAC-layer bridges are required. The bridges must support the 8802/1 spanning-tree algorithm to ensure that only one bridge is active at any one time.
- **Digital Links** - For each minor site to major site link two 64kbps digital links are required to link the remote bridges. Traffic calculations have not been performed to determine the required data rate for these lines.
- **Multiplexers and Modems** - For each outpost site a pair of modem and multiplexers are required to link up to 16 terminals and printers to an appropriate major site.

Performance Traffic figures and response times have not been calculated for this policy for the reasons discussed above. The resilience aspects of the policy are addressed by duplication of all inter-LAN links, giving virtually 100% availability of these key connections.

In order to calculate the minimum availability of the path between two devices on the network the availabilities are combined for all the components on the appropriate paths. The path considered, minor site terminal to remote minor site UNIX host, is the longest path on the network. The overall availability combines the following component availabilities, as defined in the sections above:

$$Ats*Amm*Awan*Amm*Aux=99.94\%$$

This is just within the requirement.

Costs The costs associated with this policy reflect the hardware and software listed above. In addition, there are the usual costs associated with setting up a major new network; these are outside the scope of this guide.

Skills There are significant additional skills associated with setting up a major new network; these are outside the scope of this guide.

7.4.2 Applications Policy

The operational requirement specifies that users of the UNIX hosts are to be provided with a menu based interface to the resources of the network. In the initial configuration, the OSI-based UNIX hosts on the network support the following applications:

- **File Transfer** - File transfer operates both within the local site, between bridged major and minor sites, and between sites separated by the WAN.
- **Electronic mail** - A low level of electronic mail usage is envisaged initially. As there is no significant usage of electronic mail currently, an X.400-based user agent can be adopted without causing disruption.
- **Printing** - Printers are connected to terminal server host ports; the UNIX hosts must connect to the appropriate port in order to service a print job. As no standard applications exist to support network printing, it is necessary to purchase a proprietary package (if one exists), otherwise a new program must be developed.
- **Virtual Terminal** - For virtual terminal access between terminal servers and UNIX hosts, the GOSIP profile specifies use of the VT profile. Two VTE profiles are required, The forms profile supports screen-based applications; the Telnet profile supports character-at-a-time applications. These profiles must be supported both by the terminal servers and the hosts involved in virtual terminal access.

  Depending on the resolution of the issues raised in the Connectivity Policy section regarding the availability of an OSI VTP implementation, virtual terminal access may initially be provided by the IPS Telnet protocol.

  Where terminals are connected directly to a specific UNIX host, a remote login utility is required to allow such users to access other hosts in the network.

Migration Checklist

Technology The following software components are required to support the application policy:

- **File Transfer** - Each host requires a BSFT-conformant file transfer utility to provide interactive and script-driven file transfer. In addition, each host that is to be the target of file transfer requests requires an FTAM responder which supports at least the BSFT responder profile.
- **Electronic mail** - Each host must support an X.400-based *User Agent* and *Message Transfer Agent* to support electronic mail access.
- **Virtual terminal** - A host implementation of VTP is required on each host to support connections from server-connected terminals. In addition, where there are directly connected terminals or where required to support the IPS host policy discussed below, a virtual terminal utility is required to support access to other hosts.

Performance There are no specific performance requirements associated with this policy.

Costs The only costs associated with this policy are those associated with the software components listed above.

Skills The skills required to access the applications in the new network can be incorporated into the training required for the applications themselves.

7.4.3 IPS Host Policy

At two sites there are currently groups of UNIX hosts, connected to a LAN and running IPS protocols. In the absence of information in the statement of the scenario, it is assumed that users' terminals are connected directly to the IPS hosts, the LAN being used for file transfer and electronic mail.

To access the new OSI-based network, each IPS LAN is connected, using a single MAC-layer bridge, to the new OSI-based LAN at its major site. This policy has a coexistence phase and a migration phase.

During the coexistence phase, the hosts continue to run IPS protocol stacks. Access to the OSI-based hosts is provided by dual-stack hosts on the local major site LAN acting as gateways.

Virtual terminal access to OSI-based hosts is achieved by logging into the local gateway host using the Telnet protocol and then using a VTP-based utility to connect to remote hosts. This requires the users affected to use two utilities to access certain hosts. Use of scripts to support remote access may simplify the process, however the additional steps required to access the host running the application can be incorporated into the training required.

File transfer access to file systems resident on OSI-based hosts is achieved using relay scripts executed remotely at the gateway host. Thus the file is first transferred from the remote host into the file system of the gateway host using an FTAM-based utility, it is then transferred to the IPS-only host using the FTP utility.

Migration of each IPS-based LAN is performed in a single move due to the small number of hosts involved. Each host requires software as specified above for the the connectivity and application policies.

Migration Checklist

The migration checklist for this policy reflects those for the connectivity and application policies presented above.

Technology Additional hardware and software requirements for the IPS host policy are listed here:

- **Local Bridge** - Each IPS-based LAN requires a single MAC-layer bridge to connect it to the local major site LAN.
- **Coexistent IPS Protocol Stack** - Each major site that supports IPS-based hosts requires one of the new OSI-based UNIX hosts to support both an OSI and an IPS protocol stack for the period of the coexistence.
- **Virtual Terminal Utility** - In order to support IPS-user login to OSI-based hosts, the gateway host must support a VTP-based virtual terminal utility. This may be combined with similar requirements to support users on terminals which are directly connected to OSI-based hosts.
- **File Transfer Scripts** - Scripts must be developed to support the relay of files between the IPS and OSI-based hosts.

Performance The connection of the IPS host LANs using a single local MAC-layer bridge implies that all hosts on the existing LAN lose connectivity when the bridge fails.

The availability figures for the users directly connected to UNIX hosts on the IPS LANs are affected by the lower MTBF figures for a terminal connected to a host as opposed to a terminal server. In addition, there is a single local LAN bridge in the connection path in place of the dual remote bridges used to connect major and minor sites. This gives modified availabilities as defined in the performance section of the migration checklist for the connectivity policy:

$$Aux*Abr*Awan*Amm*Aux=99.904\%$$

This does not meet the level specified by the operational requirement. It is clear that terminals need to be connected to terminal servers to achieve the availability targets. During the coexistence period, the requirement to perform remote login or file transfer using a gateway host further reduces the availability figure:

$$Aux*Abr*Aux*Awan*Amm*Aux=99.864\%$$

During the period of coexistence, virtual terminal sessions to OSI-based hosts experience slower response times. This effects character-at-a-time applications in particular. It is likely that character echo response times may be increased by 20 to 100 milliseconds.

Support of gateway connections from IPS-based hosts also has a significant affect on the performance of the host being used as a gateway. During the period of coexistence it may be necessary to reduce the number of users which the gateway hosts support to compensate.

Costs There are costs associated with developing and maintaining the scripts which support file relay.

Skills The migration of the IPS systems to OSI protocols does imply significant disruption of the systems users. Selection of IPS-like utilities for file transfer and virtual terminal access can minimise this.

7.4.4 Management Policy

The scenario states explicit requirements for a hierarchical management system allowing both centralised and local control of the network's resources. Distributed management and administration has been covered in the earlier scenarios. The discussion and migration checklists given in those sections of this chapter are equally applicable to this scenario with some additional complexity due to the need to control the WAN-based inter-site links.

7.5 SCENARIO 5 - USER CONCERNS

In this scenario, an academic research organisation has a hierarchical network of workstations, file servers, database servers and computation servers. The network is based around a optical fibre backbone LAN, with a number of workstation LANs, each bridged onto the backbone by means of its file server. The major application is file transfer, with some virtual terminal access required for running applications remotely.

This scenario reflects scenario 1 in the sense that the primary concern of any strategy proposed for introduction of OSI protocols must be to overcome the resistance of those who plan and run the network. In this case it is an academic environment rather than a commercial one; as a consequence the current network is almost exclusively based upon IPS protocols rather than being dominated by proprietary networks. As a government-funded organisation in the USA, pressure for adoption of OSI protocols is likely to come from funding departments in order to comply with federal standards. In addition, despite the opinions to the contrary, OSI-based access to business partners and other research organisation is becoming increasingly viable using OSI protocols, particularly X.400-based electronic mail.

In this environment overcoming prejudice is the key to success rather than proving the cost justification for any change. For this reason, the policy presented for this scenario does not address migration of the underlying infrastructure at all at this stage. Rather it proposes the introduction of OSI applications on top of the existing IPS infrastructure. In addition, a transport relay gateway provides access to external networks and to the emerging global X.400-based electronic mail community.

The resulting **OSI Application Policy** is described below.

7.5.1 OSI Applications Policy

The objectives of the policy proposed here are as follows:

- Provide access to business partners and other research organisations both in North America and worldwide for the exchange of data. The description of the scenario says little about how this is achieved in the current network. The OSI strategies of the organisations involved can be used as rationale for this objective.
- Gain experience in the use of OSI-based application protocols. By running OSI applications over the current IPS infrastructure, the operational aspects can be mastered within a reliable and stable environment.
- Preserve the quality of the organisation's service by retaining the proven IPS-based infrastructure.
- Position the organisation for the installation of an OSI-based infrastructure at some point in the future. Having gained experience in the use and operation of application protocols, the migration to OSI-based transport protocols can be achieved with little disruption to users. This wait and see approach avoids the need to make an early choice about whether to use a connection-oriented or connectionless network service.
- Demonstrate active commitment to government objectives for the adoption of OSI. It should be possible to delay consideration of OSI-based infrastructure until a major upgrade or purchase of new equipment is required.

To achieve these objectives the following strategy is proposed:

- Provide an OSI application environment based upon the upper layers of the OSI stack running over a TCP-based transport layer. This *hybrid stack* technique uses the RFC-1006 convergence protocol to provide an OSI TP0 transport service over TCP virtual circuits.
- Run file transfer and electronic mail applications over this hybrid stack.
- Provide a transport-level gateway to external organisations. This gateway uses the *transport relay* technique to convert from the hybrid stack used in the local environment to a full OSI stack. The relay can provide interface to OSI stacks with either connection-oriented or connectionless network layers.
- Where other organisations are running OSI applications over a compatible hybrid stack, the applications can communicate directly.
- Provide a directory service application to support the OSI applications and provide the additional addressing required to support transparent operation of the transport relay.
- Upgrade the current network management systems to support the hybrid stack and OSI applications.

The migration checklist for this policy is given below:

Migration Checklist

Technology The following software and hardware components are required to implement the OSI application policy:

- **Hybrid stack** - Each workstation and server requires an implementation of the OSI upper layers which uses the RFC-1006 convergence protocol to run on top of an IPS-based transport service. The convergence protocol uses the local transport interface (either sockets, TLI or XTI) to access the TCP service.
- **X.400 user agent** - Each workstation requires an X.400-based mail utility to allow users access to OSI-based mail.
- **X.400 message transfer agent** - Each file server runs an X.400-based mail switch to support the mail agents running on the workstations.
- **File transfer utility** - Each workstation and host which supports either interactive or script-driven file transfer requires a BSFT-based file transfer utility.
- **File transfer responder** - Each server which allows remote systems access to its files must support an FTAM responder.

- **Transport relay gateway** - For access to external OSI-based networks one or more gateways provide a transport-level relay function between the following protocol stacks:

  — TP0 over CONS, for access to X.25-based connection-oriented networks.
  — TP4 over CLNS, for access to connectionless networks.
  — TP0 over TCP (RFC-1006), for access to the local network.

  Ideally, a transport relay should be a dedicated implementation, rather like the current router implementations. As there are no dedicated implementations currently available, the most likely solution is a microcomputer with the appropriate set of network interfaces and transport stacks. The relay application and associated convergence protocols can then be run over the native transport interface.

- **Directory Service** - A simple, non-distributed implementation of the OSI directory service should support the presentation address and O/R name look-up requirements of this network. It is essential that the implementation selected supports a presentation address format which can accommodate an IP address as an NSAP. It is also desirable that the directory service supports the additional addressing requirements of the transport relay for both outgoing and incoming connections (it is not currently clear how this is achieved).

Performance The key to performance in this policy is the transport relay. Clearly the bandwidth of the gateway must support the required level of file transfer and electronic mail traffic, initially this is not expected to be high due to the lack of current access.

Costs Apart from the costs associated with the list of hardware and software components listed above, implementation of the policy should be regarded as a development project because of the novel ways in which some of the components are being combined.

Skills Operational and user skills associated with the new applications must be developed. In addition, as mentioned above in the **Costs** section, the novel ways in which some of the components are being combined suggests that protocol integration skills are required to implement the policy successfully.

Appendix A

Glossary

Abstract Syntax Notation 1
(ASN.1) A notation defined in the OSI standards that allows data to be described in a machine-independent fashion.

Advanced Research Projects Agency
See *Defense Advanced Research Projects Agency*.

Advanced Research Projects Agency Network
(ARPANET) Network created by DARPA running the Internet Protocol Suite (IPS).

API
See *Application Programming Interface*.

Application Programming Interface
(API) A set of functions used by an application programmer. For example, the TLI application programming interface is used to interface with the *transport layer* of the network.

ARPANET
See *Advanced Research Projects Agency Network*.

ASCE
See *Association Control Service Elements*.

ASN.1
See *Abstract Syntax Notation 1*.

Berkeley Software Distribution
(BSD) A popular variant of the UNIX operating system provided by the Regents of the University of California at Berkeley.

BSD
See *Berkeley Software Distribution*.

CCITT
See *Consultative Committee of International Telegraph and Telephone*.

Consultative Committee of International Telegraph and Telephone
(CCITT) An international committee whose membership is composed of government postal, telephone, and telegraph agencies (PTTs).

DARPA
See *Defense Advanced Research Projects Agency*.

Defense Advanced Research Projects Agency
(DARPA) An agency of the U.S. Department of Defense whose networking project resulted in the creation of ARPANET.

Destination Service Access Point
(DSAP) The LLC service access point value defined in ISO 8802/2 and IEEE 802.2.

Detailed Network Interface
(DNI) A network-independent application programming interface specified by the IEEE P1003.12 working group.

DNI
See *Detailed Network Interface.*

DSAP
See *Destination Service Access Point.*

Estelle
A language for specifying a state machine.

Federal Information Processing Standard
(FIPS) An information processing standard established by the National Institute for Standards and Technology.

File Transfer Protocol
(FTP) IPS application for transferring files over the network.

FIPS
See *Federal Information Processing Standard.*

FTAM
See *File Transfer, Access and Management.*

File Transfer, Access and Management
An OSI defined service for accessing and managing files across the network.

FTP
See *File Transfer Protocol*

GOSIP
See *Government OSI Profile.*

Government OSI Profile
(GOSIP) A government-defined network profile based upon a subset of the OSI protocol suite. It is designed to guarantee that all conforming implementations interconnect without problems.

IAB
See *Internet Activity Board.*

Institute of Electrical and Electronics Engineers
(IEEE) Organisation of engineers and engineering organisations that defines standards such as the 802 networking standard.

Internet Activity Board
(IAB) The organisation that guides RFCs through the standardisation process. The IPS standards were produced through this process.

IA5
International Alphabet 5 - similar to ASCII.

IEEE
See *Institute of Electrical and Electronics Engineers.*

International Standards Organization

(ISO) A standards organisation with the membership composed of the standards organisations from each participating country. OSI working groups generate the OSI protocol suite standards.

Internet

Used generically to refer to interconnected networks, particularly the ARPA internet.

Internet Protocol

(IP) Internetworking protocol used in IPS networks. It is normally used in conjunction with TCP (TCP / IP).

Internet Protocol Suite

(IPS) The protocol suite used by ARPANET that is composed of the Transmission Control Protocol (TCP) and the Internet Protocol (IP). It is often referred to as *TCP/IP*.

IP

See *Internet Protocol*.

IPS

See *Internet Protocol Suite*.

ISO

See *International Standards Organization*.

LAN

See *Local Area Network*.

LLC

See *Logical Link Control*.

Local Area Network

(LAN) A data network with high speed and low latency, usually connecting computer equipment in a defined area, such as a building.

Logical Link Control

(LLC) The network sublayer defined ISO 8802 / 2 and IEEE 802.2.

Message Handling Service

(MHS) A service defined by CCITT in the X.400 standards for the transfer of mail and other types of messages over a network. ISO has adopted this standard for the OSI protocol suite where it is referred to as MOTIS.

Message Transfer Agent

(MTA) The application that provides the message transfer service for MHS.

MHS

See *Message Handling Service*.

MTA

See *Message Transfer Agent*.

National Institute for Standards and Technology

(NTIS) The division of the U.S. Department of Commerce that creates standards for use within the U.S. Government. It was formerly known as the National Bureau of Standards.

Network File Service

(NFS) An application that provides remote file access over IPS networks.

NFS

See *Network File Service*.

NIST

See *National Institute for Standards and Technology*.

OSI

See *Open System Interconnection*.

POSIX

An IEEE proposed standard for a portable operating system.

Postal, Telephone and Telegraph

(PTT) Any government agency that controls or regulates these activities. CCITT, which creates many networking standards, is made up of PTTs from the member countries.

PTT

See *Postal, Telephone and Telegraph*.

Remote Operations Service Element

(ROSE) A service element in the applications layer of the OSI protocol suite. It manages Remote Procedure Calls (RPCs).

Remote Procedure Call

(RPC) A call by one endpoint in a communications link for the other endpoint to perform a procedure.

Request for Comments

(RFC) Documents made available by the IAB. These documents are distributed to the Internet community for comment and revision. Many of these documents are now Internet standards.

RFC

See *Request for Comments*.

ROSE

See *Remote Operations Service Entry*.

RPC

See *Remote Procedure Call*.

Simple Mail Transfer Protocol

(SMTP) Protocol used for the transfer of mail on IPS networks.

Simple Network Interface
(SNI) A network-independent application programming interface specified by the IEEE P1003.12 working group.

Simple Network Management Protocol
(SNMP) A protocol for managing IPS networks.

SMTP
See *Simple Mail Transfer Protocol*.

SNI
See *Simple Network Interface*.

SNMP
See *Simple Network Management Protocol*.

Streams
A software interface to the network provided in UNIX System V.

TCP
See *Transmission Control Protocol*.

TCP/IP
See *Transmission Control Protocol* and *Internet Protocol*.

Telecommunications Network Protocol
(TELNET) An application available on many IPS networks that provides remote login to other computers.

Telnet
See *Telecommunications Network Protocol*.

TFTP See *Trivial File Transfer Protocol*.

TLI
See *Transport Layer Interface*.

Transport Layer Interface
(TLI) A set of standard network interface functions available on UNIX System V. XTI is the X/Open standard, derived from TLI.

Trivial File Transfer Protocol
(TFTP) IPS application for transferring files over the network. It is based upon an unreliable transport service and provides no authentication.

UA
See *User Agent*.

User Agent
(UA) The application generates and receives messages for MHS. Usually, the user invokes this application to read and create mail.

X.400
Message Handling Service (MHS) defined by CCITT. This standard has been adopted by ISO under the ISO 11021 standard and is referred to as MOTIS.

X/Open Transport Interface

(XTI) A set of standard network interface functions defined by X/Open. This standard is derived from TLI.

XTI

See *X/Open Transport Interface*.

Index