

CAE Specification

Networking Services (XNS) Issue 5

The Open Group



© February 1997, The Open Group

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the copyright owners.

Portions of this document are extracted from IEEE Std 1003.1-1990, copyright © 1990 by the Institute of Electrical and Electronics Engineers, Inc. with the permission of the IEEE.

Portions of this document were extracted from IEEE Draft Standard P1003.2/D12, copyright © 1992 by the Institute of Electrical and Electronics Engineers, Inc. with the permission of the IEEE. No further reproduction of this material is permitted without the written permission of the publisher. IEEE Std 1003.2-1992, copyright © 1992 by the Institute of Electrical and Electronics Engineers, Inc., and ISO/IEC 9945-2:1993, Information Technology — Portable Operating System (POSIX) — Part 2: Shell and Utilities, are technically identical to IEEE Draft Standard P1003.2/D12 in these areas.

Portions of this document are derived from copyrighted material owned by Hewlett-Packard Company, International Business Machines Corporation, Novell Inc., The Open Software Foundation, and Sun Microsystems, Inc.

CAE Specification

Networking Services (XNS) Issue 5

ISBN: 1-85912-165-9

Document Number: C523

Published in the U.K. by The Open Group, February 1997.

Any comments relating to the material contained in this document may be submitted to:

The Open Group
Apex Plaza
Forbury Road
Reading
Berkshire, RG1 1AX
United Kingdom

or by Electronic Mail to:

OGSpecs@opengroup.org

Contents

Chapter 1	Common Information.....	1
1.1	Terminology.....	1
1.2	Use and Implementation of Interfaces.....	2
1.2.1	C Language Definition.....	2
1.3	The Compilation Environment.....	3
1.3.1	The Name Space.....	3
1.4	Relationship to the XSH Specification.....	6
1.4.1	Error Numbers.....	6
1.5	Thread Safety.....	6
1.6	Thread Cancellation Points.....	6
1.7	Relationship to Emerging Formal Standards.....	7
Chapter 2	General Introduction to the XTI.....	9
Chapter 3	Explanatory Notes for XTI.....	11
3.1	Transport Endpoints.....	11
3.2	Transport Providers.....	11
3.3	Association of a UNIX Process to an Endpoint.....	12
3.4	Use of the Same Protocol Address.....	12
3.5	Modes of Service.....	13
3.6	Error Handling.....	13
3.7	Synchronous and Asynchronous Execution Modes.....	14
3.8	Effect of Signals.....	16
3.9	Event Management.....	16
Chapter 4	XTI Overview.....	19
4.1	Overview of Connection-oriented Mode.....	19
4.1.1	Initialisation/De-initialisation Phase.....	20
4.1.2	Overview of Connection Establishment.....	21
4.1.3	Overview of Data Transfer.....	22
4.1.4	Overview of Connection Release.....	24
4.2	Overview of Connectionless Mode.....	26
4.2.1	Initialisation/De-initialisation Phase.....	26
4.2.2	Overview of Data Transfer.....	26
4.3	XTI Features.....	28
4.3.1	XTI Functions versus Protocols.....	29
Chapter 5	States and Events in XTI.....	31
5.1	Transport Interfaces States.....	32
5.2	Outgoing Events.....	33
5.3	Incoming Events.....	34
5.4	Transport User Actions.....	35

5.5	State Tables.....	36
5.6	Events and TLOOK Error Indication.....	38
Chapter 6	The Use of Options in XTI	39
6.1	Generalities	39
6.2	The Format of Options.....	40
6.3	The Elements of Negotiation.....	41
6.3.1	Multiple Options and Options Levels.....	41
6.3.2	Illegal Options	41
6.3.3	Initiating an Option Negotiation.....	42
6.3.4	Responding to a Negotiation Proposal	43
6.3.5	Retrieving Information about Options.....	44
6.3.6	Privileged and Read-only Options.....	45
6.4	Option Management of a Transport Endpoint	46
6.5	Supplements	48
6.5.1	The Option Value T_UNSPEC	48
6.5.2	The info Argument	48
6.5.3	Summary	48
6.6	Portability Aspects.....	50
Chapter 7	XTI Library Functions and Parameters	51
7.1	How to Prepare XTI Applications.....	51
7.2	Key for Parameter Arrays	51
7.3	Return of TLOOK Error.....	52
7.4	Use of “struct netbuf”	53
	<i>t_accept()</i>	54
	<i>t_alloc()</i>	57
	<i>t_bind()</i>	59
	<i>t_close()</i>	62
	<i>t_connect()</i>	63
	<i>t_error()</i>	66
	<i>t_errno</i>	68
	<i>t_free()</i>	69
	<i>t_getinfo()</i>	71
	<i>t_getprotaddr()</i>	74
	<i>t_getstate()</i>	76
	<i>t_listen()</i>	77
	<i>t_look()</i>	79
	<i>t_open()</i>	81
	<i>t_optmgmt()</i>	84
	<i>t_rcv()</i>	91
	<i>t_rcvconnect()</i>	93
	<i>t_rcvdis()</i>	95
	<i>t_rcvrel()</i>	97
	<i>t_rcvreldata()</i>	98
	<i>t_rcvudata()</i>	100
	<i>t_rcvuderr()</i>	102
	<i>t_rcvv()</i>	104

		<i>t_rcvvudata()</i>	106
		<i>t_snd()</i>	108
		<i>t_snddis()</i>	111
		<i>t_sndrel()</i>	113
		<i>t_sndreldata()</i>	114
		<i>t_sndudata()</i>	116
		<i>t_sndv()</i>	118
		<i>t_sndvudata()</i>	121
		<i>t_strerror()</i>	124
		<i>t_sync()</i>	125
		<i>t_sysconf()</i>	127
		<i>t_unbind()</i>	128
Chapter	8	Sockets Interfaces	129
	8.1	Sockets Overview	129
		<i>accept()</i>	130
		<i>bind()</i>	132
		<i>close()</i>	134
		<i>connect()</i>	135
		<i>fcntl()</i>	138
		<i>fgetpos()</i>	139
		<i>fsetpos()</i>	140
		<i>ftell()</i>	141
		<i>getpeername()</i>	142
		<i>getsockname()</i>	143
		<i>getsockopt()</i>	144
		<i>listen()</i>	147
		<i>lseek()</i>	148
		<i>poll()</i>	149
		<i>read()</i>	150
		<i>recv()</i>	151
		<i>recvfrom()</i>	153
		<i>recvmsg()</i>	156
		<i>select()</i>	159
		<i>send()</i>	160
		<i>sendmsg()</i>	162
		<i>sendto()</i>	165
		<i>setsockopt()</i>	168
		<i>shutdown()</i>	171
		<i>socket()</i>	172
		<i>socketpair()</i>	174
		<i>write()</i>	176
Chapter	9	Sockets Headers	177
		<fcntl.h>	178
		<sys/socket.h>	179
		<sys/stat.h>	183
		<sys/uio.h>	184

		<sys/un.h>.....	185
Chapter 10	IP Address Resolution Interfaces.....		187
	<i>endhostent()</i>		188
	<i>endnetent()</i>		190
	<i>endprotoent()</i>		191
	<i>endservent()</i>		192
	<i>gethostbyaddr()</i>		194
	<i>gethostbyname()</i>		195
	<i>getnetbyaddr()</i>		196
	<i>getprotobyname()</i>		197
	<i>getservbyname()</i>		198
	<i>h_errno</i>		199
	<i>htonl()</i>		200
	<i>inet_addr()</i>		201
	<i>ntohl()</i>		203
	<i>sethostent()</i>		204
	<i>setnetent()</i>		205
	<i>setprotoent()</i>		206
	<i>setservent()</i>		207
Chapter 11	IP Address Resolution Headers		209
	<arpa/inet.h>.....		210
	<netdb.h>.....		211
	<netinet/in.h>.....		213
	<unistd.h>.....		214
Appendix A	ISO Transport Protocol Information.....		215
	A.1 General.....		215
	A.2 Options.....		216
	A.2.1 Connection-mode Service		216
	A.2.1.1 Options for Quality of Service and Expedited Data.....		216
	A.2.1.2 Management Options		218
	A.2.2 Connectionless-mode Service		220
	A.2.2.1 Options for Quality of Service		220
	A.2.2.2 Management Options		221
	A.3 Functions		222
	A.4 The <xti_osi.h> Header File.....		225
Appendix B	Internet Protocol-specific Information.....		229
	B.1 General.....		229
	B.2 Options.....		230
	B.2.1 TCP-level Options		230
	B.2.2 T_UDP-level Options.....		231
	B.2.3 T_IP-level Options.....		232
	B.3 Functions		235
	B.4 The <xti_inet.h> Header File.....		238

Appendix C	Guidelines for Use of XTI.....	241
C.1	Transport Service Interface Sequence of Functions.....	241
C.2	Example in Connection-oriented Mode.....	242
C.3	Example in Connectionless Mode.....	244
C.4	Writing Protocol-independent Software.....	245
C.5	Event Management.....	246
C.5.1	Short-term Solution.....	246
C.5.2	XTI Events.....	247
C.6	The Poll Function.....	248
C.6.1	Example of Use of Poll.....	248
C.7	The Select Function.....	257
C.7.1	Example of Use of Select.....	257
Appendix D	Use of XTI to Access NetBIOS.....	266
D.1	Introduction.....	266
D.2	Objectives.....	266
D.3	Scope.....	267
D.4	Issues.....	268
D.5	NetBIOS Names and Addresses.....	269
D.6	NetBIOS Connection Release.....	271
D.7	Options.....	272
D.8	XTI Functions.....	273
Appendix E	XTI and TLI.....	279
E.1	Restrictions Concerning the Use of XTI.....	279
E.2	Relationship between XTI and TLI.....	280
Appendix F	Headers and Definitions for XTI.....	281
Appendix G	Abbreviations.....	295
Appendix H	Minimum OSI Functionality (CAE Specification).....	297
H.1	General.....	297
H.1.1	Rationale for using XTI-mOSI.....	297
H.1.2	Migrant Applications.....	297
H.1.3	OSI Functionality.....	297
H.1.4	mOSI API versus XAP.....	298
H.1.5	Upper Layers Functionality Exposed via mOSI.....	298
H.1.5.1	Naming and Addressing Information used by mOSI.....	298
H.1.5.2	XTI Options Specific to mOSI.....	298
H.2	Options.....	300
H.2.1	ACSE/Presentation Connection-mode Service.....	300
H.2.2	ACSE/Presentation Connectionless-mode Service.....	301
H.2.3	Transport Service Options.....	302
H.3	Functions.....	303
H.4	Implementors) Notes.....	307
H.4.1	Upper Layers FUs, Versions and Protocol Mechanisms.....	307
H.4.2	Mandatory and Optional Parameters.....	307

H.4.3	Mapping XTI Functions to ACSE/Presentation Services.....	308
H.4.3.1	Connection-mode Services	308
H.4.3.2	Connectionless-mode Services	312
H.5	Option Data Types and Structures.....	313
H.6	<xti_mosi.h> Header File.....	317
Appendix I	SNA Transport Provider.....	319
I.1	Introduction	319
I.2	SNA Transport Protocol Information	320
I.2.1	General.....	320
I.2.2	SNA Addresses	321
I.2.3	Options.....	322
I.2.3.1	Connection-Mode Service Options.....	322
I.2.4	Functions	324
I.3	Mapping XTI to SNA Transport Provider	327
I.3.1	General Guidelines.....	328
I.3.2	Flows Illustrating Full Duplex Mapping	329
I.3.3	Full Duplex Mapping.....	338
I.3.3.1	Parameter Mappings.....	340
I.3.4	Half Duplex Mapping.....	350
I.3.5	Return Code to Event Mapping.....	351
Appendix J	The Internet Protocols.....	353
Appendix K	IPX/SPX Transport Provider	355
K.1	General.....	355
K.2	Namespace	355
K.2.1	IPX.....	355
K.2.2	SPX.....	356
K.3	Options.....	356
K.3.1	IPX-level Options.....	356
K.3.2	SPX-level Options.....	357
K.4	Functions	359
Appendix L	ATM Transport Protocol Information for XTI.....	363
L.1	General.....	363
L.2	ATM Addresses.....	364
L.2.1	ATM Network Address	364
L.2.2	ATM Protocol Address	364
L.2.3	t_atm_sap Structure	365
L.3	Options.....	369
L.3.1	Signalling-level Options.....	369
L.3.2	Absolute Requirements	370
L.3.3	Further Remarks.....	370
L.4	Existing Functions	381
L.5	New Functions	384
	t_addleaf()	385
	t_removeleaf()	387

	<i>t_rcvleafchange()</i>	389
L.6	Implementation Notes.....	391
Appendix M	ATM Transport Protocol Information for Sockets	395
M.1	General.....	395
M.2	Existing Functions	396
M.3	Point-to-Multipoint Connections	399
M.3.1	Adding a Leaf	399
M.3.2	Removing a Leaf	400
M.3.3	Receiving Indication of a Change in Leaf Status.....	400
M.4	Implementation Notes.....	402
Appendix N	ATM Transport Headers	405
N.1	Proposed Additions to <xti.h>.....	405
N.2	Proposed Additions to <sys/socket.h>.....	405
N.3	<xti_atm.h>	405
N.4	<netatm/atm.h>.....	406
N.5	<_atm_common.h>	407
	Glossary	415
	Index.....	419
 List of Figures		
I-1	Active Connection Establishment, Blocking Version (1 of 2)	329
I-2	Active Connection Establishment, Non-blocking Version (2 of 2).....	330
I-3	Passive Connection Establishment, Instantiation Version (1 of 3)	331
I-4	Passive Connection Establishment, Blocking Version (2 of 3).....	332
I-5	Passive Connection Establishment, Non-blocking Version (3 of 3)	333
I-6	XTI Function to LU 6.2 Verb Mapping: Blocking <i>t_snd</i>	334
I-7	XTI Function to LU 6.2 Verb Mapping: Non-blocking <i>t_snd</i>	335
I-8	XTI Function to LU 6.2 Verb Mapping: Blocking <i>t_rcv</i>	336
I-9	Mapping from XTI Calls to LU 6.2 Verbs (Passive side).....	337
 List of Tables		
3-1	Events and <i>t_look()</i>	15
4-1	Classification of the XTI Functions	29
5-1	Transport Interface States	32
5-2	Transport Interface Outgoing Events	33
5-3	Transport Interface Incoming Events	34
5-4	Transport Interface User Actions.....	35
5-5	Initialisation/De-initialisation States	36
5-6	Data Transfer States: Connectionless-mode.....	36
5-7	Connection/Release/Data Transfer States: Connection-mode	37
7-1	XTI-level Options.....	88
A-1	Options for Quality of Service and Expedited Data.....	216

A-2	Management Options	218
A-3	Options for Quality of Service	220
A-4	Management Option	221
B-1	TCP-level Options	230
B-2	T_UDP-level Option.....	231
B-3	T_IP-level Options.....	232
C-1	Sequence of Transport Functions in Connection-oriented Mode	243
C-2	Sequence of Transport Functions in Connectionless Mode	244
H-1	APCO-level Options	300
H-2	APCL-level Options	301
H-3	Association Establishment.....	309
H-4	Data Transfer.....	310
H-5	Association Release	311
H-6	Connectionless-mode ACSE Service	312
I-1	SNA Options.....	323
I-2	Fields for <i>info</i> Parameter.....	324
I-3	Default Characteristics returned by <i>t_open()</i>	325
I-4	FDX LU 6.2 Verb Definitions	328
I-5	XTI Mapping to LU 6.2 Full Duplex Verbs	338
I-6	Relation Symbol Description.....	340
I-7	<i>t_accept</i> <--> FDX Verbs and Parameters	340
I-8	<i>t_bind</i> <--> FDX Verbs and Parameters	341
I-9	<i>t_close</i> <--> FDX Verbs and Parameters.....	341
I-10	<i>t_connect</i> <--> FDX Verbs and Parameters.....	342
I-11	<i>t_getprocaddr</i> <--> FDX Verbs and Parameters.....	344
I-12	<i>t_listen</i> <--> FDX Verbs and Parameters.....	345
I-13	<i>t_optmgmt</i> <--> FDX Verbs and Parameters.....	345
I-14	<i>t_rcv</i> <--> FDX Verbs and Parameters	346
I-15	<i>t_rcvconnect</i> <--> FDX Verbs and Parameters.....	347
I-16	<i>t_snd</i> <--> FDX Verbs and Parameters.....	348
I-17	<i>t_snddis</i> (Existing Connection) <--> FDX Verbs and Parameters.....	349
I-18	<i>t_snddis</i> (Incoming Connect Req.) <--> FDX Verbs and Parameters...	349
I-19	<i>t_sndrel</i> <--> FDX Verbs and Parameters	349
I-20	Mapping of XTI Events to SNA Events.....	351
L-1	Signaling-level Options	369

Preface

The Open Group

The Open Group is an international open systems organisation that is leading the way in creating the infrastructure needed for the development of network-centric computing and the information superhighway. Formed in 1996 by the merger of the X/Open Company and the Open Software Foundation, The Open Group is supported by most of the world's largest user organisations, information systems vendors and software suppliers. By combining the strengths of open systems specifications and a proven branding scheme with collaborative technology development and advanced research, The Open Group is well positioned to assist user organisations, vendors and suppliers in the development and implementation of products supporting the adoption and proliferation of open systems.

With more than 300 member companies, The Open Group helps the IT industry to advance technologically while managing the change caused by innovation. It does this by:

- consolidating, prioritising and communicating customer requirements to vendors
- conducting research and development with industry, academia and government agencies to deliver innovation and economy through projects associated with its Research Institute
- managing cost-effective development efforts that accelerate consistent multi-vendor deployment of technology in response to customer requirements
- adopting, integrating and publishing industry standard specifications that provide an essential set of blueprints for building open information systems and integrating new technology as it becomes available
- licensing and promoting the X/Open brand that designates vendor products which conform to X/Open Product Standards
- promoting the benefits of open systems to customers, vendors and the public.

The Open Group operates in all phases of the open systems technology lifecycle including innovation, market adoption, product development and proliferation. Presently, it focuses on seven strategic areas: open systems application platform development, architecture, distributed systems management, interoperability, distributed computing environment, security, and the information superhighway. The Open Group is also responsible for the management of the UNIX trade mark on behalf of the industry.

The X/Open Process

This description is used to cover the whole Process developed and evolved by X/Open. It includes the identification of requirements for open systems, development of CAE and Preliminary Specifications through an industry consensus review and adoption procedure (in parallel with formal standards work), and the development of tests and conformance criteria.

This leads to the preparation of a Product Standard which is the name used for the documentation that records the conformance requirements (and other information) to which a vendor may register a product. There are currently two forms of Product Standard, namely the Profile Definition and the Component Definition, although these will eventually be merged into one.

The X/Open brand logo is used by vendors to demonstrate that their products conform to the relevant Product Standard. By use of the X/Open brand they guarantee, through the X/Open Trade Mark Licence Agreement (TMLA), to maintain their products in conformance with the Product Standard so that the product works, will continue to work, and that any problems will be fixed by the vendor.

Open Group Publications

The Open Group publishes a wide range of technical literature, the main part of which is focused on specification development and product documentation, but which also includes Guides, Snapshots, Technical Studies, Branding and Testing documentation, industry surveys and business titles.

There are several types of specification:

- *CAE Specifications*

CAE (Common Applications Environment) Specifications are the stable specifications that form the basis for our product standards, which are used to develop X/Open branded systems. These specifications are intended to be used widely within the industry for product development and procurement purposes.

Anyone developing products that implement a CAE Specification can enjoy the benefits of a single, widely supported industry standard. In addition, they can demonstrate product compliance through the X/Open brand. CAE Specifications are published as soon as they are developed, so enabling vendors to proceed with development of conformant products without delay.

- *Preliminary Specifications*

Preliminary Specifications usually address an emerging area of technology and consequently are not yet supported by multiple sources of stable conformant implementations. They are published for the purpose of validation through implementation of products. A Preliminary Specification is not a draft specification; rather, it is as stable as can be achieved, through applying The Open Group's rigorous development and review procedures.

Preliminary Specifications are analogous to the *trial-use* standards issued by formal standards organisations, and developers are encouraged to develop products on the basis of them. However, experience through implementation work may result in significant (possibly upwardly incompatible) changes before its progression to becoming a CAE Specification. While the intent is to progress Preliminary Specifications to corresponding CAE Specifications, the ability to do so depends on consensus among Open Group members.

- *Consortium and Technology Specifications*

The Open Group publishes specifications on behalf of industry consortia. For example, it publishes the NMF SPIRIT procurement specifications on behalf of the Network Management Forum. It also publishes Technology Specifications relating to OSF/1, DCE, OSF/Motif and CDE.

Technology Specifications (formerly AES Specifications) are often candidates for consensus review, and may be adopted as CAE Specifications, in which case the relevant Technology Specification is superseded by a CAE Specification.

In addition, The Open Group publishes:

- *Product Documentation*

This includes product documentation — programmer's guides, user manuals, and so on — relating to the Pre-structured Technology Projects (PSTs), such as DCE and CDE. It also includes the Single UNIX Documentation, designed for use as common product documentation for the whole industry.

- *Guides*

These provide information that is useful in the evaluation, procurement, development or management of open systems, particularly those that relate to the CAE Specifications. The Open Group Guides are advisory, not normative, and should not be referenced for purposes of specifying or claiming conformance to a Product Standard.

- *Technical Studies*

Technical Studies present results of analyses performed on subjects of interest in areas relevant to The Open Group's Technical Programme. They are intended to communicate the findings to the outside world so as to stimulate discussion and activity in other bodies and the industry in general.

- *Snapshots*

These provide a mechanism to disseminate information on its current direction and thinking, in advance of possible development of a Specification, Guide or Technical Study. The intention is to stimulate industry debate and prototyping, and solicit feedback. A Snapshot represents the interim results of a technical activity.

Versions and Issues of Specifications

As with all *live* documents, CAE Specifications require revision to align with new developments and associated international standards. To distinguish between revised specifications which are fully backwards compatible and those which are not:

- A new *Version* indicates there is no change to the definitive information contained in the previous publication of that title, but additions/extensions are included. As such, it *replaces* the previous publication.
- A new *Issue* indicates there is substantive change to the definitive information contained in the previous publication of that title, and there may also be additions/extensions. As such, both previous and new documents are maintained as current publications.

Corrigenda

Readers should note that Corrigenda may apply to any publication. Corrigenda information is published on the World-Wide Web at <http://www.opengroup.org/public/pubs>.

Ordering Information

Full catalogue and ordering information on all Open Group publications is available on the World-Wide Web at <http://www.opengroup.org/public/pubs>.

This Document

This document is a CAE Specification, although it includes some appendices which are classed as Preliminary Specifications. The introductory section in each appendix clearly identifies its status.

The Networking Services (XNS) specification describes the industry-standard Open Systems interfaces to communications services. These include two APIs to transport-level process-to-process communications: Sockets and XTI. Both Sockets and XTI are specified for use over Internet protocols (TCP, UDP and IP) and ISO Transport protocols. They also include a set of Internet address resolution interfaces which are commonly used in conjunction with Sockets. XTI support for many other protocols is described in appendices to XNS.

This XNS Issue 5 is one of the UNIX98 specifications. Branded UNIX98 systems support the Sockets, XTI and Address Resolution interfaces described in XNS Issue 5. Sockets and Address Resolution must be supported over the Internet protocols. XTI may be supported over either the Internet or ISO Transport protocols. Other protocols may also be provided, but this is not required for the Brand.

XNS Issue 5 contains a number of new features over the previous publication¹. The most important new feature is the removal of implicit data length assumptions (for example, that an item of type *int* is 32 bits long). This enables the XNS APIs to be used effectively on new computer architectures with 64-bit (or larger) word lengths. Other new features include:

- taking account of realtime threads
- support for scatter/gather
- XTI can now be used in a way that takes less of the application's name space
- use of XTI over IPX/SPX is now specified
- use of Sockets and XTI with the broadband Asynchronous Transfer Mode (ATM) protocol is now specified.

XNS Issue 5 includes revisions to align with the IEEE Standard 1003.1: 1990 (Portable Operating System Interface). At the time of publication, the IEEE project P1003.1g is considering draft extensions, and their current draft includes some aspects with which XNS can not align (for example, it still makes implicit data length assumptions). Attempts are being made in P1003.1g to seek alignment with XNS Issue 5.

Structure

- Chapter 1 contains information comparable to that in the **XSH** specification. It applies to the Sockets and Address Resolution interfaces (see below) if the UNIX compilation environment is in effect.
- Chapter 2 is a general introduction to the Transport Interface (XTI).
- Chapter 3 provides explanatory notes.
- Chapter 4 gives an overview for XTI.
- Chapter 5 describes the states and events in XTI.
- Chapter 6 describes the use of options.

1. The previous XNS publication was CAE Specification, September 1994, Networking Services, Issue 4, (ISBN: 1-85912-049-0, C438).

- Chapter 7 contains reference manual pages for the XTI functions and parameters.
- Chapter 8 gives an overview and interfaces for Sockets.
- Chapter 9 defines the headers for Sockets.
- Chapter 10 gives an overview and interfaces for IP Address Resolution.
- Chapter 11 defines the headers for IP Address Resolution.
- Appendix A describes the protocol-specific information that is relevant for ISO transport providers, including for ISO-over-TCP (RFC 1006).
- Appendix B describes the protocol-specific information that is relevant for TCP and UDP transport providers.
- Appendix C gives guidelines for the use of XTI.
- Appendix D specifies a standard programming interface to NetBIOS transport providers in systems compliant to XTI.
- Appendix E describes how XTI provides refinement of the Transport Level Interface (TLI).
- Appendix F presents a summary of the contents of the `<xti.h>` header file.
- Appendix G lists abbreviations used in this document.
- Appendix H provides a simple API exposing a minimum set of OSI Upper Layers functionality (mOSI).
- Appendix I describes the protocol-specific information and mapping to XTI functions that is relevant for Systems Network Architecture (SNA) transport providers.
- Appendix J contains a brief explanation of the Internet Protocols.
- Appendix K describes the protocol-specific information and mapping to XTI functions that is relevant for an IPX/SPX transport provider.
- Appendix L describes the protocol-specific information for XTI that is relevant for ATM transport providers.
- Appendix M describes the protocol-specific information for Sockets that is relevant for ATM transport providers.
- Appendix N presents header files information required for ATM transport providers.

Typographical Conventions

The following typographical conventions are used throughout this document:

- **Bold** font is used in text for options to commands, filenames, keywords, type names, data structures and their members.
- *Italic* strings are used for emphasis or to identify the first instance of a word requiring definition. Italics in text also denote:
 - command operands, command option-arguments or variable names, for example, substitutable argument prototypes
 - environment variables, which are also shown in capitals
 - utility names
 - external variables, such as *errno*

- functions; these are shown as follows: *name()*. Names without parentheses are C external variables, C function family names, utility names, command operands or command option-arguments.
- Normal font is used for the names of constants and literals.
- The notation **<file.h>** indicates a header file.
- Names surrounded by braces, for example, {ARG_MAX}, represent symbolic limits or configuration values which may be declared in appropriate headers by means of the C **#define** construct.
- The notation [ABCD] is used to identify a return value ABCD, including if this is an error value.
- Syntax, code examples and user input in interactive examples are shown in *fixed width* font. Brackets shown in this font, [], are part of the syntax and do *not* indicate optional items. In syntax the | symbol is used to separate alternatives, and ellipses (. . .) are used to show that additional arguments are optional.

Trade Marks

AT&T[®] is a registered trade mark of AT&T in the U.S.A. and other countries.

HP[®] is a registered trade mark of Hewlett-Packard.

Motif[®], OSF/1[®] and UNIX[®] are registered trade marks and the “X Device”[™] and The Open Group[™] are trade marks of The Open Group.

SNA is a product of International Business Machines Corporation.

/usr/group[®] is a registered trade mark of UniForum, the International Network of UNIX System Users.

Acknowledgements

- AT&T for permission to reproduce portions of its copyrighted System V Interface Definition (SVID) and material from the UNIX System V Release 2.0 documentation.
- The Institution of Electrical and Electronics Engineers, Inc. for permission to reproduce portions of its copyrighted material.
- The IEEE Computer Society's Portable Applications Standards Committee (PASC), whose Standards contributed to our work.
- The UniForum (formerly /usr/group) Technical Committee's Internationalization Subcommittee for work on internationalised regular expressions.
- The ANSI X3J11 Committees.
- The Open Group gratefully acknowledges the valued contribution of the following people in the development of this specification, along with their corporate affiliation at the time of their contribution:

Josee Auber	Hewlett-Packard	Lori Mickelson	Unisys
Kathryn Britton	IBM	Laura Micks	IBM
Philippe Camus	Groupe Bull	Finnbarr Murphy	DIGITAL
Andrew Chandler	ICL	Alagu Periyannan	Apple
Andre Cohen	Groupe Bull	George Preoteasa	Hewlett-Packard
Paul Comstock	Hewlett-Packard	John Ronciak	Unisys
Andrew Gollan	Sun	Seth Rosenthal	Novell
Michel Habert	Groupe Bull	Eric Scoredos	Hewlett-Packard
Torez Hiley	USL	Maria Stanley	DIGITAL
Martin Jess	NCR	Lutz Temme	SNI
Mukesh Kacker	Sun	Roger Turner	IBM
Gerhard Kieselmann	SNI	Keith Weir	ICL
David Laight	Fujitsu/ICL	Robert Weirick	Unisys
Jack McCann	DIGITAL	Greg Wiley	SFC (SCO)
Hiroshi Maruta	Hitachi	Isaac Wong	Hewlett-Packard

Referenced Documents

The following documents are referenced in this specification:

ACSE

ISO 8649

ISO 8649: 1988, Information Processing Systems — Open Systems Interconnection — Service Definition for the Association Control Service Element, together with:

Technical Corrigendum 1: 1990 to ISO 8649: 1988

Amendment 1: 1990 to ISO 8649: 1988

Authentication during association establishment.

Amendment 2: 1991 to ISO 8649: 1988

Connectionless-mode ACSE Service.

ISO 8650

ISO 8650: 1988, Information Processing Systems — Open Systems Interconnection — Protocol specification for the Association Control Service Element, together with:

Technical Corrigendum 1: 1990 to ISO 8650: 1988

Amendment 1: 1990 to ISO 8650: 1988

Authentication during association establishment.

ISO/IEC 10035

ISO/IEC 10035: 1991, Information Technology — Open Systems Interconnection — Connectionless ACSE Protocol Specification.

Presentation

ISO 8822

ISO 8822: 1988, Information Processing Systems — Open Systems Interconnection — Connection-oriented Presentation Service Definition.

ISO 8823

ISO 8823: 1988, Information Processing Systems — Open Systems Interconnection — Connection-oriented Presentation Protocol Specification.

ISO 8824

ISO 8824: 1990, Information Technology — Open Systems Interconnection — Specification of Abstract Syntax Notation One (ASN.1).

BER

ISO/IEC 8825: 1990 (ITU-T Recommendation X.209 (1988)), Information Technology — Open Systems Interconnection — Specification of Basic Encoding Rules for Abstract Syntax Notation One (ASN.1).

ISO/IEC 9576

ISO/IEC 9576: 1991, Information Technology — Open Systems Interconnection — Connectionless Presentation Protocol Specification.

Session

ISO 8326

ISO 8326: 1987, Information Processing Systems — Open Systems Interconnection — Basic Connection-oriented Session Service Definition.

ISO 8327

ISO 8327: 1987, Information Processing Systems — Open Systems Interconnection — Basic Connection-oriented Session Protocol Specification.

Amendment 3: 1992 to ISO 8327: 1987 — Additional Synchronization Functionality.

ATM

ATMNAS ATM Forum: “Native ATM services: Semantic Description, Version 1”, obtainable via anonymous ftp from Internet address ftp.atmforum.com, in directory /pub/approved-specs, files af-saa-0048.000.doc (Word 6.0) or af-saa-0048.000.ps (postscript).

UNIATM Forum: “ATM User-Network Interface (UNI) Specification, Version 3.1”, published by Prentice Hall. Also obtainable electronically via anonymous ftp from Internet address ftp.atmforum.com, in directory /pub/UNI/ver3.1.

Other References

Minimal OSI

ISO/IEC DISP 11188-3, International Standardized Profile — Common Upper Layer Requirements — Part 3: Minimal OSI Upper Layers Facilities, Version 6, 1994-04-14.

ISO 7498

ISO 7498: 1984, Information Processing Systems — Open Systems Interconnection — Basic Reference Model.

ISO Transport

	Connection-Oriented	Connectionless
Protocol Definition	IS 8073-1986	IS 8602
Service Definition	IS 8072-1986	IS 8072/Add.1-1986

TCP

Transmission Control Protocol, RFC 793 (Defense Communication Agency, DDN Protocol Handbook, Volume II, DARPA Internet Protocols, (December 1985).

Also see TCP, Transmission Control Protocol, Military Standard, Mil-std-1778, Defense Communication Agency, DDN Protocol Handbook, Volume I, DOD Military Standard Protocols (December 1985).

UDP

User Datagram Protocol, RFC 768 (Defense Communication Agency, DDN Protocol Handbook, Volume II, DARPA Internet Protocols, December 1985).

TLI Specifications

Networking Services Extension, draft version of SVID Issue 2, Volume III, 1986.

NetBIOS

Mappings of NetBIOS services to OSI and IPS transport protocols are provided in the CAE Specification, October 1992, Protocols for PC Interworking: SMB, Version 2 (ISBN: 1-872630-45-6, C209).

Referenced Documents

SNA

SNA National Registry, IBM document G325-6025-0.

CURL

Common Upper Layer Requirements, Part 3: Minimal OSI Upper Layer Facilities — OIW/EWOS working document.

XSH, Issue 5

CAE Specification, January 1997, System Interfaces and Headers, Issue 5 (ISBN: 1-85912-181-0, C606).

XCU, Issue 5

CAE Specification, January 1997, Commands and Utilities, Issue 5 (ISBN: 1-85912-191-8, C604).

XBD, Issue 5

CAE Specification, January 1997, System Interface Definitions, Issue 5 (ISBN: 1-85912-186-1, C605).

Common Information

This chapter provides general information that applies to the XTI, Sockets and IP Address Resolution interfaces defined in this volume.

1.1 Terminology

The information in this section applies only to the Sockets and IP Address Resolution interfaces.

The following terms are used in this specification:

can

This describes a permissible optional feature or behaviour available to the user or application; all systems support such features or behaviour as mandatory requirements.

implementation-dependent

The value or behaviour is not consistent across all implementations. The provider of an implementation normally documents the requirements for correct program construction and correct data in the use of that value or behaviour. When the value or behaviour in the implementation is designed to be variable or customisable on each instantiation of the system, the provider of the implementation normally documents the nature and permissible ranges of this variation. Applications that are intended to be portable must not rely on implementation-dependent values or behaviour.

legacy

Certain features are *legacy*, which means that they are being retained for compatibility with older applications, but have limitations which make them inappropriate for developing portable applications. New applications should use alternative means of obtaining equivalent functionality. Legacy features are marked **LEGACY**.

may

With respect to implementations, the feature or behaviour is optional. Applications should not rely on the existence of the feature. To avoid ambiguity, the reverse sense of *may* is expressed as *need not*, instead of *may not*.

must

This describes a requirement on the application or user.

should

With respect to implementations, the feature is recommended, but it is not mandatory. Applications should not rely on the existence of the feature.

With respect to users or applications, the word means recommended programming practice that is necessary for maximum portability.

undefined

A value or behaviour is undefined if this document imposes no portability requirements on applications for erroneous program constructs or erroneous data. Implementations may specify the result of using that value or causing that behaviour, but such specifications are not guaranteed to be consistent across all implementations. An application using such behaviour is not fully portable to all systems.

unspecified

A value or behaviour is unspecified if this document imposes no portability requirements on

applications for correct program construct or correct data. Implementations may specify the result of using that value or causing that behaviour, but such specifications are not guaranteed to be consistent across all implementations. An application requiring a specific behaviour, rather than tolerating any behaviour when using that functionality, is not fully portable to all systems.

will

This means that the behaviour described is a requirement on the implementation and applications can rely on its existence.

1.2 Use and Implementation of Interfaces

Each of the following statements applies unless explicitly stated otherwise in the detailed descriptions that follow. If an argument to a function has an invalid value (such as a value outside the domain of the function, or a pointer outside the address space of the program, or a null pointer), the behaviour is undefined.

For backward compatibility purposes, any of the function names within the header file may be redefined by the implementation, using the pre-processor “#define” mechanism (or any other similar compiler supported mechanism) to another function name. If redefined, the function name so defined will be in accordance with the name space rules in Section 1.3.1 on page 3.

As a result of changes in this issue of this document, application writers are only required to include the minimum number of headers. Implementations of XSI-conformant systems will make all necessary symbols visible as described in the Headers section of this document.

1.2.1 C Language Definition

The C language that is the basis for the synopses and code examples in this document is *ISO C*, as specified in the referenced ISO C standard. *Common Usage C*, which refers to the C language before standardisation, was the basis for previous editions of the **XTI** specification.

1.3 The Compilation Environment

Applications should ensure that the feature test macro `_XOPEN_SOURCE` is defined with the value 500 before inclusion of any header. This is needed to enable the functionality described in this document, and possibly to enable functionality defined elsewhere in the Common Applications Environment.

The `_XOPEN_SOURCE` macro may be defined automatically by the compilation process, but to ensure maximum portability, applications should make sure that `_XOPEN_SOURCE` is defined by using either compiler options or `#define` directives in the source files, before any `#include` directives. Identifiers in this document may be undefined using the `#undef` directive as described in Section 1.3.1. These `#undef` directives must follow all `#include` directives of any XSI headers.

Since this specification is aligned with the ISO C standard, and since all functionality enabled by `_POSIX_C_SOURCE` set greater than zero and less than or equal to 199506L should be enabled by `_XOPEN_SOURCE` set equal to 500, there should be no need to define either `_POSIX_SOURCE` or `_POSIX_C_SOURCE` if `_XOPEN_SOURCE` is defined. Therefore if `_XOPEN_SOURCE` is set equal to 500 and `_POSIX_SOURCE` is defined, or `_POSIX_C_SOURCE` is set greater than zero and less than or equal to 199506L, the behavior is the same as if only `_XOPEN_SOURCE` is defined and set equal to 500. However, should `_POSIX_C_SOURCE` be set to a value greater than 199506L, the behaviour is undefined.

The `c89` and `cc` utilities recognise the additional `-I` operand for standard libraries:

`-I xnet` If the implementation defines `_XOPEN_UNIX`, this operand makes visible all functions referenced in this document. An implementation may search this library in the absence of this operand.

It is unspecified whether the library `libxnet.a` exists as a regular file.

If the implementation supports the utilities marked **DEVELOPMENT** in the XCU specification, the `lint` utility recognises the additional `-I` operand for standard libraries:

`-I xnet` Names the library `llib-lxnet.ln`, which will contain functions specified in this document.

It is unspecified whether the library `llib-lxnet.ln` exists as a regular file.

1.3.1 The Name Space

All identifiers in this document are defined in at least one of the headers, as shown in Chapter 9, Chapter 11 of XNS and Chapter 4 of XSH (see the referenced document **XSH**). When `_XOPEN_SOURCE` is defined, each header defines or declares some identifiers, potentially conflicting with identifiers used by the application. The set of identifiers visible to the application consists of precisely those identifiers from the header pages of the included headers, as well as additional identifiers reserved for the implementation. In addition, some headers may make visible identifiers from other headers as indicated on the relevant header pages.

The identifiers reserved for use by the implementation are described below.

1. Each identifier with external linkage described in the header section is reserved for use as an identifier with external linkage if the header is included.
2. Each macro name described in the header section is reserved for any use if the header is included.
3. Each identifier with file scope described in the header section is reserved for use as an identifier with file scope in the same name space if the header is included.

If any header in the following table is included, identifiers with the following prefixes or suffixes shown are reserved for any use by the implementation.

Header	Prefix	Suffix
<arpa/inet.h>	in_, inet_	
<netdb.h>	h_, n_, p_, s_	
<netinet/in.h>	in_, ip_, s_, sin_	
<sys/socket.h>	sa_, if_, ifc_, ifru_, infu_, ifra_, msg_, cmsg_, l_	
<sys/un.h>	sun_	
<xti.h>	l_, t_, T_	
ANY header		_t

If any header in the following table is included, macros with the prefixes shown may be defined. After the last inclusion of a given header, an application may use identifiers with the corresponding prefixes for its own purpose, provided their use is preceded by an **#undef** of the corresponding macro.

Header	Prefix
<netinet/in.h>	IMPLINK_, IN_, INADDR_, IP_, IPPORT_, IPPROTO_, SOCK_
<sys/socket.h>	AF_, CMSG_, MSG_, PF_, SCM_, SHUT_, SO
<xti.h>	OPT_, T_ ² , XTI_

The following identifiers are reserved regardless of the inclusion of headers:

1. All identifiers that begin with an underscore and either an upper-case letter or another underscore are always reserved for any use by the implementation.
2. All identifiers that begin with an underscore are always reserved for use as identifiers with file scope in both the ordinary identifier and tag name spaces.

2. The following T_ prefixes are currently used: T_INET_, T_IP_, T_ISO_, T_, T_TCL_, T_TCP_, T_TCO_, T_UDP_.

3. All identifiers in the table below are reserved for use as identifiers with external linkage.

XTI:				
t_accept	t_getinfo	t_rcv	t_rcvv	t_sndudata
t_alloc	t_getprotaddr	t_rcvconnect	t_rcvvudata	t_sndv
t_bind	t_getstate	t_rcvdis	t_snd	t_sndvudata
t_close	t_listen	t_rcvrel	t_snd	t_strerror
t_connect	t_look	t_rcvreldata	t_snddis	t_sync
t_errno	t_open	t_rcvudata	t_sndrel	t_sysconf
t_error	t_optmgmt	t_rcvuderr	t_sndreldata	t_unbind
t_free				
Sockets:				
accept	getsockname	recvfrom	sendmsg	shutdown
bind	getsockopt	recvmsg	sendto	socket
connect	listen	send	setsockopt	socketpair
getpeername	recv			
IP Address Resolution:				
endhostent	gethostname	getservbyname	inet_addr	ntohl
endnetent	getnetbyaddr	getservbyport	inet_lnaof	ntohs
endprotoent	getnetbyname	getservent	inet_makeaddr	sethostent
endservent	getnetent	h_errno	inet_netof	setnetent
gethostbyaddr	getprotobyname	htonl	inet_network	setprotoent
gethostbyname	getprotobyname	htons	inet_ntoa	setservent
gethostent	getprotoent			

All the identifiers defined in this document that have external linkage are always reserved for use as identifiers with external linkage.

No other identifiers are reserved.

Applications must not declare or define identifiers with the same name as an identifier reserved in the same context. Since macro names are replaced whenever found, independent of scope and name space, macro names matching any of the reserved identifier names must not be defined if any associated header is included.

Headers may be included in any order, and each may be included more than once in a given scope, with no difference in effect from that of being included only once.

If used, a header must be included outside of any external declaration or definition, and it must be first included before the first reference to any type or macro it defines, or to any function or object it declares. However, if an identifier is declared or defined in more than one header, the second and subsequent associated headers may be included after the initial reference to the identifier. Prior to the inclusion of a header, the program must not define any macros with names lexically identical to symbols defined by that header.

1.4 Relationship to the XSH Specification

1.4.1 Error Numbers

Some functions provide an error number in *errno*, which is either a variable or macro defined in `<errno.h>`; the macro expands to a modifiable lvalue of type `int`.

A list of valid values for *errno* and advice to application writers on the use of *errno* appears in the XSH specification.

1.5 Thread Safety

All interfaces defined by this document will be thread-safe, except for the following interfaces which need not be thread-safe:

- byaddr()*
- gethostbyname()*
- gethostent()*
- getnetbyaddr()*
- getnetbyname()*
- getnetent()*
- getprotobynumber()*
- getprotobyname()*
- getprotoent()*
- getservbyname()*
- getservbyport()*
- getservent()*
- inet_ntoa()*

1.6 Thread Cancellation Points

Cancellation points will occur when a thread is executing any of the following functions:

- accept()*
- connect()*
- recv()*
- recvfrom()*
- recvmsg()*
- send()*
- sendmsg()*
- sendto()*

- t_close()*
- t_connect()*
- t_listen()*
- t_rcv()*
- t_rcvconnect()*
- t_rcvrel()*
- t_rcvreldata()*
- t_rcvudata()*
- t_rcvv()*
- t_rcvvudata()*
- t_snd()*

t_sndrel()
t_sndreldata()
t_sndudata()
t_sndv()
t_sndvudata()

A cancellation point may also occur when a thread is executing any of the following functions:

endhostent()
endnetent()
endprotoent()
endservent()
gethostbyaddr()
gethostbyname()
gethostent()
gethostname()
getnetbyaddr()
getnetbyname()
getnetent()
getprotobynumber()
getprotobyname()
getprotoent()
getservbyport()
getservbyname()
getservent()
sethostent()
setnetent()
setprotoent()
setservent()

An implementation will not introduce cancellation points into any other function specified in this document.

See the referenced **XSH**, Section 2.8 for further information.

1.7 Relationship to Emerging Formal Standards

The IEEE 1003.1g standards committee is also developing interfaces to XTI and Sockets. X/Open is actively involved in the work of this committee.

General Introduction to the XTI

The X/Open Transport Interface (XTI) specification defines an independent transport-service interface that allows multiple users to communicate at the transport level of the OSI reference model. The specification describes transport-layer characteristics that are supported by a wide variety of transport-layer protocols. Supported characteristics include:

- connection establishment
- state change support
- event handling
- data transfer
- option manipulation.

Although all transport-layer protocols support these characteristics, they vary in their level of support and/or their interpretation and format. For example, there are transport-level options which remain constant across all transport providers while there are other options which are transport-provider specific or have different values/names for different transport providers.

The main Chapters in this specification describe interfaces, parameters and semantics constant across all transport providers. The remainder of the document consists of appendices that provide valuable information that is not an integral part of the main body since it is either descriptive or applies only to some transport providers.

Some appendices provide information pertinent to writing XTI applications over specific transport providers. The transport providers fall into three classes:

- Those corresponding to traditional transport providers, such as:
 - ISO Transport (connection-mode or connectionless-mode)
 - TCP
 - UDP
 - NetBIOS.
- Those corresponding to commonly used subsets of higher-layer protocols that provide transport-like services, such as:
 - minimal functionality OSI (mOSI), that is, OSI ACSE/Presentation with the kernel and duplex functional units
 - SNA LU6.2 subset.
- Mixed-protocol providers that provide the appearance of one protocol over a different protocol such as:
 - ISO transport appearance (connection-mode) over TCP.

The ISO appendix (Appendix A) also describes a transport provider that uses RFC 1006 to compensate for the differences between ISO transport and TCP so that a TCP provider can present an ISO transport appearance.

While XTI gives transport users considerable independence from the underlying transport provider, the differences between providers are not entirely hidden. Appendix C includes guidelines for writing transport-provider-independent software, which can be done primarily by using only functions supported by all providers, avoiding option management, and using a provider-independent means of acquiring addresses.

While the transport-provider-specific appendices are intended mostly for transport users, they are also used by implementors of transport providers. For the purposes of the implementors, some of the appendices show how XTI services can be mapped to primitives associated with the specific providers. These are provided as guidance only and do not dictate anything about a given implementation.

Some of the appendices to the XTI specification are included as vehicles for communicating information needed by implementors, or guidelines to the use of the specification in question. The Guidelines for the use of XTI (Appendix C), Minimum OSI Functionality (Appendix H), SNA Transport provider (Appendix I) and comparison of XTI to TLI (Appendix E) belong to this category.

Some other appendices, however, have evolved into a prescriptive specification, as in the case of Appendix A for the ISO transport provider, Appendix B for the Internet transport provider and Appendix D for the NetBIOS transport provider. Since not every XTI implementor would find it relevant to implement the functionality of all of these appendices, they have been kept separate from the main XTI specification, thus becoming brandable XTI options. Support for these transport providers is declared by vendors through the XTI Conformance Statement Questionnaire.

An appendix may have a different status from the overall XTI specification. Thus the appendix for a particular transport provider may be a Preliminary Specification while the document is a CAE specification. When this is the case, the status of the appendix is clearly identified in its own introduction.

Topics beyond the scope of the XTI specification include:

- Address parameters

Several functions have parameters for addresses. The structure of these addresses is beyond the scope of this document. Specific implementations specify means for transport users to get or construct addresses.

- Event management

In order for applications to use XTI in a fully asynchronous manner, it will be necessary for the application to include facilities of an Event Management (EM) interface. Such EM facility may allow the application to be notified of a number of events over a range of active transport connections. For example, one event may denote a connection is flow-controlled. While Appendix C provides some guidelines for using EM in XTI applications, a complete specification defining an EM interface is beyond the scope of this document.

Explanatory Notes for XTI

3.1 Transport Endpoints

A *transport endpoint* specifies a communication path between a transport user and a specific transport provider, which is identified by a local file descriptor (*fd*). When a user opens a transport provider identifier, a local file descriptor *fd* is returned which identifies the transport endpoint. A transport provider is defined to be the transport protocol that provides the services of the transport layer. All requests to the transport provider must pass through a transport endpoint. The file descriptor *fd* is returned by the function *t_open()* and is used as an argument to the subsequent functions to identify the transport endpoint. A transport endpoint (*fd* and local address) can support only one established transport connection at a time.

To be active, a transport endpoint must have a transport address associated with it by the *t_bind()* function. A transport connection is characterised by the association of two active endpoints, made by using the functions of establishment of transport connection. The *fd* is a communication path to a transport provider. There is no direct assignation of the processes to the transport provider, so multiple processes, which obtain the *fd* by *open()*, *fork()* or *dup()* operations, may access a given communication path. Note that the *open()* function will work only if the opened character string is a pathname.

Note that in order to guarantee portability, the only operations which the applications may perform on any *fd* returned by *t_open()* are those defined by XTI and *fcntl()*, *dup()* or *dup2()*. Other operations are permitted but these will have system-dependent results.

3.2 Transport Providers

The transport layer may comprise one or more *transport providers* at the same time. The identifier parameter of the transport provider passed to the *t_open()* function determines the required transport provider. To keep the applications portable, the identifier parameter of the transport provider should not be hard-coded into the application source code.

An application which wants to manage multiple transport providers must call *t_open()* for each provider. For example, a server application which is waiting for incoming connection indications from several transport providers must open a transport endpoint for each provider and listen for connection indications on each of the associated file descriptors.

3.3 Association of a UNIX Process to an Endpoint

One process can simultaneously open several *fd*s. However, in synchronous mode, the process must manage the different actions of the associated transport connections sequentially. Conversely, several processes can share the same *fd* (by *fork()* or *dup()* operations) but they have to synchronise themselves so as not to issue a function that is unsuitable to the current state of the transport endpoint.

It is important to remember that the transport provider treats all users of a transport endpoint as a single user. If multiple processes are using the same endpoint, they should coordinate their activities so as not to violate the state of the provider. The *t_sync()* function returns the current state of the provider to the user, thereby enabling the user to verify the state before taking further action. This coordination is only valid among cooperating processes; it is possible that a process or an incoming event could change the provider's state after a *t_sync()* is issued.

A process can listen for an incoming connection indication on one *fd* and accept the connection on a different *fd* which has been bound with the *qlen* parameter (see *t_bind()*) set to zero. This facilitates the writing of a listener application whereby the listener waits for all incoming connection indications on a given Transport Service Access Point (TSAP). The listener will accept the connection on a new *fd*, and *fork()* a child process to service the request without blocking other incoming connection indications.

3.4 Use of the Same Protocol Address

If several endpoints are bound to the same protocol address, only one at the time may be listening for incoming connections. However, others may be in data transfer state or establishing a transport connection as initiators.

3.5 Modes of Service

The transport service interface supports two modes of service: connection-mode and connectionless-mode. A single transport endpoint may not support both modes of service simultaneously.

The connection-mode transport service is circuit-oriented and enables data to be transferred over an established connection in a reliable, sequenced manner. This service enables the negotiation of the parameters and options that govern the transfer of data. It provides an identification mechanism that avoids the overhead of address transmission and resolution during the data transfer phase. It also provides a context in which successive units of data, transferred between peer users, are logically related. This service is attractive to applications that require relatively long-lived, datastream-oriented interactions.

In contrast, the connectionless-mode transport service is message-oriented and supports data transfer in self-contained units with no logical relationship required among multiple units. These units are also known as datagrams. This service requires a pre-existing association between the peer users involved, which determines the characteristics of the data to be transmitted. No dynamic negotiation of parameters and options is supported by this service. All the information required to deliver a unit of data (for example, destination address) is presented to the transport provider, together with the data to be transmitted, in a single service access which need not relate to any other service access. Also, each unit of data transmitted is entirely self-contained, and can be independently routed by the transport provider. This service is attractive to applications that involve short-term request/response interactions, exhibit a high level of redundancy, are dynamically reconfigurable or do not require guaranteed, in-sequence delivery of data.

3.6 Error Handling

Two levels of error are defined for the transport interface. The first is the library error level. Each library function has one or more error returns. Failures are indicated by a return value of `-1`. When header file `<xti.h>` is included, symbol `t_errno` is defined as a modifiable lvalue of type `int`, `t_errno`,³ and can be used to access the specific error number when such a failure occurs. Applications should not include `t_errno` in the left operand of assignment statements. This value is set when errors occur but is not cleared on successful library calls, so it should be tested only after an error has been indicated. A diagnostic function, `t_error()`, prints out information on the current transport error. The state of the transport provider may change if a transport error occurs.

The second level of error is the operating system service routine level. A special library level error number has been defined called [TSYSERR] which is generated by each library function when an operating system service routine fails or some general error occurs. When a function sets `t_errno` to [TSYSERR], the specific system error may be accessed through the external variable `errno`.

For example, a system error can be generated by the transport provider when a protocol error has occurred. If the error is severe, it may cause the file descriptor and transport endpoint to be

3. This may be implemented as a macro. In addition the name `_t_errno` is an XTI library-reserved-name for use within such a macro. A typical definition of `t_errno` for a multithreaded implementation is:

```
extern int *_t_errno(void);
#define t_errno (*_t_errno())
```

unusable. To continue in this case, all users of the *fd* must close it. Then the transport endpoint may be re-opened and initialised.

3.7 Synchronous and Asynchronous Execution Modes

The transport service interface is inherently asynchronous; various events may occur which are independent of the actions of a transport user. For example, a user may be sending data over a transport connection when an asynchronous disconnection indication arrives. The user must somehow be informed that the connection has been broken.

The transport service interface supports two execution modes for handling asynchronous events: synchronous mode and asynchronous mode. In the synchronous mode of operation, the transport primitives wait for specific events before returning control to the user. While waiting, the user cannot perform other tasks. For example, a function that attempts to receive data in synchronous mode will wait until data arrives before returning control to the user. Synchronous mode is the default mode of execution. It is useful for user processes that want to wait for events to occur, or for user processes that maintain only a single transport connection.

The asynchronous mode of operation, on the other hand, provides a mechanism for notifying a user of some event without forcing the user to wait for the event. The handling of networking events in an asynchronous manner is seen as a desirable capability of the transport interface. This would enable users to perform useful work while expecting a particular event. For example, a function that attempts to receive data in asynchronous mode will return control to the user immediately if no data is available. The user may then periodically poll for incoming data until it arrives. The asynchronous mode is intended for those applications that expect long delays between events and have other tasks that they can perform in the meantime or handle multiple connections concurrently.

The two execution modes are not provided through separate interfaces or different functions. Instead, functions that process incoming events have two modes of operation: synchronous and asynchronous. The desired mode is specified through the `O_NONBLOCK` flag, which may be set when the transport provider is initially opened, or before any specific function or group of functions is executed using the `fcntl()` operating system service routine. The effect of this flag is local to this process and is completely specified in the description of each function.

Nine (only eight if the orderly release is not supported) asynchronous events are defined in the transport service interface to cover both connection-mode and connectionless-mode service. They are represented as separate bits in a bit-mask using the following defined symbolic names:

- `T_LISTEN`
- `T_CONNECT`
- `T_DATA`
- `T_EXDATA`
- `T_DISCONNECT`
- `T_ORDREL`
- `T_UDERR`
- `T_GODATA`
- `T_GOEXDATA`.

These are described in Section 3.9 on page 16.

A process that issues functions in synchronous mode must still be able to recognise certain asynchronous events and act on them if necessary. This is handled through a special transport error [TLOOK] which is returned by a function when an asynchronous event occurs. The `t_look()` function is then invoked to identify the specific event that has occurred when this error is returned.

Another means to notify a process that an asynchronous event has occurred is polling. The polling capability enables processes to do useful work and periodically poll for one of the above asynchronous events. This facility is provided by setting `O_NONBLOCK` for the appropriate primitive(s).

Events and `t_look()`

All events that occur at a transport endpoint are stored by XTI. These events are retrievable one at a time via the `t_look()` function. If multiple events occur, it is implementation-dependent in what order `t_look()` will return the events. An event is outstanding on a transport endpoint until it is consumed. Every event has a corresponding consuming function which handles the event and consumes it. In addition, the abortive `T_DISCONNECT` consumes other pending events. Both `T_DATA` and `T_EXDATA` events are consumed when the corresponding consuming function has read all the corresponding data associated with that event. The intention of this is that `T_DATA` should always indicate that there is data to receive. Two events, `T_GODATA` and `T_GOEXDATA`, are also cleared as they are returned by `t_look()`. Table 3-1 summarises this.

Event	Cleared on <code>t_look()</code> ?	Consuming XTI functions
<code>T_LISTEN</code>	No	<code>t_listen()</code>
<code>T_CONNECT</code>	No	<code>t_{rcv}connect()</code> ⁵
<code>T_DATA</code>	No	<code>t_rcv{v}{udata}()</code>
<code>T_EXDATA</code>	No	<code>t_rcv{v}()</code>
<code>T_DISCONNECT</code>	No	<code>t_rcvdis()</code>
<code>T_UDERR</code>	No	<code>t_rcvuderr()</code>
<code>T_ORDREL</code>	No	<code>t_rcvrel{data}()</code>
<code>T_ORDRELDATA</code>	No	<code>t_rcvreldata()</code>
<code>T_GODATA</code>	Yes	<code>t_snd{v}{udata}()</code>
<code>T_GOEXDATA</code>	Yes	<code>t_snd{v}()</code>

Table 3-1 Events and `t_look()`

5. In the case of the `t_connect()` function the `T_CONNECT` event is both generated and consumed by the execution of the function and is therefore not visible to the application.

3.8 Effect of Signals

In both the synchronous and the asynchronous execution modes, XTI calls may be affected by signals. Unless specified otherwise in the description of each function, the functions behave as described below.

If a synchronous XTI call is blocking under circumstances where an asynchronous call would have returned because no event was available, then the call returns `-1` with `t_errno` set to `[TSYSERR]` and `errno` set to `[EINTR]`. The state of the endpoint is unchanged.

In addition an `[EINTR]` error may be returned by all XTI calls (except `t_error()` and `t_strerror()`) under implementation defined conditions. In these cases the state of the endpoint will not have been changed, and no data will have been sent or received. Any buffers provided by the user for return values may have been overwritten.

A “well written” application will itself mask out signals except during specific code sequences (typically only its idle point) to avoid having to handle an `[EINTR]` return from all system calls.

Application writers should be aware that XTI calls may be implemented in a library as multiple system calls. In order to maintain the endpoint and associated library data areas in a consistent state, some of these system calls may be repeated when interrupted by a signal.

Applications should not call XTI functions from within a signal handler or using the `longjmp()` or `siglongjmp()` interfaces (see reference **XSH**) to exit a signal handler, as either may leave XTI data areas in an inconsistent state.

Applications may be able to cause the XTI library itself to generate signals that interrupt its internal actions (for example, by issuing `ioctl(fd, I_SETSIG, S_INPUT)` on a UNIX system); this may cause the user’s signal handler to be scheduled, but will not stop the XTI call from completing.

3.9 Event Management

Each XTI call deals with one transport endpoint at a time. It is not possible to wait for several events from different sources, particularly from several transport connections at a time. We recognise the need for this functionality which may be available today in a system-dependent fashion.

Throughout the document we refer to an event management service called Event Management (EM) which provides those functions useful to XTI. This Event Management will allow a process to be notified of the following events:

<code>T_LISTEN</code>	A connection request from a remote user was received by a transport provider (connection-mode service only); this event may occur under the following conditions: <ol style="list-style-type: none"> 1. The file descriptor is bound to a valid address. 2. No transport connection is established at this time.
<code>T_CONNECT</code>	In connection mode only; a connection response was received by the transport provider; occurs after a <code>t_connect()</code> has been issued.
<code>T_DATA</code>	Normal data (whole or part of Transport Service Data Unit (TSDU)) was received by the transport provider.
<code>T_EXDATA</code>	Expedited data was received by the transport provider.

T_DISCONNECT	In connection mode only; a disconnection request was received by the transport provider. It may be reported on both data transfer functions and connection establishment functions and on the <i>t_snddis()</i> function.
T_ORDREL	An orderly release request was received by a transport provider (connection mode with orderly release only).
T_UDERR	In connectionless-mode only; an error was found in a previously sent datagram. It may be notified on the <i>t_rcvudata()</i> or <i>t_unbind()</i> function calls.
T_GODATA	Flow control restrictions on normal data flow that led to a [TFLOW] error have been lifted. Normal data may be sent again.
T_GOEXDATA	Flow control restrictions on expedited data flow that led to a [TFLOW] error have been lifted. Expedited data may be sent again.

4.1 Overview of Connection-oriented Mode

The connection-mode transport service consists of four phases of communication:

- Initialisation/De-initialisation
- Connection Establishment
- Data Transfer
- Connection Release.

A state machine is described in Section C.1 on page 241, and the figure in Section C.2 on page 242, which defines the legal sequence in which functions from each phase may be issued.

In order to establish a transport connection, a user (application) must:

1. supply a *transport provider identifier* for the appropriate type of transport provider (using *t_open()*); this establishes a transport endpoint through which the user may communicate with the provider
2. associate (bind) an address with this endpoint (using *t_bind()*)
3. use the appropriate connection functions (using *t_connect()*, or *t_listen()* and *t_accept()*) to establish a transport connection; the set of functions depends on whether the user is an initiator or responder
4. once the connection is established, normal, and if authorised, expedited data can be exchanged; of course, expedited data may be exchanged only if:
 - the provider supports it
 - its use is not precluded by the selection of protocol characteristics; for example, the use of ISO Transport Class 0
 - both the peer transport providers support the feature, have negotiated it if necessary, and the peer applications have agreed to its use.

The semantics of expedited data may be quite different for different transport providers. XTI's notion of expedited data has been defined as the lowest reasonable common denominator.

The transport connection can be released at any time by using the disconnection functions. Then the user can either de-initialise the transport endpoint by closing the file descriptor returned by *t_open()* (thereby freeing the resource for future use), or specify a new local address (after the old one has been unbound) or reuse the same address and establish a new transport connection.

When reusing an endpoint, the user should be aware that the local address may be different from that originally bound (for example, one of the systems' local addresses may be explicitly associated with the endpoint, or the address may have been changed during *t_accept()*). The value of certain options will also reflect the previous connect and may need to be reset (see the

definition for `t_optmgmt()` to the providers' default values.

4.1.1 Initialisation/De-initialisation Phase

The functions that support initialisation/de-initialisation tasks are described below. All such functions provide local management functions.

<code>t_open()</code>	This function creates a transport endpoint and returns protocol-specific information associated with that endpoint. It also returns a file descriptor that serves as the local identifier of the endpoint.
<code>t_bind()</code>	This function associates a protocol address with a given transport endpoint, thereby activating the endpoint. It also directs the transport provider to begin accepting connection indications if so desired.
<code>t_unbind()</code>	This function disables a transport endpoint such that no further request destined for the given endpoint will be accepted by the transport provider.
<code>t_close()</code>	This function informs the transport provider that the user is finished with the transport endpoint, and frees any local resources associated with that endpoint.

The following functions are also local management functions, but can be issued during any phase of communication:

<code>t_getprotaddr()</code>	This function returns the addresses (local and remote) associated with the specified transport endpoint.
<code>t_getinfo()</code>	This function returns protocol-specific information associated with the specified transport endpoint.
<code>t_getstate()</code>	This function returns the current state of the transport endpoint.
<code>t_sync()</code>	This function synchronises the data structures managed by the transport library with the transport provider.
<code>t_alloc()</code>	This function allocates storage for the specified library data structure.
<code>t_free()</code>	This function frees storage for a library data structure that was allocated by <code>t_alloc()</code> .
<code>t_error()</code>	This function prints out a message describing the last error encountered during a call to a transport library function.
<code>t_look()</code>	This function returns the current event associated with the given transport endpoint.
<code>t_optmgmt()</code>	This function enables the user to get or negotiate protocol options with the transport provider.
<code>t_strerror()</code>	This function maps an XTI error into a language-dependent error message string.
<code>t_sysconf()</code>	This function is used to obtain the values of configurable and implementation-dependent XTI variables.

4.1.2 Overview of Connection Establishment

This phase enables two transport users to establish a transport connection between them. In the connection establishment scenario, one user is considered active and initiates the conversation, while the second user is passive and waits for a transport user to request a connection.

In connection mode:

- The user has first to establish an endpoint; that is, to open a communications path between the application and the transport provider.
- Once established, an endpoint must be bound to an address and more than one endpoint may be bound to the same address. A transport user can determine the addresses associated with a connection using the `t_getprotaddr()` function.
- An endpoint can be associated with one, and only one, established transport connection.
- It is possible to use an endpoint to receive and enqueue incoming connection indications (only if the provider is able to accept more than one outstanding connection indication; this mode of operation is declared at the time of calling `t_bind()` by setting `qlen` greater than 0). However, if more than one endpoint is bound to the same address, only one of them may be used in this way.
- The `t_listen()` function is used to look for an enqueued connection indication; if it finds one (at the head of the queue), it returns details of the connection indication, and a local sequence number which uniquely identifies this indication, or it may return a negative value with `t_errno` set to [TNODATA]. The number of outstanding connection requests to dequeue is limited by the value of the `qlen` parameter accepted by the transport provider on the `t_bind()` call.
- If the endpoint has more than one connection indication enqueued, the user should dequeue all connection indications (and disconnection indications) before accepting or rejecting any or all of them. The number of outstanding connection indications is limited by the value of the `qlen` parameter accepted by the transport provider on the call to `t_bind()`.
- When accepting a connection indication, the transport service user may issue the accept on the same (listening) endpoint or on a different endpoint.

If the same endpoint is used, the listening endpoint cannot receive or enqueue incoming connect indications for the duration of the connection.

If a different endpoint is used, the listening endpoint can continue to receive and enqueue incoming connect indications.

- If the user issues a `t_connect()` on a listening endpoint, again, that endpoint cannot receive or enqueue incoming connect indications for the duration of the connection.
- A connection attempt failure will result in a value `-1` returned from either the `t_connect()` or `t_rcvconnect()` call, with `t_errno` set to [TLOOK] indicating that a [T_DISCONNECT] event has arrived. In this case, the reason for the failure may be identified by issuing a `t_rcvdis()` call.

The functions that support these operations of connection establishment are:

`t_connect()` This function requests a connection to the transport user at a specified destination and waits for the remote user's response. This function may be executed in either synchronous or asynchronous mode. In synchronous mode, the function waits for the remote user's response before returning control to the local user. In asynchronous mode, the function initiates connection establishment but returns control to the local user before a

	response arrives.
<code>t_rcvconnect()</code>	This function enables an active transport user to determine the status of a previously sent connection request. If the request was accepted, the connection establishment phase will be complete on return from this function. This function is used in conjunction with <code>t_connect()</code> to establish a connection in an asynchronous manner, or to continue an interrupted synchronous-mode <code>t_connect()</code> call.
<code>t_listen()</code>	This function enables the passive transport user to receive connection indications from other transport users.
<code>t_accept()</code>	This function is issued by the passive user to accept a particular connection request after an indication has been received.

4.1.3 Overview of Data Transfer

Once a transport connection has been established between two users, data may be transferred back and forth over the connection in a full duplex way. The functions that support data transfer in connection mode are as follows:

<code>t_snd()</code>	This function enables transport users to send either normal or expedited data over a transport connection.
<code>t_rcv()</code>	This function enables transport users to receive either normal or expedited data on a transport connection.
<code>t_sndv()</code>	This function enables transport users to send either normal or expedited data from non-contiguous buffers over a transport connection.
<code>t_rcvv()</code>	This function enables transport users to receive either normal or expedited data into non-contiguous buffers on a transport connection.

Throughout the rest of this section, all references of calls to `t_rcv()` include calls to `t_rcvv()`, and calls to `t_snd()` include calls to `t_sndv()`.

In data transfer phase, the occurrence of the [T_DISCONNECT] event implies an unsuccessful return from the called function (`t_snd()` or `t_rcv()`) with `t_errno` set to [TLOOK]. The user must then issue a `t_look()` call to get more details.

Receiving Data

If data (normal or expedited) is immediately available, then a call to `t_rcv()` returns data. If the transport connection no longer exists, then the call returns immediately, indicating failure. If data is not immediately available and the transport connection still exists, then the result of a call to `t_rcv()` depends on the mode:

- Asynchronous Mode

The call returns immediately, indicating failure. The user must continue to “poll” for incoming data, either by issuing repeated call to `t_rcv()`, or by using the `t_look()` or the EM interface.

- Synchronous Mode

The call is blocked until one of the following conditions becomes true:

- Data (normal or expedited) is received.
- A disconnection indication is received.

- A signal has arrived.

The user may issue a `t_look()` or use EM calls, to determine if data is available.

If a normal TSDU is to be received in multiple `t_rcv()` calls, then its delivery may be interrupted at any time by the arrival of expedited data. The application can detect this by checking the *flags* field on return from a call to `t_rcv()`; this will be indicated by `t_rcv()` returning:

- data with T_EXPEDITED flag not set and T_MORE set (this is a fragment of normal data)
- data with T_EXPEDITED set (and T_MORE set or unset); this is an expedited message (whole or part of, depending on the setting of T_MORE). The provider will continue to return the expedited data (on this and subsequent calls to `t_rcv()`) until the end of the Extended Transport Service Data Unit (ETSDU) is reached, at which time it will continue to return normal data. It is the user's responsibility to remember that the receipt of normal data has been interrupted in this way.

Sending Data

If the data can be accepted immediately by the provider, then it is accepted, and the call returns the number of octets accepted. If the data cannot be accepted because of a permanent failure condition (for example, transport connection lost), then the call returns immediately, indicating failure. If the data cannot be accepted immediately because of a transient condition (for example, lack of buffers, flow control in effect), then the result of a call to `t_snd()` depends on the execution mode:

- Asynchronous Mode

The call returns immediately indicating failure. If the failure was due to flow control restrictions, then it is possible that only part of the data will actually be accepted by the transport provider. In this case `t_snd()` will return a value that is less than the number of octets requested to be sent. The user may either retry the call to `t_snd()` or first receive notification of the clearance of the flow control restriction via either `t_look()` or the EM interface, then retry the call. The user may retry the call with the data remaining from the original call or with more (or less) data, and with the T_MORE flag set appropriately to indicate whether this is now the end of the TSDU.

- Synchronous Mode

The call is blocked until one of the following conditions becomes true:

- The flow control restrictions are cleared and the transport provider is able to accept a new data unit. The `t_snd()` function then returns successfully.
- A disconnection indication is received. In this case the `t_snd()` function returns unsuccessfully with `t_errno` set to [TLOOK]. The user can issue a `t_look()` function to determine the cause of the error. For this particular case `t_look()` will return a T_DISCONNECT event. All or part of the data that was being sent might be lost.
- An internal problem occurs. In this case the `t_snd()` function returns unsuccessfully with `t_errno` set to [TSYSERR]. Data that was being sent will be lost.

For some transport providers, normal data and expedited data constitute two distinct flows of data. If either flow is blocked, the user may nevertheless continue using the other one, but in synchronous mode a second process is needed. The user may send expedited data between the fragments of a normal TSDU, that is, a `t_snd()` call with the T_EXPEDITED flag set may follow a `t_snd()` with the T_MORE flag set and the T_EXPEDITED flag not set.

Note that XTI supports two modes of sending data, record-oriented and stream-oriented. In the record-oriented mode, the concept of TSDU is supported, that is, message boundaries are preserved. In stream-oriented mode, message boundaries are not preserved and the concept of a TSDU is not supported. A transport user can determine the mode by using the `t_getinfo()` function, and examining the `tsdu` field. If `tsdu` is greater than zero, this indicates that record-oriented mode is supported and the return value indicates the maximum TSDU size. If `tsdu` is zero, this indicates that stream-oriented transfer is supported. For more details see `t_getinfo()` on page 71.

4.1.4 Overview of Connection Release

Some communication providers (for example, the ISO connection-mode transport) support only the abortive release. However, some communication providers (for example, mOSI, NetBIOS, SNA, TCP) also support an orderly release. XTI includes functions to provide access to transports that support or require the orderly release features.

An abortive release may be invoked from either the connection establishment phase or the data transfer phase. When in the connection establishment phase, a transport user may use the abortive release to reject a connection request. In the data transfer phase, either user may abort a connection at any time. The abortive release is not negotiated by the transport users and it takes effect immediately on request. The user on the other side of the connection is notified when a connection is aborted. The transport provider may also initiate an abortive release, in which case both users are informed that the connection no longer exists. There is no guarantee of delivery of user data once an abortive release has been initiated.

Whatever the state of a transport connection, its user(s) will be informed as soon as possible of the failure of the connection through a disconnection event or an unsuccessful return from a blocking `t_snd()` or `t_rcv()` call. If the user wants to prevent loss of data by notifying the remote user of an imminent connection release, it is the user's responsibility to use an upper level mechanism. For example, the user may send specific (expedited) data and wait for the response of the remote user before issuing a disconnection request.

Some transport providers support an orderly release capability (for example, mOSI, NetBIOS, SNA, TCP). If supported by the communications provider, orderly release may be invoked from the data transfer phase to enable two users to gracefully release a connection. The procedure for orderly release prevents the loss of data that may occur during an abortive release.

When supported by the underlying protocol, some communications providers optionally allow applications to send or retrieve user data with an orderly release, through the use of the `t_sndreldata()` or `t_rcvreldata()` functions instead of the `t_sndrel()` or `t_rcvrel()` functions.

The functions that support connection release are:

- | | |
|-----------------------|--|
| <i>t_snddis()</i> | This function can be issued by either transport user to initiate the abortive release of a transport connection. It may also be used to reject a connection request during the connection establishment phase. |
| <i>t_rcvdis()</i> | This function identifies the reason for the abortive release of a connection, where the connection is released by the transport provider or another transport user. |
| <i>t_sndrel()</i> | This function can be called by either transport user to initiate an orderly release. The connection remains intact until both users call this function and <i>t_rcvrel()</i> . |
| <i>t_rcvrel()</i> | This function is called when a user is notified of an orderly release request, as a means of informing the transport provider that the user is aware of the remote user's actions. |
| <i>t_sndreldata()</i> | This function can be used instead of <i>t_sndrel()</i> to send user data with an orderly release. |
| <i>t_rcvreldata()</i> | This function can be used instead of <i>t_rcvrel()</i> to retrieve orderly release user data. |

4.2 Overview of Connectionless Mode

The connectionless-mode transport service consists of two phases of communication: initialisation/de-initialisation and data transfer. A brief description of each phase and its associated functions is presented below. A state machine is described in Section C.1 on page 241, and the figure in Section C.3 on page 244, that defines the legal sequence in which functions from each phase may be issued.

In order to permit the transfer of connectionless-mode data, a user (application) must:

1. supply a transport endpoint for the appropriate type of provider (using `t_open()`); this establishes a transport endpoint through which the user may communicate with the provider
2. associate (bind) an address with this transport endpoint (using `t_bind()`)

The user may then send and/or receive connectionless-mode data, as required, using the functions `t_sndudata()` and `t_rcvudata()`. Once the data transfer phase is finished, the application may either directly close the file descriptor returned by `t_open()` (using `t_close()`), thereby freeing the resource for future use, or start a new exchange of data after disassociating the old address and binding a new one.

4.2.1 Initialisation/De-initialisation Phase

The functions that support the initialisation/de-initialisation tasks are the same functions used in the connection-mode service.

4.2.2 Overview of Data Transfer

Once a transport endpoint has been activated, a user is free to send and receive data units through that endpoint in connectionless mode as follows:

<code>t_sndudata()</code>	This function enables transport users to send a self-contained data unit to the user at the specified protocol address.
<code>t_sndvudata()</code>	This function enables transport users to send a self-contained data unit to the user from one or more non-contiguous buffers at the specified protocol address.
<code>t_rcvudata()</code>	This function enables transport users to receive data units from
<code>t_rcvvudata()</code>	This function enables transport users to receive data units from other users into one or more non-contiguous buffers.
<code>t_rcvuderr()</code>	This function enables transport users to retrieve error information associated with a previously sent data unit.

The only possible events reported to the user are [T_UDERR], [T_DATA] and [T_GODATA]. Expedited data cannot be used with a connectionless-mode transport provider.

Throughout the rest of this section, all references of calls to `rcvudata()` include calls to `rcvvudata()`, and calls to `t_sndudata()` include calls to `t_sndvudata()`.

Receiving Data

If data is available (a datagram or a part), the `t_rcvudata()` call returns immediately indicating the number of octets received. If data is not immediately available, then the result of the `t_rcvudata()` call depends on the chosen mode:

- Asynchronous Mode

The call returns immediately indicating failure. The user must either retry the call repeatedly, or “poll” for incoming data by using the EM interface or the `t_look()` function so as not to be blocked.

- Synchronous Mode

The call is blocked until one of the following conditions becomes true:

- A datagram is received.
- An error is detected by the transport provider.
- A signal has arrived.

The application may use the `t_look()` function or the EM mechanism to know if data is available instead of issuing a `t_rcvudata()` call which may be blocking.

Sending Data

- Synchronous Mode

In order to maintain some flow control, the `t_sndudata()` function will block until the datagram has been accepted by the provider. The call returns immediately after the datagram has been sent. A process which sends data in synchronous mode may be blocked for some time.

- Asynchronous Mode

The transport provider may refuse to send a new datagram for flow control restrictions. In this case, the `t_sndudata()` call fails returning a negative value and setting `t_errno` to [TFLOW]. The user may retry later or use the `t_look()` function or EM interface to be informed of the flow control restriction removal.

If `t_sndudata()` is called before the destination user has activated its transport endpoint, the data unit may be discarded.

4.3 XTI Features

The following functions, which correspond to the subset common to connection-mode and connectionless-mode services, are always implemented:

```
t_bind()
t_close()
t_look()
t_open()
t_sync()
t_unbind()
```

If a connection-mode Transport Service is provided, then the following functions are always implemented:

```
t_accept()
t_connect()
t_listen()
t_rcv()
t_rcvv()
t_rcvconnect()
t_rcvdis()
t_snd()
t_sndv()
t_snddis()
```

If XTI supports the access to the connectionless-mode Transport Service, the following three functions are always implemented:

```
t_rcvudata()
t_rcvvudata()
t_rcvuderr()
t_sndudata()
t_sndvudata()
```

Mandatory mechanisms:

- synchronous mode
- asynchronous mode.

Utility functions:

```
t_alloc()
t_free()
t_error()
t_getprotaddr()
t_getinfo()
t_getstate()
t_optmgmt()
t_strerror()
t_sysconf()
```

The orderly release mechanism (using *t_sndrel()*, *t_sndreldata()*, *t_rcvrel()* and *t_rcvreldata()*) is supported only for T_COTS_ORD type providers. Use with other providers will cause the [TNOTSUPPORT] error to be returned.

Optional mechanisms:

- the ability to manage (enqueue) more than one incoming connection indication at any one time
- the address of the caller passed with *t_accept()* may optionally be checked by an XTI implementation

4.3.1 XTI Functions versus Protocols

Table 4-1 presents all the functions defined in XTI. The character “x” indicates that the mapping of that function is possible onto a connection-mode or connectionless-mode Transport Service. The table indicates the type of utility functions as well.

Functions	Necessary for Protocol		Utility Functions	
	Connection Mode	Connectionless Mode	General	Memory
<i>t_accept()</i>	x			
<i>t_alloc()</i>				x
<i>t_bind()</i>	x	x		
<i>t_close()</i>	x	x		
<i>t_connect()</i>	x			
<i>t_error()</i>			x	
<i>t_free()</i>				x
<i>t_getprotaddr()</i>			x	
<i>t_getinfo()</i>			x	
<i>t_getstate()</i>			x	
<i>t_listen()</i>	x			
<i>t_look()</i>	x	x		
<i>t_open()</i>	x	x		
<i>t_optmgmt()</i>			x	
<i>t_rcv()</i>	x			
<i>t_rcvv()</i>	x			
<i>t_rcvconnect()</i>	x			
<i>t_rcvdis()</i>	x			
<i>t_rcvrel()</i>	x			
<i>t_rcvreldata()</i>	x			
<i>t_rcvudata()</i>		x		
<i>t_rcvvudata()</i>		x		
<i>t_rcvuderr()</i>		x		
<i>t_snd()</i>	x			
<i>t_sndv()</i>	x			
<i>t_snddis()</i>	x			
<i>t_sndrel()</i>	x			
<i>t_sndreldata()</i>	x			
<i>t_sndudata()</i>		x		
<i>t_sndvudata()</i>		x		
<i>t_strerror()</i>			x	
<i>t_sync()</i>			x	
<i>t_sysconf()</i>			x	
<i>t_unbind()</i>	x	x		

Table 4-1 Classification of the XTI Functions

States and Events in XTI

Table 5-1 through Table 5-7 are included to describe the possible states of the transport provider as seen by the transport user, to describe the incoming and outgoing events that may occur on any connection, and to identify the allowable sequence of function calls. Given a current state and event, the transition to the next state is shown as well as any actions that must be taken by the transport user.

The allowable sequence of functions is described in Table 5-5, Table 5-6 and Table 5-7. The support functions, *t_getprotaddr()*, *t_getstate()*, *t_getinfo()*, *t_alloc()*, *t_free()*, *t_look()* and *t_sync()*, are excluded from the state tables because they do not affect the state of the interface. Each of these functions may be issued from any state except the uninitialised state. Similarly, the *t_error()*, *t_strerror()* and *t_sysconf()* functions have been excluded from the state table because they do not affect the state of the interface.

5.1 Transport Interfaces States

XTI manages a transport endpoint by using at most 8 states:

- T_UNINIT
- T_UNBND
- T_IDLE
- T_OUTCON
- T_INCON
- T_DATAXFER
- T_INREL
- T_OUTREL.

The states T_OUTREL and T_INREL are significant only if the optional orderly release function is both supported and used.

Table 5-1 describes all possible states of the transport provider as seen by the transport user. The service type may be connection-mode, connection-mode with orderly release or connectionless-mode.

State	Description	Service Type
T_UNINIT	uninitialised - initial and final state of interface	T_COTS T_CLTS T_COTS_ORD
T_UNBND	unbound	T_COTS T_COTS_ORD T_CLTS
T_IDLE	no connection established	T_COTS T_COTS_ORD T_CLTS
T_OUTCON	outgoing connection pending for active user	T_COTS T_COTS_ORD
T_INCON	incoming connection pending for passive user	T_COTS T_COTS_ORD
T_DATAXFER	data transfer	T_COTS T_COTS_ORD
T_OUTREL	outgoing orderly release (waiting for orderly release indication)	T_COTS_ORD
T_INREL	incoming orderly release (waiting to send orderly release request)	T_COTS_ORD

Table 5-1 Transport Interface States

5.2 Outgoing Events

The following outgoing events correspond to the successful return or error return of the specified user-level transport functions causing XTI to change state, where these functions send a request or response to the transport provider. In Table 5-2, some events (for example, accept1, accept2 and accept3) are distinguished by the context in which they occur. The context is based on the values of the following:

ocnt Count of outstanding connection indications (connection indications passed to the user but not accepted or rejected).

fd File descriptor of the current transport endpoint.

resfd File descriptor of the transport endpoint where a connection will be accepted.

Event	Description	Service Type
opened	successful return of <i>t_open()</i>	T_COTS, T_COTS_ORD, T_CLTS
bind	successful return of <i>t_bind()</i>	T_COTS, T_COTS_ORD, T_CLTS
optmgmt	successful return of <i>t_optmgmt()</i>	T_COTS, T_COTS_ORD, T_CLTS
unbind	successful return of <i>t_unbind()</i>	T_COTS, T_COTS_ORD, T_CLTS
closed	successful return of <i>t_close()</i>	T_COTS, T_COTS_ORD, T_CLTS
connect1	successful return of <i>t_connect()</i> in synchronous mode	T_COTS, T_COTS_ORD
connect2	TNODATA error on <i>t_connect()</i> in asynchronous mode, or TLOOK error due to a disconnection indication arriving on the transport endpoint, or TSYSERR error and errno set to EINTR.	T_COTS, T_COTS_ORD
accept1	successful return of <i>t_accept()</i> with <i>ocnt</i> == 1, <i>fd</i> == <i>resfd</i>	T_COTS, T_COTS_ORD
accept2	successful return of <i>t_accept()</i> with <i>ocnt</i> == 1, <i>fd</i> != <i>resfd</i>	T_COTS, T_COTS_ORD
accept3	successful return of <i>t_accept()</i> with <i>ocnt</i> > 1	T_COTS, T_COTS_ORD
snd	successful return of <i>t_snd()</i> or <i>t_sndv()</i>	T_COTS, T_COTS_ORD
snddis1	successful return of <i>t_snddis()</i> with <i>ocnt</i> <= 1	T_COTS, T_COTS_ORD
snddis2	successful return of <i>t_snddis()</i> with <i>ocnt</i> > 1	T_COTS, T_COTS_ORD
sndrel	successful return of <i>t_sndrel()</i> or <i>t_sndreldata()</i>	T_COTS_ORD
sndudata	successful return of <i>t_sndudata()</i> or <i>t_sndvudata()</i>	T_CLTS

Table 5-2 Transport Interface Outgoing Events

Note: *ocnt* is only meaningful for the listening transport endpoint (*fd*).

5.3 Incoming Events

The following incoming events correspond to the successful return of the specified user-level transport functions, where these functions retrieve data or event information from the transport provider. One incoming event is not associated directly with the return of a function on a given transport endpoint:

pass_conn Occurs when a user transfers a connection to another transport endpoint. This event occurs on the endpoint that is being passed the connection, despite the fact that no function is issued on that endpoint. The event *pass_conn* is included in the state tables to describe what happens when a user accepts a connection on another transport endpoint.

In Table 5-3, the *rcvdis* events are distinguished by the context in which they occur. The context is based on the value of *ocnt*, which is the count of outstanding connection indications on the current transport endpoint.

Incoming Event	Description	Service Type
listen	successful return of <i>t_listen()</i>	T_COTS T_COTS_ORD
rcvconnect	successful return of <i>t_rcvconnect()</i>	T_COTS T_COTS_ORD
rcv	successful return of <i>t_rcv()</i> or <i>rcvv()</i>	T_COTS T_COTS_ORD
rcvdis1	successful return of <i>t_rcvdis()</i> with <i>ocnt == 0</i>	T_COTS T_COTS_ORD
rcvdis2	successful return of <i>t_rcvdis()</i> with <i>ocnt == 1</i>	T_COTS T_COTS_ORD
rcvdis3	successful return of <i>t_rcvdis()</i> with <i>ocnt > 1</i>	T_COTS T_COTS_ORD
rcvrel	successful return of <i>t_rcvrel()</i> or <i>rcvreldata()</i>	T_COTS_ORD
rcvudata	successful return of <i>t_rcvudata()</i> or <i>rcvvudata()</i>	T_CLTS
rcvuderr	successful return of <i>t_rcvuderr()</i>	T_CLTS
pass_conn	receive a passed connection	T_COTS T_COTS_ORD

Table 5-3 Transport Interface Incoming Events

5.4 Transport User Actions

Some state transitions are accompanied by a list of actions the transport user must take. These actions are represented by the notation $[n]$, where n is the number of the specific action as described in Table 5-4.

[1]	Set the count of outstanding connection indications to zero.
[2]	Increment the count of outstanding connection indications.
[3]	Decrement the count of outstanding connection indications.
[4]	Pass a connection to another transport endpoint as indicated in <i>t_accept()</i> .

Table 5-4 Transport Interface User Actions

5.5 State Tables

Table 5-5, Table 5-6 and Table 5-7 describe the possible next states, given the current state and event. The state is that of the transport provider as seen by the transport user.

The contents of each box represent the next state given the current state (column) and the current incoming or outgoing event (row). An empty box represents a state/event combination that is invalid. Along with the next state, each box may include an action list (as specified in Table 5-4 on page 35). The transport user must take the specific actions in the order specified in the state table.

A separate table is shown for initialisation/de-initialisation, data transfer in connectionless-mode and connection/release/data transfer in connection-mode.

state event	T_UNINIT	T_UNBND	T_IDLE
opened	T_UNBND		
bind		T_IDLE [1]	
unbind			T_UNBND
closed		T_UNINIT	T_UNINIT

Table 5-5 Initialisation/De-initialisation States

state event	T_IDLE
sndudata	T_IDLE
rcvudata	T_IDLE
rcvuderr	T_IDLE

Table 5-6 Data Transfer States: Connectionless-mode

state event	T_UNBND	T_IDLE	T_OUTCON	T_INCON	T_DATAXFER	T_OUTREL	T_INREL
<i>connect1</i>		T_DATAXFER					
<i>connect2</i>		T_OUTCON					
<i>rcvconnect</i>			T_DATAXFER				
<i>listen</i>		T_INCON[2]		T_INCON[2]			
<i>accept1</i>				T_DATAXFER[3]			
<i>accept2</i>				T_IDLE[3][4]			
<i>accept3</i>				T_INCON[3][4]			
<i>snd</i>					T_DATAXFER		T_INREL
<i>rcv</i>					T_DATAXFER	T_OUTREL	
<i>snddis1</i>			T_IDLE	T_IDLE[3]	T_IDLE	T_IDLE	T_IDLE
<i>snddis2</i>				T_INCON[3]			
<i>rcvdis1</i>			T_IDLE		T_IDLE	T_IDLE	T_IDLE
<i>rcvdis2</i>				T_IDLE[3]			
<i>rcvdis3</i>				T_INCON[3]			
<i>sndrel</i>					T_OUTREL		T_IDLE
<i>rcvrel</i>					T_INREL	T_IDLE	
<i>pass_conn</i>	T_DATAXFER	T_DATAXFER					
<i>optmgmt</i>	T_UNBND	T_IDLE	T_OUTCON	T_INCON	T_DATAXFER	T_OUTREL	T_INREL
<i>closed</i>		T_UNINIT	T_UNINIT	T_UNINIT	T_UNINIT	T_UNINIT	T_UNINIT

Table 5-7 Connection/Release/Data Transfer States: Connection-mode

5.6 Events and TLOOK Error Indication

The following list describes the asynchronous events which cause an XTI call to return with a [TLOOK] error:

<i>t_accept()</i>	T_DISCONNECT, T_LISTEN
<i>t_connect()</i>	T_DISCONNECT, T_LISTEN ⁶
<i>t_listen()</i>	T_DISCONNECT ⁷
<i>t_rcv()</i>	T_DISCONNECT, T_ORDREL ⁸
<i>t_rcvconnect()</i>	T_DISCONNECT
<i>t_rcvrel()</i>	T_DISCONNECT
<i>t_rcvreldata()</i>	T_DISCONNECT
<i>t_rcvudata()</i>	T_UDERR
<i>t_rcvv()</i>	T_DISCONNECT, T_ORDREL
<i>t_rcvvudata()</i>	T_UDERR
<i>t_snd()</i>	T_DISCONNECT, T_ORDREL
<i>t_sndudata()</i>	T_UDERR
<i>t_sndv()</i>	T_DISCONNECT, T_ORDREL
<i>t_sndvudata()</i>	T_UDERR
<i>t_unbind()</i>	T_LISTEN, T_DATA ⁹ , T_UDERR
<i>t_sndrel()</i>	T_DISCONNECT
<i>t_sndreldata()</i>	T_DISCONNECT
<i>t_snddis()</i>	T_DISCONNECT

Once a [TLOOK] error has been received on a transport endpoint via an XTI function, subsequent calls to that and other XTI functions, to which the same [TLOOK] error applies, will continue to return [TLOOK] until the event is consumed. An event causing the [TLOOK] error can be determined by calling *t_look()* and then can be consumed by calling the corresponding consuming XTI function as defined in Table 3-1.

-
6. This occurs only when a *t_connect* is done on an endpoint which has been bound with a *qlen* > 0 and for which a connection indication is pending.
 7. This event indicates a disconnection on an outstanding connection indication.
 8. This occurs only when all pending data has been read.
 9. T_DATA may only occur for the connectionless mode.

The Use of Options in XTI

6.1 Generalities

The functions `t_accept()`, `t_connect()`, `t_listen()`, `t_optmgmt()`, `t_rcvconnect()`, `t_rcvudata()`, `t_rcvuderr()` and `t_sndudata()` contain an *opt* argument of type `struct netbuf` as an input or output parameter. This argument is used to convey options between the transport user and the transport provider.

There is no general definition about the possible contents of options. There are general XTI options and those that are specific for each transport provider. Some options allow the user to tailor his communication needs, for instance by asking for high throughput or low delay. Others allow the fine-tuning of the protocol behaviour so that communication with unusual characteristics can be handled more effectively. Other options are for debugging purposes, or to request that certain operations should be performed.

Options values have meaning to, and are defined by, the protocol level in which they apply. In some cases, option values can be negotiated by a transport user. This includes the simple case where the transport user can simply enforce its use. Often, the transport provider or even the remote transport user may have the right to negotiate a value of lesser quality than a proposed one, that is, a delay may become longer, or a throughput may become lower.

It is useful to differentiate between options with end-to-end significance (sometimes referred to as *association-related*¹⁰ and those that are not. Association-related options are intimately related to the particular transport connection or datagram transmission. If the calling user specifies such an option, some ancillary information is transferred across the network in most cases. The interpretation and further processing of this information is protocol-dependent. For instance, in an ISO connection-mode communication, the calling user may specify quality-of-service parameters on connection establishment. These are first processed and possibly lowered by the local transport provider, then sent to the remote transport provider that may degrade them again, and finally conveyed to the called user that makes the final selection and transmits the selected values back to the caller.

Options without end-to-end significance do not contain information destined for the remote transport user. Some have purely local relevance, for example, an option that enables debugging. Others influence the transmission, for instance the option that sets the `T_IP time-to-live` field, or `T_TCP_NODELAY` (see Appendix B on page 229). Local options are negotiated solely between the transport user and the local transport provider. The distinction between these two categories of options is visible in XTI through the following relationship: on output, the functions `t_listen()` and `t_rcvudata()` return options with end-to-end significance only. The functions `t_rcvconnect()` and `t_rcvuderr()` may return options of both categories. On input, options of both categories may be specified with `t_accept()` and `t_sndudata()`. The functions `t_connect()` and `t_optmgmt()` can process and return both categories of options.

10. The term “association” is used to denote a pair of communicating transport users, that is, the communication has end-to-end significance.

Options can be further distinguished between those which maintain a persistent value and those which are requests to perform actions to which the option values are effectively parameters. For each action that it supports and that maintains a persistent value, the transport provider has a default value. These defaults are sufficient for the majority of communication relations. Hence, a transport user should only request options actually needed to perform the task, and leave all others at their default value.

This chapter describes the general framework for the use of options. This framework is obligatory for all transport providers. The specific options that are legal for use with a specific transport provider are described in the provider-specific appendices (see Appendix A on page 215 and Appendix B on page 229). General XTI options are described in `t_optmgmt()` on page 84.

6.2 The Format of Options

Options are conveyed via an `opt` argument of **struct netbuf**. Each option in the buffer specified is of the form **struct t_opthdr** possibly followed by an option value.

A transport provider embodies a stack of protocols. The `level` field of **struct t_opthdr** identifies the XTI level or a protocol of the transport provider such as T_TCP or ISO 8073:1986. The `name` field identifies the option within the level, and `len` contains its total length, that is, the length of the option header **t_opthdr** plus the length of the option value. The `status` field is used by the XTI level or the transport provider to indicate success or failure of a negotiation (see Section 6.3.5 on page 44 and `t_optmgmt()` on page 84).

Several options can be concatenated. The option user has, however to ensure that each options header and value part starts at a boundary appropriate for the architecture-specific alignment rules. The macros T_OPT_FIRSTHDR(nbp), T_OPT_NXTHDR(nbp, tohp), T_OPT_DATA(tohp) are provided for that purpose. These macros are defines in the man-page for `t_optmgmt()` (see `t_optmgmt()` on page 84).

T_OPT_FIRSTHDR is useful for finding an appropriately aligned start of the option buffer. T_OPT_NEXTHDR is useful for moving to the start of the next appropriately aligned option in the option buffer. Note that OPT_NEXTHDR is also available for backward compatibility requirements. T_OPT_DATA is useful for finding the start of the data part in the option buffer where the contents of its values start on an appropriately aligned boundary.

The length of the option buffer is given by `opt.len`. The alignment characters are included in the length.

6.3 The Elements of Negotiation

This section describes the general rules governing the passing and retrieving of options and the error conditions that can occur. Unless explicitly restricted, these rules apply to all functions that allow the exchange of options.

6.3.1 Multiple Options and Options Levels

When multiple options are specified in an option buffer on input, different rules apply to the levels that may be specified, depending on the function call. Multiple options specified on input to `t_optmgmt()` must address the same option level. Options specified on input to `t_connect()`, `t_accept()` and `t_sndudata()` can address different levels.

6.3.2 Illegal Options

Only legal options can be negotiated; illegal options cause failure. An option is illegal if the following applies:

- The length specified in `t_opthdr.len` exceeds the remaining size of the option buffer (counted from the beginning of the option).
- The option value is illegal. The legal values are defined for each option. (See `t_optmgmt()` on page 84 and the protocol specific appendices where options are described for specific protocols).

If an illegal option is passed to XTI, the following will happen:

- If an illegal option is passed to `t_optmgmt()` then the function fails with `t_errno` set to [TBADOPT].
- If an illegal option is passed to `t_accept()` or `t_connect()` then either the function fails with `t_errno` set to [TBADOPT] or the connection establishment fails at a later stage, depending on when the implementation detects the illegal option.
- If an illegal option is passed to `t_sndudata()` then either the function fails with `t_errno` set to [TBADOPT] or it successfully returns but a T_UDERR event occurs to indicate that the datagram was not sent.

If the transport user passes multiple options in one call and one of them is illegal, the call fails as described above. It is, however, possible that some or even all of the submitted legal options were successfully negotiated. The transport user can check the current status by a call to `t_optmgmt()` with the T_CURRENT flag set (see `t_optmgmt()` on page 84).

Specifying an option level unknown to the transport provider does not cause failure in calls to `t_connect()`, `t_accept()` or `t_sndudata()`; the option is discarded in these cases. The function `t_optmgmt()` fails with [TBADOPT].

Specifying an option name that is unknown to or not supported by the protocol selected by the option level does not cause failure. The option is discarded in calls to `t_connect()`, `t_accept()` or `t_sndudata()`. The function `t_optmgmt()` returns T_NOTSUPPORT in the `status` field of the option.

6.3.3 Initiating an Option Negotiation

A transport user initiates an option negotiation when calling `t_connect()`, `t_sndudata()` or `t_optmgmt()` with the flag `T_NEGOTIATE` set.

The negotiation rules for these functions depend on whether an option request is an absolute requirement or not. This is explicitly defined for each option (see `t_optmgmt()` on page 84 and the protocol specific appendices where options are described for specific protocols). In case of an ISO transport provider, for example, the option that requests use of expedited data is not an absolute requirement. On the other hand, the option that requests protection could be an absolute requirement.

Note: The notion “absolute requirement” originates from the quality-of-service parameters in ISO 8072:1986. Its use is extended here to all options.

If the proposed option value is an absolute requirement, three outcomes are possible:

- The negotiated value is the same as the proposed one. When the result of the negotiation is retrieved, the *status* field in `t_opthdr` is set to `T_SUCCESS`.
- The negotiation is rejected if the option is supported but the proposed value cannot be negotiated. This leads to the following behaviour:
 - `t_optmgmt()` successfully returns, but the returned option has its *status* field set to `T_FAILURE`.
 - Any attempt to establish a connection aborts; a `T_DISCONNECT` event occurs, and a synchronous call to `t_connect()` fails with `[TLOOK]`.
 - `t_sndudata()` fails with `[TLOOK]` or successfully returns, but a `T_UDERR` event occurs to indicate that the datagram was not sent.

If multiple options are submitted in one call and one of them is rejected, XTI behaves as just described. Although the connection establishment or the datagram transmission fails, options successfully negotiated before some option was rejected retain their negotiated values. There is no roll-back mechanism (see Section 6.4 on page 46).

The function `t_optmgmt()` attempts to negotiate each option. The *status* fields of the returned options indicate success (`T_SUCCESS`) or failure (`T_FAILURE`).

- If the local transport provider does not support the option at all, `t_optmgmt()` reports `T_NOTSUPPORT` in the *status* field. The functions `t_connect()` and `t_sndudata()` ignore this option.

If the proposed option value is not an absolute requirement, two outcomes are possible:

- The negotiated value is of equal or lesser quality than the proposed one (for example, a delay may become longer).

When the result of the negotiation is retrieved, the *status* field in `t_opthdr` is set to `T_SUCCESS` if the negotiated value equals the proposed one, or set to `T_PARTSUCCESS` otherwise.

- If the local transport provider does not support the option at all, `t_optmgmt()` reports `T_NOTSUPPORT` in the *status* field. The functions `t_connect()` and `t_sndudata()` ignore this option.

Unsupported options do not cause functions to fail or a connection to abort, since different vendors possibly implement different subsets of options. Furthermore, future enhancements of XTI might encompass additional options that are unknown to earlier implementations of transport providers. The decision whether or not the missing support of an option is acceptable

for the communication is left to the transport user.

The transport provider does not check for multiple occurrences of the same option, possibly with different option values. It simply processes the options in the option buffer one after the other. However, the user should not make any assumption about the order of processing.

Not all options are independent of one another. A requested option value might conflict with the value of another option that was specified in the same call or is currently effective (see Section 6.4 on page 46). These conflicts may not be detected at once, but later they might lead to unpredictable results. If detected at negotiation time, these conflicts are resolved within the rules stated above. The outcomes may thus be quite different and depend on whether absolute or non-absolute requests are involved in the conflict.

Conflicts are usually detected at the time a connection is established or a datagram is sent. If options are negotiated with `t_optmgmt()`, conflicts are usually not detected at this time, since independent processing of the requested options must allow for temporal inconsistencies.

When called, the functions `t_connect()` and `t_sndudata()` initiate a negotiation of *all* options with end-to-end significance according to the rules of this section. Options not explicitly specified in the function calls themselves are taken from an internal option buffer that contains the values of a previous negotiation (see Section 6.4 on page 46).

6.3.4 Responding to a Negotiation Proposal

In connection-mode communication, some protocols give the peer transport users the opportunity to negotiate characteristics of the transport connection to be established. These characteristics are options with end-to-end significance. With the connection indication, the called user receives (via `t_listen()`) a proposal about the option values that should be effective for this connection. The called user can accept this proposal or weaken it by choosing values of lower quality (for example, longer delays than proposed). The called user can, of course, refuse the connection establishment altogether.

The called user responds to a negotiation proposal via `t_accept()`. If the called transport user tries to negotiate an option of higher quality than proposed, the outcome depends on the protocol to which that option applies. Some protocols may reject the option, some protocols take other appropriate action described in protocol-specific appendices. If an option is rejected, the following error occurs:

The connection fails; a T_DISCONNECT event occurs. It depends on timing and implementation conditions whether the `t_accept()` call still succeeds or fails with [TLOOK].

If multiple options are submitted with `t_accept()` and one of them is rejected, the connection fails as described above. Options that could be successfully negotiated before the erroneous option was processed retain their negotiated value. There is no roll-back mechanism (see Section 6.4 on page 46).

The response options can be specified with the `t_accept()` call. Alternatively, they can be specified by calling `t_optmgmt()` and passing it the file descriptor that will subsequently be passed as `resfd` to `t_accept()` to identify the responding endpoint (see Section 6.4 on page 46). In case of conflict between option settings made by calls to `t_optmgmt()` and `t_accept()` at different times, the latest settings when `t_accept()` is called shall prevail. Note that the response to a negotiation proposal is activated when `t_accept()` is called. A `t_optmgmt()` call with erroneous option values as described above shall succeed; the connection aborts at the time `t_accept()` is called.

The connection also fails if the selected option values lead to contradictions.

The function `t_accept()` does not check for multiple specification of an option (see Section 6.3.3 on page 42). Unsupported options are ignored.

6.3.5 Retrieving Information about Options

This section describes how a transport user can retrieve information about options. To be explicit, a transport user must be able to:

- know the result of a negotiation (for example, at the end of a connection establishment)
- know the proposed option values under negotiation (during connection establishment)
- retrieve option values sent by the remote transport user for notification only (for example, T_IP options)
- check option values currently effective for the transport endpoint.

To this end, the functions `t_connect()`, `t_listen()`, `t_optmgmt()`, `t_rcvconnect()`, `t_rcvudata()` and `t_rcvuderr()` take an output argument `opt` of **struct netbuf**. The transport user has to supply a buffer where the options shall be written to; `opt.buf` must point to this buffer, and `opt.maxlen` must contain the buffer's size. The transport user can set `opt.maxlen` to zero to indicate that no options are to be retrieved.

Which options are returned depend on the function call involved:

`t_connect()` (synchronous mode) and `t_rcvconnect()`

The functions return the values of all options with end-to-end significance that were received with the connection response and the negotiated values of those options without end-to-end significance that had been specified on input. However, options specified on input in the `t_connect()` call that are not supported or refer to an unknown option level are discarded and not returned on output.

The `status` field of each option returned with `t_connect()` or `t_rcvconnect()` indicates if the proposed value (T_SUCCESS) or a degraded value (T_PARTSUCCESS) has been negotiated. The `status` field of received ancillary information (for example, T_IP options) that is not subject to negotiation is always set to T_SUCCESS.

`t_listen()`

The received options with end-to-end significance are related to the incoming connection (identified by the sequence number), not to the listening endpoint. (However, the option values currently effective for the listening endpoint can affect the values retrieved by `t_listen()`, since the transport provider might be involved in the negotiation process, too.) Thus, if the same options are specified in a call to `t_optmgmt()` with action T_CURRENT, will usually not return the same values.

The number of received options may be variable for subsequent connection indications, since many options with end-to-end significance are only transmitted on explicit demand by the calling user (for example, T_IP options or ISO 8072:1986 throughput). It is even possible that no options at all are returned.

The `status` field is irrelevant.

`t_rcvudata()`

The received options with end-to-end significance are related to the incoming datagram, not to the transport endpoint `fd`. Thus, if the same options are specified in a call to `t_optmgmt()` with action T_CURRENT, `t_optmgmt()` will usually not return the same values.

The number of options received may vary from call to call.

The *status* field is irrelevant.

t_rcvuderr() The returned options are related to the options input at the previous *t_sndudata()* call that produced the error. Which options are returned and which values they have depend on the specific error condition.

The *status* field is irrelevant.

t_optmgmt() This call can process and return both categories of options. It acts on options related to the specified transport endpoint, not on options related to a connection indication or an incoming datagram. A detailed description is given in *t_optmgmt()* on page 84.

6.3.6 Privileged and Read-only Options

Privileged options or option values are those that may be requested by privileged users only. The meaning of privilege is hereby implementation-defined.

Read-only options serve for information purposes only. The transport user may be allowed to read the option value but not to change it. For instance, to select the value of a protocol timer or the maximum length of a protocol data unit may be too subtle to leave to the transport user, though the knowledge about this value might be of some interest. An option might be read-only for all users or solely for non-privileged users. A privileged option might be inaccessible or read-only for non-privileged users.

An option might be negotiable in some XTI states and read-only in other XTI states. For instance, the ISO quality-of-service options are negotiable in the states T_IDLE and T_INCON and read-only in all other states (except T_UNINIT).

If a transport user requests negotiation of a read-only option, or a non-privileged user requests illegal access to a privileged option, the following outcomes are possible:

- *t_optmgmt()* successfully returns, but the returned option has its *status* field set to T_NOTSUPPORT if a privileged option was requested illegally, and to T_READONLY if modification of a read-only option was requested.
- If negotiation of a read-only option is requested, *t_accept()* or *t_connect()* either fail with [TACCES], or the connection establishment aborts and a T_DISCONNECT event occurs. If the connection aborts, a synchronous call to *t_connect()* fails with [TLOOK]. It depends on timing and implementation conditions whether a *t_accept()* call still succeeds or fails with [TLOOK].

If a privileged option is illegally requested, the option is quietly ignored. (A non-privileged user shall not be able to select an option which is privileged or unsupported.)

- If negotiation of a read-only option is requested, *t_sndudata()* may return [TLOOK] or successfully return, but a T_UDERR event occurs to indicate that the datagram was not sent.

If a privileged option is illegally requested, the option is quietly ignored. (A non-privileged user shall not be able to select an option which is privileged or unsupported.)

If multiple options are submitted to *t_connect()*, *t_accept()* or *t_sndudata()* and a read-only option is rejected, the connection or the datagram transmission fails as described. Options that could be successfully negotiated before the erroneous option was processed retain their negotiated values. There is no roll-back mechanism (see also Section 6.4 on page 46).

6.4 Option Management of a Transport Endpoint

This section describes how option management works during the lifetime of a transport endpoint.

Each transport endpoint is (logically) associated with an internal option buffer. When a transport endpoint is created, this buffer is filled with a system default value for each supported option. Depending on the option, the default may be 'OPTION ENABLED', 'OPTION DISABLED' or denote a time span, etc. These default settings are appropriate for most uses. Whenever an option value is modified in the course of an option negotiation, the modified value is written to this buffer and overwrites the previous one. At any time, the buffer contains all option values that are currently effective for this transport endpoint.

The current value of an option can be retrieved at any time by calling `t_optmgmt()` with the flag `T_CURRENT` set. Calling `t_optmgmt()` with the flag `T_DEFAULT` set yields the system default for the specified option.

A transport user can negotiate new option values by calling `t_optmgmt()` with the flag `T_NEGOTIATE` set. The negotiation follows the rules described in Section 6.3 on page 41.

Some options may be modified only in specific XTI states and are read-only in other XTI states. Many options with end-to-end significance, for instance, may not be changed in the state `T_DATAXFER`, and an attempt to do so will fail (see Section 6.3.6 on page 45). The legal states for each option are specified with its definition.

As usual, options with end-to-end significance take effect at the time a connection is established or a datagram is transmitted. This is the case if they contain information that is transmitted across the network or determine specific transmission characteristics. If such an option is modified by a call to `t_optmgmt()`, the transport provider checks whether the option is supported and negotiates a value according to its current knowledge. This value is written to the internal option buffer.

The final negotiation takes place if the connection is established or the datagram is transmitted. This can result in a degradation of the option value or even in a negotiation failure. The negotiated values are written to the internal option buffer.

Some options may be changed in the state `T_DATAXFER`, for example, those specifying buffer sizes. Such changes might affect the transmission characteristics and lead to unexpected side effects (for example, data loss if a buffer size was shortened) if the user does not care.

The transport user can explicitly specify both categories of options on input when calling `t_connect()`, `t_accept()` or `t_sndudata()`. The options are at first locally negotiated option-by-option, and the resulting values written to the internal option buffer. The modified option buffer is then used if a further negotiation step across the network is required, as for instance in connection-oriented ISO communication. The newly negotiated values are then written to the internal option buffer.

At any stage, a negotiation failure can lead to an abort of the transmission. If a transmission aborts, the option buffer will preserve the content it had at the time the failure occurred. Options that could be negotiated just before the error occurred are written back to the option buffer, whether the XTI call fails or succeeds.

It is up to the transport user to decide which options it explicitly specifies on input when calling `t_connect()`, `t_accept()` or `t_sndudata()`. The transport user need not pass options at all, by setting the `len` field of the function's input `opt` argument to zero. The current content of the internal option buffer is then used for negotiation without prior modification.

The negotiation procedure for options at the time of a `t_connect()`, `t_accept()` or `t_sndudata()` call always obeys the rules in Section 6.3.3 on page 42 and Section 6.3.4 on page 43, whether the options were explicitly specified during the call or implicitly taken from the internal option buffer.

The transport user should not make assumptions about the order in which options are processed during negotiation.

A value in the option buffer is only modified as a result of a successful negotiation of this option. It is, in particular, not changed by a connection release. There is no history mechanism that would restore the buffer state existing prior to the connection establishment or the datagram transmission. The transport user must be aware that a connection establishment or a datagram transmission may change the internal option buffer, even if each option was originally initialised to its default value.

6.5 Supplements

This section contains supplementary remarks and a short summary.

6.5.1 The Option Value T_UNSPEC

Some options may not have a fully specified value all the time. An ISO transport provider, for instance, that supports several protocol classes, might not have a preselected preferred class before a connection establishment is initiated. At the time of the connection request, the transport provider may conclude from the destination address, quality-of-service parameters and other locally available information which preferred class it should use. A transport user asking for the default value of the preferred class option in state T_IDLE would get the value T_UNSPEC. This value indicates that the transport provider did not yet select a value. The transport user could negotiate another value as the preferred class, for example, T_CLASS2. The transport provider would then be forced to initiate a connection request with class 2 as the preferred class.

An XTI implementation may also return the value T_UNSPEC if it can currently not access the option value. This may happen, for example, in the state T_UNBND in systems where the protocol stacks reside on separate controller cards and not in the host. The implementation may never return T_UNSPEC if the option is not supported at all.

If T_UNSPEC is a legal value for a specific option, it may be used by the user on input, too. It is used to indicate that it is left to the provider to choose an appropriate value. This is especially useful in complex options as ISO throughput, where the option value has an internal structure (see T_TCO_THROUGHPUT in Appendix A on page 215). The transport user may leave some fields unspecified by selecting this value. If the user proposes T_UNSPEC, the transport provider is free to select an appropriate value. This might be the default value, some other explicit value, or T_UNSPEC.

Each option where it is legal to use T_UNSPEC specifies its use as part of its description.

6.5.2 The info Argument

The functions *t_open()* and *t_getinfo()* return values representing characteristics of the transport provider in the argument *info*. The value of *info→options* is used by *t_alloc()* to allocate storage for an option buffer to be used in an XTI call. The value is sufficient for all uses.

In general, *info→options* also includes the size of privileged options, even if these are not read-only for non-privileged users. Alternatively, an implementation can choose to return different values in *info→options* for privileged and non-privileged users.

The values in *info→etsdu*, *info→tsdu*, *info→connect* and *info→discon* may be modified as soon as the T_DATAXFER state is entered. Calling *t_optmgmt()* need not influence these values (see *t_optmgmt()* on page 84).

6.5.3 Summary

- The format of an option is defined by a header **struct t_opthdr**, followed by an option value.
- On input, several options can be specified in an input *opt* argument. Each option must begin on a *t_uscalar_t* boundary.
- There are options with end-to-end significance and options without end-to-end significance. On output, the functions *t_listen()* and *t_rcvudata()* return options with end-to-end significance only. The functions *t_rcvconnect()* and *t_rcvuderr()* may return options of both categories. On input, options of both categories may be specified with *t_accept()* and

t_sndudata(). The functions *t_connect()* and *t_optmgmt()* can process and return both categories of options.

- A transport endpoint is (logically) associated with an internal option buffer, where the currently effective values are stored. Each successful negotiation of an option modifies this buffer, regardless of whether the call initiating the negotiation succeeds or fails.
- When calling *t_connect()*, *t_accept()* or *t_sndudata()*, the transport user can choose to submit the currently effective option values by setting the *len* field of the input *opt* argument to zero.
- If a connection is accepted via *t_accept()*, the explicitly specified option values together with the currently effective option values of *resfd*, not of *fd*, matter in this negotiation step.
- The options returned by *t_rcvuderr()* are those negotiated with the outgoing datagram that produced the error. If the error occurred during option negotiation, the returned option might represent some mixture of partly negotiated and not-yet negotiated options.

6.6 Portability Aspects

An application programmer who writes XTI programs faces two portability aspects:

- portability across protocol profiles
- portability across different system platforms (possibly from different vendors).

Options are intrinsically coupled with a definite protocol or protocol profile. Making explicit use of them therefore degrades portability across protocol profiles.

Different vendors might offer transport providers with different option support. This is due to different implementations and product policies. The lists of options on the *t_optmgmt()* manual page and in the protocol-specific appendices are maximal sets but do not necessarily reflect common implementation practice. Vendors will implement subsets that suit their needs. Making careless use of options therefore endangers portability across different system platforms.

Every implementation of a protocol profile accessible by XTI can be used with the default values of options. Applications can thus be written that do not care about options at all.

An application program that processes options retrieved from an XTI function should discard options it does not know in order to lessen its dependence from different system platforms and future XTI releases with possibly increased option support.

XTI Library Functions and Parameters

The function synopses in this Chapter 6, and the function definitions in Appendix F, conform to ISO C standard (see **referenced documents**).

7.1 How to Prepare XTI Applications

In a software development environment, a program, for example that uses XTI functions must be compiled with the XTI Library.

An application may also have a requirement to define some feature test macro(s) as described in Section 1.3 on page 3 (The Compilation Environment) to enable the functionality described in this specification.

The generic XTI structures and constants are all defined in the `<xti.h>` header, which can be found in Appendix F on page 281.

Any protocol-specific header files described in Appendices of this specification and needed by the application must follow the inclusion of `<xti.h>`.

Note: When `_XOPEN_SOURCE` is defined to be less than 500, the definitions from some protocol specific headers are automatically exposed by the inclusion of `<xti.h>`. The individual appendices for those protocols document the headers for which this applies.

7.2 Key for Parameter Arrays

For each XTI function description, a table is given which summarises the contents of the input and output parameter. The key is given below:

- x The parameter value is meaningful. (Input parameter must be set before the call and output parameter may be read after the call.)
- (x) The content of the object pointed to by the x pointer is meaningful.
- ? The parameter value is meaningful but the parameter is optional.
- (?) The content of the object pointed to by the ? pointer is optional.
- / The parameter value is meaningless.
- = The parameter after the call keeps the same value as before the call.

7.3 Return of TLOOK Error

Many of the XTI functions contained in this chapter return a [TLOOK] error to report the occurrence of an asynchronous event. For these functions a complete list describing the function and the events is provided in Section 5.6 on page 38.

7.4 Use of “struct netbuf”

Many of the XTI functions have an argument that points to a structure with members of type **struct netbuf**. These are used to pass arbitrary length buffer items such as user data, network addresses and option buffers to and from the XTI interfaces.

struct netbuf contains the following members:

```
unsigned int maxlen
unsigned int len
void        *buf
```

When a buffer item is being passed to the provider, *buf* and *len* are the address and length in bytes of the data in the user’s buffer. *maxlen* is not used and may take any value.

When a buffer item is being passed to the user, prior to the call the user must set *buf* and *maxlen* to the address and length of the user’s buffer. On return the provider will set *len* to the number of bytes of data placed into the user’s buffer. If the user’s buffer has a non-zero *maxlen* but is not long enough for the buffer item, then the provider will fail the XTI call, setting *t_errno* to *TBUFOVFLW* and discarding the event being returned to the user. If *maxlen* is zero then the provider will discard the data that would have been returned, and ignore *len* and *buf* (which need not have been a valid address).

NAME

t_accept - accept a connection request

SYNOPSIS

```
#include <xti.h>

int t_accept(
    int fd,
    int resfd,
    const struct t_call *call)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>resfd</i>	x	/
<i>call</i> → <i>addr.maxlen</i>	=	=
<i>call</i> → <i>addr.len</i>	x	=
<i>call</i> → <i>addr.buf</i>	? (?)	=
<i>call</i> → <i>opt.maxlen</i>	=	=
<i>call</i> → <i>opt.len</i>	x	=
<i>call</i> → <i>opt.buf</i>	? (?)	=
<i>call</i> → <i>udata.maxlen</i>	=	=
<i>call</i> → <i>udata.len</i>	x	=
<i>call</i> → <i>udata.buf</i>	? (?)	=
<i>call</i> → <i>sequence</i>	x	=

This function is issued by a transport user to accept a connection request. The parameter *fd* identifies the local transport endpoint where the connection indication arrived; *resfd* specifies the local transport endpoint where the connection is to be established, and *call* contains information required by the transport provider to complete the connection. The parameter *call* points to a **t_call** structure which contains the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
int sequence;
```

In *call*, *addr* is the protocol address of the calling transport user, *opt* indicates any options associated with the connection, *udata* points to any user data to be returned to the caller, and *sequence* is the value returned by *t_listen()* that uniquely associates the response with a previously received connection indication. The address of the caller, *addr* may be null (length zero). Where *addr* is not null then it may optionally be checked by XTI.

A transport user may accept a connection on either the same, or on a different, local transport endpoint than the one on which the connection indication arrived. Before the connection can be accepted on the same endpoint (*resfd==fd*), the user must have responded to any previous connection indications received on that transport endpoint (via *t_accept()* or *t_snddis()*). Otherwise, *t_accept()* will fail and set *t_errno* to [TINDOUT].

If a different transport endpoint is specified (*resfd!=fd*), then the user may or may not choose to bind the endpoint before the *t_accept()* is issued. If the endpoint is not bound prior to the *t_accept()*, the endpoint must be in the T_UNBND state before the *t_accept()* is issued, and the transport provider will automatically bind it to an address that is appropriate for the protocol concerned. If the transport user chooses to bind the endpoint it must be bound to a protocol address with a *qlen* of zero and must be in the T_IDLE state before the *t_accept()* is issued.

Responding endpoints should be supplied to *t_accept()* in the state T_UNBND.

The call to *t_accept()* may fail with *t_errno* set to [TLOOK] if there are indications (for example connect or disconnect) waiting to be received on endpoint *fd*. Applications should be prepared for such a failure.

The *udata* argument enables the called transport user to send user data to the caller and the amount of user data must not exceed the limits supported by the transport provider as returned in the *connect* field of the *info* argument of *t_open()* or *t_getinfo()*. If the *len* field of *udata* is zero, no data will be sent to the caller. All the *maxlen* fields are meaningless.

When the user does not indicate any option (*call*→*opt.len* = 0) the connection shall be accepted with the option values currently set for the responding endpoint *resfd* (see Section 6.3.4 on page 43 and Section 6.4 on page 46).

CAVEATS

There may be transport provider-specific restrictions on address binding. See Appendix A on page 215 and Appendix B on page 229.

Some transport providers do not differentiate between a connection indication and the connection itself. If the connection has already been established after a successful return of *t_listen()*, *t_accept()* will assign the existing connection to the transport endpoint specified by *resfd* (see Appendix B on page 229).

VALID STATES

fd: T_INCON

resfd (*fd*!=*resfd*): T_IDLE, T_UNBND

ERRORS

On failure, *t_errno* is set to one of the following:

[TACCES]	The user does not have permission to accept a connection on the responding transport endpoint or to use the specified options.
[TBADADDR]	The specified protocol address was in an incorrect format or contained illegal information.
[TBADDATA]	The amount of user data specified was not within the bounds allowed by the transport provider.
[TBADF]	The file descriptor <i>fd</i> or <i>resfd</i> does not refer to a transport endpoint.
[TBADOPT]	The specified options were in an incorrect format or contained illegal information.
[TBADSEQ]	Either an invalid sequence number was specified, or a valid sequence number was specified but the connection request was aborted by the peer. In the latter case, its T_DISCONNECT event will be received on the listening endpoint.
[TINDOUT]	The function was called with <i>fd</i> == <i>resfd</i> but there are outstanding connection indications on the endpoint. Those other connection indications must be handled either by rejecting them via <i>t_snddis(3)</i> or accepting them on a different endpoint via <i>t_accept(3)</i> .
[TLOOK]	An asynchronous event has occurred on the transport endpoint referenced by <i>fd</i> and requires immediate attention.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.

[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> or <i>resfd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TPROVMISMATCH]	The file descriptors <i>fd</i> and <i>resfd</i> do not refer to the same transport provider.
[TRESADDR]	This transport provider requires both <i>fd</i> and <i>resfd</i> to be bound to the same address. This error results if they are not.
[TRESQLEN]	The endpoint referenced by <i>resfd</i> (where <i>resfd</i> != <i>fd</i>) was bound to a protocol address with a <i>qlen</i> that is greater than zero.
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_connect(), *t_getstate()*, *t_listen()*, *t_open()*, *t_optmgmt()*, *t_rcvconnect()*.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_alloc - allocate a library structure

SYNOPSIS

#include <xti.h>

```
void *t_alloc(
    int fd,
    int struct_type,
    int fields)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>struct_type</i>	x	/
<i>fields</i>	x	/

The *t_alloc()* function dynamically allocates memory for the various transport function argument structures as specified below. This function will allocate memory for the specified structure, and will also allocate memory for buffers referenced by the structure.

The structure to allocate is specified by *struct_type* and must be one of the following:

T_BIND	struct	t_bind
T_CALL	struct	t_call
T_OPTMGMT	struct	t_optmgmt
T_DIS	struct	t_discon
T_UNITDATA	struct	t_unitdata
T_UDERROR	struct	t_uderr
T_INFO	struct	t_info

where each of these structures may subsequently be used as an argument to one or more transport functions.

Each of the above structures, except T_INFO, contains at least one field of type **struct netbuf**. For each field of this type, the user may specify that the buffer for that field should be allocated as well. The length of the buffer allocated will be equal to or greater than the appropriate size as returned in the *info* argument of *t_open()* or *t_getinfo()*. The relevant fields of the *info* argument are described in the following list. The *fields* argument specifies which buffers to allocate, where the argument is the bitwise-or of any of the following:

T_ADDR	The <i>addr</i> field of the t_bind , t_call , t_unitdata or t_uderr structures.
T_OPT	The <i>opt</i> field of the t_optmgmt , t_call , t_unitdata or t_uderr structures.
T_UDATA	The <i>udata</i> field of the t_call , t_discon or t_unitdata structures.
T_ALL	All relevant fields of the given structure. Fields which are not supported by the transport provider specified by <i>fd</i> will not be allocated.

For each relevant field specified in *fields*, *t_alloc()* will allocate memory for the buffer associated with the field, and initialise the *len* field to zero and the *buf* pointer and *maxlen* field accordingly. Irrelevant or unknown values passed in *fields* are ignored. Since the length of the buffer allocated will be based on the same size information that is returned to the user on a call to *t_open()* and *t_getinfo()*, *fd* must refer to the transport endpoint through which the newly allocated structure will be passed. In the case where a T_INFO structure is to be allocated, *fd* may be set to any value. In this way the appropriate size information can be accessed. If the size

value associated with any specified field is T_INVALID (see *t_open()* or *t_getinfo()*), *t_alloc()* will be unable to determine the size of the buffer to allocate and will fail, setting *t_errno* to [TSYSERR] and *errno* to [EINVAL]. If the size value associated with any specified field is T_INFINITE (see *t_open()* or *t_getinfo()*), then the behaviour of *t_alloc()* is implementation-defined. For any field not specified in *fields*, *buf* will be set to the null pointer and *len* and *maxlen* will be set to zero.

The pointer returned if the allocation succeeds is suitably aligned so that it can be assigned to a pointer to any type of object and then used to access such an object or array of such objects in the space allocated.

Use of *t_alloc()* to allocate structures will help ensure the compatibility of user programs with future releases of the transport interface functions.

VALID STATES

ALL - apart from T_UNINIT

ERRORS

On failure, *t_errno* is set to one of the following:

- | | |
|----------------|--|
| [TBADF] | struct_type is other than T_INFO and the specified file descriptor does not refer to a transport endpoint. |
| [TNOSTRUCTYPE] | Unsupported <i>struct_type</i> requested. This can include a request for a structure type which is inconsistent with the transport provider type specified, that is, connection-mode or connectionless-mode. |
| [TPROTO] | This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>). |
| [TSYSERR] | A system error has occurred during execution of this function. |

RETURN VALUE

On successful completion, *t_alloc()* returns a pointer to the newly allocated structure. On failure, a null pointer is returned.

SEE ALSO

t_free(), *t_getinfo()*, *t_open()*.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_bind - bind an address to a transport endpoint

SYNOPSIS

```
#include <xti.h>

int t_bind(
    int fd,
    const struct t_bind *req,
    struct t_bind *ret)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>req</i> → <i>addr.maxlen</i>	=	=
<i>req</i> → <i>addr.len</i>	x≥0	=
<i>req</i> → <i>addr.buf</i>	x (x)	=
<i>req</i> → <i>qlen</i>	x≥0	=
<i>ret</i> → <i>addr.maxlen</i>	x	=
<i>ret</i> → <i>addr.len</i>	/	x
<i>ret</i> → <i>addr.buf</i>	?	(?)
<i>ret</i> → <i>qlen</i>	/	x≥0

ret→*addr.buf* (the pointer itself, not the buffer it points to) is also unchanged.

This function associates a protocol address with the transport endpoint specified by *fd* and activates that transport endpoint. In connection mode, the transport provider may begin enqueueing incoming connect indications, or servicing a connection request on the transport endpoint. In connectionless-mode, the transport user may send or receive data units through the transport endpoint.

The *req* and *ret* arguments point to a **t_bind** structure containing the following members:

```
struct netbuf addr;
unsigned qlen;
```

The *addr* field of the **t_bind** structure specifies a protocol address, and the *qlen* field is used to indicate the maximum number of outstanding connection indications.

The parameter *req* is used to request that an address, represented by the **netbuf** structure, be bound to the given transport endpoint. The parameter *len* specifies the number of bytes in the address, and *buf* points to the address buffer. The parameter *maxlen* has no meaning for the *req* argument. On return, *ret* contains an encoding for the address that the transport provider actually bound to the transport endpoint; if an address was specified in *req*, this will be an encoding of the same address. In *ret*, the user specifies *maxlen*, which is the maximum size of the address buffer, and *buf* which points to the buffer where the address is to be placed. On return, *len* specifies the number of bytes in the bound address, and *buf* points to the bound address. If *maxlen* equals zero, no address is returned. If *maxlen* is greater than zero and less than the length of the address, *t_bind()* fails with *t_errno* set to [TBUFOVFLW].

If the requested address is not available, *t_bind()* will return -1 with *t_errno* set as appropriate. If no address is specified in *req* (the *len* field of *addr* in *req* is zero or *req* is NULL), the transport provider will assign an appropriate address to be bound, and will return that address in the *addr* field of *ret*. If the transport provider could not allocate an address, *t_bind()* will fail with *t_errno* set to [TNOADDR].

The parameter *req* may be a null pointer if the user does not wish to specify an address to be bound. Here, the value of *qlen* is assumed to be zero, and the transport provider will assign an address to the transport endpoint. Similarly, *ret* may be a null pointer if the user does not care what address was bound by the provider and is not interested in the negotiated value of *qlen*. It is valid to set *req* and *ret* to the null pointer for the same call, in which case the provider chooses the address to bind to the transport endpoint and does not return that information to the user.

The *qlen* field has meaning only when initialising a connection-mode service. It specifies the number of outstanding connection indications that the transport provider should support for the given transport endpoint. An outstanding connection indication is one that has been passed to the transport user by the transport provider but which has not been accepted or rejected. A value of *qlen* greater than zero is only meaningful when issued by a passive transport user that expects other users to call it. The value of *qlen* will be negotiated by the transport provider and may be changed if the transport provider cannot support the specified number of outstanding connection indications. However, this value of *qlen* will never be negotiated from a requested value greater than zero to zero. This is a requirement on transport providers; see **CAVEATS** below. On return, the *qlen* field in *ret* will contain the negotiated value.

If *fd* refers to a connection-mode service, this function allows more than one transport endpoint to be bound to the same protocol address (however, the transport provider must also support this capability), but it is not possible to bind more than one protocol address to the same transport endpoint. If a user binds more than one transport endpoint to the same protocol address, only one endpoint can be used to listen for connection indications associated with that protocol address. In other words, only one *t_bind()* for a given protocol address may specify a value of *qlen* greater than zero. In this way, the transport provider can identify which transport endpoint should be notified of an incoming connection indication. If a user attempts to bind a protocol address to a second transport endpoint with a value of *qlen* greater than zero, *t_bind()* will return -1 and set *t_errno* to [TADDRBUSY]. When a user accepts a connection on the transport endpoint that is being used as the listening endpoint, the bound protocol address will be found to be busy for the duration of the connection, until a *t_unbind()* or *t_close()* call has been issued. No other transport endpoints may be bound for listening on that same protocol address while that initial listening endpoint is active (in the data transfer phase or in the T_IDLE state). This will prevent more than one transport endpoint bound to the same protocol address from accepting connection indications.

If *fd* refers to connectionless mode service, this function allows for more than one transport endpoint to be associated with a protocol address, where the underlying transport provider supports this capability (often in conjunction with value of a protocol-specific option). If a user attempts to bind a second transport endpoint to an already bound protocol address when such capability is not supported for a transport provider, *t_bind()* will return -1 and set *t_errno* to [TADDRBUSY].

VALID STATES

T_UNBND

ERRORS

On failure, *t_errno* is set to one of the following:

[TACCES]	The user does not have permission to use the specified address.
[TADDRBUSY]	The requested address is in use.
[TBADADDR]	The specified protocol address was in an incorrect format or contained illegal information.
[TBADF]	The specified file descriptor does not refer to a transport endpoint.

[TBUFOVFLW]	The number of bytes allowed for an incoming argument (<i>maxlen</i>) is greater than 0 but not sufficient to store the value of that argument. The provider's state will change to T_IDLE and the information to be returned in <i>ret</i> will be discarded.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TNOADDR]	The transport provider could not allocate an address.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_accept(), *t_alloc()*, *t_close()*, *t_connect()*, *t_unbind()*.

CAVEATS

The requirement that the value of *qlen* never be negotiated from a requested value greater than zero to zero implies that transport providers, rather than the XTI implementation itself, accept this restriction.

An implementation need not allow an application explicitly to bind more than one communications endpoint to a single protocol address, while permitting more than one connection to be accepted to the same protocol address. That means that although an attempt to bind a communications endpoint to some address with *qlen=0* might be rejected with [TADDRBUSY], the user may nevertheless use this (unbound) endpoint as a responding endpoint in a call to *t_accept()*. To become independent of such implementation differences, the user should supply unbound responding endpoints to *t_accept()*.

The local address bound to an endpoint may change as result of a *t_accept()* or *t_connect()* call. Such changes are not necessarily reversed when the connection is released.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_close - close a transport endpoint

SYNOPSIS

#include <xti.h>

```
int t_close(
    int fd)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/

The *t_close()* function informs the transport provider that the user is finished with the transport endpoint specified by *fd*, and frees any local library resources associated with the endpoint. In addition, *t_close()* closes the file associated with the transport endpoint.

The function *t_close()* should be called from the T_UNBND state (see *t_getstate()*). However, this function does not check state information, so it may be called from any state to close a transport endpoint. If this occurs, the local library resources associated with the endpoint will be freed automatically. In addition, *close()* will be issued for that file descriptor; if there are no other descriptors in this process or in another process which references the communication endpoint, any connection that may be associated with that endpoint is broken. The connection may be terminated in an orderly or abortive manner.

A *t_close()* issued on a connection endpoint may cause data previously sent, or data not yet received, to be lost. It is the responsibility of the transport user to ensure that data is received by the remote peer.

For advice on how to write protocol-independent applications, see Section C.4 on page 245. For information on protocol-specific behaviour of *t_close()*, see the XTI Appendix for the relevant transport provider.

VALID STATES

ALL - apart from T_UNINIT

ERRORSOn failure, *t_errno* is set to the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO*t_getstate()*, *t_open()*, *t_unbind()*.**CHANGE HISTORY****Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_connect - establish a connection with another transport user

SYNOPSIS

```
#include <xti.h>

int t_connect(
    int fd,
    const struct t_call *sndcall,
    struct t_call *rcvcall)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>sndcall</i> → <i>addr.maxlen</i>	=	=
<i>sndcall</i> → <i>addr.len</i>	x	=
<i>sndcall</i> → <i>addr.buf</i>	x (x)	=
<i>sndcall</i> → <i>opt.maxlen</i>	=	=
<i>sndcall</i> → <i>opt.len</i>	x	=
<i>sndcall</i> → <i>opt.buf</i>	x (x)	=
<i>sndcall</i> → <i>udata.maxlen</i>	=	=
<i>sndcall</i> → <i>udata.len</i>	x	=
<i>sndcall</i> → <i>udata.buf</i>	? (?)	=
<i>sndcall</i> → <i>sequence</i>	=	=
<i>rcvcall</i> → <i>addr.maxlen</i>	x	/
<i>rcvcall</i> → <i>addr.len</i>	/	x
<i>rcvcall</i> → <i>addr.buf</i>	?	(?)
<i>rcvcall</i> → <i>opt.maxlen</i>	x	=
<i>rcvcall</i> → <i>opt.len</i>	/	x
<i>rcvcall</i> → <i>opt.buf</i>	?	(?)
<i>rcvcall</i> → <i>udata.maxlen</i>	x	=
<i>rcvcall</i> → <i>udata.len</i>	/	x
<i>rcvcall</i> → <i>udata.buf</i>	?	(?)
<i>rcvcall</i> → <i>sequence</i>	=	=

This function enables a transport user to request a connection to the specified destination transport user. This function can only be issued in the T_IDLE state. The parameter *fd* identifies the local transport endpoint where communication will be established, while *sndcall* and *rcvcall* point to a **t_call** structure which contains the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
int sequence;
```

The parameter *sndcall* specifies information needed by the transport provider to establish a connection and *rcvcall* specifies information that is associated with the newly established connection.

In *sndcall*, *addr* specifies the protocol address of the destination transport user, *opt* presents any protocol-specific information that might be needed by the transport provider, *udata* points to optional user data that may be passed to the destination transport user during connection establishment, and *sequence* has no meaning for this function.

On return, in *rcvcall*, *addr* contains the protocol address associated with the responding transport endpoint, *opt* represents any protocol-specific information associated with the connection, *udata* points to optional user data that may be returned by the destination transport user during connection establishment, and *sequence* has no meaning for this function.

The *opt* argument permits users to define the options that may be passed to the transport provider. These options are specific to underlying protocol of the transport provider or XTI interface and are described in Appendix F (for the XTI interface) and other protocol-specific appendices which are part of this specification. The user may choose not to negotiate protocol options by setting the *len* field of *opt* to zero. In this case, the provider uses the option values currently set for the communications endpoint.

If used, *sndcall*→*opt.buf* must point to a buffer with the corresponding options, and *sndcall*→*opt.len* must specify its length. The *maxlen* and *buf* fields of the **netbuf** structure pointed by *rcvcall*→*addr* and *rcvcall*→*opt* must be set before the call.

The *udata* argument enables the caller to pass user data to the destination transport user and receive user data from the destination user during connection establishment. However, the amount of user data must not exceed the limits supported by the transport provider as returned in the *connect* field of the *info* argument of *t_open()* or *t_getinfo()*. If the *len* of *udata* is zero in *sndcall*, no data will be sent to the destination transport user.

On return, the *addr*, *opt* and *udata* fields of *rcvcall* will be updated to reflect values associated with the connection. Thus, the *maxlen* field of each argument must be set before issuing this function to indicate the maximum size of the buffer for each. However, *maxlen* can be set to zero, in which case no information to this specific argument is given to the user on the return from *t_connect()*. If *maxlen* is greater than zero and less than the length of the value, *t_connect()* fails with *t_errno* set to [TBUFOVFLW]. If *rcvcall* is set to NULL, no information at all is returned.

By default, *t_connect()* executes in synchronous mode, and will wait for the destination user's response before returning control to the local user. A successful return (that is, return value of zero) indicates that the requested connection has been established. However, if O_NONBLOCK is set (via *t_open()* or *fcntl()*), *t_connect()* executes in asynchronous mode. In this case, the call will not wait for the remote user's response, but will return control immediately to the local user and return -1 with *t_errno* set to [TNODATA] to indicate that the connection has not yet been established. In this way, the function simply initiates the connection establishment procedure by sending a connection request to the destination transport user. The *t_rcvconnect()* function is used in conjunction with *t_connect()* to determine the status of the requested connection.

When a synchronous *t_connect()* call is interrupted by the arrival of a signal, the state of the corresponding transport endpoint is T_OUTCON, allowing a further call to either *t_rcvconnect()*, *t_rcvdis()* or *t_snddis()*. When an asynchronous *t_connect()* call is interrupted by the arrival of a signal, the state of the corresponding transport endpoint is T_IDLE.

VALID STATES

T_IDLE

ERRORS

On failure, *t_errno* is set to one of the following:

[TACCES]	The user does not have permission to use the specified address or options.
[TADDRBUSY]	This transport provider does not support multiple connections with the same local and remote addresses. This error indicates that a connection already exists.

[TBADADDR]	The specified protocol address was in an incorrect format or contained illegal information.
[TBADDATA]	The amount of user data specified was not within the bounds allowed by the transport provider.
[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBADOPT]	The specified protocol options were in an incorrect format or contained illegal information.
[TBUFOVFLW]	The number of bytes allocated for an incoming argument (<i>maxlen</i>) is greater than 0 but not sufficient to store the value of that argument. If executed in synchronous mode, the provider's state, as seen by the user, changes to T_DATAXFER, and the information to be returned in <i>rcvcall</i> is discarded.
[TLOOK]	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
[TNODATA]	O_NONBLOCK was set, so the function successfully initiated the connection establishment procedure, but did not wait for a response from the remote user.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_accept(), *t_alloc()*, *t_getinfo()*, *t_listen()*, *t_open()*, *t_optmgmt()*, *t_rcvconnect()*.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_error - produce error message

SYNOPSIS

```
#include <xti.h>
```

```
int t_error(
    const char *errmsg)
```

DESCRIPTION

Parameters	Before call	After call
<i>errmsg</i>	x	=

The *t_error()* function produces a message on the standard error output which describes the last error encountered during a call to a transport function. The argument string *errmsg* is a user-supplied error message that gives context to the error.

The error message is written as follows: first (if *errmsg* is not a null pointer and the character pointed to by *errmsg* is not the null character) the string pointed to by *errmsg* followed by a colon and a space; then a standard error message string for the current error defined in *t_errno*. If *t_errno* has a value different from [TSYSERR], the standard error message string is followed by a newline character. If, however, *t_errno* is equal to [TSYSERR], the *t_errno* string is followed by the standard error message string for the current error defined in *errno* followed by a newline.

The language for error message strings written by *t_error()* is that of the current locale. If it is English, the error message string describing the value in *t_errno* may be derived from the comments following the *t_errno* codes defined in *xti.h*. The contents of the error message strings describing the value in *errno* are the same as those returned by the *strerror(3C)* function with an argument of *errno*.

The error number, *t_errno*, is only set when an error occurs and it is not cleared on successful calls.

EXAMPLE

If a *t_connect()* function fails on transport endpoint *fd2* because a bad address was given, the following call might follow the failure:

```
t_error("t_connect failed on fd2");
```

The diagnostic message to be printed would look like:

```
t_connect failed on fd2: incorrect addr format
```

where *incorrect addr format* identifies the specific error that occurred, and *t_connect failed on fd2* tells the user which function failed on which transport endpoint.

VALID STATES

All - apart from T_UNINIT

ERRORS

No errors are defined for the *t_error()* function.

RETURN VALUE

Upon completion, a value of 0 is returned.

CHANGE HISTORY

Issue 4

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_errno - XTI error return value

SYNOPSIS

```
#include <xti.h>
```

DESCRIPTION

t_errno is used by XTI functions to return error values.

XTI functions provide an error number in *t_errno* which has type *int* and is defined in **<xti.h>**. The value of *t_errno* will be defined only after a call to a XTI function for which it is explicitly stated to be set and until it is changed by the next XTI function call. The value of *t_errno* should only be examined when it is indicated to be valid by a function's return value. Programs should obtain the definition of *t_errno* by the inclusion of **<xti.h>**. The practice of defining *t_errno* in program as *extern* in *t_errno* is obsolescent. No XTI function sets *t_errno* to 0 to indicate an error.

It is unspecified whether *t_errno* is a macro or an identifier with external linkage. It represents a modifiable **lvalue** of type *int*. If a macro definition is suppressed in order to access an actual object or a program defines an identifier with name *t_errno*, the behavior is undefined.

The symbolic values stored in *t_errno* by an XTI function are defined in the **ERRORS** sections in all relevant XTI function definition pages.

NAME

t_free - free a library structure

SYNOPSIS

#include <xti.h>

```
int t_free(
    void *ptr,
    int struct_type)
```

DESCRIPTION

Parameters	Before call	After call
<i>ptr</i>	x	/
<i>struct_type</i>	x	/

The *t_free()* function frees memory previously allocated by *t_alloc()*. This function will free memory for the specified structure, and will also free memory for buffers referenced by the structure.

The argument *ptr* points to one of the seven structure types described for *t_alloc()*, and *struct_type* identifies the type of that structure which must be one of the following:

T_BIND	struct	t_bind
T_CALL	struct	t_call
T_OPTMGMT	struct	t_optmgmt
T_DIS	struct	t_discon
T_UNITDATA	struct	t_unitdata
T_UDERROR	struct	t_uderr
T_INFO	struct	t_info

where each of these structures is used as an argument to one or more transport functions.

The function *t_free()* will check the *addr*, *opt* and *udata* fields of the given structure (as appropriate) and free the buffers pointed to by the *buf* field of the **netbuf** structure. If *buf* is a null pointer, *t_free()* will not attempt to free memory. After all buffers are freed, *t_free()* will free the memory associated with the structure pointed to by *ptr*.

Undefined results will occur if *ptr* or any of the *buf* pointers points to a block of memory that was not previously allocated by *t_alloc()*.

VALID STATES

ALL - apart from T_UNINIT

ERRORSOn failure, *t_errno* is set to the following:

[TNOSTRUCTYPE]	Unsupported <i>struct_type</i> requested.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_alloc().

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_getinfo - get protocol-specific service information

SYNOPSIS

```
#include <xti.h>

int t_getinfo(
    int fd,
    struct t_info *info)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>info</i> → <i>addr</i>	/	x
<i>info</i> → <i>options</i>	/	x
<i>info</i> → <i>tsdu</i>	/	x
<i>info</i> → <i>etsdu</i>	/	x
<i>info</i> → <i>connect</i>	/	x
<i>info</i> → <i>discon</i>	/	x
<i>info</i> → <i>servtype</i>	/	x
<i>info</i> → <i>flags</i>	/	x

This function returns the current characteristics of the underlying transport protocol and/or transport connection associated with file descriptor *fd*. The *info* pointer is used to return the same information returned by *t_open()*, although not necessarily precisely the same values. This function enables a transport user to access this information during any phase of communication.

This argument points to a **t_info** structure which contains the following members:

```
t_scalar_t addr;      /*max size in octets of the transport protocol address*/
t_scalar_t options;  /*max number of bytes of protocol-specific options */
t_scalar_t tsdu;     /*max size in octets of a transport service data unit */
t_scalar_t etsdu;    /*max size in octets of an expedited transport service*/
                    /*data unit (ETSDU) */
t_scalar_t connect;  /*max number of octets allowed on connection */
                    /*establishment functions */
t_scalar_t discon;   /*max number of octets of data allowed on t_snddis() */
                    /*and t_rcvdis() functions */
t_scalar_t servtype; /*service type supported by the transport provider */
t_scalar_t flags;    /*other info about the transport provider */
```

The values of the fields have the following meanings:

- addr** A value greater than zero indicates the maximum size of a transport protocol address and a value of T_INVALID (-2) specifies that the transport provider does not provide user access to transport protocol addresses.
- options** A value greater than zero indicates the maximum number of bytes of protocol-specific options supported by the provider, and a value of T_INVALID (-2) specifies that the transport provider does not support user-settable options.
- tsdu** A value greater than zero specifies the maximum size in octets of a transport service data unit (TSDU); a value of T_NULL (zero) specifies that the transport provider does not support the concept of TSDU, although it does support the sending of a datastream with no logical

boundaries preserved across a connection; a value of T_INFINITE (-1) specifies that there is no limit on the size in octets of a TSDU; and a value of T_INVALID (-2) specifies that the transfer of normal data is not supported by the transport provider.

<i>etsdu</i>	A value greater than zero specifies the maximum size in octets of an expedited transport service data unit (ETSDU); a value of T_NULL (zero) specifies that the transport provider does not support the concept of ETSDU, although it does support the sending of an expedited data stream with no logical boundaries preserved across a connection; a value of T_INFINITE (-1) specifies that there is no limit on the size (in octets) of an ETSDU; and a value of T_INVALID (-2) specifies that the transfer of expedited data is not supported by the transport provider. Note that the semantics of expedited data may be quite different for different transport providers (see Appendix A on page 215 and Appendix B on page 229).
<i>connect</i>	A value greater than zero specifies the maximum number of octets that may be associated with connection establishment functions and a value of T_INVALID (-2) specifies that the transport provider does not allow data to be sent with connection establishment functions.
<i>discon</i>	If the T_ORDRELDATA bit in flags is clear, a value greater than zero specifies the maximum number of octets that may be associated with the <i>t_snddis()</i> and <i>t_rcvdis()</i> functions, and a value of T_INVALID (-2) specifies that the transport provider does not allow data to be sent with the abortive release functions. If the T_ORDRELDATA bit is set in flags, a value greater than zero specifies the maximum number of octets that may be associated with the <i>t_sndreldata()</i> , <i>t_rcvreldata()</i> , <i>t_snddis()</i> and <i>t_rcvdis()</i> functions.
<i>servtype</i>	This field specifies the service type supported by the transport provider, as described below.
<i>flags</i>	This is a bit field used to specify other information about the communications provider. If the T_ORDRELDATA bit is set, the communications provider supports sending user data with an orderly release. If the T_SENDZERO bit is set in flags, this indicates that the underlying transport provider supports the sending of zero-length TSDUs. See Appendix A on page 215 for a discussion of the separate issue of zero-length fragments within a TSDU.

If a transport user is concerned with protocol independence, the above sizes may be accessed to determine how large the buffers must be to hold each piece of information. Alternatively, the *t_alloc()* function may be used to allocate these buffers. An error will result if a transport user exceeds the allowed data size on any function. The value of each field may change as a result of protocol option negotiation during connection establishment (the *t_optmgmt()* call has no effect on the values returned by *t_getinfo()*). These values will only change from the values presented to *t_open()* after the endpoint enters the T_DATAXFER state.

The *servtype* field of *info* specifies one of the following values on return:

T_COTS	The transport provider supports a connection-mode service but does not support the optional orderly release facility.
T_COTS_ORD	The transport provider supports a connection-mode service with the optional orderly release facility.

T_CLTS The transport provider supports a connectionless-mode service. For this service type, *t_open()* will return T_INVALID (-2) for *etsdu*, *connect* and *discon*.

VALID STATES

ALL - apart from T_UNINIT

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADF] The specified file descriptor does not refer to a transport endpoint.

[TPROTO] This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (*t_errno*).

[TSYSERR] A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_alloc(), *t_open()*.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_getprotaddr - get the protocol addresses

SYNOPSIS

```
#include <xti.h>

int t_getprotaddr(
    int fd,
    struct t_bind *boundaddr,
    struct t_bind *peeraddr)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>boundaddr</i> → <i>addr.maxlen</i>	x	=
<i>boundaddr</i> → <i>addr.len</i>	/	x
<i>boundaddr</i> → <i>addr.buf</i>	?	(?)
<i>boundaddr</i> → <i>qlen</i>	=	=
<i>peeraddr</i> → <i>addr.maxlen</i>	x	=
<i>peeraddr</i> → <i>addr.len</i>	/	x
<i>peeraddr</i> → <i>addr.buf</i>	?	(?)
<i>peeraddr</i> → <i>qlen</i>	=	=

The *t_getprotaddr()* function returns local and remote protocol addresses currently associated with the transport endpoint specified by *fd*. In *boundaddr* and *peeraddr* the user specifies *maxlen*, which is the maximum size (in bytes) of the address buffer, and *buf* which points to the buffer where the address is to be placed. On return, the *buf* field of *boundaddr* points to the address, if any, currently bound to *fd*, and the *len* field specifies the length of the address. If the transport endpoint is in the T_UNBND state, zero is returned in the *len* field of *boundaddr*. The *buf* field of *peeraddr* points to the address, if any, currently connected to *fd*, and the *len* field specifies the length of the address. If the transport endpoint is not in the T_DATAXFER, T_INREL, T_OUTCON or T_OUTREL states, zero is returned in the *len* field of *peeraddr*. If the *maxlen* field of *boundaddr* or *peeraddr* is set to zero, no address is returned.

VALID STATES

ALL - apart from T_UNINIT

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBUFOVFLW]	The number of bytes allocated for an incoming argument (<i>maxlen</i>) is greater than 0 but not sufficient to store the value of that argument.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of zero is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate the error.

SEE ALSO

t_bind().

CHANGE HISTORY

Issue 4

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_getstate - get the current state

SYNOPSIS

```
#include <xti.h>

int t_getstate(
    int fd)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/

The *t_getstate()* function returns the current state of the provider associated with the transport endpoint specified by *fd*.

VALID STATES

ALL - apart from T_UNINIT

ERRORSOn failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSTATECHNG]	The transport provider is undergoing a transient state change.
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

State is returned upon successful completion. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error. The current state is one of the following:

T_UNBND	Unbound.
T_IDLE	Idle.
T_OUTCON	Outgoing connection pending.
T_INCON	Incoming connection pending.
T_DATAXFER	Data transfer.
T_OUTREL	Outgoing direction orderly release sent.
T_INREL	Incoming direction orderly release received.

If the provider is undergoing a state transition when *t_getstate()* is called, the function will fail.

SEE ALSO*t_open()*.**CHANGE HISTORY****Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_listen - listen for a connection indication

SYNOPSIS

```
#include <xti.h>

int t_listen(
    int fd,
    struct t_call *call)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>call</i> → <i>addr.maxlen</i>	x	=
<i>call</i> → <i>addr.len</i>	/	x
<i>call</i> → <i>addr.buf</i>	?	(?)
<i>call</i> → <i>opt.maxlen</i>	x	=
<i>call</i> → <i>opt.len</i>	/	x
<i>call</i> → <i>opt.buf</i>	?	(?)
<i>call</i> → <i>udata.maxlen</i>	x	=
<i>call</i> → <i>udata.len</i>	/	x
<i>call</i> → <i>udata.buf</i>	?	(?)
<i>call</i> → <i>sequence</i>	/	x

This function listens for a connection indication from a calling transport user. The argument *fd* identifies the local transport endpoint where connection indications arrive, and on return, *call* contains information describing the connection indication. The parameter *call* points to a **t_call** structure which contains the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
int sequence;
```

In *call*, *addr* returns the protocol address of the calling transport user. This address is in a format usable in future calls to *t_connect()*. Note, however that *t_connect()* may fail for other reasons, for example [TADDRBUSY]. *opt* returns options associated with the connection indication, *udata* returns any user data sent by the caller on the connection request, and *sequence* is a number that uniquely identifies the returned connection indication. The value of *sequence* enables the user to listen for multiple connection indications before responding to any of them.

Since this function returns values for the *addr*, *opt* and *udata* fields of *call*, the *maxlen* field of each must be set before issuing the *t_listen()* to indicate the maximum size of the buffer for each. If the *maxlen* field of *call*→*addr*, *call*→*opt* or *call*→*udata* is set to zero, no information is returned for this parameter.

By default, *t_listen()* executes in synchronous mode and waits for a connection indication to arrive before returning to the user. However, if O_NONBLOCK is set via *t_open()* or *fcntl()*, *t_listen()* executes asynchronously, reducing to a poll for existing connection indications. If none are available, it returns -1 and sets *t_errno* to [TNODATA].

VALID STATES

T_IDLE, T_INCON

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBADQLEN]	The argument <i>qlen</i> of the endpoint referenced by <i>fd</i> is zero.
[TBUFOVFLW]	The number of bytes allocated for an incoming argument (<i>maxlen</i>) is greater than 0 but not sufficient to store the value of that argument. The provider's state, as seen by the user, changes to T_INCON, and the connection indication information to be returned in <i>call</i> is discarded. The value of <i>sequence</i> returned can be used to do a <i>t_snddis()</i> .
[TLOOK]	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
[TNODATA]	O_NONBLOCK was set, but no connection indications had been queued.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TQFULL]	The maximum number of outstanding connection indications has been reached for the endpoint referenced by <i>fd</i> . Note that a subsequent call to <i>t_listen()</i> may block until another incoming connection indication is available. This can only occur if at least one of the outstanding connection indications becomes no longer outstanding, for example through a call to <i>t_accept()</i> .
[TSYSERR]	A system error has occurred during execution of this function.

CAVEATS

Some transport providers do not differentiate between a connection indication and the connection itself. If this is the case, a successful return of *t_listen()* indicates an existing connection (see Appendix B on page 229).

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

fcntl(), *t_accept()*, *t_alloc()*, *t_bind()*, *t_connect()*, *t_open()*, *t_optmgmt()*, *t_rcvconnect()*.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_look - look at the current event on a transport endpoint

SYNOPSIS

```
#include <xti.h>
```

```
int t_look(
    int fd)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/

This function returns the current event on the transport endpoint specified by *fd*. This function enables a transport provider to notify a transport user of an asynchronous event when the user is calling functions in synchronous mode. Certain events require immediate notification of the user and are indicated by a specific error, [TLOOK], on the current or next function to be executed. Details on events which cause functions to fail [TLOOK] may be found in Section 5.6 on page 38.

This function also enables a transport user to poll a transport endpoint periodically for asynchronous events.

VALID STATES

ALL - apart from T_UNINIT

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon success, *t_look()* returns a value that indicates which of the allowable events has occurred, or returns zero if no event exists. One of the following events is returned:

T_LISTEN	Connection indication received.
T_CONNECT	Connect confirmation received.
T_DATA	Normal data received.
T_EXDATA	Expedited data received.
T_DISCONNECT	Disconnection received.
T_UDERR	Datagram error indication.
T_ORDREL	Orderly release indication.
T_GODATA	Flow control restrictions on normal data flow that led to a [TFLOW] error have been lifted. Normal data may be sent again.
T_GOEXDATA	Flow control restrictions on expedited data flow that led to a [TFLOW] error have been lifted. Expedited data may be sent again.

On failure, -1 is returned and *t_errno* is set to indicate the error.

SEE ALSO

t_open(), *t_snd()*, *t_sndudata()*.

APPLICATION USAGE

Additional functionality is provided through the Event Management (EM) interface.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_open - establish a transport endpoint

SYNOPSIS

```
#include <xti.h> #include <fcntl.h>
```

```
int t_open(
    const char *name,
    int oflag,
    struct t_info *info)
```

DESCRIPTION

Parameters	Before call	After call
<i>name</i>	x	=
<i>oflag</i>	x	=
<i>info</i> → <i>addr</i>	/	x
<i>info</i> → <i>options</i>	/	x
<i>info</i> → <i>tsdu</i>	/	x
<i>info</i> → <i>etsdu</i>	/	x
<i>info</i> → <i>connect</i>	/	x
<i>info</i> → <i>discon</i>	/	x
<i>info</i> → <i>servtype</i>	/	x
<i>info</i> → <i>flags</i>	/	x

The *t_open()* function must be called as the first step in the initialisation of a transport endpoint. This function establishes a transport endpoint by supplying a transport provider identifier that indicates a particular transport provider (that is, transport protocol) and returning a file descriptor that identifies that endpoint.

The argument *name* points to a transport provider identifier and *oflag* identifies any open flags (as in *open()*). The argument *oflag* is constructed from O_RDWR optionally bitwise inclusive-OR'ed with O_NONBLOCK. These flags are defined by the header <fcntl.h>. The file descriptor returned by *t_open()* will be used by all subsequent functions to identify the particular local transport endpoint.

This function also returns various default characteristics of the underlying transport protocol by setting fields in the **t_info** structure. This argument points to a **t_info** which contains the following members:

```
t_scalar_t addr;      /* max size of the transport protocol address */
t_scalar_t options;  /* max number of bytes of
/* protocol-specific options
t_scalar_t tsdu;     /* max size of a transport service data
/* unit (TSDU)
t_scalar_t etsdu;   /* max size of an expedited transport
/* service data unit (ETSDU)
t_scalar_t connect; /* max amount of data allowed on
/* connection establishment functions
t_scalar_t discon;  /* max amount of data allowed on
/* t_snddis() and t_rcvdis() functions
t_scalar_t servtype; /* service type supported by the
/* transport provider
t_scalar_t flags;   /* other info about the transport provider
```

The values of the fields have the following meanings:

<i>addr</i>	A value greater than zero (T_NULL) indicates the maximum size of a transport protocol address and a value of -2 (T_INVALID) specifies that the transport provider does not provide user access to transport protocol addresses.
<i>options</i>	A value greater than zero (T_NULL) indicates the maximum number of bytes of protocol-specific options supported by the provider, and a value of -2 (T_INVALID) specifies that the transport provider does not support user-settable options.
<i>tsdu</i>	A value greater than zero (T_NULL) specifies the maximum size of a transport service data unit (TSDU); a value of zero (T_NULL) specifies that the transport provider does not support the concept of TSDU, although it does support the sending of a data stream with no logical boundaries preserved across a connection; a value of -1 (T_INFINITE) specifies that there is no limit to the size of a TSDU; and a value of -2 (T_INVALID) specifies that the transfer of normal data is not supported by the transport provider.
<i>etsdu</i>	A value greater than zero (T_NULL) specifies the maximum size of an expedited transport service data unit (ETSDU); a value of zero (T_NULL) specifies that the transport provider does not support the concept of ETSDU, although it does support the sending of an expedited data stream with no logical boundaries preserved across a connection; a value of -1 (T_INFINITE) specifies that there is no limit on the size of an ETSDU; and a value of -2 (T_INVALID) specifies that the transfer of expedited data is not supported by the transport provider. Note that the semantics of expedited data may be quite different for different transport providers (see Appendix A on page 215 and Appendix B on page 229).
<i>connect</i>	A value greater than zero (T_NULL) specifies the maximum amount of data that may be associated with connection establishment functions, and a value of -2 (T_INVALID) specifies that the transport provider does not allow data to be sent with connection establishment functions.
<i>discon</i>	If the T_ORDRELDATA bit in flags is clear, a value greater than zero (T_NULL) specifies the maximum amount of data that may be associated with the <i>t_snddis()</i> and <i>t_rcvdis()</i> functions, and a value of -2 (T_INVALID) specifies that the transport provider does not allow data to be sent with the abortive release functions. If the T_ORDRELDATA bit is set in flags, a value greater than zero (T_NULL) specifies the maximum number of octets that may be associated with the <i>t_sndreldata()</i> , <i>t_rcvreldata()</i> , <i>t_snddis()</i> and <i>t_rcvdis()</i> functions.
<i>servtype</i>	This field specifies the service type supported by the transport provider, as described below.
<i>flags</i>	This is a bit field used to specify other information about the communications provider. If the T_ORDRELDATA bit is set, the communications provider supports user data to be sent with an orderly release. If the T_SENDZERO bit is set in flags, this indicates the underlying transport provider supports the sending of zero-length TSDUs. See Appendix A on page 215 for a discussion of the separate issue of zero-length fragments within a TSDU.

If a transport user is concerned with protocol independence, the above sizes may be accessed to determine how large the buffers must be to hold each piece of information. Alternatively, the

t_alloc() function may be used to allocate these buffers. An error will result if a transport user exceeds the allowed data size on any function.

The *servtype* field of *info* specifies one of the following values on return:

T_COTS	The transport provider supports a connection-mode service but does not support the optional orderly release facility.
T_COTS_ORD	The transport provider supports a connection-mode service with the optional orderly release facility.
T_CLTS	The transport provider supports a connectionless-mode service. For this service type, <i>t_open()</i> will return -2 (T_INVALID) for <i>etsdu</i> , <i>connect</i> and <i>discon</i> .

A single transport endpoint may support only one of the above services at one time.

If *info* is set to a null pointer by the transport user, no protocol information is returned by *t_open()*.

VALID STATES

T_UNINIT

ERRORS

On failure, *t_errno* is set to the following:

[TBADFLAG]	An invalid flag is specified.
[TBADNAME]	Invalid transport provider name.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUES

A valid file descriptor is returned upon successful completion. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

open().

CHANGE HISTORY

Issue 4

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_optmgmt - manage options for a transport endpoint

SYNOPSIS

```
#include <xti.h>

int t_optmgmt(
    int fd,
    const struct t_optmgmt *req,
    struct t_optmgmt *ret)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>req</i> → <i>opt.maxlen</i>	=	=
<i>req</i> → <i>opt.len</i>	x	=
<i>req</i> → <i>opt.buf</i>	x (x)	=
<i>req</i> → <i>flags</i>	x	=
<i>ret</i> → <i>opt.maxlen</i>	x	=
<i>ret</i> → <i>opt.len</i>	/	x
<i>ret</i> → <i>opt.buf</i>	?	(?)
<i>ret</i> → <i>flags</i>	/	x

The *t_optmgmt()* function enables a transport user to retrieve, verify or negotiate protocol options with the transport provider. The argument *fd* identifies a transport endpoint.

The *req* and *ret* arguments point to a **t_optmgmt** structure containing the following members:

```
struct netbuf opt;
t_scalar_t flags;
```

The *opt* field identifies protocol options and the *flags* field is used to specify the action to take with those options.

The options are represented by a **netbuf** structure in a manner similar to the address in *t_bind()*. The argument *req* is used to request a specific action of the provider and to send options to the provider. The argument *len* specifies the number of bytes in the options, *buf* points to the options buffer, and *maxlen* has no meaning for the *req* argument. The transport provider may return options and flag values to the user through *ret*. For *ret*, *maxlen* specifies the maximum size of the options buffer and *buf* points to the buffer where the options are to be placed. If *maxlen* in *ret* is set to zero, no options values are returned. On return, *len* specifies the number of bytes of options returned. The value in *maxlen* has no meaning for the *req* argument, but must be set in the *ret* argument to specify the maximum number of bytes the options buffer can hold.

Each option in the options buffer is of the form **struct t_opthdr** possibly followed by an option value.

The *level* field of **struct t_opthdr** identifies the XTI level or a protocol of the transport provider. The *name* field identifies the option within the level, and *len* contains its total length; that is, the length of the option header **t_opthdr** plus the length of the option value. If *t_optmgmt()* is called with the action T_NEGOTIATE set, the *status* field of the returned options contains information about the success or failure of a negotiation.

Several options can be concatenated. The option user has, however to ensure that each options header and value part starts at a boundary appropriate for the architecture-specific alignment rules. The macros T_OPT_FIRSTHDR(nbp), T_OPT_NXTHDR(nbp, tohp), T_OPT_DATA(tohp)

are provided for that purpose.

- T_OPT_DATA** If argument is a pointer to a **t_opthdr** structure, this macro returns an unsigned character pointer to the data associated with the **t_opthdr**.
- T_OPT_NXTHDR** If the first argument is a pointer to a netbuf structure associated with an option buffer and second argument is a pointer to a **t_opthdr** structure within that option buffer, this macro returns a pointer to the next **t_opthdr** structure or a null pointer if this **t_opthdr** is the last **t_opthdr** in the option buffer.
- T_OPT_FIRSTHDR** If the argument is a pointer to a **netbuf** structure associated with an option buffer, this macro returns the pointer to the first **t_opthdr** structure in the associated option buffer, or a null pointer if there is no option buffer associated with this **netbuf** or if it is not possible or the associated option buffer is too small to accommodate even the first aligned option header.

T_OPT_FIRSTHDR is useful for finding an appropriately aligned start of the option buffer. **T_OPT_NEXTHDR** is useful for moving to the start of the next appropriately aligned option in the option buffer. Note that **OPT_NEXTHDR** is also available for backward compatibility requirements. **T_OPT_DATA** is useful for finding the start of the data part in the option buffer where the contents of its values start on an appropriately aligned boundary.

If the transport user specifies several options on input, all options must address the same level.

If any option in the options buffer does not indicate the same level as the first option, or the level specified is unsupported, then the *t_optmgmt()* request will fail with [TBADOPT]. If the error is detected, some options have possibly been successfully negotiated. The transport user can check the current status by calling *t_optmgmt()* with the **T_CURRENT** flag set.

Chapter 6 contains a detailed description about the use of options and should be read before using this function.

The *flags* field of *req* must specify one of the following actions:

- T_NEGOTIATE** This action enables the transport user to negotiate option values.
- The user specifies the options of interest and their values in the buffer specified by *req->opt.buf* and *req->opt.len*. The negotiated option values are returned in the buffer pointed to by *ret->opt.buf*. The *status* field of each returned option is set to indicate the result of the negotiation. The value is **T_SUCCESS** if the proposed value was negotiated, **T_PARTSUCCESS** if a degraded value was negotiated, **T_FAILURE** if the negotiation failed (according to the negotiation rules), **T_NOTSUPPORT** if the transport provider does not support this option or illegally requests negotiation of a privileged option, and **T_READONLY** if modification of a read-only option was requested. If the status is **T_SUCCESS**, **T_FAILURE**, **T_NOTSUPPORT** or **T_READONLY**, the returned option value is the same as the one requested on input.

The overall result of the negotiation is returned in *ret->flags*.

This field contains the worst single result, whereby the rating is done according to the order **T_NOTSUPPORT**, **T_READONLY**, **T_FAILURE**, **T_PARTSUCCESS**, **T_SUCCESS**. The value **T_NOTSUPPORT** is the worst result and **T_SUCCESS** is the best.

For each level, the option T_ALLOPT (see below) can be requested on input. No value is given with this option; only the **t_opthdr** part is specified. This input requests to negotiate all supported options of this level to their default values. The result is returned option by option in *ret→opt.buf*. (Note that depending on the state of the transport endpoint, not all requests to negotiate the default value may be successful.)

T_CHECK

This action enables the user to verify whether the options specified in *req* are supported by the transport provider.

If an option is specified with no option value (it consists only of a **t_opthdr** structure), the option is returned with its *status* field set to T_SUCCESS if it is supported, T_NOTSUPPORT if it is not or needs additional user privileges, and T_READONLY if it is read-only (in the current XTI state). No option value is returned.

If an option is specified with an option value, the *status* field of the returned option has the same value, as if the user had tried to negotiate this value with T_NEGOTIATE. If the status is T_SUCCESS, T_FAILURE, T_NOTSUPPORT or T_READONLY, the returned option value is the same as the one requested on input.

The overall result of the option checks is returned in *ret→flags*. This field contains the worst single result of the option checks, whereby the rating is the same as for T_NEGOTIATE.

Note that no negotiation takes place. All currently effective option values remain unchanged.

T_DEFAULT

This action enables the transport user to retrieve the default option values. The user specifies the options of interest in *req→opt.buf*. The option values are irrelevant and will be ignored; it is sufficient to specify the **t_opthdr** part of an option only. The default values are then returned in *ret→opt.buf*.

The *status* field returned is T_NOTSUPPORT if the protocol level does not support this option or the transport user illegally requested a privileged option, T_READONLY if the option is read-only, and set to T_SUCCESS in all other cases. The overall result of the request is returned in *ret→flags*. This field contains the worst single result, whereby the rating is the same as for T_NEGOTIATE.

For each level, the option T_ALLOPT (see below) can be requested on input. All supported options of this level with their default values are then returned. In this case, *ret→opt.maxlen* must be given at least the value *info→options* (see *t_getinfo()*, *t_open()*) before the call.

T_CURRENT

This action enables the transport user to retrieve the currently effective option values. The user specifies the options of interest in *req→opt.buf*. The option values are irrelevant and will be ignored; it is sufficient to specify the **t_opthdr** part of an option only. The currently effective values are then returned in *ret→opt.buf*.

The *status* field returned is T_NOTSUPPORT if the protocol level does not support this option or the transport user illegally requested a privileged option, T_READONLY if the option is read-only, and set to T_SUCCESS in all other cases. The overall result of the request is returned in

ret→*flags*. This field contains the worst single result, whereby the rating is the same as for T_NEGOTIATE.

For each level, the option T_ALLOPT (see below) can be requested on input. All supported options of this level with their currently effective values are then returned.

The option T_ALLOPT can only be used with *t_optmgmt()* and the actions T_NEGOTIATE, T_DEFAULT and T_CURRENT. It can be used with any supported level and addresses all supported options of this level. The option has no value; it consists of a **t_opthdr** only. Since in a *t_optmgmt()* call only options of one level may be addressed, this option should not be requested together with other options. The function returns as soon as this option has been processed.

Options are independently processed in the order they appear in the input option buffer. If an option is multiply input, it depends on the implementation whether it is multiply output or whether it is returned only once.

Transport providers may not be able to provide an interface capable of supporting T_NEGOTIATE and/or T_CHECK functionalities. When this is the case, the error [TNOTSUPPORT] is returned.

The function *t_optmgmt()* may block under various circumstances and depending on the implementation. The function will block, for instance, if the protocol addressed by the call resides on a separate controller. It may also block due to flow control constraints; that is, if data sent previously across this transport endpoint has not yet been fully processed. If the function is interrupted by a signal, the option negotiations that have been done so far may remain valid. The behaviour of the function is not changed if O_NONBLOCK is set.

XTI-LEVEL OPTIONS

XTI-level options are not specific for a particular transport provider. An XTI implementation supports none, all or any subset of the options defined below. An implementation may restrict the use of any of these options by offering them only in the privileged or read-only mode, or if *fd* relates to specific transport providers.

The subsequent options do not have end-to-end significance (see Chapter 6). They may be negotiated in all XTI states except T_UNINIT.

The protocol level is XTI_GENERIC. For this level, the following options are defined:

option name	type of option value	legal option value	meaning
XTI_DEBUG XTI_LINGER	array of t_uscalar_t struct t_linger	see text see text	enable debugging linger on close if data is present
XTI_RCVBUF	t_uscalar_t	size in octets	receive buffer size
XTI_RCVLOWAT	t_uscalar_t	size in octets	receive low-water mark
XTI_SNDBUF	t_uscalar_t	size in octets	send buffer size
XTI_SNDLOWAT	t_uscalar_t	size in octets	send low-water mark

Table 7-1 XTI-level Options

A request for XTI_DEBUG is an absolute requirement. A request to activate XTI_LINGER is an absolute requirement; the timeout value to this option is not. XTI_RCVBUF, XTI_RCVLOWAT, XTI_SNDBUF and XTI_SNDLOWAT are not absolute requirements.

XTI_DEBUG This option enables debugging. The values of this option are implementation-defined. Debugging is disabled if the option is specified with “no value”; that is, with an option header only.

The system supplies utilities to process the traces. Note that an implementation may also provide other means for debugging.

XTI_LINGER This option is used to linger the execution of a *t_close()* or *close()* if send data is still queued in the send buffer. The option value specifies the linger period. If a *close()* or *t_close()* is issued and the send buffer is not empty, the system attempts to send the pending data within the linger period before closing the endpoint. Data still pending after the linger period has elapsed is discarded.

Depending on the implementation, *t_close()* or *close()* either block for at maximum the linger period, or immediately return, whereupon the system holds the connection in existence for at most the linger period.

The option value consists of a structure **t_linger** declared as:

```
struct t_linger {
    t_scalar_t l_onoff; /* switch option on/off */
    t_scalar_t l_linger; /* linger period in seconds */
}
```

Legal values for the field *l_onoff* are:

```
T_NO  switch option off
T_YES activate option
```

The value *l_onoff* is an absolute requirement.

The field *l_linger* determines the linger period in seconds. The transport user can request the default value by setting the field to T_UNSPEC. The default timeout value depends on the underlying transport provider (it is often T_INFINITE). Legal values for this field are T_UNSPEC, T_INFINITE and all non-negative numbers.

The *l_linger* value is not an absolute requirement. The implementation may place upper and lower limits to this value. Requests that fall short of

the lower limit are negotiated to the lower limit.

Note that this option does not linger the execution of *t_snddis()*.

XTI_RCVBUF This option is used to adjust the internal buffer size allocated for the receive buffer. The buffer size may be increased for high-volume connections, or decreased to limit the possible backlog of incoming data.

This request is not an absolute requirement. The implementation may place upper and lower limits on the option value. Requests that fall short of the lower limit are negotiated to the lower limit.

Legal values are all positive numbers.

XTI_RCVLOWAT This option is used to set a low-water mark in the receive buffer. The option value gives the minimal number of bytes that must have accumulated in the receive buffer before they become visible to the transport user. If and when the amount of accumulated receive data exceeds the low-water mark, a T_DATA event is created, an event mechanism (for example, *poll()* or *select()*) indicates the data, and the data can be read by *t_rcv()* or *t_rcvudata()*.

This request is not an absolute requirement. The implementation may place upper and lower limits on the option value. Requests that fall short of the lower limit are negotiated to the lower limit.

Legal values are all positive numbers.

XTI_SNDBUF This option is used to adjust the internal buffer size allocated for the send buffer.

This request is not an absolute requirement. The implementation may place upper and lower limits on the option value. Requests that fall short of the lower limit are negotiated to the lower limit.

Legal values are all positive numbers.

XTI_SNDLOWAT This option is used to set a low-water mark in the send buffer. The option value gives the minimal number of bytes that must have accumulated in the send buffer before they are sent.

This request is not an absolute requirement. The implementation may place upper and lower limits on the option value. Requests that fall short of the lower limit are negotiated to the lower limit.

Legal values are all positive numbers.

VALID STATES

ALL - apart from T_UNINIT

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBADFLAG]	An invalid flag was specified.
[TBADOPT]	The specified options were in an incorrect format or contained illegal information.

[TBUFOVFLW]	The number of bytes allowed for an incoming argument (<i>maxlen</i>) is greater than 0 but not sufficient to store the value of that argument. The information to be returned in <i>ret</i> will be discarded.
[TNOTSUPPORT]	This action is not supported by the transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_accept(), *t_alloc()*, *t_connect()*, *t_getinfo()*, *t_listen()*, *t_open()*, *t_rcvconnect()*, Chapter 6.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_rcv - receive data or expedited data sent over a connection

SYNOPSIS

```
#include <xti.h>

int t_rcv(
    int fd,
    void *buf,
    unsigned int nbytes,
    int *flags)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>buf</i>	x	(x)
<i>nbytes</i>	x	/
<i>flags</i>	/	x

This function receives either normal or expedited data. The argument *fd* identifies the local transport endpoint through which data will arrive, *buf* points to a receive buffer where user data will be placed, and *nbytes* specifies the size of the receive buffer. The argument *flags* may be set on return from *t_rcv()* and specifies optional flags as described below.

By default, *t_rcv()* operates in synchronous mode and will wait for data to arrive if none is currently available. However, if `O_NONBLOCK` is set (via *t_open()* or *fcntl()*), *t_rcv()* will execute in asynchronous mode and will fail if no data is available. (See [TNODATA] below.)

On return from the call, if `T_MORE` is set in *flags*, this indicates that there is more data, and the current transport service data unit (TSDU) or expedited transport service data unit (ETSDU) must be received in multiple *t_rcv()* calls. In the asynchronous mode, or under unusual conditions (for example, the arrival of a signal or `T_EXDATA` event), the `T_MORE` flag may be set on return from the *t_rcv()* call even when the number of bytes received is less than the size of the receive buffer specified. Each *t_rcv()* with the `T_MORE` flag set indicates that another *t_rcv()* must follow to get more data for the current TSDU. The end of the TSDU is identified by the return of a *t_rcv()* call with the `T_MORE` flag not set. If the transport provider does not support the concept of a TSDU as indicated in the *info* argument on return from *t_open()* or *t_getinfo()*, the `T_MORE` flag is not meaningful and should be ignored. If *nbytes* is greater than zero on the call to *t_rcv()*, *t_rcv()* will return 0 only if the end of a TSDU is being returned to the user.

On return, the data is expedited if `T_EXPEDITED` is set in *flags*. If `T_MORE` is also set, it indicates that the number of expedited bytes exceeded *nbytes*, a signal has interrupted the call, or that an entire ETSDU was not available (only for transport protocols that support fragmentation of ETSDUs). The rest of the ETSDU will be returned by subsequent calls to *t_rcv()* which will return with `T_EXPEDITED` set in *flags*. The end of the ETSDU is identified by the return of a *t_rcv()* call with `T_EXPEDITED` set and `T_MORE` cleared. If the entire ETSDU is not available it is possible for normal data fragments to be returned between the initial and final fragments of an ETSDU.

If a signal arrives, *t_rcv()* returns, giving the user any data currently available. If no data is available, *t_rcv()* returns `-1`, sets *t_errno* to [TSYSERR] and *errno* to [EINTR]. If some data is available, *t_rcv()* returns the number of bytes received and `T_MORE` is set in *flags*.

In synchronous mode, the only way for the user to be notified of the arrival of normal or expedited data is to issue this function or check for the T_DATA or T_EXDATA events using the *t_look()* function. Additionally, the process can arrange to be notified via the EM interface.

VALID STATES

T_DATAXFER, T_OUTREL

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TLOOK]	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
[TNODATA]	O_NONBLOCK was set, but no data is currently available from the transport provider.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

On successful completion, *t_rcv()* returns the number of bytes received. Otherwise, it returns -1 on failure and *t_errno* is set to indicate the error.

SEE ALSO

fcntl(), *t_getinfo()*, *t_look()*, *t_open()*, *t_snd()*.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_rcvconnect - receive the confirmation from a connection request

SYNOPSIS

```
#include <xti.h>

int t_rcvconnect(
    int fd,
    struct t_call *call)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>call</i> → <i>addr.maxlen</i>	x	=
<i>call</i> → <i>addr.len</i>	/	x
<i>call</i> → <i>addr.buf</i>	?	(?)
<i>call</i> → <i>opt.maxlen</i>	x	=
<i>call</i> → <i>opt.len</i>	/	x
<i>call</i> → <i>opt.buf</i>	?	(?)
<i>call</i> → <i>udata.maxlen</i>	x	=
<i>call</i> → <i>udata.len</i>	/	x
<i>call</i> → <i>udata.buf</i>	?	(?)
<i>call</i> → <i>sequence</i>	=	=

This function enables a calling transport user to determine the status of a previously sent connection request and is used in conjunction with *t_connect()* to establish a connection in asynchronous mode, and to complete a synchronous *t_connect()* call that was interrupted by a signal. The connection will be established on successful completion of this function.

The argument *fd* identifies the local transport endpoint where communication will be established, and *call* contains information associated with the newly established connection. The argument *call* points to a **t_call** structure which contains the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
int sequence;
```

In *call*, *addr* returns the protocol address associated with the responding transport endpoint, *opt* presents any options associated with the connection, *udata* points to optional user data that may be returned by the destination transport user during connection establishment, and *sequence* has no meaning for this function.

The *maxlen* field of each argument must be set before issuing this function to indicate the maximum size of the buffer for each. However, *maxlen* can be set to zero, in which case no information to this specific argument is given to the user on the return from *t_rcvconnect()*. If *call* is set to NULL, no information at all is returned. By default, *t_rcvconnect()* executes in synchronous mode and waits for the connection to be established before returning. On return, the *addr*, *opt* and *udata* fields reflect values associated with the connection.

If O_NONBLOCK is set (via *t_open()* or *fcntl()*), *t_rcvconnect()* executes in asynchronous mode, and reduces to a poll for existing connection confirmations. If none are available, *t_rcvconnect()* fails and returns immediately without waiting for the connection to be established. (See [TNODATA] below.) In this case, *t_rcvconnect()* must be called again to complete the connection establishment phase and retrieve the information returned in *call*.

VALID STATES

T_OUTCON

ERRORSOn failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBUFOVFLW]	The number of bytes allocated for an incoming argument (<i>maxlen</i>) is greater than 0 but not sufficient to store the value of that argument, and the connection information to be returned in <i>call</i> will be discarded. The provider's state, as seen by the user, will be changed to T_DATAXFER.
[TLOOK]	An asynchronous event has occurred on this transport connection and requires immediate attention.
[TNODATA]	O_NONBLOCK was set, but a connection confirmation has not yet arrived.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUEUpon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.**SEE ALSO***t_accept()*, *t_alloc()*, *t_bind()*, *t_connect()*, *t_listen()*, *t_open()*, *t_optmgmt()*.**CHANGE HISTORY****Issue 4**The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_rcvdis - retrieve information from disconnection

SYNOPSIS

#include <xti.h>

```
int t_rcvdis(
    int fd,
    struct t_discon *discon)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>discon</i> → <i>udata.maxlen</i>	x	=
<i>discon</i> → <i>udata.len</i>	/	x
<i>discon</i> → <i>udata.buf</i>	?	(?)
<i>discon</i> → <i>reason</i>	/	x
<i>discon</i> → <i>sequence</i>	/	?

This function is used to identify the cause of a disconnection and to retrieve any user data sent with the disconnection. The argument *fd* identifies the local transport endpoint where the connection existed, and *discon* points to a **t_discon** structure containing the following members:

```
struct netbuf udata;
int reason;
int sequence;
```

The field *reason* specifies the reason for the disconnection through a protocol-dependent reason code, *udata* identifies any user data that was sent with the disconnection, and *sequence* may identify an outstanding connection indication with which the disconnection is associated. The field *sequence* is only meaningful when *t_rcvdis()* is issued by a passive transport user who has executed one or more *t_listen()* functions and is processing the resulting connection indications. If a disconnection indication occurs, *sequence* can be used to identify which of the outstanding connection indications is associated with the disconnection.

The *maxlen* field of *udata* may be set to zero, if the user does not care about incoming data. If, in addition, the user does not need to know the value of *reason* or *sequence*, *discon* may be set to NULL and any user data associated with the disconnection indication shall be discarded. However, if a user has retrieved more than one outstanding connection indication (via *t_listen()*) and *discon* is a null pointer, the user will be unable to identify with which connection indication the disconnection is associated.

VALID STATES

T_DATAXFER,T_OUTCON,T_OUTREL,T_INREL,T_INCON(ocnt > 0)

ERRORSOn failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBUFOVFLW]	The number of bytes allocated for incoming data (<i>maxlen</i>) is greater than 0 but not sufficient to store the data. If <i>fd</i> is a passive endpoint with <i>ocnt</i> > 1, it remains in state T_INCON; otherwise, the endpoint state is set to T_IDLE.
[TNODIS]	No disconnection indication currently exists on the specified transport endpoint.

[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_alloc(), *t_connect()*, *t_listen()*, *t_open()*, *t_snddis()*.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_rcvrel - acknowledge receipt of an orderly release indication

SYNOPSIS

```
#include <xti.h>
```

```
int t_rcvrel(
    int fd)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/

This function is used to receive an orderly release indication for the incoming direction of data transfer. The argument *fd* identifies the local transport endpoint where the connection exists. After receipt of this indication, the user may not attempt to receive more data via *t_rcv()* or *t_rcvv()*. Such an attempt will fail with *t_errno* set to [TOUTSTATE]. However, the user may continue to send data over the connection if *t_sndrel()* has not been called by the user. This function is an optional service of the transport provider, and is only supported if the transport provider returned service type T_COTS_ORD on *t_open()* or *t_getinfo()*. Any user data that may be associated with the orderly release indication is discarded when *t_rcvrel()* is called.

VALID STATES

T_DATAXFER, T_OUTREL

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TLOOK]	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
[TNOREL]	No orderly release indication currently exists on the specified transport endpoint.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_getinfo(), *t_open()*, *t_sndrel()*.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_rcvreldata - receive an orderly release indication or confirmation containing user data

SYNOPSIS

```
#include <xti.h>

int t_rcvreldata(
    int fd,
    struct t_discon *discon)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>discon</i> → <i>udata.maxlen</i>	x	=
<i>discon</i> → <i>udata.len</i>	/	x
<i>discon</i> → <i>udata.buf</i>	?	(?)
<i>discon</i> → <i>reason</i>	/	x
<i>discon</i> → <i>sequence</i>	/	=

This function is used to receive an orderly release indication for the incoming direction of data transfer and to retrieve any user data sent with the release. The argument *fd* identifies the local transport endpoint where the connection exists, and *discon* points to a **t_discon** structure containing the following members:

```
struct netbuf udata;
int reason;
int sequence;
```

After receipt of this indication, the user may not attempt to receive more data via *t_rcv()* or *t_rcvv()*. Such an attempt will fail with *t_error* set to [TOUTSTATE]. However, the user may continue to send data over the connection if *t_sndrel()* or *t_sndreldata()* has not been called by the user.

The field *reason* specifies the reason for the disconnection through a protocol-dependent *reason code*, and *udata* identifies any user data that was sent with the disconnection; the field *sequence* is not used.

If a user does not care if there is incoming data and does not need to know the value of *reason*, *discon* may be a null pointer, and any user data associated with the disconnection will be discarded.

If *discon*→*udata.maxlen* is greater than zero and less than the length of the value, *t_rcvreldata()* fails with *t_errno* set to [TBUFOVFLW].

This function is an optional service of the transport provider, only supported by providers of service type T_COTS_ORD. The flag T_ORDRELDATA in the *info*→*flag* field returned by *t_open()* or *t_getinfo()* indicates that the provider supports orderly release user data; when the flag is not set, this function behaves as *t_rcvrel()* and no user data is returned.

This function may not be available on all systems.

VALID STATES

T_DATAXFER, T_OUTREL

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBUFOVFLW]	The number of bytes allocated for incoming data (maxlen) is greater than 0 but not sufficient to store the data, and the disconnection information to be returned in <i>discon</i> will be discarded. The provider state, as seen by the user, will be changed as if the data was successfully retrieved.
[TLOOK]	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
[TNOREL]	No orderly release indication currently exists on the specified transport endpoint.
[TNOTSUPPORT]	Orderly release is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_getinfo(), *t_open()*, *t_sndreldata()*, *t_rcvrel()*, *t_sndrel()*.

NAME

t_rcvudata - receive a data unit

SYNOPSIS

```
#include <xti.h>

int t_rcvudata(
    int fd,
    struct t_unitdata *unitdata,
    int *flags)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>unitdata</i> → <i>addr.maxlen</i>	x	=
<i>unitdata</i> → <i>addr.len</i>	/	x
<i>unitdata</i> → <i>addr.buf</i>	?	(?)
<i>unitdata</i> → <i>opt.maxlen</i>	x	=
<i>unitdata</i> → <i>opt.len</i>	/	x
<i>unitdata</i> → <i>opt.buf</i>	?	(?)
<i>unitdata</i> → <i>udata.maxlen</i>	x	=
<i>unitdata</i> → <i>udata.len</i>	/	x
<i>unitdata</i> → <i>udata.buf</i>	?	(?)
<i>flags</i>	/	x

This function is used in connectionless-mode to receive a data unit from another transport user. The argument *fd* identifies the local transport endpoint through which data will be received, *unitdata* holds information associated with the received data unit, and *flags* is set on return to indicate that the complete data unit was not received. The argument *unitdata* points to a **t_unitdata** structure containing the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
```

The *maxlen* field of *addr*, *opt* and *udata* must be set before calling this function to indicate the maximum size of the buffer for each. If the *maxlen* field of *addr* or *opt* is set to zero, no information is returned in the *buf* field of this parameter.

On return from this call, *addr* specifies the protocol address of the sending user, *opt* identifies options that were associated with this data unit, and *udata* specifies the user data that was received.

By default, *t_rcvudata()* operates in synchronous mode and will wait for a data unit to arrive if none is currently available. However, if `O_NONBLOCK` is set (via *t_open()* or *fcntl()*), *t_rcvudata()* will execute in asynchronous mode and will fail if no data units are available.

If the buffer defined in the *udata* field of *unitdata* is not large enough to hold the current data unit, the buffer will be filled and `T_MORE` will be set in *flags* on return to indicate that another *t_rcvudata()* should be called to retrieve the rest of the data unit. Subsequent calls to *t_rcvudata()* will return zero for the length of the address and options until the full data unit has been received.

If the call is interrupted, *t_rcvudata()* will return `[EINTR]` and no datagrams will have been removed from the endpoint.

VALID STATES

T_IDLE

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBUFOVFLW]	The number of bytes allocated for the incoming protocol address or options (<i>maxlen</i>) is greater than 0 but not sufficient to store the information. The unit data information to be returned in <i>unitdata</i> will be discarded.
[TLOOK]	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
[TNODATA]	O_NONBLOCK was set, but no data units are currently available from the transport provider.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

fcntl(), *t_alloc()*, *t_open()*, *t_rcvuderr()*, *t_sndudata()*.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_rcvuderr - receive a unit data error indication

SYNOPSIS

```
#include <xti.h>

int t_rcvuderr(
    int fd,
    struct t_uderr *uderr)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>uderr</i> → <i>addr.maxlen</i>	x	=
<i>uderr</i> → <i>addr.len</i>	/	x
<i>uderr</i> → <i>addr.buf</i>	?	(?)
<i>uderr</i> → <i>opt.maxlen</i>	x	=
<i>uderr</i> → <i>opt.len</i>	/	x
<i>uderr</i> → <i>opt.buf</i>	?	(?)
<i>uderr</i> → <i>error</i>	/	x

This function is used in connectionless-mode to receive information concerning an error on a previously sent data unit, and should only be issued following a unit data error indication. It informs the transport user that a data unit with a specific destination address and protocol options produced an error. The argument *fd* identifies the local transport endpoint through which the error report will be received, and *uderr* points to a **t_uderr** structure containing the following members:

```
struct netbuf addr;
struct netbuf opt;
t_scalar_t error;
```

The *maxlen* field of *addr* and *opt* must be set before calling this function to indicate the maximum size of the buffer for each. If this field is set to zero for *addr* or *opt*, no information is returned in the *buf* field of this parameter.

On return from this call, the *addr* structure specifies the destination protocol address of the erroneous data unit, the *opt* structure identifies options that were associated with the data unit, and *error* specifies a protocol-dependent error code.

If the user does not care to identify the data unit that produced an error, *uderr* may be set to a null pointer, and *t_rcvuderr()* will simply clear the error indication without reporting any information to the user.

VALID STATES

T_IDLE

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBUFOVFLW]	The number of bytes allocated for the incoming protocol address or options (<i>maxlen</i>) is greater than 0 but not sufficient to store the information. The unit data error information to be returned in <i>uderr</i> will be discarded.

[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TNOUDERR]	No unit data error indication currently exists on the specified transport endpoint.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_rcvudata(), *t_sndudata()*.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_rcvv - receive data or expedited data sent over a connection and put the data into one or more non-contiguous buffers

SYNOPSIS

```
#include <xti.h>

int t_rcvv(
    int fd,
    struct t_iovec *iov,
    unsigned int iovcount,
    int *flags)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>iov</i>	x/	
<i>iovcount</i>	x	/
<i>iov[0].iov_base</i>	x(/)	=(x)
<i>iov[0].iov_len</i>	x	=
. . . .		
<i>iov[iovcount-1].iov_base</i>	x(/)	=(x)
<i>iov[iovcount-1].iov_len</i>	x	=

This function receives either normal or expedited data. The argument *fd* identifies the local transport endpoint through which data will arrive, *iov* points to an array of buffer address/buffer size pairs (*iov_base*, *iov_len*). The *t_rcvv()* function receives data into the buffers specified by *iov[0].iov_base*, *iov[1].iov_base*, through *iov[iovcount-1].iov_base*, always filling one buffer before proceeding to the next.

Note: The limit on the total number of bytes available in all buffers passed (that is, *iov(0).iov_len* + . . . + *iov(iovcount-1).iov_len*) may be constrained by implementation limits. If no other constraint applies, it will be limited by [INT_MAX]. In practice, the availability of memory to an application is likely to impose a lower limit on the amount of data that can be sent or received using scatter/gather functions.

The argument *iovcount* contains the number of buffers which is limited to T_IOV_MAX (an implementation-defined value of at least 16). If the limit is exceeded, the function will fail with [TBADDDATA].

The argument *flags* may be set on return from *t_rcvv()* and specifies optional flags as described below.

By default, *t_rcvv()* operates in synchronous mode and will wait for data to arrive if none is currently available. However, if O_NONBLOCK is set (via *t_open()* or *fcntl()*), *t_rcvv()* will execute in asynchronous mode and will fail if no data is available (see [TNODATA] below).

On return from the call, if T_MORE is set in *flags*, this indicates that there is more data, and the current transport service data unit (TSDU) or expedited transport service data unit (ETSDU) must be received in multiple *t_rcvv()* or *t_rcv()* calls. In the asynchronous mode, or under unusual conditions (for example, the arrival of a signal or T_EXDATA event), the T_MORE flag may be set on return from the *t_rcvv()* call even when the number of bytes received is less than the total size of all the receive buffers. Each *t_rcvv()* with the T_MORE flag set indicates that another *t_rcvv()* must follow to get more data for the current TSDU. The end of the TSDU is identified by the return of a *t_rcvv()* call with the T_MORE flag not set. If the transport provider

does not support the concept of a TSDU as indicated in the *info* argument on return from *t_open()* or *t_getinfo()*, the T_MORE flag is not meaningful and should be ignored. If the amount of buffer space passed in *iov* is greater than zero on the call to *t_rcvv()*, then *t_rcvv()* will return 0 only if the end of a TSDU is being returned to the user.

On return, the data is expedited if T_EXPEDITED is set in flags. If T_MORE is also set, it indicates that the number of expedited bytes exceeded *nbytes*, a signal has interrupted the call, or that an entire ETSDU was not available (only for transport protocols that support fragmentation of ETSDUs). The rest of the ETSDU will be returned by subsequent calls to *t_rcvv()* which will return with T_EXPEDITED set in flags. The end of the ETSDU is identified by the return of a *t_rcvv()* call with T_EXPEDITED set and T_MORE cleared. If the entire ETSDU is not available it is possible for normal data fragments to be returned between the initial and final fragments of an ETSDU.

If a signal arrives, *t_rcvv()* returns, giving the user any data currently available. If no data is available, *t_rcvv()* returns -1, sets *t_errno* to [TSYSERR] and *errno* to [EINTR]. If some data is available, *t_rcvv()* returns the number of bytes received and T_MORE is set in flags.

In synchronous mode, the only way for the user to be notified of the arrival of normal or expedited data is to issue this function or check for the T_DATA or T_EXDATA events using the *t_look()* function. Additionally, the process can arrange to be notified via the EM interface.

VALID STATES

T_DATAXFER, T_OUTREL

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADDDATA]	<i>iovcnt</i> is greater than T_IOV_MAX.
[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TLOOK]	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
[TNODATA]	O_NONBLOCK was set, but no data is currently available from the transport provider.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

On successful completion, *t_rcvv()* returns the number of bytes received. Otherwise, it returns -1 on failure and *t_errno* is set to indicate the error.

SEE ALSO

fcntl(), *t_getinfo()*, *t_look()*, *t_open()*, *t_rcv()*, *t_snd()*, *t_sndv()*.

NAME

t_rcvvudata — receive a data unit into one or more noncontiguous buffers

SYNOPSIS

```
#include <xti.h>

int t_rcvvudata(
    int fd,
    struct t_unitdata *unitdata,
    struct t_iovec *iov,
    unsigned int iovcount,
    int *flags)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>unitdata->addr.maxlen</i>	x	=
<i>unitdata->addr.len</i>	/	x
<i>unitdata->addr.buf</i>	?(/)	=(/)
<i>unitdata->opt.maxlen</i>	x	=
<i>unitdata->opt.len</i>	/	x
<i>unitdata->opt.buf</i>	?(/)	=(?)
<i>unitdata->udata.maxlen</i>	/	=
<i>unitdata->udata.len</i>	/	=
<i>unitdata->udata.buf</i>	/	=
<i>iov[0].iov_base</i>	x	=(x)
<i>iov[0].iov_len</i>	x	=
. . . .		
<i>iov[iovcount-1].iov_base</i>	x(/)	=(x)
<i>iov[iovcount-1].iov_len</i>	x	=
<i>iovcount</i>	x	/
<i>flags</i>	/	/

This function is used in connectionless mode to receive a data unit from another transport user. The argument *fd* identifies the local transport endpoint through which data will be received, *unitdata* holds information associated with the received data unit, *iovcount* contains the number of non-contiguous *udata* buffers which is limited to T_IOV_MAX (an implementation-defined value of at least 16), and *flags* is set on return to indicate that the complete data unit was not received. If the limit on *iovcount* is exceeded, the function fails with [TBADDDATA]. The argument *unitdata* points to a **t_unitdata** structure containing the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
```

The *maxlen* field of *addr* and *opt* must be set before calling this function to indicate the maximum size of the buffer for each. The *udata* field of **t_unitdata** is not used. The *iov_len* and *iov_base* fields of *iov[0]* through *iov[iovcount-1]* must be set before calling *t_rcvvudata()* to define the buffer where the userdata will be placed. If the *maxlen* field of *addr* or *opt* is set to zero then no information is returned in the *buf* field for this parameter.

On return from this call, *addr* specifies the protocol address of the sending user, *opt* identifies options that were associated with this data unit, and *iov[0].iov_base* through *iov[iovcount-1].iov_base* contains the user data that was received. The return value of *t_rcvvudata()* is the

number of bytes of user data given to the user.

Note: The limit on the total number of bytes available in all buffers passed (that is, $iov(0).iov_len + \dots + iov(iovcount-1).iov_len$) may be constrained by implementation limits. If no other constraint applies, it will be limited by [INT_MAX]. In practice, the availability of memory to an application is likely to impose a lower limit on the amount of data that can be sent or received using scatter/gather functions.

By default, `t_rcvvudata()` operates in synchronous mode and waits for a data unit to arrive if none is currently available. However, if `O_NONBLOCK` is set (via `t_open()` or `fcntl()`), `t_rcvvudata()` executes in asynchronous mode and fails if no data units are available.

If the buffers defined in the `iov[]` array are not large enough to hold the current data unit, the buffers will be filled and `T_MORE` will be set in flags on return to indicate that another `t_rcvvudata()` should be called to retrieve the rest of the data unit. Subsequent calls to `t_rcvvudata()` will return zero for the length of the address and options, until the full data unit has been received.

VALID STATES

T_IDLE

ERRORS

On failure, `t_errno` is set to one of the following:

[TBADDDATA]	<code>iovcount</code> is greater than <code>T_IOV_MAX</code> .
[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBUFOVFLW]	The number of bytes allocated for the incoming protocol address or options (<code>maxlen</code>) is greater than 0 but not sufficient to store the information. The unit data information to be returned in <code>unitdata</code> will be discarded.
[TLOOK]	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
[TNODATA]	<code>O_NONBLOCK</code> was set, but no data units are currently available from the transport provider.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <code>fd</code> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<code>t_errno</code>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUES

On successful completion, `t_rcvvudata()` returns the number of bytes received. Otherwise, it returns `-1` on failure and `t_errno` is set to indicate the error.

SEE ALSO

`fcntl()`, `t_alloc()`, `t_open()`, `t_rcvudata()`, `t_rcvuderr()`, `t_sndudata()`, `t_sndvudata()`.

NAME

t_snd - send data or expedited data over a connection

SYNOPSIS

```
#include <xti.h>
```

```
int t_snd(
    int fd,
    void *buf,
    unsigned int nbytes,
    int flags)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>buf</i>	x (x)	=
<i>nbytes</i>	x	/
<i>flags</i>	x	/

This function is used to send either normal or expedited data. The argument *fd* identifies the local transport endpoint over which data should be sent, *buf* points to the user data, *nbytes* specifies the number of bytes of user data to be sent, and *flags* specifies any optional flags described below:

T_EXPEDITED If set in *flags*, the data will be sent as expedited data and will be subject to the interpretations of the transport provider.

T_MORE If set in *flags*, this indicates to the transport provider that the transport service data unit (TSDU) (or expedited transport service data unit - ETSDU) is being sent through multiple *t_snd()* calls. Each *t_snd()* with the T_MORE flag set indicates that another *t_snd()* will follow with more data for the current TSDU (or ETSDU).

The end of the TSDU (or ETSDU) is identified by a *t_snd()* call with the T_MORE flag not set. Use of T_MORE enables a user to break up large logical data units without losing the boundaries of those units at the other end of the connection. The flag implies nothing about how the data is packaged for transfer below the transport interface. If the transport provider does not support the concept of a TSDU as indicated in the *info* argument on return from *t_open()* or *t_getinfo()*, the T_MORE flag is not meaningful and will be ignored if set.

The sending of a zero-length fragment of a TSDU or ETSDU is only permitted where this is used to indicate the end of a TSDU or ETSDU; that is, when the T_MORE flag is not set. Some transport providers also forbid zero-length TSDUs and ETSDUs. See Appendix A on page 215 for a fuller explanation.

T_PUSH If set in *flags*, requests that the provider transmit all data that it has accumulated but not sent. The request is a local action on the provider and does not affect any similarly named protocol flag (for example, the TCP PUSH flag). This effect of setting this flag is protocol-dependent, and it may be ignored entirely by transport providers which do not support the use of this feature.

Note: The communications provider is free to collect data in a send buffer until it accumulates a sufficient amount for transmission.

By default, *t_snd()* operates in synchronous mode and may wait if flow control restrictions prevent the data from being accepted by the local transport provider at the time the call is made. However, if `O_NONBLOCK` is set (via *t_open()* or *fcntl()*), *t_snd()* will execute in asynchronous mode, and will fail immediately if there are flow control restrictions. The process can arrange to be informed when the flow control restrictions are cleared via either *t_look()* or the EM interface.

On successful completion, *t_snd()* returns the number of bytes (octets) accepted by the communications provider. Normally this will equal the number of octets specified in *nbytes*. However, if `O_NONBLOCK` is set or the function is interrupted by a signal, it is possible that only part of the data has actually been accepted by the communications provider. In this case, *t_snd()* returns a value that is less than the value of *nbytes*. If *t_snd()* is interrupted by a signal before it could transfer data to the communications provider, it returns `-1` with *t_errno* set to `[TSYSERR]` and *errno* set to `[EINTR]`.

If *nbytes* is zero and sending of zero bytes is not supported by the underlying communications service, *t_snd()* returns `-1` with *t_errno* set to `[TBADDDATA]`.

The size of each TSDU or ETSDU must not exceed the limits of the transport provider as specified by the current values in the TSDU or ETSDU fields in the *info* argument returned by *t_getinfo()*.

The error `[TLOOK]` is returned for asynchronous events. It is required only for an incoming disconnect event but may be returned for other events.

VALID STATES

T_DATAXFER, T_INREL

ERRORS

On failure, *t_errno* is set to one of the following:

- `[TBADDDATA]` Illegal amount of data:
- A single send was attempted specifying a TSDU (ETSDU) or fragment TSDU (ETSDU) greater than that specified by the current values of the TSDU or ETSDU fields in the *info* argument.
 - A send of a zero byte TSDU (ETSDU) or zero byte fragment of a TSDU (ETSDU) is not supported by the provider (see Appendix A on page 215).
 - Multiple sends were attempted resulting in a TSDU (ETSDU) larger than that specified by the current value of the TSDU or ETSDU fields in the *info* argument — the ability of an XTI implementation to detect such an error case is implementation-dependent (see **CAVEATS**, below).
- `[TBADF]` The specified file descriptor does not refer to a transport endpoint.

[TBADFLAG]	An invalid flag was specified.
[TFLOW]	O_NONBLOCK was set, but the flow control mechanism prevented the transport provider from accepting any data at this time.
[TLOOK]	An asynchronous event has occurred on this transport endpoint.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

On successful completion, *t_snd()* returns the number of bytes accepted by the transport provider. Otherwise, -1 is returned on failure and *t_errno* is set to indicate the error.

Note: If the number of bytes accepted by the communications provider is less than the number of bytes requested, this may either indicate that O_NONBLOCK is set and the communications provider is blocked due to flow control, or that O_NONBLOCK is clear and the function was interrupted by a signal.

SEE ALSO

t_getinfo(), *t_open()*, *t_rcv()*.

CAVEATS

It is important to remember that the transport provider treats all users of a transport endpoint as a single user. Therefore if several processes issue concurrent *t_snd()* calls then the different data may be intermixed.

Multiple sends which exceed the maximum TSDU or ETSDU size may not be discovered by XTI. In this case an implementation-dependent error will result (generated by the transport provider) perhaps on a subsequent XTI call. This error may take the form of a connection abort, a [TSYSERR], a [TBADDDATA] or a [TPROTO] error.

If multiple sends which exceed the maximum TSDU or ETSDU size are detected by XTI, *t_snd()* fails with [TBADDDATA].

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_snddis - send user-initiated disconnection request

SYNOPSIS

```
#include <xti.h>

int t_snddis(
    int fd,
    const struct t_call *call)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>call</i> → <i>addr.maxlen</i>	=	=
<i>call</i> → <i>addr.len</i>	=	=
<i>call</i> → <i>addr.buf</i>	=	=
<i>call</i> → <i>opt.maxlen</i>	=	=
<i>call</i> → <i>opt.len</i>	=	=
<i>call</i> → <i>opt.buf</i>	=	=
<i>call</i> → <i>udata.maxlen</i>	=	=
<i>call</i> → <i>udata.len</i>	x	=
<i>call</i> → <i>udata.buf</i>	?(?)	=
<i>call</i> → <i>sequence</i>	?	=

This function is used to initiate an abortive release on an already established connection, or to reject a connection request. The argument *fd* identifies the local transport endpoint of the connection, and *call* specifies information associated with the abortive release. The argument *call* points to a **t_call** structure which contains the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
int sequence;
```

The values in *call* have different semantics, depending on the context of the call to *t_snddis()*. When rejecting a connection request, *call* must be non-null and contain a valid value of *sequence* to uniquely identify the rejected connection indication to the transport provider. The *sequence* field is only meaningful if the transport connection is in the T_INCON state. The *addr* and *opt* fields of *call* are ignored. In all other cases, *call* need only be used when data is being sent with the disconnection request. The *addr*, *opt* and *sequence* fields of the **t_call** structure are ignored. If the user does not wish to send data to the remote user, the value of *call* may be a null pointer.

The *udata* structure specifies the user data to be sent to the remote user. The amount of user data must not exceed the limits supported by the transport provider, as returned in the *discon* field, of the *info* argument of *t_open()* or *t_getinfo()*. If the *len* field of *udata* is zero, no data will be sent to the remote user.

VALID STATES

T_DATAXFER,T_OUTCON,T_OUTREL,T_INREL,T_INCON(ocnt > 0)

ERRORSOn failure, *t_errno* is set to one of the following:

[TBADF] The specified file descriptor does not refer to a transport endpoint.

[TBADDATA]	The amount of user data specified was not within the bounds allowed by the transport provider.
[TBADSEQ]	An invalid sequence number was specified, or a null <i>call</i> pointer was specified, when rejecting a connection request.
[TLOOK]	An asynchronous event, which requires attention, has occurred.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_connect(), *t_getinfo()*, *t_listen()*, *t_open()*.

CAVEATS

t_snddis() is an abortive disconnection. Therefore a *t_snddis()* issued on a connection endpoint may cause data previously sent via *t_snd()*, or data not yet received, to be lost (even if an error is returned).

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_sndrel - initiate an orderly release

SYNOPSIS

#include <xti.h>

```
int t_sndrel(
    int fd)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/

For transport providers of type T_COTS_ORD, this function is used to initiate an orderly release of the outgoing direction of data transfer and indicates to the transport provider that the transport user has no more data to send. The argument *fd* identifies the local transport endpoint where the connection exists. After calling *t_sndrel()*, the user may not send any more data over the connection. However, a user may continue to receive data if an orderly release indication has not been received. For transport providers of types other than T_COTS_ORD, this function fails with error [TNOTSUPPORT].

VALID STATES

T_DATAXFER,T_INREL

ERRORSOn failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TFLOW]	O_NONBLOCK was set, but the flow control mechanism prevented the transport provider from accepting the function at this time.
[TLOOK]	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO*t_getinfo()*, *t_open()*, *t_rcvrel()*.**CHANGE HISTORY****Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_sndreldata - initiate/respond to an orderly release with user data

SYNOPSIS

```
#include <xti.h>

int t_sndreldata(
    int fd,
    struct t_discon *discon)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>discon</i> → <i>udata.maxlen</i>	/	=
<i>discon</i> → <i>udata.len</i>	x	/
<i>discon</i> → <i>udata.buf</i>	?(?)	/
<i>discon</i> → <i>reason</i>	?	/
<i>discon</i> → <i>sequence</i>	/	/

This function is used to initiate an orderly release of the outgoing direction of data transfer and to send user data with the release. The argument *fd* identifies the local transport endpoint where the connection exists, and *discon* points to a **t_discon** structure containing the following members:

```
struct netbuf udata;
int reason;
int sequence;
```

After calling *t_sndreldata()*, the user may not send any more data over the connection. However, a user may continue to receive data if an orderly release indication has not been received.

The field *reason* specifies the reason for the disconnection through a protocol-dependent *reason code*, and *udata* identifies any user data that is sent with the disconnection; the field *sequence* is not used.

The *udata* structure specifies the user data to be sent to the remote user. The amount of user data must not exceed the limits supported by the transport provider, as returned in the *discon* field of the *info* argument of *t_open()* or *t_getinfo()*. If the *len* field of *udata* is zero or if the provider did not return T_ORDRELDATA in the *t_open()* flags, no data will be sent to the remote user.

If a user does not wish to send data and reason code to the remote user, the value of *discon* may be a null pointer.

This function is an optional service of the transport provider, only supported by providers of service type T_COTS_ORD. The flag T_ORDRELDATA in the *info*→*flag* field returned by *t_open()* or *t_getinfo()* indicates that the provider supports orderly release user data; when the flag is not set, this function behaves as *t_rcvrel()* and no user data is returned.

This function may not be available on all systems.

VALID STATES

T_DATAXFER, T_INREL

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADDDATA]	The amount of user data specified was not within the bounds allowed by the transport provider, or user data was supplied and the provider did not return T_ORDRELDATA in the <i>t_open()</i> flags.
[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TFLOW]	O_NONBLOCK was set, but the flow control mechanism prevented the transport provider from accepting the function at this time.
[TLOOK]	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
[TNOTSUPPORT]	Orderly release is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_getinfo(), *t_open()*, *t_rcvreldata()*, *t_rcvrel()*, *t_sndrel()*.

NAME

t_sndudata - send a data unit

SYNOPSIS

```
#include <xti.h>

int t_sndudata(
    int fd,
    const struct t_unitdata *unitdata)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>unitdata</i> → <i>addr.maxlen</i>	=	=
<i>unitdata</i> → <i>addr.len</i>	x	=
<i>unitdata</i> → <i>addr.buf</i>	x(x)	=
<i>unitdata</i> → <i>opt.maxlen</i>	=	=
<i>unitdata</i> → <i>opt.len</i>	x	=
<i>unitdata</i> → <i>opt.buf</i>	?(?)	=
<i>unitdata</i> → <i>udata.maxlen</i>	=	=
<i>unitdata</i> → <i>udata.len</i>	x	=
<i>unitdata</i> → <i>udata.buf</i>	x(x)	=

This function is used in connectionless-mode to send a data unit to another transport user. The argument *fd* identifies the local transport endpoint through which data will be sent, and *unitdata* points to a **t_unitdata** structure containing the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
```

In *unitdata*, *addr* specifies the protocol address of the destination user, *opt* identifies options that the user wants associated with this request, and *udata* specifies the user data to be sent. The user may choose not to specify what protocol options are associated with the transfer by setting the *len* field of *opt* to zero. In this case, the provider uses the option values currently set for the communications endpoint.

If the *len* field of *udata* is zero, and sending of zero octets is not supported by the underlying transport service, the *t_sndudata()* will return -1 with *t_errno* set to [TBADDDATA].

By default, *t_sndudata()* operates in synchronous mode and may wait if flow control restrictions prevent the data from being accepted by the local transport provider at the time the call is made. However, if O_NONBLOCK is set (via *t_open()* or *fcntl()*), *t_sndudata()* will execute in asynchronous mode and will fail under such conditions. The process can arrange to be notified of the clearance of a flow control restriction via either *t_look()* or the EM interface.

If the amount of data specified in *udata* exceeds the TSDU size as returned in the *tsdu* field of the *info* argument of *t_open()* or *t_getinfo()*, a [TBADDDATA] error will be generated. If *t_sndudata()* is called before the destination user has activated its transport endpoint (see *t_bind()*), the data unit may be discarded.

If it is not possible for the transport provider to immediately detect the conditions that cause the errors [TBADDADDR] and [TBADDOPT], these errors will alternatively be returned by *t_rcvuderr*. Therefore, an application must be prepared to receive these errors in both of these ways.

If the call is interrupted, *t_sndudata()* will return [EINTR] and the datagram will not be sent.

VALID STATES

T_IDLE

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADADDR]	The specified protocol address was in an incorrect format or contained illegal information.
[TBADDATA]	Illegal amount of data. A single send was attempted specifying a TSDU greater than that specified in the <i>info</i> argument, or a send of a zero byte TSDU is not supported by the provider.
[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBADOPT]	The specified options were in an incorrect format or contained illegal information.
[TFLOW]	O_NONBLOCK was set, but the flow control mechanism prevented the transport provider from accepting any data at this time.
[TLOOK]	An asynchronous event has occurred on this transport endpoint.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

fcntl(), *t_alloc()*, *t_open()*, *t_rcvudata()*, *t_rcvuderr()*.

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_sndv — send data or expedited data, from one or more non-contiguous buffers, on a connection

SYNOPSIS

```
#include <xti.h>
```

```
int t_sndv(
    int fd,
    const struct t_iovec *iov,
    unsigned int iovcount,
    int flags)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>iov</i>	x	/
<i>iovcount</i>	x	/
<i>iov[0].iov_base</i>	x(x)	=(=)
<i>iov[0].iov_len</i>	x	=
. . . .		
<i>iov[iovcount-1].iov_base</i>	x(x)	=(=)
<i>iov[iovcount-1].iov_len</i>	x	=
<i>flags</i>	x	/

This function is used to send either normal or expedited data. The argument *fd* identifies the local transport endpoint over which data should be sent, *iov* points to an array of buffer address/buffer length pairs. *t_sndv()* sends data contained in buffers *iov[0]*, *iov[1]*, through *iov[iovcount-1]*. *iovcount* contains the number of non-contiguous data buffers which is limited to T_IOV_MAX (an implementation-defined value of at least 16). If the limit is exceeded, the function fails with [TBADDDATA].

Note: The limit on the total number of bytes available in all buffers passed (that is: *iov(0).iov_len + . . . + iov(iovcount-1).iov_len*) may be constrained by implementation limits. If no other constraint applies, it will be limited by [INT_MAX]. In practice, the availability of memory to an application is likely to impose a lower limit on the amount of data that can be sent or received using scatter/gather functions.

The argument *flags* specifies any optional flags described below:

T_EXPEDITED

If set in *flags*, the data will be sent as expedited data and will be subject to the interpretations of the transport provider.

T_MORE

If set in *flags*, this indicates to the transport provider that the transport service data unit (TSDU) (or expedited transport service data unit — ETSDU) is being sent through multiple *t_sndv()* calls. Each *t_sndv()* with the T_MORE flag set indicates that another *t_sndv()* (or *t_snd()*) will follow with more data for the current TSDU (or ETSDU).

The end of the TSDU (or ETSDU) is identified by a *t_sndv()* call with the T_MORE flag not set. Use of T_MORE enables a user to break up large logical data units without losing the boundaries of those units at the other end of the connection. The flag implies nothing about how the data is packaged for transfer below the transport interface. If the transport provider does not support the concept of a TSDU as indicated in the *info* argument on

return from `t_open()` or `t_getinfo()`, the `T_MORE` flag is not meaningful and will be ignored if set.

The sending of a zero-length fragment of a TSDU or ETSDU is only permitted where this is used to indicate the end of a TSDU or ETSDU, that is, when the `T_MORE` flag is not set. Some transport providers also forbid zero-length TSDUs and ETSDUs. See Appendix A for a fuller explanation.

If set in *flags*, requests that the provider transmit all data that it has accumulated but not sent. The request is a local action on the provider and does not affect any similarly named protocol flag (for example, the TCP PUSH flag). This effect of setting this flag is protocol-dependent, and it may be ignored entirely by transport providers which do not support the use of this feature.

Note: The communications provider is free to collect data in a send buffer until it accumulates a sufficient amount for transmission.

By default, `t_sndv()` operates in synchronous mode and may wait if flow control restrictions prevent the data from being accepted by the local transport provider at the time the call is made. However, if `O_NONBLOCK` is set (via `t_open()` or `fcntl()`), `t_sndv()` executes in asynchronous mode, and will fail immediately if there are flow control restrictions. The process can arrange to be informed when the flow control restrictions are cleared via either `t_look()` or the EM interface.

On successful completion, `t_sndv()` returns the number of bytes accepted by the transport provider. Normally this will equal the total number of bytes to be sent, that is,

$$(iov[0].iov_len + \dots + iov[iovcount-1].iov_len)$$

However, the interface is constrained to send at most `INT_MAX` bytes in a single send. When `t_sndv()` has submitted `INT_MAX` (or lower constrained value, see the note above) bytes to the provider for a single call, this value is returned to the user. However, if `O_NONBLOCK` is set or the function is interrupted by a signal, it is possible that only part of the data has actually been accepted by the communications provider. In this case, `t_sndv()` returns a value that is less than the value of `nbytes`. If `t_sndv()` is interrupted by a signal before it could transfer data to the communications provider, it returns `-1` with `t_errno` set to `[TSYSERR]` and `errno` set to `[EINTR]`.

If the number of bytes of data in the *iov* array is zero and sending of zero octets is not supported by the underlying transport service, `t_sndv()` returns `-1` with `t_errno` set to `[TBADDDATA]`.

The size of each TSDU or ETSDU must not exceed the limits of the transport provider as specified by the current values in the TSDU or ETSDU fields in the *info* argument returned by `t_getinfo()`.

The error `[TLOOK]` is returned for asynchronous events. It is required only for an incoming disconnect event but may be returned for other events.

VALID STATES

`T_DATAXFER`, `T_INREL`

ERRORS

On failure, `t_errno` is set to one of the following:

`[TBADDDATA]`

Illegal amount of data:

- A single send was attempted specifying a TSDU (ETSDU) or fragment TSDU (ETSDU) greater than that specified by the current values of the TSDU or ETSDU fields in the *info* argument.
- A send of a zero byte TSDU (ETSDU) or zero byte fragment of a TSDU (ETSDU) is not supported by the provider.

- Multiple sends were attempted resulting in a TSDU (ETSDU) larger than that specified by the current value of the TSDU or ETSDU fields in the *info* argument — the ability of an XTI implementation to detect such an error case is implementation-dependent (see CAVEATS, below).
- *iovcount* is greater than T_IOV_MAX.

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBADFLAG]	An invalid flag was specified.
[TFLOW]	O_NONBLOCK was set, but the flow control mechanism prevented the transport provider from accepting any data at this time.
[TLOOK]	An asynchronous event has occurred on this transport endpoint.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUES

On successful completion, *t_sndv()* returns the number of bytes accepted by the transport provider. Otherwise, -1 is returned on failure and *t_errno* is set to indicate the error.

Notes:

1. In synchronous mode, if more than INT_MAX bytes of data are passed in the *iov* array, only the first INT_MAX bytes will be passed to the provider.
2. If the number of bytes accepted by the communications provider is less than the number of bytes requested, this may either indicate that O_NONBLOCK is set and the communications provider is blocked due to flow control, or that O_NONBLOCK is clear and the function was interrupted by a signal.

SEE ALSO

t_getinfo(), *t_open()*, *t_rcvv()*, *t_rcv()*, *t_snd()*.

CAVEATS

It is important to remember that the transport provider treats all users of a transport endpoint as a single user. Therefore if several processes issue concurrent *t_sndv()* or *t_snd()* calls, then the different data may be intermixed.

Multiple sends which exceed the maximum TSDU or ETSDU size may not be discovered by XTI. In this case an implementation-dependent error will result (generated by the transport provider), perhaps on a subsequent XTI call. This error may take the form of a connection abort, a [TSYSERR], a [TBADDDATA] or a [TPROTO] error.

If multiple sends which exceed the maximum TSDU or ETSDU size are detected by XTI, *t_sndv()* fails with [TBADDDATA].

NAME

t_sndvudata — send a data unit from one or more noncontiguous buffers

SYNOPSIS

```
#include <xti.h>

int t_sndvudata(
    int fd,
    struct t_unitdata *unitdata,
    struct t_iovec *iov,
    unsigned int iovcount)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>unitdata->addr.maxlen</i>	/	=
<i>unitdata->addr.len</i>	x	=
<i>unitdata->addr.buf</i>	x(x)	=(=)
<i>unitdata->opt.maxlen</i>	/	=
<i>unitdata->opt.len</i>	x	=
<i>unitdata->opt.buf</i>	?(?)	=(=)
<i>unitdata->udata.maxlen</i>	/	=
<i>unitdata->udata.len</i>	/	=
<i>unitdata->udata.buf</i>	/	=
<i>iov[0].iov_base</i>	x(x)	=(=)
<i>iov[0].iov_len</i>	x	=
. . . .		
<i>iov[iovcount-1].iov_base</i>	x(x)	=(=)
<i>iov[iovcount-1].iov_len</i>	x	=
<i>iovcount</i>	x	/

This function is used in connectionless mode to send a data unit to another transport user. The argument *fd* identifies the local transport endpoint through which data will be sent, *iovcount* contains the number of non-contiguous *udata* buffers and is limited to an implementation-defined value given by T_IOV_MAX which is at least 16, and *unitdata* points to a **t_unitdata** structure containing the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
```

If the limit on *iovcount* is exceeded, the function fails with [TBADDDATA].

In **unitdata**, *addr* specifies the protocol address of the destination user, and *opt* identifies options that the user wants associated with this request. The *udata* field is not used. The user may choose not to specify what protocol options are associated with the transfer by setting the *len* field of *opt* to zero. In this case, the provider may use default options.

The data to be sent is identified by *iov[0]* through *iov[iovcount-1]*.

Note: The limit on the total number of bytes available in all buffers passed (that is: *iov(0).iov_len + . . . + iov(iovcount-1).iov_len*) may be constrained by implementation limits. If no other constraint applies, it will be limited by [INT_MAX]. In practice, the availability of memory to an application is likely to impose a lower limit on the amount of data that can be sent or received using scatter/gather functions.

By default, *t_sndvudata()* operates in synchronous mode and may wait if flow control restrictions prevent the data from being accepted by the local transport provider at the time the call is made. However, if `O_NONBLOCK` is set (via *t_open()* or *fcntl()*), *t_sndvudata()* executes in asynchronous mode and will fail under such conditions. The process can arrange to be notified of the clearance of a flow control restriction via either *t_look()* or the EM interface.

If the amount of data specified in *iov[0]* through *iov[iovcount-1]* exceeds the TSDU size as returned in the *tsdu* field of the *info* argument of *t_open()* or *t_getinfo()*, or is zero and sending of zero octets is not supported by the underlying transport service, a [TBADDDATA] error is generated. If *t_sndvudata()* is called before the destination user has activated its transport endpoint (see *t_bind()*), the data unit may be discarded.

If it is not possible for the transport provider to immediately detect the conditions that cause the errors [TBADDDADDR] and [TBADDOPT], these errors will alternatively be returned by *t_rcvuderr()*. An application must therefore be prepared to receive these errors in both of these ways.

VALID STATES

T_IDLE

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADDDADDR]	The specified protocol address was in an incorrect format or contained illegal information.
[TBADDDATA]	Illegal amount of data. <ul style="list-style-type: none"> • A single send was attempted specifying a TSDU greater than that specified in the <i>info</i> argument, or a send of a zero byte TSDU is not supported by the provider. • <i>iovcount</i> is greater than T_IOV_MAX.
[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBADDOPT]	The specified options were in an incorrect format or contained illegal information.
[TFLOW]	<code>O_NONBLOCK</code> was set, but the flow control mechanism prevented the transport provider from accepting any data at this time.
[TLOOK]	An asynchronous event has occurred on this transport endpoint.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUES

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

fcntl(), *t_alloc()*, *t_open()*, *t_rcvudata()*, *t_rcvvudata()*, *t_rcvuderr()*, *t_sndudata()*.

NAME

`t_strerror` - produce an error message string

SYNOPSIS

```
#include <xti.h>
```

```
const char *t_strerror(
    int errnum)
```

DESCRIPTION

Parameters	Before call	After call
<i>errnum</i>	x	/

The `t_strerror()` function maps the error number in *errnum* that corresponds to an XTI error to a language-dependent error message string and returns a pointer to the string. The string pointed to will not be modified by the program, but may be overwritten by a subsequent call to the `t_strerror` function. The string is not terminated by a newline character. The language for error message strings written by `t_strerror()` is that of the current locale. If it is English, the error message string describing the value in *t_errno* may be derived from the comments following the *t_errno* codes defined in `<xti.h>`. If an error code is unknown, and the language is English, `t_strerror()` returns the string:

```
"<error>: error unknown"
```

where `<error>` is the error number supplied as input. In other languages, an equivalent text is provided.

VALID STATES

ALL - apart from T_UNINIT

RETURN VALUE

The function `t_strerror()` returns a pointer to the generated message string.

SEE ALSO

`t_error()`

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_sync - synchronise transport library

SYNOPSIS

#include <xti.h>

```
int t_sync(
    int fd)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/

For the transport endpoint specified by *fd*, *t_sync()* synchronises the data structures managed by the transport library with information from the underlying transport provider. In doing so, it can convert an uninitialised file descriptor (obtained via *open()*, *dup()* or as a result of a *fork()* and *exec()*) to an initialised transport endpoint, assuming that the file descriptor referenced a transport endpoint, by updating and allocating the necessary library data structures. This function also allows two cooperating processes to synchronise their interaction with a transport provider.

For example, if a process forks a new process and issues an *exec()*, the new process must issue a *t_sync()* to build the private library data structure associated with a transport endpoint and to synchronise the data structure with the relevant provider information.

It is important to remember that the transport provider treats all users of a transport endpoint as a single user. If multiple processes are using the same endpoint, they should coordinate their activities so as not to violate the state of the transport endpoint. The function *t_sync()* returns the current state of the transport endpoint to the user, thereby enabling the user to verify the state before taking further action. This coordination is only valid among cooperating processes; it is possible that a process or an incoming event could change the endpoint's state *after* a *t_sync()* is issued.

If the transport endpoint is undergoing a state transition when *t_sync()* is called, the function will fail.

VALID STATES

ALL - apart from T_UNINIT

ERRORSOn failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint. This error may be returned when the <i>fd</i> has been previously closed or an erroneous number may have been passed to the call.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSTATECHNG]	The transport endpoint is undergoing a state change.
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

On successful completion, the state of the transport endpoint is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error. The state returned is one of the following:

T_UNBND	Unbound.
T_IDLE	Idle.
T_OUTCON	Outgoing connection pending.
T_INCON	Incoming connection pending.
T_DATAXFER	Data transfer.
T_OUTREL	Outgoing orderly release (waiting for an orderly release indication).
T_INREL	Incoming orderly release (waiting for an orderly release request).

SEE ALSO

dup(), exec(), fork(), open().

CHANGE HISTORY**Issue 4**

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

NAME

t_sysconf — get configurable XTI variables

SYNOPSIS

```
#include <xti.h>

int t_sysconf(
    int name);
```

DESCRIPTION

Parameters	Before call	After call
<i>name</i>	x	/

The *t_sysconf()* function provides a method for the application to determine the current value of configurable and implementation-dependent XTI limits or options.

The *name* argument represents the XTI system variable to be queried. The following table lists the minimal set of XTI system variables from <xti.h> that can be returned by *t_sysconf()*, and the symbolic constants, defined in <xti.h> that are the corresponding values used for *name*.

Variable	Value of Name
T_IOV_MAX	_SC_T_IOV_MAX

VALID STATES

All.

ERRORS

On failure, *t_errno* is set to the following:

[TBADFLAG] *name* has an invalid value.

RETURN VALUES

If *name* is valid, *t_sysconf()* returns the value of the requested limit/option (which might be -1) and leaves *t_errno* unchanged. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_rcvv(), *t_rcvvudata()*, *t_sndv()*, *t_sndvudata()*.

NAME

t_unbind - disable a transport endpoint

SYNOPSIS

```
#include <xti.h>
```

```
int t_unbind(  
    int fd)
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/

The *t_unbind()* function disables the transport endpoint specified by *fd* which was previously bound by *t_bind()*. On completion of this call, no further data or events destined for this transport endpoint will be accepted by the transport provider. An endpoint which is disabled by using *t_unbind()* can be enabled by a subsequent call to *t_bind()*.

VALID STATES

T_IDLE

ERRORS

On failure, *t_errno* is set to one of the following:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TLOOK]	An asynchronous event has occurred on this transport endpoint.
[TOUTSTATE]	The communications endpoint referenced by <i>fd</i> is not in one of the states in which a call to this function is valid.
[TPROTO]	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (<i>t_errno</i>).
[TSYSERR]	A system error has occurred during execution of this function.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t_errno* is set to indicate an error.

SEE ALSO

t_bind().

CHANGE HISTORY

Issue 4

The **SYNOPSIS** section is placed in the form of a standard C function prototype.

Sockets Interfaces

This chapter gives an overview of the Sockets interfaces and includes functions, macros and external variables to support portability at the C-language source level.

The associated headers are documented in Chapter 9.

8.1 Sockets Overview

All network protocols are associated with a specific protocol family. A protocol family provides basic services to the protocol implementation to allow it to function within a specific network environment. These services can include packet fragmentation and reassembly, routing, addressing, and basic transport. A protocol family can support multiple methods of addressing, though the current protocol implementations do not. A protocol family normally comprises a number of protocols, one per socket type. It is not required that a protocol family support all socket types. A protocol family can contain multiple protocols supporting the same socket abstraction.

A protocol supports one of the socket abstractions detailed in the manual page for the *socket()* function. A specific protocol can be accessed either by creating a socket of the appropriate type and protocol family, or by requesting the protocol explicitly when creating a socket. Protocols normally accept only one type of address format, usually determined by the addressing structure inherent in the design of the protocol family and network architecture. Certain semantics of the basic socket abstractions are protocol specific. All protocols are expected to support the basic model for their particular socket type, but can, in addition, provide nonstandard facilities or extensions to a mechanism. For example, a protocol supporting the SOCK_STREAM abstraction can allow more than one byte of out-of-band data to be transmitted per out-of-band message.

This specification covers local UNIX connections and Internet protocols.

Addressing

Associated with each address family is an address format. All network addresses adhere to a general structure, called a **sockaddr**. The length of the structure varies according to the address family.

Routing

Sockets provides packet routing facilities. A routing information database is maintained, which is used in selecting the appropriate network interface when transmitting packets.

Interfaces

Each network interface in a system corresponds to a path through which messages can be sent and received. A network interface usually has a hardware device associated with it, though certain interfaces such as the loopback interface do not.

NAME

accept — accept a new connection on a socket

SYNOPSIS

```
#include <sys/socket.h>

int accept (int socket, struct sockaddr *address,
           socklen_t *address_len);
```

DESCRIPTION

The *accept()* function extracts the first connection on the queue of pending connections, creates a new socket with the same socket type protocol and address family as the specified socket, and allocates a new file descriptor for that socket.

The function takes the following arguments:

<i>socket</i>	Specifies a socket that was created with <i>socket()</i> , has been bound to an address with <i>bind()</i> , and has issued a successful call to <i>listen()</i> .
<i>address</i>	Either a null pointer, or a pointer to a sockaddr structure where the address of the connecting socket will be returned.
<i>address_len</i>	Points to a socklen_t which on input specifies the length of the supplied sockaddr structure, and on output specifies the length of the stored address.

If *address* is not a null pointer, the address of the peer for the accepted connection is stored in the **sockaddr** structure pointed to by *address*, and the length of this address is stored in the object pointed to by *address_len*.

If the actual length of the address is greater than the length of the supplied **sockaddr** structure, the stored address will be truncated.

If the protocol permits connections by unbound clients, and the peer is not bound, then the value stored in the object pointed to by *address* is unspecified.

If the listen queue is empty of connection requests and O_NONBLOCK is not set on the file descriptor for the socket, *accept()* will block until a connection is present. If the *listen()* queue is empty of connection requests and O_NONBLOCK is set on the file descriptor for the socket, *accept()* will fail and set *errno* to [EAGAIN] or [EWOULDBLOCK].

The accepted socket cannot itself accept more connections. The original socket remains open and can accept more connections.

RETURN VALUE

Upon successful completion, *accept()* returns the nonnegative file descriptor of the accepted socket. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The *accept()* function will fail if:

[EAGAIN] or [EWOULDBLOCK]	O_NONBLOCK is set for the socket file descriptor and no connections are present to be accepted.
[EBADF]	The <i>socket</i> argument is not a valid file descriptor.
[ECONNABORTED]	A connection has been aborted.
[EFAULT]	The <i>address</i> or <i>address_len</i> parameter can not be accessed or written.

[EINTR]	The <i>accept()</i> function was interrupted by a signal that was caught before a valid connection arrived.
[EINVAL]	The <i>socket</i> is not accepting connections.
[EMFILE]	{OPEN_MAX} file descriptors are currently open in the calling process.
[ENFILE]	The maximum number of file descriptors in the system are already open.
[ENOTSOCK]	The <i>socket</i> argument does not refer to a socket.
[EOPNOTSUPP]	The socket type of the specified socket does not support accepting connections.

The *accept()* function may fail if:

[ENOBUFS]	No buffer space is available.
[ENOMEM]	There was insufficient memory available to complete the operation.
[ENOSR]	There was insufficient STREAMS resources available to complete the operation.
[EPROTO]	A protocol error has occurred; for example, the STREAMS protocol stack has not been initialised.

APPLICATION USAGE

When a connection is available, *select()* will indicate that the file descriptor for the socket is ready for reading.

SEE ALSO

bind(), *connect()*, *listen()*, *socket()*, <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

`bind` — bind a name to a socket

SYNOPSIS

```
#include <sys/socket.h>

int bind(int socket, const struct sockaddr *address,
         socklen_t address_len);
```

DESCRIPTION

The `bind()` function assigns an *address* to an unnamed socket. Sockets created with `socket()` function are initially unnamed; they are identified only by their address family.

The function takes the following arguments:

<i>socket</i>	Specifies the file descriptor of the socket to be bound.
<i>address</i>	Points to a sockaddr structure containing the address to be bound to the socket. The length and format of the address depend on the address family of the socket.
<i>address_len</i>	Specifies the length of the sockaddr structure pointed to by the <i>address</i> argument.

The socket in use may require the process to have appropriate privileges to use the `bind()` function.

RETURN VALUE

Upon successful completion, `bind()` returns 0. Otherwise, `-1` is returned and `errno` is set to indicate the error.

ERRORS

The `bind()` function will fail if:

[EADDRINUSE]	The specified address is already in use.
[EADDRNOTAVAIL]	The specified address is not available from the local machine.
[EAFNOSUPPORT]	The specified address is not a valid address for the address family of the specified socket.
[EBADF]	The <i>socket</i> argument is not a valid file descriptor.
[EFAULT]	The <i>address</i> argument can not be accessed.
[EINVAL]	The socket is already bound to an address, and the protocol does not support binding to a new address; or the socket has been shut down.
[ENOTSOCK]	The <i>socket</i> argument does not refer to a socket.
[EOPNOTSUPP]	The socket type of the specified socket does not support binding to an address.

If the address family of the socket is `AF_UNIX`, then `bind()` will fail if:

[EACCES]	A component of the path prefix denies search permission, or the requested name requires writing in a directory with a mode that denies write permission.
[EDESTADDRREQ] or [EISDIR]	The <i>address</i> argument is a null pointer.

[EIO]	An I/O error occurred.
[ELOOP]	Too many symbolic links were encountered in translating the pathname in <i>address</i> .
[ENAMETOOLONG]	A component of a pathname exceeded {NAME_MAX} characters, or an entire pathname exceeded {PATH_MAX} characters.
[ENOENT]	A component of the pathname does not name an existing file or the pathname is an empty string.
[ENOTDIR]	A component of the path prefix of the pathname in <i>address</i> is not a directory.
[EROFS]	The name would reside on a read-only filesystem.

The *bind()* function may fail if:

[EACCES]	The specified address is protected and the current user does not have permission to bind to it.
[EINVAL]	The <i>address_len</i> argument is not a valid length for the address family.
[EISCONN]	The socket is already connected.
[ENAMETOOLONG]	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds {PATH_MAX}.
[ENOBUFS]	Insufficient resources were available to complete the call.
[ENOSR]	There were insufficient STREAMS resources for the operation to complete.

APPLICATION USAGE

An application program can retrieve the assigned socket name with the *getsockname()* function.

SEE ALSO

connect(), *getsockname()*, *listen()*, *socket()*, <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

close — close a file descriptor

Note: The **XSH** specification contains the basic definition of this interface. The following additional information pertains to Sockets.

DESCRIPTION

If *fdes* refers to a socket, *close()* causes the socket to be destroyed. If the socket is connection-mode, and the `SO_LINGER` option is set for the socket with non-zero linger time, and the socket has untransmitted data, then *close()* will block for up to the current linger interval until all data is transmitted.

CHANGE HISTORY

First released in Issue 4.

NAME

connect — connect a socket

SYNOPSIS

```
#include <sys/socket.h>

int connect(int socket, const struct sockaddr *address,
            socklen_t address_len);
```

DESCRIPTION

The *connect()* function requests a connection to be made on a socket. The function takes the following arguments:

<i>socket</i>	Specifies the file descriptor associated with the socket.
<i>address</i>	Points to a sockaddr structure containing the peer address. The length and format of the address depend on the address family of the socket.
<i>address_len</i>	Specifies the length of the sockaddr structure pointed to by the <i>address</i> argument.

If the socket has not already been bound to a local address, *connect()* will bind it to an address which, unless the socket's address family is AF_UNIX, is an unused local address.

If the initiating socket is not connection-mode, then *connect()* sets the socket's peer address, but no connection is made. For SOCK_DGRAM sockets, the peer address identifies where all datagrams are sent on subsequent *send()* calls, and limits the remote sender for subsequent *recv()* calls. If *address* is a null address for the protocol, the socket's peer address will be reset.

If the initiating socket is connection-mode, then *connect()* attempts to establish a connection to the address specified by the *address* argument.

If the connection cannot be established immediately and O_NONBLOCK is not set for the file descriptor for the socket, *connect()* will block for up to an unspecified timeout interval until the connection is established. If the timeout interval expires before the connection is established, *connect()* will fail and the connection attempt will be aborted. If *connect()* is interrupted by a signal that is caught while blocked waiting to establish a connection, *connect()* will fail and set *errno* to [EINTR], but the connection request will not be aborted, and the connection will be established asynchronously.

If the connection cannot be established immediately and O_NONBLOCK is set for the file descriptor for the socket, *connect()* will fail and set *errno* to [EINPROGRESS], but the connection request will not be aborted, and the connection will be established asynchronously. Subsequent calls to *connect()* for the same socket, before the connection is established, will fail and set *errno* to [EALREADY].

When the connection has been established asynchronously, *select()* and *poll()* will indicate that the file descriptor for the socket is ready for writing.

The socket in use may require the process to have appropriate privileges to use the *connect()* function.

RETURN VALUE

Upon successful completion, *connect()* returns 0. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The *connect()* function will fail if:

[EADDRNOTAVAIL]	The specified address is not available from the local machine.
[EAFNOSUPPORT]	The specified address is not a valid address for the address family of the specified socket.
[EALREADY]	A connection request is already in progress for the specified socket.
[EBADF]	The <i>socket</i> argument is not a valid file descriptor.
[ECONNREFUSED]	The target address was not listening for connections or refused the connection request.
[EFAULT]	The address parameter can not be accessed.
[EINPROGRESS]	O_NONBLOCK is set for the file descriptor for the socket and the connection cannot be immediately established; the connection will be established asynchronously.
[EINTR]	The attempt to establish a connection was interrupted by delivery of a signal that was caught; the connection will be established asynchronously.
[EISCONN]	The specified socket is connection-mode and is already connected.
[ENETUNREACH]	No route to the network is present.
[ENOTSOCK]	The <i>socket</i> argument does not refer to a socket.
[EPROTOTYPE]	The specified address has a different type than the socket bound to the specified peer address.
[ETIMEDOUT]	The attempt to connect timed out before a connection was made.
If the address family of the socket is AF_UNIX, then <i>connect()</i> will fail if:	
[EIO]	An I/O error occurred while reading from or writing to the file system.
[ELOOP]	Too many symbolic links were encountered in translating the pathname in <i>address</i> .
[ENAMETOOLONG]	A component of a pathname exceeded {NAME_MAX} characters, or an entire pathname exceeded {PATH_MAX} characters.
[ENOENT]	A component of the pathname does not name an existing file or the pathname is an empty string.
[ENOTDIR]	A component of the path prefix of the pathname in <i>address</i> is not a directory.
The <i>connect()</i> function may fail if:	
[EACCES]	Search permission is denied for a component of the path prefix; or write access to the named socket is denied.
[EADDRINUSE]	Attempt to establish a connection that uses addresses that are already in use.
[ECONNRESET]	Remote host reset the connection request.
[EHOSTUNREACH]	The destination host cannot be reached (probably because the host is down or a remote router cannot reach it).
[EINVAL]	The <i>address_len</i> argument is not a valid length for the address family; or invalid address family in sockaddr structure.

- [ENAMETOOLONG] Pathname resolution of a symbolic link produced an intermediate result whose length exceeds {PATH_MAX}.
- [ENETDOWN] The local interface used to reach the destination is down.
- [ENOBUFS] No buffer space is available.
- [ENOSR] There were insufficient STREAMS resources available to complete the operation.
- [EOPNOTSUPP] The socket is listening and can not be connected.

APPLICATION USAGE

If *connect()* fails, the state of the socket is unspecified. Portable applications should close the file descriptor and create a new socket before attempting to reconnect.

SEE ALSO

accept(), *bind()*, *close()*, *getsockname()*, *poll()*, *select()*, *send()*, *shutdown()*, *socket()*, <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

fcntl — file control

Note: The XSH specification contains the basic definition of this interface. The following additional information pertains to Sockets.

DESCRIPTION

The following additional values for *cmd* are defined in <fcntl.h>:

F_GETOWN If *fdes* refers to a socket, get the process or process group ID specified to receive SIGURG signals when out-of-band data is available. Positive values indicate a process ID; negative values, other than -1, indicate a process group ID. If *fdes* does not refer to a socket, the results are unspecified.

F_SETOWN If *fdes* refers to a socket, set the process or process group ID specified to receive SIGURG signals when out-of-band data is available, using the value of the third argument, *arg*, taken as type **int**. Positive values indicate a process ID; negative values, other than -1, indicate a process group ID. If *fdes* does not refer to a socket, the results are unspecified.

RETURN VALUE

Upon successful completion, the value returned depends on *cmd* as follows:

F_GETOWN Value of the socket owner process or process group; this will not be -1.

F_SETOWN Value other than -1.

CHANGE HISTORY

First released in Issue 4.

NAME

fgetpos — get current file position information

Note: The **XSH** specification contains the basic definition of this interface. The following additional information pertains to Sockets.

ERRORS

The *fgetpos()* function may fail if:

[ESPIPE] The file descriptor underlying *stream* is associated with a socket.

CHANGE HISTORY

First released in Issue 4.

NAME

fsetpos — set current file position

Note: The **XSH** specification contains the basic definition of this interface. The following additional information pertains to Sockets.

ERRORS

The *fsetpos()* function may fail if:

[ESPIPE] The file descriptor underlying *stream* is associated with a socket.

CHANGE HISTORY

First released in Issue 4.

NAME

ftell — return a file offset in a stream

Note: The **XSH** specification contains the basic definition of this interface. The following additional information pertains to Sockets.

ERRORS

The *ftell()* function may fail if:

[ESPIPE] The file descriptor underlying *stream* is associated with a socket.

CHANGE HISTORY

First released in Issue 4.

NAME

getpeername — get the name of the peer socket

SYNOPSIS

```
#include <sys/socket.h>

int getpeername(int socket, struct sockaddr *address,
                socklen_t *address_len);
```

DESCRIPTION

The *getpeername()* function retrieves the peer address of the specified socket, stores this address in the **sockaddr** structure pointed to by the *address* argument, and stores the length of this address in the object pointed to by the *address_len* argument.

If the actual length of the address is greater than the length of the supplied **sockaddr** structure, the stored address will be truncated.

If the protocol permits connections by unbound clients, and the peer is not bound, then the value stored in the object pointed to by *address* is unspecified.

RETURN VALUE

Upon successful completion, 0 is returned. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The *getpeername()* function will fail if:

- | | |
|--------------|--|
| [EBADF] | The <i>socket</i> argument is not a valid file descriptor. |
| [EFAULT] | The <i>address</i> or <i>address_len</i> parameter can not be accessed or written. |
| [EINVAL] | The socket has been shut down. |
| [ENOTCONN] | The socket is not connected or otherwise has not had the peer prespecified. |
| [ENOTSOCK] | The <i>socket</i> argument does not refer to a socket. |
| [EOPNOTSUPP] | The operation is not supported for the socket protocol. |

The *getpeername()* function may fail if:

- | | |
|-----------|--|
| [ENOBUFS] | Insufficient resources were available in the system to complete the call. |
| [ENOSR] | There were insufficient STREAMS resources available for the operation to complete. |

SEE ALSO

accept(), *bind()*, *getsockname()*, *socket()*, <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

getsockname — get the socket name

SYNOPSIS

```
#include <sys/socket.h>

int getsockname(int socket, struct sockaddr *address,
                socklen_t *address_len);
```

DESCRIPTION

The *getsockname()* function retrieves the locally-bound name of the specified socket, stores this address in the **sockaddr** structure pointed to by the *address* argument, and stores the length of this address in the object pointed to by the *address_len* argument.

If the actual length of the address is greater than the length of the supplied **sockaddr** structure, the stored address will be truncated.

If the socket has not been bound to a local name, the value stored in the object pointed to by *address* is unspecified.

RETURN VALUE

Upon successful completion, 0 is returned, the *address* argument points to the address of the socket, and the *address_len* argument points to the length of the address. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The *getsockname()* function will fail:

- | | |
|--------------|--|
| [EBADF] | The <i>socket</i> argument is not a valid file descriptor. |
| [EFAULT] | The <i>address</i> or <i>address_len</i> parameter can not be accessed or written. |
| [ENOTSOCK] | The <i>socket</i> argument does not refer to a socket. |
| [EOPNOTSUPP] | The operation is not supported for this socket's protocol. |

The *getsockname()* function may fail if:

- | | |
|-----------|--|
| [EINVAL] | The socket has been shut down. |
| [ENOBUFS] | Insufficient resources were available in the system to complete the call. |
| [ENOSR] | There were insufficient STREAMS resources available for the operation to complete. |

SEE ALSO

accept(), *bind()*, *getpeername()*, *socket()*, <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

getsockopt — get the socket options

SYNOPSIS

```
#include <sys/socket.h>

int getsockopt(int socket, int level, int option_name,
               void *option_value, socklen_t *option_len);
```

DESCRIPTION

The *getsockopt()* function retrieves the value for the option specified by the *option_name* argument for the socket specified by the *socket* argument. If the size of the option value is greater than *option_len*, the value stored in the object pointed to by the *option_value* argument will be silently truncated. Otherwise, the object pointed to by the *option_len* argument will be modified to indicate the actual length of the value.

The *level* argument specifies the protocol level at which the option resides. To retrieve options at the socket level, specify the *level* argument as `SOL_SOCKET`. To retrieve options at other levels, supply the appropriate protocol number for the protocol controlling the option. For example, to indicate that an option will be interpreted by the TCP (Transport Control Protocol), set *level* to the protocol number of TCP, as defined in the `<netinet/in.h>` header, or as determined by using *getprotobyname()* function.

The socket in use may require the process to have appropriate privileges to use the *getsockopt()* function.

The *option_name* argument specifies a single option to be retrieved. It can be one of the following values defined in `<sys/socket.h>`:

SO_DEBUG	Reports whether debugging information is being recorded. This option stores an int value. This is a boolean option.
SO_ACCEPTCONN	Reports whether socket listening is enabled. This option stores an int value.
SO_BROADCAST	Reports whether transmission of broadcast messages is supported, if this is supported by the protocol. This option stores an int value. This is a boolean option.
SO_REUSEADDR	Reports whether the rules used in validating addresses supplied to <i>bind()</i> should allow reuse of local addresses, if this is supported by the protocol. This option stores an int value. This is a boolean option.
SO_KEEPALIVE	Reports whether connections are kept active with periodic transmission of messages, if this is supported by the protocol. If the connected socket fails to respond to these messages, the connection is broken and processes writing to that socket are notified with a SIGPIPE signal. This option stores an int value. This is a boolean option.
SO_LINGER	Reports whether the socket lingers on <i>close()</i> if data is present. If <code>SO_LINGER</code> is set, the system blocks the process during <i>close()</i> until it can transmit the data or until the end of the interval indicated by the linger member, whichever comes first. If <code>SO_LINGER</code> is not specified, and <i>close()</i> is issued, the system handles the call in a way that allows the process to continue as quickly as possible. This option stores a linger structure.

SO_OOBINLINE	Reports whether the socket leaves received out-of-band data (data marked urgent) in line. This option stores an int value. This is a boolean option.
SO_SNDBUF	Reports send buffer size information. This option stores an int value.
SO_RCVBUF	Reports receive buffer size information. This option stores an int value.
SO_ERROR	Reports information about error status and clears it. This option stores an int value.
SO_TYPE	Reports the socket type. This option stores an int value.
SO_DONTROUTE	Reports whether outgoing messages bypass the standard routing facilities. The destination must be on a directly-connected network, and messages are directed to the appropriate network interface according to the destination address. The effect, if any, of this option depends on what protocol is in use. This option stores an int value. This is a boolean option.
SO_RCVLOWAT	Reports the minimum number of bytes to process for socket input operations. The default value for SO_RCVLOWAT is 1. If SO_RCVLOWAT is set to a larger value, blocking receive calls normally wait until they have received the smaller of the low water mark value or the requested amount. (They may return less than the low water mark if an error occurs, a signal is caught, or the type of data next in the receive queue is different than that returned, e.g. out of band data). This option stores an int value. Note that not all implementations allow this option to be retrieved.
SO_RCVTIMEO	Reports the timeout value for input operations. This option stores a timeval structure with the number of seconds and microseconds specifying the limit on how long to wait for an input operation to complete. If a receive operation has blocked for this much time without receiving additional data, it returns with a partial count or errno set to [EAGAIN] or [EWOULDBLOCK] if no data were received. The default for this option is zero, which indicates that a receive operation will not time out. Note that not all implementations allow this option to be retrieved.
SO_SNDLOWAT	Reports the minimum number of bytes to process for socket output operations. Non-blocking output operations will process no data if flow control does not allow the smaller of the send low water mark value or the entire request to be processed. This option stores an int value. Note that not all implementations allow this option to be retrieved.
SO_SNDTIMEO	Reports the timeout value specifying the amount of time that an output function blocks because flow control prevents data from being sent. If a send operation has blocked for this time, it returns with a partial count or with errno set to [EAGAIN] or [EWOULDBLOCK] if no data were sent. The default for this option is zero, which indicates that a send operation will not time out. The option stores a timeval structure. Note that not all implementations allow this option to be retrieved.

For boolean options, a zero value indicates that the option is disabled and a non-zero value indicates that the option is enabled.

Options at other protocol levels vary in format and name.

The socket in use may require the process to have appropriate privileges to use the *getsockopt()* function.

RETURN VALUE

Upon successful completion, *getsockopt()* returns 0. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The *getsockopt()* function will fail if:

- [EBADF] The *socket* argument is not a valid file descriptor.
- [EFAULT] The *option_value* or *option_len* parameter can not be accessed or written.
- [EINVAL] The specified option is invalid at the specified socket level.
- [ENOPROTOOPT] The option is not supported by the protocol.
- [ENOTSOCK] The *socket* argument does not refer to a socket.

The *getsockopt()* function may fail if:

- [EACCES] The calling process does not have the appropriate privileges.
- [EINVAL] The socket has been shut down.
- [ENOBUFS] Insufficient resources are available in the system to complete the call.
- [ENOSR] There were insufficient STREAMS resources available for the operation to complete.

SEE ALSO

bind(), *close()*, *endprotoent()*, *setsockopt()*, *socket()*, <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

listen — listen for socket connections and limit the queue of incoming connections

SYNOPSIS

```
#include <sys/socket.h>

int listen(int socket, int backlog);
```

DESCRIPTION

The *listen()* function marks a connection-mode socket, specified by the *socket* argument, as accepting connections, and limits the number of outstanding connections in the socket's listen queue to the value specified by the *backlog* argument.

If *listen()* is called with a *backlog* argument value that is less than 0, the function sets the length of the socket's listen queue to 0.

The implementation may include incomplete connections in the queue subject to the queue limit. The implementation may also increase the specified queue limit internally if it includes such incomplete connections in the queue subject to this limit.

Implementations may limit the length of the socket's listen queue. If *backlog* exceeds the implementation-dependent maximum queue length, the length of the socket's listen queue will be set to the maximum supported value.

The socket in use may require the process to have appropriate privileges to use the *listen()* function.

RETURN VALUE

Upon successful completions, *listen()* returns 0. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The *listen()* function will fail if:

- | | |
|----------------|---|
| [EBADF] | The <i>socket</i> argument is not a valid file descriptor. |
| [EDESTADDRREQ] | The socket is not bound to a local address, and the protocol does not support listening on an unbound socket. |
| [EINVAL] | The <i>socket</i> is already connected. |
| [ENOTSOCK] | The <i>socket</i> argument does not refer to a socket. |
| [EOPNOTSUPP] | The socket protocol does not support <i>listen()</i> . |

The *listen()* function may fail if:

- | | |
|-----------|--|
| [EACCES] | The calling process does not have the appropriate privileges. |
| [EINVAL] | The <i>socket</i> has been shut down. |
| [ENOBUFS] | Insufficient resources are available in the system to complete the call. |

SEE ALSO

accept(), *connect()*, *socket()*, <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

lseek — move read/write file offset

Note: The **XSH** specification contains the basic definition of this interface. The following additional information pertains to Sockets.

ERRORS

The *lseek()* function will fail if:

[ESPIPE] The file descriptor underlying *stream* is associated with a socket.

CHANGE HISTORY

First released in Issue 4.

NAME

poll — input/output multiplexing

Note: The **XSH** specification contains the basic definition of this interface. The following additional information pertains to Sockets.

DESCRIPTION

The *poll()* function supports sockets.

A file descriptor for a socket that is listening for connections will indicate that it is ready for reading, once connections are available. A file descriptor for a socket that is connecting asynchronously will indicate that it is ready for writing, once a connection has been established.

CHANGE HISTORY

First released in Issue 4.

NAME

read, readv — read from file

Note: The **XSH** specification contains the basic definition of this interface. The following additional information pertains to Sockets.

DESCRIPTION

If *files* refers to a socket, *read()* is equivalent to *recv()* with no flags set.

CHANGE HISTORY

First released in Issue 4.

NAME

recv — receive a message from a connected socket

SYNOPSIS

```
#include <sys/socket.h>

ssize_t recv(int socket, void *buffer, size_t length, int flags);
```

DESCRIPTION

The *recv()* function receives a message from a connection-mode or connectionless-mode socket. It is normally used with connected sockets because it does not permit the application to retrieve the source address of received data. The function takes the following arguments:

<i>socket</i>	Specifies the socket file descriptor.
<i>buffer</i>	Points to a buffer where the message should be stored.
<i>length</i>	Specifies the length in bytes of the buffer pointed to by the <i>buffer</i> argument.
<i>flags</i>	Specifies the type of message reception. Values of this argument are formed by logically OR'ing zero or more of the following values: <ul style="list-style-type: none"> MSG_PEEK Peeks at an incoming message. The data is treated as unread and the next <i>recv()</i> or similar function will still return this data. MSG_OOB Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific. MSG_WAITALL Requests that the function block until the full amount of data requested can be returned. The function may return a smaller amount of data if a signal is caught, if the connection is terminated, if MSG_PEEK was specified, or if an error is pending for the socket.

The *recv()* function returns the length of the message written to the buffer pointed to by the *buffer* argument. For message-based sockets such as SOCK_DGRAM and SOCK_SEQPACKET, the entire message must be read in a single operation. If a message is too long to fit in the supplied buffer, and MSG_PEEK is not set in the *flags* argument, the excess bytes are discarded. For stream-based sockets such as SOCK_STREAM, message boundaries are ignored. In this case, data is returned to the user as soon as it becomes available, and no data is discarded.

If the MSG_WAITALL flag is not set, data will be returned only up to the end of the first message.

If no messages are available at the socket and O_NONBLOCK is not set on the socket's file descriptor, *recv()* blocks until a message arrives. If no messages are available at the socket and O_NONBLOCK is set on the socket's file descriptor, *recv()* fails and sets *errno* to [EAGAIN] or [EWOULDBLOCK].

RETURN VALUE

Upon successful completion, *recv()* returns the length of the message in bytes. If no messages are available to be received and the peer has performed an orderly shutdown, *recv()* returns 0. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The `recv()` function will fail if:

[EAGAIN] or [EWOULDBLOCK]	The socket's file descriptor is marked <code>O_NONBLOCK</code> and no data is waiting to be received; or <code>MSG_OOB</code> is set and no out-of-band data is available and either the socket's file descriptor is marked <code>O_NONBLOCK</code> or the socket does not support blocking to await out-of-band data.
[EBADF]	The <i>socket</i> argument is not a valid file descriptor.
[ECONNRESET]	A connection was forcibly closed by a peer.
[EFAULT]	The <i>buffer</i> parameter can not be accessed or written.
[EINTR]	The <code>recv()</code> function was interrupted by a signal that was caught, before any data was available.
[EINVAL]	The <code>MSG_OOB</code> flag is set and no out-of-band data is available.
[ENOTCONN]	A receive is attempted on a connection-mode socket that is not connected.
[ENOTSOCK]	The <i>socket</i> argument does not refer to a socket.
[EOPNOTSUPP]	The specified flags are not supported for this socket type or protocol.
[ETIMEDOUT]	The connection timed out during connection establishment, or due to a transmission timeout on active connection.

The `recv()` function may fail if:

[EIO]	An I/O error occurred while reading from or writing to the file system.
[ENOBUFS]	Insufficient resources were available in the system to perform the operation.
[ENOMEM]	Insufficient memory was available to fulfill the request.
[ENOSR]	There were insufficient STREAMS resources available for the operation to complete.

APPLICATION USAGE

The `recv()` function is identical to `recvfrom()` with a zero *address_len* argument, and to `read()` if no flags are used.

The `select()` and `poll()` functions can be used to determine when data is available to be received.

SEE ALSO

`poll()`, `read()`, `recvmsg()`, `recvfrom()`, `select()`, `send()`, `sendmsg()`, `sendto()`, `shutdown()`, `socket()`, `write()`, <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

recvfrom — receive a message from a socket

SYNOPSIS

```
#include <sys/socket.h>

ssize_t recvfrom(int socket, void *buffer, size_t length, int flags,
                 struct sockaddr *address, socklen_t *address_len);
```

DESCRIPTION

The *recvfrom()* function receives a message from a connection-mode or connectionless-mode socket. It is normally used with connectionless-mode sockets because it permits the application to retrieve the source address of received data.

The function takes the following arguments:

<i>socket</i>	Specifies the socket file descriptor.
<i>buffer</i>	Points to the buffer where the message should be stored.
<i>length</i>	Specifies the length in bytes of the buffer pointed to by the <i>buffer</i> argument.
<i>flags</i>	Specifies the type of message reception. Values of this argument are formed by logically OR'ing zero or more of the following values: <ul style="list-style-type: none"> MSG_PEEK Peeks at an incoming message. The data is treated as unread and the next <i>recvfrom()</i> or similar function will still return this data. MSG_OOB Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific. MSG_WAITALL Requests that the function block until the full amount of data requested can be returned. The function may return a smaller amount of data if a signal is caught, if the connection is terminated, if MSG_PEEK was specified, or if an error is pending for the socket.
<i>address</i>	A null pointer, or points to a sockaddr structure in which the sending address is to be stored. The length and format of the address depend on the address family of the socket.
<i>address_len</i>	Specifies the length of the sockaddr structure pointed to by the <i>address</i> argument.

The *recvfrom()* function returns the length of the message written to the buffer pointed to by the *buffer* argument. For message-based sockets such as SOCK_DGRAM and SOCK_SEQPACKET, the entire message must be read in a single operation. If a message is too long to fit in the supplied buffer, and MSG_PEEK is not set in the *flags* argument, the excess bytes are discarded. For stream-based sockets such as SOCK_STREAM, message boundaries are ignored. In this case, data is returned to the user as soon as it becomes available, and no data is discarded.

If the MSG_WAITALL flag is not set, data will be returned only up to the end of the first message.

Not all protocols provide the source address for messages. If the *address* argument is not a null pointer and the protocol provides the source address of messages, the source address of the

received message is stored in the **sockaddr** structure pointed to by the *address* argument, and the length of this address is stored in the object pointed to by the *address_len* argument.

If the actual length of the address is greater than the length of the supplied **sockaddr** structure, the stored address will be truncated.

If the *address* argument is not a null pointer and the protocol does not provide the source address of messages, the the value stored in the object pointed to by *address* is unspecified.

If no messages are available at the socket and O_NONBLOCK is not set on the socket's file descriptor, *recvfrom()* blocks until a message arrives. If no messages are available at the socket and O_NONBLOCK is set on the socket's file descriptor, *recvfrom()* fails and sets *errno* to [EAGAIN] or [EWOULDBLOCK].

RETURN VALUE

Upon successful completion, *recvfrom()* returns the length of the message in bytes. If no messages are available to be received and the peer has performed an orderly shutdown, *recvfrom()* returns 0. Otherwise the function returns -1 and sets *errno* to indicate the error.

ERRORS

The *recvfrom()* function will fail if:

[EAGAIN] or [EWOULDBLOCK]

The socket's file descriptor is marked O_NONBLOCK and no data is waiting to be received; or MSG_OOB is set and no out-of-band data is available and either the socket's file descriptor is marked O_NONBLOCK or the socket does not support blocking to await out-of-band data.

[EBADF]

The *socket* argument is not a valid file descriptor.

[ECONNRESET]

A connection was forcibly closed by a peer.

[EFAULT]

The *buffer*, *address* or *address_len* parameter can not be accessed or written.

[EINTR]

A signal interrupted *recvfrom()* before any data was available.

[EINVAL]

The MSG_OOB flag is set and no out-of-band data is available.

[ENOTCONN]

A receive is attempted on a connection-mode socket that is not connected.

[ENOTSOCK]

The *socket* argument does not refer to a socket.

[EOPNOTSUPP]

The specified flags are not supported for this socket type.

[ETIMEDOUT]

The connection timed out during connection establishment, or due to a transmission timeout on active connection.

The *recvfrom()* function may fail if:

[EIO]

An I/O error occurred while reading from or writing to the file system.

[ENOBUFS]

Insufficient resources were available in the system to perform the operation.

[ENOMEM]

Insufficient memory was available to fulfill the request.

[ENOSR]

There were insufficient STREAMS resources available for the operation to complete.

APPLICATION USAGE

The *select()* and *poll()* functions can be used to determine when data is available to be received.

SEE ALSO

poll(), *read()*, *recv()*, *recvmsg()*, *select()*, *send()*, *sendmsg()*, *sendto()*, *shutdown()*, *socket()*, *write()*, **<sys/socket.h>**.

CHANGE HISTORY

First released in Issue 4.

NAME

recvmsg — receive a message from a socket

SYNOPSIS

```
#include <sys/socket.h>

ssize_t recvmsg(int socket, struct msghdr *message, int flags);
```

DESCRIPTION

The *recvmsg()* function receives a message from a connection-mode or connectionless-mode socket. It is normally used with connectionless-mode sockets because it permits the application to retrieve the source address of received data.

The function takes the following arguments:

<i>socket</i>	Specifies the socket file descriptor.
<i>message</i>	Points to a msghdr structure, containing both the buffer to store the source address and the buffers for the incoming message. The length and format of the address depend on the address family of the socket. The msg_flags member is ignored on input, but may contain meaningful values on output.
<i>flags</i>	Specifies the type of message reception. Values of this argument are formed by logically OR'ing zero or more of the following values: <ul style="list-style-type: none"> MSG_OOB Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific. MSG_PEEK Peeks at the incoming message. MSG_WAITALL Requests that the function block until the full amount of data requested can be returned. The function may return a smaller amount of data if a signal is caught, if the connection is terminated, if MSG_PEEK was specified, or if an error is pending for the socket.

The *recvmsg()* function receives messages from unconnected or connected sockets and returns the length of the message.

The *recvmsg()* function returns the total length of the message. For message-based sockets such as SOCK_DGRAM and SOCK_SEQPACKET, the entire message must be read in a single operation. If a message is too long to fit in the supplied buffers, and MSG_PEEK is not set in the *flags* argument, the excess bytes are discarded, and MSG_TRUNC is set in the **msg_flags** member of the **msghdr** structure. For stream-based sockets such as SOCK_STREAM, message boundaries are ignored. In this case, data is returned to the user as soon as it becomes available, and no data is discarded.

If the MSG_WAITALL flag is not set, data will be returned only up to the end of the first message.

If no messages are available at the socket and O_NONBLOCK is not set on the socket's file descriptor, *recvfrom()* blocks until a message arrives. If no messages are available at the socket and O_NONBLOCK is set on the socket's file descriptor, *recvfrom()* function fails and sets *errno* to [EAGAIN] or [EWOULDBLOCK].

In the **msghdr** structure, the **msg_name** and **msg_namelen** members specify the source address if the socket is unconnected. If the socket is connected, the **msg_name** and **msg_namelen**

members are ignored. The **msg_name** member may be a null pointer if no names are desired or required. The **msg_iov** and **msg_iovlen** fields are used to specify where the received data will be stored. **msg_iov** points to an array of **iovec** structures; **msg_iovlen** must be set to the dimension of this array. In each **iovec** structure, the **iov_base** field specifies a storage area and the **iov_len** field gives its size in bytes. Each storage area indicated by **msg_iov** is filled with received data in turn until all of the received data is stored or all of the areas have been filled.

On successful completion, the **msg_flags** member of the message header is the bitwise-inclusive OR of all of the following flags that indicate conditions detected for the received message:

MSG_EOR	End of record was received (if supported by the protocol).
MSG_OOB	Out-of-band data was received.
MSG_TRUNC	Normal data was truncated.
MSG_CTRUNC	Control data was truncated.

RETURN VALUE

Upon successful completion, *recvmsg()* returns the length of the message in bytes. If no messages are available to be received and the peer has performed an orderly shutdown, *recvmsg()* returns 0. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The *recvmsg()* function will fail if:

[EAGAIN] or [EWOULDBLOCK]	The socket's file descriptor is marked O_NONBLOCK and no data is waiting to be received; or MSG_OOB is set and no out-of-band data is available and either the socket's file descriptor is marked O_NONBLOCK or the socket does not support blocking to await out-of-band data.
[EBADF]	The <i>socket</i> argument is not a valid open file descriptor.
[ECONNRESET]	A connection was forcibly closed by a peer.
[EFAULT]	The <i>message</i> parameter, or storage pointed to by the <i>msg_name</i> , <i>msg_control</i> or <i>msg_iov</i> fields of the <i>message</i> parameter, or storage pointed to by the iovec structures pointed to by the <i>msg_iov</i> field can not be accessed or written.
[EINTR]	This function was interrupted by a signal before any data was available.
[EINVAL]	The sum of the iov_len values overflows an ssize_t . or the MSG_OOB flag is set and no out-of-band data is available.
[EMSGSIZE]	The msg_iovlen member of the msghdr structure pointed to by <i>message</i> is less than or equal to 0, or is greater than {IOV_MAX}.
[ENOTCONN]	A receive is attempted on a connection-mode socket that is not connected.
[ENOTSOCK]	The <i>socket</i> argument does not refer to a socket.
[EOPNOTSUPP]	The specified flags are not supported for this socket type.
[ETIMEDOUT]	The connection timed out during connection establishment, or due to a transmission timeout on active connection.

The *recvmsg()* function may fail if:

[EIO]	An IO error occurred while reading from or writing to the file system.
[ENOBUFS]	Insufficient resources were available in the system to perform the operation.
[ENOMEM]	Insufficient memory was available to fulfill the request.
[ENOSR]	There were insufficient STREAMS resources available for the operation to complete.

APPLICATION USAGE

The *select()* and *poll()* functions can be used to determine when data is available to be received.

SEE ALSO

poll(), *recv()*, *recvfrom()*, *select()*, *send()*, *sendmsg()*, *sendto()*, *shutdown()*, *socket()*, [<sys/socket.h>](#).

CHANGE HISTORY

First released in Issue 4.

NAME

select — synchronous I/O multiplexing

Note: The **XSH** specification contains the basic definition of this interface. The following additional information pertains to Sockets.

DESCRIPTION

A file descriptor for a socket that is listening for connections will indicate that it is ready for reading, when connections are available. A file descriptor for a socket that is connecting asynchronously will indicate that it is ready for writing, when a connection has been established.

CHANGE HISTORY

First released in Issue 4.

NAME

send — send a message on a socket

SYNOPSIS

```
#include <sys/socket.h>

ssize_t send(int socket, const void *buffer, size_t length, int flags);
```

DESCRIPTION

<i>socket</i>	Specifies the socket file descriptor.
<i>buffer</i>	Points to the buffer containing the message to send.
<i>length</i>	Specifies the length of the message in bytes.
<i>flags</i>	Specifies the type of message transmission. Values of this argument are formed by logically OR'ing zero or more of the following flags:
MSG_EOR	Terminates a record (if supported by the protocol)
MSG_OOB	Sends out-of-band data on sockets that support out-of-band communications. The significance and semantics of out-of-band data are protocol-specific.

The *send()* function initiates transmission of a message from the specified socket to its peer. The *send()* function sends a message only when the socket is connected (including when the peer of a connectionless socket has been set via *connect()*).

The length of the message to be sent is specified by the *length* argument. If the message is too long to pass through the underlying protocol, *send()* fails and no data is transmitted.

Successful completion of a call to *send()* does not guarantee delivery of the message. A return value of -1 indicates only locally-detected errors.

If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does not have `O_NONBLOCK` set, *send()* blocks until space is available. If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does have `O_NONBLOCK` set, *send()* will fail. The *select()* and *poll()* functions can be used to determine when it is possible to send more data.

The socket in use may require the process to have appropriate privileges to use the *send()* function.

RETURN VALUE

Upon successful completion, *send()* returns the number of bytes sent. Otherwise, -1 is returned and *errno* is set to indicate the error.

APPLICATION USAGE

The *send()* function is identical to *sendto()* with a null pointer *dest_len* argument, and to *write()* if no flags are used.

ERRORS

The *send()* function will fail if:

[EAGAIN] or [EWOULDBLOCK]

The socket's file descriptor is marked `O_NONBLOCK` and the requested operation would block.

[EBADF]

The *socket* argument is not a valid file descriptor.

[ECONNRESET]	A connection was forcibly closed by a peer.
[EDESTADDRREQ]	The socket is not connection-mode and no peer address is set.
[EFAULT]	The <i>buffer</i> parameter can not be accessed.
[EINTR]	A signal interrupted <i>send()</i> before any data was transmitted.
[EMSGSIZE]	The message is too large be sent all at once, as the socket requires.
[ENOTCONN]	The socket is not connected or otherwise has not had the peer prespecified.
[ENOTSOCK]	The <i>socket</i> argument does not refer to a socket.
[EOPNOTSUPP]	The <i>socket</i> argument is associated with a socket that does not support one or more of the values set in <i>flags</i> .
[EPIPE]	The socket is shut down for writing, or the socket is connection-mode and is no longer connected. In the latter case, and if the socket is of type SOCK_STREAM, the SIGPIPE signal is generated to the calling process.

The *send()* function may fail if:

[EACCES]	The calling process does not have the appropriate privileges.
[EIO]	An I/O error occurred while reading from or writing to the file system.
[ENETDOWN]	The local interface used to reach the destination is down.
[ENETUNREACH]	No route to the network is present.
[ENOBUFS]	Insufficient resources were available in the system to perform the operation.
[ENOSR]	There were insufficient STREAMS resources available for the operation to complete.

SEE ALSO

connect(), *getsockopt()*, *poll()*, *recv()*, *recvfrom()*, *recvmsg()*, *select()*, *sendmsg()*, *sendto()*, *setsockopt()*, *shutdown()*, *socket()*, <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

sendmsg — send a message on a socket using a message structure

SYNOPSIS

```
#include <sys/socket.h>

ssize_t sendmsg (int socket, const struct msghdr *message, int flags);
```

DESCRIPTION

The *sendmsg()* function sends a message through a connection-mode or connectionless-mode socket. If the socket is connectionless-mode, the message will be sent to the address specified by *msghdr*. If the socket is connection-mode, the destination address in *msghdr* is ignored.

The function takes the following arguments:

<i>socket</i>	Specifies the socket file descriptor.				
<i>message</i>	Points to a msghdr structure, containing both the destination address and the buffers for the outgoing message. The length and format of the address depend on the address family of the socket. The msg_flags member is ignored.				
<i>flags</i>	Specifies the type of message transmission. The application may specify 0 or the following flag: <table> <tr> <td>MSG_EOR</td> <td>Terminates a record (if supported by the protocol)</td> </tr> <tr> <td>MSG_OOB</td> <td>Sends out-of-band data on sockets that support out-of-band data. The significance and semantics of out-of-band data are protocol-specific.</td> </tr> </table>	MSG_EOR	Terminates a record (if supported by the protocol)	MSG_OOB	Sends out-of-band data on sockets that support out-of-band data. The significance and semantics of out-of-band data are protocol-specific.
MSG_EOR	Terminates a record (if supported by the protocol)				
MSG_OOB	Sends out-of-band data on sockets that support out-of-band data. The significance and semantics of out-of-band data are protocol-specific.				

The *msg_iov* and *msg_iovlen* fields of message specify zero or more buffers containing the data to be sent. *msg_iov* points to an array of **iovec** structures; *msg_iovlen* must be set to the dimension of this array. In each **iovec** structure, the *iov_base* field specifies a storage area and the *iov_len* field gives its size in bytes. Some of these sizes can be zero. The data from each storage area indicated by *msg_iov* is sent in turn.

Successful completion of a call to *sendmsg()* does not guarantee delivery of the message. A return value of -1 indicates only locally-detected errors.

If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does not have `O_NONBLOCK` set, *sendmsg()* function blocks until space is available. If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does have `O_NONBLOCK` set, *sendmsg()* function will fail.

If the socket protocol supports broadcast and the specified address is a broadcast address for the socket protocol, *sendmsg()* will fail if the `SO_BROADCAST` option is not set for the socket.

The socket in use may require the process to have appropriate privileges to use the *sendmsg()* function.

RETURN VALUE

Upon successful completion, *sendmsg()* function returns the number of bytes sent. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The *sendmsg()* function will fail if:

[EAGAIN] or [EWOULDBLOCK]

The socket's file descriptor is marked `O_NONBLOCK` and the requested operation would block.

[EAFNOSUPPORT]	Addresses in the specified address family cannot be used with this socket.
[EBADF]	The <i>socket</i> argument is not a valid file descriptor.
[ECONNRESET]	A connection was forcibly closed by a peer.
[EFAULT]	The <i>message</i> parameter, or storage pointed to by the <i>msg_name</i> , <i>msg_control</i> or <i>msg_iov</i> fields of the <i>message</i> parameter, or storage pointed to by the <i>iovec</i> structures pointed to by the <i>msg_iov</i> field can not be accessed.
[EINTR]	A signal interrupted <i>sendmsg()</i> before any data was transmitted.
[EINVAL]	The sum of the <i>iov_len</i> values overflows an <i>ssize_t</i> .
[EMSGSIZE]	The message is too large to be sent all at once (as the socket requires), or the <i>msg_iovlen</i> member of the <i>msghdr</i> structure pointed to by <i>message</i> is less than or equal to 0 or is greater than {IOV_MAX}.
[ENOTCONN]	The socket is connection-mode but is not connected.
[ENOTSOCK]	The <i>socket</i> argument does not refer a socket.
[EOPNOTSUPP]	The <i>socket</i> argument is associated with a socket that does not support one or more of the values set in <i>flags</i> .
[EPIPE]	The socket is shut down for writing, or the socket is connection-mode and is no longer connected. In the latter case, and if the socket is of type SOCK_STREAM, the SIGPIPE signal is generated to the calling process.

If the address family of the socket is AF_UNIX, then *sendmsg()* will fail if:

[EIO]	An I/O error occurred while reading from or writing to the file system.
[ELOOP]	Too many symbolic links were encountered in translating the pathname in the socket address.
[ENAMETOOLONG]	A component of a pathname exceeded {NAME_MAX} characters, or an entire pathname exceeded {PATH_MAX} characters.
[ENOENT]	A component of the pathname does not name an existing file or the pathname is an empty string.
[ENOTDIR]	A component of the path prefix of the pathname in the socket address is not a directory.

The *sendmsg()* function may fail if:

[EACCES]	Search permission is denied for a component of the path prefix; or write access to the named socket is denied.
[EDESTADDRREQ]	The socket is not connection-mode and does not have its peer address set, and no destination address was specified.
[EHOSTUNREACH]	The destination host cannot be reached (probably because the host is down or a remote router cannot reach it).
[EIO]	An I/O error occurred while reading from or writing to the file system.
[EISCONN]	A destination address was specified and the socket is already connected.
[ENETDOWN]	The local interface used to reach the destination is down.

- [ENETUNREACH] No route to the network is present.
- [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- [ENOMEM] Insufficient memory was available to fulfill the request.
- [ENOSR] There were insufficient STREAMS resources available for the operation to complete.

If the address family of the socket is AF_UNIX, then *sendmsg()* may fail if:

- [ENAMETOOLONG] Pathname resolution of a symbolic link produced an intermediate result whose length exceeds {PATH_MAX}.

APPLICATION USAGE

The *select()* and *poll()* functions can be used to determine when it is possible to send more data.

SEE ALSO

getsockopt(), *poll()*, *recv()*, *recvfrom()*, *recvmsg()*, *select()*, *send()*, *sendto()*, *setsockopt()*, *shutdown()*, *socket()*, <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

sendto — send a message on a socket

SYNOPSIS

```
#include <sys/socket.h>

ssize_t sendto(int socket, const void *message, size_t length, int flags,
               const struct sockaddr *dest_addr, socklen_t dest_len);
```

DESCRIPTION

The *sendto()* function sends a message through a connection-mode or connectionless-mode socket. If the socket is connectionless-mode, the message will be sent to the address specified by *dest_addr*. If the socket is connection-mode, *dest_addr* is ignored.

The function takes the following arguments:

<i>socket</i>	Specifies the socket file descriptor.
<i>message</i>	Points to a buffer containing the message to be sent.
<i>length</i>	Specifies the size of the message in bytes.
<i>flags</i>	Specifies the type of message transmission. Values of this argument are formed by logically OR'ing zero or more of the following flags: <ul style="list-style-type: none"> MSG_EOR Terminates a record (if supported by the protocol) MSG_OOB Sends out-of-band data on sockets that support out-of-band data. The significance and semantics of out-of-band data are protocol-specific.
<i>dest_addr</i>	Points to a sockaddr structure containing the destination address. The length and format of the address depend on the address family of the socket.
<i>dest_len</i>	Specifies the length of the sockaddr structure pointed to by the <i>dest_addr</i> argument.

If the socket protocol supports broadcast and the specified address is a broadcast address for the socket protocol, *sendto()* will fail if the SO_BROADCAST option is not set for the socket.

The *dest_addr* argument specifies the address of the target. The *length* argument specifies the length of the message.

Successful completion of a call to *sendto()* does not guarantee delivery of the message. A return value of -1 indicates only locally-detected errors.

If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does not have O_NONBLOCK set, *sendto()* blocks until space is available. If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does have O_NONBLOCK set, *sendto()* will fail.

The socket in use may require the process to have appropriate privileges to use the *sendto()* function.

RETURN VALUE

Upon successful completion, *sendto()* returns the number of bytes sent. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The *sendto()* function will fail if:

[EAFNOSUPPORT]	Addresses in the specified address family cannot be used with this socket.
[EAGAIN] or [EWOULDBLOCK]	The socket's file descriptor is marked O_NONBLOCK and the requested operation would block.
[EBADF]	The <i>socket</i> argument is not a valid file descriptor.
[ECONNRESET]	A connection was forcibly closed by a peer.
[EFAULT]	The <i>message</i> or <i>destaddr</i> parameter can not be accessed.
[EINTR]	A signal interrupted <i>sendto()</i> before any data was transmitted.
[EMSGSIZE]	The message is too large to be sent all at once, as the socket requires.
[ENOTCONN]	The socket is connection-mode but is not connected.
[ENOTSOCK]	The <i>socket</i> argument does not refer to a socket.
[EOPNOTSUPP]	The <i>socket</i> argument is associated with a socket that does not support one or more of the values set in <i>flags</i> .
[EPIPE]	The socket is shut down for writing, or the socket is connection-mode and is no longer connected. In the latter case, and if the socket is of type SOCK_STREAM, the SIGPIPE signal is generated to the calling process.

If the address family of the socket is AF_UNIX, then *sendto()* will fail if:

[EIO]	An I/O error occurred while reading from or writing to the file system.
[ELOOP]	Too many symbolic links were encountered in translating the pathname in the socket address.
[ENAMETOOLONG]	A component of a pathname exceeded {NAME_MAX} characters, or an entire pathname exceeded {PATH_MAX} characters.
[ENOENT]	A component of the pathname does not name an existing file or the pathname is an empty string.
[ENOTDIR]	A component of the path prefix of the pathname in the socket address is not a directory.

The *sendto()* function may fail if:

[EACCES]	Search permission is denied for a component of the path prefix; or write access to the named socket is denied.
[EDESTADDRREQ]	The socket is not connection-mode and does not have its peer address set, and no destination address was specified.
[EHOSTUNREACH]	The destination host cannot be reached (probably because the host is down or a remote router cannot reach it).
[EINVAL]	The <i>dest_len</i> argument is not a valid length for the address family.
[EIO]	An I/O error occurred while reading from or writing to the file system.
[EISCONN]	A destination address was specified and the socket is already connected.
[ENETDOWN]	The local interface used to reach the destination is down.
[ENETUNREACH]	No route to the network is present.

- [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- [ENOMEM] Insufficient memory was available to fulfill the request.
- [ENOSR] There were insufficient STREAMS resources available for the operation to complete.

If the address family of the socket is AF_UNIX, then *sendto()* may fail if:

- [ENAMETOOLONG] Pathname resolution of a symbolic link produced an intermediate result whose length exceeds {PATH_MAX}.

APPLICATION USAGE

The *select()* and *poll()* functions can be used to determine when it is possible to send more data.

SEE ALSO

getsockopt(), *poll()*, *recv()*, *recvfrom()*, *recvmsg()*, *select()*, *send()*, *sendmsg()*, *setsockopt()*, *shutdown()*, *socket()*, <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

setsockopt — set the socket options

SYNOPSIS

```
#include <sys/socket.h>

int setsockopt(int socket, int level, int option_name, const void
               *option_value, socklen_t option_len);
```

DESCRIPTION

The *setsockopt()* function sets the option specified by the *option_name* argument, at the protocol level specified by the *level* argument, to the value pointed to by the *option_value* argument for the socket associated with the file descriptor specified by the *socket* argument.

The *level* argument specifies the protocol level at which the option resides. To set options at the socket level, specify the *level* argument as SOL_SOCKET. To set options at other levels, supply the appropriate protocol number for the protocol controlling the option. For example, to indicate that an option will be interpreted by the TCP (Transport Control Protocol), set *level* to the protocol number of TCP, as defined in the <netinet/in.h> header, or as determined by using *getprotobyname()*.

The *option_name* argument specifies a single option to set. The *option_name* argument and any specified options are passed uninterpreted to the appropriate protocol module for interpretations. The <sys/socket.h> header defines the socket level options. The options are as follows:

SO_DEBUG	Turns on recording of debugging information. This option enables or disables debugging in the underlying protocol modules. This option takes an int value. This is a boolean option.
SO_BROADCAST	Permits sending of broadcast messages, if this is supported by the protocol. This option takes an int value. This is a boolean option.
SO_REUSEADDR	Specifies that the rules used in validating addresses supplied to <i>bind()</i> should allow reuse of local addresses, if this is supported by the protocol. This option takes an int value. This is a boolean option.
SO_KEEPALIVE	Keeps connections active by enabling the periodic transmission of messages, if this is supported by the protocol. This option takes an int value. If the connected socket fails to respond to these messages, the connection is broken and processes writing to that socket are notified with a SIGPIPE signal. This is a boolean option.
SO_LINGER	Lingers on a <i>close()</i> if data is present. This option controls the action taken when unsent messages queue on a socket and <i>close()</i> is performed. If SO_LINGER is set, the system blocks the process during <i>close()</i> until it can transmit the data or until the time expires. If SO_LINGER is not specified, and <i>close()</i> is issued, the system handles the call in a way that allows the process to continue as quickly as possible. This option takes a linger structure, as defined in the <sys/socket.h> header, to specify the state of the option and linger interval.
SO_OOBINLINE	Leaves received out-of-band data (data marked urgent) in line. This option takes an int value. This is a boolean option.

SO_SNDBUF	Sets send buffer size. This option takes an int value.
SO_RCVBUF	Sets receive buffer size. This option takes an int value.
SO_DONTROUTE	Requests that outgoing messages bypass the standard routing facilities. The destination must be on a directly-connected network, and messages are directed to the appropriate network interface according to the destination address. The effect, if any, of this option depends on what protocol is in use. This option takes an int value. This is a boolean option.
SO_RCVLOWAT	Sets the minimum number of bytes to process for socket input operations. The default value for SO_RCVLOWAT is 1. If SO_RCVLOWAT is set to a larger value, blocking receive calls normally wait until they have received the smaller of the low water mark value or the requested amount. (They may return less than the low water mark if an error occurs, a signal is caught, or the type of data next in the receive queue is different than that returned, e.g. out of band data). This option takes an int value. Note that not all implementations allow this option to be set.
SO_RCVTIMEO	Sets the timeout value that specifies the maximum amount of time an input function waits until it completes. It accepts a timeval structure with the number of seconds and microseconds specifying the limit on how long to wait for an input operation to complete. If a receive operation has blocked for this much time without receiving additional data, it returns with a partial count or <i>errno</i> set to [EAGAIN] or [EWOULDBLOCK] if no data were received. The default for this option is zero, which indicates that a receive operation will not time out. This option takes a timeval structure. Note that not all implementations allow this option to be set.
SO_SNDLOWAT	Sets the minimum number of bytes to process for socket output operations. Non-blocking output operations will process no data if flow control does not allow the smaller of the send low water mark value or the entire request to be processed. This option takes an int value. Note that not all implementations allow this option to be set.
SO_SNDTIMEO	Sets the timeout value specifying the amount of time that an output function blocks because flow control prevents data from being sent. If a send operation has blocked for this time, it returns with a partial count or with <i>errno</i> set to [EAGAIN] or [EWOULDBLOCK] if no data were sent. The default for this option is zero, which indicates that a send operation will not time out. This option stores a timeval structure. Note that not all implementations allow this option to be set.

For boolean options, 0 indicates that the option is disabled and 1 indicates that the option is enabled.

Options at other protocol levels vary in format and name.

RETURN VALUE

Upon successful completion, *setsockopt()* returns 0. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The *setsockopt()* function will fail if:

[EBADF] The *socket* argument is not a valid file descriptor.

[EDOM]	The send and receive timeout values are too big to fit into the timeout fields in the socket structure.
[EFAULT]	The <i>option_value</i> parameter can not be accessed or written.
[EINVAL]	The specified option is invalid at the specified socket level or the socket has been shut down.
[EISCONN]	The socket is already connected, and a specified option can not be set while the socket is connected.
[ENOPROTOOPT]	The option is not supported by the protocol.
[ENOTSOCK]	The <i>socket</i> argument does not refer to a socket.

The *setsockopt()* function may fail if:

[ENOMEM]	There was insufficient memory available for the operation to complete.
[ENOBUFS]	Insufficient resources are available in the system to complete the call.
[ENOSR]	There were insufficient STREAMS resources available for the operation to complete.

APPLICATION USAGE

The *setsockopt()* function provides an application program with the means to control socket behaviour. An application program can use *setsockopt()* to allocate buffer space, control timeouts, or permit socket data broadcasts. The `<sys/socket.h>` header defines the socket-level options available to *setsockopt()*.

Options may exist at multiple protocol levels. The SO_ options are always present at the uppermost socket level.

SEE ALSO

bind(), *endprotoent()*, *getsockopt()*, *socket()*, `<sys/socket.h>`.

CHANGE HISTORY

First released in Issue 4.

NAME

shutdown — shut down socket send and receive operations

SYNOPSIS

```
#include <sys/socket.h>

int shutdown(int socket, int how);
```

DESCRIPTION

<i>socket</i>	Specifies the file descriptor of the socket.
<i>how</i>	Specifies the type of shutdown. The values are as follows:
SHUT_RD	Disables further receive operations.
SHUT_WR	Disables further send operations.
SHUT_RDWR	Disables further send and receive operations.

The *shutdown()* function disables subsequent send and/or receive operations on a socket, depending on the value of the *how* argument.

RETURN VALUE

Upon successful completion, *shutdown()* returns 0. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The *shutdown()* function will fail if:

[EBADF]	The <i>socket</i> argument is not a valid file descriptor.
[EINVAL]	The <i>how</i> argument is invalid.
[ENOTCONN]	The socket is not connected.
[ENOTSOCK]	The <i>socket</i> argument does not refer to a socket.

The *shutdown()* function may fail if:

[ENOBUFS]	Insufficient resources were available in the system to perform the operation.
[ENOSR]	There were insufficient STREAMS resources available for the operation to complete.

SEE ALSO

getsockopt(), *read()*, *recv()*, *recvfrom()*, *recvmsg()*, *select()*, *send()*, *sendto()*, *setsockopt()*, *socket()*, *write()*, <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

socket — create an endpoint for communication

SYNOPSIS

```
#include <sys/socket.h>

int socket(int domain, int type, int protocol);
```

DESCRIPTION

The *socket()* function creates an unbound socket in a communications domain, and returns a file descriptor that can be used in later function calls that operate on sockets.

The function takes the following arguments:

<i>domain</i>	Specifies the communications domain in which a socket is to be created.
<i>type</i>	Specifies the type of socket to be created.
<i>protocol</i>	Specifies a particular protocol to be used with the socket. Specifying a <i>protocol</i> of 0 causes <i>socket()</i> to use an unspecified default protocol appropriate for the requested socket type.

The *domain* argument specifies the address family used in the communications domain. The address families supported by the system are implementation-dependent.

The `<sys/socket.h>` header defines at least the following values for the *domain* argument:

AF_UNIX	File system pathnames.
AF_INET	Internet address.

The *type* argument specifies the socket type, which determines the semantics of communication over the socket. The socket types supported by the system are implementation-dependent. Possible socket types include:

SOCK_STREAM	Provides sequenced, reliable, bidirectional, connection-mode byte streams, and may provide a transmission mechanism for out-of-band data.
SOCK_DGRAM	Provides datagrams, which are connectionless-mode, unreliable messages of fixed maximum length.
SOCK_SEQPACKET	Provides sequenced, reliable, bidirectional, connection-mode transmission path for records. A record can be sent using one or more output operations and received using one or more input operations, but a single operation never transfers part of more than one record. Record boundaries are visible to the receiver via the MSG_EOR flag.

If the *protocol* argument is non-zero, it must specify a protocol that is supported by the address family. The protocols supported by the system are implementation-dependent.

The process may need to have appropriate privileges to use the *socket()* function or to create some sockets.

RETURN VALUE

Upon successful completion, *socket()* returns a nonnegative integer, the socket file descriptor. Otherwise a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

The *socket()* function will fail if:

[EAFNOSUPPORT]	The implementation does not support the specified address family.
[EMFILE]	No more file descriptors are available for this process.
[ENFILE]	No more file descriptors are available for the system.
[EPROTONOSUPPORT]	The protocol is not supported by the address family, or the protocol is not supported by the implementation.
[EPROTOTYPE]	The socket type is not supported by the protocol.

The *socket()* function may fail if:

[EACCES]	The process does not have appropriate privileges.
[ENOBUFS]	Insufficient resources were available in the system to perform the operation.
[ENOMEM]	Insufficient memory was available to fulfill the request.
[ENOSR]	There were insufficient STREAMS resources available for the operation to complete.

APPLICATION USAGE

The documentation for specific address families specify which protocols each address family supports. The documentation for specific protocols specify which socket types each protocol supports.

The application can determine if an address family is supported by trying to create a socket with *domain* set to the protocol in question.

SEE ALSO

accept(), *bind()*, *connect()*, *getsockname()*, *getsockopt()*, *listen()*, *recv()*, *recvfrom()*, *recvmsg()*, *send()*, *sendmsg()*, *setsockopt()*, *shutdown()*, *socketpair()*, <netinet/in.h>, <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

socketpair — create a pair of connected sockets

SYNOPSIS

```
#include <sys/socket.h>

int socketpair(int domain, int type, int protocol,
               int socket_vector[2]);
```

DESCRIPTION

The `socketpair()` function creates an unbound pair of connected sockets in a specified *domain*, of a specified *type*, under the protocol optionally specified by the *protocol* argument. The two sockets are identical. The file descriptors used in referencing the created sockets are returned in `socket_vector[0]` and `socket_vector[1]`.

<i>domain</i>	Specifies the communications domain in which the sockets are to be created.
<i>type</i>	Specifies the type of sockets to be created.
<i>protocol</i>	Specifies a particular protocol to be used with the sockets. Specifying a <i>protocol</i> of 0 causes <code>socketpair()</code> to use an unspecified default protocol appropriate for the requested socket type.
<i>socket_vector</i>	Specifies a 2-integer array to hold the file descriptors of the created socket pair.

The *type* argument specifies the socket type, which determines the semantics of communications over the socket. The socket types supported by the system are implementation-dependent. Possible socket types include:

SOCK_STREAM	Provides sequenced, reliable, bidirectional, connection-mode byte streams, and may provide a transmission mechanism for out-of-band data.
SOCK_DGRAM	Provides datagrams, which are connectionless-mode, unreliable messages of fixed maximum length.
SOCK_SEQPACKET	Provides sequenced, reliable, bidirectional, connection-mode transmission path for records. A record can be sent using one or more output operations and received using one or more input operations, but a single operation never transfers part of more than one record. Record boundaries are visible to the receiver via the MSG_EOR flag.

If the *protocol* argument is non-zero, it must specify a protocol that is supported by the address family. The protocols supported by the system are implementation-dependent.

The process may need to have appropriate privileges to use the `socketpair()` function or to create some sockets.

RETURN VALUE

Upon successful completion, this function returns 0. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The `socketpair()` function will fail if:

[EAFNOSUPPORT]	The implementation does not support the specified address family.
[EMFILE]	No more file descriptors are available for this process.

[ENFILE]	No more file descriptors are available for the system.
[EOPNOTSUPP]	The specified protocol does not permit creation of socket pairs.
[EPROTONOSUPPORT]	The protocol is not supported by the address family, or the protocol is not supported by the implementation.
[EPROTOTYPE]	The socket type is not supported by the protocol.

The *socketpair()* function may fail if:

[EACCES]	The process does not have appropriate privileges.
[ENOBUFS]	Insufficient resources were available in the system to perform the operation.
[ENOMEM]	Insufficient memory was available to fulfill the request.
[ENOSR]	There were insufficient STREAMS resources available for the operation to complete.

APPLICATION USAGE

The documentation for specific address families specifies which protocols each address family supports. The documentation for specific protocols specifies which socket types each protocol supports.

The *socketpair()* function is used primarily with UNIX domain sockets and need not be supported for other domains.

SEE ALSO

socket(), <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

write, writev — write on a file

Note: The **XSH** specification contains the basic definition of this interface. The following additional information pertains to Sockets.

DESCRIPTION

If *files* refers to a socket, *write()* is equivalent to *send()* with no flags set.

CHANGE HISTORY

First released in Issue 4.

Sockets Headers

This chapter describes the contents of headers used by the X/Open Sockets functions, macros and external variables.

Headers contain the definition of symbolic constants, common structures, preprocessor macros and defined types. Each function in Chapter 8 specifies the headers that an application must include in order to use that function. In most cases only one header is required. These headers are present on an application development system; they do not have to be present on the target execution system.

NAME

fcntl.h — file control options

Note: The **XSH** specification contains the basic definition of this interface. The following additional information pertains to Sockets.

DESCRIPTION

The **<fcntl.h>** header defines the following additional values for *cmd* used by *fcntl()*:

F_GETOWN Get process or process group ID to receive SIGURG signals.

F_SETOWN Set process or process group ID to receive SIGURG signals.

CHANGE HISTORY

First released in Issue 4.

NAME

sys/socket.h — Internet Protocol family

SYNOPSIS

#include <sys/socket.h>

DESCRIPTION

<sys/socket.h> makes available a type, **socklen_t**, which is an unsigned opaque integral type of length of at least 32 bits¹¹.

The <sys/socket.h> header defines the unsigned integral type **sa_family_t**.

The <sys/socket.h> header defines the **sockaddr** structure that includes at least the following members:

sa_family_t	sa_family	address family
char	sa_data[]	socket address (variable-length data)

The <sys/socket.h> header defines the **msghdr** structure that includes at least the following members:

void	*msg_name	optional address
socklen_t	msg_namelen	size of address
struct iovec	*msg_iov	scatter/gather array
int	msg_iovlen	members in msg_iov
void	*msg_control	ancillary data, see below
socklen_t	msg_controllen	ancillary data buffer len
int	msg_flags	flags on received message

The <sys/socket.h> header defines the **cmsghdr** structure that includes at least the following members:

socklen_t	cmsg_len	data byte count, including the cmsghdr
int	cmsg_level	originating protocol
int	cmsg_type	protocol-specific type

Ancillary data consists of a sequence of pairs, each consisting of a **cmsghdr** structure followed by a data array. The data array contains the ancillary data message, and the **cmsghdr** structure contains descriptive information that allows an application to correctly parse the data.

The values for **cmsg_level** will be legal values for the level argument to the *getsockopt()* and *setsockopt()* functions. The system documentation should specify the **cmsg_type** definitions for the supported protocols.

Ancillary data is also possible at the socket level. The <sys/socket.h> header defines the following macro for use as the **cmsg_type** value when **cmsg_level** is SOL_SOCKET:

SCM_RIGHTS	Indicates that the data array contains the access rights to be sent or received.
------------	--

11. To forestall portability problems, it is recommended that applications should not use values larger than $2^{32} - 1$.

The **<sys/socket.h>** header defines the following macros to gain access to the data arrays in the ancillary data associated with a message header:

CMSG_DATA(*cmsg*) If the argument is a pointer to a **cmsg_hdr** structure, this macro returns an unsigned character pointer to the data array associated with the **cmsg_hdr** structure.

CMSG_NXTHDR(*mhdr, cmsg*)
If the first argument is a pointer to a **msghdr** structure and the second argument is a pointer to a **cmsg_hdr** structure in the ancillary data, pointed to by the **msg_control** field of that **msghdr** structure, this macro returns a pointer to the next **cmsg_hdr** structure, or a null pointer if this structure is the last **cmsg_hdr** in the ancillary data.

CMSG_FIRSTHDR(*mhdr*)
If the argument is a pointer to a **msghdr** structure, this macro returns a pointer to the first **cmsg_hdr** structure in the ancillary data associated with this **msghdr** structure, or a null pointer if there is no ancillary data associated with the **msghdr** structure.

The **<sys/socket.h>** header defines the **linger** structure that includes at least the following members:

int	<code>l_onoff</code>	indicates whether linger option is enabled
int	<code>l_linger</code>	linger time, in seconds

The **<sys/socket.h>** header defines the following macros, with distinct integral values:

SOCK_DGRAM	Datagram socket
SOCK_STREAM	Byte-stream socket
SOCK_SEQPACKET	Sequenced-packet socket

The **<sys/socket.h>** header defines the following macro for use as the *level* argument of *setsockopt()* and *getsockopt()*.

SOL_SOCKET Options to be accessed at socket level, not protocol level.

The **<sys/socket.h>** header defines the following macros, with distinct integral values, for use as the *option_name* argument in *getsockopt()* or *setsockopt()* calls:

SO_ACCEPTCONN	Socket is accepting connections.
SO_BROADCAST	Transmission of broadcast messages is supported.
SO_DEBUG	Debugging information is being recorded.
SO_DONTROUTE	bypass normal routing
SO_ERROR	Socket error status.
SO_KEEPALIVE	Connections are kept alive with periodic messages.
SO_LINGER	Socket lingers on close.
SO_OOBINLINE	Out-of-band data is transmitted in line.
SO_RCVBUF	Receive buffer size.
SO_RCVLOWAT	receive "low water mark"
SO_RCVTIMEO	receive timeout
SO_REUSEADDR	Reuse of local addresses is supported.
SO_SNDBUF	Send buffer size.
SO_SNDLOWAT	send "low water mark"
SO_SNDTIMEO	send timeout
SO_TYPE	Socket type.

The <sys/socket.h> header defines the following macros, with distinct integral values, for use as the valid values for the **msg_flags** field in the **msg_hdr** structure, or the flags parameter in *recvfrom()*, *recvmsg()*, *sendto()* or *sendmsg()* calls:

MSG_CTRUNC	Control data truncated.
MSG_DONTROUTE	Send without using routing tables.
MSG_EOR	Terminates a record (if supported by the protocol).
MSG_OOB	Out-of-band data.
MSG_PEEK	Leave received data in queue.
MSG_TRUNC	Normal data truncated.
MSG_WAITALL	Wait for complete message.

The <sys/socket.h> header defines the following macros, with distinct integral values:

AF_UNIX	UNIX domain sockets
AF_UNSPEC	Unspecified
AF_INET	Internet domain sockets

The <sys/socket.h> header defines the following macros, with distinct integral values:

SHUT_RD	Disables further receive operations.
SHUT_WR	Disables further send operations.
SHUT_RDWR	Disables further send and receive operations.

The following are declared as functions, and may also be defined as macros:

```
int    accept(int socket, struct sockaddr *address,
             socklen_t *address_len);
int    bind(int socket, const struct sockaddr *address,
            socklen_t address_len);
int    connect(int socket, const struct sockaddr *address,
              socklen_t address_len);
int    getpeername(int socket, struct sockaddr *address,
                  socklen_t *address_len);
int    getsockname(int socket, struct sockaddr *address,
                  socklen_t *address_len);
int    getsockopt(int socket, int level, int option_name,
                 void *option_value, socklen_t *option_len);
int    listen(int socket, int backlog);
ssize_t recv(int socket, void *buffer, size_t length, int flags);
ssize_t recvfrom(int socket, void *buffer, size_t length,
                 int flags, struct sockaddr *address, socklen_t *address_len);
ssize_t recvmsg(int socket, struct msg_hdr *message, int flags);
ssize_t send(int socket, const void *message, size_t length, int flags);
ssize_t sendmsg(int socket, const struct msg_hdr *message, int flags);
ssize_t sendto(int socket, const void *message, size_t length, int flags,
               const struct sockaddr *dest_addr, socklen_t dest_len);
int    setsockopt(int socket, int level, int option_name,
                 const void *option_value, socklen_t option_len);
int    shutdown(int socket, int how);
int    socket(int domain, int type, int protocol);
int    socketpair(int domain, int type, int protocol,
                 int socket_vector[2]);
```

SEE ALSO

accept(), bind(), connect(), getpeername(), getsockname(), getsockopt(), listen(), recv(), recvfrom(), recvmsg(), send(), sendmsg(), sendto(), setsockopt(), shutdown(), socket(), socketpair().

CHANGE HISTORY

First released in Issue 4.

NAME

sys/stat.h — data returned by the *stat()* function.

Note: The **XSH** specification contains the basic definition of this interface. The following additional information pertains to Sockets.

DESCRIPTION

The following additional symbolic name for the value of **st_mode** is defined:

File type:

S_IFMT type of file
S_IFSOCK socket

The following macro will test whether a file is of the specified type. The value *m* supplied to the macro is the value of **st_mode** from a **stat** structure. The macro evaluates to a non-zero value if the test is true, 0 if the test is false.

S_ISSOCK (*m*) test for a socket

CHANGE HISTORY

First released in Issue 4.

NAME

sys/uio.h — definitions for scatter/gather I/O

Note: The **XSH** specification contains the basic definition of this interface.

DESCRIPTION

This specification uses the **iovec** structure defined in <sys/uio.h> as definitions for scatter/gather I/O.

NAME

sys/un.h — definitions for UNIX-domain sockets

SYNOPSIS

```
#include <sys/un.h>
```

DESCRIPTION

The <sys/un.h> header defines the **sockaddr_un** structure that includes at least the following members:

```
sa_family_t  sun_family      address family
char         sun_path[]      socket pathname
```

The **sockaddr_un** structure is used to store addresses for UNIX domain sockets. Values of this type must be cast to **struct sockaddr** for use with the socket interfaces defined in this document.

The <sys/un.h> header defines the type **sa_family_t** as described in <sys/socket.h>.

The size of `sun_path` (see <sys/un.h>) has intentionally been left undefined. This was done for good reasons. Different implementations have used different sizes. For example, BSD4.3 uses a size of 108. BSD4.4 uses a size of 104. Since most of the implementations today originated from BSD versions, most of the major vendors today use a size that ranges from 92 to 108. Applications should not assume a particular length for `sun_path` or assume that it can hold `_POSIX_PATH_MAX` characters (255).

SEE ALSO

bind(), *socket()*, *socketpair()*.

IP Address Resolution Interfaces

Address Resolution refers to a set of interfaces that obtain network information and are usable in conjunction with both XTI and Sockets when using the Internet Protocol (IP).

This chapter provides reference manual pages for the address resolution API. This includes functions, macros and external variables to support application portability at the C-language source level.

NAME

endhostent, gethostbyaddr, gethostbyname, gethostent, sethostent — network host database functions

SYNOPSIS

```
#include <netdb.h>

extern int h_errno;

void endhostent(void);

struct hostent *gethostbyaddr(const void *addr, size_t len, int type);
struct hostent *gethostbyname(const char *name);
struct hostent *gethostent(void);
void sethostent(int stayopen);
```

DESCRIPTION

The *gethostent()*, *gethostbyaddr()*, and *gethostbyname()* functions each return a pointer to a **hostent** structure, the members of which contain the fields of an entry in the network host database.

The *gethostent()* function reads the next entry of the database, opening a connection to the database if necessary.

The *gethostbyaddr()* function searches the database and finds an entry which matches the address family specified by the *type* argument and which matches the address pointed to by the *addr* argument, opening a connection to the database if necessary. The *addr* argument is a pointer to the binary-format (that is, not null-terminated) address in network byte order, whose length is specified by the *len* argument. The datatype of the address depends on the address family. For an address of type AF_INET, this is an **in_addr** structure, defined in **<netinet/in.h>**.

The *gethostbyname()* function searches the database and finds an entry which matches the host name specified by the *name* argument, opening a connection to the database if necessary. If *name* is an alias for a valid host name, the function returns information about the host name to which the alias refers, and *name* is included in the list of aliases returned.

The *sethostent()* function opens a connection to the network host database, and sets the position of the next entry to the first entry. If the *stayopen* argument is non-zero, the connection to the host database will not be closed after each call to *gethostent()* (either directly, or indirectly through one of the other *gethost**() functions).

The *endhostent()* function closes the connection to the database.

RETURN VALUE

On successful completion, *gethostbyaddr()*, *gethostbyname()* and *gethostent()* return a pointer to a **hostent** structure if the requested entry was found, and a null pointer if the end of the database was reached or the requested entry was not found. Otherwise, a null pointer is returned.

On unsuccessful completion, *gethostbyaddr()* and *gethostbyname()* functions set *h_errno* to indicate the error.

ERRORS

No errors are defined for *endhostent()*, *gethostent()* and *sethostent()*.

The *gethostbyaddr()* and *gethostbyname()* functions will fail in the following cases, setting *h_errno* to the value shown in the list below. Any changes to *errno* are unspecified.

[HOST_NOT_FOUND]	No such host is known.
[NO_DATA]	The server recognised the request and the name but no address is available. Another type of request to the name server for the domain might return an answer.
[NO_RECOVERY]	An unexpected server failure occurred which can not be recovered.
[TRY_AGAIN]	A temporary and possibly transient error occurred, such as a failure of a server to respond.

APPLICATION USAGE

The *gethostent()*, *gethostbyaddr()*, and *gethostbyname()* functions may return pointers to static data, which may be overwritten by subsequent calls to any of these functions.

These functions are generally used with the Internet address family.

SEE ALSO

endservent(), *htonl()*, *inet_addr()*, <netdb.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

endnetent, getnetbyaddr, getnetbyname, getnetent, setnetent — network database functions

SYNOPSIS

```
#include <netdb.h>

void endnetent(void);

struct netent *getnetbyaddr(in_addr_t net, int type);

struct netent *getnetbyname(const char *name);

struct netent *getnetent(void);

void setnetent(int stayopen);
```

DESCRIPTION

The *getnetbyaddr()*, *getnetbyname()* and *getnetent()*, functions each return a pointer to a **netent** structure, the members of which contain the fields of an entry in the network database.

The *getnetent()* function reads the next entry of the database, opening a connection to the database if necessary.

The *getnetbyaddr()* function searches the database from the beginning, and finds the first entry for which the address family specified by *type* matches the **n_addrtype** member and the network number *net* matches the **n_net** member, opening a connection to the database if necessary. The *net* argument is the network number in host byte order.

The *getnetbyname()* function searches the database from the beginning and finds the first entry for which the network name specified by *name* matches the **n_name** member, opening a connection to the database if necessary.

The *setnetent()* function opens and rewinds the database. If the *stayopen* argument is non-zero, the connection to the net database will not be closed after each call to *getnetent()* (either directly, or indirectly through one of the other *getnet*()* functions).

The *endnetent()* function closes the database.

RETURN VALUE

On successful completion, *getnetbyaddr()*, *getnetbyname()* and *getnetent()*, return a pointer to a **netent** structure if the requested entry was found, and a null pointer if the end of the database was reached or the requested entry was not found. Otherwise, a null pointer is returned.

ERRORS

No errors are defined.

APPLICATION USAGE

The *getnetbyaddr()*, *getnetbyname()* and *getnetent()*, functions may return pointers to static data, which may be overwritten by subsequent calls to any of these functions.

These functions are generally used with the Internet address family.

SEE ALSO

<netdb.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

endprotoent, getprotobynumber, getprotobyname, getprotoent, setprotoent — network protocol database functions

SYNOPSIS

```
#include <netdb.h>

void endprotoent(void);

struct protoent *getprotobyname(const char *name);

struct protoent *getprotobynumber(int proto);

struct protoent *getprotoent(void);

void setprotoent(int stayopen);
```

DESCRIPTION

The *getprotobyname()*, *getprotobynumber()* and *getprotoent()*, functions each return a pointer to a **protoent** structure, the members of which contain the fields of an entry in the network protocol database.

The *getprotoent()* function reads the next entry of the database, opening a connection to the database if necessary.

The *getprotobyname()* function searches the database from the beginning and finds the first entry for which the protocol name specified by *name* matches the **p_name** member, opening a connection to the database if necessary.

The *getprotobynumber()* function searches the database from the beginning and finds the first entry for which the protocol number specified by *number* matches the **p_proto** member, opening a connection to the database if necessary.

The *setprotoent()* function opens a connection to the database, and sets the next entry to the first entry. If the *stayopen* argument is non-zero, the connection to the network protocol database will not be closed after each call to *getprotoent()* (either directly, or indirectly through one of the other *getproto*()* functions).

The *endprotoent()* function closes the connection to the database.

RETURN VALUES

On successful completion, *getprotobyname()*, *getprotobynumber()* and *getprotoent()* functions return a pointer to a **protoent** structure if the requested entry was found, and a null pointer if the end of the database was reached or the requested entry was not found. Otherwise, a null pointer is returned.

ERRORS

No errors are defined.

APPLICATION USAGE

The *getprotobyname()*, *getprotobynumber()* and *getprotoent()* functions may return pointers to static data, which may be overwritten by subsequent calls to any of these functions.

These functions are generally used with the Internet address family.

SEE ALSO

<netdb.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

endservent, getservbyport, getservbyname, getservent, setservent — network services database functions

SYNOPSIS

```
#include <netdb.h>

void endservent(void);

struct servent *getservbyname(const char *name, const char *proto);
struct servent *getservbyport(int port, const char *proto);
struct servent *getservent(void);

void setservent(int stayopen);
```

DESCRIPTION

The *getservbyname()*, *getservbyport()* and *getservent()* functions each return a pointer to a **servent** structure, the members of which contain the fields of an entry in the network services database.

The *getservent()* function reads the next entry of the database, opening a connection to the database if necessary.

The *getservbyname()* function searches the database from the beginning and finds the first entry for which the service name specified by *name* matches the **s_name** member and the protocol name specified by *proto* matches the **s_proto** member, opening a connection to the database if necessary. If *proto* is a null pointer, any value of the **s_proto** member will be matched.

The *getservbyport()* function searches the database from the beginning and finds the first entry for which the port specified by *port* matches the **s_port** member and the protocol name specified by *proto* matches the **s_proto** member, opening a connection to the database if necessary. If *proto* is a null pointer, any value of the **s_proto** member will be matched. The *port* argument must be in network byte order.

The *setservent()* function opens a connection to the database, and sets the next entry to the first entry. If the *stayopen* argument is non-zero, the net database will not be closed after each call to the *getservent()* function (either directly, or indirectly through one of the other *getserv*()* functions).

The *endservent()* function closes the database.

RETURN VALUES

On successful completion, *getservbyname()*, *getservbyport()* and *getservent()* return a pointer to a **servent** structure if the requested entry was found, and a null pointer if the end of the database was reached or the requested entry was not found. Otherwise, a null pointer is returned.

ERRORS

No errors are defined.

APPLICATION USAGE

The *port* argument of *getservbyport()* need not be compatible with the port values of all address families.

The *getservent()*, *getservbyname()* and *getservbyport()* functions may return pointers to static data, which may be overwritten by subsequent calls to any of these functions.

These functions are generally used with the Internet address family.

SEE ALSO

endhostent(), *endprotoent()*, *htonl()*, *inet_addr()*, **<netdb.h>**.

CHANGE HISTORY

First released in Issue 4.

NAME

gethostbyaddr, gethostbyname, gethostent — network host database functions

SYNOPSIS

```
#include <netdb.h>

struct hostent *gethostbyaddr(const void *addr, size_t len, int type);
struct hostent *gethostbyname(const char *name);
struct hostent *gethostent(void);
```

DESCRIPTION

Refer to *endhostent()*.

CHANGE HISTORY

First released in Issue 4.

NAME

gethostname — get name of current host

SYNOPSIS

```
#include <unistd.h>

int gethostname(char *name, size_t namelen);
```

DESCRIPTION

The *gethostname()* function returns the standard host name for the current machine. The *namelen* argument specifies the size of the array pointed to by the *name* argument. The returned name is null-terminated, except that if *namelen* is an insufficient length to hold the host name, then the returned name is truncated and it is unspecified whether the returned name is null-terminated.

Host names are limited to 255 bytes.

RETURN VALUE

On successful completion, 0 is returned. Otherwise, -1 is returned.

ERRORS

No errors are defined.

SEE ALSO

gethostid() (in the XSH specification), *uname()*, <unistd.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

getnetbyaddr, getnetbyname, getnetent — network database functions

SYNOPSIS

```
#include <netdb.h>

struct netent *getnetbyaddr(uint32_t net, int type);
struct netent *getnetbyname(const char *name);
struct netent *getnetent(void);
```

DESCRIPTION

Refer to *endnetent()*.

CHANGE HISTORY

First released in Issue 4.

NAME

getprotobynumber, getprotobyname, getprotoent — network protocol database functions

SYNOPSIS

```
#include <netdb.h>

struct protoent *getprotobyname(const char *name);
struct protoent *getprotobynumber(int proto);
struct protoent *getprotoent(void);
```

DESCRIPTION

Refer to *endprotoent()*.

CHANGE HISTORY

First released in Issue 4.

NAME

getservbyport, getservbyname, getservent — network services database functions

SYNOPSIS

```
#include <netdb.h>

struct servent *getservbyname(const char *name, const char *proto);
struct servent *getservbyport(int port, const char *proto);
struct servent *getservent(void);
```

DESCRIPTION

Refer to *endservent()*.

CHANGE HISTORY

First released in Issue 4.

NAME

h_errno — error return value for network database operations

SYNOPSIS

```
extern int h_errno;
```

DESCRIPTION

Refer to *endhostent()*.

The `<netdb.h>` header provides a declaration of *h_errno* as a modifiable l-value of type **int**. The above synopsis is an example; other declarations as modifiable l-values of type **int** may be used in its place.

CHANGE HISTORY

First released in Issue 4.

NAME

htonl, htons, ntohl, ntohs — convert values between host and network byte order

SYNOPSIS

```
#include <arpa/inet.h>

uint32_t htonl(uint32_t hostlong);
uint16_t htons(uint16_t hostshort);
uint32_t ntohl(uint32_t netlong);
uint16_t ntohs(uint16_t netshort);
```

DESCRIPTION

These functions convert 16-bit and 32-bit quantities between network byte order and host byte order.

The **uint32_t** and **uint16_t** types are made available by inclusion of **<inttypes.h>** (see referenced document **XSH**).

RETURN VALUES

The *htonl()* and *htons()* functions return the argument value converted from host to network byte order.

The *ntohl()* and *ntohs()* functions return the argument value converted from network to host byte order.

ERRORS

No errors are defined.

APPLICATION USAGE

These functions are most often used in conjunction with Internet addresses and ports as returned by *gethostent()* and *getservent()*.

On some architectures these functions are defined as macros that expand to the value of their argument.

SEE ALSO

endhostent(), *endservent()*, **<arpa/inet.h>**.

CHANGE HISTORY

First released in Issue 4.

NAME

inet_addr, inet_network, inet_makeaddr, inet_lnaof, inet_netof, inet_ntoa — Internet address manipulation

SYNOPSIS

```
#include <arpa/inet.h>

in_addr_t inet_addr(const char *cp);
in_addr_t inet_lnaof(struct in_addr in);
struct in_addr inet_makeaddr(in_addr_t net, in_addr_t lna);
in_addr_t inet_netof(struct in_addr in);
in_addr_t inet_network(const char *cp);
char *inet_ntoa(struct in_addr in);
```

DESCRIPTION

The *inet_addr()* function converts the string pointed to by *cp*, in the Internet standard dot notation, to an integer value suitable for use as an Internet address.

The *inet_lnaof()* function takes an Internet host address specified by *in* and extracts the local network address part, in host byte order.

The *inet_makeaddr()* function takes the Internet network number specified by *net* and the local network address specified by *lna*, both in host byte order, and constructs an Internet address from them.

The *inet_netof()* function takes an Internet host address specified by *in* and extracts the network number part, in host byte order.

The *inet_network()* function converts the string pointed to by *cp*, in the Internet standard dot notation, to an integer value suitable for use as an Internet network number.

The *inet_ntoa()* function converts the Internet host address specified by *in* to a string in the Internet standard dot notation.

All Internet addresses are returned in network order (bytes ordered from left to right).

Values specified using dot notation take one of the following forms:

- a.b.c.d When four parts are specified, each is interpreted as a byte of data and assigned, from left to right, to the four bytes of an Internet address.
- a.b.c When a three-part address is specified, the last part is interpreted as a 16-bit quantity and placed in the rightmost two bytes of the network address. This makes the three-part address format convenient for specifying Class B network addresses as *128.net.host*.
- a.b When a two-part address is supplied, the last part is interpreted as a 24-bit quantity and placed in the rightmost three bytes of the network address. This makes the two-part address format convenient for specifying Class A network addresses as *net.host*.
- a When only one part is given, the value is stored directly in the network address without any byte rearrangement.

All numbers supplied as parts in dot notation may be decimal, octal, or hexadecimal, as specified in the ISO C standard (that is, a leading 0x or 0X implies hexadecimal; otherwise, a leading 0 implies octal; otherwise, the number is interpreted as decimal).

RETURN VALUE

Upon successful completion, *inet_addr()* returns the Internet address. Otherwise, it returns (**in_addr_t**)(-1).

Upon successful completion, *inet_network()* returns the converted Internet network number. Otherwise, it returns (**in_addr_t**)(-1).

The *inet_makeaddr()* function returns the constructed Internet address.

The *inet_lnaof()* function returns the local network address part.

The *inet_netof()* function returns the network number.

The *inet_ntoa()* function returns a pointer to the network address in Internet-standard dot notation.

ERRORS

No errors are defined.

APPLICATION USAGE

The return value of *inet_ntoa()* may point to static data that may be overwritten by subsequent calls to *inet_ntoa()*.

SEE ALSO

endhostent(), *endnetent()*, **<arpa/inet.h>**.

CHANGE HISTORY

First released in Issue 4.

NAME

ntohl, ntohs — convert values between host and network byte order

SYNOPSIS

```
#include <arpa/inet.h>
uint32_t ntohl(uint32_t netlong);
uint16_t ntohs(uint16_t netshort);
```

DESCRIPTION

Refer to *htonl()*.

CHANGE HISTORY

First released in Issue 4.

sethostent()

NAME

sethostent — network host database function

SYNOPSIS

```
#include <netdb.h>

void sethostent(int stayopen);
```

DESCRIPTION

Refer to *endhostent()*.

CHANGE HISTORY

First released in Issue 4.

NAME

setnetent — network database function

SYNOPSIS

```
#include <netdb.h>

void setnetent(int stayopen);
```

DESCRIPTION

Refer to *endnetent()*.

CHANGE HISTORY

First released in Issue 4.

setprotoent()

NAME

setprotoent — network protocol database function

SYNOPSIS

```
#include <netdb.h>

void setprotoent(int stayopen);
```

DESCRIPTION

Refer to *endprotoent()*.

CHANGE HISTORY

First released in Issue 4.

NAME

setservent — network services database function

SYNOPSIS

```
#include <netdb.h>

void setservent(int stayopen);
```

DESCRIPTION

Refer to *endservent()*.

CHANGE HISTORY

First released in Issue 4.

IP Address Resolution Headers

This chapter provides reference manual pages on the headers for the Address Resolution API.

NAME

arpa/inet.h — definitions for internet operations

SYNOPSIS

```
#include <arpa/inet.h>
```

DESCRIPTION

The **<arpa/inet.h>** header makes available the type **in_port_t** and the type **in_addr_t** as defined in the description of **<netinet/in.h>**.

The **<arpa/inet.h>** header makes available the **in_addr** structure, as defined in the description of **<netinet/in.h>**.

The following may be declared as functions, or defined as macros, or both:

```
uint32_t htonl(uint32_t hostlong);
uint16_t htons(uint16_t hostshort);
uint32_t ntohl(uint32_t netlong);
uint16_t ntohs(uint16_t netshort);
```

The **uint32_t** and **uint16_t** types are made available by inclusion of **<inttypes.h>** (see referenced document **XSH**).

The following are declared as functions, and may also be defined as macros:

```
in_addr_t      inet_addr(const char *cp);
in_addr_t      inet_lnaof(struct in_addr in);
struct in_addr inet_makeaddr(in_addr_t net, in_addr_t lna);
in_addr_t      inet_netof(struct in_addr in);
in_addr_t      inet_network(const char *cp);
char           *inet_ntoa(struct in_addr in);
```

Inclusion of the **<arpa/inet.h>** header may also make visible all symbols from **<netinet/in.h>** and **<inttypes.h>**.

SEE ALSO

htonl(), *inet_addr()*, **<netinet/in.h>**.

CHANGE HISTORY

First released in Issue 4.

NAME

netdb.h — definitions for network database operations

SYNOPSIS

```
#include <netdb.h>
```

DESCRIPTION

The <netdb.h> header may make available the type **in_port_t** and the type **in_addr_t** as defined in the description of <netinet/in.h>.

The <netdb.h> header defines the **hostent** structure that includes at least the following members:

```
char  *h_name      Official name of the host.
char  **h_aliases  A pointer to an array of pointers to alternative host names,
                  terminated by a null pointer.
int    h_addrtype  Address type.
int    h_length    The length, in bytes, of the address.
char  **h_addr_list A pointer to an array of pointers to network addresses (in
                  network byte order) for the host, terminated by a null pointer.
```

The <netdb.h> header defines the **netent** structure that includes at least the following members:

```
char  *n_name      Official, fully-qualified (including the domain) name of the host.
char  **n_aliases  A pointer to an array of pointers to alternative network names,
                  terminated by a null pointer.
int    n_addrtype  The address type of the network.
uint32_t n_net     The network number, in host byte order.
```

The **uint_32_t** type is made available by inclusion of <<inttypes.h>> (see referenced document **XSH**).

The <netdb.h> header defines the **protoent** structure that includes at least the following members:

```
char  *p_name      Official name of the protocol.
char  **p_aliases  A pointer to an array of pointers to alternative protocol names,
                  terminated by a null pointer.
int    p_proto     The protocol number.
```

The <netdb.h> header defines the **servent** structure that includes at least the following members:

```
char  *s_name      Official name of the service.
char  **s_aliases  A pointer to an array of pointers to alternative service names,
                  terminated by a null pointer.
int    s_port      The port number at which the service resides, in network byte order.
char  *s_proto     The name of the protocol to use when contacting the service.
```

The <netdb.h> header defines the macro **IPPORT_RESERVED** with the value of the highest reserved Internet port number.

The <netdb.h> header provides a declaration of **h_errno** as a modifiable l-value of type **int**. For example:

```
extern int h_errno;
```

The **<netdb.h>** header defines the following macros for use as error values for *gethostbyaddr()* and *gethostbyname()*:

```
HOST_NOT_FOUND
NO_DATA
NO_RECOVERY
TRY_AGAIN
```

The following are declared as functions, and may also be defined as macros:

```
void          endhostent(void);
void          endnetent(void);
void          endprotoent(void);
void          endservent(void);
struct hostent *gethostbyaddr(const void *addr, size_t len, int type);
struct hostent *gethostbyname(const char *name);
struct hostent *gethostent(void);
struct netent *getnetbyaddr(uint32_t net, int type);
struct netent *getnetbyname(const char *name);
struct netent *getnetent(void);
struct protoent *getprotobyname(const char *name);
struct protoent *getprotobynumber(int proto);
struct protoent *getprotoent(void);
struct servent *getservbyname(const char *name, const char *proto);
struct servent *getservbyport(int port, const char *proto);
struct servent *getservent(void);
void          sethostent(int stayopen);
void          setnetent(int stayopen);
void          setprotoent(int stayopen);
void          setservent(int stayopen);
```

Inclusion of the **<netdb.h>** header may also make visible all symbols from **<netinet/in.h>** and **<inttypes.h>**.

SEE ALSO

endhostent(), *endnetent()*, *endprotoent()*, *endservent()*.

CHANGE HISTORY

First released in Issue 4.

NAME

netinet/in.h — Internet Protocol family

SYNOPSIS

```
#include <netinet/in.h>
```

DESCRIPTION

When header file <netinet.h> is included, the following types are defined through **typedef**.

in_port_t An unsigned integral type of exactly 16 bits.

in_addr_t An unsigned integral type of exactly 32 bits.

The <netinet/in.h> header defines the **in_addr** structure that includes at least the following member:

```
in_addr_t      s_addr
```

The <netinet/in.h> header defines the **sockaddr_in** structure that includes at least the following member:

```
sa_family_t    sin_family
in_port_t      sin_port
struct in_addr sin_addr
unsigned char  sin_zero[8]
```

The **sockaddr_in** structure is used to store addresses for the Internet protocol family. Values of this type must be cast to **struct sockaddr** for use with the socket interfaces defined in this document.

The <netinet/in.h> header defines the type **sa_family_t** as described in <sys/socket.h>.

The <netinet/in.h> header defines the following macros for use as values of the *level* argument of *getsockopt()* and *setsockopt()*:

```
IPPROTO_IP      Dummy for IP.
IPPROTO_ICMP    Control message protocol.
IPPROTO_TCP     TCP.
IPPROTO_UDP     User datagram protocol.
```

The <netinet/in.h> header defines the following macros for use as destination addresses for *connect()*, *sendmsg()* and *sendto()*:

```
INADDR_ANY      Local host address.
INADDR_BROADCAST Broadcast address.
```

ntohl(), *ntohs()*, *htonl()* and *htons()* as defined in the description of <arpa/inet.h> are available. Inclusion of the <netinet/in.h> header may also make visible all symbols from <arpa/inet.h>.

SEE ALSO

getsockopt(), *setsockopt()*. <sys/socket.h>.

CHANGE HISTORY

First released in Issue 4.

NAME

unistd.h — standard symbolic constants and types

Note: The **XSH** specification contains the basic definition of this interface. The following additional information pertains to IP Address Resolution.

DESCRIPTION

The following is declared as a function and may also be defined as a macro:

```
int gethostname(char *address, int address_len);
```

SEE ALSO

gethostname().

CHANGE HISTORY

First released in Issue 4.

ISO Transport Protocol Information

A.1 General

This appendix describes the protocol-specific information that is relevant for ISO transport providers. This appendix also describes the protocol-specific information that is relevant when ISO transport services are provided over a TCP network¹².

In general, this Appendix describes the characteristics that the ISO and ISO-over-TCP transport providers have in common, with notes indicating where they differ.

Notes:

1. Protocol address:

In an ISO environment, the protocol address is the transport address.

2. Sending data of zero octets:

The transport service definition, both in connection-mode and in connectionless-mode, does not permit sending a TSDU of zero octets. So, in connectionless-mode, if the *len* parameter is set to zero, the *t_sndudata()* call will always return unsuccessfully with *-1* and *t_errno* set to [TBADDDATA]. In connection-mode, if the *nbytes* parameter is set to zero, the *t_snd()* call will return with *-1* and *t_errno* set to [TBADDDATA] if either the T_MORE flag is set, or the T_MORE flag is not set and the preceding *t_snd()* call completed a TSDU or ETSDU (that is, the call has requested sending a zero byte TSDU or ETSDU).

3. An ISO-over-TCP transport provider does not provide the connectionless-mode.

This Appendix also defines the data structures and constants required for ISO and ISO-over-TCP transport providers which are exposed through the `<xti_osi.h>` header file.

Applications written to compilation environments earlier than those required by this issue of the specification (see Section 1.3 on page 3) and defining `_XOPEN_SOURCE` to be less than 500, will have these data structures and constants exposed through the inclusion of `<xti.h>`

12. The mapping for ISO-over-TCP that is referred to in this Appendix is that defined by RFC-1006: *ISO Transport Service on top of the TCP*, Version 3, May 1987, Marshall T Rose and Dwight E Cass, Network Working Group, Northrop Research & Technology Center. See also the *X/Open Guide to IPS-OSI Coexistence and Migration*. The relevant sections are 4.6.2 (Implementation of OSI Services over IPS) and 4.6.3 (Comments).

A.2 Options

Options are formatted according to the structure **t_opthdr** as described in Chapter 6. A transport provider compliant to this specification supports none, all or any subset of the options defined in Section A.2.1 and Section A.2.2 on page 220. An implementation may restrict the use of any of these options by offering them only in the privileged or read-only mode. An ISO-over-TCP provider supports a subset of the options defined in Section A.2.1.

A.2.1 Connection-mode Service

The protocol level of all subsequent options is ISO_TP.

All options are association-related, that is, they are options with end-to-end significance (see Chapter 6). They may be negotiated in the XTI states T_IDLE and T_INCON, and are read-only in all other states except T_UNINIT.

A.2.1.1 Options for Quality of Service and Expedited Data

These options are all defined in the ISO 8072:1986 transport service definition (see the ISO Transport references). The definitions are not repeated here.

Option Name	Type of Option Value	Legal Option Value	Meaning
T_TCO_THROUGHPUT	struct thrpt	octets per second	throughput
T_TCO_TRANSDEL	struct transdel	time in milliseconds	transit delay
T_TCO_RESERRORRATE	struct rate	OPT_RATIO	residual error rate
T_TCO_TRANSFFAILPROB	struct rate	OPT_RATIO	transfer failure probability
T_TCO_ESTFAILPROB	struct rate	OPT_RATIO	connection establ. failure probability
T_TCO_RELFAILPROB	struct rate	OPT_RATIO	connection release failure probability
T_TCO_ESTDELAY	struct rate	time in milliseconds	connection establ. delay
T_TCO_RELDELAY	struct rate	time in milliseconds	connection release delay
T_TCO_CONNRRESIL	struct rate	OPT_RATIO	connection resilience
T_TCO_PROTECTION	t_uscalar_t	see text	protection
T_TCO_PRIORITY	t_uscalar_t	see text	priority
T_TCO_EXPD	t_uscalar_t	T_YES/T_NO	expedited data

Table A-1 Options for Quality of Service and Expedited Data

OPT_RATIO is defined as $OPT_RATIO = -\log_{10}(\text{ratio})$. The *ratio* is dependent on the parameter, but is always composed of a number of failures divided by a total number of samples. This may be, for example, the number of TSDUs transferred in error divided by the total number of TSDU transfers (T_TCO_RESERRORRATE).

Absolute Requirements

For the options in Table A-1 on page 216, the transport user can indicate whether the request is an absolute requirement or whether a degraded value is acceptable. For the QOS options based on **struct rate** an absolute requirement is specified via the field *minacceptvalue*, if that field is given a value different from T_UNSPEC. The value specified for T_TCO_PROTECTION is an absolute requirement if the T_ABSREQ flag is set. The values specified for T_TCO_EXPD and T_TCO_PRIORITY are never absolute requirements.

Further Remarks

A detailed description of the options for Quality of Service can be found in the ISO 8072:1986 specification. The field elements of the structures in use for the option values are self-explanatory. Only the following details remain to be explained.

- If these options are returned with *t_listen()*, their values are related to the incoming connection and not to the transport endpoint where *t_listen()* was issued. To give an example, the value of T_TCO_PROTECTION is the value sent by the calling transport user, and not the value currently effective for the endpoint (that could be retrieved by *t_optmgmt()* with the flag T_CURRENT set). The option is not returned at all if the calling user did not specify it. An analogous procedure applies for the other options. See also Chapter 6.
- If, in a call to *t_accept()*, the called transport user tries to negotiate an option of higher quality than proposed, the option is rejected and the connection establishment fails (see Section 6.3.4 on page 43).
- The values of the QOS options T_TCO_THROUGHPUT, T_TCO_TRANSDEL, T_TCO_RESERRORRATE, T_TCO_TRANSFFAILPROB, T_TCO_ESTFAILPROB, T_TCO_RELFAILPROB, T_TCO_ESTDELAY, T_TCO_RELDELAY and T_TCO_CONNRRESIL have a structured format. A user requesting one of these options might leave a field of the structure unspecified by setting it to T_UNSPEC. The transport provider is then free to select an appropriate value for this field. The transport provider may return T_UNSPEC in a field of the structure to the user to indicate that it has not yet decided on a definite value for this field.

T_UNSPEC is not a legal value for T_TCO_PROTECTION, T_TCO_PRIORITY and T_TCO_EXPD.

- T_TCO_THROUGHPUT and T_TCO_TRANSDEL
If *avgthrp* (average throughput) is not defined (both fields set to T_UNSPEC), the transport provider considers that the average throughput has the same values as the maximum throughput (*maxthrp*). An analogous procedure applies to T_TCO_TRANSDEL.
- The ISO specification ISO 8073:1986 does not differentiate between average and maximum transit delay. Transport providers that support this option adopt the values of the maximum delay as input for the CR TPDU.
- T_TCO_PROTECTION
This option defines the general level of protection. The symbolic constants in the following list are used to specify the required level of protection:

T_NOPROTECT	No protection feature.
T_PASSIVEPROTECT	Protection against passive monitoring.
T_ACTIVEPROTECT	Protection against modification, replay, addition or deletion.

Both flags T_PASSIVEPROTECT and T_ACTIVEPROTECT may be set simultaneously but are exclusive with T_NOPROTECT. If the T_ACTIVEPROTECT or T_PASSIVEPROTECT

flags are set, the user may indicate that this is an absolute requirement by also setting the T_ABSREQ flag.

- T_TCO_PRIORITY

Five priority levels are defined by XTI:

T_PRIDFLT	Lower level.
T_PRILOW	Low level.
T_PRIMID	Medium level.
T_PRIHIGH	High level.
T_PRITOP	Higher level.

- It is recommended that transport users avoid expedited data with an ISO-over-TCP transport provider, since the RFC 1006 treatment of expedited data does not meet the data reordering requirements specified in ISO 8072:1986, and may not be supported by the provider.

The number of priority levels is not defined by ISO 8072:1986. The parameter only has meaning in the context of some management entity or structure able to judge relative importance.

A.2.1.2 Management Options

These options are parameters of an ISO transport protocol according to ISO 8073:1986. They are not included in the ISO transport service definition ISO 8072:1986, but are additionally offered by XTI. Transport users wishing to be truly ISO-compliant should thus not adhere to them. T_TCO_LTPDU is the only management option supported by an ISO-over-TCP transport provider.

Avoid specifying both QOS parameters and management options at the same time.

Option Name	Type of Option Value	Legal Option Value	Meaning
T_TCO_LTPDU	t_uscalar_t	length in octets	maximum length of TPDU
T_TCO_ACKTIME	t_uscalar_t	time in milliseconds	acknowledge time
T_TCO_REASTIME	t_uscalar_t	time in seconds	reassignment time
T_TCO_PREFCLASS	t_uscalar_t	see text	preferred class
T_TCO_ALTCLASS1	t_uscalar_t	see text	1st alternative class
T_TCO_ALTCLASS2	t_uscalar_t	see text	2nd alternative class
T_TCO_ALTCLASS3	t_uscalar_t	see text	3rd alternative class
T_TCO_ALTCLASS4	t_uscalar_t	see text	4th alternative class
T_TCO_EXTFORM	t_uscalar_t	T_YES/T_NO/T_UNSPEC	extended format
T_TCO_FLOWCTRL	t_uscalar_t	T_YES/T_NO/T_UNSPEC	flowctrl
T_TCO_CHECKSUM	t_uscalar_t	T_YES/T_NO/T_UNSPEC	checksum
T_TCO_NETEXP	t_uscalar_t	T_YES/T_NO/T_UNSPEC	network expedited data
T_TCO_NETRECPTCF	t_uscalar_t	T_YES/T_NO/T_UNSPEC	use of network receipt confirmation

Table A-2 Management Options

Absolute Requirements

A request for any of these options is considered an absolute requirement.

Further Remarks

- If these options are returned with `t_listen()` their values are related to the incoming connection and not to the transport endpoint where `t_listen()` was issued. That means that `t_optmgmt()` with the flag `T_CURRENT` set would usually yield a different result (see Chapter 6).
- For management options that are subject to peer-to-peer negotiation the following holds: If, in a call to `t_accept()`, the called transport user tries to negotiate an option of higher quality than proposed, the option is rejected and the connection establishment fails (see Section 6.3.4 on page 43).
- A connection-mode transport provider may allow the transport user to select more than one alternative class. The transport user may use the options `T_ALTCLASS1`, `T_ALTCLASS2`, etc. to denote the alternatives. A transport provider only supports an implementation-dependent limit of alternatives and ignores the rest.
- The value `T_UNSPEC` is legal for all options in Table A-2 on page 218. It may be set by the user to indicate that the transport provider is free to choose any appropriate value. If returned by the transport provider, it indicates that the transport provider has not yet decided on a specific value.
- Legal values for the options `T_PREFCLASS`, `T_ALTCLASS1`, `T_ALTCLASS2`, `T_ALTCLASS3` and `T_ALTCLASS4` are `T_CLASS0`, `T_CLASS1`, `T_CLASS2`, `T_CLASS3`, `T_CLASS4` and `T_UNSPEC`.
- If a connection has been established, `T_TCO_PREFCLASS` will be set to the selected value, and `T_ALTCLASS1` through `T_ALTCLASS4` will be set to `T_UNSPEC`, if these options are supported.
- *Warning* on the use of `T_TCO_LTPDU`: Sensible use of this option requires that the application programmer knows about system internals. Careless setting of either a lower or a higher value than the implementation-dependent default may degrade the performance.

Legal values for an ISO transport provider are `T_UNSPEC` and multiples of 128 up to $128 \cdot (2^{32} - 1)$ or the largest multiple of 128 that will fit in a `t_uscalar_t`. Values other than powers of 2 between 2^7 and 2^{13} are only supported by transport providers that conform to the 1992 update to ISO 8073.

Legal values for an ISO-over-TCP provider are `T_UNSPEC` and any power of 2 between 2^{**7} and 2^{**11} , and 65531.

The action taken by a transport provider is implementation-dependent if a value is specified which is not exactly as defined in ISO 8073:1986 or its addendums.

- The management options are not independent of one another, and not independent of the options defined in Section A.2.1.1 on page 216. A transport user must take care not to request conflicting values. If conflicts are detected at negotiation time, the negotiation fails according to the rules for absolute requirements (see Chapter 6). Conflicts that cannot be detected at negotiation time will lead to unpredictable results in the course of communication. Usually, conflicts are detected at the time the connection is established.

Some relations that must be obeyed are:

- If T_TCO_EXP is set to T_YES and T_TCO_PREFCLASS is set to T_CLASS2, T_TCO_FLOWCTRL must also be set to T_YES.
- If T_TCO_PREFCLASS is set to T_CLASS0, T_TCO_EXP must be set to T_NO.
- The value in T_TCO_PREFCLASS must not be lower than the value in T_TCO_ALTCLASS1, T_TCO_ALTCLASS2, and so on.
- Depending on the chosen QOS options, further value conflicts might occur.

A.2.2 Connectionless-mode Service

The protocol level of all subsequent options is ISO_TP (as in Section A.2.1 on page 216).

All options are association-related, that is, they are options with end-to-end significance (see Chapter 6). They may be negotiated in all XTI states but T_UNINIT.

A.2.2.1 Options for Quality of Service

These options are all defined in the ISO 8072/Add.1:1986 transport service definition (see the ISO Transport references). The definitions are not repeated here. None of these options are supported by an ISO-over-TCP transport provider, since it does not support connectionless-mode.

Option Name	Type of Option Value	Legal Option Value	Meaning
T_TCL_TRANSDEL	struct rate	time in milliseconds	transit delay
T_TCL_RESERRORRATE	struct rate	OPT_RATIO	residual error rate
T_TCL_PROTECTION	t_uscalar_t	see text	protection
T_TCL_PRIORITY	t_uscalar_t	see text	priority

Table A-3 Options for Quality of Service

Absolute Requirements

A request for any of these options is an absolute requirement.

Further Remarks

A detailed description of the options for Quality of Service can be found in ISO 8072/Add.1:1986. The field elements of the structures in use for the option values are self-explanatory. Only the following details remain to be explained.

- These options are negotiated only between the local user and the local transport provider.
- The meaning, type of option value, and the range of legal option values are identical for T_TCO_RESERRORRATE and T_TCL_RESERRORRATE, T_TCO_PRIORITY and T_TCL_PRIORITY, T_TCO_PROTECTION and T_TCL_PROTECTION (see Table A-1 on page 216, ISO 8072:1986).
- T_TCL_TRANSDEL and T_TCO_TRANSDEL are different. T_TCL_TRANSDEL specifies the maximum transit delay expected during a datagram transmission. Note that the type of option value is a **struct rate** contrary to the **struct transdel** of T_TCO_TRANSDEL. The range of legal option values for each field of **struct rate** is the same as that of T_TCO_TRANSDEL.

- If these options are returned with *t_rcvudata()* their values are related to the received datagram and not to the transport endpoint where *t_rcvudata()* was issued. On the other hand, *t_optmgmt()* with the flag T_CURRENT set returns the values that are currently effective for outgoing datagrams.
- The function *t_rcvuderr()* returns the option value of the data unit previously sent that produced the error.

A.2.2.2 Management Options

This option is a parameter of an ISO transport protocol, according to ISO 8602. It is not included in the ISO transport service definition ISO 8072/Add.1:1986, but is an additional offer by XTI. Transport users wishing to be truly ISO-compliant should thus not adhere to it.

Avoid specifying both QOS parameters and this management option at the same time.

Option Name	Type of Option Value	Legal Option Value	Meaning
T_TCL_CHECKSUM	t_uscalar_t	T_YES/T_NO	checksum computation

Table A-4 Management Option

Absolute Requirements

A request for this option is an absolute requirement.

Further Remarks

T_TCL_CHECKSUM

This is the option allows disabling/enabling of the checksum computation. The legal values are T_YES (checksum enabled) and T_NO (checksum disabled).

If this option is returned with *t_rcvudata()*, its value indicates whether or not a checksum was present in the received datagram.

The advisability of turning off the checksum check is controversial.

A.3 Functions

- t_accept()* The parameter *call*→*udata.len* must be in the range 0 to 32. The user may send up to 32 octets of data when accepting the connection.
- If *fd* is not equal to *resfd*, *resfd* should either be in state T_UNBND or be in state T_IDLE and be bound to the same address as *fd* with the *qlen* parameter set to 0.
- A process can listen for an incoming indication on a given *fd* and then accept the connection on another endpoint *resfd* which has been bound to the same or a different protocol address with the *qlen* parameter (of the *t_bind()* function) set to 0. The protocol address bound to the new accepting endpoint (*resfd*) should in general be the same as the listening endpoint (*fd*), because at the present time, the ISO transport service definition (ISO 8072:1986) does not authorise acceptance of an incoming connection indication with a responding address different from the called address, except under certain conditions (see ISO 8072:1986 paragraph 12.2.4, Responding Address), but it also states that it may be changed in the future.
- t_bind()* The *addr* field of the *t_bind()* structure represents the local TSAP.
- t_close()* The *t_close()* call will cause a *close()* call to be made on the descriptor of this XTI communication endpoint. If there are no other descriptors in this process or any other process which reference this communication endpoint, the *close()* call will perform an abortive release on any connection associated with this endpoint.
- t_connect()* The *sndcall*→*addr* structure specifies the remote called TSAP. In the present version, the returned address set in *rcvcall*→*addr* will have the same value.
- The setting of *sndcall*→*udata* is optional for ISO connections, but with no data, the *len* field of *udata* must be set to 0. The *maxlen* and *buf* fields of the **netbuf** structure, pointed to by *rcvcall*→*addr* and *rcvcall*→*opt*, must be set before the call.
- t_getinfo()* The information returned by *t_getinfo()* reflects the characteristics of the transport connection or, if no connection is established, the maximum characteristics a transport connection could take on using the underlying transport provider. In all possible states except T_DATAXFER, the function *t_getinfo()* returns in the parameter *info* the same information as was returned by *t_open()*. In T_DATAXFER, however, the information returned may differ from that returned by *t_open()*, depending on:
- the transport class negotiated during connection establishment (ISO transport provider only)
 - the negotiation of expedited data transfer for this connection.

In T_DATAXFER, the *etsdu* field in the **t_info** structure is set to T_INVALID (-2) if no expedited data transfer was negotiated, and to 16 otherwise. The remaining fields are set according to the characteristics of the transport protocol class in use for this connection, as defined in the following table.

Parameters	Before Call	After Call			
		Connection Class 0	Connection Class 1-4	Connectionless	ISO-over-TCP
<i>fd</i>	x	/	/	/	/
<i>info</i> → <i>addr</i>	/	x	x	x	x
<i>info</i> → <i>options</i>	/	x Note 1	x Note 1	x Note 1	x Note 1
<i>info</i> → <i>tsdu</i>	/	x Note 2	x Note 2	0→63488	x Note 2
<i>info</i> → <i>etsdu</i>	/	T_INVALID	16 / T_INVALID see Note 3	T_INVALID	16 / T_INVALID see Note 3
<i>info</i> → <i>connect</i>	/	T_INVALID	32	T_INVALID	32 / T_INVALID
<i>info</i> → <i>discon</i>	/	T_INVALID	64	T_INVALID	64 / T_INVALID
<i>info</i> → <i>servtype</i>	/	T_COTS	T_COTS	T_CLTS	T_COTS
<i>info</i> → <i>flags</i>	/	0	0	0	0

Note 1: 'x' equals T_INVALID (-2) or an integral number greater than zero.

Note 2: 'x' equals T_INFINITE (-1) or an integral number greater than zero.

Note 3: Depending on the negotiation of expedited data transfer.

For RFC 1006 (ISO over TCP) support, the value of *info*→*etsdu* depends on the negotiation of expedited data transfer.

t_listen()

The *call*→*addr* structure contains the remote calling TSAP. Since, at most, 32 octets of data will be returned with the connection indication, *call*→*udata.maxlen* should be set to 32 before the call to *t_listen()*.

If the user has set *qlen* greater than 1 (on the call to *t_bind()*), the user may queue up several connection indications before responding to any of them. The user should be forewarned that the ISO transport provider may start a timer to be sure of obtaining a response to the connection request in a finite time. So if the user queues the connection indications for too long before responding to them, the transport provider initiating the connection will disconnect it.

t_open()

The function *t_open()* is called as the first step in the initialisation of a transport endpoint. This function returns various default characteristics associated with the different classes. According to ISO 8073:1986, an OSI transport provider supports one or several out of five different transport protocols, class 0 through class 4. The default characteristics returned in the parameter *info* are those of the highest-numbered protocol class the transport provider is able to support. If, for example, a transport provider supports classes 2 and 0, the characteristics returned are those of class 2. If the transport provider is limited to class 0, the characteristics returned are those of class 0.

The following table gives the characteristics associated with the different classes.

Parameters	Before Call	After Call			
		Connection Class 0	Connection Class 1-4	Connectionless	ISO-over-TCP
<i>name</i>	x	/	/	/	/
<i>oflag</i>	x	/	/	/	/
<i>info</i> → <i>addr</i>	/	x	x	x	x
<i>info</i> → <i>options</i>	/	x Note 1	x Note 1	x Note 1	x Note 1
<i>info</i> → <i>tsdu</i>	/	x Note 2	x Note 2	0→63488	x Note 2
<i>info</i> → <i>etsdu</i>	/	T_INVALID	16	T_INVALID	16/T_INVALID
<i>info</i> → <i>connect</i>	/	T_INVALID	32	T_INVALID	32/T_INVALID
<i>info</i> → <i>discon</i>	/	T_INVALID	64	T_INVALID	64/T_INVALID
<i>info</i> → <i>servtype</i>	/	T_COTS	T_COTS	T_CLTS	T_COTS
<i>info</i> → <i>flags</i>	/	0	0	0	0

Note 1: 'x' equals T_INVALID (-2) or an integral number greater than zero.

Note 2: 'x' equals T_INFINITE (-1) or an integral number greater than zero.

- t_rcv()* If expedited data arrives after part of a TSDU has been retrieved, receipt of the remainder of the TSDU will be suspended until the ETSDU has been processed. Only after the full ETSDU has been retrieved (T_MORE not set), will the remainder of the TSDU be available to the user.
- t_rcvconnect()* On return, the *call*→*addr* structure contains the remote calling TSAP. Since, at most, 32 octets of data will be returned to the user, *call*→*udata.maxlen* should be set to 32 before the call to *t_rcvconnect()*.
- t_rcvdis()* Since, at most, 64 octets of data will be returned to the user, *discon*→*udata.maxlen* should be set to 64 before the call to *t_rcvdis()*.
- t_rcvudata()* The *unitdata*→*addr* structure specifies the remote TSAP. If the T_MORE flag is set, an additional *t_rcvudata()* call is needed to retrieve the entire TSDU. Only normal data is returned via the *t_rcvudata()* call. This function is not supported by an ISO-over-TCP transport provider.
- t_rcvuderr()* The *uderr*→*addr* structure contains the remote TSAP.
- t_snd()* Zero byte TSDUs are not supported. The T_EXPEDITED flag is not a legal flag unless expedited data has been negotiated for this connection.
- t_snddis()* Since, at most, 64 octets of data may be sent with the disconnect, *call*→*udata.len* will have a value less than or equal to 64.
- t_sndudata()* The *unitdata*→*addr* structure specifies the remote TSAP. The ISO connectionless-mode transport service does not support the sending of expedited data. This function is not supported by an ISO-over-TCP transport provider.

A.4 The <xti_osi.h> Header File

This section presents definitions for the <xti_osi.h> header. The values defined for some of the symbolic constants are designated as not mandatory for conformance purposes. These are identified by a comment in the definitions.

```
/* SPECIFIC ISO OPTION AND MANAGEMENT PARAMETERS */

/*
 * Definition of the ISO transport classes
 */

#define T_CLASS0 0
#define T_CLASS1 1
#define T_CLASS2 2
#define T_CLASS3 3
#define T_CLASS4 4

/*
 * Definition of the priorities.
 */

#define T_PRITOP 0
#define T_PRIHIGH 1
#define T_PRIMID 2
#define T_PRILOW 3
#define T_PRIDFLT 4

/*
 * Definitions of the protection levels
 */

#define T_NOPROTECT 1
#define T_PASSIVEPROTECT 2
#define T_ACTIVEPROTECT 4

/*
 * Default value for the length of TPDUs.
 */

#define T_LTPDUDFLT 128 /* define obsolete in XPG4 */

/*
 * rate structure.
 */
struct rate {
    t_scalar_t targetvalue; /* target value */
    t_scalar_t minacceptvalue; /* value of minimum acceptable quality */
};
```

```

/*
 * reqvalue structure.
 */
struct reqvalue {
    struct rate    called;    /* called rate */
    struct rate    calling;   /* calling rate */
};

/*
 * thrpt structure.
 */
struct thrpt {
    struct reqvalue    maxthrpt;    /* maximum throughput */
    struct reqvalue    avgthrpt;    /* average throughput */
};

/*
 * transdel structure
 */
struct transdel {
    struct reqvalue    maxdel;    /* maximum transit delay */
    struct reqvalue    avgdel;    /* average transit delay */
};

/*
 * Protocol Levels
 *
 * The values defined for these Protocol Levels are recommended-only,
 * not mandatory.
 */

#define T_ISO_TP    0x0100
#define ISO_TP      0x0100    (LEGACY)

/*
 * Options for Quality of Service and Expedited Data (ISO 8072:1994)
 *
 * The values defined for these QoS and Expedited Data are
 * recommended-only, not mandatory.
 */

#define T_TCO_THROUGHPUT    0x0001
#define TCO_THROUGHPUT      0x0001    (LEGACY)
#define T_TCO_TRANSDEL      0x0002
#define TCO_TRANSDEL        0x0002    (LEGACY)
#define T_TCO_RESERRORRATE  0x0003
#define TCO_RESERRORRATE    0x0003    (LEGACY)
#define T_TCO_TRANSFFAILPROB 0x0004
#define TCO_TRANSFFAILPROB  0x0004    (LEGACY)
#define T_TCO_ESTFAILPROB   0x0005
#define TCO_ESTFAILPROB     0x0005    (LEGACY)
#define T_TCO_RELFAILPROB   0x0006
#define TCO_RELFAILPROB     0x0006    (LEGACY)
#define T_TCO_ESTDELAY      0x0007
#define TCO_ESTDELAY        0x0007    (LEGACY)
#define T_TCO_RELDELAY      0x0008

```



```

#define TCO_RELDELAY          0x0008          (LEGACY)
#define T_TCO_CONNRRESIL     0x0009
#define TCO_CONNRRESIL       0x0009          (LEGACY)
#define T_TCO_PROTECTION     0x000a
#define TCO_PROTECTION       0x000a          (LEGACY)
#define T_TCO_PRIORITY       0x000b
#define TCO_PRIORITY         0x000b          (LEGACY)
#define T_TCO_EXPD           0x000c
#define TCO_EXPD             0x000c          (LEGACY)

#define T_TCL_TRANSDEL       0x000d
#define TCL_TRANSDEL         0x000d          (LEGACY)
#define T_TCL_RESERRORRATE   T_TCO_RESERRORRATE
#define TCL_RESERRORRATE     T_TCO_RESERRORRATE (LEGACY)
#define T_TCL_PROTECTION     T_TCO_PROTECTION
#define TCL_PROTECTION       T_TCO_PROTECTION (LEGACY)
#define T_TCL_PRIORITY       T_TCO_PRIORITY
#define TCL_PRIORITY         T_TCO_PRIORITY   (LEGACY)

/*
 * Management Options
 *
 * The values defined for these Management Options are
 * recommended-only, not mandatory.
 */

#define T_TCO_LTPDU          0x0100
#define TCO_LTPDU            0x0100
#define T_TCO_ACKTIME        0x0200
#define TCO_ACKTIME          0x0200
#define T_TCO_REASTIME       0x0300
#define TCO_REASTIME         0x0300
#define T_TCO_EXTFORM        0x0400
#define TCO_EXTFORM          0x0400
#define T_TCO_FLOWCTRL       0x0500
#define TCO_FLOWCTRL         0x0500
#define T_TCO_CHECKSUM       0x0600
#define TCO_CHECKSUM         0x0600
#define T_TCO_NETEXP         0x0700
#define TCO_NETEXP           0x0700
#define T_TCO_NETRECPTCF     0x0800
#define TCO_NETRECPTCF       0x0800
#define T_TCO_PREFCLASS      0x0900
#define TCO_PREFCLASS        0x0900
#define T_TCO_ALTCLASS1      0x0a00
#define TCO_ALTCLASS1        0x0a00
#define T_TCO_ALTCLASS2      0x0b00
#define TCO_ALTCLASS2        0x0b00
#define T_TCO_ALTCLASS3      0x0c00
#define TCO_ALTCLASS3        0x0c00
#define T_TCO_ALTCLASS4      0x0d00
#define TCO_ALTCLASS4        0x0d00

#define T_TCL_CHECKSUM       T_TCO_CHECKSUM
#define TCL_CHECKSUM         T_TCO_CHECKSUM

```


Internet Protocol-specific Information

B.1 General

This appendix describes the protocol-specific information that is relevant for TCP and UDP transport providers.

This Appendix also defines data structures and constants required for TCP and UDP transport providers which are exposed through the `<xti_inet.h>` header file.

Applications written to compilation environments earlier than those required by this issue of the specification (see Section 1.3 on page 3) and defining `_XOPEN_SOURCE` to be less than 500, will have these data structures and constants exposed through the inclusion of `<xti.h>`.

Notes

- T_MORE flag and TSDUs

The notion of TSDU is not supported by TCP transport provider, so the T_MORE flag will be ignored when TCP is used.

- Expedited data

TCP does not have a notion of expedited data in a sense comparable to ISO expedited data. TCP defines an urgent mechanism, by which in-line data is marked for urgent delivery. T_UDP has no urgent mechanism. See the TCP Standard for more detailed information.

- Orderly Release

The orderly release functions `t_sndrel()` and `t_rcvrel()` were defined to support the orderly release facility of TCP. They are the recommended means of releasing a TCP connection. The specification of TCP states that only established connections may be closed with orderly release: that is, on an endpoint in T_DATAXFER or T_INREL state.

- Abortive release

Functions `t_snddis()` and `t_rcvdis()` may be used to perform abortive release over TCP transport. However their use is not recommended as the abortive release primitive (RST segment) is not transmitted reliably by the TCP protocol.

- Connection establishment

TCP does not allow the possibility of refusing a connection indication. Each connection indication causes the TCP transport provider to establish the connection. Therefore, `t_listen()` and `t_accept()` have a semantic which is slightly different from that for ISO providers.

- Connection release

After a connection has been released, the local address bound to the endpoint may be qualified with the local IP address (rather than having a wildcard IP address). Also, the port number itself may have been changed during the `t_accept()` processing. If the endpoint is not being reused immediately then it is recommended that it should be unbound or closed so that other users can successfully bind to the address.

B.2 Options

Options are formatted according to the structure `t_opthdr` as described in Chapter 6. A transport provider compliant to this specification supports none, all or any subset of the options defined in Section B.2.1, Section B.2.2 and Section B.2.3. An implementation may restrict the use of any of these options by offering them only in the privileged or read-only mode.

B.2.1 TCP-level Options

The protocol level is `T_INET_TCP`. For this level, Table B-1 shows the options that are defined.

Option Name	Type of Option Value	Legal Option Value	Meaning
<code>T_TCP_KEEPALIVE</code>	<code>struct t_kpalive</code>	see text	check if connections are alive
<code>T_TCP_MAXSEG</code>	<code>t_uscalar_t</code>	length in octets	get TCP maximum segment size
<code>T_TCP_NODELAY</code>	<code>t_uscalar_t</code>	<code>T_YES/T_NO</code>	don't delay send to coalesce packets

Table B-1 TCP-level Options

These options do *not* have end-to-end significance. They may be negotiated in all XTI states except `T_UNBND` and `T_UNINIT`. They are read-only in state `T_UNBND`. See Chapter 6 for the difference between options that have end-to-end significance and those that do not.

Absolute Requirements

A request for `T_TCP_NODELAY` and a request to activate `T_TCP_KEEPALIVE` is an absolute requirement. `T_TCP_MAXSEG` is a read-only option.

Further Remarks

`T_TCP_KEEPALIVE` If this option is set, a keep-alive timer is activated to monitor idle connections that might no longer exist. If a connection has been idle since the last keep-alive timeout, a keep-alive packet is sent to check if the connection is still alive or broken.

Keep-alive packets are not an explicit feature of TCP, and this practice is not universally accepted. According to RFC 1122 (see Referenced Documents):

“a keep-alive mechanism should only be invoked in server applications that might otherwise hang indefinitely and consume resources unnecessarily if a client crashes or aborts a connection during a network failure”.

The option value consists of a structure `t_kpalive` declared as:

```
struct t_kpalive {
    t_scalar_t kp_onoff;    /* switch option on/off */
    t_scalar_t kp_timeout; /* keep-alive timeout   */
                          /* in minutes           */
}
```

Legal values for the field `kp_onoff` are:

```
T_NO    switch keep-alive timer off
T_YES   activate keep-alive timer
```

The field *kp_timeout* determines the frequency of keep-alive packets being sent, in minutes. The transport user can request the default value by setting the field to T_UNSPEC. The default is implementation-dependent, but at least 120 minutes (see the referenced RFC 1122). Legal values for this field are T_UNSPEC and all positive numbers.

The timeout value is not an absolute requirement. The implementation may pose upper and lower limits to this value. Requests that fall short of the lower limit may be negotiated to the lower limit.

The use of this option might be restricted to privileged users.

T_TCP_MAXSEG This option is read-only. It is used to retrieve the maximum TCP segment size.

T_TCP_NODELAY Under most circumstances, TCP sends data as soon as it is presented. When outstanding data has not yet been acknowledged, it gathers small amounts of output to be sent in a single packet once an acknowledgement is received. For a small number of clients, such as window systems (for example, MIT X Window System) that send a stream of mouse events which receive no replies, this packetisation may cause significant delays. T_TCP_NODELAY is used to defeat this algorithm. Legal option values are T_YES (“don’t delay”) and T_NO (“delay”).

B.2.2 T_UDP-level Options

The protocol level is T_INET_UDP. The option defined for this level is shown in Table B-2.

Option Name	Type of Option Value	Legal Option Value	Meaning
T_UDP_CHECKSUM	t_uscalar_t	T_YES/T_NO	checksum computation

Table B-2 T_UDP-level Option

This option has end-to-end significance. It may be negotiated in all XTI states except T_UNBND and T_UNINIT. It is read-only in state T_UNBND. See Chapter 6 for the difference between options that have end-to-end significance and those that do not.

Absolute Requirements

A request for this option is an absolute requirement.

Further Remarks

T_UDP_CHECKSUM The option allows disabling/enabling of the T_UDP checksum computation. The legal values are T_YES (checksum enabled) and T_NO (checksum disabled).

If this option is returned with *t_rcvdata()*, its value indicates whether a checksum was present in the received datagram or not.

Numerous cases of undetected errors have been reported when applications chose to turn off checksums for efficiency. The advisability of ever turning off the checksum check is very controversial.

B.2.3 T_IP-level Options

The protocol level is T_INET_IP. The options defined for this level are listed in Table B-3.

Option Name	Type of Option Value	Legal Option Value	Meaning
T_IP_BROADCAST	unsigned int	T_YES/T_NO	permit sending of broadcast messages
T_IP_DONTRROUTE	unsigned int	T_YES/T_NO	just use interface addresses
T_IP_OPTIONS	array of unsigned characters	see text	T_IP per-packet options
T_IP_REUSEADDR	unsigned int	T_YES/T_NO	allow local address reuse
T_IP_TOS	unsigned char	see text	T_IP per-packet type of service
T_IP_TTL	unsigned char	time in seconds	T_IP per packet time-to-live

Table B-3 T_IP-level Options

T_IP_OPTIONS and T_IP_TOS are both options with end-to-end significance. All other options do *not* have end-to-end significance. See Chapter 6 for the difference between options with end-to-end significance and options without.

T_IP_REUSEADDR may be negotiated in all XTI states except T_UNINIT. All other options may be negotiated in all other XTI states except T_UNBND and T_UNINIT; they are read-only in the state T_UNBND.

Absolute Requirements

A request for any of these options is an absolute requirement.

Further Remarks

T_IP_BROADCAST This option requests permission to send broadcast datagrams. It was defined to make sure that broadcasts are not generated by mistake. The use of this option is often restricted to privileged users.

T_IP_DONTRROUTE This option indicates that outgoing messages should bypass the standard routing facilities. It is mainly used for testing and development.

T_IP_OPTIONS This option is used to set (retrieve) the OPTIONS field of each outgoing (incoming) IP datagram. Its value is a string of octets composed of a number of T_IP options, whose format matches those defined in the IP specification with one exception: the list of addresses for the source routing options must include the first-hop gateway at the beginning of the list of gateways. The first-hop gateway address will be extracted from the option list and the size adjusted accordingly before use.

The option is disabled if it is specified with “no value”; that is, with an option header only.

The functions *t_connect()* (in synchronous mode), *t_listen()*, *t_rcvconnect()* and *t_rcvudata()* return the OPTIONS field, if any, of the received IP datagram associated with this call. The function *t_rcvuderr()* returns the OPTIONS field of the data unit previously sent that produced the error. The function *t_optmgmt()* with T_CURRENT set retrieves the currently effective T_IP_OPTIONS that is sent with outgoing datagrams.

Common applications never need this option. It is mainly used for network debugging and control purposes.

T_IP_REUSEADDR Many TCP implementations do not allow the user to bind more than one transport endpoint to addresses with identical port numbers. If **T_IP_REUSEADDR** is set to **T_YES** this restriction is relaxed in the sense that it is now allowed to bind a transport endpoint to an address with a port number and an underspecified internet address (“wild card” address) and further endpoints to addresses with the same port number and (mutually exclusive) fully specified internet addresses.

T_IP_TOS This option is used to set (retrieve) the *type-of-service* field of an outgoing (incoming) **T_IP** datagram. This field can be constructed by any OR’ed combination of one of the precedence flags and the type-of-service flags **T_LDELAY**, **T_HITHRPT** and **T_HIRES**:

— Precedence:

These flags specify datagram precedence, allowing senders to indicate the importance of each datagram. They are intended for Department of Defense applications. Legal flags are:

```
T_ROUTINE
T_PRIORITY
T_IMMEDIATE
T_FLASH
T_OVERRIDEFLASH
T_CRITIC_ECP
T_INETCONTROL
T_NETCONTROL.
```

Applications using **T_IP_TOS** but not the precedence level should use the value **T_ROUTINE** for precedence.

— Type of service:

These flags specify the type of service the IP datagram desires. Legal flags are:

```
T_NOTOS      requests no distinguished type of service
T_LDELAY     requests low delay
T_HITHRPT    requests high throughput
T_HIRES      requests high reliability
```

The option value is set using the macro **SET_TOS(*prec,tos*)**, where *prec* is set to one of the precedence flags and *tos* to one or an OR’ed combination of the type-of-service flags. **SET_TOS()** returns the option value.

The functions **t_connect()**, **t_listen()**, **t_rcvconnect()** and **t_rcvudata()** return the *type-of-service* field of the received IP datagram associated with this call. The function **t_rcvuderr()** returns the *type-of-service* field of the data unit previously sent that produced the error.

The function **t_optmgmt()** with **T_CURRENT** set retrieves the currently effective **T_IP_TOS** value that is sent with outgoing datagrams.

The requested *type-of-service* cannot be guaranteed. It is a hint to the routing algorithm that helps it choose among various paths to a destination. Note also, that most hosts and gateways in the Internet these days ignore the *type-of-service* field.

T_IP_TTL

This option¹³ is used to set the *time-to-live* field in an outgoing IP datagram. It specifies how long, in seconds, the datagram is allowed to remain in the Internet. The *time-to-live* field of an incoming datagram is not returned by any function (since it is not an option with end-to-end significance).

13. This is a simplified description. Refer to RFC 1122 (see Referenced Documents) for precise details.

B.3 Functions

- t_accept()* Issuing *t_accept()* assigns an already established connection to *resfd*. Since user data cannot be exchanged during the connection establishment phase, *call->udata.len* must be set to 0. When (*resfd != fd*), the function *t_accept()* is recommended to be called with *resfd* in T_UNBND state, though an endpoint which is bound to any local address (in T_IDLE state) can also be used. After *t_accept()* completes, the endpoint *resfd* will represent a connected TCP endpoint whose complete binding essentially has both local and remote address components. If file descriptor *resfd* was unbound before calling *t_accept()*, after the call completes its local address binding would be to the same protocol address bound to *fd*. If file descriptor *resfd* was bound to a local address before calling *t_accept()*, that local address binding is dissolved and the local address part of the binding after *t_accept()* completes would become same as the address bound to *fd*. If options with end-to-end significance (T_IP_OPTIONS, T_IP_TOS) are to be sent with the connection confirmation, the values of these options must be set with *t_optmgmt()* before the T_LISTEN event occurs. When the transport user detects a T_LISTEN, TCP has already established the connection. Association-related options passed with *t_accept()* become effective at once, but since the connection is already established, they are transmitted with subsequent IP datagrams sent out in the T_DATAXFER state.
- t_bind()* The *addr* field of the **t_bind** structure represents the local socket; that is, an address which specifically includes a port identifier. Some implementations treat port number 0 as a request to bind to any unused port. Other than that value, a port number part of the binding is specific. The IP address part of the binding can represent a single IP address or a wildcard binding to an address that could represent multiple IP addresses that are legal for the host.
- t_close()* The *t_close()* call will result in a *close* call on the descriptor of this XTI communication endpoint. If there are no other descriptors in this process or any other process which reference this communication endpoint, the *close()* call will perform an orderly connection termination according to the rules of a TCP CLOSE call on this connection endpoint as specified in standards RFC 793 and RFC 1122. If the XTI_LINGER option is supported and is used to enable the *linger* option, the linger time will affect the time an implementation lingers in the execution of *t_close()* or **close()**. A linger time of 0 specified with the XTI_LINGER option may cause an abortive release of a TCP connection, resulting in lost data.
- t_connect()* The *sndcall->addr* structure specifies the remote socket. In the present version, the returned address set in *rcvcall->addr* will have the same value. Since user data cannot be exchanged during the connection establishment phase, *sndcall->udata.len* must be set to 0. Note that the peer TCP, and not the peer transport user, confirms the connection.

t_listen() Upon successful return, *t_listen()* indicates an existing connection and not a connection indication.

Since user data cannot be exchanged during the connection establishment phase, *call->udata.maxlen* must be set to 0 before the call to *t_listen()*. The *call->addr* structure contains the remote calling socket.

t_look() As soon as a segment with the TCP urgent pointer set enters the TCP receive buffer, the event T_EXDATA is indicated. T_EXDATA remains set until all data up to the byte pointed to by the TCP urgent pointer has been received. If the urgent pointer is updated, and the user has not yet received the byte previously pointed to by the urgent pointer, the update is invisible to the user.

t_open() *t_open()* is called as the first step in the initialisation of a transport endpoint. This function returns various default characteristics of the underlying transport protocol by setting fields in the **t_info** structure.

The following should be the values returned by the call to *t_open()* and *t_getinfo()* with the indicated transport providers.

Parameters	Before call	After call	
		T_TCP/T_IP	T_UDP/T_IP
<i>name</i>	x	/	/
<i>oflag</i>	x	/	/
<i>info->addr</i>	/	x	x
<i>info->options</i>	/	x	x
<i>info->tsdu</i>	/	0	x
<i>info->etsdu</i>	/	T_INFINITE	T_INVALID
<i>info->connect</i>	/	T_INVALID	T_INVALID
<i>info->discon</i>	/	T_INVALID	T_INVALID
<i>info->servtype</i>	/	T_COTS_ORD	T_CLTS
<i>info->flags</i>	/	0	T_SNDZERO

'x' equals T_INVALID (-2) or an integral number greater than zero.

t_rcv() The T_MORE flag should be ignored if normal data is delivered. If a byte in the data stream is pointed to by the TCP urgent pointer, as many bytes as possible preceding this marked byte and the marked byte itself are denoted as urgent data and are received with the T_EXPEDITED flag set. If the buffer supplied by the user is too small to hold all urgent data, the T_MORE flag will be set, indicating that urgent data still remains to be read. Note that the number of bytes received with the T_EXPEDITED flag set is not necessarily equal to the number of bytes sent by the peer user with the T_EXPEDITED flag set.

t_rcvconnect() Since user data cannot be exchanged during the connection establishment phase, *call->udata.maxlen* must be set to 0 before the call to *t_rcvconnect()*. On return, the *call->addr* structure contains the address of the remote calling endpoint.

t_rcvdis() Since data may not be sent with a disconnect, the *discon->udata* structure will not be meaningful.

t_snd() The T_MORE flag should be ignored. If *t_snd()* is called with more than one byte specified and with the T_EXPEDITED flag set, then the last byte of the buffer will be the byte pointed to by the TCP urgent pointer. If the T_EXPEDITED flag is set, at least one byte must be sent.

Implementor's Note: Data for a `t_snd()` call with the `T_EXPEDITED` flag set may not pass data sent previously.

- `t_snddis()` Since data may not be sent with a disconnect, `call->udata.len` must be set to zero.
- `t_sndudata()` Be aware that the maximum size of a connectionless-mode TSDU varies among implementations.

B.4 The <xti_inet.h> Header File

This section presents definitions for the <xti_inet.h> header. The values defined for some of the symbolic constants are designated as not mandatory for conformance purposes. These are identified by a comment in the definitions.

```

/* INTERNET-SPECIFIC ENVIRONMENT */

/*
 * TCP level
 *
 * The values defined for the TCP Level are recommended-only,
 * not mandatory.
 */

#define T_INET_TCP 0x6
#define INET_TCP 0x6 (LEGACY)

/*
 * TCP-level Options
 *
 * The values defined for the TCP-level Options are recommended-only,
 * not mandatory.
 */

#define T_TCP_NODELAY 0x1 /* don't delay packets to coalesce */
#define TCP_NODELAY 0x1 /* (LEGACY) */
#define T_TCP_MAXSEG 0x2 /* get maximum segment size */
#define TCP_MAXSEG 0x2 /* (LEGACY) */
#define T_TCP_KEEPALIVE 0x8 /* check, if connections are alive */
#define TCP_KEEPALIVE 0x8 /* (LEGACY) */

/*
 * Structure used with TCP_KEEPALIVE option.
 */
struct t_kpalive {
    t_scalar_t kp_onoff; /* option on/off */
    t_scalar_t kp_timeout; /* timeout in minutes */
};

#define T_GARBAGE 0x02

/*
 * UDP level
 *
 * The values defined for the UDP Level are recommended-only,
 * not mandatory.
 */

#define T_INET_UDP 0x11
#define INET_UDP 0x11 /* LEGACY */

```

```

/*
 * UDP-level Options
 *
 * The values defined for the UDP-level Options are recommended-only,
 * not mandatory.
 */
#define T_UDP_CHECKSUM    T_TCO_CHECKSUM    /* checksum computation */
#define UDP_CHECKSUM      T_TCO_CHECKSUM    /* LEGACY */

/*
 * IP level
 *
 * The values defined for the IP Level are recommended-only,
 * not mandatory.
 */
#define T_INET_IP        0x0
#define INET_IP          0x0    LEGACY    */

/*
 * IP-level Options
 *
 * The values defined for the IP-level Options are recommended-only,
 * not mandatory.
 */
#define T_IP_OPTIONS     0x1    /* IP per-packet options */
#define IP_OPTIONS       0x1    /* LEGACY */
#define T_IP_TOS         0x2    /* IP per-packet type of service */
#define IP_TOS           0x2    /* LEGACY */
#define T_IP_TTL         0x3    /* IP per-packet time to live */
#define IP_TTL           0x3    /* LEGACY */
#define T_IP_REUSEADDR   0x4    /* allow local address reuse */
#define IP_REUSEADDR     0x4    /* LEGACY */
#define T_IP_DONTROUTE   0x10   /* just use interface addresses */
#define IP_DONTROUTE     0x10   /* LEGACY */
#define T_IP_BROADCAST   0x20   /* permit sending of broadcast msgs */
#define IP_BROADCAST     0x20   /* LEGACY */

/*
 * IP_TOS precedence levels
 */

#define T_ROUTINE        0
#define T_PRIORITY      1
#define T_IMMEDIATE     2
#define T_FLASH         3
#define T_OVERRIDEFLASH 4
#define T_CRITIC_ECP    5
#define T_INETCONTROL   6
#define T_NETCONTROL    7

```

```
/*
 * IP_TOS type of service
 */

#define T_NOTOS          0
#define T_LDELAY        (1 << 4)
#define T_HITHRPT       (1 << 3)
#define T_HIREL         (1 << 2)
#define T_LOCOST        (1 << 1)

#define SET_TOS(prec, tos) ((0x7 & (prec)) << 5 | (0x1c & (tos)))
```

Guidelines for Use of XTI

C.1 Transport Service Interface Sequence of Functions

In order to describe the allowable sequence of function calls, this section gives some rules regarding the maintenance of the state of the interface:

- It is the responsibility of the transport provider to keep a record of the state of the interface as seen by the transport user.
- The transport provider will not process a function that places the interface out of state.
- If the user issues a function out of sequence, the transport provider will indicate this where possible through an error return on that function. The state will not change. In this case, if any data is passed with the function when not in the T_DATAXFER state, that data will not be accepted or forwarded by the transport provider.
- The uninitialised state (T_UNINIT) of a transport endpoint is the initial state. The endpoint must be initialised and bound before the transport provider may view it as active.
- The uninitialised state is also the final state, and the transport endpoint must be viewed as unused by the transport provider. The *t_close()* function will close the transport endpoint and free the transport library resources for another endpoint.
- According to Table 5-5 on page 36, *t_close()* should only be issued from the T_UNBND state. If it is issued from any other state, and no other user has that endpoint open, the action will be abortive, the transport endpoint will be successfully closed, and the library resources will be freed for another endpoint. When *t_close()* is issued, the transport provider must ensure that the address associated with the specified transport endpoint has been unbound from that endpoint. The provider sends appropriate disconnects if *t_close()* is not issued from the unbound state.

The following rules apply only to the connection-mode transport service:

- The transport connection release phase can be initiated at any time during the connection establishment phase or data transfer phase.
- The only time the state of a transport service interface of a transport endpoint may be transferred to another transport endpoint is when the *t_accept()* function specifies such action. The following rules then apply to the cooperating transport endpoints:
 - The endpoint that is to accept the current state of the interface should either be in state T_UNBND or be in state T_IDLE with the *qlen* parameter set to 0.

C.2 Example in Connection-oriented Mode

Table C-1 on page 243 shows the allowable sequence of functions of an active user and passive user communicating using a connection-mode transport service. This example is not meant to show all the functions that must be called, but rather to highlight the important functions that request a particular service. Blank lines are used to indicate that the function would be called by another user prior to a related function being called by the remote user. For example, the active user calls *t_connect()* to request a connection and the passive user would receive an indication of the connection request (via the return from *t_listen()*) and then would call the *t_accept()*.

The state diagram in Table C-1 on page 243 shows the flow of the events through the various states. The active user is represented by a solid line and the passive user is represented by a dashed line. This example shows a successful connection being established and terminated using connection-mode transport service without orderly release. For a detailed description of all possible states and events, see Table 5-7 on page 37.

Active User	Passive User
<i>t_open()</i>	<i>t_open()</i>
<i>t_bind()</i>	<i>t_bind()</i>
	<i>t_listen()</i>
<i>t_connect()</i>	<i>t_accept()</i>
<i>t_rcvconnect()</i>	
<i>t_snd()</i>	
<i>t_sndv()</i>	<i>t_rcv()</i>
	<i>t_rcvv()</i>
<i>t_snddis()</i>	
	<i>t_rcvdis()</i>
<i>t_unbind()</i>	<i>t_unbind()</i>
<i>t_close()</i>	<i>t_close()</i>

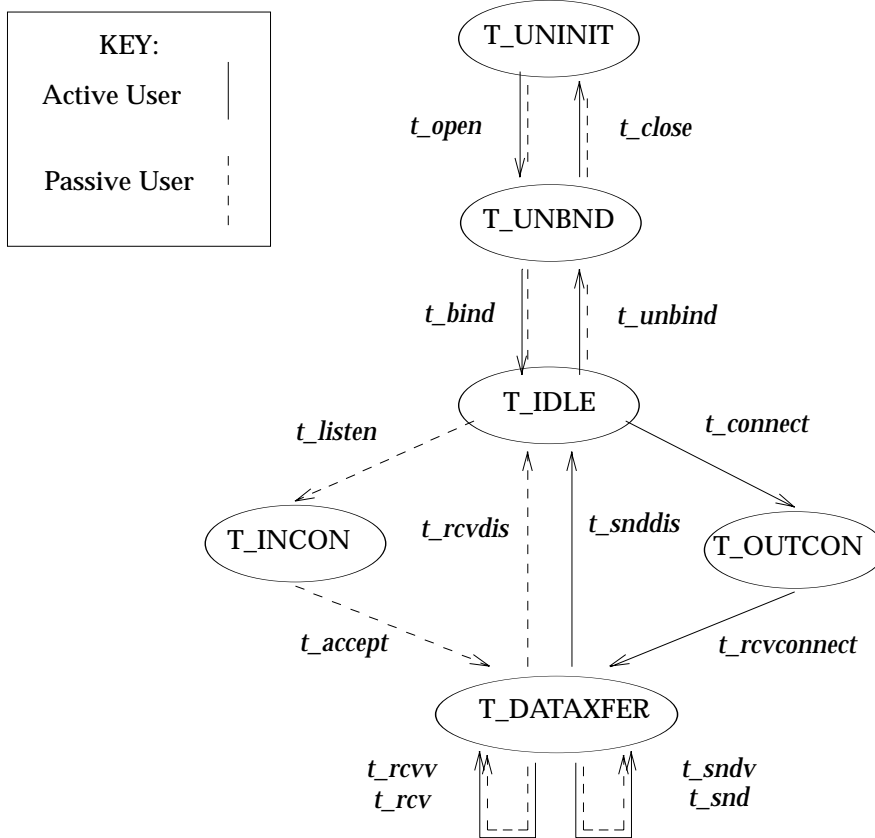


Table C-1 Sequence of Transport Functions in Connection-oriented Mode

C.3 Example in Connectionless Mode

Table C-2 shows the allowable sequence of functions of user A and user B communicating using a connectionless-mode transport service. This example is not meant to show all the functions that must be called but rather to highlight the important functions that request a particular service. Blank lines are used to indicate that a function would be called by another user prior to a related function being called by the remote user.

The state diagram that follows shows the flow of the events through the various states. This example shows a successful exchange of data between user A and user B. For a detailed description of all possible states and events, see Table 5-7 on page 37.

User A	User B
<i>t_open()</i>	<i>t_open()</i>
<i>t_bind()</i>	<i>t_bind()</i>
<i>t_sndudata()</i>	
	<i>t_rcvudata()</i>
<i>t_unbind()</i>	<i>t_unbind()</i>
<i>t_close()</i>	<i>t_close()</i>




Table C-2 Sequence of Transport Functions in Connectionless Mode

C.4 Writing Protocol-independent Software

In order to maximise portability of XTI applications between different kinds of machines and to support protocol independence, there are some general rules:

1. An application should only make use of those functions and mechanisms described as being mandatory features of XTI.
2. In the connection-mode service, the concept of a transport service data unit (TSDU) may not be supported by all transport providers. The user should make no assumptions about the preservation of logical data boundaries across a connection.
3. If an application is not intended to run only over an ISO transport provider, then the name of the device should not be hard-coded into it. While software may be written for a particular class of service (for example, connectionless-mode service), it should not be written to depend on any attribute of the underlying protocol.
4. The protocol-specific service limits returned on the *t_open()* and *t_getinfo()* functions must not be exceeded. It is the responsibility of the user to access these limits and then adhere to the limits throughout the communication process.
5. The user program should not look at or change options that are specific to the underlying protocol. The *t_optmgmt()* function enables a user to access default protocol options from the transport provider, which may then be blindly passed as an argument on the appropriate connection establishment function. Optionally, the user can choose not to pass options as an argument on connection establishment functions.
6. Protocol-specific addressing issues should be hidden from the user program. Similarly, the user must have some way of accessing destination addresses in an invisible manner, such as through a name server. However, the details for doing so are outside the scope of this interface specification.
7. The reason codes associated with *t_rcvdis()* are protocol-dependent. The user should not interpret this information if protocol independence is a concern.
8. The error codes associated with *t_rcvuderr()* are protocol-dependent. The user should not interpret this information if protocol independence is a concern.
9. The orderly release facility of the connection-mode service (that is, *t_sndrel()* and *t_rcvrel()*) should not be used by programs targeted for multiple protocol environments. This facility is not supported by all connection-based transport protocols. In particular, its use will prevent programs from successfully communicating with ISO open systems.
10. The semantics of expedited data are different across different transport providers (for example, ISO and TCP). An application intended to run over different transport providers should avoid their use.
11. The semantics of closing a connection may be different across transport providers. For example, closing a connection to ISO is abortive while closing a connection to TCP is orderly. A portable application should not assume either facility is available. If the service provider is of type *COTS_ORD*, a portable application should use *t_sndrel()* / *t_rcvrel()* prior to calling *t_close()*.

C.5 Event Management

In the absence of a standardised Event Management interface, the following guidelines are offered for the use of existing and widely available mechanisms by XTI applications.

These guidelines provide information additional to that given in Section 3.7 on page 14 and Section 3.9 on page 16.

For applications to use XTI in a fully asynchronous manner, they will need to use the facilities of an Event Management (EM) Interface. Such an EM will allow the application to be notified of a number of XTI events over a range of active endpoints. These events may be associated with:

- connection indication
- data indication
- disconnection indication
- flow control being lifted.

In the same way, the EM mechanism should allow the application to be notified of events coming from external sources, such as:

- asynchronous I/O completion
- expiration of timer
- resource availability.

When handling multiple transport connections, the application could either:

- fork a process for each new connection to be handled
- or:
- handle all connections within a single process by making use of the EM facilities.

The application will have to maintain an appropriate balance and choose the right trade-off between the number of processes and the number of connections managed per process in order to minimise the resulting overhead.

Unfortunately, the system facilities to suspend and await notification of an event are presently system-dependent, although work is in progress within standards bodies to provide a unified and portable mechanism.

Hence, for the foreseeable future, applications could use whatever underlying system facilities exist for event notification.

C.5.1 Short-term Solution

Many vendors currently provide either the System V *poll()* or BSD *select()* system calls which both give the ability to suspend until there is activity on a member of a set of file descriptors or a timeout.

Given the fact that a transport endpoint identifying a transport connection maps to a file descriptor, applications can take advantage of such EM mechanisms offered by the system (for example, *poll()* or *select()*). The design of more efficient and sophisticated applications, that make full use of all the XTI features, then becomes easily possible.

Guidelines for the use of *poll()* and *select()* are included in manual-page format, following the end of this section.

C.5.2 XTI Events

The XTI events can be divided into two classes of events.

- **Class 1:** events related to reception of data.

T_LISTEN	Connect request indication.
T_CONNECT	Connect response indication.
T_DATA	Reception of normal data indication.
T_EXDATA	Reception of expedited data indication.
T_DISCONNECT	Disconnection request indication.
T_ORDREL	Orderly release request indication.
T_UDERR	Notification of an error in a previously sent datagram.

This class of events should always be monitored by the application.

- **Class 2:** events related to emission of data (flow control).

T_GODATA	Normal data may be sent again.
T_GOEXDATA	Expedited data may be sent again.

This class of events informs the application that flow control restrictions have been lifted on a given file descriptor.

The application should request to be notified of this class of events whenever a flow control restriction has previously occurred on this endpoint (for example, [TFLOW] error has been returned on a `t_snd()` call).

Note that this class of event should not be monitored systematically otherwise the application would be notified each time a message is sent.

C.6 The Poll Function

Refer to the description of *poll()* in the referenced **XSH** specification. Moreover, Chapter 8 on page 129 of the current document gives additional information on the specific effect of *poll()* when applied to Sockets.

C.6.1 Example of Use of Poll

The following description gives the outline of an XTI server program making use of the *poll()* function.

```

/*
 * This is a simple server application example to show how poll() can
 * be used in a portable manner to wait for the occurrence of XTI events.
 * In this example, poll() is used to wait for the events T_LISTEN,
 * T_DISCONNECT, T_DATA and T_GODATA.
 *
 * A transport endpoint is opened in asynchronous mode over a
 * message-oriented transport provider (for example, ISO). The endpoint
 * is bound with qlen = 1 and the application enters an endless loop
 * to wait for all incoming XTI events on all its active endpoints.
 * For all connection indications received, a new endpoint is opened
 * with qlen = 0 and the connection request is accepted on that endpoint.
 * For all established connections, the application waits for data
 * to be received from one of its clients, sends the received data
 * back to the sender and waits for data again.
 * The cycle repeats until all the connections are released by
 * the clients. The disconnection indications are processed and the
 * endpoints closed.
 *
 * The example references two fictitious functions:
 *
 * - int get_provider(int tpid, char * tpname)
 *   Given a number as transport provider id, the function returns in
 *   tpname a string as transport provider name that can be used with
 *   t_open(). This function hides the different naming schemes of
 *   different XTI implementations.
 *
 * - int get_address(char * symb_name, struct netbuf address)
 *   Given a symbolic name symb_name and a pointer to a struct netbuf
 *   with allocated buffer space as input, the function returns a
 *   protocol address. This function hides the different addressing
 *   schemes of different XTI implementations.
 */

/*
 * General Includes
 */
#include <sys/types.h>
#include <fcntl.h>
#include <stdio.h>
#include <xti.h>

/*
 * Include files for poll()
 */
#include <poll.h>

/*
 * Various Defines
 */
/*
 * The XTI events T_CONNECT, T_DISCONNECT, T_LISTEN, T_ORDREL and T_UDERR
 * are related to one of the poll flags in INEVENTS (to which one, depends
 * on the implementation). POLLOUT means that (at least) normal data may
 * be sent, and POLLWRBAND that expedited data may be sent.
 */

```

```

#define ERREVENTS      (POLLERR | POLLHUP | POLLNVAL)
#define INEVENTS      (POLLIN | POLLRDNORM | POLLRDBAND | POLLPRI)
#define OUTEVENTS     (POLLOUT | POLLWRBAND)
#define MY_PROVIDER    1 /* transport provider id */
#define MAXSIZE       4000 /* size of send/receive buffer */
#define TPLEN         30 /* maximum length of provider name */
#define MAXCNX        10 /* maximum number of connections */

extern int      errno;

/*
 * Declaration of non-integer external functions
 */
void      exit();
void      perror();

/* ===== */
main()
{
    register int      i; /* loop variable */
    register int      num; /* return value of t_snd()
                           /* and t_rcv() */
    int               discflag = 0; /* flag to indicate a
                           /* disc indication */
    int               errflag = 0; /* flag to indicate an error */
    int               event; /* stores events returned
                           /* by t_look() */
    int               fd; /* current file descriptor */
    int               fdd; /* file descriptor
                           /* for t_accept() */
    int               flags; /* used with t_rcv() */
    char              *datbuf; /* current send/receive buffer */
    unsigned int      act = 0; /* active endpoints */
    struct t_info      info; /* used with t_open() */
    struct t_bind      *preq; /* used with t_bind() */
    struct t_call      *pcall; /* used with t_listen()
                           /* and t_accept() */
    struct t_discon     discon; /* used with t_rcvdis() */
    char              tpname[TPLEN]; /* transport provider name */
    char              buf[MAXCNX][MAXSIZE]; /* send/receive buffers */
    int               rcvdata[MAXCNX]; /* amount of data
                           /* already received */
    int               snddata[MAXCNX]; /* amount of data already sent */

    struct pollfd      fds[MAXCNX]; /* used with poll() */

    /*
     * Get name of transport provider
     */
    if (get_provider(MY_PROVIDER, tpname) == -1) {
        perror(">>> get_provider failed");
        exit(1);
    }

    /*

```



```

    * Establish a transport endpoint in asynchronous mode
    */
if ((fd = t_open(tpname, O_RDWR | O_NONBLOCK, &info)) == -1) {
    t_error(">>> t_open failed");
    exit(1);
}

/*
 * Allocate memory for the parameters passed with t_bind().
 */
if ((preq = (struct t_bind *) t_alloc(fd, T_BIND, T_ADDR)) == NULL) {
    t_error(">>> t_alloc(T_BIND) failed");
    t_close(fd);
    exit(1);
}

/*
 * Given a symbolic name ("MY_NAME"), get_address returns an address
 * and its length in preq->addr.buf and preq->addr.len.
 */
if (get_address("MY_NAME", &(preq->addr)) == -1) {
    perror(">>> get_address failed");
    t_close(fd);
    exit(1);
}
preq->qlen = 1;          /* is a listening endpoint */

/*
 * Bind the local protocol address to the transport endpoint.
 * The returned information is discarded.
 */
if (t_bind(fd, preq, NULL) == -1) {
    t_error(">>> t_bind failed");
    t_close(fd);
    exit(1);
}
if (t_free(preq, T_BIND) == -1) {
    t_error(">>> t_free failed");
    t_close(fd);
    exit(1);
}

/*
 * Allocate memory for the parameters used with t_listen.
 */
if ((pcall = (struct t_call *) t_alloc(fd, T_CALL, T_ALL)) == NULL) {
    t_error(">>> t_alloc(T_CALL) failed");
    t_close(fd);
    exit(1);
}

/*
 * Initialise entry 0 of the fds array to the listening endpoint.
 * To be portable across different XTI implementations,
 * register for INEVENTS and not for POLLIN.
 */
fds[act].fd = fd;
fds[act].events = INEVENTS;
fds[act].revents = 0;

```

```

rcvdata[act] = 0;
snddata[act] = 0;
act = 1;

/*
 * Enter an endless loop to wait for all incoming events.
 * Connect requests are accepted on new opened endpoints.
 * The example assumes that data is first sent by the client.
 * Then, the received data is sent back again and so on, until
 * the client disconnects.
 * Note that the total number of active endpoints (act) should
 * at least be 1, corresponding to the listening endpoint.
 */
fprintf(stderr, "Waiting for XTI events...\n");
while (act > 0) {
    /*
     * Wait for any events
     */
    if (poll(&fds, (size_t)act, (int) -1) == -1) {
        perror(">>> poll failed");
        exit(1);
    }
    /*
     * Process incoming events on all active endpoints
     */
    for (i = 0 ; i < act ; i++) {
        if (fds[i].revents == 0)
            continue;          /* no event for this endpoint */
        if (fds[i].revents & ERREVENTS) {
            fprintf(stderr, "[%d] Unexpected poll events: 0x%x\n",
                    fds[i].fd, fds[i].revents);
            continue;
        }
        /*
         * set the current endpoint
         * set the current send/receive buffer
         */
        fd = fds[i].fd;
        datbuf = buf[i];

        /*
         * Check for events
         */
        switch((event = t_look(fd))) {
        case T_LISTEN:
            /*
             * Must be a connection indication
             */
            if (t_listen(fd, pcall) == -1) {
                t_error(">>> t_listen failed");
                exit(1);
            }
            /*
             * If it will exceed the maximum number
             * of connections that the server can handle,
             * reject the connection indication.
             */
            if (act >= MAXCNX) {

```

```

        fprintf(stderr, ">>> Connection request rejected\n");
        if (t_snddis(fd, pcall) == -1)
            t_error(">>> t_snddis failed");
        continue;
    }
    /*
     * Establish a transport endpoint
     * in asynchronous mode
     */
    if ((fdd = t_open(tpname, O_RDWR | O_NONBLOCK, &info))
        == -1) {
        t_error(">>> t_open failed");
        continue;
    }
    /*
     * Accept connection on this endpoint.
     * fdd no longer needs to be bound,
     * t_accept() will do it.
     */
    if (t_accept(fd, fdd, pcall) == -1) {
        t_error(">>> t_accept failed");
        t_close(fdd);
        continue;
    }
    fprintf(stderr, "Connection [%d] opened\n", fdd);
    /*
     * Register for all flags that might indicate
     * a T_DATA or T_DISCONNECT event, i. e.,
     * register for INEVENTS (to be portable
     * through all XTI implementations).
     */
    fds[act].fd = fdd;
    fds[act].events = INEVENTS;
    fds[act].revents = 0;
    rcvdata[act] = 0;
    snddata[act] = 0;
    act++;
    break;

case T_DATA:
    /*
     * Must be a data indication
     */
    if ((num = t_rcv(fd, (datbuf + rcvdata[i]),
        (MAXSIZE - rcvdata[i]), &flags)) == -1) {
        switch (t_errno) {
        case TNODATA:
            /* No data is currently
             * available: repeat the loop
             */
            continue;
        case TLOOK:
            /* Must be a T_DISCONNECT event:
             * set discflag
             */
            event = t_look(fd);
            if (event == T_DISCONNECT) {
                discflag = 1;
            }
        }
    }

```

```

        break;
    }
    else
        fprintf(stderr, "Unexpected event %d\n",
                event);

default:
    /* Unexpected failure */
    t_error(">>> t_rcv failed");
    fprintf(stderr, "connection id: [%d]\n", fd);
    errflag = 1;
    break;
}
}

if (discflag || errflag)
    /* exit from the event switch */
    break;
fprintf(stderr, "[%d] %d bytes received\n", fd, num);
rcvdata[i] += num;
if (rcvdata[i] < MAXSIZE)
    continue;
if (flags & T_MORE) {
    fprintf(stderr, "[%d] TSDU too long for receive
                buffer\n", fd);
    errflag = 1;
    break; /* exit from the event switch */
}

/*
 * Send the data back:
 * Repeat t_snd() until either the whole TSDU
 * is sent back, or an event occurs.
 */
fprintf(stderr, "[%d] sending data back\n", fd);
do {
    if ((num = t_snd(fd, (datbuf + snddata[i]),
                    (MAXSIZE - snddata[i]), 0)) == -1) {
        switch (t_errno) {
            case TFLOW:
                /*
                 * Register for the flags
                 * OUTEVENTS to get awoken by
                 * T_GODATA, and for INEVENTS
                 * to get aware of T_DISCONNECT
                 * or T_DATA.
                 */
                fds[i].events |= OUTEVENTS;
                continue;

            case TLOOK:
                /*
                 * Must be a T_DISCONNECT event:
                 * set discflag
                 */
                event = t_lock(fd);
                if (event == T_DISCONNECT) {
                    discflag = 1;
                    break;
                }
        }
    }
}

```

```

        else
            fprintf(stderr, "Unexpected event %d\n",
                    event);

        default:
            t_error(">>> t_snd failed");
            fprintf(stderr, "connection id: [%d]\n", fd);
            errflag = 1;
            break;
    }
}
else {
    snddata[i] += num;
}
} while (MAXSIZE > snddata[i] && !discflag && !errflag);
/*
 * Reset send/receive counters
 */
rcvdata[i] = 0;
snddata[i] = 0;
break;

case T_GODATA:
/*
 * Flow control restriction has been lifted
 * restore initial event flags
 */
fds[i].events = INEVENTS;
continue;
case T_DISCONNECT:
/*
 * Must be a disconnection indication
 */
discflag = 1;
break;
case -1:
/*
 * Must be an error
 */
t_error(">>> t_look failed");
errflag = 1;
break;
default:
/*
 * Must be an unexpected event
 */
fprintf(stderr, "[%d] Unexpected event %d\n", fd, event);
errflag = 1;
break;
} /* end event switch */

if (discflag) {
/*
 * T_DISCONNECT has been received.
 * User data is not expected.
 */
if (t_rcvdis(fd, &discon) == -1)
    t_error(">>> t_rcvdis failed");
else

```

```

        fprintf(stderr, "[%d] Disconnection reason: 0x%x\n",
                fd, discon.reason);
    }
    if (discflag || errflag) {
        /*
         * Close transport endpoint and
         * decrement number of active connections
         */
        t_close(fd);
        act--;
        /* Move last entry of fds array to current slot,
         * adjust internal counters and flags
         */
        fds[i].events = fds[act].events;
        fds[i].revents = fds[act].revents;
        fds[i].fd = fds[act].fd;
        discflag = 0; /* clear disconnection flag */
        errflag = 0; /* clear error flag */
        i--; /* Redo the for() event loop to consider
             * events related to the last entry of
             * fds array */
        fprintf(stderr, "Connection [%d] closed\n", fd);
    }

    /* end of for() event loop */
} /* end of while() loop */
fprintf(stderr, ">>> Warning: no more active endpoints\n");
exit(1);
}

```

C.7 The Select Function

Refer to the description of `select()` in the referenced **XSH** specification. Moreover, Chapter 8 on page 129 of the current document gives additional information on the specific effect of `select()` when applied to Sockets.

C.7.1 Example of Use of Select

The following gives the outline of an XTI server program making use of `select()`.

The following describes the outline of an XTI server program making use of the `select()` function.

```

/*
 * This is a simple server application example to show how select() can
 * be used in a portable manner to wait for the occurrence of XTI events.
 * In this example, select() is used to wait for the events T_LISTEN,
 * T_DISCONNECT, T_DATA and T_GODATA.
 *
 * A transport endpoint is opened in asynchronous mode over a
 * message-oriented transport provider (for example, ISO). The endpoint is
 * bound with qlen = 1, and the application enters an endless loop to wait
 * for all incoming XTI events on all its active endpoints.
 * For all connection indications received, a new endpoint is opened with
 * qlen = 0 and the connection request is accepted on that endpoint.
 * For all established connections, the application waits for data to be
 * received from one of its clients, sends the received data back to the
 * sender and waits for data again.
 * The cycle repeats until all the connections are released by the clients.
 * The disconnection indications are processed and the endpoints closed.
 *
 * The example references two fictitious functions:
 *
 * - int get_provider(int tpid, char * tpname)
 *   Given a number as transport provider id, the function returns in
 *   tpname a string as transport provider name that can be used with
 *   t_open(). This function hides the different naming schemes of
 *   different XTI implementations.
 *
 * - int get_address(char * symb_name, struct netbuf address)
 *   Given a symbolic name symb_name and a pointer to a struct netbuf
 *   with allocated buffer space as input, the function returns a
 *   protocol address. This function hides the different addressing
 *   schemes of different XTI implementations.
 */
/*
 * General Includes
 */
#include <fcntl.h>
#include <stdio.h>
#include <xti.h>
/*
 * Include files for select().
 */
#include <sys/select.h>
#include <sys/time.h>

```

```

/*
 * Various Defines
 */

#define MY_PROVIDER      1 /* transport provider id */
#define MAXSIZE         4000 /* size of send/receive buffer */
#define TPLEN           30 /* maximum length of provider name */
#define MAXCNX          10 /* maximum number of connections */

extern int      errno;

/*
 * Declaration of non-integer external functions.
 */
void      exit();
void      perror();

/* ===== */

main()
{

    register int      i; /* loop variable */
    register int      num; /* return value of t_snd()
                           /* and t_rcv() */

    int               discflag = 0; /* flag to indicate a
                                   /* disc indication
    int               errflag = 0; /* flag to indicate an error
    int               event; /* stores events returned
                           /* by t_look()
    int               fd; /* current file descriptor
    int               fdd; /* file descriptor
                           /* for t_accept()
    int               flags; /* used with t_rcv()
    char              *datbuf; /* current send/receive
                           /* buffer
    size_t            act = 0; /* active endpoints
    struct t_info      info; /* used with t_open()
    struct t_bind      *preq; /* used with t_bind()
    struct t_call      *pcall; /* used with t_listen()
                           /* and t_accept()
    struct t_discon    discon; /* used with t_rcvdis()
    char              tpname[TPLEN]; /* transport provider name

    int               fds[MAXCNX]; /* array of file descriptors
    char              buf[MAXCNX][MAXSIZE] /* send/receive buffers
    int               rcvdata[MAXCNX]; /* amount of data
                           /* already received
    int               snddata[MAXCNX]; /* amount of data already sent

    fd_set            rfds, wfds, xfds; /* file descriptor sets
                                   /* for select()
    fd_set            rfd, wfd, xfd; /* initial values of

```



```

/* file descriptor sets */
/* rfd, wfd and xfd */

/*
 * Get name of transport provider
 */
if (get_provider(MY_PROVIDER, tpname) == -1) {
    perror(">>> get_provider failed");
    exit(1);
}

/*
 * Establish a transport endpoint in asynchronous mode
 */
if ((fd = t_open(tpname, O_RDWR | O_NONBLOCK, &info)) == -1) {
    t_error(">>> t_open failed");
    exit(1);
}

/*
 * Allocate memory for the parameters passed with t_bind().
 */
if ((preq = (struct t_bind *) t_alloc(fd, T_BIND, T_ADDR)) == NULL) {
    t_error(">>> t_alloc(T_BIND) failed");
    t_close(fd);
    exit(1);
}

/*
 * Given a symbolic name ("MY_NAME"), get_address returns an address
 * and its length in preq->addr.buf and preq->addr.len.
 */
if (get_address("MY_NAME", &(preq->addr)) == -1) {
    perror(">>> get_address failed");
    t_close(fd);
    exit(1);
}
preq->qlen = 1;          /* is a listening endpoint */

/*
 * Bind the local protocol address to the transport endpoint.
 * The returned information is discarded.
 */
if (t_bind(fd, preq, NULL) == -1) {
    t_error(">>> t_bind failed");
    t_close(fd);
    exit(1);
}
if (t_free(preq, T_BIND) == -1) {
    t_error(">>> t_free failed");
    t_close(fd);
    exit(1);
}

/*
 * Allocate memory for the parameters used with t_listen.
 */
if ((pcall = (struct t_call *) t_alloc(fd, T_CALL, T_ALL)) == NULL) {
    t_error(">>> t_alloc(T_CALL) failed");
}

```

```

        t_close(fd);
        exit(1);
    }

    /*
     * Initialise listening endpoint in descriptor set.
     * To be portable across different XTI implementations,
     * register for descriptor set rfdds and xfdds
     */
    FD_ZERO(&rfdds);
    FD_ZERO(&xfdds);
    FD_ZERO(&wfdds);
    FD_SET(fd, &rfdds);
    FD_SET(fd, &xfdds);
    fds[act] = fd;
    rcvdata[act] = 0;
    snddata[act] = 0;
    act = 1;

    /*
     * Enter an endless loop to wait for all incoming events.
     * Connect requests are accepted on a new opened endpoint.
     * The example assumes that data is first sent by the client.
     * Then, the received data is sent back again and so on, until
     * the client disconnects.
     * Note that the total number of active endpoints (act) should
     * at least be 1, corresponding to the listening endpoint.
     */
    fprintf(stderr, "Waiting for XTI events...\n");
    while (act > 0) {
        /*
         * Wait for any events
         */

        /*
         * Set the mask sets rfds, xfds and wfds to their initial values
         */
        rfds = rfdds;
        xfds = xfdds;
        wfds = wfdds;
        if (select(OPEN_MAX, &rfds, &wfds, &xfds,
            (struct timeval *) NULL) == -1) {
            perror(">>> select failed");
            exit(1);
        }
        /*
         * Process incoming events on all active endpoints
         */
        for (i = 0 ; i < act ; i++) {
            /*
             * set the current endpoint
             * set the current send/receive buffer
             */
            fd = fds[i];
            datbuf = buf[i];

            if (FD_ISSET(fd, &xfds)) {
                fprintf(stderr, "[%d] Unexpected select events\n", fd);
                continue;
            }
        }
    }
}

```

```

}
if (!FD_ISSET(fd, &rfd) && !FD_ISSET(fd, &wfd))
    continue; /* no event for this endpoint */

/*
 * Check for events
 */
switch((event = t_look(fd))) {
case T_LISTEN:
    /*
     * Must be a connection indication
     */
    if (t_listen(fd, pcall) == -1) {
        t_error(">>> t_listen failed");
        exit(1);
    }

    /*
     * If it will exceed the maximum number
     * of connections that the server can handle,
     * reject the connection indication.
     */
    if (act >= MAXCNX) {
        fprintf(stderr, ">>> Connection request
                        rejected\n");
        if (t_snddis(fd, pcall) == -1)
            t_error(">>> t_snddis failed");
        continue;
    }

    /*
     * Establish a transport endpoint
     * in asynchronous mode
     */
    if ((fdd = t_open(tpname, O_RDWR | O_NONBLOCK,
                    &info)) == -1) {
        t_error(">>> t_open failed");
        continue;
    }

    /*
     * Accept connection on this endpoint.
     * fdd no longer needs to be bound,
     * t_accept() will do it
     */
    if (t_accept(fd, fdd, pcall) == -1) {
        t_error(">>> t_accept failed");
        t_close(fdd);
        continue;
    }
    fprintf(stderr, "Connection [%d] opened\n", fdd);

    /*
     * Register for all flags that might indicate
     * a T_DATA or T_DISCONNECT event, i. e.,
     * register for rfd and xfd (to be portable
     * through all XTI implementations).
     */
    fds[act] = fdd;
    FD_SET(fdd, &rfd);
    FD_SET(fdd, &xfd);

```

```

rcvdata[act] = 0;
snddata[act] = 0;
act++;
break;

case T_DATA:
    /* Must be a data indication
    */
    if ((num = t_rcv(fd, (datbuf + rcvdata[i]),
        (MAXSIZE - rcvdata[i]), &flags)) == -1) {
        switch (t_errno) {
            case TNODATA:
                /* No data is currently
                * available: repeat the loop
                */
                continue;
            case TLOOK:
                /* Must be a T_DISCONNECT event:
                * set discflag
                */
                event = t_look(fd);
                if (event == T_DISCONNECT) {
                    discflag = 1;
                    break;
                }
                else
                    fprintf(stderr, "Unexpected event %d\n", event);

            default:
                /* Unexpected failure */
                t_error(">>> t_rcv failed");
                fprintf(stderr, "connection id: [%d]\n", fd);
                errflag = 1;
                break;
        }
    }

    if (discflag || errflag)
        /* exit from the event switch */
        break;
    fprintf(stderr, "[%d] %d bytes received\n", fd, num);
    rcvdata[i] += num;
    if (rcvdata[i] < MAXSIZE)
        continue;
    if (flags & T_MORE) {
        fprintf(stderr, "[%d] TSDU too long for receive
            buffer\n", fd);

        errflag = 1;
        break; /* exit from the event switch */
    }

    /*
    * Send the data back.
    * Repeat t_snd() until either the whole TSDU
    * is sent back, or an event occurs.
    */
    fprintf(stderr, "[%d] sending data back\n", fd);
    do {
        if ((num = t_snd(fd, (datbuf + snddata[i]),

```

```

(MAXSIZE - snddata[i]), 0)) == -1) {
switch (t_errno) {
case TFLOW:
/*
 * Register for wfds to get
 * awoken by T_GODATA, and for
 * rfds and xfds to get aware of
 * T_DISCONNECT or T_DATA.
 */
FD_SET(fd, &wfdds);
continue;

case TLOOK:
/*
 * Must be a T_DISCONNECT event:
 * set discflag
 */
event = t_look(fd);
if (event == T_DISCONNECT) {
    discflag = 1;
    break;
}
else
    fprintf(stderr, "Unexpected event
                %d\n", event);

default:
    t_error(">>> t_snd failed");
    fprintf(stderr, "connection id: [%d]\n", fd);
    errflag = 1;
    break;
}
}
else {
    snddata[i] += num;
}
} while (MAXSIZE > snddata[i] && !discflag && !errflag);
/*
 * Reset send/receive counter
 */
rcvdata[i] = 0;
snddata[i] = 0;
break;

case T_GODATA:
/*
 * Flow control restriction has been lifted
 * restore initial event flags
 */
FD_CLR(fd, &wfdds);
continue;

case T_DISCONNECT:
/*
 * Must be a disconnection indication
 */
discflag = 1;
break;

case -1:
/*

```

```

        * Must be an error
        */
        t_error(">>> t_look failed");
        errflag = 1;
        break;
default:
    /*
     * Must be an unexpected event
     */
    fprintf(stderr, "[%d] Unexpected event %d\n", fd, event);
    errflag = 1;
    break;
}    /* end event switch */

if (discflag) {
    /*
     * T_DISCONNECT has been received.
     * User data is not expected.
     */
    if (t_rcvdis(fd, &discon) == -1)
        t_error(">>> t_rcvdis failed");
    else
        fprintf(stderr, "[%d] Disconnection reason: 0x%x\n",
                fd, discon.reason);
}

if (discflag || errflag) {
    /*
     * Close transport endpoint and
     * decrement number of active connections
     */
    t_close(fd);
    act--;
    /*
     * Unregister fd from initial mask sets
     */
    FD_CLR(fd, &rfdds);
    FD_CLR(fd, &xfdds);
    FD_CLR(fd, &wfdds);
    /* Move last entry of fds array to current slot,
     * adjust internal counters and flags
     */
    fds[i] = fds[act];
    discflag = 0;    /* clear disconnection flag */
    errflag = 0;    /* clear error flag */
    i--;    /* Redo the for() event loop to consider
             * events related to the last entry of
             * fds array */
    fprintf(stderr, "Connection [%d] closed\n", fd);
}

}    /* end of for() event loop */

}    /* end of while() loop */
fprintf(stderr, ">>> Warning: no more active endpoints\n");
exit(1);
}

```


Use of XTI to Access NetBIOS

D.1 Introduction

NetBIOS represents an important *de facto* standard for networking DOS and OS/2 PCs. The X/Open Specification **Protocols for X/Open PC Interworking: SMB** (see the referenced **NetBIOS** specification) provides mappings of NetBIOS services to OSI and IPS transport protocols¹⁴.

The following CAE Specification extends that work to provide a standard programming interface to NetBIOS transport providers in X/Open-compliant systems, using an existing X/Open Common Applications Environment (CAE) interface, XTI.

The X/Open Transport Interface (XTI) defines a transport service interface that is independent of any specific transport provider.

This CAE Specification defines a standard for using XTI to access NetBIOS transport providers. Applications that use XTI to access NetBIOS transport providers are referred to as “transport users”.

This Appendix also defines data structures and constants required for NetBIOS transport providers which are exposed through `<xti_netbios.h>` header file. Applications written to compilation environments earlier than those required by this issue of the specification (see Appendix F on page 281) and defining `_XOPEN_SOURCE` to be less than 500, will have these data structures and constants exposed through the inclusion of `<xti.h>` on platforms where XTI interface use is supported for NetBIOS transports.

D.2 Objectives

The objectives of this standardisation are:

1. to facilitate the development and portability of CAE applications that interwork with the large installed base of NetBIOS applications in a Local Area Network (LAN) environment
2. to enable a single application to use the same XTI interface to communicate with remote applications through either an IPS profile, an OSI profile or a NetBIOS profile (that is, RFC 1001/1002 or TOP/NetBIOS)
3. to provide a common interface that can be used for IPC with clients using either (PC)NFS or SMB protocols for resources sharing.

This CAE Specification provides a migration step to users moving from proprietary systems in a NetBIOS environment to open systems, that is, the X/Open CAE.

14. The mappings are defined by the Specification of NetBIOS Interface and Name Service Support by Lower Layer OSI Protocols, and RFC 1001/RFC 1002 respectively. See the referenced **NetBIOS** specification. The relevant chapters are Chapter 13, NetBIOS Interface to ISO Transport Services, Chapter 14, Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Concepts and Methods and Chapter 15, Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Detailed Specification.

D.3 Scope

No extensions to XTI, as it is defined in the main body of this CAE Specification, are made in this NetBIOS CAE Specification. This NetBIOS CAE Specification is concerned only with standardisation of the mapping of XTI to the NetBIOS facilities, and not a new definition of XTI itself.

This CAE Specification applies only to the use of XTI in the single NetBIOS subnetwork case, and does not provide for the support of applications operating in multiple, non-overlapping NetBIOS name spaces.

The following NetBIOS facilities found in various NetBIOS implementations are considered outside the scope of XTI (note that this list is not necessarily definitive):

- LAN.STATUS.ALERT
- RESET
- SESSION STATUS
- TRACE
- UNLINK
- RPL (Remote Program Load)
- ADAPTER STATUS
- FIND NAME
- SEND.NOACK
- CHAIN.SEND.NOACK
- CANCEL
- receiving a datagram on any name
- receiving data on any connection.

It must also be noted that not all commands are specified in the protocols.

Omitting these does not restrict interoperability with the majority of NetBIOS implementations, since they have local significance only (RESET, SESSION STATUS), are concerned with systems management (UNLINK, RPL, ADAPTER STATUS), or are LAN- and vendor-specific (FIND NAME). If and how these functions are made available to the programmer is left to the implementor of this particular XTI implementation.

D.4 Issues

The primary issues for XTI as a transport interface to NetBIOS concern the passing of NetBIOS names and name type information through XTI, specification of restrictions on XTI functions in the NetBIOS environment, and handling the highly dynamic assignment of NetBIOS names.

D.5 NetBIOS Names and Addresses

NetBIOS uses 16-octet alphanumeric names as “transport” addresses. NetBIOS names must be exactly 16 octets, with shorter names padded with spaces to 16 octets. In addition, NetBIOS names are either unique names, group names or local names, and must be identified as such in certain circumstances. A local NetBIOS name is a name that is not depended on the network.

The following restrictions should be applied to NetBIOS names. Failure to observe these restrictions may result in unpredictable results.

1. Byte 0 of the name is not allowed to be hexadecimal 00 (0x00).
2. Byte 0 of the name is not allowed to be an asterisk, except as noted elsewhere in this specification to support broadcast datagrams.
3. Names should not begin with company names or trademarks.
4. Names should not begin with hexadecimal FF (0xFF).
5. Byte 15 of the name should not be in the range 0x00 – 0x1F.

The concept of a permanent node name, as provided in the native NetBIOS environment, is not supported in the X/Open CAE.

The following definitions are supplied with any implementation of XTI on top of NetBIOS. These definitions are exposed by the inclusion of `<xti_netbios.h>`.

```
#define T_NB_UNIQUE          0
#define NB_UNIQUE            0          (LEGACY)
#define T_NB_GROUP          1
#define NB_GROUP             1          (LEGACY)
#define T_NB_LOCAL          2
#define NB_LOCAL             2          (LEGACY)
#define T_NB_NAMELEN        16
#define NB_NAMELEN          16          (LEGACY)
#define T_NB_BCAST_NAME     "*"       /* asterisk plus 15 spaces */
#define NB_BCAST_NAME      "*"       (LEGACY)
```

The protocol addresses passed in calls to `t_bind()`, `t_connect()`, etc., are structured as follows:

```

1      2                                     17
+-----+-----+-----+-----+-----+
| Type |           NetBIOS Name           |
+-----+-----+-----+-----+-----+
```

Type The first octet specifies the type of the NetBIOS name. It may be set to `T_NB_UNIQUE`, `T_NB_GROUP` or `T_NB_LOCAL`.

NetBIOS Name Octets 2 through 17 contain the 16-octet NetBIOS name.

All NetBIOS names, complete with the name type identifier, are passed through XTI in a `netbuf` address structure (that is, `struct netbuf addr`), where `addr.buf` points to a NetBIOS protocol address as defined above. This applies to all XTI functions that pass or return a (NetBIOS) protocol address (for example, `t_bind()`, `t_connect()`, `t_rcvudata()`, etc.).

Note, however, that only the `t_bind()` and `t_getprotaddr()` functions use the name type information. All other functions ignore it.

If the NetBIOS protocol address is returned, the name type information is to be ignored since the NetBIOS transport providers do not provide the type information in the connection establishment phase.

NetBIOS names can become invalid even after they have been registered successfully due to the NetBIOS name conflict resolution process (for example, Top/NetBIOS NameConflictAdvise indication). For existing NetBIOS connections this has no effect since the connection endpoint can still be identified by the *fd*. However, in the connection establishment phase *2t_listen()* and *t_connect()* this event is indicated by setting *t_errno* to [TBADF].

D.6 NetBIOS Connection Release

Native NetBIOS implementations provide a linger-on-close release mechanism whereby a transport disconnection request (NetBIOS HANGUP) will not complete until all outstanding send commands have completed. NetBIOS attempts to deliver all queued data by delaying, if necessary, disconnection for a period of time. The period of time might be configurable; a value of 20 seconds is common practice. Data still queued after this time period may get discarded so that delivery cannot be guaranteed.

XTI, however, offers two different modes to release a connection: an abortive mode via `t_snddis()/t_rcvdis()`, and a graceful mode via `t_sndrel()/t_rcvrel()`. If a connection release is initiated by a `t_snddis()`, queued send data may be discarded. Only the use of `t_sndrel()` guarantees that the linger-on-close mechanism is enabled as described above. The support of `t_sndrel()/t_rcvrel()` is optional and only provided by implementations with servtype `T_COTS_ORD` (see `t_getinfo()` in Section D.8 on page 273).

A call to `t_sndrel()` initiates the linger-on-close mechanism and immediately returns with the XTI state changed to `T_OUTREL`. The NetBIOS provider sends all outstanding data followed by a NetBIOS Close Request. After receipt of a NetBIOS Close Response, the NetBIOS provider informs the transport user, via the event `T_ORDREL`, that is to be consumed by calling `t_rcvrel()`. If a timeout occurs, however, a `T_DISCONNECT` event with a corresponding reason code is generated.

Receive data arriving before the NetBIOS Close Request is sent is indicated by `T_DATA` and can be read by the transport user.

Calling `t_snddis()` initiates an abortive connection release and immediately returns with the XTI state changed to `T_IDLE`. Outstanding send and receive data may be discarded. The NetBIOS provider sends as many outstanding data as possible prior to closing the connection, but discards any receive data. Some outstanding data may be discarded by the `t_snddis()` mechanism, so that not all data can be sent by the NetBIOS provider. Furthermore, an occurring timeout condition could not be indicated to the transport user.

An incoming connection release will always result in a `T_DISCONNECT` event, never in a `T_ORDREL` event. To be precise, if the NetBIOS provider receives a Close Request, it discards any pending send and receive data, sends a Close Response and informs the transport user via `T_DISCONNECT`.

D.7 Options

No NetBIOS-specific options are defined. An implementation may, however, provide XTI-level options (see *t_optmgmt()* on page 84).

D.8 XTI Functions

- t_accept()*** No user data may be returned to the caller (*call*→*udata.len*=0).
 This function may only be used with connection-mode transport endpoints. The *t_accept()* function will fail if a user attempts to accept a connection request on a connectionless-mode endpoint and *t_errno* will be set to [TNOTSUPPORT].
- t_alloc()*** No special considerations for NetBIOS transport providers.
- t_bind()*** The NetBIOS name and name type values are passed to the transport provider in the *req* parameter (*req*→*addr.buf*) and the actual bound address is returned in the *ret* parameter (*ret*→*addr.buf*), as described earlier in Section D.5 on page 269. If the NetBIOS transport provider is unable to register the name specified in the *req* parameter, the call to *t_bind()* will fail with *t_errno* set to [TADDRBUSY] if the name is already in use, or to [TBADADDR] if it was an illegal NetBIOS name. If the NetBIOS name type is T_NB_LOCAL the name is not defended on the network, that is, it is not registered.
 If the *req* parameter is a null pointer or *req*→*addr.len*=0, the transport provider may assign an address for the user. This may be useful for outgoing connections on which the name of the caller is not important.
 If the name specified in *req* parameter is T_NB_BCAST_NAME, *qlen* must be zero, and the transport endpoint the name is bound to is enabled to receive broadcast datagrams. In this case, the transport endpoint must support connectionless-mode service, otherwise the *t_bind()* function will fail and *t_errno* will be set to [TBADADDR].
- t_close()*** No special considerations for NetBIOS transport providers.
 It is assumed that the NetBIOS transport provider will release the NetBIOS name associated with the closed endpoint if this is the only endpoint bound to this name and the name has not already been released as the result of a previous *t_unbind()* call on this endpoint.
- t_connect()*** The NetBIOS name of the destination transport user is provided in the *sndcall* parameter (*sndcall*→*addr.buf*), and the NetBIOS name of the responding transport user is returned in the *rcvcall* parameter (*rcvcall*→*addr.buf*), as described in Section D.5 on page 269. If the connection is successful, the NetBIOS name of the responding transport user will always be the same as that specified in the *sndcall* parameter.
 Local NetBIOS connections are supported. NetBIOS datagrams are sent, if applicable, to local names as well as remote names. No user data may be sent during connection establishment (*udata.len*=0 in *sndcall*).
 This function may only be used with connection-mode transport endpoints. The *t_connect()* function will fail if a user attempts to initiate a connection on a connectionless-mode endpoint and *t_errno* will be set to [TNOTSUPPORT].
 [TBADF] may be returned in the case that the NetBIOS name associated with the *fd* referenced in the *t_connect()* call is no longer in the CAE system name table (see Section D.5 on page 269).
- t_error()*** No special considerations for NetBIOS transport providers.

<code>t_free()</code>	No special considerations for NetBIOS transport providers.
<code>t_getinfo()</code>	The values of the parameters in the <code>t_info</code> structure will reflect NetBIOS transport limitations, as follows: <ul style="list-style-type: none"> <code>addr</code> <code>sizeof()</code> the NetBIOS protocol address, as defined in Section D.5 on page 269. <code>options</code> Equals <code>-2</code>, indicating no user-settable options. <code>tsdu</code> Equals the size returned by the transport provider. If the <code>fd</code> is associated with a connection-mode endpoint it is a positive value, not larger than 131070. If the <code>fd</code> is associated with a connectionless-mode endpoint it is a positive value not larger than 65535¹⁵. <code>etsdu</code> Equals <code>-2</code>, indicating expedited data is not supported. <code>connect</code> Equals <code>-2</code>, indicating data cannot be transferred during connection establishment. <code>discon</code> Equals <code>-2</code>, indicating data cannot be transferred during connection release. <code>servtype</code> Set to <code>T_COTS</code> if the <code>fd</code> is associated with a connection-mode endpoint, or <code>T_CLTS</code> if associated with a connectionless-mode endpoint. Optionally, may be set to <code>T_COTS_ORD</code> if the <code>fd</code> is associated with a connection-mode endpoint and the transport provider supports the use of <code>t_sndrel()/t_rcvrel()</code> as described in Section D.6 on page 271. <code>flags</code> Equals <code>T_SNDZERO</code>, indicating that zero TSDUs may be sent.
<code>t_getprotaddr()</code>	The NetBIOS name and name type of the transport endpoint referred to by the <code>fd</code> are passed in the <code>boundaddr</code> parameter (<code>boundaddr→addr.buf</code>), as described in Section D.5 on page 269; 0 is returned in <code>boundaddr→addr.len</code> if the transport endpoint is in the <code>T_UNBND</code> state. The NetBIOS name currently connected to <code>fd</code> , if any, is passed in the <code>peeraddr</code> parameter (<code>peeraddr→addr.buf</code>); the value 0 is returned in <code>peeraddr→addr.len</code> if the transport endpoint is not in the <code>T_DATAXFER</code> state.
<code>t_getstate()</code>	No special considerations for NetBIOS transport providers.
<code>t_listen()</code>	On return, the <code>call</code> parameter provides the NetBIOS name of the calling transport user (that issued the connection request), as described in Section D.5 on page 269. No user data may be transferred during connection establishment (<code>call→udata.len=0</code> on return). This function may only be used with connection-mode transport endpoints. The <code>t_listen()</code> function will fail if a user attempts to <code>listen</code> on a connectionless-mode endpoint and <code>t_errno</code> will be set to <code>[TNOTSUPPORT]</code> . <code>[TBADF]</code> may be returned in the case that the NetBIOS name associated with

15. For the mappings to OSI and IPS protocols, the value cannot exceed 512 or 1064 respectively.

- the *fd* referenced in the *t_listen()* function is no longer in the CAE system name table, as may occur as a result of the NetBIOS name conflict resolution process (for example, TOP/NetBIOS NameConflictAdvise indication).
- t_look()* Since expedited data is not supported in NetBIOS, the T_EXDATA and T_GOEXDATA events cannot be returned.
- t_open()* No special considerations for NetBIOS transport providers, other than restrictions on the values returned in the *t_info* structure. These restrictions are described in *t_getinfo()* on page 71.
- t_optmgmt()* No special considerations for NetBIOS transport providers.
- t_rcv()* This function may only be used with connection-mode transport endpoints. The *t_rcv()* function will fail if a user attempts a receive on a connectionless-mode endpoint and *t_errno* will be set to [TNOTSUPPORT].
- The *flags* parameter will never be set to T_EXPEDITED, as expedited data is not supported.
- Data transfer in the NetBIOS environment is record-oriented, and the transport user should expect to see usage of the *T_MORE* flag when the message size exceeds the available buffer size.
- t_rcvconnect()* The NetBIOS name of the transport user responding to the previous connection request is provided in the *call* parameter (*call*→*addr.buf*), as described in Section D.5 on page 269.
- No user data may be returned to the caller (*call*→*udata.len*=0 on return).
- This function may only be used with connection-mode transport endpoints. The *t_rcvconnect()* function will fail if a user attempts to establish a connection on a connectionless-mode endpoint and *t_errno* will be set to [TNOTSUPPORT].
- t_rcvdis()* The following disconnection reason codes are valid for any implementation of a NetBIOS provider under XTI:
- ```

#define T_NB_ABORT 0x18 /* session ended abnormally */
#define NB_ABORT 0x18 /* (LEGACY) */
#define T_NB_CLOSED 0x0A /* session closed */
#define NB_CLOSED 0x0A /* (LEGACY) */
#define T_NB_NOANSWER 0x14 /* no answer (cannot find */
 /* name called */
#define NB_NOANSWER 0x14 /* (LEGACY) */
#define T_NB_OPREJ 0x12 /* session open rejected */
#define NB_OPREJ 0x12 /* (LEGACY) */

```
- These definitions are exposed by the inclusion of `<xti_netbios.h>`.
- t\_rcvrel()* As described in Section D.6 on page 271, a T\_ORDREL event will never occur in the T\_DATAXFER state, but only in the T\_OUTREL state. A transport user thus has only to prepare for a call to *t\_rcvrel()* if it previously initiated a connection release by calling *t\_sndrel()*. As a side effect, the state T\_INREL is unreachable for the transport user.
- If T\_COTS\_ORD is not supported by the underlying NetBIOS transport provider, this function will fail with *t\_errno* set to [TNOTSUPPORT].
- t\_rcvudata()* The NetBIOS name of the sending transport user is provided in the *unitdata* parameter (*unitdata*→*addr.buf*), as described in Section D.5 on page 269.

The *fd* associated with the *t\_rcvudata()* function must refer to a connectionless-mode transport endpoint. The function will fail if a user attempts to receive on a connection-mode endpoint and *t\_errno* will be set to [TNOTSUPPORT]. [TBADF] may be returned in the case that the NetBIOS name associated with the *fd* referenced in the *t\_rcvudata()* function is no longer in the CAE system name table, as may occur as a result of the NetBIOS name conflict resolution process (for example, TOP/NetBIOS NameConflictAdvise indication).

To receive a broadcast datagram, the endpoint must be bound to the NetBIOS name T\_NB\_BCAST\_NAME.

*t\_rcvuderr()* If attempted on a connectionless-mode transport endpoint, this function will fail with *t\_errno* set to [TNOUDERR], as no NetBIOS unit data error codes are defined. If attempted on a connection-mode transport endpoint, this function will fail with *t\_errno* set to [TNOTSUPPORT].

*t\_snd()* The T\_EXPEDITED flag may not be set, as NetBIOS does not support expedited data transfer.

This function may only be used with connection-mode transport endpoints. The *t\_snd()* function will fail if a user attempts a send on a connectionless-mode endpoint and *t\_errno* will be set to [TNOTSUPPORT].

The maximum value of the *nbytes* parameter is determined by the maximum TSDU size allowed by the transport provider. The maximum TSDU size can be obtained from the *t\_getinfo()* call.

Data transfer in the NetBIOS environment is record-oriented. The transport user can use the *T\_MORE* flag in order to fragment a TSDU and send it via multiple calls to *t\_snd()*. See *t\_snd()* on page 108 for more details.

NetBIOS does not support the notion of expedited data. A call to *t\_snd()* with the *T\_EXPEDITED* flag will fail with *t\_errno* set to [TBADDDATA].

If the NetBIOS provider has received a HANGUP request from the remote user and still has receive data to deliver to the local user, XTI may not detect the HANGUP situation during a call to *t\_snd()*. The actions that are taken are implementation-dependent:

- *t\_snd()* might fail with *t\_errno* set to [TPROTO]
- *t\_snd()* might succeed, although the data is discarded by the transport provider, and an implementation-dependent error (generated by the NetBIOS provider) will result on a subsequent XTI call. This could be a [TSYSERR], a [TPROTO] or a connection release indication after all the receive data has been delivered.

*t\_snddis()* The *t\_snddis()* function initiates an abortive connection release. The function returns immediately. Outstanding send and receive data may be discarded. See Section D.6 on page 271 for further details.

No user data may be sent in the disconnection request (call→*u*data.len=0).

This function may only be used with connection-mode transport endpoints. The *t\_snddis()* function will fail if a user attempts a disconnection request on a connectionless-mode endpoint and *t\_errno* will be set to [TNOTSUPPORT].

- t\_sndrel()* The *t\_sndrel()* function initiates the NetBIOS release mechanism that attempts to complete outstanding sends within a timeout period before the connection is released. The function returns immediately. The transport user is informed by T\_ORDREL when all sends have been completed and the connection has been closed successfully. If, however, the timeout occurs, the transport user is informed by a T\_DISCONNECT event with an appropriate disconnection reason code. See Section D.6 on page 271 for further details.
- If the NetBIOS transport provider did not return T\_COTS\_ORD with *t\_open()*, this function will fail with *t\_errno* set to [TNOTSUPPORT].
- t\_sndudata()* The NetBIOS name of the destination transport user is provided in the *unitdata* parameter (*unitdata*→*addr.buf*), as described in Section D.5 on page 269.
- The *fd* associated with the *t\_sndudata()* function must refer to a connectionless-mode transport endpoint. The function will fail if a user attempts this function on a connection-mode endpoint and *t\_errno* will be set to [TNOTSUPPORT]. [TBADF] may be returned in the case that the NetBIOS name associated with the *fd* referenced in the *t\_sndudata()* function is no longer in the CAE system name table, as may occur as a result of the NetBIOS name conflict resolution process (for example, TOP/NetBIOS NameConflictAdvise indication).
- To send a broadcast datagram, the NetBIOS name in the NetBIOS address structure provided in *unitdata*→*addr.buf* must be T\_NB\_BCAST\_NAME.
- t\_strerror()* No special considerations for NetBIOS transport providers.
- t\_sync()* No special considerations for NetBIOS transport providers.
- t\_unbind()* No special considerations for NetBIOS transport providers.
- It is assumed that the NetBIOS transport provider will release the NetBIOS name associated with the endpoint if this is the only endpoint bound to this name.



## *XTI and TLI*

XTI is based on the SVID Issue 2, Volume III, Networking Services Extensions (see Referenced Documents).

XTI provides refinement of the Transport Level Interface (TLI) where such refinement is considered necessary. This refinement takes the form of:

- additional commentary or explanatory text, in cases where the TLI text is either ambiguous or not sufficiently detailed
- modifications to the interface, to cater for service and protocol problems which have been fully considered. In this case, it must be emphasised that such modifications are kept to an absolute minimum, and are intended to avoid any fundamental changes to the interface defined by TLI
- the removal of dependencies on specific UNIX versions and specific transport providers.

### **E.1 Restrictions Concerning the Use of XTI**

It is important to bear in mind the following points when considering the use of XTI:

- It was stated that XTI “recommends” a subset of the total set of functions and facilities defined in TLI, and also that XTI introduces modifications to some of these functions and/or facilities where this is considered essential. For these reasons, an application which is written in conformance to XTI may not be immediately portable to work over a provider which has been written in conformance to TLI.
- XTI does not address management aspects of the interface, that is:
  - how addressing may be done in such a way that an application is truly portable
  - no selection and/or negotiation of service and protocol characteristics.

For addressing, the same is also true for TLI. In this case, it is envisaged that addresses will be managed by a higher-level directory function. For options selection and/or negotiation, XTI attempts to define a basic mechanism by which such information may be passed across the transport service interface, although again, this selection/negotiation may be done by a higher-level management function (rather than directly by the user). Since address structure is not currently defined, the user protocol address is system-dependent.

## E.2 Relationship between XTI and TLI

The following features can be considered as XTI extensions to the System V Release 3 version of TLI:

- Some functions may return more error types. The use of the [TOUTSTATE] error is generalised to almost all protocol functions.
- The transport provider identifier has been generalised to remove the dependence on a device driver implementation.
- Additional events have been defined to help applications make full use of the asynchronous features of the interface.
- Additional features have been introduced to *t\_snd()*, *t\_sndrel()* and *t\_rcvrel()* to allow fuller use of TCP transport providers.
- Usage of options for certain types of transport service has been defined to increase application portability.
- Because most XTI functions require read/write access to the transport provider, the usage of flags O\_RDONLY and O\_WRONLY has been withdrawn from the XTI.
- XTI checks the value of *qlen* and prevents an application from waiting forever when issuing *t\_listen()*.
- XTI allows an application to call *t\_accept()* with a *resfd* which is not bound to a local address.
- XTI provides the additional utility functions *t\_strerror()* and *t\_getprotaddr()*.

## Headers and Definitions for XTI

Section 7.1 on page 51 contains a normative requirement that the contents and structures found in this appendix appear in the `<xti.h>` header.

Applications written to compilation environments earlier than those required by this issue of the specification (see Section 1.3 on page 3) and defining `_XOPEN_SOURCE` to be less than 500, will have these data structures and constants of certain protocol specific headers automatically exposed by the inclusion of `<xti.h>` for compatibility. The individual protocol-specific appendices document the providers for which this may be the case.

The function definitions in this appendix conform to the ISO C standard (see **referenced documents**).

Values specified for some of the symbolic constants in this XTI header definition are designated as **not mandatory** for conformance purposes. These are identified by the comment:

```
/* value is recommended-only, not mandatory */
```

in the following definitions.

`<xti.h>` defines `t_scalar_t` and `t_uscalar_t` respectively as a signed and unsigned opaque integral type of equal length of at least 32 bits.

Certain items in this appendix are marked **LEGACY**. This which means that they are being retained for compatibility with older applications, but have limitations which make them inappropriate for developing portable applications. New applications should use alternative means of obtaining equivalent functionality.

```
/*
 * The following are the error codes needed by both the kernel
 * level transport providers and the user level library.
 */

#define TBADADDR 1 /* incorrect addr format */
#define TBADOPT 2 /* incorrect option format */
#define TACCES 3 /* incorrect permissions */
#define TBADF 4 /* illegal transport fd */
#define TNOADDR 5 /* couldn't allocate addr */
#define TOUTSTATE 6 /* out of state */
#define TBADSEQ 7 /* bad call sequence number */
#define TSYSERR 8 /* system error */
#define TLOOK 9 /* event requires attention */
#define TBADDATA 10 /* illegal amount of data */
#define TBUFOVFLW 11 /* buffer not large enough */
#define TFLOW 12 /* flow control */
#define TNODATA 13 /* no data */
#define TNODIS 14 /* discon_ind not found on queue */
#define TNOUDERR 15 /* unitdata error not found */
#define TBADFLAG 16 /* bad flags */
#define TNOREL 17 /* no ord rel found on queue */
#define TNOTSUPPORT 18 /* primitive/action not supported */
```

```

#define TSTATECHNG 19 /* state is in process of changing */
#define TNOSTRUCTYPE 20 /* unsupported struct-type requested */
#define TBADNAME 21 /* invalid transport provider name */
#define TBADQLEN 22 /* qlen is zero */
#define TADDRBUSY 23 /* address in use */
#define TINDOUT 24 /* outstanding connection indications */
#define TPROVMISMATCH 25 /* transport provider mismatch */
#define TRESQLEN 26 /* resfd specified to accept w/qlen >0 */
#define TRESADDR 27 /* resfd not bound to same addr as fd */
#define TQFULL 28 /* incoming connection queue full */
#define TPROTO 29 /* XTI protocol error */

/*
 * The following are the events returned.
 */

#define T_LISTEN 0x0001 /* connection indication received */
#define T_CONNECT 0x0002 /* connection confirmation received */
#define T_DATA 0x0004 /* normal data received */
#define T_EXDATA 0x0008 /* expedited data received */
#define T_DISCONNECT 0x0010 /* disconnection received */
#define T_UDERR 0x0040 /* datagram error indication */
#define T_ORDREL 0x0080 /* orderly release indication */
#define T_GODATA 0x0100 /* sending normal data is again possible */
#define T_GOEXDATA 0x0200 /* sending expedited data is again
 /* possible */

/*
 * The following are the flag definitions needed by the
 * user level library routines.
 */

#define T_MORE 0x001 /* more data */
#define T_EXPEDITED 0x002 /* expedited data */
#define T_PUSH 0x004 /* send data immediately */
#define T_NEGOTIATE 0x004 /* set opts */
#define T_CHECK 0x008 /* check opts */
#define T_DEFAULT 0x010 /* get default opts */
#define T_SUCCESS 0x020 /* successful */
#define T_FAILURE 0x040 /* failure */
#define T_CURRENT 0x080 /* get current options */
#define T_PARTSUCCESS 0x100 /* partial success */
#define T_READONLY 0x200 /* read-only */
#define T_NOTSUPPORT 0x400 /* not supported */

/*
 * XTI error return.
 */

/* t_errno is a modifiable lvalue of type int */
/* In a single threaded environment a typical definition of t_errno is: */

extern int t_errno;

/* In a multi-threading environment a typical definition of t_errno is: */

extern int *_t_errno(void);
#define t_errno (*(_t_errno()))

```



## Headers and Definitions for XTI

```
/*
 * iov maximum
 */
#define T_IOV_MAX 16 /* maximum number of scatter/gather buffers */
 /* value is not mandatory but if present */
 /* must be at least 16. */

struct t_iovec {
 void *iov_base;
 size_t iov_len;
};

/*
 * XTI LIBRARY FUNCTIONS
 */

/* XTI Library Function: t_accept - accept a connection request*/
extern int t_accept(int, int, const struct t_call *);
/* XTI Library Function: t_alloc - allocate a library structure*/
extern void *t_alloc(int, int, int);
/* XTI Library Function: t_bind - bind an address to a transport endpoint*/
extern int t_bind(int, const struct t_bind *, struct t_bind *);
/* XTI Library Function: t_close - close a transport endpoint*/
extern int t_close(int);
/* XTI Library Function: t_connect - establish a connection */
extern int t_connect(int, const struct t_call *, struct t_call *);
/* XTI Library Function: t_error - produce error message*/
extern int t_error(const char *);
/* XTI Library Function: t_free - free a library structure*/
extern int t_free(void *, int);
/* XTI Library Function: t_getinfo - get protocol-specific service */
/* information*/
extern int t_getinfo(int, struct t_info *);
/* XTI Library Function: t_getprotaddr - get protocol addresses*/
extern int t_getprotaddr(int, struct t_bind *, struct t_bind *);
/* XTI Library Function: t_getstate - get the current state*/
extern int t_getstate(int);
/* XTI Library Function: t_listen - listen for a connection indication*/
extern int t_listen(int, struct t_call *);
/* XTI Library Function: t_look - look at current event on a transport */
/* endpoint*/
extern int t_look(int);
/* XTI Library Function: t_open - establish a transport endpoint*/
extern int t_open(const char *, int, struct t_info *);
/* XTI Library Function: t_optmgmt - manage options for a transport */
/* endpoint*/
extern int t_optmgmt(int, const struct t_optmgmt *,
 struct t_optmgmt *);
/* XTI Library Function: t_rcv - receive data or expedited data on a */
/* connection*/
extern int t_rcv(int, void *, unsigned int, int *);
/* XTI Library Function: t_rcvconnect - receive the confirmation from */
/* a connection request */
extern int t_rcvconnect(int, struct t_call *);
/* XTI Library Function: t_rcvdis - retrieve information from disconnect*/
extern int t_rcvdis(int, struct t_discon *);
/* XTI Library Function: t_rcvrel - acknowledge receipt of */
```

```

/* an orderly release indication */
extern int t_rcvrel(int);
/* XTI Library Function: t_rcvreldata - receive an orderly release */
/* indication or confirmation containing user data */
extern int t_rcvreldata(int, struct t_discon *discon)
/* XTI Library Function: t_rcvudata - receive a data unit*/
extern int t_rcvudata(int, struct t_unitdata *, int *);
/* XTI Library Function: t_rcvuderr - receive a unit data error indication*/
extern int t_rcvuderr(int, struct t_uderr *);
/* XTI Library Function: t_rcvv - receive data or expedited data sent*/
/* over a connection and put the data */
/* into one or more noncontiguous buffers*/
extern int t_rcvv(int, struct t_iovec *, unsigned int, int *);
/* XTI Library Function: t_rcvvudata - receive a data unit into one */
/* or more noncontiguous buffers*/
extern int t_rcvvudata(int, struct t_unitdata *, struct t_iovec *, \
 unsigned int, int *);
/* XTI Library Function: t_snd - send data or expedited data over a */
/* connection */
extern int t_snd(int, void *, unsigned int, int);
/* XTI Library Function: t_snddis - send user-initiated disconnect request*/
extern int t_snddis(int, const struct t_call *);
/* XTI Library Function: t_sndrel - initiate an orderly release*/
extern int t_sndrel(int);
/* XTI Library Function: t_sndreldata - initiate or respond to an */
/* orderly release with user data */
extern int t_sndreldata(int, struct t_discon *);
/* XTI Library Function: t_sndudata - send a data unit*/
extern int t_sndudata(int, const struct t_unitdata *);
/* XTI Library Function: t_sndv - send data or expedited data, */
/* from one or more noncontiguous buffers, on a connection*/
extern int t_sndv(int, const struct t_iovec *, unsigned int iovcount, int);
/* XTI Library Function: t_sndvudata - send a data unit from one or */
/* more non-contiguous buffers*/
extern int t_sndvudata(int, struct t_unitdata *, struct t_iovec *, unsigned int);
/* XTI Library Function: t_strerror - generate error message string */
extern const char *t_strerror(int);
/* XTI Library Function: t_sync - synchronise transport library*/
extern int t_sync(int);
/* XTI Library Function: t_sysconf - get configurable XTI variables */
extern int t_sysconf(int);
/* XTI Library Function: t_unbind - disable a transport endpoint*/
extern int t_unbind(int);

/*
 * Protocol-specific service limits.
 */
struct t_info {
 t_scalar_t addr; /*max size of the transport protocol address */
 t_scalar_t options; /*max number of bytes of protocol-specific options */
 t_scalar_t tsdu; /*max size of a transport service data unit */
 t_scalar_t etsdu; /*max size of expedited transport service data unit */
 t_scalar_t connect; /*max amount of data allowed on connection */
 /*establishment functions */
 t_scalar_t discon; /*max data allowed on t_snddis and t_rcvdis functions*/
 t_scalar_t servtype; /*service type supported by transport provider */
 t_scalar_t flags; /*other info about the transport provider */
};

```

## Headers and Definitions for XTI

```
/*
 * Service type defines.
 */

#define T_COTS 01 /* connection-mode transport service */
#define T_COTS_ORD 02 /* connection-mode with orderly release */
#define T_CLTS 03 /* connectionless-mode transport service */

/*
 * Flags defines (other info about the transport provider).
 */

#define T_SENDFZERO 0x001 /* supports 0-length TSDUs */
#define T_ORDRELDATA 0x002 /* supports orderly release data */

/*
 * netbuf structure.
 */

struct netbuf {
 unsigned int maxlen;
 unsigned int len;
 void *buf;
};

/*
 * t_opthdr structure
 */
struct t_opthdr {
 t_uscalar_t len; /* total length of option; that is, */
 /* sizeof (struct t_opthdr) + length */
 /* of option value in bytes */
 t_uscalar_t level; /* protocol affected */
 t_uscalar_t name; /* option name */
 t_uscalar_t status; /* status value */
/* followed by the option value */
};

/*
 * t_bind - format of the address arguments of bind.
 */

struct t_bind {
 struct netbuf addr;
 unsigned int qlen;
};

/*
 * Options management structure.
 */

struct t_optmgmt {
 struct netbuf opt;
 t_scalar_t flags;
};

/*
 * Disconnection structure.
 */
```

```

struct t_discon {
 struct netbuf udata; /* user data */
 int reason; /* reason code */
 int sequence; /* sequence number */
};

/*
 * Call structure.
 */

struct t_call {
 struct netbuf addr; /* address */
 struct netbuf opt; /* options */
 struct netbuf udata; /* user data */
 int sequence; /* sequence number */
};

/*
 * Datagram structure.
 */

struct t_unitdata {
 struct netbuf addr; /* address */
 struct netbuf opt; /* options */
 struct netbuf udata; /* user data */
};

/*
 * Unitdata error structure.
 */

struct t_uderr {
 struct netbuf addr; /* address */
 struct netbuf opt; /* options */
 t_scalar_t error; /* error code */
};

/*
 * The following are structure types used when dynamically
 * allocating the above structures via t_alloc().
 */

#define T_BIND 1 /* struct t_bind */
#define T_OPTMGMT 2 /* struct t_optmgmt */
#define T_CALL 3 /* struct t_call */
#define T_DIS 4 /* struct t_discon */
#define T_UNITDATA 5 /* struct t_unitdata */
#define T_UDERROR 6 /* struct t_uderr */
#define T_INFO 7 /* struct t_info */

/*
 * The following bits specify which fields of the above
 * structures should be allocated by t_alloc().
 */

```

## Headers and Definitions for XTI

```
#define T_ADDR 0x01 /* address */
#define T_OPT 0x02 /* options */
#define T_UDATA 0x04 /* user data */
#define T_ALL 0xffff /* all the above fields supported */

/*
 * The following are the states for the user.
 */

#define T_UNBND 1 /* unbound */
#define T_IDLE 2 /* idle */
#define T_OUTCON 3 /* outgoing connection pending */
#define T_INCON 4 /* incoming connection pending */
#define T_DATAXFER 5 /* data transfer */
#define T_OUTREL 6 /* outgoing release pending */
#define T_INREL 7 /* incoming release pending */

/*
 * General purpose defines.
 */

#define T_YES 1
#define T_NO 0
#define T_NULL 0
#define T_ABSREQ 0x8000
#define T_INFINITE (-1)
#define T_INVALID (-2)

/*
 * Definitions for t_sysconf
 */
#define _SC_T_IOV_MAX 1 /* value is recommended only, not mandatory */

/*
 * General definitions for option management
 */
#define T_UNSPEC (~0 - 2) /* applicable to u_long, t_scalar_t, char .. */
#define T_ALLOPT 0 /* value is recommended-only, not mandatory */
#define T_ALIGN(p) (((t_uscalar_t)(p) + (sizeof (t_scalar_t) - 1)) \
 & ~(sizeof (t_scalar_t) - 1))
#define T_OPT_DATA(tohp) /* definition to be added */

#define T_OPT_NEXTHDR(pbuf, buflen, popt) \
 (((char *) (popt) + T_ALIGN((popt)->len) < \
 (char *) (pbuf) + (buflen)) ? \
 (struct t_opthdr *) ((char *) (popt) + T_ALIGN((popt)->len)) : \
 (struct t_opthdr *) 0)

#define T_OPT_FIRSTHDR(nbp) /* definition to be added */

/* OPTIONS ON XTI LEVEL */

/*
 * XTI Level
 *
 * The values defined for the XTI Level are recommended-only,
 * not mandatory.
 */
```

```

*/

#define XTI_GENERIC 0xffff

/*
 * XTI-level Options
 *
 * The values defined for these XTI-level Options are recommended-only,
 * not mandatory.
 */

#define XTI_DEBUG 0x0001 /* enable debugging */
#define XTI_LINGER 0x0080 /* linger on close if data present */
#define XTI_RCVBUF 0x1002 /* receive buffer size */
#define XTI_RCVLOWAT 0x1004 /* receive low-water mark */
#define XTI_SNDBUF 0x1001 /* send buffer size */
#define XTI_SNDLOWAT 0x1003 /* send low-water mark */

/*
 * Structure used with linger option.
 */
struct t_linger {
 t_scalar_t l_onoff; /* option on/off */
 t_scalar_t l_linger; /* linger time */
};

/* SPECIFIC ISO OPTION AND MANAGEMENT PARAMETERS */

/*
 * Definition of the ISO transport classes
 */

#define T_CLASS0 0
#define T_CLASS1 1
#define T_CLASS2 2
#define T_CLASS3 3
#define T_CLASS4 4

/*
 * Definition of the priorities.
 */

#define T_PRITOP 0
#define T_PRIHIGH 1
#define T_PRIMID 2
#define T_PRILOW 3
#define T_PRIDFLT 4

/*
 * Definitions of the protection levels
 */

#define T_NOPROTECT 1

```

## Headers and Definitions for XTI

```
#define T_PASSIVEPROTECT 2
#define T_ACTIVEPROTECT 4

/*
 * Default value for the length of TPDU's.
 */

#define T_LTPDUDFLT 128 /* define obsolete in XPG4 */

/*
 * rate structure.
 */
struct rate {
 t_scalar_t targetvalue; /* target value */
 t_scalar_t minacceptvalue; /* value of minimum acceptable quality */
};

/*
 * reqvalue structure.
 */
struct reqvalue {
 struct rate called; /* called rate */
 struct rate calling; /* calling rate */
};

/*
 * thrpt structure.
 */
struct thrpt {
 struct reqvalue maxthrpt; /* maximum throughput */
 struct reqvalue avgthrpt; /* average throughput */
};

/*
 * transdel structure
 */
struct transdel {
 struct reqvalue maxdel; /* maximum transit delay */
 struct reqvalue avgdel; /* average transit delay */
};

/*
 * Protocol Levels
 *
 * The values defined for these Protocol Levels are recommended-only,
 * not mandatory.
 */

#define T_ISO_TP 0x0100
#define ISO_TP 0x0100 (LEGACY)
```

```

/*
 * Options for Quality of Service and Expedited Data (ISO 8072:1994)
 *
 * The values defined for these QoS and Expedited Data are
 * recommended-only, not mandatory.
 */

#define T_TCO_THROUGHPUT 0x0001
#define TCO_THROUGHPUT 0x0001 (LEGACY)
#define T_TCO_TRANSDEL 0x0002
#define TCO_TRANSDEL 0x0002 (LEGACY)
#define T_TCO_RESERRORRATE 0x0003
#define TCO_RESERRORRATE 0x0003 (LEGACY)
#define T_TCO_TRANSFFAILPROB 0x0004
#define TCO_TRANSFFAILPROB 0x0004 (LEGACY)
#define T_TCO_ESTFAILPROB 0x0005
#define TCO_ESTFAILPROB 0x0005 (LEGACY)
#define T_TCO_RELFAILPROB 0x0006
#define TCO_RELFAILPROB 0x0006 (LEGACY)
#define T_TCO_ESTDELAY 0x0007
#define TCO_ESTDELAY 0x0007 (LEGACY)
#define T_TCO_RELDELAY 0x0008
#define TCO_RELDELAY 0x0008 (LEGACY)
#define T_TCO_CONNRRESIL 0x0009
#define TCO_CONNRRESIL 0x0009 (LEGACY)
#define T_TCO_PROTECTION 0x000a
#define TCO_PROTECTION 0x000a (LEGACY)
#define T_TCO_PRIORITY 0x000b
#define TCO_PRIORITY 0x000b (LEGACY)
#define T_TCO_EXPD 0x000c
#define TCO_EXPD 0x000c (LEGACY)

#define T_TCL_TRANSDEL 0x000d
#define TCL_TRANSDEL 0x000d (LEGACY)
#define T_TCL_RESERRORRATE T_TCO_RESERRORRATE
#define TCL_RESERRORRATE T_TCO_RESERRORRATE (LEGACY)
#define T_TCL_PROTECTION T_TCO_PROTECTION
#define TCL_PROTECTION T_TCO_PROTECTION (LEGACY)
#define T_TCL_PRIORITY T_TCO_PRIORITY
#define TCL_PRIORITY T_TCO_PRIORITY (LEGACY)

/*
 * Management Options
 *
 * The values defined for these Management Options are
 * recommended-only, not mandatory.
 */

#define T_TCO_LTPDU 0x0100
#define TCO_LTPDU 0x0100
#define T_TCO_ACKTIME 0x0200
#define TCO_ACKTIME 0x0200
#define T_TCO_REASTIME 0x0300
#define TCO_REASTIME 0x0300

```



## Headers and Definitions for XTI

```
#define T_TCO_EXTFORM 0x0400
#define TCO_EXTFORM 0x0400
#define T_TCO_FLOWCTRL 0x0500
#define TCO_FLOWCTRL 0x0500
#define T_TCO_CHECKSUM 0x0600
#define TCO_CHECKSUM 0x0600
#define T_TCO_NETEXP 0x0700
#define TCO_NETEXP 0x0700
#define T_TCO_NETRECPTCF 0x0800
#define TCO_NETRECPTCF 0x0800
#define T_TCO_PREFCLASS 0x0900
#define TCO_PREFCLASS 0x0900
#define T_TCO_ALTCLASS1 0x0a00
#define TCO_ALTCLASS1 0x0a00
#define T_TCO_ALTCLASS2 0x0b00
#define TCO_ALTCLASS2 0x0b00
#define T_TCO_ALTCLASS3 0x0c00
#define TCO_ALTCLASS3 0x0c00
#define T_TCO_ALTCLASS4 0x0d00
#define TCO_ALTCLASS4 0x0d00

#define T_TCL_CHECKSUM T_TCO_CHECKSUM
#define TCL_CHECKSUM T_TCO_CHECKSUM

/* INTERNET-SPECIFIC ENVIRONMENT */
/*
 * TCP level
 *
 * The values defined for the TCP Level are recommended-only,
 * not mandatory.
 */

#define T_INET_TCP 0x6
#define INET_TCP 0x6 (LEGACY)

/*
 * TCP-level Options
 *
 * The values defined for the TCP-level Options are recommended-only,
 * not mandatory.
 */

#define T_TCP_NODELAY 0x1 /* don't delay packets to coalesce */
#define TCP_NODELAY 0x1 /* (LEGACY) */
#define T_TCP_MAXSEG 0x2 /* get maximum segment size */
#define TCP_MAXSEG 0x2 /* (LEGACY) */
#define T_TCP_KEEPALIVE 0x8 /* check, if connections are alive */
#define TCP_KEEPALIVE 0x8 /* (LEGACY) */
```

```

/*
 * Structure used with TCP_KEEPPALIVE option.
 */
struct t_kpalive {
 t_scalar_t kp_onoff; /* option on/off */
 t_scalar_t kp_timeout; /* timeout in minutes */
};

#define T_GARBAGE 0x02 /* LEGACY */

/*
 * UDP level
 *
 * The values defined for the UDP Level are recommended-only,
 * not mandatory.
 */

#define T_INET_UDP 0x11
#define INET_UDP 0x11 (LEGACY)

/*
 * UDP-level Options
 *
 * The values defined for the UDP-level Options are recommended-only,
 * not mandatory.
 */

#define T_UDP_CHECKSUM T_TCO_CHECKSUM /* checksum computation */
#define UDP_CHECKSUM T_TCO_CHECKSUM /* LEGACY */

/*
 * IP level
 *
 * The values defined for the IP Level are recommended-only,
 * not mandatory.
 */

#define T_INET_IP 0x0
#define INET_IP 0x0 (LEGACY)

/*
 * IP-level Options
 *
 * The values defined for the IP-level Options are recommended-only,
 * not mandatory.
 */

#define T_IP_OPTIONS 0x1 /* IP per-packet options */
#define IP_OPTIONS 0x1 /* (LEGACY) */
#define T_IP_TOS 0x2 /* IP per-packet type of service */
#define IP_TOS 0x2 /* (LEGACY) */
#define T_IP_TTL 0x3 /* IP per-packet time to live */
#define IP_TTL 0x3 /* (LEGACY) */

```

## Headers and Definitions for XTI

```
#define T_IP_REUSEADDR 0x4 /* allow local address reuse */
#define IP_REUSEADDR 0x4 /* (LEGACY) */
#define T_IP_DONTROUTE 0x10 /* just use interface addresses */
#define IP_DONTROUTE 0x10 /* (LEGACY) */
#define T_IP_BROADCAST 0x20 /* permit sending of broadcast msgs */
#define IP_BROADCAST 0x20 /* (LEGACY) */

/*
 * IP_TOS precedence levels
 */

#define T_ROUTINE 0
#define T_PRIORITY 1
#define T_IMMEDIATE 2
#define T_FLASH 3
#define T_OVERRIDEFLASH 4
#define T_CRITIC_ECP 5
#define T_INETCONTROL 6
#define T_NETCONTROL 7

/*
 * IP_TOS type of service
 */

#define T_NOTOS 0
#define T_LDELAY (1 << 4)
#define T_HITHRPT (1 << 3)
#define T_HIREL (1 << 2)
#define T_LOCOST (1 << 1)

#define SET_TOS(prec, tos) ((0x7 & (prec)) << 5 | (0x1c & (tos)))
```



## *Abbreviations*

|       |                                                |
|-------|------------------------------------------------|
| CO    | Connection-oriented                            |
| CL    | Connectionless                                 |
| EM    | Event Management                               |
| ETSDU | Expedited Transport Service Data Unit          |
| ISO   | International Organization for Standardization |
| OSI   | Open System Interconnection                    |
| SVID  | System V Interface Definition                  |
| TC    | Transport Connection                           |
| TCP   | Transmission Control Protocol                  |
| TLI   | Transport Level Interface                      |
| TSAP  | Transport Service Access Point                 |
| TSDU  | Transport Service Data Unit                    |
| UDP   | User Datagram Protocol                         |
| XTI   | X/Open Transport Interface                     |
| XEM   | X/Open Event Management Interface              |



# Minimum OSI Functionality (CAE Specification)

## H.1 General

The purpose of this specification is to provide a simple API exposing a minimum set of OSI Upper Layers functionality (mOSI).

### H.1.1 Rationale for using XTI-mOSI

This appendix uses the concept of a minimal set of OSI upper layer facilities that support basic communication applications. A Basic Communication Application simply requires the ability to open and close communications with a peer and to send and receive messages with a peer.

XTI-mOSI is designed specifically for Basic Communication Applications that are in one of these categories:

- applications that are to be migrated from the Internet world (TCP or UDP) or from a NetBIOS environment to OSI
- applications accessing the OSI transport service that wish to migrate to an OSI seven-layer, conformant environment
- applications that require a simple octet-stream connection between peer processes. The benefit of XTI-mOSI to these applications is that it extends the family of *transport services* that are available via a single, protocol independent, API.

### H.1.2 Migrant Applications

For the first kind of applications (those migrating to OSI or intended to work over a variety of *transport* mechanisms), the migration effort will be greatly simplified if they were already using XTI — mOSI offers several new options, but, as described later in this section, default values are generally provided.

In addition to applications already using XTI, the X Window System (X) and Internet Protocol Suite applications (in general) are examples of potential Migrant applications.

### H.1.3 OSI Functionality

mOSI is suited to applications that require only the Minimal Upper Layer facilities which are described in the profile ISO/IEC DISP 11188 — Common Upper Layer Requirements, Part 3: Minimal OSI upper layer facilities, expected to reach International Standardized Profile (ISP) status in the first half of 1995. These are:

- ACSE Kernel functional unit
- Presentation Kernel functional unit
- Session Kernel and Full Duplex functional units.

The XTI-mOSI interface provides access to OSI ACSE and Presentation services. With mOSI, the optional parameters available to the application have been selected with the intent of facilitating interoperability and diagnostic of problems. They are described later in this section.

Most applications only need the Kernel functionality. This is even true for most of the OSI standard applications: Remote Database Access (RDA), Directory (X.500), FTAM without recovery, OSI Distributed Transaction Processing (TP) without 2-phase commitment, OSI Management.

#### H.1.4 mOSI API versus XAP

X/Open has developed the XAP interface (ACSE/Presentation API) to provide access to ACSE and Presentation functionality. It provides an interface which spans from implementations of minimal OSI through to full ACSE/Presentation/Session with all functional units.

XTI-mOSI has been developed to support migrant transport applications, and also applications which require a simple octet stream connection between peers.

Application designers should consider which API is most appropriate to their needs.

Applications which are being ported from other transport providers, or require a simple octet string connection, should use XTI-mOSI.

Applications which require session function units outside kernel, or which require multiple presentation contexts, or which may require additional session/presentation by mOSI in the future, should use XAP.

When only minimal OSI support is required, the user should consider the availability of XTI-mOSI and XAP on target hardware platforms when selecting an interface.

#### H.1.5 Upper Layers Functionality Exposed via mOSI

These are presented as they are exposed via mOSI options and specific parameters.

##### H.1.5.1 Naming and Addressing Information used by mOSI

The **addr** structure (used in *t\_bind()*, *t\_connect()*, *t\_accept()*) is a combined naming and addressing data type, identifying one end or the other of the association.

The address part is a Presentation Address. The *calling* and *called* addresses are required parameters, while the use of a *responding* address is optional.

The name part (Application Process (AP) Title, Application Entity (AE) Qualifier and the AP and AE invocation-identifiers) is always optional.

ISO Directory facilities, when available, can relate the name parts (identifying specific applications) to the addresses of the real locations where they can be accessed.

The general format of the **addr** structure can be found in Section H.5 on page 313, while its precise structure is implementation dependent.

##### H.1.5.2 XTI Options Specific to mOSI

- Application Context Name

An application context name identifies a set of tasks to be performed by an application. It is exchanged during association establishment with the purpose of conveying a common understanding of the work to be done.

This parameter is exposed to offer some negotiation capabilities to the application and to increase the chances of interoperability.

When receiving a non suitable or unknown value from a peer application, the application may propose an alternate value or decide to terminate prematurely the association.



A default value (in the form of an Object Identifier) is provided, identifying a generic XTI-mOSI application. Its value can be found in Section H.5 on page 313.

- Presentation Contexts

A presentation context is the association of an abstract syntax with a transfer syntax. The presentation context is used by the application to identify how the data is structured and by the OSI Application Layer to identify how the data should be encoded/decoded.

A *generic* presentation context is defined for a stream-oriented, unstructured, data transfer service with *null* encoding:

*abstract syntax:* The single data type of this abstract syntax is a sequence of octets that are defined in the application protocol specification as being consecutive octets on a stream-oriented transport mechanism, without regard for any semantic or other boundaries.

*transfer syntax:* The data value shall be represented as an octet-aligned presentation data value. If two or more data values are concatenated together they are considered to be a single (longer) data value. (This is the *null* encoding rule).

The value of the Object Identifiers for this *generic* presentation context can be found in Section H.5 on page 313.

- Presentation Context Definition and Result List, Defined Context Set

As negotiation occurs between the peer OSI Application layers, the presentation context(s) proposed by the application may not be accepted.

The Presentation Context Definition and Result List indicates, for each of the proposed presentation context, if it is accepted or, if not, provides a reason code; the application may choose to terminate the association prematurely if it does not suit its requirements.

## H.2 Options

Options are formatted according to the structure `t_opthdr` as described in Chapter 6 on page 39. An OSI provider compliant to this specification supports all, none or a subset of the options defined in Section H.2.1. An implementation may restrict the use of any of the options by offering them in privileged or `read_only` mode.

An explanation of when an application may benefit from using the XTI options specific to mOSI can be found in Section H.1 on page 297.

### H.2.1 ACSE/Presentation Connection-mode Service

The protocol level for all subsequent options is `T_ISO_APCO`.

All options have end-to-end significance (see Chapter 6 on page 39). They may be negotiated in the XTI states `T_IDLE` and `T_INCON`, and are read-only in all other states except `T_UNINIT`. The structures referenced are specified in Section H.5 on page 313.

| Option Name                 | Type of Option Value                                                          | Legal Option Value            | Meaning                                         |
|-----------------------------|-------------------------------------------------------------------------------|-------------------------------|-------------------------------------------------|
| <code>T_AP_CNTX_NAME</code> | Object identifier item (see Section H.5 on page 313)                          | see text<br>default: see text | Application Context Name                        |
| <code>T_AP_PCL</code>       | Presentation Context Definition and Result list (see Section H.5 on page 313) | see text<br>default: see text | Presentation Context Definition and Result List |

**Table H-1** APCO-level Options

#### Further Remarks

- Application Context Name

A default value (for a *generic* XTI-mOSI application) is provided. It is defined in Section H.5 on page 313.

The application may choose to propose, through this option, a value different from the default one. The application may also use this option to check the value returned by the peer application and decide if the association should be kept or terminated.

- Presentation Context Definition and Result List

A default is provided: a list with one presentation context (the stream oriented, unstructured, data transfer service with *null* encoding — this is described in section Section H.1 on page 297). The abstract syntax is the default abstract syntax and the transfer syntax is the default transfer syntax, as specified in Section H.5 on page 313.

The codes for the result of negotiation and reason for rejection are defined in Section H.5 on page 313. The responding application, after reading this option, may choose to continue or terminate the association.

Only a single abstract syntax and transfer syntax can be used by XTI-mOSI. On `t_accept()`, this is assumed to be the first usable abstract syntax and the first transfer syntax for that abstract syntax.

### Negotiation of Abstract and Transfer Syntax

When initiating a connection, the application proposes one or more presentation contexts, each comprising an abstract syntax and one or more transfer syntaxes in the Presentation Context Definition and Result List option (or omits this option to select the CULR-3 defaults), and issues a `t_connect()`.

If the connection is accepted, the Presentation Context Definition and Result List is updated to reflect the results of negotiation for each element of the context list, and a single presentation context is selected.

**Note:** If the responder accepts multiple presentation contexts, the XTI-mOSI provider aborts the connection on receipt of the A-ASSOCIATE confirm.

When responding to a remote connect, the application can specifically mark presentation contexts as rejected using the `res` field, and can re-order the syntax array to select a single transfer syntax.

On calling `t_accept()`, the first presentation context marked as accepted is selected, and all other contexts omitted or not marked rejected-user are marked as by the provider as rejected (T\_PCL\_PREJ\_LMT\_DCS\_EXCEED). In an accepted context, the provider will accept the first (or only remaining) transfer syntax.

**Note:** On return to the application from `t_listen()`, all supportable presentation contexts are marked as accepted in the T\_AP\_PCL option, and all unsupported contexts are marked as *rejected-provider*. This permits the application to return the same option value on `t_accept()` (or leave it unchanged) to select the first available abstract syntax and transfer syntax.

### Management Options

No management options are defined.

## H.2.2 ACSE/Presentation Connectionless-mode Service

The protocol level for all subsequent options is T\_ISO\_APCL.

All options have end-to-end significance (see Chapter 6 on page 39). They may be negotiated in all XTI states except T\_UNINIT. The structures referenced are specified in Section H.5 on page 313.

| Option Name    | Type of Option Value                                                          | Legal Option Value            | Meaning                                         |
|----------------|-------------------------------------------------------------------------------|-------------------------------|-------------------------------------------------|
| T_AP_CNTX_NAME | Object identifier item (see Section H.5 on page 313)                          | see text<br>default: see text | Application Context Name                        |
| T_AP_PCL       | Presentation Context Definition and Result list (see Section H.5 on page 313) | see text<br>default: see text | Presentation Context Definition and Result List |

**Table H-2** APCL-level Options

**Further Remarks**

- Application Context Name

A default value (for a *generic* XTI-mOSI application) is provided. It is defined in Section H.5 on page 313.

The application may choose to propose, through this option, a value different from the default one. The application may also use this option to check the value returned by the peer application and decide if the datagram should be kept or discarded.

- Presentation Context Definition and Result List

In connectionless mode, the transfer syntaxes are not negotiated. Their use is determined by the sending application entity, and must be acceptable by the receiving application entity. A default value is provided by XTI: a list with one element, the *generic* presentation context (the stream-oriented, unstructured, data transfer service with *null* encoding described in Section H.1 on page 297). The corresponding abstract and transfer syntaxes are specified in Section H.5 on page 313.

Only a single abstract syntax and transfer syntax can be used by connectionless-mode XTI-mOSI. If more than one presentation context is present in the options list for *t\_sndudata()*, the first is used.

**Management Options**

No management options are defined.

**H.2.3 Transport Service Options**

Some of the options defined for XTI ISO Transport Connection-mode Service or Transport Connectionless-mode Service may be made available to mOSI users: the Options for Quality of Service.

These Options are defined in Section A.2.1.1 on page 216 and Section A.2.2.1 on page 220. The Quality of Service parameters are passed directly by the OSI Upper Layers to the Transport Layer. These options can thus be used to specify OSI Upper Layers quality of service parameters via XTI.

This facility is implementation dependent. An attempt to specify an unsupported option will return with the status field set to T\_NOTSUPPORT.

None of these options are available with an ISO-over-TCP transport provider.

### H.3 Functions

Functions *rcvreldata()* (see *t\_rcvreldata()* on page 98) and *sndreldata()* (see *t\_sndreldata()* on page 114) were introduced as part of this XTI-mOSI functionality. The rationale for this is that for ISO ACSE providing an orderly release mechanism, user data is a parameter of the release service, so when mapping XTI primitives to ACSE/Presentation (XTI-mOSI), disconnection user data may be received from peer applications. Although abortive release primitives (*t\_snddis*, *t\_rcvdis*) permit sending and receiving of user data, orderly release primitives (*t\_sndrel*, *t\_rcvrel*) do not. Therefore, new functions having a user data parameter *t\_rcvreldata()* and *t\_sndreldata()* were added to provide the necessary support to handle this user data.

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>t_accept()</i>  | <p>If <i>fd</i> is not equal to <i>resfd</i>, <i>resfd</i> should either be in state T_UNBND or be in state T_IDLE with the <i>qlen</i> parameter set to 0.</p> <p>The <i>addr</i> parameter passed to/returned from <i>t_bind()</i> when <i>resfd</i> is bound may be different from the <i>addr</i> parameter corresponding to <i>fd</i>.</p> <p>The <i>opt</i> parameter may be used to change the Application Context Name received.</p>                                                                                                                              |
| <i>t_alloc()</i>   | No special considerations for mOSI providers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <i>t_bind()</i>    | <p>The <i>addr</i> field of the <b>t_bind</b> structure represents the local presentation address and optionally the local AP Title, AE Qualifier, AP and AE invocation-identifiers (see Section H.1 on page 297 and Section H.5 on page 313 for more details).</p> <p>This local <i>addr</i> field is used, depending on the XTI primitive, as the calling, called or responding address, the called address being different from the responding address only when two different file descriptors (<i>fd</i>, <i>resfd</i>), bound to different addresses, are used.</p> |
| <i>t_close()</i>   | Any connections that are still active at the endpoint are abnormally terminated. The peer applications will be informed of the disconnection by a [T_DISCONNECT] event. The value of the disconnection reason will be T_AC_ABRT_NSPEC.                                                                                                                                                                                                                                                                                                                                    |
| <i>t_connect()</i> | <p>The <i>sndcall</i>→<i>addr</i> structure specifies the Called Presentation Address. The <i>rcvcall</i>→<i>addr</i> structure specifies the Responding Presentation Address. The structure may also be used to assign values for the Called AP Title, Called AE Qualifier, Called AP invocation-identifier and Called AE invocation-identifier.</p> <p>Before the call, the <i>sndcall</i>→<i>opt</i> structure may be used to request an Application Context name or Presentation Context different from the default value.</p>                                        |
| <i>t_error()</i>   | No special considerations for mOSI providers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <i>t_free()</i>    | No special considerations for mOSI providers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <i>t_getinfo()</i> | The information supported by <i>t_getinfo()</i> reflects the characteristics of the <i>transport</i> connection, or if no connection is established, the default characteristics of the underlying OSI layers. In all possible states except T_DATAXFER, the function <i>t_getinfo()</i> returns in the parameter <i>info</i> the same information as was returned by <i>t_open()</i> . In state T_DATAXFER, however, the information returned in <i>info</i> → <i>connect</i> and <i>info</i> → <i>discon</i> may differ.                                                |

The parameters of the *t\_getinfo()* function are summarised in the table below.

| Parameters    | Before call | After call      |                     |
|---------------|-------------|-----------------|---------------------|
|               |             | Connection-mode | Connectionless-mode |
| fd            | x           | /               | /                   |
| info→addr     | /           | x               | x                   |
| info→options  | /           | x               | x                   |
| info→tsdu     | /           | T_INFINITE (-1) | T_INFINITE (-1)     |
| info→etsdu    | /           | (T_INVALID (-2) | T_INVALID (-2)      |
| info→connect  | /           | x               | T_INVALID (-2)      |
| info→discon   | /           | x               | T_INVALID (-2)      |
| info→servtype | /           | T_COTS_ORD      | T_CLTS              |
| info→flags    | /           | 0               | 0                   |

x equals an integral number greater than 0.

The values of the parameters in the **t\_info** structure for the *t\_getinfo()* function reflect the mOSI provider particularities.

- connect, discon

The values returned in *info→connect* and *info→discon* in state T\_DATAXFER may differ from the values returned by *t\_open()*: negotiation takes place during association establishment and, as a result, these values may be reduced. For *info→connect*, this change of value may be indicated by the provider, but is of little use to the application.

- flags

mOSI does not support sending of TSDU of zero length, so this value equals 0.

|                        |                                                                                                                                                                                                                                                                                                        |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>t_getprotaddr()</i> | The protocol addresses are naming and addressing parameters as defined in Section H.1 on page 297 and Section H.5 on page 313.                                                                                                                                                                         |
| <i>t_getstate()</i>    | No special considerations for mOSI providers.                                                                                                                                                                                                                                                          |
| <i>t_listen()</i>      | The <i>call→addr</i> structure contains the remote Calling Presentation Address and the remote Calling AP Title, AE Qualifier, and AP and AE invocation identifiers if received.<br><br>Incoming user data encoded as multiple presentation data values will cause the TBADDDATA error to be returned. |
| <i>t_look()</i>        | Since expedited data is not supported for a mOSI provider, T_EXDATA and T_GOEXDATA events cannot occur.                                                                                                                                                                                                |
| <i>t_open()</i>        | <i>t_open()</i> is called as the first step in the initialisation of a <i>transport</i> endpoint. This function returns various default characteristics of the underlying OSI layers.                                                                                                                  |

The parameters of the `t_open()` function are summarised in the table below.

| Parameters    | Before call | After call      |                     |
|---------------|-------------|-----------------|---------------------|
|               |             | Connection-mode | Connectionless-mode |
| name          | x           | /               | /                   |
| oflag         | /           | /               | /                   |
| info→addr     | /           | x               | x                   |
| info→options  | /           | x               | x                   |
| info→tsdu     | /           | T_INFINITE (-1) | T_INFINITE (-1)     |
| info→etsdu    | /           | T_INVALID (-2)  | T_INVALID (-2)      |
| info→connect  | /           | x               | T_INVALID (-2)      |
| info→discon   | /           | x               | T_INVALID (-2)      |
| info→servtype | /           | T_COTS_ORD      | T_CLTS              |
| info→flags    | /           | 0               | 0                   |

x equals an integral number greater than 0.

The values of the parameters in the **t\_info** structure reflect mOSI limitations as follows:

- connect, discon

These values are limited by the version of the session supported by the mOSI provider, and are generally much larger than those supported by an ISO Transport or TCP provider.

- flags

mOSI does not support sending of *TSDU* of zero length, so this value equals 0.

**Note:** The name (device file) parameter passed to `t_open()` will differ when the application accesses an mOSI provider or an ISO Transport provider.

|                             |                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>t_optmgt()</code>     | The options available with mOSI providers are described in section Section H.2 on page 300.                                                                                                                                                                                                                                                                                     |
| <code>t_rcv()</code>        | The flags parameter will never be set to [T_EXPEDITED], as expedited data transfer is not supported.                                                                                                                                                                                                                                                                            |
| <code>t_rcvconnect()</code> | The <code>call→addr</code> structure specifies the remote Responding Presentation Address, and the remote responding AP Title, AE Qualifier, and AP and AE invocation identifiers if received.<br><br>The <code>call→opt</code> structure may also contain an Application Context Name and/or Presentation Context Definition Result List.                                      |
| <code>t_rcvdis()</code>     | Possible values for disconnection reason codes are specified in Section H.5 on page 313.                                                                                                                                                                                                                                                                                        |
| <code>t_rcvrel()</code>     | With this primitive, user data cannot be received on normal release: any user data in the received flow is discarded (see <code>t_rcvreldata()</code> on page 98).                                                                                                                                                                                                              |
| <code>t_rcvudata()</code>   | The <code>unitdata→addr</code> structure specifies the remote Presentation address, and optionally the remote AP Title, AE Qualifier, AP and AE invocation-identifiers. If the T_MORE flag is set, an additional <code>t_rcvudata()</code> call is needed to retrieve the entire A-UNIT-DATA service unit. Only normal data is returned via the <code>t_rcvudata()</code> call. |

|                     |                                                                                                                                                                                                                              |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>t_rcvuderr()</i> | This function is not supported by a mOSI provider since badly formed A-UNIT-DATA APDUs are discarded.                                                                                                                        |
| <i>t_snd()</i>      | Zero-length TSDUs are not supported.<br><br>Since expedited data transfer is not supported for a mOSI provider, the parameter flags shall not have [T_EXPEDITED] set.                                                        |
| <i>t_snddis()</i>   | No special considerations for mOSI providers.                                                                                                                                                                                |
| <i>t_sndrel()</i>   | With this primitive, user data cannot be sent on normal release (see <i>t_sndreldata()</i> on page 114).                                                                                                                     |
| <i>t_sndudata()</i> | The <i>unitdata→addr</i> structure specifies the remote Presentation address, and optionally the remote AP Title, AE Qualifier, AP and AE invocation-identifiers. Only normal data is sent via the <i>t_sndudata()</i> call. |
| <i>t_strerror()</i> | No special considerations for mOSI providers.                                                                                                                                                                                |
| <i>t_sync()</i>     | No special considerations for mOSI providers.                                                                                                                                                                                |
| <i>t_unbind()</i>   | No special considerations for mOSI providers.                                                                                                                                                                                |



## H.4 Implementors' Notes

### H.4.1 Upper Layers FUs, Versions and Protocol Mechanisms

The implementation negotiates:

Session: Kernel, Full Duplex, version 2, or version 1 if version 2 not supported, no segmentation.

Other session protocol mechanisms are out of scope, except Basic Concatenation which is mandatory and transparent to the application.

Presentation: Kernel, Normal Mode

ACSE: Kernel

If invalid (non-negotiable) options are requested by the peer and detected by the provider once the association is already established (such as the ACSE presentation context missing in the Defined Context Set), the association is rejected via an A-(P)-ABORT generated by the implementation.

### H.4.2 Mandatory and Optional Parameters

- If the Local Presentation Address is not passed to *t\_bind()* in *req→addr*, then it is returned in *ret→addr*.
- The remote (called) Presentation Address (in *t\_connect()*, *sndcall→addr*) parameter must be explicitly set by the application.
- The following parameters are mandatory for the protocol machine, but default values are provided. If the application does not wish to set the corresponding parameter, the default value will be used. The default value may be changed through *t\_optmgt* (see Section H.2 on page 300):
  - Application Context Name (opt parameter)
  - Presentation Context List (opt parameter).

The presentation context of ACSE is required and used. The user should not request it as the implementation will insert it automatically in the context list.

If the user does not specifically request an Application Context name via the opt parameter of *t\_accept()* (that is, for the A-Associate response), the implementation uses the Application Context name that was received in the A-Associate indication.

- The following parameters are optional for the protocol and default values of null are defined. If the application does not set them otherwise, they are omitted from the outgoing protocol stream.
  - local AP-title (in *t\_bind()*, *req→addr*)
  - called AP-title (in *t\_connect()*, *sndcall→addr*)
  - responding AP-title (if *t\_accept()* specifies a new accepting endpoint *resfd*, in the protocol address bound to *resfd*)
  - local AE-qualifier (in *t\_bind()*, *req→addr*)
  - called AE-qualifier (in *t\_connect()*, *sndcall→addr*)
  - responding AE-qualifier (if *t\_accept()* specifies a new accepting endpoint *resfd*, in the protocol address bound to *resfd*).

- local AP and AE invocation-identifiers (in *t\_bind()*, *req*→*addr*)
- called AP and AE invocation-identifiers (in *t\_connect()*, *sndcall*→*addr*)
- responding AP and AE invocation-identifiers (if *t\_accept()* specifies a new accepting endpoint *resfd*, in the protocol address bound to *resfd*).
- The following parameters are optional for the protocol machine and not supported through the XTI interface. Their handling is implementation-defined. Received values in the incoming protocol stream, if any, are discarded:
  - ACSE Protocol Version (default= version 1)
  - Presentation Protocol Version (default= version 1)
  - ACSE Implementation Information
  - Session connection identifiers.

During association establishment (that is, before the XTI-mOSI provider negotiates acceptance of a single abstract syntax/transfer syntax pair), an XTI-mOSI application initiating the association will only send a single presentation data value in the user information parameter. The XTI-mOSI provider will insure that the first abstract syntax and transfer syntax pair being negotiated is the one required for its encoding.

### H.4.3 Mapping XTI Functions to ACSE/Presentation Services

In the following tables, the definition of which parameters are mandatory and which are optional can be found in ISO/IEC DISP 11183 — Common Upper Layers Requirements, part 3 (see reference **CULR**).

#### H.4.3.1 Connection-mode Services

##### Association Establishment (successful, unsuccessful)

**Note:** XTI does not support the concept of a negative association establishment; that is, the equivalent of a negative A-ASSOCIATE response. That is, an XTI-mOSI implementation does not generate an AARE- APDU.

To reject an association request, the responding application issues *t\_snddis()*, which is mapped to a A-ABORT.

However, a negative A-ASSOCIATE confirm (AARE- APDU) may be received from a non-XTI OSI peer. The negative A-ASSOCIATE confirm event is mapped to *t\_rcvdis()*.

Table H-3 Association Establishment

| XTI call     | Parameter                                                                                                                                   | Service          | Parameter                                                                                                                                                                                                                           |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| t_connect    | sndcall→addr<br>sndcall→addr (1)<br>sndcall→addr (1)<br>sndcall→addr<br>sndcall→addr<br>sndcall→opt (2)<br>sndcall→opt (3)<br>sndcall→udata | A-ASSOCIATE req  | Called Presentation Address<br>Called AP Title<br>Called AE Qualifier<br>Called AP invocation-identifier<br>Called AE invocation-identifier<br>Application Context Name<br>P-context Definition and Result List<br>User Information |
| {t_bind}     | req   ret→addr                                                                                                                              |                  | Calling Presentation Address                                                                                                                                                                                                        |
| {t_bind}     | req   ret→addr                                                                                                                              |                  | Calling AP Title                                                                                                                                                                                                                    |
| {t_bind}     | req   ret→addr                                                                                                                              |                  | Calling AE Qualifier                                                                                                                                                                                                                |
| t_listen     | call→addr<br>call→addr (1)<br>call→addr (1)<br>call→opt<br>call→opt (4)<br>call→udata                                                       | A-ASSOCIATE ind  | Calling Presentation Address<br>Calling AP Title<br>Calling AE Qualifier<br>Application Context Name<br>P-context Definition and Result List<br>User Information                                                                    |
| {t_bind}     | req   ret→addr                                                                                                                              |                  | Called Presentation Address                                                                                                                                                                                                         |
| {t_bind}     | req   ret→addr (1)                                                                                                                          |                  | Called AP Title                                                                                                                                                                                                                     |
| {t_bind}     | req   ret→addr (1)                                                                                                                          |                  | Called AE Qualifier                                                                                                                                                                                                                 |
| {t_bind}     | req   ret→addr                                                                                                                              |                  | Calling AP invocation-identifier                                                                                                                                                                                                    |
| {t_bind}     | req   ret→addr                                                                                                                              |                  | Calling AE invocation-identifier                                                                                                                                                                                                    |
| t_accept     | call→addr<br>call→opt<br>call→opt<br>call→udata                                                                                             | A-ASSOCIATE rsp+ | not used: Calling Presentation Address<br>Application Context Name<br>P-context Definition and Result List<br>User Information                                                                                                      |
| {internal}   | ::="accepted"                                                                                                                               |                  | Result                                                                                                                                                                                                                              |
| {t_bind}     | req   ret→addr                                                                                                                              |                  | Responding Presentation Address                                                                                                                                                                                                     |
| {t_bind}     | req   ret→addr (1)                                                                                                                          |                  | Responding AP Title                                                                                                                                                                                                                 |
| {t_bind}     | req   ret→addr (1)                                                                                                                          |                  | Responding AE Qualifier                                                                                                                                                                                                             |
| {t_bind}     | req   ret→addr                                                                                                                              |                  | Responding AP invocation-identifier                                                                                                                                                                                                 |
| {t_bind}     | req   ret→addr                                                                                                                              |                  | Responding AE invocation-identifier                                                                                                                                                                                                 |
| not sent     |                                                                                                                                             | A-ASSOCIATE rsp- |                                                                                                                                                                                                                                     |
| t_connect    | (synchronous mode)                                                                                                                          | A-ASSOCIATE cnf+ |                                                                                                                                                                                                                                     |
|              | rcvcall→addr                                                                                                                                |                  | Responding Presentation Address                                                                                                                                                                                                     |
|              | rcvcall→addr                                                                                                                                |                  | Responding AP Title                                                                                                                                                                                                                 |
|              | rcvcall→addr                                                                                                                                |                  | Responding AE Qualifier                                                                                                                                                                                                             |
|              | rcvcall→addr                                                                                                                                |                  | Responding AP invocation-identifier                                                                                                                                                                                                 |
|              | rcvcall→addr                                                                                                                                |                  | Responding AE invocation-identifier                                                                                                                                                                                                 |
|              | rcvcall→opt                                                                                                                                 |                  | Application Context Name                                                                                                                                                                                                            |
|              | rcvcall→opt                                                                                                                                 |                  | P-context Definition and Result List                                                                                                                                                                                                |
|              | rcvcall→udata                                                                                                                               |                  | User Information                                                                                                                                                                                                                    |
| {internal}   | ::="accepted"                                                                                                                               |                  | Result                                                                                                                                                                                                                              |
| {internal}   | ::="ACSE service-user"                                                                                                                      |                  | Result Source                                                                                                                                                                                                                       |
| t_rcvconnect | (asynchronous mode)                                                                                                                         | A-ASSOCIATE cnf+ |                                                                                                                                                                                                                                     |
|              | call→addr                                                                                                                                   |                  | Responding Presentation Address                                                                                                                                                                                                     |
|              | call→addr                                                                                                                                   |                  | Responding AP Title                                                                                                                                                                                                                 |

| XTI call    | Parameter                       | Service          | Parameter                            |
|-------------|---------------------------------|------------------|--------------------------------------|
|             | call→addr                       |                  | Responding AE Qualifier              |
|             | call→addr                       |                  | Responding AP invocation-identifier  |
|             | call→addr                       |                  | Responding AE invocation-identifier  |
|             | call→opt                        |                  | Application Context Name             |
|             | call→opt                        |                  | P-context Definition and Result List |
|             | call→udata                      |                  | User Information                     |
| {discarded} | ::="accepted"                   |                  | Result                               |
| {discarded} | ::="ACSE service-user"          |                  | Result Source-diagnostic             |
| t_rcvdis    |                                 | A-ASSOCIATE cnf- |                                      |
|             | discon→udata                    |                  | User Information                     |
|             | discon→reason (5)               |                  | Result                               |
| {internal}  | ACSE serv-user   pres serv-prov |                  | Result Source-diagnostic             |
|             | {discarded}                     |                  | Application Context Name             |
|             | {discarded}                     |                  | P-context Definition and Result List |

**Notes:**

- (1) if either the AP title or AE qualifier is selected for sending, the other must be selected.
- (2) *sndcall→opt* or, if no option specified, default value
- (3) *sndcall→opt* or, if no option specified, default value, with ACSE added by provider
- (4) *call→opt* with ACSE context removed from the list passed to user
- (5) combines Result and Result Source-diagnostic

**Data Transfer**

| XTI call | Parameter | Service    | Parameter |
|----------|-----------|------------|-----------|
| t_snd    | buf       | P-DATA req | User Data |
| t_rcv    | buf       | P-DATA ind | User Data |

**Table H-4** Data Transfer

**Association Release (orderly, abortive)**

This table makes the assumption that the XTI-mOSI provider supports the orderly release facility with user data (*t\_sndreldata()* — see *t\_rcvreldata()* on page 98, and *t\_rcvreldata()* — see *t\_rcvreldata()* on page 98). When this is not the case, User Information is not sent, Reason is supplied via an internal mechanism with A-RELEASE request and response, User Information and Reason received in A-RELEASE indication and confirmation are discarded.

| <b>XTI call</b>     | <b>Parameter</b>                                                | <b>Service</b> | <b>Parameter</b>               |
|---------------------|-----------------------------------------------------------------|----------------|--------------------------------|
| <i>t_sndrel2</i>    | <i>reldata</i> → <i>reason</i><br><i>reldata</i> → <i>udata</i> | A-RELEASE req  | Reason<br>User Information     |
| <i>t_rcvreldata</i> | <i>reldata</i> → <i>reason</i><br><i>reldata</i> → <i>udata</i> | A-RELEASE ind  | Reason<br>User Information     |
| <i>t_sndrel2</i>    | <i>reldata</i> → <i>reason</i><br><i>reldata</i> → <i>udata</i> | A-RELEASE rsp  | Reason<br>User Information     |
| <i>t_rcvreldata</i> | <i>reldata</i> → <i>reason</i><br><i>reldata</i> → <i>udata</i> | A-RELEASE cnf  | Reason<br>User Information     |
| <i>t_snddis</i>     | n/s<br><i>call</i> → <i>udata</i>                               | A-ABORT req    | Diagnostic<br>User Information |
| <i>t_rcvdis</i>     | <i>discon</i> → <i>reason</i><br><i>discon</i> → <i>udata</i>   | A-ABORT ind    | Diagnostic<br>User Information |
| <i>t_rcvdis</i>     | <i>discon</i> → <i>reason</i>                                   | A-P-ABORT ind  | Diagnostic                     |

**Table H-5** Association Release

## H.4.3.2 Connectionless-mode Services

| XTI call   | Parameter        | Service            | Parameter                            |
|------------|------------------|--------------------|--------------------------------------|
| t_sndudata |                  | A-UNIT-DATA source |                                      |
|            | unitdata→addr    |                    | Called Presentation Address          |
|            | unitdata→addr    |                    | Called AP Title                      |
|            | unitdata→addr    |                    | Called AE Qualifier                  |
|            | unitdata→addr    |                    | Called AP invocation-identifier      |
|            | unitdata→addr    |                    | Called AE invocation-identifier      |
|            | unitdata→opt (1) |                    | Application Context Name             |
|            | unitdata→opt (2) |                    | P-context Definition and Result List |
|            | unitdata→udata   |                    | User Information                     |
| {t_bind}   | req   ret→addr   |                    | Calling Presentation Address         |
| {t_bind}   | req   ret→addr   |                    | Calling AP Title                     |
| {t_bind}   | req   ret→addr   |                    | Calling AE Qualifier                 |
| {t_bind}   | req   ret→addr   |                    | Calling AP invocation-identifier     |
| {t_bind}   | req   ret→addr   |                    | Calling AE invocation-identifier     |
| t_rcvudata |                  | A-UNIT-DATA sink   |                                      |
|            | unitdata→addr    |                    | Calling Presentation Address         |
|            | unitdata→addr    |                    | Calling AP Title                     |
|            | unitdata→addr    |                    | Calling AE Qualifier                 |
|            | unitdata→addr    |                    | Calling AP invocation-identifier     |
|            | unitdata→addr    |                    | Calling AE invocation-identifier     |
|            | unitdata→opt     |                    | Application Context Name             |
|            | unitdata→opt (3) |                    | P-context Definition and Result List |
|            | unitdata→udata   |                    | User Information                     |
| {t_bind}   | req   ret→addr   |                    | Called Presentation Address          |
| {t_bind}   | req   ret→addr   |                    | Called AP Title                      |
| {t_bind}   | req   ret→addr   |                    | Called AE Qualifier                  |
| {t_bind}   | req   ret→addr   |                    | Called AP invocation-identifier      |
| {t_bind}   | req   ret→addr   |                    | Called AE invocation-identifier      |

Table H-6 Connectionless-mode ACSE Service

**Notes:**

- (1) unitdata→opt or, if no option specified, default value
- (2) unitdata→opt or, if no option specified, default value, with ACSE added by provider
- (3) unitdata→opt with ACSE context removed from the list passed to user

## H.5 Option Data Types and Structures

### SPECIFIC ISO ACSE/PRESENTATION OPTIONS

#### Naming and Addressing Datatype

The `buf[]` part of the `addr` structure is an `mosiaddr` structure defined in the following way:

```
struct t_mosiaddr {
 t_uscalar_t flags;
 t_scalar_t osi_ap_inv_id;
 t_scalar_t osi_ae_inv_id;
 unsigned int osi_apt_len;
 unsigned int osi_aeq_len;
 unsigned int osi_paddr_len;
 unsigned char osi_addr[T_AP_MAX_ADDR];
};
```

where:

- the `flags` field indicates the validity of the contents of the invocation identifier fields within the structure. One or more of the following bits may be set:

|                  |                                                               |
|------------------|---------------------------------------------------------------|
| T_OSI_AP_IID_BIT | The contents of the <code>osi_ap_inv_id</code> field is valid |
| T_OSI_AE_IID_BIT | The contents of the <code>osi_ae_inv_id</code> field is valid |

Unused bit in `flags` must be set zero by the user when creating a `t_mosiaddr` structure for sending, and should be ignore by the user on receipt.

In a `t_mosiaddr` structure, a bit set in `flags` indicates the presence of the corresponding invocation identifier in the PDU. Similarly a bit not set indicates absence of the corresponding invocation identifier in the PDU.

- the AP Title starts at  
`osi_addr[0]`
- the AE Qualifier starts at  
`osi_addr[T_ALIGN(osi_apt_len)]`
- the Presentation Address is at  
`osi_addr[T_ALIGN(osi_apt_len)+T_ALIGN(osi_aeq_len)]`
- T\_AP\_MAX\_ADDR is an implementation-defined constant.

The application is responsible for encoding/decoding the AP title and AE qualifier; alternatively, a lookup routine may be provided (outside the scope of this specification).

#### ACSE/Presentation Option Levels and Names

```
#define T_ISO_APCO 0x0200
#define T_ISO_APCL 0x0300
#define T_AP_CNTX_NAME 0x1
#define T_AP_PCL 0x2
```

**Object Identifier Representation within Options**

The application context, abstract syntax and transfer syntax all utilise object identifiers. An object identifier is held in the BER encoded form as a variable length item, whose length can be inferred from the length of the option, as in the following macro:

```
#define T_OPT_VALLEN(opt) (opt->len - sizeof(struct t_opthdr)).
```

The application is responsible for encoding/decoding the object identifier value. Alternatively, a lookup routine may be provided (outside the scope of this specification).

**Application Context Name Option**

The application context name option consists of an object identifier item as defined above.

**Presentation Context Definition and Result List Option**

The Presentation Context Definition and Result List option is used to propose a presentation context, giving its abstract and transfer syntax, and to hold the result of negotiation of a presentation context.

The presentation context definition and result list option is a variable sized option consisting of a `t_scalar_t` giving the number of presentation contexts followed by an array of that number of presentation context item offset elements. Each element is defined as:

```
struct t_ap_pco_el {
 t_scalar_t count;
 t_scalar_t offset;
}
```

Each presentation context items offset element gives: the *count* of syntax entries in the presentation context item (including the abstract syntax entry) and the *offset* of the presentation context item from the beginning of the Presentation Context Definition and Result List option value.

Each presentation context item consists of a header followed by an array.

The header is defined as:

```
struct t_ap_pc_item {
 t_scalar_t pci;
 t_scalar_t res;
}
```

where *pci* is a unique odd integer (if this is zero in the *sndcall* argument of *t\_connect()*, the provider will substitute an appropriate value), and *res* is the result of negotiation. The array of syntax offset elements immediately follows the header. Each element is defined as:

```
struct t_ap_syn_off {
 t_scalar_t size;
 t_scalar_t offset;
}
```

where *size* is the length of the syntax object identifier contents, and *offset* is the offset of the object identifier for the syntax from the beginning of the Presentation Context Definition and Result List option value.

The first element in the array of syntax offset elements refers to the abstract syntax the second to the first transfer syntax, and so on.



The following values are used in the *res* field of a presentation context item:

```
#define T_PCL_ACCEPT 0x0000
 /*pres. context accepted */
#define T_PCL_USER_REJ 0x0100
 /*pres. context rejected by peer application */
#define T_PCL_PREJ_RSN_NSPEC 0x0200
 /*prov. reject: no reason specified */
#define T_PCL_PREJ_A_SYTX_NSUP 0x0201
 /*prov. reject: abstract syntax not supported*/
#define T_PCL_PREJ_T_SYTX_NSUP 0x0202
 /*prov. reject:transfer syntax not supported */
#define T_PCL_PREJ_LMT_DCS_EXCEED 0x0203
 /*prov. reject: local limit on DCS exceeded */
```

For the default abstract syntax, transfer syntax and application context, this Appendix uses object identifiers which are specified in the profile (ISO/IEC pDISP 11188 - Common Upper Layer Requirements (CULR), Part 3: Minimal OSI upper layer facilities - OIW/EWOS working document). Thus the descriptions provided in this Appendix are informative only.

### Default Abstract Syntax for mOSI

The following OBJECT IDENTIFIER have been defined in CULR part 3:

```
{iso(1) standard(0) curl(11188) mosi(3) default-abstract-syntax(1) version(1)}
```

This object identifier can be used as the abstract syntax when the application protocol (above ACSE) can be treated as single presentation data values (PDVs). Each PDV is a sequence of consecutive octets without regard for semantic or other boundaries. The object identifier may also be used when, for pragmatic reasons, the actual abstract syntax of the application is not identified in Presentation Layer negotiation.

#### Notes:

1. Applications specified using ASN.1 should not use the default abstract syntax.
2. As this object identifier is used by all applications using the default abstract syntax for mOSI, it cannot be used to differentiate between applications. One of the ACSE parameters; for example, AE Title or Presentation address, may be used to differentiate between applications.

### Default Transfer Syntax for mOSI

If the default transfer syntax and the abstract syntax are identical, the OBJECT IDENTIFIER for the default abstract syntax is used. If they are not identical, the OBJECT identifier for the default transfer syntax is:

```
{iso(1) standard(0) curl(11188) mosi(3) default-transfer-syntax(2) version(1)}
```

**Note:** In the presentation data value of the PDV list of Presentation Protocol or in the encoding of User Information of ACSE Protocol, only *octet-aligned* or *arbitrary* can be used for default transfer syntax for mOSI. *Single-ASN1-type* cannot be used for default transfer syntax for mOSI.

**Default Application Context for mOSI**

The following OBJECT IDENTIFIER has been defined in Common Upper Layer Requirements (CULR) part 3:

```
{iso(1) standard(0) curl(11188) mosi(3) default-application-context(3) version(1)}
```

This application context supports the execution of any application using the default abstract syntax for mOSI.

**Reason Codes for Disconnections**

```
#define T_AC_U_AARE__NONE 0x0001 /* connection rejected by */
/* peer user: no reason given */
#define T_AC_U_AARE_ACN 0x0002 /* connection rejected: */
/* application context name */
/* not supported */
#define T_AC_U_AARE_APT 0x0003 /* connection rejected: */
/* AP title not recognised */
#define T_AC_U_AARE_AEQ 0x0005 /* connection rejected: */
/* AE qualifier not recognised */
#define T_AC_U_AARE_PEER_AUTH 0x000e /* connection rejected: */
/* authentication required */
#define T_AC_P_ABRT_NSPEC 0x0011 /* aborted by peer provider: */
/* no reason given */
#define T_AC_P_AARE_VERSION 0x0012 /* connection rejected: */
/* no common version */
```

Other reason codes may be specified as implementation defined constants. In order to be portable, an application should not interpret such information, which should only be used for troubleshooting purposes.

## H.6 <xti\_mosi.h> Header File

This section presents the additional header file information for XTI-mOSI.

Implementations supporting XTI-mOSI will provide equivalent definitions in <xti\_mosi.h>. XTI-mOSI programs should include <xti\_mosi.h> as well as <xti.h>.

```

/* mosi address structure */

struct t_mosiaddr {
 t_uscalar_t flags;
 t_scalar_t osi_ap_inv_id;
 t_scalar_t osi_ae_inv_id;
 unsigned int osi_apt_len;
 unsigned int osi_aeq_len;
 unsigned int osi_paddr_len;
 unsigned char osi_addr[MAX_ADDR];
};

#define T_ISO_APCO 0x0200
#define T_ISO_APCL 0x0300
#define T_AP_CNTX_NAME 0x1
#define T_AP_PCL 0x2

#define T_OPT_VALLEN(opt) (opt->len - sizeof(struct t_opthdr)).

/* presentation context definition and result list option */

struct t_ap_pco_el {
 t_scalar_t count;
 t_scalar_t offset;
}

/* presentation context item header */

struct t_ap_pc_item {
 t_scalar_t pci; /* unique odd integer */
 t_scalar_t res; /* result of negotiation */
}

/* presentation context item element */

struct t_ap_syn_off {
 t_scalar_t size; /* length of syntax object identifier contents */
 t_scalar_t offset; /* offset of object identifier for the syntax */
}

/* values for res of a presentation context item */

#define T_PCL_ACCEPT 0x0000 /* pres. context accepted */
#define T_PCL_USER_REJ 0x0100 /* pres. context rejected */
/* by peer application */
#define T_PCL_PREJ_RSN_NSPEC 0x0200 /* prov. reject: */
/* no reason specified */
#define T_PCL_PREJ_A_SYTX_NSUP 0x0201 /* prov. reject: abstract */
/* syntax not supported */
#define T_PCL_PREJ_T_SYTX_NSUP 0x0202 /* prov. reject: transfer */
/* syntax not supported */

```

```
#define T_PCL_PREJ_LMT_DCS_EXCEED 0x0203 /* prov. reject: local */
 /* limit on DCS exceeded */

/* Reason Codes for Disconnections */

#define T_AC_U_AARE__NONE 0x0001 /*connection rejected by */
 /*peer user: no reason given */
#define T_AT_C_U_AARE_ACN 0x0002 /*connection rejected: */
 /*application context name */
 /*not supported */
#define T_AC_U_AARE_APT 0x0003 /*connection rejected: */
 /*AP title not recognised */
#define T_AC_U_AARE_AEQ 0x0005 /*connection rejected: */
 /*AE qualifier not recognised */
#define T_AC_U_AARE_PEER_AUTH 0x000e /*connection rejected: */
 /*authentication required */
#define T_AC_P_ABRT_NSPEC 0x0011 /*aborted by peer provider: */
 /*no reason given */
#define T_AC_P_AARE_VERSION 0x0012 /*connection rejected: */
 /*no common version */
```

# SNA Transport Provider

## I.1 Introduction

This Appendix includes:

- Protocol-specific information that is relevant for Systems Network Architecture (SNA) transport providers.

It assumes native SNA users, that is, those prepared to use SNA addresses and other SNA transport characteristics (for example, mode name for specifying quality of service).

- Information on the mapping of XTI functions to Full Duplex (FDX) LU 6.2.

Systems that do not support LU 6.2 full duplex can simulate them using twin-opposed half-duplex conversations. Protocols for doing so will be published separately.

The half-duplex verbs have been published for several years. The full duplex verbs are documented in the CPI-C Reference, Version 2.1<sup>16</sup>.

This Appendix also defines data structures and constants required for NetBIOS transport providers which are exposed through `<xti_sna.h>` header file.

Applications written to compilation environments earlier than those required by this issue of the specification (see Appendix F on page 281) and defining `_XOPEN_SOURCE` to be less than 500, will have these data structures and constants exposed through the inclusion of `<xti.h>` on platforms where XTI interface use is supported for SNA transports.

---

16. CPI-C Reference, Version 2.1; publicly available from IBM branch offices under the IBM reference SC26-4399.

## I.2 SNA Transport Protocol Information

This section describes the protocol-specific information that is relevant for Systems Network Architecture (SNA) transport providers.

### I.2.1 General

1. Protocol address

For information about SNA addresses, see Section I.2.2 on page 321.

2. Connection establishment

Native SNA has no confirmed allocation protocol for full duplex conversations. When a conversation is allocated, the connection message is buffered and sent with the first data that is sent on the conversation. When the `t_connect()` or `t_rcvconnect()` function completes, connectivity has been established to the partner node, but not to the partner program. Since notification that the partner is not available may occur later, the disconnection reasons returned on `t_rcvdis()` include `[T_SNA_CONNECTION_SETUP_FAILURE]`, indicating that the connection establishment never completed successfully.

An SNA program that needs to know that the partner is up and running before it proceeds sending data must have its own user-level protocol to determine if this is so.

3. Parallel connections

LU 6.2 allows multiple, simultaneous connections between the same pair of addresses. The number of connections possible between two systems depends on limits defined by system administrators.

4. Sending data of zero octets is supported.

5. Expedited data

In connection-oriented mode, expedited data transfer can be negotiated by the two transport providers during connection establishment. Expedited data transfer is supported if both transport providers support it. However negotiation between transport users is not supported. Therefore the expedited option is read-only.

6. `t_close()`

The semantics of `t_close()` on an SNA Transport Provider is simplex orderly, that is, the send pipe of the XTI application issuing the `t_close()` is closed, but the receive pipe remains open. Any data sent prior to the `t_close()` will be delivered to the partner.

7. SNA buffers data from multiple `t_snd()` functions until the SNA send buffer is full, allowing multiple records to be sent in one transmission. However, users sometimes have reasons for ensuring that a record is sent immediately. By setting the `T_PUSH` flag on the `t_snd()` function, the transport user causes data to be transferred without waiting for the buffer to be filled.

In order to take advantage of the performance improvement that SNA buffering offers, the XTI user must set the `T_SNA_ALWAYS_PUSH` option to `T_NO` (default is `T_YES`). If this option is not set to `T_NO`, a push will be done for every `t_snd()` and the `T_PUSH` flag will have no effect.

8. Programs migrated to SNA from other transport providers may want every `t_snd()` to cause a message to be sent immediately in order to match behaviour on the original provider. The default of this option is `T_YES`; thus the default is that a `t_snd()` will always be sent out immediately.

## I.2.2 SNA Addresses

In an SNA environment, the protocol address always includes a network-ID-qualified logical unit (LU) name. This is the address of the node where the program resides.

For the `t_connect()` and `t_sndudata()` functions, the address also contains a transaction program name (TPN), identifying the program addressed in the partner node. A file descriptor used to accept incoming connection requests should have a complete SNA name, including TPN, bound to it with `t_bind()`.

A file descriptor used for outgoing connection requests may optionally have only a network-id-qualified LU name bound to it.

Since the `t_listen()` returns only the LU name part of the address, this address is not adequate for opening up a connection back to the source. The transport user must know the TPN of its partner by some mechanism other than XTI services.

However, `t_rcvudata()` returns the complete address of the partner that can be used to send a datagram back to it.

An SNA address has the following structure. When the TPN is not included, the TPN length (`sna_tpn_length`) is set to zero, and the string that follows is null.

```

/* The definitions for maximum LU name and netid lengths have specific */
/* values because these maxima are a fixed SNA characteristic, */
/* not an implementation option. Maximum TP length is a implementation */
/* option, although the maximum maximum is 64. */
/* */

#define T_SNA_MAX_NETID_LEN 8
#define SNA_MAX_NETID_LEN 8 (LEGACY)
#define T_SNA_MAX_LU_LEN 8
#define SNA_MAX_LU_LEN 8 (LEGACY)
#define T_SNA_MAX_TPN_LEN
#define SNA_MAX_TPN_LEN (LEGACY)

struct sna_addr{
 u_char sna_netid (T_SNA_MAX_NETID_LEN),
 u_char sna_lu (T_SNA_MAX_LU_LEN),
 u_short sna_tpn_len, /* less than or equal to T_SNA_MAX_TPN_LEN */
 u_char sna_tpn (sna_tpn_len)
}

```

### Notes:

1. network-identifier (`sna_netid`): The address can contain either an SNA network identifier or the defined value, `SYS_NET`, which indicates that the predefined network identifier associated with the local system should be used.
2. IBM Corporation provides a registration facility for SNA network identifiers to guarantee global uniqueness. (See IBM document G325-6025-0, SNA Network Registry).
3. LU name (`sna_lu`): The address can contain either a specific LU name or the defined value, `SYS_LU`, which indicates that the system default LU name is to be used.
4. LU name and network identifier fields are fixed length. For values shorter than 8 characters, they are blank filled to the right.
5. Transaction program name (`sna_tpn`): This field can take one of three values:

- **Null value:** No transaction program name is to be associated with the file descriptor.

This is adequate for file descriptors used for outgoing connection requests.

If no transaction program is associated with a file descriptor when a *t\_listen()*, *t\_rcvudata()*, or *t\_sndudata()* is issued, the function will return a TPROTO error.

- **Specified value:** A value that will be known by a partner program; for example, a well-known transaction program name used by a server.
  - **Defined value, DYNAMIC\_TPN:** An indication that the system should generate a TP name for the file descriptor.
6. The values SYS\_NET, SYS\_LU and DYNAMIC\_TPN may not be used as real values of the sna\_netid, sna\_lu or sna\_tpn fields, respectively.

### I.2.3 Options

Options are formatted according to the structure **t\_opthdr** as described in Chapter 6. A transport provider compliant to this specification supports none, all, or any subset of the options defined in Section I.2.3.1.

#### I.2.3.1 Connection-Mode Service Options

The protocol level of all subsequent options is SNA.

All options have end-to-end significance. Some may be negotiated in the XTI states T\_IDLE and T\_INCON, and all are read-only in all other states except T\_UNINIT.



**Options for Service Quality and Expedited Data**

Table I-1 shows the SNA options that affect the quality of a connection and the transport service level provided.

| Option Name       | Type of Option Value | Legal Option Value                                                                                               | Meaning                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------|----------------------|------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| T_SNA_MODE        | char                 | T_SNA_BATCH<br>T_SNA_BATCHSC<br>T_SNA_INTER<br>T_SNA_INTERSC<br>T_SNA_DEFAULT<br>any user-defined SNA mode value | SNA mode, which controls the underlying class of service selected for the connection. The SNA mode is specified only by the active side of the connection.<br><br>If not specified, the default mode is T_SNA_DEFAULT.<br><br>The default mode characteristics may vary from system to system.                                                                                                                                                                                                                                                                                                                         |
| T_SNA_ALWAYS_PUSH | t_uscalar_t          | T_YES / T_NO                                                                                                     | If T_YES, every <i>t_snd()</i> operation will cause the message to be sent immediately.<br><br>If T_NO, the data from a <i>t_snd()</i> operation may be buffered and sent later. The transport user can set the T_PUSH flag on a <i>t_snd()</i> function call to cause the data to be sent immediately.<br><br>Default value is T_NO.<br><br>This option is primarily for programs migrated to SNA from other protocol stacks that always send data immediately. It allows them to request behaviour similar to that on the original provider. However, setting T_SNA_ALWAYS_PUSH to T_YES may affect its performance. |

**Table I-1** SNA Options

## I.2.4 Functions

- t\_accept()* Since user data is not exchanged during connection establishment, the parameter *call→udata.len* must be 0.
- t\_bind()* The *addr* field of the *t\_bind* structure represents the local network-id-qualified LU name of the local logical unit and the transaction program name of the program issuing the *t\_bind()* function.
- If the endpoint was bound in the passive mode (that is, *qlen* > 0) and the requested address has a null transaction program subfield, the function completes with the T\_BADADDR error.
- t\_connect()* The *sndcall→addr* specifies the network-ID-qualified LU name and transaction program name of the remote connection partner.
- An SNA transport provider allows more than one connection between the same address pair.
- Since user data cannot be exchanged during the connection establishment phase, *sndcall→udata.len* must be set to 0. On return, *rcvcall→udata.maxlen* should be set to 0.
- t\_getinfo()* In all states except T\_DATAXFER, the function *t\_getinfo()* returns in the parameter *info* the same information that was returned by *t\_open()*. In T\_DATAXFER state, however, the information returned may differ from that returned by *t\_open()*, depending on whether the remote transport provider supports expedited data transfer. The fields of *info* are set as defined in the table below.

| Parameters           | Before Call | After Call                       |
|----------------------|-------------|----------------------------------|
| <i>fd</i>            | <b>x</b>    | /                                |
| <i>info→addr</i>     | /           | 82                               |
| <i>info→options</i>  | /           | <b>x</b> <sup>1</sup>            |
| <i>info→tsdu</i>     | /           | T_INFINITE (-1)                  |
| <i>info→etsdu</i>    | /           | T_INVALID (-2) / 86 <sup>2</sup> |
| <i>info→connect</i>  | /           | T_INVALID (-2)                   |
| <i>info→discon</i>   | /           | T_INVALID (-2)                   |
| <i>info→servtype</i> | /           | T_COTS_ORD                       |
| <i>info→flags</i>    | /           | T_SNDZERO                        |

**Table I-2** Fields for *info* Parameter

**Notes:**

1. **x** means an integral number greater than zero.
2. Depending on the negotiation of expedited data transfer.

- t\_getprotaddr()* The *boundaddr* value includes the transaction program name of the local program.
- The *peeraddr* value (if any) includes only the network-ID-qualified LU name of the partner.
- t\_listen()* The *call→addr* structure contains the network-ID-qualified LU name of the remote partner.

*t\_open()* The default characteristics returned by *t\_open()* are shown in the table below.

| Parameters    | Before Call | After Call                       |
|---------------|-------------|----------------------------------|
| name          | x           | /                                |
| oflag         | x           | /                                |
| info→addr     | /           | 82                               |
| info→options  | /           | x <sup>1</sup>                   |
| info→tsdu     | /           | T_INFINITE (-1)                  |
| info→etsdu    | /           | T_INVALID (-2) / 86 <sup>2</sup> |
| info→connect  | /           | T_INVALID (-2)                   |
| info→discon   | /           | T_INVALID (-2)                   |
| info→servtype | /           | T_COTS_ORD                       |
| info→flags    | /           | T_SNDZERO                        |

**Table I-3** Default Characteristics returned by *t\_open()*

**Notes:**

1. x means an integral number greater than 0.
2. Expedited data may or may not be supported by the local transport provider.

*t\_rcv()* If expedited data arrives after part of a TSDU (logical record) has been retrieved, receipt of the remainder of the TSDU will be suspended until after the ETSDU has been processed. Only after the full ETSDU has been retrieved (T\_MORE not set), will the remainder of the TSDU be available to the user.

*t\_rcvconnect()* Since no user data can be returned on *t\_rcvconnect()*, the *call→udata.len* should be set to 0 before the function is invoked.

*t\_rcvdis()* Since user data is not sent during disconnection, the value *discon→udata.len* should be set to 0 before *t\_rcvdis()* is called.

The following disconnection reason codes are valid for any implementation of an SNA transport provider under XTI:

```
#define T_SNA_CONNECTION_SETUP_FAILURE.
#define SNA_CONNECTION_SETUP_FAILURE. (LEGACY)
#define T_SNA_USER_DISCONNECT
#define SNA_USER_DISCONNECT (LEGACY)
#define T_SNA_SYSTEM_DISCONNECT
#define SNA_SYSTEM_DISCONNECT (LEGACY)
#define T_SNA_TIMEOUT
#define SNA_TIMEOUT (LEGACY)
#define T_SNA_CONNECTION_OUTAGE
#define SNA_CONNECTION_OUTAGE (LEGACY)
```

These definitions are exposed by the inclusion of `<xti_sna.h>`.

*t\_snd()* Unless the T\_SNA\_ALWAYS\_PUSH option is set to T\_YES or the T\_PUSH flag on the *t\_snd()* function is set, the SNA transport provider may collect data in a send buffer until it accumulates a sufficient amount for transmission. The amount of data that is accumulated can vary from one connection to another.

In order to take advantage of the performance improvement that SNA buffering offers, the XTI user must set the T\_SNA\_ALWAYS\_PUSH option to T\_NO (Default is T\_YES). If this option is not set to T\_NO, a push will be done for every *t\_snd()* and the T\_PUSH flag will have no effect.

*t\_snddis()* Since no user data is sent during a disconnection operation, *call→udata.len* should be set to 0 before the call to *t\_snddis()*.

*t\_sndudata()* The *unitdata→addr* field contains the full SNA address, including network-id-qualified LU name and transaction program identifier, of the remote partner.

If address associated with the file descriptor has a null transaction program name subfield, the function completes with the TPROTO error.

The *unitdata→opt* structure may contain an SNA mode governing the transmission of the data. For example, the program may confine data transmission to secure lines by selecting the T\_SNA\_INTERSC or T\_SNA\_BATCHSC modes.

### **I.3 Mapping XTI to SNA Transport Provider**

This section presents the mapping of XTI functions to Full Duplex (FDX) LU 6.2.

First, several flow diagrams are given to illustrate the function of the XTI Mapper. Following this, mapping tables are given that show the FDX LU 6.2 verbs that the XTI Mapper needs to generate for each XTI function. For each XTI function that maps to a FDX verb, an additional table is referenced that gives the mappings of each parameter. Finally, a table shows mapping of LU 6.2 FDX return codes to XTI events.

The use of FDX LU 6.2 verbs in this section is for illustrative purposes only and is analogous to OSI's use of service primitives, that is, as a way to explain the semantics provided by the protocol. The FDX LU6.2 verbs are used only to help in understanding the SNA protocol, and are not a required part of an implementation.

### I.3.1 General Guidelines

General guidelines for mapping XTI to an SNA transport provider are listed below:

- In the following flow diagrams, notice that the XTI mapper always has a RECEIVE\_AND\_WAIT posted. This is done so that when data comes into the SNA transport provider, the XTI mapper is able to set an event indicator. Then, when a T\_LOOK is issued, the XTI application can be informed that there is data to be received.
- The XTI mapper keeps a table that maps an XTI fd to a RESOURCE(variable) on the FDX verbs.
- In this section, we assume the XTI mapper will be using the FDX LU 6.2 basic conversation verb interface. The following table Table I-4 gives an explanation for each FDX verb that is used in these mappings.

**Table I-4** FDX LU 6.2 Verb Definitions

| FDX Verb               | Description                                                                                                                                                                                                                                                                      |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ALLOCATE               | Allocates a conversation between the local transaction program and a remote (partner) transaction program.                                                                                                                                                                       |
| DEALLOCATE             | DEALLOCATE with TYPE(FLUSH) closes the local program's send queue. Both the local and remote program must close their send queues independently.<br><br>DEALLOCATE with TYPE(ABEND_PROG) is an abrupt termination that will close both sides of the conversation simultaneously. |
| FLUSH                  | Flushes the local LU's send buffer.                                                                                                                                                                                                                                              |
| GET_ATTRIBUTES         | Returns information pertaining to the specified conversation.                                                                                                                                                                                                                    |
| GET_TP_PROPERTIES      | Returns information pertaining to the transaction program issuing the verb.                                                                                                                                                                                                      |
| RECEIVE_ALLOCATE       | Receives a new conversation with a partner transaction program that issued ALLOCATE.                                                                                                                                                                                             |
| RECEIVE_AND_WAIT       | Waits for data to arrive on the specified conversation and then receives the data. If data is already available, the program receives it without waiting.                                                                                                                        |
| RECEIVE_EXPEDITED_DATA | Receives data sent by the remote transaction program in an expedited manner, via the SEND_EXPEDITED_DATA verb.                                                                                                                                                                   |
| SEND_DATA              | Sends data to the remote transaction program.                                                                                                                                                                                                                                    |
| SEND_EXPEDITED_DATA    | Sends data to the remote transaction program in an expedited manner. This means that it may arrive at the remote transaction program before data sent earlier via a send queue verb; for example, SEND_DATA.                                                                     |
| WAIT_FOR_COMPLETION    | Waits for posting to occur on one or more non-blocking operations represented in the specified wait objects. Posting of a non-blocking operation occurs when the LU has completed the associated non-blocking verb and filled all the return values.                             |

### I.3.2 Flows Illustrating Full Duplex Mapping

The following diagrams show mappings from the XTI function calls for active connection establishment to SNA verb sequences. The first Figure I-1 is used for blocking XTI calls; the second Figure I-2 is used for non-blocking calls.

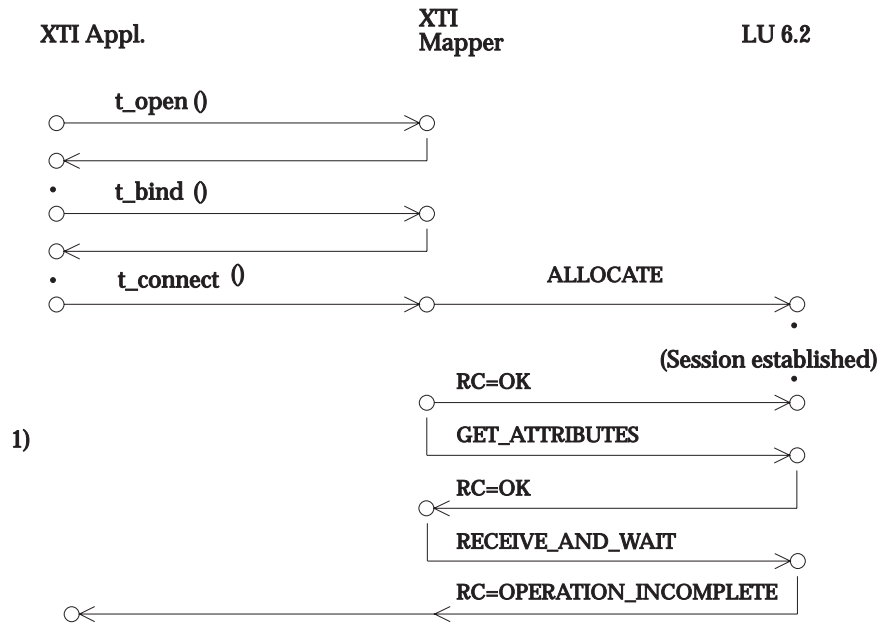


Figure I-1 Active Connection Establishment, Blocking Version (1 of 2)

#### Annotations

1. `GET_ATTRIBUTES` is issued after the session is established and before the return for the `t_connect`. This is only done if the mode name, or partner LU name are required on the return to `t_connect`. This would be indicated by a non-zero value in either the `rcvcall→addr.buf` or `eercvcall→opt.buf` fields on `t_connect`.

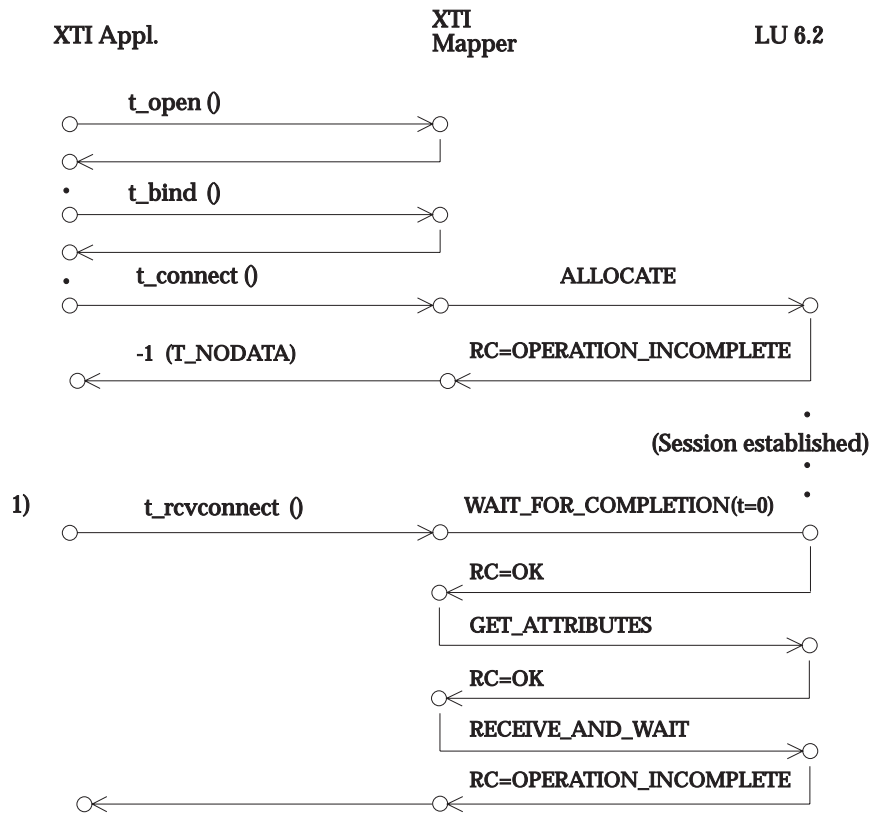


Figure I-2 Active Connection Establishment, Non-blocking Version (2 of 2)

**Annotations**

1. The XTI application will issue a *t\_rcvconnect* as a poll to see if the *t\_connect* has completed. The *t\_rcvconnect* will cause a *WAIT\_FOR\_COMPLETION*, with time=0, to be issued. The *WAIT\_FOR\_COMPLETION* will check on the wait object from the previous non-blocking *ALLOCATE*.

When the *t\_connect* has completed successfully a *GET\_ATTRIBUTES* is issued if the mode name, or partner LU name are required on the return of the *t\_rcvconnect*.

After the *GET\_ATTRIBUTES*, a non-blocking *RECEIVE\_AND\_WAIT* is issued to post a receive for any incoming data.



The next three diagrams show possible mappings of SNA Attach processing for an incoming connection to the series of XTI calls on the passive side of a connection.

The first Figure I-3 uses the native SNA instantiation mechanism; that is, programs are instantiated when the connection request arrives. This requires that the TP name (that is, the XTI application name) is known as part of the LU definition. This is **before** the *t\_bind* is issued.

The second Figure I-4 is a blocking use of the interface, where the SNA transport provider allows a connection request to be received by an existing program. This model, although not described in the architecture, is supported by many SNA products.

The third Figure I-5 is a non-blocking use of the interface, where the SNA transport provider allows a connection request to be received by an existing program. This model is described as part of the FDX architecture.

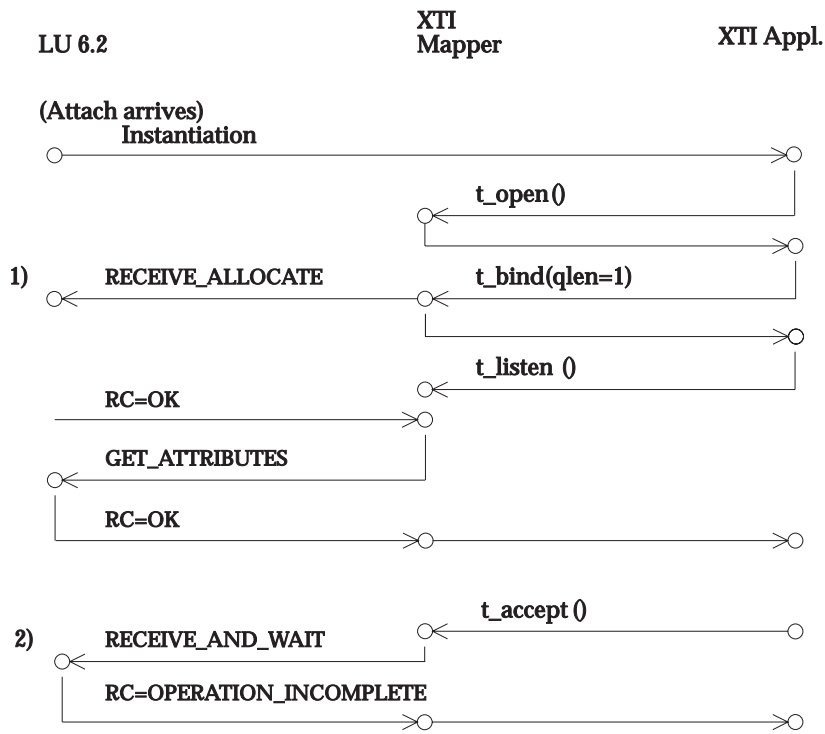


Figure I-3 Passive Connection Establishment, Instantiation Version (1 of 3)

**Annotations**

1. If *qlen* in *t\_bind* is > 0, a RECEIVE\_ALLOCATE will be issued for each connection request that can be queued. When the RECEIVE\_ALLOCATE completes successfully a GET\_ATTRIBUTES is issued only if the mode name, or partner LU name are required on the return to *t\_listen*.
2. The *t\_accept* will cause a RECEIVE\_AND\_WAIT to be issued. The RECEIVE\_AND\_WAIT is issued to post a receive for any incoming data.

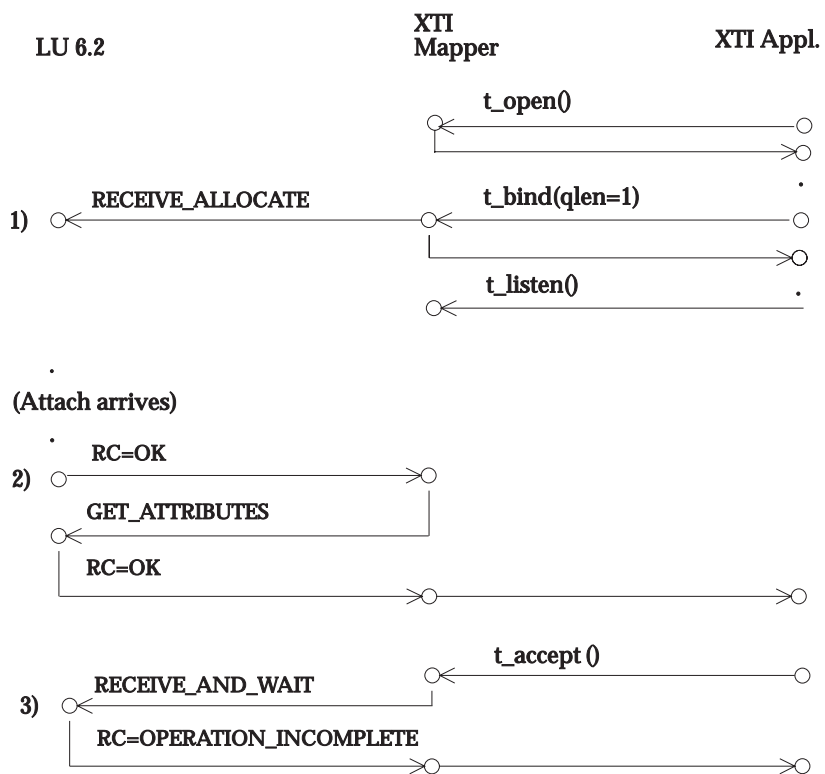


Figure I-4 Passive Connection Establishment, Blocking Version (2 of 3)

**Annotations**

1. The *t\_bind* will cause a blocking RECEIVE\_ALLOCATE to be issued for each connection request that can be queued.
2. When the RECEIVE\_ALLOCATE completes successfully a GET\_ATTRIBUTES is issued only if the mode name, or partner LU name are required on the return to *t\_listen*.
3. The *t\_accept* will cause a RECEIVE\_AND\_WAIT to be issued. The RECEIVE\_AND\_WAIT is issued to post a receive for any incoming data.

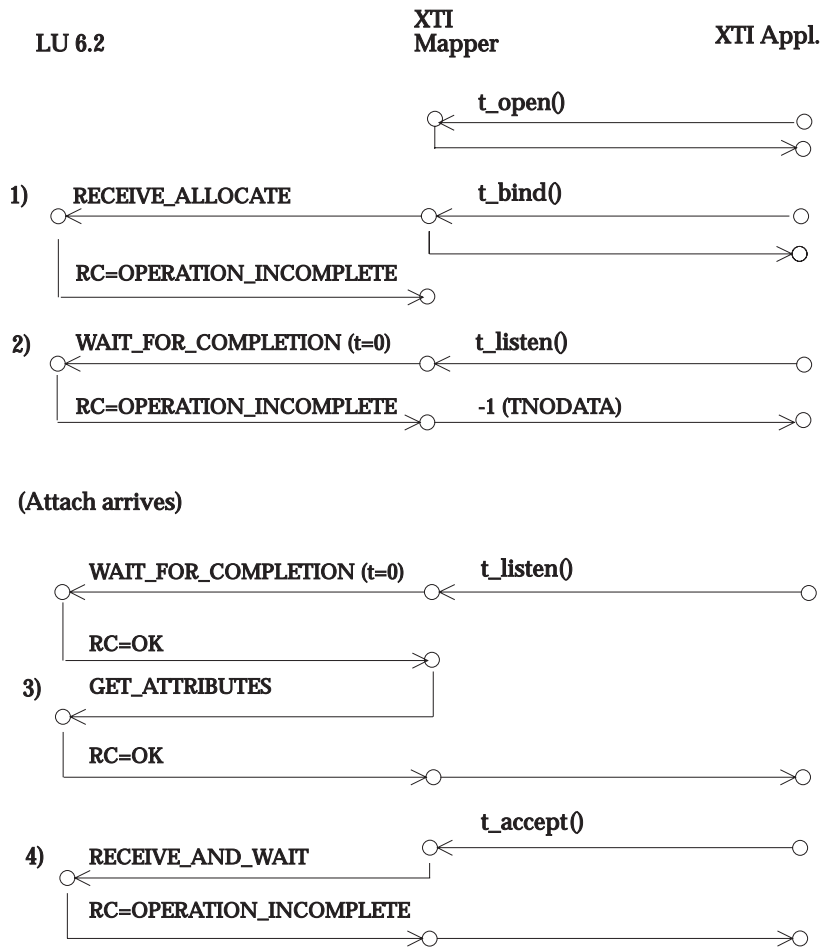
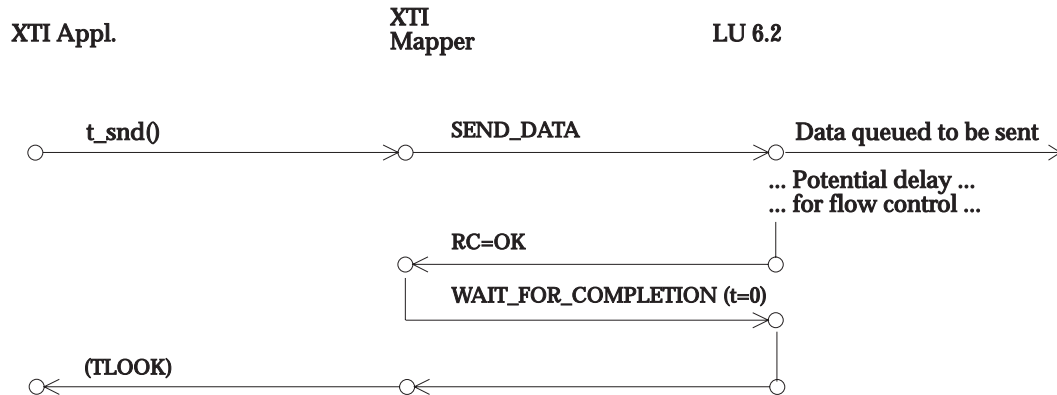


Figure I-5 Passive Connection Establishment, Non-blocking Version (3 of 3)

**Annotations**

1. The *t\_bind* will cause a non-blocking RECEIVE\_ALLOCATE to be issued for each connection request that can be queued.
2. A *t\_listen* is used as a poll to see if a connection request has been received. The *t\_listen* will cause a WAIT\_FOR\_COMPLETION, with time=0, to be issued. The WAIT\_FOR\_COMPLETION will check on the wait object from the previous non-blocking RECEIVE\_ALLOCATE. In this example, when the first *t\_listen* is issued, the RECEIVE\_ALLOCATE is still outstanding; but the RECEIVE\_ALLOCATE has completed before the second *t\_listen* is issued.
3. When the WAIT\_FOR\_COMPLETION indicates that the RECEIVE\_ALLOCATE has completed successfully, a GET\_ATTRIBUTES is issued only if the mode name, or partner LU name are required on the return to *t\_listen*.
4. The *t\_accept* will cause a RECEIVE\_AND\_WAIT to be issued. The RECEIVE\_AND\_WAIT is issued to post a receive for any incoming data.

The next diagram, Figure I-6, shows the mapping for the blocking XTI *t\_snd()* call. The diagram after this, Figure I-7, shows the non-blocking mapping of the XTI *t\_snd()* call.



**Figure I-6** XTI Function to LU 6.2 Verb Mapping: Blocking *t\_snd*

1. The blocking *t\_snd* will cause a blocking SEND\_DATA to be issued. This will block until the LU accepts and queues all the data being sent.

If EXPEDITED=YES, the mapper will issue a SEND\_EXPEDITED\_DATA verb rather than the SEND\_DATA.

2. When the SEND\_DATA returns, a WAIT\_FOR\_COMPLETION, with time=0, is issued to see if the wait object for any outstanding non-blocking LU 6.2 verbs have been posted. At a minimum, there will be an outstanding RECEIVE\_AND\_WAIT, waiting for any incoming data, that needs to be checked. If any wait objects have been posted, the return code on the *t\_snd* is set to TLOOK. This will inform the XTI application to issue a *t\_look* to see what has been posted.

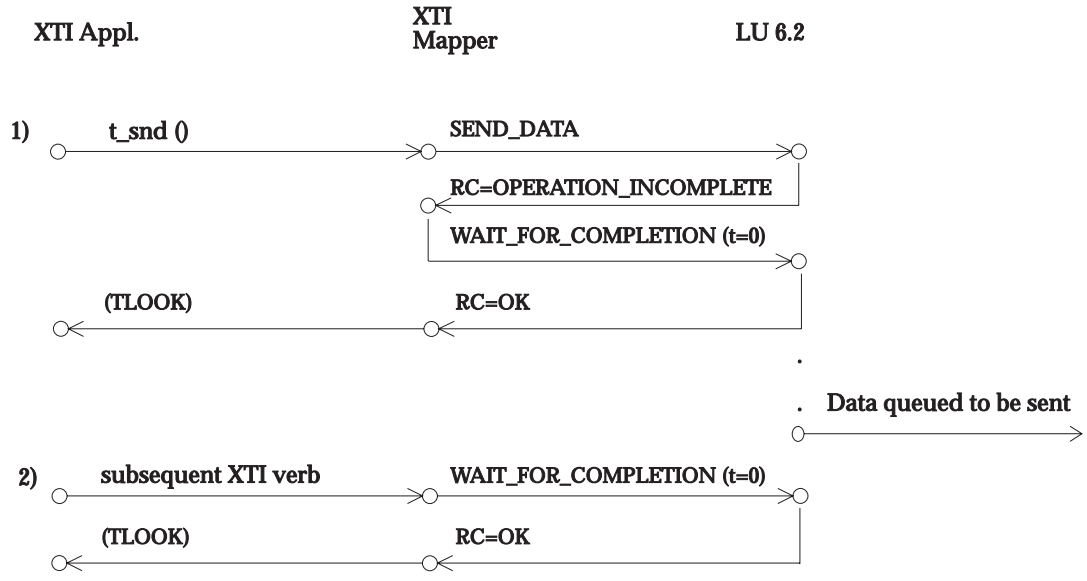


Figure I-7 XTI Function to LU 6.2 Verb Mapping: Non-blocking t\_snd

1. The XTI mapper needs to either accept all the data being sent, or none of it. In this case, all the data is accepted, thus the non-blocking *t\_snd* causes a non-blocking SEND\_DATA to be issued.

The XTI mapper then needs to issue a WAIT\_FOR\_COMPLETION to see if any other blocking LU 6.2 verbs have completed. This case is not shown in this diagram.

If EXPEDITED=YES, the mapper will issue a SEND\_EXPEDITED\_DATA verb rather than the SEND\_DATA.

2. When a subsequent XTI verb is issued (for example, *t\_rcv* or *t\_send*), a WAIT\_FOR\_COMPLETION, with time=0, is issued to see if the wait object for any outstanding non-blocking LU 6.2 verbs have been posted. In this case, one of the wait objects will be the one associated with the non-blocking SEND\_DATA. If any wait objects have been posted, the return code on the *t\_snd* is set to TLOOK. This will inform the XTI application to issue a *t\_look* to see what has been posted.

There may be an additional LU 6.2 verb issued due to the *subsequent XTI* verb that was issued. This is not shown in the above diagram.

The next diagram, Figure I-8, shows the mapping for blocking XTI receive call, and the diagram after this, Figure I-9, shows the mapping for non-blocking XTI receive call.

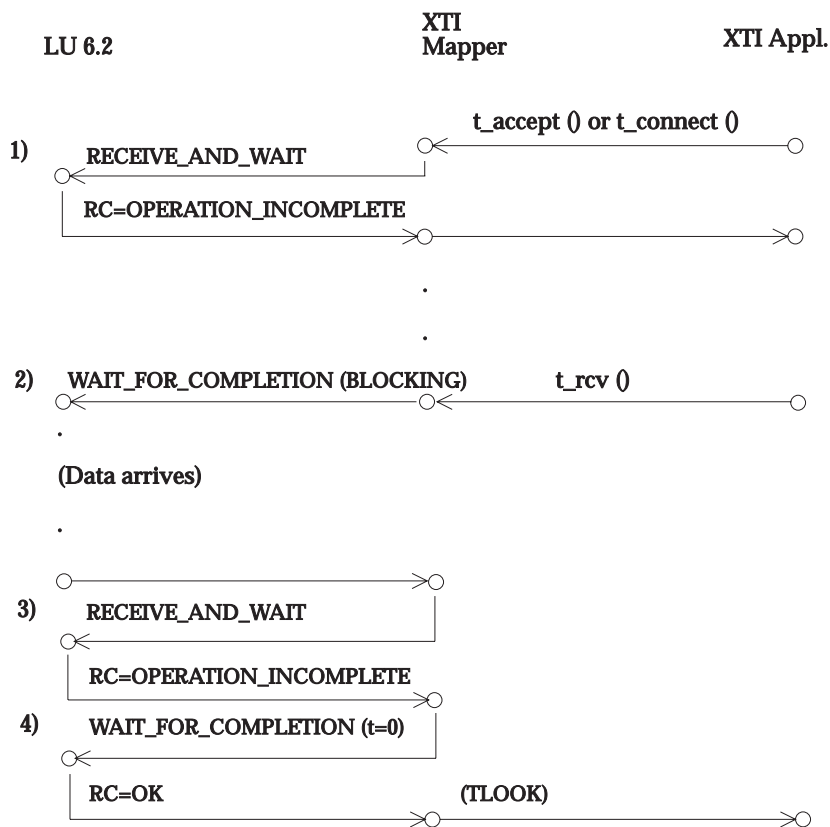


Figure I-8 XTI Function to LU 6.2 Verb Mapping: Blocking t\_rcv

1. There is always an outstanding non-blocking RECEIVE\_AND\_WAIT, this is true whether the XTI application is using blocking or non-blocking mode. This is to post a receive for any incoming data.  
 In this diagram, the outstanding RECEIVE\_AND\_WAIT was issued when the connection was setup. This could be as a result of either a t\_accept or t\_connect.
2. When the XTI issues a blocking t\_rcv, the XTI mapper will issue a blocking WAIT\_FOR\_COMPLETION to wait on the wait object associated with the outstanding RECEIVE\_AND\_WAIT. This will block until data is received on this connection.
3. When data is received, the mapper needs to issue a non-blocking RECEIVE\_AND\_WAIT to replace the one that just completed.
4. Issue a WAIT\_FOR\_COMPLETION, with time=0, to see if any other outstanding wait objects have been posted.

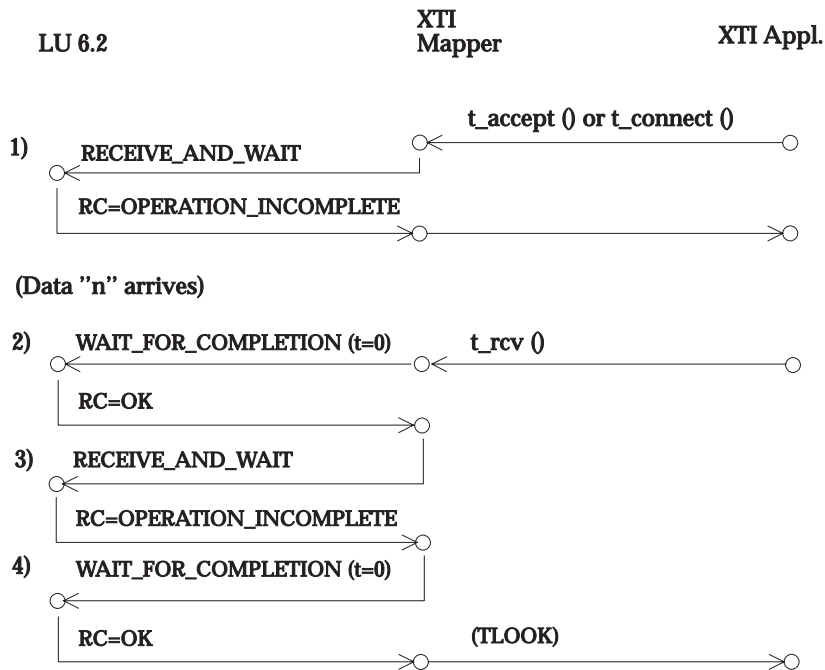


Figure I-9 Mapping from XTI Calls to LU 6.2 Verbs (Passive side)

1. There is always an outstanding non-blocking RECEIVE\_AND\_WAIT, this is true whether the XTI application is using blocking or non-blocking mode. This is to post a receive for any incoming data.

In this diagram, the outstanding RECEIVE\_AND\_WAIT was issued when the connection was setup. This could be as a result of either a *t\_accept* or *t\_connect*.

2. When the XTI application issues a non-blocking *t\_rcv*, the XTI mapper will issue a WAIT\_FOR\_COMPLETION, with T=0, to see if the wait object for the outstanding RECEIVE\_AND\_WAIT has been posted. When the wait object has been posted, the XTI mapper needs to pass the data to the XTI application buffer.

It is possible that the amount of incoming data in the XTI mapper buffer is more than the XTI application stated on the *t\_rcv*. In this case, the XTI Mapper will set the TMORE flag. Then, when the next *t\_rcv* is issued, the remaining data will be passes to the XTI application **BEFORE** issuing the WAIT\_FOR\_COMPLETION to check the wait object on the outstanding RECEIVE\_AND\_WAIT.

3. When data is received, the mapper needs to issue a non-blocking RECEIVE\_AND\_WAIT to replace the one that just completed.
4. Issue a WAIT\_FOR\_COMPLETION, with time=0, to see if any other outstanding wait objects have been posted.

### I.3.3 Full Duplex Mapping

The following table shows the mapping from XTI function calls to full duplex LU 6.2 verbs.

**Table I-5 XTI Mapping to LU 6.2 Full Duplex Verbs**

| <b>XTI Function</b>    | <b>SNA FDX LU6.2 verb</b>                                                                                                                                                                                                                                                                                                                | <b>Comments</b>                                                                                                                                                        |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>t_accept()</i>      | RECEIVE_AND_WAIT                                                                                                                                                                                                                                                                                                                         | User data is not exchanged during connection establishment.<br>Refer to Table I-7 on page 340.                                                                         |
| <i>t_alloc()</i>       | Local                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                        |
| <i>t_bind()</i>        | If qlen>0: RECEIVE_ALLOCATE for each connection request that can be queued.<br><br>Optionally: DEFINE_TP<br><br>With the instantiation model, the TP name (that is, XTI application name) must be known by the LU before the TP can be instantiated. This is prior to the <i>t_bind</i> being issued. (Refer to Figure I-3 on page 331.) | Refer to Table I-8 on page 341.                                                                                                                                        |
| <i>t_close()</i>       | If connection still up issue DEALLOCATE TYPE(FLUSH)                                                                                                                                                                                                                                                                                      | Refer to Table I-9 on page 341.                                                                                                                                        |
| <i>t_connect()</i>     | ALLOCATE RETURN_CONTROL<br>GET_ATTRIBUTES<br>RECEIVE_AND_WAIT                                                                                                                                                                                                                                                                            | Refer to Table I-10 on page 342.                                                                                                                                       |
| <i>t_error()</i>       | Local                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                        |
| <i>t_free()</i>        | Local                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                        |
| <i>t_getinfo()</i>     | Local                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                        |
| <i>t_getprotaddr()</i> | GET_TP_PROPERTIES to get OWN_FULLY_QUALIFIED_LU_NAME and OWN_TP_NAME<br><br>GET_ATTRIBUTES to get PARTNER_FULLY_QUALIFIED_LU_NAME                                                                                                                                                                                                        | The partner's TP name must be learned by some mechanism other than XTI services. In connectionless mode, there is no partner name.<br>Refer to Table I-11 on page 344. |
| <i>t_getstate()</i>    | Local                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                        |
| <i>t_listen()</i>      | WAIT_FOR_COMPLETION                                                                                                                                                                                                                                                                                                                      | Refer to Table I-12 on page 345.                                                                                                                                       |
| <i>t_look()</i>        | WAIT_FOR_COMPLETION                                                                                                                                                                                                                                                                                                                      |                                                                                                                                                                        |
| <i>t_open()</i>        | Local                                                                                                                                                                                                                                                                                                                                    | <i>t_open</i> sets blocking mode (that is, blocking or non-blocking)                                                                                                   |
| <i>t_optmgmt()</i>     | GET_ATTRIBUTES                                                                                                                                                                                                                                                                                                                           | To get Mode name<br>Refer to Table I-13 on page 345.                                                                                                                   |
| <i>t_rcv()</i>         | WAIT_FOR_COMPLETION<br>RECEIVE_AND_WAIT<br>[RECEIVE_EXPEDITED_DATA]                                                                                                                                                                                                                                                                      | Refer to Table I-14 on page 346.                                                                                                                                       |



| <b>XTI Function</b>   | <b>SNA FDX LU6.2 verb</b>                                      | <b>Comments</b>                                                                                                                                                                                                                                                                                                       |
|-----------------------|----------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>t_rcvconnect()</i> | WAIT_FOR_COMPLETION<br>GET_ATTRIBUTES<br>RECEIVE_AND_WAIT      | Refer to Figure I-2 on page 330.                                                                                                                                                                                                                                                                                      |
| <i>t_rcvdis()</i>     | Local                                                          | Event caused by DEALLOCATE_ABEND_* or RESOURCE_FAILURE_* return code on any verb                                                                                                                                                                                                                                      |
| <i>t_rcvrel()</i>     | Local                                                          | Event caused by DEALLOCATE_NORMAL return code on RECEIVE_* verb                                                                                                                                                                                                                                                       |
| <i>t_snd()</i>        | SEND_DATA (expedited data)<br>[FLUSH]<br>[SEND_EXPEDITED_DATA] | Every <i>t_snd</i> causes a SEND_DATA to be issued - even if T_MORE set. If T_MORE is set, the LL continuation bit is set.<br><br>A zero-length TSDU causes the following LL to be sent: hex 0002. This can be used to turn off the LL continuation set on the previous send.<br><br>Refer to Table I-16 on page 348. |
| <i>t_snddis()</i>     | DEALLOCATE<br>TYPE(ABEND_PROG)                                 | Takes down both directions of the connection<br><br>Refer to Table I-17 on page 349.                                                                                                                                                                                                                                  |
| <i>t_sndrel()</i>     | DEALLOCATE TYPE(FLUSH)                                         | Takes down send direction of conversation only.<br><br>Refer to Table I-19 on page 349.                                                                                                                                                                                                                               |
| <i>t_sndudata()</i>   | SEND_DATA on datagram server conversation                      | Refer to Table I-19 on page 349.                                                                                                                                                                                                                                                                                      |
| <i>t_stterror()</i>   | local                                                          |                                                                                                                                                                                                                                                                                                                       |
| <i>t_sync()</i>       | local                                                          |                                                                                                                                                                                                                                                                                                                       |
| <i>t_unbind()</i>     | local                                                          |                                                                                                                                                                                                                                                                                                                       |

## I.3.3.1 Parameter Mappings

**Table I-6** Relation Symbol Description

| Relation Symbol | Meaning                                                                                            |
|-----------------|----------------------------------------------------------------------------------------------------|
| Used Locally    | Value is used locally by XTI Mapper                                                                |
| Created Locally | XTI Mapper creates the value                                                                       |
| Constant        | Only one value is acceptable in this field. It is an error condition if any other value is passed. |
| <---            | XTI Application parameter maps directly to FDX Verb parameter.                                     |
| --->            | FDX Verb parameter maps directly to XTI Application parameter.                                     |

**Table I-7** *t\_accept* <--> FDX Verbs and Parameters

| XTI Function and Parameters | <--Relation-->  | FDX Verb & Parameter     |
|-----------------------------|-----------------|--------------------------|
| <b>t_accept</b>             |                 | <b>RECEIVE_AND_WAIT</b>  |
| <b>Input</b>                |                 |                          |
| fd                          | Used Locally    |                          |
| resfd                       | ---->           | RESOURCE(variable)       |
| call-->addr.len             | Used Locally    |                          |
| call-->addr.buf             | Used Locally    |                          |
| call-->opt.len              | Used Locally    |                          |
| call-->opt.buf              | Used Locally    |                          |
| call-->udata.len            | Constant        | =0                       |
| call-->udata.buf            | Constant        | =nullptr                 |
|                             | Created Locally | LENGTH(variable)         |
|                             | Created Locally | FILL(addr of local bufr) |
|                             | Constant        | WAIT_OBJECT(BLOCKING)    |
| <b>Output</b>               |                 |                          |
| t_errno                     | <----           | RETURN_CODE(variable)    |

**Table I-8** *t\_bind* <--> FDX Verbs and Parameters

| <b>XTI Function and Parameters</b> | <b>&lt;--Relation--&gt;</b> | <b>FDX Verb &amp; Parameter</b>                                                   |
|------------------------------------|-----------------------------|-----------------------------------------------------------------------------------|
| <b>t_bind</b> with qlen>0          |                             | <b>RECEIVE_ALLOCATE</b>                                                           |
| <b>Input</b>                       |                             |                                                                                   |
| fd                                 | Used Locally                |                                                                                   |
| req-->addr.len                     | Used Locally                |                                                                                   |
| req-->addr.buf                     | <----                       | LOCAL_LU_NAME(variable)<br>TP_NAME(variable)                                      |
| req-->qlen                         | Used Locally                | >0, RECEIVE_ALLOCATE<br>issued for each connection<br>request that can be queued. |
| ret-->addr.maxlen                  | Used Locally                |                                                                                   |
|                                    | Constant                    | RETURN_CONTROL<br>(WHEN_ALLOCATE_RECEIVED)                                        |
|                                    | Constant                    | SCOPE(ALL)                                                                        |
|                                    | Created Locally             | WAIT_OBJECT(BLOCKING)                                                             |
| <b>Output</b>                      |                             |                                                                                   |
| ret-->addr.len                     | Created Locally             |                                                                                   |
| ret-->addr.buf                     | ---->                       | LOCAL_LU_NAME(variable)<br>TP_AL_LU_NAME(variable)                                |
| ret-->addr.qlen                    | Created Locally             |                                                                                   |
|                                    | Used Locally                | RESOURCE(variable)                                                                |
| t_errno                            | ---->                       | RETURN_CODE(variable)                                                             |

**Table I-9** *t\_close* <--> FDX Verbs and Parameters

| <b>XTI Function and Parameters</b>                              | <b>&lt;--Relation--&gt;</b> | <b>FDX Verb &amp; Parameter</b> |
|-----------------------------------------------------------------|-----------------------------|---------------------------------|
| <b>t_close</b><br>If connection is still in<br>T_DATAXFER state |                             | <b>DEALLOCATE</b>               |
| <b>Input</b>                                                    |                             |                                 |
| fd                                                              | ---->                       | RESOURCE(variable)              |
|                                                                 | Constant                    | TYPE(FLUSH)                     |
| <b>Output</b>                                                   |                             |                                 |
| t_errno                                                         | <----                       | RETURN_CODE(variable)           |

Table I-10 *t\_connect* <--> FDX Verbs and Parameters

| XTI Function and Parameters | <--Relation-->  | FDX Verb & Parameter                                                                                                               |
|-----------------------------|-----------------|------------------------------------------------------------------------------------------------------------------------------------|
| <b>t_connect</b>            |                 | <b>ALLOCATE RETURN_CONTROL</b>                                                                                                     |
| <b>Input</b>                |                 |                                                                                                                                    |
| fd                          |                 |                                                                                                                                    |
| sndcall-->addr.len          | Used Locally    |                                                                                                                                    |
| sndcall-->addr.buf          | ---->           | LU_NAME(), TP_NAME()                                                                                                               |
| sndcall-->opt.len           | Used Locally    |                                                                                                                                    |
| sndcall-->opt.buf           | ---->           | MODE_NAME()                                                                                                                        |
| sndcall-->udata.len         | Constant        | =0, user data not allowed                                                                                                          |
| sndcall-->udata.buf         | Constant        | =nullptr                                                                                                                           |
| rcvcall-->addr.maxlen       | Used Locally    |                                                                                                                                    |
| rcvcall-->addr.buf          | Used Locally    |                                                                                                                                    |
| rcvcall-->opt.maxlen        | Constant        | =0, user data not allowed                                                                                                          |
| rcvcall-->opt.buf           | Used Locally    |                                                                                                                                    |
| rcvcall-->udata.maxlen      | Constant        | =0, user data not allowed                                                                                                          |
| rcvcall-->udata.buf         | Constant        | =nullptr                                                                                                                           |
|                             | Constant        | TYPE(FULL_DUPLEX_BASIC_CONV)                                                                                                       |
|                             | Created Locally | RETURN_CODE<br>(WHEN_SESSION_FREE)<br><br>If platform does not support this<br>tower (Tower 205), use<br>(WHEN_SESSION_ALLOCATED). |
|                             | Created Locally | WAIT_OBJECT(BLOCKING) if<br>blocking<br><br>WAIT_OBJECT(VALUE(variable)) if<br>non-blocking                                        |
| <b>Output</b>               |                 |                                                                                                                                    |
|                             | Used Locally    | RESOURCE(variable)                                                                                                                 |
| t_errno                     | <----           |                                                                                                                                    |
| <b>t_connect</b>            |                 | <b>GET_ATTRIBUTES</b>                                                                                                              |
| <b>Input</b>                |                 |                                                                                                                                    |
| fd                          | Used Locally    |                                                                                                                                    |
|                             | Created Locally | RESOURCE(variable)                                                                                                                 |
| <b>Output</b>               |                 |                                                                                                                                    |
| rcvcall-->addr.len          | <----           | PARTNER_FULLY_QUALIFIED_<br>LU_NAME(variable)                                                                                      |
| rcvcall-->addr.buf          | <-----          | PARTNER_FULLY_QUALIFIED_<br>LU_NAME(variable)                                                                                      |
| rcvcall-->opt.len           | <----           | MODE_NAME(variable)                                                                                                                |
| rcvcall-->opt.buf           | <----           | MODE_NAME(variable)                                                                                                                |
| t_errno                     | <----           | RETURN_CODE(variable)                                                                                                              |
| <b>t_connect</b>            |                 | <b>RECEIVE_AND_WAIT</b>                                                                                                            |
| <b>Input</b>                |                 |                                                                                                                                    |

| <b>XTI Function and Parameters</b> | <b>&lt;--Relation--&gt;</b> | <b>FDX Verb &amp; Parameter</b> |
|------------------------------------|-----------------------------|---------------------------------|
| fd                                 | Used Locally                |                                 |
|                                    | Created Locally             | RESOURCE(variable)              |
| <b>Output</b>                      |                             |                                 |
|                                    | Created Locally             | LENGTH(variable)                |
|                                    | Created Locally             | FILL(addr of local bufr)        |
|                                    | Constant                    | WAIT_OBJECT(BLOCKING)           |

Table I-11 *t\_getprocaddr* <--> FDX Verbs and Parameters

| XTI Function and Parameters | <--Relation-->  | FDX Verb & Parameter                                           |
|-----------------------------|-----------------|----------------------------------------------------------------|
| <b>t_getprotaddr</b>        |                 | <b>GET_ATTRIBUTES to get partner LU name</b>                   |
| <b>Input</b>                |                 |                                                                |
| fd                          | Used Locally    |                                                                |
|                             | Created Locally | RESOURCE(variable)                                             |
| boundaddr-->maxlen          | Used Locally    |                                                                |
| boundaddr-->addr.buf        | Used Locally    |                                                                |
| peeraddr-->maxlen           | Used Locally    |                                                                |
| peeraddr-->addr.buf         | Used Locally    |                                                                |
| <b>Output</b>               |                 |                                                                |
| peeraddr-->addr.len         | Created Locally |                                                                |
| buf(peeraddr-->addr.buf)    | <----           | PARTNER_FULLY_QUALIFIED_LU_NAME(variable)                      |
| t_errno                     | <----           | RETURN_CODE(variable)                                          |
| <b>t_getprotaddr</b>        |                 | <b>GET_TP_PROPERTIES to get local TP name</b>                  |
| <b>Input</b>                |                 |                                                                |
| fd                          | Used Locally    |                                                                |
|                             | Created Locally | RESOURCE(variable)                                             |
| <b>Output</b>               |                 |                                                                |
| boundaddr-->addr.len        | Created Locally |                                                                |
| buf(boundaddr-->addr.buf)   | <----           | OWN_FULLY_QUALIFIED_LU_NAME(variable)<br>OWN_TP_NAME(variable) |
| t_errno                     | <----           | RETURN_CODE(variable)                                          |

Table I-12 *t\_listen* <--> FDX Verbs and Parameters

| XTI Function and Parameters    | <--Relation-->  | FDX Verb & Parameter                      |
|--------------------------------|-----------------|-------------------------------------------|
| <b>t_listen</b>                |                 | <b>WAIT_FOR_COMPLETION</b>                |
| <b>Input</b>                   |                 |                                           |
|                                | Created Locally | WAIT_OBJECT_LIST(variable)                |
|                                | Constant        | TIMEOUT(VALUE(variable=0))                |
| <b>Output</b>                  |                 |                                           |
| t_errno                        | <----           | RETURN_CODE(variable)                     |
|                                | Used Locally    | STATUS_LIST(variable)                     |
| <b>t_listen</b>                |                 | <b>GET_ATTRIBUTES</b>                     |
| <b>Input</b>                   |                 |                                           |
| fd                             | Used Locally    |                                           |
|                                | Created Locally | RESOURCE(variable)                        |
| <b>Output</b>                  |                 |                                           |
| call-->addr.len                | <-----          | PARTNER_FULLY_QUALIFIED_LU_NAME(variable) |
| bufr&larrow.(call-->addr.bufr) | <-----          | PARTNER_FULLY_QUALIFIED_LU_NAME(variable) |
| call-->opt.len                 | <-----          | MODE_NAME(variable)                       |
| bufr&larrow.(call-->opt.bufr)  | <-----          | MODE_NAME(variable)                       |

Table I-13 *t\_optmgmt* <--> FDX Verbs and Parameters

| XTI Function and Parameters | <--Relation-->  | FDX Verb & Parameter  |
|-----------------------------|-----------------|-----------------------|
| <b>t_optmgmt</b>            |                 | <b>GET_ATTRIBUTES</b> |
| <b>Input</b>                |                 |                       |
| fd                          | Used Locally    |                       |
|                             | Created Locally | RESOURCE(variable)    |
| req-->opt.maxlen            | Used Locally    |                       |
| req-->opt.len               | Used Locally    |                       |
| req-->opt.bufr              | Used Locally    |                       |
| req-->opt.flags             | Used Locally    |                       |
| ret-->opt.maxlen            | Used Locally    |                       |
| ret-->opt.bufr              | Used Locally    |                       |
| <b>Output</b>               |                 |                       |
| ret-->opt.len               | <----           | MODE_NAME(variable)   |
| ret-->opt.bufr              | <----           | MODE_NAME(variable)   |
| ret-->flags                 | Created Locally |                       |
| t_errno                     | <----           | RETURN_CODE(variable) |

Table I-14 *t\_rcv* <--> FDX Verbs and Parameters

| XTI Function and Parameters                                                                                                                                                                                                                                                          | <--Relation-->  | FDX Verb & Parameter                                                                                                              |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <b>t_rcv</b>                                                                                                                                                                                                                                                                         |                 | <b>RECEIVE_AND_WAIT</b>                                                                                                           |
| A RECEIVE_AND_WAIT has been issued prior to this t_rcv. The data received on this RECEIVE_AND_WAIT will be returned to the XTI application via the t_rcv.<br>Before the return to the t_rcv, the mapper will issue another RECEIVE_AND_WAIT to post a receive for any incoming data. |                 |                                                                                                                                   |
| <b>Input</b>                                                                                                                                                                                                                                                                         |                 | <b>This is the RECEIVE_AND_WAIT that will be issued before the return to t_rcv.</b>                                               |
| fd                                                                                                                                                                                                                                                                                   | Used Locally    |                                                                                                                                   |
|                                                                                                                                                                                                                                                                                      | Created Locally | RESOURCE(variable)                                                                                                                |
| nbytes                                                                                                                                                                                                                                                                               | Used Locally    |                                                                                                                                   |
|                                                                                                                                                                                                                                                                                      | Created Locally | LENGTH(variable)                                                                                                                  |
|                                                                                                                                                                                                                                                                                      | Created Locally | FILL(XTI mapper buffer)                                                                                                           |
|                                                                                                                                                                                                                                                                                      | Created Locally | WAIT_OBJECT(VALUE(variable))                                                                                                      |
| <b>t_rcv</b>                                                                                                                                                                                                                                                                         |                 | <b>RECEIVE_AND_WAIT</b>                                                                                                           |
| <b>Output</b>                                                                                                                                                                                                                                                                        |                 | <b>These are fields from the previously issued RECEIVE_AND_WAIT</b>                                                               |
| buf                                                                                                                                                                                                                                                                                  | <----           | Data from FILL buffer                                                                                                             |
| Return Value for function                                                                                                                                                                                                                                                            | <----           | LENGTH(variable)                                                                                                                  |
| flags<br>• T_MORE<br>• T_EXPEDITED=NO/YES                                                                                                                                                                                                                                            | Created Locally | WHAT_RECEIVED(variable)<br>If there is expedited data to be received, a RECEIVE_EXPEDITED_DATA verb will be issued to receive it. |
| t_errno                                                                                                                                                                                                                                                                              | <----           | RETURN_CODE                                                                                                                       |
| <b>t_rcv</b>                                                                                                                                                                                                                                                                         |                 | <b>WAIT_FOR_COMPLETION</b>                                                                                                        |
| <b>Input</b>                                                                                                                                                                                                                                                                         |                 |                                                                                                                                   |
|                                                                                                                                                                                                                                                                                      | Created Locally | WAIT_OBJECT_LIST(variable)                                                                                                        |
|                                                                                                                                                                                                                                                                                      | Constant        | TIMEOUT(VALUE(variable=0))                                                                                                        |
| <b>Output</b>                                                                                                                                                                                                                                                                        |                 |                                                                                                                                   |
| errno                                                                                                                                                                                                                                                                                | <----           | RETURN_CODE(variable)                                                                                                             |
| T_LOOK                                                                                                                                                                                                                                                                               | Used Locally    | STATUS_LIST(variable)                                                                                                             |



Table I-15 *t\_rcvconnect* <--> FDX Verbs and Parameters

| XTI Function and Parameters | <--Relation-->  | FDX Verb & Parameter                      |
|-----------------------------|-----------------|-------------------------------------------|
| <b>t_rcvconnect</b>         |                 | <b>GET_ATTRIBUTES</b>                     |
| <b>Input</b>                |                 |                                           |
| fd                          | Used Locally    |                                           |
|                             | Created Locally | RESOURCE(variable)                        |
| <b>Output</b>               |                 |                                           |
| call-->addr.len             | <----           | PARTNER_FULLY_QUALIFIED_LU_NAME(variable) |
| call-->addr.buf             | <----           | PARTNER_FULLY_QUALIFIED_LU_NAME(variable) |
| call-->opt.len              | <----           | MODE_NAME(variable)                       |
| call-->opt.buf              | <----           | MODE_NAME(variable)                       |
| t_errno                     | <----           | RETURN_CODE(variable)                     |
| <b>t_rcvconnect</b>         |                 | <b>RECEIVE_AND_WAIT</b>                   |
| <b>Input</b>                |                 |                                           |
| fd                          | Used Locally    |                                           |
|                             | Created Locally | RESOURCE(variable)                        |
|                             | Created Locally | LENGTH(variable)                          |
|                             | Created Locally | FILL(addr of local bufr)                  |
|                             | Constant        | WAIT_OBJECT(VALUE(variable))              |
| <b>t_rcvconnect</b>         |                 | WAIT_FOR_COMPLETION                       |
| <b>Input</b>                |                 |                                           |
|                             | Created Locally | WAIT_OBJECT_LIST(variable)                |
|                             | Constant        | TIMEOUT(VALUE(variable=0))                |
| <b>Output</b>               |                 |                                           |
| t_errno                     | <----           | RETURN_CODE(variable)                     |
| T_LOOK                      | Used Locally    | STATUS_LIST(variable)                     |

Table I-16 *t\_snd* <--> FDX Verbs and Parameters

| XTI Function and Parameters                                                                                    | <--Relation-->  | FDX Verb & Parameter                                                              |
|----------------------------------------------------------------------------------------------------------------|-----------------|-----------------------------------------------------------------------------------|
| <b>t_snd()</b>                                                                                                 |                 | <b>SEND_DATA</b>                                                                  |
| <b>Input</b>                                                                                                   |                 |                                                                                   |
| fd                                                                                                             | Used Locally    |                                                                                   |
|                                                                                                                | Created Locally | RESOURCE(variable)                                                                |
| buf                                                                                                            | ---->           | DATA(variable)                                                                    |
| nbytes                                                                                                         | ---->           | LENGTH(variable)                                                                  |
| flags <ul style="list-style-type: none"> <li>• T_EXPEDITED=NO</li> <li>• T_MORE</li> <li>• T_FLUSH</li> </ul>  | ---->           | LL continuation bit                                                               |
|                                                                                                                | Created Locally | WAIT_OBJECT(BLOCKING) if blocking<br>WAIT_OBJECT(VALUE(variable)) if non-blocking |
| <b>Output</b>                                                                                                  |                 |                                                                                   |
| (TLOOK)                                                                                                        | <----           | EXPEDITED_DATA_RECEIVED                                                           |
| t_errno                                                                                                        | <----           | RETURN_CODE                                                                       |
| <b>t_snd()</b>                                                                                                 |                 | <b>SEND_EXPEDITED_DATA</b>                                                        |
| <b>Input</b>                                                                                                   |                 |                                                                                   |
| fd                                                                                                             | Used Locally    |                                                                                   |
|                                                                                                                | Created Locally | RESOURCE(variable)                                                                |
| buf                                                                                                            | ---->           | DATA(variable)                                                                    |
| nbytes                                                                                                         | ---->           | LENGTH(variable)                                                                  |
| flags <ul style="list-style-type: none"> <li>• T_EXPEDITED=YES</li> <li>• T_MORE</li> <li>• T_FLUSH</li> </ul> | ---->           | LL continuation bit<br>FLUSH Verb                                                 |
|                                                                                                                | Created Locally | WAIT_OBJECT(BLOCKING) if blocking                                                 |
| WAIT_OBJECT(VALUE(variable)) if non-blocking                                                                   |                 |                                                                                   |
| <b>Output</b>                                                                                                  |                 |                                                                                   |
| (TLOOK)                                                                                                        | <----           | EXPEDITED_DATA_RECEIVED                                                           |
| t_errno                                                                                                        | <----           | RETURN_CODE                                                                       |

**Table I-17** *t\_snddis* (Existing Connection) <--> FDX Verbs and Parameters

| <b>XTI Function and Parameters</b>           | <b>&lt;--Relation--&gt;</b> | <b>FDX Verb &amp; Parameter</b> |
|----------------------------------------------|-----------------------------|---------------------------------|
| <b>t_snddis()</b><br>For existing connection |                             | DEALLOCATE                      |
| <b>Input</b>                                 |                             |                                 |
| fd                                           | Used Locally                |                                 |
|                                              | Created Locally             | RESOURCE(variable)              |
| call                                         | Constant                    | =nullptr                        |
|                                              | Constant                    | TYPE(ABEND_PROG)                |
| <b>Output</b>                                |                             |                                 |
| t_errno                                      | <----                       | RETURN_CODE(variable)           |

**Table I-18** *t\_snddis* (Incoming Connect Req.) <--> FDX Verbs and Parameters

| <b>XTI Function and Parameters</b>                         | <b>&lt;--Relation--&gt;</b> | <b>FDX Verb &amp; Parameter</b> |
|------------------------------------------------------------|-----------------------------|---------------------------------|
| <b>t_snddis()</b><br>To reject incoming connection request |                             | DEALLOCATE                      |
| <b>Input</b>                                               |                             |                                 |
| call-->sequence                                            | Used Locally                |                                 |
| <b>Output</b>                                              |                             |                                 |
| t_errno                                                    | <----                       | RETURN_CODE(variable)           |

**Table I-19** *t\_sndrel* <--> FDX Verbs and Parameters

| <b>XTI Function and Parameters</b> | <b>&lt;--Relation--&gt;</b> | <b>FDX Verb &amp; Parameter</b> |
|------------------------------------|-----------------------------|---------------------------------|
| <b>t_sndrel()</b>                  |                             | DEALLOCATE                      |
| <b>Input</b>                       |                             |                                 |
| fd                                 | Used Locally                |                                 |
|                                    | Created Locally             | RESOURCE(variable)              |
|                                    | Created Locally             | TYPE(FLUSH)                     |
| <b>Output</b>                      |                             |                                 |
| t_errno                            | <----                       | RETURN_CODE(variable)           |

**I.3.4 Half Duplex Mapping**

The interface to the SNA transport provider is the FDX LU 6.2 interface. If the SNA transport provider does not support FDX LU 6.2, the FDX verbs can be mapped to two half-duplex LU 6.2 connections, one used to send in each direction. This gives the appearance of a full duplex connection without requiring any conversation direction turn-arounds on the half-duplex conversations.

The mapping of FDX LU 6.2 verbs to two half-duplex connections is described in FDX Kit.

### I.3.5 Return Code to Event Mapping

The following table Table I-20 shows how the return codes on LU 6.2 verbs are mapped to XTI events.

Any return code for which there no mapping given in this table will create a disconnect.

**Table I-20** Mapping of XTI Events to SNA Events

| <b>XTI Event</b> | <b>Full Duplex SNA Event</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| T_CONNECT        | ALLOCATE completes with RETURN_CODE=OK                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| T_DATA           | RECEIVE_AND_WAIT completes with OK return code and WHAT_RECEIVED = <ul style="list-style-type: none"> <li>• DATA_COMPLETE</li> <li>• DATA_INCOMPLETE</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| T_DISCONNECT     | One of the following has occurred:<br>SEND_DATA issued with RETURN_CODE = ERROR_INDICATION with subcode from list below: <ul style="list-style-type: none"> <li>• ALLOCATION_ERROR</li> <li>• DEALLOCATE_ABEND_PROG</li> <li>• DEALLOCATE_ABEND_SVC</li> <li>• DEALLOCATE_ABEND_TIMER</li> <li>• RESOURCE_FAILURE_NO_RETRY</li> <li>• RESOURCE_FAILURE_RETRY</li> </ul> Any other verb issued with RETURN_CODE of <ul style="list-style-type: none"> <li>• ALLOCATION_ERROR</li> <li>• DEALLOCATE_ABEND_PROG</li> <li>• DEALLOCATE_ABEND_SVC</li> <li>• DEALLOCATE_ABEND_TIMER</li> <li>• RESOURCE_FAILURE_NO_RETRY</li> <li>• RESOURCE_FAILURE_RETRY</li> </ul> |
| T_EXDATA         | RECEIVE_EXPEDITED_DATA completes with OK return code                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| T_GODATA         | Flow control restrictions on normal data flow that lead to a [TFLOW] error have been lifted. Normal data may be sent again.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| T_GOEXDATA       | Flow control restrictions on expedited data flow that lead to a [TFLOW] error have been lifted. Expedited data may be sent again.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| T_LISTEN         | When the partner program is instantiated when the connection request arrives (the typical LU 6.2 model), this event is posted as soon as the program issues the <i>t_listen()</i> function call.<br><br>When the partner program already exists, this event is posted when the connection request arrives and is matched with the program.                                                                                                                                                                                                                                                                                                                       |
| T_ORDREL         | Set when a RECEIVE_* verb completes with RETURN_CODE = DEALLOCATE_NORMAL.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |



## The Internet Protocols

The Internet Protocol (IP) family is a collection of protocols designed for use in the Internet and using the Internet address format. The Internet family provides protocol support for the **SOCK\_STREAM** and **SOCK\_DGRAM** socket types.

Internet addresses are 4-byte quantities, stored in network byte order (on “little-endian” machines, these are word and byte reversed). The `<netinet/in.h>` header defines this address as a discriminated union.

The address `INADDR_ANY` can be given in a `bind()` call on a socket that uses the TCP or UDP protocol. For TCP, this lets the socket accept connections targeted at any of the host’s IP addresses. For UDP, it lets the socket accept packets addressed to any of the host’s IP addresses. The address `INADDR_BROADCAST` is the IP broadcast address. If `INADDR_BROADCAST` is used in a `sendto()` call on a UDP socket, and the host has one or more network interfaces that support the broadcast feature, a broadcast packet will be sent via one or more of those interfaces.

The Internet protocol family includes the following protocols:

- the Internet protocol (IP)
- the Internet Control Message Protocol (ICMP)
- the Transmission Control Protocol (TCP)
- the User Datagram Protocol (UDP).

TCP is used to support the **SOCK\_STREAM** abstraction while UDP is used to support the **SOCK\_DGRAM** abstraction.

The 32-bit Internet address contains both network and host parts. It is frequency-encoded; the most-significant bit is clear in Class A addresses, in which the high-order 8 bits are the network number. Class B addresses use the high-order 16 bits as the network field, and Class C addresses have a 24-bit network part. Sites with a cluster of local networks and a connection to the Internet may choose to use a single network number for cluster; this is done by using subnet addressing. The local (host) portion of the address is further subdivided into subnet and host parts. Within a subnet, each subnet appears to be an individual network; externally, the entire cluster appears to be a single, uniform network requiring only a single routing entry. The above description of classes of addresses need not apply in a network where the (relatively recent) Classless Internet Domain Routing (CIDR) described in RFC1517-RFC1518 is deployed.





# IPX/SPX Transport Provider

## K.1 General

This appendix is a Preliminary Specification. It specifies protocol-specific information that is relevant for mapping XTI functions to SPX and IPX transport providers.

The description given here is limited to the IPX protocol and the enhanced SPX (or SPXII) protocol. All references to the SPX protocol refer to this version unless specifically noted. In compliance with the X/Open Interface Adoption Criteria, this protocol is obtainable from multiple sources.

### Notes:

1. Neither IPX nor SPX supports expedited data. All data is handled on a first-come, first-served basis.
2. The protocol-specific data structures used by IPX and SPX (most notably, the `T_SPX2_OPTIONS` structure) are likely to grow in the future.

## K.2 Namespace

### K.2.1 IPX

If the header `xti_ipx.h` is included, identifiers with the prefixes, suffixes or complete names shown are reserved for any use by the implementation.

| Header                         | Prefix             |
|--------------------------------|--------------------|
| <code>&lt;xti_ipx.h&gt;</code> | <code>t_ipx</code> |

If the header `xti_ipx.h` is included, macros with the following prefixes may be defined. After the last inclusion of `xti_ipx.h` an application may use identifiers with the following prefixes for its own purpose, provided their use is preceded by an `#undef` of the macro.

| Header                         | Prefix              |
|--------------------------------|---------------------|
| <code>&lt;xti_ipx.h&gt;</code> | <code>T_IPX_</code> |

### K.2.2 SPX

If the header **x<sub>ti</sub>\_sp<sub>x</sub>.h** is included identifiers with the prefixes, suffixes or complete names shown are reserved for any use by the implementation.

| Header                                | Prefix                    |
|---------------------------------------|---------------------------|
| <x <sub>ti</sub> _sp <sub>x</sub> .h> | t_sp <sub>x</sub> , T_SPX |

If the header **x<sub>ti</sub>\_sp<sub>x</sub>.h** is included, macros with the prefixes may be defined. After the last inclusion of **x<sub>ti</sub>\_sp<sub>x</sub>.h** an application may use identifiers with the following prefixes for its own purpose, provided there use is preceded by an **#undef** of the macro.

| Header                                | Prefix |
|---------------------------------------|--------|
| <x <sub>ti</sub> _sp <sub>x</sub> .h> | T_SPX_ |

## K.3 Options

### K.3.1 IPX-level Options

IPX options are association related. IPX options may be negotiated in all XTI states except T\_UNBND and T\_UNINIT. The level associated with each option is T\_IPX\_OPT. The name member should be set to T\_IPX\_OPTS\_V1. IPX options are stored in a structure of type **t\_ipxOptions** that follows the XTI **t\_opthdr** structure containing at least the following members:

```
typedef struct t_ipxOptions {
 unsigned short ipx_checksum /* Checksum */
 unsigned char ipx_packet_type /* IPX Packet type */
} t_ipxOpts_t;
```

This option may only be manipulated using the *t\_sndudata()* or *t\_rcvudata()* functions.

Other options may be defined in the future.

At least the following packet types are known to IPX:

T\_IPX\_NULL\_PACKET\_TYPE Used for all packets not classified by any other type  
T\_IPX\_NCP\_PACKET\_TYPE Used for NCP packets (that is, Netware Control Protocol)  
T\_IPX\_SPX\_PACKET\_TYPE Sequenced packet protocol used for SPX packets

### K.3.2 SPX-level Options

Some values of the SPX options have end-to-end significance as defined below. Unless otherwise noted, they may be negotiated in all XTI states except T\_UNBND and T\_UNINIT. The level associated with each option is T\_SPX\_OPT. The name member is set to T\_SPX\_OPTS\_V1. SPX options are stored following the XTI **t\_opthdr** structure in a format defined by the following data structure:

```
typedef struct t_spx2_options {
 unsigned int versionNumber;
 unsigned int spxIIOptionNegotiate;
 unsigned int spxIIRetryCount;
 unsigned int spxIIMinimumRetryDelay;
 unsigned int spxIIMaximumRetryDelta;
 unsigned int spxIIWatchdogTimeout;
 unsigned int spxIIConnectionTimeout;
 unsigned int spxIILocalWindowSize;
 unsigned int spxIIRemoteWindowSize;
 unsigned int spxIIConnectionID;
 unsigned int spxIIInboundPacketSize;
 unsigned int spxIIOutboundPacketSize;
 unsigned int spxIISessionFlags;
} T_SPX2_OPTIONS;
```

An application passing a **t\_opthdr** structure followed by a **T\_SPX2\_OPTIONS** structure to a function that can support options information will get information about all legal options on each call.

Each of this structure's members is described below.

#### versionNumber

The application must set this member to T\_SPX\_OPTIONS\_VERSION. This member is used to manage forward compatibility as this structure is extended in the future. Access is through the *t\_optmgmt()* or other functions that use the **t\_call** data structure. T\_SPX\_OPTIONS\_VERSION is currently set to 1.

#### spxIIOptionNegotiate

This value is association-related (that is, has end-to-end significance). This member can be set to T\_SPX\_NEGOTIATE\_OPTIONS (default) to indicate that an endpoint wishes to exchange option data with a remote endpoint, or to T\_SPX\_NO\_NEGOTIATE\_OPTIONS to indicate that it does not wish to do this. This is a negotiable value, and may be manipulated through the *t\_optmgmt()* function or through functions that use the **t\_call** data structure.

#### spxIIRetryCount

This value specifies the number of unsuccessful transmission attempts that SPX will make before aborting a connection. Setting this value to 0 causes the default value (10) to be used. This is a negotiable value, and may be manipulated through the *t\_optmgmt()* function or through functions that use the **t\_call** data structure.

#### spxIIMinimumRetryDelay

An internal round-trip time algorithm normally calculates the delay before a transmission retry. Setting this member to a non-zero value overrides this algorithm and uses the value as the fixed number of milliseconds before a retransmit. The default value is 300 milliseconds. This is a negotiable value, and may be manipulated through the *t\_optmgmt()* function or through functions that use the **t\_call** data structure.

**spxIIMaximumRetryDelta**

The value of this member is added to *spxIIMinimumRetryDelay* or to the current round-trip time to determine the maximum retry delay. Setting this value to 0 causes the default delay (5 seconds) to be used. This is a negotiable value, and may be manipulated through the *t\_optmgmt()* function or through functions that use the **t\_call** data structure.

**spxIIWatchdogTimeout**

This value determines the number of seconds that SPX will wait on an inactive connection before sending a still-alive query to the remote endpoint. This is a negotiable value, and may be manipulated through the *t\_optmgmt()* function or through functions that use the **t\_call** data structure.

**spxIIConnectionTimeout**

This value is the number of seconds that SPX will wait after a successful connect request before the first session packet must arrive. If a packet does not arrive in this interval, the connection is aborted. This is a negotiable value, and may be manipulated through the *t\_optmgmt()* function or through functions that use the **t\_call** data structure.

**spxIILocalWindowSize**

This value is association-related (that is, has end-to-end significance). This member sets the size of the local endpoint's receive window in packets. A zero setting requests the driver to negotiate the window size. The local driver default is 8. This is a negotiable value, and may be manipulated through the *t\_optmgmt()* function or through functions that use the **t\_call** data structure.

**spxIIRemoteWindowSize**

This value is association-related (that is, has end-to-end significance). This member is meaningful only after a connection has been established, that is, in states T\_OUTCON, T\_INCON, T\_DATAXFER. It contains the number of packets in the remote endpoint's receive window. This is a non-negotiable value and is retrieved through functions that use the **t\_call** data structure. Attempts to set this value with the T\_NEGOTIATE flag set will have no effect.

**spxIIConnectionID**

This value is association-related (that is, has end-to-end significance). This member is meaningful only after a connection has been established, that is, in states T\_OUTCON, T\_INCON, T\_DATAXFER. It contains the local endpoint connection ID. This is a non-negotiable value and is retrieved through functions that use the **t\_call** data structure. Attempts to set this value with the T\_NEGOTIATE flag set will have no effect.

**spxIIIInboundPacketSize**

This value is association-related (that is, has end-to-end significance). This member is meaningful only after a connection has been established, that is, in states T\_OUTCON, T\_INCON, T\_DATAXFER. It contains the number of bytes in each incoming data packet. This value may change due to a route change in mid-connection. There is no way to notify an application that this has occurred. This is a non-negotiable value and is retrieved through functions that use the **t\_call** data structure. Attempts to set this value with the T\_NEGOTIATE flag set will have no effect.

**spxIIOutboundPacketSize**

This value is association-related (that is, has end-to-end significance). This member is meaningful only after a connection has been established, that is, in states T\_OUTCON, T\_INCON, T\_DATAXFER. It contains the number of bytes in each outgoing data packet. This value may change due to a route change in mid-connection. There is no way to notify an application that this has occurred. This is a non-negotiable value and is retrieved through functions that use the **t\_call** data structure. Attempts to set this value with the

T\_NEGOTIATE flag set will have no effect.

#### spxIISessionFlags

This value is association-related (that is, has end-to-end significance). This bit member contains flags that control physical layer characteristics of SPX packets. These may include checksums, data encryption, or data signing. The following flags have been defined:

| Flag                  | Value | Description         |
|-----------------------|-------|---------------------|
| T_SPX_SF_NONE         | 0x00  | Options off         |
| T_SPX_SF_IPX_CHECKSUM | 0x01  | Packet checksums on |
| T_SPX_SF_SPX2_SESSION | 0x02  | Compatibility flag  |

The T\_SPX\_SF\_NONE and T\_SPX\_SF\_IPX\_CHECKSUM flags are read/write, and accessible through the *t\_optmgmt()* function and functions using the **t\_call** data structure. The T\_SPX\_SF\_SPX2\_SESSION flag is read-only and is accessible through functions that use the **t\_call** data structure.

## K.4 Functions

*t\_accept()* No special considerations.

*t\_bind()* **For IPX and SPX:**

IPX and SPX support static and dynamic port assignment. If the application does not wish to chose a port, it sets *req* to NULL or *req*→*addr.len* to 0. In this case, the provider assigns a dynamic port in the range 0x4000 to 0x7fff.

If the application wishes to bind to a specific port (that is, a static port), *req*→*addr.buf* must point to an addressing structure.

A process without sufficient privilege can only request a port in the range 0x8000 to 0xffff. If the request is for a port outside the static port range, *t\_bind()* will fail and *t\_errno* will be set to TACCESS. A process with appropriate privilege can request any port number.

**For SPX:**

SPX allows only a single transport endpoint to be bound to a port. If a requested port is already bound to another endpoint, *t\_bind()* will fail and set *t\_errno* to [TADDRBUSY].

*t\_connect()* SPX does not support connect user data.

*t\_getinfo()* The default characteristics returned by *t\_getinfo()* in the **info** structure for an IPX endpoint are:

| Parameters                    | Before call | After call     |
|-------------------------------|-------------|----------------|
| <i>info</i> → <i>addr</i>     | /           | x              |
| <i>info</i> → <i>options</i>  | /           | x              |
| <i>info</i> → <i>tsdu</i>     | /           | x (1)          |
| <i>info</i> → <i>etsdu</i>    | /           | T_INVALID (-2) |
| <i>info</i> → <i>connect</i>  | /           | T_INVALID (-2) |
| <i>info</i> → <i>discon</i>   | /           | T_INVALID (-2) |
| <i>info</i> → <i>servtype</i> | /           | T_CLTS         |
| <i>info</i> → <i>flags</i>    | /           | 0              |

(1) ‘x’ is the smallest of the maximum size of transport data units supported by connected LANs.

The default characteristics returned by *t\_getinfo()* in the **info** structure for an SPX endpoint are:

| Parameters                     | Before call | After call      |
|--------------------------------|-------------|-----------------|
| <i>info</i> -> <i>addr</i>     | /           | x               |
| <i>info</i> -> <i>options</i>  | /           | x               |
| <i>info</i> -> <i>tsdu</i>     | /           | T_INFINITE (-1) |
| <i>info</i> -> <i>etsdu</i>    | /           | T_INVALID (-2)  |
| <i>info</i> -> <i>connect</i>  | /           | T_INVALID (-2)  |
| <i>info</i> -> <i>discon</i>   | /           | T_INVALID (-2)  |
| <i>info</i> -> <i>servtype</i> | /           | T_COTS_ORD      |
| <i>info</i> -> <i>flags</i>    | /           | 0               |

- t\_listen()* SPX does not support connect user data.  
An option buffer, as described in Section K.3.2, ‘‘SPX-level Options’’ will be passed to the user (unless *call*→*opt.maxlen* is zero).
- t\_open()* The default characteristics returned by *t\_open()* in the **info** structure for an IPX endpoint are the same as those listed in the table above for *t\_getinfo()*.  
Similarly, the default characteristics returned by *t\_open()* in the **info** structure for an SPX endpoint are the same as those listed in the table above for *t\_getinfo()*.
- t\_optmgmt()* **For IPX:**  
No IPX options are currently available through the *t\_optmgmt()* function.  
**For SPX:**  
Options may be set or retrieved using the *t\_optmgmt()* function only when the endpoint is in the state T\_IDLE.  
SPX supports the standard options buffer using the **t\_opthdr** structure. The **t\_opthdr** structure is followed by a **T\_SPX2\_OPTIONS** structure, as described in Section K.3, Options.  
SPX supports options that may be examined and negotiated with the *t\_optmgmt()* function. The options are listed and described in Section K.3, Options.
- t\_rcv()* SPX does not support expedited data, so the T\_EXPEDITED flag will not be set.  
SPX allows logical units of data to be unlimited in length.

If the SPX watchdog (see Section K.3.2, “SPX-level Options”) determines that the remote transport endpoint is no longer participating in the connection, the SPX watchdog generates a disconnect indication which causes `t_rcv()` to return with a [T\_LOOK] error.

- `t_rcvdis()` SPX does not support disconnect user data.
- On return, the `discon→reason` member is set to one of the following:
- `T_SPX_CONNECTION_TERMINATED`  
Indicates that no error occurred and an SPX terminate connection packet was received from the remote endpoint. This reason code indicates success.
- `T_SPX_CONNECTION_FAILED`  
Indicates that the remote endpoint failed to acknowledge a transmission.
- `t_rcvdata()` The `t_rcvdata` function is issued on an IPX endpoint.
- On return from the function, `unitdata→addr.buf` points to the address information for the remote endpoint, and `unitdata→opt.buf` points to an options buffer as described in Section K.3.1, “IPX-level Options”.
- If a packet received by a `t_rcvdata()` request did not have the checksum calculated and verified, the `t_ipx_checksum` member of the **struct t\_ipxOptions**, is set to 0xFFFF. Any other value indicates that the provider has calculated and verified the checksum of the received packet. For all received packets, the `t_ipx_packet_type` member is set to the packet type of the received packet.
- `t_snd()` SPX does not support expedited data so the T\_EXPEDITED flag must not be set.
- `t_snddis()` SPX does not support the sending of user data or options with a disconnect request.
- `t_sndudata()` If `unitdata→opt.len` is zero, then a packet type of T\_IPX\_NULL\_PACKET\_TYPE is used and no checksum is generated.
- To send any other packet type, `unitdata->opt` must reference an options buffer as defined in Section K.3.1. “IPX-level Options”. The IPX packet type must be defined. A checksum will be generated if `t_ipx_checksum` is set to T\_IPX\_CHECKSUM\_TRIGGER.





# ATM Transport Protocol Information for XTI

## L.1 General

This appendix is a Preliminary Specification. It describes the protocol-specific information that is relevant for ATM transport providers.

The following notes apply:

- This version of XTI supports a subset of the functions specified by the ATM Forum as the User-Network Interface (see reference **UNI**), versions 3.0 and 3.1.
- Frequent reference is made to the ATM Forum's "Native ATM Services: Semantic Description, Version 1.0", identified in this Appendix as reference **ATMNAS**.
- The ATM transport provider does not support Connectionless Transport Service, so use of the functions *t\_rcvudata()*, *t\_rcvuderr()* and *t\_sndudata()* always cause the [TNOTSUPPORT] error to be returned.
- The ATM transport provider supports both reliable and unreliable data transport services. This is selected via parameter *name* of function *t\_open()*. Note that both services are categorized as a T\_COTS type transport provider.
- The ATM transport provider does not support an orderly release mechanism, so use of the functions *t\_sndrel()* and *t\_rcvrel()* always cause the [TNOTSUPPORT] error to be returned.
- The ATM transport provider does not support expedited data transfer.
- The ATM transport provider does not support sending user data during connection setup or release.
- The ATM transport provider does not support ATM PVCs.
- The ATM transport provider does not support a mechanism to specify the congestion indication bit and the user-user byte in AAL5.
- At the current time, only AAL-5 message mode is supported; each TSDU is carried across the network in the payload field of a single AAL-5 PDU. The ATM transport provider does not support AAL1 and User-defined AAL.
- When a transport user passively waits for incoming connect indications through a transport endpoint, the ATM protocol address bound to the endpoint must conform to the ATM Forum guidelines (referenced document **ATMNAS**) for the specification of a SAP address. The SAP address is a vector that includes fields for the ATM network address (with selector byte), identification of a layer-2 protocol, identification of a layer-3 protocol, and identification of an application
- At the current time, alternate BLLI negotiation is not supported; only a single BLLI information element is used in connection setup.
- A new event, T\_LEAFCHANGE, is defined which is used by event management to notify a transport user when the status of a leaf has changed on a point-to-multipoint connection. Specifically, event T\_LEAFCHANGE is flagged whenever any one of the following occurs:
  - an attempt to add a new leaf via *t\_addleaf()* in asynchronous mode was successful
  - an attempt to add a new leaf via *t\_addleaf()* in asynchronous mode was unsuccessful

- either the leaf or the ATM network caused the leaf to be removed from the connection.

## L.2 ATM Addresses

This section describes the two kinds of ATM addresses mentioned in this appendix: network and protocol.

### L.2.1 ATM Network Address

An ATM network address specifies an ATM device that resides on the user side of the User Network Interface (UNI). There are three typical uses of an ATM network address:

- A transport user actively initiating a connection can specify the ATM network address from which the connection is originating. This is helpful when the originating device supports multiple ATM network addresses over the same UNI.
- A transport user passively receiving a connection can inspect the ATM network address from which the connection originated.
- A transport user acting as the root of a point-to-multipoint connection specifies the ATM network address of any leaves to be added to the connection.

The ATM network address is expressed by either the `t_atm_addr` structure or the `t_atm_sap` structure, depending on the context. The `t_atm_addr` structure is used to query and negotiate options with the transport provider, whereas the `t_atm_sap` structure is used for any function that requires an ATM address as a parameter.

The `t_atm_addr` structure is defined in the options section. The `t_atm_sap` structure is defined in Section L.2.3 on page 365. Note that when the `t_atm_sap` structure is used to convey a network address, the following fields of the structure must have a value of TATM\_ABSENT:

```
t_atm_sap_layer2.SVE_tag
t_atm_sap_layer3.SVE_tag
t_atm_sap_appl.SVE_tag
```

### L.2.2 ATM Protocol Address

An ATM protocol address specifies an ATM device that resides on the user side of the UNI, plus a service access point (SAP) within the said ATM device. There are two typical uses of an ATM protocol address:

- A transport user actively initiating a connection specifies an ATM protocol address as the destination of the connection.
- A transport user passively receiving a connection specifies an ATM protocol address as the address at which the transport user is awaiting the incoming connection.

### L.2.3 t\_atm\_sap Structure

The ATM address (network or protocol) is expressed via the **t\_atm\_sap** structure, defined below. Note that an implementation may overlay this structure onto a character array, prepended by other information (for example, address family identifier).

```

struct t_atm_sap {

 struct t_atm_sap_addr {
 int8_t SVE_tag_addr;
 int8_t SVE_tag_selector;
 uint8_t address_format;
 uint8_t address_length;
 uint8_t address [20];
 } t_atm_sap_addr;

 struct t_atm_sap_layer2 {
 int8_t SVE_tag;
 uint8_t ID_type;
 union {
 uint8_t simple_ID;
 uint8_t user_defined_ID;
 } ID;
 } t_atm_sap_layer2;

 struct t_atm_sap_layer3 {
 int8_t SVE_tag;
 uint8_t ID_type;
 union {
 uint8_t simple_ID;
 int32_t IPI_ID;
 struct {
 uint8_t OUI [3];
 uint8_t PID [2];
 } SNAP_ID;
 uint8_t user_defined_ID;
 } ID;
 } t_atm_sap_layer3;

 struct t_atm_sap_appl {
 int8_t SVE_tag;
 uint8_t ID_type;
 union {
 uint8_t T_ISO_ID [8];
 struct {
 uint8_t OUI [3];
 uint8_t app_ID [4];
 } vendor_ID;
 uint8_t user_defined_ID[8];
 } ID;
 } t_atm_sap_appl;
}

```

Legal values for the field *t\_atm\_sap\_addr.SVE\_tag\_addr* are T\_ATM\_PRESENT and T\_ATM\_ANY. The semantic meaning of this field is found in section 4.4 of referenced document **ATMNAS**.

Legal values for the field *t\_atm\_sap\_addr.SVE\_tag\_selector* are T\_ATM\_PRESENT, T\_ATM\_ABSENT, and T\_ATM\_ANY. Note that T\_ATM\_PRESENT is valid only for ATM Endsystem addresses, and T\_ATM\_ABSENT is valid only for E.164 addresses. The semantic meaning of this field is found in section 4.4 of referenced document **ATMNAS**.

Legal values for the field *t\_atm\_sap\_addr.address\_format* are T\_ATM\_ENDSYS\_ADDR and T\_ATM\_E164\_ADDR. This field is mapped to octet 5 of the Q.2931 “Called Party Number” information element.

Legal values for the field *t\_atm\_sap\_addr.address\_length* are 0 thru 20. This field is not mapped to any octets of a Q.2931 information element. Instead, it specifies the valid number of array elements in field *t\_atm\_sap\_addr.address*.

Legal values for the field *t\_atm\_sap\_addr.address* can be found in section 5.1.3 of the referenced UNI specification (versions 3.0 and 3.1). This field is mapped to octets 6 and beyond of the Q.2931 “Called Party Number” information element.

Legal values for the field *t\_atm\_sap\_layer2.SVE\_tag* are T\_ATM\_PRESENT, T\_ATM\_ABSENT, and T\_ATM\_ANY. The semantic meaning of this field is found in section 4.4 of referenced document **ATMNAS**.

Legal values for the field *t\_atm\_sap\_layer2.ID\_type* are:

|                 |                                              |
|-----------------|----------------------------------------------|
| T_ATM_SIMPLE_ID | identification via ITU encoding              |
| T_ATM_USER_ID   | identification via a user-defined codepoint. |

This field is not mapped to any octets of a Q.2931 information element. Instead, it specifies the proper union member in the **t\_atm\_layer2** structure.

Legal values for the field *t\_atm\_sap\_layer2.ID.simple\_ID* are:

|                       |                  |
|-----------------------|------------------|
| T_ATM_BLLI2_I1745     | I.1745           |
| T_ATM_BLLI2_Q921      | Q.921            |
| T_ATM_BLLI2_X25_LINK  | X.25, link layer |
| T_ATM_BLLI2_X25_MLINK | X.25, multilink  |
| T_ATM_BLLI2_LAPB      | Extended LAPB    |
| T_ATM_BLLI2_HDLC_ARM  | I.4335, ARM      |
| T_ATM_BLLI2_HDLC_NRM  | I.4335, NRM      |
| T_ATM_BLLI2_HDLC_ABM  | I.4335, ABM      |
| T_ATM_BLLI2_I8802     | I.8802           |
| T_ATM_BLLI2_X75       | X.75             |
| T_ATM_BLLI2_Q922      | Q.922            |
| T_ATM_BLLI2_I7776     | I.7776           |

This field is mapped to octet 6 (bits 1 thru 5) of the Q.2931 BLLI (“Broadband Low Layer Information”) information element.

Legal values for the field *t\_atm\_sap\_layer2.ID.user\_defined\_ID* are 0 thru 127. This field is mapped to octet 6a (bits 1 thru 7) of the Q.2931 BLLI information element.

Legal values for the field *t\_atm\_sap\_layer3.SVE\_tag* are T\_ATM\_PRESENT, T\_ATM\_ABSENT, and T\_ATM\_ANY. The semantic meaning of this field is found in section 4.4 of referenced document **ATMNAS**.

Legal values for the field *t\_atm\_sap\_layer3.ID\_type* are:

|                 |                                 |
|-----------------|---------------------------------|
| T_ATM_SIMPLE_ID | identification via ITU encoding |
|-----------------|---------------------------------|

|               |                                                              |
|---------------|--------------------------------------------------------------|
| T_ATM_IPI_ID  | identification via ISO/IEC TR 9577 (during connection setup) |
| T_ATM_SNAP_ID | identification via SNAP                                      |
| T_ATM_USER_ID | identification via a user-defined codepoint                  |

This field is not mapped to any octets of a Q.2931 information element. Instead, it specifies the proper union member in the **t\_atm\_layer3** structure.

Legal values for the field *t\_atm\_sap\_layer3.ID.simple\_ID* are:

|                   |        |
|-------------------|--------|
| T_ATM_BLLI3_X25   | X.25   |
| T_ATM_BLLI3_I8208 | I.8208 |
| T_ATM_BLLI3_X223  | X.223  |
| T_ATM_BLLI3_I8473 | I.8473 |
| T_ATM_BLLI3_T70   | T.70   |
| T_ATM_BLLI3_I9577 | I.9577 |

Note that a value of T\_ATM\_BLLI3\_I9577 in this field indicates that the identification of the layer 3 protocol is done in the user (data) plane, as specified in ISO/IEC TR 9577. This field is mapped to octet 7 (bits 1 thru 5) of the Q.2931 BLLI information element.

Legal values for the field *t\_atm\_sap\_layer3.ID.IPI\_ID* are those values defined by ISO/IEC TR 9577. This field is mapped to octet 7a (bits 1 thru 7) and octet 7b (bit 7) of the Q.2931 BLLI information element.

Legal values for the field *t\_atm\_sap\_layer3.ID.SNAP\_ID.OUI* are the 24-bit Organization Unique Identifiers assigned by the IEEE. This field is mapped to octets 8.1 thru 8.3 of the Q.2931 BLLI information element.

Legal values for the field *t\_atm\_sap\_layer3.ID.SNAP\_ID.PID* are defined by the organization identified in the preceding field. This field is mapped to octets 8.4 thru 8.5 of the Q.2931 BLLI information element.

Legal values for the field *t\_atm\_sap\_layer3.ID.user\_defined\_ID* are 0 thru 127. This field is mapped to octet 7a (bits 1 thru 7) of the Q.2931 BLLI information element.

Legal values for the field *t\_atm\_sap\_appl.SVE\_tag* are T\_ATM\_PRESENT, T\_ATM\_ABSENT, and T\_ATM\_ANY. The semantic meaning of this field is found in section 4.4 of referenced document **ATMNAS**.

Legal values for the field *t\_atm\_sap\_appl.ID\_type* are:

|                     |                                             |
|---------------------|---------------------------------------------|
| T_ATM_ISO_APP_ID    | ISO codepoint                               |
| T_ATM_VENDOR_APP_ID | vendor-specific codepoint                   |
| T_ATM_USER_APP_ID   | identification via a user-defined codepoint |

This field is mapped to octet 5 (bits 1 thru 7) of the Q.2931 BHLI information element. It also specifies the proper union member in the **t\_atm\_sap\_appl** structure.

Legal values for the field *t\_atm\_sap\_appl.ID.T\_ISO\_ID* are reserved for specification by ISO. At the time of publication, this was an area of further study for ISO. This field is mapped to octets 6 thru 13 of the Q.2931 BHLI information element.

Legal values for the field *t\_atm\_sap\_appl.ID.vendor\_ID.OUI* are the 24-bit Organizationally Unique Identifiers assigned by the IEEE. This field is mapped to octets 6 thru 8 of the Q.2931 BHLI information element.

Legal values for the field *t\_atm\_sap\_appl.ID.vendor\_ID.app\_ID* are specified by the vendor identified in the *vendor\_ID.OUI* field. The *vendor\_ID.app\_ID* field is mapped to octets 9 thru 12 of the Q.2931 BHLI information element.

Legal values for the field *t\_atm\_sap\_appl.ID.user\_defined\_ID* are 0 through 127. This field is mapped to octets 6 thru 13 of the Q.2931 BHLI information element.

## L.3 Options

Options are formatted according to the structure **t\_opthdr** as described in Chapter 6. A compliant ATM transport provider supports all of the options defined in this appendix. An implementation may restrict the use of any of these options by offering them only in the privileged or read-only mode.

### L.3.1 Signalling-level Options

The protocol level is T\_ATM\_SIGNALING. For this level, Table L-1 shows the options that are defined.

| Option Name      | Type of Option Value          | Legal Option Value | Meaning                            |
|------------------|-------------------------------|--------------------|------------------------------------|
| T_ATM_AAL5       | struct <b>t_atm_aal5</b>      | see text           | ATM adaptation layer 5             |
| T_ATM_TRAFFIC    | struct <b>t_atm_traffic</b>   | see text           | data traffic descriptor            |
| T_ATM_BEARER_CAP | struct <b>t_atm_bearer</b>    | see text           | ATM service capabilities           |
| T_ATM_BHLI       | struct <b>t_atm_bhli</b>      | see text           | higher-layer protocol              |
| T_ATM_BLLI       | struct <b>t_atm_blli</b>      | see text           | lower-layer protocol (1st choice)  |
| T_ATM_DEST_ADDR  | struct <b>t_atm_addr</b>      | see text           | call responders network address    |
| T_ATM_DEST_SUB   | struct <b>t_atm_addr</b>      | see text           | call responder's subaddress        |
| T_ATM_ORIG_ADDR  | struct <b>t_atm_addr</b>      | see text           | call initiators network address    |
| T_ATM_ORIG_SUB   | struct <b>t_atm_addr</b>      | see text           | call initiator's subaddress        |
| T_ATM_CALLER_ID  | struct <b>t_atm_caller_id</b> | see text           | caller's identification attributes |
| T_ATM_CAUSE      | struct <b>t_atm_cause</b>     | see text           | cause of disconnection             |
| T_ATM_QOS        | struct <b>t_atm_qos</b>       | see text           | desired quality of service         |
| T_ATM_TRANSIT    | struct <b>t_atm_transit</b>   | see text           | public carrier transit network     |

Table L-1 Signaling-level Options

These options are all association-related. See Chapter 6 for the difference between options that are association-related and those that are not.

With the exception of option T\_ATM\_CAUSE, each of these options may be negotiated for a connection initiated by the transport user. The results of any such negotiation are signalled to the remote device during connection establishment. The negotiation could be done in either of the following ways:

- the *t\_connect()* function, plus *t\_rcvconnect()* if needed for asynchronous mode
- the *t\_optmgmt()* function, if called before *t\_connect()*.

The first of the above two methods is recommended, since some of the option values can change during connection establishment. Use of the first method returns the updated options to the transport user as parameters of the *t\_connect()* or *t\_rcvconnect()* function call.

For the case of a transport user passively awaiting incoming calls, only the T\_ATM\_AAL5 and T\_ATM\_BLLI options may be negotiated. The results of any such negotiation are signalled to the remote device during connection establishment. The negotiation could be done in either of the following ways:

- the *t\_accept()* function
- the *t\_optmgmt()* function on the responding (not listening) transport endpoint, if called before *t\_accept()*.

The first of the above two methods is recommended, since the options proposed by the initiating party are not associated with the responding transport endpoint prior to *t\_accept()*. Thus, the

`t_optmgmt()` function call could succeed with invalid options, but the connection would be aborted when `t_accept()` is called.

### L.3.2 Absolute Requirements

A request for option T\_ATM\_AAL5, T\_ATM\_TRAFFIC, T\_ATM\_BLLI, or T\_ATM\_QOS is not an absolute requirement<sup>17</sup>. Any other option request is an absolute requirement.

The above listed options can be negotiated down by the ATM transport provider or peer endsystem if necessary and when appropriate.

### L.3.3 Further Remarks

#### T\_ATM\_AAL5

This option is used to signal end-to-end ATM adaptation layer (AAL5) parameters. The description of the information element that is signalled across the ATM network can be found in section 5.4.5.5 of the referenced UNI specification (versions 3.0 and 3.1).

When an incoming connection indication is present, the transport user optionally negotiates this option, which is signalled to the ATM device originating the ATM call. The specific fields within this option that may be modified by the transport user are *forward\_max\_SDU\_size* and *backward\_max\_SDU\_size*. The negotiation could be done in either of the following ways:

- the `t_accept()` function
- the `t_optmgmt()` function, if called before `t_accept()`.

The option value consists of a structure `t_atm_aal5` declared as:

```
struct t_atm_aal5 {
 int32_t forward_max_SDU_size;
 int32_t backward_max_SDU_size;
 int32_t SCS_type;
}
```

Legal values for the field *forward\_max\_SDU\_size* are T\_ATM\_ABSENT, and 0 through  $(2^{16} - 1)$ . This field is mapped to octets 6.1 and 6.2 of the Q.2931 information element.

Legal values for the field *backward\_max\_SDU\_size* are T\_ATM\_ABSENT, and 0 through  $(2^{16} - 1)$ . This field is mapped to octets 7.1 and 7.2 of the Q.2931 information element.

Legal values for the field *SCS\_type* are:

|                       |                             |
|-----------------------|-----------------------------|
| T_ATM_ABSENT          | no indication               |
| T_ATM_NULL            | no SCS layer on top of AAL5 |
| T_ATM_SCS_SSCOP_REL   | SSCOP (assured) SCS         |
| T_ATM_SCS_SSCOP_UNREL | SSCOP (unassured) SCS       |
| T_ATM_SCS_FR          | frame relay SCS             |

This field is mapped to octet 8.1 of the Q.2931 information element. If, as a default, the transport provider causes this Q.2931 information element field to be present in the connection setup, then this field shall default to a value consistent with the parameter *name*

<sup>17</sup> The definition of *absolute requirement* is given in Section 6.3.3 on page 42.



that was specified for `t_open()`:

| <code>t_open()</code> name parameter | default SCS_type     |
|--------------------------------------|----------------------|
| AAL5                                 | T_ATM_NULL           |
| SSCOP/AAL5                           | T_ATM_SSCS_SSCOP_REL |

Note that if all fields of the option have a value of T\_ATM\_ABSENT, then this denotes that the entire information element is not present in the Q.2931 network message.

#### T\_ATM\_TRAFFIC

This option is used to signal the data bandwidth parameters. The description of the information element that is signalled across the ATM network can be found in section 5.4.5.6 of the referenced UNI specification (versions 3.0 and 3.1).

The option value consists of a structure `t_atm_traffic` declared as:

```

struct t_atm_traffic_substruct {
 int32_t PCR_high_priority;
 int32_t PCR_all_traffic;
 int32_t SCR_high_priority;
 int32_t SCR_all_traffic;
 int32_t MBS_high_priority;
 int32_t MBS_all_traffic;
 int32_t tagging;
}

struct t_atm_traffic {
 struct t_atm_traffic_substruct forward;
 struct t_atm_traffic_substruct backward;
 uint8_t best_effort;
}

```

Legal values for the field `forward.PCR_high_priority` are T\_ATM\_ABSENT, and 0 thru  $(2^{**24} - 1)$ . This field is mapped to octets 5.1 thru 5.3 of the Q.2931 information element.

Legal values for the field `forward.PCR_all_traffic` are 0 thru  $(2^{**24} - 1)$ . This field is mapped to octets 7.1 thru 7.3 of the Q.2931 information element.

Legal values for the field `forward.SCR_high_priority` are T\_ATM\_ABSENT, and 0 thru  $(2^{**24} - 1)$ . This field is mapped to octets 9.1 thru 9.3 of the Q.2931 information element.

Legal values for the field `forward.SCR_all_traffic` are T\_ATM\_ABSENT, and 0 thru  $(2^{**24} - 1)$ . This field is mapped to octets 11.1 thru 11.3 of the Q.2931 information element.

Legal values for the field `forward.MBS_high_priority` are T\_ATM\_ABSENT, and 0 thru  $(2^{**24} - 1)$ . This field is mapped to octets 13.1 thru 13.3 of the Q.2931 information element.

Legal values for the field `forward.MBS_all_traffic` are T\_ATM\_ABSENT, and 0 thru  $(2^{**24} - 1)$ . This field is mapped to octets 15.1 thru 15.3 of the Q.2931 information element.

Legal values for the field `forward.tagging` are T\_YES and T\_NO. This field is mapped to octet 18.1 (bit 1) of the Q.2931 information element.

Legal values for the field `backward.PCR_high_priority` are T\_ATM\_ABSENT, and 0 thru  $(2^{**24} - 1)$ . This field is mapped to octets 6.1 thru 6.3 of the Q.2931 information element.

Legal values for the field `backward.PCR_all_traffic` are 0 thru  $(2^{**24} - 1)$ . This field is mapped to octets 8.1 thru 8.3 of the Q.2931 information element.

Legal values for the field *backward.SCR\_high\_priority* are T\_ATM\_ABSENT, and 0 thru  $(2^{**24} - 1)$ . This field is mapped to octets 10.1 thru 10.3 of the Q.2931 information element.

Legal values for the field *backward.SCR\_all\_traffic* are T\_ATM\_ABSENT, and 0 thru  $(2^{**24} - 1)$ . This field is mapped to octets 12.1 thru 12.3 of the Q.2931 information element.

Legal values for the field *backward.MBS\_high\_priority* are T\_ATM\_ABSENT, and 0 thru  $(2^{**24} - 1)$ . This field is mapped to octets 14.1 thru 14.3 of the Q.2931 information element.

Legal values for the field *backward.MBS\_all\_traffic* are T\_ATM\_ABSENT, and 0 thru  $(2^{**24} - 1)$ . This field is mapped to octets 16.1 thru 16.3 of the Q.2931 information element.

Legal values for the field *backward.tagging* are T\_YES and T\_NO. This field is mapped to octet 18.1 (bit 2) of the Q.2931 information element.

Legal values for the field *best\_effort* are T\_YES and T\_NO. This field is mapped to octet 17 of the Q.2931 information element.

#### T\_ATM\_BEARER\_CAP

This option is used to signal the capabilities of the ATM service. The description of the information element that is signalled across the ATM network can be found in section 5.4.5.7 of the referenced UNI specification (versions 3.0 and 3.1).

The option value consists of a structure **t\_atm\_bearer** declared as:

```

struct t_atm_bearer {
 uint8_t bearer_class;
 uint8_t traffic_type;
 uint8_t timing_requirements;
 uint8_t clipping_susceptibility;
 uint8_t connection_configuration;
}

```

Legal values for the field *bearer\_class* (see ITU Recommendation F.811) are:

|               |                |
|---------------|----------------|
| T_ATM_CLASS_A | bearer class A |
| T_ATM_CLASS_C | bearer class C |
| T_ATM_CLASS_X | bearer class X |

This field is mapped to octet 5 (bits 1 thru 5) of the Q.2931 information element.

Legal values for the field *traffic\_type* are:

|            |                               |
|------------|-------------------------------|
| T_ATM_NULL | no indication of traffic type |
| T_ATM_CBR  | constant bit rate             |
| T_ATM_VBR  | variable bit rate             |

This field is mapped to octet 5a (bits 3 thru 5) of the Q.2931 information element.

Legal values for the field *timing\_requirements* are:

|                     |                                |
|---------------------|--------------------------------|
| T_ATM_NULL          | no indication of requirements  |
| T_ATM_END_TO_END    | end-to-end timing required     |
| T_ATM_NO_END_TO_END | end-to-end timing not required |

This field is mapped to octet 5a (bits 1 and 2) of the Q.2931 information element.

Legal values for the field *clipping\_susceptibility* are:

|       |                             |
|-------|-----------------------------|
| T_NO  | not susceptible to clipping |
| T_YES | susceptible to clipping     |

This field is mapped to octet 6 (bits 6 and 7) of the Q.2931 information element.

Legal values for the field *connection\_configuration* are:

|                 |                                |
|-----------------|--------------------------------|
| T_ATM_1_TO_1    | point-to-point connection      |
| T_ATM_1_TO_MANY | point-to-multipoint connection |

This field is mapped to octet 6 (bits 1 and 2) of the Q.2931 information element.

#### T\_ATM\_BHLI

This option is used to signal information about the application that will communicate across the connection. The description of the information element that is signalled across the ATM network can be found in section 5.4.5.8 of the referenced UNI specification (versions 3.0 and 3.1).

For the transport user initiating the connection, the option values pertaining to the specification of the remote ATM protocol address are overwritten with the corresponding parameters of the *t\_connect()* function call. The following fields of this option are set to values found in analogous fields of **t\_atm\_sap\_appl** structure in the destination protocol address when function *t\_connect()* is invoked:

*ID\_type*,  
and all members of union  
*ID*.

The option value consists of a structure **t\_atm\_bhli** declared as:

```
struct t_atm_bhli {
 int32_t ID_type;
 union {
 uint8_t T_ISO_ID [8];
 struct {
 uint8_t OUI [3];
 uint8_t app_ID [4];
 } vendor_ID;
 uint8_t user_defined_ID[8];
 } ID;
}
```

Legal values for the field *ID\_type* are:

|                     |                                             |
|---------------------|---------------------------------------------|
| T_ATM_ABSENT        | application ID is not indicated             |
| T_ATM_ISO_APP_ID    | ISO codepoint                               |
| T_ATM_VENDOR_APP_ID | vendor-specific codepoint                   |
| T_ATM_USER_APP_ID   | identification via a user-defined codepoint |

This field is mapped to octet 5 (bits 1 thru 7) of the Q.2931 information element. The value T\_ATM\_ABSENT denotes that the entire information element is not present in the Q.2931 network message.

Legal values for the field *ID.T\_ISO\_ID* are reserved for specification by ISO. At the time of publication, this was an area of further study for ISO. This field is mapped to octets 6 thru 13 of the Q.2931 information element.

Legal values for the field *ID.vendor\_ID.OUI* are the 24-bit Organization Unique Identifiers assigned by the IEEE. This field is mapped to octets 6 thru 8 of the Q.2931 information element.

Legal values for the field *ID.vendor\_ID.app\_ID* are specified by the vendor identified in the *ID.vendor\_ID.OUI* field. The *ID.vendor\_ID.app\_ID* field is mapped to octets 9 thru 12 of the Q.2931 information element.

Legal values for the field *ID.user\_defined\_ID* are 0 through 127. This field is mapped to octets 6 thru 13 of the Q.2931 BHLI information element.

#### T\_ATM\_BLLI

This option is used to signal information about the layer 2 and layer 3 protocols (if any) that will communicate across the connection. The description of the information element that is signalled across the ATM network can be found in section 5.4.5.9 of the referenced UNI specification (versions 3.0 and 3.1). Note that this option represents the first choice of a transport user initiating a connection across the ATM network. Up to three such choices can be signalled across the ATM network (however, XTI supports only one choice at this time).

For the transport user initiating the connection, the option values pertaining to the specification of the remote ATM protocol address are overwritten with the corresponding parameters of the *t\_connect()* function call. The following fields of this option are set to values found in analogous fields of **t\_atm\_sap\_layer2** and **t\_atm\_sap\_layer3** structures in the destination protocol address when function *t\_connect()* is invoked:

- *layer\_2\_protocol.ID\_type*, and all members of union *layer\_2\_protocol.ID*
- *layer\_3\_protocol.ID\_type*, and all members of union *layer\_3\_protocol.ID*.

When an incoming connection indication is present, the transport user optionally negotiates this option, which is signalled to the ATM device originating the ATM call. The specific fields within this option that may be modified by the transport user are:

- *layer\_2\_protocol.mode*, *layer\_2\_protocol.window\_size*
- *layer\_3\_protocol.mode*, *layer\_3\_protocol.packet\_size*
- *layer\_3\_protocol.window\_size*.

The option negotiation could be done in either of the following ways:

- the *t\_accept()* function
- the *t\_optmgmt()* function, if called before *t\_accept()*.

The transport user accepting the incoming connection indication must perform any negotiation according to the guidelines described in section C.3, Annex C of the referenced UNI specification, versions 3.0 and 3.1. Support of negotiation described in section C.4 of Annex C is for further study.

The option value consists of a structure **t\_atm\_blli** declared as:

```

struct t_atm_blli {
 struct {
 int8_t ID_type;
 union {
 uint8_t simple_ID;
 uint8_t user_defined_ID;
 } ID;
 int8_t mode;
 int8_t window_size;
 } layer_2_protocol;
 struct {
 int8_t ID_type;
 union {
 uint8_t simple_ID;
 int32_t IPI_ID;
 struct {
 uint8_t OUI [3];
 uint8_t PID [2];
 } SNAP_ID;
 uint8_t user_defined_ID;
 } ID;
 int8_t mode;
 int8_t packet_size;
 int8_t window_size;
 } layer_3_protocol;
}

```

Legal values for the field *layer\_2\_protocol.ID\_type* are:

|                 |                                             |
|-----------------|---------------------------------------------|
| T_ATM_ABSENT    | layer 2 identification is not present       |
| T_ATM_SIMPLE_ID | identification via ITU encoding             |
| T_ATM_USER_ID   | identification via a user-defined codepoint |

This field is not mapped to any octets of a Q.2931 information element. Instead, it specifies the proper union member in the **t\_atm\_layer2** structure.

Legal values for the field *layer\_2\_protocol.ID.simple\_ID* are:

|                       |                  |
|-----------------------|------------------|
| T_ATM_BLLI2_I174      | I.1745           |
| T_ATM_BLLI2_Q921      | Q.921            |
| T_ATM_BLLI2_X25_LINK  | X.25, link layer |
| T_ATM_BLLI2_X25_MLINK | X.25, multilink  |
| T_ATM_BLLI2_LAPB      | Extended LAPB    |
| T_ATM_BLLI2_HDLC_ARM  | I.4335, ARM      |
| T_ATM_BLLI2_HDLC_NRM  | I.4335, NRM      |
| T_ATM_BLLI2_HDLC_ABM  | I.4335, ABM      |
| T_ATM_BLLI2_I8802     | I.8802           |
| T_ATM_BLLI2_X75       | X.75             |
| T_ATM_BLLI2_Q922      | Q.922            |
| T_ATM_BLLI2_I7776     | I.7776           |

This field is mapped to octet 6 (bits 1 thru 5) of the Q.2931 information element.

Legal values for the field *layer\_2\_protocol.ID.user\_defined\_ID* are 0 thru 127. This field is mapped to octet 6a (bits 1 thru 7) of the Q.2931 BLLI information element.

Legal values for the field *layer\_2\_protocol.mode* are T\_ATM\_ABSENT, T\_ATM\_BLLI\_NORMAL\_MODE, and T\_ATM\_BLLI\_EXTENDED\_MODE. This field is mapped to octet 6a (bits 6 and 7) of the Q.2931 information element.

Legal values for the field *layer\_2\_protocol.window\_size* are T\_ATM\_ABSENT, and 1 thru 127. This field is mapped to octet 6b (bits 1 thru 7) of the Q.2931 information element.

Legal values for the field *layer\_3\_protocol.ID\_type* are:

|               |                                             |
|---------------|---------------------------------------------|
| T_ATM_ABSENT  | no layer 3 protocol identification          |
| T_ATM_SIMPLE  | ID identification via ITU encoding          |
| T_ATM_IPI_ID  | identification via ISO/IEC TR 9577          |
| T_ATM_SNAP_ID | identification via SNAP                     |
| T_ATM_USER_ID | identification via a user-defined codepoint |

This field is not mapped to any octets of the Q.2931 information element. Instead, it specifies the proper union member in structure **t\_atm\_blli**.

Legal values for the field *layer\_3\_protocol.ID.simple\_ID* are:

|                   |                                  |
|-------------------|----------------------------------|
| T_ATM_BLLI3_X25   | X.25                             |
| T_ATM_BLLI3_I8208 | I.8208                           |
| T_ATM_BLLI3_X223  | X.223                            |
| T_ATM_BLLI3_I8473 | I.8473                           |
| T_ATM_BLLI3_T70   | 70                               |
| T_ATM_BLLI3_I9577 | I.9577 (during connection setup) |

Note that a value of T\_ATM\_BLLI3\_I9577 in this field indicates that the identification of the layer 3 protocol is done in the user (data) plane, as specified in ISO/IEC TR 9577. This field is mapped to octet 7 (bits 1 thru 5) of the Q.2931 information element.

Legal values for the field *layer\_3\_protocol.ID.IPI\_ID* are T\_ATM\_ABSENT, and those values defined by ISO/IEC TR 9577. This field is mapped to octet 7a (bits 1 thru 7) and octet 7b (bit 7) of the Q.2931 information element. Note that the value T\_ATM\_ABSENT is used to signal that the identification of the network-layer protocol is carried with each TSDU in the data plane, according to codepoints defined by ISO/IEC TR 9577.

Legal values for the field *layer\_3\_protocol.ID.SNAP\_ID.OUI* are the 24-bit Organization Unique Identifiers assigned by IEEE. This field is mapped to octets 8.1 thru 8.3 of the Q.2931 information element.

Legal values for the field *layer\_3\_protocol.ID.SNAP\_ID.PID* are defined by the organization identified in the preceding field. This field is mapped to octets 8.4 thru 8.5 of the Q.2931 information element.

Legal values for the field *layer\_3\_protocol.ID.user\_defined\_ID* are 0 thru 127. This field is mapped to octet 7a (bits 1 thru 7) of the Q.2931 BLLI information element.

Legal values for the field *layer\_3\_protocol.mode* are T\_ATM\_ABSENT, T\_ATM\_BLLI\_NORMAL\_MODE, and T\_ATM\_BLLI\_EXTENDED\_MODE. This field is mapped to octet 7a (bits 6 and 7) of the Q.2931 information element.

Legal values for the field *layer\_3\_protocol.packet\_size* are:

|                       |
|-----------------------|
| T_ATM_ABSENT          |
| T_ATM_PACKET_SIZE_16  |
| T_ATM_PACKET_SIZE_32  |
| T_ATM_PACKET_SIZE_64  |
| T_ATM_PACKET_SIZE_128 |

T\_ATM\_PACKET\_SIZE\_256  
 T\_ATM\_PACKET\_SIZE\_512  
 T\_ATM\_PACKET\_SIZE\_1024  
 T\_ATM\_PACKET\_SIZE\_2048  
 T\_ATM\_PACKET\_SIZE\_4096

This field is mapped to octet 7b (bits 1 thru 4) of the Q.2931 information element.

Legal values for the field *layer\_3\_protocol.window\_size* are T\_ATM\_ABSENT, and 1 thru 127. This field is mapped to octet 7c (bits 1 thru 7) of the Q.2931 information element.

#### T\_ATM\_DEST\_ADDR

This option is used to signal the ATM network address of the connections destination. The description of the information element that is signalled across the ATM network can be found in section 5.4.5.11 of the referenced UNI specification (versions 3.0 and 3.1).

For the transport user initiating the connection, the option values pertaining to the specification of the remote ATM protocol address are overwritten with the corresponding parameters of the *t\_connect()* function call. The following fields of this option are set to values found in analogous fields of the destination protocol address when function *t\_connect()* is invoked:

- *address\_format*
- *address\_length*
- *address*.

The option value consists of a structure **t\_atm\_addr** declared as:

```
struct t_atm_addr {
 int8_t address_format;
 uint8_t address_length;
 uint8_t address [20];
}
```

Legal values for the field *address\_format* are T\_ATM\_ENDSYS\_ADDR and T\_ATM\_E164\_ADDR. This field is mapped to octet 5 of the Q.2931 information element.

Legal values for the field *address\_length* are 0 thru 20. This field is not mapped to any octets of a Q.2931 information element. Instead, it specifies the valid number of array elements in field *address*.

Legal values for the field *address* can be found in section 5.1.3 of the referenced UNI specification (versions 3.0 and 3.1). This field is mapped to octets 6 and beyond of the Q.2931 information element.

#### T\_ATM\_DEST\_SUB

This option is used to signal the ATM subaddress of the connections destination. The description of the information element that is signalled across the ATM network can be found in section 5.4.5.12 of the referenced UNI specification (versions 3.0 and 3.1).

The option value consists of a structure **t\_atm\_addr** (see option T\_ATM\_DEST\_ADDR for the structure's declaration). Note that for this option, field *address\_format* must have a value of either T\_ATM\_NSAP\_ADDR or T\_ATM\_ABSENT.

A value of T\_ATM\_ABSENT in field *address\_format* of structure **t\_atm\_addr** indicates that the information element is not present in the Q.2931 signalling message.

#### T\_ATM\_ORIG\_ADDR

This option is used to signal the ATM network address of the connections originator. The

description of the information element that is signalled across the ATM network can be found in section 5.4.5.13 of the referenced **UNI** specification (versions 3.0 and 3.1).

The option value consists of a structure **t\_atm\_addr** (see option T\_ATM\_DEST\_ADDR for the structure's declaration).

A value of T\_ATM\_ABSENT in field *address\_format* of structure **t\_atm\_addr** indicates that the information element is not present in the Q.2931 signalling message.

#### T\_ATM\_ORIG\_SUB

This option is used to signal the ATM subaddress of the connections originator. The description of the information element that is signalled across the ATM network can be found in section 5.4.5.14 of the referenced **UNI** specification (versions 3.0 and 3.1).

The option value consists of a structure **t\_atm\_addr** (see option T\_ATM\_DEST\_ADDR for the structure's declaration). Note that for this option, field *address\_format* must have a value of either T\_ATM\_NSAP\_ADDR or T\_ATM\_ABSENT.

A value of T\_ATM\_ABSENT in field *address\_format* of structure **t\_atm\_addr** indicates that the information element is not present in the Q.2931 signalling message.

#### T\_ATM\_CALLER\_ID

This option is used to signal additional attributes concerning the network address of the connection's originator. The description of the information element that is signalled across the ATM network can be found in section 5.4.5.13 of the referenced **UNI** specification (versions 3.0 and 3.1).

The option value consists of a structure **t\_atm\_caller\_id** declared as:

```
struct t_atm_caller_id {
 int8_t presentation;
 uint8_t screening;
}
```

Legal values for the field *presentation* are:

```
T_ATM_ABSENT
T_ATM_PRES_ALLOWED
T_ATM_PRES_RESTRICTED
T_ATM_PRES_UNAVAILABLE
```

This field is mapped to octet 5a (bits 6 and 7) of the Q.2931 information element.

Legal values for the field *screening* are:

```
T_ATM_ABSENT
T_ATM_USER_ID_NOT_SCREENED
T_ATM_USER_ID_PASSED_SCREEN
T_ATM_USER_ID_FAILED_SCREEN
T_ATM_NETWORK_PROVIDED_ID
```

This field is mapped to octet 5a (bits 1 and 2) of the Q.2931 information element.

#### T\_ATM\_CAUSE

This option is used to signal the cause of a disconnection. The description of the information element that is signalled across the ATM network can be found in section 5.4.5.15 of the referenced **UNI** specification (versions 3.0 and 3.1).

Upon disconnection, the transport user optionally negotiates this option, the results of which are signalled to the remote ATM device. Any such negotiation must be performed



via the `t_optmgmt()` function.

The option value consists of a structure `t_atm_cause` declared as:

```
struct t_atm_cause {
 int8_t coding_standard;
 uint8_r location;
 uint8_r cause_value;
 uint8_r diagnostics [4];
}
```

Legal values for the field `coding_standard` are `T_ATM_ABSENT`, `T_ATM_ITU_CODING`, and `T_ATM_NETWORK_CODING`. This field is mapped to octet 2 (bits 6 and 7) of the Q.2931 information element. The value of `T_ATM_ABSENT` denotes that the entire information element is not present in the Q.2931 network message.

Legal values for the field `location` are:

```
T_ATM_LOC_USER
T_ATM_LOC_LOCAL_PRIVATE_NET
T_ATM_LOC_LOCAL_PUBLIC_NET
T_ATM_LOC_TRANSIT_NET
T_ATM_LOC_REMOTE_PUBLIC_NET
T_ATM_LOC_REMOTE_PRIVATE_NET
T_ATM_LOC_INTERNATIONAL_NET
T_ATM_LOC_BEYOND_INTERWORKING
```

This field is mapped to octet 5 (bits 1 thru 4) of the Q.2931 information element.

Legal values for the field `cause_value` are listed in a full-page table in the referenced UNI specification. This field is mapped to octet 6 (bits 1 thru 7) of the Q.2931 information element.

Legal values for the field `diagnostics` are beyond the scope of this specification. This field is mapped to octets 7 and beyond of the Q.2931 information element.

#### T\_ATM\_QOS

This option is used to signal the desired quality of service. The description of the information element that is signalled across the ATM network can be found in section 5.4.5.18 of the referenced UNI specification (versions 3.0 and 3.1).

The option value consists of a structure `t_atm_qos` declared as:

```
struct t_atm_qos_substruct {
 int32_t coding_standard;
}

struct t_atm_qos {
 int8_t coding_standard;
 struct t_atm_qos_substruct forward;
 struct t_atm_qos_substruct backward;
}
```

Legal values for the field `coding_standard` are `T_ATM_ABSENT`, `T_ATM_ITU_CODING`, and `T_ATM_NETWORK_CODING`. This field is mapped to octet 2 (bits 6 and 7) of the Q.2931 information element. The value of `T_ATM_ABSENT` denotes that the entire information element is not present in the Q.2931 network message.

Legal values for the field *forward.qos\_class* are:

```
T_ATM_QOS_CLASS_0
T_ATM_QOS_CLASS_1
T_ATM_QOS_CLASS_2
T_ATM_QOS_CLASS_3
T_ATM_QOS_CLASS_4
```

This field is mapped to octet 5 of the Q.2931 information element.

Legal values for the field *backward.qos\_class* are the same as those specified for field *forward.qos\_class* above. This field ( *backward.qos\_class*) is mapped to octet 6 of the Q.2931 information element.

#### T\_ATM\_TRANSIT

This option is used to signal the selection of the inter-exchange public carrier. The description of the information element that is signalled across the ATM network can be found in section 5.4.5.22 of the referenced **UNI** specification (versions 3.0 and 3.1).

The option value consists of a structure **t\_atm\_transit** declared as:

```
struct t_atm_transit {
 uint8_t length;
 uint8_t network_id[];
}
```

The field *length* specifies how many characters in array *network\_id* are valid.

Legal values for the field *length* are 0 thru 255. The value of 0 denotes that the entire information element is not present in the Q.2931 network message.

Legal values for the field *network\_id* are beyond the scope of this document. This field is mapped to octets 6 and beyond of the Q.2931 information element.

## L.4 Existing Functions

### *t\_accept()*

The parameter *call*→*addr* is ignored by the ATM transport provider.

The parameter *call*→*opt* is used only to set negotiable ATM connection attributes. These attributes (and XTI options) are T\_ATM\_AAL5 and T\_ATM\_BLLI.

Note that the transport provider must queue data sent via *t\_snd()* until the end-to-end data path is operational.

### *t\_bind()*

When a transport user wishes to passively wait for incoming connect indications through a transport endpoint, then *t\_bind()* is called with parameter *req*→*qlen* having a non-zero value. For ATM, the following additional restrictions apply:

- The ATM protocol address bound to a transport endpoint must conform to the ATM Forum guidelines (referenced document **ATMNAS**) for the specification of a SAP address. The SAP address is a vector that includes fields for the ATM network address (with selector byte), identification of a layer-2 protocol, identification of a layer-3 protocol, and identification of an application.
- The ATM Forum guidelines (referenced document **ATMNAS**) determines uniqueness of an ATM protocol address. If *t\_bind()* is called with a protocol address considered equivalent to an existing bound address, then [TADDRBUSY] is returned.

An implementation may optionally allow at most one application to bind to the “wildcard catch-all” SAP of (T\_ATM\_ANY, T\_ATM\_ANY, T\_ATM\_ANY, T\_ATM\_ANY, T\_ATM\_ANY) for the 5 SAP vector elements.

- ATM transport providers must support a maximum of one outstanding connect indication; therefore, parameter *ret*→*qlen* must contain a value of 1 to reflect the downward negotiation. The transport provider must queue incoming calls that have not been indicated to the transport user, up to some maximum queue size that is implementation dependent.

When a transport user wishes to initiate outgoing connect requests through a transport endpoint, it is recommended that *t\_bind()* be called with parameters *req* and *ret* being null pointers.

If *t\_bind()* is called with parameter *req*→*addr.buf* pointing to an ATM network or protocol address, and parameter *req*→*qlen* having a value of zero; then the ATM provider must perform the following:

- If *t\_bind()* is called with a non-null pointer for parameter *ret*, then upon return, parameter *ret*→*addr.buf* contains the bound address.
- The transport provider saves the bound address for future use in case *t\_getprotaddr()* is called.
- The bound address is overwritten with the destination protocol SAP address if *t\_accept()* is called and this transport endpoint becomes the endpoint upon which the connection is established.
- The bound address is not passed to the peer entity when making an outbound connection.

### *t\_close()*

If the transport user wishes to convey the cause of the disconnection to the remote user, then option T\_ATM\_CAUSE must be negotiated via *t\_optmgmt()* prior to calling *t\_close()*.

*t\_connect()*

Parameter *sndcall*→*addr* specifies the ATM protocol address of the destination transport user. This ATM protocol address must conform to the ATM Forum guidelines (see referenced document **ATMNAS**) for the specification of a SAP address. The SAP address is a vector that includes fields for the ATM network address (with selector byte), identification of a layer-2 protocol, identification of a layer-3 protocol, and identification of an application.

*t\_getinfo()*

The following information parameters are returned:

| Parameters               | Before Call | After Call             |                        |
|--------------------------|-------------|------------------------|------------------------|
|                          |             | AAL-5                  | SSCOP / AAL-5          |
| <i>fd</i>                | x           | /                      | /                      |
| <i>info-&gt;addr</i>     | /           | x                      | x                      |
| <i>info-&gt;options</i>  | /           | x                      | x                      |
| <i>info-&gt;tsdu</i>     | /           | $1 \leq x \leq 65,536$ | $1 \leq x \leq 65,528$ |
| <i>info-&gt;etsdu</i>    | /           | -2                     | -2                     |
| <i>info-&gt;connect</i>  | /           | -2                     | -2                     |
| <i>info-&gt;discon</i>   | /           | -2                     | -2                     |
| <i>info-&gt;servtype</i> | /           | T_COTS                 | T_COTS                 |
| <i>info-&gt;flags</i>    | /           | 0                      | 0                      |

*t\_listen()*

The parameter *call*→*addr* is not the ATM protocol address of the calling transport user; it is merely the network address of the calling user's ATM device. This address may not be sufficient to be the destination protocol address in the *t\_connect()* function.

Upon successful return of *t\_listen()*, the options contained in parameter *call* are those received in the Q.2931 signalling SETUP message from the ATM network.

*t\_look()*

A new event, T\_LEAFCHANGE, is defined which is used by event management to notify a transport user when the status of a leaf has changed on a point-to-multipoint connection.

*t\_open()*

The choice of whether or not the transport provider is requested to implement the SSCOP protocol above AAL-5 in the data plane is conveyed via parameter name. The SSCOP protocol provides a reliable data delivery service. See the notes on function *t\_getinfo()* in this section for a discussion of parameter *info*.

*t\_rcvdis()*

The parameter reason is an 8-bit cause value that is sent across the ATM network in octet 6 of the Q.2931 Cause information element. Additional cause information may be obtained from the T\_ATM\_CAUSE option via *t\_optmgmt()*.

*t\_snd()*

ATM does not support the transport of expedited data. If the transport user sets the T\_EXPEDITED flag in parameter *flags*, then error [TBADFLAG] is returned.

ATM does not support a zero-length TSDU. If parameter *nbytes* is zero, then [TBADDATA] is returned if either:

- the T\_MORE flag (in parameter *flags*) is set

- the T\_MORE flag is not set and this is the first *t\_snd()* call in state T\_DATAXFER
- the T\_MORE flag is not set and the preceding *t\_snd()* call completed a TSDU (T\_MORE flag was not set in preceding call).

When the connection present at the transport endpoint is a leaf on a point-to-multipoint connection, the transport provider returns [TNOTSUPPORT] for *t\_snd()*.

#### *t\_snddis()*

If the transport user wishes to convey the cause of the disconnection to the remote user, then option T\_ATM\_CAUSE must be negotiated via *t\_optmgmt()* prior to calling *t\_snddis()*.

## **L.5 New Functions**

The following new functions are defined in the following man-pages:

*t\_addleaf()*

*t\_removeleaf()*

*t\_rcvleafchange()*

**NAME**

t\_addleaf - add a leaf to a point-to-multipoint connection

**SYNOPSIS**

```
#include <xti.h>

int32_t t_addleaf (fd, leafid, addr)
int32_t fd;
int32_t leafid;
struct netbuf *addr;
```

**DESCRIPTION**

| Parameters         | Before call | After call |
|--------------------|-------------|------------|
| <i>fd</i>          | x           | /          |
| <i>leafid</i>      | x           | /          |
| <i>addr.maxlen</i> | /           | /          |
| <i>addr.len</i>    | x           | /          |
| <i>addr.buf</i>    | x (x)       | /          |

This function enables a transport user to add a leaf to a point-to-multipoint connection. This function can only be issued in the T\_DATAXFER state. The parameter *fd* identifies the local transport endpoint that serves as the root of the point-to-multipoint connection. The parameter *leafid*, provided by the transport user, will be used by subsequent functions *t\_removeleaf()* and *t\_rcvleafchange()* to identify the particular leaf being added. The parameter *addr* is the address of the device being added as a leaf.

The values used for *leafid* may be used simultaneously on point-to-multipoint connections other than the connection indicated by *fd*.

By default, *t\_addleaf()* executes in synchronous mode, and will wait for the remote leaf user's response before returning control to the local user. A successful return (that is, return value of zero) indicates that the requested leaf has been added to the connection. However, if O\_NONBLOCK is set (via *t\_open()* or *fcntl()*), *t\_addleaf()* executes in asynchronous mode. In this case, the call will not wait for the remote leaf user's response, but will return control immediately to the local user and return -1 with *t\_errno* set to [TNODATA] to indicate that the leaf has not yet been added. In this way, the function simply initiates the leaf addition procedure by sending an "add leaf" request to the remote leaf user. The *t\_rcvleafchange()* function is used in conjunction with *t\_addleaf()* to determine the status of the requested leaf addition.

**VALID STATES**

T\_DATAXFER

**ERRORS**

On failure, *t\_errno* is set to one of the following:

[TBADF]

The specified file descriptor does not refer to a transport endpoint. Also used when the connection is not point-to-multipoint.

[TNODATA]

O\_NONBLOCK was set, so the function successfully initiated the leaf addition procedure, but did not wait for a response from the remote leaf user.

**[TBADADDR]**

The specified address was in an incorrect format or contained illegal information.

**[TADDRBUSY]**

This transport provider does not support more than one instance of a particular leaf on a given point-to-multipoint connection. This error indicates that the leaf is already a participant in the point-to-multipoint connection.

**[TOUTSTATE]**

The function was issued in the wrong sequence on the transport endpoint referenced by *fd*.

**[TNOTSUPPORT]**

This function is not supported by the underlying transport provider.

**[TSYSERR]**

A system error has occurred during execution of this function.

**[TLOOK]**

An asynchronous event, which requires attention, has occurred.

**[TPROTO]**

This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI *t\_errno* return code.

**RETURN VALUE**

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t\_errno* is set to indicate the error.

**SEE ALSO**

*t\_removeleaf()*, *t\_rcvleafchange()*.

**ATM PROTOCOL SPECIFICS**

The parameter *addr* is filled with a **t\_atm\_addr** structure and represents the ATM network address of the leaf being added.



**NAME**

t\_removeleaf - drop a leaf from a point-to-multipoint connection

**SYNOPSIS**

```
#include <xti.h>

int32_t t_removeleaf (fd, leafid, reason)
int32_t fd;
int32_t leafid;
int32_t reason;
```

**DESCRIPTION**

| Parameters    | Before call | After call |
|---------------|-------------|------------|
| <i>fd</i>     | x           | /          |
| <i>leafid</i> | x           | /          |
| <i>reason</i> | x           | /          |

This function is used to initiate an abortive removal of a leaf from an already established point-to-multipoint connection. This function can only be issued in the T\_DATAXFER state. The parameter *fd* identifies the local connection endpoint that serves as the root of the point-to-multipoint connection, and parameter *leafid* identifies the leaf that is being removed from the connection. The parameter *reason* specifies the reason for the leaf removal through a protocol-dependent reason code.

**VALID STATES**

T\_DATAXFER

**ERRORS**

On failure, *t\_errno* is set to one of the following:

**[TBADF]**

The specified file descriptor does not refer to a transport endpoint. Also used when the connection is not point-to-multipoint.

**[TOUTSTATE]**

The function was issued in the wrong sequence on the transport endpoint referenced by *fd*.

**[TNOTSUPPORT]**

This function is not supported by the underlying transport provider.

**[TSYSERR]**

A system error has occurred during execution of this function.

**[TLOOK]**

An asynchronous event, which requires attention, has occurred.

**[TPROTO]**

This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI *t\_errno* return code.

**RETURN VALUE**

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t\_errno* is set to indicate the error.

**SEE ALSO**

*t\_addleaf()*, *t\_rcvleafchange()*.

**CAVEATS**

*t\_removeleaf()* is an abortive removal of the leaf. Therefore, *t\_removeleaf()* may cause data previously sent via *t\_snd()* to not be received by the leaf (even if an error is returned).

**ATM PROTOCOL SPECIFICS**

The parameter *reason* is an 8-bit cause value that is sent across the ATM network in octet 6 of the Q.2931 Cause information element.

**NAME**

t\_rcvleafchange - receive an indication about a leaf in a point-to-multipoint connection

**SYNOPSIS**

```
#include <xti.h>

int32_t t_rcvleafchange (fd, change)
int32_t fd;
struct t_leaf_status *change;
```

**DESCRIPTION**

| Parameters           | Before call | After call |
|----------------------|-------------|------------|
| <i>fd</i>            | x           | /          |
| <i>change.leafid</i> | /           | x          |
| <i>change.status</i> | /           | x          |
| <i>change.reason</i> | /           | ?          |

This function is used to determine the status change of a leaf on a point-to-multipoint connection. This function can only be issued in the T\_DATAXFER state. The parameter *fd* identifies the local connection endpoint that serves as the root of the point-to-multipoint connection, and parameter *change* points to a **t\_leaf\_status** structure containing the following members:

```
int32_t leafid;
int32_t status;
int32_t reason;
```

The field *leafid* identifies the leaf whose status has changed, and field *status* specifies the change (either T\_CONNECT or T\_DISCONNECT). When *status* has a value of T\_CONNECT, field *reason* is meaningless. When *status* has a value of T\_DISCONNECT, field *reason* specifies the reason why the leaf was removed from the point-to-multipoint connection or why a pending addleaf failed, through a protocol-dependent reason code.

**VALID STATES**

T\_DATAXFER

**ERRORS**

On failure, *t\_errno* is set to one of the following:

[TBADF]

The specified file descriptor does not refer to a transport endpoint. Also used when the connection is not point-to-multipoint.

[TNODATA]

No leaf change indication currently exists on the specified point-to-multipoint connection.

[TOUTSTATE]

The function was issued in the wrong sequence on the transport endpoint referenced by *fd*.

[TNOTSUPPORT]

This function is not supported by the underlying transport provider.

[TSYSERR]

A system error has occurred during execution of this function.

[TLOOK]

An asynchronous event, which requires attention, has occurred.

**[TPROTO]**

This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI *t\_errno* return code.

**RETURN VALUE**

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *t\_errno* is set to indicate the error.

**SEE ALSO**

*t\_addleaf()*, *t\_removeleaf()*.

**ATM PROTOCOL SPECIFICS**

The parameter *reason* is an 8-bit cause value that is sent across the ATM network in octet 6 of the Q.2931 Cause information element.

## L.6 Implementation Notes

This section maps the functions of XTI onto the primitives specified in referenced document ATMNAS). This mapping information is provided as guidance for the design and development of ATM Transport Providers.

### *t\_accept()*

This function implements the ATM Forums *ATM\_accept\_incoming\_call* primitive. If any options are present in parameter *call*→*opt*, then each of these options is an implementation of the ATM Forum's *ATM\_set\_connection\_attributes* primitive.

The return of this function can be mapped to the following ATM Forum primitives:

- If successful for the setup of a point-to-point connection, the function's return implements the ATM Forums *ATM\_P2P\_call\_active* primitive.
- If successful for the setup of a point-to-multipoint connection, the function's return implements the ATM Forums *ATM\_P2MP\_call\_active* primitive.

### *t\_bind()*

For transport users passively waiting for incoming connect indications, this function implements the ATM Forum primitives *ATM\_prepare\_incoming\_call* and *ATM\_wait\_on\_incoming\_call*. For transport users initiating connect requests, this function implements the ATM Forums *ATM\_prepare\_outgoing\_call* primitive.

Note that parameter *queue\_size* of the ATM Forum's *ATM\_prepare\_incoming\_call* primitive is different to XTI's *qlen*: The parameter *queue\_size* is the maximum number of incoming calls that have arrived but have not yet been presented to the transport user, while the parameter *qlen* is the maximum number of incoming calls that have been presented to the transport user but not yet accepted. Therefore, *queue\_size* cannot be set via the *t\_bind()* function; the value is determined by the transport provider..

### *t\_close()*

Normally, this function is called from the T\_UNBND state; in this case, there is no needed mapping to the ATM Forum primitives. XTI also allows this function to be legally called from other states as well. When that occurs, this function implements the ATM Forum's *ATM\_abort\_connection* primitive. However, the transport user cannot specify the cause of the aborted connection.

### *t\_connect()*

The invocation of this function implements the ATM Forum's *ATM\_connect\_outgoing\_call* primitive. If any options are present in parameter *sndcall*→*opt*, then each of these options is an implementation of the ATM Forum's *ATM\_set\_connection\_attributes* primitive. If any options are present in parameter *rcvcall*→*optc*, then each of these options is an implementation of the ATM Forum's *ATM\_query\_connection\_attributes* primitive.

Additionally, when the transport endpoint is in synchronous mode, the return of this function can be mapped to the following ATM Forum primitives:

- If successful for the setup of a point-to-point connection, the function's return implements the ATM Forums *ATM\_P2P\_call\_active* primitive.
- If successful for the setup of a point-to-multipoint connection, the function's return implements the ATM Forums *ATM\_P2MP\_call\_active* primitive.
- If unsuccessful for the setup of a connection, the function's return has no needed mapping to an ATM Forum primitive.

*t\_listen()*

The successful return of this function is an implementation of the ATM Forum's *ATM\_arrival\_of\_incoming\_call* primitive.

The ATM network address (of the calling user) returned in parameter *call→addr* is an implementation of the ATM Forum's *ATM\_query\_connection\_attributes* primitive, where the connection attribute (and XTI option) being queried is T\_ATM\_ORIG\_ADDR. If any options are present in parameter *call→opt*, then each of these options is an implementation of the ATM Forum's *ATM\_query\_connection\_attributes* primitive.

*t\_open()*

This function implements the ATM Forums *ATM\_associate\_endpoint* primitive.

*t\_optmgmt()*

For any option values that the transport user attempts to negotiate (present in parameter *req→opt* when parameter *req→flags* equals T\_NEGOTIATE), each of these options is an implementation of the ATM Forum's *ATM\_set\_connection\_attributes* primitive. For any options present in parameter *ret→opt*, each of these options is an implementation of the ATM Forum's *ATM\_query\_connection\_attributes* primitive.

*t\_rcv()*

The invocation of this function implements the ATM Forum's *ATM\_receive\_data* primitive. Note that the ATM Forum's primitive can be further classified as either a "polling implementation" or a "blocking implementation". XTI's synchronous mode corresponds to "blocking implementation"; asynchronous mode corresponds to "polling implementation".

*t\_rcvconnect()*

The return of this function can be mapped to the following ATM Forum primitives:

- If successful for the setup of a point-to-point connection, the function's return implements the ATM Forum's *ATM\_P2P\_call\_active* primitive.
- If successful for the setup of a point-to-multipoint connection, the function's return implements the ATM Forum's *ATM\_P2MP\_call\_active* primitive.
- If unsuccessful for the setup of a connection, the function's return has no needed mapping to an ATM Forum primitive.

If any options are present in parameter *call→opt*, then each of these options is an implementation of the ATM Forum's *ATM\_query\_connection\_attributes* primitive.

*t\_rcvdis()*

The successful return of this function implements the ATM Forums *ATM\_call\_release* (indication) primitive.

*t\_snd()*

The invocation of this function implements the ATM Forums *ATM\_send\_data* primitive.

*t\_snddis()*

Depending on the XTI state, invocation of this function implements the following ATM Forum primitives:

|                         |                                                                  |
|-------------------------|------------------------------------------------------------------|
| state T_DATAXFER:       | <i>ATM_call_release</i> (request) or <i>ATM_abort_connection</i> |
| state T_INCON:          | <i>ATM_reject_incoming_call</i>                                  |
| all other valid states: | <i>ATM_abort_connection</i> .                                    |

*t\_unbind()*

There is no ATM Forum primitive to which this function can be mapped. From the

perspective of the ATM Forum state machine, this function does the following:

- moves any future outgoing connections on this endpoint from state A2 to A1
- moves any future incoming connections on this endpoint from state A4 to A1.

*t\_addleaf()*

This function implements the ATM Forum's *ATM\_add\_party* primitive. If successful in the addition of the leaf, the function's return implements the ATM Forum's *ATM\_add\_party\_success* primitive.

*t\_removeleaf()*

This function implements the ATM Forum's *ATM\_drop\_party* (Request) primitive.

*t\_rcvleafchange()*

This function implements the following ATM Forum primitives:

*ATM\_add\_party\_success*

*ATM\_add\_party\_reject*

*ATM\_drop\_party*(Indication)





# ATM Transport Protocol Information for Sockets

## M.1 General

This appendix is a Preliminary Specification. It describes the protocol-specific information that is relevant for ATM transport providers.

The following general notes apply:

- This version of Sockets supports a subset of the functions specified by the ATM Forum as the User-Network Interface (see referenced document **UNI**), versions 3.0 and 3.1.
- Frequent reference is made to the ATM Forum's "Native ATM Services: Semantic Description, Version 1.0", identified as referenced document **ATMNAS**.
- ATM supports connection-oriented sockets only. Specifically, the socket type is `SOCK_SEQPACKET`. At the current time, only AAL-5 message mode is supported. Each sockets message is carried across the network in the payload field of a single AAL-5 PDU. AAL1 and User-defined AAL are not supported.
- ATM supports both reliable and unreliable data transport services. Reliable service is obtained by specifying `ATM_PROTO_SSCOP` as the protocol when calling `socket()`. Unreliable service is obtained by specifying `ATM_PROTO_AAL5` as the protocol when calling `socket()`.
- ATM addresses (both network and protocol) are defined in Appendix L( ATM Transport Protocol Information for XTI) of this specification.
- When a transport user passively waits for incoming connect indications, the ATM protocol address bound to the socket must conform to the ATM Forum guidelines (see referenced document **ATMNAS**) for the specification of a SAP address. The SAP address is a vector that includes fields for the ATM network address (with selector byte), identification of a layer-2 protocol, identification of a layer-3 protocol, and identification of an application
- AAL and BLLI negotiation is not supported since `accept()` contains no provision for examining and modifying connection attributes.
- The ATM transport provider does not support ATM PVCs.
- The ATM transport provider does not support a mechanism to specify the congestion indication bit and the user-user byte in AAL5.

## M.2 Existing Functions

### *accept()*

The parameter *address*, if non-null, points to the ATM network address of the peer connecting socket. The ATM protocol address of the peer connecting socket is unavailable.

Note that the transport protocol must queue data sent via *send()*, *sendmsg()*, or *sendto()* until the end-to-end data path is operational.

### *bind()*

When a transport user wishes to passively wait for incoming connections through a socket, the following additional restrictions apply for ATM transports:

- The ATM protocol address bound to a socket must conform to the ATM Forum guidelines (referenced document **ATMNAS**) for the specification of a SAP address. The SAP address is a vector that includes fields for the ATM network address (with selector byte), identification of a layer-2 protocol, identification of a layer-3 protocol, and identification of an application.
- The ATM Forum guidelines (referenced document **ATMNAS**) determines uniqueness of an ATM protocol address. If *bind()* is called with a protocol address considered equivalent to an existing bound address, then [EADDRINUSE] is returned.

When a transport user wishes to initiate outgoing connections through a socket, then *bind()* should not be called.

### *close()*

If the transport user wishes to convey the cause of the disconnection to the peer socket, then socket option T\_ATM\_CAUSE must be set before function *close()* is invoked.

### *connect()*

Parameter *address* specifies the ATM protocol address of the destination transport user. This ATM protocol address must conform to the ATM Forum guidelines (referenced document **ATMNAS**) for the specification of a SAP address. The SAP address is a vector that includes fields for the ATM network address (with selector byte), identification of a layer-2 protocol, identification of a layer-3 protocol, and identification of an application.

### *getpeername()*

The parameter *address*, if meaningful, points to one of the following:

- an ATM network address, for sockets in which the transport user passively waits for an incoming connection
- an ATM protocol address, for sockets in which the transport user actively initiates an outgoing connection.

### *getsockname()*

The parameter *address*, if meaningful, points to one of the following:

- an ATM protocol address, for sockets in which the transport user passively waits for an incoming connection
- an ATM network address, for sockets in which the transport user actively initiates an outgoing connection.

*getsockopt()*

ATM transport protocol options use T\_ATM\_SIGNALING for the parameter level. The ATM-specific options are listed in the table below. Most of these options are defined in Appendix L in this specification. Option T\_ATM\_LEAF\_IND is defined in Section M.3 on page 399.

| Option Name      | Type of Option Value   | Meaning                            |
|------------------|------------------------|------------------------------------|
| T_ATM_AAL5       | struct t_atm_aal5      | ATM adaptation layer 5             |
| T_ATM_TRAFFIC    | struct t_atm_traffic   | data traffic descriptor            |
| T_ATM_BEARER_CAP | struct t_atm_bearer    | ATM service capabilities           |
| T_ATM_BHLI       | struct t_atm_bhli      | higher-layer protocol              |
| T_ATM_BLLI       | struct t_atm_blli      | lower-layer protocol (1st choice)  |
| T_ATM_DEST_ADDR  | struct t_atm_addr      | call responders network address    |
| T_ATM_DEST_SUB   | struct t_atm_addr      | call responder's subaddress        |
| T_ATM_ORIG_ADDR  | struct t_atm_addr      | call initiators network address    |
| T_ATM_ORIG_SUB   | struct t_atm_addr      | call initiator's subaddress        |
| T_ATM_CALLER_ID  | struct t_atm_caller_id | caller's identification attributes |
| T_ATM_CAUSE      | struct t_atm_cause     | cause of disconnection             |
| T_ATM_QOS        | struct t_atm_qos       | desired quality of service         |
| T_ATM_TRANSIT    | struct t_atm_transit   | public carrier transit network     |
| T_ATM_LEAF_IND   | struct t_atm_leaf_ind  | indication of leaf status change   |

*recv()*

The parameter *flags* must equal a value of 0 or MSG\_PEEK.

*recvfrom()*

Note that ATM protocol supports only connection-oriented sockets.

*recvmsg()*

Note that ATM protocol supports only connection-oriented sockets.

*send()*

The parameter *flags* must equal a value of 0.

When the connection present at the socket is a leaf on a point-to-multipoint connection, any invocation of function *send()* results in [EOPNOTSUPP].

*sendmsg()*

Note that ATM protocol supports only connection-oriented sockets.

*sendto()*

Note that ATM protocol supports only connection-oriented sockets.

*setsockopt()*

ATM transport protocol options use SOL\_ATM\_SIGNALING for the parameter level. The ATM-specific options are listed in the table below. Most of these options are defined in Appendix L. Options T\_ATM\_ADD\_LEAF and T\_ATM\_DROP\_LEAF are defined in Section M.3 on page 399.

| Option Name      | Type of Option Value   | Meaning                            |
|------------------|------------------------|------------------------------------|
| T_ATM_AAL5       | struct t_atm_aal5      | ATM adaptation layer 5             |
| T_ATM_TRAFFIC    | struct t_atm_traffic   | data traffic descriptor            |
| T_ATM_BEARER_CAP | struct t_atm_bearer    | ATM service capabilities           |
| T_ATM_BHLI       | struct t_atm_bhli      | higher-layer protocol              |
| T_ATM_BLLI       | struct t_atm_blli      | lower-layer protocol (1st choice)  |
| T_ATM_DEST_ADDR  | struct t_atm_addr      | call responders network address    |
| T_ATM_DEST_SUB   | struct t_atm_addr      | call responder's subaddress        |
| T_ATM_ORIG_ADDR  | struct t_atm_addr      | call initiators network address    |
| T_ATM_ORIG_SUB   | struct t_atm_addr      | call initiator's subaddress        |
| T_ATM_CALLER_ID  | struct t_atm_caller_id | caller's identification attributes |
| T_ATM_CAUSE      | struct t_atm_cause     | cause of disconnection             |
| T_ATM_QOS        | struct t_atm_qos       | desired quality of service         |
| T_ATM_TRANSIT    | struct t_atm_transit   | public carrier transit network     |
| T_ATM_ADD_LEAF   | struct t_atm_add_leaf  | add a leaf to a connection         |
| T_ATM_DROP_LEAF  | struct t_atm_drop_leaf | remove a leaf from a connection    |

*shutdown()*

This function has a local scope only; no control information is transferred across the ATM network.

*socket()*

The following parameter values should be used for ATM protocol transports:

- *domain* = AF\_ATM
- *type* = SOCK\_SEQPACKET
- *protocol* = ATM\_PROTO\_AAL5 or ATM\_PROTO\_SSCOP

If parameter *protocol* signifies use of the SSCOP protocol, it is recommended that the transport provider use the default of T\_ATM\_SSCS\_SSCOP\_REL in field *SSCS\_type* of option T\_ATM\_AAL5.

*socketpair()*

This function is not supported by an ATM domain. If the function is invoked, the transport provider returns [EOPNOTSUPP].

## M.3 Point-to-Multipoint Connections

To support point-to-multipoint connections, procedures are defined for the addition and removal of other parties (called *leaves* in this specification) to the connection. These procedures are performed only by the root of the connection. This section specifies how socket options are used within X/Sockets to accomplish these functions. For more information, see section 4.2 (“SVC Provisioning”) of referenced document **ATMNAS**.

### M.3.1 Adding a Leaf

#### T\_ATM\_ADD\_LEAF

This option is used to request that a leaf be added to an existing point-to-multipoint connection. Note that this option is “write-only”; it is a valid parameter of *setsockopt()*, but not *getsockopt()*.

The option value consists of a structure **t\_atm\_add\_leaf** declared as:

```
struct t_atm_add_leaf {
 int32_t leaf_ID;
 struct t_atm_addr leaf_address;
}
```

The field *leaf\_ID* is a unique identifier for the leaf on the point-to-multipoint connection. When the connection is initially setup via *connect()*, the remote device is implicitly assigned a leaf identifier of zero. Other leaves added to the connection are assigned a non-zero value by the transport user. The leaf identifier is used as a correlator for the following subsequent operations:

- receiving indications regarding changes in the leaf’s status
- removing the leaf from the point-to-multipoint connection.

The values used for *leaf\_ID* may be used simultaneously on point-to-multipoint connections other than the connection indicated by parameter *socket*.

Legal values for the field *leaf\_ID* are 1 thru  $(2^{15} - 1)$ . This field is mapped to octets 6 thru 6.1 of the ATM Forum’s “Endpoint Reference” information element, defined in section 5.4.8.1 in referenced document **ATMNAS**.

The field *leaf\_address* specifies the network address of the device being added as a leaf of the connection. Legal values for the field *leaf\_address* are the same as those defined for option T\_ATM\_DEST\_ADDR in Appendix L (ATM Transport Protocol Information for XTI).

A successful return of this function does not imply that the leaf has been successfully added; rather, it means that the leaf addition request is being processed by the network. Completion (success or failure) of the leaf addition is indicated at a later time. An unsuccessful return of this function with *errno* set to [EINVAL] is additionally used to indicate the following error conditions:

- the socket type of the specified socket is not a point-to-multipoint connection
- the connection has been aborted
- the connection was forcibly closed by a peer.

### M.3.2 Removing a Leaf

#### T\_ATM\_DROP\_LEAF

This option is used to initiate an abortive removal of a leaf from an already established point-to-multipoint connection. When this option is invoked, data previously sent via *send()* may not be received by the leaf. Note that this option is “write-only”; it is a valid parameter of *setsockopt()*, but not *getsockopt()*.

The option value consists of a structure **t\_atm\_drop\_leaf** declared as:

```
struct t_atm_drop_leaf {
 int32_t leaf_ID;
 int32_t reason;
}
```

The field *leaf\_ID* identifies the leaf that is being removed from the connection. When the connection is initially setup via *connect()*, the remote device is implicitly assigned a leaf identifier of zero. Other leaves added to the connection are assigned a non-zero value by the transport user when the leaf is added to the connection.

Legal values for the field *leaf\_ID* are 0 thru  $(2^{15} - 1)$ . This field is mapped to octets 6 thru 6.1 of the ATM Forum's "Endpoint Reference" information element, defined in section 5.4.8.1 in the referenced document **ATMNAS**.

The field *reason* specifies the reason for the leaf removal. This field is mapped to octet 6 of the Q.2931 Cause information element.

A successful return of this function does not imply that the leaf has been completely removed; rather, it means that the leaf removal request is being processed by the network. Eventual completion of the leaf removal is assumed. An unsuccessful return of this function with *errno* set to [EINVAL] is additionally used to indicate the following error conditions:

- the socket type of the specified socket is not a point-to-multipoint connection
- the connection has been aborted
- the connection was forcibly closed by a peer.

### M.3.3 Receiving Indication of a Change in Leaf Status

There must exist an operating system-specific method of notifying the transport user that a status change has occurred for at least one leaf on a given point-to-multipoint connection. When the transport user detects that such a change has occurred for socket, the transport user invokes *getsockopt()* with option parameter T\_ATM\_LEAF\_IND . This action should be repeated until field *status* of the returned option value has a value of T\_LEAF\_NOCHANGE.

#### T\_ATM\_LEAF\_IND

This option is used to receive indications of status changes for leaves on an already established point-to-multipoint connection. Note that this option is "read-only"; it is a valid parameter of *getsockopt()*, but not *setsockopt()*.

The option value consists of a structure **t\_atm\_leaf\_ind** declared as:

```
struct t_atm_leaf_ind {
 int32_t status;
 int32_t leaf_ID;
 int32_t reason;
}
```

The field *status* will contain one of these values:

- **T\_LEAF\_NOCHANGE**  
no status change has occurred for any leaf associated with this socket.
- **T\_LEAF\_CONNECTED**  
a previous request to add a leaf has been successfully completed
- **T\_LEAF\_DISCONNECTED**  
signifies that one of these occurred:
  - a previous request to add a leaf was unsuccessful
  - the leaf has initiated a release from the connection
  - the network has dropped the leaf from the connection.

The field *leaf\_ID* identifies the leaf whose status has changed. When the connection is initially setup via *connect()*, the remote device is implicitly assigned a leaf identifier of zero. Other leaves added to the connection are assigned a non-zero value by the transport user when the leaf is added to the connection.

Legal values for the field *leaf\_ID* are 0 thru  $(2^{15} - 1)$ . This field is mapped to octets 6 thru 6.1 of the ATM Forum's "Endpoint Reference" information element, defined in section 5.4.8.1 of referenced document **ATMNAS**.

When field *status* has a value of **T\_LEAF\_DISCONNECTED**, the field *reason* specifies the reason for the leaf removal. This field is mapped to octet 6 of the Q.2931 Cause information element.

An unsuccessful return of this function with *errno* set to [EINVAL] is additionally used to indicate the following error conditions:

- the socket type of the specified socket is not a point-to-multipoint connection
- the connection has been aborted
- the connection was forcibly closed by a peer.

## M.4 Implementation Notes

This section maps the functions of X/Sockets onto the primitives specified in referenced document *ATMNAS*. This mapping is provided as guidance for the design and development of ATM transport providers.

### *accept()*

The successful return of this function can be mapped to the following ATM Forum primitives:

- *ATM\_wait\_on\_incoming\_call* primitive
- *ATM\_arrival\_of\_incoming\_call* primitive
- *ATM\_accept\_incoming\_call* primitive
- *ATM\_P2P\_call\_active* primitive (for point-to-point connections)
- *ATM\_P2MP\_call\_active* primitive (for point-to-multipoint connections)

### *bind()*

For transport users passively waiting for incoming connections, this function implements the ATM Forum's *ATM\_prepare\_incoming\_call* primitive. Note that parameter *queue\_size* of the ATM Forum's *ATM\_prepare\_incoming\_call* primitive is provided via X/Socket's *listen()* function.

### *close()*

The invocation of this function implements the ATM Forum's *ATM\_abort\_connection* or *ATM\_call\_release* (request) primitive.

### *connect()*

The invocation of this function implements the ATM Forum's *ATM\_connect\_outgoing\_call* primitive. Additionally, the successful return of this function can be mapped to the following ATM Forum primitives:

- *ATM\_P2P\_call\_active* primitive (point-to-point connection)
- *ATM\_P2MP\_call\_active* primitive (point-to-multipoint connection)

### *getpeername()*

This function implements the ATM Forum's *ATM\_query\_connection\_attributes* primitive.

### *getsockname()*

This function implements the ATM Forum's *ATM\_query\_connection\_attributes* primitive.

### *getsockopt()*

This function implements the ATM Forum's *ATM\_query\_connection\_attributes* primitive. For point-to-multipoint connections, this function also function implements the ATM Forum's *ATM\_add\_party\_success*, *ATM\_add\_party\_reject* and *ATM\_drop\_party* (Indication) primitives.

### *listen()*

The invocation of this function partially implements the ATM Forum's *ATM\_wait\_on\_incoming\_call* primitive, since the queue is enabled. Note that the ATM Forum's primitive can be further classified in this case as a "polling implementation".

### *recv()*

The invocation of this function implements the ATM Forum's *ATM\_receive\_data* primitive.

### *recvfrom()*

The invocation of this function implements the ATM Forum's *ATM\_receive\_data* primitive.



*recvmsg()*

The invocation of this function implements the ATM Forum's *ATM\_receive\_data* primitive.

*send()*

The invocation of this function implements the ATM Forum's *ATM\_send\_data* primitive.

*sendmsg()*

The invocation of this function implements the ATM Forum's *ATM\_send\_data* primitive.

*sendto()*

The invocation of this function implements the ATM Forum's *ATM\_send\_data* primitive.

*setsockopt()*

This function implements the ATM Forum's *ATM\_set\_connection\_attributes* primitive. For point-to-multipoint connections, this function also function implements the ATM Forum's *ATM\_add\_party* and *ATM\_drop\_party* (Request) primitives.

*socket()*

This function implements the ATM Forum's *ATM\_associate\_endpoint* primitive.



# ATM Transport Headers

This Appendix is a Preliminary Specification. It presents the header files `<xti_atm.h>`, `<netatm/atm.h>` and `<_atm_common.h>`, and the proposed additions to `<xti.h>` and `<sys/socket.h>`.

## N.1 Proposed Additions to `<xti.h>`

```
#define T_LEAFCHANGE 0x0400 /* status of a leaf has changed */

/*
 * The following are new XTI library functions:
 */
/* XTI Library Function: t_addleaf - add a leaf */
extern int t_addleaf(int32_t, int32_t, struct netbuf *);
/* XTI Library Function: t_removeleaf - remove a leaf */
extern int t_removeleaf(int32_t, int32_t, int32_t);
/* XTI Library Function: t_rcvleafchange - acknowledge
 * receipt of a leaf change indication */
extern int t_rcvleafchange(int32_t, struct t_leaf_status *);
```

## N.2 Proposed Additions to `<sys/socket.h>`

Add the following macro with distinct integral value in `<sys/socket.h>` in the same name space as `AF_UNSPEC`:

```
#define AF_ATM /* ATM transport sockets */
```

## N.3 `<xti_atm.h>`

```
#include <_atm_common.h>

/*
 * ATM Levels
 */
#define T_ATM_SIGNALING 0x5301 /* options signalled across UNI */
/* value is not mandatory */
```

## N.4 <netatm/atm.h>

```
#include <_atm_common.h>

/*
 * First, the protocol constants.
 */
#define ATM_PROTO_AAL5 0x5301 /* AAL type 5 protocol */
#define ATM_PROTO_SSCOP 0x5302 /* SSCOP protocol */

/*
 * ATM commonly used constants
 */
#define T_YES 1
#define T_NO 0
#define T_LEAF_NOCHANGE 0 /* value not mandatory */
#define T_LEAF_CONNECTED 1 /* value not mandatory */
#define T_LEAF_DISCONNECTED 2 /* value not mandatory */

/* ATM-SPECIFIC OPTIONS */

/*
 * ATM signalling-level options
 */
#define T_ATM_SIGNALING 0x5301 /* options signalled across UNI */
/* value is not mandatory */
#define T_ATM_ADD_LEAF 0x21 /* add leaf to connection */
/* value is not mandatory */
#define T_ATM_DROP_LEAF 0x22 /* remove leaf from connection */
/* value is not mandatory */
#define T_ATM_LEAF_IND 0x23 /* indication of leaf status */
/* value is not mandatory */

/*
 * ATM data structures used for point-to-multipoint connection support
 */
struct t_atm_add_leaf {
 int32_t leaf_ID;
 struct t_atm_addr leaf_address;
};

struct t_atm_drop_leaf {
 int32_t leaf_ID;
 int32_t reason;
};

struct t_atm_leaf_ind {
 int32_t status;
 int32_t leaf_ID;
 int32_t reason;
};
```

**N.5 <\_atm\_common.h>**

This file should not be included by applications. It is provided so that symbols in it can be exposed through XTI and Socket headers.

```

/*
 * For the purposes of conformance testing, it may be assumed that any
 * constant values defined in these header files are mandatory, unless
 * the constant:
 * 1. defines an option or options level
 * 2. is accompanied by a comment that specifies the value is
 * not mandatory.
 */

/*
 * Leaf status structure.
 */
struct t_leaf_status {
 int32_t leafid; /* leaf identifier */
 int32_t status; /* current status */
 int32_t reason; /* reason for leaf removal */
};

/*
 * ATM commonly used constants
 */
#define T_ATM_ABSENT (-1)
#define T_ATM_PRESENT (-2)
#define T_ATM_ANY (-3)
/*
 * In the 3 constants defined immediately above, the specific value
 * is not mandatory, but any conforming value must be negative.
 */
#define T_ATM_NULL 0
#define T_ATM_ENDSYS_ADDR 1 /* value is not mandatory */
#define T_ATM_NSAP_ADDR 2 /* value is not mandatory */
#define T_ATM_E164_ADDR 3 /* value is not mandatory */
#define T_ATM_ITU_CODING 0
#define T_ATM_NETWORK_CODING 3

/* ATM-SPECIFIC ADDRESSES */

/*
 * ATM protocol address structure
 */
struct t_atm_sap {
 struct t_atm_sap_addr {
 int8_t SVE_tag_addr;
 int8_t SVE_tag_selector;
 uint8_t address_format;
 uint8_t address_length;
 uint8_t address [20];
 } t_atm_sap_addr;

 struct t_atm_sap_layer2 {
 int8_t SVE_tag;
 uint8_t ID_type;
 }
};

```

```

 union {
 uint8_t simple_ID;
 uint8_t user_defined_ID;
 } ID;
 } t_atm_sap_layer2;

 struct t_atm_sap_layer3 {
 int8_t SVE_tag;
 uint8_t ID_type;
 union {
 uint8_t simple_ID;
 int32_t IPI_ID;
 struct {
 uint8_t OUI [3];
 uint8_t PID [2];
 } SNAP_ID;
 uint8_t user_defined_ID;
 } ID;
 } t_atm_sap_layer3;

 struct t_atm_sap_appl {
 int8_t SVE_tag;
 uint8_t ID_type;
 union {
 uint8_t ISO_ID [8];
 struct {
 uint8_t OUI [3];
 uint8_t app_ID [4];
 } vendor_ID;
 uint8_t user_defined_ID [8];
 } ID;
 } t_atm_sap_appl;
}

/* ATM-SPECIFIC OPTIONS */

/*
 * ATM signalling-level options
 */
#define T_ATM_AAL5 0x1 /* ATM adaptation layer 5 */
#define T_ATM_TRAFFIC 0x2 /* data traffic descriptor */
#define T_ATM_BEARER_CAP 0x3 /* ATM service capabilities */
#define T_ATM_BHLI 0x4 /* higher-layer protocol */
#define T_ATM_BLLI 0x5 /* lower-layer protocol */
#define T_ATM_DEST_ADDR 0x6 /* call responder's address */
#define T_ATM_DEST_SUB 0x7 /* call responder's subaddress */
#define T_ATM_ORIG_ADDR 0x8 /* call initiator's address */
#define T_ATM_ORIG_SUB 0x9 /* call initiator's subaddress */
#define T_ATM_CALLER_ID 0xa /* caller's ID attributes */
#define T_ATM_CAUSE 0xb /* cause of disconnection */
#define T_ATM_QOS 0xc /* desired quality of service */
#define T_ATM_TRANSIT 0xd /* choice of public carrier */

/*
 * T_ATM_AAL5 structure
 */
struct t_atm_aal5 {
 int32_t forward_max_SDU_size;

```

```

 int32_t backward_max_SDU_size;
 int32_t SCS_type;
};

/*
 * T_ATM_AAL5 values
 */
#define T_ATM_SSCS_SSCOP_REL 1
#define T_ATM_SSCS_SSCOP_UNREL 2
#define T_ATM_SSCS_FR 4

/*
 * T_ATM_TRAFFIC structure
 */
struct t_atm_traffic_substruct {
 int32_t PCR_high_priority;
 int32_t PCR_all_traffic;
 int32_t SCR_high_priority;
 int32_t SCR_all_traffic;
 int32_t MBS_high_priority;
 int32_t MBS_all_traffic;
 int32_t tagging;
}

struct t_atm_traffic {
 struct t_atm_traffic_substruct forward;
 struct t_atm_traffic_substruct backward;
 uint8_t best_effort;
}

/*
 * T_ATM_BEARER_CAP structure
 */
struct t_atm_bearer {
 uint8_t bearer_class;
 uint8_t traffic_type;
 uint8_t timing_requirements;
 uint8_t clipping_susceptibility;
 uint8_t connection_configuration;
}

/*
 * T_ATM_BEARER_CAP values
 */
#define T_ATM_CLASS_A 0x01 /* bearer class A */
#define T_ATM_CLASS_C 0x03 /* bearer class C */
#define T_ATM_CLASS_X 0x10 /* bearer class X */
#define T_ATM_CBR 0x01 /* constant bit rate */
#define T_ATM_VBR 0x02 /* variable bit rate */
#define T_ATM_END_TO_END 0x01 /* end-to-end timing required */
#define T_ATM_NO_END_TO_END 0x02 /* end-to-end timing not required */
#define T_ATM_1_TO_1 0x00 /* point-to-point connection */
#define T_ATM_1_TO_MANY 0x01 /* point-to-multipoint connection */

/*
 * T_ATM_BHLI structure
 */
struct t_atm_bhli {
 int32_t ID_type;

```

```

 union {
 uint8_t ISO_ID [8];
 struct {
 uint8_t OUI [3];
 uint8_t app_ID [4];
 } vendor_ID;
 uint8_t user_defined_ID [8];
 } ID;
}

/*
 * T_ATM_BHLI values
 */
#define T_ATM_ISO_APP_ID 0 /* ISO codepoint */
#define T_ATM_VENDOR_APP_ID 3 /* vendor-specific codepoint */
#define T_ATM_USER_APP_ID 1 /* user-specific codepoint */

/*
 * T_ATM_BLLI structure
 */
struct t_atm_blli {
 struct {
 int8_t ID_type;
 union {
 uint8_t simple_ID;
 uint8_t user_defined_ID;
 } ID;
 int8_t mode;
 int8_t window_size;
 } layer_2_protocol;
 struct {
 int8_t ID_type;
 union {
 uint8_t simple_ID;
 int32_t IPI_ID;
 struct {
 uint8_t OUI [3];
 uint8_t PID [2];
 } SNAP_ID;
 uint8_t user_defined_ID;
 } ID;
 int8_t mode;
 int8_t packet_size;
 int8_t window_size;
 } layer_3_protocol;
}

/*
 * T_ATM_BLLI values
 */
#define T_ATM_SIMPLE_ID 1 /* ID via ITU encoding */
#define T_ATM_IPI_ID 2 /* ID via ISO/IEC TR 9577 */
#define T_ATM_SNAP_ID 3 /* ID via SNAP */
#define T_ATM_USER_ID 4 /* ID via user codepoints */
/* Constant values in the above 4 definitions are not mandatory */

#define T_ATM_BLLI_NORMAL_MODE 1
#define T_ATM_BLLI_EXTENDED_MODE 2

```



```

#define T_ATM_BLLI2_I1745 1 /* I.1745 */
#define T_ATM_BLLI2_Q921 2 /* Q.921 */
#define T_ATM_BLLI2_X25_LINK 6 /* X.25, link layer */
#define T_ATM_BLLI2_X25_MLINK 7 /* X.25, multilink */
#define T_ATM_BLLI2_LAPB 8 /* Extended LAPB */
#define T_ATM_BLLI2_HDLC_ARM 9 /* I.4335, ARM */
#define T_ATM_BLLI2_HDLC_NRM 10 /* I.4335, NRM */
#define T_ATM_BLLI2_HDLC_ABM 11 /* I.4335, ABM */
#define T_ATM_BLLI2_I8802 12 /* I.8802 */
#define T_ATM_BLLI2_X75 13 /* X.75 */
#define T_ATM_BLLI2_Q922 14 /* Q.922 */
#define T_ATM_BLLI2_I7776 17 /* I.7776 */

#define T_ATM_BLLI3_X25 6 /* X.25 */
#define T_ATM_BLLI3_I8208 7 /* I.8208 */
#define T_ATM_BLLI3_X223 8 /* X.223 */
#define T_ATM_BLLI3_I8473 9 /* I.8473 */
#define T_ATM_BLLI3_T70 10 /* T.70 */
#define T_ATM_BLLI3_I9577 11 /* I.9577 */

#define T_ATM_PACKET_SIZE_16 4
#define T_ATM_PACKET_SIZE_32 5
#define T_ATM_PACKET_SIZE_64 6
#define T_ATM_PACKET_SIZE_128 7
#define T_ATM_PACKET_SIZE_256 8
#define T_ATM_PACKET_SIZE_512 9
#define T_ATM_PACKET_SIZE_1024 10
#define T_ATM_PACKET_SIZE_2048 11
#define T_ATM_PACKET_SIZE_4096 12

/*
 * ATM network address structure
 */
struct t_atm_addr {
 int8_t address_format;
 uint8_t address_length;
 uint8_t address [20];
}

/*
 * T_ATM_CALLER_ID structure
 */
struct t_atm_caller_id {
 int8_t presentation;
 uint8_t screening;
}

/*
 * T_ATM_CALLER_ID values
 */
#define T_ATM_PRES_ALLOWED 0
#define T_ATM_PRES_RESTRICTED 1
#define T_ATM_PRES_UNAVAILABLE 2
#define T_ATM_USER_ID_NOT_SCREENED 0
#define T_ATM_USER_ID_PASSED_SCREEN 1
#define T_ATM_USER_ID_FAILED_SCREEN 2
#define T_ATM_NETWORK_PROVIDED_ID 3

/*

```

```

* T_ATM_CAUSE structure
*/
struct t_atm_cause {
 int8_t coding_standard;
 uint8_r location;
 uint8_r cause_value;
 uint8_r diagnostics [4];
}

/*
* T_ATM_CAUSE values
*/
#define T_ATM_LOC_USER 0
#define T_ATM_LOC_LOCAL_PRIVATE_NET 1
#define T_ATM_LOC_LOCAL_PUBLIC_NET 2
#define T_ATM_LOC_TRANSIT_NET 3
#define T_ATM_LOC_REMOTE_PUBLIC_NET 4
#define T_ATM_LOC_REMOTE_PRIVATE_NET 5
#define T_ATM_LOC_INTERNATIONAL_NET 7
#define T_ATM_LOC_BEYOND_INTERWORKING 10

#define T_ATM_CAUSE_UNALLOCATED_NUMBER 1
#define T_ATM_CAUSE_NO_ROUTE_TO_TRANSIT_NETWORK 2
#define T_ATM_CAUSE_NO_ROUTE_TO_DESTINATION 3
#define T_ATM_CAUSE_NORMAL_CALL_CLEARING 16
#define T_ATM_CAUSE_USER_BUSY 17
#define T_ATM_CAUSE_NO_USER_RESPONDING 18
#define T_ATM_CAUSE_CALL_REJECTED 21
#define T_ATM_CAUSE_NUMBER_CHANGED 22
#define T_ATM_CAUSE_ALL_CALLS_WITHOUT_CALLER_ID_REJECTED 23
#define T_ATM_CAUSE_DESTINATION_OUT_OF_ORDER 27
#define T_ATM_CAUSE_INVALID_NUMBER_FORMAT 28
#define T_ATM_CAUSE_RESPONSE_TO_STATUS_ENQUIRY 30
#define T_ATM_CAUSE_UNSPECIFIED_NORMAL 31
#define T_ATM_CAUSE_REQUESTED_VPCI_VCI_NOT_AVAILABLE 35
#define T_ATM_CAUSE_VPCI_VCI_ASSIGNMENT_FAILURE 36
#define T_ATM_CAUSE_USER_CELL_RATE_NOT_AVAILABLE 37
#define T_ATM_CAUSE_NETWORK_OUT_OF_ORDER 38
#define T_ATM_CAUSE_TEMPORARY_FAILURE 41
#define T_ATM_CAUSE_ACCESS_INFO_DISCARDED 43
#define T_ATM_CAUSE_NO_VPCI_VCI_AVAILABLE 45
#define T_ATM_CAUSE_UNSPECIFIED_RESOURCE_UNAVAILABLE 47
#define T_ATM_CAUSE_QUALITY_OF_SERVICE_UNAVAILABLE 49
#define T_ATM_CAUSE_BEARER_CAPABILITY_NOT_AUTHORIZED 57
#define T_ATM_CAUSE_BEARER_CAPABILITY_UNAVAILABLE 58
#define T_ATM_CAUSE_SERVICE_OR_OPTION_UNAVAILABLE 63
#define T_ATM_CAUSE_BEARER_CAPABILITY_NOT_IMPLEMENTED 65
#define T_ATM_CAUSE_INVALID_TRAFFIC_PARAMETERS 73
#define T_ATM_CAUSE_AAL_PARAMETERS_NOT_SUPPORTED 78
#define T_ATM_CAUSE_INVALID_CALL_REFERENCE_VALUE 81
#define T_ATM_CAUSE_IDENTIFIED_CHANNEL_DOES_NOT_EXIST 82
#define T_ATM_CAUSE_INCOMPATIBLE_DESTINATION 88
#define T_ATM_CAUSE_INVALID_ENDPOINT_REFERENCE 89
#define T_ATM_CAUSE_INVALID_TRANSIT_NETWORK_SELECTION 91
#define T_ATM_CAUSE_TOO_MANY_PENDING_ADD_PARTY_REQUESTS 92
#define T_ATM_CAUSE_MANDATORY_INFO_ELEMENT_MISSING 96
#define T_ATM_CAUSE_MESSAGE_TYPE_NOT_IMPLEMENTED 97
#define T_ATM_CAUSE_INFO_ELEMENT_NOT_IMPLEMENTED 99
#define T_ATM_CAUSE_INVALID_INFO_ELEMENT_CONTENTS 100

```

```
#define T_ATM_CAUSE_MESSAGE_INCOMPATIBLE_WITH_CALL_STATE 101
#define T_ATM_CAUSE_RECOVERY_ON_TIMER_EXPIRY 102
#define T_ATM_CAUSE_INCORRECT_MESSAGE_LENGTH 104
#define T_ATM_CAUSE_UNSPECIFIED_PROTOCOL_ERROR 111

/*
 * T_ATM_QOS structure
 */
struct t_atm_qos_substruct {
 int32_t coding_standard;
}

struct t_atm_qos {
 int8_t coding_standard;
 struct t_atm_qos_substruct forward;
 struct t_atm_qos_substruct backward;
}

/*
 * T_ATM_QOS values
 */
#define T_ATM_QOS_CLASS_0 0
#define T_ATM_QOS_CLASS_1 1
#define T_ATM_QOS_CLASS_2 2
#define T_ATM_QOS_CLASS_3 3
#define T_ATM_QOS_CLASS_4 4

/*
 * T_ATM_TRANSIT structure
 */
struct t_atm_transit {
 uint8_t length;
 uint8_t network_id[]; /* variable-sized array */
}
```



# Glossary

**abortive release**

An abrupt termination of a transport connection, which may result in the loss of data.

**asynchronous mode**

The mode of execution in which transport service functions do not wait for specific asynchronous events to occur before returning control to the user, but instead return immediately if the event is not pending.

**connection establishment**

The phase in connection-mode that enables two transport users to create a transport connection between them.

**connection-mode**

A mode of transfer where a logical link is established between two endpoints. Data is passed over this link by a sequenced and reliable way.

**connectionless-mode**

A mode of transfer where different units of data are passed through the network without any relationship between them.

**connection release**

The phase in connection-mode that terminates a previously established transport connection between two users.

**datagram**

A unit of data transferred between two users of the connectionless-mode service.

**data transfer**

The phase in connection-mode or connectionless-mode that supports the transfer of data between two transport users.

**expedited data**

Data that are considered urgent. The specific semantics of expedited data are defined by the transport provider that provides the transport service.

**expedited transport service data unit**

The amount of expedited user data, the identity of which is preserved from one end of a transport connection to the other (that is, an expedited message).

**host byte order**

The implementation-dependent byte order supported by the local host machine (see **network byte order** on page 416). Functions are provided to convert 16 and 32-bit values between network and host byte order (see *htonl()*).

**initiator**

An entity that initiates a connection request.

**legacy**

An item marked “LEGACY” in this specification means that it is being retained for compatibility with older applications, but has limitations which makes it inappropriate for developing portable applications. New applications should use alternative means of obtaining equivalent functionality.

**network byte order**

The byte order in which the most significant byte of a multibyte integer value is transmitted first. This byte order is the standard byte order for Internet protocols.

**network host database**

A database whose entries define the names and network addresses of host machines. See *gethostent()*.

**network net database**

A database whose entries define the names and network numbers of networks. See *getnetent()*.

**network protocol database**

A database whose entries define the names and protocol numbers of protocols. See *getprotoent()*.

**network service database**

A database whose entries define the names and local port numbers of services. See *getservent()*.

**orderly release**

A procedure for gracefully terminating a transport connection with no loss of data.

**responder**

An entity with whom an initiator wishes to establish a transport connection.

**socket**

A communications endpoint associated with a file descriptor that provides communications services using a specified communications protocol.

**synchronous mode**

The mode of execution in which transport service functions wait for specific asynchronous events to occur before returning control to the user.

**transport address**

The identifier used to differentiate and locate specific transport endpoints in a network.

**transport connection**

The communication circuit that is established between two transport users in connection-mode.

**transport endpoint**

The communication path, which is identified by a file descriptor, between a transport user and a specific transport provider. A transport endpoint is called passive before, and active after, a relationship is established, with a specific instance of this transport provider, identified by the TSAP.

**transport provider identifier**

A character string used by the function to identify the transport service provider.

**transport service access point**

A TSAP is a uniquely identified instance of the transport provider. A TSAP is used to identify a transport user on a certain endsystem. In connection-mode, a single TSAP may have more than one connection established to one or more remote TSAPs; each individual connection then is identified by a transport endpoint at each end.

**transport service data unit**

A unit of data transferred across the transport service with boundaries and content preserved unchanged. A TSDU may be divided into sub-units passed between the user and XTI. The T\_MORE flag is set in all but the last fragment of a TSDU sequence constituting a TSDU. The T\_MORE flag implies nothing about how the data is handled and passed to the lower level by the transport provider, and how they are delivered to the remote user.

**transport service provider**

A transport protocol providing the service of the transport layer.

**transport service user**

An abstract representation of the totality of those entities within a single system that make use of the transport service.

**user application**

The set of user programs, implemented as one or more process(es) in terms of UNIX semantics, written to realise a task, consisting of a set of user required functions.





# Index

|                                           |                                                                                                |
|-------------------------------------------|------------------------------------------------------------------------------------------------|
| <arpa/inet.h>.....                        | 210                                                                                            |
| <fcntl.h> .....                           | 178                                                                                            |
| <netatm/atm.h> .....                      | 405                                                                                            |
| <netdb.h> .....                           | 211                                                                                            |
| <netinet/in.h>.....                       | 213                                                                                            |
| <sys/socket.h>.....                       | 179, 405                                                                                       |
| <sys/stat.h> .....                        | 183                                                                                            |
| <sys/uio.h>.....                          | 184                                                                                            |
| <sys/un.h>.....                           | 185                                                                                            |
| <unistd.h>.....                           | 214                                                                                            |
| <xti.h> .....                             | 51, 77, 281, 405                                                                               |
| <xti_atm.h>.....                          | 405                                                                                            |
| <xti_mosi.h> .....                        | 317                                                                                            |
| _XOPEN_SOURCE.....                        | 3                                                                                              |
| abortive release.....                     | 24, 111, 415                                                                                   |
| accept .....                              | 222                                                                                            |
| accept().....                             | 130                                                                                            |
| accept1 .....                             | 33, 37                                                                                         |
| accept2 .....                             | 33, 37                                                                                         |
| accept3 .....                             | 33, 37                                                                                         |
| additions to <sys/socket.h> for ATM ..... | 405                                                                                            |
| additions to <xti.h> for ATM .....        | 405                                                                                            |
| address.....                              | 11-12, 19-21, 26, 29, 55<br>.....59-60, 63-66, 71, 74, 77<br>.....81-82, 100-101, 116-117, 215 |
| address Class.....                        | 353                                                                                            |
| application .....                         | 11-14, 245, 280                                                                                |
| applications.....                         | 51                                                                                             |
| portability .....                         | 245, 279                                                                                       |
| applications portability.....             | 11, 280                                                                                        |
| association-related .....                 | 220                                                                                            |
| association-related options.....          | 39                                                                                             |
| asynchronous .....                        | 38, 79                                                                                         |
| asynchronous events .....                 | 14                                                                                             |
| asynchronous mode .....                   | 14, 93, 415                                                                                    |
| ATM addresses.....                        | 364                                                                                            |
| ATM levels .....                          | 405                                                                                            |
| ATM options .....                         | 369                                                                                            |
| ATM transport headers .....               | 405                                                                                            |
| ATM transport protocol for sockets .....  | 395                                                                                            |
| ATM transport protocol for XTI .....      | 363                                                                                            |
| ATM transport provider.....               | 363                                                                                            |
| bind.....                                 | 11, 28, 33, 36, 59, 128                                                                        |
| bind() .....                              | 132                                                                                            |
| broadcast .....                           | 353                                                                                            |
| buffer.....                               | 57, 69, 91                                                                                     |
| byte order of Internet address.....       | 353                                                                                            |
| C language                                |                                                                                                |
| Issue 4 environment.....                  | 2                                                                                              |
| Call structure .....                      | 286                                                                                            |
| caller.....                               | 29, 64, 77                                                                                     |
| can.....                                  | 1                                                                                              |
| character string.....                     | 11                                                                                             |
| checksum check .....                      | 231                                                                                            |
| child process .....                       | 12                                                                                             |
| CL.....                                   | 295                                                                                            |
| Class of address .....                    | 353                                                                                            |
| close .....                               | 28, 36-37, 62                                                                                  |
| close() .....                             | 134                                                                                            |
| closed.....                               | 33, 36-37                                                                                      |
| cluster of local networks .....           | 353                                                                                            |
| cmsghdr .....                             | 179                                                                                            |
| CO .....                                  | 295                                                                                            |
| Common Usage C.....                       | 2                                                                                              |
| compatibility                             |                                                                                                |
| future.....                               | 58                                                                                             |
| compilation environment .....             | 3                                                                                              |
| connect() .....                           | 135                                                                                            |
| connect1.....                             | 33, 37                                                                                         |
| connect2.....                             | 33, 37                                                                                         |
| connection.....                           | 11, 63, 91, 93, 111, 113                                                                       |
| connection establishment .....            | 19, 21, 63-64<br>.....93, 229, 241, 415                                                        |
| connection indication .....               | 35, 59-60, 77, 95                                                                              |
| connection mode .....                     | 216                                                                                            |
| connection release.....                   | 19, 24-25, 229, 241, 415                                                                       |
| connection request.....                   | 54, 77, 93, 111                                                                                |
| connection-mode .....                     | 13, 19, 21, 28, 37, 241-242, 415                                                               |
| connection-mode service .....             | 37                                                                                             |
| connectionless .....                      | 224                                                                                            |
| connectionless-mode .....                 | 13, 26, 28, 36, 100<br>.....102, 116, 220, 244, 415                                            |
| constants.....                            | 51                                                                                             |
| control message protocol.....             | 353                                                                                            |
| create                                    |                                                                                                |
| transport endpoint .....                  | 20                                                                                             |
| current event .....                       | 31, 36, 79                                                                                     |
| current state .....                       | 12, 31, 36, 76, 125                                                                            |
| data .....                                | 91, 95, 100, 108, 111, 114, 116                                                                |
| data transfer .....                       | 19, 22, 26, 36-37, 241, 415                                                                    |
| data unit.....                            | 23, 26, 98, 100, 116                                                                           |
| discarded.....                            | 27                                                                                             |

|                         |                                  |
|-------------------------|----------------------------------|
| datagram                | 13, 27, 415                      |
| datagram structure      | 286                              |
| de-initialisation       | 19-20, 26, 36                    |
| default                 | 81, 245                          |
| device                  | 245                              |
| device driver           | 280                              |
| discarded data unit     | 27                               |
| discon                  | 98, 114                          |
| disconnect              | 15, 95                           |
| indication              | 33                               |
| disconnection           | 19, 38, 95                       |
| request                 | 111                              |
| disconnection structure | 286                              |
| dup                     | 11-12, 125-126                   |
| duplex                  | 22                               |
| EBADF                   |                                  |
| in recvmsg()            | 157                              |
| EM                      | 246, 295                         |
| end-to-end significance | 39                               |
| endhostent()            | 188                              |
| endnetent()             | 190                              |
| endprotoent()           | 191                              |
| endservent()            | 192                              |
| enqueue                 | 21, 29                           |
| errmsg                  | 66                               |
| errno                   | 66                               |
| errnum                  | 124                              |
| error                   | 282                              |
| error code              | 102, 124                         |
| error codes             | 281-282                          |
| TACCES                  | 281                              |
| TADDRBUSY               | 282                              |
| TBADADDR                | 281                              |
| TBADDATA                | 281                              |
| TBADF                   | 281                              |
| TBADFLAG                | 281                              |
| TBADNAME                | 282                              |
| TBADOPT                 | 281                              |
| TBADQLEN                | 282                              |
| TBADSEQ                 | 281                              |
| TBUFOVFLW               | 281                              |
| TFLOW                   | 281                              |
| TINDOUT                 | 282                              |
| TLOOK                   | 281                              |
| TNOADDR                 | 281                              |
| TNODATA                 | 281                              |
| TNODIS                  | 281                              |
| TNOREL                  | 281                              |
| TNOSTRUCTYPE            | 282                              |
| TNOTSUPPORT             | 281                              |
| TNOUDERR                | 281                              |
| TOUTSTATE               | 281                              |
| TPROTO                  | 282                              |
| TPROVMISMATCH           | 282                              |
| TQFULL                  | 282                              |
| TRESADDR                | 282                              |
| TRESQLEN                | 282                              |
| TSTATECHNG              | 282                              |
| TSYSERR                 | 281                              |
| error handling          | 13                               |
| error indication        | 102                              |
| error message           | 66, 124                          |
| error number            | 124                              |
| error numbers           | 6                                |
| established             |                                  |
| connection              | 242                              |
| ETSDU                   | 23, 81-82, 91, 105, 108-110, 295 |
| event                   | 31, 36, 244                      |
| current                 | 31, 36, 79                       |
| event management        | 16-17                            |
| Event Management        | 246                              |
| events                  | 247, 282                         |
| accept1                 | 33, 37                           |
| accept2                 | 33, 37                           |
| accept3                 | 33, 37                           |
| bind                    | 33, 36                           |
| closed                  | 33, 36-37                        |
| connect1                | 33, 37                           |
| connect2                | 33, 37                           |
| incoming                | 34                               |
| listen                  | 34, 37, 77                       |
| opened                  | 33, 36                           |
| optmgmt                 | 33, 36                           |
| outgoing                | 33                               |
| pass_conn               | 34, 37                           |
| rcv                     | 34, 37, 91                       |
| rcvconnect              | 34, 37                           |
| rcvdis1                 | 34, 37                           |
| rcvdis2                 | 34, 37                           |
| rcvdis3                 | 34, 37                           |
| rcvrel                  | 34, 37                           |
| rcvudata                | 34, 36                           |
| rcvuderr                | 34, 36                           |
| snd                     | 33, 37                           |
| snddis1                 | 33, 37                           |
| snddis2                 | 33, 37                           |
| sndrel                  | 33, 37                           |
| sndudata                | 33, 36                           |
| T_CONNECT               | 282                              |
| T_DATA                  | 282                              |
| T_DISCONNECT            | 282                              |
| T_EXDATA                | 282                              |

## Index

|                                             |                              |                                              |                   |
|---------------------------------------------|------------------------------|----------------------------------------------|-------------------|
| T_GODATA .....                              | 282                          | getnetbyaddr() .....                         | 190, 196          |
| T_GOEXDATA .....                            | 282                          | getnetbyname() .....                         | 190               |
| T_LISTEN .....                              | 282                          | getnetent() .....                            | 190               |
| T_ORDREL .....                              | 282                          | getpeername() .....                          | 142               |
| T_UDERR .....                               | 282                          | getprotobyname() .....                       | 191               |
| unbind .....                                | 33, 36                       | getprotobynumber() .....                     | 191, 197          |
| events and t_look .....                     | 15                           | getprotoent() .....                          | 191               |
| example .....                               | 249                          | getservbyname() .....                        | 192               |
| example for select .....                    | 257                          | getservbyport() .....                        | 192, 198          |
| exec .....                                  | 125-126                      | getservent() .....                           | 192               |
| execution mode .....                        | 14, 23                       | getsockname() .....                          | 143               |
| expedited data .....                        | 19, 23, 82, 91, 105, 108     | getsockopt() .....                           | 144               |
| .....                                       | 216, 224, 226, 229, 290, 415 | headers                                      |                   |
| expedited transport service data unit ..... | 415                          | <xti.h> .....                                | 281               |
| ETSDU .....                                 | 82, 91                       | host byte order .....                        | 415               |
| family of protocols .....                   | 353                          | host part of address .....                   | 353               |
| fcntl .....                                 | 11, 14, 64, 77-78, 91-93     | hostent .....                                | 211               |
| .....                                       | 100-101, 109, 116-117        | htonl() .....                                | 200               |
| fcntl() .....                               | 138                          | htons() .....                                | 200               |
| fcntl.h .....                               | 81                           | h_errno .....                                | 199               |
| fd .....                                    | 11, 33                       | h_errno() .....                              | 188               |
| features .....                              | 28-29                        | ICMP .....                                   | 353               |
| fgetpos() .....                             | 139                          | implementation-dependent .....               | 1                 |
| file descriptor .....                       | 11, 62, 71, 81, 125          | INADDR_ANY .....                             | 213, 353          |
| file.c .....                                | 51                           | INADDR_BROADCAST .....                       | 213               |
| flag .....                                  | 81, 84, 91                   | incoming events .....                        | 34                |
| flags .....                                 | 81, 84, 91, 224, 282, 285    | inet_addr() .....                            | 201               |
| T_CHECK .....                               | 282                          | inet_lnaof() .....                           | 201               |
| T_CURRENT .....                             | 282                          | inet_makeaddr() .....                        | 201               |
| T_DEFAULT .....                             | 282                          | inet_netof() .....                           | 201               |
| T_EXPEDITED .....                           | 282                          | inet_network() .....                         | 201               |
| T_FAILURE .....                             | 282                          | inet_ntoa() .....                            | 201               |
| T_MORE .....                                | 282                          | initialisation .....                         | 19-20, 26, 36, 81 |
| T_NEGOTIATE .....                           | 282                          | initiator .....                              | 19, 415           |
| T_NOTSUPPORT .....                          | 282                          | interfaces                                   |                   |
| T_PARTSUCCESS .....                         | 282                          | implementation .....                         | 2                 |
| T_PUSH .....                                | 282                          | use .....                                    | 2                 |
| T_READONLY .....                            | 282                          | Internet address                             |                   |
| T_SUCCESS .....                             | 282                          | byte order .....                             | 353               |
| flow control .....                          | 27                           | Internet protocol .....                      | 353               |
| fork .....                                  | 11-12, 125-126               | Internet Protocol family .....               | 353               |
| fsetpos() .....                             | 140                          | Internet protocol family .....               | 353               |
| ftell() .....                               | 141                          | Internet protocol-specific information ..... | 229               |
| full duplex .....                           | 22                           | in_addr .....                                | 210, 213          |
| F_GETOWN .....                              | 138, 178                     | in_addr_t .....                              | 210-211           |
| F_SETOWN .....                              | 138, 178                     | in_port_t .....                              | 210-211, 213      |
| General purpose defines .....               | 287                          | IOV_MAX .....                                | 283               |
| gethostbyaddr() .....                       | 188, 194                     | IP .....                                     | 353               |
| gethostbyname() .....                       | 188                          | IP-level Options .....                       | 239, 292          |
| gethostent() .....                          | 188                          | IPPORT_RESERVED .....                        | 211               |
| gethostname() .....                         | 195                          |                                              |                   |

- IPPROTO\_ macros
  - defined in <netinet/in.h> .....213
- IP\_TOS type of service .....240, 293
- ISO .....215, 225, 288, 295
  - priorities.....225, 288
  - protection levels.....225, 288
  - transport classes.....225, 288
- ISO C .....2
- language-dependent .....66
- legacy .....1, 415
- library functions.....283
- library structure .....57
- linger .....**180**
- listen.....60, 77, 222
- listen().....**147**
- listener application.....12
- little-endian.....353
- lseek() .....**148**
- management options .....218, 221, 227, 290
- mandatory features.....245
- maximum size
  - address.....82
  - address buffer.....59, 74
  - buffer.....64, 77, 93, 100, 102
  - ETSDU.....82, 110
  - TSDU.....24, 82, 110, 237
- may .....1
- memory
  - allocate.....57, 69
- mode
  - asynchronous .....14
  - connection.....19, 21, 28, 37, 216, 241-242
  - connectionless.....26, 28, 36, 100, 102
  - .....116, 220, 244
  - record-oriented .....24
  - stream-oriented.....24
  - synchronous .....14, 79
- modes of service .....13
- mosi Header File.....317
- msghdr.....**179**
- MSG\_ macros
  - defined in <sys/socket.h> .....**181**
- multiple options.....45
- must .....1
- name space
  - X/Open.....3
- native ATM services .....363, 395
- netbuf structure .....40, 57, 84
- netent.....**211**
- network byte order.....416
- network host database .....416
- network net database.....416
- network protocol database.....416
- network service database .....416
- next state.....36
- ntohl().....**200**, 203
- ntohs() .....**200**
- NULL.....59
- null
  - call.....111
- null pointer .....58, 60, 66, 69, 83, 95, 102
- ocnt .....**33**
- open .....81
- opened .....33, 36
- option
  - value.....46
- option management .....287
- option negotiation
  - initiate .....42
  - response.....43
- option values .....216
- options
  - association-related .....39
  - connection mode .....216
  - connectionless-mode .....220
  - expedited data.....216
  - format.....48
  - generalities.....39
  - illegal.....41
  - ISO-specific .....225-226, 288-289
  - management.....218
  - multiple .....45
  - privileged .....45
  - quality of service .....216
  - read-only .....45
  - retrieving information.....44
  - TCP-level.....230
  - transport endpoint .....84
  - transport level .....9
  - transport provider.....64
  - T\_IP-level .....232
  - T\_UDP-level .....231
  - unsupported.....42
  - XTI-level.....288
- Options management structure .....**285**
- options with end-to-end significance.....39
- optmgmt.....28, 33, 36
- orderly release .....24, 97, 113, 416
- OSI .....295
  - transport classes .....223
- outgoing events.....33
- outstanding connection indications.....35, 60, 95

## Index

|                                       |                                         |
|---------------------------------------|-----------------------------------------|
| O_NONBLOCK flag .....                 | 14                                      |
| pass_conn.....                        | 34, 37                                  |
| poll .....                            | 79                                      |
| poll() .....                          | <b>149</b>                              |
| polling.....                          | 15                                      |
| portability.....                      | 50                                      |
| portable .....                        | 11, 245, 279                            |
| precedence levels                     |                                         |
| IP.....                               | 239, 293                                |
| primitives.....                       | 14-15                                   |
| process .....                         | 12                                      |
| program .....                         | 51                                      |
| programs                              |                                         |
| multiple protocol.....                | 245                                     |
| protocol.....                         | 11, 29, 39, 59, 63, 71                  |
| .....                                 | 74, 81, 84, 102, 215, 245               |
| protocol independence .....           | 72, 82, 245                             |
| protocol-specific servicelimits ..... | 284                                     |
| protocols in Internet family.....     | 353                                     |
| protoent .....                        | <b>211</b>                              |
| quality of service.....               | 216, 220, 226, 290                      |
| queue.....                            | 21, 29, 78                              |
| rate .....                            | 216                                     |
| rate structure .....                  | <b>225, 289</b>                         |
| rcv .....                             | 34, 37                                  |
| rcvconnect.....                       | 34, 37                                  |
| rcvdis1.....                          | 34, 37                                  |
| rcvdis2.....                          | 34, 37                                  |
| rcvdis3.....                          | 34, 37                                  |
| rcvrel .....                          | 34, 37                                  |
| rcvreldata .....                      | 34                                      |
| rcvudata.....                         | 34                                      |
| rcvuderr .....                        | 34                                      |
| rcvudata .....                        | 284                                     |
| read() .....                          | <b>150</b>                              |
| readv() .....                         | <b>150</b>                              |
| reason                                |                                         |
| disconnection .....                   | 95                                      |
| receipt.....                          | 97                                      |
| receive.....                          | 91, 98, 100                             |
| Receiving Data .....                  | 22, 27                                  |
| record-oriented.....                  | 24                                      |
| recv().....                           | <b>151</b>                              |
| recvfrom().....                       | <b>153</b>                              |
| recvmsg().....                        | <b>156</b>                              |
| release .....                         | 19, 24, 37, 97, 111, 113                |
| reliable.....                         | 13                                      |
| remote user .....                     | 16, 21, 24-25, 62, 64-65, 111, 242, 244 |
| reqvalue .....                        | 216, 226, 289                           |
| reqvalue structure .....              | <b>226, 289</b>                         |
| resfd.....                            | <b>33</b>                               |
| responder .....                       | 19, 416                                 |
| safety .....                          | 6                                       |
| sa_family_t.....                      | <b>179</b>                              |
| select() .....                        | <b>159</b>                              |
| send() .....                          | <b>160</b>                              |
| Sending Data .....                    | 23, 27                                  |
| sendmsg() .....                       | <b>162</b>                              |
| sendto() .....                        | <b>165</b>                              |
| servent.....                          | <b>211</b>                              |
| server program .....                  | 248, 257                                |
| service definition                    |                                         |
| ISO.....                              | 24, 215, 222                            |
| TCP .....                             | 24                                      |
| service type defines.....             | 285                                     |
| sethostent() .....                    | <b>188, 204</b>                         |
| setnetent().....                      | <b>190, 205</b>                         |
| setprotoent() .....                   | <b>191, 206</b>                         |
| setservent().....                     | <b>192, 207</b>                         |
| setsockopt() .....                    | <b>168</b>                              |
| should.....                           | 1                                       |
| shutdown() .....                      | <b>171</b>                              |
| snd .....                             | 33, 37-38, 215                          |
| snddis1.....                          | 33, 37                                  |
| snddis2.....                          | 33, 37                                  |
| sndrel.....                           | 33, 37-38                               |
| sndreldata.....                       | 33                                      |
| sndudata .....                        | 33, 36, 38, 215                         |
| sockaddr .....                        | <b>179</b>                              |
| sockaddr_in .....                     | <b>213</b>                              |
| sockaddr_un .....                     | <b>185</b>                              |
| socket .....                          | 235-236, 416                            |
| socket() .....                        | <b>172</b>                              |
| socketpair().....                     | <b>174</b>                              |
| SOCK_DGRAM .....                      | 353                                     |
| SOCK_RAW.....                         | 353                                     |
| SOCK_STREAM .....                     | 353                                     |
| SO_ macros                            |                                         |
| defined in <sys/socket.h> .....       | <b>180</b>                              |
| standard error.....                   | 66                                      |
| state .....                           | 31-32, 36, 287                          |
| current.....                          | 31, 36, 76, 125                         |
| next .....                            | 36                                      |
| T_DATAXFER .....                      | <b>32, 287</b>                          |
| T_IDLE .....                          | <b>32, 287</b>                          |
| T_INCON .....                         | <b>32, 287</b>                          |
| T_INREL .....                         | <b>32, 287</b>                          |
| T_OUTCON .....                        | <b>32, 287</b>                          |
| T_OUTREL .....                        | <b>32, 287</b>                          |
| T_UNBIND .....                        | <b>32</b>                               |
| T_UNBND .....                         | 287                                     |
| T_UNIT.....                           | <b>32</b>                               |

|                             |                 |                                      |                              |
|-----------------------------|-----------------|--------------------------------------|------------------------------|
| state table.....            | 36-37, 244      | TFLOW.....                           | 27, 247, 281                 |
| status                      |                 | thread cancellation point .....      | 6                            |
| connection.....             | 64              | thread safety .....                  | 6                            |
| connection request.....     | 22, 93          | threads.....                         | 6                            |
| stream-oriented.....        | 24              | thrpt.....                           | 216, 226, 289                |
| strerror(3C).....           | 66              | thrpt structure.....                 | 226, 289                     |
| struct netbuf.....          | 285             | TINDOUT .....                        | 282                          |
| struct rate.....            | 225, 289        | TLI.....                             | 279-280, 295                 |
| struct reqvalue.....        | 226, 289        | TLOOK.....                           | 15, 23, 38, 281              |
| struct thrpt.....           | 226, 289        | TNOADDR .....                        | 281                          |
| struct transdel.....        | 226, 289        | TNODATA .....                        | 281                          |
| struct t_bind.....          | 285             | TNODIS.....                          | 281                          |
| struct t_call.....          | 286             | TNOREL.....                          | 281                          |
| struct t_discon.....        | 286             | TNOSTRUCTYPE.....                    | 282                          |
| struct t_info.....          | 284             | TNOTSUPPORT .....                    | 281                          |
| struct t_kpalive.....       | 238, 292        | TNOUDERR.....                        | 281                          |
| struct t_linger.....        | 288             | TOS precedence levels .....          | 239, 293                     |
| struct t_opthdr.....        | 285             | TOUTSTATE.....                       | 281                          |
| struct t_optmgmt.....       | 285             | TPDU lengths.....                    | 225, 289                     |
| struct t_uderr.....         | 286             | TPROTO .....                         | 282                          |
| struct t_unitdata.....      | 286             | TPROVMISMATCH.....                   | 282                          |
| structure types .....       | 286             | TQFULL.....                          | 282                          |
| T_BIND.....                 | 286             | transdel.....                        | 216                          |
| T_CALL.....                 | 286             | transdel structure .....             | 226, 289                     |
| T_DIS.....                  | 286             | Transmission Control Protocol.....   | 353                          |
| T_INFO.....                 | 286             | transport address .....              | 11, 215, 416                 |
| T_OPTMGMT .....             | 286             | transport classes.....               | 223, 225, 288                |
| T_UDERROR.....              | 286             | transport connection.....            | 11, 21, 62, 71, 113, 416     |
| T_UNITDATA .....            | 286             | transport endpoint .....             | 11, 32-33, 46, 59, 62-63     |
| subnet part of address..... | 353             | .....                                | 79, 81, 83-84, 125, 128, 416 |
| SVID.....                   | 295             | Transport Level Interface (TLI)..... | 279-280                      |
| synchronise.....            | 125             | transport level options.....         | 9                            |
| synchronous mode .....      | 14, 79, 93, 416 | transport provider.....              | 11, 19, 31-32, 36, 40        |
| t-opthdr.....               | 216             | .....                                | 71, 81, 125, 215, 241, 285   |
| TACCES.....                 | 281             | transport provider identifier.....   | 11, 19, 81, 416              |
| TADDRBUSY.....              | 282             | transport service.....               | 9, 215, 241                  |
| TBADADDR.....               | 281             | transport service access point.....  | 416                          |
| TBADDATA .....              | 281             | TSAP.....                            | 12                           |
| TBADF.....                  | 281             | transport service data unit.....     | 416                          |
| TBADFLAG .....              | 281             | TSDU.....                            | 16, 82, 91, 245              |
| TBADNAME.....               | 282             | transport service provider .....     | 417                          |
| TBADOPT .....               | 281             | transport service user.....          | 11, 19, 21, 24               |
| TBADQLEN.....               | 282             | .....                                | 31-32, 63, 241, 417          |
| TBADSEQ.....                | 281             | transport user actions .....         | 35                           |
| TBUFOVFLW .....             | 281             | TRESADDR.....                        | 282                          |
| TC.....                     | 295             | TRESQLEN .....                       | 282                          |
| TCP.....                    | 24, 295, 353    | TSAP.....                            | 12, 295                      |
| TCP-level options .....     | 230, 238, 291   | TSDU.....                            | 16, 23, 81-82, 91, 108-110   |
| terminated                  |                 | .....                                | 116, 215, 224, 229, 295      |
| connection.....             | 242             | TSTATECHNG.....                      | 282                          |
| terminology.....            | 1               | TSYSERR.....                         | 13, 23, 66, 281              |

## Index

|                     |                                |                    |                       |
|---------------------|--------------------------------|--------------------|-----------------------|
| T_ABSREQ            | 287                            | t_getinfo          | 28, 71, 222, 283      |
| t_accept            | 28, 54, 222, 229, 235, 283     | t_getinfo()        | 71                    |
| t_accept()          | 54                             | t_getprotaddr      | 28, 74, 283           |
| T_ACTIVEPROTECT     | 225, 289                       | t_getprotaddr()    | 74                    |
| t_addleaf()         | 385                            | t_getstate         | 28, 76, 283           |
| T_ADDR              | 287                            | t_getstate()       | 76                    |
| T_ALIGN             | 287                            | T_GODATA           | 15, 17, 26, 247, 282  |
| T_ALL               | 287                            | T_GOEXDATA         | 15, 17, 247, 282      |
| t_alloc             | 28, 57, 69, 283, 286           | T_HIRES            | 240, 293              |
| t_alloc()           | 57                             | T_HITHRPT          | 240, 293              |
| T_ALLOPT            | 287                            | T_IDLE             | 32, 36-37, 63, 287    |
| t_atm_sap structure | 365                            | T_IMMEDIATE        | 239, 293              |
| t_bind              | 11, 28, 59, 128, 222, 235, 283 | T_INCON            | 32, 37, 287           |
| T_BIND              | 286                            | T_INETCONTROL      | 239, 293              |
| t_bind()            | 59                             | T_INET_IP          | 239, 292              |
| t_call              | 54                             | T_INET_TCP         | 238, 291              |
| T_CALL              | 286                            | T_INET_UDP         | 238, 292              |
| T_CHECK             | 282                            | T_INFINITE         | 287                   |
| T_CLASS0            | 225, 288                       | T_INFO             | 286                   |
| T_CLASS1            | 225, 288                       | T_INREL            | 32, 37, 287           |
| T_CLASS2            | 225, 288                       | T_INVALID          | 287                   |
| T_CLASS3            | 225, 288                       | T_IP-level options | 232                   |
| T_CLASS4            | 225, 288                       | T_IP_BROADCAST     | 232, 239, 293         |
| t_close             | 28, 62, 222, 241, 283          | T_IP_DONTROUTE     | 232, 239, 293         |
| t_close()           | 62                             | T_IP_OPTIONS       | 232, 239, 292         |
| T_CLTS              | 285                            | T_IP_REUSEADDR     | 233, 239, 293         |
| T_CONNECT           | 15                             | T_IP_TOS           | 233, 239, 292         |
| t_connect           | 15                             | T_IP_TTL           | 234, 239, 292         |
| T_CONNECT           | 16                             | T_ISO_TP           | 226, 289              |
| t_connect           | 28, 63, 93, 222, 235           | t_kpalive          | 238, 292              |
| T_CONNECT           | 247, 282                       | T_LDELAY           | 240, 293              |
| t_connect           | 283                            | T_LISTEN           | 15                    |
| t_connect()         | 44, 63                         | t_listen           | 15                    |
| T_COTS              | 285                            | T_LISTEN           | 16                    |
| T_COTS_ORD          | 285                            | t_listen           | 28                    |
| T_CRITIC_ECP        | 239, 293                       | T_LISTEN           | 38                    |
| T_CURRENT           | 282                            | t_listen           | 38, 77, 223, 229, 236 |
| T_DATA              | 15-16, 26, 37-38, 247, 282     | T_LISTEN           | 247, 282              |
| T_DATAXFER          | 32, 37, 287                    | t_listen           | 283                   |
| T_DEFAULT           | 282                            | t_listen()         | 44, 77                |
| T_DIS               | 15, 38, 286                    | T_LOCOST           | 240, 293              |
| T_DISCONNECT        | 15, 17, 23, 38, 247, 282       | t_look             | 15, 28, 79, 236, 283  |
| t_errno             | 13, 68, 124, 282               | t_look()           | 79                    |
| t_error             | 13, 28, 66, 124, 283           | T_LTPDUDFLT        | 225, 289              |
| t_error()           | 66                             | T_MORE             | 23, 91, 105, 108, 282 |
| T_EXDATA            | 15-16, 236, 247, 282           | T_MORE flag        | 229                   |
| T_EXPEDITED         | 23, 91, 105, 282               | T_NB_ABORT         | 275                   |
| T_FAILURE           | 282                            | T_NB_BCAST_NAME    | 269                   |
| T_FLASH             | 239, 293                       | T_NB_CLOSED        | 275                   |
| t_free              | 28, 69, 283                    | T_NB_GROUP         | 269                   |
| t_free()            | 69                             | T_NB_LOCAL         | 269                   |

|                        |                               |                                     |                                |
|------------------------|-------------------------------|-------------------------------------|--------------------------------|
| T_NB_NAMELEN.....      | 269                           | t_rcvvudata.....                    | 28                             |
| T_NB_NOANSWER.....     | 275                           | t_rcvvudata().....                  | <b>106</b>                     |
| T_NB_OPREJ.....        | 275                           | T_READONLY.....                     | <b>282</b>                     |
| T_NB_UNIQUE.....       | 269                           | t_removeleaf().....                 | <b>387</b>                     |
| T_NEGOTIATE.....       | <b>282</b>                    | T_ROUTINE.....                      | <b>239, 293</b>                |
| T_NETCONTROL.....      | <b>239, 293</b>               | T_SENDZERO.....                     | <b>285</b>                     |
| T_NEXTHDR.....         | <b>287</b>                    | T_SNA_CONNECTION_OUTAGE.....        | 325                            |
| T_NO.....              | <b>287</b>                    | T_SNA_CONNECTION_SETUP_FAILURE..... | 325                            |
| T_NOPROTECT.....       | <b>225, 288</b>               | T_SNA_MAX_LU_LEN.....               | 321                            |
| T_NOTOS.....           | <b>240, 293</b>               | T_SNA_MAX_NETID_LEN.....            | 321                            |
| T_NOTSUPPORT.....      | <b>282</b>                    | T_SNA_MAX_TPN_LEN.....              | 321                            |
| T_NULL.....            | <b>287</b>                    | T_SNA_SYSTEM_DISCONNECT.....        | 325                            |
| t_open.....            | 11, 28, 81, 223, 236, 283     | T_SNA_TIMEOUT.....                  | 325                            |
| t_open().....          | <b>81</b>                     | T_SNA_USER_DISCONNECT.....          | 325                            |
| T_OPT.....             | <b>287</b>                    | t_snd.....                          | 15, 28, 108, 224, 236-237, 284 |
| t_optmgmt.....         | 28, 84, 283                   | t_snd().....                        | <b>108</b>                     |
| T_OPTMGMT.....         | <b>286</b>                    | t_snddis.....                       | 28, 111, 224, 237, 284         |
| t_optmgmt().....       | 45, 84                        | t_snddis().....                     | <b>111</b>                     |
| T_OPT_DATA.....        | <b>287</b>                    | t_sndrel.....                       | <b>113, 284</b>                |
| T_OPT_FIRSTHDR.....    | <b>287</b>                    | t_sndrel().....                     | <b>113</b>                     |
| T_ORDREL.....          | 15, 17, 38, 247, 282          | t_sndreldata.....                   | <b>114, 284</b>                |
| T_ORDRELDATA.....      | <b>285</b>                    | t_sndreldata().....                 | <b>114</b>                     |
| T_OUTCON.....          | <b>32, 37, 287</b>            | t_sndudata.....                     | 28, 116, 224, 237, 284         |
| T_OUTREL.....          | <b>32, 37, 287</b>            | t_sndudata().....                   | <b>116</b>                     |
| T_OVERRIDEFLASH.....   | <b>239, 293</b>               | t_sndv.....                         | 28, 284                        |
| T_PARTSUCCESS.....     | <b>282</b>                    | t_sndv().....                       | <b>118</b>                     |
| T_PASSIVEPROTECT.....  | <b>225, 289</b>               | t_sndvudata.....                    | 28, 284                        |
| T_PRIDFLT.....         | <b>225, 288</b>               | t_sndvudata().....                  | <b>121</b>                     |
| T_PRIHIGH.....         | <b>225, 288</b>               | t_sterror.....                      | 28, 124, 284                   |
| T_PRILOW.....          | <b>225, 288</b>               | t_sterror().....                    | <b>124</b>                     |
| T_PRIMID.....          | <b>225, 288</b>               | T_SUCCESS.....                      | <b>282</b>                     |
| T_PRIORITY.....        | <b>239, 293</b>               | t_sync.....                         | 12, 28, 125, 284               |
| T_PRITOP.....          | <b>225, 288</b>               | t_sync().....                       | <b>125</b>                     |
| T_PUSH.....            | 109, 119, 282                 | t_sysconf.....                      | 28, 284                        |
| t_rcv.....             | 15, 28, 38, 91, 224, 236, 283 | t_sysconf().....                    | <b>127</b>                     |
| t_rcv().....           | <b>91</b>                     | T_TCL_CHECKSUM.....                 | <b>227, 291</b>                |
| t_rcvconnect.....      | 28, 38, 93, 224, 236          | T_TCL_PRIORITY.....                 | <b>227, 290</b>                |
| t_rcvconnect().....    | 44, 93, 283                   | T_TCL_PROTECTION.....               | <b>227, 290</b>                |
| t_rcvdis.....          | 15, 28, 95, 224, 236, 283     | T_TCL_RESERRORRATE.....             | <b>227, 290</b>                |
| t_rcvdis().....        | <b>95</b>                     | T_TCL_TRANSDEL.....                 | <b>227, 290</b>                |
| t_rcvleafchange()..... | <b>389</b>                    | T_TCO_ACKTIME.....                  | <b>227, 290</b>                |
| t_rcvrel.....          | 15, 38, 97, 284               | T_TCO_ALTCLASS1.....                | <b>227, 291</b>                |
| t_rcvrel().....        | <b>97</b>                     | T_TCO_ALTCLASS2.....                | <b>227, 291</b>                |
| t_rcvreldata.....      | <b>98, 284</b>                | T_TCO_ALTCLASS3.....                | <b>227, 291</b>                |
| t_rcvreldata().....    | <b>98</b>                     | T_TCO_ALTCLASS4.....                | <b>227, 291</b>                |
| t_rcvudata.....        | 28, 38, 100, 224, 284         | T_TCO_CHECKSUM.....                 | <b>227, 291</b>                |
| t_rcvudata().....      | 44, 100                       | T_TCO_CLASS.....                    | <b>227, 291</b>                |
| t_rcvuderr.....        | 15, 28, 102, 224, 284         | T_TCO_CONNRRESIL.....               | <b>227, 290</b>                |
| t_rcvuderr().....      | 45, 102                       | T_TCO_ESTDELAY.....                 | <b>226, 290</b>                |
| t_rcvv.....            | 28, 284                       | T_TCO_ESTFAILPROB.....              | <b>226, 290</b>                |
| t_rcvv().....          | <b>104</b>                    | T_TCO_EXPD.....                     | <b>227, 290</b>                |



## Index

|                                |                          |                                            |                       |
|--------------------------------|--------------------------|--------------------------------------------|-----------------------|
| T_TCO_EXTFORM .....            | 227, 291                 | XEM .....                                  | 295                   |
| T_TCO_FLOWCTRL .....           | 227, 291                 | XTI .....                                  | 9, 295                |
| T_TCO_LTPDU .....              | 227, 290                 | applications .....                         | 51                    |
| T_TCO_NETEXP .....             | 227, 291                 | features .....                             | 28-29                 |
| T_TCO_NETRECPTCF .....         | 227, 291                 | library .....                              | 51                    |
| T_TCO_PRIORITY .....           | 227, 290                 | XTI error return .....                     | 282                   |
| T_TCO_PROTECTION .....         | 227, 290                 | XTI level .....                            | 288                   |
| T_TCO_REASTIME .....           | 227, 290                 | XTI-level options .....                    | 88, 288               |
| T_TCO_RELDELAY .....           | 227, 290                 | XTI_DEBUG .....                            | 288                   |
| T_TCO_RELFAILPROB .....        | 226, 290                 | XTI_GENERIC .....                          | 288                   |
| T_TCO_RESERRORRATE .....       | 226, 290                 | XTI_LINGER .....                           | 288                   |
| T_TCO_THROUGHPUT .....         | 226, 290                 | XTI_RCVBUF .....                           | 288                   |
| T_TCO_TRANSDEL .....           | 226, 290                 | XTI_RCVLOWAT .....                         | 288                   |
| T_TCO_TRANSFFAILPROB .....     | 226, 290                 | XTI_SNDBUF .....                           | 288                   |
| T_TCP .....                    | 229                      | XTI_SNTLOWAT .....                         | 288                   |
| T_TCP_KEEPALIVE .....          | 230, 238, 291            | Zero-length TSDUs and TSDU fragments ..... | 72                    |
| T_TCP_MAXSEG .....             | 231, 238, 291            | .....                                      | 82, 108-109, 215, 224 |
| T_TCP_NODELAY .....            | 231, 238, 291            |                                            |                       |
| T_UDATA .....                  | 287                      |                                            |                       |
| T_UDERR .....                  | 15, 17, 26, 38, 247, 282 |                                            |                       |
| T_UDERROR .....                | 286                      |                                            |                       |
| T_UDP .....                    | 229                      |                                            |                       |
| T_UDP-level options .....      | 231                      |                                            |                       |
| T_UDP_CHECKSUM .....           | 231, 239, 292            |                                            |                       |
| t_unbind .....                 | 28                       |                                            |                       |
| T_UNBIND .....                 | 32                       |                                            |                       |
| t_unbind .....                 | 128, 284                 |                                            |                       |
| t_unbind() .....               | 128                      |                                            |                       |
| T_UNBND .....                  | 36-37, 287               |                                            |                       |
| T_UNIT .....                   | 32                       |                                            |                       |
| T_UNITDATA .....               | 286                      |                                            |                       |
| T_UNSPEC .....                 | 48, 287                  |                                            |                       |
| T_YES .....                    | 287                      |                                            |                       |
| UDP .....                      | 295, 353                 |                                            |                       |
| UDP-level options .....        | 239, 292                 |                                            |                       |
| unbind .....                   | 28, 33, 36, 38           |                                            |                       |
| undefined .....                | 1                        |                                            |                       |
| unitdata .....                 | 100, 116                 |                                            |                       |
| Unitdata error structure ..... | 286                      |                                            |                       |
| UNIX                           |                          |                                            |                       |
| process .....                  | 12                       |                                            |                       |
| versions .....                 | 279                      |                                            |                       |
| unspecified .....              | 1                        |                                            |                       |
| user application .....         | 19, 26, 417              |                                            |                       |
| user data .....                | 95, 114                  |                                            |                       |
| User Datagram Protocol .....   | 353                      |                                            |                       |
| User-Network Interface .....   | 395                      |                                            |                       |
| will .....                     | 2                        |                                            |                       |
| write() .....                  | 176                      |                                            |                       |
| writev() .....                 | 176                      |                                            |                       |
| X/Open name space .....        | 3                        |                                            |                       |

