



X/Open CAE Specification

X/Open Transport Interface (XTI)

X/Open Company, Ltd.



© January 1992, X/Open Company Limited

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.

X/Open CAE Specification

X/Open Transport Interface (XTI)

ISBN: 1 872630 29 4

X/Open Document Number: XO/CAE/91/600

Set in Palatino by X/Open Company Ltd., U.K.

Printed by Maple Press, U.K.

Published by X/Open Company Ltd., U.K.

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

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

or by Electronic Mail to:

XoSpecs@xopen.co.uk

Contents

3X/OPEN TRANSPORT INTERFACE (XTI)^p

Chapter	1	GENERAL INTRODUCTION TO THE XTI
Chapter	2	EXPLANATORY NOTES
	2.1	TRANSPORT ENDPOINTS
	2.2	TRANSPORT PROVIDERS
	2.3	ASSOCIATION OF A UNIX PROCESS TO AN ENDPOINT
	2.4	USE OF THE SAME PROTOCOL ADDRESS
	2.5	MODES OF SERVICE
	2.6	ERROR HANDLING
	2.7	SYNCHRONOUS AND ASYNCHRONOUS EXECUTION MODES
	2.8	EVENT MANAGEMENT
	2.9	CHANGE HISTORY FROM XPG3 VERSION
	2.9.1	Changes Appearing in Revised XTI (1990)
	2.9.2	Changes Appearing in Addendum to Revised XTI (1991)
Chapter	3	XTI OVERVIEW
	3.1	OVERVIEW OF CONNECTION-ORIENTED MODE
	3.1.1	Initialisation/De-initialisation Phase
	3.1.2	Overview of Connection Establishment
	3.1.3	Overview of Data Transfer
	3.1.4	Overview of Connection Release
	3.2	OVERVIEW OF CONNECTIONLESS MODE
	3.2.1	Initialisation/De-initialisation Phase
	3.2.2	Overview of Data Transfer
	3.3	XTI FEATURES
	3.3.1	XTI Functions Versus Protocols
Chapter	4	STATES AND EVENTS IN XTI
	4.1	TRANSPORT INTERFACES STATES
	4.2	OUTGOING EVENTS

	4.3	INCOMING EVENTS
	4.4	TRANSPORT USER ACTIONS
	4.5	STATE TABLES
	4.6	EVENTS AND TLOOK ERROR INDICATION
Chapter	5	THE USE OF OPTIONS
	5.1	GENERALITIES
	5.2	THE FORMAT OF OPTIONS
	5.3	THE ELEMENTS OF NEGOTIATION
	5.3.1	Multiple Options and Options Levels
	5.3.2	Illegal Options
	5.3.3	Initiating an Option Negotiation
	5.3.4	Responding to a Negotiation Proposal
	5.3.5	Retrieving Information about Options
	5.3.6	Privileged and Read-only Options
	5.4	OPTION MANAGEMENT OF A TRANSPORT ENDPOINT
	5.5	SUPPLEMENTS
	5.5.1	The Option Value T_UNSPEC
	5.5.2	The info Argument
	5.5.3	Summary
	5.6	PORTABILITY ASPECTS
Chapter	6	XTI LIBRARY FUNCTIONS AND PARAMETERS
	6.1	HOW TO PREPARE XTI APPLICATIONS
	6.2	KEY FOR PARAMETER ARRAYS
	6.3	RETURN OF TLOOK ERROR
		<i>t_acceptt_ac</i>
		<i>t_alloct_al</i>
		<i>t_bindt_bi</i>
		<i>t_closet_cl</i>
		<i>t_connectt_co</i>
		<i>t_errort_er</i>
		<i>t_freet_fr</i>
		<i>t_getinfot_ge</i>
		<i>t_getprotaddrt_ge</i>
		<i>t_getstatet_ge</i>
		<i>t_listent_li</i>
		<i>t_lookt_lo</i>
		<i>t_opent_op</i>
		<i>t_optmgmtt_op</i>
		<i>t_rcvt_rc</i>

Contents

t_rcvconnectt_rc
t_rcvdist_rc
t_rcvreft_rc
t_rcvudatat_rc
t_rcvuderrt_rc
t_sndt_sn
t_snddist_sn
t_sndreft_sn
t_sndudatat_sn
t_strrerrort_st
t_syncnt_sy
t_unbindt_un

Appendix	A	ISO TRANSPORT PROTOCOL INFORMATION
	A.1	GENERAL
	A.2	OPTIONS
	A.2.1	Connection-mode Service
	A.2.1.1	Options for Quality of Service and Expedited Data (ISO 8072:1986)
	A.2.1.2	Management Options
	A.2.2	Connectionless-mode Service
	A.2.2.1	Options for Quality of Service (ISO 8072/Add.1:1986)
	A.2.2.2	Management Options
	A.3	FUNCTIONS
Appendix	B	INTERNET PROTOCOL-SPECIFIC INFORMATION
	B.1	GENERAL
	B.2	OPTIONS
	B.2.1	TCP-level Options
	B.2.2	UDP-level Options
	B.2.3	IP-level Options
	B.3	FUNCTIONS
Appendix	C	GUIDELINES FOR USE OF XTI
	C.1	TRANSPORT SERVICE INTERFACE SEQUENCE OF FUNCTIONS
	C.2	EXAMPLE IN CONNECTION-ORIENTED MODE
	C.3	EXAMPLE IN CONNECTIONLESS MODE
	C.4	WRITING PROTOCOL-INDEPENDENT SOFTWARE
	C.5	EVENT MANAGEMENT
	C.5.1	Introduction
	C.5.2	Short-term Solution

	C.5.3	XTI Events
	C.5.4	Guidelines for Use of System V poll()
	C.5.5	Guidelines for Use of BSD select()
Appendix	D	USE OF XTI TO ACCESS NETBIOS
	D.1	INTRODUCTION
	D.2	OBJECTIVES
	D.3	SCOPE
	D.4	ISSUES
	D.5	NetBIOS NAMES AND ADDRESSES
	D.6	NetBIOS CONNECTION RELEASE
	D.7	OPTIONS
	D.8	XTI FUNCTIONS
Appendix	E	XTI AND TLI
	E.1	RESTRICTIONS CONCERNING THE USE OF XTI
	E.2	RELATIONSHIP BETWEEN XTI AND TLI
Appendix	F	HEADERS AND DEFINITIONS
	F.1	THE <xti.h> HEADER
Appendix	G	ABBREVIATIONS
Appendix	H	GLOSSARY



Preface

X/Open

X/Open is an independent, worldwide, open systems organisation supported by most of the world's largest information systems suppliers, user organisations and software companies. Its mission is to bring to users greater value from computing, through the practical implementation of open systems.

X/Open's strategy for achieving this goal is to combine existing and emerging standards into a comprehensive, integrated, high-value and usable system environment, called the Common Applications Environment (CAE). This environment covers the standards, above the hardware level, that are needed to support open systems. It provides for portability and interoperability of applications, and allows users to move between systems with a minimum of retraining.

The components of the Common Applications Environment are defined in X/Open CAE Specifications. These contain, among other things, an evolving portfolio of practical application programming interfaces (APIs), which significantly enhance portability of application programs at the source code level, and definitions of, and references to, protocols and protocol profiles, which significantly enhance the interoperability of applications.

The X/Open CAE Specifications are supported by an extensive set of conformance tests and a distinct X/Open trademark - the XPG brand - that is licensed by X/Open and may be carried only on products that comply with the X/Open CAE Specifications.

The XPG brand, when associated with a vendor's product, communicates clearly and unambiguously to a procurer that the software bearing the brand correctly implements the corresponding X/Open CAE Specifications. Users specifying XPG-conformance in their procurements are therefore certain that the branded products they buy conform to the CAE Specifications.

X/Open is primarily concerned with the selection and adoption of standards. The policy is to use formal approved *de jure* standards, where they exist, and to adopt widely supported *de facto* standards in other cases.

Where formal standards do not exist, it is X/Open policy to work closely with standards development organisations to assist in the creation of formal standards covering the needed functions, and to make its own work freely available to such organisations. Additionally, X/Open has a commitment to align its definitions with formal approved standards.

X/Open Specifications

There are two types of X/Open specification:

- *CAE Specifications*

CAE (Common Applications Environment) Specifications are the long-life specifications that form the basis for conformant and branded X/Open systems. They are intended to be used widely within the industry for product development and procurement purposes.

Developers who base their products on a current CAE Specification can be sure that either the current specification or an upwards-compatible version of it will be referenced by a future XPG brand (if not referenced already), and that a variety of compatible, XPG-branded systems capable of hosting their products will be available, either immediately or in the near future.

CAE Specifications are not published to coincide with the launch of a particular XPG brand, but are published as soon as they are developed. By providing access to its specifications in this way, X/Open makes it possible for products that conform to the CAE (and hence are eligible for a future XPG brand) to be developed as soon as practicable, enhancing the value of the XPG brand as a procurement aid to users.

- *Preliminary Specifications*

These are specifications, usually addressing an emerging area of technology, and consequently not yet supported by a base of conformant product implementations, that are released in a controlled manner for the purpose of validation through practical implementation or prototyping. A Preliminary Specification is not a “draft” specification. Indeed, it is as stable as X/Open can make it, and on publication has gone through the same rigorous X/Open development and review procedures as a CAE Specification.

Preliminary Specifications are analogous with the “trial-use” standards issued by formal standards organisations, and product development teams are intended to develop products on the basis of them. However, because of the nature of the technology that a Preliminary Specification is addressing, it is untried in practice and may therefore change before being published as a CAE Specification. In such a case the CAE Specification will be made as upwards-compatible as possible with the corresponding Preliminary Specification, but complete upwards-compatibility in all cases is not guaranteed.

In addition, X/Open periodically publishes:

- *Snapshots*

Snapshots are “draft” documents, which provide a mechanism for X/Open to disseminate information on its current direction and thinking to an interested audience, in advance of formal publication, with a view to soliciting feedback and comment.

Preface

A Snapshot represents the interim results of an X/Open technical activity. Although at the time of publication X/Open intends to progress the activity towards publication of an X/Open Preliminary or CAE Specification, X/Open is a consensus organisation, and makes no commitment regarding publication.

Similarly, a Snapshot does not represent any commitment by any X/Open member to make any specific products available.

X/Open Guides

X/Open Guides provide information that X/Open believes is useful in the evaluation, procurement, development or management of open systems, particularly those that are X/Open-compliant.

X/Open Guides are not normative, and should not be referenced for purposes of specifying or claiming X/Open-conformance.

This Document

This document is a CAE Specification (see above). It defines the X/Open Transport Interface (XTI), a transport service interface that is independent of any specific transport provider. XTI is concerned primarily with the ISO Transport Service Definition (connection-oriented or connectionless). However, it may be adapted for use over other types of provider. In particular, XTI has been extended to include TCP and UDP, since these types are widely supported within the X/Open community.

This document merges the following two previous publications:

- **X/Open Developers' Specification (1990)**
Revised XTI (X/Open Transport Interface)
ISBN 1 872630 05 7
- **X/Open Addendum (August 1991)**
Addendum to Revised XTI
ISBN 1 872630 21 9

into a single publication.

It also contains a revised **Appendix D, Use of XTI to Access NETBIOS**. This Appendix was published as a Preliminary Specification in the **Revised XTI** Specification, and is now revised and upgraded to CAE status.

This XTI CAE Specification contains a number of changes and additions compared with the version published in the **X/Open Portability Guide, Issue 3**, arising principally from implementation experience by X/Open member companies. These changes and additions are detailed in **Chapter 2, Explanatory Notes**.

A compliant system shall meet the definitive requirements described in this XTI CAE Specification.

Preface



Trademarks

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

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

Palatino is a trademark of Linotype AG and/or its subsidiaries.

Referenced Documents

The following documents are referenced in this specification:

- The OSI model is described in:

ISO 7498, Information Processing Systems, Open Systems Interconnection, Basic Reference Model (IS: 1984)

- The reference documents for ISO transport are summarised below:

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

- The reference document for TCP protocol is:

TCP, Transmission Control Protocol, Military Standard, Mil-std-1778 (Source A*) and RFC 793 (Source B**)

- The reference document for UDP protocol is:

UDP, User Datagram Protocol, RFC 768 (Source B**)

- The reference document for the TLI specifications is:

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

- Mappings of NetBIOS services to OSI and IPS transport protocols are provided in the X/Open specification entitled **Protocols for X/Open PC Interworking: SMB**, published by X/Open Company Ltd., 1991.

* Source A:

Defense Communication Agency, DDN Protocol Handbook (Volume I), DOD Military Standard Protocols (December 1985).

** Source B:

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

General Introduction to the XTI

The XTI (X/Open Transport Interface) specification defines a transport service interface which is independent of any specific transport provider. The interface is provided by way of a set of library functions for the C programming language. XTI serves three main purposes:

- XTI describes a wide set of functions and facilities which vary in importance and/or usefulness;
- XTI is concerned primarily with the ISO Transport Service Definition (connection-oriented or connectionless). However, it may be adapted for use over other types of provider. In particular, XTI has been extended to include TCP and UDP since these types are widely supported within the X/Open community (see **Referenced Documents**);
- XTI is UNIX version-independent.

Note that 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. The EM facilities will allow the application to be notified of a number of events, including those events associated with flow control, over a range of active transport connections.

Explanatory Notes

2.1 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.

2.2 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 connect indications from several transport providers must open a transport endpoint for each provider and listen for connect indications on each of the associated file descriptors.

2.3 ASSOCIATION OF A UNIX PROCESS TO AN ENDPOINT

One process can simultaneously open several *fds*. 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 connect 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 connect 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 connect indications.

2.4 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.

2.5 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 (e.g., 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.

2.6 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`. An external integer, `t_errno`, which is defined in the header `<xti.h>`, holds the specific error number when such a failure occurs. 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. If implemented, 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 unusable. To continue in this case, all users of the `fd` must close it. Then the transport endpoint may be re-opened and initialised.

2.7 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 disconnect 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. Note that if a signal arrives, blocking calls are interrupted and return a negative return code with `t_errno` set to `[TSYSERR]` and `errno` set to `[EINTR]`. In this case the call will have no effect.

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 2.8, Event Management**.

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 the 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 clears it. 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 2-1 summarises this.

Event	Cleared on <i>t_look()</i> ?	Consuming XTI functions
T_LISTEN	No	<i>t_listen()</i>
T_CONNECT	No	<i>t_rcvconnect()</i> ¹
T_DATA	No	<i>t_rcv{udata}()</i>
T_EXDATA	No	<i>t_rcv()</i>
T_DISCONNECT	No	<i>t_rcvdis()</i>
T_UDERR	No	<i>t_rcvuderr()</i>
T_ORDREL	No	<i>t_rcvrel()</i>
T_GODATA	Yes	<i>t_snd{udata}()</i>
T_GOEXDATA	Yes	<i>t_snd()</i>

Table 2-1. Events and *t_look()*

2.8 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:

- **T_LISTEN:**
A connect request from a remote user was received by a transport provider (connection-mode service only); this event may occur under the following conditions:
 1. the file descriptor is bound to a valid address;
 2. no transport connection is established at this time.
- **T_CONNECT:**
In connection mode only; a connect response was received by the transport provider; occurs after a *t_connect()* has been issued.
- **T_DATA:**
Normal data (whole or part of Transport Service Data Unit (TSDU)) was received by the transport provider.
- **T_EXDATA:**
Expedited data was received by the transport provider.
- **T_DISCONNECT:**
In connection mode only; a disconnect request was received by the transport provider. It

1. 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.

may be reported on both data transfer functions and connection establishment functions and on the `t_snddis()` 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 `t_rcvudata()` or `t_unbind()` 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.

2.9 CHANGE HISTORY FROM XPG3 VERSION

This section summarises the main revisions to the XTI since publication of the **X/Open Portability Guide, Issue 3**, known as **XPG3**.

For ease of tracking, the changes are presented in two stages:

- those which appeared in the **Revised XTI Developers' Specification (1990)**
- those which appeared in the follow-up **Addendum to Revised XTI (August 1991)**.

These changes arise principally from implementation experience gathered by X/Open member companies.

2.9.1 Changes Appearing in Revised XTI (1990)

The major changes which appeared in the **Revised XTI (1990)** are:

Delete optional functions

The concept of mandatory versus optional functions is contrary to the goal of portability. All XTI functions are now mandatory; [TNOTSUPPORT] should be returned if the transport provider does not support the function requested.

Error messages

The format of messages produced by the `t_error()` function has been clarified. See also the additional function `t_strerror()`.

Multiple use of addresses

More stringent recommendations about multiple use of addresses have been made. This enhances portability across different transport providers.

State behaviour

The state machine behaviour of XTI has been clarified by the addition of a T_UNBND column in Table 4-7 of **Chapter 4, States and Events in XTI**, and by the identification of a number of additional cases where asynchronous events

result in the return of the TLOOK error.

Zero length TSDUs and TSDU fragments

The extent of support for zero length TSDUs and zero length TSDU fragments has been set out more clearly. See the descriptions of functions *t_snd()* and *t_getinfo()* in **Chapter 6, XTI Library Functions and Parameters**.

T_MORE

The significance of the T_MORE flag for asynchronously received data has been clarified. See the description of *t_rcv()* in **Chapter 6, XTI Library Functions and Parameters**.

Protocol options

The description of protocol options for both OSI and TCP has been much enhanced (see **Appendix A, ISO Transport Protocol Information** and **Appendix B, Internet Protocol-specific Information**).

Options and management structures

These have been extensively revised, especially those covering connection-oriented OSI (see **Appendix F, Headers and Definitions**).

Expedited data

The different significance of expedited data in the OSI and TCP cases has been clarified.

Connect semantics

Differences in underlying protocol semantics between OSI and TCP at connection establishment have been clarified. See **Appendix B, Internet Protocol-specific Information** and the descriptions of *t_accept()* and *t_listen()* in **Chapter 6, XTI Library Functions and Parameters**.

The main additions are:

Additional function: *t_getprotaddr*

This function yields the local and remote protocol addresses currently associated with a transport endpoint.

Additional function: *t_strerror*

This function maps an error number into a language-dependent error message string. The functionality corresponds to the error message changes in the *t_error()* function.

Addition of Valid States to function descriptions

All function descriptions now contain an indication of the interface states for which they are valid.

Addition of new error codes

A number of new error codes have been added (see **Appendix F, Headers and Definitions** for summary).

A number of minor changes also appeared, including:

- Clarification of the use of the term “socket” in the TCP case.

- Clarification of support for automatic generation of addresses.
- Clarification of the management of flow control.
- Clarification of the significant differences between transport providers.
- Clarification of the issue of non-guaranteed delivery of data at connection close.
- Clarification of the ways in which error indications may be received in connectionless working.
- Enhancement of `t_optmngmt()` to allow retrieval of current value of transport provider options.
- Addition of *extern* definitions for all XTI functions in **Appendix F, Headers and Definitions**.
- Numerous small corrections and clarifications.

2.9.2 Changes Appearing in Addendum to Revised XTI (1991)

The following list itemises the updates to the **Revised XTI (X/Open Transport Interface)** document, which appeared in the **Addendum to Revised XTI (August 1991)**. This list refers to the chapter, section and appendix references in the 1990 **Revised XTI** document.

All these changes are integrated into this **X/Open Transport Interface (XTI) CAE Specification**.

- **Chapter 2, Section 2.9.1, Changes**

The “Protocol options” and “Options and management structures” paragraphs are deleted and replaced with the following:

Option management

The management and usage of options have been completely revised. The changes affect **Chapter 5**, the `t_optmngmt()` manual pages in **Chapter 6**, and **Appendices A, B and F**.

- **Chapter 4, Section 4.5, State Tables**

Delete the row concerning “optmngmt” from **Figure 5**, and add a new row to **Figure 7** for the event “optmngmt”, as follows:

optmngmt	T_IDLE	T_OUTCON	T_INCON	T_DATAXFER	T_OUTREL	T_INREL	T_UNBND
----------	--------	----------	---------	------------	----------	---------	---------

- **Chapter 5, Transport Protocol-specific Options**

Chapter 5 is renamed to **The Use of Options** and is completely replaced with new text.

- **Chapter 6, `t_accept()`**

In the second paragraph, “protocol-specific parameters” is replaced with “options”.

In the sixth paragraph, the following sentence is removed:

“The values of parameters specified by *opt* and the syntax of those values are protocol-specific.”

In the seventh paragraph, “protocol-specific option” is replaced with “option”.

- **Chapter 6, *t_connect()***

The half-sentence in the sixth paragraph:

“If used, *sndcall->opt.buf* must point to the corresponding options structures (**isoco_options** or **tcp_options**);”

is replaced with:

“If used, *sndcall->opt.buf* must point to a buffer with the corresponding options;”

- **Chapter 6, *t_listen()***

In the second paragraph, “protocol-specific parameters” is replaced with “options”.

- **Chapter 6, *t_optmgmt()***

The manual pages for *t_optmgmt()* in **Chapter 6** are completely replaced with new text.

- **Chapter 6, *t_rcvconnect()***

In the third paragraph, “protocol-specific information” is replaced with “options”.

- **Chapter 6, *t_rcvudata()* and *t_rcvuderr()***

In the third paragraph, “protocol-specific options” is replaced with “options”.

- **Chapter 6, *t_sndudata()***

In the second paragraph, “protocol-specific options” is replaced with “options”.

- **Appendix A, ISO Transport Protocol Information**

The text in **Appendix A** is completely replaced with new text.

- **Appendix B, Internet Protocol-specific Information**

The text in **Appendix B** is completely replaced with new text.

- **Appendix F, Headers and Definitions**

The text in **Appendix F** is completely replaced with new text.

XTI Overview

3.1 OVERVIEW OF CONNECTION-ORIENTED MODE

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

- Initialisation/De-initialisation;
- Connection Establishment;
- Data Transfer, and
- Connection Release.

A state machine is described in **Section C.1, Transport Service Interface Sequence of Functions** and the figure in **Section C.2, Example in Connection-oriented Mode**, 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, e.g., the use of Class 0;
 - negotiation as to its use has been agreed between the two peer transport providers.

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.

5. The transport connection can be released at any time by using the disconnect 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.

3.1.1 Initialisation/De-initialisation Phase

The functions that support initialisation/de-initialisation tasks are described below. All such functions provide local management functions; no information is sent over the network.

<i>t_open()</i>	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.
<i>t_bind()</i>	This function associates a protocol address with a given transport endpoint, thereby activating the endpoint. It also directs the transport provider to begin accepting connect indications if so desired.
<i>t_optmgmt()</i>	This function enables the user to get or negotiate protocol options with the transport provider.
<i>t_unbind()</i>	This function disables a transport endpoint such that no further request destined for the given endpoint will be accepted by the transport provider.
<i>t_close()</i>	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:

<i>t_getprotaddr()</i>	This function returns the addresses (local and remote) associated with the specified transport endpoint.
<i>t_getinfo()</i>	This function returns protocol-specific information associated with the specified transport endpoint.
<i>t_getstate()</i>	This function returns the current state of the transport endpoint.
<i>t_sync()</i>	This function synchronises the data structures managed by the transport library with the transport provider.
<i>t_alloc()</i>	This function allocates storage for the specified library data structure.
<i>t_free()</i>	This function frees storage for a library data structure that was allocated by <i>t_alloc()</i> .
<i>t_error()</i>	This function prints out a message describing the last error encountered during a call to a transport library function.
<i>t_look()</i>	This function returns the current event(s) associated with the given transport endpoint.
<i>t_strerror()</i>	This function maps an XTI error into a language-dependent error message string.

3.1.2 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 firstly to establish an endpoint, i.e., 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 connect indications (only if the provider is able to accept more than one outstanding connect 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 connect indication; if it finds one (at the head of the queue), it returns details of the connect 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 connect 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 connect indication enqueued, the user should dequeue all connect indications (and disconnect indications) before accepting or rejecting any or all of them. The number of outstanding connect indications is limited by the value of the *qlen* parameter accepted by the transport provider on the call to *t_bind()*;
- when accepting a connect 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 can no longer be used to receive and enqueue incoming connect indications. The bound protocol address will be found to be busy for the duration of the active transport endpoint. No other transport endpoints may be bound for listening to the same protocol address while the listening endpoint is in the data transfer or disconnect phase (i.e., until a *t_unbind()* call is issued).

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

- if the user issues a *t_connect()* on a listening endpoint, again, that endpoint can no longer be used to receive and enqueue incoming connect requests;
- a connect 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:

<i>t_connect()</i>	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.
<i>t_rcvconnect()</i>	This function enables an active transport user to determine the status of a previously sent connect 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 <i>t_connect()</i> to establish a connection in an asynchronous manner.
<i>t_listen()</i>	This function enables the passive transport user to receive connect indications from other transport users.
<i>t_accept()</i>	This function is issued by the passive user to accept a particular connect request after an indication has been received.

3.1.3 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. Two functions have been defined to support data transfer in connection mode as follows:

<i>t_snd()</i>	This function enables transport users to send either normal or expedited data over a transport connection.
<i>t_rcv()</i>	This function enables transport users to receive either normal or expedited data on a transport connection.

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 disconnect indication is received, or
- 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 (e.g., transport connection lost), then the call returns immediately, indicating failure. If the data cannot be accepted immediately because of a transient condition (e.g., 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 disconnect 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. Data that was being sent will be lost, or
- 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 **Chapter 6, *t_getinfo()***.

3.1.4 Overview of Connection Release

The ISO Connection-oriented Transport Service Definition supports only the abortive release. However, the TCP Transport Service Definition also supports an orderly release. Some XTI implementations may support this orderly release.

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 connect 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 disconnect 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 disconnect request.

The orderly release capability is an optional feature of TCP. If supported by the TCP transport 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.

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 connect 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.

3.2 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, Transport Service Interface Sequence of Functions** and the figure in **Section C.3, Example in Connectionless Mode**, that defines the legal sequence in which functions from each phase may be issued.

In order to permit the transfer of connectionless 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()`), and
3. the user may then send and/or receive connectionless 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.

3.2.1 Initialisation/De-initialisation Phase

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

3.2.2 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_rcvudata()</code>	This function enables transport users to receive data units from other users.
<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 transport provider.

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, or
- 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 returns when sending a new datagram becomes possible again. 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.

3.3 XTI FEATURES

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

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

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

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

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

```
t_rcvudata()
t_rcvuderr()
t_sndudata()
```

Mandatory mechanisms:

- Synchronous mode
- Asynchronous mode

Utility functions:

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

The orderly release mechanism (using *t_sndrel()* and *t_rcvrel()*), is supported only for T_COTS_ORD type providers. Use with other providers will cause the [TNOTSUPPORT] error to be returned. The use of orderly release is definitely not recommended in order to make applications using TCP portable onto the ISO Transport Layer.

Optional mechanisms:

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

3.3.1 XTI Functions Versus Protocols

Table 3-1 presents all the functions defined in XTI. The character 'x' indicates that the mapping of that function is possible onto a Connection-oriented or Connectionless Transport Service. The table indicates the type of utility functions as well.

Functions	Necessary for Protocol		Utility Functions	
	Connection Oriented	Connectionless	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_rcvconnect()</i>	x			
<i>t_rcvdis()</i>	x			
<i>t_rcvrel()</i>	x			
<i>t_rcvudata()</i>		x		
<i>t_rcvuderr()</i>		x		
<i>t_snd()</i>	x			
<i>t_snddis()</i>	x			
<i>t_sndrel()</i>	x			
<i>t_sndudata()</i>		x		
<i>t_strerror()</i>			x	
<i>t_sync()</i>			x	
<i>t_unbind()</i>	x	x		

Table 3-1. Classification of the XTI Functions

States and Events in XTI

Tables 4-1 to 4-7 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 Tables 4-5, 4-6 and 4-7. 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()* and *t_strerror()* functions have been excluded from the state table because they do not affect the state of the interface.

4.1 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 4-1 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 4-1. Transport Interface States

4.2 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 4-2, some events (e.g., *acceptX*) are distinguished by the context in which they occur. The context is based on the values of the following:

ocnt Count of outstanding connect indications (connect 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 disconnect 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> resfd	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>	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>	T_COTS_ORD
sndudata	successful return of <i>t_sndudata()</i>	T_CLTS

Table 4-2. Transport Interface Outgoing Events

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

4.3 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*, which 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 4-3, 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 connect 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>	T_COTS T_COTS_ORD
rcvdis1	successful return of <i>t_rcvdis()</i> with <i>ocnt</i> == 0	T_COTS T_COTS_ORD
rcvdis2	successful return of <i>t_rcvdis()</i> with <i>ocnt</i> == 1	T_COTS T_COTS_ORD
rcvdis3	successful return of <i>t_rcvdis()</i> with <i>ocnt</i> > 1	T_COTS T_COTS_ORD
rcvrel	successful return of <i>t_rcvrel()</i>	T_COTS_ORD
rcvudata	successful return of <i>t_rcvudata()</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 4-3. Transport Interface Incoming Events

4.4 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 4-4.

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

Table 4-4. Transport Interface User Actions

4.5 STATE TABLES

Tables 4-5, 4-6 and 4-7 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 4-4). 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 4-5. Initialisation/De-initialisation State Table

state event	T_IDLE
sndudata	T_IDLE
rcvudata	T_IDLE
rcvuderr	T_IDLE

Table 4-6. Data Transfer State Table for Connectionless-Mode Service

state event	T_IDLE	T_OUTCON	T_INCON	T_DATAXFER	T_OUTREL	T_INREL	T_UNBND
<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_IDLE	T_OUTCON	T_INCON	T_DATAXFER	T_OUTREL	T_INREL	T_UNBIND
<i>closed</i>	T_UNINIT	T_UNINIT	T_UNINIT	T_UNINIT	T_UNINIT	T_UNINIT	

Figure 4-7. Connection/Release/Data Transfer State Table for Connection-mode Service

4.6 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 ¹
<i>t_listen()</i>	T_DISCONNECT ²
<i>t_rcv()</i>	T_DISCONNECT, T_ORDREL ³
<i>t_rcvconnect()</i>	T_DISCONNECT
<i>t_rcvrel()</i>	T_DISCONNECT
<i>t_rcvudata()</i>	T_UDERR
<i>t_snd()</i>	T_DISCONNECT, T_ORDREL
<i>t_sndudata()</i>	T_UDERR
<i>t_unbind()</i>	T_LISTEN, T_DATA ⁴
<i>t_sndrel()</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 **Chapter 2, Table 2-1, Events and t_look()**.

-
1. This occurs only when a *t_connect* is done on an endpoint which has been bound with a *qlen* > 0 and for which a connect indication is pending.
 2. This event indicates a disconnect on an outstanding connect indication.
 3. This occurs only when all pending data has been read.
 4. T_DATA may only occur for the connectionless mode.

The Use of Options

5.1 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.

All options have default values. Their values have meaning to and are defined by the protocol level in which they apply. However, their 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 the proposed one, i.e., a delay may become longer, or a throughput may become lower.

It is useful to differentiate between options that are *association-related*¹ 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-oriented 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 that are not association-related do not contain information destined for the remote transport user. Some have purely local relevance, e.g., an option that enables debugging. Others influence the transmission, for instance the option that sets the IP *time-to-live* field, or TCP_NODELAY (see **Appendix B, Internet Protocol-specific Information**). 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 association-related options only. The functions `t_rcvconnect()` and `t_rcvuderr()` may return options of both

1. The term “association” is used to denote a pair of communicating transport users.

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.

The transport provider has a default value for each option it supports. 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, ISO Transport Protocol Information**, and **Appendix B, Internet Protocol-specific Information**). General XTI options are described in *t_optmgmt()* in **Chapter 6, XTI Library Functions and Parameters**.

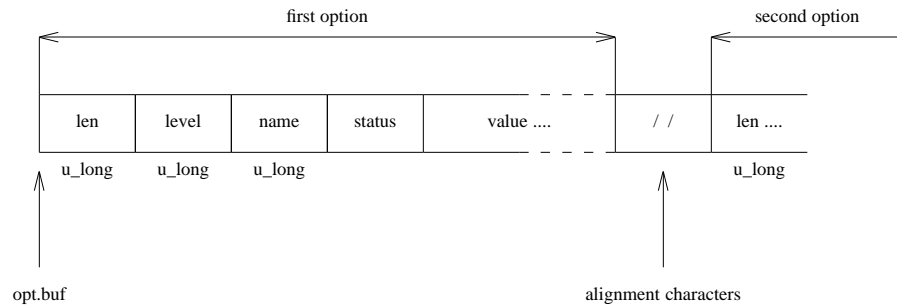
5.2 THE FORMAT OF OPTIONS

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 as TCP or ISO 8073:1986. The *name* field identifies the option within the level, and *len* contains its total length, i.e., 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 5.3.5, Retrieving Information about Options** and *t_optmgmt()* in **Chapter 6, XTI Library Functions and Parameters**).

Several options can be concatenated. The transport user has, however, to ensure that each option starts at a long-word boundary. The macro **OPT_NEXTHDR(pbuf, buflen, poption)** can be used for that purpose. The parameter *pbuf* denotes a pointer to an option buffer *opt.buf*, and *buflen* is its length. The parameter *pooption* points to the current option in the option buffer. **OPT_NEXTHDR** returns a pointer to the position of the next option, or returns a null pointer if the option buffer is exhausted. The macro is helpful for writing and reading. See **<xti.h>** in **Appendix F, Headers and Definitions** for the exact definition.

The option buffer thus has the following form (unsigned long is abbreviated to *u_long*):



The length of the option buffer is given by *opt.len*.

5.3 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.

5.3.1 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.

5.3.2 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()` in **Chapter 6, XTI Library Functions and Parameters, Appendix A, ISO Transport Protocol Information** and **Appendix B, Internet Protocol-specific Information**.)

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

- A call to `t_optmgmt()` fails with [TBADOPT].
- `t_accept()` or `t_connect()` fail either with [TBADOPT], or the connection establishment aborts, depending on the implementation and the time the illegal option is detected. If the connection aborts, a T_DISCONNECT event occurs, and 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] in that case.
- A call to `t_sndudata()` either fails with [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()` in **Chapter 6, XTI Library Functions and Parameters**).

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 *level* field of the option.

5.3.3 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()` in **Chapter 6, XTI Library Functions and Parameters, Appendix A, ISO Transport Protocol Information** and **Appendix B, Internet Protocol-specific Information**). 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 5.4, Option Management of a Transport Endpoint**).

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 (e.g., 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 5.4, Option Management of a Transport Endpoint**). 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* association-related options 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 5.4, Option Management of a Transport Endpoint**).

5.3.4 Responding to a Negotiation Proposal

In connection-oriented communication, some protocols give the peer transport users the opportunity to negotiate characteristics of the transport connection to be established. These characteristics are association-related options. With the connect 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 (e.g., 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 5.4, Option Management of a Transport Endpoint**).

The response options can either be specified with the *t_accept()* call, or can be preset for the *r_resfd* in a *t_optmgmt()* call (action T_NEGOTIATE) prior to *t_accept()* (see **Section 5.4, Option Management of a Transport Endpoint**). 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 will 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 5.3.3, Initiating an Option Negotiation**). Unsupported options are ignored.

5.3.5 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 (e.g., 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 (e.g., 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 association-related options that were received with the connection response and the negotiated values of those non-association-related options 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 (e.g., IP options) that is not subject to negotiation is always set to T_SUCCESS.

t_listen() The received association-related options 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, *t_optmgmt()* will usually not return the same values.

The number of received options may be variable for subsequent connect indications, since many association-related options are only transmitted on explicit demand by the calling user (e.g., 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 association-related options 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 connect indication or an incoming datagram. A detailed description is given in *t_optmgmt()* in **Chapter 6, XTI Library Functions and Parameters**.

5.3.6 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_optmgt()* 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]. 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.) It depends on timing and implementation conditions whether a *t_accept()* call still succeeds or fails with [TLOOK].
- 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 5.4, Option Management of a Transport Endpoint**).

5.4 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 in **Section 5.3, The Elements of Negotiation**.

Some options may be modified only in specific XTI states and are read-only in other XTI states. Many association-related options, for instance, may not be changed in the state `T_DATAXFER`, and an attempt to do so will fail (see **Section 5.3.6, Privileged and Read-only Options**). The legal states for each option are specified with its definition.

As usual, association-related options 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`, e.g., those specifying buffer sizes. Such changes might affect the transmission characteristics and lead to unexpected side effects (e.g., 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 of **Section 5.3.3, Initiating an Option Negotiation**, and **Section 5.3.4, Responding to a Negotiation Proposal**, 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.

5.5 SUPPLEMENTS

This section contains supplementary remarks and a short summary.

5.5.1 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, e.g., T_CLASS2. The transport provider would then be forced to initiate a connect 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 TCO_THROUGHPUT in **Appendix A, ISO Transport Protocol Information**). 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.

For each option, it is specified whether or not T_UNSPEC is a legal value for negotiation purposes.

5.5.2 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* possibly diminish as soon as the T_DATAXFER state is entered. Calling *t_optmgmt()* does not influence these values (see **Chapter 6, t_optmgmt()**).

5.5.3 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 long-word boundary.
- There are options that are association-related and options that are not. On output, the functions *t_listen()* and *t_rcvudata()* return association-related options 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.

5.6 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

6.1 HOW TO PREPARE XTI APPLICATIONS

In a software development environment, a program, for example *file.c*, that uses XTI functions must be compiled with the XTI Library. This can be done using the following command:

```
cc file.c -lxti           (e.g., for normal library)
```

The syntax for shared libraries is implementation-dependent.

The XTI structures and constants are all defined in the `<xti.h>` header, which can be found in **Appendix F, Headers and Definitions**.

6.2 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.

6.3 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 4.6, Events and TLOOK Error Indication**.

NAME

t_accept – accept a connect request

SYNOPSIS

```
#include <xti.h>

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

DESCRIPTION

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

This function is issued by a transport user to accept a connect request. The parameter *fd* identifies the local transport endpoint where the connect 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 connect 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 connect indication arrived. Before the connection can be accepted on the same endpoint (*resfd==fd*), the user must have responded to any previous connect 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()*, then the transport provider will automatically bind it to the same protocol address *fd* is bound to. 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.

The call to *t_accept()* will fail with *t_errno* set to [TLOOK] if there are indications (e.g., connect or disconnect) waiting to be received on the endpoint *fd*.

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) it is assumed that the connection is to be accepted unconditionally. The transport provider may choose options other than the defaults to ensure that the connection is accepted successfully.

CAVEATS

There may be transport provider-specific restrictions on address binding. See **Appendix A, ISO Transport Protocol Information** and **Appendix B, Internet Protocol-specific Information**.

Some transport providers do not differentiate between a connect 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, Internet Protocol-specific Information**).

VALID STATES

fd: T_INCON
r T_IDLE

ERRORS

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

[TBADF]	The file descriptor <i>fd</i> or <i>resfd</i> does not refer to a transport endpoint.
[TOUTSTATE]	The function was called in the wrong sequence on the transport endpoint referenced by <i>fd</i> , or the transport endpoint referred to by <i>resfd</i> is not in the appropriate state.
[TACCES]	The user does not have permission to accept a connection on the responding transport endpoint or to use the specified options.
[TBADOPT]	The specified options were 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.
[TBADADDR]	The specified protocol address was in an incorrect format or contained illegal information.

[TBADSEQ]	An invalid sequence number was specified.
[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.
[TSYSERR]	A system error has occurred during execution of this function.
[TINDOUT]	The function was called with <i>fd==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> .
[TPROVMISMATCH]	The file descriptors <i>fd</i> and <i>resfd</i> do not refer to the same transport provider.
[TRESQLEN]	The endpoint referenced by <i>resfd</i> (where <i>resfd</i> fd) was bound to a protocol address with a <i>qlen</i> that is greater than zero.
[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 (<i>t_errno</i>).
[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.

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()*.

t_alloc()

XTI Library Functions and Parameters

NAME

t_alloc – allocate a library structure

SYNOPSIS

```
#include <xti.h>

char *t_alloc(fd, struct_type, fields)
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 *fd* 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 this way the appropriate size information can be accessed. If the size value associated with any specified field is -1 or -2 (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]. For any field not specified in *fields*, *buf* will be set to the null pointer and *len* and *maxlen* will be set to zero.

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:

[TBADEF]	The specified file descriptor does not refer to a transport endpoint.
[TSYSERR]	A system error has occurred during execution of this function.
[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, i.e., connection-oriented or connectionless.
[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 (<i>t_errno</i>).

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()*.

t_bind()

XTI Library Functions and Parameters

NAME

t_bind – bind an address to a transport endpoint

SYNOPSIS

```
#include <xti.h>

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

DESCRIPTION

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

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 connect 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 the address that the transport provider actually bound to the transport endpoint; this is the same as the address specified by the user in *req*. 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* is not large enough to hold the returned address, an error will result.

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 connect indications that the transport provider should support for the given transport endpoint. An outstanding connect 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 connect 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 connect 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 connect 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 connect indications.

If *fd* refers to a connectionless-mode service, only one endpoint may be associated with a protocol address. If a user attempts to bind a second transport endpoint to an already bound protocol address, *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:

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TOUTSTATE]	The function was issued in the wrong sequence.
[TBADADDR]	The specified protocol address was in an incorrect format or contained illegal information.
[TNOADDR]	The transport provider could not allocate an address.
[TACCES]	The user does not have permission to use the specified address.
[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.
[TSYSERR]	A system error has occurred during execution of this function.
[TADDRBUSY]	The requested address is in use.
[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 (<i>t_errno</i>).

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_close()*, *t_open()*, *t_optmgmt()*, *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.

A transport provider may not allow an explicit binding of more than one transport endpoint to the same protocol address, although it allows more than one connection to be accepted for the same protocol address. To ensure portability, it is, therefore, recommended not to bind transport endpoints that are used as responding endpoints (*resfd*) in a call to *t_accept()*, if the responding address is to be the same as the called address.

NAME

t_close – close a transport endpoint

SYNOPSIS

#include <xti.h>

int t_close(fd)

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; the *close()* will be abortive if there are no other descriptors in this, or in another process which references the transport endpoint, and in this case will break any transport connection that may be associated with that endpoint.

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.

VALID STATES

ALL - apart from T_UNINIT

ERRORS

On 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 (<i>t_errno</i>).

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()*.

t_connect()

XTI Library Functions and Parameters

NAME

t_connect – establish a connection with another transport user

SYNOPSIS

```
#include <xti.h>

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

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>sndcall->addr.maxlen</i>	/	/
<i>sndcall->addr.len</i>	x	/
<i>sndcall->addr.buf</i>	x (x)	/
<i>sndcall->opt.maxlen</i>	/	/
<i>sndcall->opt.len</i>	x	/
<i>sndcall->opt.buf</i>	x (x)	/
<i>sndcall->udata.maxlen</i>	/	/
<i>sndcall->udata.len</i>	x	/
<i>sndcall->udata.buf</i>	? (?)	/
<i>sndcall->sequence</i>	/	/
<i>rcvcall->addr.maxlen</i>	x	/
<i>rcvcall->addr.len</i>	/	x
<i>rcvcall->addr.buf</i>	?	(?)
<i>rcvcall->opt.maxlen</i>	x	/
<i>rcvcall->opt.len</i>	/	x
<i>rcvcall->opt.buf</i>	?	(?)
<i>rcvcall->udata.maxlen</i>	x	/
<i>rcvcall->udata.len</i>	/	x
<i>rcvcall->udata.buf</i>	?	(?)
<i>rcvcall->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 the underlying protocol of the transport provider and are described for ISO and TCP protocols in **Appendix A, ISO Transport Protocol Information**, **Appendix B, Internet Protocol-specific Information** and **Appendix F, Headers and Definitions**. The user may choose not to negotiate protocol options by setting the *len* field of *opt* to zero. In this case, the provider may use default options.

If used, *sndcall->opt.buf* must point to a buffer with the corresponding options; 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, *rcvcall* may be a null pointer, in which case no information is given to the user on return from *t_connect()*.

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 (i.e., 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 connect 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()*.

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.
[TOUTSTATE]	The function was issued in the wrong sequence.
[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.
[TBADADDR]	The specified protocol address was in an incorrect format or contained illegal information.
[TBADOPT]	The specified protocol options were 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.
[TACCES]	The user does not have permission to use the specified address or options.
[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.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TSYSERR]	A system error has occurred during execution of this function.
[TADDRBUSY]	This transport provider does not support multiple connections with the same local and remote addresses. This error indicates that a connection already exists.
[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 (<i>t_errno</i>).

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()*.

NAME

t_error – produce error message

SYNOPSIS

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

```
int t_error(errmsg)
```

```
char *errmsg;
```

DESCRIPTION

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

The *t_error()* function produces a language-dependent 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 implementation-defined. If it is in English, the error message string describing the value in *t_errno* is identical to 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.

t_error()

XTI Library Functions and Parameters

RETURN VALUE

Upon completion, a value of 0 is returned.

NAME

t_free – free a library structure

SYNOPSIS

```
#include <xti.h>

int t_free(ptr, struct_type)
char *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

ERRORS

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

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

t_free()

XTI Library Functions and Parameters

suitable XTI (*t_errno*).

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().

NAME

t_getinfo – get protocol-specific service information

SYNOPSIS

```
#include <xti.h>

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

DESCRIPTION

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

```
long addr;          /* max size of the transport protocol address */
long options;       /* max number of bytes of protocol-specific options */
long tsdu; /* max size of a transport service data unit (TSDU) */
long etsdu; /* max size of an expedited transport service */
                /* data unit (ETSDU) */
long connect; /* max amount of data allowed on connection */
                /* establishment functions */
long discon; /* max amount of data allowed on t_snddis() */
                /* and t_rcvdis() functions */
long servtype; /* service type supported by the transport provider */
long 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 -2 specifies that the transport provider does not provide user access to transport protocol addresses.

<i>options</i>	A value greater than zero indicates the maximum number of bytes of protocol-specific options supported by the provider, and a value of -2 specifies that the transport provider does not support user-settable options.
<i>tsdu</i>	A value greater than zero specifies the maximum size of a transport service data unit (TSDU); a value of 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 -1 specifies that there is no limit on the size of a TSDU; and a value of -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 of an expedited transport service data unit (ETSDU); a value of 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 -1 specifies that there is no limit on the size of an ETSDU; and a value of -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, ISO Transport Protocol Information and Appendix B, Internet Protocol-specific Information).
<i>connect</i>	A value greater than zero specifies the maximum amount of data that may be associated with connection establishment functions and a value of -2 specifies that the transport provider does not allow data to be sent with connection establishment functions.
<i>discon</i>	A value greater than zero 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 specifies that the transport provider does not allow data to be sent with the abortive release 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 transport provider. 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, ISO Transport Protocol Information 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 affect 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, <i>t_open()</i> will return -2 for <i>etsdu</i> , <i>connect</i> and <i>discon</i> .

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.
[TSYSERR]	A system error has occurred during execution of this function.
[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 (<i>t_errno</i>).

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()*.

t_getprotaddr()

XTI Library Functions and Parameters

NAME

t_getprotaddr – get the protocol addresses

SYNOPSIS

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

```
int t_getprotaddr(fd, boundaddr, peeraddr)
```

```
int fd;
```

```
struct t_bind *boundaddr;
```

```
struct t_bind *peeraddr;
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>boundaddr->maxlen</i>	x	/
<i>boundaddr->addr.len</i>	/	x
<i>boundaddr->addr.buf</i>	?	(?)
<i>boundaddr->qlen</i>	/	/
<i>peeraddr->maxlen</i>	x	/
<i>peeraddr->addr.len</i>	/	x
<i>peeraddr->addr.buf</i>	?	(?)
<i>peeraddr->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 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 state, zero is returned in the *len* field of *peeraddr*.

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.
[TSYSERR]	A system error has occurred during execution of this function.
[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 (<i>t_errno</i>).

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().

t_getstate()

XTI Library Functions and Parameters

NAME

t_getstate – get the current state

SYNOPSIS

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

```
int t_getstate(fd)
```

```
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

ERRORS

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

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TSTATECHNG]	The transport provider is undergoing a transient state change.
[TSYSERR]	A system error has occurred during execution of this function.
[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 (<i>t_errno</i>).

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 orderly release (waiting for an orderly release indication)
T_INREL	incoming orderly release (waiting to send an orderly release request)

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

SEE ALSO

t_open().

NAME

t_listen – listen for a connect indication

SYNOPSIS

```
#include <xti.h>

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

DESCRIPTION

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

This function listens for a connect request from a calling transport user. The argument *fd* identifies the local transport endpoint where connect indications arrive, and on return, *call* contains information describing the connect 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 connect request, *udata* returns any user data sent by the caller on the connect request, and *sequence* is a number that uniquely identifies the returned connect indication. The value of *sequence* enables the user to listen for multiple connect 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.

By default, *t_listen()* executes in synchronous mode and waits for a connect 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 connect 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:

[TBADEF]	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 <code>T_INCON</code> , and the connect indication information to be returned in <i>call</i> is discarded. The value of <i>sequence</i> returned can be used to do a <code>t_snddis()</code> .
[TNODATA]	<code>O_NONBLOCK</code> was set, but no connect indications had been queued.
[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 function was issued in the wrong sequence on the transport endpoint referenced by <i>fd</i> .
[TSYSERR]	A system error has occurred during execution of this function.
[TQFULL]	The maximum number of outstanding indications has been reached for the endpoint referenced by <i>fd</i> .
[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 (<i>t_errno</i>).

CAVEATS

Some transport providers do not differentiate between a connect indication and the connection itself. If this is the case, a successful return of `t_listen()` indicates an existing connection (see **Appendix B, Internet Protocol-specific Information**).

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()`.

NAME

t_look – look at the current event on a transport endpoint

SYNOPSIS

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

```
int t_look(fd)
```

```
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 [T_LOOK] may be found in **Section 4.6, Events and TLOOK Error Indication**.

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.
[TSYSERR]	A system error has occurred during execution of this function.
[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 (<i>t_errno</i>).

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	disconnect received

t_look()

XTI Library Functions and Parameters

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.

NAME

t_open – establish a transport endpoint

SYNOPSIS

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

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

DESCRIPTION

Parameters	Before call	After call
<i>name</i>	x	/
<i>oflag</i>	x	/
<i>info->addr</i>	/	x
<i>info->options</i>	/	x
<i>info->tsdu</i>	/	x
<i>info->etsdu</i>	/	x
<i>info->connect</i>	/	x
<i>info->discon</i>	/	x
<i>info->servtype</i>	/	x
<i>info->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 (i.e., 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:

```
long addr;          /* max size of the transport protocol */
                   /* address */
long options;       /* max number of bytes of */
                   /* protocol-specific options */
long tsdu; /* max size of a transport service data */
                   /* unit (TSDU) */
long etsdu;         /* max size of an expedited transport */
```

```

/* service data unit (ETSDU) */
long connect; /* max amount of data allowed on */
/* connection establishment functions */
long discon; /* max amount of data allowed on */
/* t_snddis() and t_rcvdis() functions */
long servtype; /* service type supported by the */
/* transport provider */
long flags; /* other info about the transport provider */

```

The values of the fields have the following meanings:

<i>addr</i>	A value greater than zero indicates the maximum size of a transport protocol address and a value of -2 specifies that the transport provider does not provide user access to transport protocol addresses.
<i>options</i>	A value greater than zero indicates the maximum number of bytes of protocol-specific options supported by the provider and a value of -2 specifies that the transport provider does not support user-settable options.
<i>tsdu</i>	A value greater than zero specifies the maximum size of a transport service data unit (TSDU); a value of zero 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 specifies that there is no limit to the size of a TSDU; and a value of -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 of an expedited transport service data unit (ETSDU); a value of 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 -1 specifies that there is no limit on the size of an ETSDU; and a value of -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, ISO Transport Protocol Information and Appendix B, Internet Protocol-specific Information).
<i>connect</i>	A value greater than zero specifies the maximum amount of data that may be associated with connection establishment functions and a value of -2 specifies that the transport provider does not allow data to be sent with connection establishment functions.
<i>discon</i>	A value greater than zero 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 specifies that the transport provider does not allow data to be sent with the abortive release functions.
<i>servtype</i>	This field specifies the service type supported by the transport provider, as described below.

flags This is a bit field used to specify other information about the transport provider. 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, ISO Transport Protocol Information** 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 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.
[TSYSERR]	A system error has occurred during execution of this function.
[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 (<i>t_errno</i>).

RETURN VALUE

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().

t_optmgmt()

XTI Library Functions and Parameters

NAME

t_optmgmt - manage options for a transport endpoint

SYNOPSIS

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

```
int t_optmgmt(fd, req, ret)
```

```
int fd;
```

```
struct t_optmgmt *req;
```

```
struct t_optmgmt *ret;
```

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>req->opt.maxlen</i>	/	/
<i>req->opt.len</i>	x	/
<i>req->opt.buf</i>	x (x)	/
<i>req->flags</i>	x	/
<i>ret->opt.maxlen</i>	x	/
<i>ret->opt.len</i>	/	x
<i>ret->opt.buf</i>	?	(?)
<i>ret->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;  
long 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. 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, i.e., 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.

Each option in the input or output option buffer must start at a long-word boundary. The macro **OPT_NEXTHDR(pbuf, buflen, poption)** can be used for that purpose. The parameter *pbuf* denotes a pointer to an option buffer *opt.buf*, and *buflen* is its length. The parameter *poption* points to the current option in the option buffer. **OPT_NEXTHDR** returns a pointer to the position of the next option or returns a null pointer if the option buffer is exhausted. The macro is helpful for writing and reading. See <xti.h> in **Appendix F, Headers and Definitions** for the exact definition.

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 5, The Use of Options 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, i.e., 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 are not association-related (see **Chapter 5, The Use of Options**). 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	array of unsigned longs	see text	enable debugging
XTI_LINGER	struct linger	see text	linger on close if data is present
XTI_RCVBUF	unsigned long	size in octets	receive buffer size
XTI_RCVLOWAT	unsigned long	size in octets	receive low-water mark
XTI_SNDBUF	unsigned long	size in octets	send buffer size
XTI_SNDLOWAT	unsigned long	size in octets	send low-water mark

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’, i.e., 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 {
    long l_onoff;      /* switch option on/off */
    long 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 (e.g., *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.
[TOUTSTATE]	The function was issued in the wrong sequence.
[TACCES]	The user does not have permission to negotiate the specified options.
[TBADOPT]	The specified options were in an incorrect format or contained illegal information.
[TBADFLAG]	An invalid flag was specified.
[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.
[TSYSERR]	A system error has occurred during execution of this function.
[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 (<i>t_errno</i>).
[TNOTSUPPORT]	This action is not supported by the transport provider.

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 5, The Use of Options.**

NAME

t_rcv – receive data or expedited data sent over a connection

SYNOPSIS

```
#include <xti.h>

int t_rcv(fd, buf, nbytes, flags)
int fd;
char *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, 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 returned is expedited data if T_EXPEDITED is set in *flags*. If the number of bytes of expedited data exceeds *nbytes*, *t_rcv()* will set T_EXPEDITED and T_MORE on return from the initial call. Subsequent calls to retrieve the remaining ETSDU will have T_EXPEDITED set on return. The end of the ETSDU is identified by the return of a *t_rcv()* call with the T_MORE flag not set.

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.
[TNODATA]	O_NONBLOCK was set, but no data is currently available from the transport provider.
[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 function was issued in the wrong sequence on the transport endpoint referenced by <i>fd</i> .
[TSYSERR]	A system error has occurred during execution of this function.
[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 (<i>t_errno</i>).

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()*.

NAME

t_rcvconnect – receive the confirmation from a connect request

SYNOPSIS

```
#include <xti.h>

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

DESCRIPTION

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

This function enables a calling transport user to determine the status of a previously sent connect request and is used in conjunction with *t_connect()* to establish a connection in asynchronous mode. 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, *call* may be a null pointer, in which case no information is given to the user on return from *t_rcvconnect()*. 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 connect 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

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, and the connect 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.
[TNODATA]	<code>O_NONBLOCK</code> was set, but a connect confirmation has not yet arrived.
[TLOOK]	An asynchronous event has occurred on this transport connection and requires immediate attention.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The function was issued in the wrong sequence on the transport endpoint referenced by <i>fd</i> .
[TSYSERR]	A system error has occurred during execution of this function.
[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 (<i>t_errno</i>).

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_bind()*, *t_connect()*, *t_listen()*, *t_open()*, *t_optmgmt()*.

NAME

t_rcvdis – retrieve information from disconnect

SYNOPSIS

```
#include <xti.h>

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

DESCRIPTION

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

This function is used to identify the cause of a disconnect and to retrieve any user data sent with the disconnect. 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 disconnect through a protocol-dependent reason code, *udata* identifies any user data that was sent with the disconnect, and *sequence* may identify an outstanding connect indication with which the disconnect 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 connect indications. If a disconnect indication occurs, *sequence* can be used to identify which of the outstanding connect indications is associated with the disconnect.

If a user does not care if there is incoming data and does not need to know the value of *reason* or *sequence*, *discon* may be a null pointer and any user data associated with the disconnect will be discarded. However, if a user has retrieved more than one outstanding connect indication (via *t_listen()*) and *discon* is a null pointer, the user will be unable to identify with which connect indication the disconnect is associated.

VALID STATES

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

ERRORS

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

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TNODIS]	No disconnect indication currently exists on the specified 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.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TSYSERR]	A system error has occurred during execution of this function.
[TOUTSTATE]	The function was issued in the wrong sequence on the transport endpoint referenced by <i>fd</i> .
[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 (<i>t_errno</i>).

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()*.

NAME

t_rcvrel – acknowledge receipt of an orderly release indication

SYNOPSIS

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

```
int t_rcvrel(fd)
```

```
int fd;
```

DESCRIPTION

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

This function is used to acknowledge receipt of an orderly release indication. 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 because such an attempt will block forever. 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()*.

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.
[TNOREL]	No orderly release indication currently exists on the specified transport endpoint.
[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.
[TSYSERR]	A system error has occurred during execution of this function.
[TOUTSTATE]	The function was issued in the wrong sequence on the transport endpoint referenced by <i>fd</i> .
[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 (<i>t_errno</i>).

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.

t_rcvrel()

XTI Library Functions and Parameters

SEE ALSO

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

NAME

t_rcvudata – receive a data unit

SYNOPSIS

```
#include <xti.h>

int t_rcvudata(fd, unitdata, flags)
int fd;
struct t_unitdata *unitdata;
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>	x	/
<i>unitdata->udata.len</i>	/	x
<i>unitdata->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.

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.

VALID STATES

T_IDLE

ERRORS

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

[TBADEF]	The specified file descriptor does not refer to a transport endpoint.
[TNODATA]	O_NONBLOCK was set, but no data units are currently available from the transport provider.
[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.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TOUTSTATE]	The function was issued in the wrong sequence on the transport endpoint referenced by <i>fd</i> .
[TSYSERR]	A system error has occurred during execution of this function.
[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 (<i>t_errno</i>).

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()*.

NAME

t_rcvuderr – receive a unit data error indication

SYNOPSIS

```
#include <xti.h>

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

DESCRIPTION

Parameters	Before call	After call
<i>fd</i>	x	/
<i>uderr->addr.maxlen</i>	x	/
<i>uderr->addr.len</i>	/	x
<i>uderr->addr.buf</i>	?	(?)
<i>uderr->opt.maxlen</i>	x	/
<i>uderr->opt.len</i>	/	x
<i>uderr->opt.buf</i>	?	(?)
<i>uderr->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;
long 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.

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:

[TBADEF]	The specified file descriptor does not refer to a transport endpoint.
[TNOUDERR]	No unit data error indication currently exists on the specified 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.
[TSYSERR]	A system error has occurred during execution of this function.
[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 (<i>t_errno</i>).

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()*.

NAME

t_snd – send data or expedited data over a connection

SYNOPSIS

```
#include <xti.h>

int t_snd(fd, buf, nbytes, flags)
int fd;
char *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, i.e., when the T_MORE flag is not set. Some transport providers also forbid zero-length TSDUs and ETSDUs. See **Appendix A, ISO Transport Protocol Information** for a fuller explanation.

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 accepted by the transport provider. Normally this will equal the number of bytes specified in *nbytes*. However, if *O_NONBLOCK* is set, 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 value of *nbytes*. If *nbytes* is zero and sending of zero octets is not supported by the underlying transport service, *t_snd()* will return -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] may be returned to inform the process that an event (e.g., a disconnect) has occurred.

VALID STATES

T_DATAXFER, T_INREL

ERRORS

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

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TBADDDATA]	Illegal amount of data: <ul style="list-style-type: none"> — 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 <i>info</i> 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, ISO Transport Protocol Information), or — 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 <i>info</i> argument - the ability of an XTI implementation to detect such an error case is implementation-dependent (see CAVEATS, below).
[TBADFLAG]	An invalid flag was specified.
[TFLOW]	<i>O_NONBLOCK</i> was set, but the flow control mechanism prevented the transport provider from accepting any data at this time.
[TNOTSUPPORT]	This function is not supported by the underlying transport provider.
[TLOOK]	An asynchronous event has occurred on this transport endpoint.
[TOUTSTATE]	The function was issued in the wrong sequence on the transport endpoint referenced by <i>fd</i> .

[TSYSERR]	A system error has occurred during execution of this function.
[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 (<i>t_errno</i>).

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 that in asynchronous mode, if the number of bytes accepted by the transport provider is less than the number of bytes requested, this may indicate that the transport provider is blocked due to flow control.

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].

t_snddis()

XTI Library Functions and Parameters

NAME

t_snddis – send user-initiated disconnect request

SYNOPSIS

```
#include <xti.h>

int t_snddis(fd, call)
int fd;
struct t_call *call;
```

DESCRIPTION

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

This function is used to initiate an abortive release on an already established connection, or to reject a connect 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 connect request, *call* must be non-null and contain a valid value of *sequence* to uniquely identify the rejected connect 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 disconnect 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)

ERRORS

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

[TBADF]	The specified file descriptor does not refer to a transport endpoint.
[TOUTSTATE]	The function was issued in the wrong sequence on the transport endpoint referenced by <i>fd</i> .
[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 connect request.
[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 (<i>t_errno</i>).

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 disconnect. 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).

t_sndrel()

XTI Library Functions and Parameters

NAME

t_sndrel – initiate an orderly release

SYNOPSIS

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

```
int t_sndrel(fd)
```

```
int fd;
```

DESCRIPTION

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

This function is used to initiate an orderly release of a transport connection 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. 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()*.

VALID STATES

T_DATAXFER, T_INREL

ERRORS

On 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 function was issued in the wrong sequence on the transport endpoint referenced by <i>fd</i> .
[TSYSERR]	A system error has occurred during execution of this function.
[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 (<i>t_errno</i>).

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()*.

t_sndudata()

XTI Library Functions and Parameters

NAME

t_sndudata – send a data unit

SYNOPSIS

```
#include <xti.h>

int t_sndudata(fd, unitdata)
int fd;
struct t_unitdata *unitdata;
```

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>	x	/
<i>unitdata->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 may use default options.

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 [TBADOPT]. These errors will alternatively be returned by *t_rcvuderr*. Therefore, an application must 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:

[TBADDDATA]	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.
[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 function was issued in the wrong sequence on the transport endpoint referenced by <i>fd</i> .
[TSYSERR]	A system error has occurred during execution of this function.
[TBADADDR]	The specified protocol address was in an incorrect format or contained illegal information.
[TBADOPT]	The specified options were in an incorrect format or contained illegal information.
[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 (<i>t_errno</i>).

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()*.

t_strerror()

XTI Library Functions and Parameters

NAME

t_strerror - produce an error message string

SYNOPSIS

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

```
char *t_strerror(errnum)
```

```
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 implementation-defined. If it is English, the error message string describing the value in *t_errno* is identical to 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()

NAME

t_sync – synchronise transport library

SYNOPSIS

#include <xti.h>

int t_sync(fd)

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. |
| [TSTATECHNG] | The transport endpoint is undergoing a state change. |
| [TSYSERR] | A system error has occurred during execution of this function. |

t_sync()

XTI Library Functions and Parameters

[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 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()*.

NAME

t_unbind – disable a transport endpoint

SYNOPSIS

```
#include <xti.h>

int t_unbind(fd)
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. |
| [TOUTSTATE] | The function was issued in the wrong sequence. |
| [TLOOK] | An asynchronous event has occurred on this transport endpoint. |
| [TSYSERR] | A system error has occurred during execution of this function. |
| [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 (<i>t_errno</i>). |

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().

ISO Transport Protocol Information

A.1 GENERAL

This appendix describes the protocol-specific information that is relevant for ISO transport providers.

Notes

- Protocol address

In an ISO environment, the protocol address is the transport address.

- Sending data of zero octets

The transport service definition, both in connection-oriented 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-oriented 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 (i.e., the call has requested sending a zero byte TSDU or ETSDU).

- Expedited data

In connection-oriented mode, and when the transport class permits it, the expedited data option must be negotiated during the connection establishment phase. In connectionless mode this feature is not supported.

A.2 OPTIONS

Options are formatted according to the structure **t_opthdr** as described in **Chapter 5, The Use of Options**. A transport provider compliant to this specification supports none, all or any subset of the options defined in **Section A.2.1, Connection-mode Service** and **Section A.2.2, Connectionless-mode Service**. An implementation may restrict the use of any of these options by offering them only in the privileged or read-only mode.

A.2.1 Connection-mode Service

The protocol level of all subsequent options is ISO_TP.

All options are association-related (see **Chapter 5, The Use of Options**). 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 (ISO 8072:1986)

These options are all defined in the ISO 8072:1986 transport service definition. The definitions are not repeated here.

Option Name	Type of Option Value	Legal Option Value	Meaning
TCO_THROUGHPUT	struct thrpt	octets per second	throughput
TCO_TRANSDEL	struct transdel	time in milliseconds	transit delay
TCO_RESERRORRATE	struct rate	OPT_RATIO	residual error rate
TCO_TRANSFFAILPROB	struct rate	OPT_RATIO	transfer failure probability
TCO_ESTFAILPROB	struct rate	OPT_RATIO	connection establ. failure probability
TCO_RELFAILPROB	struct rate	OPT_RATIO	connection release failure probability
TCO_ESTDELAY	struct rate	time in milliseconds	connection establ. delay
TCO_RELDELAY	struct rate	time in milliseconds	connection release delay
TCO_CONNRRESIL	struct rate	OPT_RATIO	connection resilience
TCO_PROTECTION	unsigned long	see text	protection
TCO_PRIORITY	unsigned long	see text	priority
TCO_EXPD	unsigned long	T_YES/T_NO	expedited data

Table A-1. Options for Quality of Service and Expedited Data (ISO 8072:1986)

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 (TCO_RESERRORRATE).

Absolute Requirements

For the options in **Table A-1**, 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 TCO_PROTECTION is an absolute requirement if the T_ABSREQ flag is set. The values specified for TCO_EXPD and 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 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 5, The Use of Options**.
- 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 5.3.4, Repending to a Negotiation Proposal**).
- The values of the QOS options TCO_THROUGHPUT, TCO_TRANSDEL, TCO_RESERRORRATE, TCO_TRANSFFAILPROB, TCO_ESTFAILPROB, TCO_RELFAILPROB, TCO_ESTDELAY, TCO_RELDELAY and 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 TCO_PROTECTION, TCO_PRIORITY and TCO_EXPD.

- TCO_THROUGHPUT and TCO_TRANSDEL
If *avgthrpt* (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 (*maxthrpt*). An analogous procedure applies to 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.
- 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.

- 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

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.

Avoid specifying both QOS parameters and management options at the same time.

Option Name	Type of Option Value	Legal Option Value	Meaning
TCO_LTPDU	unsigned long	length in octets	maximum length of TPDU
TCO_ACKTIME	unsigned long	time in milliseconds	acknowledge time
TCO_REASTIME	unsigned long	time in seconds	reassignment time
TCO_PREFCLASS	unsigned long	see text	preferred class
TCO_ALTCLASS1	unsigned long	see text	1st alternative class
TCO_ALTCLASS2	unsigned long	see text	2nd alternative class
TCO_ALTCLASS3	unsigned long	see text	3rd alternative class
TCO_ALTCLASS4	unsigned long	see text	4th alternative class
TCO_EXTFORM	unsigned long	T_YES/T_NO/T_UNSPEC	extended format
TCO_FLOWCTRL	unsigned long	T_YES/T_NO/T_UNSPEC	flowctr
TCO_CHECKSUM	unsigned long	T_YES/T_NO/T_UNSPEC	checksum
TCO_NETEXP	unsigned long	T_YES/T_NO/T_UNSPEC	network expedited data
TCO_NETRECPTCF	unsigned long	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 5, The Use of Options**).
- 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 5.3.4, Responding to a Negotiation Proposal**).
- 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**. 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, 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 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 are T_UNSPEC and all positive values.

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 **Table A-1**. 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 5, The Use of Options**). 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 TCO_EXP is set to T_YES and TCO_PREFCLASS is set to T_CLASS2, TCO_FLOWCTRL must also be set to T_YES.
- If TCO_PREFCLASS is set to T_CLASS0, TCO_EXP must be set to T_NO.
- The value in TCO_PREFCLASS must not be lower than the value in TCO_ALTCLASS1, 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, Connection-mode Service**).

All options are association-related (see **Chapter 5, The Use of Options**). They may be negotiated in all XTI states but T_UNINIT.

A.2.2.1 Options for Quality of Service (ISO 8072/Add.1:1986)

These options are all defined in the ISO 8072/Add.1:1986 transport service definition. The definitions are not repeated here.

Option Name	Type of Option Value	Legal Option Value	Meaning
TCL_TRANSDEL	struct rate	time in milliseconds	transit delay
TCL_RESERRORRATE	struct rate	OPT_RATIO	residual error rate
TCL_PROTECTION	unsigned long	see text	protection
TCL_PRIORITY	unsigned long	see text	priority

Table A-3. Options for Quality of Service (ISO 8072/Add.1:1986)**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 TCO_RESERRORRATE and TCL_RESERRORRATE, TCO_PRIORITY and TCL_PRIORITY, TCO_PROTECTION and TCL_PROTECTION (see **Section A.2.1.1, Options for Quality of Service and Expedited Data (ISO 8072:1986)**).
- TCL_TRANSDEL and TCO_TRANSDEL are different. 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 TCO_TRANSDEL. The range of legal option values for each field of **struct rate** is the same as that of 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
TCL_CHECKSUM	unsigned long	T_YES/T_NO	checksum computation

Table A-4. Management Option**Absolute Requirements**

A request for this option is an absolute requirement.

Further Remarks

TCL_CHECKSUM 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 have been bound to the same address as *fd* with the *qlen* parameter set to 0 when the *t_bind()* was called for that *resfd*.

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_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 the connection establishment, and
- the negotiation of expedited data transfer for this connection.

In T_DATAXFER, the *etsdu* field in the **t_info** structure is set to -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 table below.

Parameters	Before Call	After Call		
		Connection Class 0	Connection Class 1-4	Connectionless
<i>fd</i>	x	/	/	/
<i>info->addr</i>		x	x	x
<i>info->options</i>	/	x (1)	x (1)	x (1)
<i>info->tsdu</i>	/	x (2)	x (2)	0->63488
<i>info->etsdu</i>	/	-2	16/-2 (3)	-2
<i>info->connect</i>	/	-2	32	-2
<i>info->discon</i>	/	-2	64	-2
<i>info->servtype</i>	/	T_COTS	T_COTS	T_CLTS
<i>info->flags</i>	/	0	0	0

1. 'x' equals -2 or an integral number greater than zero.
2. 'x' equals -1 or an integral number greater than 0.
3. Depending 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 connect 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 connect 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 connect request in a finite time. So if the user queues the connect 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 table below gives the characteristics associated with the different classes.

Parameters	Before Call	After Call		
		Connection Class 0	Connection Class 1-4	Connectionless
<i>name</i>	x	/	/	/
<i>oflag</i>	x	/	/	/
<i>info->addr</i>	/	x	x	x
<i>info->options</i>	/	x (1)	x (1)	x (1)
<i>info->tsdu</i>	/	x (2)	x (2)	0->63488
<i>info->etsdu</i>	/	-2	16	-2
<i>info->connect</i>	/	-2	32	-2
<i>info->discon</i>	/	-2	64	-2
<i>info->servtype</i>	/	T_COTS	T_COTS	T_CLTS
<i>info->flags</i>	/	0	0	0

1. 'x' equals -2 or an integral number greater than zero.
2. 'x' equals -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.

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 transport service does not support the sending of expedited data.

Internet Protocol-specific Information

B.1 GENERAL

This appendix describes the protocol-specific information that is relevant for TCP and UDP transport providers.

Notes

- T_MORE flag and TSDUs

The notion of TSDU is not supported by a TCP transport provider, so the T_MORE flag will be ignored when TCP is used. The TCP PUSH flag cannot be used through the XTI interface because the TCP Military Standard (see **Referenced Documents**) states that:

“Successive pushes may not be preserved because two or more units of pushed data may be joined into a single pushed unit by either the sending or receiving TCP. Pushes are not visible to the receiving Upper Level Protocol and are not intended to serve as a record boundary marker”.

- 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. UDP has no urgent mechanism. See the TCP Military 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. However, its use is not recommended so that applications using TCP may be ported to use ISO Transport. The specification of TCP states that only established connections may be closed with orderly release, i.e., on an endpoint in T_DATAXFER or T_INREL state.

- Connection establishment

TCP does not allow the possibility of refusing a connection indication. Each connect 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.

B.2 OPTIONS

Options are formatted according to the structure **t_opthdr** as described in **Chapter 5, The Use of Options**. A transport provider compliant to this specification supports none, all or any subset of the options defined in **Section B.2.1, TCP-level Options** to **Section B.2.3, IP-level Options**. 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 INET_TCP. For this level, Table B-1 shows the options that are defined.

Option Name	Type of Option Value	Legal Option Value	Meaning
TCP_KEEPAIVE	struct t_kpalive	see text	check if connections are alive
TCP_MAXSEG	unsigned long	length in octets	get TCP maximum segment size
TCP_NODELAY	unsigned long	T_YES/T_NO	don't delay send to coalesce packets

Table B-1. TCP-level Options

These options are *not* association-related. They may be negotiated in all XTI states except T_UNBND and T_UNINIT. They are read-only in state T_UNBND. See **Chapter 5, The Use of Options** for the difference between options that are association-related and those that are not.

Absolute Requirements

A request for TCP_NODELAY and a request to activate TCP_KEEPAIVE is an absolute requirement. TCP_MAXSEG is a read-only option.

Further Remarks

TCP_KEEPAIVE 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:

“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 {
    long kp_onoff;    /* switch option on/off */
    long 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
T_YES T_GARBAGE	activate keep-alive timer and send garbage octet

Usually, an implementation should send a keep-alive packet with no data (T_GARBAGE not set). If T_GARBAGE is set, the keep-alive packet contains one garbage octet for compatibility with erroneous TCP implementations.

An implementation is, however, not obliged to support T_GARBAGE (see RFC 1122). Since the *kp_onoff* value is an absolute requirement, the request “T_YES | T_GARBAGE” may therefore be rejected.

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 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.

TCP_MAXSEG	This option is read-only. It is used to retrieve the maximum TCP segment size.
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 (e.g., MIT X Window System) that send a stream of mouse events which receive no replies, this packetisation may cause significant delays. TCP_NODELAY is used to defeat this algorithm. Legal option values are T_YES (“don’t delay”) and T_NO (“delay”).

B.2.2 UDP-level Options

The protocol level is INET_UDP. The option defined for this level is shown in Table B-2.

Option Name	Type of Option Value	Legal Option Value	Meaning
UDP_CHECKSUM	unsigned long	T_YES/T_NO	checksum computation

Table B-2. UDP-level Option

This option is association-related. It may be negotiated in all XTI states except T_UNBND and T_UNINIT. It is read-only in state T_UNBND. See **Chapter 5, The Use of Options** for the difference between options that are association-related and those that are not.

Absolute Requirements

A request for this option is an absolute requirement.

Further Remarks

UDP_CHECKSUM The option allows disabling/enabling of the UDP 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 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 IP-level Options

The protocol level is INET_IP. The options defined for this level are listed in Table B-3.

Option Name	Type of Option Value	Legal Option Value	Meaning
IP_BROADCAST	unsigned int	T_YES/T_NO	permit sending of broadcast messages
IP_DONTROUTE	unsigned int	T_YES/T_NO	just use interface addresses
IP_OPTIONS	array of unsigned characters	see text	IP per-packet options
IP_REUSEADDR	unsigned int	T_YES/T_NO	allow local address reuse
IP_TOS	unsigned char	see text	IP per-packet type of service
IP_TTL	unsigned char	time in seconds	IP per packet time-to-live

Table B-3. IP-level Options

IP_OPTIONS and IP_TOS are both association-related options. All other options are *not* association-related. See **Chapter 5, The Use of Options** for the difference between association-related options and options that are not.

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

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.
IP_DONTROUTE	This option indicates that outgoing messages should bypass the standard routing facilities. It is mainly used for testing and development.
IP_OPTIONS	<p>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 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.</p> <p>The option is disabled if it is specified with “no value”, i.e., with an option header only.</p> <p>The functions <i>t_connect()</i> (in synchronous mode), <i>t_listen()</i>, <i>t_rcvconnect()</i> and <i>t_rcvudata()</i> return the OPTIONS field, if any, of the received IP datagram associated with this call. The function <i>t_rcvuderr()</i> returns the OPTIONS field of the data unit previously sent that produced the error. The function <i>t_optmgmt()</i> with T_CURRENT set retrieves the currently effective IP_OPTIONS that is sent with outgoing datagrams.</p> <p>Common applications never need this option. It is mainly used for network debugging and control purposes.</p>
IP_REUSEADDR	Many TCP implementations do not allow the user to bind more than one transport endpoint to addresses with identical port numbers. If 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.

IP_TOS

This option is used to set (retrieve) the *type-of-service* field of an outgoing (incoming) 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_HIREL:

— 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 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_HIREL 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 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.

IP_TTL

This option 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 association-related option).

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. Also, *resfd* must be bound to the same address as *fd*. A potential restriction on binding of endpoints to protocol addresses is described under *t_bind()* below.
- If association-related options (IP_OPTIONS, IP_TOS) are to be sent with the connect 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, i.e., an address which specifically includes a port identifier.
- In the connection-oriented mode (i.e., TCP), the *t_bind()* function may only bind one transport endpoint to any particular protocol address. If that endpoint was bound in passive mode, i.e., *qlen* > 0, then other endpoints will be bound to the passive endpoint's protocol address via the *t_accept()* function only; that is, if *fd* refers to the passive endpoint and *resfd* refers to the new endpoint on which the connection is to be accepted, *resfd* will be bound to the same protocol address as *fd* after the successful completion of the *t_accept()* function.
- 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	
		TCP/IP	UDP/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>	/	-1	-2
<i>info->connect</i>	/	-2	-2
<i>info->discon</i>	/	-2	-2
<i>info->servtype</i>	/	T_COTS/T_COTS_ORD	T_CLTS
<i>info->flags</i>	/	T_SNDZERO	T_SNDZERO

'x' equals -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 remote calling socket.

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 TSDU varies among implementations.

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 4-5 in **Chapter 4, States and Events in XTI**, *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 must be bound to an appropriate protocol address and must be in the T_IDLE state.
 - The user transferring the current state of an endpoint must have correct permissions for the use of the protocol address bound to the accepting transport endpoint.

- The endpoint that transfers the state of the transport interface is placed into the T_IDLE state by the transport provider after the completion of the transfer if there are no more outstanding connect indications.

C.2 EXAMPLE IN CONNECTION-ORIENTED MODE

Figure C-1 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 connect request (via the return from *t_listen()*) and then would call the *t_accept()*.

The state diagram in Figure C-1 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 4-7 in **Chapter 4, States and Events in XTI**.

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_rcv()</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>

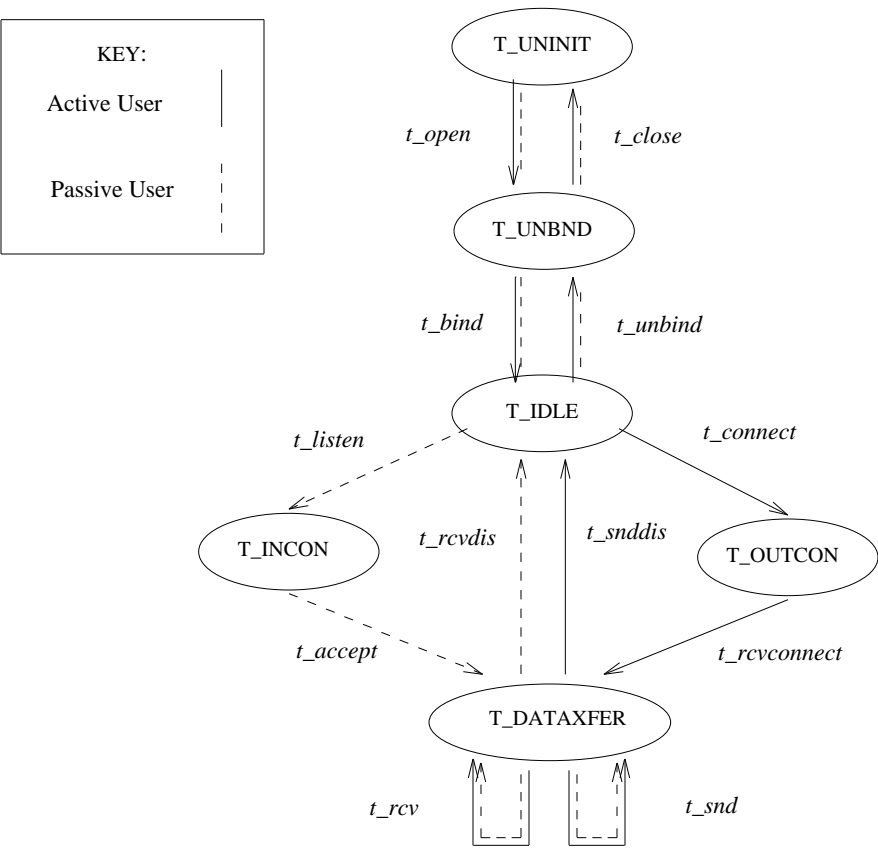


Figure C-1. Example of a Sequence of Transport Functions in Connection-oriented Mode

C.3 EXAMPLE IN CONNECTIONLESS MODE

Figure C-2 shows the allowable sequence of functions of user A and user B communicating using a connectionless 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 4-7 in Chapter 4, States and Events in XTI**.

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>




Figure C-2. Example of a 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 machine 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 (e.g., 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 connect establishment function. Optionally, the user can choose not to pass options as an argument on connect 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 optional orderly release facility of the connection-mode service (i.e., *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 (e.g., ISO and TCP). An application intended to run over different transport providers should avoid their use.

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 2.7, Synchronous and Asynchronous Execution Modes** and **Section 2.8, Event Management**.

C.5.1 Introduction

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.5.2 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 (e.g., *poll()* or *select()*). The design of more efficient and sophisticated applications, that make full use of all the XTI features, then becomes easily possible.

C.5.3 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	Disconnect 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 (e.g., [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.5.4 Guidelines for Use of System V poll()

poll() is defined in the System V Interface Definition, Third Edition as follows. Note that this definition may vary slightly in other systems.

NAME

poll - input/output multiplexing

SYNOPSIS

```
#include <poll.h>
int poll(struct pollfd fds[], unsigned long nfd, int timeout);
```

DESCRIPTION

poll() provides users with a mechanism for multiplexing input/output over a set of file descriptors. *poll()* identifies those file descriptors on which a user can read or write data, or on which certain events have occurred. A user can read data using *read()* and write data using *write()*. For STREAMS file descriptors, a user can also receive messages using *getmsg()* and *getpmsg()*, and send messages using *putmsg()* and *putpmsg()*.

fds specifies the file descriptors to be examined and the events of interest for each file descriptor. It is a pointer to an array with one element for each open file descriptor of interest. The array's elements are *pollfd* structures which contain the following members:

```
int fd;           /* file descriptor */
short events;     /* requested events */
short revents;    /* returned events */
```

where *fd* specifies an open file descriptor and *events* and *revents* are bit-masks constructed by OR'ing a combination of the following event flags:

POLLIN	Data other than high-priority data may be read without blocking. For STREAMS, this flag is set even if the message is of zero length.
POLLRDNORM	Normal data (priority band equals 0) may be read without blocking. For STREAMS, this flag is set even if the message is of zero length.
POLLRDBAND	Data from a non-zero priority band may be read without blocking. For STREAMS, this flag is set even if the message is of zero length.
POLLPRI	High-priority data may be received without blocking. For STREAMS, this flag is set even if the message is of zero length.
POLLOUT	Normal data may be written without blocking.
POLLWRBAND	Priority data (priority band greater than 0) may be written.
POLLER	An error has occurred on the device or STREAM. This flag is only valid in the <i>revents</i> bitmask; it is not used in the <i>events</i> field.

POLLUP	The device has been disconnected. This event and POLLOUT are mutually exclusive; a STREAM can never be writable if a hangup has occurred. However, this event and POLLIN, POLLRDNORM, POLLRDBAND or POLLPRI are not mutually exclusive. This flag is only valid in the <i>revents</i> bitmask; it is not used in the <i>events</i> field.
POLLNVAL	The specified <i>fd</i> value is invalid. This flag is only valid in the <i>revents</i> field; it is not used in the <i>events</i> field.

For each element of the array pointed to by *fds*, *poll()* examines the given file descriptor for the event(s) specified in *events*. The number of file descriptors to be examined is specified by *nfds*.

If the value of *fd* is less than zero, *events* is ignored and *revents* is set to zero in that entry on return from *poll()*.

The results of the *poll()* query are stored in the *revents* field in the *pollfd* structure. Bits are set in the *revents* bitmask to indicate which of the requested events are true. If none of the requested events are true, none of the specified bits is set in *revents* when the *poll()* call returns. The events flags POLLUP, POLLERR and POLLNVAL, are always set in the *revents* if the conditions they indicate are true; this occurs even though these flags were not present in *events*.

If none of the defined events have occurred on any selected file descriptor, *poll()* waits at least *timeout* milliseconds for an event to occur on any of the selected file descriptors. On a computer where millisecond timing accuracy is not available, *timeout* is rounded up to the nearest legal value available on that system. If the value of *timeout* is 0, *poll()* returns immediately. If the value of *timeout* is -1 *poll()* blocks until a requested event occurs or until the call is interrupted. *poll()* is not affected by the O_NDELAY and O_NONBLOCK flags.

RETURN VALUES

Upon successful completion, the function *poll()* returns a non-negative value. A positive value indicates the total number of file descriptors that have been selected (i.e., file descriptors for which the *revents* field is non-zero). A value of 0 indicates that the call timed out and no file descriptors have been selected. Upon failure, the function *poll()* returns a value -1 and sets *errno* to indicate an error.

ERRORS

Under the following conditions, the function *poll()* fails and sets *errno* to:

EAGAIN	If the allocation of internal data structures failed but the request should be attempted again.
EINTR	If a signal was caught during the <i>poll()</i> system call.
EINVAL	If the argument <i>nfds</i> is less than zero or greater than {OPEN_MAX}.

For an application to be notified of any XTI events on each of its active endpoints, the array pointed to by *fds* should contain as many elements as active endpoints identified by the file descriptor *fd*, and the *events* member of those elements should be set to the combination of event flags as specified below:

- **For Class 1 events:**

POLLIN | POLLPRI (for System V Release 3)

or

POLLIN | POLLRDNORM | POLLRDBAND | POLLPRI (for System V Release 4)

- **For Class 2 events:**

POLLOUT (for System V Release 3)

or

POLLOUT | POLLWRBAND (for System V Release 4)

In a System V Release 3, the meaning of POLLOUT may differ for different XTI implementations. It could either mean:

- that both normal and expedited data may be sent

or

- that normal data may be sent and the flow of expedited data cannot be monitored via *poll()*.

A truly portable XTI application should, therefore, not assume that the flow of expedited data is monitored by *poll()*. This is not a serious restriction, since an application usually only sends small amounts of expedited data and flow restrictions are not a major problem.

In a System V Release 4, the meaning of POLLOUT and POLLWRBAND is intended to be the same for all XTI implementations.

POLLOUT Normal data may be sent.

POLLWRBAND Expedited data may be sent.

Hereafter we describe the outline of an XTI server program making use of the System V *poll()*.

```

/*
 * 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.
 * The number of poll flags has increased from System V Release 3 to System V
 * Release 4. Hence, if this program is to be used in a System V Release 3,
 * the constant SVR3 must be defined during compile time.
 *
 * A transport endpoint is opened in asynchronous mode over a message-oriented
 * transport provider (e.g., 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 connect indications received, a new endpoint is opened with qlen = 0
 * and the connect 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 disconnect 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 <stropts.h>
#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.
 */
#ifdef SVR3
#define ERREVENTS      (POLLERR | POLLHUP | POLLNVAL)
#define INEVENTS       (POLLIN | POLLPRI)
#define OUTEVENTS      POLLOUT
#else
#define ERREVENTS      (POLLERR | POLLHUP | POLLNVAL)
#define INEVENTS       (POLLIN | POLLRDNORM | POLLRDBAND | POLLPRI)
#define OUTEVENTS      (POLLOUT | POLLWRBAND)
#endif
#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 connect 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 connect 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_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 (MAXSIZEag);
/*
 * 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 disconnect 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] Disconnect 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 disconnect 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.5.5 Guidelines for Use of BSD select()

select() is defined in the 4.3 Berkeley Software Distribution as follows. Note that this definition may vary slightly in other systems.

NAME

select - synchronous I/O multiplexing

SYNOPSIS

```
#include <sys/types.h>
#include <sys/time.h>

nfound = select(nfds, readfds, writefds, exceptfds, timeout)
int nfound, nfds;
fd_set *readfds, *writefds, *exceptfds;
struct timeval *timeout;

FD_SET(fd, &fdset)
FD_CLR(fd, &fdset)
FD_ISSET(fd, &fdset)
FD_ZERO(&fdset)
int fd;
fd_set fdset;
```

DESCRIPTION

select() examines the I/O descriptor sets whose addresses are passed in *readfds*, *writefds* and *exceptfds* to see if some of their descriptors are ready for reading, ready for writing, or have an exceptional condition pending, respectively. The first *nfds* descriptors are checked in each set; i.e., the descriptors from 0 through *nfds* - 1 in the descriptor sets are examined. On return, *select()* replaces the given descriptor sets with subsets consisting of those descriptors that are ready for the requested operation. The total number of ready descriptors in all the sets is returned in *nfound*.

The descriptor sets are stored as bit fields in arrays of integers. The following macros are provided for manipulating such descriptor sets: *FD_ZERO(&fdset)* initialises a descriptor set *fdset* to the null set. *FD_SET(fd, &fdset)* includes a particular descriptor *fd* in *fdset*. *FD_CLR(fd, &fdset)* removes *fd* from *fdset*. *FD_ISSET(fd, &fdset)* is non-zero if *fd* is a member of *fdset*, zero otherwise. The behaviour of these macros is undefined if a descriptor value is less than zero or greater than or equal to *FD_SETSIZE*, which is normally at least equal to the maximum number of descriptors supported by the system.

If *timeout* is a non-zero pointer, it specifies a maximum interval to wait for the selection to complete. If *timeout* is a zero pointer, the select blocks indefinitely. To affect a poll, the *timeout* argument should be non-zero, pointing to a zero-valued *timeval* structure.

Any of *readfds*, *writefds* and *exceptfds* may be given as zero pointers if no descriptors are of interest.

RETURN VALUES

select() returns the number of ready descriptors that are contained in the descriptor sets, or -1 if an error occurred. If the time limit expires then *select()* returns 0. If *select()* returns with an error, including one due to an interrupted call, the descriptor sets will be unmodified.

ERRORS

An error return from *select()* indicates:

- | | |
|----------|--|
| [EBADF] | One of the descriptor sets specified an invalid descriptor. |
| [EINTR] | A signal was delivered before the time limit expired and before any of the selected events occurred. |
| [EINVAL] | The specified time limit is invalid. One of its components is negative or too large. |

Many systems provide the macros *FD_SET*, *FD_CLR*, *FD_ISSET* and *FD_ZERO* in `<sys/types.h>` or other header files to manipulate these bit masks. If not available they should be defined by the user (see the program example below).

For an application to be notified of any XTI events on each of its active endpoints identified by a file descriptor *fd*, this file descriptor *fd* should be included in the appropriate descriptor sets *readfds*, *exceptfds* or *writefds* as specified below:

- **For Class 1 events:**

Set the bit masks *readfds* and *exceptfds* by *FD_SET(fd, readfds)* and *FD_SET(fd, exceptfds)*.

- **For Class 2 events:**

Set the bit mask *writefds* by *FD_SET(fd, writefds)*.

If, on return of *select()*, the bit corresponding to *fd* is set in *writefds*, this can have a different meaning for different XTI implementations. It could either mean:

- that both normal and expedited data may be sent, or
- that normal data may be sent and the flow of expedited data cannot be monitored via *select()*.

A truly portable XTI application should, therefore, not assume that the flow of expedited data is monitored by *select()*. This is not a serious restriction, since an application usually only sends small amounts of expedited data and flow restrictions are not a major problem.

Hereafter we describe the outline of an XTI server program making use of the BSD *select()*.

```

/*
 * 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 (e.g., 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 connect indications received, a new endpoint is opened with qlen = 0
 * and the connect 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 disconnect 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(). Some UNIX derivatives use other includes,
 * e.g., <sys/times.h> instead of <sys/time.h>.
 * <sys/select.h> instead of <sys/types.h>.
 */
#include <sys/types.h>
#include <time.h>

```

```

/*
 * Includes that are only relevant, if the type fd_set and the macros FD_SET,
 * FD_CLR, FD_ISSET and FD_ZERO have to be explicitly defined in this program.
 */
#include <limits.h>
#include <string.h> /* for memset() */

/*
 * 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 */
/*
 * Select uses bit masks of file descriptors in longs. Most systems provide a
 * type "fd_set" and macros in <sys/types.h> or <sys/select.h> to ease the use
 * of select().
 * They are explicitly defined below in case that they are not defined in
 * <sys/types.h> or <sys/select.h>.
 */
/*
 * OPEN_MAX should be >= number of maximum open files per process
 */
#ifndef OPEN_MAX
#define OPEN_MAX      256
#endif
#ifndef NFDBITS
#define NFDBITS (sizeof(long) * CHAR_BIT) /* bits per mask */
#endif
#ifndef howmany
#define howmany(x, y) (((x)+((y)-1))/(y))
#endif
#ifndef FD_SET
typedef struct fd_set {
    long  fds_bits[howmany(OPEN_MAX, NFDBITS)];
} fd_set;
#define FD_SET(n, p) ((p)->fds_bits[(n)/NFDBITS] |= (1 << ((n) % NFDBITS)))
#define FD_CLR(n, p) ((p)->fds_bits[(n)/NFDBITS] &= ~(1 << ((n) % NFDBITS)))
#define FD_ISSET(n, p) ((p)->fds_bits[(n)/NFDBITS] & (1 << ((n) % NFDBITS)))
#define FD_ZERO(p)    memset((p), (u_char) 0, sizeof(*(p)))
#endif /* FD_SET */

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_revdis() */
    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 rfdds, wfdds, xfdds;       /* initial values of file descriptor */
    /* sets rfds, wfds and xfds */

```

```

/*
 * 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 rfd, xfd and wfd to their initial values
     */
    rfd = rfd;
    xfd = xfd;
    wfd = wfd;
    if (select(OPEN_MAX, &rfd, &wfd, &xfd,
              (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, &xfd)) {
            fprintf(stderr, "[%d] Unexpected select events\n", fd);
            continue;
        }
        if (FD_ISSET(fd, &wfd))
            continue; /* no event for this endpoint */
    }
}

```

```

/*
 * Check for events
 */
switch((event = t_look(fd))) {
case T_LISTEN:
/*
 * Must be a connect 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 connect 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 rfdds and xfdds (to be portable
 * through all XTI implementations).
 */
fds[act] = fdd;
FD_SET(fdd, &rfdds);
FD_SET(fdd, &xfdds);
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 (MAXSIZEag);
/*
 * 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 disconnect 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] Disconnect 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, &rfdsets);
            FD_CLR(fd, &xfds);
            FD_CLR(fd, &wfds);
            /* Move last entry of fds array to current slot,
             * adjust internal counters and flags
             */
            fds[i] = fds[act];
            discflag = 0; /* clear disconnect 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 **Referenced Documents**) provides mappings of NetBIOS services to OSI and IPS transport protocols.¹

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”.

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 (i.e., 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, i.e., the X/Open CAE.

1. 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 **X/Open Developers' Specification, Protocols for X/Open PC Interworking: SMB**. 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 or group names, and must be identified as such in certain circumstances.

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. They should be included in `<xti.h>`.

```
#define NB_UNIQUE      0
#define NB_GROUP       1
#define NB_NAMELEN     16
#define NB_BCAST_NAME  "*"      /* asterisk plus 15 spaces */
```

The protocol addresses passed in calls to `t_bind()`, `t_connect()`, etc., are structured as follows:



Type The first octet specifies the type of the NetBIOS name. It may be set to NB_UNIQUE or NB_GROUP.

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 (i.e., **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 (e.g., `t_bind()`, `t_connect()`, `t_rcvdata()`, 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 (e.g., 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 (*t_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 disconnect 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, XTI Functions**).

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_DISIN 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()** in **Chapter 6, Library Functions and Parameters**).

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-oriented transport endpoints. The *t_accept()* function will fail if a user attempts to accept a connection request on a connectionless 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, NetBIOS Names and Addresses**. 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 *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 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 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, NetBIOS Names and Addresses**. 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-oriented transport endpoints. The `t_connect()` function will fail if a user attempts to initiate a connection on a connectionless 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, NetBIOS Names and Addresses**).

`t_error()`

No special considerations for NetBIOS transport providers.

`t_free()`

No special considerations for NetBIOS transport providers.

`t_getinfo()`

The values of the parameters in the `t_info` structure will reflect NetBIOS transport limitations, as follows:

<code>addr</code>	<code>sizeof()</code> the NetBIOS protocol address, as defined in Section D.5, NetBIOS Names and Addresses .
<code>options</code>	Equals -2, 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-oriented endpoint it is a positive value, not larger than 131070. If the <code>fd</code> is associated with a connectionless endpoint it is a positive value not larger than 65535 ² .
<code>etsdu</code>	Equals -2, indicating expedited data is not supported.
<code>connect</code>	Equals -2, indicating data cannot be transferred during connection establishment.
<code>discon</code>	Equals -2, 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-oriented endpoint, or <code>T_CLTS</code> if associated with a connectionless endpoint. Optionally, may be set to <code>T_COTS_ORD</code> if the <code>fd</code> is associated with a connection-oriented endpoint and the transport provider supports the use of <code>t_sndrel()</code> / <code>t_rcvrel()</code> as described in Section D.6, NetBIOS Connection Release .
<code>flags</code>	Equals <code>T_SNDZERO</code> , indicating that zero TSDUs may be sent.

2. For the mappings to OSI and IPS protocols, the value cannot exceed 512 or 1064 respectively.

t_getprotaddr()

The NetBIOS name and name type of the transport endpoint referred to by the *fd* are passed in the *boundaddr* parameter (*boundaddr->addr.buf*), as described in **Section D.5, NetBIOS Names and Addresses**; 0 is returned in *boundaddr->addr.len* if the transport endpoint is in the T_UNBND state. The NetBIOS name currently connected to *fd*, if any, is passed in the *peeraddr* parameter (*peeraddr->addr.buf*); the value 0 is returned in *peeraddr->addr.len* if the transport endpoint is not in the T_DATAXFER state.

t_getstate()

No special considerations for NetBIOS transport providers.

t_listen()

On return, the *call* parameter provides the NetBIOS name of the calling transport user (that issued the connection request), as described in **Section D.5, NetBIOS Names and Addresses**.

No user data may be transferred during connection establishment (*call->udata.len*=0 on return).

This function may only be used with connection-oriented transport endpoints. The *t_listen()* function will fail if a user attempts to “listen” on a connectionless 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_listen()* function is no longer in the CAE system name table, as may occur as a result of the NetBIOS name conflict resolution process (e.g., 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()**.

t_optmgmt()

No special considerations for NetBIOS transport providers.

t_rcv()

This function may only be used with connection-oriented transport endpoints. The *t_rcv()* function will fail if a user attempts a receive on a connectionless 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, NetBIOS Names and Addresses**.

No user data may be returned to the caller (*call->udata.len*=0 on return).

This function may only be used with connection-oriented transport endpoints. The *t_rcvconnect()* function will fail if a user attempts to establish a connection on a connectionless endpoint and *t_errno* will be set to [TNOTSUPPORT].

t_rcvdis()

The following disconnect reason codes are valid for any implementation of a NetBIOS provider under XTI:

```
#define NB_ABORT      0x18    /* session ended abnormally */
#define NB_CLOSED     0x0A    /* session closed */
#define NB_NOANSWER    0x14    /* no answer (cannot find name called */
#define NB_OPREJ      0x12    /* session open rejected */
```

These definitions should be included in **<xti.h>**.

t_rcvrel()

As described in **Section D.6, NetBIOS Connection Release**, 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, NetBIOS Names and Addresses**.

The *fd* associated with the *t_rcvudata()* function must refer to a connectionless transport endpoint. The function will fail if a user attempts to receive on a connection-oriented 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 (e.g., TOP/NetBIOS NameConflictAdvise indication).

To receive a broadcast datagram, the endpoint must be bound to the NetBIOS name NB_BCAST_NAME.

t_revuderr()

If attempted on a connectionless 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-oriented 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-oriented transport endpoints. The *t_snd()* function will fail if a user attempts a send on a connectionless 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()** in **Chapter 6, XTI Library Functions and Parameters** 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, NetBIOS Connection Release** for further details.

No user data may be sent in the disconnect request (call->udata.len=0).

This function may only be used with connection-oriented transport endpoints. The *t_snddis()* function will fail if a user attempts a disconnect request on a connectionless 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 T_DISIN and an appropriate disconnect reason code. See **Section D.6, NetBIOS Connection Release** 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, NetBIOS Names and Addresses**.

The *fd* associated with the *t_sndudata()* function must refer to a connectionless transport endpoint. The function will fail if a user attempts this function on a connection-oriented 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 (e.g., TOP/NetBIOS NameConflictAdvise indication).

To send a broadcast datagram, the NetBIOS name in the NetBIOS address structure provided in *unitdata->addr.buf* must be 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 which is not bound to a local address.
- XTI provides the additional utility functions *t_strerror()* and *t_getprotaddr()*.

Headers and Definitions

F.1 THE <xti.h> HEADER

```

/*
 * 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 TYSERR        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 /* connect confirmation received */
#define T_DATA        0x0004 /* normal data received */
#define T_EXDATA      0x0008 /* expedited data received */
#define T_DISCONNECT  0x0010 /* disconnect 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_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.
 */

extern int t_errno;
/* XTI LIBRARY FUNCTIONS */
/* XTI Library Function: t_accept – accept a connect request*/
extern int t_accept();
/* XTI Library Function: t_alloc – allocate a library structure*/
extern char *t_alloc();
/* XTI Library Function: t_bind – bind an address to a transport endpoint*/
extern int t_bind();
/* XTI Library Function: t_close – close a transport endpoint*/
extern int t_close();
/* XTI Library Function: t_connect – establish a connection */
extern int t_connect();
/* XTI Library Function: t_error – produce error message*/
extern int t_error();
/* XTI Library Function: t_free – free a library structure*/

```

```
extern int t_free();
/* XTI Library Function: t_getprotaddr – get protocol addresses*/
extern int t_getprotaddr();
/* XTI Library Function: t_getinfo – get protocol-specific service information*/
extern int t_getinfo();
/* XTI Library Function: t_getstate – get the current state*/
extern int t_getstate();
/* XTI Library Function: t_listen – listen for a connect indication*/
extern int t_listen();
/* XTI Library Function: t_look – look at current event on a transport endpoint*/
extern int t_look();
/* XTI Library Function: t_open – establish a transport endpoint*/
extern int t_open();
/* XTI Library Function: t_optmgmt – manage options for a transport endpoint*/
extern int t_optmgmt();
/* XTI Library Function: t_rcv – receive data or expedited data on a connection*/
extern int t_rcv();
/* XTI Library Function: t_rcvdis – retrieve information from disconnect*/
extern int t_rcvdis();
/* XTI Library Function: t_rcvrel – acknowledge receipt of */
/* an orderly release indication */
extern int t_rcvrel();
/* XTI Library Function: t_rcvudata – receive a data unit*/
extern int t_rcvudata();
/* XTI Library Function: t_rcvuderr – receive a unit data error indication*/
extern int t_rcvuderr();
/* XTI Library Function: t_snd – send data or expedited data over a connection*/
extern int t_snd();
/* XTI Library Function: t_snddis – send user-initiated disconnect request*/
extern int t_snddis();
/* XTI Library Function: t_sndrel – initiate an orderly release*/
extern int t_sndrel();
/* XTI Library Function: t_sndudata – send a data unit*/
extern int t_sndudata();
/* XTI Library Function: t_strerror – generate error message string */
extern char *t_strerror();
/* XTI Library Function: t_sync – synchronise transport library*/
extern int t_sync();
/* XTI Library Function: t_unbind – disable a transport endpoint*/
extern int t_unbind();
```

```

/*
 * Protocol-specific service limits.
 */

struct t_info {
    long    addr;        /* max size of the transport protocol address */
    long    options;     /* max number of bytes of protocol-specific options */
    long    tsdu;        /* max size of a transport service data unit */
    long    etsdu;       /* max size of expedited transport service data unit */
    long    connect;     /* max amount of data allowed on connection */
                    /* establishment functions */
    long    discon;      /* max data allowed on t_snddis and t_rcvdis functions */
    long    servtype;    /* service type supported by transport provider */
    long    flags;       /* other info about the transport provider */
};

/*
 * Service type defines.
 */

#define T_COTS        01    /* connection-oriented transport service */
#define T_COTS_ORD    02    /* connection-oriented with orderly release */
#define T_CLTS        03    /* connectionless transport service */

/*
 * Flags defines (other info about the transport provider).
 */

#define T_SENDDIS 0x001    /* supports 0-length TSDUs */

/*
 * netbuf structure.
 */

struct netbuf {
    unsigned int    maxlen;
    unsigned int    len;
    char            *buf;
};

```



```

/*
 * t_opthdr structure
 */

struct t_opthdr {
    unsigned long len;          /* total length of option; i.e.
                                sizeof (struct t_opthdr) + length of
                                option value in bytes */

    unsigned long level;        /* protocol affected */
    unsigned long name;         /* option name */
    unsigned long status;       /* status value */
    /* followed by the option value */
};

/*
 * t_bind – format of the address and options arguments of bind.
 */

struct t_bind {
    struct netbuf    addr;
    unsigned         qlen;
};

/*
 * Options management structure.
 */

struct t_optmgmt {
    struct netbuf    opt;
    long             flags;
};

/*
 * Disconnect 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 */
    long          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().
 */

#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_UNUSED     -1
#define T_NULL        0
#define T_ABSREQ      0x8000
#define T_INFINITE    -1
#define T_INVALID     -2

/* T_INFINITE and T_INVALID are values of t_info */

/*
 * General definitions for option management
 */

#define T_UNSPEC      (~0 -- 2)    /* applicable to u_long, long, char ... */
#define T_ALLOPT      0
#define T_ALIGN(p)    (((unsigned long) (p) + (sizeof(long) - 1)) \
                        &~ (sizeof(long) -- 1))

#define OPT_NEXTHDR(pbuf, buflen, popt) \
    (((char *) (popt) + T_ALIGN((popt)->len) < \
    pbuf + buflen) ? \
    (struct t_opthdr *) ((char *) (popt) + T_ALIGN((popt)->len)) : \
    (struct t_opthdr *) NULL)

/* OPTIONS ON XTI LEVEL */

/* XTI-level */

#define XTI_GENERIC    0xffff

```

```

/*
 * XTI-level Options
 */

#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 {
    long    l_onoff;  /* option on/off */
    long    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
#define T_PASSIVEPROTECT 2
#define T_ACTIVEPROTECT 4

```

```

/*
 * Default value for the length of TPDUs.
 */

#define    T_LTPDUFLT    128    /* define obsolete in XPG4 */

/*
 * rate structure.
 */

struct rate {
    long targetvalue;    /* target value */
    long 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
 */

#define    ISO_TP    0x0100

```

```

/*
 * Options for Quality of Service and Expedited Data (ISO 8072:1986)
 */

#define TCO_THROUGHPUT      0x0001
#define TCO_TRANSDEL        0x0002
#define TCO_RESERRORRATE    0x0003
#define TCO_TRANSFFAILPROB  0x0004
#define TCO_ESTFAILPROB     0x0005
#define TCO_RELFAILPROB     0x0006
#define TCO_ESTDELAY        0x0007
#define TCO_RELDELAY        0x0008
#define TCO_CONNRRESIL      0x0009
#define TCO_PROTECTION      0x000a
#define TCO_PRIORITY        0x000b
#define TCO_EXPD            0x000c

#define TCL_TRANSDEL        0x000d
#define TCL_RESERRORRATE    TCO_RESERRORRATE
#define TCL_PROTECTION      TCO_PROTECTION
#define TCL_PRIORITY        TCO_PRIORITY

/*
 * Management Options
 */

#define TCO_LTPDU           0x0100
#define TCO_ACKTIME         0x0200
#define TCO_REASTIME        0x0300
#define TCO_EXTFORM         0x0400
#define TCO_FLOWCTRL        0x0500
#define TCO_CHECKSUM        0x0600
#define TCO_NETEXP          0x0700
#define TCO_NETRECPTCF      0x0800
#define TCO_PREFCLASS       0x0900
#define TCO_ALTCLASS1       0x0a00
#define TCO_ALTCLASS2       0x0b00
#define TCO_ALTCLASS3       0x0c00
#define TCO_ALTCLASS4       0x0d00

#define TCL_CHECKSUM        TCO_CHECKSUM

/* INTERNET SPECIFIC ENVIRONMENT */

/*
 * TCP level
 */

#define INET_TCP            0x6

```

```

/*
 *TCP-level Options
 */

#define    TCP_NODELAY    0x1    /* don't delay packets to coalesce */
#define    TCP_MAXSEG    0x2    /* get maximum segment size */
#define    TCP_KEEPAIVE    0x8    /* check, if connections are alive */

/*
 * Structure used with TCP_KEEPAIVE option.
 */

struct t_kpalive {
    long    kp_onoff;    /* option on/off */
    long    kp_timeout;    /* timeout in minutes */
};

#define    T_GARBAGE    0x02

/*
 * UDP level
 */

#define    INET_UDP    0x11

/*
 * UDP-level Options
 */

#define    UDP_CHECKSUM    TCO_CHECKSUM    /* checksum computation */

/*
 * IP level
 */

#define    INET_IP    0x0

/*
 * IP-level Options
 */

#define    IP_OPTIONS    0x1    /* IP per-packet options */
#define    IP_TOS    0x2    /* IP per-packet type of service */
#define    IP_TTL    0x3    /* IP per-packet time to live /
#define    IP_REUSEADDR    0x4    /* allow local address reuse */
#define    IP_DONTROUTE    0x10    /* just use interface addresses */
#define    IP_BROADCAST    0x20    /* permit sending of broadcast msgs */

```

```
/*
 * 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 SET_TOS(prec, tos) ((0x7 & (prec)) << 5 | (0x1c & (tos)))
```


Appendix G

Abbreviations

CO	Connection-oriented
CL	Connectionless
EM	Event Management
ETSDU	Extended 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

Abbreviations



Appendix H

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).

Initiator

An entity that initiates a connect request.

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.

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 `t_open()` 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

abortive release: 18, 100, 197
accept: 119
accept1: 27, 31
accept2: 27, 31
accept3: 27, 31
address: 3-4, 13-15, 20, 23, 49-51, 54-56, 59-61, 65, 68, 71, 75-76, 93-94, 104-105, 111
addresses:
 multiple use: 8
application: 3-5, 136, 182
applications: 47
 portability: 136, 181
applications portability: 3, 182
association-related: 116
association-related options: 33
asynchronous: 32, 73
asynchronous events: 6
asynchronous mode: 5, 87, 197
bind: 3, 22, 27, 30, 50, 54, 56, 109
buffer: 52-53, 63, 85
Call structure: 187
caller: 23, 49-50, 59, 71
character string: 3, 198
checksum check: 126
child process: 4
CL: 195
close: 22, 30-31, 57
closed: 27, 30-31
CO: 195
compatibility:
 future: 53
connect indication: 29, 49-50, 54-55, 71, 89
connect request: 49, 71, 87, 100
Connect semantics: 9
connect1: 27, 31
connect2: 27, 31
connection: 3, 49, 59, 85, 87, 100, 102
connection establishment: 13, 15-16, 59, 87, 123, 131, 197
connection mode: 4, 13, 15, 22, 31, 112, 129, 131-132, 197
connection mode service: 31
connection release: 13, 18-19, 131, 197
connection-oriented: 22, 131-132
connection-oriented mode: 13, 129
connectionless: 22, 121, 135
connectionless mode: 4, 20, 22, 30, 93, 95, 104, 116, 135, 197
constants: 47
create:
 transport endpoint: 14
current event: 0, 30, 73
current state: 0, 4, 30, 70, 107
data: 85, 89, 93, 97, 100, 104
data transfer: 13, 16, 20, 30-31, 131, 197
data unit: 18, 20, 93, 104
 discarded: 21
datagram: 4, 20, 197
datagram structure: 188
de-initialisation: 13-15, 20, 30
default: 75, 136
definitions: 11
Delete optional functions: 8
device: 136
device driver: 182
discarded data unit: 21
disconnect: 7, 13, 32, 89
 indication: 27
 request: 100
disconnect structure: 187
dup: 3, 107-108
duplex: 16
EINVAL: 53
EM: 137, 195
enqueue: 15, 23
errmsg: **61**
errno: 61
errnum: 106
error: 184
error code: 95, 106
error codes: 183
 TACCES: 183
 TADDRBUSY: 183
 TBADADDR: 183

TBADDDATA: 183
 TBADF: 183
 TBADFLAG: 183
 TBADNAME: 183
 TBADOPT: 183
 TBADQLEN: 183
 TBADSEQ: 183
 TBUFOVFLW: 183
 TFLOW: 183
 TINDOUT: 183
 TLOOK: 183
 TNOADDR: 183
 TNODATA: 183
 TNODIS: 183
 TNOREL: 183
 TNOSTRUCTYPE: 183
 TNOTSUPPORT: 183
 TNOUDERR: 183
 TOUTSTATE: 183
 TPROTO: 183
 TPROVMISMATCH: 183
 TQFULL : 183
 TRESADDR: 183
 TRESQLEN: 183
 TSTATECHNG: 183
 TSYSERR: 183
 error handling: 4
 error indication: 95
 error message: 8, 61, 106
 error number: 5, 106
 established:
 connection: 132
 ETSDU: 17, 75-76, 85, 97-99, 195
 event: 0, 30, 135
 current: 0, 30, 73
 event management: 7-8
 Event Management: 137
 events: 138, 184
 accept1: 27, 31
 accept2: 27, 31
 accept3: 27, 31
 bind: 27, 30
 closed: 27, 30-31
 connect1: 27, 31
 connect2: 27, 31
 incoming: 28
 listen: 28, 31, 71
 opened: 27, 30
 optmgmt: 27, 30
 outgoing: 27
 pass_conn: 28, 31
 rcv: 28, 31, 85
 rcvconnect: 28, 31
 rcvdis1: 28, 31
 rcvdis2: 28, 31
 rcvdis3: 28, 31
 rcvrel: 28, 31
 rcvudata: 28, 30
 rcvuderr: 28, 30
 snd: 27, 31
 snddis1: 27, 31
 snddis2: 27, 31
 sndrel: 27, 31
 sndudata: 27, 30
 T_CONNECT: 184
 T_DATA: 184
 T_DISCONNECT: 184
 T_EXDATA: 184
 T_GODATA: 184
 T_GOEXDATA: 184
 T_LISTEN: 184
 T_ORDREL: 184
 T_UDERR: 184
 unbind: 27, 30
 events and t_look: 6-7
 example: 142, 156
 exceptfds: 155
 exec: 107-108
 execution mode: 5, 17
 Expedited data: 9
 expedited data: 13, 17-18, 76, 85, 97, 111-112, 121, 123, 141, 155, 192, 197
 expedited transport service data unit (ETSDU): 76, 85, 197
 fcntl: 3, 6, 59, 71-72, 85-86, 88, 93-94, 98, 104-105
 fcntl.h: 75
 fd: 3, 27
 features: 22-23
 file descriptor: 3, 57, 65, 75, 107, 141, 155
 file.c: 47
 flag: 75, 78, 85
 flags: 75, 78, 85, 121, 184, 186
 T_CHECK: 184

Index

- T_CURRENT: 184
- T_DEFAULT: 184
- T_EXPEDITED: 184
- T_FAILURE: 184
- T_MORE: 184
- T_NEGOTIATE: 184
- T_NOTSUPPORT: 184
- T_PARTSUCCESS: 184
- T_READONLY: 184
- T_SUCCESS: 184
- flow control: 21
- fork: 3-4, 107-108
- full duplex: 16
- General purpose defines: 189
- headers: 11
- incoming events: 28
- INET_IP: **193**
- initialisation: 13-15, 20, 30, 75
- initiator: 13, 197
- Internet protocol-specific information: 123
- Internet Transport-specific Information: 11
- IP-level options: 126
- IP-level Options: 193
- IP_BROADCAST: 127, 193
- IP_DONTROUTE: 127, 193
- IP_OPTIONS: 127, 193
- IP_REUSEADDR: 127, 193
- IP_TOS: 128, 193
- IP_TOS type of service: 194
- IP_TTL: 128, 193
- ISO: 111, 190, 195
 - priorities: 190
 - protection levels: 190
 - transport classes: 190
- ISO Transport Protocol: 11
- ISO_TP: **191**
- language-dependent: 61
- library functions: 184
- library structure: 52
- listen: 55, 71, 119
- listener application: 4
- management options: 114, 117, 192
- mandatory features: 136
- maximum size:
 - address: 65, 76
 - address buffer: 54, 68
 - buffer: 59, 71, 88, 93, 95
- ETSDU: 66, 76, 99
- TSDU: 18, 66, 76, 99, 130
- memory:
 - allocate: 52, 63
- mode:
 - asynchronous: 5
 - connection-oriented: 13, 15, 22, 31, 112, 129, 131-132
 - connectionless: 20, 22, 30, 93, 95, 104, 116, 135
 - record-oriented: 18
 - stream-oriented: 18
 - synchronous: 5, 73
- modes of service: 4
- multiple options: 41
- Multiple use of addresses: 8
- netbuf structure: 35, 52, 78
- next state: 30
- NEXTHDR: **189**
- NULL: 55
- null:
 - call: 100
 - caller address: 49
- null pointer: 53, 55, 59, 61, 63, 77, 88-89, 95
- ocnt: **27**
- open: 75
- opened: 27, 30
- option:
 - buffer: 35
 - value: 42
- Option management: 10
- option management: 189
- option negotiation:
 - initiate: 37
 - response: 38
- option values: 112
- options:
 - association-related: 33
 - connection mode: 112
 - connectionless mode: 116
 - expedited data: 112
 - format: 45
 - generalities: 33
 - illegal: 36
 - IP-level: 126
 - ISO-specific: 190-191
 - management: 114

- multiple: 41
- privileged: 40
- quality of service: 112
- read-only: 40
- retrieving information: 39
- TCP-level: 124
- transport endpoint: 78
- transport provider: 59
- UDP-level: 126
- unsupported: 38
- Options and management structures: 9
- Options management structure: **187**
- optmgmt: 22, 27, 30
- orderly release: 18-19, 91, 102, 123, 197
- OSI: 9, 195
 - transport classes: 120
- outgoing events: 27
- outstanding connect indications: 29, 55, 89
- O_NONBLOCK flag: 5
- pass_conn: 28, 31
- permissions: 131
- poll: 73
- poll(): 139
- POLLIN: 141
- polling: 6
- POLLOUT: 141
- POLLPRI: 141
- POLLRDBAND: 141
- POLLRDNORM: 141
- POLLWRBAND: 141
- portability: 46
- portable: 3, 22, 136, 141, 155, 181
- precedence levels:
 - IP: 194
- primitives: 5-6
- process: 3
- program: 47
- programs:
 - multiple protocol: 136
- protocol: 3, 23, 33, 54-55, 59, 65, 68, 75, 78, 95, 111, 136
- protocol independence: 66, 77, 136
- Protocol options: 9
- protocol-specific servicelimits: 186
- PUSH flag: 123
- quality of service: 112, 116, 192
- queue: 15, 23, 72
- rate: 112
- rate structure: **191**
- rcv: 28, 31
- rcvconnect: 28, 31
- rcvdis1: 28, 31
- rcvdis2: 28, 31
- rcvdis3: 28, 31
- rcvrel: 28, 31
- rcvudata: 28
- rcvuderr: 28
- readfds: 155
- reason:
 - disconnect: 89
- receipt: 91
- receive: 85, 93
- Receiving Data: 16, 20
- record-oriented: 18
- release: 13, 18, 31, 91, 100, 102
- reliable: 4, 197
- remote user: 7, 16, 18-19, 57, 59-60, 100, 132, 135, 198
- reqvalue: 112, 191
- reqvalue structure: **191**
- resfd: **27**
- responder: 13, 197
- select(): 153
- Sending Data: 17, 21
- server program: 141, 155
- service definition: 14
 - ISO: 1, 18, 111, 119
 - TCP: 18
- service type defines: 186
- snd: 27, 31-32, 111
- snddis1: 27, 31
- snddis2: 27, 31
- sndrel: 27, 31-32
- sndudata: 27, 30, 32, 111
- socket: 9, 129-130
- standard error: 61
- state: 0, 25-26, 30, 189
 - current: 0, 30, 70, 107
 - next: 30
 - T_DATAXFER: **26**, 189
 - T_IDLE: **26**, 189
 - T_INCON: **26**, 189
 - T_INREL: **26**, 189
 - T_OUTCON: **26**, 189

Index

- T_OUTREL: **26**, 189
- T_UNBIND: **26**
- T_UNBND: 189
- T_UNIT: **26**
- State behaviour: 8
- state table: 30-31, 135
- State tables:
 - optmgmt: 10
- status:
 - connect request: 16, 87
 - connection: 59
- stream-oriented: 18
- sterror(3C): 61
- struct netbuf: 186
- struct rate: 191
- struct reqvalue: 191
- struct thrpt: 191
- struct transdel: 191
- struct t_bind: 187
- struct t_call: 188
- struct t_discon: 187
- struct t_info: 186
- struct t_kpalive: 193
- struct t_linger: 190
- struct t_opthdr: 187
- struct t_optmgmt: 187
- struct t_uderr: 188
- struct t_unitdata: 188
- structure types: 188
 - T_BIND: 188
 - T_CALL: 188
 - T_DIS: 188
 - T_INFO: 188
 - T_OPTMGMT: 188
 - T_UDERROR: 188
 - T_UNITDATA: 188
- SVID: 195
- synchronise: 107
- synchronous mode: 5, 73, 198
- t-opthdr: 112
- TACCES: **183**
- TADDRBUSY: **183**
- TBADADDR: **183**
- TBADDATA: **183**
- TBADF: **183**
- TBADFLAG: **183**
- TBADNAME: **183**
- TBAADOPT: **183**
- TBADQLEN: **183**
- TBADSEQ: **183**
- TBUFOVFLW: **183**
- TC: 195
- TCL_CHECKSUM: **192**
- TCL_PRIORITY: **192**
- TCL_PROTECTION: **192**
- TCL_RESERRORRATE: **192**
- TCL_TRANSDel: **192**
- TCO_ACKTIME: **192**
- TCO_ALTCLASS1: **192**
- TCO_ALTCLASS2: **192**
- TCO_ALTCLASS3: **192**
- TCO_ALTCLASS4: **192**
- TCO_CHECKSUM: **192**
- TCO_CLASS: **192**
- TCO_CONNRRESIL: **192**
- TCO_ESTDELAY: **192**
- TCO_ESTFAILPROB: **192**
- TCO_EXPd: **192**
- TCO_EXTFORM: **192**
- TCO_FLOWCTRL: **192**
- TCO_LTPDU: **192**
- TCO_NETEXP: **192**
- TCO_NETRECPTCF: **192**
- TCO_PRIORITY: **192**
- TCO_PROTECTION: **192**
- TCO_REASTIME: **192**
- TCO_RELDELAY: **192**
- TCO_RELFAILPROB: **192**
- TCO_RESERRORRATE: **192**
- TCO_THROUGHPUT: **192**
- TCO_TRANSDel: **192**
- TCO_TRANSFFAILPROB: **192**
- TCP: 9, 18-19, 123, 129, 195
- TCP-level options: 124, 193
- TCP_KEEPAIVE: 124, 193
- TCP_MAXSEG: 125, 193
- TCP_NODELAY: 125, 193
- terminated:
 - connection: 132
- TFlow: 21, 138, 183
- thrpt: 112, 191
- thrpt structure: **191**
- TINDOUT: **183**
- TLI: 181, 195

- TLOOK: 6, 18, 32, 183
- TNOADDR: **183**
- TNODATA: **183**
- TNODIS: **183**
- TNOREL: **183**
- TNOSTRUCTYPE: **183**
- TNOTSUPPORT: 8, 183
- TNOUDERR: **183**
- TOS precedence levels: 194
- TOUTSTATE: **183**
- TPDU lengths: 191
- TPROTO: **183**
- TPROVMISMATCH: **183**
- TQFULL: **183**
- transdel: 112
- transdel structure: 191
- transport address: 3, 111, 198
- transport classes: 120, 190
- transport connection: 3, 15, 57, 65, 102, 197-198
- transport endpoint: 3, 26-27, 42, 49, 54, 56-58, 73, 75, 77-78, 107, 109, 198
- Transport Level Interface (TLI): 181
- transport provider: 3, 13, 25-26, 30, 35, 59, 65, 75, 107, 111, 131, 186
- transport provider identifier: 3, 13, 75, 198
- transport service: 0, 111, 131
- transport service access point (TSAP): 4, 198
- transport service data unit (TSDU): 7, 76, 85, 136, 198
- transport service provider: **198**
- transport service user: 3, 13, 15, 18, 25-26, 58, 131, 198
- transport user actions: 29
- TRESADDR: **183**
- TRESQLEN: **183**
- TSAP: 4, 195, 198
- TSDU: 7, 17, 75-76, 85, 97-99, 104, 111, 121, 123, 195, 198
- TSTATECHNG: **183**
- TSYSERR: 5, 18, 61, 183
- T_ABSREQ: **189**
- t_accept: 22, 49, 119, 123, 129, 184
- t_accept(): 10
- T_ACTIVEPROTECT: **190**
- T_ADDR: **188**
- T_ALIGN: **189**
- T_ALL: **188**
- t_alloc: 22, 52, 63, 184, 188
- T_ALLOPT: **189**
- t_bind: 3, 22, 54, 109, 119, 129, 184
- T_BIND: **188**
- T_CALL: **188**
- T_CHECK: **184**
- T_CLASS0: **190**
- T_CLASS1: **190**
- T_CLASS2: **190**
- T_CLASS3: **190**
- T_CLASS4: **190**
- t_close: 22, 57, 131, 184
- T_CLTS: **186**
- T_CONNECT: 7
- t_connect: 7, 22, 58, 87, 119, 129
- T_CONNECT: 138, 184
- t_connect: 184
- t_connect(): 11, 39
- T_COTS: **186**
- T_COTS_ORD: **186**
- T_CRITIC_ECP: **194**
- T_CURRENT: **184**
- T_DATA: 7, 20, 31-32, 138, 184
- T_DATAXFER: **26**, 31, 189
- T_DEFAULT: **184**
- T_DIS: 7, 32, 188
- T_DISCONNECT: 7, 18, 32, 138, 184
- t_errno: 5, 106, 184
- t_error: 5, 22, 61, 106, 184
- T_EXDATA: 7, 129, 138, 184
- T_EXPEDITED: 17, 85, 184
- T_FAILURE: **184**
- T_FLASH: **194**
- t_free: 22, 63, 185
- t_getinfo: 22, 65, 119, 185
- t_getprotaddr: 22, 68, 185
- t_getstate: 22, 70, 185
- T_GODATA: 7-8, 20, 138, 184
- T_GOEXDATA: 7-8, 138, 184
- T_HIRES: **194**
- T_HITHRPT: **194**
- T_IDLE: **26**, 30-31, 58, 189
- T_IMMEDIATE: **194**
- T_INCON: **26**, 31, 189
- T_INETCONTROL: **194**
- T_INFINITE: **189**

Index

T_INFO: **188**
T_INREL: **26**, 31, 189
T_INVALID: **189**
t_kpalive: **193**
T_LDELAY: **194**
T_LISTEN: 7
t_listen: 7, 22
T_LISTEN: 32
t_listen: 32, 71, 120, 123, 129
T_LISTEN: 138, 184
t_listen: 185
t_listen(): 11, 40
t_look: 6, 22, 73, 129, 185
T_LTPDUFLT: **191**
T_MORE: 17, 85, 97, 184
T_MORE flag: 123
T_NEGOTIATE: **184**
T_NETCONTROL: **194**
T_NO: **189**
T_NOPROTECT: **190**
T_NOTOS: **194**
T_NOTSUPPORT: **184**
T_NULL: **189**
t_open: 3, 22, 75, 120, 129, 185
T_OPT: **188**
t_optmgmt: 10-11, 22, 78, 185
T_OPTMGMT: **188**
t_optmgmt(): 40
T_ORDREL: 7-8, 32, 138, 184
T_OUTCON: **26**, 31, 189
T_OUTREL: **26**, 31, 189
T_OVERRIDEFLASH: **194**
T_PARTSUCCESS: **184**
T_PASSIVEPROTECT: **190**
T_PRIDFLT: **190**
T_PRIHIGH: **190**
T_PRILOW: **190**
T_PRIMID: **190**
T_PRIORITY: **194**
T_PRITOP: **190**
t_rcv: 7, 22, 32, 85, 121, 130, 185
t_rcvconnect: 22, 32, 87, 121, 130
t_rcvconnect(): 11, 39
t_rcvdis: 7, 22, 89, 121, 130, 185
t_rcvrel: 7, 32, 91, 123, 185
t_rcvudata: 22, 32, 93, 121, 185
t_rcvudata(): 11, 40
t_rcvuderr: 7, 22, 95, 121, 185
t_rcvuderr(): 11, 40
T_READONLY: **184**
T_ROUTINE: **194**
T_SENDZERO: **186**
t_snd: 7, 22, 97, 121, 130, 185
t_snddis: 22, 100, 121, 130, 185
t_sndrel: **102**, 123, 185
t_sndudata: 22, 104, 121, 130, 185
t_sndudata(): 11
t_strerror: 22, 106, 185
T_SUCCESS: **184**
t_sync: 4, 22, 107, 185
T_UDATA: **188**
T_UDERR: 7-8, 20, 32, 138, 184
T_UDERROR: **188**
t_unbind: 22
T_UNBIND: **26**
t_unbind: **109**, 185
T_UNBND: 30-31, 189
T_UNIT: **26**
T_UNITDATA: **188**
T_UNSPEC: 44, 189
T_UNUSED: **189**
T_YES: **189**
UDP: 123, 195
UDP-level options: 126
UDP-level Options: 193
UDP_CHECKSUM: 126, 193
unbind: 22, 27, 30, 32
unitdata: 93, 104
Unitdata error structure: 188
UNIX:
 process: 3
 version-independent: 1
 versions: 181
Use of Options: 10
user application: 13, 20, 198
user data: 89
writefds: 155
XEM: 195
XTI: 0, 195
 applications: 47
 features: 22-23
 library: 47
XTI error return: 184
XTI-level options: 82, 189

xti.h: 47, 71, 183
XTI_DEBUG: **190**
XTI_GENERIC: **189**
XTI_LINGER: **190**
XTI_RCVBUF: **190**
XTI_RCVLOWAT: **190**
XTI_SNDBUF: **190**
XTI_SNTLOWAT: **190**
Zero length TSDUs and TSDU fragments: 9,
66, 77, 97-98, 111, 121