

BS ISO 18440:2016



BSI Standards Publication

Space data and information transfer systems — Space link extension — Internet protocol for transfer services

National foreword

This British Standard is the UK implementation of ISO 18440:2016. It supersedes BS ISO 18440:2013 which is withdrawn.

The UK participation in its preparation was entrusted to Technical Committee ACE/68, Space systems and operations.

A list of organizations represented on this committee can be obtained on request to its secretary.

This publication does not purport to include all the necessary provisions of a contract. Users are responsible for its correct application.

© The British Standards Institution 2016.
Published by BSI Standards Limited 2016

ISBN 978 0 580 93378 3

ICS 49.140

Compliance with a British Standard cannot confer immunity from legal obligations.

This British Standard was published under the authority of the Standards Policy and Strategy Committee on 30 November 2016.

Amendments/Corrigenda issued since publication

Date	Text affected
------	---------------

INTERNATIONAL
STANDARD

BS ISO 18440:2016

ISO
18440

Second edition
2016-11-15

**Space data and information transfer
systems — Space link extension —
Internet protocol for transfer services**

*Systèmes de transfert des informations et données spatiales — Extension
de liaisons spatiales — Protocole Internet pour services de transfert*



Reference number
ISO 18440:2016(E)

© ISO 2016



COPYRIGHT PROTECTED DOCUMENT

© ISO 2016

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Ch. de Blandonnet 8 • CP 401
CH-1214 Vernier, Geneva, Switzerland
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
copyright@iso.org
Web www.iso.org

Published in Switzerland

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular the different approval criteria needed for the different types of ISO documents should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2. www.iso.org/directives

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received. www.iso.org/patents

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation on the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the WTO principles in the Technical Barriers to Trade (TBT) see the following URL: [Foreword - Supplementary information](#)

ISO 18440 was prepared by the Consultative Committee for Space Data Systems (CCSDS) (as CCSDS 913.1-B-2, September 2015) and was adopted (without modifications except those stated in clause 2 of this International Standard) by Technical Committee ISO/TC 20, *Aircraft and space vehicles*, Subcommittee SC 13, *Space data and information transfer systems*.

This second edition cancels and replaces the first edition (ISO 18440:2013), which has been technically revised.

STATEMENT OF INTENT

The Consultative Committee for Space Data Systems (CCSDS) is an organization officially established by the management of its members. The Committee meets periodically to address data systems problems that are common to all participants, and to formulate sound technical solutions to these problems. Inasmuch as participation in the CCSDS is completely voluntary, the results of Committee actions are termed **Recommended Standards** and are not considered binding on any Agency.

This **Recommended Standard** is issued by, and represents the consensus of, the CCSDS members. Endorsement of this **Recommendation** is entirely voluntary. Endorsement, however, indicates the following understandings:

- o Whenever a member establishes a CCSDS-related **standard**, this **standard** will be in accord with the relevant **Recommended Standard**. Establishing such a **standard** does not preclude other provisions which a member may develop.
- o Whenever a member establishes a CCSDS-related **standard**, that member will provide other CCSDS members with the following information:
 - The **standard** itself.
 - The anticipated date of initial operational capability.
 - The anticipated duration of operational service.
- o Specific service arrangements shall be made via memoranda of agreement. Neither this **Recommended Standard** nor any ensuing **standard** is a substitute for a memorandum of agreement.

No later than five years from its date of issuance, this **Recommended Standard** will be reviewed by the CCSDS to determine whether it should: (1) remain in effect without change; (2) be changed to reflect the impact of new technologies, new requirements, or new directions; or (3) be retired or canceled.

In those instances when a new version of a **Recommended Standard** is issued, existing CCSDS-related member standards and implementations are not negated or deemed to be non-CCSDS compatible. It is the responsibility of each member to determine when such standards or implementations are to be modified. Each member is, however, strongly encouraged to direct planning for its new standards and implementations towards the later version of the Recommended Standard.

FOREWORD

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. CCSDS has processes for identifying patent issues and for securing from the patent holder agreement that all licensing policies are reasonable and non-discriminatory. However, CCSDS does not have a patent law staff, and CCSDS shall not be held responsible for identifying any or all such patent rights.

Through the process of normal evolution, it is expected that expansion, deletion, or modification of this document may occur. This Recommended Standard is therefore subject to CCSDS document management and change control procedures, which are defined in *Organization and Processes for the Consultative Committee for Space Data Systems* (CCSDS A02.1-Y-4). Current versions of CCSDS documents are maintained at the CCSDS Web site:

<http://www.ccsds.org/>

Questions relating to the contents or status of this document should be sent to the CCSDS Secretariat at the e-mail address indicated on page i.

At time of publication, the active Member and Observer Agencies of the CCSDS were:

Member Agencies

- Agenzia Spaziale Italiana (ASI)/Italy.
- Canadian Space Agency (CSA)/Canada.
- Centre National d'Etudes Spatiales (CNES)/France.
- China National Space Administration (CNSA)/People's Republic of China.
- Deutsches Zentrum für Luft- und Raumfahrt (DLR)/Germany.
- European Space Agency (ESA)/Europe.
- Federal Space Agency (FSA)/Russian Federation.
- Instituto Nacional de Pesquisas Espaciais (INPE)/Brazil.
- Japan Aerospace Exploration Agency (JAXA)/Japan.
- National Aeronautics and Space Administration (NASA)/USA.
- UK Space Agency/United Kingdom.

Observer Agencies

- Austrian Space Agency (ASA)/Austria.
- Belgian Federal Science Policy Office (BFSP0)/Belgium.
- Central Research Institute of Machine Building (TsNIIMash)/Russian Federation.
- China Satellite Launch and Tracking Control General, Beijing Institute of Tracking and Telecommunications Technology (CLTC/BITTT)/China.
- Chinese Academy of Sciences (CAS)/China.
- Chinese Academy of Space Technology (CAST)/China.
- Commonwealth Scientific and Industrial Research Organization (CSIRO)/Australia.
- Danish National Space Center (DNSC)/Denmark.
- Departamento de Ciência e Tecnologia Aeroespacial (DCTA)/Brazil.
- Electronics and Telecommunications Research Institute (ETRI)/Korea.
- European Organization for the Exploitation of Meteorological Satellites (EUMETSAT)/Europe.
- European Telecommunications Satellite Organization (EUTELSAT)/Europe.
- Geo-Informatics and Space Technology Development Agency (GISTDA)/Thailand.
- Hellenic National Space Committee (HNSC)/Greece.
- Indian Space Research Organization (ISRO)/India.
- Institute of Space Research (IKI)/Russian Federation.
- KFKI Research Institute for Particle & Nuclear Physics (KFKI)/Hungary.
- Korea Aerospace Research Institute (KARI)/Korea.
- Ministry of Communications (MOC)/Israel.
- National Institute of Information and Communications Technology (NICT)/Japan.
- National Oceanic and Atmospheric Administration (NOAA)/USA.
- National Space Agency of the Republic of Kazakhstan (NSARK)/Kazakhstan.
- National Space Organization (NSPO)/Chinese Taipei.
- Naval Center for Space Technology (NCST)/USA.
- Scientific and Technological Research Council of Turkey (TUBITAK)/Turkey.
- South African National Space Agency (SANSA)/Republic of South Africa.
- Space and Upper Atmosphere Research Commission (SUPARCO)/Pakistan.
- Swedish Space Corporation (SSC)/Sweden.
- Swiss Space Office (SSO)/Switzerland.
- United States Geological Survey (USGS)/USA.

DOCUMENT CONTROL

Document	Title	Date	Status
CCSDS 913.1-B-1	Space Link Extension—Internet Protocol for Transfer Services, Recommended Standard, Issue 1	September 2008	Original issue, superseded
CCSDS 913.1-B-2	Space Link Extension—Internet Protocol for Transfer Services, Recommended Standard, Issue 2	September 2015	Current issue: – changes the recommended algorithm for secure one-way hash function from SHA-1 to SHA-256; – updates references.

NOTE – Substantive changes from the previous issue are marked with change bars in the inside margin.

CONTENTS

<u>Section</u>	<u>Page</u>
1 INTRODUCTION	1-1
1.1 PURPOSE.....	1-1
1.2 SCOPE.....	1-1
1.3 APPLICABILITY.....	1-2
1.4 RATIONALE.....	1-2
1.5 DOCUMENT STRUCTURE.....	1-2
1.6 DEFINITIONS, NOMENCLATURE, AND CONVENTIONS.....	1-4
1.7 REFERENCES.....	1-8
2 DESCRIPTION OF THE INTERNET SLE PROTOCOL	2-1
2.1 INTRODUCTION.....	2-1
2.2 ARCHITECTURAL MODEL.....	2-1
2.3 AUTHENTICATION LAYER.....	2-3
2.4 DATA ENCODING LAYER.....	2-4
2.5 TRANSPORT MAPPING LAYER.....	2-5
2.6 INTERFACES.....	2-11
2.7 SECURITY ASPECTS OF THE INTERNET SLE PROTOCOL.....	2-22
3 ISP1 MESSAGES AND PROCEDURES	3-1
3.1 AUTHENTICATION LAYER.....	3-1
3.2 DATA ENCODING LAYER.....	3-3
3.3 TRANSPORT MAPPING LAYER.....	3-4
4 TML STATE TABLE	4-1
4.1 INTRODUCTION.....	4-1
4.2 NOTATION.....	4-1
4.3 STATES.....	4-2
4.4 EVENTS.....	4-2
4.5 PREDICATES.....	4-4
4.6 ACTIONS.....	4-4
4.7 STATE TABLE.....	4-6
ANNEX A TML DIAGNOSTIC CODES (NORMATIVE)	A-1
ANNEX B DIFFERENCES WITH EARLIER IMPLEMENTATIONS (INFORMATIVE)	B-1
ANNEX C INDEX TO DEFINITIONS (INFORMATIVE)	C-1
ANNEX D ACRONYMS (INFORMATIVE)	D-1

ANNEX E INFORMATIVE REFERENCES (INFORMATIVE)E-1

CONTENTS (continued)

<u>Figure</u>		<u>Page</u>
1-1	SLE Services and SLE API Documentation	1-3
2-1	ISP1 Architecture Model	2-2
3-1	ASN.1 Type for Generation of ‘the Protected’	3-2
3-2	ASN.1 Type for the Credentials Parameter	3-3
3-3	Layout of a TML Message	3-4
3-4	Layout of a TML Context Message	3-5

Table

2-1	Primitives of the AL Interface	2-11
2-2	Parameters of the Primitive AL-SLE-PDU	2-11
2-3	Primitives of the DEL Interface	2-12
2-4	Parameters of the Primitive DEL-SLE-PDU	2-12
2-5	Primitives of the TML Data Transfer Interface	2-12
2-6	Parameters of the Primitive TML-SLE-PDU	2-13
2-7	Primitives of the TML Association Control Interface	2-13
2-8	Parameters of the Primitive TML-CONNECT	2-13
2-9	Parameters of the Primitive TML-PEER-ABORT	2-14
2-10	Parameters of the Primitive TML-PROTOCOL-ABORT	2-15
2-11	Primitives of the TML Listener Interface	2-15
2-12	Parameters of the Primitive TML-START-LISTEN	2-15
2-13	Parameters of the Primitive TML-STOP-LISTEN	2-16
2-14	Primitives Used for the TCP-Interface	2-19
2-15	Parameters of the Primitive TCP-PASSIVE-OPEN	2-19
2-16	Parameters of the Primitive TCP-CONNECT	2-20
2-17	Parameters of the Primitive TCP-DATA	2-21
3-1	TML Message Type Identifiers	3-5

1 INTRODUCTION

1.1 PURPOSE

The Space Link Extension (SLE) Reference Model (reference [1]) identifies a set of SLE Transfer Services that enable missions to send forward space link data units to a spacecraft and to receive return space link data units from a spacecraft. A subset of these services is specified by the SLE Transfer Service Recommended Standards (references [2], [3], [4], [5], and [6]). The SLE Transfer Service Recommended Standards specify

- a) the operations necessary to provide the transfer service;
- b) the parameter data associated with each operation;
- c) the behaviors that result from the invocation of each operation; and
- d) the relationship between, and the valid sequence of, the operations and resulting behaviors.

However, they deliberately do not specify the methods or technologies required for communications.

The purpose of this Recommended Standard is to define a protocol for transfer of SLE Protocol Data Units (PDUs) defined in the SLE Transfer Service Recommended Standards using the Internet protocols TCP (Transmission Control Protocol, reference [7]) and IP (Internet Protocol, reference [8]) for data transfer and the Abstract Syntax Notation One (ASN.1, references [9] and [10]) for data encoding. This protocol is referred to as the Internet SLE Protocol One (ISP1).

1.2 SCOPE

This Recommended Standard defines a protocol for transfer of SLE PDUs between an SLE user and an SLE provider system in terms of:

- a) the procedures used to establish and release associations;
- b) the messages exchanged on an established association;
- c) the procedures used to monitor the status of data communication connections; and
- d) the methods used to ensure that data are converted between different formats and representations on different platforms.

It does not specify:

- a) individual designs, implementations, or products;
- b) the configuration of the data communications infrastructure, including configuration of the TCP and IP protocols;

- c) the means by which addresses (IP addresses and TCP port numbers) are agreed, assigned, and communicated.

This Recommended Standard responds to the requirements imposed by the Recommended Standards for SLE transfer services that were available when this Recommended Standard was released. The protocol specified in this Recommended Standard conforms to the requirements on data communication services set forth in those Recommended Standards.

1.3 APPLICABILITY

1.3.1 APPLICABILITY OF THIS RECOMMENDED STANDARD

This Recommended Standard provides a basis for the development of real systems that implement the Internet SLE Protocol. It is applicable for systems acting as an SLE service user or SLE service provider.

1.3.2 LIMITS OF APPLICABILITY

This Recommended Standard specifies the Internet SLE Protocol that may be applied by an SLE System for inter-Agency cross support. It is neither a specification of, nor a design for, real systems that may be implemented for the control and monitoring of existing or future missions.

1.4 RATIONALE

The goal of this Recommended Standard is to create a standard for interoperability between the tracking stations and/or ground data handling systems of various agencies and the users of SLE transfer services based on the technologies TCP/IP and ASN.1.

1.5 DOCUMENT STRUCTURE

1.5.1 ORGANIZATION

This document is organized as follows:

- a) section 1 presents the purpose, scope, applicability and rationale of this Recommended Standard and lists the definitions, conventions, and references used throughout the Recommended Standard;
- b) section 2 describes the Internet SLE Protocol by means of an architectural model identifying individual protocol layers and the interfaces to higher layers;
- c) section 3 specifies the messages exchanged via ISP1 and the procedures to be applied for connection establishment and release and for data transfer;

SPACE LINK EXTENSION—INTERNET PROTOCOL FOR TRANSFER SERVICES

- d) section 4 specifies the state table for the protocol;
- e) annex A provides ISP1-specific diagnostic codes for the SLE PEER-ABORT operation;
- f) annex B describes differences with earlier implementations of ISP1;
- g) annex C lists all terms used in this document and identifies where they are defined;
- h) annex D lists all acronyms used within this document;
- i) annex E contains a list of informative reference documents.

1.5.2 SLE SERVICES DOCUMENTATION TREE

This Recommended Standard is part of a suite of documents specifying the SLE services. The SLE services constitute one of the three types of Cross Support Services:

- a) Part 1: SLE Services;
- b) Part 2: Ground Domain Services;
- c) Part 3: Ground Communications Services.

The basic organization of the SLE services documentation is shown in figure 1-1. The various documents are described in the following subsections.

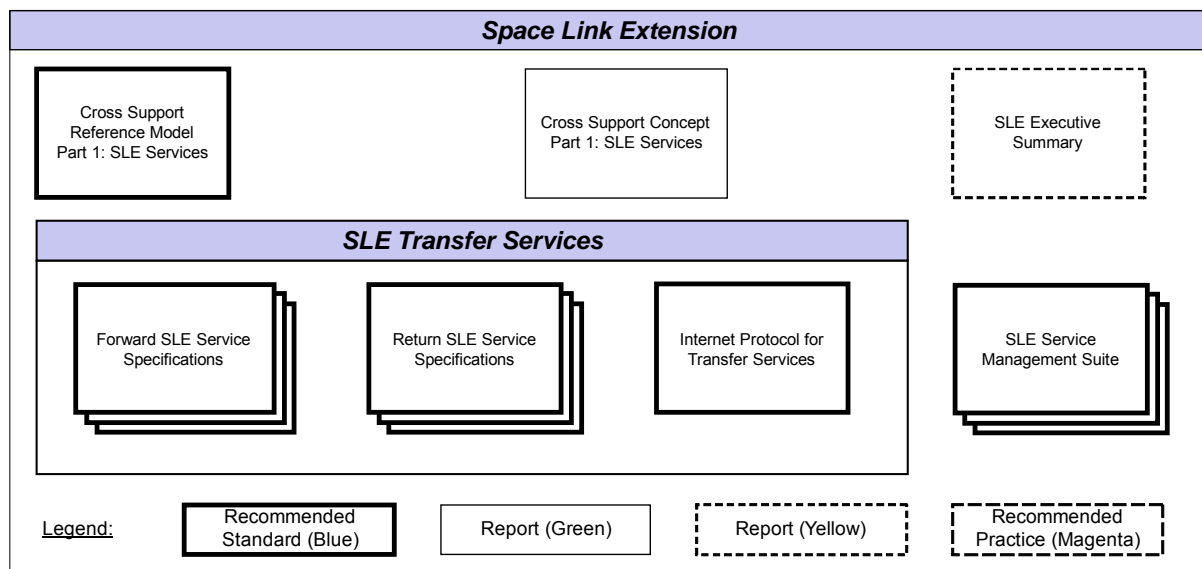


Figure 1-1: SLE Services and SLE API Documentation

- a) *Cross Support Concept—Part 1: Space Link Extension Services* (reference [E1]), a Report introducing the concepts of cross support and the SLE services;

- b) *Cross Support Reference Model—Part 1: Space Link Extension Services* (reference [1]), a Recommended Standard that defines the framework and terminology for the specification of SLE services;
- c) *Forward SLE Service Specifications*, a set of Recommended Standards that will provide specification of all forward link SLE services;
- d) *Return SLE Service Specifications*, a set of Recommended Standards that will provide specification of all return link SLE services;
- e) *Internet Protocol for Transfer Services*, this Recommended Standard;
- f) *SLE Service Management Specifications*, a set of Recommended Standards that establish the basis of SLE service management.

1.6 DEFINITIONS, NOMENCLATURE, AND CONVENTIONS

1.6.1 DEFINITIONS

1.6.1.1 Definitions from the SLE Reference Model

This Recommended Standard makes use of the following terms defined in reference [1]:

- a) initiator;
- b) operation;
- c) responder;
- d) service user (user);
- e) service provider (provider);
- f) SLE protocol data unit (SLE-PDU);
- g) SLE transfer service instance (service instance).

1.6.1.2 Definitions from SLE Transfer Service Specifications

This Recommended Standard makes use of the following terms defined in references [2], [3], [4], [5], and [6]:

- a) association;
- b) communications service;
- c) confirmed operation;
- d) invocation;

- e) parameter (of an operation);
- f) port identifier;
- g) return;
- h) unconfirmed operation.

1.6.1.3 Definitions from TCP/IP

This Recommended Standard makes use of the following terms defined in references [7] and [8]:

- a) Internet Protocol (IP);
- b) IP address;
- c) port (of TCP);
- d) port number;
- e) Transmission Control Protocol (TCP);
- f) segment (of TCP);
- g) socket.

1.6.1.4 Definitions from Abstract Syntax Notation One

This Recommended Standard makes use of the following terms defined in references [9] and [10]:

- a) Abstract Syntax Notation One (ASN.1);
- b) Basic Encoding Rules (BER);
- c) Distinguished Encoding Rules (DER);
- d) encoding rules (of ASN.1);
- e) encoding;
- f) module (of ASN.1).

1.6.1.5 Definitions from OSI Basic Reference Model

This Recommended Standard makes use of the following terms defined in reference [14]:

- a) abstract syntax;
- b) primitive;

- c) (protocol) layer;
- d) transfer syntax.

1.6.1.6 Additional Definitions

1.6.1.6.1 General

For the purposes of this Recommended Standard, the following definitions also apply.

1.6.1.6.2 (SLE) Application

An application is a software entity in an SLE user system or an SLE provider system that makes use of the ISP1 protocol, as distinguished from the software implementing the protocol layers defined in this Recommended Standard. The application is considered to implement the ‘higher layers’ defined in the architectural model in section 2.

1.6.1.6.3 Local Application

The local application is the application implementing the higher layers interfacing with a given instance of an ISP1 implementation.

1.6.1.6.4 Peer Application

The peer application is the application that communicates with the local application via the ISP1 protocol. The peer application is typically located on a remote network, but may also be located on the local network, or even on the same host as the local application.

1.6.1.6.5 Application Identifier

The application identifier is the authority identifier of the application by which the application is identified in the BIND invocation and the BIND return.

NOTE – For an initiating SLE application, the application identifier is the initiator-identifier, and for a responding SLE application, the application identifier is the responder-identifier.

1.6.2 NOMENCLATURE

1.6.2.1 Normative Text

The following conventions apply for the normative specifications in this Recommended Standard:

- a) the words ‘shall’ and ‘must’ imply a binding and verifiable specification;

- b) the word ‘should’ implies an optional, but desirable, specification;
- c) the word ‘may’ implies an optional specification;
- d) the words ‘is’, ‘are’, and ‘will’ imply statements of fact.

NOTE – These conventions do not imply constraints on diction in text that is clearly informative in nature.

1.6.2.2 Informative Text

In the normative sections of this document, informative text is set off from the normative specifications either in notes or under one of the following subsection headings:

- Overview;
- Background;
- Rationale;
- Discussion.

1.6.3 CONVENTIONS

The Internet SLE Protocol is specified by a layered architecture model in which the interfaces between the layers are defined using abstract service primitives, roughly following the concepts in the OSI Basic Reference Model (reference [14]).

A service primitive is a signal optionally associated with a set of parameters that is generated by one layer and consumed by another (adjacent) layer. The direction in which the information is passed is defined by one of the following:

request	a primitive that is passed from the higher layer to the lower layer;
indication	a primitive that is passed from the lower layer to the higher layer;
response	a primitive generated by the higher layer in response to an indication of the same type;
confirmation	a primitive generated by the lower layer in response to a request of the same type.

For every service primitive, the following specifications are provided:

- a) the name of the primitive;
- b) the uses of the primitive (request, indication, response, confirmation);
- c) the parameters associated with each of these uses.

In addition, the conditions under which the source emits a primitive and the tasks that the receiver shall perform are described.

In this Recommended Standard, primitive names are capitalized and are followed by the specific use in lowercase letters. The primitive name is prefixed by the identifier of the layer that defines the interface, e.g., ‘TML-CONNECT-indication’. In the state table in section 4, the prefix is omitted in order to avoid clashes with the notation used in these tables.

1.7 REFERENCES

The following documents contain provisions which, through reference in this text, constitute provisions of this Recommended Standard. At the time of publication, the editions indicated were valid. All documents are subject to revision, and users of this Recommended Standard are encouraged to investigate the possibility of applying the most recent editions of the documents indicated below. The CCSDS Secretariat maintains a register of currently valid CCSDS Recommended Standards.

NOTE – A list of informative references is provided in annex E.

- [1] *Cross Support Reference Model—Part 1: Space Link Extension Services*. Issue 2. Recommendation for Space Data System Standards (Blue Book), CCSDS 910.4-B-2. Washington, D.C.: CCSDS, October 2005.
- [2] *Space Link Extension—Return All Frames Service Specification*. Issue 3. Recommendation for Space Data System Standards (Blue Book), CCSDS 911.1-B-3. Washington, D.C.: CCSDS, January 2010.
- [3] *Space Link Extension—Return Channel Frames Service Specification*. Issue 2. Recommendation for Space Data System Standards (Blue Book), CCSDS 911.2-B-2. Washington, D.C.: CCSDS, January 2010.
- [4] *Space Link Extension—Return Operational Control Fields Service Specification*. Issue 2. Recommendation for Space Data System Standards (Blue Book), CCSDS 911.5-B-2. Washington, D.C.: CCSDS, January 2010.
- [5] *Space Link Extension—Forward CLTU Service Specification*. Issue 3. Recommendation for Space Data System Standards (Blue Book), CCSDS 912.1-B-3. Washington, D.C.: CCSDS, July 2010.
- [6] *Space Link Extension—Forward Space Packet Service Specification*. Issue 2. Recommendation for Space Data System Standards (Blue Book), CCSDS 912.3-B-2. Washington, D.C.: CCSDS, July 2010.
- [7] J. Postel. *Transmission Control Protocol*. STD 7. Reston, Virginia: ISOC, September 1981.
- [8] J. Postel. *Internet Protocol*. STD 5. Reston, Virginia: ISOC, September 1981.

- [9] *Information Technology—Abstract Syntax Notation One (ASN.1): Specification of Basic Notation*. 4th ed. International Standard, ISO/IEC 8824-1:2008. Geneva: ISO, 2008.
- [10] *Information Technology—ASN.1 Encoding Rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)*. 4th ed. International Standard, ISO/IEC 8825-1:2008. Geneva: ISO, 2008.
- [11] *Information Technology—Open Systems Interconnection—The Directory—Part 8: Public-Key and Attribute Certificate Frameworks*. 7th ed. International Standard, ISO/IEC 9594-8:2014. Geneva: ISO, 2014.
- [12] *Secure Hash Standard (SHS)*. Federal Information Processing Standards Publication 180-4. Gaithersburg, Maryland: NIST, March 2012.
- [13] *Time Code Formats*. Issue 4. Recommendation for Space Data System Standards (Blue Book), CCSDS 301.0-B-4. Washington, D.C.: CCSDS, November 2010.
- [14] *Information Technology—Open Systems Interconnection—Basic Reference Model: The Basic Model*. 2nd ed. International Standard, ISO/IEC 7498-1:1994. Geneva: ISO, 1994.

2 DESCRIPTION OF THE INTERNET SLE PROTOCOL

2.1 INTRODUCTION

This section describes a layered model of a system supporting the Internet SLE Protocol One and introduces the main concepts by describing each layer and the interfaces between the layers. Detailed specifications are provided in section 3, which references the concepts, layers, and interfaces described by this model. It is stressed that this model has the sole purpose of supporting the specifications in this Recommended Standard and is not intended to suggest any specific design of a real implementation.

The discussion in this section assumes that the reader is familiar with the CCSDS Recommended Standards for Space Link Extension transfer services provided by references [2], [3], [4], [5], and [6], and assumes a general background on the Internet protocols, especially on TCP (reference [7]).

2.2 ARCHITECTURAL MODEL

The architectural model used to specify the Internet SLE Protocol is shown in figure 2-1.

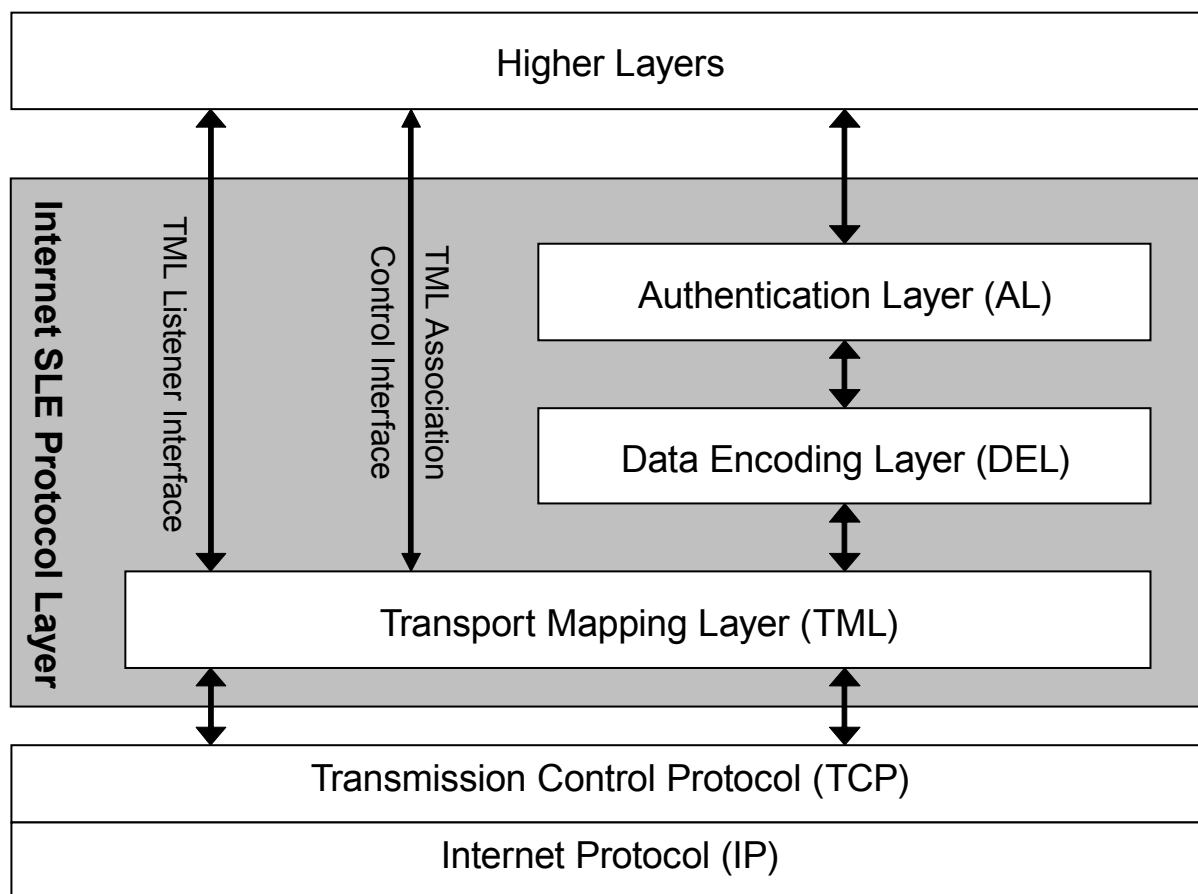


Figure 2-1: ISP1 Architecture Model

The Higher Layers represent the functionality specified by the Recommended Standards for SLE transfer services (references [2], [3], [4], [5], and [6]), including in particular:

- a) preparation of SLE protocol data units to be sent to the peer application;
- b) analysis and processing of the SLE protocol data units received from the peer application;
- c) implementation of the state tables defined in the applicable Recommended Standard for the SLE transfer service being provided or used.

The Internet SLE Protocol Layer is further decomposed into the following sub-layers:

- a) The Authentication Layer (AL) is responsible for generating and analyzing the credentials specified in the Recommended Standards for SLE transfer services. For this purpose this Recommended Standard specifies use of the simple authentication scheme defined in reference [11] and the Secure Hash Function (SHA-256) defined in reference [12].
- b) The Data Encoding Layer (DEL) is responsible for encoding of SLE protocol data units received from higher layers and decoding of protocol data units received from the peer application. For this purpose, this Recommended Standard specifies use of the ASN.1 syntax defined by reference [9] and of the ASN.1 Basic Encoding Rules (BER) defined by reference [10].
- c) The Transport Mapping Layer (TML) handles the interface to the Transmission Control Protocol (TCP) specified by reference [7].

ISP1 maps one TCP connection to one ISP1 association, which is used by the higher layers for one SLE association as specified by the Recommended Standards for SLE transfer services. ISP1 provides data encoding and decoding services but does not process SLE protocol data units in any other respect.

For implementations assuming the responder role, the Transport Mapping Layer provides a service to listen for incoming TCP connection requests. It therefore provides an interface by which higher layers can request to start and stop listening for a specified SLE responder port identifier, which the TML maps to an IP address and a TCP port number.

Except for the interface to control listening, all interfaces refer to one association only. ISP1 association establishment and release involves only the Transport Mapping Layer. Once the association has been established, all SLE protocol data units are passed through the Data Encoding Layer and the Authentication Layer.

ISP1 requires a number of configuration parameters for its operation, which are identified in the following subsections. This Recommended Standard does not specify how these parameters are defined, agreed, stored, and made available to the implementation of the ISP1.

2.3 AUTHENTICATION LAYER

The ISP1 Authentication Layer (AL) receives SLE protocol data units from the higher layers and adds the credentials parameter if required. Likewise, it receives decoded protocol data units from the Data Encoding Layer and verifies that the credentials match the security attributes of the peer application. If authentication is not used for a given service instance or for a given PDU, the Authentication Layer passes the protocol data unit to the next layer without modification or analysis.

For generation and analysis of the credentials, the Authentication Layer uses the simple authentication scheme specified by reference [11], which is based on a secret password and a message digest generated from that password, the time at which the credentials were generated, and a random number. For the secure one-way hash function, required for this scheme, the Secure Hash Function (SHA-256) specified by reference [12] is used.

If authentication fails for a PDU received from the peer application, the Authentication Layer informs the higher layers, which are expected to handle this event as specified by the Recommended Standards for SLE transfer services.

The Authentication Layer requires the following service instance configuration parameter for its operation:

- authentication-level, defined by the Recommended Standards for SLE transfer services, which can have the values
 - a) ‘all’: all invocations and returns, except the invocation of PEER-ABORT, shall be authenticated;
 - b) ‘bind’: only the BIND invocation and return shall be authenticated;
 - c) ‘none’: no invocation or return shall be authenticated.

If the authentication-level is set to ‘all’ or ‘bind’ the Authentication Layer requires the following additional configuration parameters:

- a) the identifier of the local application;
- b) the password of the local application;
- c) the identifier of the peer application;
- d) the password of the peer application;
- e) the maximum time allowed between generation of the credentials by the invoker of an SLE operation and verification of the credentials by the performer.

NOTE – If the local application assumes the initiator role for the BIND operation, the identifier of the local application is the initiator-identifier and the identifier of the peer application is the responder-identifier defined by the Recommended Standards for SLE transfer services. Otherwise, the identifier of the local application is the responder-identifier and the identifier of the peer application is the initiator-identifier.

2.4 DATA ENCODING LAYER

The ISPI Data Encoding Layer (DEL) is responsible for the encoding and decoding of all SLE protocol data units, except for the PEER-ABORT invocation PDU. For this purpose it uses the ASN.1 types defined by the Recommended Standards for SLE transfer services and applies the Basic Encoding Rules (BER) defined by reference [10].

NOTE – This Recommended Standard assumes that the PEER-ABORT operation is implemented by direct interaction between the TML and higher layers. Handling of the PEER-ABORT operation by the TML is described in 2.5.5.

The DEL receives SLE PDUs from the Authentication Layer, encodes the PDU and passes the encoded PDU as a data buffer to the Transport Mapping Layer. Likewise, it receives encoded PDUs from the Transport Mapping Layer, decodes them, and passes the decoded PDU to the Authentication Layer.

When receiving a BIND invocation PDU from either the AL or the TML, DEL uses the parameters service-type and version-number to identify the ASN.1 modules to use for all PDUs subsequently exchanged on the same association. If the service type or version number is not supported, the DEL informs the higher layers, which are expected to apply the procedures specified by the Recommended Standards for SLE transfer services.

NOTE – This Recommended Standard assumes that the DEL uses the parameters of the BIND operation to identify the ASN.1 modules to use for encoding and decoding, because no additional configuration parameters are then needed. However, the intention is not to prevent an implementation from applying any other suitable approach as long as it is ensured that the ASN.1 specification used is the one specified in the Recommended Standard that is identified by the service-type and version-number parameters in the BIND invocation.

In case of errors (e.g., decoding failure), the DEL informs the higher layers, which are expected to abort the association using the PEER-ABORT operation.

2.5 TRANSPORT MAPPING LAYER

2.5.1 INTRODUCTION

The ISP1 Transport Mapping Layer handles the interface to TCP/IP. It provides the following services:

- a) mapping of logical port identifiers to TCP sockets (IP addresses and TCP port numbers);
- b) establishment and release of TCP connections;
- c) transfer of SLE PDUs;
- d) execution of the PEER-ABORT using features available with TCP;
- e) monitoring of the status of TCP connections.

In addition, this Recommended Standard defines a special connection establishment procedure, which is intended to support fast recovery from failures, provided that a responder uses redundant hosts.

2.5.2 TML MESSAGES

2.5.2.1 General

As the TCP user interface is based on the concept of a byte stream, this Recommended Standard defines a simple message format for transmission of SLE PDUs and exchange of protocol control information between two communicating TMLs. The following types of messages are defined:

- a) an SLE PDU message for transfer of SLE PDUs except for the PEER-ABORT PDU, for which the procedure explained in 2.5.5 is used;
- b) a TML context message, used to initialize a connection between two TMLs;
- c) a TML heartbeat message used to monitor the status of a TCP connection.

A TML message consists of a standard message header of eight octets and a message body. The message header contains an identifier of the message-type and the number of octets in the message body. The message body is defined in the following subsections.

2.5.2.2 SLE PDU Message

The body of the SLE PDU message consists of the encoded SLE PDU.

2.5.2.3 TML Context Message

The body of the TML context message contains the following information:

- a) the identification of the protocol, i.e., the characters 'ISP1' in ASCII encoding;
- b) the version of this Recommended Standard (the number 1);
- c) the heartbeat-interval and dead-factor, used for monitoring of the status of the TCP connection (see 2.5.3).

2.5.2.4 TML Heartbeat Message

The body of the TML heartbeat message is empty. A TML heartbeat message consists of the message header only, where the length field is set to 'zero'.

2.5.3 MONITORING OF THE STATUS OF THE TCP CONNECTION

The Recommended Standards for SLE transfer services require that SLE PDUs be transferred

- a) in sequence;
- b) completely and with integrity;
- c) without duplication;
- d) with flow control that notifies backpressure to the application layer in the event of congestion; and
- e) with notification to the application layer in the event that communications between the SLE service user and the SLE service provider are disrupted, possibly resulting in a loss of data.

While the TCP fully supports requirements a) to d), it does not intrinsically provide a means of distinguishing between an idle connection and a dead connection. Therefore this Recommended Standard specifies the special procedure described in this section.

NOTE – Some TCP implementations provide a configurable 'keep alive' mechanism to periodically probe idle connections. However, reference [E7] specifically requires that the default interval between such probes must be at least two hours.

The procedure makes use of two parameters, a heartbeat-interval and a dead-factor. The heartbeat-interval defines the maximum duration in seconds in which no message might be transmitted by a TML. A value of 'zero' indicates that the heartbeat mechanism shall not be used. The dead-factor defines the number of heartbeat intervals after which a TCP connection is considered dead when no message has been received from the peer TML.

The TML makes use of two timers for a heartbeat mechanism, a heartbeat-transmit-timer (HBT timer) and a heartbeat-receive-timer (HBR timer). It sets the timeout value of the HBT timer to the number of seconds defined by the heartbeat-interval and the timeout value of the HBR timer to the number of seconds defined by the heartbeat-interval multiplied by the dead-factor.

Whenever the TML transmits a message, it restarts the HBT timer. If the HBT timer expires, the TML transmits a TML heartbeat message and restarts the HBT timer.

Whenever the TML receives a message, it restarts the HBR timer. When the HBR timer expires, the TML assumes that the connection has failed, aborts using TCP-ABORT and issues the TML-PROTOCOL-ABORT-indication primitive.

Following successful establishment of the TCP connection, the TML initiating the association transmits the proposed values for the heartbeat-interval and the dead-factor as part of the TML context message, and then starts the HBT timer and the HBR timer.

When receiving the TML context message, the responding TML verifies that the values of the heartbeat-interval and the dead-factor are in the acceptable range. If the values are acceptable, the responding TML starts the HBT timer and the HBR timer. If the values are not acceptable, the TML aborts the connection using PEER-ABORT with the TML diagnostics 'heartbeat unacceptable'.

The heartbeat mechanism remains in effect for the complete lifetime of the TCP connection. As explained in 2.5.4.2, the HBR timer is finally used to ensure that the TCP connection is released in an orderly manner.

2.5.4 CONNECTION ESTABLISHMENT AND RELEASE

2.5.4.1 Initiating TML

When receiving the primitive TML-CONNECT-request, the TML uses the port identifier to derive the IP address and the TCP port number used by the responder. It then requests the TCP to establish a connection, specifying the foreign socket. It is recommended to leave the local socket unspecified and let TCP derive the local IP address and select an ephemeral port.

If the TCP connection succeeds, the TML transmits the TML context message, issues the primitive TML-CONNECT-confirmation, and is ready to transfer SLE PDUs. If the TCP connection fails, the TML issues the primitive TML-PROTOCOL-ABORT with the appropriate diagnostics.

NOTE – This Recommended Standard assumes that the TCP implementation takes care of temporary network problems and retransmissions and that the TCP implementation can be configured as needed. Therefore, any failure reported by the TCP is assumed to be a hard error. However, this Recommended Standard does not exclude that an implementation retries connection establishment for a

configurable period, or until it receives TML-RESET because the return timer has expired.

When receiving the primitive TML-DISCONNECT-request, the TML closes the TCP connection.

2.5.4.2 Responding TML

When receiving the primitive TML-START-LISTEN-request, the TML derives the TCP socket using the logical port identifier passed as argument. It then performs a ‘passive open’ (see 2.6.6.2), specifying the local socket.

NOTE – It is recommended not to specify the foreign socket, with the effect that the TCP accepts connections from any address. Although reference [7] specifies that the foreign socket can be defined by a call to passive open and that the TCP will then accept connections only from that socket, TCP APIs generally do not provide this option. A responder might nevertheless filter the IP addresses from which it is prepared to accept a connection, e.g., by means of a firewall.

When the TML is notified that a new connection has been accepted, it starts a configurable TML start-up timer and waits for the first message to arrive. If the timer expires before data arrive, the TML aborts the connection using TCP-ABORT. The TML checks that the first message received on a connection is a valid TML context message. If that is not the case, it aborts the connection with TCP-ABORT.

If the first message received is a valid TML context message, the TML performs the following checks:

- a) The protocol identifier must be ‘ISP1’; otherwise, the TML aborts using TCP-ABORT.
- b) The version number must be supported by the responding TML. If the version is not supported, the TML aborts using TCP-ABORT.
- c) The heartbeat parameters must be in an acceptable range; otherwise, the TML aborts using TCP-ABORT.

When these checks are passed, the TML issues the primitive TML-CONNECT-indication and is ready to receive and transmit SLE PDUs on the new connection.

When receiving TML-DISCONNECT-request, the TML terminates the HBT timer, restarts the HBR timer, and waits for an indication that the peer has closed its side of the connection. When the peer has closed the connection, the TML closes its side of the TCP connection. If the HBR timer expires before the peer has closed the connection, the TML aborts the connection using TCP-ABORT. The TML also aborts the connection if it receives any data after receiving TML-DISCONNECT-request.

2.5.5 ABORTING CONNECTIONS

The TML provides two methods to abort a TCP connection:

- a) the PEER-ABORT operation by which the diagnostic code is transmitted to the peer TML; and
- b) a ‘silent abort’ by which no further information is made available to the peer.

The latter method is requested by the primitive TML-RESET and is mapped directly to TCP-ABORT.

A ‘Close after Peer Abort’ timer, CPA timer, is used to check that the peer TML reacts to the PEER-ABORT in due time.

When receiving the primitive TML-PEER-ABORT-request, the TML transmits the diagnostic parameter as one byte of TCP urgent data. It then stops the HBT and HBR timers, starts the CPA timer, and waits for the peer to close its side of the connection. The TML silently discards any data arriving after peer abort. When the peer closes the connection, the TML closes its side of the connection. If the peer does not close the connection before the CPA timer expires, the TML aborts the connection with TCP-ABORT.

When the TML is notified that urgent data are present, it reads and discards all data from the TCP receive buffer up to but excluding the byte identified by the TCP urgent pointer. It then reads the PEER-ABORT diagnostics, closes the connection, and issues the primitive TML-PEER-ABORT-indication.

2.5.6 ERROR REPORTING

Except for the initial start-up phase on the responding side (see 2.5.4.2), the TML uses the peer abort procedure to report errors on the TML level such a badly formatted TML message or reception of a TML context message by an initiating TML. In these cases it uses one of the TML diagnostic codes defined in annex A. When a TML receives urgent data with a TML diagnostic code (i.e., a code in the range 128 to 199), it reports the problem using the primitive TML-PROTOCOL-ABORT-indication instead of TML-PEER-ABORT-indication.

NOTE – The Recommended Standards for SLE transfer services reserve PEER-ABORT diagnostic codes 0 to 127 for reporting of errors defined by these Recommended Standards. This Recommended Standard reserves the codes 128 to 199 for errors reported by the TML, leaving the codes 200 to 255 for use by implementations.

2.5.7 REDUNDANT HOSTS

If redundant hosts are used on the responding side, it is desirable to allow connection establishment, without the need to know which host is operational. Such a feature is considered particularly useful for fast recovery after failure of a connection. This version of the specification defines a very basic procedure for this purpose, which avoids the need for additional infrastructure (e.g., a directory service). More reliable and flexible procedures are left for further study and might be included in future versions.

The procedure requires an initiator to associate a set of TCP sockets with one logical responder port identifier. If the TML is requested to connect to a responder port that maps to more than one socket, it attempts to establish a TCP connection to all sockets in the set. It is expected that connection establishment will succeed for only a single socket and that all other connections will fail. If one of the connections succeeds, that TML requests TCP-ABORT on all other sockets in the set.

The procedure requires that only a single responding host listen for incoming connection requests at any time. The means by which this is achieved are outside the scope of this Recommended Standard.

2.5.8 TML CONFIGURATION

2.5.8.1 Initiating TML

An initiating TML requires the following configuration parameter for its operation:

- mapping from responder-port-identifiers to one or more TCP sockets each identified by IP address and TCP port number.

The following configuration parameters are needed for every association:

- a) proposed heart-beat interval;
- b) proposed dead-factor;
- c) close-after-peer-abort timeout.

2.5.8.2 Responding TML

A responding TML requires the following configuration parameter for its operation:

- mapping from responder-port-identifiers to local TCP sockets.

The following configuration parameters are needed for every association:

- a) acceptable range of the heartbeat-interval;
- b) acceptable range of the dead-factor;

- c) TML-startup-timeout;
- d) close-after-peer-abort timeout.

2.6 INTERFACES

2.6.1 AUTHENTICATION LAYER INTERFACES

2.6.1.1 Overview

The primitives defined for the interface between the higher layers and the AL are listed in table 2-1.

Table 2-1: Primitives of the AL Interface

Primitive	Request	Indication	Confirmation
AL-SLE-PDU	X	X	

2.6.1.2 AL-SLE-PDU

The primitive AL-SLE-PDU is used to pass SLE operation invocations and returns between the higher layers and the AL. Its parameters are shown in table 2-2.

Table 2-2: Parameters of the Primitive AL-SLE-PDU

Parameter	Request	Indication
SLE PDU	X	X
authentication-result		X
decoding-result		X

Request The request is used by the higher layers to pass PDUs to the AL for processing and transmission. If applicable, the AL generates and adds the credentials and forwards the PDU to the DEL using DEL-SLE-PDU-request.

Indication The indication is issued by the AL when a PDU received from the DEL via DEL-SLE-PDU-indication has been authenticated. The result of the authentication process is reported in the parameter authentication-result. If authentication of a PDU is not required, the effect is as if the PDU were successfully authenticated. The decoding result is copied from the DEL-SLE-PDU-indication.

2.6.2 DATA ENCODING LAYER INTERFACE

2.6.2.1 Overview

The primitives defined for the interface between the AL and the DEL are listed in table 2-3.

Table 2-3: Primitives of the DEL Interface

Primitive	Request	Indication	Confirmation
DEL-SLE-PDU	X	X	

2.6.2.2 DEL-SLE-PDU

The primitive DEL-SLE-PDU is used to pass SLE operation invocations and returns between the AL and the DEL. Its parameters are shown in table 2-4.

Table 2-4: Parameters of the Primitive DEL-SLE-PDU

Parameter	Request	Indication
SLE PDU	X	X
decoding-result		X

Request The request is used by the AL to forward PDUs received from higher layers to the DEL. The DEL encodes the PDU and forwards the encoded PDU to the TML using TML-SLE-PDU-request.

Indication The indication is issued by the DEL when a PDU received from the TML via TML-SLE-PDU-indication has been decoded or when decoding has failed. The result of the decoding process is reported in the parameter decoding-result.

2.6.3 TML DATA TRANSFER INTERFACE

2.6.3.1 Overview

The primitives defined for the interface between the DEL and the TML are listed in table 2-5.

Table 2-5: Primitives of the TML Data Transfer Interface

Primitive	Request	Indication	Confirmation
TML-SLE-PDU	X	X	

2.6.3.2 TML-SLE-PDU

The primitive TML-SLE-PDU is used to pass encoded SLE PDUs between the TML and the DEL. The parameters of the primitive are shown in table 2-6.

Table 2-6: Parameters of the Primitive TML-SLE-PDU

Parameter	Request	Indication
Encoded PDU	X	X

Request The request is issued by the DEL when receiving DEL-SLE-PDU-request. The TML adds the TML message header and transmits the PDU as a TML SLE PDU message.

Indication The indication is issued by the TML when receiving a TML SLE PDU message.

2.6.4 TML ASSOCIATION CONTROL INTERFACE

2.6.4.1 Overview

The primitives defined for the interface between the TML and the higher layers for establishment and release of associations are listed in table 2-7.

Table 2-7: Primitives of the TML Association Control Interface

Primitive	Request	Indication	Confirmation
TML-CONNECT	X	X	X
TML-DISCONNECT	X		
TML-PEER-ABORT	X	X	
TML-PROTOCOL-ABORT		X	
TML-RESET	X		

2.6.4.2 TML-CONNECT

The primitive TML-CONNECT is used to request establishment of a TCP connection, report an incoming connection request, or confirm successful connection establishment. The parameters of the primitive are shown in table 2-8.

Table 2-8: Parameters of the Primitive TML-CONNECT

Parameter	Request	Indication	Confirmation
Responder Port Identifier	X		

- Request** The request is issued by the higher layers on the initiating side to request a connection to the peer application. The TML attempts to establish a connection to the destination defined by the responder port identifier, as described in 2.5.4 and 2.5.7.
- Indication** The indication is issued by a responding TML when a new TCP connection has been accepted and the initial TML context message has been received and accepted.
- Confirmation** The confirmation is issued by an initiating TML when establishment of a TCP connection succeeds.

2.6.4.3 TML-DISCONNECT

The primitive TML-DISCONNECT is issued by the higher layers to request that a TCP connection be disconnected. An initiating TML initializes the TCP connection release procedure as described in 2.5.4. A responding TML prepares for connection release as described in 2.5.4.2. The primitive does not have any parameters.

2.6.4.4 TML-PEER-ABORT

The primitive TML-PEER-ABORT is used to request abortion of the association via PEER-ABORT or to report that the association has been aborted. Its parameters are listed in table 2-9.

Table 2-9: Parameters of the Primitive TML-PEER-ABORT

Parameter	Request	Indication
Diagnostic (one octet)	X	X
Originator (Local TML or Peer)		X

- Request** The primitive is used by the higher layers to request that the association be aborted using PEER-ABORT. The TML performs the peer abort procedure as described in 2.5.5.
- Indication** The indication is issued by the TML when it receives urgent data containing diagnostics in the range reserved by the Recommended Standards for SLE transfer services.

2.6.4.5 TML-PROTOCOL-ABORT

The primitive TML-PROTOCOL-ABORT is issued by the TML if the connection fails and an error is reported by TCP or if the connection has been aborted because of an error reported on the level of the TML.

Table 2-10: Parameters of the Primitive TML-PROTOCOL-ABORT

Parameter	Indication
Diagnostics	X

2.6.4.6 TML-RESET

The primitive TML-RESET is issued by the higher layers to request abortion of the TCP connection without transmitting diagnostics to the peer. It does not carry any parameter. The TML terminates the connection using TCP-ABORT.

2.6.5 TML LISTENER INTERFACE

2.6.5.1 Overview

The TML listener interface supports the two primitives displayed in table 2-11.

Table 2-11: Primitives of the TML Listener Interface

Primitive	Request
TML-START-LISTEN	X
TML-STOP-LISTEN	X

2.6.5.2 TML-START-LISTEN

The primitive TML-START-LISTEN is used by the SPL to request the TML to accept TCP connections on the socket associated with the responder port parameter passed as argument. For this interface the responder port identifier must map to a single socket and the IP address must be valid for the local host.

Table 2-12: Parameters of the Primitive TML-START-LISTEN

Parameter	Request
Responder Port Identifier	X

2.6.5.3 TML-STOP-LISTEN

The primitive TML-STOP-LISTEN is used by the SPL to instruct the TML that no more connections should be accepted on the socket associated with the responder port parameter passed as argument.

Table 2-13: Parameters of the Primitive TML-STOP-LISTEN

Parameter	Request
Responder Port Identifier	X

2.6.6 TCP INTERFACE

2.6.6.1 Introduction

The TCP specification in reference [7] does not define an application program interface. It provides only a functional description of user commands that all TCP implementations must provide. In practice, two different APIs have evolved:

- a) The ‘socket interface’ was originally developed for BSD UNIX. It is the most popular interface and is available on most platforms. Unfortunately there is no commonly agreed specification of the socket interface and implementations vary in a number of subtle details. POSIX has developed a draft standard P3001.1g for the socket interface (reference [E8]), but this standard has not yet been widely supported at the time of writing.
- b) The Transport Layer Interface (TLI) was developed for UNIX System V as a general-purpose transport interface supporting Internet and OSI protocols. X/Open has adopted and extended TLI as the X/Open Transport Interface (XTI). The POSIX standard also includes XTI.

This Recommended Standard does not assume use of any specific TCP API. In order to specify how TCP shall be used, it defines a set of abstract service primitives based on the functional description of user commands in reference [7]. This subsection first provides a summary discussion of these user commands and then a definition of the service primitives used for the specification.

2.6.6.2 TCP User Commands Defined in RFC 793

2.6.6.2.1 General

Reference [7] defines the following user commands for TCP:

- a) OPEN to establish a new connection or listen for incoming connect requests;
- b) SEND to send data;
- c) RECEIVE to read data that have been received by the TCP;
- d) CLOSE to terminate a TCP connection;
- e) ABORT to abort a TCP connection.

In addition, reference [7] requires that TCP provide ‘TCP-to-User messages’ to inform the user of certain events. These messages are not further detailed.

2.6.6.2.2 OPEN

Input: local port number
foreign socket (IP address and port number)
active/passive flag
timeout - optional

Output: local connection name

Reference [7] additionally defines arguments related to security and precedence. These have been omitted, as they are not generally supported by TCP APIs.

If the active/passive flag is set to ‘active’, the TCP attempts to establish a connection to the foreign socket (active open). If the flag is set to ‘passive’, the TCP accepts incoming connection requests (passive open). According to reference [7], a user can specify only the local port number, but TCP APIs allow specification of the IP address as well (the IP address must be a valid address on the host). However, TCP APIs generally do not allow passive open to specify the foreign socket. Therefore, this specification does not consider that option. Whether the timeout option is supported depends on the specific API.

2.6.6.2.3 SEND

Input: local connection name
buffer address
byte count
PUSH flag
URGENT flag
timeout - optional

Output: none

The SEND command copies the number of bytes defined by ‘byte count’ from ‘buffer address’ to the TCP transmit buffer. The description in reference [7] refers to a blocking interface where the command would return only when all data have been copied. APIs supporting a non-blocking interface return the number of bytes that have been copied. To send the remaining data the command must be reissued when the TCP reports it has sufficient room in its transmit buffer.

The PUSH flag instructs the TCP actually to transmit the data as soon as possible. TCP APIs generally do not support specification of the flag by the user. Therefore, it is not further considered in this Recommended Standard.

When the URGENT flag is set, TCP sets the urgent flag in the next segment transmitted indicating to the receiver that urgent data are present. In addition it sets an urgent pointer in the segment header that points to the end of the urgent data.

Whether the timeout option is supported depends on the specific API.

2.6.6.2.4 RECEIVE

Input: local connection name
buffer address
byte count

Output: byte count
PUSH flag
URGENT flag

The RECEIVE command copies a maximum of ‘byte count’ bytes from the TCP receive buffer to the ‘buffer address’ and returns the actual number of bytes copied. If present, it also returns the PUSH flag and the URGENT flag.

2.6.6.2.5 CLOSE

Input: local connection name

Output: none

The CLOSE command terminates transmission of data on a connection and signals to the peer TCP that the connection is closing. Data can still be received on this connection. This is called a ‘half-close’ by reference [7]. The TCP connection terminates when both users have closed the connection.

NOTE – The CLOSE command must not be confused with the close function of the socket interface, which actually performs a ‘full close’.

2.6.6.2.6 ABORT

Input: local connection name

Output: none

When ABORT is issued, the TCP sends a special segment with the flag RST set (RESET segment) and immediately releases all resources allocated to the connection. Reception of a RESET segment causes a hard TCP error.

2.6.6.3 Service Primitives Used in this Recommended Standard

2.6.6.3.1 Overview

This Recommended Standard uses the abstract primitives listed in table 2-14 to describe how the TCP interface is used. These primitives are defined in terms of the TCP user commands in reference [7] as far as possible. Additional specifications are added where necessary.

Table 2-14: Primitives Used for the TCP-Interface

Primitive	Request	Indication	Confirmation
TCP-PASSIVE-OPEN	X		
TCP-CONNECT	X	X	X
TCP-DISCONNECT	X	X	
TCP-DATA	X	X	
TCP-URGENT-DATA	X	X	
TCP-ABORT	X	X	

2.6.6.3.2 TCP-PASSIVE-OPEN

The TCP-PASSIVE-OPEN primitive is used to request TCP to accept connections on a specified local port. The parameters are listed in table 2-15.

Table 2-15: Parameters of the Primitive TCP-PASSIVE-OPEN

Parameter	Request
Local IP address	X
Local port number	X

Request Call to the command OPEN with the active/passive flag set to 'passive'. The local IP address can remain unspecified. In this case, TCP will accept a connection arriving on any interface of the host.

2.6.6.3.3 TCP-CONNECT

The TCP-CONNECT primitive request is used to request establishment of a TCP connection to a specified socket. The indication refers to an indication of TCP that a connection has been accepted.

Table 2-16: Parameters of the Primitive TCP-CONNECT

Parameter	Request	Indication	Confirmation
Foreign IP address	X	X	
Foreign port number	X	X	
Local IP address	X		
Local port number	X		

Request Request of the OPEN command with the active/passive flag set to ‘active’. The foreign IP address and port number are specified. The local port number remains unspecified such that TCP can select an ephemeral port. It is recommended to leave the local IP address and port number unspecified, but supply of these parameters is not excluded.

Indication TCP-to-User message, indicating that a new connection has been accepted by the TCP on a socket for which TCP-PASSIVE-OPEN-request had been issued. TCP APIs generally require the application program to accept the new connection before data can be exchanged.

Confirmation Indication by the TCP that a non-blocking connect request succeeded.

2.6.6.3.4 TCP-DISCONNECT

The TCP-DISCONNECT service request is used to request a half-close of the connection. The indication refers to an indication by TCP that the peer has closed its side of the connection.

Request Request of the command CLOSE.

Indication TCP-to-User message, indicating that the peer has closed its side of the connection. (In the socket API this is indicated by returning a zero byte count from a request to read data from the TCP receive buffer.)

As mentioned earlier in 2.6.6.2, the TCP user command CLOSE is supposed to perform a ‘half close’, which indicates that no more data will be transmitted but still allows receiving data. Standard interfaces provided by TCP APIs to disconnect a TCP connection terminate both transmission and reception of data and perform the TCP connection release procedure autonomously without user interaction. This Recommended Standard has been specifically designed to allow use of such interfaces as well as interfaces that support the original concepts expressed in reference [7] to implement the primitive TCP-DISCONNECT.

2.6.6.3.5 TCP-DATA

The primitive TCP-DATA is used to describe data transfer across a TCP connection.

Table 2-17: Parameters of the Primitive TCP-DATA

Parameter	Request	Indication
Data buffer	X	X

Request Call of the command SEND with the URGENT flag not set, repeating the call if necessary to send a complete message.

Indication TCP-to-User message, indicating that data are present in the TCP receive buffer and subsequent call of the command RECEIVE. For the purpose of this Recommended Standard it is assumed that RECEIVE is called once to read the TML message header followed by one or more calls to read the message body.

2.6.6.3.6 TCP-URGENT-DATA

The primitive TCP-URGENT-DATA is used to describe urgent data transfer across a TCP connection.

Request Call of the command SEND with the URGENT flag set.

Indication TCP-to-User message, indicating that urgent data are present.

2.6.6.3.7 TCP-ABORT

The service request is used to abort a connection. The indication refers to a TCP error report indicating that the connection can no longer be used.

Request Call of the command ABORT.

Indication TCP-to-User message, indicating a hard TCP error.

This Recommended Standard recognizes that some implementations of specific TCP APIs do not support the user command ABORT or do not allow to issue this command in certain states, e.g., after closing of the connection has been requested. In such cases, the primitive TCP-ABORT should be implemented using the most efficient method for termination of the connection that is supported by the TCP API. Such alternative methods might invoke the standard orderly connection release procedure and not transmit a RESET segment, which is considered a conforming behavior.

2.7 SECURITY ASPECTS OF THE INTERNET SLE PROTOCOL

2.7.1 SECURITY BACKGROUND / INTRODUCTION

The SLE Internet Protocol specified in this Recommended Standard provides a communication service for SLE services and therefore any discussion of security aspects must consider the combination of the applicable SLE transfer service specification and the SLE Internet Protocol.

The SLE transfer services are defined as layered application services operating over underlying communication services that must meet certain requirements but which are otherwise unspecified. The SLE Internet Protocol specifies a communications service based on the Transmission Control Protocol (TCP). As such it meets the requirements on the communications service set forth by the SLE Transfer Service specification but does not provide explicit support for security capabilities other than authentication and access control. Such capabilities can however be provided by establishment of a secure communications channel for ISP1 traffic using secure protocols at the transport or network layer and appropriate tunneling techniques. Missions and service providers are expected to select from the available technologies to acquire the performance and security capabilities appropriate to the mission. Specification of the various tunneling technologies, and in particular their associated security provisions, are outside the scope of this Recommended Standard.

SLE transfer services transfer data that originate on a mission spacecraft or are destined for a mission spacecraft. As such, the SLE transfer services have custody of the data for only a portion of the end-to-end data path between MDOS and mission spacecraft. Consequently the ability of an SLE transfer service to secure the transfer of mission spacecraft data is limited to that portion of the end-to-end path that is provided by the SLE transfer service (i.e., the terrestrial link between the MDOS and the ground termination of the ground-space link to the mission spacecraft). End-to-end security must also involve securing the data as it crosses the ground-space link, which can be provided by some combination of securing the mission data itself (e.g., encryption of the mission data within CCSDS space packets) and securing the ground-space link (e.g., encryption of the physical ground-space link). Thus while SLE transfer services play a role in the end-to-end security of the data path, they do not control and cannot ensure that end-to-end security. This component perspective is reflected in the security provisions of the SLE transfer services and of the SLE Internet Protocol.

2.7.2 STATEMENTS OF SECURITY CONCERNS

2.7.2.1 General

This subsection identifies ISP1 support for capabilities that responds to these security concerns in the areas of data privacy, data integrity, authentication, access control, availability of resources, and auditing.

2.7.2.2 Data Privacy (also known as Confidentiality)

This Recommended Standard and the SLE transfer service specifications do not define explicit data privacy requirements or capabilities to ensure data privacy. Data privacy is expected to be ensured by the mission application processes that communicate over the SLE transfer service, by a secure communications channel through which ISP1 traffic is tunneled, or some combination of both. For example, mission application processes might apply end-to-end encryption to the contents of the CCSDS space link data units carried as data by the SLE transfer service. Alternatively or in addition, the network connection between the SLE entities might be encrypted when tunneled through a secure communications channel.

2.7.2.3 Data Integrity

This Recommended Standard and the SLE transfer service specifications do not define explicit data integrity requirements or capabilities to ensure data integrity beyond the protection provided by the TCP. Explicit protection of data integrity, if required, must be provided by a secure communications channel through which ISP1 traffic is tunneled.

2.7.2.4 Authentication

This Recommended Standard requires use of the protected simple authentication procedure (Protected 1) defined in reference [11] for implementation of the authentication requirements specified in the SLE transfer service specifications (see 3.1).

The SLE transfer service authentication capability can be selectively set to authenticate at one of three levels: authenticate every invocation and return, authenticate only the BIND operation invocation and return, or perform no authentication. Depending upon the security environment in which the SLE service user and provider are operating, and the security requirements of the spaceflight mission, the SLE transfer service authentication level can be adapted by choosing the SLE operation invocation and returns that shall be authenticated.

If threat analysis results entail that the simple authentication procedure specified by this Recommended Standard does not provide sufficient protection, then strong authentication must be provided by a secure communications channel through which ISP1 traffic is tunneled.

2.7.2.5 Access Control

SLE transfer service specifications define access control requirements, and define `initiator-identifier` and `responder-identifier` parameters of the service operation invocations and returns that are used to perform SLE transfer service access control. The procedure by which access to SLE transfer services is controlled is described in annex F of the Cross Support Service Green Book (reference [E1]).

2.7.2.6 Availability of Resources

The SLE transfer services are provided via communication networks that have some limit to the resources available to support those SLE transfer services. If these resources can be diverted from their support of the SLE transfer services (in what is commonly known as ‘denial of service’) then the performance of the SLE transfer services may be curtailed or inhibited.

This Recommended Standard specifies use of a start-up timer to prevent unauthenticated and unauthorized users from blocking resources (see 3.3.4.2) but beyond that does not define explicit capabilities to prevent denial of service. Further protection, if required, must be provided by a secure communications channel through which ISP1 traffic is tunneled.

2.7.2.7 Auditing

This SLE Internet Protocol specification and the SLE transfer service specifications do not define explicit security auditing requirements or capabilities. Security auditing is expected to be negotiated and implemented bilaterally between the spaceflight mission and the service provider.

2.7.3 POTENTIAL THREATS AND ATTACK SCENARIOS

The SLE Internet Protocol and SLE transfer services depend on unspecified mechanisms operating above the SLE transfer service (between a mission spacecraft application process and its peer application process on the ground), in a secure communications channel through which ISP1 traffic is tunneled, or some combination of both, to ensure data privacy (confidentiality). If no such mechanisms are actually implemented, or the mechanisms selected are inadequate or inappropriate to the network environment in which the mission is operating, an attacker could read the spacecraft telemetry or command data contained in the SLE protocol data units as they traverse the WAN between service user and service provider.

The SLE Internet Protocol and SLE transfer services constrain the ability of a third party to seize control of an active SLE transfer service instance, but do not specify mechanisms that would prevent an attacker from intercepting the protocol data units and replacing the contents of the `data` parameter. The prevention of such a replacement attack depends on unspecified mechanisms operating above the SLE transfer service (between a mission spacecraft application process and its peer application process on the ground), in a secure communications channel through which ISP1 traffic is tunneled, in bilaterally-agreed extra capabilities applied to the SLE transfer service (e.g., encryption of the `data` parameter) or some combination of the three. If no such mechanisms are actually implemented, or the mechanisms selected are inadequate or inappropriate to the network environment in which the mission is operating, an attacker could substitute telemetry data without detection.

If the SLE transfer service authentication capability is not used or if the implementation of the SLE authentication capability by the simple authentication procedure specified in this

Recommended Standard is inadequate or inappropriate to the network environment in which the mission is operating, and if authentication is not ensured by a secure communications channel through which ISP1 traffic is tunneled, then attackers may somehow obtain valid `initiator-identifier` values and use them to initiate SLE transfer service instances by which they could gain access to spacecraft commanding or to telemetry data.

The SLE Internet Protocol provides basic protection against denial of service attacks in order to ensure that the supporting network has sufficient resources to provide sufficient support to legitimate users. If the mechanisms provided for this purpose are inadequate or inappropriate to the network environment in which the mission is operating and no additional protection is provided by a secure communications channel through which ISP1 traffic is tunneled, then an attacker could prevent legitimate users from receiving telemetry from their spacecraft or from commanding their spacecraft.

If the provider of SLE transfer service provides no security auditing capabilities or if a user chooses not to employ auditing capabilities that do exist, then attackers may delay or escape detection long enough to do serious (or increasingly serious) harm to the mission.

2.7.4 CONSEQUENCES OF NOT APPLYING SECURITY

The consequences of not applying security to an SLE transfer service are possible degradation and loss of ability to receive telemetry from the spacecraft, substitution of altered telemetry data, degradation and loss of ability to command the spacecraft, and even loss of the spacecraft itself.

3 ISP1 MESSAGES AND PROCEDURES

3.1 AUTHENTICATION LAYER

3.1.1 Authentication shall be based on the following parameters:

- a) User name. The following rules shall apply for this parameter:
 - 1) the user name shall be a string of 3 to 16 characters;
 - 2) the user name shall be identical to the authority identifier of the application by which the application is identified in the BIND invocation and the BIND return.

NOTE – For an initiating SLE application, the user name is the initiator-identifier, and for a responding SLE application the user name is the responder-identifier.

- b) Password. The password shall be an octet string of 6 to 16 octets.

3.1.2 When the Authentication Layer receives an SLE PDU requiring authentication from the higher layers it shall generate the credentials and insert them into the PDU passed to the Data Encoding Layer.

NOTE – Whether a PDU requires authentication is determined by the authentication level agreed upon for a service instance; it can be

- ‘all’ — all PDUs except PEER-ABORT-invocation shall be authenticated;
- ‘bind’ — only the BIND invocation and return shall be authenticated;
- ‘none’ — no PDU shall be authenticated.

3.1.2.1 Generation of credentials shall be performed according to the protected simple authentication procedure (Protected 1) defined in reference [11] and detailed by the following specifications.

3.1.2.1.1 The following information shall be encoded using the ASN.1 syntax defined in reference [9] and the Distinguished Encoding Rules (DER) specified in reference [10]:

- a) the current time, using the CCSDS day segmented time code without the P-field (see reference [13]);
- b) a random number;
- c) the user name of the local SLE application;
- d) the password of the local SLE application.

NOTE – Encoding the information with DER provides a platform-independent bit pattern from which a hash code can be generated. Use of ASN.1 and DER for generation of credentials does not imply that ASN.1 or DER is used for encoding of data exchanged between the service user and the service provider.

3.1.2.1.2 The ASN.1 type used for encoding shall be the one defined in figure 3-1.

```

HashInput ::= SEQUENCE
{
  time          OCTET STRING (SIZE(8))
  , randomNumber INTEGER (0 .. 2147483647)
  , userName    VisibleString
  , password    OCTET STRING
}

```

Figure 3-1: ASN.1 Type for Generation of ‘the Protected’

3.1.2.1.3 The output of the encoder shall be passed through a one-way hash function to obtain a message digest, also referred to as ‘the protected’.

3.1.2.1.4 The following information shall be inserted into the parameter credentials of the PDU before passing it to the Data Encoding Layer:

- a) the time parameter used for generation of the message digest;
- b) the random number used for generation of the message digest;
- c) the message digest itself.

NOTE – The user name used for authentication is transmitted as part of the BIND invocation or return. Therefore it is not necessary to include it into the credentials parameter.

3.1.2.2 Authentication of credentials received from the peer application shall be performed according to the protected simple authentication procedure (Protected 1) defined in reference [11] and detailed by the following specifications.

3.1.2.2.1 The time in the credentials shall be checked against the current time. If the time difference is larger than acceptable, authentication shall fail.

NOTE – This Recommended Standard assumes that the acceptable delay is provided to the Authentication Layer as a configuration parameter.

3.1.2.2.2 The following information shall be encoded using the ASN.1 type defined in figure 3-1 and the Distinguished Encoding Rules:

- a) the time obtained from the credentials, in the CCSDS day segment time format without P-Field (see reference [13]);
- b) the random number obtained from the credentials;

SPACE LINK EXTENSION—INTERNET PROTOCOL FOR TRANSFER SERVICES

- c) the user name of the peer application;
- d) the password of the peer application.

3.1.2.2.3 The encoded data shall be passed through a one-way hash function to obtain a message digest.

3.1.2.2.4 The message digest shall be compared with the message digest in the credentials. If these match, authentication shall be regarded as successful. Otherwise, authentication shall fail.

3.1.2.3 The one-way hash function used shall be SHA-256 defined by reference [12].

3.2 DATA ENCODING LAYER

3.2.1 The DEL shall encode and decode SLE PDUs using the ASN.1 types defined in the Recommended Standard for the applicable transfer service type and version number and the Basic Encoding Rules defined in reference [10].

NOTE – The service type and version number can be extracted from the BIND invocation PDU, which must be the first PDU passed to the DEL after establishment of an ISP1 association.

3.2.2 If decoding of a PDU fails, the DEL shall inform the higher layers, which shall apply the procedures foreseen by the Recommended Standards for SLE transfer services.

3.2.3 The DEL shall encode the credentials parameter defined for SLE PDUs using the ASN.1 type shown in figure 3-2. It shall encode the resulting bytes as an OCTET STRING according to the definitions in the Recommended Standards for SLE transfer services.

```

ISP1Credentials ::= SEQUENCE
{
  time          OCTET STRING (SIZE (8)) -- CCSDS CDS time code
  , randomNumber INTEGER (0 .. 2147483647)
  , theProtected OCTET STRING (SIZE (20..32))
}

```

Figure 3-2: ASN.1 Type for the Credentials Parameter

NOTES

- 1 For the purpose of authentication, the constituents of the credentials must be distinguishable for the receiver. This implies the need for a refined ASN.1 type.
- 2 The length of the ‘the protected’ field is sized to a maximum of 32 bytes to hold the output SHA-256 hash function. However, to allow backward compatible implementations, the ‘the protected’ field can also be 20 bytes long, which is the

length of the output generated by the SHA-1 function specified by the previous version of this Recommended Standard (reference [E10]).

3.3 TRANSPORT MAPPING LAYER

3.3.1 TECHNOLOGY

The TML shall use the Transmission Control Protocol (TCP) specified by reference [7] for communication with a remote TML. It shall map one SLE association to exactly one TCP connection.

3.3.2 TML MESSAGES

3.3.2.1 The TML shall use three types of messages for exchange of data over an established TCP connection:

- a) an SLE PDU message for transfer of SLE PDUs;
- b) a TML context message to transmit TML initialization parameters;
- c) a TML heartbeat message to periodically probe an idle TCP connection.

3.3.2.2 TML messages shall be formatted as defined by the specifications in 3.3.2.2.1 to 3.3.2.2.6.

NOTES

- 1 The layout of a TML message is shown in figure 3-3.
- 2 The message body contains any number of octets, not necessarily a multiple of four.




Figure 3-3: Layout of a TML Message

3.3.2.2.1 A TML message consists of an eight-octet message header and a message body.

3.3.2.2.2 The message header consists of the following fields in the sequence indicated:

- a) a message type field of one octet containing the message type identifier as defined by table 3-1;

Table 3-1: TML Message Type Identifiers

Type ID	Message type
1	SLE PDU Message
2	TML Context Message
3	TML Heartbeat Message

- b) a reserved field of 3 octets set to 'zero';
- c) a length field of four octets, which holds the number of octets in the message body represented as a binary unsigned integer value.

3.3.2.2.3 The body of an SLE PDU Message consists of the encoded SLE PDU.

3.3.2.2.4 The body of a TML context message consists of 12 octets and contains the following fields in the sequence indicated:

- a) the protocol identification field (four octets) containing the characters 'I' 'S' 'P' '1' in ASCII encoding (hexadecimal 49 53 50 31);
- b) a reserved field of 3 octets, set to 'zero';
- c) the protocol version (1 octet) set to the binary value '1';
- d) the heartbeat-interval (2 octets) represented as a binary unsigned integer value;
- e) the dead-factor (2 octets) represented as a binary unsigned integer value.

NOTE – The parameters heartbeat interval and dead-factor are defined in 3.3.3. The layout of the TML context message is shown in figure 3-4.

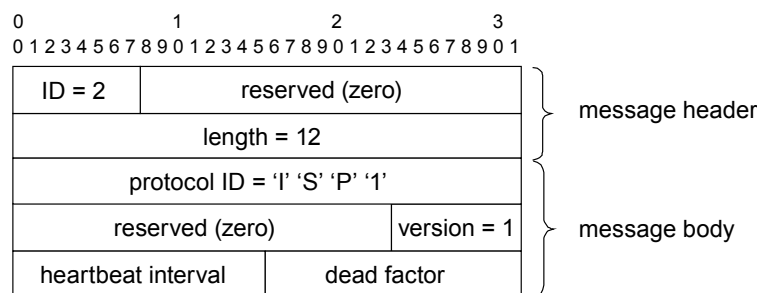


Figure 3-4: Layout of a TML Context Message

3.3.2.2.5 The body of a TML heartbeat message is empty.

NOTE – A TML heartbeat message consists of a message header only, where the length field is set to 'zero'.

3.3.2.2.6 All integer values in TML messages shall be in network octet order ('big endian').

NOTE – Because the interface to TCP is defined in terms of bytes, the TML must ensure that the byte order of integer values is correct. The order in which the bits of individual bytes are transmitted depends on lower-level protocols and hardware interfaces and cannot be influenced by the application software. If there is a need to reverse the bit order within a byte, it is usually performed either by the hardware or low-level device drivers.

3.3.3 HEARTBEAT MECHANISM

NOTE – The TML uses the procedure specified in this section to detect failure of a TCP connection during periods in which no messages are transmitted. This procedure is described first, because the parameters and timers defined by the procedure are frequently referred to by the specification of connection establishment and release.

3.3.3.1 Detection of failed connections shall be based on the following parameters:

- a) the heartbeat-interval specifies the maximum period, measured in seconds, in which a TML might not transmit any message;
- b) the dead-factor defines the time, measured in units of heartbeat-intervals, after which the TML shall assume that a connection has failed, if it does not receive any message from the peer TML.

3.3.3.2 If the heartbeat-interval is set to 'zero' the procedure defined in this section shall not be applied.

3.3.3.3 The TML shall use a heartbeat transmit-timer (HBT timer), set to the value of the heartbeat-interval according to the following specifications:

3.3.3.3.1 Whenever the TML transmits a message, it shall restart the HBT timer.

3.3.3.3.2 When the HBT timer expires, the TML shall transmit a TML heartbeat message and restart the HBT timer.

3.3.3.4 The TML shall use a heartbeat receive-timer (HBR timer), set to the value of the heartbeat-interval multiplied by the value of the dead-factor, according to the following specifications:

3.3.3.4.1 Whenever the TML receives a message, it shall restart the HBR timer.

3.3.3.4.2 When the HBR timer expires, the TML shall assume that the connection has failed. It shall issue TCP-ABORT and inform the higher layers by TML-PROTOCOL-ABORT-indication with the TML diagnostic set to 'heartbeat timeout'.

3.3.4 CONNECTION ESTABLISHMENT

3.3.4.1 Initiator Side

NOTE – This subsection specifies the connection establishment procedure used by the TML when a logical port identifier maps to a single TCP socket. The procedure for connecting to multiple hosts concurrently is defined in 3.3.8.

3.3.4.1.1 When receiving TML-CONNECT-request, the TML shall establish a TCP connection as defined by the specifications in 3.3.4.1.1.1 to 3.3.4.1.1.3.

3.3.4.1.1.1 The TML shall issue TCP-CONNECT-request using the IP address and port number derived from the logical port identifier to specify the foreign IP address and port number.

NOTE – TCP APIs generally allow the local address and port number to remain unspecified. They will then automatically use an IP address by which the socket can be reached and assign an ephemeral port number.

3.3.4.1.1.2 If TCP-CONNECT-request succeeds, the TML shall:

- a) start the HBT timer and the HBR timer; and
- b) issue TML-CONNECT-confirmation.

NOTE – The TML context message is defined in 3.3.2. The heartbeat mechanism is defined in 3.3.3. If the heartbeat interval is ‘zero’, the heartbeat timers are not started.

3.3.4.1.1.3 If TCP-CONNECT-request fails, or any other error occurs, the TML shall issue TML-PROTOCOL-ABORT-indication with the appropriate diagnostics.

NOTE – This Recommended Standard assumes that the TCP implementation takes care of temporary network problems and re-transmissions and that the TCP implementation can be configured as needed. Therefore, failure of TCP-CONNECT-request is assumed to be a hard error. This Recommended Standard does not exclude that an implementation retries connection establishment for a configurable period or until it receives TML-RESET-request because the return timer expires.

3.3.4.2 Responder Side

3.3.4.2.1 When receiving TML-START-LISTEN-request, the TML shall request TCP-PASSIVE-OPEN using the IP address and port number derived from the logical responder port identifier to specify the local address and port number.

3.3.4.2.2 When receiving TML-STOP-LISTEN-request, the TML shall instruct TCP to stop accepting connection requests for the IP address and port number identified by the argument ‘responder port identifier’.

NOTE – This effect can generally be achieved by closing the listen socket (Socket API) or transport endpoint (XTI/TLI).

3.3.4.2.3 When receiving TCP-CONNECT-indication, the TML shall create a new association in the state ‘TML starting’ and perform the following procedure:

3.3.4.2.3.1 It shall start a TML start-up timer (TMS timer).

3.3.4.2.3.2 If the TML receives at least the amount of data specified for a TML message header before the TMS timer expires, it shall read the first eight octets of the data and perform the following checks:

3.3.4.2.3.2.1 If the first eight octets of the data do not present a valid TML message header, the TML shall stop the TMS timer and abort the connection using TCP-ABORT.

NOTE – As explained in 2.6.6.3.7, this specification allows TCP-ABORT to have the same effect as TCP-DISCONNECT-request, if the TCP user command ABORT is not supported by the TCP API.

3.3.4.2.3.2.2 If the message received is a TML context message, the TML shall retrieve the complete message and perform the checks in 3.3.4.2.3.2.3 to 3.3.4.2.3.2.5.

NOTE – The TML context message is defined in 3.3.2.

3.3.4.2.3.2.3 If the protocol identifier is not correct, the TML shall stop the TMS timer and abort the connection using TCP-ABORT.

3.3.4.2.3.2.4 If the version number does not match the version of this specification supported by the TML, the TML shall stop the TMS timer and abort the connection using TCP-ABORT.

3.3.4.2.3.2.5 If the heartbeat interval and dead-factor values are acceptable, the TML shall start the HBT timer.

NOTES

1 The heartbeat interval and dead-factor are explained in 3.3.3.

2 If the heartbeat interval is ‘zero’, the heartbeat transmit timer shall not be not started.

3.3.4.2.3.3 When all checks have been passed, the TML shall change the state of the association to ‘data transfer’ and issue TML-CONNECT-indication.

3.3.4.2.3.4 If the first message received is a valid TML message other than the TML context message, the TML shall stop the TMS timer and abort the connection using TCP-ABORT.

3.3.4.2.3.5 When receiving the first valid SLE PDU message in the state ‘data transfer’, the TML shall stop the TMS timer and start the HBR timer.

NOTE – If the heartbeat interval is ‘zero’, the heartbeat receive timer shall not be started.

3.3.4.2.3.6 If the TMS timer expires before a valid context message and the first valid SLE PDU message arrive, the TML shall terminate the connection using TCP-ABORT.

3.3.4.2.3.7 If the TMS timer expires in the state ‘data transfer’, i.e., after reception of the context message, but before reception of a valid SLE PDU message, the TML shall additionally inform the higher layers using TML-PROTOCOL-ABORT indication.

3.3.5 ORDERLY CONNECTION RELEASE

3.3.5.1 When receiving TML-DISCONNECT-request, an initiating TML shall close the TCP connection.

NOTE – If the peer TML does not close its side of the connection, actual release of resources allocated to the TCP socket depends on the policy applied by the TCP implementation. Implementations can ensure release of resources using the following procedure if it is supported by the TCP API and the TCP implementation:

- a) perform a half-close for the sending side and start a timer;
- b) if any data arrive, abort the connection using TCP-ABORT;
- c) if the peer does not close the connection within the timeout, abort the connection using TCP-ABORT.

3.3.5.2 When receiving TML-DISCONNECT-request, a responding TML shall perform the steps described in 3.3.5.2.1 to 3.3.5.2.4:

3.3.5.2.1 The TML shall stop the HBT timer, restart the HBR timer, and wait for TCP-DISCONNECT-indication.

NOTE – If the heartbeat interval is ‘zero’, the heartbeat mechanism shall be disabled. In this case, no timers shall be used for the connection release procedure.

3.3.5.2.2 The TML shall abort the TCP connection using TCP-ABORT-request if it receives further data.

NOTE – No data can arrive if both sides correctly implement the SLE operations as defined by the applicable Recommended Standard.

3.3.5.2.3 When receiving TCP-DISCONNECT-indication, the TML shall close its side of the connection using TCP-DISCONNECT-request.

3.3.5.2.4 If the HBR timer expires before the TML receives TCP-DISCONNECT-indication, it shall abort the connection with TCP-ABORT-request.

3.3.6 ABORTING CONNECTIONS

3.3.6.1 Peer Abort

3.3.6.1.1 When receiving TML-PEER-ABORT-request, the TML shall perform the procedure specified by 3.3.6.1.3.

3.3.6.1.2 Except when explicitly defined otherwise in this Recommended Standard, the TML shall use the peer abort procedure defined by 3.3.6.1.3 to report errors to the peer TML, using one of the TML diagnostic codes defined in annex A. In this case, it shall issue TML-PROTOCOL-ABORT-indication with the appropriate diagnostic to inform the higher layers that it has aborted the connection.

3.3.6.1.3 The TML shall implement the PEER-ABORT operation defined by the Recommended Standards for SLE transfer services as defined by the specifications in 3.3.6.1.3.1 to 3.3.6.1.3.3:

3.3.6.1.3.1 The aborting TML shall perform the following steps:

- a) the TML shall send the PEER-ABORT diagnostics as one byte of urgent data using TCP-URGENT-DATA-request;
- b) the TML shall stop the HBT and HBR timers, and start the CPA timer;
- c) the TML shall wait for TCP-DISCONNECT-indication;
- d) the TML shall silently discard all data that arrive after it has sent the urgent data;
- e) when receiving TCP-DISCONNECT-indication, the TML shall close its side of the connection using TCP-DISCONNECT-request;
- f) if the CPA timer expires before the TML receives TCP-DISCONNECT-indication, the TML shall abort the connection with TCP-ABORT-request.

NOTE – As explained in 2.6.6.3.7, this specification allows TCP-ABORT to have the same effect as TCP-DISCONNECT-request, if the TCP user command ABORT is not supported by the TCP API.

3.3.6.1.3.2 When the TML is notified by TCP-URGENT-DATA-indication that urgent data are pending, it shall perform the following steps:

- a) the TML shall read and discard all data from the TCP receive buffer up to but excluding the byte referred to by the TCP urgent pointer;
- b) the TML shall read and analyze the byte referenced by the TCP urgent pointer;
- c) if the diagnostic code is in the range reserved by the Recommended Standards for SLE transfer services, the TML shall issue the primitive TML-PEER-ABORT-indication passing the diagnostic as a parameter;

- d) otherwise the TML shall issue TML-PROTOCOL-ABORT-indication primitive with the diagnostics retrieved from the urgent byte;
- e) the TML shall close the TCP connection.

NOTES

- 1 When the receiving TML has been notified of the presence of urgent data, that TML might receive requests to transfer data from the DEL. It is expected that such requests will either be rejected with an indication that the connection is aborting or will be silently ignored. Requests to disconnect or abort the connection must be accepted, but should be ignored. The only exception is a TML-RESET request, which should cause an immediate TCP-ABORT.
- 2 If the peer TML does not close its side of the connection, actual release of resources allocated to the TCP socket depends on the policy applied by the TCP implementation. Implementations can ensure release of resources using the following procedure if it is supported by the TCP API and the TCP implementation:
 - a) perform a half-close for the sending side and start a timer;
 - b) if any data arrive, abort the connection using TCP-ABORT;
 - c) if the peer does not close the connection within the timeout, abort the connection using TCP-ABORT.

3.3.6.1.3.3 If the aborting TML is notified by TCP-URGENT-DATA-indication that urgent data are pending after transmission of the PEER-ABORT diagnostics, it shall read and discard all data up to and including the byte referred to by the TCP urgent pointer and then close the TCP connection.

NOTE – This situation can occur if both sides request PEER-ABORT nearly simultaneously. Both sides must close the TCP connection in order to avoid a deadlock situation where each side is waiting for the other side to close the connection. In this situation, neither side will be transmitting data following the PEER-ABORT diagnostic, such that a ‘full close’ can be issued without invoking TCP ABORT. Because the PEER-ABORT procedure was already invoked locally, the PEER-ABORT diagnostic received will not be passed to the higher layers.

3.3.6.2 Silent Abort

When receiving TML-RESET-request, the TML shall abort the connection with TCP-ABORT-request.

NOTE – As explained in 2.6.6.3.7, this Recommended Standard allows TCP-ABORT to have the same effect as TCP-DISCONNECT-request, if the TCP user command ABORT is not supported by the TCP API.

3.3.6.3 Protocol Abort

The TML shall issue TML-PROTOCOL-ABORT-indication in the following cases:

- a) TCP reports an unrecoverable error;
- b) TCP reports a send timeout;
- c) TCP reports that the peer has closed its side of the connection by TCP-DISCONNECT-indication before TML-DISCONNECT-request or TCP-URGENT-DATA-indication has been received.

NOTE – As indicated in 3.3.6.1 the TML additionally issues TML-PROTOCOL-ABORT-indication when receiving peer abort with a diagnostic code in the range defined by annex A.

3.3.7 SLE PDU TRANSFER

3.3.7.1 When receiving TML-SLE-PDU-request, the TML shall add the message header for the SLE-PDU message and transmit the message using TCP-DATA-request.

3.3.7.2 When receiving TCP-DATA-indication, the TML shall check the first eight octets of the data and proceed as defined in 3.3.7.2.1 to 3.3.7.2.4.

NOTE – The very first message received on a TCP connection by a responding TML shall be handled differently, as specified in 3.3.4.2.

3.3.7.2.1 If the first eight octets of the data do not present a valid TML message header, the TML shall abort the connection with the peer abort procedure and the diagnostics ‘badly formatted TML message’.

3.3.7.2.2 If the message is a TML heartbeat message, the TML shall restart the HBR timer.

3.3.7.2.3 If the message is an SLE PDU message, the TML shall retrieve the message body from the TCP, forward it using the primitive TML-SLE-PDU-indication, and restart the HBR timer.

3.3.7.2.4 If the message is a TML context message, the TML shall abort the connection with the peer abort procedure and the diagnostics ‘TML protocol error’.

3.3.8 CONNECTION ESTABLISHMENT TO REDUNDANT HOSTS

3.3.8.1 If an initiating TML receives TML-CONNECT-request with a port identifier that maps to a set of sockets, it shall perform the modified connection establishment procedure defined in 3.3.8.1.1 to 3.3.8.1.3.

3.3.8.1.1 The TML shall call TCP-CONNECT-request for each of the sockets in the set.

3.3.8.1.2 When the TCP connection establishment succeeds for one of the sockets, the TML shall immediately call TCP-ABORT on all other sockets in the set. It shall then proceed with the connection establishment as specified by 3.3.4.1.1.2, using the socket for which TCP connection establishment succeeded.

3.3.8.1.3 When TCP connection fails for all sockets in the set, the TML shall issue TML-PROTOCOL-ABORT-indication with the appropriate diagnostics.

4 TML STATE TABLE

4.1 INTRODUCTION

This section defines state tables for the Transport Mapping Layer. The state tables for the Data Encoding Layer and the Authentication Layer are trivial and are therefore not detailed in this document.

The state table reflects the baseline requirements specified in section 3 and does not cover implementation considerations discussed in additional comments on these requirements. It specifies connection establishment with a single provider and does not cover the procedure for concurrent connection establishment with redundant provider systems described in 3.3.8.

4.2 NOTATION

NOTE – The notation used for the state tables is the one specified by UML for state diagrams (see reference [E9]). This notation has been slightly extended to adapt it to state tables. It is summarized below in Extended Backus-Naur Form (EBNF) notation together with the extensions. Extensions are highlighted by underlining. For formulation of conditions, the Object Constraint Language (OCL) specified by UML is used.

An incoming event in the event column is defined by

<origin> ‘.’ <event-name> [‘(’ <arguments> ‘)’]

Processing of the event is described by the following sequence:

```
[<guard-condition>] [<action-expression>]* [<send-clause>]* [<state-transition>]
<guard-condition> ::= [ ‘ ’ <condition> ‘ ’ ]
<condition> ::= conditional expression formulated in OCL
<action-expression> ::= ‘/’ <action-name> [ ‘(’ <arguments> ‘)’ ]
<send-clause> ::= ‘^’ <target> ‘.’ <event-name> [ ‘(’ <arguments> ‘)’ ]
<state-transition> ::= ‘→’ <new-state>
```

Transition to self is not shown in the tables.

Actions can be simple actions or compound actions. Compound actions are displayed in capital letters and are expanded using simple pseudo-code (IF THEN ELSE END IF) together with the notational elements shown above.

For events passed between the layers defined in this Recommended Standard, the names of the service primitives as defined in 2.6 are used. Because of the limited space, the use of a primitive is abbreviated as shown below:

- request req
- response rsp

- indication ind
- confirmation cnf

State / event combinations that cannot occur when the TML is correctly implemented and the higher layers correctly implement the SLE transfer service are marked by ‘N/A’.

4.3 STATES

The state table uses the following states for an ISP1 association:

- | | | |
|----|---------------|--|
| S0 | CLOSED | The initial and final state, in which no TCP connection exists. |
| S1 | TML STARTING | Initialization state for TML; depending on the role (initiator or responder) the TML establishes the TCP connection, or waits for the initial message on an accepted TCP connection. |
| S2 | DATA TRANSFER | SLE PDUs can be exchanged. |
| S3 | PEER ABORTING | The TML has received an indication that urgent data are present and waits for the PEER-ABORT diagnostic to arrive. |
| S4 | TML CLOSING | The TML releases the TCP connection in response to a PEER-ABORT or a DISCONNECT request from the higher layers. |

4.4 EVENTS

4.4.1 EVENTS RECEIVED FROM THE HIGHER LAYERS

The events received from the higher layers are the service primitives defined in 2.6.4 and 2.6.5, namely:

- CONNECT-request
- DISCONNECT-request
- PEER-ABORT-request
- RESET-request

4.4.2 EVENTS SENT TO THE HIGHER LAYERS

The events sent to the higher layers are the service primitives defined in 2.6.4 and 2.6.5, namely:

- CONNECT-indication
- CONNECT-confirmation
- PEER-ABORT-indication
- PROTOCOL-ABORT-indication

4.4.3 EVENTS RECEIVED FROM THE DATA ENCODING LAYER

The event received from the DEL is the service primitive defined in 2.6.3, namely

- SLE-PDU-request

4.4.4 EVENTS SENT TO THE DATA ENCODING LAYER

The event sent to the DEL is the service primitive defined in 2.6.3, namely

- SLE-PDU-indication

4.4.5 EVENTS RECEIVED FROM THE TCP

The events received from the TCP are the service primitives defined in 2.6.6.3, namely

- CONNECT-indication
- CONNECT-confirmation
- DISCONNECT-indication
- DATA-indication
- URGENT-DATA-indication (signal that urgent data are present in the incoming data stream)
- ABORT-indication

In addition, ‘TCP Error’ and ‘Timeout’ are considered events that are received from the TCP. ‘Timeout’ refers to either a connect timeout or send timeout indicated by the TCP.

4.4.6 EVENTS SENT TO THE TCP

The events sent to the TCP are the service primitives defined in 2.6.6.3, namely

- CONNECT-request
- DISCONNECT-request
- DATA-request
- URGENT-DATA-request
- ABORT-request

4.4.7 INTERNAL EVENTS

TMS Timeout	the TML start-up timer expires
HBT Timeout	the heartbeat-transmit-timer expires
HBR Timeout	the heartbeat-receive-timer expires
CPA Timeout	the ‘Close after Peer Abort’ timer expires

4.5 PREDICATES

diagnostics = SLE diagnostics	The diagnostic code received for PEER-ABORT is in the range reserved by the Recommended Standards for SLE transfer services.
diagnostics = TML diagnostics	The diagnostic code received for PEER-ABORT is in the range defined in annex A of this Recommended Standard.
first PDU	The PDU is the first one received on the association.
heartbeat acceptable	The parameters received with the TML context message are acceptable to the receiver.
local peer abort	The PEER-ABORT procedure was requested by the local application.
msg = TML Context	The message received is a correctly formatted TML context message.
msg <> TML Context	The message received is not a TML context message.
msg = TML PDU	The message received is a PDU message.
msg = TML Heartbeat	The message received is a TML heartbeat message.
msg <> TML Message	The message received is not a valid TML message.
role = initiator	The TML is in the role of the initiator; i.e., it initializes the TCP connection.
role = responder	The TML is in the role of the responder; i.e., it accepts the TCP connection.
urgent byte available	The urgent byte was read from the TCP interface.
protocol supported	The protocol identifier is 'ISP1' and the version number is supported by the TML.

4.6 ACTIONS

4.6.1 SIMPLE ACTIONS

/build TML message	Add the TML message header to the encoded PDU.
/cleanup	Stop HBR timer and HBT timer.
/discard	Silently discard the data received.
/discard pending data	If unsent data exist discard them, even if part of a TML message has already been transmitted.
/discard normal data	Discard all received data except for the urgent byte.
/extract PDU	Remove the TML message header from the encoded PDU.
/reject	Reject the request received from higher layers.
/restart HBR timer	Stop the heartbeat-receive-timer and start it again.
/restart HBT timer	Stop the heartbeat-transmit-timer and start it again.
/start HBR timer	Start the heartbeat-receive-timer.
/start HBT timer	Start the heartbeat-transmit-timer.
/start TMS timer	Start the TML start-up timer.
/stop HBR timer	Stop the heartbeat-receive-timer.
/stop HBT timer	Stop the heartbeat-transmit-timer.
/stop TMS timer	Stop the TML start-up timer.

4.6.2 COMPOUND ACTIONS

/PEER ABORT(diagnostic) is defined as

```
^TCP.URGENT-DATA-req(diagnostic)
/discard pending data
/set local peer abort = true
/stop HBT timer
/stop HBR timer
/start CPA timer
```

4.7 STATE TABLE

	S0 - CLOSED	S1 - TML STARTING	S2 - DATA TRANSFER	S3 - PEER ABORTING	S4 - TML CLOSING
HL: CONNECTreq	^TCP.CONNECTreq → S1 (N/A for the responder)	N/A	N/A	N/A	N/A
TCP: CONNECTcnf	N/A	^TCP.DATAreq(TML-Context) /start HBR-Timer /start HBT-Timer ^HL.ConnectConf → S2 (N/A for the responder)	N/A	N/A	N/A
TCP: CONNECTind	/start TMS-timer /set first PDU = false → S1 (N/A for the initiator)	N/A	N/A	N/A	N/A

	S0 - CLOSED	S1 - TML STARTING	S2 - DATA TRANSFER	S3 - PEER ABORTING	S4 - TML CLOSING
TCP: DATAind(msg)	N/A	[role = initiator] ^TCP.ABORTreq ^HL.PROTOCOL-ABORTind → S0 [role = responder] [msg = TML Context] [not protocol supported] /stop TMS timer ^TCP.ABORTreq → S0 [not heartbeat acceptable] /stop TMS timer /PEER ABORT (heartbeat values not acceptable) → S4 /start HBT-Timer /set first PDU = true ^HL.CONNECTind → S2 [msg <> TML Context] /stop TMS timer ^TCP.ABORTreq → S0	[msg = TML-PDU] [role = responder] [not first PDU] /restart HBR timer [first PDU] /stop TMS timer /start HBR Timer /set first PDU = false [role = initiator] /restart HBR timer /extract PDU ^DEL.SLE-PDUind(PDU) [msg = TML-Heartbeat] restart HBR-Timer [msg = TML-Context] [role=responder and first PDU] /stop TMS timer ^HL.PROTOCOL-ABORTind /PEER ABORT(protocol err) → S4 [msg <> TML message] ^HL.PROTOCOL-ABORTind [role=responder and first PDU] /stop TMS timer /PEER ABORT (badly formatted message) → S4	/discard normal data [urgent byte available] [not local peer abort] [diagnostics = SLE diagnostics] ^HL.PEER-ABORTind [diagnostics = TML diagnostics] ^HL.PROTOCOL-ABORTind ^TCP.DISCONNECT-req → S0	[not local peer abort] ^TCP.ABORTreq /stop HBR Timer → S0 [local peer abort] /discard
HL: DISCONNECT req	N/A	N/A	[role = initiator] ^TCP.DISCONNECTreq /cleanup → S0 [role = responder] /set local peer abort = false /stop HBT-timer /restart HBR timer → S4	/reject(aborting)	N/A

	S0 - CLOSED	S1 - TML STARTING	S2 - DATA TRANSFER	S3 - PEER ABORTING	S4 - TML CLOSING
TCP: DISCONNECT ind	N/A	^TCP.DISCONNECTreq /stop TMS-Timer → S0 (N/A for the initiator)	^TCP.DISCONNECTreq ^HL.PROTOCOL-ABORTind /cleanup → S0	^TCP.DISCONNECTreq [not local peer abort] ^HL.PROTOCOL- ABORTind /stop HBR-Timer → S0	^TCP.DISCONNECTreq /stop HBR-Timer → S0
DEL: SLE-PDUreq	N/A	N/A	/build TML message ^TCP.DATAreq(message) /restart HBT-Timer	/reject(aborting)	N/A
HL: PEER- ABORTreq (diagnostics)	N/A	N/A	/PEER ABORT (diagnostics) → S4	/set local peer abort = true	N/A
TCP: URGENT- DATAind	N/A	^TCP.ABORTreq /stop TMS-Timer → S0 (N/A for the initiator)	/set local peer abort = false → S3	N/A	[not local peer abort] ^TCP.ABORTreq /stop HBR-Timer → S0 [local peer abort] → S3
HL: RESETreq	N/A	^TCP.ABORTreq → S0 (N/A for the responder)	^TCP.ABORTreq /cleanup → S0	^TCP.ABORTreq /stop HBR-Timer → S0	N/A
TCP: ABORTind	N/A	[role = initiator] ^HL.PROTOCOL-ABORTind → S0	^HL.PROTOCOL-ABORTind /cleanup → S0	[not local peer abort] ^HL.PROTOCOL- ABORTind /stop HBR-Timer → S0	/stop HBR-Timer → S0
TCP: Timeout	N/A	^TCP.ABORTreq ^HL. PROTOCOL-ABORTind → S0 (N/A for the responder)	^TCP.ABORTreq ^HL. PROTOCOL-ABORTind /cleanup → S0	N/A	N/A

	S0 - CLOSED	S1 - TML STARTING	S2 - DATA TRANSFER	S3 - PEER ABORTING	S4 - TML CLOSING
TCP: Error	N/A	[role = initiator] ^HL.PROTOCOL- ABORTind [role = responder] /stop TMS-Timer → S0	^HL.PROTOCOL-ABORTind /cleanup → S0	[not local peer abort] ^HL.PROTOCOL- ABORTind /stop HBR-Timer → S0	/stop HBR-Timer → S0
TMS Timeout	N/A	^TCP.ABORTreq → S0 (N/A for the initiator)	^TCP.ABORTreq ^HL.PROTOCOL-ABORTind → S0 (N/A for the initiator)	N/A	N/A
HBT Timeout	N/A	N/A	^TCP.DATAreq(TML heartbeat) /start HBT-Timer	N/A	N/A
HBR Timeout	N/A	N/A	^TCP.ABORTreq ^HL.PROTOCOL-ABORTind /cleanup → S0	^TCP.ABORTreq [not local peer abort] ^HL.PROTOCOL- ABORTind → S0	^TCP.ABORTreq → S0
CPA Timeout	N/A	N/A	N/A	N/A	^TCP.ABORTreq → S0

ANNEX A**TML DIAGNOSTIC CODES****(NORMATIVE)**

The diagnostic codes listed in this annex are used to report errors detected on the level of the Transport Mapping Layer.

The error codes are allocated in the range 128 to 255, which the Recommended Standards for SLE transfer services make available for technology specific use. Errors related to this Recommended Standard use the range 128 to 199, leaving the range 200 to 255 for use by implementations.

Code	Meaning
128	TML protocol error
129	Badly formatted TML message
130	Heartbeat parameters not acceptable
131	Association establishment timeout
132	Heartbeat receive timeout
133	Unexpected disconnect by peer
134	Premature disconnect during peer abort
135	Timeout during peer abort
199	Other reason

ANNEX B

DIFFERENCES WITH EARLIER IMPLEMENTATIONS

(INFORMATIVE)

B1 INTRODUCTION

The precursor of this Recommended Standard is a document called ‘Specification of an SLE API Proxy for TCP/IP and ASN.1’ (reference [E5]), which was published by the European Space Agency. This document adopts the specification of the authentication scheme and of the service instance identifier from another ESA document called ‘SLE C++ Application Program Interface for Transfer Services’ (reference [E6]).

All known implementations of SLE transfer services at the time this Recommended Standard was published were based on reference [E5]. This annex describes the differences between this Recommended Standard and reference [E5] in order to support development of implementations that can interoperate with implementations based on reference [E5].

Two differences must be considered:

- a) support for earlier versions of the SLE transfer services RAF, RCF, and CLTU; and
- b) structure of the service instance identifier.

B2 EARLIER VERSIONS OF THE SERVICES RAF, RCF, AND CLTU

Reference [E5] was based on earlier versions of the Recommended Standards for the services RAF, RCF, and CLTU, namely

- reference [E2] for the Return All Frames service specification;
- reference [E3] for the Return Channel Frames specification; and
- reference [E4] for the Forward CLTU service specification.

All these draft Recommended Standards specify the version number ‘1’, whereas references [2], [3], and [5] specify version number 2. Some of the data types specified in annex A of references [E2], [E3], and [E4] differ from those defined in references [2], [3], and [5]. An implementation wishing to interoperate with earlier implementations will have to use the ASN.1 specification in references [E2], [E3], and [E4].

Implementations invoking the BIND operation will have to know whether the responder conforms to this Recommended Standard or to reference [E5]. If the responder conforms to reference [E5], the initiator must specify version 1 in the BIND operation and use the ASN.1 types specified in the applicable draft Recommended Standard.

Implementations performing the BIND operations can derive the ASN.1 specification to use from the version number parameter in the BIND invocation.

NOTE – The ASN.1 encoding of the BIND invocation has not been modified and it is therefore possible to decode the BIND invocation with either ASN.1 specification.

B3 STRUCTURE OF THE SERVICE INSTANCE IDENTIFIER

Because references [E2], [E3], and [E4] did not specify the structure of the service instance identifier, a preliminary specification was provided by reference [E6]. This specification defined a set of attribute names that could be used within a service instance identifier and specified that all attribute values should be character strings.

The specification of the service instance identifier in annex A of references [2], [3], and [5] adopts a small subset of the attribute names defined by reference [E6] and additionally specifies the sequence of attributes in a service instance identifier. It adopts the specification that all attribute values shall be character strings. Therefore, a service instance identifier constructed according to the specification in references [2], [3], and [5] generally conforms to the specification given in reference [E6].

However, there is one exception, which must be taken into account by implementations wishing to support earlier implementations. For historical reasons the service type entry in the service instance identifier originally distinguished between Virtual Channel Frames (VCF) and Master Channel Frames (MCF), whereas the new specification has been aligned with the service definitions and allows only the attribute RCF (Return Channel Frames). Therefore, a service instance identifier for an RCF service instance with timely online delivery mode constructed according to the specification in reference [E6] would contain the attribute value pair ‘vcf=onlt1’ or ‘mcf=onlt1’, whereas the same identifier constructed according to reference [3] must contain the attribute-value pair ‘rcf=onlt1’.

ANNEX C

INDEX TO DEFINITIONS

(INFORMATIVE)

This annex lists terms used in this Recommended Standard and, for each term, provides a reference to the definition of that term.

Term	Reference
abstract syntax	reference [14]
Abstract Syntax Notation One (ASN.1)	reference [9]
application (of SLE)	subsection 1.6.1.6.2
application identifier	subsection 1.6.1.6.5
association	references [2] to [6]
Basic Encoding Rules (BER)	reference [10]
communications service	references [2] to [6]
confirmed operation	references [2] to [6]
Distinguished Encoding Rules (DER)	reference [10]
encoding	reference [10]
encoding rules (of ASN.1)	reference [10]
initiator	reference [1]
Internet Protocol (IP)	reference [8]
invocation	references [2] to [6]
IP address	reference [8]
layer (of a protocol)	reference [14]
local application	subsection 1.6.1.6.3
module (of ASN.1)	reference [9]
operation	reference [1]
parameter (of an operation)	references [2] to [6]
peer application	subsection 1.6.1.6.4
port (of TCP)	reference [7]
port identifier	references [2] to [6]
port number	reference [7]

Term	Reference
primitive	reference [14]
responder	reference [1]
return	references [2] to [6]
segment (of TCP)	reference [7]
service provider (provider)	reference [1]
service user (user)	reference [1]
SLE protocol data unit (SLE-PDU)	reference [1]
SLE transfer service instance (service instance)	reference [1]
socket	reference [7]
transfer syntax	reference [14]
Transmission Control Protocol (TCP)	reference [7]
unconfirmed operation	references [2] to [6]

ANNEX D**ACRONYMS****(INFORMATIVE)**

This annex lists the acronyms used in this Recommended Standard.

AL	Authentication Layer
API	Application Program Interface
ASN.1	Abstract Syntax Notation One
BER	Basic Encoding Rules
CCSDS	Consultative Committee for Space Data Systems
CPA	Close after Peer Abort (Timer)
DEL	Data Encoding Layer
DER	Distinguished Encoding Rules
EBNF	Extended Backus-Naur Form
HBR	Heartbeat Receive (Timer)
HBT	Heartbeat Transmit (Timer)
ICMP	Internet Control Message Protocol
IP	Internet Protocol
ISO	International Standardization Organization
ISP1	Internet SLE Protocol One
MSL	Maximum Segment Lifetime
MTU	Maximum Transmission Unit
OCL	Object Constraint Language
OSI	Open Systems Interconnection
PDU	Protocol Data Unit

RFC	Request For Comments
SLE	Space Link Extension
TCP	Transmission Control Protocol
TLI	Transport Layer Interface
TML	Transport Mapping Layer
TMS	TML Startup (Timer)
UML	Unified Modeling Language
XTI	X/Open Transport Interface

ANNEX E

INFORMATIVE REFERENCES

(INFORMATIVE)

- [E1] *Cross Support Concept—Part 1: Space Link Extension Services*. Issue 3. Report Concerning Space Data System Standards (Green Book), CCSDS 910.3-G-3. Washington, D.C.: CCSDS, March 2006.
- [E2] *Space Link Extension—Return All Frames Service Specification*. Issue 1.7. Proposed Draft Recommendation for Space Data System Standards (Draft Red Book), CCSDS 911.1-R-1.7. n.p.: n.p., September 1999.
- [E3] *Space Link Extension—Return Virtual Channel Frames Service Specification*. Issue 1.7. Proposed Draft Recommendation for Space Data System Standards (Draft Red Book), CCSDS 911.2-R-1.7. n.p.: n.p., September 1999.
- [E4] *Space Link Extension—Forward CLTU Service*. Issue 1.99h. Proposed Draft Recommendation for Space Data System Standards (Draft Red Book), CCSDS 912.1-R-1.99h. n.p.: n.p., February 2000.
- [E5] *Specification of a SLE API Proxy for TCP/IP and ASN.1*. Issue 1.1. SLES-SW-API-0002-TOS-GCI. Paris: ESA, February 2001.
- [E6] *SLE C++ Application Program Interface for Transfer Services*. SLES-SW-API-0001-TOS-GCI. Paris: ESA, December 2000.
- [E7] R. Braden, ed. *Requirements for Internet Hosts*. STD 3. Reston, Virginia: ISOC, October 1989.
- [E8] *Standard for Information Technology—Portable Operating System Interface (POSIX(R))*. IEEE Std 1003.1-2008. Piscataway, NJ: IEEE, 2008.
- [E9] *Information Technology—Open Distributed Processing—Unified Modeling Language (UML) Version 1.4.2*. 1st ed. International Standard, ISO/IEC 19501:2005. Geneva: ISO, 2005.
- [E10] *Space Link Extension—Internet Protocol for Transfer Services*. Issue 1-S. Recommendation for Space Data System Standards (Historical), CCSDS 913.1-B-1-S. Washington, D.C.: CCSDS, (September 2008) September 2015.

British Standards Institution (BSI)

BSI is the national body responsible for preparing British Standards and other standards-related publications, information and services.

BSI is incorporated by Royal Charter. British Standards and other standardization products are published by BSI Standards Limited.

About us

We bring together business, industry, government, consumers, innovators and others to shape their combined experience and expertise into standards-based solutions.

The knowledge embodied in our standards has been carefully assembled in a dependable format and refined through our open consultation process. Organizations of all sizes and across all sectors choose standards to help them achieve their goals.

Information on standards

We can provide you with the knowledge that your organization needs to succeed. Find out more about British Standards by visiting our website at bsigroup.com/standards or contacting our Customer Services team or Knowledge Centre.

Buying standards

You can buy and download PDF versions of BSI publications, including British and adopted European and international standards, through our website at bsigroup.com/shop, where hard copies can also be purchased.

If you need international and foreign standards from other Standards Development Organizations, hard copies can be ordered from our Customer Services team.

Copyright in BSI publications

All the content in BSI publications, including British Standards, is the property of and copyrighted by BSI or some person or entity that owns copyright in the information used (such as the international standardization bodies) and has formally licensed such information to BSI for commercial publication and use.

Save for the provisions below, you may not transfer, share or disseminate any portion of the standard to any other person. You may not adapt, distribute, commercially exploit, or publicly display the standard or any portion thereof in any manner whatsoever without BSI's prior written consent.

Storing and using standards

Standards purchased in soft copy format:

- A British Standard purchased in soft copy format is licensed to a sole named user for personal or internal company use only.
- The standard may be stored on more than 1 device provided that it is accessible by the sole named user only and that only 1 copy is accessed at any one time.
- A single paper copy may be printed for personal or internal company use only.

Standards purchased in hard copy format:

- A British Standard purchased in hard copy format is for personal or internal company use only.
- It may not be further reproduced – in any format – to create an additional copy. This includes scanning of the document.

If you need more than 1 copy of the document, or if you wish to share the document on an internal network, you can save money by choosing a subscription product (see 'Subscriptions').

Reproducing extracts

For permission to reproduce content from BSI publications contact the BSI Copyright & Licensing team.

Subscriptions

Our range of subscription services are designed to make using standards easier for you. For further information on our subscription products go to bsigroup.com/subscriptions.

With **British Standards Online (BSOL)** you'll have instant access to over 55,000 British and adopted European and international standards from your desktop. It's available 24/7 and is refreshed daily so you'll always be up to date.

You can keep in touch with standards developments and receive substantial discounts on the purchase price of standards, both in single copy and subscription format, by becoming a **BSI Subscribing Member**.

PLUS is an updating service exclusive to BSI Subscribing Members. You will automatically receive the latest hard copy of your standards when they're revised or replaced.

To find out more about becoming a BSI Subscribing Member and the benefits of membership, please visit bsigroup.com/shop.

With a **Multi-User Network Licence (MUNL)** you are able to host standards publications on your intranet. Licences can cover as few or as many users as you wish. With updates supplied as soon as they're available, you can be sure your documentation is current. For further information, email subscriptions@bsigroup.com.

Revisions

Our British Standards and other publications are updated by amendment or revision.

We continually improve the quality of our products and services to benefit your business. If you find an inaccuracy or ambiguity within a British Standard or other BSI publication please inform the Knowledge Centre.

Useful Contacts

Customer Services

Tel: +44 345 086 9001

Email (orders): orders@bsigroup.com

Email (enquiries): cservices@bsigroup.com

Subscriptions

Tel: +44 345 086 9001

Email: subscriptions@bsigroup.com

Knowledge Centre

Tel: +44 20 8996 7004

Email: knowledgecentre@bsigroup.com

Copyright & Licensing

Tel: +44 20 8996 7070

Email: copyright@bsigroup.com

BSI Group Headquarters

389 Chiswick High Road London W4 4AL UK