



ISO/IEC JTC1/SC22
Languages
Secretariat: CANADA (SCC)

ISO/IEC JTC1/SC22
N 1251

SEPTEMBER 1992

TITLE: Draft Paper on: Programming Language
Independent Specification Methods

SOURCE: Secretariat ISO/IEC JTC1/SC22

WORK ITEM: N/A

STATUS: New

CROSS REFERENCE: N/A

DOCUMENT TYPE: Information document

ACTION: For information to SC22 Member Bodies.

Address reply to: ISO/IEC JTC1/SC22 Secretariat
J.L. Côté
Treasury Board Secretariat
140 O'Connor St., 10th Floor, Ottawa, Ontario, Canada, K1A 0R5
Tel.: (613)957-2496 Telex: 053-3336 Fax: (613)996-2690

to SC22
for info

Draft TCOS-SSC Technical Report — Programming Language Independent Specification Methods

Sponsor

Technical Committee on Operating Systems
and Application Environments
of the
IEEE Computer Society

Work Item Number: JTC 1.22.21.x.y

Abstract: *Programming Language Independent Specification Methods* provides guidance to the working groups in the TCOS Standards Subcommittee who are preparing language-independent specifications and language bindings for approval as IEEE standards and eventually ISO/IEC standards. It describes an interface model with abstract datatypes and abstract procedure calls, specifies the notational conventions associated with the model, and provides guidelines for the use of the model in programming-language-independent specifications and language bindings to them. Although this document has been purposely designed to look somewhat like a real standard, it is not an official IEEE or ISO publication.

Keywords: language-independent specifications, language bindings, datatype, procedure call

TCOS-LIS / D4 May 1991

This is an unapproved draft and is subject to change.
All rights reserved by the Institute of
Electrical and Electronics Engineers.
Do not specify or claim conformance to this document.

**The Institute of Electrical and Electronics Engineers, Inc.
345 East 47th Street, New York, NY 10017, USA**

Copyright IEEE, all rights reserved. Permission is granted for unlimited reproduction in any form as long as the document or parts of the document are copied exactly and the copies are used for IEEE purposes and related ISO/IEC JTC1 standards activities only. Reproductions are not to be sold and the material in the document is not to be published in any book, magazine, or trade journal.

May 1991

SHxxxx

**UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.**

Contents

| | PAGE |
|---|------|
| Foreword | v |
| Document Status | vi |
| Section 1: Background | 1 |
| Section 2: Scope and Purpose | 3 |
| 2.1 Goals | 3 |
| 2.2 Non-Goals | 4 |
| Section 3: References | 7 |
| Section 4: Definitions | 9 |
| Section 5: The Model | 11 |
| 5.1 Using the Model | 11 |
| 5.2 Execution Sequence and Concurrency | 12 |
| 5.3 Datatypes | 13 |
| 5.3.1 Value Spaces | 13 |
| 5.3.2 Properties and Operations | 14 |
| 5.3.3 Base Datatypes and Datatype Derivation | 15 |
| 5.3.4 Constructed Datatypes | 15 |
| 5.3.5 Common Datatypes | 16 |
| 5.4 Value Names | 19 |
| 5.5 Procedures | 19 |
| Section 6: Conventions | 21 |
| 6.1 Notation | 21 |
| 6.2 Identifiers | 22 |
| 6.3 Datatypes | 22 |
| 6.4 Value Names | 23 |
| 6.5 Procedures | 24 |
| Section 7: Conformance | 25 |
| Section 8: Guidelines for Language-Independent Specifications | 27 |
| 8.1 General Guidelines | 27 |
| 8.1.1 Document Organization | 27 |
| 8.1.2 Terminology | 27 |
| 8.1.3 Documentation Requirements | 28 |
| 8.1.4 Language-Specific Features | 28 |
| 8.1.5 Atomicity | 28 |

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

| | | |
|---|---|----|
| 8.2 | Identifiers | 29 |
| 8.3 | Datatypes | 29 |
| 8.3.1 | Opaque Datatypes | 29 |
| 8.3.2 | Named Datatypes | 29 |
| 8.3.3 | Order Datatypes | 30 |
| 8.3.4 | Numeric Datatypes | 30 |
| 8.3.5 | Special Datatypes | 30 |
| 8.3.6 | Derived Datatypes | 30 |
| 8.3.7 | Constructed Datatypes | 30 |
| 8.3.8 | Handling 'flag words' | 30 |
| 8.3.9 | Handling Sets and Lists | 31 |
| 8.4 | Value Names | 31 |
| 8.5 | Procedures | 31 |
| 8.5.1 | How Big are Atomic Procedures? | 31 |
| 8.5.2 | Atomic Set-and-Return-Previous-Value Procedures | 32 |
| 8.5.3 | Atomic Compound Procedures | 32 |
| 8.5.4 | Avoid Overloaded Procedures | 32 |
| 8.5.5 | No Side Effects on Procedure Failure | 32 |
| 8.5.6 | Procedures that 'Can't Fail' | 33 |
| 8.5.7 | Boolean Procedures | 33 |
| 8.5.8 | Undefined vs Ignored Input Parameters | 33 |
| 8.5.9 | Undefined Output Parameters | 34 |
| 8.5.10 | Procedure Descriptions | 34 |
| Section 9: Guidelines for Language Bindings | | 35 |
| 9.1 | General Guidelines | 35 |
| 9.1.1 | Identify Language-Specific Interfaces | 35 |
| 9.1.2 | Atomicity | 36 |
| 9.1.3 | Provide Cross-References | 36 |
| 9.2 | Identifiers | 36 |
| 9.3 | Datatypes | 37 |
| 9.4 | Value Names | 38 |
| 9.5 | Interface Objects | 38 |
| 9.6 | Procedures | 38 |
| Annex A (informative) Examples | | 41 |
| A.1 | Directories | 41 |
| A.1.1 | Directory Operations | 42 |
| A.2 | Working Directory | 45 |
| A.2.1 | Change Current Working Directory | 45 |
| A.2.2 | Get Working Directory Pathname | 46 |
| Identifier Index | | 47 |
| Alphabetic Topical Index | | 49 |

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

TABLES

Table 1 - Revision History vi

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

Foreword

1 This TCOS-SS Technical Report describes methods for producing programming-
2 language-independent functional specifications, and for producing programming
3 language bindings to those functional specifications. These methods include
4 abstract models and notations, together with guidelines for their use. These
5 guidelines are intended to be followed by all functional specifications and language
6 binding standards developed within the IEEE Technical Committee on Operating
7 Systems.

8 The models and guidelines described in this document are adapted from current
9 ISO/IEC JTC1/SC22/WG11 draft technical reports and guidelines concerning
10 language-independent specifications and language bindings. Because there are no
11 adopted standards or guidelines in this area, this document provides normative
12 text to be included within all TCOS-SS functional specifications, concerning
13 Definitions, Conventions, and Conformance. When ISO/IEC standards for language-
14 independent specifications are adopted, these TCOS-SS guidelines can be amended
15 to conform to those standards, where appropriate, and the normative text can be
16 replaced by references to those standards. In the interim, successive revisions of
17 these guidelines will be informed by changes in the ISO/IEC draft guidelines, as
18 these become available.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

Document Status

1 This document is a draft version of a proposed TCOS-SS Technical Report contain-
2 ing guidelines to be used by the TCOS-SS working groups to develop functional
3 specifications in language-independent form and language bindings to those func-
4 tional specifications.

5 Draft 4 is the first draft produced using the document build tools developed by Hal 4
6 Jespersen and Donn Terry. Shaded margins are used in this draft to indicate 4
7 Rationale or Editor's Notes. Text introduced or substantially modified in this draft 4
8 is indicated by a small numeral (the draft number) in the right-hand margin. A 4
9 topical index is provided in this draft, but is it buggy and substantially incomplete. 4
10 The identifier index is empty. These shortcomings will be remedied in future 4
11 drafts. 4

12 Section 6 (Conventions) has been revised to show the notation actually used in 4
13 1003.1LIS/D1. New material has been added to Sections 5, 8, and 9. The Annex A 4
14 examples from previous drafts have been removed, and replaced by sample text 4
15 from Section 5 of 1003.1LIS/D1. The list of open issues in Annex B has been 4
16 removed, since it is no longer current. In some cases, relevant text has been 4
17 added to the appropriate sections in the body of the document. Further material 4
18 addressing previously identified issues will be added in subsequent drafts.

19 **Table 1 - Revision History**

| Date | Revision Number | Author(s) | Summary of Changes |
|----------|-----------------|-----------------|--|
| 88-12-07 | 0.01 | Barker, Connors | Initial Draft |
| 89-03-17 | 0.02 | Connors, Kimmel | First revision, based on feedback received during January 1989 POSIX 1003 meetings. Added front matter. |
| 89-08-03 | 0.03 | Kimmel | Updated for use by other TCOS-SS working groups. Content specific to the conversion of 1003.1 has been removed. |
| 90-07-05 | 1.0 | Rabin, Walli | Reorganized, revised, and enlarged. The datatype model was replaced and adapted from JTC1/SC22/WG11 N162 Guidelines on Language-Independent Datatypes. |
| 90-10-15 | 2.0 | Rabin, Walli | Revised, based on feedback received during the July, 1990 POSIX meetings. |
| 91-01-04 | 3.0 | Rabin, Walli | Minor revisions, including some alignment of terminology with WG11. |
| 91-05-10 | 4.0 | Rabin | Revised, based on experience developing 1003.1LIS/D1.0. |

36 Please direct written comments to the author using the following address:

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

37 Paul Rabin
38 Open Software Foundation
39 11 Cambridge Center
40 Cambridge, Massachusetts
41 USA 02142

42 rabin@osf.org or uunet!osf.org!rabin
43 FAX: +1 617 225 2782
44 TEL: +1 617 621 8873

**UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.**

Programming Language Independent Specification Methods

Section 1: Background

1 "Since in principle there is no reason why a particular system facil-
2 ity should not be used by a program, regardless of the language in
3 which it is written, it is the practice of system facility specifiers to
4 define an 'abstract' functional interface that is language independ-
5 ent. In this way, the concepts in a particular system facility may be
6 refined by experts in that area without regard for language peculiar-
7 ities. An internally coherent view of a particular system facility is
8 defined, relating system functions to each other in a consistent way
9 and relating the system functions to other layers within the system
10 facility, including protocols for communication with other objects in
11 the total system."¹⁾

12 Language-independent specifications are components of an architectural approach
13 to the specification of open systems. This approach seeks to promote the free
14 interoperation and substitution of system components by providing rigorous and
15 public specifications of functional interfaces.

16 The ISO, together with its committees and subcommittees, defines the highest pre-
17 cedence standards for open systems, and also provides guidelines for the develop-
18 ment of standards for open systems. The ISO has mandated that language-
19 independent specifications be provided together with language bindings for all

20
21

1) ISO/IEC JTC1/SC22 document N754 - Guidelines for Language Bindings, page 2.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

22 system services. These language-independent specifications play a central role in
23 the standards arena.

24 • A language-independent service specification specifies those requirements that
25 are common to all language bindings to that service, and hence provides a func-
26 tional specification to which language bindings may conform.

27 • A language-independent service specification is a re-usable component for con-
28 structing language bindings.

29 • A language-independent service specification provides an abstract specification of
30 a service, and hence facilitates more rigorous modelling of services.

31 • A language-independent service specification facilitates interoperability between
32 applications using different languages to share a common service implementation.

33 • A language-independent service specification facilitates the specification of rela-
34 tionships between one service and another.

35 All TCOS-SS interface specifications that are forwarded to ISO/IEC
36 JTC22/WG15 for adoption as International Standards will consist of a
37 language-independent functional specification plus one or more language bindings.

38 Because a functional specification is abstract, a degree of divergence among its
39 language bindings is permitted, even encouraged to the extent that the languages
40 themselves differ from each other. It is precisely the purpose of a functional
41 specification to define those aspects in which the various language bindings should
42 conform to each other and those aspects in which they may differ. The methods
43 described in this document aim to provide the greatest possible flexibility for
44 language bindings.

45 New features may arise in a functional specification first, or as language-specific
46 extensions in one of the language bindings. Once a new feature is incorporated in
47 a revision of a functional specification, it will be propagated to revisions of the
48 language bindings. Development should be coordinated to prevent unnecessary
49 divergence of the language bindings from each other.

50 The specification scheme described in this document assumes that a language
51 binding is dependent on prior specifications for both the language and the service,
52 and has the obligation of conformance to them, rather than the reverse. In the
53 case of TCOS-SS services, the language binding has come first, and been the source
54 for the specifications of the language and the service. This naturally causes a
55 conflict between the authority of the primary language binding and that of the
56 language-independent service specification. Initially, the primary language bind-
57 ing will be authoritative. However, as other language bindings are developed and
58 used, the language-independent service specification will take on a more indepen-
59 dent role.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

Section 2: Scope and Purpose

1 This document is addressed to members of TCOS-SS working groups developing
2 functional specifications or language bindings to functional specifications. Its pur-
3 pose is to assist and coordinate the development of functional specifications and
4 language bindings by defining an abstract model, and providing guidelines for the
5 use of that model in the development of new functional specifications, the deriva-
6 tion of a base standard from an existing language binding, and the development of
7 new language bindings to a functional specification.

8 The model presented in this document is not intended to formalize all parts of an
9 interface specification, but only those aspects that are of concern to language bind-
10 ings. The semantics of the underlying services are not specifically addressed,
11 although they may be addressed by extensions of the current methods.

12 The model is primarily intended for use in developing language-independent
13 specifications for operating system and related services, and language bindings for
14 procedural programming languages. Whether the same or other methods are
15 suitable for other types of services and programming languages has not been
16 investigated.

17 This document provides technical guidelines for the development of language-
18 independent specifications and language bindings. This document does not pro-
19 vide organizational or administrative guidelines for the management of the
20 development process.²⁾

21 2.1 Goals

22 The proposals contained in this document have the following objectives:

- 23 1. To meet requirements to provide language-independent functional
24 specifications, and to conform to ISO guidelines for the development of these
25 specifications and of language bindings to them.
- 26 2. To facilitate the development of language-independent functional
27 specifications that have sufficient richness and precision to ensure common,
28 recognizable base functionality across all language bindings for application
29 developers.

30 _____
31 2) See [ISO1], Guidelines 1-5.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

- 32 3. To facilitate the development of language bindings to these functional
33 specifications and to allow these language bindings the greatest possible
34 flexibility to exploit the strengths of each particular language.

35 2.2 Non-Goals

36 The following related objectives are consciously excluded as being outside of the
37 scope of the immediate requirements:

- 38 1. To ensure interoperability between programs or program modules written in
39 different programming languages. Language-independent specifications
40 should identify those system datatypes whose interchange is required for a
41 minimum level of interoperability. However, the mechanism to be used to
42 support interchange of system of other datatypes will not be specified.

43 Interoperability is a difficult and poorly understood problem. At least parts
44 of this problem seem to fall within the scope of language and networking con-
45 cerns rather than operating system concerns. The appropriate division of
46 labor needs to be explored more fully. As the solutions become better under-
47 stood, and the requirements that properly belong to operating system inter-
48 face specifications become better defined, these requirements will be added
49 to the scope of this document.

50 Briefly, there are three related problems under the umbrella of interopera-
51 bility that seem to be most closely tied to language-independent
52 specifications:

- 53 • First, the system exports private datatypes to applications, mainly for use as
54 identifiers or attributes of system objects. A distributed application needs to
55 be able to share values of these system datatypes among its separate parts.
- 56 • Second, applications define their own datatypes for internal use. A distri-
57 buted application needs to export these datatypes to the system, when it
58 uses system services to store or transport values of these application data-
59 types among its separate parts.
- 60 • Third, applications need to use standard datatypes defined by an external
61 authority, including characters, time values, mathematical structures, etc.
62 Values of such standard datatypes need to be freely transportable between
63 the application and the system, and among different parts of an application.
- 64 2. To define complete service or interface semantics through the use of formal,
65 verifiable specification languages.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

66 Some formalization of functional specifications is needed in order to facilitate
67 both the derivation and the validation of language bindings. This require-
68 ment primarily concerns interface syntax and a thin layer of semantics for
69 the datatypes passed across the interface. Beyond this, it is certainly desir-
70 able that the underlying semantics (for instance, of the objects implemented
71 by the service) be clear, explicit, and unambiguous. For the foreseeable
72 future, functional specifications will have a "formalization gradient"; that is,
73 the level of formalization will decrease with depth from the surface (inter-
74 face). Language bindings that attempt to map to features below the surface
75 in order to gain flexibility, will have greater difficulty the deeper they go.

76 3. To ensure the portability of language or binding implementations, for
77 instance by defining an actual system interface.

78 A functional specification defines an abstraction of the same application
79 interface that the language binding defines more concretely, not a lower-level
80 interface on top of which language support may be implemented.

81 4. To permit direct binding to the language-independent specification, by speci-
82 fying interfaces in an implemented (or easily implementable) notation that
83 might permit binding through embedded alien syntax or other methods.

84 A functional specification defines an abstraction of an API, not an API. Each
85 language binding will define a concrete API that uses native syntax.

86 The first two excluded objectives are simply beyond the scope of the current pro-
87 ject, but do not conflict with the models or guidelines described in this document.
88 Future work should extend the first tentative steps taken here in these areas.
89 The second two excluded objectives do fundamentally conflict with the approach
90 taken here, and are not possible extensions of this work.

91 Note that this document does not specifically address the related issue of interna-
92 tionalization, which shares a concern for the treatment of character datatypes.
93 The term "language-independent" in this document always refers to programming
94 languages and not to so-called natural languages.

95 Note also that this document is not a standard, and cannot be referenced by nor-
96 mative text within a standard. Functional specifications need to specify their own
97 service interfaces, including datatypes and procedure calls. The model presented
98 in this document describes the information needed to specify datatypes and pro-
99 cedure calls, together with sample specifications for the datatypes likely to be
100 needed for the specification of TCOS-SS interfaces. These should be imitated
101 rather than referenced by functional specifications. When ISO standards for
102 Language-Independent Datatypes and Procedure Calls become available, func-
103 tional specifications may reference these as appropriate.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

Section 3: References

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

- 1 [POSIX.1] IS 9945-1: 1990
2 IEEE Std 1003.1-1990
3 Portable Operating System Interface for Computer
4 Environments
- 5 [PSSG] POSIX Standards Style Guide
6 Draft 4
- 7 [ISO1] ISO/IEC JTC1/SC22 N754
8 Work Item JTC1.22.14
9 Proposed DTR 10182 on: Information Processing Systems -
10 Guidelines for Language Bindings
- 11 [ISO2] ISO/IEC JTC1/SC22 N842
12 Work Item JTC1.22.17
13 ANSI X3T2/90-087
14 Common Language-Independent Datatypes: Working Draft #4
- 15 [ISO3] ISO/IEC JTC1/SC22/WG11 N188
16 Work Item JTC1.22.16
17 ANSI X3T2/90-074
18 Common Language-Independent Procedure Calling Mechanism:
19 Working Draft Version 2.1
- 20 [ISO4] ISO/IEC JTC1/SC22/WG11 N194R
21 Language-Independent Standards: Second Draft
- 22 [ISO5] ISO/IEC JTC1/SC21 N4927
23 Information Processing Systems - Open Systems Interconnection -
24 Remote Procedure Call
- 25 [ECMA] ECMA-127 Remote Procedure Call Using OSI
- 26 [PCTE] ECMA-xxx Portable Common Tool Environment
27 Draft 6 - October, 1989
28 = ISO/IEC JTC1/SC22/WG15 N70
- 29 [DM] OSI Object Management API Specification
30 Version 2.0 Draft 5
31 X.400 API Association and X/Open Company Limited
- 32 [ASN.1] ISO 8824: 1987
33 OSI Abstract Syntax Notation One

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

Section 4: Definitions

- 1 *[Some or all of these definitions should appear in each language-independent functional*
 2 *specification.]*
- 3 **Abstract Datatype:** A datatype defined by its properties, rather than by the
 4 representation of its values.
- 5 **Base Datatype:** A datatype whose definition is used in the definition of a derived
 6 datatype, especially a datatype whose values are, or are components of, the values
 7 of a constructed datatype.
- 8 **Concrete Datatype:** A datatype directly supported by a programming language.
- 9 **Constructed Datatype:** A datatype defined by a standard construction method from
 10 previously defined or primitive datatypes.
- 11 **Datatype:** A collection of distinguished values, together with a collection of charac-
 12 terizing operations on those values.
- 13 **Datatype Family:** A set of related datatypes sharing a common, parameterized data-
 14 type definition.
- 15 **Derived Datatype:** A datatype that is defined in terms of one or more previously
 16 defined base datatypes.
- 17 **Functional Specification:** A specification, using a language-independent notation, of
 18 a service, or set of services, that are available to applications written in several
 19 programming languages.
- 20 **Language Binding:** A specification of the standard interface to a service, or set of
 21 services, for applications written in a particular programming language.
- 22 **Language Standard:** A specification of the syntax and semantics of applications
 23 written in a particular programming language.
- 24 **Procedure:** A program abstraction with formal input and output parameters that
 25 are bound to objects or values in the application that invokes it.
- 26 **Overloading:** The encoding of multiple service features in a single interface feature.
- 27 **Status Datatype:** An abstract datatype whose values may be bound to "control"
 28 values as well as "data" values.
- 29 **Value:** A member of the value space of a datatype.
- 30 **Value Space:** The set of values associated with a datatype.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
 Preliminary—Subject to Revision.



Section 5: The Model

1 The model described in this section provides a conceptual foundation for the terms
 2 and notation used to specify procedural interfaces in a language-independent
 3 manner. Because the model will be used to specify a variety of interfaces that will
 4 in turn have language bindings from a variety of programming languages, it is
 5 subject to seemingly conflicting requirements. The model must be rich enough to
 6 allow the specification of all service interfaces of interest, and also weak enough so
 7 that it can be easily mapped onto the facilities provided by a wide range of pro-
 8 gramming languages. Of course, in any particular case, only a limited binding
 9 may be possible between a service interface and a programming language. The
 10 purpose of the model is to facilitate language bindings to the extent that they are
 11 possible, and to provide a tool that may be used during the definition of new ser-
 12 vices and languages to anticipate the effect of various choices on future language
 13 bindings.

14 The model includes those interface features that are intended to be shared by all 4
 15 language bindings to a language-independent specification, and excludes those 4
 16 features that may differ. The features included in the model are abstract data- 4
 17 types and abstract procedure interfaces. All other features of language bindings, 4
 18 such as memory allocation, parameter passing, exception handling, implementa- 4
 19 tion of datatypes, choice of identifiers, and packaging and visibility control are 4
 20 excluded from the model.

21 The model does not include elements whose values can be altered by assignment.
 22 The definition of objects, and of associated allocation and reference mechanisms
 23 are features delegated to the language binding.

24 The model does not include mechanisms for grouping interface elements and con-
 25 trolling their visibility to applications. Such mechanisms may be defined by par-
 26 ticular language bindings.

27 5.1 Using the Model

28 The model is intended to provide an conceptual framework in which language-
 29 independent functional specifications can be defined. In this framework, interface
 30 features are abstract, that is, they are defined in terms of their properties and
 31 behavior rather than their implementation. Each language binding binds these
 32 abstract features to concrete features that are expressible in the language, and
 33 hence implementable by the language implementation. If the language itself sup-
 34 ports deferred binding, some interface features may remain abstract in the
 35 language binding. Such features will be finally bound to concrete features at

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
 Preliminary—Subject to Revision.

36 compile time, or even at run time.

37 Interoperability requires the ability to transfer values between program modules
38 written in different programming languages. This seems to be precluded by hav-
39 ing abstract datatypes separately bound by each language binding. Instead, the
40 binding must occur in the functional specification itself, or in the implementation.
41 Specifying concrete datatypes in the functional specification defeats the purpose of
42 allowing maximum autonomy to each language binding, yet many programming
43 languages cannot defer value and datatype binding to the implementation.

44 The model is intended to be used to specify the interfaces to services, but not to
45 specify the semantics of these services. The model does not impose a means for
46 specifying the semantics of the underlying objects accessed via the abstract inter-
47 faces defined by a language-independent specification. Various informal or formal
48 methods may be used. Developers of new functional specifications are encouraged
49 to approach the problem of specifying semantics with the same techniques of func-
50 tional and datatype abstraction on which the interface model is based.

51 To what extent is it possible to specify a service interface without specifying the
52 underlying object model implemented by the service? Should different models and
53 notations be used to specify interface syntax and semantics?

54 5.2 Execution Sequence and Concurrency

55 The model specifies abstract procedural interfaces. Unless otherwise specified, a
56 procedure call transfers a single thread of control from the application to the
57 implementation and back to the application. In certain interfaces, the thread of
58 control does not return to the application, or returns together with a new thread
59 of control. In other cases, the thread of control passes from the implementation to
60 the application and then back to the implementation.

61 In certain cases, there may be restrictions on the order in which operations may
62 be performed sequentially on certain objects. In certain cases, operations by mul-
63 tiple threads of execution on the same object are guaranteed to be serialized. In
64 other cases, the application is responsible for ensuring serialization. Since the
65 model only concerns the interface itself, it places no restrictions on the con-
66 currency behavior above or below the interface. Hence, all requirements relating
67 to execution sequence or concurrency must be specified explicitly in each func-
68 tional specification.

69 This is an area where the model should be extended. Programming languages
70 may have difficulty binding to certain control structures, so functional
71 specifications should be as explicit as possible about their requirements in this
72 regard.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

73 5.3 Datatypes

74 The datatype model presented here is adapted from the model presented in
 75 JTC1/SC22 N842 Common Language-Independent Datatypes Working Draft #4
 76 [ISO2]. Features were omitted that were not needed for developing language-
 77 independent specifications of TCOS-SS interfaces. That document should be refer-
 78 enced for further explanation of the model presented here, and for guidance in
 79 extending this model.

80 The datatype model is intended to include all datatypes needed to specify TCOS-SS
 81 interfaces, and to support the range of abstractions needed to permit flexible
 82 language binding, while preserving the essential characteristics of each datatype.

83 For the purpose of this model, a datatype is considered to be a set of values
 84 together with a set of operations on those values. The identity of values of a data-
 85 type is dependent on the operations associated with the datatype. The identities
 86 of the operations themselves is assumed to be unproblematic. All values of a data-
 87 type must be the result of some set of operations that do not themselves depend on
 88 any value. For any set of operations, identical values yield identical results.

89 The set of all possible abstract datatypes is itself a complex datatype (fortunately
 90 outside the scope of this datatype model!), with mappings between the value
 91 spaces and operations of related datatypes. There are many possible ways to
 92 describe these relationships, and hence to categorize and specify particular data-
 93 types. The current approach is partly analytical and partly descriptive. A variety
 94 of datatype attributes are defined in order to permit flexible specification of data-
 95 types.

96 Because the model does not include datatype-valued expressions, a simple
 97 approach to datatype identity is possible: Datatypes with distinct names are dis-
 98 tinct datatypes, regardless of their definitions (or implementations). Each value is
 99 considered to be associated with a unique datatype. In certain cases the natural
 100 mapping between the value spaces of two or more related datatypes is sufficiently
 101 obvious that references to corresponding values may be substituted for each other
 102 without harm.

103 5.3.1 Value Spaces

104 The value space of an abstract datatype has characteristics that are independent
 105 of the relationships between values. These include:

106 • Atomicity: A datatype is atomic if its values have no visible components that can
 107 be operated on independently. An atomic datatype value appears simple, regard-
 108 less of its possible implementation. Non-atomic datatypes are constructed from
 109 atomic datatypes by standard construction operations.

110 • Nameability: Completely opaque datatypes have no names for any of their
 111 values. Other datatypes may have names for some or all values. Names may be
 112 defined by external standards or conventions, or may be declared in a functional
 113 specification.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
 Preliminary—Subject to Revision.

- 198 The following fundamental datatype construction methods are available:
- 199 • Array: An array datatype is derived from two datatypes: a base datatype and an
200 index datatype. A value of an array datatype is a mapping from the index data-
201 type to the base datatype.
 - 202 • Choice: A choice datatype is derived from one or more base datatypes. A value 4
203 of a choice datatype is a set of mappings from names to base datatypes, exactly 4
204 one of which is valid.
 - 205 • Record: A record datatype is derived from one or more base datatypes. A value 4
206 of a record datatype is a set of mappings from names to base datatypes, all of 4
207 which are valid.
- 208 Identity of values of constructed datatypes is defined as identity of corresponding 4
209 component values, recursively applied. For choice datatypes, corresponding com- 4
210 ponent values must be either both defined or both undefined.

211 5.3.5 Common Datatypes

212 The model defines a set of commonly used datatypes and datatype families, both
213 simple and constructed. The further information needed to define a particular
214 member of a datatype family, including base datatypes or numeric parameters, is
215 described in each case. The syntax associated with each of these datatypes is
216 defined in Section 6 - Conventions.

217 The following list is provisional. A rationale needs to be provided for the particu-
218 lar selection of basic datatypes. The definitions for the basic datatypes should be
219 structured according to the previously defined concepts. The character and time
220 datatypes, in particular, may require additional concepts.

221 5.3.5.1 Boolean

222 A datatype with two values, named "TRUE" and "FALSE", and a set of logical 4
223 operations: NOT, OR, AND, XOR, etc.

224 Parameters: None

225 5.3.5.2 Bit

226 A datatype with two values, named "0" ("ZERO") and "1" ("ONE"), and a set of bit- 4
227 wise operations: NOT, OR, AND, XOR, etc. The values are ordered, with $0 < 1$.

228 Parameters: None

229 5.3.5.3 Opaque

230 This is a family of datatypes, with no order or other operations defined. An 4
231 opaque datatype may have associated names that identify distinguished values.

232 Parameters: Number of distinct values

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

233 5.3.5.4 State

234 This is a family of unordered, finite datatypes. Each value is identified by an asso- 4
235 ciated name.

236 Parameters: List of names 4

237 The name of this datatype family was changed from "Enumerated" to "State" for
238 alignment with [ISO2]. Previous drafts raised the issue of whether ordered or
239 unordered datatypes (or both) should be supported. This draft includes the (unor-
240 dered) State datatype family, but not the (ordered) Enumerated datatype family,
241 because no requirement for an ordered type with defined names has been found in
242 POSIX.1].

243 5.3.5.5 Ordinal

244 This is a family of discrete, ordered datatypes, each of which is a bounded
245 subrange of the (ideal) infinite ordinal datatype. Ordinal datatypes are pure order 4
246 datatypes. Normally, the first value of an ordinal datatype has the name "1", and
247 successive values have the same names as the successive integers.

248 Parameters: Number of values

249 5.3.5.6 Integer

250 This is a family of ordered datatypes, each of which is a bounded subrange of the
251 (ideal) infinite integer datatype. Each integer datatype inherits the operations of
252 the base datatype: add, subtract, multiply, divide.

253 Parameters: Lower bound, Upper bound

254 Although arbitrary subranges are possible, in practice particular datatypes are
255 specified as signed or unsigned, and by a number of values equal to 2^n . Integer
256 datatypes that map efficiently onto hardware-supported datatypes are the most
257 commonly used. In most cases the range includes zero as midpoint or endpoint. A
258 recent exception is the datatype `ssize_t` (defined in [POSIX.1]), which has the range 4
259 `[-1, SSIZE_MAX]`.

260 5.3.5.7 Array

261 This is a family of datatypes constructed from a base datatype and an index data-
262 type. The base datatype may be any datatype. The index datatype must be a
263 finite, ordered datatype. A value of an array datatype contains a value of the base
264 datatype corresponding to each value of the index datatype.

265 Parameters: Base datatype, Index datatype 4

266 5.3.5.8 Choice 4

267 This is a family of datatypes constructed from a sequence of base datatypes, each 4
268 associated with a name. A value of a choice datatype contains, for exactly one 4
269 name, a value of the corresponding base datatype. 4

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

270 Parameters: List of <name, datatype> pairs

271 5.3.5.9 Record

272 This is a family of datatypes constructed from a sequence of base datatypes, each 4
273 associated with a name. A value of a record datatype contains, for each name, a 4
274 value of the corresponding base datatypes.

275 Parameters: List of <name, datatype> pairs

276 5.3.5.10 List

277 This is a family of datatypes constructed from a base datatype. A value of a list
278 datatype contains a sequence of zero or more values of the base datatype. Apply-
279 ing an operation to all members of a list may be supported through either of two
280 datatypes of programming paradigm. In the first, the sequencing control is pro-
281 vided by the application. In the second, it is provided by the implementation.

282 Parameters: Base datatype, optional maximum length

283 Is there a need for a datatype that supports concatenation like a list, but supports
284 indexing like an array?

285 5.3.5.11 Set

286 This is a family of datatypes constructed from a base datatype. A value of a set
287 datatype contains an unordered collection of values of the base datatype, each
288 occurring at most once. The `is_member` operation returns a boolean value that
289 depends on whether the specified value is a member of the set. Applying an
290 operation to all members of a set may be supported through either of two data-
291 types of programming paradigm. In the first, the sequencing control is provided
292 by the application. In the second, it is provided by the implementation.

293 Parameters: Base datatype

294 5.3.5.12 Character

295 This is a family of datatypes whose values are either "glyphs" or special non-
296 printing control characters. All character datatypes used in TCOS-SS functional
297 specifications include the set of Portable Filename Characters, defined in POSIX.1.

298 Character datatypes are considered unordered; all collation properties are exter-
299 nal to the model.

300 Parameters: Character Set

301 A character set can be specified by designating an externally specified character
302 set, or by enumeration, or by a combination of these.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

303 Character datatypes depend only on the character set, and not on the encoding of
304 the character set. 4

305 5.3.5.13 Character String

306 This is a family of datatypes constructed from a base character datatype. A value
307 of a character string datatype is a sequence of values of the base character data-
308 type.

309 Parameters: Base character datatype, optional maximum length

310 5.3.5.14 Time

311 This is a family of scaled, one-dimensional vector datatypes used to measure time 4
312 intervals. Time units are specified by designating an externally specified time 4
313 unit (e.g., "second", "day"). Time scales are specified as multiples or fractions of 4
314 the time unit. 4

315 Parameters: Unit, Scale, Maximum value

316 5.4 Value Names

317 The model includes symbolic names that are used to represent specific values of
318 specific datatypes. The model does not specify when a name is bound to its value:
319 at compile time, link time, or load time. The use of a value name as a parameter
320 in a datatype specification does not imply that its value is defined at compile time.

321 5.5 Procedures

322 Procedures are modelled as abstract interfaces. Unlike most real programming
323 languages, the model allows input and output parameters of any datatype. All
324 parameters are passed by value. There are no in-out parameters, or parameters
325 passed by name or by reference. Information communicated between a caller and
326 a service is contained wholly within the procedure parameters (except for the 4
327 status value). There is no use of "global variables" in the model. Of course, real
328 language bindings may choose other means for passing values across the interface.

329 Procedures do not have a distinguished return parameter. All returned values are
330 modelled as output parameters.

331 Procedures also return an abstract status value. In a language binding, the status
332 value may be mapped in various ways, including to an output parameter, a global
333 status variable, or an exception mechanism. Status values are also typed, so that
334 each procedure can be associated with a particular status value space. All status
335 datatypes will have a distinguished value indicating successful completion, and
336 other values indicating specific reasons for failure. Output parameters are valid if
337 the status value indicates success, and are invalid otherwise. There is no provi-
338 sion for informational status. If the status value is not the unique value

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

339 indicating success, the procedure is assumed to have failed with no side effects,
340 unless specified otherwise.

341 This procedure model follows the success/failure procedure model used in 4
342 [POSIX.1]. On success, all output parameters have defined values and all expected 4
343 side effects occur. On failure, no output parameters have defined values and no 4
344 side effects occur, except that a diagnostic status value is made available to the 4
345 application. This status value is not considered an output parameter because 4
346 language bindings might map the status value to control features such as excep- 4
347 tion handlers, as well as to data features. How the success or failure of each pro- 4
348 cedure is indicated to applications is defined by each language binding. 4

349 Members of WG11 have suggested a more general model that does not distinguish 4
350 between success and failure, and allows arbitrary subsets of the output parame- 4
351 ters to have defined values for each status. 4

352 The interface model includes procedures exported by the application as well as
353 those exported by the service implementation. Examples of the former from the C
354 language include "main()" and signal handlers.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

Section 6: Conventions

1 *[This Conventions section will be included in every language-independent functional*
 2 *specification.*

3 *The POSIX Standards Style Guide [PSSG] describes macros that can be used to format the*
 4 *notation defined below.]*

5 This subclause defines the notation used within this standard to specify the infor-
 6 mation that is communicated through interfaces to the defined services. The
 7 notation assumes a model of an abstract procedural interface, although language
 8 bindings to this standard may take other forms.

9 The key elements of the notation are abstract data datatype definitions, value
 10 name definitions, and procedure definitions.

11 It is expected that in each case the notations specified here will be supplemented
 12 by additional descriptions using other notational schemes. For instance, for the
 13 specification of each procedure, the notation described below is to be used in the
 14 Synopsis section, while the semantics are described in the Description section.

15 6.1 Notation

16 A modified BNF notation is used in this section. The characters ":", "=", "<", ">",
 17 "|", "{", and "}" are always part of this notation and are never used in terminal
 18 strings.

19 Non-terminal elements consist of a "<" character followed by an identifier followed
 20 by a ">" character. Each production consists of a non-terminal followed by the
 21 characters "::=" followed by a right-hand-side expression. In a right-hand-side
 22 expression, the "|" character separates alternate elements, and the "{" and "}"
 23 characters delimit optional elements.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
 Preliminary—Subject to Revision.

24 **6.2 Identifiers**

```

25 <identifier> ::= <initial-char> <identifier-body>
26 <initial-char> ::= <upper-case-letter> | <lower-case-letter>
27 <identifier-body> ::= <body-char> { <identifier-body> }
28 <body-char> ::= <upper-case-letter> | <lower-case-letter> |
29 <digit> | <separator>
30 <lower-case-letter> ::= a | b | c | d | e | f | g | h | i |
31 j | k | l | m | n | o | p | q | r | s | t |
32 u | v | w | x | y | z
33 <upper-case-letter> ::= A | B | C | D | E | F | G | H | I |
34 J | K | L | M | N | O | P | Q | R | S | T |
35 U | V | W | X | Y | Z
36 <digit> ::= 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0
37 <separator> ::= - | _

```

38 The following typographic conventions are used as a mnemonic aid to indicate the
39 syntactic datatype of identifiers.

40 1. Lowercase identifiers in the Helvetica Italic font are used for datatype and pro- 4
41 cedure names. Datatype names are distinguished by the suffix "_type". 4

42 Examples: *directory_handle_type, open_directory* 4

43 2. Roman uppercase identifiers are used for value names. 4

44 Examples: TRUE, MAX_DIRECTORY_HANDLE 4

45 3. Lowercase identifiers in the New Century Schoolbook Italic font are used for 4
46 procedure parameters. 4

47 Examples: *directory_name* 4

48 **6.3 Datatypes**

49 The following notation is used to specify a datatype synopsis: 4

50 **Datatype:** <datatype-name> 4

51 **Definition:** <datatype-definition> 4

52 **Description:** 4

53 **where** 4

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

```

54     <datatype-name> ::= <identifier>
55     <datatype-definition> ::= <atomic-type> | <constructed-type>
56     <atomic-type> ::= boolean_type/<opaque-type>
57     <opaque-type> ::= ordinal_type | <integer-type>
58     <state-type> ::= opaque/<size>
59     <state-list> ::= state/<state-list>
60     <state-value> ::= <state-value> { , <state-value> }
61     <ordinal-type> ::= <ordinal_type/<size>
62     <size> ::= <number>
63     <integer-type> ::= integer_type/<minimum>
64     <minimum> ::= <signed-number>
65     <maximum> ::= <signed-number>
66
67     <constructed-type> ::= <array-type> | <choice-type> | <record-type>
68     <array-type> ::= array<index-type>of
69     <index-type> ::= <ordinal-type> | <integer-type> | <state-type>
70     <base-type> ::= <type>
71     <choice-type> ::= choice/<member-list>
72     <record-type> ::= record/<member-list>
73     <member-list> ::= <member-definition> { , <member-definition> }
74     <member-definition> ::= <member-name> : <member-type>
75     <member-name> ::= <identifier>
76     <member-type> ::= <type>

```

77 This notation scheme is incomplete. The notations for specifying character data-
78 types, time datatypes, and for specifying a datatype "from scratch" have not yet
79 been formulated. Suggestions are welcome.

80 6.4 Value Names

81 The following notation is used to specify a value name synopsis: 4

```

82 Name:     <value-name> 4
83 Definition: <datatype> 4
84 Description: 4

```

85 where

```

86     <value-name> ::= <identifier>
87     <datatype> ::= <datatype-name> | <datatype-definition>

```

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

88 **6.5 Procedures**

89 The following notation is used to specify a procedure: 4

90 **Procedure:** <procedure-name> 491 **Status Type:** <datatype> 492 **Input:** <parameter-name> : <datatype> 4

93 <parameter-name> : <datatype> 4

94 **Output:** <parameter-name> : <datatype> 4

95 <parameter-name> : <datatype> 4

96 **Description:** 497 **where**

98 <procedure-name> ::= <identifier>

99 <datatype> ::= <datatype-name> | <datatype-definition>

100 <parameter-name> ::= <identifier>

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
 Preliminary—Subject to Revision.

```

54     <datatype-name> ::= <identifier>
55     <datatype-definition> ::= <atomic-type> | <constructed-type>
56     <atomic-type> ::= boolean_type/<opaque-type>
57                     <ordinal-type> | <integer-type>
58     <opaque-type> ::= opaque/<size>
59     <state-type> ::= state/<state-list>
60     <state-list> ::= <state-value> { , <state-value> }
61     <state-value> ::= <identifier>
62     <ordinal-type> ::= ordinal_type/<size>
63     <size> ::= <number>
64     <integer-type> ::= integer_type/<minimum>
65     <minimum> ::= <signed-number>
66     <maximum> ::= <signed-number>

67     <constructed-type> ::= <array-type> | <choice-type> | <record-type>
68     <array-type> ::= array<index-type>of
69     <index-type> ::= <ordinal-type> | <integer-type> | <state-type>
70     <base-type> ::= <type>
71     <choice-type> ::= choice/<member-list>
72     <record-type> ::= record/<member-list>
73     <member-list> ::= <member-definition> { , <member-definition> }
74     <member-definition> ::= <member-name> : <member-type>
75     <member-name> ::= <identifier>
76     <member-type> ::= <type>

```

77 This notation scheme is incomplete. The notations for specifying character data-
78 types, time datatypes, and for specifying a datatype "from scratch" have not yet
79 been formulated. Suggestions are welcome.

80 6.4 Value Names

81 The following notation is used to specify a value name synopsis: 4

```

82 Name:      <value-name>
83 Definition: <datatype>
84 Description:

```

85 where

```

86     <value-name> ::= <identifier>
87     <datatype> ::= <datatype-name> | <datatype-definition>

```

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

88 **6.5 Procedures**

| | | |
|-----|--|---|
| 89 | The following notation is used to specify a procedure: | 4 |
| 90 | Procedure: <procedure-name> | 4 |
| 91 | Status Type: <datatype> | 4 |
| 92 | Input: <parameter-name> : <datatype> | 4 |
| 93 | <parameter-name> : <datatype> | 4 |
| 94 | Output: <parameter-name> : <datatype> | 4 |
| 95 | <parameter-name> : <datatype> | 4 |
| 96 | Description: | 4 |
| 97 | where | |
| 98 | <procedure-name> ::= <identifier> | |
| 99 | <datatype> ::= <datatype-name> <datatype-definition> | |
| 100 | <parameter-name> ::= <identifier> | |

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

Section 7: Conformance

1 *[The following clauses should be included in every language-independent functional*
2 *specification.]*

3 *Editor's Note: This section will define the general requirements for conformance of a*
4 *language binding to a language-independent functional specification.*

5 *Conformance depends on properties of the mapping between the datatypes, value names, and*
6 *procedures of the language binding to those of the language-independent specification. For*
7 *instance, the mapping must preserve equality and other relations between each datatype's*
8 *values. How should these properties be specified?*

9 *Possible conformance levels include:*

- 10 • *direct conformance*
- 11 • *partial conformance*
- 12 • *conformance with extensions*

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

Section 8: Guidelines for Language-Independent Specifications

1 This chapter provides guidelines for developing a new abstract interface
2 specification or deriving one from an existing language binding.

3 8.1 General Guidelines

4 8.1.1 Document Organization 4

5 In developing a language-independent standard from a language-specific standard,
6 it may happen that the changes made in the translation process suggest improve-
7 ments to the document organization of the language-independent standard. The 4
8 benefits of such changes need to be balanced against the value for purposes of 4
9 review of maintaining a close correspondence between the sections of the 4
10 language-independent standard and those of the language-specific standard from 4
11 which it was derived. 4

12 8.1.2 Terminology 4

13 The development of a language-independent standard from existing language- 4
14 specific system service specifications involves the articulation of a new, abstract 4
15 interface based on coherent data and procedure models. The entire specification 4
16 must be reviewed carefully to determine which requirements are language-specific 4
17 and which are language-independent, and to invent new points of articulation 4
18 between them where necessary. Thus, the development process imposes the 4
19 requirement and provides the opportunity for a greater rigor in the use of termi- 4
20 nology. 4

21 It must be ensured that all normative terms and phrases are defined, and used 4
22 consistently, and that any use of undefined terms does not lead to normative 4
23 ambiguity. 4

24 Clear and consistent terminology should be used when referring to system objects 4
25 and their attributes, to associations between objects, and to internal procedures 4
26 that are invoked directly or indirectly.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

27 8.1.3 Documentation Requirements

28 The translation to language-independent form introduces the new term
29 "language-binding defined", analogous to "implementation defined", for those
30 features and behaviors that are not specified in a language-independent standard
31 but are required to be specified in each language binding.

32 The terms "unspecified" and "undefined" also shift in meaning, since any feature
33 or behavior that is unspecified or undefined in a language-independent standard
34 may be specified in a language binding, or specified to be implementation defined.

35 As a result of these changes in the framework of documentation requirements, all
36 uses of these terms must be reconsidered in the translation process. Some
37 features and behaviors that are specified in a language-specific standard become
38 language-binding specified in a language-independent standard. Some features
39 and behaviors that are implementation-defined in a language-specific standard
40 become unspecified in a language-independent standard.

41 8.1.4 Language-Specific Features

42 Interfaces should be generalized, removing restrictions imposed by a particular
43 language binding. Language-independent interface specifications should abstract
44 from the implementation to the purpose of the interface.

45 • Features that are part of a specific language binding, but for which the primary
46 specification is the language standard, should not be included in the language-
47 independent specification.

48 • Features that depend on, or are made necessary by features of the specific
49 language, should not be included in the language-independent specification.

50 • Features that are included in a specific language binding for historical reasons
51 only, or are deprecated, should not be included in the language-independent
52 specification.

53 Features that originate in a specific language binding but are supportable and
54 useful to other language bindings should be included in the language-independent
55 specification, though perhaps in a more general form.

56 8.1.5 Atomicity

57 In general, the language-independent specification should describe atomic
58 features that will not be further subdivided by language bindings. Features
59 should be combined only where it is desirable for all language bindings to group
60 features the same way.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

61 8.2 Identifiers

62 Since the language-independent service specification is intended for human reada-
63 bility, identifiers should be descriptive without being cumbersome. Identifiers
64 should adhere to a consistent style, and use typographic and other conventions to
65 distinguish syntactic categories, such as procedures, parameters, datatypes, and
66 value names. Commonly used abbreviations, such as "id" for "identifier", are
67 acceptable if used consistently.

68 It is useful for language-independent specifications to recommend for language 4
69 bindings any natural groupings of identifiers that should be capable of being 4
70 independently made visible by applications.

71 8.3 Datatypes

72 Select a datatype that is as abstract as possible, yet provides the required opera-
73 tions.

74 New datatypes should be specified by recursively specifying constructed datatypes.
75 Eventually, the construction will rest on one or more simple datatypes. Each con-
76 structed datatype should be specified by its method of construction and the data-
77 types of its component values. Each simple datatype that is a member of a data-
78 type family should be specified by its distinguishing attributes within that data-
79 type family.

80 When specifying opaque datatypes, any names representing values of that data-
81 type should be specified with it.

82 For each datatype, the standard should specify whether the language binding
83 should fully specify the corresponding concrete datatype, or instead defer this to
84 the implementation.

85 8.3.1 Opaque Datatypes

86 Use opaque datatypes to indicate that values have no visible semantics. Be sure
87 to specify the required number of different values that must be supported by an
88 opaque datatype.

89 Use opaque datatypes for attributes of, and handles for, externally defined objects.

90 8.3.2 Named Datatypes

91 For each state datatype used, the functional specification should specify whether
92 the datatype is extensible.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

93 **8.3.3 Order Datatypes**

94 Use order datatypes for sequencing and array indexing.

95 **8.3.4 Numeric Datatypes**

96 Use numeric datatypes only for counting and measuring. Be sure to specify the
97 required range of values.

98 **8.3.5 Special Datatypes**

99 **8.3.6 Derived Datatypes**

100 **8.3.6.1 Abstracted Datatypes**

101 **8.3.6.2 Restricted Datatypes**

102 If a parameter is a restricted datatype, it may be treated as if it were a less res-
103 tricted version of the base datatype, with an exception status if the actual parame-
104 ter is outside of the declared datatype. This is permitted because in general,
105 language bindings cannot be expected to perform range checking for restricted
106 datatypes.

107 **8.3.7 Constructed Datatypes**

108 Constructed datatypes should be used only where a specific data representation is
109 required, and where the datatype values have only those properties implied by the
110 construction method. Aggregate datatypes may prove restrictive if access control
111 mechanisms do not affect all fields equally.

112 For each union or record datatype used, the functional specification should specify
113 whether the datatype is closed or extensible.

114 **8.3.8 Handling 'flag words'**

115 If a group of flags represents an object state, this should be specified as a record of
116 boolean_type values. If a group of flags specifies command options, this should be
117 specified as separate boolean_type values.

4
4

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

| | | |
|-----|--|---|
| 118 | 8.3.9 Handling Sets and Lists | 4 |
| 119 | Concentrate on defining the datatypes. Don't constrain bindings to one set of | 4 |
| 120 | operations: <code>get_next_element</code> vs <code>callback_for_each_element</code> . A choice that is | 4 |
| 121 | natural for one language binding may not be so for another. | 4 |
| 122 | 8.4 Value Names | 4 |
| 123 | Most value names fall into one of the following categories: | 4 |
| 124 | • configuration parameters: boolean and integer values | 4 |
| 125 | • state datatype values | 4 |
| 126 | • opaque datatype values | 4 |
| 127 | Names that represent values of state datatypes or opaque datatypes should be | 4 |
| 128 | specified as part of the definition of the datatype. Unlike these, the meaning of | 4 |
| 129 | configuration parameters is not closely associated with the datatypes of their | 4 |
| 130 | values, so these should be specified separately. | 4 |
| 131 | Where a functional specification uses an abstract datatype to represent a concrete | - |
| 132 | datatype in a language binding, literal value names should be represented by | 4 |
| 133 | corresponding symbolic value names. | 4 |
| 134 | Functional specifications should indicate when names are bound to their values. | 4 |
| 135 | An important use of value names is to announce configuration options. Guidelines | 4 |
| 136 | should be provided for this. | 4 |
| 137 | 8.5 Procedures | 4 |
| 138 | 8.5.1 How Big are Atomic Procedures? | 4 |
| 139 | When a language-independent specification is developed to support two or more | 4 |
| 140 | existing language bindings, identification of the appropriate atomic procedures | 4 |
| 141 | may be difficult. | 4 |
| 142 | The procedures in language bindings must be specifiable in terms of combinations | |
| 143 | of procedures in the associated language-independent standard. Hence, the | 4 |
| 144 | atomic procedures selected in a language-independent standard should be no | 4 |
| 145 | larger than the functional intersections of the procedures of all language bindings | 4 |
| 146 | that must be supported. | 4 |
| 147 | On the other hand, a language binding is incomplete to the extent that all possible | 4 |
| 148 | combinations of the atomic procedures in the associated language-independent | 4 |
| 149 | standard are not supported. Hence, the atomic procedures selected in a | 4 |
| 150 | language-independent standard should be no smaller than the corresponding pro- | 4 |
| 151 | cedures of all language bindings that need to be supported. | 4 |

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

152 In case of conflict between these rules, the former takes precedence, since it is 4
 153 more important for a language binding to be well-defined than for it to be com- 4
 154 plete. 4

155 8.5.2 Atomic Set-and-Return-Previous-Value Procedures 4

156 For procedures that get and set service or service object attributes, two 4
 157 approaches to defining atomic procedures are possible. One is to have a single 4
 158 procedure that sets a new value of the attribute and returns the previous value. 4
 159 The other is to have two separate procedures: one to set a new value of the attri- 4
 160 bute, and one to get the current value of the attribute. In the first approach it is 4
 161 not possible to atomically get the current value of the attribute without tem- 4
 162 porarily modifying that value in a way that might be externally visible. In the 4
 163 second approach it is not possible to atomically set a new value and get the 4
 164 current value. Each approach provides useful atomic procedures for a different 4
 165 program design. A compromise approach would define two procedures: one to get 4
 166 the current value, and one to both set a new value and return the current value. 4

167 8.5.3 Atomic Compound Procedures 4

168 A procedure might atomically perform some subset of a group of actions, with the 4
 169 subset depending on the value of some parameters. In this case, breaking the pro- 4
 170 cedure up into elementary procedures to undo overloading would change the 4
 171 semantics, since atomic combined actions would no longer be supported. 4
 172 The choice of atomic operations affects the object model. If the a group of attri- 4
 173 butes can only be set together, then they constitute a single attribute. 4

174 8.5.4 Avoid Overloaded Procedures

175 Break up parameters that encode multiple values. Break up procedures whose 4
 176 parameters use different value subranges with different meanings, or which use 4
 177 different values to indicate different functions. Care must be taken, however, to 4
 178 avoid the introduction of race conditions. Overloaded procedures that use 4
 179 different parameter values, or the presence or absence of a parameter value, to 4
 180 indicate alternate underlying procedures, may be safely broken up to expose the 4
 181 underlying procedures. However, procedures that use overloading to indicate a 4
 182 selection of underlying procedures that are to be performed atomically, cannot be 4
 183 broken up without introducing race conditions. 4

184 8.5.5 No Side Effects on Procedure Failure 4

185 Any exceptions to the general rule that procedures that fail shall have no side 4
 186 effects should be carefully documented. 4
 187 The specification of side effects on failure, or valid output parameters on failure, 4
 188 should be avoided. Whenever an error condition would result in a determinate 4
 189 side effect or a determinate output parameter value, this condition should be 4

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
 Preliminary—Subject to Revision.

114 • **Cardinality:** A datatype may have a finite or infinite number of different values.
115 Infinite datatypes are seldom used directly in abstract interface specifications, but
116 are useful for the definition of other datatypes. The cardinality of finite datatypes
117 may be fixed in a functional specification or deferred to the language binding or
118 implementation. 4

119 5.3.2 Properties and Operations

120 Other characteristics of a datatype depend on relationships among members of the
121 value space of the datatype. Completely opaque datatypes have no visible rela-
122 tionships among their values. The value spaces of non-opaque datatypes have a
123 visible structure that is defined by functions associated with the datatype.

124 Since the functions associated with a datatype can be composed to yield new func-
125 tions, the complete set of functions defined on the value space of a datatype may
126 be quite large. It is usually sufficient (and preferable) to list only a small set of
127 basic functions. Where alternate basic function sets are possible, the choice of a
128 specific basic function set can have a significant effect on the programming style of
129 applications and the efficiency of implementations.

130 Because the model includes "pure" values, and does not (currently) include objects
131 whose value can be altered by assignment, all datatypes support operations to test
132 for value identity, but need not support assignment of values. 4

133 • **Order:** The most common relation between values is order. A datatype whose
134 values have only order relations is an order datatype. Order datatypes support an
135 `in_order` function that takes two values of the datatype and returns a boolean
136 value: `TRUE` if its arguments are in order and `FALSE` otherwise. If the order data-
137 type is discrete, it will support functions that return the predecessor or successor
138 of a value. If the order datatype has a least or greatest value, it will support func-
139 tions that return them. The family of order datatypes alone is quite varied,
140 including linear and cyclic orders. A datatype whose relations include order rela-
141 tions is called an ordered datatype, to distinguish it from the pure order data-
142 types.

143 • **Numeric:** Most order datatypes have numeric analogs that supplement the
144 order functions with arithmetic functions. Non-standard order datatypes may
145 require non-standard arithmetic functions. Numeric datatypes may be pure, non-
146 dimensional quantities (scalars) that may be freely combined using arithmetic
147 operations, or they may have one or more associated dimensions, such as length or
148 time, that must be "conserved" by arithmetic operations. Among the numeric
149 datatypes are integer datatypes, scaled datatypes, rational datatypes, real data-
150 types, complex datatypes, etc. Scaled datatypes (defined in [ISO2]) behave like
151 integers multiplied by a constant integral or fractional scale factor. Of the
152 numeric datatypes, only integer and scaled datatypes are likely to be used in
153 TCOS-SS functional specifications. 4

154 • **Special:** This group of datatypes includes all those that support functions other
155 than those supported by order datatypes or numeric datatypes. These functions
156 may be defined in various ways, including axiomatically, or by enumeration of
157 their values.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.

Preliminary—Subject to Revision.

158 5.3.3 Base Datatypes and Datatype Derivation

159 Although datatypes may be defined "from scratch", by defining the value space and
160 the operations on that value space, datatypes are more commonly defined by their
161 relation to similar, already defined datatypes. For this purpose, a set of base data-
162 types is predefined in the model (see below).

163 Datatype derivation can be either abstract or concrete. Abstract datatype deriva-
164 tion is simply a specification convenience, which does not imply a similar relation-
165 ship between the implementations of the base and the derived datatypes. Con-
166 crete datatype derivation, on the other hand, does imply an analogous relationship
167 between the implementations of the base and derived datatypes, and implies sup-
168 port for "casting" operations to generate values of the derived datatype from those
169 of the base datatype, and conversely.

170 The following fundamental datatype derivation methods are available:

171 • **Restriction:** A restricted datatype is one that is derived from another datatype
172 by selecting a subset of its value space, but keeping the same associated functions
173 (with suitably adjusted domain and range). Examples are subrange datatypes and
174 constructed datatypes with invariants.

175 • **Abstraction:** An abstracted datatype is one that is derived from another data-
176 type by keeping the same value space, but removing some of the supported func-
177 tions on values. An example is the order datatype derived from the integer data-
178 type by keeping the value space and value names but removing the arithmetic
179 operations.

180 • **Extension:** An extended datatype is one that is derived from another datatype
181 by the addition of new values to its value space.

182 • **Enrichment:** An enriched datatype is one that is derived from another datatype
183 by defining additional functions on its value space.

184 Of course, the change of value space that occurs in the derivation of a restricted or
185 extended datatype usually involves changes in the domain and range of the func-
186 tions defined on that value space. These changes may introduce exception condi-
187 tions. A function on a restricted datatype may yield a value that lies outside the
188 value subset associated with that datatype. For example, the last value of a
189 bounded order datatype has no successor.

190 How should such exception conditions be specified?

191 5.3.4 Constructed Datatypes

192 Constructed datatypes use special kinds of concrete datatype derivation to create
193 a non-atomic value space from one or more base datatypes. Constructed data-
194 types support no operations except the casting operations, so their component
195 values are separately manipulable. This makes them particularly useful as
196 representations within a module that defines and exports new functions, and
197 exports the datatype as an atomic datatype.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

190 specified as a success condition. If necessary, additional output parameters can be 4
 191 defined to provide diagnostic information. 4

192 8.5.6 Procedures that 'Can't Fail' 4

193 Error conditions can be defined in three ways: by the language-independent stan- 4
 194 dard, by the language binding, or by the implementation. The current distinction, 4
 195 between procedures that shall not fail and procedures that simply have no error 4
 196 conditions defined, refers to the implementation. In the former case, implementa- 4
 197 tions shall not define error conditions; in the latter case implementations may 4
 198 define error conditions. However, language bindings may always define error con- 4
 199 ditions. Even a procedure that returns an integer might raise an exception in a 4
 200 language that supports dynamic typing if called with an argument of the wrong 4
 201 datatype. The language-independent standard should not distinguish between 4
 202 procedures where language-specific errors are likely to be defined and procedures 4
 203 where language-specific errors are not likely to be defined. 4

204 For now, we are assuming that error conditions defined by language bindings and 4
 205 implementations will share the same status type. This implies that Status Type 4
 206 should be specified for all procedures, whether or not the language-independent 4
 207 specification specifies any conditions in which the procedure will fail. 4

208 8.5.7 Boolean Procedures 4

209 In creating a language-independent specification, there is often a choice to be 4
 210 made between alternate versions of procedures whose purpose is to perform a test. 4
 211 In the procedure model, test procedures can be specified to return a `boolean_type` 4
 212 output parameter that depends on the result of the test, or they can be specified to 4
 213 succeed or fail, depending on the result of the test. In certain language bindings, 4
 214 these versions may require very different programming styles. 4

215 If the `boolean_type` version is used, it should be carefully specified which condi- 4
 216 tions result in success with an output value of `FALSE` and which conditions result 4
 217 in failure. Depending on the definitions of these conditions, the `boolean_type` ver- 4
 218 sion might require an order for the detection of error conditions that is not 4
 219 required by the other version. 4

220 8.5.8 Undefined vs Ignored Input Parameters 4

221 In the procedure model, all input parameters must have defined values on a pro- 4
 222 cedure call. This model permits a parameter value to be ignored in certain cases, 4
 223 but this should be avoided in general as a form of overloading. 4

224 A different case is when the presence or absence of a value for an input parameter 4
 225 has a controlling effect on a procedure. This should be modelled by a `boolean_type` 4
 226 input parameter accompanying the optional input parameter. If the `boolean_type` 4
 227 parameter has the value `TRUE`, the associated input parameter will be used; if 4
 228 `FALSE`, the associated input parameter will be ignored. Language bindings that 4
 229 support optional procedure parameters, or provide a special `NULL` value for the 4

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
 Preliminary—Subject to Revision.

230 parameter datatype, can eliminate the `boolean_type` parameter. 4

231 8.5.9 Undefined Output Parameters 4

232 In the procedure model described in TCOS LIS, all output parameters have defined 4
233 values after a successful procedure call. In a few cases in POSIX.1LIS, output 4
234 parameters are undefined for successful procedures. In these cases, each such 4
235 output parameter is accompanied by a new boolean parameter whose value indi- 4
236 cates whether the accompanied parameter is defined or not.

237 8.5.10 Procedure Descriptions

238 In addition to the interface specifications captured by the model, further 4
239 specifications are needed, including:

240 • Restrictions on the use of this interface due to access controls or required 4
241 privileges.

242 • Restrictions on the sequential or concurrent order in which this interface may be 4
243 called.

244 • Restrictions dependent on the internal state of the service implementation.

245 • The precise dependency of the values of the output parameters on the values of 4
246 the input parameters and the internal state of the service implementation.

247 • The initial state of the service implementation, and for each possible state, the 4
248 new state resulting from a successful call to the procedure.

249 The specification of conditions that will or may result in the unsuccessful comple- 4
250 tion of the procedure, and the status value that will be returned in each case, is 4
251 part of the informal interface description, and not part of the interface definition.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

Section 9: Guidelines for Language Bindings

1 This section provides guidelines for specifying programming language bindings to
2 a language-independent specification.

3 9.1 General Guidelines

4 A language binding defines a mapping between a language standard and a func-
5 tional specification. This mapping should be harmonious, so that correct and
6 natural access to the service features is provided to programs that are not only
7 correct, but have good style in the particular programming language³⁾.

8 Features should not be duplicated. Service features should be bound to existing
9 language features, where possible, rather than to newly invented features.

10 Language bindings should reference the associated functional specification and
11 language standard whenever feasible rather than duplicating specifications con-
12 tained in those standards.

13 It is not necessary that a language binding provide access to all features of a func-
14 tional specification. In particular, features of a base standard may be omitted
15 from a language binding to that base standard where these conflict with features
16 of the associate language standard. The omission from a language binding of any
17 feature of the associated functional specification should be documented in that
18 language binding together with a rationale for the omission.

19 9.1.1 Identify Language-Specific Interfaces

20 All features of the language binding that have no equivalents in the language-
21 independent specification must be identified as language-specific extensions. This
22 includes both interfaces that are defined in the language standard, and interfaces
23 that are unique to the language binding.

24 _____
25 3) It is desirable that both language standards and functional specifications be accompanied by
26 guidelines for language bindings. This will facilitate the development of new language
27 bindings, and will encourage consistency among bindings to different functional specifications
28 and to different languages, respectively. See [ISO1], Guideline 4, p. 18.

29 9.1.2 Atomicity

30 The language binding may combine groups of features specified in the functional
 31 specification into a single feature or interface. However, a language binding
 32 should not map multiple features onto a single feature of the functional
 33 specification. In particular, interfaces that change the state of the service imple-
 34 mentation shall not be subdivided. It is permissible, though discouraged, to subdi-
 35 vide non-state-changing interfaces⁴⁾.

36 The mapping of abstract and opaque datatypes by datatypes supported by the
 37 language may require exposing representations, and therefore parts, of values
 38 intended to be treated as wholes. Care must be taken so that the interfaces pro-
 39 vided in the language binding, including language-specific extensions, preserve the
 40 integrity of these values.

41 9.1.3 Provide Cross-References

42 It is preferable for language bindings to a functional specification to have the same
 43 structure and organization as the functional specification, to facilitate comparison
 44 and verification of conformance⁵⁾. Wherever a language binding follows a different
 45 document organization, adequate indices and cross-references must be provided to
 46 allow correlation with the functional specification.

47 9.2 Identifiers

48 Each language binding must map the identifiers in the language-independent
 49 specification to identifiers in the programming language. These identifiers should
 50 conform to the conventions for identifier style in the language, as well as the for-
 51 mal requirements on character set, maximum length, etc.

52 The language binding must specify any effects of application constructs on the visi- 4
 53 bility of identifiers, and any requirements or options for the presence of such 4
 54 application constructs. The language binding must also specify which identifiers 4
 55 are reserved in each case, and the effect of application use of defined or reserved 4
 56 identifiers. 4

57 Since applications and language binding implementations may both introduce 4
 58 identifiers into an applications identifier name-space, language bindings should 4
 59 provide mechanisms for identifier name-space management. Where the program- 4
 60 ming language provides the necessary facilities (and they should be encouraged to 4
 61 do so), a modular approach to name-space management should be taken, including 4
 62 support for the following features: 4

63 _____
 64 4) See [ISO1] Guideline 8, page 20.
 65 5) See [ISO1] Guideline 48, page 37.

- 66 • Where language binding identifiers fall into functional clusters that might be 4
67 separately useful to applications, mechanisms should be provided for applications 4
68 to control the visibility of each cluster separately. 4
- 69 • Applications should be able to freely define identifiers, without interference from 4
70 indentifiers imported from language bindings. 4
- 71 • In environments where language bindings for multiple system services are avail- 4
72 able to applications, it is desirable to provide mechanisms for applications to selec- 4
73 tively make visible the identifiers exported by any set of language bindings, and to 4
74 qualify identifiers in order to avoid conflicts between identifiers imported from 4
75 more than one language binding, or between an imported and an application- 4
76 defined identifier. 4
- 77 • In environments where some language binding implementations may export 4
78 identifiers to other language binding implementations, it is desirable to provide 4
79 mechanisms to ensure that identifiers defined by an application do not interfere 4
80 with identifiers exported from one language binding implementation to another, 4
81 and that the visibility of identifiers exported by a language binding does not entail 4
82 the visibility to applications of identifiers imported by that language binding. 4
- 83 A less desirable approach is to share identifier name-spaces among the application
84 and all language bindings in an environment. This entails allocating reserved 4
85 name-spaces to each language binding, with the remainder available for use by 4
86 applications. Aside from restricting the name-spaces available to applications 4
87 (and to language binding implementations), this approach creates a coordination 4
88 problem among language bindings.

89 9.3 Datatypes

- 90 For each datatype included in a language binding, all required operations must be
91 supported. Additional operations may also be supported where these do not
92 conflict with the required operations. Such additional operations must be docu-
93 mented as language-specific extensions. Where such extensions are of general
94 utility, they should be incorporated in future revisicns of the language-
95 independent specification. In general, language-specific extensions should be
96 avoided, to minimize the divergence of the different language bindings to a service.
- 97 • use language-supported datatypes where possible
98 • use explicit procedures for datatypes not supported by the language
99 • take advantage of datatype-checking mechanisms
- 100 Different instances of an abstract datatype family need not be bound to members
101 of the same concrete datatype family. For instance, different list datatypes need
102 not be implemented the same way.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

103 9.4 Value Names

104 Value names in the abstract interface specification should be preserved in the
105 language binding. That is, they should not be replaced by literal values, however
106 the abstract datatype of the value name is represented in the language binding.

107 9.5 Interface Objects

108 Where a language-independent specification specifies all interfaces in terms of
109 values, and keeps all system objects hidden below the interface, language bindings
110 may expose certain objects at the interface. Examples in [POSIX.1] are the DIR
111 and FILE structures.

112 In such cases, language bindings must include additional requirements on applica-
113 tions regarding the initialization of objects, the effect of using copies of objects, and
114 restrictions on concurrent access to objects.

115 Language bindings that might be used in multi-threaded environments should
116 avoid the use of static objects for parameter passing, and should specify which pro-
117 cedure calls will behave as if serialized. Language bindings may provide addi-
118 tional mechanisms (i.e., locking) for application control of serialization.

119 9.6 Procedures

120 For each procedure in the language-independent functional specification, several
121 decisions need to be made:

122 • Choose the parameter passing method. Examples include: by value, by refer-
123 ence, by implicit reference using a global object.

124 • If the language supports functions, the use of the return value must be decided.

125 • Choose the mechanism for reporting status. Examples include: return value,
126 global value, or exception mechanism.

127 Where possible, parameters should be passed by value. Many languages impose
128 restrictions on the datatypes that can be passed by value. Nearly all languages
129 allow only a single output parameter to be passed by value. For output param-
130 eters passed by reference, it must be decided whether the caller or the service will
131 allocate memory to store the value. Language bindings may map a parameter
132 passed by reference to a global object, although this is generally undesirable.

133 The actual parameter passing mechanism used in a language binding may intro-
134 duce language-specific error conditions. For instance, if an out parameter is
135 passed by reference, the procedure might fail if the reference supplied by the
136 caller is invalid, or the target object is too small.

137 Where the interface elements present in the language binding map directly to
138 those of the abstract service specification, it is not necessary to repeat the inter-
139 face semantics. Instead, a cross-reference to the appropriate part of the abstract

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

140 services specification should be given. Where the mapping is not direct, such
141 specification should be provided as is necessary to indicate the relation between
142 the behavior of the interfaces in the language binding to that of the interfaces in
143 the abstract service specification.

144 For each interface element in the language-independent specification, a type of
145 mapping must be chosen.

146 • high-level: use existing language services or create new services at the same
147 level as existing language services

148 • low-level: provide direct access to service interfaces

149 • no mapping: if service feature conflicts with language model

150 Wherever there is substantial overlap between the service features and the
151 features of the language binding, the relation between them must be specified. If
152 the mapping is direct, the correspondence between the identifiers used in the ser-
153 vice specification and those used in the language binding must be given.

154 If the language binding does not provide a direct mapping to a service interface,
155 additional information is needed.

156 • What is the relation between the states of the language binding features and the
157 underlying service features? How are they synchronized?

158 • What are the restrictions on interoperability between the (high-level) language
159 interfaces and the (low-level) service interfaces?

160 • How is control transferred between the high-level and the low-level interfaces?

161 A non-direct mapping is one that uses different datatypes or values, or depends on
162 state information outside of the service implementation.

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

Annex A
(informative)
Examples

1 The following examples are taken from the initial draft of the language- 4
2 independent version of [POSIX.1].

3 **A.1 Directories**

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

4 A.1.1 Directory Operations 4

5 A.1.1.1 Synopsis 4

| | | |
|----|---|---|
| 6 | Datatype: <i>directory_handle_type</i> | 4 |
| 7 | Definition: <i>opaque_type</i> [MAX_DIRECTORY_HANDLE] | 4 |
| 8 | Description: Directory stream. See Description. | 4 |
| 9 | Name: MAX_DIRECTORY_HANDLE | 4 |
| 10 | Datatype: <i>integer_type</i> [<implementation defined>] | 4 |
| 11 | Description: Number of possible values in <i>directory_handle_type</i> . | 4 |
| 12 | Procedure: <i>open_directory</i> | 4 |
| 13 | Status Type: <i>posix_status_type</i> | 4 |
| 14 | Input: <i>directory_name</i> : <i>pathname_type</i> | 4 |
| 15 | Output: <i>directory_handle</i> : <i>directory_handle_type</i> | 4 |
| 16 | Procedure: <i>read_a_directory_entry</i> | 4 |
| 17 | Status Type: <i>posix_status_type</i> | 4 |
| 18 | Input: <i>directory_handle</i> : <i>directory_handle_type</i> | 4 |
| 19 | Output: <i>directory_entry_name</i> : <i>filename_type</i> | 4 |
| 20 | <i>end_directory_flag</i> : <i>boolean_type</i> | 4 |
| 21 | Procedure: <i>rewind_directory</i> | 4 |
| 22 | Status Type: <i>posix_status_type</i> | 4 |
| 23 | Input: <i>directory_handle</i> : <i>directory_handle_type</i> | 4 |
| 24 | Procedure: <i>close_directory</i> | 4 |
| 25 | Status Type: <i>posix_status_type</i> | 4 |
| 26 | Input: <i>directory_handle</i> : <i>directory_handle_type</i> | 4 |

27 A.1.1.2 Description 4

28 A value of type *directory_handle_type* represents a *directory stream*, which is a 4
 29 sequence of all the directory entries in a particular directory. A conforming appli- 4
 30 cation shall only attempt to access directory entries using a value of type 4
 31 *directory_handle_type* that has been returned from a successful call to the *read_a_* 4
 32 *directory_entry* procedure and shall not attempt to access directory entries after a 4
 33 successful call to the *close_directory* procedure on that directory stream. 4

34 Directory entries represent files; directory entries may be removed from a direc- 4
 35 tory or added to a directory asynchronously to the operations described in this 4
 36 subclause (A.1.1). The *directory_handle_type* may be implemented using a file 4
 37 descriptor. In that case, applications can only open up to a total of {OPEN_MAX} 4
 38 files and directories; see 4

39 A successful call to any of the *overlay_process_image* procedures shall close any 4
 40 directory streams that are open in the calling process. The result of using a direc- 4
 41 tory stream after one of the *overlay_process_image* family of procedures is 4
 42 undefined. After a call to the *fork_a_process* procedure either the parent or the 4
 43 child (but not both) can continue processing the directory stream using the *read_* 4
 44 *a_directory_entry* procedure or *rewind_directory* procedure or both. If both the 4

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
 Preliminary—Subject to Revision.

45 parent and child processes use these procedures, the result is undefined. Either 4
 46 or both processes can use the *close_directory* procedure. 4

47 The *open_directory* procedure shall open a directory stream corresponding to the 4
 48 directory named by *directory_name*. The directory stream shall be positioned at 4
 49 the first entry. 4

50 The *read_a_directory_entry* procedure shall return the directory entry at the 4
 51 current position in the directory stream to which *directory_handle* refers, and 4
 52 position the directory stream at the next entry. After the last entry in the direc- 4
 53 tory has been returned, subsequent calls to *read_a_directory_entry* shall set *end_-* 4
 54 *directory_flag* to TRUE and the value of *directory_entry_name* is undefined. 4

55 The *read_a_directory_entry* procedure shall not return an empty *directory_entry_-* 4
 56 *name*. It is unspecified whether entries are returned for dot or dot-dot. 4

57 The *read_a_directory_entry* procedure may buffer several directory entries per 4
 58 actual read operation; the *read_a_directory_entry* procedure shall mark for update 4
 59 the *time_file_last_accessed* field of the directory each time the directory is actually 4
 60 read. 4

61 The *rewind_directory* procedure shall reset the position of the directory stream to 4
 62 which *directory_handle* refers to the beginning of the directory. It also shall cause 4
 63 the directory stream to refer to the current state of the corresponding directory, 4
 64 as a call to *open_directory* procedure would have done. 4

65 If a directory entry that has not been returned is removed from or added to the 4
 66 directory after the most recent call to the *open_directory* or *rewind_directory* pro- 4
 67 cedures, whether a subsequent call to *read_a_directory_entry* procedure returns 4
 68 that entry is unspecified. 4

69 The *close_directory* procedure shall close the directory stream referred to by 4
 70 *directory_handle*. If a file descriptor is used to implement type *directory_handle_-* 4
 71 *type*, that file descriptor shall be closed. 4

72 If the *directory_handle* value passed to any of these procedures does not refer to a 4
 73 currently open directory stream, the result is undefined. 4

74 A.1.1.3 Errors 4

75 If any of the following conditions occur, the *open_directory* procedure shall fail and 4
 76 post the corresponding status value: 4

77 *error_access_denied* 4
 78 Search permission is denied for a component of the path prefix of 4
 79 *directory_name*, or read permission is denied for the directory itself. 4

80 *error_name_length_limit* 4
 81 The length of the *directory_name* argument exceeds {PATH_MAX}, or a 4
 82 pathname component is longer than {NAME_MAX} while {POSIX_NO_- 4

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
 Preliminary—Subject to Revision.

| | | |
|-----|--|---|
| 83 | TRUNC) is in effect. | 4 |
| 84 | <i>error_file_does_not_exist</i> | 4 |
| 85 | The named directory does not exist or <i>directory_name</i> is empty. | 4 |
| 86 | <i>error_is_not_a_directory</i> | 4 |
| 87 | A component of <i>directory_name</i> is not a directory. | 4 |
| 88 | For each of the following conditions, when the condition is detected, the <i>open_</i> - | 4 |
| 89 | <i>directory</i> procedure shall fail and post the corresponding status value: | 4 |
| 90 | <i>error_process_open_file_limit</i> | 4 |
| 91 | Too many file descriptors are currently open for the process. | 4 |
| 92 | <i>error_system_open_file_limit</i> | 4 |
| 93 | Too many file descriptors are currently open in the system. | 4 |
| 94 | For each of the following conditions, when the condition is detected, the <i>read_a_</i> - | 4 |
| 95 | <i>directory_entry</i> procedure shall fail and post the corresponding status value: | 4 |
| 96 | <i>error_invalid_file_descriptor</i> | 4 |
| 97 | The <i>directory_handle</i> argument does not refer to an open directory | 4 |
| 98 | stream. | 4 |
| 99 | For each of the following conditions, when the condition is detected, the <i>close_</i> - | 4 |
| 100 | <i>directory</i> procedure shall fail and post the the corresponding status value: | 4 |
| 101 | <i>error_invalid_file_descriptor</i> | 4 |
| 102 | The <i>directory_handle</i> argument does not refer to an open directory | 4 |
| 103 | stream. | 4 |
| 104 | A.1.1.4 Cross-References | 4 |
| 105 | None. | 4 |

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

| | | |
|-----|--|---|
| 106 | A.2 Working Directory | |
| 107 | A.2.1 Change Current Working Directory | 4 |
| 108 | A.2.1.1 Synopsis | 4 |
| 109 | Procedure: <i>change_current_working_directory</i> | 4 |
| 110 | Status Type: <i>posix_status_type</i> | 4 |
| 111 | Input: <i>target_directory_name</i> : <i>pathname_type</i> | 4 |
| 112 | A.2.1.2 Description | 4 |
| 113 | The <i>change_current_working_directory</i> procedure shall cause the named directory | 4 |
| 114 | to become the current working directory, that is, the starting point for resolutions | 4 |
| 115 | of pathnames not beginning with slash. The <i>target_directory_name</i> argument is | 4 |
| 116 | the name of the directory to change to. | 4 |
| 117 | If the <i>change_current_working_directory</i> procedure fails, the current working direc- | 4 |
| 118 | tory shall remain unchanged by this procedure. | 4 |
| 119 | A.2.1.3 Errors | 4 |
| 120 | If any of the following conditions occur, the <i>change_current_working_directory</i> pro- | 4 |
| 121 | cedure shall fail and post the corresponding status value: | 4 |
| 122 | <i>error_access_denied</i> | 4 |
| 123 | Search permission is denied for any component of the pathname. | 4 |
| 124 | <i>error_name_length_limit</i> | 4 |
| 125 | The <i>target_directory_name</i> argument exceeds {PATH_MAX} in length, or | 4 |
| 126 | a pathname component is longer than {NAME_MAX} while {POSIX_NO_- | 4 |
| 127 | TRUNC} is in effect. | 4 |
| 128 | <i>error_is_not_a_directory</i> | 4 |
| 129 | A component of the pathname is not a directory. | 4 |
| 130 | <i>error_file_does_not_exist</i> | 4 |
| 131 | The named directory does not exist or <i>target_directory_name</i> is empty. | 4 |
| 132 | A.2.1.4 Cross-References | |
| 133 | <i>get_current_working_directory</i> , A.2.2.1. | 4 |

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

| | | |
|-----|---|---|
| 134 | A.2.2 Get Working Directory Pathname | 4 |
| 135 | A.2.2.1 Synopsis | 4 |
| 136 | Procedure: <i>get_current_working_directory</i> | 4 |
| 137 | Status Type: <i>posix_status_type</i> | 4 |
| 138 | Output: <i>directory_name:pathname_type</i> | 4 |
| 139 | A.2.2.2 Description | 4 |
| 140 | The <i>get_current_working_directory</i> procedure shall return an absolute pathname of | 4 |
| 141 | the current working directory. | 4 |
| 142 | A.2.2.3 Errors | 4 |
| 143 | For each of the following conditions, if the condition is detected, the <i>get_current_</i> | 4 |
| 144 | <i>working_directory</i> procedure shall fail and post the corresponding status value: | 4 |
| 145 | <i>error_access_denied</i> | 4 |
| 146 | Read or search permission was denied for a component of the pathname. | 4 |
| 147 | A.2.2.4 Cross-References | 4 |
| 148 | <i>change_current_working_directory</i> , A.2.1.1. | 4 |

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

Identifier Index

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.



Alphabetic Topical Index

A

/ ... 23
 / ... 23
 Abstracted Datatypes ... 30
 AND ... 16
 ANSI ... 8
 Array ... 17
 ASN.1 ... 8
 Atomic Compound Procedures ... 32
 Atomic Set-and-Return-Previous-Value Procedures ... 32
 Atomicity ... 28, 36
 Avoid Overloaded Procedures ... 32

B

Background ... 1
 background ... 1
 Base Datatypes and Datatype Derivation ... 15
 Bit ... 16
 BNF ... 21
 Boolean Procedures ... 33
 Boolean ... 16
boolean_type ... 42

C

Change Current Working Directory ... 45
change_current_working_directory ... 45-46
 definition of procedure ... 45
 Character String ... 19
 Character ... 18
 Choice ... 17
close_directory ... 42-44
 definition of procedure ... 42
 Common Datatypes ... 16
 conformance
 Conformance ... 25
 conformance ... 25, 36

Constructed Datatypes ... 15, 30
 Conventions ... 21
 Cross-References ... 44-46
 current working directory
 change ... 45

D

<datatype-name> ... 22
 Datatypes ... 13, 22, 29, 37
 datatypes
 directory_handle_type ... 42
 Definitions ... 9
 Derived Datatypes ... 30
 Directories ... 41
 directory entry ... 42-43
 Directory Operations ... 42
 directory ... 22, 41-46
 change current working ... 45
 working pathname ... 46
directory_entry_name ... 42-43
directory_handle ... 42-44
directory_handle_type ... 22, 42-43
 definition of datatype ... 42
directory_name ... 22, 42-44, 46
 DIR ... 38
 Document Organization ... 27
 Document Status
 document
 Documentation Requirements ... 28
 dot ... 43
 dot-dot ... 43
 DTR ... 8

E

ECMA-127 ... 8
 ECMA ... 8
end_directory_flag ... 42-43
error_access_denied ... 43, 45-46
error_file_does_not_exist ... 44-45

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
 Preliminary—Subject to Revision.

error_invalid_file_descriptor ... 44
error_is_not_a_directory ... 44-45
error_name_length_limit ... 43, 45
error_process_open_file_limit ... 44
error_system_open_file_limit ... 44
Examples ... 41
Execution Sequence and Concurrency ... 12

F

FALSE ... 14, 16, 33
FAX
file_descriptor ... 42-44
filename_type ... 42
FILE ... 38
Foreword
fork_a_process ... 42

G

General Guidelines ... 27, 35
Get Working Directory Pathname ... 46
get_current_working_directory ... 45-46
 definition of procedure ... 46
Goals ... 3
Guidelines for Language Bindings ... 35
Guidelines for Language-Independent Specifications ... 27

H

Handling 'flag words' ... 30
Handling Sets and Lists ... 31
How Big are Atomic Procedures? ... 31

I

Identifiers ... 22, 29, 36
Identify Language-Specific Interfaces ... 35
IEEE Std 1003.1 ... 8
IEEE
implementation defined ... 28
<index-type> ... 23
Integer ... 17
integer_type ... 42

Interface Objects ... 38
ISO 8824 ... 8
ISO1 ... 3, 8, 35-36
ISO2 ... 8, 13-14, 17
ISO3 ... 8
ISO4 ... 8
ISO5 ... 8

J

JTC1

L

language binding
language-binding defined ... 28
Language-Specific Features ... 28
LIS ... 34
LIS/D1
List ... 18

M

MAX_DIRECTORY_HANDLE ... 22, 42
 definition of ... 42
Model ... 11

N

Named Datatypes ... 29
(NAME_MAX) ... 43, 45
No Side Effects on Procedure Failure ... 32
Non-Goals ... 4
Notation ... 21
NOT ... 16
NULL ... 33
Numeric Datatypes ... 30

O

ONE ... 16
Opaque Datatypes ... 29
Opaque ... 16
opaque_type ... 42
open_a_file
 definition of procedure ... 42

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
 Preliminary—Subject to Revision.

open_directory ... 22, 42-44
 definition of procedure ... 42
 (OPEN_MAX) ... 42
 Order Datatypes ... 30
 Ordinal ... 17
 OSI ... 8
overlay_process_image ... 42

P

pathname
 get working directory ... 46
 pathname ... 42-43, 45-46
pathname_type ... 42, 45-46
 (PATH_MAX) ... 43, 45
 PCTE ... 8
 POSIX.1 ... 8, 17-18, 20, 38, 41
 POSIX.1LIS ... 34
 (POSIX_NO_TRUNC) ... 43, 45
posix_status_type ... 42, 45-46
 Procedure Descriptions ... 34
 Procedures that 'Can't Fail' ... 33
 Procedures ... 19, 24, 31, 38
 procedures
 change_current_working_directory ... 45
 close_directory ... 42
 get_current_working_directory ... 46
 open_a_file ... 42
 open_directory ... 42
 read_a_directory_entry ... 42
 rewind_directory ... 42
 Properties and Operations ... 14
 Provide Cross-References ... 36
 PSSG ... 8, 21

R

read_a_directory_entry ... 42-44
 definition of procedure ... 42
 Record ... 18
 Restricted Datatypes ... 30
rewind_directory ... 42-43
 definition of procedure ... 42

S

Scope and Purpose ... 3

SC21 ... 8
 SC22
 Set ... 18
 Special Datatypes ... 30
 SSIZE_MAX ... 17
 State ... 17

T

target_directory_name ... 45
 TCOS ... 34
 TCOS-SS
 TEL
 terminal ... 21
 Terminology ... 27
 Time ... 19
time_file_last_accessed ... 43
 TRUE ... 14, 16, 22, 33

U

Undefined Output Parameters ... 34
 Undefined vs Ignored Input Parameters
 ... 33
 undefined ... 16, 27-28, 33-34, 42-43
 unspecified ... 28, 43
 USA
 Using the Model ... 11

V

Value Names ... 19, 23, 31, 38
 Value Spaces ... 13

W

WG11
 WG15 ... 2, 8
 Working Directory ... 45
 working directory
 change ... 45

X

X.400 ... 8

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
 Preliminary—Subject to Revision.

TCOS-LIS/D4

X/Open ... 8

XOR ... 16

Z

ZERO ... 16

UNAPPROVED DRAFT. All Rights Reserved by IEEE.
Preliminary—Subject to Revision.

