
**Information technology — Open
distributed processing — Reference
model: Architecture**

*Technologies de l'information — Traitement réparti ouvert — Modèle de
référence: Architecture*

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 10746-3:2009

PDF disclaimer

This PDF file may contain embedded typefaces. In accordance with Adobe's licensing policy, this file may be printed or viewed but shall not be edited unless the typefaces which are embedded are licensed to and installed on the computer performing the editing. In downloading this file, parties accept therein the responsibility of not infringing Adobe's licensing policy. The ISO Central Secretariat accepts no liability in this area.

Adobe is a trademark of Adobe Systems Incorporated.

Details of the software products used to create this PDF file can be found in the General Info relative to the file; the PDF-creation parameters were optimized for printing. Every care has been taken to ensure that the file is suitable for use by ISO member bodies. In the unlikely event that a problem relating to it is found, please inform the Central Secretariat at the address given below.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 10746-3:2009



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2009

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published by ISO in 2010

Published in Switzerland

CONTENTS

	<i>Page</i>
Foreword.....	v
Introduction	vi
1 Scope	1
2 Normative references	1
2.1 Identical Recommendations International Standards	1
2.2 Paired Recommendations International Standards equivalent in technical content.....	1
3 Definitions	2
3.1 Descriptive definitions	2
3.2 Abbreviations	3
4 Framework	3
4.1 Viewpoints	4
4.2 ODP viewpoint languages	5
4.3 ODP functions	5
4.4 ODP distribution transparencies	5
4.5 Standards derived from the framework	6
4.6 Conformance	7
5 Enterprise language	7
5.1 Concepts	7
5.2 Structuring rules	7
5.3 Conformance and reference points	8
6 Information language	8
6.1 Concepts	9
6.2 Structuring rules	9
6.3 Conformance and reference points	9
7 Computational language	10
7.1 Concepts	10
7.2 Structuring rules	12
7.3 Conformance and reference points	18
8 Engineering language	18
8.1 Concepts	19
8.2 Structuring rules	20
8.3 Conformance and reference points	28
9 Technology language	29
9.1 Concepts	29
9.2 Structuring rules	29
9.3 Conformance and reference points	29
10 Consistency rules	29
10.1 Computational and information specification correspondences	30
10.2 Engineering and computational specification correspondences	30
10.3 Technology and engineering specification correspondences	31
11 ODP functions	31
12 Management functions	32
12.1 Node management function	32
12.2 Object management function	33
12.3 Cluster management function	33
12.4 Capsule management function	34
13 Coordination functions	35
13.1 Event notification function	35
13.2 Checkpoint and recovery function	35

	<i>Page</i>
13.3 Deactivation and reactivation function.....	36
13.4 Group function	37
13.5 Replication function	37
13.6 Migration function	37
13.7 Transaction function.....	38
13.8 ACID transaction function	38
13.9 Engineering interface reference tracking function	39
14 Repository functions.....	39
14.1 Storage function	39
14.2 Information organization function.....	39
14.3 Relocation function	40
14.4 Type repository function	40
14.5 Trading function.....	41
15 Security functions.....	41
15.1 Concepts.....	41
15.2 Access control function.....	42
15.3 Security audit function	42
15.4 Authentication function.....	42
15.5 Integrity function.....	43
15.6 Confidentiality function	43
15.7 Non-repudiation function	43
15.8 Key management function	44
16 ODP distribution transparency	44
16.1 Access transparency	45
16.2 Failure transparency	45
16.3 Location transparency	46
16.4 Migration transparency	46
16.5 Persistence transparency	46
16.6 Relocation transparency	46
16.7 Replication transparency.....	47
16.8 Transaction transparency.....	47
Annex A – Formal computational supertype/subtype rules.....	48
A.1 Notations and conventions	48
A.2 Type system.....	48
A.3 Signal interface signature types.....	51
A.4 Operation interface signature types.....	52
A.5 Stream interface types	52
A.6 Example	52
Annex B – Human-computer interactions	54
B.1 Specifying human/system interactions	54

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of the joint technical committee is to prepare International Standards. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

ISO/IEC 10746-3 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 7, *Software and systems engineering*, in collaboration with ITU-T. The identical text is published as Rec. ITU-T X.903 (10/2009).

This second edition cancels and replaces the first edition (ISO/IEC 10746-3:1996), which has been technically revised.

ISO/IEC 10746 consists of the following parts, under the general title *Information technology — Open distributed processing — Reference model*:

- *Part 1: Overview*
- *Part 2: Foundations*
- *Part 3: Architecture*
- *Part 4: Architectural semantics*

Introduction

The rapid growth of distributed processing has led to a need for a coordinating framework for the standardization of open distributed processing (ODP). This reference model provides such a framework. It creates an architecture within which support of distribution, interworking and portability can be integrated.

The reference model of open distributed processing, Recommendations ITU-T X.901 | ISO/IEC 10746-1 to X.904 | ISO/IEC 10746-4, is based on precise concepts derived from current distributed processing developments and, as far as possible, on the use of formal description techniques for specification of the architecture.

The reference model consists of:

- Recommendation ITU-T X.901 | ISO/IEC 10746-1: Overview: Contains a motivational overview of ODP giving scoping, justification and explanation of key concepts, and an outline of the ODP architecture. It contains explanatory material on how this reference model is to be interpreted and applied by its users, who may include standards writers and architects of ODP systems. It also contains a categorization of required areas of standardization expressed in terms of the reference points for conformance identified in this Recommendation | International Standard. This part is not normative.
- Recommendation ITU-T X.902 | ISO/IEC 10746-2: Foundations: Contains the definition of the concepts and analytical framework for normalized description of (arbitrary) distributed processing systems. It introduces the principles of conformance to ODP standards and the way in which they are applied. This is only to a level of detail sufficient to support this Recommendation | International Standard and to establish requirements for new specification techniques. This part is normative.
- Recommendation ITU-T X.903 | ISO/IEC 10746-3: Architecture: Contains the specification of the required characteristics that qualify distributed processing as open. These are the constraints to which ODP standards must conform. It uses the descriptive techniques from Recommendation ITU-T X.902 | ISO/IEC 10746-2. This part is normative.
- Recommendation ITU-T X.904 | ISO/IEC 10746-4: Architectural semantics: Contains a formalization of the ODP modelling concepts defined in Recommendation ITU-T X.902 | ISO/IEC 10746-2, clauses 8 and 9. The formalization is achieved by interpreting each concept in terms of the constructs of the different standardized formal description techniques. This part is normative.

This Recommendation | International Standard contains two annexes:

- Annex A – Formal computational supertype/subtype rules.
- Annex B – Human-computer interactions

Annex A forms an integral part of the reference model. Annex B does not form an integral part of the reference model.

INTERNATIONAL STANDARD RECOMMENDATION ITU-T

Information technology – Open distributed processing – Reference model: Architecture

1 Scope

This Recommendation | International Standard:

- defines how ODP systems are specified, making use of concepts in Rec. ITU-T X.902 | ISO/IEC 10746-2;
- identifies the characteristics that qualify systems as ODP systems.

It establishes a framework for coordinating the development of existing and future standards for ODP systems and is provided for reference by those standards.

2 Normative references

The following Recommendations and International Standards contain provisions which, through reference in this text, constitute provisions of this Recommendation | International Standard. At the time of publication, the editions indicated were valid. All Recommendations and Standards are subject to revision, and parties to agreements based on this Recommendation | International Standard are encouraged to investigate the possibility of applying the most recent edition of the Recommendations and Standards listed below. Members of IEC and ISO maintain registers of currently valid International Standards. The Telecommunication Standardization Bureau of the ITU maintains a list of currently valid ITU-T Recommendations.

2.1 Identical Recommendations | International Standards

- Recommendation ITU-T X.200 (1994) | ISO/IEC 7498-1:1994, *Information technology – Open Systems Interconnection – Basic Reference Model: The basic model*.
- Recommendation ITU-T X.810 (1995) | ISO/IEC 10181-1:1996, *Information technology – Open Systems Interconnection – Security frameworks for open systems: Overview*.
- Recommendation ITU-T X.811 (1995) | ISO/IEC 10181-2:1996, *Information technology – Open Systems Interconnection – Security frameworks for open systems: Authentication framework*.
- Recommendation ITU-T X.812 (1995) | ISO/IEC 10181-3:1996, *Information technology – Open Systems Interconnection – Security frameworks for open systems: Access control framework*.
- Recommendation ITU-T X.813 (1996) | ISO/IEC 10181-4:1997, *Information technology – Open Systems Interconnection – Security frameworks for open systems: Non-repudiation framework*.
- Recommendation ITU-T X.814 (1995) | ISO/IEC 10181-5:1996, *Information technology – Open Systems Interconnection – Security frameworks for open systems: Confidentiality framework*.
- Recommendation ITU-T X.815 (1995) | ISO/IEC 10181-6:1996, *Information technology – Open Systems Interconnection – Security frameworks for open systems: Integrity framework*.
- Recommendation ITU-T X.816 (1995) | ISO/IEC 10181-7:1996, *Information technology – Open Systems Interconnection – Security frameworks for open systems: Security audit and alarms framework*.
- Recommendation ITU-T X.851 (1997) | ISO/IEC 9804:1998, *Information technology – Open Systems Interconnection – Service definition for the Commitment, Concurrency and Recovery service element*.
- Recommendation ITU-T X.902 (2009) | ISO/IEC 10746-2:2010, *Information technology – Open distributed processing – Reference model: Foundations*.
- ISO/IEC 11770-1:1996, *Information technology – Security techniques – Key management – Part 1: Framework*.

2.2 Paired Recommendations | International Standards equivalent in technical content

- Recommendation ITU-T X.800 (1991), *Security architecture for Open Systems Interconnection for CCITT applications*.
ISO 7498-2:1989, *Information processing systems – Open Systems Interconnection – Basic Reference Model – Part 2: Security Architecture*.

3 Definitions

For the purposes of this Recommendation | International Standard, the following definitions apply.

3.1 Descriptive definitions

This reference model makes use of the following term defined in Rec. ITU-T X.200 | ISO/IEC 7498-1:

- transfer syntax.

This Recommendation | International Standard makes use of the following terms defined in Rec. ITU-T X.810 | ISO/IEC 10181-1:

- trusted third party.

This Recommendation | International Standard makes use of the following terms defined in Rec. ITU-T X.811 | ISO/IEC 10181-2:

- claimant;
- exchange authentication information;
- principal.

This Recommendation | International Standard makes use of the following terms defined in Rec. ITU-T X.812 | ISO/IEC 10181-3:

- access control information;
- access control decision function;
- access control enforcement function;
- initiator;
- target.

This Recommendation | International Standard makes use of the following terms defined in Rec. ITU-T X.813 | ISO/IEC 10181-4:

- evidence generator;
- evidence user;
- evidence verifier;
- (non-repudiable data) originator;
- (non-repudiable data) recipient;
- non-repudiation evidence;
- non-repudiation service requester;
- notary;
- originator;
- recipient.

This Recommendation | International Standard makes use of the following terms defined in Rec. ITU-T X.814 | ISO/IEC 10181-5:

- confidentiality-protected data;
- hide;
- reveal.

This Recommendation | International Standard makes use of the following terms defined in Rec. ITU-T X.815 | ISO/IEC 10181-6:

- integrity-protected data;
- shield;
- validate.

This Recommendation | International Standard makes use of the following terms defined in Rec. ITU-T X.816 | ISO/IEC 10181-7:

- audit recorder function;
- audit trail examiner function;
- audit trail collector function.

This Recommendation | International Standard makes use of the following terms defined in Rec. ITU-T X.851 | ISO/IEC 9804:

- atomicity;
- consistency;
- durability;
- isolation.

This Recommendation | International Standard makes use of the following terms defined in ISO/IEC 11770-1 Key Management Framework:

- certification authority;
- key distribution centre;
- key management;
- key translation centre;
- public key.

This reference model makes use of the terms defined in Rec. ITU-T X.902 | ISO/IEC 10746-2, shown in Figure 1.

3.2 Abbreviations

For the purposes of this Recommendation | International Standard, the following abbreviations apply:

ACID	Atomicity, Consistency, Isolation and Durability
ODP	Open Distributed Processing
OSI	Open Systems Interconnection

4 Framework

This reference model defines a framework comprising:

- five *viewpoints*, called enterprise, information, computational, engineering and technology which provide a basis for the specification of ODP systems;
- a *viewpoint language* for each viewpoint, defining concepts and rules for specifying ODP systems from the corresponding viewpoint;
- specifications of the *functions* required to support ODP systems;
- *transparency prescriptions* showing how to use the ODP functions to achieve distribution transparency.

The architecture for ODP systems and the composition of functions is determined by the combination of the computational language, the engineering language and the transparency prescriptions.

abstraction;	fault;	persistence;
action;	factory;	policy;
action signature;	<x> group;	producer object;
activity;	identifier;	programmatic reference point;
architecture;	information;	prohibition;
atomicity;	initiating object;	Quality of Service;
behaviour;	instance;	reference point;
binding;	instantiation;	refinement;
class;	interaction;	role;
client object;	interchange reference point;	rule;
communication;	interface;	server object;
communication management;	interface signature;	spawn action;
composition;	interworking reference point;	specification;
component;	introduction;	stability;
configuration;	invariant;	state;
conformance point;	liaison;	subdomain;
consumer object;	location in space;	subtype;
contract;	location in time;	supertype;
creation;	name;	system;
data;	naming context;	<x> template;
decomposition;	naming domain;	term;
deletion;	notification;	thread;
distributed processing;	object;	trading;
distribution transparency;	obligation;	type;
<x> domain;	ODP standards;	viewpoint.
entity;	ODP system;	
environment;	open distributed processing;	
error;	perceptual reference point;	
establishing behaviour;	permission;	
failure;		

Figure 1 – Terms taken from Rec. ITU-T X.902 | ISO/IEC 10746-2

4.1 Viewpoints

4.1.1 Concepts

4.1.1.1 enterprise viewpoint: A viewpoint on an ODP system and its environment that focuses on the purpose, scope and policies for that system.

4.1.1.2 information viewpoint: A viewpoint on an ODP system and its environment that focuses on the semantics of information and information processing.

4.1.1.3 computational viewpoint: A viewpoint on an ODP system and its environment which enables distribution through functional decomposition of the system into objects which interact at interfaces.

4.1.1.4 engineering viewpoint: A viewpoint on an ODP system and its environment that focuses on the mechanisms and functions required to support distributed interaction between objects in the system.

4.1.1.5 technology viewpoint: A viewpoint on an ODP system and its environment that focuses on the choice of technology in that system.

4.1.2 Using viewpoints

The enterprise, information, computational, engineering and technology viewpoints have been chosen as a necessary and sufficient set to meet the needs of ODP standards. Viewpoints can be applied, at an appropriate level of abstraction,

to a complete ODP system, in which case the environment defines the context in which the ODP system operates. Viewpoints can also be applied to individual components of an ODP system, in which case the component's environment will include some abstraction of both the system's environment and other system components.

NOTE – The process of abstraction might be such that the system's environment and the other system components are composed into a single object.

4.2 ODP viewpoint languages

4.2.1 Concepts

4.2.1.1 <Viewpoint> language: Definitions of concepts and rules for the specification of an ODP system from the <viewpoint> viewpoint; thus: **engineering language:** definitions of concepts and rules for the specification of an ODP system from the engineering viewpoint.

4.2.2 Using viewpoint languages

This reference model defines a set of five languages, each corresponding to one of the viewpoints defined in 4.1.1. Each language is used for the specification of an ODP system from the corresponding viewpoint. These languages are:

- the enterprise language (defined in clause 5);
- the information language (defined in clause 6);
- the computational language (defined in clause 7);
- the engineering language (defined in clause 8);
- the technology language (defined in clause 9).

Each language uses concepts taken from Rec. ITU-T X.902 | ISO/IEC 10746-2, and introduces refinements of those concepts, prescriptive rules and additional viewpoint-specific concepts relevant to the nature of the specifications concerned. These additional concepts are, in turn, defined using concepts from Rec. ITU-T X.902 | ISO/IEC 10746-2.

A system specification comprises one or more viewpoint specifications. These specifications must be mutually consistent. Rules for the consistent structuring of viewpoint specifications are given in clause 10. The specifier must demonstrate by other means that terms in the specifications are used consistently. A specification of a system using several viewpoint specifications will often restrict implementations more than a specification using fewer viewpoint specifications. Objects identified in one viewpoint can be specified using the viewpoint language associated with that viewpoint or using the viewpoint languages associated with other viewpoints. It is not necessary to specify an object fully from every viewpoint in order to achieve a mutually consistent set of viewpoint specifications.

NOTE 1 – The qualification of a term from Rec. ITU-T X.902 | ISO/IEC 10746-2 by the name of a viewpoint (e.g., as in "**computational** object") is interpreted as using the term from Rec. ITU-T X.902 | ISO/IEC 10746-2, subject to whatever additional provisions are specified in the identified viewpoint language.

NOTE 2 – The unqualified use of a term from Rec. ITU-T X.902 | ISO/IEC 10746-2 in a viewpoint specification (e.g., "interface") is interpreted as if the term had been qualified by the name of the viewpoint (i.e., "computational interface"), if the associated viewpoint language places additional constraints on the term.

4.3 ODP functions

4.3.1 ODP function: A function required to support open distributed processing.

4.3.2 Using ODP functions

This reference model specifies, in clauses 11 to 15, the functions required to achieve open distributed processing.

Each ODP function description contains:

- an explanation of the use of the function for open distributed processing;
- prescriptive statements, about the structure and behaviour of the function, sufficient to ensure the overall integrity of the reference model;
- a statement of other ODP functions upon which it depends.

4.4 ODP distribution transparencies

4.4.1 Concepts

4.4.1.1 access transparency: A distribution transparency which masks differences in data representation and invocation mechanisms to enable interworking between objects.

4.4.1.2 failure transparency: A distribution transparency which masks, from an object, the failure and possible recovery of other objects (or itself), to enable fault tolerance.

4.4.1.3 location transparency: A distribution transparency which masks the use of information about location in space when identifying and binding to interfaces.

4.4.1.4 migration transparency: A distribution transparency which masks, from an object, the ability of a system to change the location of that object. Migration is often used to achieve load balancing and reduce latency.

4.4.1.5 relocation transparency: A distribution transparency which masks relocation of an interface from other interfaces bound to it.

4.4.1.6 replication transparency: A distribution transparency which masks the use of a group of mutually behaviourally compatible objects to support an interface. Replication is often used to enhance performance and availability.

4.4.1.7 persistence transparency: A distribution transparency which masks, from an object, the deactivation and reactivation of other objects (or itself). Deactivation and reactivation are often used to maintain the persistence of an object when a system is unable to provide it with processing, storage and communication functions continuously.

4.4.1.8 transaction transparency: A distribution transparency which masks coordination of activities amongst a configuration of objects to achieve consistency.

4.4.2 Using distribution transparency

Distribution transparency is an important end-user requirement in distributed systems. This reference model defines a set of distribution transparencies which make it possible to implement ODP systems which are distribution transparent from the point of view of users of those systems. Distribution transparency is selective; the reference model includes rules for selecting and combining distribution transparencies in ODP systems.

This reference model requires, for each distribution transparency defined in 4.4.1.1 to 4.4.1.8, definitions of both:

- a transparency schema for expressing requirements for the particular transparency – the schema is specified in the enterprise, information computational specifications;
- a refinement process for transforming a specification which contains requirements for the particular distribution transparency to a specification which explicitly realizes the masking implied by that transparency – the refinement process is specified in the engineering specification.

NOTE 1 – In some cases (e.g., access transparency) the schema is null; in others (e.g., transaction transparency) the schema contains one or more parameters dictating the precise form of transparency required.

NOTE 2 – The refinement process typically involves introducing additional behaviour, including the use of one or more ODP functions, into the specification.

The specifications of the refinement processes in clause 16 are prescriptive to the level required to ensure overall integrity of the reference model.

4.5 Standards derived from the framework

This reference model provides a framework for the definition of new standards and the use of existing standards as ODP standards.

ODP standards are any of:

- standards for components of ODP systems;
- standards for composing ODP system components;
- standards for modelling and specifying ODP systems.

ODP standards:

- use the enterprise language to specify policies;
- use the information language to specify consistent use and interpretation of information in, and between, standards;
- use the computational language to specify the configuration and behaviour of interfaces;
- use the engineering language to specify the infrastructures they require;
- use the technology language to specify conformance to international, private, or consensual specifications.

Standards for methodology, modelling, programming, implementation and testing of ODP systems use the framework as a whole.

ODP standards can be based on a subset of this reference model (e.g., by excluding some forms of interaction, particular functions or transparencies). Such standards can also extend this reference model, provided that the extensions they introduce do not change or contradict its provisions. Extensions will relate new terms to terms defined in this reference model: for example, by introducing new types and new type rules.

ODP standards comply with all the prescriptive statements in this reference model.

4.6 Conformance

The enterprise, information, computational and engineering languages are used to specify the conformance requirements for ODP systems. The technology language can be used to assert conformance to ODP standards in ODP systems. Each interface which is defined as a conformance point has an information specification to enable interpretation of interactions of that interface. The rules for identifying conformance points are given in the computational and engineering languages.

An ODP system conforms to an ODP standard if it satisfies the conformance requirements of that standard.

5 Enterprise language

The enterprise language comprises concepts, rules and structures for the specification of an ODP system from the enterprise viewpoint.

An enterprise specification defines the purpose, scope and policies of an ODP system.

In this reference model, prescription in the enterprise viewpoint is restricted to a basic set of concepts and rules addressing the scope and nature of enterprise specifications.

5.1 Concepts

The enterprise language contains the concepts of Rec. ITU-T X.902 | ISO/IEC 10746-2 and those defined here, subject to the rules of 5.2.

5.1.1 community: A configuration of objects formed to meet an objective. The objective is expressed in a contract, which expresses how this objective can be met by defining roles and interactions required, assignments of objects to the roles, and policies governing the collective behaviour.

5.1.2 <X> federation: A community of <x> domains.

5.2 Structuring rules

An enterprise specification defines, and the enterprise language is able to express, the purpose, scope and policies of an ODP system in terms of each of the following items:

- roles played by the system;
- activities undertaken by the system;
- policy statements about the system, including those relating to environment contracts.

In an enterprise specification, an ODP system and the environment in which it operates are represented as a community. At some level of description the ODP system is represented as an enterprise object in the community. The objectives and scope of the ODP system are defined in terms of the roles it fulfils within the community of which it is part, and policy statements about those roles. A community is defined in terms of each of the following elements:

- the enterprise objects comprising the community;
- the roles fulfilled by each of those objects;
- policies governing interactions between enterprise objects fulfilling roles;
- policies governing the creation, usage and deletion of resources by enterprise objects fulfilling roles;
- policies governing the configuration of enterprise objects and assignment of roles to enterprise objects;
- policies relating to environment contracts governing the system.

NOTE 1 – Enterprise policies may constrain some of the distribution transparencies defined in the computational specifications. For example, an enterprise policy on avoidance of single points of hardware failure might imply a requirement for migration transparency for all computational objects, irrespective of their individual constraints in the computational specification.

A role is defined in terms of the permissions, obligations, prohibitions and behaviour of the enterprise object fulfilling the role. An enterprise object can fulfil one or more roles in a community, and the roles which it can fulfil are determined by the contract on which the community is based. While it is part of one community, the enterprise object can continue to fulfil roles in other communities, subject to the provisions in the contracts of the communities involved. The enterprise object can fulfil different roles in different communities. Interactions between enterprise objects fulfilling appropriate roles within different communities can be considered as interactions between those communities.

NOTE 2 – Examples of roles include policy administrator, president, service provider, owner, manager, shareholder, consumer.

NOTE 3 – Examples of environment contracts in enterprise specifications include safety requirements, legislative requirements and codes of practice.

NOTE 4 – In an enterprise specification the term "<x> object", where <x> is a role, is interpreted as meaning "an enterprise object fulfilling an <x> role"; where an enterprise object fulfils multiple roles, the names can be concatenated, e.g., "owner driver object".

When fulfilling a role, an object becomes subject to permissions, obligations and prohibitions by delegation or transfer. In some roles, objects are permitted to change policy. There are five fundamental types of actions with respect to contractual matters:

- an object incurs an obligation to another object (it must currently be permitted to incur the obligation);
- an object fulfils an obligation to another object;
- an object waives an obligation to another object;
- an object acquires permission from another object to perform some action it was previously forbidden to perform;
- an object is forbidden to perform an action it was previously permitted to perform.

NOTE 5 – An important special case of acquisition is where the permitted action is performative, i.e., when an object in a subordinate role is enabled to issue further permissions or obligations on behalf of an object fulfilling a superior role. This leads to the notion of agency or delegation.

Obligations include accounting and charging for the use of resources. Billing and payment is modelled as the reassignment of resources between objects in accordance with the roles they fulfil.

A resource is either consumable or non-consumable. A consumable resource deletes itself after some amount of use. In <x> federation, the objective defines the resources each <x> domain in the federation shares with the other members of the federation. The objective can leave each domain with a defined degree of autonomy in the use of its own resources. The establishing behaviour for an <x> federation can allow autonomy for each participating <x> domain in deciding whether or not to become part of the federation.

5.3 Conformance and reference points

Conformance statements in the enterprise language require that the behaviour of an ODP system is conformant to a particular set of objectives and policies.

An implementor claiming conformance must identify the engineering reference points which give access to the system and the engineering, computational and information specifications which apply at them. By this act the identified reference points become conformance points. The interactions at these conformance points can then be interpreted in enterprise language terms to check that the enterprise specification is not violated.

Enterprise specifications can be applied to all four classes of reference point (programmatic, perceptual, interworking and interchange reference points) identified in Rec. ITU-T X.902 | ISO/IEC 10746-2.

6 Information language

The information language comprises concepts, rules and structures for the specification of an ODP system from the information viewpoint.

An information specification defines the semantics of information and the semantics of information processing in an ODP system.

In this reference model, prescription in the information viewpoint is restricted to a basic set of concepts and rules addressing the scope and nature of information specifications.

NOTE – The information specification includes the specification of all information held by the systems: however, it does not specify the representation of that information in the system itself, although there is, necessarily, a notation for representation of the information specification. Some representations will be determined by the nature of the platform (e.g., transfer formats); others will be to do with representation for human consumption.

6.1 Concepts

The information language contains the concepts of Rec. ITU-T X.902 | ISO/IEC 10746-2 and those defined here, subject to the rules of 6.2.

6.1.1 invariant schema: A set of predicates on one or more information objects which must always be true. The predicates constrain the possible states and state changes of the objects to which they apply.

NOTE 1 – Thus, an invariant schema is the specification of the types of one or more information objects that will always be satisfied by whatever behaviour the objects might exhibit.

NOTE 2 – An invariant schemata may constrain some of the distribution transparencies defined in the computational specifications. For example, a particular information object type may place a constraint that all the information objects that are instances of it must be accessed using transaction transparency, without the need to so qualify all such interactions in the computational specification.

6.1.2 static schema: A specification of the state of one or more information objects, at some point in time, subject to the constraints of any invariant schemata.

NOTE – Thus, a static schema is the specification of the types of one or more information objects at some particular point in time. These types are subtypes of the types specified in the invariant schema.

6.1.3 dynamic schema: A specification of the allowable state changes of one or more information objects, subject to the constraints of any invariant schemata.

NOTE 1 – Behaviour in an information system can be modelled as transitions from one static schema to another, i.e., reclassification of instances from one type to another.

NOTE 2 – In the information language, a state change involving a set of objects can be regarded as an interaction between those objects. Not all of the objects involved in the interaction need change state; some of the objects may be involved in a read-only manner.

6.2 Structuring rules

An information specification defines the semantics of information and the semantics of information processing in an ODP system in terms of a configuration of information objects, the behaviour of those objects and environment contracts for the system.

An information object template references static, invariant and dynamic schemata. The relationships between information objects can be modelled as part of the state of those information objects. Information objects are either atomic or are represented as a composition of component information objects. The state of the composite object is represented by the combined state of its component information objects. An atomic information object template represents a concept for which there is no model at a particular level of abstraction. A composite information object represents a derived concept expressed in terms of other concepts. Since object composition includes encapsulation, an information object that is a component of one composite object cannot be a component of another. Therefore, information objects resulting from the instantiation of a composite information object template only exist as part of the instantiated composite object and have no meaning outside it.

Allowable state changes specified by a dynamic schema can include the creation of new information objects and the deletion of information objects involved in the dynamic schema. Allowable state changes can be subject to ordering and temporal constraints.

NOTE 1 – The result of accessing the state of one or more information objects can be modelled as the creation of a new information object.

In an information specification, the configuration of information objects and the behaviour of those objects need not be suitable for distribution (e.g., there need not be any concept of failure or location for information interactions).

NOTE 2 – If an information notation uses the concept of interface, any interfaces defined cannot themselves be reference points; thus there is no commitment to the interfaces appearing in an implementation.

6.3 Conformance and reference points

Conformance statements in information specifications require that the behaviour of an ODP system is conformant to a particular set of invariant, static and dynamic schemata.

An implementor claiming conformance must list the relevant engineering reference points which give access to the system and the engineering and computational specifications which apply to them. By this act the identified reference points become conformance points. The interactions at these conformance points can then be interpreted in information language terms to check that they are consistent with the invariant, static and dynamic schemata.

Information specifications can be applied to all four classes of reference point (programmatic, perceptual, interworking, and interchange reference points) identified in Rec. ITU-T X.902 | ISO/IEC 10746-2.

7 Computational language

The computational language comprises concepts, rules and structures for the specification of an ODP system from the computational viewpoint.

NOTE – The computational language defines an abstract object model – the virtual machine that interprets the computational design and thus has to be realized by any supporting middleware.

A computational specification defines the functional decomposition of an ODP system into objects which interact at interfaces.

In the computational viewpoint, applications and ODP functions consist of configurations of interacting computational objects.

The interactions defined by the specification are constrained by a set of distribution transparencies that are defined in transparency schemata contained in the specification.

7.1 Concepts

The computation language contains the concepts of Rec. ITU-T X.902 | ISO/IEC 10746-2 and those defined here, subject to the structuring rules of 7.2. The trading function (see 14.5) is an integral part of the computational language; the other ODP functions (see 11) are included in the language as needed.

7.1.1 signal: An atomic shared action resulting in one-way communication from an initiating object to a responding object.

NOTE 1 – A signal is an interaction.

NOTE 2 – An example of a signal is the initiation of an event notification by the sending interface or the receipt of the event notification by the receiving interface. This assumes that a compound binding is being considered, so that these two cases are distinct. If the binding were primitive, sending and receipt could not be distinguished.

7.1.2 operation: An interaction between a client object and a server object which is either an interrogation or an announcement.

NOTE – An operation can be refined in terms of signals (see 7.2.2.5).

7.1.3 announcement: An interaction – the invocation – initiated by a client object resulting in the conveyance of information from that client object to a server object, requesting a function to be performed by that server object.

NOTE – An example of an announcement is the initiation of the sending of a message in a messaging system.

7.1.4 interrogation: An interaction consisting of

- one interaction – the invocation – initiated by a client object, resulting in the conveyance of information from that client object to a server object, requesting a function to be performed by the server object, followed by
- a second interaction – the termination – initiated by the server object, resulting in the conveyance of information from the server object to the client object in response to the invocation.

NOTE 1 – In interrogations, invocations and terminations are always paired. Announcements do not have terminations. Thus there is no possibility of an operation consisting of an invocation followed by a sequence of associated terminations.

NOTE 2 – An example of an interrogation is a request/response exchange.

7.1.5 flow: An abstraction of a sequence of interactions, resulting in conveyance of information from a producer object to a consumer object.

NOTE 1 – A flow may be used to abstract over, for example, the exact structure of a sequence of interactions, or over a continuous interaction including the special case of an analogue information flow.

NOTE 2 – A flow can be refined in terms of signals (see 7.2.2.5). However, a flow would not be defined as a composition of signals and operations.

NOTE 3 – An example of a flow is a multimedia data broadcast.

7.1.6 signal interface: An interface in which all the interactions are signals.

7.1.7 operation interface: An interface in which all the interactions are operations.

7.1.8 stream interface: An interface in which all the interactions are flows.

7.1.9 computational object template: An object template which comprises a set of computational interface templates which the object can instantiate, a behaviour specification and an environment contract specification.

7.1.10 computational interface template: An interface template for either a signal interface, a stream interface or an operation interface. A computational interface template comprises a signal, a stream or an operation interface signature as appropriate, a behaviour specification and an environment contract specification.

7.1.11 signal interface signature: An interface signature for a signal interface. A signal interface signature comprises a finite set of action signatures, one for each signal type in the interface. Each action signature comprises the name for the signal, the number, names and types of its parameters and an indication of causality (initiating or responding, but not both) with respect to the object which instantiates the signal interface.

NOTE – The inclusion, in an action signature, of information about the number of parameters is optional.

7.1.12 operation interface signature: An interface signature for an operation interface. An operation interface signature comprises a set of announcement and interrogation signatures as appropriate, one for each operation type in the interface, together with an indication of causality (client or server, but not both) for the interface as a whole, with respect to the object which instantiates the operation interface.

Each announcement signature is an action signature containing both the name of the invocation and the number, names and types of its parameters.

Each interrogation signature comprises an action signature with the following elements:

- the name of the invocation;
- the number, names and types of its parameters;
- a finite, non-empty set of action signatures, one for each possible termination type of the invocation, each containing both the name of the termination and the number, names and types of its parameters.

NOTE – Inclusion, in announcement and interrogation signatures, of the information about the number of parameters is optional.

7.1.13 stream interface signature: An interface signature for a stream interface. A stream interface comprises a finite set of action signatures, one for each flow in the stream interface. Each action signature for a flow contains the name of the flow, the information type of the flow, and an indication of causality for the flow (i.e., producer or consumer but not both) with respect to the object which instantiates the stream interface.

NOTE 1 – The phrase "complementary interface signature to *X*", where *X* is itself an interface signature describes an interface signature identical to *X* in all respects except causality, which is opposite to that in *X*.

NOTE 2 – Many interface definition languages (IDLs) do not capture the indication of causality and depend upon the context in which the IDL is used to determine the causality that is to be applied.

NOTE 3 – Causalities are defined for objects, to categorize their roles in interactions (see clause 14 of Rec. ITU-T X.902 | ISO/IEC 10746-2), and not for individual actions. Therefore, the causalities indicated in the interaction signatures (for signal and flows) and interface signatures (for operations) refer to the causality of the object instantiating the interface: initiator/responder, producer/consumer, or client/server (see also 7.2.2).

7.1.14 binding object: A computational object which supports a binding between a set of other computational objects.

NOTE 1 – Binding objects are subject to special provisions (see 7.2.3).

NOTE 2 – A signal may be introduced in the context of interfaces to binding objects where there is a requirement to localize interfaces for the purpose of specifying QoS.

7.1.15 computational factory: A factory that returns an interface reference to the computational object it creates.

7.1.16 computational component: A component specified in a computational viewpoint that has a control interface and declared sets of:

- operation interfaces in which it plays a server role (*facets*);
- operation interfaces in which it plays a client role (*receptacles*);
- operation interfaces originating announcements carrying notifications of typed events (*event sources*);
- operation interfaces consuming announcements carrying notifications of typed events (*event sinks*);
- operation interfaces supporting accessors and mutators for attributes (*attributes*).

The interfaces in these sets are called *features* of the component.

The control interface is an interface that provides at least operations that access the component template and retrieve a list of features and their interface references (e.g., to allow component introspection and feature navigation).

NOTE – The shorthand names in italics above correspond to the terms used in the CORBA component model (CCM)

7.1.17 computational container: A container for a declared set of computational component types that has a management interface at which it can be requested to:

- create a computational component of one of the types and add it to the container's content;
- delete a computational component from the container;
- list the computational components it currently contains;
- list the computational factories it provides.

7.2 Structuring rules

A computational specification describes the functional decomposition of an ODP system, in distribution transparent terms, as:

- a configuration of computational objects (including binding objects);
- the internal actions of those objects;
- the interactions that occur among those objects;
- environment contracts for those objects and their interfaces;
- the selection of distribution transparencies required.

A computational specification is constrained by the rules of the computational language. These comprise:

- interaction rules (see 7.2.2), binding rules (see 7.2.3) and type rules (see 7.2.4) that provide distribution transparent interworking;
- template rules (see 7.2.5) that apply to all computational objects;
- failure rules (see 7.2.6) that apply to all computational objects and identify the potential points of failure in computational activities.

Portability rules (see 7.2.7) are provided to give guidance to developers of ODP portability standards.

A computational specification defines an initial set of computational objects and their behaviour. The configuration will change as the computational objects:

- instantiate further computational objects;
- instantiate further computational interfaces;
- perform binding actions;
- effect control functions upon binding objects;
- delete computational interfaces;
- delete computational objects.

NOTE – Except in failure action, computational objects only delete themselves or their own computational interfaces.

7.2.1 Naming rules

Each kind of name defined in the computational language has an associated context, as follows:

- a signal name in a signal interface signature is an identifier in the context of that signature;
- a flow name in a stream interface signature is an identifier in the context of that signature;
- an invocation name in an operation interface signature is an identifier in the context of that signature;
- a termination name in an operation interface signature is an identifier in the context of the operation signature in which it appears;
- the name of a parameter in a signal signature is an identifier in the context of that signature;
- the name of a parameter in an invocation signature in an operation interface signature is an identifier in the context of that signature;
- the name of a parameter in a termination signature in an operation interface signature is an identifier in the context of that signature;
- the name of a parameter in a signal signature in a signal interface signature is an identifier in the context of that signature.

NOTE 1 – Thus signal names are distinct in any signal interface signature but signals in different signatures can have the same name, and so forth.

A computational interface identifier is unambiguous within its context (i.e., cannot be associated with more than one computational interface in that context). The choice of contexts for computational interface identifiers is a matter of language design and therefore beyond the scope of this reference model. Hence the reference model places no constraints on the extent of contexts for computational interface identifiers. Thus no reliance can be placed on:

- the extent of naming contexts for computational interface identifiers (e.g., assumptions about them being related to engineering language structures such as nodes or communications domains);
- the uniqueness of computational interface identifiers (i.e., synonyms are permitted);
- a computational interface identifier denoting the same computational interface in all places where the identifier appears (i.e., names need not be "global").

NOTE 2 – A particular computational notation may not have explicit terms denoting computational identifiers; therefore in that notation computational interface identifiers are implicit; however they remain subject to the rules given above.

7.2.2 Interaction rules

Each interaction of a computational object occurs at one of its computational interfaces. The computational language imposes constraints on the behaviour permitted at a computational interface. Interaction at an unbound interface fails. The binding rules (see 7.2.3) impose constraints on how interfaces can be bound.

The interaction part of the computational language supports three models of interaction, each of which has an associated kind of computational interface:

- signals and signal interfaces;
- flows and stream interfaces;
- operations and operation interfaces.

In addition to the different kinds of interaction supported, the interaction models differ in their failure properties. The participants in a flow or operation can have an inconsistent view of an interaction at different times, especially when failures have occurred. In contrast to streams and operations, there is no concept of partial failure of a signal – a signal either succeeds or fails identically for both participants in the interaction.

7.2.2.1 Signal interaction rules

A computational object offering a signal interface of a given signal interface type:

- initiates signals that have initiating causality in the interface's signature;
- responds to signals that have responding causality in the interface's signature.

NOTE – Signals are generally used when it is necessary to have a less abstract representation of interfaces, including stream and operation interfaces (see 7.2.2.5).

7.2.2.2 Stream interaction rules

A computational object offering a stream interface:

- generates flows that have producer causality in the interface's signature;
- receives flows that have consumer causality in the interface's signature.

7.2.2.3 Operation interaction rules

A client object using an operation interface invokes the operations named in the interface's signature. A server object offering an operation interface expects any of the operations named in the interface's signature. In the case of an interrogation, the server responds to the invocation by initiating any one of the terminations named for the operation in the server interface signature. The client expects any of the terminations named for the operation in the client interface signature. The duration of the operation is arbitrary unless required otherwise by environment contracts applicable to the objects and interfaces involved.

NOTE 1 – If a client thread invokes a chain of interrogations, the two way "handshake" of invocation and termination ensures that operations will be responded to by the server in the same order as the client initiated them. If the client invokes a chain of announcements (or a chain containing both announcements and invocations) there is no handshake to guarantee the order in which the server responds to the announcement, unless this is implicit in the environment contracts applicable to the interaction. There are no ordering guarantees for either interrogations or announcements that are in different descendant activities of a previous dividing action.

NOTE 2 – The behaviour of an object using an operation interface may vary, depending on the type of the interface. A client object may be blocked until it gets one of the terminations (e.g., in synchronous messaging), or a client object may continue its activity (e.g., in asynchronous messaging), and expect to receive one of the terminations some time later with an appropriate mechanism, such as by an introduction of a message management object with a message management interface, providing e.g., the following functions:

- passing invocations/terminations in a message management object;
- asking for delivery of invocations/terminations via a message management object;
- retrieving invocations/terminations from a message management object;
- receiving signals indicating an arrival of an invocation/termination from a message management object.

7.2.2.4 Parameter rules

The parameters for signals, invocations and terminations can include identifiers for computational interfaces and computational interface signature types.

NOTE 1 – The possibility of parameters expressing computational interface types makes the computational signature type systems higher order. Explicit representation of signature types is required, for example in trading, where the parameters of import and export operations include signature types:

```
trader.import (T: Type, ....) : (service: T) -> failed (reason: String)
trader.export (T: Type, service: T) : (....) -> failed (reason: String)
```

This introduces a need for dynamic signature subtype checking (see 7.2.5.1).

A formal parameter that is an identifier for a computational interface is qualified by a computational interface signature type. The corresponding actual parameter must reference an interface with that interface signature type (or one of its subtypes). The actual parameter can only be used as if it referenced a computational interface with the same signature type as the formal parameter (or one of the formal parameter's supertypes). After an interaction, both initiator and responder can reference the identified interface, although possibly by different computational interface identifiers.

NOTE 2 – This rule prevents the user of the interface referenced by the actual parameter being able to perform additional interactions beyond those in the formal interface signature type, even if the interface referenced by the actual parameter is a subtype of the interface signature type associated with the formal parameter.

7.2.2.5 Flows, operations and signals

Flows and operations can be defined in terms of signals. This enables signal interfaces to be used as a basis for explaining multi-party binding; end-to-end Quality of Service characteristics and compound bindings between different kinds of interface (e.g., stream to operation interface bindings).

A definition of flows using signals depends upon the details of the interactions abstracted in the specification of the stream interface concerned and is, therefore, beyond the scope of this reference model.

NOTE 1 – The level of detail at which interactions are presented is a choice to be made by the creator of the computational specification. In general, a flow is seen as essentially unstructured. The data transfer process is seen as a single stream, with the open/close hidden in the binding process. The structure of the flow may be seen in the engineering viewpoint as a structure of channels and their usage.

NOTE 2 – There can be cases where the model is of the creation of a binding to support negotiation as a sequence of operations, the last of which creates a flow, followed by some closing operation. Similarly, suspend/resume could be controlled by the application or it could be subject to selected transparencies.

Operations can be modelled by introducing signal interfaces corresponding to both the client operation interface and server operation interface involved:

- in a signal interface corresponding to a client operation interface there is a signal (invocation submit) corresponding to each invocation with the same parameters, and, in the case of an interface containing interrogations, a signal (termination deliver) corresponding to each possible termination with the same parameters as that termination;
- in the signal interface corresponding to a server operation interface there is a signal (invocation deliver) corresponding to each invocation with the same parameters, and, in the case of an interface containing interrogations, a signal (termination submit) corresponding to each possible termination with the same parameters as that termination.

This creates an equivalence between the resulting set of signals and the set of invocations and terminations in the operation interfaces being described.

7.2.3 Binding rules

In this reference model, binding is defined with reference to binding actions. Use of such actions is called explicit binding. There are two kinds of binding actions: primitive binding actions and compound binding actions.

A primitive binding action binds two computational objects directly. A compound binding action can be expressed in terms of primitive binding actions linking two or more computational objects via a binding object. The presence of a binding object in a computational binding gives the means to express configuration and quality of service control (see 7.2.3.3).

NOTE 1 – Definition of appropriate behaviour for a binding object enables the provision of different communication styles such as remote invocation, message transfer, etc.

In notations which have no terms for expressing binding actions, binding is implicit. Implicit binding for other than server operation interfaces is not defined in the reference model because in other cases it is not self-evident where the initiative in binding should be placed relative to subsequent interaction. The additional information needed can be supplied in an explicit binding action.

NOTE 2 – The nature of computational binding and binding objects, and the engineering mechanisms that support them are further specified in Rec. ITU-T X.930 (1998) | ISO/IEC 14753:1999, *Information technology – Open Distributed Processing – Interface references and binding*.

7.2.3.1 Implicit binding rules for server operation interfaces

If an invocation by a client object references a server operation interface to which the client is not bound, implicit binding is required. Setting up an implicit binding involves the following procedure if a suitable client operation interface bound to the server does not exist:

- create a client operation interface of complementary signature type to the server interface;
- bind the client operation interface to the server operation interface;
- invoke the server object using the client operation interface;
- (optionally) when the operation completes, delete the client interface.

7.2.3.2 Primitive binding rules

Primitive binding actions enable binding of an interface of the object which initiates the action to another interface (of another object, or of itself). The binding action is parameterized by two identifiers, one for each interface involved. The pre-conditions for a primitive binding action are that both interfaces involved are of the same kind (viz., signal, stream or operation), that they are of complementary causality and that their signature types are complementary.

Primitive binding either establishes a binding between the two interfaces concerned, or fails.

Deleting an interface that has been bound to another interface using a primitive binding action deletes the binding as well as the interface.

7.2.3.3 Compound binding rules

Compound binding actions enable a set of interfaces to be bound, using a binding object to support the binding. Except for the provisions of this clause, a binding object is an ordinary computational object. In a binding object template, the behaviour specification is expressed in terms of a set of formal role parameters each of which is associated with an interface template.

Compound binding actions are parameterized by a binding object template and a set of interfaces to be bound. If successful, the binding action allows interaction between some or all of the members of the set of interfaces.

The pre-conditions for compound binding are that, for each formal role in the binding object template:

- the corresponding interface parameter must be of the same kind (viz., signal, stream or operation) as the interface template associated with the formal role in the binding object template;
- the corresponding interface parameter must be of complementary causality to the interface template associated with the formal role in the binding object template;
- the corresponding interface parameter must be a subtype of the signature type of the interface template associated with the formal role in the binding object template.

A compound binding action comprises the following steps:

- a binding object is instantiated from the binding object template;
- each interface template within the binding object associated with a formal role parameter in the binding object template is instantiated;
- the binding object uses primitive binding actions to bind each such interface to the interface referenced in the corresponding actual parameter of the binding action;
- a set of control interfaces are instantiated and identifiers for these interfaces returned as results of the binding action (i.e., becoming part of the state of the object which performed the action – this object can subsequently pass on the identifier by interaction with other computational objects).

The control interfaces of a binding object provide some or all of the following functions:

- monitoring the use of the binding;
- monitoring changes to the binding;
- authorizing changes to the binding;
- changing the membership of the binding;
- changing the pattern of communication enabled by the binding;
- changing the quality of service of the binding;
- deleting the binding as a whole.

The effect of deleting the primitive binding between a bound object and the binding object is determined by the behaviour of the binding object.

7.2.4 Type rules

This reference model specifies signature type rules for computational interfaces. Signature subtyping rules define minimum requirements for one interface to substitute for another. The rules are based on the interaction semantics of computational interfaces (viz., signal, stream and operation interfaces). They are sufficient to ensure that a substituted interface can consistently interpret the structure of any interactions that occur.

Signature subtyping rules for interfaces with alternative interaction semantics can be defined in terms of signals; such definitions can be introduced by an ODP standard.

7.2.4.1 Signature subtyping rules for signal interfaces

The definition of signal interface signature subtyping is given in Annex A. For signal interface types that are not defined recursively, the rules are summarized below.

Signal interface signature type X is a subtype of signal interface signature type Y if the conditions below are met:

- for every initiating signal signature in Y there is a corresponding initiating signal signature in X with the same name, with the same number and names of parameters, and that each parameter type in X is a subtype of the corresponding parameter type in Y ;
- for every responding signal signature in X there is a corresponding responding signal signature in Y with the same name, with the same number and names of parameters, and that each parameter type in Y is a subtype of the corresponding parameter type in X .

7.2.4.2 Signature subtyping rules for stream interfaces

Stream signature subtyping rules depend upon the details of the interactions abstracted in the definition of the stream interfaces concerned. In particular these details will clarify whether or not the subtyping rules will permit incomplete correspondences between the set of flows in the two interfaces. Therefore complete signature subtyping rules are beyond the scope of this reference model. The constraints on stream signature subtyping are given in Annex A. For stream interface types that are not defined recursively, the constraints are summarized below.

Stream interface X is a signature subtype of stream interface Y if the conditions below are met for all flows which have identical names:

- if the causality is producer, the information type in X is a subtype of the information type in Y ;
- if the causality is consumer, the information type in Y is a subtype of the information type in X .

7.2.4.3 Signature subtyping rules for operation interfaces

The definition of operation interface signature subtyping is given in Annex A. For interface types that are not defined recursively, the rules are summarized below.

Operation interface X is a signature subtype of interface Y if the conditions below are met:

- for every interrogation in Y , there is an interrogation signature in X (the corresponding signature in X) which defines an interrogation with the same name;
- for each interrogation signature in Y , the corresponding interrogation signature in X has the same number and names of parameters;
- for each interrogation signature in Y , every parameter type is a subtype of the corresponding parameter type of the corresponding interrogation signature in X ;
- the set of termination names of an interrogation signature in Y contains the set of termination names of the corresponding interrogation signature in X ;
- for each interrogation signature in Y , a given termination in the corresponding interrogation signature in X has the same number and names of result parameters in the termination of the same name in the interrogation signature in Y ;
- for each interrogation signature in Y , every result type associated with a given termination in the corresponding interrogation signature in X is a subtype of the result type (with the same name) in the termination with the same name in Y ;
- for every announcement in Y , there is an announcement signature in X (the corresponding signature in X) which defines an announcement with the same name;
- for each announcement signature in Y , the corresponding announcement signature in X has the same number and names of parameters;
- for each announcement signature in Y , every parameter type is a subtype of the corresponding parameter type in the corresponding announcement signature in X .

7.2.5 Template rules

7.2.5.1 Computational object template rules

A computational object (including the special case of a binding object) can:

- initiate or respond to signals;
- produce or consume flows;
- initiate operation invocations;
- respond to operation invocations;
- initiate operation terminations;
- respond to operation terminations;
- instantiate interface templates;
- instantiate object templates;
- bind interfaces;
- access and modify its state;
- delete one or more of its interfaces;
- delete itself;
- spawn, fork and join activities;
- obtain a computational interface identifier for an instance of the trading function;
- test if one computational interface signature is a subtype of another.

Any of these actions can fail.

7.2.5.2 Computational interface instantiation

Computational interface instantiation establishes one or more computational interface identifiers for the new interface in the object performing the instantiation.

7.2.5.3 Computational object template instantiation

The behaviour expression in a computational object template includes a description of the behaviour to occur when the template is instantiated (the instantiated behaviour). The environment contract specification describes the contract to be established between the instantiated object and its environment when the template is instantiated. Where the instantiation behaviour includes interface instantiation, the instantiation establishes identifiers for these interfaces in the object which initiated the instantiation.

7.2.6 Failure rules

The failure modes visible to an object are determined by its behaviour and environment contract specifications.

Any of the computational actions in 7.2.5.1 can fail and that failure can be observed by the object performing the action. Interaction can be disrupted by failure of the objects involved, or the binding between them, or both. In the case of signals, the failure is identical for, and visible to, all participants in the interaction. In the case of flows and operations, failure need not occur in all the participants, and may occur at different times and with different parameters for each failing participant.

NOTE – Examples of interaction failure include security failure, communication failure and resource failure.

For operations, the failure of the server computational object to respond to invocations or to initiate terminations can be observed by the client computational object involved.

The instantiation of an object template or a computational interface template fails if the environment contract cannot be satisfied. A binding action fails if any of the environment contracts in the interfaces being bound cannot be satisfied.

7.2.7 Portability rules

Standards for portability in ODP systems specify action templates for the actions described in 7.2.5.1. The specification of such templates is a matter of language design and is therefore beyond the scope of this reference model. In addition to syntactic concerns, a portability standard must address specific semantic issues including:

- composition rules for action templates, including templates for forking and joining actions to enable concurrency and synchronization;

- the terms available for specification of object and interface templates, together with rules for their composition;
- ordering and delivery guarantees for announcements.

A portability standard may represent the permitted actions directly (e.g., as library functions) or indirectly through syntactic structures. There can be alternative portability standards both in terms of style (e.g., an event-based processing model versus a threads based model) and content (e.g., the number of computational actions supported). This reference model identifies two.

A basic portability standard is one which includes at least:

- interrogations;
- implicit binding;
- computational object instantiation;
- computational interface instantiation;
- access to and modification of state;
- support for threads with spawn, fork and join actions;
- obtaining an identifier for a computational interface at which the trading function is provided (to enable subsequent binding to and use of the function);
- interface signature subtype testing.

An extended portability standard is one which includes all the actions described in 7.2.5.1.

7.3 Conformance and reference points

In the computational language there exists a reference point at any interface of any object. Each reference point can become either a programmatic conformance point, a perceptual conformance point, an interworking conformance point or an interchange conformance point, depending upon the requirements set when the reference point is designated as a conformance point by a specific standard or in the system specification.

In the computational language these requirements are specified in terms of interface and object templates which determine the interface of the conforming object.

An implementor claiming conformance to a computational specification must list the engineering reference points which correspond to the required computational reference point and state which transparency and engineering structures apply to them. By this act the identified reference points become conformance points. The set of interactions at these conformance points can then be interpreted in computational language terms to determine that the computational specification is not violated.

Conformance of an object at a programmatic conformance point can be tested in terms of a standardized interface specification language and a language binding which satisfies the portability rules. Conformance of an object at an interworking conformance point can be tested in terms of interactions visible in communications protocols.

8 Engineering language

The engineering language comprises concepts, rules and structures for the specification of an ODP system from the engineering viewpoint.

An engineering specification is primarily concerned with two areas – the structuring responsibilities of objects within a system (cluster, capsule and node rules) and the structuring of communication responsibilities between systems (channel rules).

An engineering specification defines the mechanisms and functions required to support distributed interaction between objects in an ODP system as a series of templates for the computational interactions, parameterized to support a range of different policies selected either for the enterprise, information or computational specifications through the transparency schema.

NOTE – The boundary between the computational and engineering viewpoints is not fixed by this reference model and depends on the distribution transparencies selected for the computational specification. An engineering specification that corresponds to a computational specification defines engineering structures to provide the distribution transparencies selected in the computational specification.

8.1 Concepts

The engineering language contains the concepts of Rec. ITU-T X.902 | ISO/IEC 10746-2 and those defined in 7.1, together with those defined here, subject to the rules of 8.2. The node management function (see 12.1), object management function (see 12.2), cluster management function (see 12.3), capsule management function (see 12.4) and engineering interface reference tracking function (see 13.9) are an integral part of the engineering language; the other ODP functions (see 11) are included in the language as needed.

8.1.1 basic engineering object: An engineering object that requires the support of a distributed infrastructure.

8.1.2 cluster: A configuration of basic engineering objects forming a single unit for the purposes of deactivation, checkpointing, reactivation, recovery and migration.

NOTE – A segment of virtual memory containing objects is an example of a cluster.

8.1.3 cluster manager: An engineering object which manages the basic engineering objects in a cluster.

8.1.4 capsule: A configuration of engineering objects forming a single unit for the purpose of encapsulation of processing and storage.

NOTE – A virtual machine (e.g., a process) is an example of a capsule.

8.1.5 capsule manager: An engineering object which manages the engineering objects in a capsule.

8.1.6 nucleus: An engineering object which coordinates processing, storage and communications functions for use by other engineering objects within the node to which it belongs.

NOTE – An operating system (kernel) is an example of a nucleus.

8.1.7 node: A configuration of engineering objects forming a single unit for the purpose of location in space, and which embodies a set of processing, storage and communication functions.

NOTE 1 – A computer and the software it supports (operating system and applications) is an example of a node.

NOTE 2 – A node may have internal structure which is not of concern in an engineering specification. Thus a node can be, for example, a parallel computer under the control of a single operating system.

8.1.8 channel: A configuration of stubs, binders, protocol objects and interceptors providing a binding between a set of interfaces to basic engineering objects, through which interaction can occur.

NOTE – Bindings that require channels are referred to as distributed bindings in the engineering language; bindings between engineering objects that do not require channels (e.g., between engineering objects in the same cluster) are referred to as local bindings.

8.1.9 stub: An engineering object in a channel, which interprets the interactions conveyed by the channel, and performs any necessary transformation or monitoring based on this interpretation.

NOTE – For example, a stub can perform marshalling of parameters into communications buffers, or the recording of activity for audit purposes.

8.1.10 binder: An engineering object in a channel, which maintains a distributed binding between interacting basic engineering objects.

8.1.11 <x> interceptor: An engineering object in a channel, placed at a boundary between <x> domains. An <x> interceptor:

- performs checks to enforce or monitor policies on permitted interactions between basic engineering objects in different domains;
- performs transformations to mask differences in interpretation of data by basic engineering objects in different domains.

NOTE – An inter-subnetwork relay is an example of an interceptor.

8.1.12 protocol object: An engineering object in a channel, which communicates with other protocol objects in the same channel to achieve interaction between basic engineering objects (possibly in different clusters, possibly in different capsules, possibly in different nodes).

8.1.13 communications domain: A set of protocol objects capable of interworking.

8.1.14 communication interface: An interface of a protocol object that can be bound to an interface of either an interceptor object or another protocol object at an interworking reference point.

8.1.15 binding endpoint identifier: An identifier, in the naming context of a capsule, used by a basic engineering object to select one of the bindings in which it is involved, for the purpose of interaction.

NOTE 1 – A memory address (for a data structure representing an engineering interface) is an example of a binding endpoint identifier.

NOTE 2 – The same form of binding endpoint identifier can be used, whether the binding involved is either local or distributed.

8.1.16 engineering interface reference: An identifier, in the context of an engineering interface reference management domain, for an engineering object interface that is available for distributed binding.

NOTE – An engineering interface reference is necessary to establish distributed bindings, and is distinct from the binding endpoint identifiers used by a basic engineering object for the purposes of interaction.

8.1.17 engineering interface reference management domain: A set of nodes forming a naming domain for the purpose of assigning engineering interface references.

8.1.18 engineering interface reference management policy: A set of permissions and prohibitions that govern the federation of engineering interface reference management domains.

8.1.19 cluster template: An object template for a configuration of objects and any activity required to instantiate those objects and establish initial bindings.

8.1.20 checkpoint: An object template derived from the state and structure of an engineering object that can be used to instantiate another engineering object, consistent with the state of the original object at the time of checkpointing.

8.1.21 checkpointing: Creating a checkpoint. Checkpoints can only be created when the engineering object involved satisfies a pre-condition stated in a checkpointing policy.

8.1.22 cluster checkpoint: A cluster template containing checkpoints of the basic engineering objects in a cluster.

8.1.23 deactivation: Checkpointing a cluster, followed by deletion of the cluster.

8.1.24 cloning: Instantiating a cluster from a cluster checkpoint.

8.1.25 recovery: Cloning a cluster after cluster failure or deletion.

8.1.26 reactivation: Cloning a cluster following its deactivation.

8.1.27 migration: Moving a cluster to a different capsule.

8.2 Structuring rules

An engineering specification defines the infrastructure required to support functional distribution of an ODP system, by:

- identifying the ODP functions required to manage physical distribution, communication, processing and storage;
- identifying the roles of different engineering objects supporting the ODP functions (for example, the nucleus).

An engineering specification is expressed in terms of:

- a configuration of engineering objects, structured as clusters, capsule and nodes;
- the activities that occur within those engineering objects;
- the interactions of those engineering objects.

An engineering specification is constrained by the rules of the engineering language. These comprise:

- channel rules (see 8.2.1), interface reference rules (see 8.2.2), distributed binding rules (see 8.2.3) and relocation rules (see 8.2.4) for the provision of distribution transparent interaction among engineering objects;
- cluster rules (see 8.2.5), capsule rules (see 8.2.6) and node rules (see 8.2.7) governing the configuration of engineering objects;
- failure rules (see 8.2.9).

8.2.1 Channel rules

A channel supports distribution transparent interaction between engineering objects. This includes, for example:

- operation execution between a client object and a server object;
- a group of objects multicasting to another group of objects;
- stream interactions involving multiple producer objects and multiple consumer objects.

NOTE 1 – The nature of channels and the engineering mechanisms that support them, and how these support computational interactions, are further specified in Rec. ITU-T X.931 (1999) | ISO/IEC 14752:2000, *Information technology – Open Distributed Processing – Protocol support for computational interactions*.

Interaction between engineering objects entails the transfer of some or all of:

- engineering interface references;
- cluster templates;
- data.

A channel is a configuration of stubs, binders, protocol objects and interceptors interconnecting a set of engineering objects. The configuration is an acyclic graph with stub objects at the outermost vertices as illustrated in Figures 2 and 3. Each path through the graph between stub objects consists of, in sequence, either


- a binder, a protocol object, a protocol object and a binder; or
- a binder, a protocol object, an interceptor, a protocol object and a binder.

The behaviour of a channel with respect to channel configuration or quality of service management is controlled via control interfaces of stubs, binders, protocol objects and interceptors. Such control interfaces are optional.

NOTE 2 – Stubs, binders, protocol objects and interceptors in a channel can have (local or distributed) bindings to other engineering objects outside the channel providing, for example, relocation or coordination functions.

NOTE 3 – Depending upon the kind of transformation involved, an interceptor may be decomposed into stubs, binders, protocol objects and basic engineering objects, mirroring the channel structure.

Objects in a channel can themselves be basic engineering objects supported by channels.



X.903(09)_F02

Figure 2 – An example of a basic client/server channel




Figure 3 – An example of a multi-endpoint channel

8.2.1.1 Stubs

Basic engineering objects which interact via channels are locally bound to stubs. In a channel, stubs provide the conversion of data carried by interactions. Stubs can apply controls and keep records (e.g., for security and accounting). Stubs can interact with engineering objects outside the channel (e.g., a security function) if required. A stub in a channel has an interface for use by the basic engineering object it supports and an interface for interaction with a binder. It can also have a control interface.

When interconnected stubs use different transfer syntaxes, there must be an interceptor in the path between them which can transform data from one syntax to the other.

A stub can be of one of the following forms:

- specific to the basic engineering object *interface instance* to which it is bound;
- specific to the *interface type* of the basic engineering object interface to which it is bound (therefore the stub may be shared between a number of channels of the same type);

- c) generic (i.e., not specific to any single interface type); such stubs may be shared between a number of channels of different types.

NOTE 1 – In a) interactions between the engineering object and the stub only convey interaction data (e.g., operation names and parameters in the case of an invocation). The stub acts as a local proxy for other basic engineering objects bound to the channel.

NOTE 2 – In b) interactions must additionally include an identifier for the channel to be used.

NOTE 3 – In c) interactions must additionally include an identifier and type for the channel to be used, to enable the stub to ensure the interaction data are compatible with the channel type.

8.2.1.2 Binders

The binders in a channel manage the end-to-end integrity of that channel. When required, binders provide relocation transparency by monitoring communication failures and repairing broken distributed bindings. Binders in a channel can interact with engineering objects outside the channel to obtain additional data needed to perform their function (e.g., a relocater to obtain location data). A binder in a channel has at least an interface for interaction with a stub and one or more interfaces for interaction with protocol objects. A binder can have a control interface. If present, the control interface enables changes in the configuration of the channel and deletion of the channel.

8.2.1.3 Protocol objects

Protocol objects provide communication functions. Protocol objects can interact with engineering objects outside a channel (e.g., directory functions) to obtain the information they need. A protocol object has an interface for interaction with a binder and at least one communications interface for interaction with other protocol objects (via interceptors if required). A protocol object can have a control interface. When protocol objects in a channel are of different type, they require an interceptor which provides protocol conversion. All protocol objects in a communication domain can communicate directly using facilities of the communications domain (which fall outside the scope of this reference model).

NOTE – ISO/IEC 19500-2:2003, *Information technology – Open Distributed Processing – Part 2: General Inter-ORB Protocol (GIOP)/Internet Inter-ORB Protocol (IIOP)* defines a specific protocol for use by protocol objects.

At any given location in time, a protocol object is identified by its location in space, but different protocol objects can occupy the same location in space at different times (i.e., network addresses can be recycled). When protocol objects in a channel are of the same type, but in disjoint communications domains, naming conflicts are possible (e.g., names for communications interfaces might be ambiguous). In this circumstance, an interceptor is required to transform names exchanged in the process of establishing and maintaining the integrity of the channel.

8.2.1.4 Interceptors

An <x> interceptor in a channel stands at the boundary between <x> domains and provides checks and transformations on interactions that cross <x> domain boundaries. Depending upon the boundary crossed, interceptors require different information to perform their task. Some interceptors will need to know the signature types of the basic engineering object interfaces bound to the channel in which the interceptor is located so that they can interpret interactions supported by the channel. An interceptor has at least two communication interfaces. It can have a control interface.

8.2.2 Interface reference rules

For the purpose of distributed binding, engineering interfaces are located in space and time by engineering interface references. Engineering interface references are defined relative to an engineering interface reference management domain which determines policy for the content, allocation, tracking and validation of engineering interface references within that domain. An engineering interface reference management domain consists of a set of nodes. Engineering interface reference domains can be federated if their engineering interface reference management policies do not conflict.

An engineering interface reference contains information which enables bindings to be established to engineering object interfaces. This information both enables nucleus objects to create channels and also enables binders within channels to maintain a distributed binding between interconnected engineering objects. The information within an engineering interface reference can take the form of:

- data;
- identifiers for interfaces giving access to such data;
- a combination of data and identifiers.

The data necessary for binding can include any or all of the following items:

- the interface type of the referenced interface;
- a channel template describing the interceptors, protocol objects, binders and stubs that can be selected when configuring a channel to support the distributed binding;

- the location in space and time (e.g., a network address) of the communication interfaces at which the binding process can be initiated;
- information to enable the detection and repair of distributed bindings invalidated by engineering object relocation.

An interface reference management domain can be divided into subdomains. In this case interface references within the domain are organized as a set of alternative sets of information, one for each subdomain in which binding is to be enabled.

NOTE 1 – If the nucleus supporting the engineering interface supports different protocols, binding processes and transfer syntaxes, the engineering interface reference will indicate the valid combinations that can be selected in any particular distributed binding; different bindings might make different selections.

NOTE 2 – This reference model does not prescribe the method by which a channel template and the location in space and time of related interfaces are derived from an engineering interface reference.

Engineering interface references are allocated by nuclei, at interfaces supporting the node management function (see 12.1.3). Engineering interface references are tracked by the engineering interface reference tracking function (see 13.9) for the purposes of detecting engineering interfaces which are no longer referenced. Policy for rebinding to engineering object interfaces which have been relocated, and updated engineering interface references for engineering interfaces which have been relocated, are recorded by the relocation function (see 14.3). These three functions (node management, engineering interface reference tracking and relocation) may be coordinated by use of a shared information organization function (see 14.2).

Engineering interface references are unambiguous in the naming context of an engineering interface reference management domain. To achieve this unambiguity, the nodes within the engineering interface reference management domain must allocate engineering interface references in a coordinated way. Engineering interfaces must be allocated in a manner which prevents an engineering interface reference from referencing the wrong interface, even in the presence of failures and interface relocation. At worst an interface reference will refer to an interface which does not exist (e.g., immediately after failure of the engineering object supporting the interface).

NOTE 3 – In ODP systems in which most interfaces do not change location, the management of interface references can be optimized: the nucleus can allocate engineering interface references autonomously; the channel type and communications interface identifier associated with the interface can be stored and transmitted within the engineering interface reference; the relocation function can be used to validate and update engineering interface references for interfaces that have been relocated.

Before issuing an engineering interface reference, the nucleus constructs a channel template which defines a configuration of stub, binder and protocol objects suitable to support interactions in the interface. In addition the nucleus establishes sufficient local structure to make binding to the interface possible, and associates the template and the local structure with a communications interface. The engineering interface reference makes this information available.

An interceptor that stands on a boundary between engineering interface reference management domains maintains mappings between engineering interface references in those domains. When an engineering interface reference or a cluster template containing engineering interface references crosses an engineering interface reference domain boundary, the engineering interface references involved must be transformed so as to be valid in the new domain.

Media interchange of engineering interface references from one interface reference management domains to another is only possible when there is a defined procedure for mapping the references to avoid ambiguity.

8.2.3 Distributed binding rules

Establishing a channel requires the creation of appropriate stubs, binders, protocol objects and interceptors. Channel establishment can be initiated by any engineering object. It is provided by each nucleus as a function of its node management interface. Distributed binding involves interaction with the nuclei of the nodes at which the interfaces to be bound are located. Channel establishment is parameterized by a channel template and a set of interface references each assigned to a particular role in the channel template. The channel template must be compatible with the channel types nominated by the engineering interface references for the interfaces to be bound. The nucleus for each object to be bound creates a configuration of stubs, binders and protocol objects at its node to support the interfaces of that object being bound. This includes configuration of their control interfaces. The protocol objects that support the channel are connected (possibly via interceptors) at their communication interfaces. The selection and configuration of stubs, binders, protocol objects and interceptors is determined by the channel template and channel types of the interface references involved. Each basic engineering object bound by the channel is assigned a binding endpoint identifier for each interface it has to the channel. Basic engineering objects use binding endpoint identifiers to nominate at which of their interfaces a distributed interaction is to occur.

NOTE 1 – Channels can be established by any engineering object irrespective of whether or not the object has an interface that is to be bound by the channel.

NOTE 2 – A basic engineering object initiating a distributed binding requires a set of interface references. These may be obtained in any of the following ways:

- a) on initialization of the object;
- b) by interaction between the initiating object and the nucleus as part of the instantiation of the initiating object's interfaces;
- c) through some chain of interactions with the other objects concerned (e.g., by parameter passing or trading).

NOTE 3 – A channel template can contain alternative configurations to be applied in selected circumstances. For example, if communication paths are insecure, encrypting stubs might be required.

8.2.4 Relocation rules

Engineering objects can be relocated as a result of:

- reactivation and deactivation;
- checkpointing and recovery;
- migration;
- communications domain management functions (e.g., changing a communication interface identifier).

NOTE 1 – A communications interface may have its identifier changed as a consequence of changing the network address of a node.

NOTE 2 – When channels are re-established, stubs, binders and protocol objects might be used which are different to those used prior to relocation. Thus, an identifier for a communications interface is not necessarily sufficient to identify an engineering object interface.

Relocation can cause channels to fail and can invalidate engineering interface references. Failed channels can be repaired if the activity which changes the location of the engineering object interface notifies the relocation function appropriately (see 14.3).

The binders in a channel detect when relocation has invalidated the channel. Either the binders collaborate to correct the mapping between engineering interface references and the channel structure (i.e., relocation transparency is required – see 16.6), or the channel fails. When relocation transparency is required the information available through the engineering interface reference enables the binders to use the relocation function to determine the new locations of the basic engineering objects involved.

8.2.5 Cluster rules

A cluster contains a set of basic engineering objects, associated with a cluster manager. Each member of a cluster can have an interface supporting the object management function. Each such object management interface is bound to the cluster's cluster manager. A basic engineering object in a cluster is always bound to its nucleus, at an interface providing the node management function and to its cluster manager. In addition, a basic engineering object in a cluster can be bound to other basic engineering objects in the same cluster, or in other clusters. Each cluster manager in a capsule is bound to the capsule manager for that capsule. This structure is illustrated in Figure 4.




Figure 4 – Example structure supporting a basic engineering object

A cluster is always contained in a single capsule. A cluster is responsible for its own security, but can be assisted by security functions. Any assisting security function must either be provided by an object in the same capsule as that cluster, or be accessed by secure interactions if it is outside the capsule. Engineering objects in the same cluster can interact using a local binding within the cluster, or using a distributed binding supported by a channel. Engineering objects in different clusters interact using distributed bindings supported by channels.

NOTE 1 – Interactions that are at perceptual or interchange reference points are not precluded.

NOTE 2 – Although no channel is needed to support local bindings between objects in the same cluster, the identifier used to perform invocations is of the same kind as is used between engineering objects in different clusters, and is still called a binding endpoint identifier.

Cluster instantiation (including cloning as a special case) is performed by a capsule manager.

If the template is a cluster checkpoint, the instantiation (i.e., cloning) enables the new cluster to act as a substitute for the original cluster from which the cluster template was derived. When required for distribution transparency reasons, the cloning process includes re-establishing any distributed bindings that were held by the original cluster.

A cluster has an associated cluster manager. The cluster manager provides the cluster management function. The cluster manager embodies the management policy for the engineering objects in its cluster. The cluster management policy might lead the cluster manager to interact with other ODP functions to complete cluster management activities.

A cluster manager assists its capsule manager in the management of the engineering interface references of objects. This may require access to the engineering interface reference tracking function.

8.2.6 Capsule rules

A capsule consists of

- cluster(s);
- cluster managers, one for each cluster in the capsule;
- a capsule manager to which each of the cluster managers in the capsule is bound.

Stubs, binders and protocol objects for a channel bound to an interface of a basic engineering object within a cluster in a capsule can be included within that capsule. All the engineering objects in a capsule are bound to the same node management interface. Engineering objects in other capsules are bound to different node management interfaces. A capsule is contained within a node. A capsule has a capsule manager. The capsule manager is bound, at an interface providing the cluster management function, to each cluster manager in the capsule. This structure is illustrated in Figure 5 (the figure abstracts out details of the nucleus).




Figure 5 – Example structure of a capsule

Capsule instantiation is performed by the nucleus using a capsule template which specifies the initial configuration of engineering objects in the capsule, including clusters, cluster managers, stubs, binders, protocols and a capsule manager.

A capsule is a naming context for binding endpoint identifiers. The reference model does not require that such identifiers be valid in any larger context. Engineering interface references are used to communicate knowledge of engineering object interfaces between capsules (for the purposes of binding).

The capsule manager embodies the management policy for the clusters in its capsule. The capsule management policy might lead the capsule manager to interact with other functions to complete capsule management activities. The capsule manager has an interface providing the capsule management function. The structures supporting interaction between cluster managers, capsule managers and the nucleus within a node are a matter of implementation detail outside the scope of this reference model.

8.2.7 Node rules

A node consists of one nucleus and a set of capsules. All of the engineering objects in a node share common processing, storage and communication functions.

A node is a member of one or more engineering interface reference management domains.

The nucleus provides a set of node management interfaces, one to each capsule within the node.

The structure of a node is illustrated in Figure 6.




Figure 6 – Example structure of a node

The procedure for node instantiation is outside the scope of this reference model. It must result in:

- introduction of the node's nucleus and its associated processing, storage and communications functions, including introduction of node management functions to enable distributed binding of engineering interface references;
- introduction of any trading function needed by the instantiation process;
- instantiation of any channels required as part of the initial configuration of the node (e.g., to supporting engineering objects such as a relocater).

The set of protocol objects introduced during node instantiation determine the initial set of communication domains to which the node belongs. The nucleus provides the node management function and embodies a node management policy. This policy might lead the nucleus to interact with other ODP functions to complete node management activities. A capsule is the basic unit for the application of node management policy; although an individual object can use the node management function, it is subject to the policies applicable to its capsule. Distinct capsules can be subject to different node management policies.

8.2.8 Application management rules

The management of the lifecycle (creation, migration, deactivation, reactivation, checkpointing, failure, recovery and deletion) of coordinated sets of clusters is driven by application-specific management policies. Application management policies can apply to individual clusters or to coordinated sets of clusters. A set of managed clusters form an application management domain. Applications management policy can be implemented by applications-specific management functions which effect changes using the mechanisms provided by the coordination and management functions defined in this reference model, e.g., the cluster management function.

Where appropriate (i.e., depending upon which distribution transparencies apply), application-specific management functions can receive notifications of significant events affecting the clusters they manage and takes action in response to those notifications. For example binding failure reports can lead to reactivation of a cluster, or excessive workload reports can lead to migration of clusters. Requests and notifications relating to one cluster in an application management

domain can lead to application-specific management functions initiating lifecycle actions on other clusters in the domain.

The specifics of application domain management are beyond the scope of this reference model.

8.2.9 Failure rules

Failures can be categorized as involving clusters, capsules, nodes or communications domains. The analysis of failures is based on the fact that:

- a failure localized to a cluster can be detected by its cluster manager;
- a failure localized to a capsule can be detected by its capsule manager;
- the failure of a node may be detected by protocol objects at other nodes to which it is interconnected;
- the failure of a communications domain may be detected by protocol objects in other communication domains to which it is interconnected.

NOTE – There is an inherent ambiguity in communication systems which may prevent a protocol object from distinguishing between communication failure and remote node failure.

8.3 Conformance and reference points

There is a programmatic reference point at an interaction point between a cluster manager and a basic engineering object.

There is a programmatic reference point at an interaction point between an engineering object and a nucleus.

There is a programmatic reference point at an interaction point between basic engineering objects.

There can be a perceptual or an interchange reference point at an interface of a basic engineering object.

Where interaction between basic engineering objects is provided by a channel, then, for each basic engineering object involved, there is a programmatic reference point at the interaction points between those objects and the corresponding stubs in the channel.

Within the configuration of objects that comprise a channel, there is a programmatic reference point at each of the following interaction points within the channel:

- between stubs (abstracting the binders, protocol objects and interceptors in the channel between the stubs);
- between stubs and binders;
- between binders (abstracting the protocol objects and interceptors in the channel between the binders);
- between binders and protocol objects;
- between protocol objects and other protocol objects within the same node (abstracting the interceptors, if any, between the protocol objects),

and there is an interworking reference point at any interaction point between protocol objects and other protocol objects or interceptors in different nodes.

Control interfaces of stubs, binders, protocol objects and interceptors are programmatic reference points.

Where engineering objects within a channel interact with other engineering objects (either within that channel or outside that channel) via interfaces which are not within that channel, then the reference points applicable to such interfaces are determined by the recursive application of these rules.

By defining conformance at the interworking reference points, interworking of systems is made possible.

By defining conformance at the programmatic reference points, portability of engineering objects between systems is made possible.

Conformance of individual engineering objects at the programmatic reference points does not itself guarantee that the engineering object will be portable to all systems or will interwork with matching engineering objects bound to other nuclei.

9 Technology language

The technology language comprises concepts, rules and structures for the specification of an ODP system from the technology viewpoint. A technology specification defines the choice of technology for an ODP system.

NOTE – In a system specification in which the choice of technology is left open, the technology specification does not provide full details of specific technology resources.

In this reference model, prescription in the technology viewpoint is restricted to a basic set of concepts and rules addressing the scope and nature of technology specifications

9.1 Concepts

The technology language consists of the concepts of Rec. ITU-T X.902 | ISO/IEC 10746-2 and the concepts defined here subject to the structuring rules in 9.2.

9.1.1 implementable standard: A template for a technology object.

9.1.2 implementation: A process of instantiation whose validity can be subject to test.

9.1.3 IXIT: Implementation extra information for testing.

9.2 Structuring rules

A technology specification defines the choice of technology for an ODP system in terms of:

- a configuration of technology objects that represent hardware and software components of the implemented system, and technology objects that represent the communication links between these components; and
- the interfaces between these technology objects.

A technology specification:

- expresses how the specifications for an ODP system are implemented;
- identifies specifications for technology relevant to the construction of ODP systems;
- provides a taxonomy for such specifications;
- specifies the information required from implementors to support testing.

When applying a specification written in another viewpoint language, a technology specification is constructed to give an interpretation of the atomic terms in the other viewpoint specification.

The technology specification for an ODP function can reference the specifications of other ODP functions.

A technology specification consists of statements that technology objects are instances of named implementable standards.

The technology specification gives a proforma for the conformance IXIT by listing the required set of templates and the descriptive names for all necessary reference points.

All implementable standards are introduced by reference to other specifications. The technology language does not define any other rules constraining the behaviour of technology objects or for constructing implementable standards.

9.3 Conformance and reference points

The technology language is used to assert that technology objects are instances of implementable standards; these standards will, in general, contain conformance statements.

10 Consistency rules

A set of specifications of an ODP system written in different viewpoint languages should not make mutually contradictory statements (see 4.2.2), i.e., they should be mutually consistent. Thus, a complete specification of a system includes statements of correspondences between terms and language constructs relating one viewpoint specification to another viewpoint specification, showing that the consistency requirement is met. The minimum requirements for consistency in a set of specifications for an ODP system is that they should exhibit the correspondences defined in this reference model and those defined within the set of specifications itself. This reference model does not declare generic correspondences between every pair of viewpoint languages. This clause is restricted to the specification of correspondences between a computational specification and an information specification, and between a computational

specification and an engineering specification. In each case, the correspondences are expressed as interpretation relationships linking terms in one viewpoint language to terms in the other viewpoint language. A set of specifications based on this reference model will, in general, need to relate all the viewpoint specifications.

The key to consistency is the idea of correspondences between specifications, i.e., a statement that some terms or structures in one specification correspond to other terms and specifications in a second specification. Correspondences can be established between two different specifications in a single language or in two different languages. Statements of correspondences between two languages imply equivalent correspondences between any pair of specifications expressed in those languages.

Analysis of consistency depends on the application of specific consistency techniques. Most of these are based on checks for particular kinds of inconsistency, and thus cannot prove absolute consistency. One form of consistency involves a set of correspondence rules to steer a transformation from one language to another. Thus, given a specification S_1 in viewpoint language L_1 and specification S_2 in viewpoint language L_2 , where S_1 and S_2 both specify the same system, a transformation T can be applied to S_1 resulting in a new specification $T(S_1)$ in viewpoint language L_2 which can be compared directly to S_2 to check, for example, for behavioural compatibility between allegedly equivalent objects or configurations of objects.

10.1 Computational and information specification correspondences

This reference model does not prescribe exact correspondences between information objects and computational objects. In particular, not all states of a computational specification need correspond to states of an information specification. There may exist transitional computational states within pieces of computational behaviour which are abstracted as atomic transitions in the information specification.

Where an information object corresponds to a set of computational objects, static and invariant schemata of an information object correspond to possible states of the computational objects. Every change in state of an information object corresponds either to some set of interactions between computational objects or to an internal action of a computational object. The invariant and dynamic schemata of the information object correspond to the behaviour and environment contract of the computational objects.

NOTE – If the concept of information interface is used in an information specification, there is no necessary correspondence between an information interface and any computational interface.

10.2 Engineering and computational specification correspondences

Each computational object which is not a binding object corresponds to a set of one or more basic engineering objects (and any channels which connect them). All the basic engineering objects in the set correspond only to that computational object.

NOTE 1 – If no transparency is selected, each computational object which is not a binding object will correspond to a basic engineering object.

Except where transparencies which replicate objects are involved, each computational interface corresponds exactly to one engineering interface, and that engineering interface corresponds only to that computational interface.

NOTE 2 – The engineering interface is supported by one of the basic engineering objects which corresponds to the computational object supporting the computational interface.

Where transparencies which replicate objects are involved, each computational interface of the objects being replicated corresponds to a set of engineering interfaces, one for each of the basic engineering objects resulting from the replication. These engineering interfaces each correspond only to the original computational interface.

Each computational interface is identified by any member of a set of one or more computational interface identifiers. Each engineering interface is identified by any member of a set of one or more engineering interface references. Thus, since a computational interface corresponds to an engineering interface, an identifier for a computational interface can be represented unambiguously by an engineering interface reference from the corresponding set.

Each computational binding (either primitive bindings or compound bindings with associated binding objects) corresponds to either an engineering local binding or an engineering channel. This engineering local binding or channel corresponds only to that computational binding. If the computational binding supports operations, the engineering local binding or channel must support the interchange of at least:

- computational signature names;
- computational operation names;
- computational termination names;
- invocation and termination parameters (including computational interface identifiers and computational interface signatures).

Except where transparencies which replicate objects are involved, each computational binding object control interface has a corresponding engineering interface and there exists a chain of engineering interactions linking that interface to any stubs, binders, protocol objects or interceptors to be controlled in support of the computational binding.

NOTE 3 – The set of control interfaces involved depends on the type of the binding object.

Each computational interaction corresponds to some chain of engineering interactions, starting and ending with an interaction involving one or more of the basic engineering objects corresponding to the interacting computational objects.

Each computational signal corresponds either to an interaction at an engineering local binding or to a chain of engineering interactions which provides the necessary consistent view of the computational interaction.

The transparency prescriptions in clause 16 specify additional correspondences.

NOTE 4 – Basic engineering objects corresponding to different computational objects can be members of the same cluster.

NOTE 5 – In an entirely object-based computational language, data are represented as abstract data types (i.e., interfaces to computational objects).

NOTE 6 – Computational interface parameters (including those for abstract data types) can be passed by reference, such parameters correspond to engineering interface references.

NOTE 7 – Computational interface parameters (including those for abstract data types) can be passed by migrating or replicating the object supporting the interface. In the case of migration such parameters correspond to cluster templates.

NOTE 8 – If the abstract state of a computational object supporting an interface parameter is invariant, the object can be cloned rather than migrated.

NOTE 9 – Cluster templates can be represented as abstract data types, thus strict correspondences between computational parameters and engineering interface references are sufficient. The use of cluster templates or data are important engineering optimizations and therefore not excluded.

10.3 Technology and engineering specification correspondences

There is a technology object corresponding to every atomic or composite engineering object.

11 ODP functions

The ODP functions defined in this Recommendation | International Standard are those that are either fundamental or widely applicable to the construction of ODP systems.

The specifications for individual ODP functions can be combined to form specifications for components of ODP systems. The identification of such components is a matter for the standardization process and is not prescribed by this reference model: it contains outline function descriptions written using concepts from Rec. ITU-T X.902 | ISO/IEC 10746-2.

Some function descriptions in this reference model introduce objects as a simplifying modelling construct. Except where explicit constraints are given on the distribution of these objects, they do not define necessary structure in an implementation.

The functions defined in this reference model are listed below. Those forming an integral part of the computational language are tagged with a "*", those forming an integral part of the engineering language are tagged with a "+":

- a) *Management functions:*
 - 1) node management function⁺;
 - 2) object management function⁺;
 - 3) cluster management function⁺;
 - 4) capsule management function⁺.
- b) *Coordination functions:*
 - 1) event notification function;
 - 2) checkpointing and recovery function;
 - 3) deactivation and reactivation function;
 - 4) group function;
 - 5) replication function;
 - 6) migration function;
 - 7) engineering interface reference tracking function⁺;

- 8) transaction function.
- c) *Repository functions*:
 - 1) storage function;
 - 2) information organization function;
 - 3) relocation function;
 - 4) type repository function;
 - 5) trading function^{†*}.
- d) *Security functions*:
 - 1) access control function;
 - 2) security audit function;
 - 3) authentication function;
 - 4) integrity function;
 - 5) confidentiality function;
 - 6) non-repudiation function;
 - 7) key management function.

12 Management functions

12.1 Node management function

The node management function controls processing, storage and communication functions within a node.

The node management function is provided by each nucleus at one or more node management interfaces.

Each capsule uses a node management interface distinct from the node management interfaces used by other capsules in the same node.

The node management function:

- manages threads;
- accesses clocks and manages timers;
- creates channels and locates interfaces.

Within the architecture defined by this reference model, the node management function is used by all other functions.

12.1.1 Thread management

Node management interfaces provide functions to spawn and fork threads within a capsule and to join, delay and synchronize threads within a capsule.

12.1.2 Clock access and timer management

Node management interfaces provide functions to determine the current time in a specified clock management domain and to start, monitor and cancel timers.

12.1.3 Channel creation and interface location

Node management interfaces provide functions to:

- a) enable binding between an engineering object within a capsule and an instance of the trading function;
- b) make an engineering interface available for binding to objects in other capsules;
- c) establish a binding between an engineering object within a capsule and a set of other engineering objects (identified by engineering interface references);
- d) from a specified engineering interface reference, determine the channel type and communications interface it implies.

NOTE 1 – The use of interactions b) and c) is explained in 8.2.3.

NOTE 2 – Interaction d) exposes the information content of the interface reference, enabling it to be stored or transformed for use in a different interface reference management domain.

Making an interface available for binding to objects in other capsules – interaction b) above – consists of the following actions:

- assigning an engineering interface reference within a nominated engineering interface reference management domain;
- assigning a communications interface through which bindings to the interface can be established;
- assigning a channel type to the interface.

12.1.4 Capsule template instantiation and capsule deletion

Node management interfaces provide functions to instantiate capsule templates and to delete capsules.

Capsule template instantiation consists of the following steps:

- allocation of processing, storage and communication functions for a new capsule in the same node as the nucleus providing the node management interface;
- creating a capsule manager for the new capsule;
- creating a capsule *management* interface in the new capsule manager;
- producing an identifier for the capsule management interface;
- creating a capsule *control* interface for the new capsule in the nucleus;
- producing an identifier for the capsule control interface.

The capsule control interface produced by capsule template instantiation enables deletion of the capsule (e.g., when its manager has failed).

Deletion of a capsule deletes all objects in the capsule.

12.2 Object management function

The object management function checkpoints and deletes objects.

When an object belongs to a cluster that can be deactivated, checkpointed or migrated, the object must have an object management interface in which it provides one or more of the following functions:

- checkpointing the object;
- deleting the object.

NOTE 1 – Thus different object management interfaces can have different interface types depending upon which functions they support.

NOTE 2 – Checkpointing an object produces information suitable for incorporation into a cluster checkpoint.

NOTE 3 – When an object is deleted, stubs, binders, protocol objects and interceptors supporting bindings to the object can be deleted.

The object management function is used by the cluster management function.

12.3 Cluster management function

The cluster management function checkpoints, recovers, migrates, deactivates or deletes clusters and is provided by each cluster manager at a cluster management interface, comprising one or more of the following functions with respect to the managed cluster:

- modifying cluster management policy (e.g., for the location of checkpoints of its cluster, for the use of the relocation function to trigger reactivation or recovery of the cluster);
- deactivating the cluster;
- checkpointing the cluster;
- replacing the cluster with a new one instantiated from a cluster checkpoint (i.e., deletion followed by recovery);
- migrating the cluster to another capsule (using the migration function);
- deleting the cluster.

NOTE – Thus different cluster management interfaces can have different interface types depending upon which functions they support.

The behaviour of a cluster manager is constrained by the management policy for its cluster. Cluster checkpointing and deactivation is only possible if all objects in the cluster have object management interfaces supporting the object

checkpointing function. Deactivation and cluster deletion both require that the objects in the cluster support object deletion.

Within the architecture defined by this reference model:

- the cluster management function is used by the capsule management function, the deactivation and reactivation function, the checkpoint and recovery function, the migration function and the engineering interface reference management function;
- the cluster management function uses the storage function for keeping checkpoints.

12.3.1 Cluster checkpoint

A cluster checkpoint contains information needed to re-establish a cluster and includes the following elements:

- a) checkpoints of the objects in the cluster;
- b) the configuration of the objects in the cluster;
- c) sufficient information to re-establish distributed bindings that involve objects in the cluster.

NOTE – Engineering interface references are an essential part of the information required to establish bindings. A cluster checkpoint will contain all of the engineering references arising from a) and c) above.

12.3.2 Cluster deletion, deactivation and failure

Cluster deletion deletes all the objects in the cluster, the cluster manager and any objects solely supporting the cluster or its cluster manager (e.g., stubs and binders). Cluster deactivation is coordinated by the deactivation and reactivation function; it produces a checkpoint of the cluster, then deletes the cluster and its supporting structure. Cluster failure causes deletion of all objects in the cluster and, in some cases, leads to deleting the cluster's supporting structure.

12.3.3 Cluster reactivation and recovery

A deactivated cluster can be reactivated from one of its checkpoints. Cluster reactivation is necessarily a capsule management function because the cluster manager is deleted as part of cluster deactivation. A cluster can be recovered from one of its checkpoints. If the associated cluster manager has not been deleted, it can initiate recovery; otherwise recovery is a capsule management function and includes creation of a new cluster manager.

12.3.4 Cluster migration

Cluster migration consists of the cloning of a source cluster in a target capsule followed by deletion of the source cluster. It is coordinated by the migration function and is parameterized by a capsule management interface for the target capsule for the migrated cluster.

12.4 Capsule management function

The capsule management function instantiates clusters (including recovery and reactivation), checkpoints all the clusters in a capsule, deactivates all the clusters in a capsule and deletes capsules. It is provided by each capsule manager at a capsule management interface comprising one or more of the following functions with respect to the managed capsule:

- instantiation (within the capsule) of a cluster template;
NOTE 1 – This includes reactivation and recovery.
- deactivation of the capsule by deactivating all the clusters within it (using the cluster management function);
- checkpointing the capsule by checkpointing all the clusters in the capsule (using the cluster management function);
- deleting the capsule, by deleting all the clusters within it, followed by deletion of the capsule manager for the capsule.

NOTE 2 – Thus different capsule management interfaces can have different types, depending upon which functions they support.

The behaviour of a capsule manager is constrained by the management policy for its capsule.

Within the architecture defined by this reference model, the capsule management function is used by the deactivation and reactivation function, the checkpoint and recovery function and the migration function.

12.4.1 Cluster template instantiation

Cluster template instantiation is parameterized by a cluster template and consists of the following steps:

- instantiating a cluster from a cluster template;
- introducing a cluster manager for the new cluster;
- producing an identifier for a cluster management interface in the new cluster manager;
- binding the new cluster to other objects according to the engineering language rules and the binding information in the cluster template.

NOTE – Cluster reactivation and cluster recovery are special cases of cluster instantiation, in which case the cluster template is a cluster checkpoint.

A cluster template can contain information specific to a domain. If the template is to be instantiated in another domain, then this information must be transformed. In particular, engineering interface references contained in the cluster must be transformed if the cluster is being instantiated in a different engineering interface reference management domain.

12.4.2 Capsule deletion

Capsule deletion entails deletion of the capsule manager and can lead to deletion of stubs, binders, protocol objects or interceptors that supported objects in the capsule or its manager.

13 Coordination functions

13.1 Event notification function

The event notification function records and makes available event histories.

13.1.1 Concepts

13.1.1.1 event history: An object representing significant actions.

13.1.2 Rules

Event producer objects interact with the event notification function to create event histories. The event notification function notifies event consumer objects of the availability of event histories.

The event notification function supports one or more event history types and has an event notification policy which determines the behaviour of the function, in particular:

- which objects can create event histories;
- which objects are notified of the creation of a new event history;
- the means by which such notifications occur;
- persistence and stability requirements for event histories;
- ordering relationships between interactions with event producer and event consumer objects.

An event consumer interacts with the event notification function to register for notification of new event histories. Depending upon the event notification policy, the interaction can:

- establish bindings to currently available event histories;
- enable communication about event histories created subsequent to the interaction.

NOTE – Event histories with stringent persistence and stability requirements can be supported using the transaction and replication functions. Ordering and multi-casting notifications can be supported using the group function.

13.2 Checkpoint and recovery function

The checkpoint and recovery function coordinates the checkpointing and recovery of failed clusters.

The checkpoint and recovery function embodies policies governing:

- when clusters should be checkpointed;
- when clusters should be recovered;
- where clusters should be recovered;

- where checkpoints should be stored;
- which checkpoint is recovered.

The checkpointing and recovery of clusters is subject to any security policy associated with those clusters, in particular, where the checkpoint is stored and where the checkpoint is recovered.

Within the architecture defined by this reference model, the checkpoint and recovery function uses the cluster management function and the capsule management function.

13.2.1 Checkpointing

Checkpointing is the responsibility of the object management function and the cluster management function. The checkpointing of a cluster is coordinated by its cluster manager: first the cluster manager uses the object management function to obtain a checkpoint of each object in the cluster; from these object checkpoints the cluster manager constructs a cluster checkpoint which it then makes persistent, using the storage function.

Depending upon the checkpointing policy, checkpointing a cluster can lead to checkpointing of other clusters which participate in common activities with the checkpointed cluster, subject to the following consistency rules:

- the initial cluster must be in a consistent state before it is checkpointed;
- there must be consistency between all of the cluster checkpoints of the various clusters being jointly checkpointed (e.g., all of the checkpoints reflect the same set of interactions having occurred between the clusters);
- when a cluster is checkpointed, all other clusters having checkpoint constraints with the cluster must be checkpointed. This rule must be applied recursively to obtain a closed set of clusters which can be checkpointed consistently.

13.2.2 Recovery

A cluster can be recovered either:

- in the capsule from which it was previously checkpointed; or
- in another capsule (e.g., when the capsule in which the checkpoint occurred has subsequently failed).

Recovery of a cluster is the responsibility of the cluster manager for that cluster, or in the absence of that cluster manager, a capsule manager. The checkpoint and recovery function interacts with the cluster manager or the capsule manager as appropriate to instantiate the cluster checkpoint. Before recovering a cluster, the checkpoint and recovery function must ensure that the cluster has been removed (e.g., as a consequence of failure). The objects bound to recovered clusters must be able to detect that the cluster has recovered from a checkpoint (e.g., so that they may re-do interactions that occurred after the checkpoint was taken).

The recovery of one cluster can lead to the recovery of other clusters, for example those forming a joint checkpoint with the cluster that is recovered.

13.3 Deactivation and reactivation function

The deactivation and reactivation function coordinates the deactivation and reactivation of clusters. It embodies policies governing:

- when clusters should be deactivated;
- where the checkpoint associated with a deactivation should be stored;
- when clusters should be reactivated;
- which checkpoint should be reactivated (e.g., the most recent);
- where clusters should be reactivated.

The deactivation and reactivation of clusters is subject to any security policy associated with those clusters, in particular within the architecture defined by this reference model, the deactivation and reactivation function uses the object management function, the cluster management function and the capsule management function. The deactivation and reactivation function is used by the migration function.

13.3.1 Deactivation

Deactivation of a cluster is a cluster management function and comprises the following steps:

- the cluster manager for the cluster involved interacts with each object in its cluster to obtain a checkpoint that can be used to make a cluster checkpoint;

- the cluster manager makes the cluster checkpoint persistent, using the storage function;
- the cluster manager deletes the cluster (and can itself be deleted).

13.3.2 Reactivation

The deactivation and reactivation function reactivates a cluster by using the capsule management function to instantiate a checkpoint of the cluster in the target capsule (and includes creation of a cluster manager for the cluster). The target capsule can either be the capsule from which it was previously deactivated or another capsule (e.g., to balance infrastructure load across several nodes).

13.4 Group function

The group function provides the necessary mechanisms to coordinate the interactions of objects in multi-party binding.

13.4.1 Concepts

13.4.1.1 interaction group: A subset of the objects participating in a binding managed by the group function.

13.4.2 Rules

For each set of objects that is bound together in an interaction group, the group function manages:

- interaction: deciding which members of the group participate in which interactions, according to an interaction policy;
- collation: derivation of a consistent view of interactions (including failed interactions), according to a collation policy;
- ordering: ensuring that interactions between group members are correctly ordered with respect to an ordering policy;
- membership: dealing with member failure and recovery, and addition and removal of members according to a membership policy.

NOTE – The behaviour of the binding object linking members of the group determines how interaction is to be effected.

13.5 Replication function

The replication function is a special case of the group function in which the members of a group are behaviourally compatible (e.g., because they are replicas from the same object template). The replication function ensures the group appears to other objects as if it were a single object by ensuring that all members participate in all interactions and that all members participate in all interactions in the same order.

The membership policy for a replica group can allow for the number of members in a replica group to be increased or decreased. Increasing the size of a replica group achieves the same effect as if a member of the group had been cloned and then added to the group in a single atomic action.

For the replication function to be applied to a cluster, the objects comprising the cluster are replicated and configured into a set of identical clusters. The corresponding objects in each such replicated cluster form replica groups. Thus, a replicated cluster is a coordinated set of replica groups.

The replication function is used by the migration function.

13.6 Migration function

The migration function coordinates the migration of a cluster from one capsule to another. It uses the cluster management function and the capsule management function and embodies policies governing when clusters should be migrated and where they can be located.

Two possible means of migration are:

- replication; or
- deactivation in one capsule followed by reactivation in another.

13.6.1 Replication

Migration of a cluster by use of the replication function comprises the following sequence of actions:

- the old cluster is treated as a cluster replica group of size one;
- a copy of the original cluster is created in the destination capsule, together with a cluster manager;

- the objects in both clusters are formed into replica groups (of size two);
- the objects in the old cluster are removed from the object groups (leaving groups of size one);
- the old cluster (and its manager) is deleted.

13.6.2 Deactivation and reactivation

Migration of a cluster by deactivation and reactivation is coordinated by the cluster's manager, and comprises deactivating the cluster at its old location, followed by reactivating the cluster at its new location.

13.7 Transaction function

13.7.1 Concepts

13.7.1.1 transaction: An activity which leads to a set of object state changes consistent with a dynamic schema (and its constraining invariant schema).

13.7.1.2 action of interest: An action in a transaction which leads to a state change of significance to the transaction.

13.7.1.3 visibility: The degree to which a transaction can access object state concurrently with other transactions.

13.7.1.4 recoverability: The degree to which object state changes resulting from failed transactions are cancelled.

13.7.1.5 permanence: The degree to which failures can affect object state changes due to completed transactions.

13.7.2 Rules

The transaction function coordinates and controls a set of transactions to achieve a specified level of visibility, recoverability and permanence.

The transaction function:

- interacts with objects to monitor the occurrence of actions of interest, cancellation of the effects of actions of interest and the causality of actions of interest;
- decides if actions of interest are in conflict;
- interacts with objects to schedule the occurrence of actions of interest to prevent conflicts;
- interacts with objects to cancel the effects of actions of interest that have occurred, in order to resolve conflicts.

It is subject to policies determining:

- which actions are actions of interest;
- which actions of interest conflict;
- which actions of interest are to be cancelled to resolve conflicts.

13.8 ACID transaction function

The ACID transaction function is a special case of the general transaction function in which:

- visibility is specified as isolation of transactions from one another;
- recoverability is specified as the requirement that transactions be atomic;
- permanence is specified as the requirement that state changes arising from transactions be durable (i.e., stable); and
- consistency is achieved by a correct execution of transactions with the atomicity, isolation and durability properties in conformance with the associated dynamic and invariant schemata.

The actions of interest for the ACID transaction function are:

- start transaction;
- commit transaction;
- abort transaction;
- access state;
- modify state;
- cancellation of any of these.

Transaction policies are expressed as serializability rules for transactions.

13.9 Engineering interface reference tracking function

The engineering interface reference tracking function monitors the transfer of engineering interface references between engineering objects in different clusters to determine when the infrastructure associated with engineering interfaces is no longer required (i.e., when no object in another cluster is in a position to bind to the referenced interface).

The engineering interface reference tracking function maintains, within its scope:

- information on the possession of engineering interface references;
- information on the existence of interfaces.

The engineering interface reference tracking function:

- is notified by stubs when an engineering interface reference is passed between clusters;
- is notified by a cluster manager when all copies of an engineering interface reference are deleted in its cluster;
- detects when no copies of an engineering interface reference are held outside the cluster supporting the engineering interface referred to and notifies the corresponding cluster manager;
- notifies the managers of capsules in which there are holders of engineering interface references to an interface which has failed, or is deleted.

NOTE – The notification of engineering interface reference tracking events can be achieved using the event notification function.

The engineering interface reference management function is constrained by the engineering interface reference management policies of the engineering interface reference management domains to which it applies.

14 Repository functions

14.1 Storage function

The storage function stores data.

14.1.1 Concepts

14.1.1.1 data repository: An object providing the storage function.

14.1.1.2 container interface: An interface of a data repository allowing access to data.

14.1.2 Rules

A data repository stores sets of data. Each set of data is associated with a container interface created when the data are stored. A container interface provides functions to:

- obtain a copy of the data stored by the interface;
- modify the data associated with the interface;
- delete the container interface and its associated data.

NOTE – Objects embody both actions and state (i.e., process and data); they are inherently persistent. The checkpoint and recovery function or replication function can be used in an infrastructure of the storage function to provide stability. The functions can use other instances of the storage function as a repository for cluster checkpoints.

14.2 Information organization function

The information organization function manages a repository of information described by an information schema and includes some or all of the following elements:

- modifying and updating the information schema;
- querying the repository, using a query language;
- modifying and updating the repository.

The form and nature of queries and responses to queries depend upon the query language. The information organization function does not permit modifications or updates to the information repository that are inconsistent with its schema.

The information organization function can be modelled as a repository of objects (with computational interfaces) standing in correspondence to entities and relationships in an ODP system. These objects support operations to:

- define attributes, properties and relationships for objects in the repository;
- add and delete objects in the repository;
- assert and delete attributes, properties and relations for selected objects in the repository;
- select objects which satisfy a predicate (i.e., a type) specified in terms of attributes, properties and relations.

NOTE 1 – The information organization function might derive additional relationships to those which are instantiated directly.

NOTE 2 – The information organization function can be used to maintain the relationship between a cluster and its checkpoint (i.e., a storage interface at which the checkpoint can be created or accessed).

NOTE 3 – In order to allow recovery in response to interactions with a failed cluster, the information organization function can be used to maintain a relationship between a cluster and the interface at which recovery of the cluster can be requested.

NOTE 4 – In order to allow reactivation in response to interactions with a deactivated cluster, the information organization function can be used to maintain a relationship between a deactivated cluster and the interface at which reactivation of the cluster can be requested.

NOTE 5 – The information needed by a relocater to validate or update an interface reference can be maintained by the information organization function.

14.3 Relocation function

The relocation function manages a repository of locations for interfaces, including locations of management functions for the cluster supporting those interfaces.

14.3.1 Concepts

14.3.1.1 relocater: An object providing the relocation function.

14.3.2 Rules

A relocater has a directory of locations for interfaces that have had their locations changed as a consequence of either communications domain management (e.g., changing a node's network address) or cluster management activities such as deactivation, migration, replication or recovery of a cluster checkpoint.

An interface can be associated with a relocater; relocaters can be specific to a single interface or common to several interfaces. When the scope of a relocater spans more than one engineering interface reference management domain, the means by which the relocater is accessed must make clear the interface reference management domain that applies.

When a relocater is associated with an interface, activities which change the location of interfaces must inform that relocater of the new location, in particular a cluster manager must notify the relocater for each interface of each object in the cluster when the cluster is relocated; only interfaces of objects which have had their locations changed need be recorded. When an engineering interface reference is to remain valid even though the object supporting it is in a deactivated cluster or failed but previously checkpointed cluster, the relocater must embody a policy for use of the coordination functions to restore the cluster, by reactivation or recovery as appropriate, when another object attempts to validate the reference.

A relocater supports interactions that:

- record a change in the location of an interface identified by an engineering interface reference;
- validate the location of an interface identified by an engineering interface reference (including, if needed, restoration of the cluster containing the object supporting the interface);
- set policy for interacting with coordination functions (e.g., the deactivation and reactivation function) when validating the location of an interface identified by an engineering interface reference.

NOTE – An object performing engineering interface reference validation can retain the results of the validation to optimize future use of the reference.

14.4 Type repository function

The type repository function manages a repository of type specifications and type relationships. It has an interface for each type specification it stores.

NOTE – The nature of the type repository function, and the engineering mechanisms that support it are further specified in Rec. ITU-T X.960 (1999) | ISO/IEC 14769:2001, *Information technology – Open Distributed Processing – Type repository function*.

14.4.1 Rules

The type repository function can be informed of type relationships in addition to those it can derive from comparison of type specifications. A type repository will not permit inconsistent relationships to be established.

Type specifications are immutable.

The type repository function supports creating types and their associated type interfaces.

A type repository interface for a specific type provides functions to:

- query the type's specification;
- assert relationships between the type and other types;
- query relationships between the type and other types.

NOTE – Subtyping relationships for computational signature types are defined by the signature subtyping rules in 7.2.4. There is no requirement for a type repository to enable signature subtypes relationships to be computed. Where signature types are included in a type repository, the repository is not permitted to impose additional signature typing rules or assertions that contradict the provisions of 7.2.4.

14.5 Trading function

The trading function mediates advertisement and discovery of interfaces.

NOTE – The nature of the trading function, and the engineering mechanisms that support it are further specified in Rec. ITU-T X.950 (1997) | ISO/IEC 13235-1:1998, *Information technology – Open Distributed Processing – Trading Function: Specification*.

14.5.1 Concepts

14.5.1.1 service offer: Information describing an interface, how to bind to it, and the service that can be invoked using it. A service offer includes the interface type information needed to perform invocation, information needed to access the specific service instance, and any attributes describing the properties of the service.

NOTE 1 – The information in a service offer may be included within the offer explicitly, or a reference to it may be provided (see 8.2.2). The expression of the information depends on the engineering solutions in use. For example, an ODP trader service offer includes an interface reference, a type name and a set of attributes in name value form (see Rec. ITU-T X.952 | ISO/IEC 13235-3). A Web Services service offer is in the form of a URL that can be used to invoke the service or to retrieve a WSDL description of it. In Java RMI, the service offer also includes a complete proxy for client use.

NOTE 2 – Service offers provide, directly or indirectly, sufficient information to enable binding to the interfaces they describe.

NOTE 3 – The trading function uses information in the service offers it manages to select one or more services that satisfy constraints given by its clients on the interface type or service properties required. The language in which these constraints are expressed depends on the specific engineering solution in use.

14.5.1.2 service export: An interaction with the trading function in which a service offer is advertised, by adding the service offer to an identified set of service offers.

14.5.1.3 service import: An interaction with the trading function which searches an identified set of service offers to discover interfaces at which a service satisfying a specified type is available.

14.5.2 Rules

The trading function provides service import and service export and its behaviour is governed by a trading policy which states rules on how the sets identified in service export are related to the sets in service import. On service import, a trading function must select only offers that satisfy the policy of the trading function itself, the policy of the exporter of the service offer and the policy of the importer of the service offer. Service import involves computational interface signature sub/supertype checking. It can, in addition, involve further levels of checks, including checks on behavioural capability and environmental constraints.

15 Security functions

15.1 Concepts

The following concepts are common to all of the security functions.

15.1.1 security policy: A set of rules that constrains one or more sets of activities of one or more sets of objects.

15.1.2 security authority: The administrator responsible for the implementation of a security policy.

15.1.3 security domain: A domain in which the members are obliged to follow a security policy established and administered by a security authority.

NOTE – The security authority is the controlling object for the security domain.

15.1.4 security interaction policy: Those aspects of the security policies of different security domains that are necessary in order for interactions to take place between those domains.

15.2 Access control function

The access control function prevents unauthorized interactions with an object. It includes both an access control decision function and an access control enforcement function. Within the context of access control, objects fulfil the roles of either target or initiator. The function requires access control information about the target, the initiator and the interaction.

The initiator requests an interaction with the target from the access control function. The access control decision function decides whether access is permitted or denied on the basis of the access control information and the decision is enforced by the access control enforcement function.

NOTE – The access control decision function and the access control enforcement function can be provided by the object which has the role of target, or by other objects.

15.3 Security audit function

The security audit function provides monitoring and collection of information about security-related actions, and subsequent analysis of the information to review security policies, controls and procedures.

The security audit function includes each of the following elements:

- audit trail analyser function;
- audit recorder function;
- audit trail examiner function;
- audit trail collector function.

15.4 Authentication function

The authentication function provides assurance of the claimed identity of an object. In the context of authentication, objects fulfil one or more of the following roles:

- principal;
- claimant;
- trusted third party.

Authentication requires use of exchange authentication information.

NOTE 1 – Any identifiable object in an ODP system can be the principal for authentication, including both objects that model people and those that model computer systems.

NOTE 2 – The object initiating authentication is not necessarily the claimant.

There are two forms of authentication:

- peer entity authentication, providing corroboration of the identity of a principal within the context of a communication relationship;
- data origin authentication, providing corroboration of the identity of the principal responsible for a specific data unit.

NOTE 3 – Authentication mechanisms are categorized in Rec. ITU-T X.811 | ISO/IEC 10181-2.

In an authentication involving two objects, either or both objects can have the role of claimant. Where both objects have the role of claimant the style of authentication is known as mutual authentication. Exchange authentication information is passed from the initiating object to the responding object and further exchange authentication information may then be passed in the reverse direction. Additional exchanges may also take place: different authentication mechanisms require different numbers of exchanges. Peer entity authentication always involves interaction with the claimant. Data origin authentication need not involve interaction with the claimant.

A claimant supports operations to acquire information needed for an instance of authentication and to generate exchange authentication information. A verifier supports operations to acquire information needed for an instance of authentication, and to verify received exchange authentication information and/or to generate it. Information may be

exchanged with an authentication server and either the claimant or the verifier (or both) either prior to or during authentication exchanges.

The authentication function may use the key management function.

15.5 Integrity function

The integrity function detects and/or prevents the unauthorized creation, alteration or deletion of data.

The integrity function includes all the following functions:

- shield;
- validate;
- unshield.

In the context of integrity, objects fulfil one or more of the following roles:

- integrity-protected data originator;
- integrity-protected data recipient.

Integrity-protected data is passed from originator to recipient. An integrity-protected data originator supports an interface providing the shield function. An integrity-protected data recipient supports an interface providing the validate or unshield functions.

The integrity function may use the key management function.

15.6 Confidentiality function

The confidentiality function prevents the unauthorized disclosure of information.

The confidentiality function includes the functions hide and reveal.

In the context of confidentiality, objects fulfil either or both of the following roles:

- confidentiality-protected information originator;
- confidentiality-protected information recipient.

Confidentiality-protected information is passed from originator to recipient. A confidentiality-protected information originator supports an interface providing the hide function. A confidentiality-protected information recipient supports an interface providing the reveal function.

The confidentiality function may use the key management function.

15.7 Non-repudiation function

The non-repudiation function prevents the denial by one object involved in an interaction of having participated in all or part of the interaction.

In the context of non-repudiation, objects fulfil one or more of the following roles:

- (non-repudiable data) originator;
- (non-repudiable data) recipient;
- evidence generator;
- evidence user;
- evidence verifier;
- non-repudiation service requester;
- notary;
- adjudicator.

The non-repudiation function makes use of non-repudiation evidence. In non-repudiation with proof of origin, the originator has the role of non-repudiation evidence generator for the origination interaction and includes this evidence in an acknowledgement of participation in the interaction. The recipient has the role of evidence user and uses the services of an evidence verifier (which may be itself) to gain confidence in the adequacy of the evidence. In non-repudiation with proof of delivery, the recipient has the role of non-repudiation evidence generator for the delivery interaction and includes this evidence in an acknowledgement of participation in the interaction. The originator has the