
Information technology — Open Distributed Processing — Naming framework

*Technologies de l'information — Traitement distribué ouvert — Cadre
de dénomination*

Reference number
ISO/IEC 14771:1999(E)



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.

© ISO/IEC 1999

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 22 734 10 79
E-mail copyright@iso.ch
Web www.iso.ch

Printed in Switzerland

Contents

	<i>Page</i>
1 Scope	1
2 Normative references.....	1
2.1 Identical Recommendations International Standards.....	1
3 Definitions	1
3.1 Terms defined in other International Standards	2
4 Abbreviations	2
5 Overview of ODP naming	2
5.1 Introduction	2
5.2 Structure of naming contexts.....	3
5.3 Processes involving names	5
6 Basic Naming Concepts.....	5
7 Name resolution model.....	7
7.1 The name resolution step.....	7
7.2 The complete name resolution process.....	8
8 Name communication	9
9 Name comparison	10
10 Federation	10
10.1 Controlling the names exported – export contexts	10
10.2 Uniform naming of federation partners – federation contexts	11
10.3 Providing a uniform user view of federated naming	12
10.4 Joining a federation	13
10.5 Evolution and optimization	13
11 Compliance	14
11.1 Basis for compliance	14
11.2 Compliance requirements	15
Annex A – Relationship with Related Naming Services and Frameworks.....	16
A.1 Relationship with OMG Object Naming Service	16
A.2 Relationship with X/Open Federated Naming API	16
A.3 Relationship with OSI Naming and Addressing	17
Annex B – Configuring name services	18
B.1 Quality of name service concerns.....	18
B.2 Distributing the name service.....	18
B.3 Configuring name servers	19
Annex C – Application of Naming	23
Annex D – Bibliography	24
Index	25

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.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 3.

In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1. 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 International Standard may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

International Standard ISO/IEC 14771 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 7, *Software engineering*, in collaboration with ITU-T. The identical text is published as ITU-T Recommendation X.910.

Annexes A to C of this International Standard are for information only.

Introduction

Names and naming are central concepts for the design and construction of open distributed systems. Most existing systems are built upon the assumption of a naming scheme in which only one selected naming convention applies to all the entities of concern (a global naming scheme). This assumption proves unsatisfactory in the context of large, evolving, heterogeneous distributed systems, managed by different authorities. Therefore, ITU-T Recs. X.90x series | ISO/IEC 10746, the Reference Model for Open Distributed Processing (ODP-RM), defines names to be context-relative.

This Recommendation | International Standard expands on the naming concepts introduced in the ODP-RM. It provides a general framework for naming, giving concepts and procedures which support fully general context-relative naming. These concepts can be applied in any ODP viewpoint. They can be applied to any function which uses naming and is subject to distribution or federation.

www.iso.org

INTERNATIONAL STANDARD

ITU-T RECOMMENDATION

INFORMATION TECHNOLOGY – OPEN DISTRIBUTED PROCESSING –
NAMING FRAMEWORK**1 Scope**

This Recommendation | International Standard:

- defines a general framework for context-relative naming, refining and elaborating on the naming concepts defined in Part 2 of the ODP-RM;
- identifies and characterizes functions necessary to handle names in the context of a federation of different naming systems; and
- clarifies the relationship between the concepts of name management (i.e. federation and naming) in distributed computing systems.

It provides a general framework for the naming of entities of interest in ODP systems, which includes naming in the infrastructure of an ODP system, naming in the applications built on the infrastructure, and naming in the enterprise the system serves.

2 Normative references

The following ITU-T 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 Telecommunications Standards Bureau of the ITU maintains a list of currently valid ITU-T Recommendations.

2.1 Identical Recommendations | International Standards

- ITU-T Recommendations X.90x series | ISO/IEC 10746, *Open Distributed Processing*.
- ITU-T Recommendations X.500 series | ISO/IEC 9594 (all parts): *The Directory*.
- ITU-T Recommendation X.650 (1996) | ISO/IEC 7498-3:1997, *Information technology – Open Systems Interconnection – Basic reference model: Naming and addressing*.
- ITU-T Recommendation X.902 (1995) | ISO/IEC 10746-2:1996, *Information technology – Open Distributed Processing – Reference Model: Foundations*.
- ITU-T Recommendation X.903 (1995) | ISO/IEC 10746-3:1996, *Information technology – Open Distributed Processing – Reference Model: Architecture*.

3 Definitions

For the purposes of this Recommendation | International Standard, the following terms and definitions apply:

3.1 Terms defined in other International Standards

This Recommendation | International Standard makes use of the following terms defined in ITU-T Rec. X.902 | ISO/IEC 10746-2:

- ODP system;
- viewpoint;
- name;
- identifier;
- name space;
- naming context;
- naming action;
- naming domain;
- naming graph;
- name resolution;
- <X> domain.

This Recommendation | International Standard makes use of the following terms defined in ITU-T Rec. X.903 | ISO/IEC 10746-3:

- ODP function;
- ODP system;
- location transparency;
- migration transparency;
- community;
- <X> federation;
- operation.

4 Abbreviations

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

ODP	Open Distributed Processing
OSI	Open Systems Interconnection
XFN	X/Open Federated Naming

5 Overview of ODP naming

5.1 Introduction

ITU-T Rec. X.902 | ISO/IEC 10746-2 introduces the basic ODP naming concepts. This clause provides an informal discussion of the naming structures supported by those concepts.

In principle, using the ODP concepts, it would be possible to construct a single universal naming context for distributed systems – a *global naming scheme*. In such a naming scheme only one selected naming convention would apply to all the entities of concern in the ODP system. Names in such a scheme are usually referred to as *absolute names*, although they are, in fact, relative to a single agreed context from which all name resolution starts.

Such a scheme is based on the assumption that all the parties participating in distributed processing follow the same naming scheme. This implies that:

- all the parties in the ODP systems agree to use the same naming convention;
- there is a relevant name authority to administer the scheme.

These conditions can apply in some specialized environments or for systems administered by a single administration. However, a truly global scheme would require agreement on a universal name space to be used in the naming of any entities that require naming, and a universal mapping between the names from the name space and the entities. Many

independently developed naming schemes have already been implemented and many more are being proposed. The ODP naming model must accommodate these different naming schemes and allow their interconnection. A single universal naming context would not support the necessary heterogeneity, and could not be managed.

Moreover, in a very large naming system, naming must remain unambiguous and the name space must therefore be equally large. Each naming action must be checked for consistency with the effects of all previous naming actions. The time required for a naming action will thus become related to the number of previous naming actions, making the availability of the naming action process unpredictable. The performance of the naming action process can only be increased at the cost of reduced consistency; some names may be used more than once (creating homonyms) and name resolution becomes ambiguous.

A context-relative naming scheme allows the federation of naming contexts and solves many of these problems.

In a context-relative naming scheme, multiple naming contexts can apply to entities in different administrative domains of the ODP system, but these naming contexts can be related to one another, so that it is possible to refer from one naming context to an entity in another naming context. In order to achieve this, in addition to associating a name with an entity, a naming action can also associate a name with another naming context. Since a naming context is something of interest, it is an entity and can be named.

Such an approach provides:

- the ability to avoid the need for unmanageably large naming systems;
- autonomy of the elements of open distributed processing systems, so that elements administered under different naming schemes continue to work on their own, but can work together;
- ease in reaching agreements and getting commitments on a naming scheme, since smaller communities are involved, which can then be brought together to form a global ODP community by federating their naming schemes;
- efficient and compact representation of names within local communities;
- the ability to incorporate different existing "global" naming schemes.

Thus, the naming process and management of consistency of names become manageable problems and existing naming contexts can be accommodated.

5.2 Structure of naming contexts

Any large distributed system is likely to comprise a number of administrative and technology domains. As a result, it is likely that the system also comprises a number of naming contexts, each relating to a name space and a set of target entities. At any given time, not all of the names from the name space, and not all of the entities from the set of target entities will be involved in a naming context.

Where a number of naming contexts exist, it may at times be necessary for an entity in one naming context to name an entity in some other naming context. Denoting an entity in another naming context requires a name for the entity, and the identification of the naming context in which the name resolves to the entity. To support such identification, a naming context can name other naming contexts. However, not all naming contexts can be named from any particular naming context, Figure 1 shows an example of how naming contexts relate names to entities.




Figure 1 – Example of naming contexts

In the example in Figure 1, there are two naming contexts, each associated with:

- a set (TE₁, TE₂) of target entities. The target entities in TE₁ are: Ep, Eq, Er, Eu, and Ev; the target entities in TE₂ are Es, Er, Eu, and Ev;
- a Name Space (NS), generated by the small letters of the Roman Alphabet and shared by both contexts;
- a set of relationships between names and entities. The entity Eu is a context with relationships of: the name a with Eq, n with Er, and g with Ev. The entity Ev is a context with relationships of: the name h with Eu, n with Es, and z with Er. However, because of the relationship of a name with the naming context entity itself (g with Ev and h with Eu), context relative names of entities are possible. Thus, Eq in TE₁ can be named h.a from the context Ev, and Es in TE₂ can be named g.n from the context Eu. An entity that is in both TE₁ and TE₂, e.g. Er, can be named either directly or with a context relative name.

Table 1 shows how the names for the entities vary with the naming context in which the names are resolved. The separator syntax (the character after the "." in the example) is determined by the naming conventions of the context in which the name being terminated is defined.

Table 1– Context relative names for entities in Figure 1

Entity	name from context 1	name from context 2
Eq	a	h.a
Er	n or g.z	h.n or z
Es	g.n	n
Eu	– (or g.h)	h
Ev	g	– (or h.g)

From both contexts, one entity (Eq as a from Ev and Es as n from Ev) is only named directly, one entity (Eq as h.a from Ev and Es as g.n from Eu) is only named indirectly via the other context, and one entity (Er) has two names, one direct and the other indirect. One entity, Ep, is left unnamed.

NOTE – Redundant names for Eq such as g.h.a from context 1 can be generated, but those can be optimized.