INTERNATIONAL STANDARD

ISO/IEC 14771

First edition
1999-12-01

Information technology — Open Distributed Processing — Naming framework

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

Reference number
ISO/IEC 14771:1999(E)

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

		Page
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.

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.

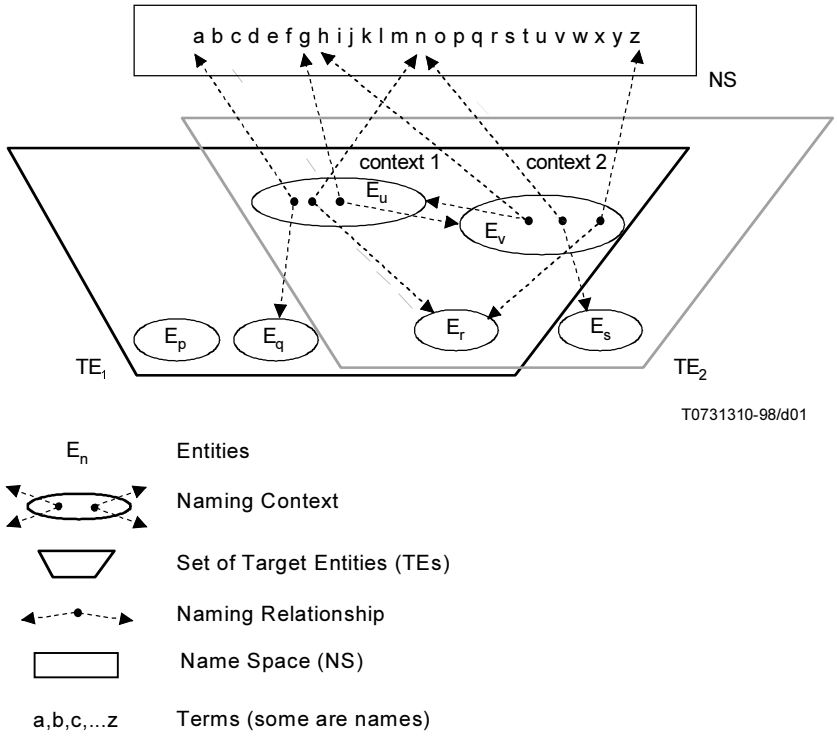

Figure 1 – Example of naming contexts

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

- a set (TE_1, TE_2) of target entities. The target entities in TE_1 are: Ep, Eq, Er, Eu, and Ev; the target entities in TE_2 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 TE_1 can be named h.a from the context Ev, and Es in TE_2 can be named g.n from the context Eu. An entity that is in both TE_1 and TE_2, 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.

5.3 Processes involving names

This Recommendation | International Standard defines a naming model comprising concepts, rules and structures governing naming in ODP systems. It places constraints and gives guidance to the specifiers of ODP-compliant naming systems. The remainder of this Recommendation | International Standard defines the processes involving names in the following clauses:

- basic naming concepts (clause 6), including the management of names which involves naming and unnaming;
- resolution of names (clause 7), in which a name is interpreted in order to make it possible to interact with the entity named;
- communication of names (clause 8), which may involve the transfer of a name to an entity which interprets names in a different context from that used by the sender. In general, the communication of names will be an active process which involves transformation of the name so that, when resolved, it continues to refer to the same entity;
- comparison of names (clause 9), to determine whether two names are known to refer to the same entity (are synonyms). However, if the naming system is sufficiently complex (for example, involving loose federations), in some cases comparison may fail to identify synonyms;
- federation of naming systems (clause 10), which involves the definition of the abstract processes for name resolution, name communication, and name comparison necessary to handle names in the federation of different naming schemes.

6 Basic Naming Concepts

The basic naming concepts comprise concepts from ITU-T Rec. X.902 | ISO/IEC 10746-2 and concepts defined in this clause.

Definitions from ITU-T Rec. X.902 | ISO/IEC 10746-2 are reproduced here for completeness. In some cases these definitions are also refined and clarified.

6.1 name: A name is a term (a linguistic construct) which, in a given naming context, refers to an entity (see 12.1 of ITU-T Rec. X.902 | ISO/IEC 10746-2).

NOTE 1 – Names can be used in a number of ways. These include:

- A name as a basis for interaction: a name may be required in the interaction with some entity. The name is used to distinguish the entity which is the target of an interaction from all other entities. Such a name is known as an invocation name. An invocation name conveys an opportunity (not a right) to interact with the entity denoted by that name. In the computational model, for instance, interfaces, operations and terminations are given invocation names.
- A name as an attribute: a name that cannot be used to interact with an entity may be used to denote that entity. An attributive name is used in the interaction between two entities to refer to a third entity, which is itself kept outside the interaction pattern.
- A name as an entity: a name can itself be seen as an entity of interest in a system. In the information viewpoint, for instance, a customer name is treated as a particular piece of information.

NOTE 2 – A name can be used within a predicate; a name may be used as part of a predicate. For instance, in "the papers written by Salzer", Salzer is a name. In query languages, predicates are formed to specify a set of responses. In constructing the query, names may be used as a shorthand to determine information about an entity.

NOTE 3 – Names can exist at different levels of abstraction, e.g. an address is a name, an interface reference is a name, and an identifier is a name.

6.2 identifier: An unambiguous name, in a given naming context (see 12.2 of ITU-T Rec. X.902 | ISO/IEC 10746-2).

6.3 name space: A set of terms usable as names (see 12.3 of ITU-T Rec. X.902 | ISO/IEC 10746-2).

Each name space is defined by a naming convention. This Recommendation | International Standard places no constraints on the nature of the symbols that may be used to construct names.

6.4 naming convention: The specification of the syntax to generate a set of terms to be used as names and of the algorithm to be used to parse those names.

A naming convention generates a name space. In general, the syntax comprises a part called a **handle** which can be used immediately in resolving the name and a part called a **remainder**, the resolution of which depends on the interpretation of the handle.

A naming convention influences the way in which name resolution proceeds with respect to a given name:

- left to right, as in "/usr/etc/ping";
- right to left, as in "support@iso.org.ch";
- no specific order, as in /S=X/P=SA/A=Telememo/C=AU;
- in some other fashion.

Individual naming systems may each impose a different constraint on the symbol set that may be used to construct names. Where entities in one naming system can be referred to from another naming system, a name translation might be required.

6.5 naming context: A relation between names and entities. Each name is drawn from a name space (a set of terms generated by a naming convention) that can be used as names in the context. Each entity is drawn from a set of target entities that can be named in the context (see 12.4 of ITU-T Rec. X.902 | ISO/IEC 10746-2).

This Recommendation | International Standard places no constraints on what may be named. Any entity, of any type, including entities external to a system, can be named.

NOTE 1 – This definition is a refinement of the definition in 12.4 ITU-T Rec. X.902 | ISO/IEC 10746-2 that supports the definition of a naming graph.

In general, not all terms are used, and not all entities are named in any particular context.

NOTE 2 – Not all names in the name space have to be bound to an entity. Names that are unbound, but nevertheless used, are unresolveable.

Each naming action and each name resolution is relative to some context.

A naming context may specify predicates to be satisfied by all the entities which are to be named from it (such as a constraint that they be of a certain type).

A naming context may specify constraints on the number of names from that context which can be associated with a single entity. However, an entity can still be given alternative names in other contexts.

A naming context may specify constraints on the number of entities that can be associated with a particular name. The constraint that at most one entity can be associated with a single name is common. If more than one entity is associated with a single name, the context may include rules for selection of a single entity as the result of name resolution.

A naming context shall not place any constraint on the length or content of the remainder resulting from name resolution.

6.6 synonym: A member of a set of names which all resolve to the same entity.

6.7 homonym: A name which, in a particular context, applies to more than one entity.

6.8 naming action: An action which associates a term from a name space with a given entity.

All naming actions are relative to a naming context (see 12.5 of ITU-T Rec. X.902 | ISO/IEC 10746-2).

NOTE – An entity is not necessarily involved when it is named, and so is not necessarily able to determine all its own names.

6.9 unnaming action: An action which removes the relationship between an entity and a name in a given context.

6.10 naming domain: A subset of a naming context such that all naming actions are performed by the controlling object of the domain (the name authority object) (see 12.6 of ITU-T Rec. X.902 | ISO/IEC 10746-2).

NOTE 1 – The set of entities in a distributed system can be referenced from a number of domains. There are many reasons for the introduction of multiple domains. For instance, domains may be used to delineate authorities, or to distinguish areas in which different security, management, or naming policies apply. Where multiple domains arise, different relationships between domains may exist. There is no requirement for different kinds of domain to coincide with regard to their membership.

NOTE 2 – The controlling object may perform other functions, in addition to performing the naming and unnaming actions.

6.11 name authority A name authority is the controlling object of a naming domain.

NOTE – Name authorities can form hierarchies to reflect the authority/sub-authority relationship, a domain may be a sub-domain of another domain. The sub-domain may have further sub-domains.

6.12 naming graph: A directed graph where each vertex denotes a naming context, and where each edge denotes an association between:

- a name appearing in the source naming context, and
- the target naming context

(see 12.7 of ITU-T Rec. X.902 | ISO/IEC 10746-2).