Interoperable Naming Service Specification

ISSUE REPORTING

All OMG specifications are subject to continuous review and improvement. As part of this process we encourage readers to report any ambiguities, inconsistencies, or inaccuracies they may find by completing the issue reporting form at *http://www.omg.org/library/issuerpt.htm.*

Contents

Contents

Preface

About This Document

Under the terms of the collaboration between OMG and X/Open Co Ltd, this document is a candidate for endorsement by X/Open, initially as a Preliminary Specification and later as a full CAE Specification. The collaboration between OMG and X/Open Co Ltd ensures joint review and cohesive support for emerging object-based specifications.

X/Open Preliminary Specifications undergo close scrutiny through a review process at X/Open before publication and are inherently stable specifications. Upgrade to full CAE Specification, after a reasonable interval, takes place following further review by X/Open. This further review considers the implementation experience of members and the full implications of conformance and branding.

Object Management Group

The Object Management Group, Inc. (OMG) is an international organization supported by over 800 members, including information system vendors, software developers and users. Founded in 1989, the OMG promotes the theory and practice of object-oriented technology in software development. The organization's charter includes the establishment of industry guidelines and object management specifications to provide a common framework for application development. Primary goals are the reusability, portability, and interoperability of object-based software in distributed, heterogeneous environments. Conformance to these specifications will make it possible to develop a heterogeneous applications environment across all major hardware platforms and operating systems.

OMG's objectives are to foster the growth of object technology and influence its direction by establishing the Object Management Architecture (OMA). The OMA provides the conceptual infrastructure upon which all OMG specifications are based.

What is CORBA?

The Common Object Request Broker Architecture (CORBA), is the Object Management Group's answer to the need for interoperability among the rapidly proliferating number of hardware and software products available today. Simply stated, CORBA allows applications to communicate with one another no matter where they are located or who has designed them. CORBA 1.1 was introduced in 1991 by Object Management Group (OMG) and defined the Interface Definition Language (IDL) and the Application Programming Interfaces (API) that enable client/server object interaction within a specific implementation of an Object Request Broker (ORB). CORBA 2.0, adopted in December of 1994, defines true interoperability by specifying how ORBs from different vendors can interoperate.

X/Open

X/Open is an independent, worldwide, open systems organization supported by most of the world's largest information system suppliers, user organizations and software companies. Its mission is to bring to users greater value from computing, through the practical implementation of open systems.

Intended Audience

The specifications described in this manual are aimed at software designers and developers who want to produce applications that comply with OMG standards for object services; the benefits of compliance are outlined in the following section, "Need for Object Services."

Need for Object Services

To understand how Object Services benefit all computer vendors and users, it is helpful to understand their context within OMG's vision of object management. The key to understanding the structure of the architecture is the Reference Model, which consists of the following components:

- **Object Request Broker**, which enables objects to transparently make and receive requests and responses in a distributed environment. It is the foundation for building applications from distributed objects and for interoperability between applications in hetero- and homogeneous environments. The architecture and specifications of the Object Request Broker are described in *CORBA: Common Object Request Broker Architecture and Specification.*

- **Object Services**, a collection of services (interfaces and objects) that support basic functions for using and implementing objects. Services are necessary to construct any distributed application and are always independent of application domains.

- **Common Facilities**, a collection of services that many applications may share, but which are not as fundamental as the Object Services. For instance, a system management or electronic mail facility could be classified as a common facility.

The Object Request Broker, then, is the core of the Reference Model. Nevertheless, an Object Request Broker alone cannot enable interoperability at the application semantic level. An ORB is like a telephone exchange: it provides the basic mechanism for making and receiving calls but does not ensure meaningful communication between subscribers. Meaningful, productive communication depends on additional interfaces, protocols, and policies that are agreed upon outside the telephone system, such as telephones, modems and directory services. This is equivalent to the role of Object Services.

What Is an Object Service Specification?

A specification of an Object Service usually consists of a set of interfaces and a description of the service's behavior. The syntax used to specify the interfaces is the OMG Interface Definition Language (OMG IDL). The semantics that specify a services's behavior are, in general, expressed in terms of the OMG Object Model. The OMG Object Model is based on objects, operations, types, and subtyping. It provides a standard, commonly understood set of terms with which to describe a service's behavior.

(For detailed information about the OMG Reference Model and the OMG Object Model, refer to the *Object Management Architecture Guide).*

Associated OMG Documents

The CORBA documentation is organized as follows:

- *Object Management Architecture Guide* defines the OMG's technical objectives and terminology and describes the conceptual models upon which OMG standards are based. It defines the umbrella architecture for the OMG standards. It also provides information about the policies and procedures of OMG, such as how standards are proposed, evaluated, and accepted.

- CORBA Platform Technologies
  - *CORBA: Common Object Request Broker Architecture and Specification* contains the architecture and specifications for the Object Request Broker.
  - *CORBA Languages*, a collection of language mapping specifications. See the individual language mapping specifications.
  - *CORBA Services,* a collection of specifications for OMG's Object Services. See the individual service specifications.
  - *CORBA Facilities,* a collection of specifications for OMG's Common Facilities. See the individual facility specifications.

- CORBA Domain Technologies
  - *CORBA Manufacturing*, a collection of specifications that relate to the manufacturing industry. This group of specifications defines standardized object-oriented interfaces between related services and functions.
  - *CORBA Med*, a collection of specifications that relate to the healthcare industry and represents vendors, healthcare providers, payers, and end users.

- *CORBA Finance*, a collection of specifications that target a vitally important vertical market: financial services and accounting. These important application areas are present in virtually all organizations: including all forms of monetary transactions, payroll, billing, and so forth.
- *CORBA Telecoms*, a collection of specifications that relate to the OMG-compliant interfaces for telecommunication systems.

The OMG collects information for each book in the documentation set by issuing Requests for Information, Requests for Proposals, and Requests for Comment and, with its membership, evaluating the responses. Specifications are adopted as standards only when representatives of the OMG membership accept them as such by vote. (The policies and procedures of the OMG are described in detail in the *Object Management Architecture Guide*.)

To obtain print-on-demand books in the documentation set or other OMG publications, contact the Object Management Group, Inc. at:

OMG Headquarters
250 First Avenue, Suite 201
Needham, MA 02494
USA
Tel: +1-781-444-0404
Fax: +1-781-444-0320
pubs@omg.org
http://www.omg.org

Service Design Principles

Build on CORBA Concepts

The design of each Object Service uses and builds on CORBA concepts:
- Separation of interface and implementation
- Object references are typed by interfaces
- Clients depend on interfaces, not implementations
- Use of multiple inheritance of interfaces
- Use of subtyping to extend, evolve and specialize functionality

Other related principles that the designs adhere to include:
- Assume good ORB and Object Services implementations. Specifically, it is assumed that CORBA-compliant ORB implementations are being built that support efficient local and remote access to "fine-grain" objects and have performance characteristics that place no major barriers to the pervasive use of distributed objects for virtually all service and application elements.
- Do not build non-type properties into interfaces

A discussion and rationale for the design of object services was included in the HP-SunSoft response to the OMG Object Services RFI (OMG TC Document 92.2.10).

Basic, Flexible Services

The services are designed to do one thing well and are only as complicated as they need to be. Individual services are by themselves relatively simple yet they can, by virtue of their structuring as objects, be combined together in interesting and powerful ways.

For example, the event and life cycle services, plus a future relationship service, may play together to support graphs of objects. Object graphs commonly occur in the real world and must be supported in many applications. A functionally-rich Folder compound object, for example, may be constructed using the life cycle, naming, events, and future relationship services as "building blocks."

Generic Services

Services are designed to be generic in that they do not depend on the type of the client object nor, in general, on the type of data passed in requests. For example, the event channel interfaces accept event data of any type. Clients of the service can dynamically determine the actual data type and handle it appropriately.

Allow Local and Remote Implementations

In general the services are structured as CORBA objects with OMG IDL interfaces that can be accessed locally or remotely and which can have local library or remote server styles of implementations. This allows considerable flexibility as regards the location of participating objects. So, for example, if the performance requirements of a particular application dictate it, objects can be implemented to work with a Library Object Adapter that enables their execution in the same process as the client.

Quality of Service is an Implementation Characteristic

Service interfaces are designed to allow a wide range of implementation approaches depending on the quality of service required in a particular environment. For example, in the Event Service, an event channel can be implemented to provide fast but unreliable delivery of events or slower but guaranteed delivery. However, the interfaces to the event channel are the same for all implementations and all clients. Because rules are not wired into a complex type hierarchy, developers can select particular implementations as building blocks and easily combine them with other components.

Objects Often Conspire in a Service

Services are typically decomposed into several distinct interfaces that provide different views for different kinds of clients of the service. For example, the Event Service is composed of *PushConsumer*, *PullSupplier* and *EventChannel* interfaces. This simplifies the way in which a particular client uses a service.

A particular service implementation can support the constituent interfaces as a single CORBA object or as a collection of distinct objects. This allows considerable implementation flexibility. A client of a service may use a different object reference to communicate with each distinct service function. Conceptually, these "internal" objects *conspire* to provide the complete service.

As an example, in the Event Service an event channel can provide both *PushConsumer* and *EventChannel* interfaces for use by different kinds of client. A particular client sends a request not to a single "event channel" object but to an object that implements either the *PushConsumer* and *EventChannel* interface. Hidden to all the clients, these objects interact to support the service.

The service designs also use distinct objects that implement specific service interfaces as the means to distinguish and coordinate different clients without relying on the existence of an object equality test or some special way of identifying clients. Using the event service again as an example, when an event consumer is connected with an event channel, a new object is created that supports the *PullSupplier* interface. An object reference to this object is returned to the event consumer which can then request events by invoking the appropriate operation on the new "supplier" object. Because each client uses a different object reference to interact with the event channel, the event channel can keep track of and manage multiple simultaneous clients. An event channel as a collection of objects conspiring to manage multiple simultaneous consumer clients.

Use of Callback Interfaces

Services often employ callback interfaces. Callback interfaces are interfaces that a client object is required to support to enable a service to *call back* to it to invoke some operation. The callback may be, for example, to pass back data asynchronously to a client.

Callback interfaces have two major benefits:

- They clearly define how a client object participates in a service.
- They allow the use of the standard interface definition (OMG IDL) and operation invocation (object reference) mechanisms.

Assume No Global Identifier Spaces

Several services employ identifiers to label and distinguish various elements. The service designs do not assume or rely on any global identifier service or global id spaces in order to function. The scope of identifiers is always limited to some context. For example, in the naming service, the scope of names is the particular naming context object.

In the case where a service generates ids, clients can assume that an id is unique within its scope but should not make any other assumption.

Finding a Service is Orthogonal to Using It

Finding a service is at a higher level and orthogonal to using a service. These services do not dictate a particular approach. They do not, for example, mandate that all services must be found via the naming service. Because services are structured as objects there does not need to be a special way of finding objects associated with services - general purpose finding services can be used. Solutions are anticipated to be application and policy specific.

Interface Style Consistency

Use of Exceptions and Return Codes

Throughout the services, exceptions are used exclusively for handling exceptional conditions such as error returns. Normal return codes are passed back via output parameters. An example of this is the use of a DONE return code to indicate iteration completion.

Explicit Versus Implicit Operations

Operations are always explicit rather than implied (e.g., by a flag passed as a parameter value to some "umbrella" operation). In other words, there is always a distinct operation corresponding to each distinct function of a service.

Use of Interface Inheritance

Interface inheritance (subtyping) is used whenever one can imagine that client code should depend on less functionality than the full interface. Services are often partitioned into several unrelated interfaces when it is possible to partition the clients into different roles. For example, an administrative interface is often unrelated and distinct in the type system from the interface used by "normal" clients.

Acknowledgments

Service Description   *1*

Contents

This chapter contains the following topics.

Topic	Page
"Overview"	1-1
"Names"	1-2
"Example Scenarios"	1-3
"Design Principles"	1-4

1.1 Overview

A name-to-object association is called a *name binding*. A name binding is always defined relative to a *naming context*. A naming context is an object that contains a set of name bindings in which each name is unique. Different names can be bound to an object in the same or different contexts at the same time. There is no requirement, however, that all objects must be named.

To *resolve a name* is to determine the object associated with the name in a given context. To *bind a name* is to create a name binding in a given context. A name is always resolved relative to a context — there are no absolute names.

Because a context is like any other object, it can also be bound to a name in a naming context. Binding contexts in other contexts creates a *naming graph* — a directed graph with nodes and labeled edges where the nodes are contexts. A naming graph allows more complex names to reference an object. Given a context in a naming graph, a

sequence of names can reference an object. This sequence of names (called a *compound name*) defines a path in the naming graph to navigate the resolution process. Figure 1-1 shows an example of a naming graph.
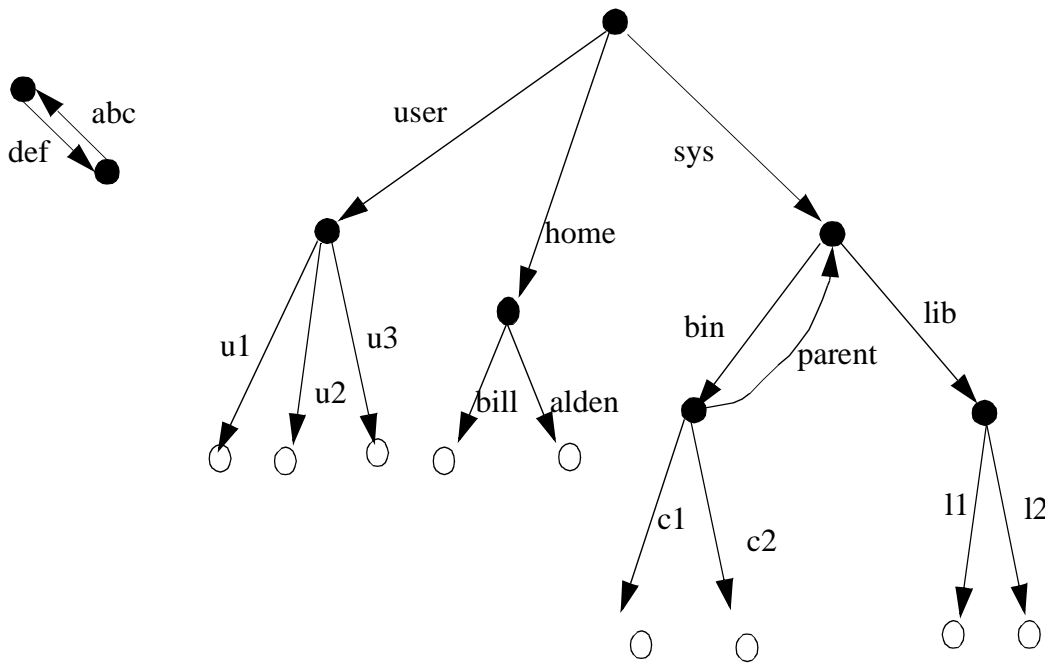
Figure 1-1 A Naming Graph

1.2 Names

Many of the operations defined on a naming context take names as parameters. Names have structure. A name is an ordered sequence of *components*.

A name with a single component is called a *simple name*; a name with multiple components is called a *compound name*. Each component except the last is used to name a context; the last component denotes the bound object. The notation:

```
component1/component2/component3
```

indicates a sequence of components.

Note – The slash (/) characters are simply a notation used here and are not intended to imply that names are sequences of characters separated by slashes.

A name component consists of two attributes: the *id attribute* and the *kind attribute*. Both the **id** attribute and the **kind** attribute are represented as IDL strings.

The **kind** attribute adds descriptive power to names in a syntax-independent way. Examples of the value of the **kind** attribute include *c_source*, *object_code*, *executable*, *postscript*, or *" "*. The naming system does not interpret, assign, or manage these

values in any way. Higher levels of software may make policies about the use and management of these values. This feature addresses the needs of applications that use syntactic naming conventions to distinguish related objects. For example Unix uses suffixes such as **.c** and **.o**. Applications (such as the C compiler) depend on these syntactic convention to make name transformations (for example, to transform **foo.c** to **foo.o**).

A sequence of **id** and **kind** pairs forming a name can be expressed as a single string using the syntax described in Section 2.3, "The BindingIterator Interface," on page 2-10. This allows names to be written down easily or to be presented as a strings in user interfaces. In addition, Section 2.4, "Stringified Names," on page 2-11 describes a way to express a name relative to a particular naming context in URL format. The URL representation provides a human-readable form of an object reference that is named in some naming context.

1.3 Example Scenarios

This section provides two short scenarios that illustrate how the naming service specification can be used by two fairly different kinds of systems -- systems that differ in the kind of implementations used to build the Naming Service and that differ in models of how clients might use the Naming Service with other object services to locate objects.

In one system, the Naming Service is implemented using an underlying enterprise-wide naming server such as DCE CDS. The Naming Service is used to construct large, enterprise-wide naming graphs where NamingContexts model "directories" or "folders" and other names identify "document" or "file" kinds of objects. In other words, the naming service is used as the backbone of an enterprise-wide filing system. In such a system, non-object-based access to the naming service may well be as commonplace as object-based access to the naming service.

The Naming Service provides the principal mechanism through which most clients of an ORB-based system locate objects that they intend to use (make requests of). Given an initial naming context, clients navigate naming contexts retrieving lists of the names bound to that context. In conjunction with properties and security services, clients look for objects with certain "externally visible" characteristics, for example, for objects with recognized names or objects with a certain time-last-modified (all subject to security considerations). All objects used in such a scheme register their externally visible characteristics with other services (a name service, a properties service, and so on).

Conventions are employed in such a scheme that meaningfully partition the name space. For example, individuals are assigned naming contexts for personal use, groups of individuals may be assigned shared naming contexts while other contexts are organized in a public section of the naming graph. Similarly, conventions are used to identify contexts that list the names of services that are available in the system (e.g., that locate a translation or printing service).

In an alternative system, the Naming Service can be used in a more limited role and can have a less sophisticated implementation. In this model, naming contexts represent the types and locations of services that are available in the system and a much shallower naming graph is employed. For example, the Naming Service is used to register the object references of a mail service, an information service, a filing service.

Given a handful of references to "root objects" obtained from the Naming Service, a client uses the Relationship and Query Services to locate objects contained in or managed by the services registered with the Naming Service. In such a system, the Naming Service is used sparingly and instead clients rely on other services such as query services to navigate through large collections of objects. Also, objects in this scheme rarely register "external characteristics" with another service - instead they support the interfaces of Query or Relationship Services.

Of course, nothing precludes the Naming Service presented here from being used to provide both models of use at the same time. These two scenarios demonstrate how this specification is suitable for use in two fairly different kinds of systems with potentially quite different kinds of implementations. The service provides a basic building block on which higher-level services impose the conventions and semantics which determine how frameworks of application and facilities objects locate other objects.

1.4 Design Principles

Several principles have driven the design of the Naming Service:

1. The design imparts no semantics or interpretation of the names themselves; this is up to higher-level software.

2. The design supports distributed, heterogeneous implementation and administration of names and name contexts.

3. Naming service clients need not be aware of the physical site of name servers in a distributed environment, or which server interprets what portion of a compound name, or of the way that servers are implemented.

4. The Naming Service is a fundamental object service, with no dependencies on other interfaces.

5. Name contexts of arbitrary and unknown implementation may be utilized together as nested graphs of nodes that cooperate in resolving names for a client. No "universal" root is needed for a name hierarchy.

6. Existing name and directory services employed in different network computing environments can be transparently encapsulated using name contexts. All of the above features contribute to making this possible.

7. The design does not address namespace administration. It is the responsibility of higher-level software to administer the namespace.

Modules and Interfaces   *2*

Contents

This chapter contains the following topics.

2.1   The CosNaming Module

The **CosNaming** module is a collection of interfaces that together define the Naming Service. This module contains three interfaces:

1. **NamingContext** interface

2. **BindingIterator** interface

3. **NamingContextExt** interface

This section describes these interfaces and their operations in detail.

The **CosNaming** module is shown below.

> **Note – Istring** was a "placeholder for a future IDL internationalized string data type" in the original specification. It is maintained solely for compatibility reasons.

```
// File: CosNaming.idl
#ifndef COSNAMING_IDL_
#define COSNAMING_IDL_

#pragma prefix "omg.org"

module CosNaming {
    typedef string Istring;

    struct NameComponent {
        Istring id;
        Istring kind;
    };
    typedef sequence<NameComponent> Name;

    enum BindingType { nobject, ncontext };

    struct Binding {
        Name            binding_name;
        BindingType     binding_type;
    };

    // Note: In struct Binding, binding_name is incorrectly defined
    // as a Name instead of a NameComponent. This definition is
    // unchanged for compatibility reasons.
    typedef sequence <Binding> BindingList;

    interface BindingIterator;

    interface NamingContext {
        enum NotFoundReason {
            missing_node, not_context, not_object
        };

        exception NotFound {
            NotFoundReason  why;
            Name            rest_of_name;
        };

        exception CannotProceed {
            NamingContext   cxt;
            Name            rest_of_name;
        };

        exception InvalidName{};
```

```
exception AlreadyBound {};

exception NotEmpty{};

void    bind(in Name n, in Object obj)
            raises(
                NotFound, CannotProceed,
                InvalidName, AlreadyBound
            );

void    rebind(in Name n, in Object obj)
            raises(NotFound, CannotProceed, InvalidName);

void    bind_context(in Name n, in NamingContext nc)
            raises(
                NotFound, CannotProceed,
                InvalidName, AlreadyBound
            );

void    rebind_context(in Name n, in NamingContext nc)
            raises(NotFound, CannotProceed, InvalidName);

Object  resolve (in Name n)
        raises(NotFound, CannotProceed, InvalidName);

void    unbind(in Name n)
            raises(NotFound, CannotProceed, InvalidName);

NamingContext    new_context();
NamingContext    bind_new_context(in Name n)
            raises(
                NotFound, AlreadyBound,
                CannotProceed, InvalidName);

void    destroy() raises(NotEmpty);

void    list(
            in unsigned long    how_many,
            out BindingList     bl,
            out BindingIterator bi
        );
};

interface BindingIterator {
    boolean next_one(out Binding b);
    boolean next_n(in unsigned long how_many, out BindingList bl);
        void    destroy();
};

interface NamingContextExt: NamingContext {
```

```
typedef string StringName;
typedef string Address;
typedef string URLString;

StringName    to_string(in Name n) raises(InvalidName);
Name          to_name(in StringName sn)
                  raises(InvalidName);

exception InvalidAddress {};

URLString     to_url(in Address addr, in StringName sn)
                  raises(InvalidAddress, InvalidName);

Object        resolve_str(in StringName n)
                  raises(
                      NotFound, CannotProceed,
                      InvalidName
                  );
    };
};
#endif // _COSNAMING_IDL_
```

2.1.1 Resolution of Compound Names

In this specification operations that are performed on compound names recursively perform the *equivalent* of a **resolve** operation on all but the last component of a name before performing the operation on the final name component. The general form is defined as follows:

ctx->op(<c1; c2; ...; cn>) equiv

ctx->resolve(<c1>)->resolve(<c2; cn-1>)->op(<cn>)

where ctx is a naming context, <c1; ...; cn> a compound name, and op a naming context operation.

Note – The intermediate components, <c1: ...; cn> of the compound name must have been bound using **bind_context** or **rebind_context** to take part in the resolve.

2.2 NamingContext Interface

The following sections describe the naming context data types and interface in detail.

2.2.1 Structures

2.2.1.1 NameComponent

```
struct NameComponent {
```

```
        Istring Id;
        Istring kind;
    };
```

A name component consists of two attributes: the identifier attribute (**id**) and the kind attribute (**kind**).

Both of these attributes are arbitrary-length strings of ISO Latin-1 characters, excluding the ASCII **NUL** character.

When comparing two **NameComponents** for equality both the **id** and the **kind** field must match in order for two **NameComponents** to be considered identical. This applies for zero-length (empty) fields as well.

An implementation may place limitations on the characters that may be contained in a name component, as well as the length of a name component. For example, an implementation may disallow certain characters, may not accept the empty string as a legal name component, or may limit name components to some maximum length.

2.2.1.2 *Name*

A name is a sequence of **NameComponent**s. The empty sequence is not a legal name. An implementation may limit the length of the sequence to some maximum. When comparing **Name**s for equality, each **NameComponent** in the first name must match the corresponding **NameComponent** in the second **Name** for the names to be considered identical.

2.2.1.3 *Binding*

```
enum BindingType { nobject, ncontext };
struct Binding {
    Name binding_name;
    BindingType binding_type;
};
typedef sequence<Binding> BindingList;
```

This types are used by the **NamingContext::list**, **BindingIterator::next_n** and **BindingIterator::next_one** operations. A **Binding** contains a **Name** in the member **binding_name**, together with the **BindingType** of that **Name** in the member **binding_type.**

Note – The **binding_name** member is incorrectly typed as a **Name** instead of a **NameComponent.** For compatibility with the original **CosNaming** specification this incorrect definition has been retained. The **binding_name** is used as a **NameComponent** and will always be a **Name** with length of 1.

The value of **_binding_type_** is **_ncontext_** if a **Name** denotes a binding created with one of the following operations:

• **bind_context**

- **rebind_context**

- **bind_new_context**

For bindings created with any other operation, the value of **BindingType** is **nobject**.

2.2.2 Exceptions

The Naming Service exceptions are defined below.

2.2.2.1 NotFound

```
exception NotFound {
    NotFoundReason  why;
    Name rest_of_name;
};
```

This exception is raised by operations when a component of a name does not identify a binding or the type of the binding is incorrect for the operation being performed. The **why** member explains the reason for the exception and the **rest_of_name** member contains the remainder of the non-working name:

- **missing_node**

  The first name component in **rest_of_name** denotes a binding that is not bound under that name within its parent context.

- **not_context**

  The first name component in **rest_of_name** denotes a binding with a type of **nobject** when the type **ncontext** was required.

- **not_object**

  The first name component in **rest_of_name** denotes a binding with a type of **ncontext** when the type **nobject** was required.

2.2.2.2 CannotProceed

```
exception CannotProceed {
    NamingContext    cxt;
    Name rest_of_name;
};
```

This exception is raised when an implementation has given up for some reason. The client, however, may be able to continue the operation at the returned naming context.

The **cxt** member contains the context that the operation may be able to retry from.

The **rest_of_name** member contains the remainder of the non-working name.

2.2.2.3 *InvalidName*

exception InvalidName {};

This exception is raised if a **Name** is invalid. A name of length zero is invalid (containing no name components). Implementations may place further limitations on what constitutes a legal name and raise this exception to indicate a violation.

2.2.2.4 *AlreadyBound*

exception AlreadyBound {};

Indicates an object is already bound to the specified name. Only one object can be bound to a particular **Name** in a context.

2.2.2.5 *NotEmpty*

exception NotEmpty {};

This exception is raised by **destroy** if the **NamingContext** contains bindings. A **NamingContext** must be empty to be destroyed.

2.2.3 *Binding Objects*

The binding operations name an object in a naming context. Once an object is bound, it can be found with the **resolve** operation. The Naming Service supports four operations to create bindings: **bind**, **rebind**, **bind_context** and **rebind_context**. **bind_new_context** also creates a binding, see Section 2.2.6, "Creating Naming Contexts," on page 2-9.

```
void bind(in Name n, in Object obj)
        raises(NotFound, CannotProceed, InvalidName, AlreadyBound);
void rebind(in Name n, in Object obj)
        raises(NotFound, CannotProceed, InvalidName);
void bind_context(in Name n, in NamingContext nc)
        raises(NotFound, CannotProceed, InvalidName, AlreadyBound);
void rebind_context(in Name n, in NamingContext nc)
        raises(NotFound, CannotProceed, InvalidName);
```

2.2.3.1 *bind*

Creates an **nobject** binding in the naming context.

2.2.3.2 *rebind*

Creates an **nobject** binding in the naming context even if the name is already bound in the context.

If already bound, the previous binding must be of type **nobject**; otherwise, a NotFound exception with a **why** reason of **not_object** is raised.

2.2.3.3 *bind_context*

Creates an **ncontext** binding in the parent naming context. Attempts to bind a nil context raise a BAD_PARAM exception.

2.2.3.4 *rebind_context*

Creates an **ncontext** binding in the naming context even if the name is already bound in the context.

If already bound, the previous binding must be of type **ncontext**; otherwise, a NotFound exception with a **why** reason of **not_context** will be raised.

2.2.3.5 *Usage*

If a binding with the specified name already exists, **bind** and **bind_context** raise an AlreadyBound exception.

If an implementation places limits on the number of bindings within a context, **bind** and **bind_context** raise the IMP_LIMIT system exception if the new binding cannot be created.

Naming contexts bound using **bind_context** and **rebind_context** participate in name resolution when compound names are passed to be resolved; naming contexts bound with **bind** and **rebind** do not.

Use of **rebind_context** may leave a potential orphaned context (one that is unreachable within an instance of the Name Service). Policies and administration tools regarding potential orphan contexts are implementation-specific.

If **rebind** or **rebind_context** raise a **NotFound** exception because an already existing binding is of the wrong type, the **rest_of_name** member of the exception has a sequence length of 1.

2.2.4 *Resolving Names*

The **resolve** operation is the process of retrieving an object bound to a name in a given context. The given name must exactly match the bound name. The naming service does not return the type of the object. Clients are responsible for "narrowing" the object to the appropriate type. That is, clients typically cast the returned object from **Object** to a more specialized interface. The IDL definition of the resolve operation is:

Object resolve (in Name n)
        raises (NotFound, CannotProceed, InvalidName);

Names can have multiple components; therefore, name resolution can traverse multiple contexts. These contexts can be federated between different Naming Service instances.

2.2.5 Unbinding Names

The **unbind** operation removes a name binding from a context. The definition of the `unbind` operation is:

void unbind(in Name n)
raises (NotFound, CannotProceed, InvalidName);

2.2.6 Creating Naming Contexts

The Naming Service supports two operations to create new contexts: **new_context** and **bind_new_context**.

NamingContext new_context();
NamingContext bind_new_context(in Name n)
raises(NotFound, AlreadyBound, CannotProceed, InvalidName);

2.2.6.1 new_context

This operation returns a new naming context. The new context is not bound to any name.

2.2.6.2 bind_new_context

This operation creates a new context and creates an **ncontext** binding for it using the name supplied as an argument.

2.2.6.3 Usage

If an implementation places limits on the number of naming contexts, both **new_context** and **bind_new_context** can raise the IMP_LIMIT system exception if the context cannot be created. **bind_new_context** can also raise IMP_LIMIT if the bind would cause an implementation limit on the number of bindings in a context to be exceeded.

2.2.7 Deleting Contexts

The **destroy** operation deletes a naming context.

void destroy()
raises(NotEmpty);

This operation destroys its naming context. If there are bindings denoting the destroyed context, these bindings are *not* removed. If the naming context contains bindings, the operation raises NotEmpty.

2.2.8 Listing a Naming Context

The **list** operation allows a client to iterate through a set of bindings in a naming context.

void list (in unsigned long how_many,
        out BindingList bl, out BindingIterator bi);
};

list returns the bindings contained in a context in the parameter **bl**. The **bl** parameter is a sequence where each element is a **Binding** containing a **Name** of length 1 representing a single **NameComponent**.

The **how_many** parameter determines the maximum number of bindings to return in the parameter **bl**, with any remaining bindings to be accessed through the returned **BindingIterator bi.**

- A non-zero value of **how_many** guarantees that **bl** contains at most **how_many** elements. The implementation is free to return fewer than the number of bindings requested by **how_many**. However, for a non-zero value of **how_many**, it may not return a **bl** sequence with zero elements unless the context contains no bindings.

- If **how_many** is set to zero, the client is requesting to use only the **BindingIterator bi** to access the bindings and **list** returns a zero length sequence in **bl**.

- The parameter **bi** returns a reference to an iterator object.
  - If the **bi** parameter returns a non-nil reference, this indicates that the call to **list** may not have returned all of the bindings in the context and that the remaining bindings (if any) must be retrieved using the iterator. This applies for all values of **how_many**.
  - If the **bi** parameter returns a nil reference, this indicates that the **bl** parameter contains all of the bindings in the context. This applies for all values of **how_many**.

2.3 The BindingIterator Interface

The **BindingIterator** interface allows a client to iterate through the bindings using the **next_one** or **next_n** operations:

If a context is modified in between calls to **list**, **next_one**, or **next_n**, the behavior of further calls to **next_one** or **next_n** is implementation-dependent.

interface BindingIterator {
    boolean next_one(out Binding b);
    boolean next_n(in unsigned long how_many,
                  out BindingList bl);
    void destroy();
};

2.3.1 next_one

The **next_one** operation returns the next binding. It returns true if it is returning a binding, false if there are no more bindings to retrieve. If **next_one** returns false, the returned **Binding** is indeterminate.

Further calls to **next_one** after it has returned false have undefined behavior.

2.3.2 next_n

next_n returns, in the parameter **bl**, bindings not yet retrieved with **list** or previous calls to **next_n** or **next_one**. It returns true if **bl** is a non-zero length sequence; it returns false if there are no more bindings and **bl** is a zero-length sequence.

The **how_many** parameter determines the maximum number of bindings to return in the parameter **bl**:

- A non-zero value of **how_many** guarantees that **bl** contains at most **how_many** elements. The implementation is free to return fewer than the number of bindings requested by **how_many**. However, it may not return a **bl** sequence with zero elements unless there are no bindings to retrieve.

- A zero value of **how_many** is illegal and raises a BAD_PARAM system exception.

next_n returns false with a **bl** parameter of length zero once all bindings have been retrieved. Further calls to **next_**n after it has returned a zero-length sequence have undefined behavior.

2.3.3 destroy

The **destroy** operation destroys its iterator. If a client invokes any operation on an iterator after calling **destroy**, the operation raises OBJECT_NOT_EXIST.

2.3.4 Garbage Collection of Iterators

Clients that create iterators but never call **destroy** can cause an implementation to permanently run out of resources. To protect itself against this scenario, an implementation is free to destroy an iterator object at any time without warning, using whatever algorithm it considers appropriate to choose iterators for destruction. In order to be robust in the presence of garbage collection, clients should be written to expect OBJECT_NOT_EXIST from calls to an iterator and handle this exception gracefully.

2.4 Stringified Names

Names are sequences of name components. This representation makes it difficult for applications to conveniently deal with names for I/O purposes, human or otherwise. This specification defines a syntax for stringified names and provides operations to

convert a name in stringified form to its equivalent sequence form and vice-versa (see Section 2.5.4, "Converting between CosNames, Stringified Names, and URLs," on page 2-16).

A stringified name represents one and only one **CosNaming::Name**. If two names are equal, their stringified representations are equal (and vice-versa).

The stringified name representation reserves use of the characters '/', '.', and '\'. The forward slash '/' is a name component separator; the dot '.' separates **id** and **kind** fields. The backslash '\' is an escape character (see Section 2.4.2, "Escape Mechanism," on page 2-13).

2.4.1 Basic Representation of Stringified Names

A stringified name consists of the name components of a name separated by a '/' character. For example, a name consisting of the components "a," "b," and "c" (in that order) is represented as

a/b/c

Stringified names use the '.' character to separate **id** and **kind** fields in the stringified representation. For example, the stringified name

a.b/c.d/.

represents the **CosNaming::Name**:

Index	id	kind
0	**a**	**b**
1	**c**	**d**
2	<empty>	<empty>

The single '.' character is the only representation of a name component with empty **id** and **kind** fields.

If a name component in a stringified name does not contain a '.' character, the entire component is interpreted as the **id** field, and the **kind** field is empty. For example:

a/./c.d/.e

corresponds to the **CosNaming::Name**:

Index	id	kind
0	**a**	**<empty>**
1	**<empty>**	**<empty>**
2	**c**	**d**
3	<empty>	e

If a name component has a non-empty **id** field and an empty **kind** field, the stringified representation consists only of the **id** field. A trailing '.' character is not permitted.

2.4.2 Escape Mechanism

The backslash '\' character escapes the reserved meaning of '/', '.', and '\' in a stringified name. The meaning of any other character following a '\' is reserved for future use.

2.4.2.1 NameComponent Separators

If a name component contains a '/' slash character, the stringified representation uses the '\' character as an escape. For example, the stringified name

a/x\/y\/z/b

represents the name consisting of the name components "a," "x/y/z," and "b."

2.4.2.2 Id and kind Fields

The backslash escape mechanism is also used for '.', so **id** and **kind** fields can contain a literal '.'. To illustrate, the stringified name

a\.b.c\.d/e.f

represents the **CosNaming::Name**:

Index	id	kind
0	**a.b**	**c.d**
1	**e**	**f**

2.4.2.3 The Escape Character

The escape character '\' must be escaped if it appears in a name component. For example, the stringified name:

a/b\\/c

represents the name consisting of the components "**a**," "**b**," and "**c**."

2.5 URL Schemes

This section describes the Uniform Resource Locator (URL) schemes available to represent a CORBA object and a CORBA object bound in a **NamingContext**.

2.5.1 IOR

The string form of an IOR (IOR:<hex_octets>) is a valid URL. The scheme name is IOR and the text after the ':' is defined in the CORBA 2.3 specification, Section 13.6.6. The IOR URL is robust and insulates the client from the encapsulated transport information and object key used to reference the object. This URL format is independent of Naming Service.

2.5.2 corbaloc

It is difficult for humans to exchange IORs through non-electronic means because of their length and the text encoding of binary information. The **corbaloc** URL scheme provides URLs that are familiar to people and similar to **ftp** or **http** URLs.

The **corbaloc** URL is described in the CORBA 2.3.1 Specification, Section 13.6.6. This URL format is independent of the Naming Service.

2.5.3 corbaname

A **corbaname** URL is similar to a **corbaloc** URL. However a **corbaname** URL also contains a stringified name that identifies a binding in a naming context.

Note – Use of '#' as a separator character as used in this document is under review by the INS FTF.

2.5.3.1 corbaname Examples

corbaname://555xyz.com/dev/NContext1#a/b/c

This example denotes a naming context that can be contacted in the same manner as a **corbaloc** URL at 555xyz.com with a key of "dev/NContext1." The "#" character denotes the start of the stringified name "**a/b/c**." This name is resolved against the context to yield the final object.

corbaname://555xyz.com#a/b/c

When an object key is not specified, as in the above example, the default key of "NameService" is used to contact the naming context.

corbaname:///#x/y/z

This URL refers to a software agent supporting IIOP version 1.0 on the local host and default port. The naming context associated with the object key **NameService** is used to resolve the name **x/y/z**, which yields the object reference denoted by the URL.

corbaname:///

This URL represents the naming context returned by the agent running on the local host at the default port. It is equivalent to **corbaloc:///NameService**.

corbaname:atm:00033...#a/b/c

corbaname://55xyz.com,atm:00033.../dev/NCtext#a/b/c

These last URLs illustrate support of multiple protocols as allowed by **corbaloc** URLs. **atm**: is an example only and is not a defined URL protocol at this time.

Note – Unlike stringified names, **corbaname**s cannot be compared directly for equality as the address specification can differ for **corbaname** URLs with the same meaning.

2.5.3.2 *corbaname Syntax*

The full **corbaname** BNF is:

<corbaname>	= **"corbaname:"<corbaloc_obj>["#"<string_name>]**
<corbaloc_obj>	= **<obj_addr_list> ["/"<key_string>]**
<obj_addr_list>	= as defined in a **corbaloc** URL
<key_string>	= as defined in a **corbaloc** URL
<string_name>	= **stringified Name	empty_string**

Where:

corbaloc_obj: portion of a **corbaname** URL that identifies the naming context. The syntax is identical to its use in a **corbaloc** URL.

obj_addr_list: as defined in a **corbaloc** URL

key_string: as defined in a **corbaloc** URL.

string_name: a stringified **Name** with URL escapes as defined below.

2.5.3.3 *corbaname Character Escapes*

corbaname URLs use the escape mechanism described in the Internet Engineering Task Force (IETF) RFC 2396. These escape rules insure that URLs can be transferred via a variety of transports without undergoing changes. The character escape rules for the stringified name portion of a **corbaname** are:

US-ASCII alphanumeric characters are not escaped. Characters outside this range are escaped, except for the following:

";" | "/" | ":" | "?"| ":" | "@" | "&" | "=" | "+" | "$" |

"," | "-" | "_" | "." | "!" | "~" | "*" | "'" | "(" | ")"

2.5.3.4 *corbaname Escape Mechanism*

The percent '%' character is used as an escape. If a character that requires escaping is present in a name component, it is encoded as two hexadecimal digits following a "%" character to represent the octet. (The first hexadecimal character represents the high-order nibble of the octet, the second hexadecimal character represents the low-order nibble.) If a '%' is not followed by two hex digits, the stringified name is syntactically invalid.

2.5.3.5 Examples

Stringified Name	After URL Escapes	Comment
a.b/c.d	a.b/c.d	URL form identical
\<a\>.b/c.d	%3ca%3e.b/c.d	Escaped "\<" and "\>"
a.b/ c.d	a.b/%20%20c.d	Escaped two " " spaces
a%b/c%d	a%25b/c%25d	Escaped two "%" percents
a\\\\b/c.d	a%5c%5c/c.d	Escaped "\\" character, which is already escaped in the stringified name.

2.5.3.6 corbaname Resolution

corbaname resolution can be implemented as a simple extension to **corbaloc** URL processing. Given a **corbaname**:

> **corbaname:\<corbaloc_obj\>["#" \<string_name\>]**

The **corbaname** is resolved by:

1. First constructing a **corbaloc** URL of the form: **corbaloc:\<corabloc_obj\>**.

   If the **\<corbaloc_obj\>** does not contain a key string, a default key of "NameService" is used.

2. This is converted to a naming context object reference with **CORBA::ORB::string_to_object**.

3. The **\<string_name\>** is converted to a **CosNaming::Name**.

4. The resulting name is passed to a **resolve** operation on the naming context.

5. The object reference returned by the **resolve** is the result.

Implementations are not required to use the method described and may make optimizations appropriate to their environment.

2.5.4 Converting between CosNames, Stringified Names, and URLs

The **NamingContextExt** interface, derived from **NamingContext**, provides the operations required to use URLs and stringified names.

```
module CosNaming {
    // ...
    interface NamingContextExt: NamingContext {
        typedef string StringName;
        typedef string Address;
        typedef string URLString;

        StringName    to_string(in Name n) raises(InvalidName);
```

```
Name          to_name(in StringName sn)
    raises(InvalidName);

exception InvalidAddress {};

URLString     to_url(in Address addrkey, in StringName sn)
    raises(InvalidAddress, InvalidName);

Object        resolve_str(in StringName n)
    raises(
        NotFound, CannotProceed,
        InvalidName
    );
    };
};
```

2.5.4.1 *to_string*

This operation accepts a **Name** and returns a stringified name. If the **Name** is invalid, an InvalidName exception is raised.

2.5.4.2 *to_name*

This operation accepts a stringified name and returns a **Name**. If the stringified name is syntactically malformed or violates an implementation limit, an InvalidName exception is raised.

2.5.4.3 *resolve_str*

This is a convenience operation that performs a resolve in the same manner as **NamingContext**::**resolve**. It accepts a stringified name as an argument instead of a **Name**.

2.5.4.4 *to_url*

This operation takes a corbaloc URL **<address>** and **<key_string>** component such as

- **//myhost.555xyz.com**

- **//myhost.555xyz.com/a/b/c**

- **atm:00002112...,//myhost.xyz.com/a/b/c**

for the first parameter, and a stringified name for the second. It then performs any escapes necessary on the parameters and returns a fully formed URL string. An exception is raised if either the corbaloc address and key parameter or name parameter are malformed.

It is legal for the address and/or stringified_name to be empty. If the address is empty, it means the local host and the iiop protocol.

2.5.4.5 *URL to Object Reference*

Conversions from URLs to objects are handled by **CORBA::ORB::string_to_object** as described in the CORBA 2.3 Specification, Section 13.6.6.

2.6 *Initial Reference to a NamingContextExt*

An initial reference to an instance of this interface can be obtained by calling **resolve_initial_references** with an **ObjectID** of **NameService**.

OMG IDL *A*

```
// File: CosNaming.idl
#ifndef _COSNAMING_IDL_
#define _COSNAMING_IDL_

#pragma prefix "omg.org"

module CosNaming {

    typedef string Istring;

    struct NameComponent {
        Istring id;
        Istring kind;
    };
    typedef sequence<NameComponent> Name;

    enum BindingType { nobject, ncontext };

    struct Binding {
        Name           binding_name;
        BindingType    binding_type;
    };

    // Note: In struct Binding, binding_name is incorrectly defined
    // as a Name instead of a NameComponent. This definition is
    // unchanged for compatibility reasons.
    typedef sequence <Binding> BindingList;

    interface BindingIterator;

    interface NamingContext {

        enum NotFoundReason {
```

```
        missing_node, not_context, not_object
};

exception NotFound {
    NotFoundReason  why;
    Name            rest_of_name;
};

exception CannotProceed {
    NamingContext   cxt;
    Name            rest_of_name;
};

exception InvalidName{};

exception AlreadyBound {};

exception NotEmpty{};

void    bind(in Name n, in Object obj)
    raises(
        NotFound, CannotProceed,
        InvalidName, AlreadyBound
    );

void    rebind(in Name n, in Object obj)
    raises(NotFound, CannotProceed, InvalidName);

void    bind_context(in Name n, in NamingContext nc)
    raises(NotFound, CannotProceed, InvalidName, AlreadyBound);

void    rebind_context(in Name n, in NamingContext nc)
    raises(NotFound, CannotProceed, InvalidName);

Object  resolve (in Name n)
    raises(NotFound, CannotProceed, InvalidName);

void    unbind(in Name n)
    raises(NotFound, CannotProceed, InvalidName);

NamingContext   new_context();
NamingContext   bind_new_context(in Name n)
    raises(
        NotFound, AlreadyBound,
        CannotProceed, InvalidName
    );

void    destroy() raises(NotEmpty);

void    list(
    in unsigned long        how_many,
```

```
                out BindingList        bl,
                out BindingIterator    bi
        );
};

interface BindingIterator {
    boolean next_one(out Binding b);
    boolean next_n(in unsigned long how_many, out BindingList bl);
    void    destroy();
};

interface NamingContextExt: NamingContext {
    typedef string StringName;
    typedef string Address;
    typedef string URLString;

    StringName      to_string(in Name n) raises(InvalidName);
    Name            to_name(in StringName sn)
            raises(InvalidName);

    exception InvalidAddress {};

    URLString       to_url(in Address addr, in StringName sn)
            raises(InvalidAddress, InvalidName);

    Object          resolve_str(in StringName n)
            raises(
                NotFound, CannotProceed,
                InvalidName, AlreadyBound
    );
    };
};
#endif // _COSNAMING_IDL_
```

Requirements   *B*

B.1   Conformance Requirements

B.1.1   Optional Interfaces

There are no optional interfaces in this specification. A compliant implementation must implement all of the functionality and interfaces described.

B.1.2   Documentation Requirements

A compliant implementation must document all of the following:

- any limitations to the character values or character sequences that may be used in a name component

- any limitations to the length (including minimum or maximum) of a name component

- any limitations to number of name components in a name

- any limitations to the maximum number of bindings in a context

- any limitations to the total number of bindings (implementation-wide)

- any limitations to the maximum number of contexts

- the means provided to deal with orphaned contexts and bindings

- Any policy for dealing with potentially orphaned naming contexts. Orphaned contexts are contexts that are not bound in any other context within a naming server.

- Any policy for destroying binding iterators that are considered to be no longer in use.

- Under what circumstances, if any, a CannotProceed exception is raised.