Spectra ORB
C Edition

Lightweight Naming Service
Guide

PRISMTECH

Spectra ORB
C Edition

Lightweight Naming Service Guide

Part Number: EORBC-NAMLWG

Doc Issue 29, 4 June 2013

Copyright Notice

Contents

Table of Contents

Preface

About the Lightweight Naming Service Guide

The *Lightweight Naming Service Guide* explains how to use the Spectra ORB Lightweight Naming Service C Edition product.

Intended Audience

The *Lightweight Naming Service Guide* is intended to be used by developers who wish to integrate the Spectra ORB Lightweight Naming Service C Edition into products which comply with OMG standards for object services. Readers who use this guide should have a good understanding of the relevant programming languages (for example C, IDL) and of the relevant underlying technologies (such as CORBA).

Organisation

The *Lightweight Naming Service Guide* provides:

- a high level description and list of main features
- an explanation of the OMG Naming Service architecture and concepts
- information about how to configure, deploy and run the Spectra ORB Lightweight Naming Service C Edition
- information about how to create applications
- other information about the service.

Conventions

The conventions listed below are used to guide and assist the reader in understanding the *Lightweight Naming Service Guide*.


Item of special significance or where caution needs to be taken.


Item contains helpful hint or special information.

WIN Information applies to Windows (*e.g.* XP, Vista, Windows 7) only.

UNIX Information applies to Unix-based systems (*e.g.* Solaris) only.

C C language specific.

C++ C++ language specific.

Java Java language specific.

Hypertext links are shown as *blue italic underlined.*

On-Line (PDF) versions of this document: Items shown as cross references to other parts of the document, *e.g. Contacts* on page viii, behave as hypertext links: readers can jump to that section of the document by clicking on the cross reference.

```
%  Commands or input which the user enters on the
   command line of their computer terminal
```

`Courier` fonts indicate programming code and file names.

Extended code fragments are shown in shaded boxes:

```
NameComponent newName[] = new NameComponent[1];

// set id field to “example” and kind field to an empty string
newName[0] = new NameComponent (“example”, ““);
```

Italics and ***Italic Bold*** indicate new terms, or emphasise an item.

Arial Bold indicate users related actions, *e.g.* **File > Save** from a menu.

Step 1: Indicates that this item is a step or stage of completing a task by a user.

Contacts

PrismTech can be reached at the following contact points for information and technical support.

USA Corporate Headquarters

PrismTech Corporation
400 TradeCenter
Suite 5900
Woburn, MA
01801
USA

Tel: +1 781 569 5819

European Head Office

PrismTech Limited
PrismTech House
5th Avenue Business Park
Gateshead
NE11 0NG
UK

Tel: +44 (0)191 497 9900
Fax: +44 (0)191 497 9901

Web: *http://www.prismtech.com*
Technical questions: *crc@prismtech.com* (Customer Response Center)
Sales enquiries: *sales@prismtech.com*

Introduction

Description

The Spectra ORB Lightweight Naming Service C Edition provides a straightforward way of finding and using objects by associating meaningful, human-understandable names to those objects. The Spectra ORB Lightweight Naming Service is used like the white pages of a telephone directory to find an object and obtain its object reference, without the need to resort to complex programming or proprietary ORB mechanisms.

The Spectra ORB Lightweight Naming Service is compliant with the OMG's Lightweight Naming Service Specification and satisfies the Naming Service requirements of the Software Communications Architecture (SCA).

The OMG's Lightweight Naming Service Specification is a sub-set of the full OMG specification and is intended to provide a service which has a minimal footprint and uses a minimum of resources. This service is designed for use in restricted deployment environments, such as in embedded systems and in the Software Defined Radio (SDR) domain.

This implementation has been designed and implemented to be as small and as proficient as possible: it is designed for use in highly demanding environments where footprint size and resources are limited.

The Spectra ORB Lightweight Naming Service C Edition does ***not*** support the `kind` property of the `CosNaming::NameComponent`. The Service will ignore this property when binding or resolving names. Developers should ensure that they use unique `id` values creating name components.

OMG Standard Features

The Spectra ORB Lightweight Naming Service provides the following OMG specified features:

- enable objects to be located by using human-intelligible names by binding names to objects
- add and remove name bindings
- change the object which a name is bound to
- group names in logical hierarchies

The Spectra ORB Lightweight Naming

SERVICE

CHAPTER

1 Basic Concepts

This section describes the basic concepts and architecture of the OMG's Lightweight Naming Service as defined in the OMG's Lightweight Services Specification.

The OMG's Lightweight Naming Service is a sub-set of the full OMG Naming Service Specification. The Lightweight Specification omits features which are considered not to be essential in a restricted environment. The purpose of this is to reduce footprint size and enable the service to be deployed on memory or storage constrained platforms. The features provided by the Lightweight Naming Service only include those defined as part of the `NamingContext` interface. The features of this interface are described under Section 2, *Specific Features*, on page 11.

1.1 OMG Standard Features

The Lightweight Naming Service has the ability to:

- give meaningful names to objects (*name bindings*)
- allow objects, which have been bound to names, to be easily found (*resolve*)
- organise names in logical hierarchies (*naming contexts*)

1.1.1 Names

The Naming Service associates meaningful names with objects. These names can be used to retrieve or reference the object, or in other words, obtain the object's IOR. An association between a name and an object is known as a *name binding*.

A *Naming Service name* is an object in its own right. A name contains a sequence of one or more *name components*. A name component has two string-type attributes, *id* and *kind*. The id and kind attributes identify the name.

A name contains a sequence of one or more name components. A name with a single name component is called a *simple name*. A name which contains a sequence of two or more name components is called a *compound name*. Compound names are used to access name bindings when they are in a hierarchy of *naming contexts*, described below.

A name binding is held in or otherwise associated with a *naming context*. A name binding cannot exist outside of a naming context. Names are bound to naming contexts, as well as to objects. It is possible to have orphaned contexts if the name binding is removed without keeping a reference to the context being unbound.

An object can be bound to one or more names, but a name can only be bound to one object. If an object is bound to more than one name, then any of those names can be used to locate the object.

Resolving is the process of locating an object or naming context by using its name.

1.1.2 Naming Contexts

A naming context is an object which contains name bindings. Each name in a given naming context must be unique, in other words, the combination of a name's name component *id* and *kind* values must be unique. The name can be used in other naming contexts.

Naming hierarchies can be created by binding a naming context to another naming context. A simple naming context hierarchy is shown in *Figure 1*. Names in a naming context can refer to other naming contexts as well as to objects. A hierarchy of naming contexts is called a *naming graph*.

The top level of the naming graph is called the *root context*. The root context is also the default context, that is, it does not need to be explicitly created whereas all other contexts, *child contexts* of the root context, must be explicitly created.

Objects, and contexts themselves, are referenced by following the hierarchy of naming contexts, starting from the root context and ending with the desired object or naming context. Technically, the object or context can be referenced in one of two ways:

1. Obtain an object reference to the context which contains the required object or context. The object's name can then be used to obtain the object's IOR.
2. Construct a compound name which contains a sequence of name components, where each name component identifies each successive naming context in the naming graph and where the last name component identifies the required object or context.

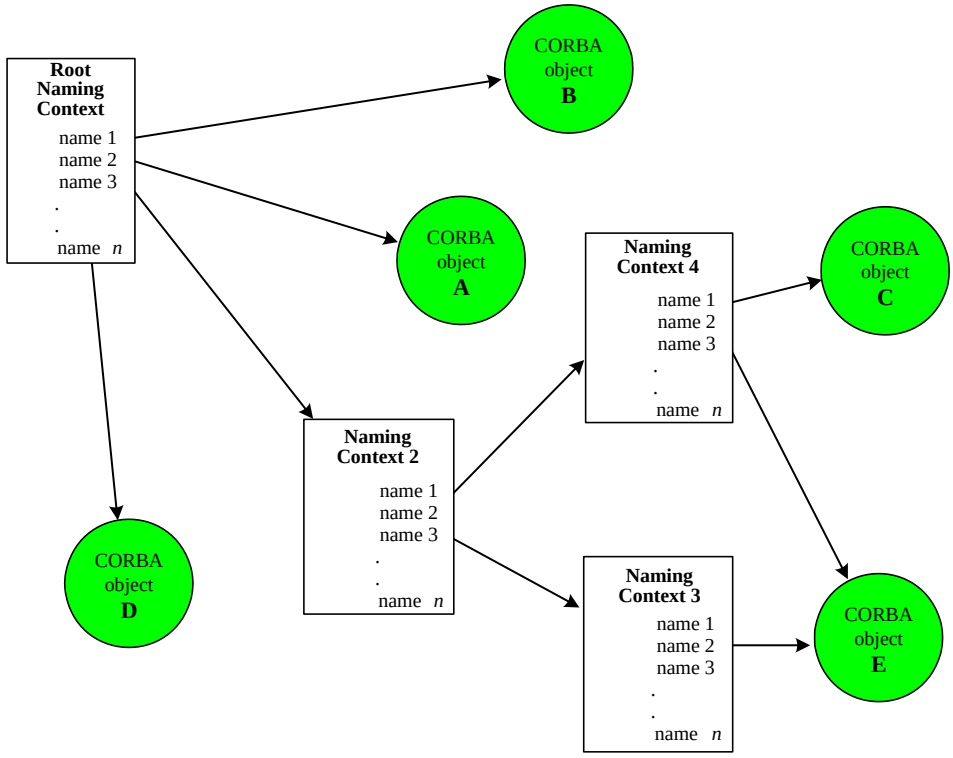

Figure 1 Example Naming Graph

For example, in *Figure 1* objects *A*, *B* and *NamingContext2* are bound directly to the root context: they can be referenced using a simple name (containing the name component which identifies the objects themselves). Objects *C* and *E* are bound to child contexts lower down the hierarchy: in order to access objects *C* and *E* from the root context, the name component for each successive naming context must be provided as a compound name. For example, the compound name for referencing object C from the root context will contain a sequence which looks like this in pseudo-code:

```
name[0] = NamingContext2.nameComponent
name[1] = NamingContext4.nameComponent
name[2] = C.nameComponent
```

The root context is always implicit in a compound name; a special operation, `resolve_initial_references`, is performed once to obtain the root context, and all subsequent `resolve` operations depend on that.

CHAPTER

2 Specific Features

The Spectra ORB Lightweight Naming Service's features are described below. As mentioned previously, the service conforms to the OMG's Lightweight Services Specification and is a subset of the full OMG Naming Service Specification.

The Spectra ORB Lightweight Naming Service supports a subset of the `NamingContext` *interface.*

The interfaces, datatypes, methods and exceptions supported by the Spectra ORB Lightweight Naming Service are listed below.

Interfaces and Datatypes

- `Name` (datatype)
- `NameComponent` (datatype)
- `NamingContext` (interface)

NamingContext Methods

- `bind()`
- `rebind()`
- `resolve()`
- `unbind()`
- `bind_new_context()`
- `destroy()`

NamingContext Exceptions

- `NotFoundReason`
- `NotFound`
- `CannotProceed`
- `InvalidName`
- `AlreadyBound`
- `NotEmpty`

The *`BindingIterator`* or *`NamingContextExt`* interfaces are not supported by the Lightweight Naming Service.

2.1 Naming Context Creation and Destruction

The `NamingContext` interface provides a `NamingContext` creation operation and a destroy operation, defined in IDL as:

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

void destroy () raises (NotEmpty);
```

The `bind_new_context` operation creates a new Naming Context and binds it using the supplied name.

The `destroy` operation requests the destruction of a `NamingContext` object. The naming context must be empty. After `destroy` is invoked, no further operations can be invoked on the object reference of the naming context object.

Bindings to a destroyed context are not removed. To do so would require a context to know about all of its parents as well as its children. An attempt to resolve a binding to a destroyed context will throw the `CORBA.INV_OBJREF` exception. Accordingly, bindings to a naming context should be removed before it is destroyed.

When a hierarchical name is used to create a new context, all the contexts that constitute the path to the new context must already exist or the *`NotFound`* exception will be raised

2.2 Object Binding and Unbinding Operations

The `NamingContext` interface provides the following object binding and unbinding operations, defined in IDL as:

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

The bind operations allow binding to occur between a name and either a generic CORBA object or a Naming Context. In order to bind a CORBA object, the name to bind against must be correctly constructed. Given a name with *n* components, the first *n - 1* components must resolve to a bound `NamingContext`.

The `rebind` operation is identical to the `bind` operation except that the `AlreadyBound` exception is not thrown; an existing binding with the same name is replaced by the new binding.

2.3 Accessing Objects and Naming Contexts

The resolve operation is used to obtain the object references of naming contexts and named objects, defined in IDL as:

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

The `resolve` operation takes a name and returns the object, if any, bound to that name.

CHAPTER

3 Using the Service

*This section describes the specific procedures and requirements for creating and running CORBA-based applications with the Spectra ORB Lightweight Naming Service C Edition. (Please note that this section is **not** intended as a tutorial of how to write CORBA-based applications with the Naming Service.)*

3.1 Running the Service

The Naming Service server instances can be run from the command line or by embedding them into application or module code. Instructions for running the service using these methods is described in the following sections, *Embedding the Service*, below, and *Running from the Command Line* on page 18.

3.1.1 Embedding the Service

The following basic tasks must be performed in order to embed a Naming Service server instance into an executable or code module:

Step 1: Include the following *`#include`* statements in your code:

```
#include "CosNaming.h"
#include "eOrbC/EORB/NamingService.h"
```

i Ensure your build system has *`$(EORBHOME)/include/services/lw`* on its include path.

Step 2: Configure the service instance by setting the property fields in the naming service's configuration structure: these properties are used to determine specific aspects of your instance's behaviour and are described later.

Step 3: Initialise a Naming Service instance by using the Naming Service's *`init()`* method:

```
root_ctx = EORB_NamingService_init (orb, poa, &config, &env)
```

where:

`orb` is the orb hosting the service
`poa` is the POA which the service is to run in (see *POA Argument Choices*)
`config` is the service instance's configuration (described under *Configuration Structure* below)
`env` is the CORBA Environment

Step 4: Link with the *`ec_lwnaming_s`* library.

3.1.1.1 POA Argument Choices

The *poa* agrument allows the user to create their own poa for the Naming Service that has properties configured to meet the demands of their system or implementation.

There are two choices for the poa argument: `NULL` or User Defined POA

1. NULL

   In this case the `EORB_NamingService` will create a POA where:

   `id assignment policy` is set to *`PortableServer_USER_ID`*
   `id_uniqueness_policy` is set to *`PortableServer_MULTIPLE_ID`*

   `lifespan_policy` is set to `PortableServer_PERSISTENT` if the config option `qosPersistent` is set to `TRUE`, otherwise it will be set to `PortableServer_TRANSIENT`.

2. User Defined POA

The User Defined POA ***MUST*** have the following policies set:

`id assignment policy` set to *`PortableServer_USER_ID`*
`id_uniqueness_policy` set to *`PortableServer_MULTIPLE_ID`*

The config option *`qosPersistent`* will be ignored when a user defined POA is provided. If persistence is required, then the appropriate policies must be set on the user defined POA.

All other policies can be set as needed.

3.1.1.2 Configuration Structure

The structure mentioned in *Step 2:* above defines the property fields described below under Table 1, *Configuration Property Descriptions*. An example of setting the configuration structure fields is shown in Example 1, *Setting the Configuration Structure Fields*, on page 17.

Table 1 Configuration Property Descriptions

Property	Description
`qosContextLocking`	Controls the read and write locking protection for concurrent access to naming contexts. *`ctx_rwLock`* should normally only be set to *`false`*, unlocked, when the service is used in read-only mode.
`qosMaxContexts`	The number of naming contexts that this service instance is expected to create.
`qosPersistent`	Creates a persistent IOR for the Naming Service and causes it to listen on a fixed port. This enables the Naming Service to always be resolved using the same IOR or CORBALOC, even when the service has been shutdown and restarted.

Example 1 Setting the Configuration Structure Fields

The configuration structure is initialised by passing the four parameters listed in Table 1, *Configuration Property Descriptions*. For example:

```
EORB_NamingService_Config config;

config.qosContextLocking = TRUE;
config.qosMaxContexts    = 100;
config.qosPersistent     = TRUE;
```

3.1.1.3 Example Server Module

The following C code example shows how a Naming Service instance can be created.

For clarity, this example omits the error-checking code which would be required in a real system.

```
int main (int argc, char ** argv)
{
   CORBA_ORB orb;
   CORBA_Environment env;
   CosNaming_NamingContext context;
   EORB_NamingService_Config config;

   config.qosPersistent     = TRUE;
   config.qosContextLocking = TRUE;
   config.qosMaxContexts    = 100;

   /* Install plugins */

   EORB_POA_plugin ();
   EORB_IIOP_plugin ();

   /* Initialize the ORB */

   orb = CORBA_ORB_init (&argc, argv, "eorb-ce", &env);

   /* Initialize naming service */

   context = EORB_NamingService_init (orb, NULL, &config, &env);

   /* Register name service reference */

  CORBA_ORB_register_initial_reference (orb, "NameService", context,
&env);

   CORBA_Object_release (context, &env);

   /* Run ORB */

   CORBA_ORB_run (orb, &env);

   return 0;
}
```

3.1.2 Running from the Command Line

An alternative to writing a module which creates and runs a Naming Service server is to run the Spectra ORB Lightweight Naming Service executable, *`lwnamingc`*, from the command line. `lwnamingc` is located in the `$EORBHOME/bin/$EORBENV` directory. The `lwnamingc` executable can be run with zero or more of the command line options listed in *Table 2*.

Table 2 Command Line Options

Option	Description
`-NameServiceContextLocking <on|off>`	Sets context locking state (see above). The default is *on*
`-NameServiceMaxContexts <num>`	Sets expected number of contexts to be created. The default is *100*.
`-NameServicePersistent <yes|no>`	Sets whether the lifespan policy is set to `PortableServer_PERSISTENT` or `PortableServer_TRANSIENT` The default is *yes* (= *_PERSISTENT*).
`-NameServiceUIOP`	Runs the service on a UIOP endpoint. Only available on systems where UIOP is a supported transport. The default is `no`.

3.1.2.1 Example

The following command line example demonstrates how to start a Naming Service instance, on a UNIX operating system, where:

- the expected number of contexts to be created is *250*

```
%  lwnamingc -NameServiceMaxContexts 250
```

3.1.3 Running on a Fixed Endpoint

The server can be run on a fixed endpoint by running with the `-ORBPOAEndpoints` argument. The Name Service servant is created within a child POA *NameService* so for example can be run on a fixed IIOP endpoint with:

```
-ORBPOAEndpoints NameService:iiop:<host>:<port>
```

and resolved as an initial reference by a client using:

```
-ORBInitRef NameService=corbaloc:iiop:<host>:<port>/NameService
```

CHAPTER

4 Creating Applications

This section describes how to create applications by using specific features of the Spectra ORB Lightweight Naming Service, including:

- *creating and destroying naming contexts and name bindings*
- *retrieving the contents of a naming context*
- *how to resolve a binding to an object.*

Examples are provided which demonstrate how these features can be implemented.

Note:

This release of the Spectra ORB Lightweight Naming Service C Edition does ***not*** support the `kind` property of the `CosNaming::NameComponent`. The Service will ignore this property when binding or resolving names. Developers should ensure that they use unique `id` values creating name components.

- No CORBA system exceptions are caught in any of these examples; code to deal with them has been omitted for the sake of clarity and brevity. These exceptions must of course be properly caught and handled in a working system.
- The `CosNaming.h` header must be included in any application using the Spectra ORB Lightweight Naming Service and ensure your build system has `$(EORBHOME)/include/services/lw` on its include path.
- The `eOrbC/EORB/NamingService.h` must be included in any application using the Lightweight Naming Service.
- The linker must link with the `ec_lwnaming_c.lib` library file.

The Spectra ORB Lightweight Naming Service exceptions are listed under Section 5.1, *Exceptions*, on page 27.

4.1 Obtaining the Root Context

Before any objects or naming contexts can be added to (bound) or located (resolved) in the Spectra ORB Lightweight Naming Service, the root or initial context must be resolved. This is achieved by first obtaining the root context's IOR, then passing it to the `CORBA_ORB_string_to_object` function. This IOR is located, as a string, in a file created by the Spectra ORB Lightweight Naming Service server. The following example function, `obtainRootContext`, shows how the IOR can be obtained and used to resolve the root context.

```
static void obtainRootContext (CORBA_ORB orb,
CORBA_Environment *env)
{
   FILE *iorfile;
   char *filename = "naming.ior";
   char ior_str[256];

   iorfile = fopen (filename, "r");
   fgets (ior_str, 256, iorfile);

   root_ctx = CORBA_ORB_string_to_object (orb, ior_str, env);
   assert (root_ctx);
   fclose (iorfile);
}
```

4.2 Naming Context Creation and Destruction

The lightweight `NamingContext` interface provides one `NamingContext` creation operation and a single destroy operation, defined in IDL as:

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

void destroy () raises (NotEmpty);
```

The *`bind_new_context()`* operation creates a new Naming Context and binds it using the supplied name.

The *`destroy`* operation requests the destruction of a `NamingContext` object. The naming context must be empty. After `destroy` is invoked, no further operations can be invoked on the object reference of the naming context object.

Bindings to a destroyed context are not removed. To do so would require a context to know about all of its parents as well as its children. An attempt to resolve a binding to a destroyed context will throw the `CORBA.INV_OBJREF` exception. Accordingly, bindings to a naming context should be removed before it is destroyed.

4.3 Binding and Unbinding Operations

The lightweight NamingService `NamingContext` interface provides three bind operations and a single unbind operation, defined in IDL as:

```
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);

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

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

The bind operations allow binding to occur between a name and either a generic CORBA object or a Naming Context. In order to bind a CORBA object, the name to bind against must be correctly constructed. Given a name with *n* components, the first *n - 1* components must resolve to a bound `NamingContext`.

Note:

- The implementation for this version of the Spectra ORB Lightweight Naming Service does not support the *`NamingContextExt`* interface. However, in order to facilitate the creation of *`CosNaming_Name`* (a `NameComponent`), a utility function, *`EORB_NamingService_to_name()`*, is provided and can be used to create name components. Code which uses this utility must contain the following *`include`* statement:

  ```
  #include "eOrbC/EORB/NamingService.h"
  ```

  The *`EORB_NamingService_to_name()`* function signature is:

  ```
  CosNaming_Name *EORB_NamingService_to_name (char *name,
     CORBA_Environment *env)
  ```

- the `NameComponent` *`kind`* element is not currently supported: name components can only be specified using the *`id`* element, where each `id` value is unique within a naming context
- stringified names are supported, together with the use of the "\" escape character, as in accordance with the OMG specification

---

The `rebind` operation is identical to the `bind` operation except that the `AlreadyBound` exception is not thrown; an existing binding with the same name is replaced by the new binding.

The `bind_new_context` operation is equivalent to creating a new `NamingContext` and then adding it using `bind_context`:

The following example uses a function called *example* to create ten naming contexts, assigns a unique name to each, then binds them to the root context using *CosNaming_NamingContext_bind_new_context*. (The names created in the example, *context_N* where *N* is a sequential value, are generated by the example function, *create_name_string*.

```
static CosNaming_NamingContext root_ctx;

/* utility */
static void create_name_string (int i, char * name, char **
result)
{
   char num [8];

   sprintf (num, "%d", i);
   strcpy (*result, name);
   strcat (*result, num);
}
static void example (CORBA_Environment *env)
{
   CosNaming_Name *cos_name;
   CosNaming_Binding *cos_binding;
   CosNaming_NamingContext ctx, recovered;
   CORBA_Object object;
   char *name;
   int i;

   /* bind 10 contexts into root context */
   for (i = 0; i < 10; i++)
   {
      name = malloc (32);
      create_name_string (i + 1, "context_", &name);

      cos_name = (CosNaming_Name*)
         EORB_NamingService_to_name (name, env);
      ctx = CosNaming_NamingContext_bind_new_context
         (root_ctx, cos_name, env);

      free (name);
      CORBA_free (cos_name);
      CORBA_Object_release (ctx, &env);
   }
```

The unbind operation, *CosNamingContext_unbind(CosNamingContext, CORBA_SequenceNameComponent, CORBA_Environment)*, removes a name binding. It does not matter which of the bind operations was used to create the binding. The following example destroys bindings created with the previous example:

```
static void bind_unbind_example (CORBA_Environment *env)
```

```
{
   CosNaming_Name *cos_name;
   CosNaming_NamingContext object;
   char *name = "entry_1";

   cos_name = (CosNaming_Name*)
      EORB_NamingService_to_name (name, env);
   object = CosNaming_NamingContext_bind_new_context
      (root_ctx, cos_name, env);

   CosNaming_NamingContext_unbind (root_ctx, cos_name, env);

   CORBA_free (cos_name);
   CORBA_Object_release (object, &env);
}
```

4.4 Accessing Naming Context Contents

In the Lightweight Naming Service, only one operation is available for accessing the contents of naming contexts, defined in IDL as:

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

The `resolve` operation returns the object, if any, bound to the name passed to it.

CHAPTER

5 Supplemental Information

5.1 Exceptions

The exceptions raised by the Spectra ORB Lightweight Naming Service are listed in *Table 3*.

Table 3 Spectra ORB Lightweight Naming Service Exceptions

Name	Purpose
AlreadyBound	Indicates an object is already bound to the specified name. Only one object can be bound to a particular name in a context.
CannotProceed	Indicates that the implementation has given up for some reason. The client, however, may be able to continue the operation at the returned naming context. One possible reason for this exception is that a Name Server holding one or more of the name bindings within a compound name is currently unavailable.
InvalidName	Indicates that the name is invalid. This implementation disallows zero length names only.
NotEmpty	Indicates that a naming context has bindings.
NotFound	Indicates that the name does not identify a binding or that the binding is not of the type required for the requested operations.

Index

Index