
Life Cycle Service Specification

Version 1.1 (updated due to slight changes in the OMG IDL based on the CORBA 2.3 specification)
New Edition: April 2000

Copyright 1994 AT&T/NCR
Copyright 1994 BNR Europe Limited
Copyright 1994 Digital Equipment Corporation
Copyright 1994 Groupe Bull
Copyright 1994 Hewlett-Packard Company
Copyright 1994 HyperDesk Corporation
Copyright 1994 ICL plc
Copyright 1994 International Business Machines Corporation
Copyright 1994 Itasca Systems, Inc.
Copyright 1994 Novell, Inc.
Copyright 1994 O2 Technologies
Copyright 1994 Object Design, Inc.
Copyright 1994 Objectivity, Inc.
Copyright 1994 Ontos, Inc.
Copyright 1994 Oracle Corporation
Copyright 1994 Persistence Software
Copyright 1994 Servio, Corp.
Copyright 1994 SunSoft, Inc.
Copyright 1994 Teknekron Software Systems, Inc.
Copyright 1994 Tivoli Systems, Inc.
Copyright 1994 Versant Object Technology Corporation

The companies listed above have granted to the Object Management Group, Inc. (OMG) a nonexclusive, royalty-free, paid up, worldwide license to copy and distribute this document and to modify this document and distribute copies of the modified version. Each of the copyright holders listed above has agreed that no person shall be deemed to have infringed the copyright in the included material of any such copyright holder by reason of having used the specification set forth herein or having conformed any computer software to the specification.

PATENT

The attention of adopters is directed to the possibility that compliance with or adoption of OMG specifications may require use of an invention covered by patent rights. OMG shall not be responsible for identifying patents for which a license may be required by any OMG specification, or for conducting legal inquiries into the legal validity or scope of those patents that are brought to its attention. OMG specifications are prospective and advisory only. Prospective users are responsible for protecting themselves against liability for infringement of patents.

NOTICE

The information contained in this document is subject to change without notice. The material in this document details an Object Management Group specification in accordance with the license and notices set forth on this page. This document does not represent a commitment to implement any portion of this specification in any company's products.

WHILE THE INFORMATION IN THIS PUBLICATION IS BELIEVED TO BE ACCURATE, THE OBJECT MANAGEMENT GROUP AND THE COMPANIES LISTED ABOVE MAKE NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THIS MATERIAL INCLUDING, BUT NOT LIMITED TO ANY WARRANTY OF TITLE OR OWNERSHIP, IMPLIED WARRANTY OF MERCHANTABILITY OR WARRANTY OF FITNESS FOR PARTICULAR PURPOSE OR USE. In no event shall The Object Management Group or any of the companies listed above be liable for errors contained herein or for indirect, incidental, special, consequential, reliance or cover damages, including loss of profits, revenue, data or use, incurred by any user or any third party. The copyright holders listed above acknowledge that the Object Management Group (acting itself or through its designees) is and shall at all times be the

sole entity that may authorize developers, suppliers and sellers of computer software to use certification marks, trademarks or other special designations to indicate compliance with these materials. This document contains information which is protected by copyright. All Rights Reserved. No part of this work covered by copyright herein may be reproduced or used in any form or by any means--graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems--without permission of the copyright owner.

RESTRICTED RIGHTS LEGEND. Use, duplication, or disclosure by government is subject to restrictions as set forth in subdivision (c) (1) (ii) of the Right in Technical Data and Computer Software Clause at DFARS 252.227.7013 OMG[®] and Object Management are registered trademarks of the Object Management Group, Inc. Object Request Broker, OMG IDL, ORB, CORBA, CORBAfacilities, CORBAservices, COSS, and IIOP are trademarks of the Object Management Group, Inc. X/Open is a trademark of X/Open Company Ltd.

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

Preface	iii
About the Object Management Group	iii
What is CORBA?	iii
Associated OMG Documents	iv
Acknowledgments	iv
1. Service Description	1-1
1.1 Overview	1-1
1.1.1 The problem of creation	1-2
1.1.2 The problem of moving or copying an object	1-2
1.1.3 The problem of operating on a graph of distributed objects	1-3
1.2 Client's Model of Object Life Cycle	1-3
1.2.1 Client's Model of Creation	1-4
1.2.2 Client's model of deleting an object	1-5
1.2.3 Client's model of copying or moving an object	1-5
1.3 Factory Finders	1-6
1.3.1 Multiple Factory Finders	1-7
1.4 Design Principles	1-7
1.5 Resolution of Technical Issues	1-8
2. Life Cycle Interfaces	2-1
2.1 The CosLifeCycle Module	2-1
2.1.1 The LifeCycleObject Interface	2-3
2.1.2 The FactoryFinder Interface	2-4
2.1.3 The GenericFactory Interface	2-5

Contents

2.1.4	Criteria	2-8
2.2	Implementing Factories	2-10
2.2.1	Minimal Factories	2-10
2.2.2	Administered Factories	2-11
2.3	Target's Use of Factories and Factory Finders	2-12
2.4	Summary of Life Cycle Service	2-12
2.4.1	Summary of Life Cycle Service Structure	2-13
Appendix A - Compound Life Cycle Specification		A-1
Appendix B - References		B-1
Appendix C - Filters		C-1
Appendix D - Administration		D-1
Appendix E - Supports for PCTE Objects		E-1

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

The following companies submitted and/or supported parts of the *Life Cycle Service* specification:

- AT&T/NCR
- BNR Europe Limited
- Digital Equipment Corporation
- Groupe Bull
- Hewlett-Packard Company
- HyperDesk Corporation
- ICL plc
- International Business Machines Corporation
- Itasca Systems, Inc.

-
- Novell, Inc.
 - 02 Technologies
 - Object Design, Inc.
 - Objectivity, Inc.
 - Ontos, Inc.
 - Oracle Corporation
 - Persistence Software
 - Servio, Corp.
 - SunSoft, Inc.
 - Teknekron Software Systems, Inc.
 - Tivoli Systems, Inc.
 - Versant Object Technology Corporation

Service Description

Contents

This chapter contains the following sections.

Section Title	Page
“Overview”	1-1
“Client’s Model of Object Life Cycle”	1-3
“Factory Finders”	1-7
“Design Principles”	1-7
“Resolution of Technical Issues”	1-8

Note – Appendix A contains an addendum to the Life Cycle Service; the addendum provides a specification for compound life cycle operations. Appendix B provides a list of References. This specification also includes additional appendices that are not part of the Life Cycle Service specification: they are included as background material. Appendix C suggests a filtering language for the filter criteria. Appendix D discusses administration of generic factories. Appendix E discusses support for PCTE objects.

1.1 Overview

Life Cycle Service defines services and conventions for creating, deleting, copying and moving objects. Because CORBA-based environments support distributed objects, the Life Cycle Service defines conventions that allow clients to perform life cycle operations on objects in different locations.

This overview describes the life cycle problem for distributed object systems.

1.1.1 The problem of creation

Figure 1-1 illustrates the problem of a client in one location creating an object in another.



Figure 1-1 Life Cycle service defines how a client can create an object “over there.”

To create an object in a different location, the following questions must be answered:

- Can the client control the location for the new object?
- On the other hand, can the location be determined according to some administered policy?
- What entity does the client communicate with in order that a new object is created?
- How does the client find that entity?
- How much control does the client have over deciding the implementation of the created object?
- Can the client influence the initial values of the newly created object?
- Can the client create an object in an implementation specific fashion?

1.1.2 The problem of moving or copying an object

Figure 1-2 illustrates the problem of moving or copying an object in a distributed object system.




Figure 1-2 Life Cycle Service defines how a client can move or copy an object over there.

To support moving or copying an object, the following questions must be answered:

- Can the client control the location for the copied or migrated object?

- On the other hand, can the location be determined according to some administered policy?
- What entity does the client communicate with to copy or migrate the object?
- How does the client find that entity?
- What happens to the implementation code of a copied or migrated object?

1.1.3 The problem of operating on a graph of distributed objects

Distributed objects do not *float* in space; they are connected to one another. The connections are called *relationships*. Relationships allow semantics to be added to references between objects. For example, relationships allow one object to *contain* another. Life Cycle services must work in the presence of graphs of related objects.

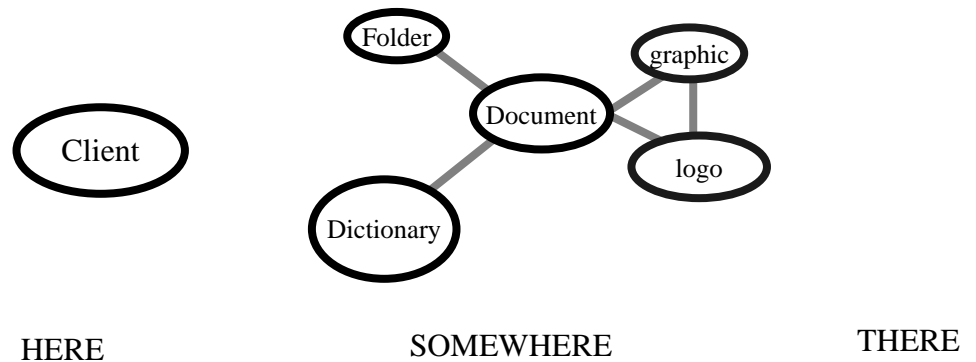


Figure 1-3 The object life cycle problem for graphs of objects is to determine the boundaries of a graph of objects and operate on that graph. In the above example, a document contains a graphic and a logo, refers to a dictionary and is contained in a folder.

Figure 1-3 illustrates the object life cycle problem for graphs of objects. In the example, the folder *contains* a document, the document *contains* a graphic and a logo and *references* a dictionary. The graphic *references* the logo that is *contained* in the document. For graphs of objects, life cycle services must answer the following questions:

- What are the boundaries of the graph? For example, if a client copies the document, which objects are affected?
- If multiple objects are affected, how is the life cycle operation actually applied to those objects?
- Are cycles in the graph preserved? For example, if copying the document results in copying the graphic and the logo, is the cycle preserved in the copy?

1.2 Client's Model of Object Life Cycle

A client is any piece of code that initiates a life cycle operation for some object. A client has a simple view of the life cycle operations.

1.2.1 Client's Model of Creation

The client's model of creation is defined in terms of factory objects. A factory is an object that creates another object. Factories are *not* special objects. As with any object, factories have well-defined IDL interfaces and implementations in some programming language.




Figure 1-4 To create an object “over there” a client must possess an object reference to a factory over there. The client simply issues a request on the factory.

There is no standard interface for a factory. Factories provide the client with specialized operations to create and initialize new instances in a natural way for the implementation. The following illustrates a factory for a document.

```
interface DocFactory {
    Document create();
    Document create_with_title(in string title);
    Document create_for(in natural_language nl);
};
```

This interface is defined for clients as a part of application development.

Factories are object implementation dependent. A different implementation of the document could define a different factory interface.

While there is no standard interface for a factory, a generic factory interface is defined by the life cycle service in Section 2.1.3, “The GenericFactory Interface,” on page 2-5. A generic factory is a creation service. It provides a generic operation for creation. Instead of invoking an object specific operation on a factory with statically defined parameters, the client invokes a standard operation whose parameters can include information about resource filters, state initialization, policy preferences, etc.

To create an object, a client must possess an object reference for a factory, which may be either a generic factory or an object-specific factory, and issue an appropriate request on the factory. As a result, a new object is created and typically an object reference is returned.

There is nothing special about this interaction.

A factory assembles the resources necessary for the existence of an object it creates. Therefore, the factory represents a scope of resource allocation, which is the set of resources available to the factory. A factory may support an interface that enables its clients to constrain the scope.

Clients find factory objects in the same fashion they find any object. Two common scenarios for clients to find factories are:

- Clients use a finding mechanism, such a naming context, drag-and-drop, or a trader, to find factories.
- Clients are passed factory objects as a parameter to an operation the client supports.

Various implementation strategies for factories are discussed in detail in Section 2.2, “Implementing Factories,” on page 2-10.

1.2.2 Client’s model of deleting an object

A client that wishes to delete an object issues a **remove**¹ request on an object supporting the `LifeCycleObject` interface. (The `LifeCycleObject` interface is defined in Section 2.1, “The `CosLifeCycle` Module,” on page 2-1.) The object receiving the request is called the target.




Figure 1-5 Illustrates a client deleting the document

To delete an object, a client must possess an object reference supporting the **LifeCycleObject** interface and issues a **remove** request on the object.

1.2.3 Client’s model of copying or moving an object

A client that wishes to move or copy an object issues a **move** or **copy** request on an object supporting the **LifeCycleObject** interface. The object receiving the request is called the target.

1. The operation is named `remove`, rather than `delete`, because `delete` collides with the `delete` operator in C++.

The move and copy operations expect an object reference supporting the **FactoryFinder** interface. The factory finder represents the “THERE” in Figure 1-6. The client is indicating to move or copy the target using a factory within the scope of the factory finder. Section 1.3, “Factory Finders,” on page 1-7 describes factory finders in more detail.

The implementations of move and copy can use the factory finder to find appropriate factories “over there.” Section 2.3, “Target’s Use of Factories and Factory Finders,” on page 2-12 describes how objects can implement move and copy using the factory finder. This is invisible to the client.




Figure 1-6 Life cycle services define how a client can move or copy an object from here to there.

In the example of Figure 1-6, client code would simply issue a **copy** request on the document and pass it an object supporting the **FactoryFinder** interface as an argument.

When a client issues a **copy** request on a target, it is assumed that the target, the factory finder, and the newly created object can all communicate via the ORB. With externalization/internalization there is no such assumption. In the presence of a future externalization service, the externalized form of the object can exist outside of the ORB for arbitrary amounts of time, be transported by means outside of the ORB and can be internalized in a different, disconnected ORB.

Note – In general, a client is unaware of how a target and a factory finder are implemented. The target may represent a simple object or it may represent a graph of objects. Similarly, a factory finder may represent a very concrete location, such as a specific storage device, or it may represent a more abstract location, such as a group of machines. The client uses the same interface in all of these cases.

1.3 Factory Finders

Factory finders support an operation, **find_factories**, which returns a sequence of factories. Clients pass factory finders to the move and copy operations, which typically invoke this operation to find a factory to interact with. (This is described in detail in Section 2.3, “Target’s Use of Factories and Factory Finders,” on page 2-12.) The new copy or the migrated object will then be within the scope of the factory finder.

Some examples of locations that a factory finder might represent are:

- somewhere on a work group’s local area network
- storage device A on machine X
- Susan’s notebook computer

1.3.1 Multiple Factory Finders

The factory finder interface given in Section 2.1, “The CosLifeCycle Module,” on page 2-1 represents the minimal functionality supported by all factory finders. Target implementations can depend on this operation being available. More sophisticated factory finding facilities can be provided by extended finding services.

Currently, the only finding service being considered for standardization by the OMG is the naming service. Others are likely to be standardized in the future. It is likely that there will always be multiple finding services, of different expressive powers, in distributed object systems.

As demonstrated in Figure 1-7, the `FactoryFinder` interface can be mixed-in with interfaces for finding services, allowing multiple finding services. Many clients simply pass factory finders on to target objects. However, objects that need the services of a more powerful finding mechanism can narrow the factory finder to an appropriate, more specific interface.




Figure 1-7 The `FactoryFinder` interface can be “mixed in” with interfaces of more powerful finding services.

The power of a factory finder is determined by the power of the finding service.

1.4 Design Principles

Several principles have driven the design of the Life Cycle Service:

1. A factory object registered at a factory finder represents an implementation at that location. Thus, a factory finder allows clients to query a location for an implementation.
2. Object implementations can embody knowledge of finding a factory, relative to a location. Object implementations usually do not embody knowledge of location.
3. The desired result for life cycle operations such as copy and move depends on relationships between the target object and other objects. The design given in Appendix A has built-in support for the two most basic kinds of relationships, *containment* and *reference*, and supports the definition of new kinds of relationships and propagation semantics.
4. The Life Cycle Service is not dependent on any particular model of persistence and is suitable for distributed, heterogeneous environments.
5. The design does not include an object equivalence service nor rely on global object identifiers.

1.5 Resolution of Technical Issues

This specification addresses the following issues that were identified for the Life Cycle Service in the *OMG Object Services Architecture*² :

- *Creation*: Many of the parameters supplied to an object `create` operator will be implementation-dependent, so that a standardized universal IDL signature for object creation is not possible. IDL signatures for object creation will be defined for various kinds of object factories, but the signatures will be specific to type, implementation, and persistent storage mechanism of the object to be created.
- *Deletion*: A `remove` operator is defined on any object supporting the `LifeCycleObject` interface. This model for deletion supports any desired paradigm for referential integrity. Appendix A describes support for the two most common paradigms, based on reference and containment relationships. Only one type of deletion is supported; a different operation should be used for archiving an object. This interface can support many paradigms for storage management (e.g., garbage collection and reference counts). Since storage management is implementation-dependent, its interface does not belong in the generalized life cycle interfaces.
- *Copying*: Appendix A describes support for shallow and deep copy, and referential integrity. A scheme based on reference and containment relationships defines scopes for operations such as copy. The concept of a factory finder is used for object location. This paradigm for copying, deleting, and moving objects works regardless of an object's ORB, persistent storage mechanism, and implementation. This design is extensible because objects participate in the traversal algorithm, and the relationship service presented in the appendix supports the definition of new kinds of relationships with different behavior.

2. *Object Services Architecture*, Document Number 92-8-4, Object Management Group, Framingham, MA, 1992.

- *Equivalence*: There was no need for an object equivalence service or global object identifiers in the design of the Life Cycle Service to support real world applications or other object services.

Life Cycle Interfaces

This chapter was updated based on slight changes to the OMG IDL due to CORBA 2.3. Changes are marked with changebars.

Contents

This chapter contains the following sections.

Section Title	Page
“The CosLifeCycle Module”	2-1
“Implementing Factories”	2-10
“Target’s Use of Factories and Factory Finders”	2-12
“Summary of Life Cycle Service”	2-12

2.1 The CosLifeCycle Module

Client code accesses the basic life cycle functionality via the **CosLifeCycle** module. This module defines the **FactoryFinder**, **LifeCycleObject**, and **GenericFactory** interfaces and describes the operations of these interfaces in detail.

```
#include <CosNaming.idl>  
  
#pragma prefix “omg.org”  
  
module CosLifeCycle{  
  
    typedef CosNaming::Name Key;  
    typedef Object Factory;  
    typedef sequence <Factory> Factories;  
  
}
```

```

typedef struct NVP {
    CosNaming::lstring  name;
    any                value;
} NameValuePair;
typedef sequence <NameValuePair> Criteria;

exception NoFactory {
    Key search_key;
};
exception NotCopyable { string reason; };
exception NotMovable { string reason; };
exception NotRemovable { string reason; };
exception InvalidCriteria{ Criteria invalid_criteria; };
exception CannotMeetCriteria { Criteria unmet_criteria; };

interface FactoryFinder {
    Factories find_factories(in Key factory_key)
        raises(NoFactory);
};

interface LifecycleObject {
    LifecycleObject copy(in FactoryFinder there,
                        in Criteria the_criteria)
        raises(NoFactory, NotCopyable, InvalidCriteria,
              CannotMeetCriteria);
    void move(in FactoryFinder there,
             in Criteria the_criteria)
        raises(NoFactory, NotMovable, InvalidCriteria,
              CannotMeetCriteria);
    void remove()
        raises(NotRemovable);
};

interface GenericFactory {
#ifdef NO_ESCAPED_IDENTIFIERS
    boolean _supports(in Key k);
#else
    boolean _supports(in Key k);
#endif
    Object create_object (
        in Key          k,
        in Criteria     the_criteria)
        raises (NoFactory, InvalidCriteria,
              CannotMeetCriteria);
};
};

```

2.1.1 The *LifeCycleObject* Interface

The **LifeCycleObject** interface defines copy, move, and remove operations. Objects participate in the life cycle service by supporting this interface.

2.1.1.1 *copy*

**LifeCycleObject copy(in FactoryFinder there,
in Criteria the_criteria)
raises(NoFactory, NotCopyable, InvalidCriteria,
CannotMeetCriteria);**

The **copy** operation makes a copy of the object. The copy is located in the scope of the factory finder passed as the first parameter. The **copy** operation returns an object reference to the new object. The new object is initialized from the existing object.

The first parameter, **there**, may be a nil object reference. If passed a nil object reference, the target object can determine the location or fail with the **NoFactory** exception.

The second parameter, **the_criteria**, allows for a number of optional parameters to be passed. Typically, the target simply passes this parameter to the factory used in creating the new object. The criteria parameter is explained in detail in Section 2.1.4, “Criteria,” on page 2-8

If the target cannot find an appropriate factory to create a copy “over there,” the **NoFactory** exception is raised. An implementation that refuses to copy itself should raise the **NotCopyable** exception. If the target does not understand the criteria, the **InvalidCriteria** exception is raised. If the target understands the criteria but cannot satisfy the criteria, the **CannotMeetCriteria** exception is raised.

In addition to these exceptions, implementations may raise standard CORBA exceptions. For example, if resources cannot be acquired for the copied object, **NO_RESOURCES** will be raised. Similarly, if a target does not implement the **copy** operation, the **NO_IMPLEMENT** exception will be raised.

It is implementation dependent whether this operation is atomic.

2.1.1.2 *move*

**void move(in FactoryFinder there,
in Criteria the_criteria)
raises(NoFactory, NotMovable, InvalidCriteria,
CannotMeetCriteria);**

The **move** operation on the target moves the object to the scope of the factory finder passed as the first parameter. The object reference for the target object remains valid after **move** has successfully executed.

The first parameter, **there**, may be a nil object reference. If passed a nil object reference, the target object can determine the location or fail with the **NoFactory** exception.

The second parameter, **the_criteria**, allows for a number of optional parameters to be passed. Typically, the target simply passes this parameter to the factory used in migrating the new object. The criteria parameter is explained in detail in Section 2.1.4, “Criteria,” on page 2-8.

If the target cannot find an appropriate factory to support migration of the object “over there,” the **NoFactory** exception is raised. An implementation that refuses to move itself should raise the **NotMovable** exception. If the target does not understand the criteria, the **InvalidCriteria** exception is raised. If the target understands the criteria but cannot satisfy the criteria, the **CannotMeetCriteria** exception is raised.

In addition to these exceptions, implementations may raise standard CORBA exceptions. For example, if resources cannot be acquired for migrating the object, **NO_RESOURCES** will be raised. Similarly, if a target does not implement the **move** operation, the **NO_IMPLEMENT** exception will be raised.

It is implementation dependent whether this operation is *atomic*.

2.1.1.3 *remove*

```
void remove()  
raises(NotRemovable);
```

Remove instructs the object to cease to exist. The object reference for the target is no longer valid after **remove** successfully completes. The client is not responsible for cleaning up any resources the object uses. An implementation that refuses to remove itself should raise the **NotRemovable** exception. In addition to this exception, implementations may raise standard CORBA exceptions.

2.1.2 *The FactoryFinder Interface*

Factory finders support an operation, **find_factories**, which returns a sequence of factories. Clients pass factory finders to the **move** and **copy** operations, which typically invoke this operation to find a factory to interact with. (This is described in detail in Section 2.3, “Target’s Use of Factories and Factory Finders,” on page 2-12.)

The factory finder interface represents the *minimal* functionality supported by all factory finders.

2.1.2.1 *find_factories*

```
Factories find_factories(in Key factory_key)  
raises(NoFactory);
```

The **find_factories** operation is passed a key used to identify the desired factory. The key is a name, as defined by the naming service. More than one factory may match the key. As such, the factory finder returns a sequence of factories. If there are no matches, the **NoFactory** exception is raised.

The scope of the key is the factory finder. The factory finder assigns no semantics to the key. It simply matches keys. It makes no guarantees about the interface or implementation of the returned factories or objects they create.

It is beyond the scope of this specification to standardize the key space. The space of keys is established by *convention* in particular environments. The *kind* field¹ of the key is useful for partitioning the key space. Suggested values for the **id** and **kind** fields are given in Table 2-1.

Table 2-1 Suggested conventions for factory finder keys.

id field	kind field	meaning
name of object interface	“object interface”	Find factories that create objects supporting the named interface.
name of equivalent implementations	“implementation equivalence class”	Find factories that create objects with implementations in a named equivalence class of implementations. ¹
name of object implementation	“object implementation”	Find factories that create objects of a particular implementation.
name of factory interface	“factory interface”	Find factories supporting the named factory interface.

1. An example of an implementation equivalence class is a set of object implementations that have compatible externalized forms.

2.1.3 The GenericFactory Interface

In many environments, management of a set of resources that are allocated to objects at creation time is required. This needs to be done in a coordinated fashion for all types of objects. The Life Cycle Service provides a framework for this which is intended to be usable in a variety of administrative environments. However, the differing environments will administer a variety of resources and it is beyond the scope of this framework to identify all the possible types of resource.

1. See the naming service specification.

While there is no standard interface for a factory, a **GenericFactory** interface is defined. The GenericFactory interface defines a generic creation operation, **create_object**. By defining a generic interface for creation, a creation service can be implemented. This is particularly useful in environments where administering a set of resources is important.

Such a generic factory can implement resource policies and represent multiple locations. In administered environments, object specific factories, such as the document factory described in Section 1.2, “Client’s Model of Object Life Cycle,” on page 1-3, may delegate the creation process to the generic factory. This is described in detail in Section 2.2.2, “Administered Factories,” on page 2-11.

The job of the generic factory is to match the creation criteria specified by clients of the **GenericFactory** interface with offers made on behalf of implementation specific factories.

The Life Cycle service provides a generic creation capability. Ultimately, implementation specific creation code is invoked by the creation service. The implementation specific code also supports the **GenericFactory** interface.

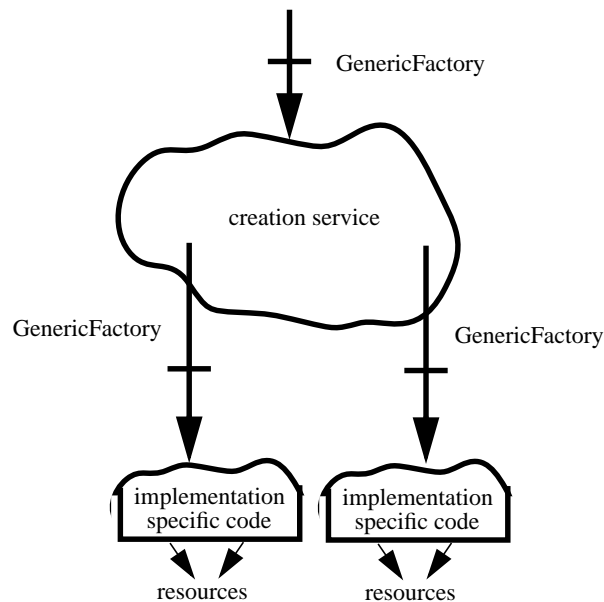


Figure 2-1 illustrates the structure of a creation service

The client of the GenericFactory interface invokes the **create_object** operation and can express criteria for creation.

Ultimately, this request will be passed to an implementation specific factory which supports the **GenericFactory** interface. To get there, the request may travel through a number of generic factories. However, all of this is transparent to the client.

2.1.3.1 *create_object*

```

boolean _supports(in Key k);
                in Key k,
                in Criteria the_criteria)
raises (NoFactory, InvalidCriteria,
        CannotMeetCriteria);

```

The **create_object** operation is passed a key used to identify the desired object to be created. The key is a name, as defined by the Naming Service.

The scope of the key is the generic factory. The generic factory assigns no semantics to the key. It simply matches keys. It makes no guarantees about the interface or implementation of the created object.

It is beyond the scope of this specification to standardize the key space. The space of keys is established by *convention* in particular environments. The kind field² of the key is useful for partitioning the key space. Suggested values for the **id** and **kind** fields are given in Table 2-2.

Table 2-2 Suggested conventions for generic factory keys.

id field	kind field	meaning
name of object interface	“object interface”	Create an object that supports the named interface.
name of equivalent implementations	“implementation equivalence class”	Create an object whose implementation is in a named equivalence class of implementations. ¹
name of object implementation	“object implementation”	Create objects of a particular implementation.

1. An example of an implementation equivalence class is a set of object implementations that have compatible externalized forms

The second parameter, **the_criteria**, allows for a number of optional parameters to be passed. Criteria are explained in detail in Section 2.1.4, “Criteria,” on page 2-8

If the generic factory cannot create an object specified by the key, then **NoFactory** is raised.

If the target does not understand the criteria, the **InvalidCriteria** exception is raised. If the target understands the criteria but cannot satisfy the criteria, the **CannotMeetCriteria** exception is raised.

2. See the naming service specification.

2.1.3.2 *supports*

boolean _supports(in Key k);

The **supports** operation returns **true** if the generic factory can create an object, given the key; otherwise, **false** is returned.

2.1.4 *Criteria*

The **create_object** operation of the GenericFactory interface expects a parameter specifying the creation criteria. The **move** and **copy** operations of the LifecycleObject interface also expects this parameter; typically they pass it through to a factory. This section documents this parameter.

The criteria parameter is expressed as an IDL sequence of name-value pairs. In particular, it is described by the following data structure given in the **CosLifeCycle** module:

```
typedef struct NVP {  
    CosNaming::Istring name;  
    any value;  
} NameValuePair;  
typedef sequence <NameValuePair> Criteria;
```

The parameter is given as a sequence of name-value pairs in order to be extensible and support “pass-through;” that is, new name-value pairs can be defined in the future and objects can be written that do not interpret the name-value pairs, but just pass them on to other objects.

Note – It is beyond the scope of this specification to standardize particular criteria. Supporting criteria is optional. Furthermore, supporting different criteria is acceptable. The criteria given here are suggestions.

Table 2-3 suggests criteria to be supported by the generic factory. Detailed descriptions follow.

Table 2-3 Suggested criteria.

critierion name	type of criterion value	interpretation
“initialization”	sequence<NameValuePair>	initialization parameters, given as a sequence of name-value pairs.
“filter”	string	allows clients of the generic factory to express a constraint on the created object.
“logical location”	sequence<NameValuePair>	allows clients of the generic factory to express a connection for the object, for example a PCTE relationship.
“preferences”	string	a way for clients to influence the policies that a generic factory may use when creating an object

“initialization”

The “initialization” criterion is a sequence of name-value pairs which is intended to contain application specific initialization values. Typically, the generic factory will pay no attention to the initialization criterion and simply passes it on to application specific factory code.

“filter”

The filter criterion is a constraint expression which provides the client with a powerful way of expressing its requirements on creation. The generic factory will use the constraint expression to make decisions about the allocation of particular resources. For example, a client could give a constraint “operating system” != “windows nt”.

These constraints are expressed in some Constraint Language. A constraint language is suggested in Appendix B.

Filters are potentially complex and **InvalidCriteria** will be raised if the filter is too complex for the factory or is syntactically incorrect.

“logical location”

The “logical location” criterion allows a client to express where a created/copied/migrated object is logically created. For example, in PCTE an object is always in a relationship with another object. In such an environment, the logical location would specify another object and a relationship.

“preferences”

The “preferences” criterion allows the client to influence the policies which the generic factory uses to make decisions. For example, a generic factory might arbitrarily choose a machine from a set of machines. Using the preferences criterion, a client could express its preference for a particular machine. Policies and preferences are described in more detail in Appendix B.

2.2 *Implementing Factories*

As defined under Section 1.2, “Client’s Model of Object Life Cycle,” on page 1-3, any object that creates another object in response to some request is called a factory. Clients depend only on the definitions in that section.

The client’s model of object life cycle has intentionally been defined abstractly. This allows a wide variety of implementation strategies.

Factories are not special objects. They have well-defined IDL interfaces and implementations in programming languages. Defining factory interfaces and implementing them are a normal part of application development.

Ultimately, the creation process requires implementation dependent code that assembles resources for the storage and execution of an object. The act of creating an object requires assembling and initializing all of the resources required to support the execution and storage of the object. The resources typically include:

- the allocation of one or more BOA object references, and
- resources related to persistence storage.

2.2.1 *Minimal Factories*

Figure 2-2 illustrates a minimal implementation of a factory that assembles resources in a single factory object.




Figure 2-2 Factories assemble resources for the execution of an object. A minimal implementation achieves this with a single factory implementation.

2.2.2 Administered Factories

Factories can delegate the creation process to a generic factory that administers a set of resources. The generic factory may apply policies to all creation requests.

Eventually such a generic creation service, needs to communicate with implementation specific code that actually assembles the resources for the object. Figure 2-3 illustrates an object specific factory, such as the document factory of Figure 2-4 on page 2-12 that delegates the creation problem to the generic creation service. The object-specific factory effectively adds a statically typed wrapper around the generic factory.

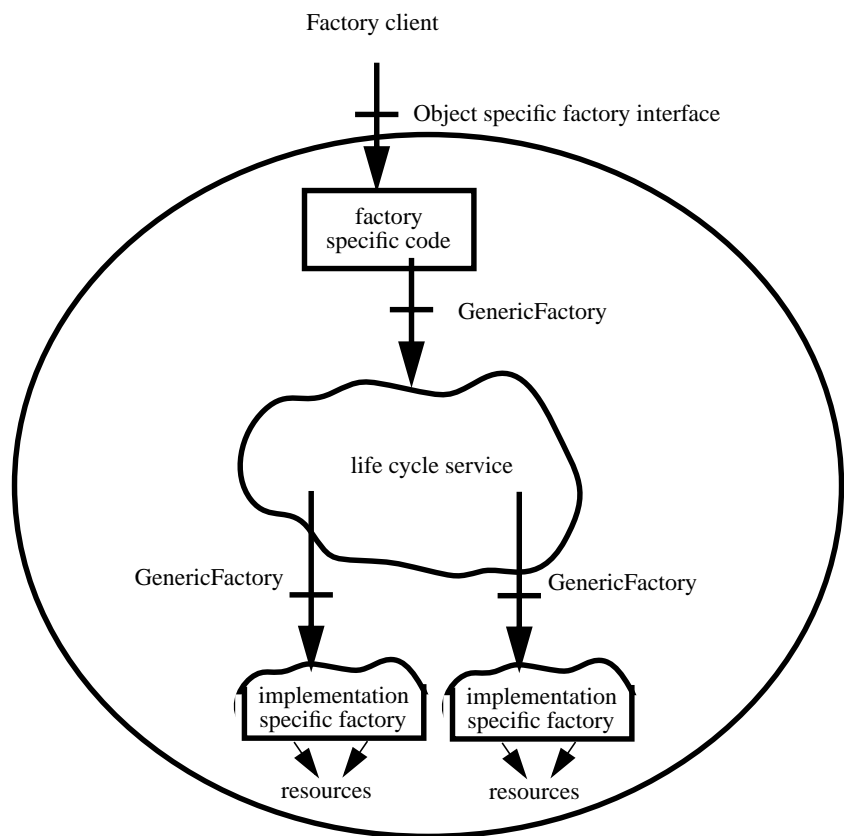


Figure 2-3 In an administered environment, factory *implementations* can delegate the creation problem to a generic factory. The generic factory can apply resource allocation policies. Ultimately the creation service communicates with implementation specific code that assembles resources for the object.

2.3 Target's Use of Factories and Factory Finders

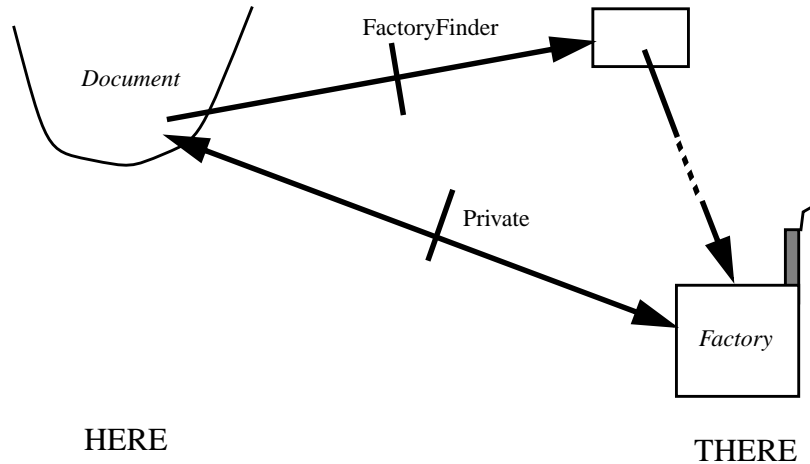


Figure 2-4 The copy and move operations are passed a `FactoryFinder` to represent “there.” The implementation of the target uses the `FactoryFinder` to find a factory object for creation over there. The protocol between the object and the factory is private. They can communicate and transfer state according to any implementation-defined protocol.

A client passes a factory finder as a parameter to a **copy** or **move** request.

Clients do not generally understand the implementation constraints of the object being copied. Clients cannot express what the target object needs in order to copy itself to the new location.

Target object implementations, on the other hand, put constraints on factories based on implementation concerns. It is unlikely that target implementation code is interested in further constraining location.

To find an appropriate factory, the target object implementation may use the factory finder with its minimal interface defined in Section 2.1.2, “The `FactoryFinder` Interface,” on page 2-4 or it may attempt to narrow the factory finder to a more sophisticated finding service with more expressive power. The target object implementation can always depend on the existence of the minimal interface.

Once the target object implementation finds a factory, it communicates with the factory using a private, implementation-defined, interface.

2.4 Summary of Life Cycle Service

The problem of distributed object life cycle is the problem of

- Creating an object

- Deleting an object
- Moving and copying an object
- Operating on a graph of distributed objects.

The client's model of object life cycle is based on *factories* and target objects supporting the **LifeCycleObject** interface. Factories are objects that create other objects. The **LifeCycleObject** interface defines operations to delete an object, to move an object and to copy an object.

A **GenericFactory** interface is defined. The generic factory interface is sufficient to create objects of different types. By defining a **GenericFactory** interface, implementations that administer resources are enabled.

2.4.1 Summary of Life Cycle Service Structure

The Life Cycle Service specification consists of these interfaces:

- **LifeCycleObject**
- **FactoryFinder**
- **GenericFactory**
- Interfaces described in Appendix A, an addendum to the Life Cycle Service

A.1 Overview

This appendix contains the specification for the compound life cycle component of the Life Cycle Service. The compound life cycle specification depends on the Life Cycle Service for the definition of the client view of Life Cycle operations. Moreover, it extends the Life Cycle Service to support compound life cycle operations on graphs of related objects. In addition, the compound life cycle specification depends on the Relationship Service for the definition of object graphs.

The Life Cycle Service specification describes a client's view of object life cycle. It describes how a client can create, copy, move and remove objects in a distributed object system. To create objects, clients find factory objects and issue create requests on factories. To copy, move and remove objects, clients issue requests on target objects supporting the **LifeCycleObject** interface.

If the target object represents a simple object, that is an object that is not part of a graph of related objects, the target provides an implementation for each of the operations in the **LifeCycleObject** interface.

If, on the other hand, the target object uses the Relationship Service for representing relationships with other objects, additional services are available to implement the compound life cycle operations. The specification in this appendix describes those services.

A.2 Key Features

The compound life cycle specification:

- Addresses the issues of copying, moving and removing objects that are related to other objects. Depending on the semantics of the relationships, these life cycle operations are applied to:
 - the object, to the relationship and to the related objects

- the object and to the relationship
- the object
- Coordinates compound life cycle operations on graphs of related objects, thus relieving object developers from implementing compound operations.
- Illustrates a general model for applying compound operations to graphs of related objects. The Externalization Service also illustrates the model.

A.3 Service Structure

The specification in this appendix defines a service that applies a compound life cycle operation to a graph of related objects, given a starting node. Compound operations traverse a graph of related objects and apply the operation to the relevant nodes, roles and relationships of the graph. The service supports the **CosCompoundLifeCycle::Operations** interface. Implementations of the service depend on the **CosCompoundLifeCycle::Node** , **CosCompoundLifeCycle::Role** and **CosCompoundLifeCycle::Relationship** interfaces which are subtypes of the Node, Role and Relationship interfaces defined in the Relationship Service. The **CosCompoundLifeCycle::Node**, **CosCompoundLifeCycle::Role** and **CosCompoundLifeCycle::Relationship** interfaces add operations to copy, remove and move nodes, roles and relationships.

The Relationship Service defines interfaces for containment and reference relationships and their roles. This appendix defines interfaces that inherit those interfaces and the compound life cycle interfaces.

A.4 Interface Overview

Table A-1 and Table A-2 summarize the interfaces defined in the **CosCompoundLifeCycle** module. The **CosCompoundLifeCycle** module is described in detail in Section A.6, “The CosCompoundLifeCycle Module,” on page A-9

Table A-1 Interfaces Defined in the *CosCompoundLifeCycle* Module for initiating compound life cycle operations

Interface	Purpose
Operations	Defines compound life cycle operations on graphs of related objects.
OperationsFactory	Defines an operation to create an object that supports the Operations interface.

Table A-2 Interfaces defined in the *CosCompoundLifeCycle* module that are used by implementations of compound life cycle operations

Interface	Inherits	Purpose
Node	CosGraphs::Node	Defines life cycle operations on nodes in graphs of related objects.
Relationship	CosRelationships::Relationship	Defines life cycle operations on relationships.
Role	CosGraphs::Role	Defines life cycle operations on roles.
PropagationCriteria Factory		Creates an object that supports the CosGraphs::TraversalCriteria interface that uses relationship propagation values.

Table A-3 and Table A-4 summarize the interfaces that combine the specific relationships defined by the Relationship Service and the life cycle interfaces defined in this appendix.

Table A-3 Interfaces defined in the *CosLifeCycleContainment* module

Interface	Inherits	Purpose
Relationship	CosContainment::Containment and CosCompoundLifeCycle::Relationship	Combines both interfaces. No additional operations are defined.
ContainsRole	CosContainment::ContainsRole and CosCompoundLifeCycle::Role	Combines both interfaces. No additional operations are defined.
ContainedInRole	CosContainment::ContainedInRole and CosCompoundLifeCycle::Role	Combines both interfaces. No additional operations are defined.

Table A-4 Interfaces defined in the *CosLifeCycleReference* module

Interface	Inherits	Purpose
Relationship	CosContainment::Reference and CosCompoundLifeCycle::Relationship	Combines both interfaces. No additional operations are defined.
ReferencesRole	CosContainment::ReferencesRole and CosCompoundLifeCycle::Role	Combines both interfaces. No additional operations are defined.
ReferencedByRole	CosContainment::ReferencedByRole and CosCompoundLifeCycle::Role	Combines both interfaces. No additional operations are defined.

A.5 Compound Life Cycle Operations

The Life Cycle specification describes a client's view of object life cycle. It describes how a client can create, copy, move and remove objects in a distributed object system. To create objects, clients find factory objects and issue create requests on factories. To copy, move and remove objects, clients issue requests on target objects supporting the **LifeCycleObject** interface.

If the target object represents a simple object, that is an object that is not part of a graph of related objects, the target provides an implementation for each of the operations in the **LifeCycleObject** interface.

If the target participates as a node in a graph of related objects, the target can delegate the life cycle operation to a service that implements the compound life cycle operation. In particular, the target simply creates an object that supports the **CosCompoundLifeCycle::Operations** interface and issues the corresponding life cycle request on it. The compound life cycle operations expect a **CosCompoundLifeCycle::Node** object reference as a starting node. The target simply passes its **CosCompoundLifeCycle::Node** object reference as the starting node.

When the life cycle object has completed issuing compound life cycle requests, it simply issues the **destroy** request to destroy the compound operation.

Figure A-1 illustrates the target's delegation of the life cycle request to compound operation.

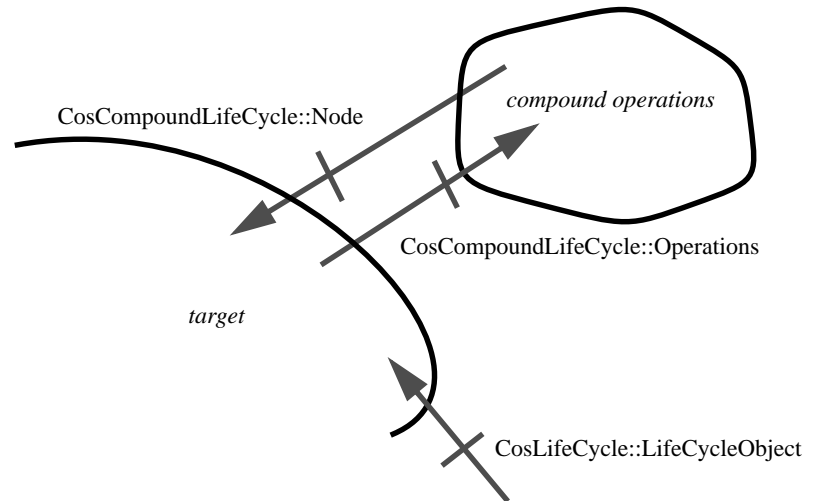


Figure A-1 A life cycle object that is part of a graph of related objects delegates the orderly operation on the graph to an object that implements the compound life cycle operation.

A.5.1 Applying the Copy Operation to the Example

We now use the example in the Relationship Service Specification to illustrate applying the **copy** operation to a graph. Figure A-2 on page A-6 illustrates the graph and the compound operation prior to applying the **copy** operation. Recall that the folder contains the document; the document is contained in the folder. The document contains the figure; the figure is contained in the document. The document contains the logo and the logo is contained in the document. On the other hand, the document references the book; the book is referenced by the document. Finally, the figure references the logo; the logo is referenced by the figure.




Figure A-2 Prior to applying copy to the graph.

In this example, the copy is performed in two passes. The first pass creates a list representation of the relevant edges of the graph. The second pass takes the list as input, copies the relevant nodes and roles, then creates all the necessary links by copying the relevant relationships.

A compound **copy** request is initiated by issuing a **LifeCycleObject::copy** request on the folder. Since the folder participates in a graph of related objects, it creates an object supporting the **CosCompoundLifeCycle::Operations** interface (the Operations object). Then the folder issues a **CosCompoundLifeCycle::Operations::copy** request on the Operations object, passing in its own **CosCompoundLifeCycle::Node** object reference as the starting node. The **copy** operation will copy the graph of related objects and return an object reference for the copy of the folder object.

The remainder of this section provides a description of how the Operations object might implement the **copy** operation.

First Pass of the Compound Copy Operation

The first pass consists of creating a list representation of the relevant edges of the graph. The Operations object uses an object supporting the **CosGraphs::Traversal** interface to do most of the work.

The Operations object creates an object supporting the **CosGraphs::TraversalCriteria** interface by calling **CosCompoundLifeCycle::PropagationCriteriaFactory::create**.

The Operations object then creates a **CosGraphs::Traversal** object by calling **CosGraphs::TraversalFactory::create_traversal_on**, passing in the object supporting the **CosGraphs::TraversalCriteria** interface. Calls on the **CosGraphs::Traversal** object yield an unordered list of **CosGraphs::Traversal::ScopedEdges** containing the following information.

(folder, ContainsRole, Containment, ContainedInRole, document)
(document, ReferencesRole, Reference, ReferencedByRole, book)
(document, ContainedInRole, Containment, ContainsRole, folder)
(document, ContainsRole, Containment, ContainedInRole, figure)
(document, ContainsRole, Containment, ContainedInRole, logo)
(figure, ReferencesRole, Reference, ReferencedByRole, logo)
(figure, ContainedInRole, Containment, ContainsRole, document)
(logo, ContainedInRole, Containment, ContainsRole, document)

This list will be referred to as the **OriginalEdgeList**.

Since the propagation value for copy from the document to the book is shallow, the traversal did not visit the book. As such, the edge:

(book, ReferencedByRole, Reference, References, document)

is not included. Although the traversal did visit the logo, the edge

(logo, ReferencedByRole, Reference, ReferencesRole, figure)

is not included because the propagation value for copy from the logo to the figure is none.

For more detailed information regarding the output of the **CosGraphs::Traversal** object with respect to the use of propagation semantics, see “Compound Operations” of the Relationship Service.

Second Pass of the Compound Copy Operation

The second pass copies all the relevant nodes and then relates them by copying the relevant relationships.

First, the set of nodes to be copied must be determined. This consists of all the distinct nodes in the left column of the **OriginalEdgeList**. Since a node may be involved in multiple edges, it may appear multiple times in the list; it should only be copied once. Each node in this set is copied by issuing a

CosCompoundLifeCycle::Node::copy_node request. This request will cause the node and all of its roles to be copied; the new node and its roles will be returned.

For each returned role of the copied node, an entry is made in a table of new roles. Each entry consists of:

- the role object is the data, and
- the node’s **CosGraphs::Traversal::TraversalScopedId** and the role’s **CORBA::InterfaceDef** together serve as a key.

The final step is to create all the relationships for the copied graph. All of the distinct relationships in the center column of the **OriginalEdgeList** need to be copied. Although a relationship may appear multiple times in the list, it should only be copied once. Each relationship is copied by issuing a **CosCompoundLifeCycle::Relationship::copy_relationship** request. The arguments to **CosCompoundLifeCycle::Relationship::copy_relationship** include the list of roles to be included in the new relationship. Some of these roles will be copies that were created as a result of processing deep propagation values; others will be roles in the original graph.

Thus, copy each unique relationship in the **OriginalEdgeList**, using **NamedRoles** as follows:

- For each role in an entry in the **OriginalEdgeList**, make a role key using the node's **TraversalScopedId** and the role's **CORBA::InterfaceDef** to search the table of new roles.
 - If the role was copied, the key will find the role's copy. The role's **RoleName** is obtained from the entry in the **OriginalEdgeList**. The role's copy and the **RoleName** are combined to form a **CosGraphs::NamedRole** which will then be included in the list of **CosGraphs::NamedRoles** passed to the **CosCompoundLifeCycle::Relationship::copy_relationship** method.
 - If no copy is found, the original **CosGraphs::NamedRole** is used instead.

Once all the Relationships have been copied, the **CosCompoundLifeCycle::Operations::copy** method is done.

Figure A-3 illustrates the result of applying copy to the graph, starting at the folder.




Figure A-3 The result of applying copy to the graph, starting at the folder.

When the **copy** operation propagates to a node because of a deep propagation value, other shallow propagation values to that node are promoted. That is, they are processed as if they were deep; relationships are formed with the copied node, not with the

original. This happened in the example; the shallow propagation value from the figure to the logo was promoted to deep because the logo was copied. As such, the new figure references the new logo, not the original logo.

A.6 The *CosCompoundLifeCycle* Module

The **CosCompoundLifeCycle** module defines:

- The **Operations** interface for initiating compound life cycle operations on graphs of related objects,
- **OperationsFactory** interface for creating compound operations,
- The **Node**, **Role**, **Relationship**, and **PropagationCriteriaFactory** interfaces for use by implementations of compound life cycle operations.

The **CosCompoundLifeCycle** module is given below. Detailed descriptions of the interfaces follow.

```
#include <CosLifeCycle.idl>
#include <CosRelationships.idl>
#include <CosGraphs.idl>

#pragma prefix "omg.org"

module CosCompoundLifeCycle {
    interface OperationsFactory;
    interface Operations;
    interface Node;
    interface Role;
    interface Relationship;
    interface PropagationCriteriaFactory;

    enum Operation {copy, move, remove};

    struct RelationshipHandle {
        CosRelationships::Relationship the_relationship;
        CosObjectIdentity::ObjectIdentifier constant_random_id;
    };

    interface OperationsFactory {
        Operations create_compound_operations();
    };

    interface Operations {
        Node copy (
            in Node starting_node,
            in CosLifeCycle::FactoryFinder there,
            in CosLifeCycle::Criteria the_criteria)
        raises (CosLifeCycle::NoFactory,
            CosLifeCycle::NotCopyable,
```

```
        CosLifeCycle::InvalidCriteria,
        CosLifeCycle::CannotMeetCriteria);
void move (
    in Node starting_node,
    in CosLifeCycle::FactoryFinder there,
    in CosLifeCycle::Criteria the_criteria)
    raises (CosLifeCycle::NoFactory,
           CosLifeCycle::NotMovable,
           CosLifeCycle::InvalidCriteria,
           CosLifeCycle::CannotMeetCriteria);
void remove (in Node starting_node)
    raises (CosLifeCycle::NotRemovable);
void destroy();
};

interface Node : CosGraphs::Node {
    exception NotLifeCycleObject {};
    void copy_node (
        in CosLifeCycle::FactoryFinder there,
        in CosLifeCycle::Criteria the_criteria,
        out Node new_node,
        out Roles roles_of_new_node)
        raises (CosLifeCycle::NoFactory,
               CosLifeCycle::NotCopyable,
               CosLifeCycle::InvalidCriteria,
               CosLifeCycle::CannotMeetCriteria);
    void move_node (
        in CosLifeCycle::FactoryFinder there,
        in CosLifeCycle::Criteria the_criteria)
        raises (CosLifeCycle::NoFactory,
               CosLifeCycle::NotMovable,
               CosLifeCycle::InvalidCriteria,
               CosLifeCycle::CannotMeetCriteria);
    void remove_node ()
        raises (CosLifeCycle::NotRemovable);
        CosLifeCycle::LifeCycleObject get_life_cycle_object()
        raises (NotLifeCycleObject);
};

interface Role : CosGraphs::Role {
    Role copy_role (
        in CosLifeCycle::FactoryFinder there,
        in CosLifeCycle::Criteria the_criteria)
        raises (CosLifeCycle::NoFactory,
               CosLifeCycle::NotCopyable,
               CosLifeCycle::InvalidCriteria,
               CosLifeCycle::CannotMeetCriteria);
    void move_role (
        in CosLifeCycle::FactoryFinder there,
        in CosLifeCycle::Criteria the_criteria)
        raises (CosLifeCycle::NoFactory,
```



```

        CosLifecycle::NotMovable,
        CosLifecycle::InvalidCriteria,
        CosLifecycle::CannotMeetCriteria);
CosGraphs::PropagationValue life_cycle_propagation (
    in Operation op,
    in RelationshipHandle rel,
    in CosRelationships::RoleName to_role_name,
    out boolean same_for_all);
};

interface Relationship : CosRelationships::Relationship {

    Relationship copy_relationship (
        in CosLifecycle::FactoryFinder there,
        in CosLifecycle::Criteria the_criteria,
        in CosGraphs::NamedRoles new_roles)
        raises (CosLifecycle::NoFactory,
            CosLifecycle::NotCopyable,
            CosLifecycle::InvalidCriteria,
            CosLifecycle::CannotMeetCriteria);
    void move_relationship (
        in CosLifecycle::FactoryFinder there,
        in CosLifecycle::Criteria the_criteria)
        raises (CosLifecycle::NoFactory,
            CosLifecycle::NotMovable,
            CosLifecycle::InvalidCriteria,
            CosLifecycle::CannotMeetCriteria);
    CosGraphs::PropagationValue life_cycle_propagation (
        in Operation op,
        in CosRelationships::RoleName from_role_name,
        in CosRelationships::RoleName to_role_name,
        out boolean same_for_all);
};

interface PropagationCriteriaFactory {
    CosGraphs::TraversalCriteria create(in Operation op);
};
};

```

A.6.1 The OperationsFactory Interface

Creating a Compound Life Cycle Operation

Operations create_compound_operations();

The **create_compound_operations** operation creates an object that implements the compound life cycle operations, that is, the factory creates and returns an object that supports the **CosCompoundLifeCycle::Operations** interface.

A.6.2 The Operations Interface

The **Operations** interface defines compound life cycle operations to copy, move and remove objects, given a starting node in a graph.

Applying the Copy Operation to a Graph of Related Objects

```
Node copy (  
  in Node starting_node,  
  in CosLifeCycle::FactoryFinder there,  
  in CosLifeCycle::Criteria the_criteria)  
raises (CosLifeCycle::NoFactory,  
  CosLifeCycle::NotCopyable,  
  CosLifeCycle::InvalidCriteria,  
  CosLifeCycle::CannotMeetCriteria);
```

The **copy** operation applies the copy operation to a graph of related objects. The starting node is provided as the **starting_node** parameter. The copy should be collocated with the factory finder given by the **there** parameter. The final parameter, **the_criteria**, allows unspecified values to be passed. This is explained in the Life Cycle specification in detail.

If a node, role or relationship in the graph refuses to be copied, the **NotCopyable** exception is raised with the node, role or relationship object reference returned as a parameter to the exception.

If appropriate factories to create a copies of the nodes and roles cannot be found, the **NoFactory** exception is raised. The exception value indicates the key used to find the factory.

In addition to the **NoFactory** and **NotCopyable** exceptions, implementations may raise standard CORBA exceptions. For example, if resources cannot be acquired for the copied graph, **NO_RESOURCES** will be raised.

It is implementation dependent whether this operation is atomic.

Applying the Move Operation to a Graph of Related Objects

```
void move (  
  in Node starting_node,  
  in CosLifeCycle::FactoryFinder there,  
  in CosLifeCycle::Criteria the_criteria)  
raises (CosLifeCycle::NoFactory,  
  CosLifeCycle::NotMovable,  
  CosLifeCycle::InvalidCriteria,  
  CosLifeCycle::CannotMeetCriteria);
```

The **move** operation applies the move operation to a graph of related objects. The starting node is provided as the **starting_node** parameter. The migrated graph should be collocated with the factory finder given by the **there** parameter. The final parameter, **the_criteria**, allows unspecified values to be passed. This is explained in the Life Cycle specification in detail.

If a node, role, or relationship in the graph refuses to be moved, the **NotMovable** exception is raised with the node, role, or relationship object reference returned as a parameter to the exception.

If appropriate factories to migrate the nodes and roles cannot be found, the **NoFactory** exception is raised. The exception value indicates the key used to find the factory.

In addition to the **NoFactory** and **NotMovable** exceptions, implementations may raise standard CORBA exceptions. For example, if resources cannot be acquired for the migrated graph, **NO_RESOURCES** will be raised.

It is implementation-dependent whether this operation is atomic.

Applying the Remove Operation to a Graph of Related Objects

**void remove (in Node starting_node)
raises (CosLifeCycle::NotRemovable);**

The remove operation applies the remove operation to a graph of related objects. The starting node is provided as the **starting_node** parameter.

If a node, role, or relationship in the graph refuses to be removed, the **NotRemovable** exception is raised with the node, role, or relationship object reference returned as a parameter to the exception.

It is implementation dependent whether this operation is atomic.

Destroying the Compound Operation

void destroy();

The **destroy** operation indicates to the compound operation that the client has completed operating on the graph. The compound operation object is destroyed.

A.6.3 The Node Interface

The **Node** interface defines operations to copy, move, and remove a node.

Copying a Node

**void copy_node (
in CosLifeCycle::Criteria the_criteria,
out Node new_node,**

```

out Roles roles_of_new_node)
raises (CosLifeCycle::NoFactory,
CosLifeCycle::NotCopyable,
CosLifeCycle::InvalidCriteria,
CosLifeCycle::CannotMeetCriteria);

```

The **copy** operation makes a copy of the node and its roles. The new node and roles should be collocated with the factory finder given by the **there** parameter. The final input parameter, **the_criteria**, allows unspecified values to be passed. This is explained in the Life Cycle specification in detail.

The result of a copy operation is a:

- node object reference for the new node, and
- sequence of roles.

Figure A-4 illustrates the result of a copy. A node, when it is born, is not in any relationships with other objects. That is, the roles in the new node are “disconnected”. It is the compound copy operation’s job to correctly establish new relationships.




Figure A-4 Copying a node returns the new object and the corresponding roles.

If the node or one of its roles refuses to be copied, the **NotCopyable** exception is raised with the node or role object reference returned as a parameter to the exception.

If an appropriate factory to create a copy cannot be found, the **NoFactory** exception is raised. The exception value indicates the key used to find the factory.

In addition to the **NoFactory** and **NotCopyable** exceptions, implementations may raise standard CORBA exceptions. For example, if resources cannot be acquired for the copied node, **NO_RESOURCES** will be raised.

Moving a Node

```

void move_node (
in CosLifeCycle::Criteria the_criteria)
raises (CosLifeCycle::NoFactory,
CosLifeCycle::NotMovable,
CosLifeCycle::InvalidCriteria,
CosLifeCycle::CannotMeetCriteria);

```

The **move** operation transfers some or all of the node's resources from "here" to "there." The move operation migrates a the node and its roles. The migrated node and roles should be collocated with the factory finder given by the **there** parameter. The final parameter, **the_criteria**, allows unspecified values to be passed. This is explained in the Life Cycle specification in detail.

If the node or one of its roles refuses to be moved, the **NotMovable** exception is raised with the node or role object reference returned as a parameter to the exception.

If an appropriate factory to support migration "over there" cannot be found, the **NoFactory** exception is raised. The exception value indicates the key used to find the factory.

In addition to the **NoFactory** and **NotMovable** exceptions, implementations may raise standard CORBA exceptions. For example, if resources cannot be acquired for the migrated node, **NO_RESOURCES** will be raised.

Removing a Node

```
void remove_node ()  
raises (CosLifeCycle::NotRemovable);
```

The **remove** operation removes the node and its roles.

If the node or one of its roles refuses to be removed, the **NotRemovable** exception is raised with the node or role object reference returned as a parameter to the exception.

Getting the Node's Life Cycle Object

```
CosLifeCycle::LifeCycleObject get_life_cycle_object()  
raises (NotLifeCycleObject);
```

Some nodes not only participate in the life cycle protocols for graphs of related objects but they also support the client's view of life cycle services. That is, the node also supports the **CosLifeCycle::LifeCycleObject** interface described in the Life Cycle Service specification. The **get_life_cycle_object** operation returns the **CosLifeCycle::LifeCycleObject** object reference for the node.

If the node does not support the **CosLifeCycle::LifeCycleObject** interface, the **NotLifeCycleObject** exception is raised.

A.6.4 The Role Interface

The **Role** interface defines operations to copy and move a role. (The **destroy** operation is defined by the base Relationship Service. As such, there is no need to define a **remove** operation.) The **Role** interface also defines an operation to return the propagation values for the **copy**, **move**, and **remove** operations.

The implementation of a **CompoundLifeCycle::Node** operation can call these operations on roles. For example, an implementation of copy on a node can call the **copy** operation on the Role.

Copying a Role

```
Role copy_role (  
    in CosLifeCycle::Criteria the_criteria)  
raises (CosLifeCycle::NoFactory,  
        CosLifeCycle::NotCopyable,  
        CosLifeCycle::InvalidCriteria,  
        CosLifeCycle::CannotMeetCriteria);
```

The **copy** operation makes a copy of the role. The new role should be collocated with the factory finder given by the `there` parameter. The final parameter, **the_criteria**, allows unspecified values to be passed. This is explained in the Life Cycle specification in detail.

The result of a **copy** operation is an object reference for the new object supporting the **Role** interface.

If the role refuses to be copied, the **NotCopyable** exception is raised with the role object reference returned as a parameter to the exception.

If an appropriate factory to create a copy cannot be found, the **NoFactory** exception is raised. The exception value indicates the key used to find the factory.

In addition to the **NoFactory** and **NotCopyable** exceptions, implementations may raise standard CORBA exceptions. For example, if resources cannot be acquired for the copied role, **NO_RESOURCES** will be raised.

Moving a Role

```
void move_role (  
    in CosLifeCycle::Criteria the_criteria)  
raises (CosLifeCycle::NoFactory,  
        CosLifeCycle::NotMovable,  
        CosLifeCycle::InvalidCriteria,  
        CosLifeCycle::CannotMeetCriteria);
```

The **move** operation transfers some or all of the role's resources. The **move** operation migrates the role. The migrated role should be collocated with the factory finder given by the **there** parameter. The final parameter, **the_criteria**, allows unspecified values to be passed. This is explained in the Life Cycle specification in detail.

If the role refuses to be moved, the **NotMovable** exception is raised with the role object reference returned as a parameter to the exception.

If an appropriate factory to support migration cannot be found, the **NoFactory** exception is raised. The exception value indicates the key used to find the factory.

In addition to the **NoFactory** and **NotMovable** exceptions, implementations may raise standard CORBA exceptions. For example, if resources cannot be acquired for the migrated role, **NO_RESOURCES** will be raised.

Getting a Propagation

```
CosGraphs::PropagationValue life_cycle_propagation (
    in Operation op,
    in RelationshipHandle rel,
    in CosRelationships::RoleName to_role_name,
    out boolean same_for_all);
```

Value

The **life_cycle_propagation** operation returns the propagation value to the role **to_role_name** for the life cycle operation **op** and the relationship **rel**. If the role can guarantee that the propagation value is the same for all relationships in which it participates, **same_for_all** is true.

A.6.5 The Relationship Interface

The Relationship interface defines operations to copy and move a relationship. (The **destroy** operation is defined by the Relationship Service. As such, there is no need to define a **remove** operation.) The Relationship interface also defines an operation to return the propagation values for the **copy**, **move** and **remove** operations.

Copying the Relationship

```
Relationship copy_relationship (
    in CosLifeCycle::FactoryFinder there,
    in CosLifeCycle::Criteria the_criteria,
    in CosGraphs::NamedRoles new_roles)
raises (CosLifeCycle::NoFactory,
    CosLifeCycle::NotCopyable,
    CosLifeCycle::InvalidCriteria,
    CosLifeCycle::CannotMeetCriteria);
```

The **copy** operation creates a new relationship. The new relationship should be collocated with the factory finder given by the **there** parameter. The second parameter, **the_criteria**, allows unspecified values to be passed. This is explained in the beginning of the Life Cycle specification in detail.

The values of the newly created relationship's attributes are defined by the implementation of this operation. However, the **named_roles** attribute of the newly created relationship must match **new_roles**. That is, the newly created relationship relates objects represented by **new_roles** parameter, not the by the original relationship's named roles.

The result of a **copy** operation is an object reference for the new object supporting the Relationship interface.

If the relationship refuses to be copied, the **NotCopyable** exception is raised with the relationship object reference returned as a parameter to the exception.

If an appropriate factory to create a copy cannot be found, the **NoFactory** exception is raised. The exception value indicates the key used to find the factory.

In addition to the **NoFactory** and **NotCopyable** exceptions, implementations may raise standard CORBA exceptions. For example, if resources cannot be acquired for the copied role, **NO_RESOURCES** will be raised.

Moving the Relationship

```
void move_relationship (  
  in CosLifeCycle::Criteria the_criteria)  
raises (CosLifeCycle::NoFactory,  
  CosLifeCycle::NotMovable,  
  CosLifeCycle::InvalidCriteria,  
  CosLifeCycle::CannotMeetCriteria);
```

The **move** operation transfers some or all of the relationship's resources. The move operation migrates the relationship. The migrated relationship should be collocated with the factory finder given by the **there** parameter. The final parameter, **the_criteria**, allows unspecified values to be passed. This is explained in the Life Cycle specification in detail.

If the relationship refuses to be moved, the **NotMovable** exception is raised with the relationship object reference returned as a parameter to the exception.

If an appropriate factory to support migration cannot be found, the **NoFactory** exception is raised. The exception value indicates the key used to find the factory.

In addition to the **NoFactory** and **NotMovable** exceptions, implementations may raise standard CORBA exceptions. For example, if resources cannot be acquired for the migrated relationship, **NO_RESOURCES** will be raised.

Getting a Propagation

```
CosGraphs::PropagationValue life_cycle_propagation (  
  in Operation op,  
  in CosRelationships::RoleName from_role_name,  
  in CosRelationships::RoleName to_role_name,  
  out boolean same_for_all);
```

Value

The **life_cycle_propagation** operation returns the relationship's propagation value from the role **from_role** to the role **to_role_name** for the life cycle operation **op**. If the role named by **from_role_name** can guarantee that the propagation value is the same for all relationships in which it participates, **same_for_all** is true.

A.6.6 *The PropagationCriteriaFactory Interface*

The **CosGraphs** module in the Relationship Service defines a general service for traversing a graph of related objects. The service accepts a “call-back” object supporting the `CosGraphs::TraversalCriteria` interface. Given a node, this object defines which edges to emit and which nodes to visit next.

The **PropagationCriteriaFactory** creates a **TraversalCriteria** object that determines which edges to emit and which nodes to visit based on propagation values for the compound life cycle operations.

Create a Traversal Criteria Based on Life Cycle Propagation Values

CosGraphs::TraversalCriteria create(in Operation op);

The **create** operation returns a **TraversalCriteria** object for an operation `op` that determines which edges to emit and which nodes to visit based on propagation values for `op`. For a more detailed discussion see Section A.5.1, “Applying the Copy Operation to the Example,” on page A-5 of this appendix and “Traversing Graphs of Related Objects” of the Relationship specification.

A.6.7 *Specific Life Cycle Relationships*

The Relationship service defines two important relationships:

- Containment is a one-to-many relationship. A container can contain many containees; a containee is contained by one container. Containment is represented by a relationship with two roles: the **ContainsRole**, and the **ContainedInRole**.
- Reference, on the other hand, is a many-to-many relationship. An object can reference many objects; an object can be referenced by many objects. Reference is represented by a relationship with two roles: **ReferencesRole** and **ReferencedByRole**.

The compound life cycle specification adds life cycle semantics to these specific relationships. That is, it defines propagation values for containment and reference.

A.7 *The CosLifeCycleContainment Module*

The **CosLifeCycleContainment** module defines three interfaces

- the **Relationship** interface
- the **ContainsRole** interface and
- the **ContainedInRole** interface.

```
#include <CosContainment.idl>
#include <CosCompoundLifeCycle.idl>

# pragma prefix “omg.org”
```

```

module CosLifeCycleContainment {

interface Relationship :
    CosCompoundLifeCycle::Relationship,
    CosContainment::Relationship {};

interface ContainsRole :
    CosCompoundLifeCycle::Role,
    CosContainment::ContainsRole {};

interface ContainedInRole :
    CosCompoundLifeCycle::Role,
    CosContainment::ContainedInRole {};
};

```

The **CosLifeCycleContainment** module does not define new operations. It merely “mixes in” interfaces from the **CosCompoundLifeCycle** and **CosContainment** modules. Although it does not add any new operations, it refines the semantics of these attributes and operations:

RelationshipFactory attribute	value
relationship_type	<code>CosLifeCycleContainment::Relationship</code>
degree	2
named_role_types	<code>"ContainsRole", CosLifeCycleContainment::ContainsRole;</code> <code>"ContainedInRole", CosLifeCycleContainment::ContainedInRole</code>

The **CosRelationships::RelationshipFactory::create operation** will raise:

- **DegreeError** if the number of roles passed as arguments is not 2.
- **RoleTypeError** if the roles are not **CosLifeCycleContainment::ContainsRole** and **CosLifeCycleContainment::ContainedInRole**.
- **MaxCardinalityExceeded** if the **CosLifeCycleContainment::ContainedInRole** is already participating in a relationship.

RoleFactory attribute for ContainsRole	value
role_type	<code>CosLifeCycleContainment::ContainsRole</code>
maximum_cardinality	unbounded
minimum_cardinality	0
related_object_types	<code>CosCompoundLifeCycle::Node</code>

- The **CosRelationships::RoleFactory::create_role** operation will raise **RelatedObjectTypeError** if the related object passed as a parameter does not support the **CosCompoundLifeCycle::Node** interface.
- The **CosRelationships::RoleFactory::link_role** operation will raise **RelationshipTypeError** if the **rel** parameter does not conform to the **CosLifeCycleContainment::Relationship** interface..

RoleFactory attribute for ContainedInRole	value
role_type	CosLifeCycleContainment::ContainedInRole
maximum_cardinality	1
minimum_cardinality	1
related_object_types	CosCompoundLifeCycle::Node

- The **CosRelationships::RoleFactory::create_role** operation will raise **RelatedObjectTypeError** if the related object passed as a parameter does not support the **CosCompoundLifeCycle::Node** interface.
- The **CosRelationships::RoleFactory::link** operation will raise **RelationshipTypeError** if the **rel** parameter does not conform to the **CosLifeCycleContainment::Relationship** interface.
- The **CosRelationships::RoleFactory::link** operation will raise **MaxCardinalityExceeded** if it is already participating in a containment relationship.

The **CosLifeCycleContainment::ContainsRole::life_cycle_propagation** operation returns the following:

operation	ContainsRole to ContainedInRole
copy	deep
move	deep
remove	deep

The **CosLifeCycleContainment::ContainedInRole::life_cycle_propagation** operation returns the following::

operation	ContainedInRole to ContainsRole
copy	shallow
move	shallow
remove	shallow

A.8 The *CosLifeCycleReference* Module

The **CosLifeCycleReference** module defines three interfaces

- the **Relationship** interface,
- the **ReferencesRole** interface, and
- the **ReferencedByRole** interface.

```
#include <CosReference.idl>
#include <CosCompoundLifeCycle.idl>

# pragma prefix "omg.org"

module CosLifeCycleReference {

interface Relationship :
    CosCompoundLifeCycle::Relationship,
    CosReference::Relationship {};

interface ReferencesRole :
    CosCompoundLifeCycle::Role,
    CosReference::ReferencesRole {};

interface ReferencedByRole :
    CosCompoundLifeCycle::Role,
    CosReference::ReferencedByRole {};
};
```

The **CosLifeCycleReference** module does not define new operations. It merely “mixes in” interfaces from the **CosCompoundLifeCycle** and **CosReference** modules.

Although it does not add any new operations, it refines the semantics of these attributes and operations:

RelationshipFactory attribute	value
relationship_type	CosLifeCycleReference::Relationship
degree	2
named_role_types	"ReferencesRole", CosLifeCycleReference::ReferencesRole; "ReferencedByRole", CosLifeCycleReference::ReferencedByRole

The **CosRelationships::RelationshipFactory::create** operation will raise:

- **DegreeError** if the number of roles passed as arguments is not 2.

- **RoleTypeError** if the roles are not **CosReference::ReferencesRole** and **CosReference::ReferencedByRole**.

RoleFactory attribute for ReferencesRole	value
role_type	CosLifeCycleReference::ReferencesRole
maximum_cardinality	unbounded
minimum_cardinality	0
related_object_types	CosCompoundLifeCycle::Node

- The **CosRelationships::RoleFactory::create_role** operation will raise the **RelatedObjectTypeError** if the related object passed as a parameter does not support the **CosCompoundLifeCycle::Node** interface.
- The **CosRelationships::RoleFactory::link** operation will raise **RelationshipTypeError** if the **rel** parameter does not conform to the **CosLifeCycleReference::Relationship** interface.

RoleFactory attribute for ReferencedByRole	value
role_type	CosLifeCycleReference::ReferencedByRole
maximum_cardinality	unbounded
minimum_cardinality	0
related_object_types	CosCompoundLifeCycle::Node

- The **CosRelationships::RoleFactory::create_role** operation will raise the **RelatedObjectTypeError** if the related object passed as a parameter does not support the **CosCompoundLifeCycle::Node** interface.
- The **CosRelationships::RoleFactory::link** operation will raise **RelationshipTypeError** if the **rel** parameter does not conform to the **CosLifeCycleRelationship::Relationship** interface.

The **CosLifeCycleReference::ReferencesRole::life_cycle_propagation** operation returns the following:

operation	ReferencesRole to ReferencedByRole
copy	shallow
move	shallow
remove	shallow

The **CosLifeCycleReference::ReferencedByRole::life_cycle_propagation** operation returns the following::

operation	ReferencedByRole to ReferencesRole
copy	none
move	shallow
remove	shallow

- The **CosRelationships::RoleFactory::create_role** operation will raise the **RelatedObjectTypeError** if the related object passed as a parameter does not support the **CosCompoundLifeCycle::Node** interface.
- The **CosRelationships::RelationshipFactory::create** operation will raise **DegreeError** if the number of roles passed as arguments is not 2. It will raise **RoleTypeError** if the roles are not **CosLifeCycleReference::ReferencesRole** and **CosLifeCycleReference::ReferencedByRole**.

References

B

1. James Rumbaugh, "Controlling Propagation of Operations using Attributes on Relations." *OOPSLA 1988 Proceedings*, pg. 285-296
2. James Rumbaugh, Michael Blaha, William Premerlani, Frederick Eddy and William Lorenzen, "Object-oriented Modeling and Design." Prentice Hall, 1991.

Note – This appendix is not part of the Life Cycle Services specification. It sketches a mechanism for expressing filters. This appendix is included to provide an example of how a filter might be provided.

A factory represents a scope of resource allocation, which is the set of resources available to the factory. Whenever it receives a creation request, a factory will allocate resources according to any policies which are in operation.

Clearly, by choosing a particular factory upon which to issue a create request, a client is exerting some control over the allocation of resources. Therefore, a client can limit the scope of resource allocation, by issuing the request on a different factory which represents a smaller set of resources.

However, there are two problems with this. Firstly, the granularity of resources may be much smaller than the granularity represented by the factories in a system. For example, there are unlikely to be factories which represent individual disk segments.

Secondly, the client may wish to rule out the use of particular resources within a scope, but avoid having a general reduction in scope. For example, the client might not be concerned with which machine within a LAN an object is created on, providing it is not on machine X.

Both of these needs can be addressed by providing a filter. In the first case, the filter is relatively simple; it will simply limit the scope of resource allocation. In the second case, the filter will need to be more sophisticated.

This appendix describes one way of providing filters using properties and constraint expressions. These concepts appear in the development of Trading in the ISO/IEC/CCITT Open Distributed Processing standards. Service providers register their service with the Trader and use properties to describe the service offer. Potential clients may then use a constraint expressions to describe the requirements which service offers must satisfy.

Similarly, the life cycle service may define a number of properties to represent the different kinds of resources available within in a system and clients may use constraint expressions to place the restrictions upon the use of those resources.

Note – The Object Services Architecture identifies an Object Properties Service which enables an object to have a set of arbitrary named values associated with it. These are very similar to the concept of properties as used in Trading and in this appendix.

C.1 Resources as Properties

Resource properties are application and generic factory implementation dependent and it is beyond the scope of this specification to identify standard properties which all generic factory implementations will recognize. The properties described in this appendix are given as examples only. Table C-1 gives some examples of properties that might be supported by a generic factory.

Table C-1 Examples of properties supported by a generic factory

Property Name	Meaning
Host	Host name of the machine
Architecture	Machine architecture, e.g. “intel”, “sparc”
OSArchitecture	Operating system architecture e.g. “solaris”, “hpux”

C.2 Constraint Expressions

Constraints are expressed in a Constraint Language which provides a set of operators which allow arbitrarily complex expressions involving properties and potential values to be specified. A property lists *satisfies* a constraint if the constraint expression is true when evaluated with respect to the property list.

Constraint expressions are very flexible. For example, if a client has an object executing on a machine called ‘Host1’ and wishes to create another object which is *not* on the same machine, the client can specify the constraint “Host != ‘Host1’”.

The constraint expression described here works with properties for which the value can be a string, a number, or a set of values.

The constraint language consists of:

- comparative functions: ==, !=, >, >=, <, <=, in
- constructors: and, or, not
- property names
- numeric and string constants
- mathematical operators: +, -, *, /
- grouping operators: (,), [,]

The following precedence relations hold in the absence of parentheses, in the order of lowest to highest:

- + and -
- * and /
- or
- and
- not

The comparative operator **in** checks for the inclusion of a particular string constant in the list which is the value of a property.

C.3 BNF for Constraint Expressions

<ConstraintExpr>	:=	[<Expr>]
<Expr>	:=	<Expr> "or" <Expr> <Expr> "and" <Expr> "not" <Expr> "(" <Expr> ")" <SetExpr> <SetOp> <SetExpr> <StrExpr> <StrOp> <StrExpr> <NumExpr> <NumOp> <NumExpr> <NumExpr> "in" <SetExpr> <StrExpr> "in" <SetExpr>
<NumOp>	:=	"==" "!=" "<" "<=" ">" ">="
<StrOp>	:=	"==" "!="
<SetOp>	:=	"==" "!="
<NumExpr>	:=	<NumTerm> <NumExpr> "+" <NumTerm> <NumExpr> "-" <NumTerm>
<NumTerm>	:=	<NumFactor> <NumTerm> "*" <NumFactor> <NumTerm> "/" <NumFactor>
<NumFactor>	:=	<Identifier> <Number> "(" <NumExpr> ")" "-" <NumFactor>
<StrExpr>	:=	<StrTerm> <StrExpr> "+" <StrTerm>

<StrTerm>	:=	<Identifier> <String> "(" <StrExpr> ")"
<SetExpr>	:=	<SetTerm> <SetExpr> "+" <SetTerm>
<SetTerm>	:=	<Identifier> <Set> "(" <SetExpr> ")"
<Identifier>	:=	<Word>
<Number>	:=	<Integer> <Float>
<Integer>	:=	{ <Digit> }+
<Float>	:=	<Mantissa> [<Sign>] [<Exponent>]
<Mantissa>	:=	<Integer> ["." [<Integer>]] "." <Integer>
<Sign>	:=	"_" "+"
<Exponent>	:=	"e" <Integer> "E" <Integer>
<Word>	:=	<Letter> { <AlphaNum> }*
<AlphaNum>	:=	<Letter> <Digit> "_"
<String>	:=	"" { <Char> }* ""
<Char>	:=	<Letter> <Digit> <Other>
<Set>	:=	"{" <Elements> "}"
<Elements>	:=	[<Element> { <Sp>+ <Element> }*]

```

<Element>      :=      <Number>
                  |
                  |
                  |      <Word>
                  |
                  |      <String>

<Letter>       :=      a|b|c|d|e|f|g|h|i|j|k
                  |
                  |      l|m|n|o|p|q|r|s|t|u|v
                  |
                  |      w|x|y|z|A|B|C|D|E|F|G
                  |
                  |      H|I|J|K|L|M|N|O|P|Q|R
                  |
                  |      S|T|U|V|W|X|Y|Z

<Digit>        :=      0|1|2|3|4|5|6|7|8|9

<Other>        :=      <Sp>|~|!|@|#|$|%|^|&|*|(
                  |
                  |      )|-|_|=|+|[{]|}|;|:
                  |
                  |      “|\|||,|<|.|>|/|?

<Sp>           :=      ” ”

```


Administration

D

Note – This appendix is not part of the Life Cycle Services specification. This description is included as a suggested way of administering generic factories.

The specification for the life cycle service includes the **GenericFactory** interface. There will be at least two styles of object which support that interface:

- implementation specific factories that actually assemble the resources for a new object, and
- generic factories which pass requests on to either implementation specific factories or other generic factories.

By configuring generic factories and implementation specific factories into a graph, a creation service can be built which administers the allocation of a large number of resources and can use them to create a wide variety of objects.

To ensure that the creation service is scalable, it is essential that the principle of federation is adopted – each component retains its autonomy rather than becoming subordinate to another.

Whenever the creation service receives a creation request, the request will need to traverse the graph until it reaches an implementation specific factory which can satisfy the request. As the request traverses the graph, each non-terminal node in the graph (i.e., the generic factories) will decide which link the request will traverse next. Decisions will be based upon information about each available link, any policies in force at that node and, of course, the actual request.

Clearly, the configuration and policies of such a creation service will need to be administered. However, the specification does not include the specification of an administration interface. This is because the principle of federation is not only important to the life cycle service. It will be essential to a number of other services,

notably trading, and the OMG plans to address the issue of federation for all object services, rather than making a premature specification addressing the needs of just one service.

The remainder of this appendix describes the principle of federation in more detail, outlines the use of policies and preferences to support federation, and then concludes with a suggestion for how an administration interface might look.

D.1 Federation

Federation is essential in large-scale distributed systems where the existence of centralized ownership and universal control cannot be assumed. In these systems the only way to achieve cooperation between autonomous systems without creating a hierarchical structure is to use federation. Federation is also beneficial to smaller systems which can exploit the high degree of flexibility which federation provides.

Federation differs from the more conventional approach of adopting a strictly hierarchical organization in a number of ways. Firstly, components can provide their service to any number of others, not just the single component which is its “parent” in the hierarchy. Secondly, components can establish peer-to-peer relationships, eliminating the need for a single component at the top of the hierarchy. Finally, this approach avoids the necessity of maintaining a global namespace. Instead, all names are relative to the context in which they are used.

Federation enables previously distinct systems to be unified without requiring global changes to their naming structures and system management hierarchies. The administration functions must ensure the systems are configured appropriately (e.g., avoiding circular references in those graphs which must be kept acyclic).

D.1.1 Federation in Object Services

In addition to the use of federation in configuring generic factories, federation is also applicable to a number of other services.

Trading is a notable example. A global offer space is neither practical nor desirable. Consequently, there will be multiple traders, each representing a different portion of the offer space. Offers held by one trader can be made available to the clients of another trader through federation.

The naming service specification also demonstrates attributes of federation. Naming contexts can be bound to other naming contexts and requests for name resolution can be passed across the links. However, it is entirely the concern of the naming context how it resolves the name within its domain (i.e., it is autonomous).

D.1.2 Federation Issues

There are a number of issues which need to be addressed for federation to be used in a cohesive fashion across all object services.

Visibility of the Federation Graph

The naming service makes the configuration of naming contexts into a graph very visible to the clients. This is essential, because the naming service must provide clients with a structured namespace.

On the other hand, it is not clear that a client should ever be able to see the internal structure of a life cycle creation service built with generic and implementation specific factories.

The trading service falls in between the two extremes. It may be useful for a client to be able to navigate the structure of a trading service graph in order to have more control over the visibility of offers. However, this may make clients too dependent upon the organization of the trading service and limit the flexibility of the system administrator in reorganizing the trading service to provide the most effective service.

Service Interface vs. Administration Interface

In general, it is desirable to federate using the service interface for the links and reserve the administration interface for the administrators. This approach ensures that autonomy is retained. However, this precludes the use of compound names in the administration functions because the administration functions cannot traverse the graph; only simple names can be used in administration only functions.

However, this is inappropriate for services where graph manipulation is an essential part of the service. For example, the naming service specification does not distinguish between administration functions for manipulating the graph and service functions. This is clearly correct; the clients need to be able to manipulate the graph by creating, binding and destroying contexts.

Multiple Service Interfaces

A node in a federation graph may be a conspiracy and offer multiple service interfaces, perhaps one for each point it is bound into the graph. However, for services where the administration is kept distinct from the service, it is likely that the conspiracy will support only one administration interface.

In these situations, it becomes necessary for an administrator to be able to match service interfaces to conspiracies, i.e. to match one or more service interfaces to an administrative interface. The example in Section D.2, “An Example LifeCycleService Module,” on page D-5 provides a solution to this which, in theory, will scale, but there may be better ways of doing this.

Cycles and Peer-to-Peer Relationships

The introduction of cycles into a federation graph is a contentious issue. Since peer-to-peer relationships are a degenerate form of cycle, any service which supports peer-to-peer relationships must be capable of handling cycles. The major impact of this is to provide loop detection on operations which would otherwise go out of control. Both trading and naming services are examples of this kind of service.

However, some services may not be able to handle cycles effectively and will wish to proscribe them. This probably covers peer-to-peer relationships, although that might be an acceptable special case. An example of this might be the life cycle creation service, where information about the current usage of the available resources must percolate up the graph in order to make informed decisions, but the introduction of cycles would make this information unclear or even meaningless.

Policies

It is frequently necessary to configure the way in which operations are performed in order to tune the performance, e.g how long a search operation may take, how many matches can be returned, or how much memory to use for a cache.

The same problems exist in distributed systems except that such configuration parameters must be explicitly passed around. Where different administrative domains are connected, such configuration parameters cannot be enforced by one domain on the other. Similarly, users may want to control the configuration but must be prevented from hogging resources (e.g., memory, disk space, etc.). Some configuration elements must be enforced (e.g., disk quotas), some elements may specify defaults which can be changed, and some elements may be requests which may or may not clash with hard limits (e.g., max memory per process).

Policies are used as a generic solution to this problem – wherever some kind of choice needs to be made, policies may be used to guide the decision making process.

Table D-1 provides some examples of policies. which a federated service might support.

Table D-1 Example policies

Policy Name	Meaning
search_algorithm	determines whether the federation graph should be traversed in a depth first or breadth first fashion.
cross_boundaries	determines whether administrative boundaries should be crossed.
maximum_distance	how far to traverse a graph before failing a request.

When invoking operations, clients can specify preferences for particular policies. Providing the service has no value set for that policy, the preference will be simply added to the policy list for the duration of the request. However, if a service policy is already specified then the preference will either be ignored or, for policies such as “maximum_distance”, the more constraining value will be adopted.

As a request traverses a graph, each node will pass its current policy set as preferences. In this way, the autonomy of individual administrative domains is preserved.

When an object doesn’t implement all choices of a policy, it should not allow its policy to be modified to an unsupported value. This means that implementation limitations are handled as Administrative hard limits which provides the correct semantics.

Where no policy is specified by either administrator or client, the implementation determines its own behavior. However, this decision would not be propagated through the graph (as a preference), leaving it to each node in the graph to make its own decision.

D.2 An Example *LifeCycleService* Module

Administrators access the administration functions via the *LifeCycleService* module, which defines the *LifeCycleServiceAdmin* interface. This example is intended to work with the *GenericFactory* interface in the specification. As a result, the administration functions cannot make use of compound names.

Note – This is only an example of the *LifeCycle* service and is NOT part of the *LifeCycle* service.

```

#include <CosLifeCycle.idl>

#pragma prefix "omg.org"

module LifeCycleService {

    typedef sequence <CosLifeCycle::NameValuePair> PolicyList;
    typedef sequence <CosLifeCycle::Key> Keys;
    typedef sequence <CosLifeCycle::NameValuePair> PropertyList;
    typedef sequence <CosNaming::NameComponent> NameComponents;

    interface LifeCycleServiceAdmin {

        attribute PolicyList policies;

        void bind_generic_factory(
            in CosLifeCycle::GenericFactory gf,
            in CosNaming::NameComponent name,
            in Keys key_set,
            in PropertyList other_properties)
            raises (CosNaming::NamingContext::AlreadyBound,
                CosNaming::NamingContext::InvalidName);

        void unbind_generic_factory(
            in CosNaming::NameComponent name)
            raises (CosNaming::NamingContext::NotFound,
                CosNaming::NamingContext::InvalidName);

        CosLifeCycle::GenericFactory resolve_generic_factory(
            in CosNaming::NameComponent name)
            raises (CosNaming::NamingContext::NotFound,
                CosNaming::NamingContext::InvalidName);

        NameComponents list_generic_factories();
    }
}

```

```
boolean match_service (in CosLifeCycle::GenericFactory f);

string get_hint();

void get_link_properties(
    in CosNaming::NameComponent name,
    out Keys key_set,
    out PropertyList other_properties)
    raises (CosNaming::NamingContext::NotFound,
           CosNaming::NamingContext::InvalidName);
};

};
```

D.2.1 The *LifeCycleServiceAdmin* Interface

The **LifeCycleServiceAdmin** interface provides the basic administration operations required to enable the lifecycle service to be administered by a set of tools or an administration service. The operations enable configuration of factories supporting the **GenericFactory** interface into a graph and setting of policies for those factories.

bind_generic_factory

```
void bind_generic_factory(
    in CosLifeCycle::GenericFactory gf,
    in CosNaming::NameComponent name,
    in Keys key_set,
    in PropertyList other_properties)
    raises (CosNaming::NamingContext::AlreadyBound,
           CosNaming::NamingContext::InvalidName);
```

This operation binds a factory supporting the **GenericFactory** interface into a graph. The name must be unique within the context of the target of the operation. From then on, that factory can be identified by that name.

In order to make a good decision about which link to choose for a request, the node needs to be provided with additional information about those factories. This information may be fairly dynamic (e.g., the current usage of the resources available through the link), or more static (e.g., the Keys for which the link can provide support).

The **key_set** parameter is a list of the keys for which the factory can provide support. In the case of an implementation specific factory, this list will often only have one member.

The **other_properties** parameter can be used to provide other static properties associated with the factory. For example, an “Architectures” property would indicate the type(s) of machine which the factory could create objects on.

Changes to the static information as well as more dynamic information can be monitored through the Events service. Each factory would generate events whenever the information changed significantly (e.g., a new **GenericFactory** interface with new keys is bound to the factory, or there is a change in the usage of resources available to the factory) and these can then be passed to those factories which need to know.

unbind_generic_factory

```
void unbind_generic_factory(  
    in CosNaming::NameComponent name)  
    raises (CosNaming::NamingContext::NotFound,  
           CosNaming::NamingContext::InvalidName);
```

This operation unbinds the generic factory identified by the name.

resolve_generic_factory

```
CosLifecycle::GenericFactory resolve_generic_factory(  
    in CosNaming::NameComponent name)  
    raises (CosNaming::NamingContext::NotFound,  
           CosNaming::NamingContext::InvalidName);
```

This operation takes the name supplied and returns the reference to the **GenericFactory** object.

list_generic_factories

```
NameComponents list_generic_factories();
```

This operation returns a list of the names of all the bound factories.

match_service

```
boolean match_service (in CosLifecycle::GenericFactory f);
```

This operation returns `true` if the generic factory interface is supported by the target.

get_hint

```
string get_hint();
```

This operation returns a hint associated with the target, see “Building a Map of a Graph” on page D-8.

get_link_properties

```
void get_link_properties(  
    in CosNaming::NameComponent name,  
    out Keys key_set,
```

```
out PropertyList other_properties)
raises (CosNaming::NamingContext::NotFound,
        CosNaming::NamingContext::InvalidName);
```

This operation returns the **key_set** and **other_properties** associated with the name.

Building a Map of a Graph

Administration tools may wish to build a map of a federation graph from scratch and some of the operations above are provided for that purpose.

First of all, the tool must obtain the set of administration interfaces for all the factories to be administered. These might be obtained from a number of sources (e.g., a well-known trading context).

For each interface, the **list_generic_factories** operation obtains a list of all the links for each node. Using **resolve_generic_factory**, a service interface can be obtained for each link. These can then be matched to an administration interface using **match_service**.

Clearly, this does not scale well if there are many nodes involved because of the average number of invocations of **match_service** required. This problem can be solved if one of the **other_properties** associated with each service interface is a *hint* and a hint is available for each administration interface. If the hints are the same, there may be a match and **match_service** is called to check. If the hints could be guaranteed to be unambiguous, the invocation could be avoided altogether, but this requires a global namespace for the hints. The best that can reasonably be achieved is to reduce the chance of a clash to a minimum.

The **get_hint** and **get_link_properties** can be used for this purpose.

Supports for PCTE Objects

E

Note – This appendix is not part of the Life Cycle Services specification. This appendix defines a set of criteria¹ suitable for supporting PCTE objects.

It is intended that objects in a PCTE repository be among those objects that can be managed through this lifecycle interface. It is reasonable to expect that applications written for PCTE will use the PCTE APIs to manage the life-cycle of PCTE objects. It is also reasonable to expect that clients not specifically written for relationship-oriented objects will not be able to manipulate the life-cycles of PCTE objects. However, between these two, one can envision clients which desire to be flexible, working on objects which may or may not be stored in the PCTE repository. One can also envision object factories, constructed to make use of PCTE which provide services to clients that are not PCTE applications because they do not have the appropriate working schemas, etc.

Support for these clients employs a series of conventional interpretations of the lifecycle operations. This appendix provides one such set of conventions to demonstrate the feasibility of the use of these interfaces in a context supporting PCTE.

Object references appear in constraint expressions in the form of character strings. Any implementation of PCTE as a CORBA Object Adapter has to establish a relationship between these and the corresponding CORBA types, and be able to convert between them.

1. As defined in Section 2.1.4, "Criteria," on page 2-8 of the life cycle specification.

E.1 Overview

A PCTE repository can be viewed as a generic factory. Using whatever naming or trading services are appropriate, a client wishing to use the PCTE factory obtains an object reference to it. To support the simple applications intending to operate within the context of a single PCTE repository, the PCTE factory supports the operations defined by both the **GenericFactory** and **FactoryFinder** interfaces. The client can then invoke the PCTE factory's **create_object** operation, or pass the factory as the "factory finder" when invoking the move or copy operations to move or copy within the same PCTE repository. These clients include the servers implementing the move and copy operations for various PCTE objects as well.

Lifecycle creation, copy, and move operations are influenced by a sequence of criteria. Criteria are specified as a sequence of name/value pairs. Certain criteria are of interest to the PCTE factories:

"logical location"

The logical location is used to express the logical connection information that must be specified when creating or copying a PCTE object. Logical location is a sequence of name/value pairs expressing a connection for the object. The PCTE factory supports and requires two logical locations:

- **ORIGIN**: A string representation of the reference to the object to which the newly created object is to be connected.
- **ORIGINLINK**: The name of the origin object's link which is to hold the link from the origin object to the newly created object.

"filter"

The filter is used to express the fact that an object being created, copied, or moved should reside on the same volume as some other, nearby, object. A filter is an expression as described in Section C.3, "BNF for Constraint Expressions," on page C-3. For PCTE, the term "NEAR=" followed by an object reference to the designated nearby object indicates that the new object is to be located at least as near as the same volume to the specified object.

"authorization"

Please note that no proposal on authorization has yet been accepted by OMG; however, this lifecycle criterion is required to create PCTE objects.

E.2 Object Creation

The **LifeCycle::GenericFactory::create_object** operation in this specification is borne by factory objects. It has two parameters:

1. a key used to identify the desired object to be created and
2. a set of criteria expressed in an NVP-list.

The corresponding PCTE operation is called OBJECT_CREATE. The parameters to OBJECT_CREATE are obtained from the **LifeCycle::GenericFactory::create_object** parameters.

The PCTE operation OBJECT_CREATE has six parameters:

1. the type of object to be created This is the “key” from LifeCycle **create_object**.
2. the origin object of the relation anchoring the new object This is the object identified as the named “ORIGIN” of the logical location criterion.
3. the name of the link from that origin object to the new object This is the string identified as the named “ORIGINLINK” of the logical location criterion.
4. an optional key for that link. This is the string identified as the named “LINKKEY” of the initialization criteria.
5. an object near whose location the object is to be created This is the string value of a required filter expression value by the qualifier “NEAR.”
6. an access mask. This is the string identified as the named “ACCESS” of the authorization criteria. This string is a simple mapping of the granted and denied access rights.

Exceptions raised by PCTE are mapped to suitable LifeCycle exceptions.

E.3 Object Deletion

The **LifeCycle::LifeCycleObject::remove** operation in this specification is borne by all life-cycle objects. It has no parameters.

The corresponding PCTE operation is called OBJECT_DELETE. The parameters to OBJECT_DELETE are obtained from the object to be deleted using information about that object defined in PCTE’s schema information about the object.

The PCTE operation OBJECT_DELETE has two parameters:

1. the origin object of a relation anchoring the object to be deleted, and
2. the name of the link from that origin object to the object to be deleted.

To both ensure that the controlling object is actually deleted and maintain the PCTE referential integrity constraints the following steps are performed for each reversible link emanating from the controlling object:

1. Determine the object, o, that the link refers to.
2. Determine the name, r&prime., of the reverse link back from o.
3. Perform PCTE OBJECT_DELETE(o, r&prime.)

The objective is accomplished when all outgoing, reversible links have been dealt with thus, or before that if one of the OBJECT_DELETE calls fails because the object has already been deleted.

Exceptions raised by PCTE are mapped to suitable LifeCycle exceptions.

E.4 Object Copying

The **LifeCycle::LifeCycleObject::copy** operation in this specification is borne by all life-cycle objects. It has two parameters:

1. a factory-finder to assist in locating a factory that provides resources for the copied object, and
2. a set of criteria expressed in an NVP-list

The corresponding PCTE operation is called OBJECT_COPY. Some of the parameters to OBJECT_COPY can be obtained directly from the LifeCycle copy parameters. Other required information is obtained from the constraint expression parameter of the LifeCycle copy.

The PCTE operation OBJECT_COPY has six parameters:

1. The object to be copied. This is the bearer object of LifeCycle copy operation.
2. The origin object of the relation anchoring the new object. This is the object identified as the named "ORIGIN" of the logical location criterion.
3. The name of the link from that origin object to the new object. This is the string identified as the named "ORIGINLINK" of the logical location criterion.
4. An optional key for that link. This is the string identified as the named "LINKKEY" of the initialization criteria.
5. An object near whose location the object is to be created. This is the string value of a required filter expression value by the qualifier "NEAR."
6. An access mask. This is the string identified as the named "ACCESS" of the authorization criteria This string is a simple mapping of the granted and denied access rights.

The semantics of the copy operation corresponds to the PCTE OBJECT_COPY semantics. They are based upon details of the object types involved, including which attributes, links and destination objects are "duplicable."

Exceptions raised by PCTE are mapped to suitable CORBA standard exceptions.

E.5 Object Moving

The **LifeCycle::LifeCycleObject::move** operation in this specification is borne by all life-cycle objects. It has two parameters:

1. A factory-finder to assist in locating a factory that provide resources for the moved object, and
2. a set of criteria expressed in an NVP-list.

The corresponding PCTE operation is called OBJECT_MOVE. The parameters to OBJECT_MOVE can be obtained directly from the LifeCycle copy parameters or from defaults.

The PCTE operation OBJECT_MOVE has three parameters:

1. The object to be copied. This is the bearer object of LifeCycle move operation.
2. An object near whose location the object is to be created. This is the string value of a required filter expression value by the qualifier "NEAR."
3. Scope - whether to move the object itself or the object and all its components. This will be defaulted to ATOMIC.

