



THE ENTERPRISE MIDDLEWARE SOLUTION

BEA WebLogic Enterprise

C++ Programming Reference

BEA WebLogic Enterprise 4.2
Document Edition 4.2
July 1999

Copyright

Copyright © 1999 BEA Systems, Inc. All Rights Reserved.

Restricted Rights Legend

This software and documentation is subject to and made available only pursuant to the terms of the BEA Systems License Agreement and may be used or copied only in accordance with the terms of that agreement. It is against the law to copy the software except as specifically allowed in the agreement. This document may not, in whole or in part, be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine-readable form without prior consent, in writing, from BEA Systems, Inc.

Use, duplication or disclosure by the U.S. Government is subject to restrictions set forth in the BEA Systems License Agreement and in subparagraph (c)(1) of the Commercial Computer Software-Restricted Rights Clause at FAR 52.227-19; subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013, subparagraph (d) of the Commercial Computer Software--Licensing clause at NASA FAR supplement 16-52.227-86; or their equivalent.

Information in this document is subject to change without notice and does not represent a commitment on the part of BEA Systems. THE SOFTWARE AND DOCUMENTATION ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. FURTHER, BEA Systems DOES NOT WARRANT, GUARANTEE, OR MAKE ANY REPRESENTATIONS REGARDING THE USE, OR THE RESULTS OF THE USE, OF THE SOFTWARE OR WRITTEN MATERIAL IN TERMS OF CORRECTNESS, ACCURACY, RELIABILITY, OR OTHERWISE.

Trademarks or Service Marks

BEA, ObjectBroker, TOP END, and TUXEDO are registered trademarks of BEA Systems, Inc. BEA Builder, BEA Connect, BEA Manager, BEA MessageQ, Jolt, M3, and WebLogic are trademarks of BEA Systems, Inc.

All other company names may be trademarks of the respective companies with which they are associated.

C++ Programming Reference

Document Edition	Date	Software Version
4.2	July 1999	BEA WebLogic Enterprise 4.2

Contents

Preface

Purpose of This Document	xvii
How to Use This Document	xviii
Related Documentation	xxi
Contact Information.....	xxiv

1. OMG IDL Syntax

OMG IDL Extensions.....	1-2
-------------------------	-----

2. Implementation Configuration File (ICF)

ICF Syntax.....	2-2
Sample ICF File.....	2-3
Creating the ICF File	2-4

3. TP Framework

Definition of Terms	3-2
A Simple Programming Model	3-5
Control Flow	3-6
Object State Management	3-7
Transaction Integration	3-7
Object Housekeeping	3-7
High-level Services	3-8
State Management	3-8
Activation Policy.....	3-9
Application-controlled Activation and Deactivation	3-11
Servant Lifetime	3-14
Saving and Restoring Object State.....	3-16

Transactions	3-16
Transaction Policies	3-17
Transaction Initiation	3-18
Transaction Termination	3-18
Transaction Suspend and Resume	3-19
Restrictions on Transactions	3-20
SQL and Global Transactions	3-21
Voting on Transaction Outcome	3-22
Transaction Time-outs	3-23
TP Framework API	3-23
Server Interface	3-24
Server::create_servant	3-25
Server::initialize()	3-28
Server::release()	3-31
Tobj_ServantBase Interface	3-33
Tobj_ServantBase:: activate_object()	3-35
Tobj_ServantBase::deactivate_object()	3-38
TP Interface	3-44
TP::application_responsibility	3-46
TP::bootstrap()	3-47
TP::close_xa_rm()	3-48
TP::create_active_object_reference()	3-50
TP::create_object_reference()	3-53
TP::deactivateEnable	3-56
TP::get_object_id ()	3-58
TP::get_object_reference()	3-59
TP::open_xa_rm()	3-60
TP::orb()	3-62
TP::register_factory()	3-63
TP::unregister_factory()	3-65
TP::userlog()	3-67
CosTransactions::TransactionalObject Interface Not Enforced	3-68
Error Conditions, Exceptions, and Error Messages	3-69
Exceptions Raised by the TP Framework	3-69
Exceptions in the Server Application Code	3-69

Exceptions and Transactions	3-70
Restriction of Nested Calls on Corba Objects.....	3-70

4. Bootstrap Object

Why Bootstrap Objects Are Needed	4-1
How Bootstrap Objects Work	4-1
Types of Remote Clients Supported.....	4-4
Capabilities and Limitations.....	4-5
Bootstrap Object API	4-6
Tobj Module	4-7
C++ Mapping	4-8
Java Mapping	4-8
Mappings for Microsoft Desktop Clients	4-9
C++ Member Functions and Java Methods.....	4-10
Tobj_Bootstrap.....	4-11
Tobj_Bootstrap::register_callback_port.....	4-14
Tobj_Bootstrap::resolve_initial_references	4-15
Tobj_Bootstrap::destroy_current().....	4-16
Automation Methods	4-17
Initialize	4-18
CreateObject.....	4-20
DestroyCurrent.....	4-22
Programming Examples	4-23
Java Client Example: Getting a SecurityCurrent Object.....	4-23
Visual Basic Client Example: Using the Bootstrap Object.....	4-25

5. FactoryFinder Interface

Capabilities, Limitations, and Requirements.....	5-2
Functional Description	5-3
Locating a FactoryFinder	5-3
Registering a Factory	5-4
Locating a Factory.....	5-5
Creating Application Factory Keys.....	5-11
C++ Member Functions and Java Methods.....	5-19
CosLifeCycle::FactoryFinder::find_factories	5-20

Tobj::FactoryFinder::find_one_factory.....	5-22
Tobj::FactoryFinder::find_one_factory_by_id.....	5-24
Tobj::FactoryFinder::find_factories_by_id.....	5-26
Tobj::Factoryfinder::list_factories	5-28
Automation Methods	5-29
DITobj_FactoryFinder.find_one_factory.....	5-30
DITobj_FactoryFinder.find_one_factory_by_id.....	5-32
DITobj_FactoryFinder.find_factories_by_id	5-34
DITobj_FactoryFinder.find_factories	5-36
DITobj_FactoryFinder.list_factories.....	5-37
Programming Examples	5-38
Using the FactoryFinder Object	5-38
Using WLE Extensions to the FactoryFinder Object	5-40

6. Security Service

Capabilities and Limitations.....	6-2
Getting Initial References to the SecurityCurrent Object.....	6-2
Basic Security-level Requirements for WLE Clients	6-3
Functional Components.....	6-4
Security Model	6-4
Authentication of Principals.....	6-4
Controlling Access to Objects.....	6-5
Administrative Control	6-5
Security Model Functional Description.....	6-6
Description	6-6
Authentication	6-8
Description of Authentication	6-8
Client Security API.....	6-17
CORBA Module.....	6-17
TimeBase Module	6-18
Security Module	6-20
Security Level 1 Module	6-22
Security Level 2 Module	6-22
Tobj Module	6-24
Method Descriptions.....	6-25

SecurityLevel1::Current::get_attributes	6-26
SecurityLevel2::PrincipalAuthenticator::authenticate	6-27
SecurityLevel2::PrincipalAuthenticator::continue_authentication ...	6-29
SecurityLevel2::Credentials::get_attributes	6-30
SecurityLevel2::Credentials::is_valid	6-31
SecurityLevel2::Current::set_credentials	6-32
SecurityLevel2::Current::get_credentials	6-33
SecurityLevel2::Current::principal_authenticator	6-34
Tobj::PrincipalAuthenticator::get_auth_type	6-35
Tobj::PrincipalAuthenticator::logon	6-36
Tobj::PrincipalAuthenticator::logoff	6-38
Tobj::PrincipalAuthenticator::build_auth_data	6-39
Automation Method Descriptions	6-40
DISecurityLevel2_Current	6-40
DISecurityLevel2_Current.get_attributes	6-41
DISecurityLevel2_Current.set_credentials	6-42
DISecurityLevel2_Current.get_credentials	6-43
DISecurityLevel2_Current.principal_authenticator	6-44
DITobj_PrincipalAuthenticator	6-44
DITobj_PrincipalAuthenticator.authenticate	6-45
DITobj_PrincipalAuthenticator.build_auth_data	6-47
DITobj_PrincipalAuthenticator.continue_authentication	6-49
DITobj_PrincipalAuthenticator.get_auth_type	6-50
DITobj_PrincipalAuthenticator.logon	6-51
DITobj_PrincipalAuthenticator.logoff	6-53
DISecurityLevel2_Credentials	6-53
DISecurityLevel2_Credentials.get_attributes	6-54
DISecurityLevel2_Credentials.is_valid	6-55
Programming Examples	6-56
C++ Example: Using WLE Extensions to Log on	6-56
C++ Example: Using CORBA Security to Log on	6-60
Java Example: Using WLE Extensions to Log on	6-60
Java Example: Getting Information from Privileges	6-62
Java Example: Checking the Validity of the Credentials	
Expiration Time	6-63

Java Example: Authentication Using SecurityLevel2::PrincipalAuthenticator	6-64
Java Example: Authentication Using Tobj::PrincipalAuthenticator	6-66
Java Example: Logging Off Using Tobj::PrincipalAuthenticator.....	6-68
Java Example: Checking the Validity of Credentials.....	6-69
Java Example: Getting Principal's Privileges	6-69
Java Example: Copying a Credentials Object	6-70
Java Example: Destroying a Credentials Object	6-71
Java Example: Getting the Principal Authenticator Object.....	6-71
Java Example: Getting Credentials	6-72
Java Example: Setting Default Credentials	6-72
Java Example: Getting a Principal's Privileges.....	6-73
Java Example: Removing a Credentials Object from the "Own" List	6-74
Java Example: Getting Credentials of the Requesting Principal.....	6-74
Java Example: Getting the Principal's Privileges from Credentials	6-75
Java Example: Getting the Principal's Privileges from the SecurityCurrent Object	6-76
Java Example: Obtaining the SecurityCurrent Object.....	6-77
Java Example: Getting Association Options	6-77
Java Example: Getting Delegation State	6-77
Java Example: Getting Delegation Mode.....	6-78

7. Transaction Service

Capabilities and Limitations	7-1
Lightweight Clients with Delegated Commit.....	7-1
Transaction Propagation	7-2
Transaction Integrity	7-2
Transaction Termination	7-3
Flat Transactions	7-3
Interoperability Between Remote Clients and the WLE Domain	7-3
Intradomain Interoperability.....	7-3
Network Interoperability	7-4
Relationship of the Transaction Service to Transaction Processing	7-4
Process Failure.....	7-5
Multithreaded Support.....	7-5

OMG Interface Definition Language (IDL).....	7-5
General Constraints	7-6
Getting Initial References to the TransactionCurrent Object	7-7
Transaction Service API.....	7-7
Data Types.....	7-8
Exceptions	7-9
Current Interface	7-10
Control Interface	7-16
TransactionalObject Interface	7-17
Other CORBAServices Object Transaction Service Interfaces.....	7-18
Transaction Service API Extensions	7-18
Exception.....	7-18
TransactionCurrent Interface.....	7-19

8. Interface Repository Interfacest

Structure and Usage.....	8-2
Programming Information.....	8-3
Performance Implications	8-4
Building Client Applications	8-5
Getting Initial References to the InterfaceRepository Object	8-5
Interface Repository Interfaces.....	8-6
Supporting Type Definitions	8-6
IObject Interface	8-7
Contained Interface	8-7
Container Interface	8-9
IDLType Interface.....	8-11
Repository Interface	8-11
ModuleDef Interface	8-12
ConstantDef Interface	8-12
TypedefDef Interface	8-13
StructDef	8-14
UnionDef	8-14
EnumDef	8-15
AliasDef	8-15
PrimitiveDef	8-16

ExceptionDef	8-16
AttributeDef	8-17
OperationDef	8-18
InterfaceDef	8-20

9. Joint Client/Servers

Main Program and Server Initialization	9-2
Servants	9-2
Servant Inheritance from Skeletons	9-3
Callback Object Models Supported	9-4
Preparing Callback Objects Using CORBA	9-5
Preparing Callback Objects Using BEAWrapper Callbacks	9-7
Callbacks	9-10
start_transient	9-11
start_persistent_systemid	9-13
restart_persistent_systemid	9-15
start_persistent_userid	9-17
stop_object	9-19
stop_all_objects	9-20
get_string_oid	9-21
~Callbacks	9-22

10. Development Commands

buildobjclient	10-3
buildobjserver	10-7
genicf	10-11
idl	10-12

11. Mapping of OMG IDL Statements to C++

Mappings	11-1
Data Types	11-2
Strings	11-4
Constants	11-4
Enums	11-5
Structs	11-6
Unions	11-8

Sequences	11-13
Arrays	11-18
Exceptions	11-20
Mapping of Pseudo-objects to C++	11-22
Usage	11-23
Mapping Rules	11-23
Relation to the C PIDL Mapping	11-25
Typedefs	11-26
Implementing Interfaces	11-27
Implementing Operations	11-29
PortableServer Functions	11-31
Modules	11-32
Interfaces	11-33
Generated Static Member Functions	11-34
Object Reference Types	11-35
Attributes	11-35
Any Type	11-38
Fixed-Length Versus Variable-Length User-Defined Types	11-49
Using var Classes	11-50
Sequence vars	11-53
Array vars	11-53
String vars	11-54
Using out Classes	11-56
Object Reference out Parameter	11-58
Sequence outs	11-59
Array outs	11-59
String outs	11-60
Argument Passing Considerations	11-61
Operation Parameters and Signatures	11-64

12. CORBA API

Global Classes	12-1
Pseudo-objects	12-2
Any Class Member Functions	12-2
CORBA::Any::Any()	12-4

CORBA::Any::Any(const CORBA::Any & InitAny)	12-5
CORBA::Any::Any(TypeCode_ptr TC, void * Value, Boolean Release)	12-6
CORBA::Any::~~Any()	12-7
CORBA::Any & CORBA::Any::operator=(const CORBA::Any & InitAny)	12-8
void CORBA::any::operator<<=()	12-9
CORBA::Boolean CORBA::Any::operator>>=()	12-10
CORBA::Any::operator<<=()	12-11
CORBA::Boolean CORBA::Any::operator>>=()	12-12
CORBA::TypeCode_ptr CORBA::Any::type() const	12-13
void CORBA::Any::replace()	12-14
Context Member Functions	12-15
Memory Management	12-15
CORBA::Context::context_name	12-16
CORBA::Context::create_child	12-17
CORBA::Context::delete_values	12-18
CORBA::Context::get_values	12-19
CORBA::Context::parent	12-21
CORBA::Context::set_one_value	12-22
CORBA::Context::set_values	12-23
ContextList Member Functions	12-24
CORBA::ContextList::count	12-25
CORBA::ContextList::add	12-26
CORBA::ContextList::add_consume	12-27
CORBA::ContextList::item	12-28
CORBA::ContextList::remove	12-29
NamedValue Member Functions	12-30
Memory Management	12-30
CORBA::NamedValue::flags	12-31
CORBA::NamedValue::name	12-32
CORBA::NamedValue::value	12-33
NVList Member Functions	12-34
Memory Management	12-34
CORBA::NVList::add	12-36

CORBA::NVList::add_item.....	12-37
CORBA::NVList::add_value.....	12-38
CORBA::NVList::count.....	12-39
CORBA::NVList::item	12-40
CORBA::NVList::remove	12-41
Object Member Functions	12-42
CORBA::Object::_create_request.....	12-44
CORBA::Object::_duplicate	12-46
CORBA::Object::_get_interface	12-47
CORBA::Object::_is_a	12-48
CORBA::Object::_is_equivalent	12-49
CORBA::Object::_nil.....	12-50
CORBA::Object::_non_existent	12-51
CORBA::Object::_request	12-52
CORBA Member Functions	12-53
CORBA::release.....	12-54
CORBA::is_nil.....	12-55
CORBA::hash	12-56
CORBA::resolve_initial_references	12-57
ORB Member Functions.....	12-58
CORBA::ORB::create_environment	12-60
CORBA::ORB::create_list.....	12-61
CORBA::ORB::create_named_value	12-62
CORBA::ORB::create_exception_list	12-63
CORBA::ORB::create_context_list	12-64
CORBA::ORB::create_policy.....	12-65
CORBA::ORB::create_operation_list.....	12-68
CORBA::ORB::get_default_context	12-69
CORBA::ORB::get_next_response	12-70
CORBA::ORB::perform_work	12-71
CORBA::ORB::run.....	12-72
CORBA::ORB::shutdown.....	12-73
CORBA::ORB::object_to_string	12-74
CORBA::ORB::poll_next_response	12-75
CORBA::ORB::work_pending	12-76

CORBA::ORB::send_multiple_requests_deferred	12-77
CORBA::ORB::send_multiple_requests_oneway	12-78
CORBA::ORB::string_to_object	12-79
ORB Initialization Member Function	12-80
CORBA::ORB_init	12-81
Policy Member Functions	12-84
CORBA::Policy::copy	12-85
CORBA::Policy::destroy	12-86
PortableServer Member Functions	12-87
PortableServer::POA::activate_object	12-88
PortableServer::POA::activate_object_with_id	12-89
PortableServer::POA::create_id_assignment_policy	12-90
PortableServer::POA::create_lifespan_policy	12-91
PortableServer::POA::create_POA	12-93
PortableServer::POA::create_reference	12-95
PortableServer::POA::create_reference_with_id	12-96
PortableServer::POA::deactivate_object	12-97
PortableServer::POA::destroy	12-98
PortableServer::POA::find_POA	12-99
PortableServer::POA::reference_to_id	12-100
PortableServer::POA::the_POAManager	12-101
PortableServer::ServantBase::_default_POA	12-102
POA Current Member Functions	12-103
PortableServer::Current::get_object_id	12-104
PortableServer::Current::get_POA	12-105
POA Manager Member Functions	12-106
PortableServer::POAManager::activate	12-107
PortableServer::POAManager::deactivate	12-108
POA Policy Member Objects	12-109
PortableServer::LifespanPolicy	12-110
PortableServer::IdAssignmentPolicy	12-111
Request Member Functions	12-112
CORBA::Request::arguments	12-113
CORBA::Request::ctx(Context_ptr)	12-114
CORBA::Request::get_response	12-115

CORBA::Request::invoke.....	12-116
CORBA::Request::operation	12-117
CORBA::Request::poll_response	12-118
CORBA::Request::result.....	12-119
CORBA::Request::env	12-120
CORBA::Request::ctx.....	12-121
CORBA::Request::contexts	12-122
CORBA::Request::exceptions	12-123
CORBA::Request::target	12-124
CORBA::Request::send_deferred.....	12-125
CORBA::Request::send_oneway.....	12-126
Strings.....	12-127
CORBA::string_alloc.....	12-128
CORBA::string_dup.....	12-129
CORBA::string_free	12-130
TypeCode Member Functions	12-131
Memory Management	12-132
CORBA::TypeCode::equal	12-133
CORBA::TypeCode::id.....	12-134
CORBA::TypeCode::kind.....	12-135
CORBA::TypeCode::param_count.....	12-137
CORBA::TypeCode::parameter.....	12-138
Exception Member Functions.....	12-139
Standard Exceptions	12-141
Exception Definitions.....	12-142
Object Nonexistence	12-143
Transaction Exceptions	12-143
ExceptionList Member Functions	12-145
CORBA::ExceptionList::count	12-146
CORBA::ExceptionList::add	12-147
CORBA::ExceptionList::add_consume	12-148
CORBA::ExceptionList::item	12-149
CORBA::ExceptionList::remove	12-150

13. Server-side Mapping

Implementing Interfaces	151
Inheritance-based Interface Implementation	152
Delegation-based Interface Implementation.....	155
Implementing Operations	159

Preface

Purpose of This Document

This document describes the BEA WebLogic Enterprise (sometimes referred to as WLE) C++ application programming interface (API).

Note: Effective February 1999, the BEA M3 product is renamed. The new name of the product is BEA WebLogic Enterprise (WLE).

Who Should Read This Document

This document is intended for application developers interested in using the WLE C++ API to write client and joint client/server applications and object implementations. It assumes a familiarity with CORBA, and with C++ and Java programming.

How This Document Is Organized

The *C++ Programming Reference* is organized as follows:

- ◆ Chapter 1, “OMG IDL Syntax,” describes the Object Management Group (OMG) Interface Definition Language (IDL), OMG IDL extensions, and includes a reference to OMG IDL coding style guidelines.
- ◆ Chapter 2, “Implementation Configuration File (ICF),” describes the Implementation Configuration File (ICF).

-
- ◆ Chapter 3, “TP Framework,” describes the WLE TP Framework application programming interface (API).
 - ◆ Chapter 4, “Bootstrap Object,” describes the Bootstrap object.
 - ◆ Chapter 5, “FactoryFinder Interface,” describes the FactoryFinder interface.
 - ◆ Chapter 6, “Security Service,” describes the Security Service.
 - ◆ Chapter 7, “Transaction Service,” describes the Transaction Service.
 - ◆ Chapter 8, “Interface Repository Interface,” describes the Interface Repository interfaces.
 - ◆ Chapter 9, “Joint Client/Servers,” describes how to program joint client/server applications and the BEAWrapper Callbacks API.
 - ◆ Chapter 10, “Development Commands,” describes the build and administration commands for UNIX and Windows NT platforms.
 - ◆ Chapter 11, “Mapping of OMG IDL Statements to C++,” describes mapping of OMG IDL statements to C++.
 - ◆ Chapter 12, “CORBA API,” describes the CORBA API.
 - ◆ Chapter 13, “Server-side Mapping,” describes server-side mapping of OMG IDL statements to C++.

How to Use This Document

This document, *C++ Programming Reference*, is designed primarily as an online, hypertext document. If you are reading this as a paper publication, note that to get full use from this document you should access it as an online document via the Online Documentation CD for the BEA WebLogic Enterprise 4.2 release.

The following sections explain how to view this document online, and how to print a copy of this document.

Opening the Document in a Web Browser

To access the online version of this document, open the following:

`\doc\wle\v42\index.htm`

Note: The online documentation requires Netscape Communicator version 4.0 or later, or Microsoft Internet Explorer version 4.0 or later.

Printing from a Web Browser

You can print a copy of this document, one file at a time, from the Web browser. Before you print, make sure that the chapter or appendix you want is displayed and *selected* in your browser. To select a chapter or appendix, click anywhere inside the chapter or appendix you want to print.

The Online Documentation CD also includes Adobe Acrobat PDF files of all of the online documents. You can use the Adobe Acrobat Reader to print all or a portion of each document. On the CD Home Page, click the PDF Files button and scroll to the entry for the document you want to print.

Documentation Conventions

The following documentation conventions are used throughout this document.

Convention	Item
boldface text	Indicates terms defined in the glossary.
Ctrl+Tab	Indicates that you must press two or more keys simultaneously.
<i>italics</i>	Indicates emphasis or book titles.

Convention	Item
monospace text	<p>Indicates code samples, commands and their options, data structures and their members, data types, directories, and file names and their extensions. Monospace text also indicates text that you must enter from the keyboard.</p> <p><i>Examples:</i></p> <pre>#include <iostream.h> void main () the pointer psz chmod u+w * .doc BITMAP float</pre>
monospace boldface text	<p>Identifies significant words in code.</p> <p><i>Example:</i></p> <pre>void commit ()</pre>
<i>monospace italic text</i>	<p>Identifies variables in code.</p> <p><i>Example:</i></p> <pre>String <i>expr</i></pre>
UPPERCASE TEXT	<p>Indicates device names, environment variables, and logical operators.</p> <p><i>Examples:</i></p> <pre>LPT1 SIGNON OR</pre>
{ }	<p>Indicates a set of choices in a syntax line. The braces themselves should never be typed.</p>
[]	<p>Indicates optional items in a syntax line. The brackets themselves should never be typed.</p> <p><i>Example:</i></p> <pre>buildobjclient [-v][-o name] [-f firstfile-syntax] [-l lastfile-syntax]-P</pre>
	<p>Separates mutually exclusive choices in a syntax line. The symbol itself should never be typed.</p>

Convention	Item
...	Indicates one of the following in a command line: <ul style="list-style-type: none">◆ That an argument can be repeated several times in a command line◆ That the statement omits additional optional arguments◆ That you can enter additional parameters, values, or other information The ellipsis itself should never be typed. <i>Example:</i> genicf [options] idl-filename...
.	Indicates the omission of items from a code example or from a syntax line. The vertical ellipsis itself should never be typed.

Related Documentation

The following sections list the documentation provided with the BEA WebLogic Enterprise software, related BEA publications, and other publications related to the technology.

BEA WebLogic Enterprise Documentation

The BEA WebLogic Enterprise information set consists of the following documents:

Installation Guide

C++ Release Notes

Java Release Notes

Getting Started

Guide to the University Sample Applications

Guide to the Java Sample Applications

Creating Client Applications

Creating C++ Server Applications

Creating Java Server Applications

Administration Guide

Using Server-to-Server Communication

C++ Programming Reference (this document)

Java Programming Reference

Java API Reference

JDBC Driver Programming Reference

System Messages

Glossary

Technical Articles

Note: The Online Documentation CD also includes Adobe Acrobat PDF files of all of the online documents. You can use the Adobe Acrobat Reader to print all or a portion of each document.

BEA Publications

Selected BEA TUXEDO Release 6.5 for BEA WebLogic Enterprise version 4.2 documents are available on the Online Documentation CD.

To access these documents:

1. Click the Other Reference button from the main menu.
2. Click the TUXEDO Documents option.

Other Publications

For more information about CORBA and related technologies, refer to the following books and specifications:

Cobb, E. 1997. *The Impact of Object Technology on Commercial Transaction Processing*. VLDB Journal, Volume 6. 173-190.

Edwards, J. with DeVoe, D. 1997. *3-Tier Client/Server At Work*. Wiley Computer Publishing.

Edwards, J., Harkey, D., and Orfali, R. 1996. *The Essential Client/Server Survival Guide*. Wiley Computer Publishing.

Flanagan, David. May 1997. *Java in a Nutshell*, 2nd Edition. O'Reilly & Associates, Incorporated.

Flanagan, David. September 1997. *Java Examples in a Nutshell*. O'Reilly & Associates, Incorporated.

Fowler, M. with Scott, K. 1997. *UML Distilled, Applying the Standard Object Modeling Language*. Addison-Wesley.

Gamma, E., Helm, R., Johnson, R., and Vlissides, J. 1995. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley Professional Computing Series.

Jacobson, I. 1994. *Object-Oriented Software Engineering: A Use Case Driven Approach*. Addison-Wesley.

Mowbray, Thomas J. and Malveau, Raphael C. (Contributor). 1997. *CORBA Design Patterns*, Paper Back and CD-ROM Edition. John Wiley & Sons, Inc.

Orfali, R., Harkey, D., and Edwards, J. 1997. *Instant CORBA*. Wiley Computer Publishing.

Orfali, R. and Harkey, D. February 1998. *Client/Server Programming with Java and CORBA*, 2nd Edition. John Wiley & Sons, Inc.

Otte, R., Patrick, P., and Roy, M. 1996. *Understanding CORBA*. Prentice Hall PTR.

Rosen, M. and Curtis, D. 1998. *Integrating CORBA and COM Applications*. Wiley Computer Publishing.

Rumbaugh, J., Blaha, M., Premerlani, W., Eddy, F., and Loresen, W. 1991.
Object-Oriented Modeling and Design. Prentice Hall.

The Common Object Request Broker: Architecture and Specification. Revision 2.2,
February 1998. Published by the Object Management Group (OMG).

CORBA services: Common Object Services Specification. Revised Edition. Updated:
November 1997. Published by the Object Management Group (OMG).

Contact Information

The following sections provide information about how to obtain support for the documentation and the software.

Documentation Support

If you have questions or comments on the documentation, you can contact the BEA Information Engineering Group by e-mail at **docsupport@beasys.com**. (For information about how to contact Customer Support, refer to the following section.)

Customer Support

If you have any questions about this version of the BEA WebLogic Enterprise product, or if you have problems installing and running the BEA WebLogic Enterprise software, contact BEA Customer Support through BEA WebSupport at www.beasys.com. You can also contact Customer Support by using the contact information provided on the Customer Support Card, which is included in the product package.

When contacting Customer Support, be prepared to provide the following information:

- ◆ Your name, e-mail address, phone number, and fax number
- ◆ Your company name and company address

-
- ◆ Your machine type and authorization codes
 - ◆ The name and version of the product you are using
 - ◆ A description of the problem and the content of pertinent error messages



1 **OMG IDL Syntax**

The Object Management Group (OMG) Interface Definition Language (IDL) is used to describe the interfaces that client objects call and that object implementations provide. An OMG IDL interface definition fully specifies each operation's parameters and provides the information needed to develop client applications that use the interface's operations.

Client applications are written in languages for which mappings from OMG IDL statements have been defined. How an OMG IDL statement is mapped to a client language construct depends on the facilities available in the client language. For example, an OMG IDL exception might be mapped to a structure in a language that has no notion of exception, or to an exception in a language that does.

OMG IDL statements obey the same lexical rules as C++ statements, although new keywords are introduced to support distribution concepts. OMG IDL statements also provide full support for standard C++ preprocessing features and OMG IDL-specific pragmas.

The OMG IDL grammar is a subset of ANSI C++ with additional constructs to support the operation invocation mechanism. OMG IDL is a declarative language; it supports C++ syntax for constant, type, and operation declarations; it does not include any algorithmic structures or variables.

For a description of OMG IDL grammar, see Chapter 3 of the *Common Object Request Broker: Architecture and Specification* Revision 2.2 “OMG IDL Syntax and Semantics.”

All OMG IDL grammar is supported, with the exception of the following type declarations and associated literals:

◆ `native`

Note: Because CORBA 2.2 states that the `native` type declaration is intended for use in Object Adapters, not user interfaces, this type is available in the `PortableServer` module only for clients that support callbacks, that is, joint client/servers.

◆ `long long`

◆ `unsigned long long`

◆ `long double`

◆ `wstring`

◆ `wchar`

◆ `fixed`

OMG IDL Extensions

The IDL compiler defines preprocessor macros specific to the platform. All macros predefined by the preprocessor that you are using can be used in the OMG IDL file, in addition to the user-defined macros. You can also define your own macros when you are compiling or loading OMG IDL files.

Table 1-1 describes the predefined macros for each platform.

Table 1-1 Predefined Macros

Macro Identifier	Platform on Which the Macro Is Defined
<code>__unix__</code>	Sun Solaris, HP-UX, Tru64 UNIX, and IBM AIX
<code>__osf1__</code>	Tru64 UNIX
<code>__sun__</code>	Sun Solaris
<code>__hpux__</code>	HP-UX

Table 1-1 Predefined Macros

Macro Identifier	Platform on Which the Macro Is Defined
__aix__	IBM AIX
__win_nt__	Microsoft Windows NT

2 Implementation Configuration File (ICF)

The WLE TP Framework application programming interface (API) provides callback methods for object activation and deactivation. These methods provide the ability for application code to implement flexible state management schemes for CORBA objects.

State management is the way you control the saving and restoring of object state during object deactivation and activation. State management also affects the duration of object activation, which influences the performance of servers and their resource usage. The external API of the TP Framework includes the `activate_object()` and `deactivate_object()` methods, which provide a possible location for state management code. Additionally, the TP Framework API includes the `deactivateEnable()` method to enable the user to control the timing of object deactivation. The default duration of object activation is controlled by policies assigned to implementations at OMG IDL compile time.

While CORBA objects are active, their state is contained in a servant. This state must be initialized when objects are first invoked (that is, the first time a method is invoked on a CORBA object after its object reference is created) and on subsequent invocations after objects have been deactivated.

While a CORBA object is deactivated, its state must be saved outside the process in which the servant was active. When an object is activated, its state must be restored. The object's state can be saved in shared memory, in a file, in a database, and so forth. It is up to the programmer to determine what constitutes an object's state, and what must be saved before an object is deactivated and restored when an object is activated.

You can use the Implementation Configuration File (ICF) to set activation policies to control the duration of object activations in each implementation. The ICF file manages object state by specifying the activation policy. The activation policy

determines the in-memory activation duration for a CORBA object. A CORBA object is active in a Portable Object Adapter (POA) if the POA's active object map contains an entry that associates an object ID with an existing servant. Object deactivation removes the association of an object ID with its active servant.

ICF Syntax

ICF syntax is as follows:

```
[#pragma activation_policy method|transaction|process]
[#pragma transaction_policy never|ignore|optional|always]
[Module module-name {]
    implementation [implementation-name]
    {
        implements (module-name::interface-name);
        [activation_policy (method|transaction|process);]
        [transaction_policy (never|ignore|optional|always);]
    };
[};]
```

pragmas

The two optional pragmas allow you to set a specific policy as the default policy for the entire ICF for all implementations that do not have an explicit `activation_policy` or `transaction_policy` statement. This feature relieves the programmer from having to specify policies for each implementation and/or allows overriding of the defaults.

Module module-name

The `module-name` variable is optional if it is optional in the OMG IDL file. This variable is used for scoping and grouping. Its use must be consistent with the way it is used inside the OMG IDL file.

implementation-name

This variable is optional and is used as the name of the servant or as the class name in the server. It is constructed using `interface-name` with an `_i` appended if it is not specified by the programmer.

implements (module-name::interface-name)

This variable identifies the module and the interface to which the activation policy and the transaction policy apply.

`activation_policy`

For a description of the activation policies, see “Activation Policy” on page 3-9.

`transaction_policy`

For a description of the transaction policies, see “Transaction Policies” on page 3-17.

Sample ICF File

Listing 2-1 shows a sample ICF file.

Listing 2-1 Sample ICF

```
module POA_University1
{
    implementation CourseSynopsisEnumerator_i
    {
        activation_policy ( process );
        transaction_policy ( optional );
        implements ( University1::CourseSynopsisEnumerator );
    };
};

module POA_University1
{
    implementation Registrar_i
    {
        activation_policy ( method );
        transaction_policy ( optional );
        implements ( University1::Registrar );
    };
};

module POA_University1
{
    implementation RegistrarFactory_i
    {
```

```
activation_policy ( process );  
transaction_policy ( optional );  
implements ( University1::RegistrarFactory );  
};
```

```
};
```

Creating the ICF File

You have the option of either coding the ICF file manually or using the `genicf` command to generate it from the OMG IDL file. The syntax and options for the `genicf` command are described in the section “`genicf`” on page 10-11.

3 TP Framework

The WLE TP Framework provides a programming TP Framework that enables users to create servers for high-performance TP applications. This chapter describes the TP Framework programming model and the TP Framework application programming interface (API) in detail. Additional information about how to use this API can be found in *Creating C++ Server Applications*.

The TP Framework is required when developing WLE servers. Later releases will relax this requirement, though it is expected that most customers will use the TP Framework as an integral part of their applications.

WLE uses BEA TUXEDO as the underlying infrastructure for providing load balancing, transactional capabilities, and administrative infrastructure. The base API used by the TP Framework is the CORBA API with BEA extensions. The TP Framework API is exposed to customers. The BEA TUXEDO ATMI is an optional API that can be mixed in with TP Framework APIs, allowing a customer to deploy distributed applications using a mix of BEA TUXEDO servers and WLE servers.

Before WLE, ORB products did not approach BEA TUXEDO's performance in large-scale environments. BEA TUXEDO systems support applications that can process hundreds of transactions per second. These applications are built using the BEA TUXEDO stateless-service programming model that minimizes the amount of system resources used for each request, and thus maximizes throughput and price performance.

Now, WLE and its TP Framework give customers a way to develop CORBA applications with performance similar to BEA TUXEDO applications. WLE servers that use the TP Framework provide throughput, response time, and price performance approaching the BEA TUXEDO stateless-service programming model, while using the CORBA programming model.

Definition of Terms

The following terms are used in the discussion of the TP Framework:

Active object—a CORBA object that can service an invocation by a client on an object reference by invoking a method of a servant (a C++ or Java object in memory). This implies that a Portable Object Adapter has an entry in an Active Object Map that maps the object ID for the object to a servant.

Active Object Map—a table maintained by a POA that maps the association of object references to servants.

application code—code that is written by the user, as opposed to system code that is provided by BEA System, Inc.

callback method—a virtual method that is implemented by application code and that is invoked by system code when needed to perform a specific function. Callback methods are never intended to be invoked directly by application code.

client—a program that invokes methods on a CORBA object. A client can be in the same process as the server for that object, a different process on the same machine, or a different machine.

CORBA—Common Object Request Broker Architecture. A multivendor standard published by the Object Management Group for distributed object-oriented computing.

CORBA object—an entity that complies with the CORBA standard upon which operations are performed. An object is defined by its interface. In this document CORBA object means an object that is invoked using either generated stubs and skeletons or CORBA-defined dynamic invocation mechanisms, as opposed to a C++ object that cannot be invoked in such a way. A CORBA object is a logical object represented by an object reference. The C++ object that is the target of an invocation on a CORBA object invocation is called a “servant” to distinguish it from ordinary C++ objects, which, in this document are called “objects.” In WLE, all CORBA objects implemented in application code are under control of the TP Framework.

CORBA ORB—any Object Request Broker (ORB) that complies with the CORBA standard. A CORBA ORB is a communications intermediary between client and server applications that typically are distributed across a network. The WLE ORB is a CORBA ORB.

Factory—any CORBA object that returns an object reference to other CORBA objects. A factory is located in the server application.

factory-based routing—a feature of WLE that permits the routing of requests on an object reference to a specific server group based on criteria supplied at the time the object reference is created by a factory.

global transaction—a transaction that can execute in more than one server, accessing data from more than one resource manager. A global transaction may be composed of several local transactions, each accessing a single resource manager.

ICF file—Implementation Configuration File. A file that is optional input to the WLE IDL compiler. It specifies properties of implementations of interfaces described in the IDL input to the compiler.

implementation code—executable code that can service a request (method invocation) on a CORBA object. An implementation of an interface described in CORBA IDL. In WLE an implementation is the executable code for a C++ class. An implementation corresponds to a C++ class; a servant corresponds to an instance of that class. A C++ class implements the interface described in OMG IDL.

interface—a declaration in OMG IDL of an interface to a CORBA object. The interface declaration contains IDL operations and attributes. The OMG IDL interface declaration is used to generate stubs and skeletons for WLE CORBA objects.

local transaction—a transaction that accesses a single database or file and is controlled by the resource manager responsible for performing concurrency control and atomicity of updates at that database.

object—a C++ object that is not an implementation of an interface. When such an object is the implementation of an interface it is called a “servant.” See also CORBA Object and Servant.

object activation—the association of a servant with an object ID in the Active Object Map of a POA and the TP Framework. The result of object activation is that an invocation can be made on a servant to service a client invocation of a method on an object reference.

object deactivation—the removal of the association of an object ID to a servant in the Active Object Map of a POA and the TP Framework. The result of object deactivation is that no client invocation on an object reference that contains this object ID can be satisfied without first performing object activation.

object ID—a value that is used by the POA and by the implementation to identify a CORBA object. Object IDs are specified when object references are created. Object ID values are hidden from clients by being encapsulated in object references. Object IDs are the unique identifiers for a CORBA object's state in WLE. The assignment and interpretation of object ID values is the responsibility of the application programmer.

object reference—an unambiguous identifier that associates an object definition with an instance of the object, such as a bank account number or an employee identification number. Client applications use object references to invoke operations upon that CORBA objects.

OID—object ID.

ORB—Object Request Broker. See CORBA ORB.

OTS—Object Transaction Service.

POA—a run-time library of functions that are built in to the server application executable image. The POA creates and manages object references to all objects used by the application. In addition, the POA manages object state and provides the infrastructure for the support of persistent objects and the portability of object implementations between different ORB products.

The WLE server application procedure automatically builds the POA into the server application. The WLE TP Framework automatically handles all the WLE server application interactions with the POA.

Note that server applications interact directly with the POA.

request—a message that is passed to an ORB. In this document, this is assumed to be as a result of a client invocation of a method on an object reference.

servant—a programming language object (a C++ or Java object in WLE) that services invocations on one or more CORBA objects. Servants always exist within the context of a server process. Requests made on an object's references are mediated by the ORB and are transformed into invocations on a particular servant. In the course of an object's lifetime, it may be associated with multiple servants; that is, requests on its references will be targeted at multiple servants.

server—a program that contains the implementations of one or more CORBA objects. A server is a process on a computer.

state—the time-varying properties of an object that affect that object's behavior. In WLE, state for a CORBA object is the state associated with the corresponding servant when the CORBA object is active. When inactive, the state is assumed to be saved in an application-specific way, such as in a database.

stroid—An Object ID represented as a string.

TP—Transaction Processing.

TP Framework—a feature of WLE software that makes it easier to implement mission-critical applications using CORBA. The TP (Transaction Processing) Framework is a run-time library of default implementations that the WLE server application build procedure links to the server application executable image. The TP Framework consists of a set of convenience functions that make it easy for you to write code that does the following:

1. Initializes the server application and executes startup and shutdown routines
2. Ties the server application to WLE domain resources
3. Manages objects, bringing them into memory when needed, flushing them from memory when no longer needed, and managing reading and writing of data for persistent objects
4. Performs object housekeeping

A Simple Programming Model

The TP Framework provides a simple, useful subset of the wide range of possible CORBA object implementation choices. You use it for the development of server-side object implementations only. When using any client-side CORBA ORB, clients interact with CORBA objects whose server-side implementations are managed by the TP Framework. Clients are unaware of the existence of the TP Framework—a client written to access a CORBA object executing in a non-BEA WLE server environment will be able to access that same CORBA object executing in a WLE server environment without any changes or restrictions to the client interface.

The TP Framework provides a server environment and an API that is easier to use and understand than the CORBA Portable Object Adapter (POA) API, and is specifically geared towards enterprise applications. It is a simple server programming model and an orthodox implementation of the CORBA model, which will be familiar to programmers using ORBs such as ORBIX or VisiBroker.

The TP Framework simplifies the programming of WLE servers by reducing the complexity of the server environment in the following ways:

- ◆ The TP Framework does all interactions with the POA and the Naming Service. The application programmer requires no knowledge of POA or Naming Service interfaces.
- ◆ The TP Framework is single threaded—only one request on one CORBA object will be processed at a time, obviating the need to write thread-safe implementations.
- ◆ A CORBA object may be involved in only one transaction at a time (consistent with the association of one object ID to one servant).

The TP Framework provides the following functionality:

- ◆ Control Flow
- ◆ Object State Management
- ◆ Transaction Integration
- ◆ Object Housekeeping
- ◆ High-level Services

Control Flow

The TP Framework, in conjunction with the ORB and the POA, controls the flow of the application program by doing the following:

- ◆ Controlling the server mainline and invoking callback methods on TP Framework-defined classes at appropriate times for server startup and shutdown. This relieves the application programmer from complex interactions related to ORB and POA initialization and coordination of transactions, resource managers, and object state on shutdown.

- ◆ Scheduling objects for activation and deactivation when client requests arrive and are completed. This removes the complexity of management of object activation and deactivation from the realm of the application programmer and enables the use of the TP monitor infrastructure's powerful load-balancing capabilities, crucial to performance of mission-critical tasks.

Object State Management

The TP Framework API provides callback methods for application code to implement flexible state management schemes for CORBA objects. State management involves the saving and restoring of object state on object deactivation and activation. It also concerns the duration of activation of objects, which influences the performance of servers and their resource usage. The default duration of object activation is controlled by policies assigned to implementations at IDL compile time.

Transaction Integration

TP Framework transaction integration provides the following features:

- ◆ CORBA objects can participate in global transactions.
- ◆ Objects participating in transactions can be implemented as stateful objects that remain in memory for the duration of a transaction (by using the transaction activation policy), to decrease client response time.
- ◆ CORBA objects that participate in transactions can affect transaction outcome either during their transactional work or just prior to the system's execution of the two-phase commit algorithm after all transactional work has been completed.
- ◆ Transactions can be automatically initiated on the server transparent to the client.

Object Housekeeping

When a server is shut down, the TP Framework rolls back any transactions that the server is involved in and deactivates any CORBA objects that are currently active.

High-level Services

The TP interface in the TP Framework API provides methods for performing object registrations and utility functions. The following services are provided:

- ◆ Object reference creation
- ◆ Factory-based routing support
- ◆ Accessors for system objects, such as the ORB
- ◆ Registration and unregistration of factories with the Factory Finder
- ◆ Application-controlled activation and deactivation
- ◆ User logging

The purpose of these high-level service methods is to eliminate the need for developers to understand the CORBA POA, CORBA Naming Service, and BEA TUXEDO APIs, which they use for their underlying implementations. By encapsulating the underlying API calls with a high-level set of methods, programmers can focus their efforts on providing business logic rather than understanding and using the more complex underlying facilities.

State Management

State management involves the saving and restoring of object state on object deactivation and activation. It also concerns the duration of activation of objects, which influences the performance of servers and their resource usage. The external API of the TP Framework provides `activate_object` and `deactivate_object` methods, which are a possible location for state management code.

Activation Policy

State management is provided in the TP Framework by the activation policy. This policy controls the activation and deactivation of servants for a particular IDL interface (as opposed to the creation and destruction of the servants). This policy is applicable only to CORBA objects using the TP Framework.

The activation policy determines the default in-memory activation duration for a CORBA object. A CORBA object is active in a POA if the POA's active object map contains an entry that associates an object ID with an existing servant. Object deactivation removes the association of an object ID with its active servant. You can choose from one of three activation policies: `method` (the default), `transaction`, or `process`.

Note: The activation policies are set in an ICF file that is configured at OMG IDL compile time. For a description of the ICF file, refer to Chapter 2, "Implementation Configuration File (ICF)."

The activation policies are described below:

◆ `method` (This is the default activation policy.)

The activation of the CORBA object (that is, the association between the object ID and the servant) lasts until the end of the method. At the completion of a method, the object is deactivated. When the next method is invoked on the object reference, the CORBA object is activated (the object ID is associated with a new servant). This behavior is similar to that of a BEA TUXEDO stateless service.

◆ transaction

The activation of the CORBA object (that is, the association between the object ID and the servant) lasts until the end of the transaction. During the transaction, multiple object methods can be invoked. The object is activated before the first method invocation on the object and is deactivated in one of the following ways:

- ◆ If a user-initiated transaction is in effect when the object is activated, the object is deactivated when the first of the following occurs: the transaction is committed or rolled back, or the server is shut down in an orderly fashion. The latter is done using either the `tmshutdown(1)` or `tmadmin(1)` command. These commands are described in the *BEA TUXEDO Reference* online document.
- ◆ If a user-initiated transaction is not in effect when the TP object is activated, the TP object is deactivated when the method completes.

The transaction activation policy provides a means for an object to vote on the outcome of the transaction prior to the execution of the two-phase commit algorithm. An object votes to roll back the transaction by calling

`Current.rollback_only()` in the

`Tobj_ServantBase::deactivate_object` method. It votes to commit the transaction by not calling `Current.rollback_only()` in the method.

Note: This is a model of resource allocation that is similar to that of a BEA TUXEDO conversational service. However, this model is less expensive than the BEA TUXEDO conversational service in that it uses fewer system resources. This is because of the WLE ORB's multicontexted dispatching model (that is, the presence of many servants in memory at the same time for one server), which makes it possible for a single server process to be shared by many concurrently active servants that service many clients. In the BEA TUXEDO system, the process would be dedicated to a single client and to only one service for the duration of a conversation.

◆ process

The activation of the CORBA object begins when it is invoked while in an inactive state and, by default, lasts until the end of the process.

Note: The TP Framework API provides an interface method (`TP::deactivateEnable`) that allows the application to control the timing of object deactivation for objects that have the `activation` policy set to `process`. For a description of this method, see the section "TP::deactivateEnable" on page 3-56.

Application-controlled Activation and Deactivation

Ordinarily, activation and deactivation decisions are made by the TP Framework, as discussed earlier in this chapter. The techniques in this section show how to use alternate mechanisms. The application can control the timing of activation and deactivation explicitly for objects with particular policies.

Explicit Activation

Application code can bypass the on-demand activation feature of the TP Framework for objects that use the `process` activation policy. The application can “preactivate” an object (that is, activate it before any invocation) using the `TP::create_active_object_reference` call.

Preactivation works as follows. Before the application creates an object reference, the application instantiates a servant and initializes that servant’s state. The application uses `TP::create_active_object_reference` to put the object into the Active Object Map (that is, associate the servant with an `ObjectId`). Then, when the first invocation is made, the TP Framework immediately directs the request to the process that created the object reference and then to the existing servant, bypassing the necessity to call `Server::create_servant` and then the servant’s `activate_object` method (just as if this were the second or later invocation on the object). Note that the object reference for such an object will not be directed to another server and the object will never go through on-demand activation as long as the object remains activated.

Since the preactivated object has the `process` activation policy, it will remain active until one of two events occurs: 1) the ending of the process or 2) a `TP::deactivateEnable` call.

USAGE NOTES

Preactivation is especially useful if the application needs to establish the servant with an initial state in the same process, perhaps using shared memory to initialize state. Waiting to initialize state until a later time and in a potentially different process may be very difficult if that state includes pointers, object references, or complex data structures. `TP::create_active_object_reference` guarantees that the preactivated object is in the same process as the code that is doing the preactivation.

While this is convenient, preactivation should be used sparingly, as should all process objects, because it preallocates precious resources. However, when needed and used properly, preallocation is more efficient than alternatives.

Examples of such usage might be an object using the “iterator” pattern. For example, there might a potentially long list of items that could be returned (in an unbound IDL sequence) from a “database_query” method (for example, the contents of the telephone book). Returning all such items in the sequence is impractical because the message size and the memory requirements would be too large.

On an initial call to get the list, an object using the iterator pattern returns only a limited number of items in the sequence and also returns a reference to an “iterator” object that can be invoked to receive further elements. This iterator object is initialized by the initial object; that is, the initial object creates a servant and sets its state to keep track of where in the long list of items the iteration currently stands (the pointer to the database, the query parameters, the cursor, and so forth).

The initial object preactivates this iterator object by using `TP::create_active_object_reference`. It also creates an object reference to that object to return to the client. The client then invokes repeatedly on the iterator object to receive, say, the next 100 items in the list each time. The advantage of preactivation in this situation is that the state might be complex. It is often easiest to set such state initially, from a method that has all the information in its context (call frame), when the initial object still has control.

CAUTION TO USERS

For objects to be preactivated in this fashion, the state usually cannot be recovered if a crash occurs. (This is because the state was considered too complex or inconvenient to set upon initial, delayed activation.) This is a valid object technique, essentially stating that the object is valid only for a single activation period.

However, a problem may arise because of the “one-time” usage. Since a client still holds an object reference that leads to the process containing that state, and since the state cannot be recreated after the crash, care must be taken that the client’s next invocation does not automatically provoke a new activation of the object, because that object would have inapplicable state.

The solution is to refuse to allow the object to be activated automatically by the TP Framework. If the user provides the `TobjS::ActivateObjectFailed` exception to the TP Framework as a result of the `activate_object` call, the TP Framework will not complete the activation and will return an exception to the client,

`CORBA::OBJECT_NOT_EXIST`. The client has presumably been warned about this possibility, since it knows about the iterator (or similar) pattern. The client must be prepared to restart the iteration.

Note: This defensive measure may not be necessary in the future; the TP Framework itself may detect that the object reference is no longer valid. In particular, you should not depend on the possibility that the `activate_object` method might be called. If the TP Framework does in fact change, `activate_object` will not be called and the framework itself will generate the `OBJECT_NOT_EXIST` exception.

Self Deactivation

Just as it is possible to preactivate an object with the `process` activation policy, it is possible to request the deactivation of an object with the `process` activation policy. The ability to preactivate and the ability to request deactivation are independent; regardless of how an object was activated, it can be deactivated explicitly.

A method in the application can request (via `TP::deactivateEnable`) that the object be deactivated. When `TP::deactivateEnable` is called and the object is subsequently deactivated, no guarantee is made that subsequent invocations on the CORBA object will result in reactivation in the same process as a previous activation. The association between the `ObjectId` and the servant exists from the activation of the CORBA object until one of the following events occurs: 1) the shutdown of the server process or 2) the application calls `TP::deactivateEnable`. After the association is broken, when the object is invoked again, it can be re-activated anywhere that is allowed by the WLE configuration parameters.

There are two forms of `TP::deactivateEnable`. In the first form (with no parameters), the object currently executing will be deactivated after completion of the method in which the call is made. The object itself makes the decision that it should be deactivated. This is often done during a method call that acts as a "signoff" signal.

The second form of `TP::deactivateEnable` allows a server to request deactivation of any active object, whether it is the object that is executing or not; that is, any part of the server can ask that the object be deactivated. This form takes parameters identifying the object to be deactivated. Explicit deactivation is not allowed for objects with an activation policy of `transaction`, because such objects cannot be safely deactivated until the end of a transaction.

In the `TP::deactivateEnable` call, the TP Framework calls the servant's `deactivate_object` method. Exactly when the TP Framework invokes `deactivate_object` depends on the state of the object to be deactivated. If the object is not currently in execution, the TP Framework deactivates it before returning to the caller. The object might be currently executing a method; this is always the case for `TP::deactivateEnable` with no parameters (since it refers to the currently executing object). In this case, `TP::deactivateEnable` is not told whether the object was deactivated immediately or not.

Servant Lifetime

A servant is a C++ class that contains methods to implement an IDL interface's operations. The user writes the servant code. The TP Framework invokes methods in the servant code to satisfy requests. The servant is created by the C++ "new" statement and is destroyed by the C++ "delete" statement. Exactly who does the creation and who does the deletion, and the timing of creation and deletion, is the subject of this section.

The Normal Case

In the normal case, the TP Framework completely controls the lifetime of a servant. The basic model is that, when a request for an inactive object arrives, the TP Framework obtains a servant and then activates it (by calling its `activate_object` method). At deactivation time, the TP Framework calls the servant's `deactivate_object` method and then disposes of the servant.

For this release of WLE, two phrases in the basic model above need to be further explained. The phrase "the TP Framework obtains a servant" means that when the TP Framework needs a servant to be created, it calls the user-written `Server::create_servant` method. At that time, the application code must return a pointer to the requested servant. The application almost always does this by using the C++ "new" statement to create a new instance of a servant. The phrase "disposes of the servant" means that the TP Framework deletes it.

The application must be aware that this current behavior of always creating and deleting a servant may change in future versions of this product. The application should not depend on the current behavior, but should write servant code that allows re-use of a servant. Specifically, the servant code must work even if the servant has not been freshly created (by the C++ "new" statement). The TP Framework reserves the right not to delete a servant after it has been deactivated and then to reactivate it. This

means that the servant must completely initialize itself at the time of the callback on the servant's `activate_object` method, not at the time of servant creation (not in the constructor).

Special Cases

There are two techniques an application can use to alter the normal TP Framework use of servants. The first has to do with obtaining a servant and the second has to do with disposing of the servant.

The application can alter the “obtaining” mechanism by using explicit preactivation. In this case, the application creates and initializes a servant before asking the TP Framework to declare it activated. Once such a servant has been turned over to the TP Framework (by the `TP::create_active_object_reference` call), that servant is treated by the TP Framework just like every other servant. The only difference is in its method of creation and initialization.

The application can alter the “disposing” mechanism by taking the responsibility for disposing of a servant instead of leaving that responsibility with the TP Framework. Once a servant is known to the TP Framework (through `Server::create_servant` or `TP::create_active_object_reference`), the TP Framework's default behavior is to dispose of that servant itself. In this case, the application code must no longer use references to the servant after deactivation.

However, the application may tell the TP Framework not to dispose of the servant (not to delete or re-use it) after the TP Framework deactivates it. Taking responsibility for a servant is done on an individual servant basis, not for a whole class of servants, by calling `TP::application_responsibility` with a parameter identifying the servant. In this case, the TP Framework does nothing further with the servant; the TP Framework does not delete, save, or make any further references to the servant.

The advantage of taking responsibility for the servant is that the servant does not have to be created anew. If obtaining the servant is an expensive proposition, the application may choose to save the servant and re-use it later. This is especially likely to be true for servants for preactivated objects, but is true in general. For example, the next time the TP Framework makes a call on `Server::create_servant`, the application might return a previously saved servant. It should be remembered that any time a servant is given to the TP Framework (even if it had been previously saved) the TP Framework assumes it has responsibility. Thus, even if the application saved the servant one time after giving the servant to the TP Framework, if the application gives the servant to the TP Framework again and want to save the servant again, the application must again call `TP::application_responsibility` to save the servant after that use.

Once an application has taken responsibility for a servant, the application must take care to delete the servant when the servant is no longer needed, the same as for any other C++ instance.

The `TP::application_responsibility` call can only be used after the TP Framework has possession of the servant. It cannot be used, for example, during the servant's `activate_object` callback because the TP Framework does not yet know about the servant (the servant has not been returned yet).

Saving and Restoring Object State

While CORBA objects are active, their state is contained in a servant. Unless an application uses `TP::create_active_object_reference`, state must be initialized when the object is first invoked (that is, the first time a method is invoked on a CORBA object after its object reference is created), and on subsequent invocations after they have been deactivated. While a CORBA object is deactivated, its state must be saved outside the process in which the servant was active. The object's state can be saved in shared memory, in a file, or in a database. Before a CORBA object is deactivated, its state must be saved, and when it is activated, its state must be restored.

The programmer determines what constitutes an object's state and what must be saved before an object is deactivated, and restored when an object is activated.

NOTE ON USE OF CONSTRUCTORS AND DESTRUCTORS FOR CORBA OBJECTS

The state of CORBA objects must not be initialized, saved, or restored in the constructors or destructors for the servant classes. This is because the TP Framework may reuse an instance of a servant rather than deleting it at deactivation. No guarantee is made as to the timing of the creation and deletion of servant instances.

Transactions

The following sections provide information about transaction policies and how to use transactions.

Transaction Policies

Eligibility of CORBA objects to participate in global transactions is controlled by the transaction policies assigned to implementations at compile time. The following policies can be assigned.

Note: The transaction policies are set in an ICF file that is configured at OMG IDL compile time. For a description of the ICF file, refer to Chapter 2, “Implementation Configuration File (ICF).”

◆ **never**

The implementation is not transactional. Objects created for this interface can never be involved in a transaction. The system generates an exception (`INVALID_TRANSACTION`) if an implementation with this policy is involved in a transaction. An `AUTOTRAN` policy specified in the `UBBCONFIG` file for the interface is ignored.

◆ **ignore**

The implementation is not transactional. This policy instructs the system to allow requests within a transaction to be made of this implementation. An `AUTOTRAN` policy specified in the `UBBCONFIG` file for the interface is ignored.

◆ **optional** (This is the default `transaction_policy`.)

The implementation may be transactional. Objects can be involved in a transaction if the request is transactional. Servers containing transactional objects must be configured within a group associated with an XA-compliant resource manager. If the `AUTOTRAN` parameter is specified in the `UBBCONFIG` file for the interface, `AUTOTRAN` is on.

◆ **always**

The implementation is transactional. Objects are required to always be involved in a transaction. If a request is made outside a transaction, the system automatically starts a transaction before invoking the method. The transaction is committed when the method ends. (This is the same behavior that results from specifying `AUTOTRAN` for an object with the option transaction policy, except that no administrative configuration is necessary to achieve this behavior, and it cannot be overridden by administrative configuration.) Servers containing transactional objects must be configured within a group that is associated with an XA-compliant resource manager.

Note: The `optional` policy is the only transaction policy that can be influenced by administrative configuration. If the system administrator sets the `AUTOTRAN` attribute for the interface by means of the `UBBCONFIG` file or by using administrative tools, the system automatically starts a transaction upon invocation of the object, if it is not already infected with a transaction (that is, the behavior is as if the `always` policy were specified).

Transaction Initiation

Transactions are initiated in one of two ways:

- ◆ By the application code via use of the `CoTransactions::Current::begin()` operation. This can be done in either the client or the server. For a description of this operation, refer to the section “Current Interface” on page 7-10.
- ◆ By the system when an invocation is done on an object which has either:
 - ◆ Transaction policy `always`
 - ◆ Transaction policy `optional` and a setting of `AUTOTRAN` for the interface

For more information, refer to the *Administration Guide*.

Transaction Termination

In general, the handling of the outcome of a transaction is the responsibility of the initiator. Therefore, the following are true:

- ◆ If the client or server application code initiates transactions, the TP Framework never commits a transaction. The WLE system may roll back the transaction if server processing tries to return to the client while the transaction is in an illegal state.
- ◆ If the system initiates a transaction, the commit or rollback will always be handled by the WLE system.

The following behavior is enforced by the WLE system:

- ◆ If no transaction is active when a method on a CORBA object is invoked and that method begins a transaction, the transaction must be either committed,

rolled back, or suspended when the method invocation returns. If none of these actions is taken, the transaction is rolled back by the TP Framework, and the `CORBA::OBJ_ADAPTER` exception is raised to the client application. This exception is raised because the transaction was initiated in the server application; therefore, the client application would not expect a transactional error condition such as `TRANSACTION_ROLLEDBACK`.

Transaction Suspend and Resume

The CORBA object must follow strict rules with respect to suspending and resuming a transaction within a method invocation. These rules and the error conditions that result from their violation are described below.

When a CORBA object method begins execution, it can be in one of the following three states with respect to transactions:

- ◆ The client application began the transaction.
 - ◆ *Legal server application behavior:* Suspend and resume the transaction within the method execution.
 - ◆ *Illegal server application behavior:* Return from the method with the transaction in the suspended state (that is, return from the method without invoking resume if suspend was invoked).
 - ◆ *Error Processing:* If illegal behavior occurs, the TP Framework raises the `CORBA::TRANSACTION_ROLLEDBACK` exception to the client application and the transaction is rolled back by the WLE system.
- ◆ The system began a transaction to provide `AUTOTRAN` or transaction policy always behavior.

Note: For each CORBA interface, set `AUTOTRAN` to `Yes` if you want a transaction to start automatically when an operation invocation is received. Setting `AUTOTRAN` to `Yes` has no effect if the interface is already in transaction mode. For more information about `AUTOTRAN`, refer to the *Administration Guide*.

- ◆ *Legal server behavior:* Suspend and resume the transaction within the method execution.

Note: Not recommended. The transaction may be timed out and aborted before the method causes the transaction to be resumed.

- ◆ *Illegal server behavior:* Return from the method with the transaction in the suspended state (that is, return from the method without invoking resume if suspend was invoked).
- ◆ *Error Processing:* If illegal behavior occurs, the TP Framework raises the CORBA::OBJ_ADAPTER exception to the client, and the transaction is rolled back by the system. The CORBA::OBJ_ADAPTER exception is raised because the client application did not initiate the transaction, and, therefore, does not expect transaction error conditions to be raised.
- ◆ The CORBA object is not involved in a transaction when it starts executing.
- ◆ *Legal server behavior:*
 - ◆ Begin and commit a transaction within the method execution.
 - ◆ Begin and roll back a transaction within the method execution.
 - ◆ Begin and suspend a transaction within the method execution.
- ◆ *Illegal server behavior:* Begin a transaction and return from the method with the transaction active.
- ◆ *Error Processing:* If illegal behavior occurs, the TP Framework raises the CORBA::OBJ_ADAPTER exception to the client application and the transaction is rolled back by the WLE system. The CORBA::OBJ_ADAPTER exception is raised because the client application did not initiate the transaction, and, therefore, does not expect transaction error conditions to be raised.

Restrictions on Transactions

The following restrictions apply to WLE transactions:

- ◆ A CORBA object in the WLE system must have the same transaction context when it returns from a method invocation that it had when the method was invoked.
- ◆ A CORBA object can be infected by only one transaction at a time. If an invocation tries to infect an already infected object, a CORBA::INVALID_TRANSACTION exception is returned.
- ◆ If a CORBA object is infected with a transaction and a nontransactional request is made on it, a CORBA::OBJ_ADAPTER exception is raised.

- ◆ If the application begins a transaction in `Server::initialize()`, it must either commit or roll back the transaction before returning from the method. If the application does not, the TP Framework shuts down the server. This is because the application has no predictable way of regaining control after completing the `Server::initialize` method.
- ◆ If a CORBA object is infected by a transaction and with an activation policy of transaction, and if the reason code passed to the method is either `DR_TRANS_COMMITTING` or `DR_TRANS_ABORTED`, no invocation on any CORBA object can be done from within the `Tobj_ServantBase::deactivate_object` method. Such an invocation results in a `CORBA::BAD_INV_ORDER` exception.

SQL and Global Transactions

Adhere to the following guidelines when using SQL and Global Transactions:

- ◆ Care should be taken when executing SQL statements outside the scope of a global transaction. The SQL standard specifies that a local transaction should be started implicitly by the database resource manager whenever an SQL statement that needs the context of a transaction is executed and no transaction is active. The standard also says that a transaction that is implicitly started by the database resource manager must then be explicitly terminated by executing a `COMMIT` or `ROLLBACK` SQL statement; the TP Framework is not responsible for terminating transactions that are started by the resource manager.

Note: This is not an issue when an application uses the XA library to connect to the Oracle server because those applications can operate only on global transactions. The Oracle server does not allow local transactions when it is using XA.
- ◆ The SQL `COMMIT` and `ROLLBACK` statements cannot be used to terminate a global transaction that has been either started explicitly using `Current.begin()` or started implicitly by the system. Check the database vendor documentation for each database product for other possible restrictions when using global transactions.
- ◆ SQL cursors may be closed when transactions are terminated. Consult your database product documentation for exact information about cursor handling rules. Application programmers should be careful to use cursors only with

CORBA objects with appropriate activation policies and within appropriate transaction boundaries.

Voting on Transaction Outcome

CORBA objects can affect transaction outcome during two stages of transaction processing:

◆ *During transactional work*

The `Current.rollback_only` method can be used to ensure that the only possible outcome is to roll back the current transaction. `Current.rollback_only()` can be invoked from any CORBA object method.

◆ *After completion of transactional work*

CORBA objects that have the transaction activation policy are given a chance to vote whether the transaction should commit or roll back after transactional work is completed. These objects are notified of the completion of transactional work prior to the start of the two-phase commit algorithm when the TP Framework invokes their `deactivate_object` method.

Note that this behavior does not apply to objects with `process` or `method` activation policies. If the CORBA object wants to roll back the transaction, it can call `Current::rollback_only`. If it wants to vote to commit the transaction, it does not make that call. Note, however, that a vote to commit does not guarantee that the transaction is committed, since other objects may subsequently vote to roll back the transaction.

Note: Users of SQL cursors must be careful when using an object with the `method` or `process` activation policy. A process opens an SQL cursor within a client-initiated transaction. For typical SQL database products, once the client commits the transaction, all cursors that were opened within that transaction are automatically closed; however, the object will not receive any notification that its cursor has been closed.

Transaction Time-outs

When a transaction time-out occurs, the transaction is marked so that the only possible outcome is to roll back the transaction, and the `CORBA::TRANSACTION_ROLLEDBACK` standard exception is returned to the client. Any attempts to send new requests will also fail with the `CORBA::TRANSACTION_ROLLEDBACK` exception until the transaction has been aborted.

TP Framework API

This section describes the TP Framework API. Additional information about how to use this API can be found in *Creating C++ Server Applications*.

The TP Framework comprises the following components:

- ◆ The `Server C++` class, which has virtual methods for application-specific server initialization and termination logic
- ◆ The `Tobj_ServantBase C++` class, which has virtual methods for object state management
- ◆ The `TP C++` class, which provides methods to:
 - ◆ Create object references for CORBA objects
 - ◆ Register (and unregister) factories with the `FactoryFinder` object
 - ◆ Initiate user-controlled preactivation and deactivation of objects
 - ◆ Initiate user-controlled deactivation of the CORBA object currently being invoked
 - ◆ Obtain an object reference to the CORBA object currently being invoked
 - ◆ Open and close XA resource managers
 - ◆ Log messages to a user log (`ULOG`) file
 - ◆ Obtain object references to the ORB and to Bootstrap objects
- ◆ Header files for these classes

- ◆ Libraries that are used by server applications

The parts of the TP Framework that are visible to programmers consist of two categories of operations:

- ◆ Service methods that can be called by user application code. These are in the TP interface.
- ◆ Callback methods that are written by the user and that are invoked by the TP Framework. This includes methods in the `Tobj_ServantBase` and `Server` classes. These operations are intended to be called by TP Framework code only. Application code should never call methods of these classes. If such calls are made, unpredictable results may occur.

Server Interface

The Server interface provides callback methods that can be used for application-specific server initialization and termination logic. This interface also provides a callback method that is used to create servants when servants are required for object activation.

The Server interface has the following characteristics:

- ◆ The `Server` class is a C++ native class.
- ◆ The `Server.h` file contains the declarations and definitions for the `Server` class.

C++ DECLARATIONS

The C++ mapping is as follows:

```
typedef Tobj_ServantBase* Tobj_Servant;

class Server {
public:
    CORBA::Boolean    initialize(int argc, char** argv);
    void              release();
    Tobj_Servant       create_servant(const char* interfaceName);
};
```

Note: Programmers must provide definitions for the `Server::initialize()`, `Server::release()`, and `Server::create_servant` methods.

Server::create_servant

Synopsis Creates a servant to instantiate a C++ object.

C++ Binding

```
class Server {
public:
    Tobj_Servant      create_servant(const char* interfaceName);
};
```

Argument *interfaceName*
 Specifies a character string that contains the fully qualified interface name for the object. This will be the same interface name that was supplied when the object reference was created (`TP::create_object_reference()` or `TP::create_active_object_reference()`) for the object reference used for this invocation. This name can be used to determine which servant needs to be constructed.

Return Value *Tobj_ServantBase*
 The pointer to the newly created servant (instance) for the specified interface. A NULL value should be returned if `create_servant()` is invoked with an interface name that it does not recognize or if the servant cannot be created for some reason.
 If the `create_servant` method returns a NULL pointer, activation fails. A `CORBA::OBJECT_NOT_EXIST()` exception is raised back to the client. Also, the following message is written to the user log (ULOG):

```
"TPFW_CAT:23: ERROR: Activating object - application raised
TobjS::CreateServantFailed. Reason = Application's
Server::create_servant returned NULL. Interface =
interfaceName, OID = oid"
```

Where *interfaceName* is the interface ID of the invoked interface and *oid* is the corresponding object ID.

Note: The restriction on the length of the `ObjectId` has been removed in this release.

Description The `create_servant` method is invoked by the TP Framework when a request arrives at the server and there is no available servant to satisfy the request. The TP Framework calls the `create_servant` method with the interface name for the servant to be created. The server application instantiates an appropriate C++ object and returns a pointer to it. Typically, the method contains a switch statement on the interface name and creates a new object, depending on the interface name.

Caution: The server application must not depend on this method being invoked for every activation of a CORBA object. The server application must not do any handling of CORBA object state in the constructors or destructors of any servant classes for CORBA objects. This is because the TP Framework may possibly reuse servants on activation and may possibly not destroy servants on deactivation.

Exception If an exception is thrown in `Server::create_servant()`, the TP Framework catches the exception. Activation fails. A `CORBA::OBJECT_NOT_EXIST()` exception is raised back to the client. In addition, an error message is written to the user log (ULOG) file, as follows, for each exception type:

```
TobjS::CreateServantFailed
```

```
"TPFW_CAT:23: ERROR: Activating object - application
raised TobjS::CreateServantFailed. Reason = reason.
Interface = interfaceName, OID = oid"
```

Where *reason* is a user-supplied reason, and *interfaceName* and *oid* are the interface ID and object ID, respectively, of the invoked CORBA object.

```
TobjS::OutOfMemory
```

```
"TPFW_CAT:22: ERROR: Activating object - application
raised TobjS::OutOfMemory. Reason = reason. Interface
= interfaceName, OID = oid"
```

Where *reason* is a user-supplied reason, and *interfaceName* and *oid* are the interface ID and object ID, respectively, of the invoked CORBA object.

```
CORBA::Exception
```

```
"TPFW_CAT:28: ERROR: Activating object - CORBA
Exception not handled by application. Exception ID =
exceptionID. Interface = interfaceName, OID = oid"
```

Where *exceptionID* is the interface ID of the exception, and *interfaceName* and *oid* are the interface ID and object ID, respectively, of the invoked CORBA object.

Other Exception

```
"TPFW_CAT:29: ERROR: Activating object - Unknown  
Exception not handled by application. Exception ID =  
exceptionID. Interface = interfaceName, OID = oid"
```

Where *exceptionID* is the interface ID of the exception, and *interfaceName* and *oid* are the interface ID and object ID, respectively, of the invoked CORBA object.

Server::initialize()

Synopsis Allows the application to perform application-specific initialization procedures, such as logging into a database, creating and registering well-known object factories, initializing global variables, and so forth.

C++ Binding

```
class Server {
public:
    CORBA::Boolean initialize(int argc, char** argv);
};
```

Arguments The `argc` and `argv` arguments are passed from the command line. The `argc` argument contains the name of the server. The `argv` argument contains the first command-line option that is specific to the application, if there are any.

Command line options are specified in the `UBBCONFIG` file using the `CLOPT` parameter in the entry for the server in the `SERVERS` section. System-recognized options come first in the `CLOPT` parameter, followed by a double-hyphen (`--`), followed by the application-specific options. The value of `argc` is one greater than the number of application-specific options. For details, see `ubbconfig(5)` in the *BEA TUXEDO Reference*.

Return Value Boolean `TRUE` or `FALSE`. `TRUE` indicates success. `FALSE` indicates failure. If an error occurs in `initialize()`, the application code should return `FALSE`. The application code should not call the system call `exit()`. Calling `exit()` does not give the TP Framework a chance to release resources allocated during startup and may cause unpredictable results.

If the return value is `FALSE`:

- ◆ `Server::release()` is not invoked.
- ◆ Any transactions that are started in the `initialize()` method and are not terminated will eventually time out; they are not automatically rolled back.

Description The `initialize` callback method, which is invoked as the last step in server initialization, allows the application to perform application-specific initialization.

Typically, a server application does the following tasks in `Server::initialize`:

- ◆ Creates references for CORBA object factories implemented in the server application and registers them with the `FactoryFinder` using the `TP::register_factory()` operation.

- ◆ Initializes global variables, if any are used.
- ◆ Opens XA resource managers if any are used by the server application.

It is the responsibility of the server application to open any required XA resource managers. This is done by invoking either of the following methods:

- ◆ `TP::open_xa_rm()`
This is the preferred technique for server applications, since it can be done on a static function, without the need to obtain an object reference.
- ◆ `Tobj::TransactionCurrent::open_xa_rm()`
A reference to the `TransactionCurrent` object can be obtained from the `Bootstrap` object. For an explanation of how to obtain a reference to the `Bootstrap` object, see the section “`TP::bootstrap()`” on page 3-47. For more information about the `TransactionCurrent` object, see Chapter 4, “`Bootstrap Object`,” and Chapter 7, “`Transaction Service`.”
- ◆ Transactions may be started in the `initialize` method after invoking the `Tobj::TransactionCurrent::open_xa_rm()` or `TP::open_xa_rm` method. However, any transactions that are started in `initialize()` must be terminated by the server application before `initialize()` returns. If the transactions are still active when control is returned, the server application fails to boot, and it exits gracefully. This happens because the server application has no logical way of either committing or rolling back the transaction after `Server::initialize()` returns. This condition is an error.

Exceptions If an exception is raised in `Server::initialize()`, the TP Framework catches the exception. The TP Framework behavior is the same as if `initialize()` returned `FALSE` (that is, an exception is considered to be a failure). In addition, an error message is written to the user log (ULOG) file, as follows, for each exception type:

```
TobjS::InitializeFailed
```

```
"TPFW_CAT:1: ERROR: Exception in
Server::initialize():IDL:beasys.com/TobjS/Initialize
Failed:1.0. Reason = reason"
```

Where *reason* is a string supplied by application code. For example:

```
Throw TobjS::InitializeFailed(
    "Couldn't register factory");
```

CORBA::Exception

```
"TPFW_CAT:1: ERROR: Exception in  
Server::initialize(): exception. Reason = unknown"
```

Where *exception* is the interface ID of the CORBA exception that was raised.

Other Exceptions

```
TPFW_CAT:1: ERROR: Exception in Server::initialize():  
unknown exception. Reason = unknown"
```


Server::release()

Synopsis	Allows the application to perform any application-specific cleanup, such as logging off a database, unregistering well-known factories, or deallocating resources.
C++ Binding	<pre>typedef Tobj_ServantBase* Tobj_Servant; class Server { public: void release(); };</pre>
Arguments	None.
Return Value	None.
Description	The <code>release</code> callback method, which is invoked as the first step in server shutdown, allows the server application to perform any application-specific cleanup. The user must override the virtual function definition.

Typical tasks performed by the application in this method are as follows:

- ◆ Close XA resource managers.
- ◆ Unregister CORBA object factories that were registered with the Factory Finder in `Server::initialize()`.
- ◆ Deallocate any server resources not yet released.

This method is normally called in response to a `tmshutdown` command from the administrator or operator.

The TP Framework provides a default implementation of `Server::release()`. The default implementation closes XA resource managers for the server. The implementation does this by issuing a `tx_close()` invocation, which uses the default `CLOSEINFO` configured for the server's group in the `UBBCONFIG` file.

It is the responsibility of the application to close any open XA resource managers. This is done by issuing either of the following calls:

- ◆ `TP::close_xa_rm`
- ◆ `Tobj::TransactionCurrent::close_xa_rm()`. A reference to the `TransactionCurrent` object can be obtained from the Bootstrap object. For an explanation of how to obtain a reference to the Bootstrap object, see the section “TP::bootstrap()” on page 3-47. For more information about the

TransactionCurrent object, see Chapter 4, “Bootstrap Object,” and Chapter 7, “Transaction Service.”

Note: Once a server receives a request from the `tmshutdown(1)` command to shut down, it can no longer receive requests from other remote objects. This may require servers to be shut down in a specific order. For example, if the `Server::release()` method in Server 1 needs to access a method of an object that resides in Server 2, Server 2 should be shut down after Server 1 is shut down. In particular, the `TP::unregister_factory()` method accesses the FactoryFinder Registrar object that resides in a separate server. The `TP::unregister_factory()` method is typically invoked from the `release()` method; therefore, the FactoryFinder server should be shut down after all servers that call `TP::unregister_factory()` in their `Server::release()` method.

Exceptions If an exception is raised in `release()`, the TP Framework catches the exception. Each exception causes an error message to be written to the user log (ULOG) file, as follows:

```
TobjS::ReleaseFailed
```

```
"TPFW_CAT:2: WARN: Exception in Server::release():
IDL:beasys.com/TobjS/ReleaseFailed:1.0. Reason =
reason"
```

Where *reason* is a string supplied by application code. For example:

```
Throw TobjS::ReleaseFailed(
    "Couldn't unregister factory");
```

```
CORBA::Exception
```

```
"TPFW_CAT:2: WARN: Exception in Server::release():  
exception. Reason = unknown"
```

Where `exception` is the interface ID of the CORBA exception that was raised.

```
Other Exceptions
```

```
"TPFW_CAT:2: WARN: Exception in Server::release():  
unknown exception. Reason = unknown"
```

In all cases, the server continues to exit.

Tobj_ServantBase Interface

The `Tobj_ServantBase` interface defines operations that allow a CORBA object to assist in the management of its state. Every implementation skeleton generated by the IDL compiler automatically inherits from the `Tobj_ServantBase` class. The `Tobj_ServantBase` class contains two virtual methods, `activate_object()` and `deactivate_object()`, that may be optionally implemented by the programmer.

Whenever a request comes in for an inactive CORBA object, the object is activated and the `activate_object()` method is invoked on the servant. When the CORBA object is deactivated, the `deactivate_object()` method is invoked on the servant. The timing of deactivation is driven by the implementation's activation policy. When the `deactivate_object()` method is invoked, the TP Framework passes in a reason code to indicate why the call was made.

Note: `Tobj_ServantBase::activate_object()` and `Tobj_ServantBase::deactivate_object()` are the only methods that the TP Framework guarantees will be invoked for CORBA object activation and deactivation. The servant class constructor and destructor may or may not be invoked at activation or deactivation time (through the `Server::create_servant` call for C++ or directly by Java). Therefore, the server application code must not do any state handling for CORBA objects in either the constructor or destructor of the servant class.

Note: The programmer does not need to use a cast or reference to `Tobj_ServantBase` directly. The `Tobj_ServantBase` methods show up as part of the skeleton and, therefore, in the implementation class for a servant. The programmer may provide definitions for the `activate_object` and `deactivate_object` methods, but the programmer should never make direct invocations on those methods; only the TP Framework should call those methods.

C++ DECLARATION (IN `Tobj_ServantBase.h`)

The C++ mapping for the `Tobj_servantBase` interface is as follows:

```
class Tobj_ServantBase : public PortableServer::ServantBase {
public:
    virtual void activate_object(const char * stroid) {}
    virtual void deactivate_object(const char*,
                                   TobjS::DeactivateReasonValue) {}
};
```

Tobj_ServantBase:: activate_object()

Synopsis Associates an object ID with a servant. This method gives the application an opportunity to restore the object's state when the object is activated. The state may be restored from shared memory, from an ordinary flat file, or from a database file.

C++ Binding

```
class Tobj_ServantBase : public PortableServer::ServantBase {
public:
    virtual void activate_object(const char * stroid) {}
};
```

Argument `stroid`

Specifies the object ID in string format. The object ID uniquely identifies this instance of the class. This is the same object ID that was specified when the object reference was created (using `TP::create_object_reference()`) or in the `TP::create_active_object_reference()` for the object reference used for this invocation.

Note: The restriction on the length of the object ID has been removed in this release.

Return Value None.

Description Object activation is triggered by a client invoking a method on an inactive CORBA object. This causes the Portable Object Adapter (POA) to assign a servant to the CORBA object. The `activate_object()` method is invoked before the method invoked by the client is invoked. If `activate_object()` returns successfully, that is, without raising an exception, the requested method is executed on the servant.

The `activate_object()` and `deactivate_object()` methods and the method invoked by the client can be used by the programmer to manage object state. The particular way these methods are used to manage object state may vary according to the needs of the application. For a discussion of how these methods might be used, see *Creating C++ Server Applications*.

If the object is currently infected with a global transaction, `activate_object()` executes within the scope of that same global transaction.

It is the responsibility of the programmer of the object to check that the stored state of the object is consistent. In other words, it is up to the application code to save a persistent flag that indicates whether or not `deactivate_object()` successfully saved the state of the object. That flag should be checked in `activate_object()`.

Exceptions If an error occurs while executing `activate_object()`, the application code should raise either a CORBA standard exception or a `TobjS::ActivateObjectFailed` exception. When an exception is raised, the TP Framework catches the exception, and the following events occur:

- ◆ The activation fails.
- ◆ The method invoked by the client is not executed.
- ◆ If `activate_object()` is executing within a transaction and the client initiated the transaction, the transaction is *not* rolled back.
- ◆ A `CORBA::OBJECT_NOT_EXIST` exception is raised back to the client.

Note: For each CORBA interface, set `AUTOTRAN` to `Yes` if you want a transaction to start automatically when an operation invocation is received. Setting `AUTOTRAN` to `Yes` has no effect if the interface is already in transaction mode. For more information about `AUTOTRAN`, refer to the *Administration Guide*.

- ◆ Based on the exception is raised, a message is written to the user log (ULOG) file, as follows:

```
TobjS::ActivateObjectFailed
```

```
"TPFW_CAT:24: ERROR: Activating object - application
raised TobjS::ActivateObjectFailed. Reason = reason.
Interface = interfaceName, OID = oid"
```

Where *reason* is a user-supplied reason, and *interfaceName* and *oid* are the interface ID and object ID, respectively, of the invoked CORBA object.

```
TobjS::OutOfMemory
```

```
"TPFW_CAT:22: ERROR: Activating object - application
raised TobjS::OutOfMemory. Reason = reason. Interface
= interfaceName, OID = oid"
```

Where *reason* is a user-supplied reason, and *interfaceName* and *oid* are the interface ID and object ID, respectively, of the invoked CORBA object.

CORBA::Exception

```
"TPFW_CAT:25: ERROR: Activating object - CORBA  
Exception not handled by application. Exception ID =  
exceptionID. Interface = interfaceName, OID = oid"
```

Where *exceptionID* is the interface ID of the exception, and *interfaceName* and *oid* are the interface ID and object ID, respectively, of the invoked CORBA object.

Other exception

```
"TPFW_CAT:26: ERROR: Activating object - Unknown  
Exception not handled by application. Exception ID =  
exceptionID. Interface = interfaceName, OID = oid"
```

Where *exceptionID* is the interface ID of the exception, and *interfaceName* and *oid* are the interface ID and object ID, respectively, of the invoked CORBA object.

Tobj_ServantBase::deactivate_object()

Synopsis Removes the association of an object ID with its servant. This method gives the application an opportunity to save all or part of the object's state before the object is deactivated. The state may be saved in shared memory, in an ordinary flat file, or in a database file.

C++ Binding

```
class Tobj_ServantBase : public PortableServer::ServantBase {
public:
    virtual void deactivate_object(const char* stroid,
                                  TobjS::DeactivateReasonValue reason) {}
};
```

Arguments `stroid` Specifies the object ID in string format. The object ID uniquely identifies this instance of the class.

Note: The restriction on the length of the object ID has been removed in this release.

`reason` Indicates the event that caused this method to be invoked. The `reason` code can be one of the following:

`DR_METHOD_END`

Indicates that the object is being deactivated after the completion of a method. It is used if the object's deactivation policy is:

- ◆ `method`
- ◆ `transaction` (only if there is no transaction in effect)
- ◆ `process` (if `TP::deactivateEnable()` called)

`DR_SERVER_SHUTDOWN`

Indicates that the object is being deactivated because the server is being shut down in an orderly fashion. It is used if the object's deactivation policy is:

- ◆ `transaction` (only if transaction is active)
- ◆ `process`

Note that when a server is shut down in an orderly fashion, all transactions that the server is involved in are marked for rollback. For more information about restrictions on processing that this causes, see the section "Description" on page 3-39.

DR_TRANS_ABORTED

This reason code is used only for objects that have the transaction activation policy. It can occur when the transaction is started by either the client or automatically by the system. When the `deactivate_object()` method is invoked with this reason code, the transaction is marked for rollback only. For more information about restrictions about processing that this causes, see the section “Description” on page 3-39.

DR_TRANS_COMMITTING

This reason code is used only for objects that have the transaction activation policy. It can occur when the transaction is started by either the client or the TP Framework. It indicates that a `Current.commit()` operation was invoked for the transaction in which the object is involved. The `deactivate_object()` method is invoked just before the transaction manager’s two-phase commit algorithm begins; that is, before `prepare` is sent to the resource managers. For more information about restrictions on processing that this causes, see the section “Description” on page 3-39.

The CORBA object is allowed to vote on the outcome of the transaction when the `deactivate_object()` method is invoked with the `DR_TRANS_COMMITTING` reason code. By invoking `Current.rollback_only()`, the method can force the transaction to be rolled back; otherwise, the two-phase commit algorithm continues. The transaction is not necessarily committed just because the `Current.rollback_only()` is not invoked in this method. Any other CORBA object or resource manager involved in the transaction could also vote to roll back the transaction.

DR_EXPLICIT_DEACTIVATE

Indicates that the object is being deactivated because the application executed a `TP::deactivateEnable(-,-,-)` on this object. This can happen only for objects that have the process activation policy.

Return Value None.

Description Object deactivation is initiated either by the system or by the application, depending on the activation policy of the implementation for the CORBA object. The `deactivate_object()` method is invoked before the CORBA object is deactivated. For details of these policies and their use, see the section “ICF Syntax” on page 2-2.

Deactivation may occur after an execution of a method invoked by a client if the activation policy for the CORBA object implementation is `method`, or as a result of the end of transactional work if the activation policy is `transaction`. It may also occur as the result of server shutdown if the activation policy is `transaction` or `process`.

In addition, the WLE software supports the use of user-controlled deactivation of CORBA objects having an activation policy of `process` or `method` via the use of the `TP::deactivateEnable()` and `TP::deactivateEnable(-,-,-)` methods. `TP::deactivateEnable` can be called inside a method of an object to cause the object to be deactivated at the end of the method. If `TP::deactivateEnable` is called in an object with the `transaction` activation policy, an exception is raised (`TobjS::IllegalOperation`) and the TP Framework takes no action. `TP::deactivateEnable(-,-,-)` can be called to deactivate any object that has a `process` activation policy. For more information, see the section “`TP::deactivateEnable`” on page 3-56.

Note: The `deactivate_object` method will be called at server shutdown time for every object remaining in the Active Object Map, whether it was entered there implicitly by the TP Framework (the activation-on-demand technique: `TP::create_servant` and the servant’s `activate_object` method) or explicitly by the user with `TP::create_active_object_reference`.

The `activate_object()` and `deactivate_object()` methods and explicit methods invoked by the client can be used by the programmer to manage object state. The manner in which these methods are used to manage object state may vary according to the needs of the application. For a discussion of how these methods might be used, see *Creating C++ Server Applications*.

The CORBA object with `transaction` activation policy gets to vote on the outcome of the transaction when the `deactivate_object()` method is invoked with the `DR_TRANS_COMMITTING` reason code. By calling `Current.rollback_only()` the method can force the transaction to be rolled back; otherwise, the two-phase commit algorithm continues. The transaction will not necessarily be committed just because `Current.rollback_only()` is not called in this method. Any other CORBA object or resource manager involved in the transaction could also vote to roll back the transaction.

Restriction Note that if the object is involved in a transaction when this method is invoked, there are restrictions on what type of processing can be done based on the reason the object is invoked. If the object was involved in a transaction, the activation policy is `transaction` and the reason code for the call is:

DR_TRANS_ABORTED

No invocations on any CORBA objects are allowed in the method. No `tpcall()` is allowed. Transactions cannot be suspended or begun.

DR_TRANS_COMMITTING

No invocations on any CORBA objects are allowed in the method. No `tpcall()` is allowed. Transactions cannot be suspended or begun.

The reason for these restrictions is that the deactivation of objects with activation policy transaction is controlled by a call to the TP Framework from the transaction manager for the transaction. When the call with reason code `DR_TRANS_COMMITTING` is made, the transaction manager is executing phase 1 (prepare) of the two-phase commit. At this stage, it is not possible to issue a call to suspend a transaction nor to initiate a new transaction. Since a call to a CORBA object that was in another process would require that process to join the transaction, and the transaction manager is already executing the prepare phase, this would cause an error¹. Since a call to a CORBA object that had no transactional properties would require that the current transaction be suspended, this would also cause an error. The same is true of a `tpcall()`.

Similarly, when the invocation with reason code `DR_TRANS_ABORTED` is made, the transaction manager is already aborting. While the transaction manager is aborting, it is not possible to either suspend a transaction or initiate a new transaction. The same restrictions apply as for `DR_TRANS_COMMITTING`.

1. In theory, this would mean that an invocation on a transactional CORBA object in the same process would be valid since it would not require a new process to be registered with the transaction manager. However, it is not possible for the programmer to guarantee that an invocation on a CORBA object will occur in-proc, therefore, this practice is discouraged.

Exceptions If the CORBA object method that is invoked by the client raises an exception, that exception is caught by the TP Framework and is eventually returned to the client. This is true even if `deactivate_object()` is invoked and raises an exception.

The client will never be notified about exceptions that are raised in `deactivate_object()`. It is the responsibility of the application code to check that the stored state of the CORBA object is consistent. For example, the application code could save a persistent flag that indicates whether or not `deactivate_object()` successfully saved the state. That flag can then be checked in `activate_object()`.

If an error occurs while executing `deactivate_object()`, the application code should raise either a CORBA standard exception or a `DeactivateObjectFailed` exception. If `deactivate_object()` was invoked by the TP Framework, the TP Framework catches the exception and the following events occur:

- ◆ The object is deactivated.
- ◆ If the client initiated a transaction, the transaction is not rolled back.
- ◆ The client is not notified of the exception that is raised in `deactivate_object()`.
- ◆ Based on which exception is raised, a message is logged to the user log (ULOG) file, as follows:

```
TobjS::DeactivateObjectFailed
```

```
"TPFW_CAT:27: ERROR: De-activating object -
application raised TobjS::DeactivateObjectFailed.
Reason = reason. Interface = interfaceName, OID = oid"
```

Where *reason* is a user-supplied reason, and *interfaceName* and *oid* are the interface ID and object ID, respectively, of the invoked CORBA object.

```
CORBA::Exception
```

```
"TPFW_CAT:28: ERROR: De-activating object - CORBA
Exception not handled by application. Exception ID =
exceptionID. Interface = interfaceName, OID = oid"
```

Where *exceptionID* is the interface ID of the exception, and *interfaceName* and *oid* are the interface ID and object ID, respectively, of the invoked CORBA object.

Other exception

```
"TPFW_CAT:29: ERROR: De-activating object - Unknown  
Exception not handled by application. Exception ID =  
exceptionID. Interface = interfaceName, OID = oid"
```

Where *exceptionID* is the interface ID of the exception, and *interfaceName* and *oid* are the interface ID and object ID, respectively, of the invoked CORBA object.

TP Interface

The TP interface supplies a set of service methods that can be invoked by application code. This is the *only* interface in the TP Framework that can safely be invoked by application code. All other interfaces have callback methods that are intended to be invoked only by system code.

The purpose of this interface is to provide high-level calls that application code can call, instead of calls to underlying APIs provided by the Portable Object Adapter (POA), the CORBA Naming Service, and the BEA TUXEDO system. By using these calls, programmers can learn a simpler API and are spared the complexity of the underlying APIs.

The TP interface implicitly uses two features of the WLE software that extend the CORBA APIs:

- ◆ Factories and the FactoryFinder object
- ◆ Factory-based routing

For information about the FactoryFinder object, see Chapter 5, “FactoryFinder Interface.” For more information about Factory-based routing, see the *Administration Guide*.

USAGE NOTES

- ◆ During server application initialization, the application constructs the object reference for an application factory. It then invokes the `register_factory()` method, passing in the factory's object reference together with a factory `id` field. On server release (shutdown), the application uses the `unregister_factory()` method to unregister the factory.
- ◆ The `TP` class is a C++ native class.
- ◆ The `TP.h` file contains the declarations and definitions for the `TP` class.

C++ Declarations (in `TP.h`)

The C++ mapping is as follows:

```
class TP {
public:
    static CORBA::Object_ptr create_object_reference(
        const char* interfaceName,
        const char* stroid,
        CORBA::NVList_ptr criteria);
    static CORBA::Object_ptr create_active_object_reference(
        const char* interfaceName,
        const char* stroid,
        Tobj_Servant servant);
    static CORBA::Object_ptr get_object_reference();
    static void register_factory(
        CORBA::Object_ptr factory_or,
        const char* factory_id);
    static void unregister_factory(
        CORBA::Object_ptr factory_or,
        const char* factory_id);
    static void deactivateEnable();
    static void deactivateEnable(
        const char* interfaceName,
        const char* stroid,
        Tobj_Servant servant);
    static CORBA::ORB_ptr orb();
    static Tobj_Bootstrap* bootstrap();
    static void open_xa_rm();
    static void close_xa_rm();
    static int userlog(char*, ...);
    static char* get_object_id(CORBA::Object_ptr obj);
    static void application_responsibility(
        Tobj_Servant servant);
};
```

TP::application_responsibility

Synopsis	Tells the TP Framework that the application is taking responsibility for the servant's lifetime.
C++ Binding	<pre>static void application_responsibility(Tobj_Servant servant);</pre>
Arguments	<div>servant</div> <div>A pointer to a servant that is already known to the TP Framework.</div>
Return Values	None.
Description	<p>This method tells the TP Framework that the application is taking responsibility for the servant's lifetime. As a result of this call, when the TP Framework has completed deactivating the object (that is, after invoking the servant's <code>deactivate_object</code> method), the TP Framework does nothing more with the object.</p> <p>Once an application has taken responsibility for a servant, the application must take care to delete servant when it is no longer needed, the same as for any other C++ instance.</p> <p>If the servant is not known to the TP Framework (that is, it is not active), this call has no effect.</p>
Exceptions	<div><code>TobjS::InvalidServant</code></div> <div>Indicates that the specified servant is Null.</div>

TP::bootstrap()

Synopsis	Returns a pointer to a <code>Tobj::Tobj_Bootstrap</code> object. The Bootstrap object is used to access initial object references for the <code>FactoryFinder</code> object, the <code>Interface Repository</code> , the <code>TransactionCurrent</code> , and the <code>SecurityCurrent</code> objects.
C++ Binding	<code>static Tobj_Bootstrap* TP::bootstrap();</code>
Arguments	None.
Return Value	Upon successful completion, <code>bootstrap()</code> returns a pointer to the <code>Tobj::Tobj_Bootstrap</code> object that is created by the TP Framework when the server application is started.
Description	<p>The TP Framework creates a <code>Tobj::Tobj_Bootstrap</code> object as part of initialization; it is not necessary for the application code to create any other <code>Tobj::Tobj_Bootstrap</code> objects in the server.</p> <p>Caution: Because the TP Framework owns the <code>Tobj::Tobj_Bootstrap</code> object, server application code must not dispose of the Bootstrap object.</p>
Exceptions	None.

TP::close_xa_rm()

Synopsis	Closes the XA resource manager to which the invoking process is linked.
C++ Binding	<code>static void TP::close_xa_rm ();</code>
Arguments	None.
Return Values	None.
Description	The <code>close_xa_rm()</code> method closes the XA resource manager to which the invoking process is linked. XA resource managers are provided by database vendors, such as Oracle and Informix.

Note: The functionality of this call is also provided by `Tobj::TransactionCurrent::close_xa_rm()`. The `TP::close_xa_rm()` method provides a more convenient way for a server application to close a resource manager because there is no need to obtain an object reference to the `TransactionCurrent` object. A reference to the `TransactionCurrent` object can be obtained from the `Bootstrap` object. See “`TP::bootstrap()`” on page 3-47 for an explanation of how to obtain a reference to the `Bootstrap` object. For more information about the `TransactionCurrent` object, see Chapter 4, “Bootstrap Object,” and Chapter 7, “Transaction Service.”

This method should be invoked once from the `Server::release()` method for each server that is involved in global transactions. This includes servers that are linked with an XA resource manager, as well as servers that are involved in global transactions, but are not actually linked with an XA-compliant resource manager.

The `close_xa_rm()` method should be invoked in place of a close invocation that is specific to the resource manager. Because resource managers differ in their initialization semantics, the specific information needed to close a particular resource manager is placed in the `CLOSEINFO` parameter in the `GROUPS` section of the `WLE` system `UBBCONFIG` file.

The format of the `CLOSEINFO` string is dependent on the requirements of the database vendor providing the underlying resource manager. For more information about the `CLOSEINFO` parameter, see the *Administration Guide* and `ubbconfig(5)` reference page in the *BEA TUXEDO Reference*. Also, refer to database vendor documentation for information about how to develop and install applications that use the XA libraries.

Exceptions `CORBA::BAD_INV_ORDER`
 There is an active transaction. The resource manager cannot be closed while a transaction is active.

`Tobj::RMFailed`
 The `tx_close()` call returned an error return code.

Note: Unlike other exceptions returned by the TP Framework, the `Tobj::RMFailed` exception is defined in `tobj_c.h` (which is derived from `tobj.idl`), not `TobjS_c.h` (which is derived from `TobjS.idl`). This is because native clients can also open XA resource managers. Therefore, the exception returned is consistent with the exception expected by native client code and by `Server::release()` if it uses the alternate mechanism, `TransactionCurrent::close_xa_rm`, which is shared with native clients.

TP::create_active_object_reference()

Synopsis Creates an object reference and preactivates an object.

C++ Binding

```
static CORBA::Object_ptr create_active_object_reference(
                                const char*      interfaceName,
                                const char*      stroid,
                                Tobj_Servant     servant);
```

Arguments `interfaceName`

Specifies a character string that contains the fully qualified interface name for the object.

`stroid`

Specifies the `ObjectId` in string format. The `ObjectId` uniquely identifies this instance of the class. The programmer decides what information to place in the `ObjectId`. One possibility would be to use it to hold a database key. Choosing the value of an object identifier, and the degree of uniqueness, is part of the application design. The WLE software cannot guarantee any uniqueness in object references, since these may be legitimately copied and shared outside the WLE environment, for example by stringifying the object reference.

`servant`

A pointer to a servant that the application has already created and initialized.

Return Value The newly created object reference.

Description This method creates an object reference and preactivates an object. The resulting object reference may be passed to clients who will use it to access the object.

Ordinarily, the application will call this method in two places:

- ◆ In `Server::initialize()` to preactivate process objects so that they do not need activation on the first invocation
- ◆ In any method that creates object references to be returned to clients

This method allows an application to activate an object explicitly before its first invocation. (For reasons you might want to do this, refer to the section “Explicit Activation” on page 3-11.) The user first creates a servant and sets its state before calling `create_active_object_reference`. The TP Framework then enters the servant and string `ObjectId` in the Active Object Map. The result is exactly the same as if the TP Framework had previously invoked `Server::create_servant`, received back the servant pointer, and then had invoked `servant::activate_object`.

The object so activated must be for an interface that was declared with the process activation policy; otherwise, an exception is raised.

If the object is deactivated, an object reference held by a client might cause the object to be activated again in some other process. For a discussion about situations in which this might be a problem, refer to the section “Explicit Activation” on page 3-11.

Caution When you preactivate objects in an interface, you must specify an activation policy of `process` in the ICF file for that interface. However, when you specify the `process` activation policy for an interface in the ICF file, this can lead to the following problem.

PROBLEM STATEMENT

1. You write `SERVER1` such that all objects on interface A are preactivated. Therefore, interface A must be process bound. However, this implies that `SERVER1` can activate objects on demand.
2. `SERVER2` also implements objects of interface A. However, instead of preactivating the objects, `SERVER2` lets the TP Framework activate them on demand (that is, they are activated as normal objects). `SERVER2` also generates object references of interface A that can be handed to clients.
3. If the administrator configures `SERVER1` and `SERVER2` in the same group, then a client can get an interface A object reference from `SERVER2` and invoke on it. Then, due to load balancing, `SERVER1` could be asked to activate an object on interface A. However, `SERVER1` is not able to activate an object on interface A on demand because its interface A objects were preactivated.

WORKAROUND

You can avoid this problem by having the administrator configure `SERVER1` and `SERVER2` in different groups. The administrator uses the `SERVERS` section of the `UBBCONFIG` file to define groups.

Exceptions:

- `TobjS::InvalidInterface`
Indicates that the specified `interfaceName` is Null.
- `TobjS::InvalidObjectId`
Indicates the specified `stroid` is NULL.
- `TobjS::ServantAlreadyActive`
The object could not be activated explicitly because the servant is already being used with another `ObjectId`. A servant can be used only with a single

`ObjectId`. To preactivate objects containing different `ObjectIds`, the application must create multiple servants and preactivate them separately, one per `ObjectId`.

`TobjS::ObjectAlreadyActive`

The object could not be activated explicitly because the `ObjectId` is already being used in the Active Object Map. A given `ObjectId` can have only one servant associated with it. To change to a different servant, the application must first deactivate the object and activate it again.

`TobjS::IllegalOperation`

The object could not be activated explicitly because it does not have the process activation policy.

TP::create_object_reference()

Synopsis Creates an object reference. The resulting object reference may be passed to clients who use it to access the object.

C++ Binding `static CORBA::Object_ptr TP::create_object_reference (`
 `const char* interfaceName,`
 `const char* stroid,`
 `CORBA::NVList_ptr criteria);`

Arguments `interfaceName`

Specifies a character string that contains the fully qualified interface name for the object.

The interface name can be retrieved by making a call on the following interface typecode id function:

```
const char* _tc_<CORBA interface name>::id();
```

where `<CORBA interface name>` is any object class name. For example:

```
char* idlname = _tc_Simple->id();
```

`stroid`

Specifies the `ObjectId` in string format. The `ObjectId` *uniquely* identifies this instance of the class. It is up to the programmer to decide what information to place in the `ObjectId`. One possibility would be to use the `ObjectId` to hold a database key. Choosing the value of an object identifier, and the degree of uniqueness, is part of the application design. The WLE software cannot guarantee any uniqueness in object references, since object references may be legitimately copied and shared outside the WLE domain (for example, by passing the object reference as a string). It is strongly recommended the you choose a unique `ObjectId` in order to allow parallel execution of invokes on object references.

Note: The restriction on the length of the object ID has been removed in this release.

`criteria`

Specifies a list of named values that can be used to provide factory-based routing for the object reference. The list is optional and is of type `CORBA::NVList`. The use of factory-based routing is optional and is dependent on the use of this argument. If you do not want to use factory-based routing, you can pass a value of 0 (zero) for this argument.

The WLE system administrator configures factory-based routing by specifying routing rules in the `UBBCONFIG` file. See the *Administration Guide* online document for details on this facility.

Return Value	<p>Object</p> <p>The newly created object reference.</p>
Description	<p>The server application is responsible for invoking the <code>create_object_reference()</code> method. This method creates an object reference. The resulting object reference may be passed to clients who will use it to access the object.</p> <p>Ordinarily, the server application calls this method in two places:</p> <ul style="list-style-type: none"> ◆ In <code>Server::initialize()</code> to create factories for the server. ◆ In factory methods to create object references to be returned to clients. <p>For examples of how and when to call the <code>create_object_reference()</code> method, see <i>Creating C++ Server Applications</i>.</p>
Exceptions	<p>The following exceptions can be raised by the <code>create_object_reference()</code> method:</p> <p><code>InvalidInterface</code> Indicates that the specified <code>interfaceName</code> is Null.</p> <p><code>InvalidObjectId</code> Indicates that the specified <code>stroid</code> is Null.</p>
Example	<p>The following example shows how to use the criteria argument:</p> <pre> CORBA::NVList_ptr criteria; CORBA::Long branch_id = 7; CORBA::Long account_id = 10001; CORBA::Any any_val; // Create the list and assign to _var to cleanup on exit CORBA::ORB::create_list (2, criteria); CORBA::NVList_var criteria_var(criteria); // Add the BRANCH_ID any_val <=& branch_id; criteria->add_value("BRANCH_ID", any_val, 0); // Add the ACCOUNT_ID any_val <=& account_id; criteria->add_value("ACCOUNT_ID", any_val, 0); </pre>


```
// Create the object reference.  
TP::create_object_reference ("IDL:BankApp/Teller:1.0",  
"Teller_01", criteria);
```

TP::deactivateEnable

Synopsis Enables application-controlled deactivation of CORBA objects.

C++ Binding Current-object format:

```
static void TP::deactivateEnable();
```

Any-object format:

```
static void TP::deactivateEnable(
    const char* interfaceName,
    const char* stroid,
    Tobj_Servant servant);
```

Arguments `interfaceName`
Specifies a character string that contains the fully qualified interface name for the object.

`stroid`
Specifies the `ObjectId` in string format for the object to be deactivated.

`servant`
A pointer to the servant associated with the stroid.

Return Value None.

Description This method can be used to cause deactivation of an object, either the object currently executing (upon completion of the method in which it is called) or another object. It can only be used for objects with the process activation policy. It provides additional flexibility for objects with the process activation policy.

Depending on which of the overloaded functions are called, the actions are as follows.

Current-object format

When called from within a method of an object with process activation policy, the object currently executing will be deactivated after completing the method being executed.

When called from within a method of an object with method activation, the effect is the same as the normal behavior of such objects (effectively, a NOOP).

When the object is deactivated, the TP Framework first removes the object from the Active Object Map. It then calls the associated servant's `deactivate_object` method with a reason of `DR_METHOD_END`.

Any-object format

The application can request deactivation of an object by specifying its `ObjectId` and the associated servant.

If the object is currently executing, the TP Framework marks it for deactivation and waits until the object's method completes before deactivating the object (as with the "current-object format"). If the object is not currently executing, the TP Framework may deactivate it immediately.

No indication is given to the caller as to the status of the deactivation. When the object is deactivated, the TP Framework first removes the object from the Active Object Map. It then calls the associated servant's

`deactivate_object` method with a reason of `DR_EXPLICIT_DEACTIVATE`.

If the object for which the deactivate is requested has a `transaction` activation policy, an `IllegalOperation` exception is raised. This is because deactivation of such objects may interfere with their correct notification of transaction completion by the WLE transaction manager.

Exceptions The following exceptions can be raised by the `deactivateEnable()` method:

`IllegalOperation`

Indicates that the `TP::deactivateEnable` method was invoked by an object with the activation policy set to `transaction`.

`TobjS::ObjectNotActive`

In the Any-object format, the object specified could not be deactivated because it was not active (the `stroid` and `servant` parameters did not identify an object that was in the Active Object Map).

TP::get_object_id ()

Synopsis	This method allows a server to retrieve the string <code>ObjectId</code> contained in an object reference that was created in the TP Framework.
C++ Binding	<code>char* TP::get_object_id(Corba::Object_ptr obj);</code>
Arguments	<code>obj</code> The object reference from which to get the <code>ObjectId</code> .
Return Value	The string <code>ObjectId</code> passed to <code>TP::create_object_reference</code> or <code>TP::create_active_object_reference</code> when the object reference was created.
Description	<p>This method allows a server to retrieve the string <code>ObjectId</code> contained in an object reference that was created in the TP Framework. If the object reference was not created in the TP Framework (for example, it was created by a client ORB), an exception is raised.</p> <p>The caller must call <code>CORBA::string_free</code> on the returned value when the object reference is no longer needed.</p>
Exception	<code>TobjS::InvalidObject</code> The object is nil or was not created by the TP Framework

TP::get_object_reference()

Synopsis	Returns a pointer to the current object.
C++ Binding	<pre>static CORBA::Object_ptr TP::get_object_reference ();</pre>
Arguments	None.
Return Value	<p>The <code>get_object_reference()</code> method returns a <code>CORBA::Object_ptr</code> for the current object when invoked within the scope of a CORBA object execution. Otherwise, the <code>TobjS::NilObject</code> exception is raised.</p> <p>Note that if <code>get_object_reference()</code> is invoked from within either <code>Server::initialize()</code> or <code>Server::release()</code>, it is considered to be invoked outside the scope of an application's TP object execution; therefore, the <code>TobjS::NilObject</code> exception is raised.</p>
Description	This method returns a pointer to the current object. The <code>CORBA::Object_ptr</code> pointer that is returned can be passed to a client.
Exceptions	<p>The following exception can be raised by the <code>get_object_reference()</code> method:</p> <p><code>NilObject</code></p> <p>Indicates that the method was invoked outside the scope of an application's CORBA object execution. The <code>reason</code> string contains <code>OutOfScope</code>.</p>

TP::open_xa_rm()

Synopsis	Opens the XA resource manager to which the invoking process is linked.
C++ Binding	<code>static void TP::open_xa_rm();</code>
Arguments	None.
Return Values	None.
Description	The <code>open_xa_rm()</code> method opens the XA resource manager to which the invoking process is linked. XA resource managers are provided by database vendors, such as Oracle and Informix.
Note:	The functionality of this method is also provided by <code>Tobj::TransactionCurrent::close_xa_rm()</code> . However, <code>TP::open_xa_rm()</code> provides a more convenient way for a server application to close a resource manager because there is no need to obtain an object reference to the <code>TransactionCurrent</code> object. A reference to the <code>TransactionCurrent</code> object can be obtained from the <code>Bootstrap</code> object. See “ <code>TP::bootstrap()</code> ” on page 3-47 for an explanation of how to obtain a reference to the <code>Bootstrap</code> object. For more information about the <code>TransactionCurrent</code> object, see Chapter 4, “Bootstrap Object,” and Chapter 7, “Transaction Service.”

This method should be invoked once from the `Server::initialize()` method for each server that participates in a global transaction. This includes servers that are linked with an XA resource manager, as well as servers that participate in a global transaction, but are not actually linked with an XA-compliant resource manager.

The `open_xa_rm()` method should be invoked in place of an open invocation that is specific to a resource manager. Because resource managers differ in their initialization semantics, the specific information needed to open a particular resource manager is placed in the `OPENINFO` parameter in the `GROUPS` section of the `UBBCONFIG` file.

The format of the `OPENINFO` string is dependent on the requirements of the database vendor providing the underlying resource manager. For more information about the `CLOSEINFO` parameter, see the *Administration Guide* and the `ubbconfig(5)` reference page in the *BEA TUXEDO Reference*. Also, refer to database vendor documentation for information about how to develop and install applications that use the XA libraries.

Note: Only one resource manager can be linked to the invoking process.

Exceptions `Tobj::RMFailed`

The `tx_open()` call returned an error return code.

Note: Unlike other exceptions returned by the TP Framework, this exception is defined in `tobj_c.h` (which is derived from `tobj.idl`), not in `TobjS_c.h` (which is derived from `TobjS.idl`). This is because native clients can also open XA resource managers. Therefore, the exception returned is consistent with the exception expected by native client code and by `Server::release()` if it uses the alternate mechanism, `TransactionCurrent::close_xa_rm`, which is shared with native clients.

TP::orb()

Synopsis	Returns a pointer to an ORB object.
C++ Binding	<pre>static CORBA::ORB_ptr TP::orb();</pre>
Arguments	None.
Return Value	Upon successful completion, <code>orb()</code> returns a pointer to the ORB object that is created by the TP Framework when the server program is started.
Description	<p>Access to the ORB object allows the application to invoke ORB operations, such as <code>string_to_object()</code> and <code>object_to_string()</code>.</p> <p>Note: Because the TP Framework owns the ORB object, the application must not delete it.</p>
Exceptions	None.

TP::register_factory()

Synopsis	Locates the WLE FactoryFinder object and registers a WLE factory.
C++ Binding	<pre>static void TP::register_factory(CORBA::Object_ptr factory_or, const char* factory_id);</pre>
Arguments	<p><code>factory_or</code> Specifies the object reference that was created for an application factory using the <code>TP::create_object_reference()</code> method.</p> <p><code>factory_id</code> Specifies a string identifier that is used to identify the application factory. For some suggestions as to the composition of this string, see <i>Creating C++ Server Applications</i>.</p>
Return Value	None.
Description	<p>This method locates the WLE FactoryFinder object and registers a WLE factory. Typically, <code>TP::register_factory()</code> is invoked from <code>Server::initialize()</code> when the server creates its factories. The <code>register_factory()</code> method locates the WLE FactoryFinder object and registers the WLE factory.</p> <p>Caution: Callback objects (that is, those created by a joint client/server directly through the POA) should not be registered with a FactoryFinder.</p>
Exceptions	<p>The following exceptions can be raised by the <code>register_factory()</code> method:</p> <p><code>TobjS::CannotProceed</code> Indicates that the FactoryFinder encountered an internal error during the search, with the error being written to the user log (ULOG). Notify the operations staff immediately if this exception is raised. Depending on the severity of the internal error, the server running the FactoryFinder or the NameManager may have terminated. If a FactoryFinder service has terminated, start a new FactoryFinder service. If the NameManager has terminated, and there is another NameManager running, start a new one. If no NameManagers are running, restart the application.</p> <p><code>TobjS::InvalidName</code> Indicates that the <code>id</code> string is empty. It is also raised if the field contains blank spaces or control characters.</p>

`TobjS::InvalidObject`

Indicates that the `factory` value is `nil`.

`TobjS::RegistrarNotAvailable`

Indicates that the `FactoryFinder` object cannot locate the `NameManager`.
Notify the operations staff immediately if this exception is raised. If no
naming services servers are running, restart the application.

`TobjS::Overflow`

Indicates that the `id` string is longer than 128 bytes (currently the maximum
allowable length).

TP::unregister_factory()

Synopsis	Locates the WLE FactoryFinder object and removes a factory.
C++ Binding	<pre>static void TP::unregister_factory (CORBA::Object_ptr factory_or, const char* factory_id);</pre>
Arguments	<p><code>factory_or</code> Specifies the object reference that was created for an application factory using the <code>TP::create_object_reference()</code> method.</p> <p><code>factory_id</code> Specifies a string identifier that is used to identify the application factory. For some suggestions as to the composition of this string, see <i>Creating C++ Server Applications</i>.</p>
Return Value	None.
Description	This method locates the WLE FactoryFinder object and removes a factory. Typically <code>TP::unregister_factory()</code> is invoked from <code>Server::release()</code> to unregister server factories.
Exceptions	<p>The following exceptions can be raised by the <code>unregister_factory()</code> method:</p> <p><code>CannotProceed</code> Indicates that the FactoryFinder encountered an internal error during the search, with the error being written to the user log (ULOG). Notify the operations staff immediately if this exception is raised. Depending on the severity of the internal error, the server running the FactoryFinder or the NameManager may have terminated. If a FactoryFinder service has terminated, start a new FactoryFinder service. If the NameManager has terminated, and there is another NameManager running, start a new one. If no NameManagers are running, restart the application.</p> <p><code>InvalidName</code> Indicates that the <code>id</code> string is empty. It is also raised if the field contains blank spaces or control characters.</p> <p><code>RegistrarNotAvailable</code> Indicates that the FactoryFinder object cannot locate the NameManager. Notify the operations staff immediately if this exception is raised. If no naming services servers are running, restart the application.</p>

`TobjS::OverFlow`

Indicates that the `id` string is longer than 128 bytes (currently the maximum allowable length).

TP::userlog()

Synopsis	Writes a message to the user log (ULOG) file.
C++ Binding	<pre>static int TP::userlog(char*, ...);</pre>
Arguments	The first argument is a <code>printf(3S)</code> style format specification. The <code>printf(3S)</code> argument is described in a C or C++ reference manual.
Return Value	The <code>userlog()</code> method returns the number of characters that were output, or a negative value if an output error was encountered. Output errors include the inability to open or write to the current log file.
Description	<p>The <code>userlog()</code> method writes a message to the user log (ULOG) file. Messages are appended to the ULOG file with a tag made up of the time (hhmmss), system name, process name, and process-id of the invoking process. The tag is terminated with a colon.</p> <p>We recommend that server applications limit their use of <code>userlog()</code> messages to messages that can be used to help debug application errors; flooding the ULOG file with incidental information can make it difficult to spot actual errors.</p>
Exceptions	None.
Example	<p>The following example shows how to use the <code>TP::userlog()</code> method:</p> <pre>userlog ("System exception caught: %s", e.get_id());</pre>

CosTransactions::TransactionalObject Interface Not Enforced

Use of this interface is now deprecated. Therefore, the use of this interface is now optional and no enforcement of descent from this interface is done for objects infected with transactions. The programmer can specify that an object is not to be infected by transactions by specifying the `never` or `ignore` transaction policies. There is no interface enforcement for eligibility for transactions. The only indicator is the transaction policy.

Note: The CORBAServices Object Transaction Service does not require that all requests be performed within the scope of a transaction. It is up to each object to determine its behavior when invoked outside the scope of a transaction; an object that requires a transaction context can raise a standard exception.

Error Conditions, Exceptions, and Error Messages

Exceptions Raised by the TP Framework

The following exceptions are raised by the TP Framework and are returned to clients when error conditions occur in, or are detected by, the TP Framework:

```
CORBA::INTERNAL  
CORBA::OBJECT_NOT_EXIST  
CORBA::OBJ_ADAPTER  
CORBA::INVALID_TRANSACTION  
CORBA::TRANSACTION_ROLLEDBACK
```

Since the reason for these exceptions may be ambiguous, each time one of these exceptions is raised, the TP Framework also writes a descriptive error message that explains the reason to the user log file.

Exceptions in the Server Application Code

Exceptions raised within a method invoked by a client are always raised back to the client exactly as they were raised in the method invoked by the client.

The following TP Framework callback methods are initiated by events other than client requests on the object:

```
Tobj_ServantBase::activate_object()  
Tobj_ServantBase::deactivate_object()  
Server::create_servant()
```

If exception conditions are raised in these methods, those exact exceptions are not reported back to the client. However, each of these methods is defined to raise an exception that includes a reason string. The TP Framework will catch the exception raised by the callback and log the reason string to the user log file. The TP Framework may raise an exception back to the client. Refer to the descriptions of the individual TP Framework callback methods for more information about these exceptions.

Example

For `Tobj_ServantBase::deactivate_object()` the following line of code throws a `DeactivateObjectFailed` exception:

```
throw TobjS::DeactivateObjectFailed( "deactivate failed to save  
state!");
```

This message is appended to the user log file with a tag made up of the time (hhmmss), system name, process name, and process-id of the calling process. The tag is terminated with a colon. The preceding throw statement causes the following line to appear in the user log file:

```
151104.T1!simpapps.247: APPEXC: deactivate failed to save state!
```

Where 151104 is the time (3:11:04pm), T1 is the system name, simpapps is the process name, 247 is the process-id, and APPEXC identifies the message as an application exception message.

Exceptions and Transactions

Exceptions that are raised in either CORBA object methods or in TP Framework callback methods will not automatically cause a transaction to be rolled back unless the TP Framework started the transaction. It is up to the application code to call `Current.rollback_only()` if the condition that caused the exception to be raised should also cause the transaction to be rolled back.

Restriction of Nested Calls on Corba Objects

The TP Framework restricts nested calls on CORBA objects. The restriction is as follows:

- ◆ During a client invocation of a method of CORBA object A, CORBA object A cannot be invoked by another CORBA object B that is acting as a client of CORBA object A.

The TP Framework will detect the fact that a second CORBA object is acting as a client to an object that is already processing a method invocation, and will return a `CORBA::OBJ_ADAPTER` exception to the caller.

Note: Application code should not depend on this behavior; that is, users should not make any processing dependent on this behavior. This restriction may be lifted in a future release.

4 Bootstrap Object

This chapter describes the Bootstrap object.

Why Bootstrap Objects Are Needed

The Problem: To communicate with WLE objects, a client application must obtain object references. The client application uses the Bootstrap object to obtain initial object references to four key objects in a WLE domain: the `FactoryFinder` (which is used to locate factory objects), `SecurityCurrent` (which is used to log on to the system), `TransactionCurrent` (which is used to manage transactions), and the `Interface Repository` (which is used to obtain information about available interfaces). However, this poses a problem: *How does the client application access the Bootstrap object?*

The solution: Bootstrap objects are local programming objects, not remote CORBA objects, in both the client and the server. When Bootstrap objects are created, their constructor requires the network address of a WLE IIOP Server Listener/Handler. Given this information, the Bootstrap object can generate object references for the above-mentioned remote objects in the WLE domain. These object references can then be used to access services available in the WLE domain.

How Bootstrap Objects Work

Bootstrap objects are created by a client or a server application that must access object references for the `Interface Repository`, or for the `FactoryFinder`, `TransactionCurrent`, or `SecurityCurrent` objects. Bootstrap objects represent the first connection to a

specific WLE domain. For a WLE remote client, the Bootstrap object is created with the host and the port for the WLE IIOP Server Listener/Handler. However, for WLE native client and server applications, there is no need to specify a host and port because they execute in a specific WLE domain. The IIOP Server Listener/Handler host and the port ID are included in the WLE domain configuration information.

Once created, Bootstrap objects satisfy requests for object references for objects in a particular WLE domain. Different Bootstrap objects allow the application to use multiple domains.

Using the Bootstrap object, you can obtain four different references, as follows:

◆ **SecurityCurrent**

The SecurityCurrent object is used to establish a security context within a WLE domain. The client can then obtain the PrincipalAuthenticator from the `principal_authenticator` attribute of the SecurityCurrent object.

◆ **TransactionCurrent**

The TransactionCurrent object is used to participate in a WLE transaction. The basic operations are as follows:

◆ **Begin**

Begin a transaction. Future operations take place within the scope of this transaction.

◆ **Commit**

End the transaction. All operations on this client have completed successfully.

◆ **Roll back**

Abort the transaction. Tell all other participants to roll back.

◆ **Suspend**

Suspend participation in the current transaction. This operation returns an object that identifies the transaction and allows the client to resume the transaction later.

◆ Resume

Resume participation in the specified transaction.

◆ FactoryFinder

The FactoryFinder object is used to obtain a factory. In the WLE system, factories are used to create application objects. The FactoryFinder provides the following different methods to find factories:

- ◆ Get a list of all available factories that match a factory object reference (`find_factories`).
- ◆ Get the factory that matches a name component consisting of id and kind (`find_one_factory`).
- ◆ Get the first available factory of a specific kind (`find_one_factory_by_id`).
- ◆ Get a list of all available factories of a specific kind (`find_factories_by_id`).
- ◆ Get a list of all registered factories (`list_factories`).

◆ InterfaceRepository

The Interface Repository contains the interface descriptions of the CORBA objects that are implemented within the WLE domain. Clients using the Dynamic Invocation Interface (DII) need a reference to the Interface Repository to be able to build CORBA request structures. The ActiveX Client is a special case of this. Internally, the implementation of the COM/IOP Bridge uses DII, so it must get the reference to the Interface Repository, although this is transparent to the desktop client.

The FactoryFinder and Interface Repository objects are not actually implemented in the environmental objects library. However, they are specific to a WLE domain and are thus conceptually similar to the SecurityCurrent and TransactionCurrent objects in usage.

The Bootstrap object implies an association or "session" between the client and the WLE domain. Within the context of this association, the Bootstrap object imposes a containment relationship with the other Current objects (or contained objects), that is, the SecurityCurrent and TransactionCurrent. Current objects are valid only for this domain and only while the Bootstrap object exists.

In addition, a client can have only one instance of each of the Current objects at any time. If a Current object already exists, an attempt to create another Current object does not fail. Instead, another reference to the already existing object is handed out; that is, a client may have more than one reference to the single instance of the Current object.

To create a new instance of a Current object, the application must first invoke the `destroy_current()` method on the Bootstrap object. This invalidates all of the Current objects, but will not destroy the session with the WLE domain. After invoking `destroy_current()`, new instances of the Current objects can be created within the WLE domain using the existing Bootstrap object.

To obtain Current objects for another domain, a different Bootstrap object must be constructed. Although it is possible to have multiple Bootstrap objects at one time, only one Bootstrap object may be "active," that is, have Current objects associated with it. Thus, an application must first invoke `destroy_current()` on the "active" Bootstrap object before obtaining new Current objects on another Bootstrap object, which then becomes the active Bootstrap object.

Servers and native clients are inside of the WLE domain; therefore, no "session" is established. However, the same containment relationships are enforced. Servers and native clients access the domain they are currently in by specifying an empty string, rather than `//host:port`. Client and server applications must use the `Tobj_Bootstrap::resolve_initial_references()` method, not the `ORB::resolve_initial_references()` method.

Types of Remote Clients Supported

Table 4-1 shows the types of remote clients that can use the Bootstrap object to access the other environmental objects, such as `FactoryFinder`, `SecurityCurrent`, `TransactionCurrent`, and `InterfaceRepository`.

Table 4-1 Remote Clients Supported

Remote Client	Description
CORBA C++	CORBA C++ client applications use the WLE C++ environmental objects to access the CORBA objects in a WLE domain, and the WLE Object Request Broker (ORB) to process from CORBA objects. Use the WLE system development commands to build these client applications (see Chapter 10, “Development Commands”).
CORBA Java	CORBA Java client applications use the Java environmental objects to access CORBA objects in a WLE domain. However, these client applications use an ORB product other than the WLE ORB to process requests from CORBA objects. These client applications are built using the ORB product’s Java development tools. Version 4.2 of the WLE software supports interoperability with Netscape Communicator Version 4.07 and 4.5 depending on the platform.
ActiveX	Use the WLE Automation environmental objects to access CORBA objects in a WLE domain, and the ActiveX Client to process requests from CORBA objects. Use the Application Builder to create bindings for CORBA objects so that they can be accessed from ActiveX client applications. ActiveX client applications are built using a development tool such as Visual Basic, Delphi, or PowerBuilder.

Capabilities and Limitations

Bootstrap objects have the following capabilities and limitations:

- ◆ Multiple Bootstrap objects can coexist in a client, although only one Bootstrap object can own the Current objects (Transaction and Security) at one time. Client applications must invoke `destroy_current()` on the Bootstrap object associated with one domain before obtaining the Current objects on another domain. Although it is possible to have multiple Bootstrap objects that establish connections to different WLE domains, only one set of Current objects is valid. Attempts to obtain other Current objects without destroying the existing Current objects fail.

- ◆ Method invocations to any WLE domain other than the domain that provides the valid `SecurityCurrent` object fail and return a `CORBA::NO_PERMISSION` exception.
- ◆ Method invocations to any WLE domain other than the domain that provides the valid `TransactionCurrent` object do not execute within the scope of a transaction.
- ◆ The transaction and security objects returned by the Bootstrap objects are BEA implementations of the Current objects. If other ("native") Current objects are present in the environment, they are ignored.

Bootstrap Object API

The Bootstrap object application programming interface (API) is described first in terms of the OMG Interface Definition Language (IDL) (for portability), and then in C++, Java, and ActiveX. C++ and Java add the necessary constructor to build a Bootstrap object for a given WLE domain.

Tobj Module

Table 4-2 shows the object reference that is returned for each type ID.

Table 4-2 Returned Object References

ID	Returned Object Reference
FactoryFinder	FactoryFinder object (Tobj::FactoryFinder)
InterfaceRepository	InterfaceRepository object (CORBA::Repository)
SecurityCurrent	SecurityCurrent object (SecurityLevel2::Current)
TransactionCurrent	OTS Current object (Tobj::TransactionCurrent)

Table 4-3 describes the Tobj module exceptions.

Table 4-3 Tobj Module Exceptions

Exception	Description
Tobj::InvalidName	Raised if id is not one of the names specified in Table 4-2. On the server, resolve_initial_references also raises Tobj::InvalidName when SecurityCurrent is passed.
Tobj::InvalidDomain	On the server application, raised if the WLE server environment is not booted.
CORBA::NO_PERMISSION	Raised if id is TransactionCurrent or SecurityCurrent and another Bootstrap object in the client owns the Current objects.
BAD_PARAM	Raised for the register_callback_port method if the object is nil or if the host contained in the object does not match the connection.
IMP_LIMIT	Raised if the register_callback_port method is called more than once.

C++ Mapping

Listing 4-1 shows the C++ declarations in the `Tobj_bootstrap.h` file.

Listing 4-1 Tobj_bootstrap.h Declarations

```
#include <CORBA.h>

class Tobj_Bootstrap {
public:
    Tobj_Bootstrap(CORBA::ORB_ptr orb, const char* address);
    CORBA::Object_ptr resolve_initial_references(
        const char* id);
    void register_callback_port(CORBA::Object_ptr objref);
    void destroy_current( );
};
```

Java Mapping

Listing 4-2 shows the `Tobj_Bootstrap.java` mapping.

Listing 4-2 Tobj_Bootstrap.java Mapping

```
package com.beasys;

public class Tobj_Bootstrap {
    public Tobj_Bootstrap(org.omg.CORBA.ORB orb,
        String address)
        throws org.omg.CORBA.SystemException;
    public class Tobj_Bootstrap {
        public Tobj_Bootstrap(org.omg.CORBA.ORB orb, String address,
            java.applet.Applet applet)
            throws org.omg.CORBA.SystemException;

        public void register_callback_port(org.omg.CORBA.Object objref)
            throws org.omg.CORBA.SystemException;
    }
}
```

```
public org.omg.CORBA.Object
    resolve_initial_references(String id)
        throws Tobj.InvalidName,
            org.omg.CORBA.SystemException;
public void destroy_current()
    throws org.omg.CORBA.SystemException;
}
```

Mappings for Microsoft Desktop Clients

The Bootstrap object is provided in the BEA ActiveX Client software for use by clients that are implemented on Microsoft desktops. There are two possible interfaces that desktop clients may use:

- ◆ The automation interface for Visual Basic (VB), Delphi, or PowerBuilder clients.
- ◆ The Dual interface that provides both the automation interfaces required by dynamic clients (Visual Basic) and the Vtable interfaces required by statically linked clients (C++). The Bootstrap object in the ActiveX Client provides the hybrid DUAL interface.

Automation Mapping

Listing 4-3 shows Automation Bootstrap interface mapping.

Listing 4-3 Automation (Dual) Bootstrap Interface Mapping

```
interface DITobj_Bootstrap : IDispatch
{
    HRESULT Initialize(
        [in] BSTR address);

    HRESULT CreateObject(
        [in] BSTR progid,
        [out, retval] IDispatch** rtn);

    HRESULT destroy_current();
};
```

C++ Member Functions and Java Methods

This section describes the C++ member functions and Java methods for Bootstrap objects.

Tobj_Bootstrap

Synopsis The Bootstrap object constructor.

C++ Mapping `Tobj_Bootstrap(CORBA::ORB_ptr orb, const char* address);`
 throws `Tobj::InvalidDomain`
 `org.omg.CORBA.SystemException;`

Java Mapping `public Tobj_Bootstrap(org.omg.CORBA.ORB orb, String address,`
 `java.applet.Applet applet)`
 throws `com.beasys.Tobj.InvalidDomain,`
 throws `org.omg.CORBA.SystemException;`

Parameters `orb`

This is a pointer to the ORB object in the client. The Bootstrap object uses the `string_to_object` method of `orb` internally.

`address`

The address of the WLE domain IIOP Server Listener/Handler. The address is specified differently depending on the type of client. There can be three types of clients, as follows:

◆ **Remote Client**

For a description of the remote clients supported by WLE systems, see the section “Types of Remote Clients Supported” on page 4-4.

For remote clients, `address` specifies the network address of an IIOP Server Listener/Handler through which client applications gain access to a WLE domain.

The address may be specified in either of the following formats:

```
//hostname:port_number"
"//#. #. #. #:port_number"
```

In the first format, the domain finds an address for *hostname* using the local name resolution facilities (usually DNS). The *hostname* must be the local machine, and the local name resolution facilities must unambiguously resolve *hostname* to the address of the local machine.

Note: The hostname must begin with a letter character.

In the second example, the `#.#.#.#` is in dotted decimal format. In dotted decimal format, each `#` should be a number from 0 to 255. This dotted decimal number represents the IP address of the local machine.

In both of the above formats, *port_number* is the TCP port number at which the domain process listens for incoming requests. The *port_number* should be a number between 0 and 65535.

Note: The network address that is specified by programmers in the Bootstrap constructor or in `TOBJADDR` must exactly match the network address in the server application's `UBBCONFIG` file. The format of the address as well as the capitalization must match. If the addresses do not match, the invocation to the Bootstrap constructor will fail with the following seemingly unrelated error message:

```
ERROR: Unofficial connection from client at
<tcp/ip address>/<port-number>
```

For example, if the network address is specified as `//TRIXIE:3500` in the ISL command line option string in the server application's `UBBCONFIG` file, specifying either `//192.12.4.6:3500` or `//trixie:3500` in the Bootstrap constructor or in `TOBJADDR` will cause the connection attempt to fail. On UNIX systems, use the `uname -n` command on the host system to determine the capitalization used. On Windows NT systems, see the host system's Network control panel to determine the capitalization used.

One or more TCP/IP addresses can be specified. Multiple addresses are specified using a comma-separated list. For example:

```
//m1.acme.com:3050
```

```
//m1.acme.com:3050, //m2.acme.com:3050, //m3.acme.com:3051
```

If multiple addresses are specified, the addresses are tried in order until a connection is established. Any member of an address list can be specified as a parenthesized grouping of pipe-separated network addresses. For example:

```
(//m1.acme.com:3050|//m2.acme.com:3050)
```

The WLE system randomly selects one of the parenthesized addresses. This strategy distributes the load randomly across a set of listener processes.

The address string can be specified either in the `TOBJADDR` environment variable or in the address parameter of the `Tobj_Bootstrap` constructor. (For information about the `TOBJADDR` environment variable, see the chapter that describes how to manage remote client applications in *Administration Guide*.) However, the address specified in `Tobj_Bootstrap` always take precedence over the `TOBJADDR` environment variable. To use the `TOBJADDR` environment variable to specify an address string, you must specify an empty string in the `Tobj_Bootstrap` address parameter.

Note: For C++ applications, `TOBJADDR` is an environment variable; for Java applications, it is a property; for Java applets, it is an HTML parameter.

◆ **Native Client**

For a native client, the address parameter in the `Tobj_Bootstrap` constructor must always be an empty string (not a null pointer). The native client connects to the application that is specified in the `TUXCONFIG` environment variable. The constructor raises `CORBA::BAD_PARAM` if the address is not empty.

◆ **Server Acting as a Client**

For a server, the address parameter in the `Tobj_Bootstrap` constructor must always be an empty string (not a null pointer). The server always connects to the application in which it is booted. The constructor raises `CORBA::BAD_PARAM` if the address is not empty.

applet (Applies to Java method only)

This is a pointer to the client applet. If the client applet does not explicitly pass the ISH host and port to the Bootstrap constructor, you can pass this argument, which causes the Bootstrap object to search for the `TOBJADDR` definition in the HTML file for the applet.

Exception `InvalidDomain`

For a remote client, raised if the Bootstrap object cannot connect to the WLE domain. The address of the WLE domain IIOP Server Listener/Handler is specified in the constructor that is specific to the programming language. For a native client or server, raised if the domain is not booted.

The constructor throws `CORBA::BAD_PARAM` if `orb` is null or if address is not in a valid format.

Description This C++ member function (or Java method) creates Bootstrap objects.

Return Values A pointer to a newly created Bootstrap object.

Tobj_Bootstrap::register_callback_port

Synopsis	Registers joint client/server's listening port in ISH.
C++ Mapping	<code>void register_callback_port(CORBA::Object_ptr objref);</code>
Java Mapping	<code>public void register_callback_port(orb.omg.CORBA.Object objref) throws org.omg.CORBA.SystemException;</code>
Parameter	<code>objref</code> The object reference created by the client.
Exceptions	<code>BAD_PARAM</code> Raised if the object is nil or if the host contained in the object does not match the connection. <code>IMP_LIMIT</code> Raised if the <code>register_callback_port</code> method is called more than once.
Description	This C++ member function (or Java method) is called to notify the ISH of a listening port in the joint client/server. This method should only be used for joint client/server ORBs that do not support GIOP 1.2 bidirectional capabilities (that is GIOP 1.0 and 1.1 client ORBs). For GIOP 1.0 and 1.1, the ISH supports only one listening port per joint client/server; therefore, the <code>register_callback_port</code> method should only be called once per connected joint client/server.
Usage Notes	If the <code>register_callback_port</code> method is <i>not</i> invoked by the joint client/server, the callback port is not registered with the ISH and the server defaults to Asymmetric Outbound IIOP. Also, if the <code>register_callback_port</code> method is not invoked by the joint client/server, you <i>must</i> start the server's IIOP Server Listener (ISL) with the <code>-o</code> option. The <code>-o</code> option enables Asymmetric outbound IIOP; otherwise, server-to-client invocations will not be allowed by the ISL/ISH.
Return Values	None.

Tobj_Bootstrap::resolve_initial_references

Synopsis Acquires CORBA object references.

C++ Mapping `CORBA::Object_ptr resolve_initial_references(
 const char* id);
 throws Tobj::InvalidName,
 org.omg.CORBA.SystemException;`

Java Mapping `public org.omg.CORBA.Object
 resolve_initial_references(String id)
 throws Tobj.InvalidName,
 org.omg.CORBA.SystemException;`

Parameter `id`
 This parameter must be one of the following:

`"FactoryFinder"
"SecurityCurrent"
"TransactionCurrent"
"InterfaceRepository"`

Exceptions `InvalidName`
 Raised if `id` is not one of the names specified above. On the server, `resolve_initial_references` also raises `Tobj::InvalidName` when `SecurityCurrent` is passed.

`CORBA::NO_PERMISSION`
 Raised if `id` is `TransactionCurrent` or `SecurityCurrent` and another Bootstrap object in the client owns the Current objects.

Description This C++ member function (or Java method) acquires CORBA object references for the `FactoryFinder`, `SecurityCurrent`, `TransactionCurrent`, and `InterfaceRepository` objects. For the specific object reference, invoke the `_narrow` function. For example, for `FactoryFinder`, invoke `Tobj::FactoryFinder::_narrow`.

Return Values Table 4-2, "Returned Object References," on page 4-7 shows the object reference that is returned for each type `id`.

Tobj_Bootstrap::destroy_current()

Synopsis Destroys the Current objects for the domain represented by the Bootstrap object.

C++ Mapping `void destroy_current();`

Java Mapping `public void destroy_current()
throws org.omg.CORBA.SystemException;`

Exception Raises `CORBA::NO_PERMISSION` if the Bootstrap object is not the owner of the Current objects.

Description This C++ member function invalidates the Current objects for the domain represented by the Bootstrap object. After invoking the `destroy_current()` method, the Current objects are marked as invalid. Any attempt to use the old Current objects from then on throws the exception `CORBA::BAD_INV_ORDER`. Good programming practice is to release all Current objects before invoking `destroy_current()`.

Note: The `destroy_current()` method must be invoked on the Bootstrap object for the domain that currently owns the two Current objects (Transaction and Security). This also results in an implicit invocation to `logoff` for security and implicitly rolls back any transaction that was begun by the client.

The application must invoke `destroy_current()` before invoking `resolve_initial_references` for `TransactionCurrent` or `SecurityCurrent` on another domain; otherwise, `resolve_initial_references` raises `CORBA::NO_PERMISSION`.

Return Values None.

Automation Methods

This section describes the Automation methods for Bootstrap objects.

Initialize

Synopsis Initializes the Bootstrap object into a WLE domain.

MIDL Mapping `HRESULT Initialize(
 [in] BSTR host);`

Automation Mapping `Sub Initialize(address As String)`

Parameter `address`
The host name and port of the WLE domain IIOP Server Listener/Handler. One or more TCP/IP addresses can be specified. Multiple addresses are specified using a comma-separated list, just as in the C++ mappings. If no address is specified, the value of the `TOBJADDR` environmental variable is used.

Note: The network address that is specified by programmers in the Bootstrap constructor or in `TOBJADDR` must exactly match the network address in the application's `UBBCONFIG` file. The format of the address as well as the capitalization must match. If the addresses do not match, the invocation to the Bootstrap constructor will fail with the following seemingly unrelated error message:

```
ERROR: Unofficial connection from client at  
<tcp/ip address>/<port-number>
```

For example, if the network address is specified as `//TRIXIE:3500` in the ISL command line option string, specifying either `//192.12.4.6:3500` or `//trixie:3500` in the Bootstrap constructor or in `TOBJADDR` will cause the connection attempt to fail. On UNIX systems, use the `uname -n` command on the host system to determine the capitalization used. On Windows NT systems, see the host system's Network control panel to determine the capitalization used.

Return Values None.

Exceptions Table 4-4 describes the exceptions.

Table 4-4 Initialize Exceptions

HRESULT	Description	Meaning
ITF_E_NO_PERMISSION_YES	Bootstrap already initialized	The Bootstrap object has already been initialized. To connect to a new WLE domain, you must create a new Bootstrap object.
E_INVALIDARG	Invalid 'address' parameter	The address supplied is not valid.
E_OUTOFMEMORY	Memory allocation failed	The required memory could not be allocated.
E_FAIL	Invalid domain	Unable to communicate with the WLE domain at the address specified or TOBJADDR is not defined.
<SYSTEM_ERROR>	Unable to obtain initial object	Unable to initialize the Bootstrap object. The system error causing the failure is returned in the "Number" member of the error object

CreateObject

Synopsis	Creates an instance of a Current environmental object.
MIDL Mapping	<pre>HRESULT CreateObject([in] BSTR progid, [out, retval] IDispatch** rtrn);</pre>
Automation Mapping	Function CreateObject(progid As String) As Object
Parameter	<p>progid</p> <p>The progid of the environmental object to create. Valid progids are:</p> <p>Tobj.FactoryFinder Tobj.SecurityCurrent Tobj.TransactionCurrent</p>
Return Value	A reference to the interface pointer of the created environmental object.
Exceptions	Table 4-5 describes the exceptions.

Table 4-5 CreateObject Exceptions

Exception	Description	Meaning
ITF_E_NO_PERMISSION_YES	Bootstrap must initialized	The Bootstrap object has not been initialized.
ITF_E_NO_PERMISSION_NO	No permission.	If the progid specifies a transaction or security current and another Bootstrap object in the client owns the current objects.
E_INVALIDARG	Invalid 'progid' parameter	The progid specified is not valid.
E_INVALIDARG	Invalid name	The requested progid is not one of the valid parameter values specified above.

Table 4-5 CreateObject Exceptions (Continued)

Exception	Description	Meaning
E_INVALIDARG	Unknown object	The requested <code>progid</code> is not registered on your system.
<SYSTEM ERROR>	CoCreate Instance() failed	The Bootstrap object could not create an instance of the requested object. The system error is returned in the "Number" member of the error object.

DestroyCurrent

Synopsis Logs out of the WLE domain and invalidates the TransactionCurrent and SecurityCurrent objects.

MIDL Mapping HRESULT destroy_current();

Automation Mapping Sub destroy_current()

Parameters None.

Return Value None.

Exceptions None.

Programming Examples

This section provides programming examples that use Bootstrap objects.

Java Client Example: Getting a SecurityCurrent Object

Listing 4-4 shows how to program a Java client to get a SecurityCurrent object.

Listing 4-4 Programming a Java Client to Get a SecurityCurrent Object

```
import org.omg.CORBA.*;
import com.beasys.*;

class client {
    public static void main(String[] args)
    {
        Tobj.PrincipalAuthenticator auth = null;

        try {
            // Initialize ORB
            ORB orb = ORB.init();

            // Create Bootstrap object
            Tobj_Bootstrap bs = new Tobj_Bootstrap(orb,
                                                    "://host:1234");

            // Get security current
            org.omg.CORBA.Object ocur =

            bs.resolve_initial_references("SecurityCurrent");
            SecurityLevel2.Current cur =
                SecurityLevel2.CurrentHelper.narrow(ocur);
        }
        catch (Tobj.InvalidName e){
            System.out.println("Invalid name: " + e);
            System.exit(1);
        }
        catch (Tobj.InvalidDomain e){
            System.out.println("Invalid domain address:
                               //host:port: " + e);

            System.exit(1);
        }
        catch (SystemException e){
            System.out.println("Exception getting security
                               current: " + e);

            System.exit(1);
        }
    }
}
```

Visual Basic Client Example: Using the Bootstrap Object

Listing 4-5 shows how to program a Visual Basic client to use the Bootstrap object.

Listing 4-5 Programming a Client in Visual Basic

```
'Declare the Bootstrap object
Public oBootstrap As DITobj_Bootstrap

'Declare the FactoryFinder object
Public oBSFactoryFinder As DITobj_FactoryFinder

'Declare factory for Registrar object
Public oRegistrarFactory As DIUniversityB_RegistrarFactory

'Declare actual Registrar object
Public oRegistrarFactory As DIUniversityB_RegistrarFactory

....

'Create the Bootstrap object
Set oBootstrap = CreateObject("Tobj.Bootstrap")

'Connect to the WLE Domain
oBootstrap.Initialize "//host:port"

'Get the FactoryFinder for the WLE Domain

Set oBSFactoryFinder =
oBootstrap.CreateObject("Tobj.FactoryFinder")

'Get a factory for the Registrar object
'using the FactoryFinder method find_one_factory_by_id

Set oRegistrarFactory =
oBSFactoryFinder.find_one_factory_by_id("RegistrarFactoryID")

'Create a Registrar object

Set oRegistrar = oRegistrarFactory.find_registrar(exc)
```

5 FactoryFinder Interface

The FactoryFinder interface provides clients with one object reference that serves as the single point of entry into the WLE domain. The WLE NameManager provides the mapping of factory names to object references for the FactoryFinder. Multiple FactoryFinders and NameManagers together provide increased availability and reliability. In this release the level of functionality has been extended to support multiple domains.

Note: The NameManager is not a naming service, such as CORBAServices Naming Service, but is merely a vehicle for storing registered factories.

In the WLE environment, application factory objects are used to create objects that clients interact with to perform their business operations (for example, TellerFactory and Teller). Application factories are generally created during server initialization and are accessed by both remote clients and clients located within the server application.

The FactoryFinder interface and the NameManager services are contained in separate (nonapplication) servers. A set of application programming interfaces (APIs) is provided so that both client and server applications can access and update the factory information.

The support for multiple domains in this release benefits customers that need to scale to a large number of machines or who want to partition their application environment. To support multiple domains, the mechanism used to find factories in a WLE environment has been enhanced to allow factories in one domain to be visible in another. The visibility of factories in other domains is under the control of the system administrator.

Capabilities, Limitations, and Requirements

During server application initialization, application factories need to be registered with the NameManager. Clients can then be provided with the object reference of a FactoryFinder to allow them to retrieve a factory object reference based on associated names that were created when the factory was registered.

The following functional capabilities, limitations, and requirements apply to this release:

- ◆ The FactoryFinder interface is in compliance with the `CosLifeCycle::FactoryFinder` interface.
- ◆ Server applications can register and unregister application factories with the CORBAServices Naming Service.
- ◆ Clients can access objects using a single point of entry—the FactoryFinder.
- ◆ Clients can construct names for objects using a simplified BEA scheme made possible by WLE extensions to the CORBAServices interface or the more general CORBA scheme.
- ◆ Multiple FactoryFinders and NameManagers can be used to increase availability and reliability in the event that one FactoryFinder or NameManager should fail.
- ◆ Support for multiple domains. Factories in one domain can be configured to be visible in another domain under administrative control.
- ◆ Two NameManager services, at a minimum, must be configured, preferably on different machines, to maintain the factory-to-object reference mapping across process failures. If both NameManagers fail, the master NameManager, which has been keeping a persistent journal of the registered factories, recovers the previous state by processing the journal so as to re-establish its internal state.
- ◆ One NameManager must be designated as the Master and the Master NameManager must be started before the Slave. If the master NameManager is started after one or more Slaves, the Master assumes that it is in recovery mode instead of in initializing mode.

Functional Description

The WLE environment promotes the use of the factory design pattern as the primary means for a client to obtain a reference to an object. Through the use of this design pattern, client applications require a mechanism to obtain a reference to an object that acts as a factory for another object. Because the WLE environment has chosen CORBA as its visible programming model, the mechanism used to locate factories is modeled after the FactoryFinder as described in the CORBAservices Specification, Chapter 6 “Life Cycle Service,” December 1997, published by the Object Management Group.

In the CORBA FactoryFinder model, application servers register active factories with a FactoryFinder. When an application server’s factory becomes inactive, the application server removes the corresponding registration from the FactoryFinder. Client applications locate factories by querying a FactoryFinder. The client application can control the references to the factory object returned by specifying criteria that is used to select one or more references.

Locating a FactoryFinder

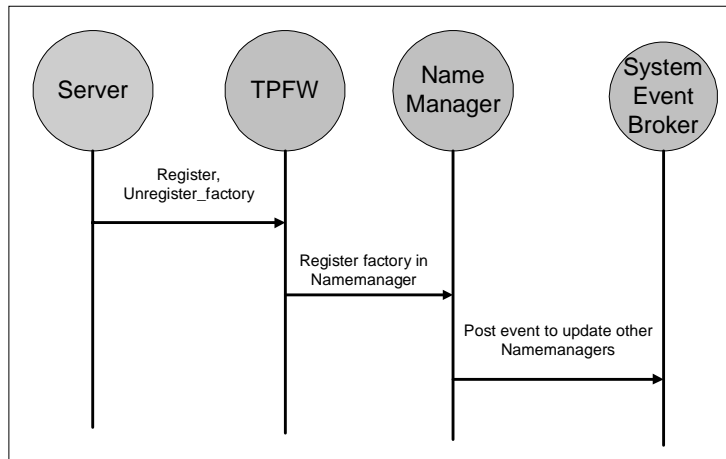
A client application must obtain a reference to a FactoryFinder before it can begin locating an appropriate factory. To obtain a reference to a FactoryFinder in the domain to which a client application is associated, the client application must invoke the `Tobj_Bootstrap::resolve_initial_references` operation with a value of “FactoryFinder”. This operation returns a reference to a FactoryFinder that is in the domain to which the client application is currently attached. For more information, see the section “Tobj_Bootstrap::register_callback_port” on page 4-14.

Note: The references to the FactoryFinder that are returned to the client application can be references to factory objects that are registered on the same machine as the FactoryFinder, on a different machine than the FactoryFinder, or possibly in a different domain than the FactoryFinder.

Registering a Factory

For a client application to be able to obtain a reference to a factory, an application server must register a reference to any factory object for which it provides an implementation with the FactoryFinder (See Figure 5-1). Using the WLE TP Framework, the registration of the reference for the factory object can be accomplished using the `TP::register_factory` operation, once a reference to a factory object has been created. The reference to the factory object, along with a value that identifies the factory, is passed to this operation. The registration of references to factory objects is typically done as part of initialization of the application (normally as part of the implementation of the operation `Server::initialize`).

Figure 5-1 Registering a Factory Object



When the server application is shutting down, it must unregister any references to factory objects that it has previously registered in the application server. This is done by passing the same reference to the factory object, along with the corresponding value used to identify the factory, to the `TP::unregister_factory` operation. Once unregistered, the reference to the factory object can then be destroyed. The process of unregistering a factory with the FactoryFinder is typically done as part of the implementation of the `Server::release` operation. For more information about these operations, see the section “Server Interface” on page 3-24.

C++ Mapping

Listing 5-1 shows the C++ class (static) methods. For more information about these methods, see the sections “TP::register_factory()” on page 3-63 and “TP::unregister_factory()” on page 3-65.

Listing 5-1 C++ Mappings for the Factory Registration Pseudo OMG IDL

```
#include <TP.h>

static void TP::register_factory(
    CORBA::Object_ptr factory_or, const char* factory_id);


static void TP::unregister_factory (
    CORBA::Object_ptr factory_or, const char* factory_id);
```

The `TP.h` header file contains the two method declarations and is to be included in any server application that wants to use these methods.

A server application generally includes this header file within the application file that contains the methods for application server initialization and release.

Locating a Factory

For a client application to request a factory to create a reference to an object, it must first obtain a reference to the factory object. The reference to the factory object is obtained by querying a `FactoryFinder` with specific selection criteria (see Figure 5-2). The criteria are determined by the format of the particular `FactoryFinder` interface and method used.

Figure 5-2 Locating a Factory Object

The WLE software extends the `CosLifeCycle::FactoryFinder` interface by introducing four methods in addition to the `find_factories()` method declared for the `FactoryFinder`. Therefore, using the `Tobj` extensions, a client can use either the `find_factories()` or `find_factories_by_id()` methods to obtain a list of application factories. A client can also use the `find_one_factory()` or `find_one_factory_by_id()` method to obtain a single application factory, and `list_factories()` to obtain a list of all registered factories.

The `CosLifeCycle::FactoryFinder` interface defines a `factory_key`, which is a sequence of `id` and `kind` strings conforming to the `CosNaming` Name shown below. The `kind` field of the `NameComponent` for all WLE application factories is set to the string `FactoryInterface` by the TP Framework when an application factory is registered. Applications supply their own value for the `id` field.

Assuming that the CORBAServices Life Cycle Service modules are contained in their own file (`ns.idl` and `lcs.idl`, respectively), only the OMG IDL code for that subset of both files that is relevant for using the WLE `FactoryFinder` is shown in the following listings.

CORBAservices Naming Service Module OMG IDL

Listing 5-2 shows the portions of the `ns.idl` file that are relevant to the `FactoryFinder`.

Listing 5-2 CORBAservices Naming OMG IDL

```
// ----- ns.idl -----

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

// This information is taken from CORBAservices: Common Object
// Services Specification, page 3-6. Revised Edition:
// March 31, 1995. Updated: November 1997. Used with permission by
// OMG.
```

CORBAservices Life Cycle Service Module OMG IDL

Listing 5-3 shows the portions of the `lcs.idl` file that are relevant to the `FactoryFinder`.

Listing 5-3 Life Cycle Service OMG IDL

```
// ----- lcs.idl -----

#include "ns.idl"

module CosLifeCycle{
    typedef CosNaming::Name Key;
    typedef Object Factory;
    typedef sequence<Factory> Factories;

    exception NoFactory{ Key search_key; }
```

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

        };

};

// This information is taken from CORBAServices: Common Object
// Services Specification, pages 6-10, 11. Revised Edition:
// March 31, 1995. Updated: November 1997. Used with permission by
// OMG.
```

Tobj Module OMG IDL

Listing 5-4 shows the Tobj Module OMG IDL.

Listing 5-4 Tobj Module OMG IDL

```
// ----- Tobj.idl -----

module Tobj {

    // Constants

    const string FACTORY_KIND = "FactoryInterface";

    // Exceptions

    exception CannotProceed { };
    exception InvalidDomain { };
    exception InvalidName { };
    exception RegistrarNotAvailable { };

    // Extension to LifeCycle Service

    struct FactoryComponent {
        CosLifeCycle::Key factory_key;
        CosLifeCycle::Factory factory_iior;
    };

    typedef sequence<FactoryComponent> FactoryListing;

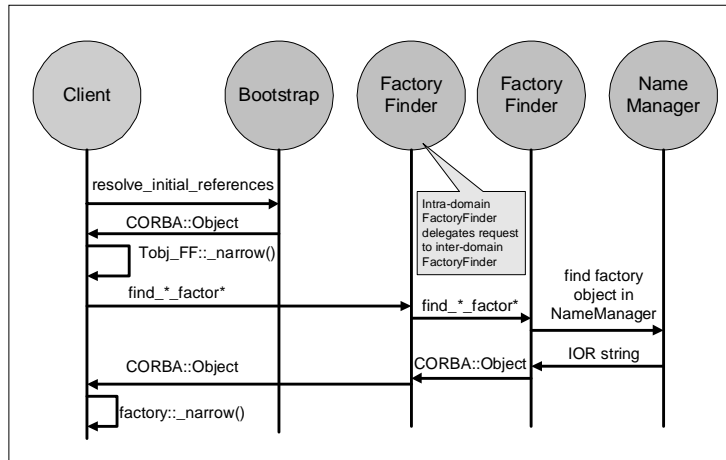
    interface FactoryFinder : CosLifeCycle::FactoryFinder {
        CosLifeCycle::Factory find_one_factory(in CosLifeCycle::Key
                                                factory_key)
    };
};
```

```
        raises (CosLifeCycle::NoFactory,
                CannotProceed,
                RegistrarNotAvailable);
    CosLifeCycle::Factory find_one_factory_by_id(in string
                                                factory_id)
        raises (CosLifeCycle::NoFactory,
                CannotProceed,
                RegistrarNotAvailable);
    CosLifeCycle::Factories find_factories_by_id(in string
                                                factory_id)
        raises (CosLifeCycle::NoFactory,
                CannotProceed,
                RegistrarNotAvailable);
    FactoryListing list_factories()
        raises (CannotProceed,
                RegistrarNotAvailable);
    };
};
```

Locating Factories in Another Domain

Typically, a `FactoryFinder` returns references to factory objects that are in the same domain as the `FactoryFinder` itself. However, it is possible to return references to factory objects in domains other than the domain in which a `FactoryFinder` exists. This can occur if a `FactoryFinder` contains information about factories that are resident in another domain (See Figure 5-3). A `FactoryFinder` finds out about these interdomain factory objects through configuration information that describes the location of these other factory objects.

When a `FactoryFinder` receives a request to locate a factory object, it must first determine if a reference to a factory object that meets the specified criteria exists. If there is registration information for a factory object that matches the criteria, the `FactoryFinder` must then determine if the factory object is local to the current domain or needs to be imported from another domain. If the factory object is from the local domain, the `FactoryFinder` returns the reference to the factory object to the client.

Figure 5-3 Inter-Domain FactoryFinder Interaction

If, on the other hand, the information indicates that the actual factory object is from another domain, the FactoryFinder delegates the request to an interdomain FactoryFinder in the appropriate domain. As a result, only a FactoryFinder in the same domain as the factory object will contain an actual reference to the factory object. The interdomain FactoryFinder is responsible for returning the reference of the factory object to the local FactoryFinder, which subsequently returns it to the client.

Why use WLE extensions?

The WLE software extends the interfaces defined in the CORBA services specification, Chapter 6 “Life Cycle Service,” December 1997, published by the Object Management Group, for the following reasons:

- ◆ Although the CORBA-defined approach is powerful and allows various selection criteria, the interface used to query a FactoryFinder can be complicated to use.
- ◆ Additionally, if the selection criterion specified by the client application is not specific enough, it is possible that more than one reference to a factory object may be returned. If this occurs, it is not immediately obvious what a client application should do next.
- ◆ Finally, the CORBA services specification did not specify a standardized mechanism through which an application server is to register a factory object.

Therefore, WLE extends the interfaces defined in the CORBA services specification to make using a `FactoryFinder` easier. The extensions are manifested as refined interfaces to the `FactoryFinder` that are derived from the interfaces specified in the CORBA services specification.

Creating Application Factory Keys

Two of the five methods provided by the `FactoryFinder` interface accept `CosLifeCycle::Keys`, which corresponds to `CosNaming::Name`. A client must be able to construct these keys.

The `CosNaming` Specification describes two interfaces that constitute a Names Library interface that can be used to create and manipulate `CosLifeCycle::Keys`. The pseudo OMG IDL statements for these interfaces is described in the following section.

Names Library Interface Pseudo OMG IDL

Note: This information is taken from the *CORBA services: Common Object Services Specification*, pp. 3-14 to 18. Revised Edition: March 31, 1995. Updated: November 1997. Used with permission by OMG.

To allow the representation of names to evolve without affecting existing client applications, it is desirable to hide the representation of names from the client application. Ideally, names themselves would be objects; however, names must be lightweight entities that are efficient to create, manipulate, and transmit. As such, names are presented to programs through the names library.

The names library implements names as pseudo-objects. A client application makes calls on a pseudo-object in the same way it makes calls on an ordinary object. Library names are described in pseudo-IDL (to suggest the appropriate language binding). C++ client applications use the same client language bindings for pseudo-IDL (PIDL) as they use for IDL.

Pseudo-object references cannot be passed across OMG IDL interfaces. As described in Chapter 3 of the *CORBA services: Common Object Services Specification*, in the section “The `CosNaming` Module,” the CORBA services Naming Service supports the `NamingContext` OMG IDL interface. The names library supports an operation to convert a library name into a value that can be passed to the name service through the `NamingContext` interface.

Note: It is not a requirement to use the names library in order to use the CORBAServices Naming Service.

The names library consists of two pseudo-IDL interfaces, the LNameComponent interface and the LName interface, as shown in Listing 5-5.

Listing 5-5 Names Library Interfaces in Pseudo-IDL

```
interface LNameComponent { // PIDL
    const short MAX_LNAME_STRLEN = 128;

    exception NotSet{ };
    exception OverFlow{ };

    string get_id
        raises (NotSet);
    void set_id(in string i)
        raises (OverFlow);
    string get_kind()
        raises(NotSet);
    void set_kind(in string k)
        raises (OverFlow);
    void destroy();
};

interface LName { // PIDL
    exception NoComponent{ };
    exception OverFlow{ };
    exception InvalidName{ };
    LName insert_component(in unsigned long i,
        in LNameComponent n)
        raises (NoComponent, OverFlow);
    LNameComponent get_component(in unsigned long i)
        raises (NoComponent);
    LNameComponent delete_component(in unsigned long i)
        raises (NoComponent);
    unsigned long num_components();
    boolean equal(in LName ln);
    boolean less_than(in LName ln);
    Name to_idl_form()
        raises (InvalidName);
    void from_idl_form(in Name n);
    void destroy();
};

LName create_lname();// C/C++
LNameComponent create_lname_component();// C/C++
```


CREATING A LIBRARY NAME COMPONENT

To create a library name component pseudo-object, use the following C/C++ function:

```
LNameComponent create_lname_component();    // C/C++
```

The returned pseudo-object can then be operated on using the operations shown in Listing 5-5.

CREATING A LIBRARY NAME

To create a library name pseudo-object, use the following C/C++ function:

```
LName create_lname();    // C/C++
```

The returned pseudo-object reference can then be operated on using the operations shown in Listing 5-5.

THE LNAMECOMPONENT INTERFACE

A name component consists of two attributes: `identifier` and `kind`. The `LNameComponent` interface defines the operations associated with these attributes, as follows:

```
string get_id()
raises(NotSet);
void set_id(in string k);
string get_kind()
raises(NotSet);
void set_kind(in string k);
```

`get_id`

The `get_id` operation returns the `identifier` attribute's value. If the attribute has not been set, the `NotSet` exception is raised.

`set_id`

The `set_id` operation sets the `identifier` attribute to the string argument.

`get_kind`

The `get_kind` operation returns the `kind` attribute's value. If the attribute has not been set, the `NotSet` exception is raised.

`set_kind`

The `set_kind` operation sets the `kind` attribute to the string argument.

THE LNAME INTERFACE

The following operations are described in this section:

- ◆ Destroying a library name component pseudo-object
- ◆ Inserting a name component
- ◆ Getting the i^{th} name component
- ◆ Deleting a name component
- ◆ Number of name components
- ◆ Testing for equality
- ◆ Testing for order
- ◆ Producing an OMG IDL form
- ◆ Translating an OMG IDL form
- ◆ Destroying a library name pseudo-object

DESTROYING A LIBRARY NAME COMPONENT PSEUDO-OBJECT

The `destroy` operation destroys library name component pseudo-objects.

```
void destroy();
```

INSERTING A NAME COMPONENT

A name has one or more components. Each component except the last is used to identify names of subcontexts. (The last component denotes the bound object.) The `insert_component` operation inserts a component after position `i`.

```
LName insert_component(in unsigned long i, in LNameComponent lnc)
raises(NoComponent, Overflow);
```

If component `i-1` is undefined and component `i` is greater than 1 (one), the `insert_component` operation raises the `NoComponent` exception.

If the library cannot allocate resources for the inserted component, the `Overflow` exception is raised.

GETTING THE i^{th} NAME COMPONENT

The `get_component` operation returns the i^{th} component. The first component is numbered 1 (one).

```
LNameComponent get_component(in unsigned long i)
raises(NoComponent);
```

If the component does not exist, the `NoComponent` exception is raised.

DELETING A NAME COMPONENT

The `delete_component` operation removes and returns the i^{th} component.

```
LNameComponent delete_component(in unsigned long i)
raises(NoComponent);
```

If the component does not exist, the `NoComponent` exception is raised.

After a `delete_component` operation has been performed, the compound name has one fewer component and components previously identified as $i+1\dots n$ are now identified as $i\dots n-1$.

NUMBER OF NAME COMPONENTS

The `num_components` operation returns the number of components in a library name.

```
unsigned long num_components();
```

TESTING FOR EQUALITY

The `equal` operation tests for equality with library name `ln`.

```
boolean equal(in LName ln);
```

TESTING FOR ORDER

The `less_than` operation tests for the order of a library name in relation to library name `ln`.

```
boolean less_than(in LName ln);
```

This operation returns true if the library name is less than the library name `ln` passed as an argument. The library implementation defines the ordering on names.

PRODUCING AN OMG IDL FORM

Pseudo-objects cannot be passed across OMG IDL interfaces. The library name is a pseudo-object; therefore, it cannot be passed across the OMG IDL interface for the CORBAservices Naming Service. Several operations in the NamingContext interface have arguments of an OMG IDL-defined structure, `Name`. The following PIDL operation on library names produces a structure that can be passed across the OMG IDL request.

```
Name to_idl_form()  
    raises(InvalidName);
```

If the name is of length 0 (zero), the `InvalidName` exception is returned.

TRANSLATING AN IDL FORM

Pseudo-objects cannot be passed across OMG IDL interfaces. The library name is a pseudo-object; therefore, it cannot be passed across the OMG IDL interface for the CORBAservices Naming Service. The NamingContext interface defines operations that return an IDL struct of type `Name`. The following PIDL operation on library names sets the components and `kind` attribute for a library name from a returned OMG IDL defined structure, `Name`.

```
void from_idl_form(in Name n);
```

DESTROYING A LIBRARY NAME PSEUDO-OBJECT

The `destroy` operation destroys library name pseudo-objects.

```
void destroy();
```

C++ Mapping

The Names Library pseudo OMG IDL interface maps to the C++ classes shown in Listing 5-6, which can be found in the `NamesLib.h` header file.

Two WLE extensions to CORBA are included to support scalability. Specifically, the `LNameComponent::set_id()` and `LNameComponent::set_kind()` methods raise an `Overflow` exception if the length of the input string exceeds `MAX_LNAME_STRLEN`. This length coincides with the maximum length of the WLE object ID (OID) and interface name. For a detailed description of the Library Name class, see the section “Names Library Interface Pseudo OMG IDL” on page 5-11.

Listing 5-6 Library Name Class

```
const short MAX_LNAME_STRLEN = 128;

class LNameComponent {
public:
    class NotSet{ };
    class Overflow{ };
    static LNameComponent* create_lname_component();
    void destroy();
    const char* get_id() const throw (NotSet);
    void set_id(const char* i) throw (Overflow);
    const char* get_kind() const throw (NotSet);
    void set_kind(const char* k) throw (Overflow);
};

class LName {
public:
    class NoComponent{ };
    class Overflow{ };
    class InvalidName{ };
    static LName* create_lname();
    void destroy();
    LName* insert_component(const unsigned long i,
                           LNameComponent* n)
        throw (NoComponent, Overflow);
    const LNameComponent* get_component(
        const unsigned long i) const
        throw (NoComponent);
    const LNameComponent* delete_component(
        const unsigned long i)
        throw (NoComponent);
    unsigned long num_components() const;
    CORBA::Boolean equal(const LName* ln) const;
    CORBA::Boolean less_than(
        const LName* ln) const; // not implemented
    CosNaming::Name* to_idl_form()
        throw (InvalidName);
};
```

```
        void from_idl_form(const CosNaming::Name& n);  
    };  
};
```

Java Mapping

The Names Library pseudo OMG IDL interface maps to the Java classes contained in the `com.beasys.Tobj` package, shown in Listing 5-7. All exceptions are contained in the same package.

For a detailed description of the Library Name class, refer to Chapter 3 in the *CORBA services: Common Object Services Specification*.

Listing 5-7 Java Mapping for LNameComponent

```
public class LNameComponent {  
    public static LNameComponent create_lname_component();  
    public static final short MAX_LNAME_STRING = 128;  
    public void destroy();  
    public String get_id() throws NotSet;  
    public void  set_id(String i) throws OverFlow;  
    public String get_kind() throws NotSet;  
    public void  set_kind(String k) throws OverFlow;  
};  
  
public class LName {  
  
    public static LName create_lname();  
    public void destroy();  
    public LName insert_component(long i, LNameComponent n)  
        throws NoComponent, OverFlow;  
    public LNameComponent get_component(long i)  
        throws NoComponent;  
    public LNameComponent delete_component(long i)  
        throws NoComponent;  
    public long num_components();  
    public boolean equal(LName ln);  
    public boolean less_than(LName ln); // not implemented  
    public org.omg.CosNaming.NameComponent[] to_idl_form()  
        throws InvalidName;  
    public void from_idl_form(org.omg.CosNaming.NameComponent[] nr);  
};
```

C++ Member Functions and Java Methods

This section describes the `FactoryFinder` C++ member functions and Java methods.

Note: All `FactoryFinder` member functions, except the `less_than` member function in `LName`, are implemented in both C++ and Java.

The following methods are described in this section:

- ◆ `CosLifeCycle::FactoryFinder::find_factories`
- ◆ `Tobj::Factoryfinder::find_one_factory`
- ◆ `Tobj::Factoryfinder::find_one_factory_by_id`
- ◆ `Tobj::Factoryfinder::find_factories_by_id`
- ◆ `Tobj::Factoryfinder::list_factories`

Note: The `CosLifeCycle::FactoryFinder::find_factories` method is the standard CORBA `CosLifeCycle` method. The four `Tobj` methods are extensions to the `CosLifeCycle` interface and, therefore, inherit the attributes of the `CosLifeCycle` interface.

CosLifeCycle::FactoryFinder::find_factories

Synopsis Obtains a sequence of factory object references.

C++ Mapping `CosLifeCycle::Factories *`
 `CORBA::Object_ptr CosLifeCycle::FactoryFinder::find_factories(
 const CosNaming::Name& factory_key)
 throw (CosLifeCycle::NoFactory);`

Java Mapping `import org.omg.CosLifeCycle.*;`

 `public org.omg.CORBA.Object[] find_factories(
 org.omg.CosNaming.NameComponent[] factory_key)
 throws org.omg.CosLifeCycle.NoFactory;`

Parameter `factory_key`
 This parameter is an unbounded sequence of NameComponents (tuple of <id, kind> pairs) that uniquely identifies a factory object reference.
 A NameComponent is defined as a having two members: an `id` and a `kind`, both of type string. The `id` field is used to represent the identity of factory object. The `kind` field is used to indicate how the value of the `id` field should be interpreted.
 References to factory object registered using the operation
 `TP::register_factory` will have a `kind` value of "FactoryInterface".

Exception `CORBA::BAD_PARAM`
 Indicates that the value of an input parameter has an inappropriate value or is invalid. Of particular importance, the exception is raised if no value or a NULL value for the parameter `factory_key` is specified.

`CosLifeCycle::NoFactory`
 Indicates that there are no factories registered that match the information in the `factory_key` parameter.

Description The `find_factories` method is called by an application to obtain a sequence of factory object references. The operation is passed a key used to identify the desired factory. The Key is a name, as defined by the CORBA services Naming service. More than one factory may match the key, and, if that is the case, the FactoryFinder returns a sequence of factories.

 The scope of the key is the FactoryFinder. The FactoryFinder 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.

Key values are considered equal if they are of equal length (same number of elements in the sequence), and if every NameComponent value in the Key matches the corresponding NameComponent value at the exact same location in the Key that was specified when the reference to the factory object was registered.

Return Values An unbounded sequence of references to Factory objects that match the information specified as the value of the `factory_key` parameter. In C++, the method returns a sequence of object references of type `CosLifeCycle::Factory`. In Java, the method returns an unbounded array of object references of type `org.omg.CORBA.Object`.

If the operation raises an exception, the return value is invalid and does not need to be released by the caller.

Tobj::FactoryFinder::find_one_factory

Synopsis Obtains a reference to a single factory object.

C++ Mapping `virtual CosLifecycle::Factory_ptr
 find_one_factory(const CosNaming::Name& factory_key) = 0;`

Java Mapping `public org.omg.CORBA.Object
 find_one_factory(org.omg.CosNaming.NameComponent[] factory_key)
 throws
 org.omg.CosLifecycle.NoFactory,
 com.beasys.Tobj.CannotProceed,
 com.beasys.Tobj.RegistrarNotAvailable;`

Parameter `factory_key`
 This parameter is an unbounded sequence of NameComponents (tuple of <id, kind> pairs) that uniquely identifies a factory object reference.
 A NameComponent is defined as a having two members: an `id` and a `kind`, both of type string. The `id` field is used to represent the identity of factory object. The `kind` field is used to indicate how the value of the `id` field should be interpreted.
 References to factory object registered using the operation `TP::register_factory` will have a `kind` value of "FactoryInterface".

Exceptions `CORBA::BAD_PARAM`
 Indicates that the value of an input parameter has an inappropriate value or is invalid. Of particular importance, the exception is raised if no value or a NULL value for the parameter `factory_key` is specified.

`CosLifecycle::NoFactory`
 Indicates that there are no factories registered that match the information in the `factory_key` parameter.

`Tobj::CannotProceed`
 Indicates that the FactoryFinder or NameManager encountered an internal error while attempting to locate a reference for a factory object.
 Error information is written to the user log.

`Tobj::RegistrarNotAvailable`
 Indicates that the FactoryFinder could not communicate with the NameManager.
 Error information is written to the user log.

Description The `find_one_factory` method is called by an application to obtain a reference to a single factory object whose key matches the value of the key specified as input to the method. If more than one factory object is registered with the specified key, the `FactoryFinder` selects one factory object based on the `FactoryFinder`'s load balancing scheme. As a result, invoking the `find_one_factory` method multiple times using the same key may return different object references.

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

Key values are considered equal if they are of equal length (same number of elements in the sequence), and if every `NameComponent` value in the `Key` matches the corresponding `NameComponent` value at the exact same location in the `Key` that was specified when the reference to the factory object was registered.

Return Values An object reference for a factory object. In C++, the method returns an object reference of type `CosLifeCycle::Factory`. In Java, the method returns an object reference of type `org.omg.CORBA.Object`.

If the operation raises an exception, the return value is invalid and does not need to be released by the caller.

Tobj::FactoryFinder::find_one_factory_by_id

Synopsis Obtains a reference to a single factory object.

C++ Mapping `virtual CosLifeCycle::Factory_ptr
 find_one_factory_by_id(const char * factory_id) = 0;`

Java Mapping `public org.omg.CORBA.Object
 find_one_factory_by_id(java.lang.String factory_id)
 throws
 org.omg.CosLifeCycle.NoFactory,
 com.beasys.Tobj.CannotProceed,
 com.beasys.Tobj.RegistrarNotAvailable;`

Parameter `factory_id`
 A NULL-terminated string that contains a value that is used to identify the registered factory object to be found.
 The value of the `factory_id` parameter is used as the value of the `id` field of a `NameComponent` that has a `kind` field with the value "FactoryInterface" when comparing against registered references for factory objects.

Exceptions `CORBA::BAD_PARAM`
 Indicates that the value of an input parameter has an inappropriate value or is invalid. Of particular importance, the exception is raised if no value or a NULL value for the parameter `factory_key` is specified.

`CosLifeCycle::NoFactory`
 Indicates that there are no factories registered that match the information in the `factory_key` parameter.

`Tobj::CannotProceed`
 Indicates that the `FactoryFinder` or `NameManager` encountered an internal error while attempting to locate a reference for a factory object.
 Error information is written to the user log.

`Tobj::RegistrarNotAvailable`
 Indicates that the `FactoryFinder` could not communicate with the `NameManager`.
 Error information is written to the user log.

Description The `find_one_factory_by_id` method is called by an application to obtain a reference to a single factory object whose registration ID matches the value of the ID specified as input to the method. If more than one factory object is registered with the

specified ID, the `FactoryFinder` selects one factory object based on the `FactoryFinder`'s load balancing scheme. As a result, invoking the `find_one_factory_by_id` operation multiple times using the same ID may return different object references.

The `find_one_factory_by_id` method behaves the same as the `find_one_factory` operation that was passed a `Key` that contains a single `NameComponent` with an `id` field that contains the same value as the `factory_id` parameter and a `kind` field that contains the value `"FactoryInterface"`.

The registered identifier for a factory is considered equal to the value of the `factory_id` parameter if the result of constructing a `CosLifeCycle::Key` structure containing a single `NameComponent` that has the `factory_id` parameter as the value of the `id` field and the value `"FactoryInterface"` as the value of the `kind` field. The values must match exactly in all respects (case, location, etc.).

Return Values An object reference for a factory object. In C++, the method returns an object reference of type `CosLifeCycle::Factory`. In Java, the method returns an object reference of type `org.omg.CORBA.Object`.

If the operation raises an exception, the return value is invalid and does not need to be released by the caller

Tobj::FactoryFinder::find_factories_by_id

Synopsis Obtains a sequence of one or more factory object references.

C++ Mapping `virtual CosLifeCycle::Factories *
 find_factories_by_id(const char * factory_id) = 0;`

Java Mapping `public org.omg.CORBA.Object[]
 find_factories_by_id(java.lang.String factory_id)
 throws
 org.omg.CosLifeCycle.NoFactory,
 com.beasys.Tobj.CannotProceed,
 com.beasys.Tobj.RegistrarNotAvailable;`

Parameter `factory_id`
 A NULL-terminated string that contains a value that is used to identify the registered factory object to be found.
 The value of the `factory_id` parameter is used as the value of the `id` field of a `NameComponent` that has a `kind` field with the value "FactoryInterface" when comparing against registered references for factory objects.

Exceptions `CORBA::BAD_PARAM`
 Indicates that the value of an input parameter has an inappropriate value or is invalid. Of particular importance, the exception is raised if no value or a NULL value for the parameter `factory_key` is specified.

`CosLifeCycle::NoFactory`
 Indicates that there are no factories registered that match the information in the `factory_key` parameter.

`Tobj::CannotProceed`
 Indicates that the `FactoryFinder` or `NameManager` encountered an internal error while attempting to locate a reference for a factory object.
 Error information is written to the user log.

`Tobj::RegistrarNotAvailable`
 Indicates that the `FactoryFinder` could not communicate with the `NameManager`.
 Error information is written to the user log.

Description The `find_factories_by_id` method is called by an application to obtain a sequence of one or more factory object references. The method is passed a NULL terminated string that contains the identifier of the factory to be located. If more than one factory object is registered with the specified ID, the `FactoryFinder` will return a list of object references for the matching registered factory objects.

The `find_factories_by_id` method behaves the same as the `find_factory` operation that was passed a `Key` that contains a single `NameComponent` with an `id` field that contains the same value as the `factory_id` parameter and a `kind` field that contains the value `"FactoryInterface"`.

The registered identifier for a factory is considered equal to the value of the `factory_id` parameter if the result of constructing a `CosLifeCycle::Key` structure containing a single `NameComponent` that has the `factory_id` parameter as the value of the `id` field and the value `"FactoryInterface"` as the value of the `kind` field. The values must match exactly in all respects (case, location, etc.).

Return Values An unbounded sequence of references to factory objects that match the information specified as the value of the `factory_key` parameter. In C++, the method returns a sequence of object references of type `CosLifeCycle::Factory`. In Java, the method returns an unbounded array of object references of type `org.omg.CORBA.Object`.

If the operation raises an exception, the return value is invalid and does not need to be released by the caller.

Tobj::Factoryfinder::list_factories

Synopsis	Obtains a lists of factory objects currently registered with the FactoryFinder.
C++ Mapping	<code>virtual FactoryListing * list_factories() = 0;</code>
Java Mapping	<pre>public com.beasys.Tobj.FactoryComponent[] list_factories() throws com.beasys.Tobj.CannotProceed, com.beasys.Tobj.RegistrarNotAvailable;</pre>
Exception	<p><code>Tobj::CannotProceed</code> Indicates that the FactoryFinder or NameManager encountered an internal error while attempting to locate a reference for a factory object. Error information is written to the user log.</p> <p><code>Tobj::RegistrarNotAvailable</code> Indicates that the FactoryFinder could not communicate with the NameManager. Error information is written to the user log.</p>
Description	The <code>list_factories</code> method is called by an application to obtain a list of the factory objects currently registered with the FactoryFinder. The method returns both the Key used to register the factory, as well as a reference to the factory object.
Return Values	<p>An unbounded sequence of <code>Tobj::FactoryComponent</code>. Each occurrence of a <code>Tobj::FactoryComponent</code> in the sequence contains a reference to the registered factory object, as well as the <code>CosLifeCycle::Key</code> that was used to register that factory object.</p> <p>If the operation raises an exception, the return value is invalid and does not need to be released by the caller.</p>

Automation Methods

This section describes the DITobj_FactoryFinder Automation methods.

DITobj_FactoryFinder.find_one_factory

Synopsis Obtains a single application factory.

MIDL Mapping `HRESULT find_one_factory(
 [in] VARIANT factory_key,
 [in,out,optional] VARIANT* exceptionInfo,
 [out,retval] IDispatch** returnValue);`

Automation Mapping `Function find_one_factory(factory_key, [exceptionInfo]) As Object`

Parameters `factory_key`
 This parameter contains a safe array of DICosNaming_NameComponent (<id, kind> value pairs) that uniquely identifies a factory object reference.

`exceptionInfo`
 An optional input argument that enables the application to get additional exception data if an error occurred.

Exceptions `NoFactory`
 This exception is raised if the FactoryFinder cannot find an application factory object reference that corresponds to the input `factory_key`.

`CannotProceed`
 This exception is raised if the FactoryFinder or CORBAServices Naming Service encounter an internal error during the search with the error being written to the user log (ULOG). Notify the operations staff immediately if this exception is raised. Depending on the severity of the internal error, the server running the FactoryFinder or CORBAServices Naming Service may have terminated. If a FactoryFinder service has terminated, start a new FactoryFinder service. If a CORBAServices Naming Service has terminated and there is another CORBAServices Naming Service running, start a new CORBAServices Naming Service. If no naming services servers are running, restart the application.

`RegistrarNotAvailable`
 This exception is raised if the FactoryFinder object cannot locate the CORBAServices Naming Service object. Notify the operations staff immediately if this exception is raised. If no naming services servers are running, restart the application.

Description This member function instructs the `FactoryFinder` to return one application factory object reference whose key matches the input `factory_key`. To accomplish this, the member function performs an equality match; that is, every `NameComponent <id, kind>` pair in the input `factory_key` must exactly match each `<id, kind>` pair in the application factory's key. If multiple factory keys contain the input `factory_key`, the `FactoryFinder` selects one factory key, based on an internally defined load balancing scheme. Invoking `find_one_factory` multiple times using the same `id` may return different object references.

Return Values Returns a reference to an interface pointer for the application factory.

DIObj_FactoryFinder.find_one_factory_by_id

Synopsis Obtains a single application factory.

MIDL Mapping `HRESULT find_one_factory_by_id(
 [in] BSTR factory_id,
 [in,out,optional] VARIANT* exceptionInfo,
 [out,retval] IDispatch** returnValue);`

Automation Mapping `Function find_one_factory_by_id(factory_id As String,
 [exceptionInfo] As Object`

Parameters `factory_id`
 This parameter represents a string identifier that is used to identify the kind or type of application factory. For some suggestions as to the composition of this string, see *Creating C++ Server Applications*.

`exceptionInfo`
 An optional input argument that enables the application to get additional exception data if an error occurred.

Exceptions `NoFactory`
 This exception is raised if the FactoryFinder cannot find an application factory object reference that corresponds to the input `factory_id`.

`CannotProceed`
 This exception is raised if the FactoryFinder or CORBAservices Naming Service encounter an internal error during the search, with the error being written to the user log (ULOG). Notify the operations staff immediately if this exception is raised. Depending on the severity of the internal error, the server running the FactoryFinder or the CORBAservices Naming Service may have terminated. If a FactoryFinder service has terminated, start a new FactoryFinder service. If a CORBAservices Naming Service has terminated and there is another CORBAservices Naming Service running, start a new CORBAservices Naming Service. If there are no naming services running, restart the application.

`RegistrarNotAvailable`
 This exception is raised if the FactoryFinder object cannot locate the CORBAservices Naming Service object. Notify the operations staff immediately if this exception is raised. If no naming service servers are running, restart the application.

Description This member function instructs the `FactoryFinder` to return one application factory object reference whose `id` in the key matches the method's input `factory_id`. To accomplish this, the member function performs an equality match (that is, the input `factory_id` must exactly match the `id` in the `<id,kind>` pair in the application factory's key). If multiple factory keys contain the input `factory_id`, the `FactoryFinder` selects one factory key, based on an internally defined load balancing scheme. Invoking `find_one_factory_by_id` multiple times using the same `id` may return different object references.

Return Values Returns a reference to an interface pointer for the application factory.

DIObj_FactoryFinder.find_factories_by_id

Synopsis Obtains a list of application factories.

```

MIDL Mapping      HRESULT find_factories_by_id(
                   [in] BSTR factory_id,
                   [in,out,optional] VARIANT* exceptionInfo,
                   [out,retval] VARIANT* returnValue);

```

Automation Mapping	Function find_factories_by_id(factory_id As String, [exceptionInfo])
--------------------	--

Parameters factory_id

This parameter represents a string identifier that will be used to identify the kind or type of application factory. The *Creating Client Applications* online document provides some suggestions as to the composition of this string.

```
exceptionInfo
```

An optional input argument that enables the application to get additional exception data if an error occurred.

Exceptions	NoFactory
------------	-----------

This exception is raised if the `FactoryFinder` cannot find an application factory object reference that corresponds to the input `factory_key` or `factory_id`.

CannotProceed

This exception is raised if the FactoryFinder or CORBAServices Naming Service encounter an internal error during the search with the error being written to the user log (ULOG). Notify the operations staff immediately if this exception is raised. Depending on the severity of the internal error, the server running the FactoryFinder or CORBAServices Naming Service may have terminated. If a FactoryFinder service has terminated, start a new FactoryFinder service. If a CORBAServices Naming Service has terminated and there is another CORBAServices Naming Service running, start a new CORBAServices Naming Service. If no naming services servers are running, restart the application.

RegistrarNotAvailable

This exception is raised if the FactoryFinder object cannot locate the CORBAServices Naming Service object. Notify the operations staff immediately if this exception is raised. If no naming services servers are running, restart the application.

Description	This member function instructs the <code>FactoryFinder</code> to return a list of application factory object references whose <code>id</code> in the keys match the method's input <code>factory_id</code> . To accomplish this, the member function performs an equality match (that is, the input <code>factory_id</code> must exactly match each <code>id</code> in the <code><id,kind></code> pair in the application factory's keys).
Return Values	Returns a variant containing an array of interface pointers to application factories.

DIObj_FactoryFinder.find_factories

Synopsis Obtains a list of application factories.

MIDL Mapping `HRESULT find_factories(
 [in] VARIANT factory_key,
 [in,out,optional] VARIANT* exceptionInfo,
 [out,retval] VARIANT* returnValue);`

Automation Mapping `Function find_factories(factory_key, [exceptionInfo])`

Parameters `factory_key`
 This parameter contains a safe array of DICosNaming_NameComponents (<id, kind> value pairs) that uniquely identifies a factory object reference.

`exceptionInfo`
 An optional input argument that enables the application to get additional exception data if an error occurred.

Exception `NoFactory`
 This exception is raised if the FactoryFinder cannot find an application factory object reference that corresponds to the input `factory_key`.

Description The `find_factories` method instructs the FactoryFinder to return a list of server application factory object references whose keys match the method's input key. The WLE system assumes that an equality match is to be performed. This means that for the two sequences of <id,kind> pairs (those corresponding to the input key and those in the application factory's keys), each are of equal length; for every pair in one sequence, there is an identical pair in the other.

Return Values Returns a variant containing an array of interface pointers to application factories.

DITobj_FactoryFinder.list_factories

Synopsis Lists all of the application factory names and object references.

MIDL Mapping `HRESULT list_factories(
 [in,out,optional] VARIANT* exceptionInfo,
 [out,retval] VARIANT* returnValue);`

Automation Mapping `Function list_factories([exceptionInfo])`

Parameter `exceptionInfo`
 An optional input argument that enables the application to get additional exception data if an error occurred.

Exception `CannotProceed`
 This exception is raised if the FactoryFinder or the CORBAservices Naming Service encounter an internal error during the search with the error being written to the user log (ULOG). Notify the operations staff immediately if this exception is raised. Depending on the severity of the internal error, the server running the FactoryFinder or the CORBAservices Naming Service may have terminated. If a FactoryFinder service has terminated, start a new FactoryFinder service. If a CORBAservices Naming Service has terminated and there is another CORBAservices Naming Service running, start a new CORBAservices Naming Service. If there are no naming service servers running, restart the application.

`RegistrarNotAvailable`
 This exception is raised if the FactoryFinder object cannot locate the CORBAservices Naming Service object. Notify the operations staff immediately if this exception is raised. It is possible that no naming service servers are running. Restart the application.

Description This method instructs the FactoryFinder to return a list containing all of the factory keys and associated object references for application factories registered with the CORBAservices Naming Service.

Return Values Returns a variant containing an array of DITobj_FactoryComponent objects. The FactoryComponent object consists of a variant containing an array of DICosNaming_NameComponent objects and an interface pointer to the application factory.

Programming Examples

This section describes how to program using the FactoryFinder interface.

Note: Remember to check for exceptions in your code.

Using the FactoryFinder Object

A FactoryFinder object is used by programmers to locate a reference to a factory object. The FactoryFinder object provides operations to obtain one or more references to factory objects based on the criteria specified.

There can be more than one FactoryFinder object in a process address space. Multiple references to a FactoryFinder object must be supported. A FactoryFinder object is semi-stateful in that it maintains state about the association between FactoryFinder objects within a domain and a particular IIOP Server Listener/Handler (ISL/ISH) through which to access the domain.

All FactoryFinder objects support the `CosLifeCycle::FactoryFinder` interface as defined in CORBA services Specification, Chapter 6 “Life Cycle Service,” December 1997, published by the Object Management Group. The interface contains one operation that is used to obtain one or more references to factory objects that meet the criteria specified.

Registering a Reference to a Factory Object

The following code fragment (Listing 5-8) shows how to use the TP Framework interface to register a reference to a factory object with a FactoryFinder.

Listing 5-8 Server Application: Registering a Factory

```
// Server Application: Registering a factory.  
// C++ Example.  
  
TP::register_factory( factory_obj.in( ), "TellerFactory" );
```

Obtaining a Reference to a FactoryFinder Object Using the CosLifecycle::FactoryFinder Interface

The following code fragment (Listing 5-9) shows how to use of the CORBA-compliant interface to obtain one or more references to factory objects.

Listing 5-9 Client Application: Getting a FactoryFinder Object Reference

```
// Client Application: Obtaining the object reference
// to factory objects.

CosLifecycle::Key_var factory_key = new CosLifecycle::Key( );
factory_key ->length(1);
factory_key[0].id = string_dupalloc( "strlen("TellerFactory") + 1 );
factory_key[0].kind = string_dupalloc(
                                strlen("FactoryInterface") + 1 );
strcpy( factory_key[0].id, "TellerFactory" );
strcpy( factory_key[0].kind, "FactoryInterface" );
CosLifecycle::Factories_var * flp = ff_np ->
                                find_factories( factory_key.in( ) );
```

Obtaining a Reference to a FactoryFinder Object Using the WLE Extensions Bootstrap object

The following code fragment (Listing 5-10) shows how to use of the WLE extensions Bootstrap object to obtain a reference to a FactoryFinder object.

Listing 5-10 Client Application: Finding One Factory Using the Tobj Approach

```
// Client Application: Finding one factory using the Tobj
// approach.

Tobj_Bootstrap * bsp = new Tobj_Bootstrap(
                                orb_ptr.in( ), host_port );
CORBA::Object_varptr ff_op = bsp ->
                                resolve_initial_references( "FactoryFinder" );
Tobj::FactoryFinder_ptrvar ff_np =
                                Tobj::FactoryFinder::_narrow( ff_op );
```

Using WLE Extensions to the FactoryFinder Object

WLE extends the FactoryFinder object with functionality to support similar capabilities to those provided by the operations defined by CORBA, but with a much simpler and more restrictive signature. The enhanced functionality is provided by defining the `Tobj::FactoryFinder` interface. The operations defined for the `Tobj::FactoryFinder` interface are intended to provide a focused, simplified form of the equivalent capability defined by CORBA. An application developer can choose to use the CORBA-defined or WLE extensions when developing an application. The interface `Tobj::FactoryFinder` is derived from the `CosLifeCycle::FactoryFinder` interface.

WLE extensions to the FactoryFinder object adhere to all the same rules as the FactoryFinder object defined in the CORBAservices Specification, Chapter 6 “Life Cycle Service,” December 1997, published by the Object Management Group.

The implementation of the extended FactoryFinder object requires users to supply either a `CosLifeCycle::Key`, as in the CORBA-defined `CosLifeCycle::FactoryFinder` interface, or a NULL-terminated string containing the identifier of a factory object to be located.

Obtaining One Factory using Tobj::FactoryFinder

The following code fragment (Listing 5-11) shows how to use the WLE extensions interface to obtain one reference to a factory object based on an identifier.

Listing 5-11 Client Application: Finding Factories Using the WLE Extensions Approach

```
CosLifeCycle::Factory_ptrvar fp_obj = ff_np ->
    find_one_factory_by_id( "TellerFactory" );
```

Obtaining One or More Factories using Tobj::FactoryFinder

The following code fragment (Listing 5-12) shows how to use the WLE extensions to obtain one or more references to factory objects based on an identifier.

Listing 5-12 Client Application: Finding One or More Factories Using the WLE Extensions Approach

```
CosLifeCycle::Factories * _var flp = ff_np ->  
                                find_factories_by_id( "TellerFactory" );
```

6 Security Service

This chapter describes the client security as provided by the WLE environmental objects.

The purpose of client security is to enable WLE clients to authenticate themselves via the IIOP Server Listener/Handler and to pass the WLE security checks. The purpose of client security is to enable WLE clients to authenticate themselves via the IIOP Server Listener/Handler and to pass the WLE security checks.

WLE client security provides two types of security authentication:

- ◆ An implementation of the security environmental objects for a CORBA Object Request Broker (ORB) environment. Client authentication is achieved using application programming interfaces (APIs) defined by CORBA security, although the authentication is performed by the IIOP Server Listener/Handler, not by the client ORB. Client security provides helper methods to create the data structures needed to call the standard CORBA authentication methods.
- ◆ An implementation of authentication similar to that found in the BEA TUXEDO system. Logon and logoff functions are provided that are easier to use than their CORBA counterparts. Logon passwords and data are secure traversing the network.

You can use either method to implement client security.

Capabilities and Limitations

This implementation of WLE client security has the following capabilities and limitations:

- ◆ Supports two types of authentication, as described above.
- ◆ Provides add-on methods to help generate the information needed for CORBA security from information specific to the WLE client, such as client name, client application password, user password (or user authentication data), and so forth.
- ◆ Implements authentication only.
- ◆ Allows remote WLE clients to authenticate themselves to WLE domains via the IIOP Server Listener/Handler so that clients can connect to a WLE domain with BEA TUXEDO style security activated.

Getting Initial References to the SecurityCurrent Object

You use the Bootstrap object to get an initial reference to the SecurityCurrent object. For a description of the Bootstrap object method, refer to “Tobj_Bootstrap::register_callback_port” on page 4-14.

Basic Security-level Requirements for WLE Clients

A client that connects to a WLE domain must provide security information according to the security level required by the WLE domain. Table 6-1 defines the security levels supported by WLE domains.

Table 6-1 Security Levels Supported by WLE Domains

Security Level	Definition
TOBJ_NOAUTH	No authentication is needed; however, the client can still authenticate itself, and must specify a user name and a client name, but no password.
TOBJ_SYSAUTH	The client must authenticate itself to the WLE domain, and must specify a user name, client name, and client application password.
TOBJ_APPAUTH	The client must provide information in addition to that which is required by TOBJ_SYSAUTH. If the default WLE authentication service is used in the WLE domain configuration, the client must provide a user password; otherwise, the client provides authentication data that is interpreted by the custom authentication service in the WLE domain.

Functional Components

This section describes the functional components of the Security Service.

Security Model

The security model in the WLE software defines the overall framework for security. The WLE product provides the flexibility to support different security mechanisms and policies that can be used to achieve the appropriate level of functionality and assurance. The WLE security model defines:

- ◆ Under what conditions clients may access objects
- ◆ What authentication of users and other principals is required to provide, who they are, and what they can do

The WLE security model is a combination of the security reference model defined in the CORBAservices Security Service specification¹ and the value-added extensions that provide a focused, simplified form of the security found in BEA TUXEDO. The WLE security model allows application developers to choose to use the security model defined by CORBA or the BEA extensions when developing an application.

Authentication of Principals

Authentication of principals, typically a human user or system entity, provides security officers with the ability to ensure that only registered principals have access to the objects in the system. An authenticated principal is used as the primary mechanism to control access to objects. The act of authenticating principals allows the security mechanisms to:

1. All references to the CORBAservices Security Service specification in this document are to the Revision 1.5, December 1998 edition, published by the Object Management Group.

- ◆ Control access to protected objects
- ◆ Identify the originator of a request

Controlling Access to Objects

The WLE security model provides a simple framework through which a security officer can limit access to authorized users only. Limiting access to objects allows security officers to prohibit access to objects by unauthorized principals. The access control framework consists of two parts:

- ◆ The object invocation policy that is enforced automatically on object invocation
- ◆ An application access policy that the user-written application can enforce

Administrative Control

The system administrator is responsible for setting security policies for client machines, server machines, and IIOP Listener/Handlers that interact with applications in their WLE domain. While the administrator sets the general policies, another person or group of people may be responsible for managing security (users, permissions, and so forth).

To provide system administrators the ability to define and enforce authentication of principals, the software provides a set of configuration parameters and utilities. Through these features, a system administrator can configure the WLE software to force the principals to be authenticated to access a system on which WLE software is installed.

Security Model Functional Description

This section provides a functional description of the security model.

Description

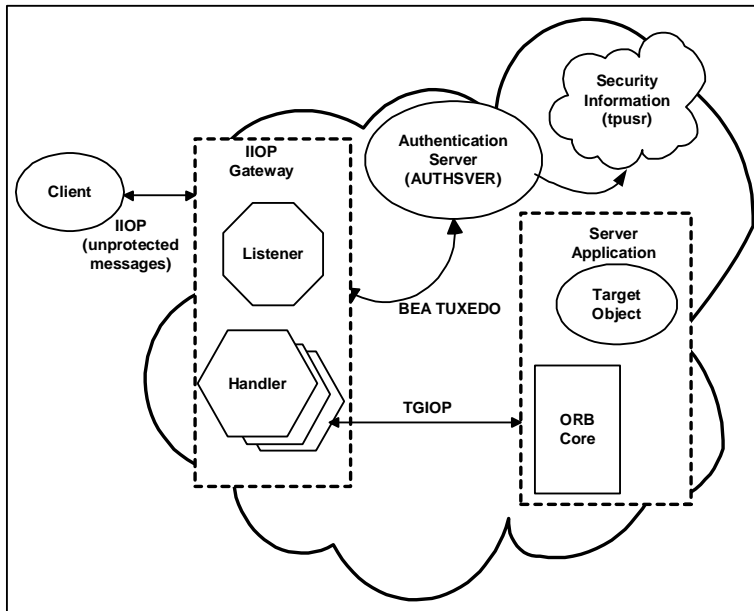
The security model adopted in the WLE software is based largely on the CORBA security model defined in the CORBAServices Security Service specification. Consequently, many of the concepts found in the CORBA security model apply to WLE security.

In addition to many of the interfaces defined by the CORBAServices Security Service, BEA provides extensions, in the form of derived interfaces. These extensions expose the security functionality found in the BEA TUXEDO system as CORBA interfaces that are found in the `Tobj` namespace. The interfaces in the `Tobj` namespace are intended to be familiar to developers of BEA TUXEDO applications and provide a focused, simplified form of the equivalent capability defined by CORBA. A programmer can choose to use the CORBA-defined security model or the BEA extensions when developing an application.

In a security model, there are usually defined sets of specific security policies. The WLE security model defines policies that specify whether a principal must be authenticated to use the system.

WLE implements a delegated trust authentication model. In this model, clients authenticate to a trusted system gateway process. In the case of WLE, the trusted system gateway process is the ISL/ISH. As part of a successful authentication, a security association, called a security context, is established between the client application and the ISL/ISH that is used to mediate access to objects. The WLE software associates the security context with the network connection over which the principal was authenticated. Except for the authentication exchange, this is currently the default behavior of the WLE system.

Figure 6-1 shows the security the security environment components.

Figure 6-1 WLE Security Environment

Logging on to the System

When a user or other principal wants to use the WLE system, the principal usually needs to authenticate and obtain credentials. The credentials obtained by the principal contain identity attributes that are used to control access to WLE servers.

The WLE Principal Authentication object provides a delegation mechanism to provide security to non-BEA branded clients. As illustrated in Figure 6-1, remote client applications perform authentication with the ISL/ISH, instead of with the server application itself. Consequently, the establishment of a security association is performed in the ISL/ISH, rather than in the server-side ORB.

In terms of CORBA security, the ISL/ISH acts as a CORBA-defined half-gateway into the WLE domain, and is, therefore, responsible for providing the security mechanisms that will be used in secure invocations for a given object.

Example of a Secure Object Invocation

The following is a description of what happens when a client invokes on a target object in a WLE environment:

- ◆ The client application obtains credentials for the user by authenticating itself with the WLE domain using a Principal Authenticator object. The request for authentication is sent to the ISL/ISH that relays the requests to an authentication server, which verifies the supplied information. If the verification process succeeds, the security system constructs a Credentials object that is used in all future invocations. The Credentials object for the principal is associated with the SecurityCurrent object that represents the security context for the current thread of execution.
- ◆ The client application invokes an object in the domain using its object reference. The request is packaged into an IIOP request and forwarded to the ISL/ISH that associates the request with the previously established security association. At this point, the ISL/ISH forwards the request, along with the credentials of the initiating principal, to an appropriate server process.

Authentication

This section provides a description of authentication.

Description of Authentication

When an active entity wants to use a secure object system, it authenticates itself and obtains credentials. The credentials contain its certified identity, and, optionally its privilege attributes, and control where and when the credentials can be used. In the WLE security model, an active entity must establish its rights to access objects in the system. The active entity must be either a principal, or a client that is acting on behalf of a principal.

A principal is defined to be either a user or a system entity that is registered in and authenticatable to the security system. Authentication may be accomplished in a number of ways. The most common way is for a user to supply a password. When a user or other principal is authenticated, the principal usually supplies:

- ◆ The principal's security name
- ◆ The authentication information needed by the particular authentication method used

Once authenticated, the principal's security attributes are maintained in the security system in a credential. The credentials provide a means for the security system to provide the principal's certified identity and describe the privileges granted to the particular principal.

Principals who initiate activities have one identity that may be used. The identity is represented in the system as attributes.

Authentication Mechanisms

As stated in "Logging on to the System" on page 6-7, the lack of interoperable security amongst the ORB vendors has resulted in it being necessary to utilize a delegation mechanism to provide authentication to client environments. The delegation mechanism used is similar to the mechanism found in BEA TUXEDO. Consequently, an authentication mechanism known as BEA TUXEDO-based security is supported in WLE domains. The implementation of this mechanism is layered on top of the security mechanism provided in BEA TUXEDO.

As in BEA TUXEDO, remote client applications perform authentication with the ISL/ISH instead of with the server application itself. Consequently, the establishment of a security association is actually performed in the ISL/ISH, rather than in the server-side ORB. Machines and server applications within a WLE domain are considered trusted. This trust is a result of a defined trust model that is based on the assumption that the machines and applications that make up the domain are under the control of administrators only.

Authentication of principals in an environment based on BEA TUXEDO requires the use of user names and passwords. Unlike most operating systems, BEA TUXEDO security defines three different authentication levels:

- ◆ TOBJ_NOAUTH—no authentication is needed.

- ◆ **TOBJ_SYSAUTH**—the principal must authenticate itself to the domain, and specify a user name, client name, and user password.
- ◆ **TOBJ_APPAUTH**—same as **TOBJ_SYSAUTH**, except that the principal must provide extra information. If the default authentication service provided in WLE software is configured, the principal must provide an application password; otherwise, the principal provides authentication data that is interpreted by a custom authentication service.

The level of authentication required is administratively controlled and is defined in the application's configuration. Because a client application is typically unaware of the level of authentication configured, it must query the security system to determine the authentication level required.

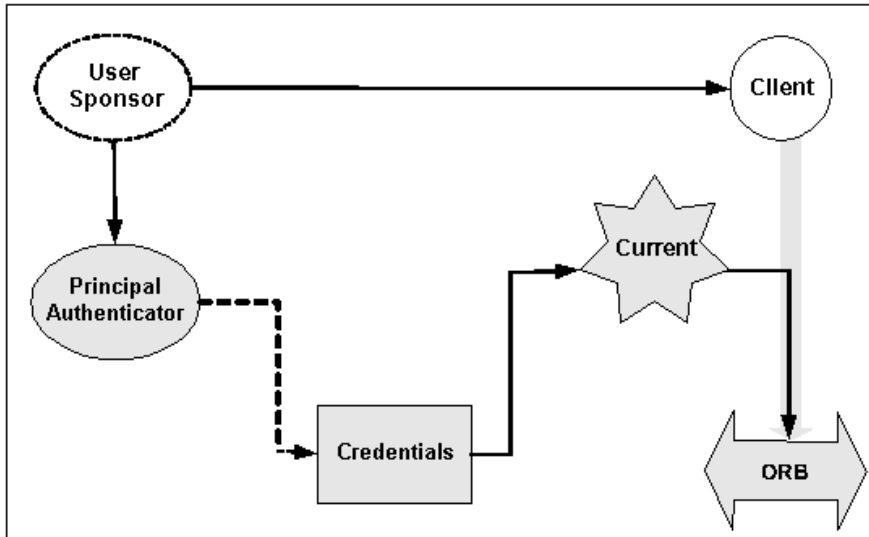
The configuration of the authentication level required is specified in an application's configuration, not on an object or method level. Consequently, if an application is configured to require authentication, all objects in the application require certified credentials for the principals. Applications can be configured to support either unauthenticated or authenticated principals. In unauthenticated scenarios, application developers may use a Principal Authenticator to assert a user name and client name, neither of which will be verified.

Because the BEA TUXEDO-based authentication mechanism is layered on top of the security mechanisms provided in BEA TUXEDO, it is possible for advanced customers to replace the Authentication Server that provides the default authentication mechanism with a custom implementation. A description of how to replace the Authentication Server in BEA TUXEDO is described in the BEA TUXEDO manuals and is outside the scope of this document.

Authentication Process

The process of authenticating a principal is done by a user sponsor (see Figure 6-2). A user sponsor is the code that calls the security interfaces for user authentication. In a WLE domain configured to use BEA TUXEDO-based security, the client application is the user sponsor.

In either case, the user provides identity and authentication data, such as a password, to the user sponsor. The user sponsor uses the Principal Authenticator object provided as part of the security implementation to make the calls necessary to authenticate the principal. The credentials for the authenticated principal are associated with the security system's implementation of the `SecurityCurrent` object and are represented by a `Credentials` object.

Figure 6-2 Authentication

PRINCIPAL AUTHENTICATOR OBJECT

The Principal Authenticator object is the object visible to the application that is responsible for the creation of Credentials for a given principal. A user or principal that requires authentication but has not been authenticated uses the Principal Authenticator object.

CREDENTIALS OBJECT

A Credentials object holds the security attributes of a principal. These security attributes include the principal's authenticated or unauthenticated identities. It also contains information for establishing security associations. The Credentials object provides methods to obtain the security attributes of the principals it represents.

SECURITYCURRENT OBJECT

The SecurityCurrent object represents the current execution context at the client and target object. The SecurityCurrent object provides methods to give access to security information associated with the execution context. The SecurityCurrent object gives access to the Credentials associated with the execution environment.

At any stage, a client or target object can find the default credentials for subsequent invocations by calling the `Current::get_credentials` method, asking for the invocation credentials. These default credentials are used in all invocations using object references.

Principal Authenticator Object

The Principal Authenticator object is used by a user or principal that requires authentication but has not been authenticated prior to calling the object system. The act of authenticating a principal results in the creation of a Credentials object that is made available as the default credentials for the application. The Credentials object is returned so it can be used for other methods on the Credentials.

The Principal Authenticator object is a singleton object; there is only a single instance allowed in a process address space. Multiple references to the Principal Authenticator object must be supported. The Principal Authenticator object is also stateless. A Credentials object is not associated with the Principal Authenticator object that created it.

All Principal Authenticator objects support the `SecurityLevel2::PrincipalAuthenticator` interface defined in the CORBAServices Security Service specification. This interface contains two methods that are used to accomplish the authentication of the principal. This is because authentication of principals may require more than one step. The `authenticate` method allows the caller to authenticate, and optionally select, attributes for the principal of this session.

Any invocation that fails because the security infrastructure does not permit the invocation will raise the standard exception `CORBA::NO_PERMISSION`. A method that fails because the feature requested is not supported by the security infrastructure implementation will raise the `CORBA::NO_IMPLEMENT` standard exception. Any parameter that has inappropriate values will raise the `CORBA::BAD_PARAM` standard exception.

The Principal Authenticator object is a locality-constrained object. Therefore, a Principal Authenticator object may not be used through the DII/DSI facilities of CORBA. Any attempt to pass a reference to this object outside of the current process, or any attempt to externalize it using `CORBA::ORB::object_to_string`, will result in the raising of the `CORBA::MARSHAL` exception.

WLE Extensions to the Principal Authenticator Object

BEA extends the Principal Authenticator object with functionality to support similar security to that found in BEA TUXEDO. The enhanced functionality is provided by defining the `Tobj::PrincipalAuthenticator` interface. This interface contains methods to provide similar capability to that available from BEA TUXEDO through the `tpinit` function.

The methods defined for the `Tobj::PrincipalAuthenticator` interface are intended to be familiar to developers of BEA TUXEDO applications and provide a focused, simplified form of the equivalent CORBA-defined capability. An application developer can choose to use the CORBA-defined or BEA extensions when developing an application. The interface `Tobj::PrincipalAuthenticator` is derived from the `CORBA SecurityLevel2::PrincipalAuthenticator` interface.

The extended Principal Authenticator object adheres to all the same rules as the Principal Authenticator object defined in the CORBA services Security Service specification.

The implementation of the extended Principal Authenticator object requires users to supply a user name, client name, and additional authentication data (for example, passwords) used for authentication. Because the information needs to be transmitted over the network to the ISL/ISH, it is protected to ensure confidentiality. The protection must include encryption of any information provided by the user.

An extended Principal Authenticator object that supports the `Tobj::PrincipalAuthenticator` interface provides the same functionality as if the `SecurityLevel2::PrincipalAuthenticator` interface was used to perform the authentication of the principal. However, unlike the `SecurityLevel2::PrincipalAuthenticator::authenticate` method, the `logon` method defined on the `Tobj::PrincipalAuthenticator` interface does not return a `Credentials` object. As a result, multithreaded applications that need to use more than one principal identity are required to call the `Current::get_credentials` method immediately after the `logon` method to retrieve the `Credentials` object as a

result of the `logon` method. Retrieval of the Credentials object directly after a `logon` method should be protected with serialized access since it is possible for another thread to also perform a `logon` method.


Credentials Object

A Credentials object (see Figure 6-3) holds the security attributes of a principal. The Credentials object provides methods to obtain and the set security attributes of the principals it represents. These security attributes include its authenticated or unauthenticated identities and privileges. It also contains information for establishing security associations.

Credentials objects are created as the result of:

- ◆ Authentication
- ◆ Copying an existing Credentials object
- ◆ Asking for a Credentials object via the SecurityCurrent object

Figure 6-3 Credentials Object



There can be more than one Credentials object in a process address space. Multiple references to a Credentials object must be supported. A Credentials object is stateful. It maintains state on behalf of the principal for which it was created. This state includes any information necessary to determine the identity and privileges of the principal it

represents. Credentials objects are not associated with the Principal Authenticator object that created it, but must contain some indication of the authentication authority that certified the principal's identity.

All Credentials objects support the `SecurityLevel2::Credentials` interface. Any invocation that fails as a result of the security infrastructure determining that the client does not have permission raises the standard exception `CORBA::NO_PERMISSION`. A method that fails because the feature requested is not supported by the security infrastructure implementation raises the `CORBA::NO_IMPLEMENT` standard exception. Any parameter that has inappropriate values raises the `CORBA::BAD_PARAM` standard exception.

The Credentials object is a locality-constrained object. Therefore, a Credentials object may not be used through the DII/DSI facilities. Any attempt to pass a reference to this object outside of the current process, or any attempt to externalize it using `CORBA::ORB::object_to_string`, will result in the raising of the `CORBA::MARSHAL` exception.

SecurityCurrent Object

The SecurityCurrent object (see Figure 6-4) represents the current execution context at both client and target objects. The SecurityCurrent object represents service-specific state information associated with the current execution context. Both clients and servers have SecurityCurrent objects that represent state associated with the thread of execution and the capsule (process) in which the thread is executing (their execution contexts).

The security SecurityCurrent object is a singleton object; there is only a single instance allowed in a process address space. Multiple references to the SecurityCurrent object must be supported.

The SecurityCurrent object is stateful. The methods of the SecurityCurrent object are intended to return information about the state associated with the current execution context. This includes information specific to both the thread of execution that is used to make the call on the SecurityCurrent object, as well as the capsule (process) to which the thread belongs. Changes in state associated purely with the thread, and not with any broader execution context, will remain until the thread terminates or until more state changes are made. State changes associated with a broader execution context (like a process) persist across multiple invocations of methods in the target object, until it is further modified through methods of the SecurityCurrent object or by other means.

Consequently, thread-specific methods called on the SecurityCurrent object are performed on the state associated with the calling thread. The thread in which the SecurityCurrent object was obtained has no influence on the behavior of these methods.


The CORBAservices Security Service specification defines two interfaces for the SecurityCurrent object associated with security:

- ◆ SecurityLevel1::Current, which derives from CORBA::Current
- ◆ SecurityLevel2::Current, which derives from the SecurityLevel1::Current interface

Both interfaces give access to security information associated with the execution context.

At any stage, a client or target object can find the default credentials for subsequent invocations by calling the Current::get_credentials method, asking for the invocation credentials. These default credentials are used in all invocations that use object references.

Figure 6-4 SecurityCurrent object



When the Current::get_attributes method is invoked by a client application, the attributes returned from the Credentials object are those of the user (for example, the one that was created by the Principal Authenticator object).

The SecurityCurrent object is a locality-constrained object. Therefore, a SecurityCurrent object may not be used through the DII/DSI facilities. Any attempt to pass a reference to this object outside of the current process, or any attempt to externalize it using CORBA::ORB::object_to_string, will result in the raising of the CORBA::MARSHAL exception.

Client Security API

The following client security application programming interface (API) modules are implemented as pseudo-objects on the client:

- ◆ CORBA module
- ◆ TimeBase module
- ◆ Security module
- ◆ Security Level 1 module
- ◆ Security Level 2 module
- ◆ Tobj module

The OMG Interface Definition Language (IDL) definitions for these modules are provided in the following sections.

CORBA Module

The Object Management Group (OMG) added the `CORBA::Current` interface to the CORBA module to support the Current pseudo-object. The change enables the CORBA module to support Security Replaceability and Security Level 2.

Listing 6-1 shows the `CORBA::Current` interface OMG IDL statements.

Listing 6-1 CORBA::Current Interface OMG IDL Statements

```
module CORBA {  
    // Extensions to CORBA  
    interface Current {  
    };  
};  
  
// This information is taken from CORBAServices: Common Object  
// Services Specification, page 15-230. Revised Edition:
```

// March 31, 1995. Updated: November 1997. Used with permission by
OMG.

TimeBase Module

All data structures pertaining to the basic Time Service, Universal Time Object, and Time Interval Object are defined in the TimeBase module. This allows other services to use these data structures without requiring the interface definitions. The interface definitions and associated enums and exceptions are encapsulated in the TimeBase module.

Listing 6-2 shows the TimeBase module OMG IDL statements.

Listing 6-2 TimeBase Module OMG IDL Statements

```
// From time service
module TimeBase {
    // interim definition of type ulonglong pending the
    // adoption of the type extension by all client ORBs.
    struct ulonglong {
        unsigned long    low;
        unsigned long    high;
    };
    typedef ulonglong    TimeT;
    typedef short        Tdft;
    struct UtcT {
        TimeT            time;        // 8 octets
        unsigned long    inacclo;    // 4 octets
        unsigned short    inacchi;    // 2 octets
        Tdft              tdf;        // 2 octets
                                // total 16 octets
    };
};

// This information is taken from CORBAServices: Common Object
// Services Specification, p. 14-5. Revised Edition:
// March 31, 1995. Updated: November 1997. Used with permission by
// OMG.
```

Table 6-2 defines the TimeBase module data types.

Note: This information is taken from *CORBAservices: Common Object Services Specification*, p. 14-6. Revised Edition: March 31, 1995. Updated: November 1997. Used with permission by OMG.

Table 6-2 TimeBase Module Data Type Definitions

Data Type	Definition
Time ulonglong	<p>OMG IDL does not at present have a native type representing an unsigned 64-bit integer. The adoption of technology submitted against that RFP will provide a means for defining a native type representing unsigned 64-bit integers in OMG IDL.</p> <p>Pending the adoption of that technology, you can use this structure to represent unsigned 64-bit integers, understanding that when a native type becomes available, it may not be interoperable with this declaration on all platforms. This definition is for the interim, and is meant to be removed when the native unsigned 64-bit integer type becomes available in OMG IDL.</p>
Time TimeT	TimeT represents a single time value, which is 64 bit in size, and holds the number of 100 nanoseconds that have passed since the base time. For absolute time, the base is 15 October 1582 00:00.
Time Tdft	<p>Tdft is of size 16 bits short type and holds the time displacement factor in the form of seconds of displacement from the Greenwich Meridian. Displacements east of the meridian are positive, while those to the west are negative.</p>
Time UtcT	<p>UtcT defines the structure of the time value that is used universally in the service. When the UtcT structure is holding, a relative or absolute time is determined by its history. There is no explicit flag within the object holding that state information. The <code>inacclo</code> and <code>inacchi</code> fields together hold a value of type <code>InaccuracyT</code> packed into 48 bits. The <code>tdf</code> field holds time zone information. Implementation must place the time displacement factor for the local time zone in this field whenever it creates a Universal Time Object (UTO).</p> <p>The content of this structure is intended to be opaque; to be able to marshal it correctly, the types of fields need to be identified.</p>

Security Module

The Security module defines the OMG IDL for security data types common to the other security modules. This module depends on the TimeBase module and must be available with any ORB that claims to be security ready.

Listing 6-3 shows the data types supported by the Security module.

Listing 6-3 Security Module OMG IDL Statements

```
module Security {
    typedef sequence<octet>    Opaque;

    // Extensible families for standard data types
    struct ExtensibleFamily {
        unsigned short    family_definer;
        unsigned short    family;
    };

    //security attributes
    typedef unsigned long    SecurityAttributeType;

    // identity attributes; family = 0
    const SecurityAttributeType    AuditId = 1;
    const SecurityAttributeType    AccountingId = 2;
    const SecurityAttributeType    NonRepudiationId = 3;

    // privilege attributes; family = 1
    const SecurityAttributeType    Public = 1;
    const SecurityAttributeType    AccessId = 2;
    const SecurityAttributeType    PrimaryGroupId = 3;
    const SecurityAttributeType    GroupId = 4;
    const SecurityAttributeType    Role = 5;
    const SecurityAttributeType    AttributeSet = 6;
    const SecurityAttributeType    Clearance = 7;
    const SecurityAttributeType    Capability = 8;

    struct AttributeType {
        ExtensibleFamily    attribute_family;
        SecurityAttributeType    attribute_type;
    };

    typedef sequence <AttributeType>    AttributeTypeLists;
    struct SecAttribute {
        AttributeType    attribute_type;
```

```
        Opaque          defining_authority;
        Opaque          value;
        // The value of this attribute can be
        // interpreted only with knowledge of type
    };

    typedef sequence<SecAttribute>  AttributeList;

    // Authentication return status
    enum AuthenticationStatus {
        SecAuthSuccess,
        SecAuthFailure,
        SecAuthContinue,
        SecAuthExpired
    };

    // Authentication method
    typedef unsigned long    AuthenticationMethod;

    enum CredentialType {
        SecInvocationCredentials;
        SecOwnCredentials;
        SecNRCredentials

    // Pick up from TimeBase
    typedef TimeBase::UtcT    UtcT;
};

// This information is taken from CORBAServices: Common Object
// Services Specification, pp. 15-193 to195. Revised Edition:
// March 31, 1995. Updated: November 1997. Used with permission by
// OMG.
```

Table 6-3 describes the Security module data type.

Table 6-3 Security Module Data Type Definition

Data Type	Definition
sequence<octet>	Data whose representation is known only to the Security Service implementation.

Security Level 1 Module

This section defines those interfaces available to client application objects that use only Level 1 Security functionality. This module depends on the CORBA module and the Security and TimeBase modules. The Current interface is implemented by the ORB.

Listing 6-4 shows the Security Level 1 module OMG IDL statements.

Listing 6-4 Security Level 1 Module OMG IDL Statements

```
module SecurityLevel1 {
    interface Current : CORBA::Current { // PIDL
        Security::AttributeList get_attributes(
            in Security::AttributeTypeList attributes
        );
    };
};

// This information is taken from CORBAServices: Common Object
// Services Specification, p. 15-198. Revised Edition:
// March 31, 1995. Updated: November 1997. Used with permission by
// OMG.
```

Security Level 2 Module

This section defines the additional interfaces available to client application objects that use Level 2 Security functionality. This module depends on the CORBA and Security modules.

Listing 6-5 shows the Security Level 2 module OMG IDL statements.

Listing 6-5 Security Level 2 Module OMG IDL Statements

```
module SecurityLevel2 {
    // Forward declaration of interfaces
    interface PrincipalAuthenticator;
```

```

interface Credentials;
interface Current;

// Interface Principal Authenticator
interface PrincipalAuthenticator {
    Security::AuthenticationStatus authenticate(
        in Security::AuthenticationMethod method,
        in string security_name,
        in Security::Opaque auth_data,
        in Security::AttributeList privileges,
        out Credentials creds,
        out Security::Opaque continuation_data,
        out Security::Opaque auth_specific_data
    );

    Security::AuthenticationStatus
        continue_authentication(
            in Security::Opaque response_data,
            inout Credentials creds,
            out Security::Opaque continuation_data,
            out Security::Opaque auth_specific_data
        );
};

// Interface Credentials
interface Credentials {
    Security::AttributeList get_attributes(
        in Security::AttributeTypeList attributes
    );
    boolean is_valid(
        out Security::UtcT expiry_time
    );
};

// Interface Current derived from SecurityLevel::Current
// providing additional operations on Current at this
// security level. This is implemented by the ORB.
interface Current : SecurityLevel::Current { // PIDL
    void set_credentials(
        in Security::CredentialType cred_type,
        in Credentials cred
    );

    Credentials get_credentials(
        in Security::CredentialType cred_type
    );
    readonly attribute PrincipalAuthenticator
        principal_authenticator;
};

```

```
};

// This information is taken from CORBAServices: Common Object
// Services Specification, pp. 15-198 to 200. Revised Edition:
// March 31, 1995. Updated: November 1997. Used with permission by
// OMG.
```

Tobj Module

This section defines the Tobj module interfaces.

This module provides the interfaces you use to program the BEA TUXEDO style of authentication.

Listing 6-6 shows the Tobj module OMG IDL statements.

Listing 6-6 Tobj Module OMG IDL Statements

```
//Tobj Specific definitions

module Tobj {
    const Security::AuthenticationMethod    TuxedoSecurity =
                                            0x54555800;

    //get_auth_type () return values
    enum AuthType {
        TOBJ_NOAUTH,
        TOBJ_SYSAUTH,
        TOBJ_APPAUTH
    };

    typedef sequence<octet>    UserAuthData;

    interface PrincipalAuthenticator :
        SecurityLevel2::PrincipalAuthenticator { // PIDL
        AuthType get_auth_type();

        Security::AuthenticationStatus logon(
            in string        user_name,
            in string        client_name,
            in string        system_password,
            in string        user_password,
```

```
        in UserAuthData      user_data
    );
    void logoff();

    void build_auth_data(
        in string              user_name,
        in string              client_name,
        in string              system_password,
        in string              user_password,
        in UserAuthData        user_data,
        out Security::Opaque   auth_data,
        out Security::AttributeList privileges
    );
};
```

Method Descriptions

This section describes the Security Service methods.

SecurityLevel1::Current::get_attributes

Synopsis	Returns attributes for the Current interface.
OMG IDL Definition	<pre>Security::AttributeList get_attributes(in Security::AttributeTypeList attributes); };</pre>
Arguments	<p>attributes</p> <p>The set of security attributes (privilege attribute types) whose values are desired. If this list is empty, all attributes are returned.</p>
Description	This method gets privilege (and other) attributes from the client's credentials for the Current interface.
Return Values	Table 6-4 describes the attribute values returned by get_attributes.

Table 6-4 Security Level 1 Returned Attribute Values

Attribute Type	Value
Security::Public	Empty (Public is returned when no authentication was performed.)
Security::AccessId	Null terminated ASCII string containing the WLE user name
Security::PrimaryGroupId	Null terminated ASCII string containing the WLE client name

Note: The other attribute types are never returned. The defining_authority field is always empty.

Note This information is taken from *CORBAservices: Common Object Services Specification*, pp. 15-103, 104. Revised Edition: March 31, 1995. Updated: November 1997. Used with permission by OMG.

SecurityLevel2::PrincipalAuthenticator::authenticate

Synopsis	Authenticates the client.
OMG IDL Definition	<pre>Security::AuthenticationStatus authenticate(in Security::AuthenticationMethod method, in string security_name, in Security::Opaque auth_data, in Security::AttributeList privileges, out Credentials creds, out Security::Opaque continuation_data, out Security::Opaque auth_specific_data);</pre>
Arguments	<p>method Must be Tobj::TuxedoSecurity. If method is invalid, authenticate raises CORBA::BAD_PARAM.</p> <p>security_name The WLE user name.</p> <p>auth_data As returned by Tobj::PrincipalAuthenticator::build_auth_data. If auth_data is invalid, authenticate raises CORBA::BAD_PARAM.</p> <p>privileges As returned by Tobj::PrincipalAuthenticator::build_auth_data. If privileges is invalid, authenticate raises CORBA::BAD_PARAM.</p> <p>creds Placed into the SecurityCurrent object.</p> <p>continuation_data Always empty.</p> <p>auth_specific_data Always empty.</p>
Description	This method authenticates the client via the IIOP Server Listener/Handler so that it can access a WLE domain.

Return Values Table 6-5 describes the values returned by `authenticate`.

Table 6-5 Authenticate Return Values

Return Value	Meaning
<code>Security::AuthenticationStatus::SecAuthSuccess</code>	The authentication succeeded.
<code>Security::AuthenticationStatus::SecAuthFailure</code>	The authentication failed, or the client was already authenticated and did not invoke <code>Tobj::PrincipalAuthenticator::logout</code> or <code>Tobj_Bootstrap::destroy_current</code> .

Note: The client can use BEA TUXEDO style authentication and call `Tobj::PrincipalAuthenticator::login` instead of `authenticate`.

Note This information is taken from *CORBAservices: Common Object Services Specification*, pp. 15-91, 92. Revised Edition: March 31, 1995. Updated: November 1997. Used with permission by OMG.

SecurityLevel2::PrincipalAuthenticator::continue_authentication

Synopsis Always fails.

OMG IDL `Security::AuthenticationStatus continue_authentication(
Definition in Security::Opaque response_data,
 inout Credentials creds,
 out Security::Opaque continuation_data,
 out Security::Opaque auth_specific_data
);`

Description Because the WLE software does authentication in one step, this method always fails and returns `Security::AuthenticationStatus::SecAuthFailure`.

Return Values Always returns `Security::AuthenticationStatus::SecAuthFailure`.

Note This information is taken from *CORBA services: Common Object Services Specification*, pp. 15-92, 93. Revised Edition: March 31, 1995. Updated: November 1997. Used with permission by OMG.

SecurityLevel2::Credentials::get_attributes

Synopsis Gets the attribute list attached to the credentials.

OMG IDL Definition

```
Security::AttributeList get_attributes(  
    in AttributeTypeList attributes  
);
```

Argument attributes
The set of security attributes (privilege attribute types) whose values are desired. If this list is empty, all attributes are returned.

Description This method returns the attribute list attached to the client's credentials. In the list of attribute types, you are required to include only the type value(s) for the attributes you want returned in the AttributeList. Attributes are not currently returned based on attribute-family or identities. In most cases, this is the same result you would get if you called `SecurityLevel1::Current::get_attributes()`, since there is only one valid set of credentials in the client at any instance in time. The results could be different if the credentials are not currently in use.

Return Values Returns attribute list.

Note This is information taken from *CORBA services: Common Object Services Specification*, p. 15-97. Revised Edition: March 31, 1995. Updated: November 1997. Used with permission by OMG.

SecurityLevel2::Credentials::is_valid

Synopsis Checks status of credentials.

OMG IDL Definition

```
boolean is_valid(  
    out Security::UtcT expiry_time  
);
```

Description This method returns TRUE if the credentials used are active at the time; that is, you did not call `Tobj::PrincipalAuthenticator::logoff` or `Tobj_Bootstrap::destroy_current`. If this method is called after `Tobj::PrincipalAuthenticator::logoff()`, FALSE is returned. If this method is called after `Tobj_Bootstrap::destroy_current()`, the `CORBA::BAD_INV_ORDER` exception is raised.

Return Values The expiration date returned will contain the maximum unsigned long long value in C++ and maximum long in Java. Until the unsigned long long datatype is adopted, the ulonglong datatype is substituted. The ulonglong datatype is defined as follows:

```
// interim definition of type ulonglong pending the  
// adoption of the type extension by all client ORBs.  
struct ulonglong {  
    unsigned long    low;  
    unsigned long    high;  
};
```

Note This information is taken from *CORBA services: Common Object Services Specification*, p. 15-97. Revised Edition: March 31, 1995. Updated: November 1997. Used with permission by OMG.

SecurityLevel2::Current::set_credentials

Synopsis Sets credentials type.

OMG IDL Definition

```
void set_credentials(  
    in Security::CredentialType cred_type,  
    in Credentials creds  
);
```

Arguments

`cred_type`
The type of credentials to be set; that is, invocation, own, or non-repudiation.

`creds`
The object reference to the Credentials object, which is to become the default.

Description This method can be used only to set `SecInvocationCredentials`; otherwise, `set_credentials` raises `CORBA::BAD_PARAM`. The credentials must have been obtained from a previous call to `SecurityLevel2::Current::get_credentials` or `SecurityLevel2::PrincipalAuthenticator::authenticate`.

Return Values None.

Note This information is taken from *CORBA services: Common Object Services Specification*, p. 15-104. Revised Edition: March 31, 1995. Updated: November 1997. Used with permission by OMG.

SecurityLevel2::Current::get_credentials

Synopsis Gets credentials type.

OMG IDL Definition

```
Credentials get_credentials(  
    in Security::CredentialType cred_type  
);
```

Arguments `cred_type`
The type of credentials to get.

Description This call can be used only to get `SecInvocationCredentials`; otherwise, `get_credentials` raises `CORBA::BAD_PARAM`. If no credentials are available, `get_credentials` raises `CORBA::BAD_INV_ORDER`.

Return Values Returns the active credentials in the client only.

Note This information is taken from *CORBA services: Common Object Services Specification*, p. 15-105. Revised Edition: March 31, 1995. Updated: November 1997. Used with permission by OMG.

SecurityLevel2::Current::principal_authenticator

Synopsis Returns the `PrincipalAuthenticator`.

OMG IDL `readonly attribute PrincipalAuthenticator`
Definition `principal_authenticator;`

Description The `PrincipalAuthenticator` returned by the `principal_authenticator` attribute is of actual type `Tobj::PrincipalAuthenticator`. Therefore, it can be used both as a `Tobj::PrincipalAuthenticator` and as a `SecurityLevel2::PrincipalAuthenticator`.

Note: This method raises `CORBA::BAD_INV_ORDER` if it is called on an invalid `SecurityCurrent` object.

Return Values Returns the `PrincipalAuthenticator`.

Tobj::PrincipalAuthenticator::get_auth_type

Synopsis Gets the type of authentication expected by the WLE domain.

OMG IDL Definition `AuthType get_auth_type();`

Description This method returns the type of authentication expected by the WLE domain. For the definition of `Tobj::AuthType` values, see the section “Basic Security-level Requirements for WLE Clients” on page 6-3.

Note: This method raises `CORBA::BAD_INV_ORDER` if it is called with an invalid `SecurityCurrent` object.

Return Values Returns the type of authentication required to access the WLE domain.

Tobj::PrincipalAuthenticator::logon

Synopsis Authenticates the client.

OMG IDL Definition

```
Security::AuthenticationStatus logon(  
    in string          user_name,  
    in string          client_name,  
    in string          system_password,  
    in string          user_password,  
    in UserAuthData    user_data  
);
```

Arguments user_name

The WLE user name. The authentication level is TOBJ_NOAUTH. If user_name is NULL or empty, or exceeds 30 characters, logon raises CORBA::BAD_PARAM.

client_name

The WLE client name. The authentication level is TOBJ_NOAUTH. If the client_name is NULL or empty, or exceeds 30 characters, logon raises the CORBA::BAD_PARAM exception.

system_password

The WLE client application password. The authentication level is TOBJ_SYSAUTH. If the client name is NULL or empty, or exceeds 30 characters, logon raises the CORBA::BAD_PARAM exception.

Note: The system_password must not exceed 30 characters.

user_password

The user password (needed for use by the default WLE authentication service). The authentication level is TOBJ_APPAUTH.

user_data

Data that is specific to the client application (needed for use by a custom WLE authentication service). The authentication level is TOBJ_APPAUTH.

Note: TOBJ_SYSAUTH includes the requirements of TOBJ_NOAUTH plus a client application password. TOBJ_APPAUTH includes the requirements of TOBJ_SYSAUTH plus additional information, such as a user password or user data.

Note: The user_password and user_data parameters are mutually exclusive, depending on the requirements of the authentication service used in the configuration of the WLE domain. The WLE default authentication service expects a user password. A customized authentication service may

require user data. The logon call raises the CORBA::BAD_PARAM exception if both `user_password` and `user_data` are specified.

Description For remote WLE clients, this method authenticates the client via the IIOP Server Listener/Handler so that the remote client can access a WLE domain. This method is functionally equivalent to `SecurityLevel2::PrincipalAuthenticator::authenticate`, but the parameters are oriented to WLE security.

Note: This method raises `CORBA::BAD_INV_ORDER` if it is called with an invalid `SecurityCurrent` object.

Return Values Table 6-6 describes the values returned by `logon`.

Table 6-6 Logon Return Values

Return Value	Meaning
<code>Security::AuthenticationStatus::SecAuthSuccess</code>	The authentication succeeded.
<code>Security::AuthenticationStatus::SecAuthFailure</code>	The authentication failed, or the client was already authenticated and did not call <code>Tobj::PrincipalAuthenticator::logout</code> or <code>Tobj_Bootstrap::destroy_current</code> .

Tobj::PrincipalAuthenticator::logoff

Synopsis Discards the WLE client authentication context.

OMG IDL Definition

```
void logoff();
```

Description This call discards the WLE client context, but does not close the network connections to the WLE domain. `Logoff` also invalidates the current credentials. After logging off, invocations using existing object references fail if the authentication type is not `TOBJ_NOAUTH`.

If the client is currently authenticated to a WLE domain, calling `Tobj_Bootstrap::destroy_current()` calls `logoff` implicitly.

Note: This method raises `CORBA::BAD_INV_ORDER` if it is called with an invalid `SecurityCurrent` object.

Return Values None.

Tobj::PrincipalAuthenticator::build_auth_data

Synopsis Creates authentication data and attributes for use by
 SecurityLevel2::PrincipalAuthenticator::authenticate.

OMG IDL

Definition `void build_auth_data(
 in string user_name,
 in string client_name,
 in string system_password,
 in string user_password,
 in UserAuthData user_data,
 out Security::Opaque auth_data,
 out Security::AttributeList privileges
);`

Arguments `user_name`
 The WLE user name.

`client_name`
 The WLE client name.

`system_password`
 The WLE client application password.

`user_password`
 The user password (default WLE authentication service).

`user_data`
 Client application-specific data (custom WLE authentication service).

`auth_data`
 For use by authenticate.

`privileges`
 For use by authenticate.

Note: If `user_name`, `client_name`, or `system_password` is NULL or empty, or exceeds 30 characters, the subsequent `authenticate` method invocation raises the CORBA::BAD_PARAM exception.

Note: The `user_password` and `user_data` parameters are mutually exclusive, depending on the requirements of the authentication service used in the configuration of the WLE domain. The WLE default authentication service expects a user password. A customized authentication service may require user data. If both `user_password` and `user_data` are specified, the subsequent authentication call raises the `CORBA::BAD_PARAM` exception.

Description This method is a helper function that creates authentication data and attributes to be used by `SecurityLevel2::PrincipalAuthenticator::authenticate`.

Note: This method raises `CORBA::BAD_INV_ORDER` if it is called with an invalid `SecurityCurrent` object.

Return Values None.

Automation Method Descriptions

This section describes the Automation Security Service methods.

DI`SecurityLevel2_Current`

The `DISecurityLevel2_Current` object is a BEA implementation of the CORBA Security model. In this release of the WLE software, the `get_attributes()`, `set_credentials()`, `get_credentials()`, and `PrincipalAuthenticator()` methods are supported.

DISecurityLevel2_Current.get_attributes

Synopsis Returns attributes for the Current interface.

MIDL Mapping

```
HRESULT get_attributes(  
    [in] VARIANT attributes,  
    [in,out,optional] VARIANT* exceptionInfo,  
    [out,retval] VARIANT* returnValue);
```

Automation Mapping Function get_attributes(attributes, [exceptionInfo])

Parameters attributes

The set of security attributes (privilege attribute types) whose values are desired. If this list is empty, all attributes are returned.

exceptioninfo

An optional input argument that allows the client application to get additional exception data if an error occurs. For the ActiveX client, all exception data is returned in the OLE Automation Error Object.

Description This method gets privilege (and other) attributes from the client's credentials for the Current interface.

Returns A variant containing an array of DISecurity_SecAttribute objects. Table 6-4 describes the attribute values returned by get_attributes.

DISecurityLevel2_Current.set_credentials

Synopsis Sets credentials type.

MIDL Mapping HRESULT set_credentials(
 [in] Security_CredentialType cred_type,
 [in] DISecurityLevel2_Credentials* cred,
 [in,out,optional] VARIANT* exceptionInfo);

Automation Mapping Sub set_credentials(cred_type As Security_CredentialType,
 cred As DISecurityLevel2_Credentials,
 [exceptionInfo])

Description This method can be used only to set SecInvocationCredentials; otherwise, set_credentials raises CORBA::BAD_PARAM. The credentials must have been obtained from a previous call to DISecurityLevel2_Current.get_credentials.

Arguments

 cred_type
 The type of credentials to be set; that is, invocation, own, or nonrepudiation.

 cred
 The object reference to the Credentials object, which is to become the default.

 exceptioninfo
 An optional input argument that allows the client application to get additional exception data if an error occurs. For the ActiveX client, all exception data is returned in the OLE Automation Error Object.

Returns None.

DISecurityLevel2_Current.get_credentials

Synopsis Gets credentials type.

MIDL Mapping

```
HRESULT get_credentials(  
    [in] Security_CredentialType cred_type,  
    [in,out,optional] VARIANT* exceptionInfo,  
    [out,retval] DISecurityLevel2_Credentials** returnValue);
```

Automation Mapping

```
Function get_credentials(cred_type As Security_CredentialType,  
    [exceptionInfo]) As DISecurityLevel2_Credentials
```

Description This call can be used only to get SecInvocationCredentials; otherwise, get_credentials raises CORBA::BAD_PARAM. If no credentials are available, get_credentials raises CORBA::BAD_INV_ORDER.

Arguments

cred_type
The type of credentials to get.

exceptioninfo
An optional input argument that allows the client application to get additional exception data if an error occurs. For the ActiveX client, all exception data is returned in the OLE Automation Error Object.

Returns A DISecurityLevel2_Credentials object for the active credentials in the client only.

DISecurityLevel2_Current.principal_authenticator

Synopsis Returns the `PrincipalAuthenticator`.

MIDL Mapping `HRESULT principal_authenticator([out, retval]
DITobj_PrincipalAuthenticator** returnValue);`

Automation Mapping Property `principal_authenticator` As `DITobj_PrincipalAuthenticator`

Description The `PrincipalAuthenticator` returned by the `principal_authenticator` property is of actual type `DITobj_PrincipalAuthenticator`. Therefore, it can be used as a `DISecurityLevel2_PrincipalAuthenticator`.

Note: This method raises `CORBA::BAD_INV_ORDER` if it is called on an invalid `SecurityCurrent` object.

Returns A `DITobj_PrincipalAuthenticator` object.

DITobj_PrincipalAuthenticator

The `DITobj_PrincipalAuthenticator` object is used to log in and log out of the WLE domain. In this release of the WLE software, the `authenticate`, `build_auth_data()`, `continue_authentication()`, `get_auth_type()`, `login()`, and `logout()` methods are implemented

DITobj_PrincipalAuthenticator.authenticate

Synopsis Authenticates the client application.

MIDL Mapping

```
HRESULT authenticate(  
    [in] long method,  
    [in] BSTR security_name,  
    [in] VARIANT auth_data,  
    [in] VARIANT privileges,  
    [out] DISecurityLevel2_Credentials** creds,  
    [out] VARIANT* continuation_data,  
    [out] VARIANT* auth_specific_data,  
    [in,out,optional] VARIANT* exceptionInfo,  
    [out,retval] Security_AuthenticationStatus* returnValue);
```

Automation Mapping

```
Function authenticate(method As Long, security_name As String,  
    auth_data, privileges, creds As DISecurityLevel2_Credentials,  
    continuation_data, auth_specific_data,  
    [exceptionInfo]) As Security_AuthenticationStatus
```

Arguments `method`
 Must be `Tobj::TuxedoSecurity`. If method is invalid, `authenticate` raises `CORBA::BAD_PARAM`.

`security_name`
 The WLE user name.

`auth_data`
 As returned by `DITobj_PrincipalAuthenticator.build_auth_data`. If `auth_data` is invalid, `authenticate` raises `CORBA::BAD_PARAM`.

`privileges`
 As returned by `DITobj_PrincipalAuthenticator.build_auth_data`. If `privileges` is invalid, `authenticate` raises `CORBA::BAD_PARAM`.

`creds`
 Placed into the `SecurityCurrent` object.

`continuation_data`
 Always empty.

`auth_specific_data`

Always empty.

`exceptioninfo`

An optional input argument that allows the client application to get additional exception data if an error occurs. For the ActiveX client, all exception data is returned in the OLE Automation Error Object.

Description This method authenticates the client via the IIOP Server Listener/Handler so that it can access a WLE domain.

Returns A `Security_AuthenticationStatus` Enum value. Table 6-5 describes the values returned by `authenticate`.

DITobj_PrincipalAuthenticator.build_auth_data

Synopsis Creates authentication data and attributes for use by
 DITobj_PrincipalAuthenticator.authenticate.

MIDL Mapping HRESULT build_auth_data(
 [in] BSTR user_name,
 [in] BSTR client_name,
 [in] BSTR system_password,
 [in] BSTR user_password,
 [in] VARIANT user_data,
 [out] VARIANT* auth_data,
 [out] VARIANT* privileges,
 [in,out,optional] VARIANT* exceptionInfo);

Automation Mapping Sub build_auth_data(user_name As String, client_name As String,
 system_password As String, user_password As String, user_data,
 auth_data, privileges, [exceptionInfo])

Arguments

user_name
 The WLE user name.

client_name
 The WLE client name.

system_password
 The WLE client application password.

user_password
 The user password (default WLE authentication service).

user_data
 Client application-specific data (custom WLE authentication service).

auth_data
 For use by authenticate.

privileges
 For use by authenticate.

exceptioninfo
 An optional input argument that allows the client application to get additional exception data if an error occurs. For the ActiveX client, all exception data is returned in the OLE Automation Error Object.

Note: If `user_name`, `client_name`, or `system_password` is NULL or empty, or exceeds 30 characters, the subsequent `authenticate` method invocation raises the `CORBA::BAD_PARAM` exception.

Note: The `user_password` and `user_data` parameters are mutually exclusive, depending on the requirements of the authentication service used in the configuration of the WLE domain. The WLE default authentication service expects a user password. A customized authentication service may require user data. If both `user_password` and `user_data` are specified, the subsequent authentication call raises the `CORBA::BAD_PARAM` exception.

Description This method is a helper function that creates authentication data and attributes to be used by `DITObj_PrincipalAuthenticator.authenticate`.

Note: This method raises `CORBA::BAD_INV_ORDER` if it is called with an invalid `SecurityCurrent` object.

Returns None.

DIObj_PrincipalAuthenticator.continue_authentication

Synopsis Always returns `Security::AuthenticationStatus::SecAuthFailure`.

MIDL Mapping

```
HRESULT continue_authentication(  
    [in] VARIANT response_data,  
    [in,out] DISecurityLevel2_Credentials** creds,  
    [out] VARIANT* continuation_data,  
    [out] VARIANT* auth_specific_data,  
    [in,out,optional] VARIANT* exceptionInfo,  
    [out,retval] Security_AuthenticationStatus* returnValue);
```

Automation Mapping

```
Function continue_authentication(response_data,  
    creds As DISecurityLevel2_Credentials, continuation_data,  
    auth_specific_data, [exceptionInfo]) As  
    Security_AuthenticationStatus
```

Description Because the WLE software does authentication in one step, this method always fails and returns `Security::AuthenticationStatus::SecAuthFailure`.

Returns Always returns `SecAuthFailure`.

DITobj_PrincipalAuthenticator.get_auth_type

Synopsis Gets the type of authentication expected by the WLE domain.

MIDL Mapping

```
HRESULT get_auth_type(  
    [in, out, optional] VARIANT* exceptionInfo,  
    [out, retval] Tobj_AuthType* returnValue);
```

Automation Mapping

```
Function get_auth_type([exceptionInfo]) As Tobj_AuthType
```

Argument `exceptioninfo`
An optional input argument that allows the client application to get additional exception data if an error occurs. For the ActiveX client, all exception data is returned in the OLE Automation Error Object.

Description This method returns the type of authentication expected by the WLE domain. For the definition of `Tobj::AuthType` values, see the section “Basic Security-level Requirements for WLE Clients” on page 6-3.

Note: This method raises `CORBA::BAD_INV_ORDER` if it is called with an invalid `SecurityCurrent` object.

Returned Result A reference to the `Tobj_AuthType` enumeration. Table 6-1 defines the valid values.

DITobj_PrincipalAuthenticator.logon

Synopsis Logs in to the WLE domain. The correct input parameters depend on the authentication level.

MIDL Mapping

```
HRESULT logon(  
    [in] BSTR user_name,  
    [in] BSTR client_name,  
    [in] BSTR system_password,  
    [in] BSTR user_password,  
    [in] VARIANT user_data,  
    [in,out,optional] VARIANT* exceptionInfo,  
    [out,retval] Security_AuthenticationStatus* returnValue);
```

Automation Mapping

```
Function logon(user_name As String, client_name As String,  
    system_password As String, user_password As String,  
    user_data, [exceptionInfo]) As Security_AuthenticationStatus
```

Description For remote WLE clients, this method authenticates the client via the IIOP Server Listener/Handler so that the remote client can access a WLE domain. This method is functionally equivalent to `DITobj_PrincipalAuthenticator.authenticate`, but the parameters are oriented to WLE security.

Arguments

`user_name`
The WLE user name. This parameter is required for `TOBJ_NOAUTH`, `TOBJ_SYSAUTH`, and `TOBJ_APPAUTH` authentication levels.

`client_name`
The WLE client name. This parameter is required for `TOBJ_NOAUTH`, `TOBJ_SYSAUTH`, and `TOBJ_APPAUTH` authentication levels.

`system_password`
The WLE client application password. This parameter is required for `TOBJ_SYSAUTH` and `TOBJ_APPAUTH` authentication levels.

`user_password`
The user password (default WLE authentication service). This parameter is required for the `TOBJ_APPAUTH` authentication level.

`user_data`

Application-specific data (custom authentication service). This parameter is required for the `TOBJ_APPAUTH` authentication level.

Note: If `user_name`, `client_name`, or `system_password` is `NULL` or empty, or exceeds 30 characters, the subsequent `authenticate` method invocation raises the `CORBA::BAD_PARAM` exception.

Note: If the authorization Level is `TOBJ_APPAUTH`, only one of `user_password` or `user_data` may be supplied.

`exceptioninfo`

An optional input argument that allows the client application to get additional exception data if an error occurs. For the ActiveX client, all exception data is returned in the OLE Automation Error Object.

Returns Table 6-6 describes the values returned by `logon`.

DIObj_PrincipalAuthenticator.logoff

Synopsis Discards the WLE client authentication context.

MIDL Mapping HRESULT logoff([in, out, optional] VARIANT* exceptionInfo);

Automation Mapping Sub logoff([exceptionInfo])

Description This call discards the WLE client context, but does not close the network connections to the WLE domain. Logoff also invalidates the current credentials. After logging off, calls using existing object references fail if the authentication type is not TOBJ_NOAUTH.

 If the client is currently authenticated to a WLE domain, calling Tobj_Bootstrap.destroy_current() calls logoff implicitly.

Argument exceptioninfo
 An optional input argument that allows the client application to get additional exception data if an error occurs. For the ActiveX client, all exception data is returned in the OLE Automation Error Object.

Returns None.

DISecurityLevel2_Credentials

The DISecurityLevel2_Credentials object is a BEA implementation of the CORBA Security model. In this release of the WLE software, the get_attributes() and is_valid() methods are supported.

DISecurityLevel2_Credentials.get_attributes

Synopsis Gets the attribute list attached to the credentials.

MIDL Mapping

```
HRESULT get_attributes(  
    [in] VARIANT attributes,  
    [in,out,optional] VARIANT* exceptionInfo,  
    [out,retval] VARIANT* returnValue);
```

Automation Mapping Function get_attributes(attributes, [exceptionInfo])

Argument attributes

The set of security attributes (privilege attribute types) whose values are desired. If this list is empty, all attributes are returned.

exceptioninfo

An optional input argument that allows the client application to get additional exception data if an error occurs. For the ActiveX client, all exception data is returned in the OLE Automation Error Object.

Description This method returns the attribute list attached to the client's credentials. In the list of attribute types, you are required to include only the type value(s) for the attributes you want returned in the AttributeList. Attributes are not currently returned based on attribute-family or identities. In most cases, this is the same result you would get if you called DISecurityLevel2.Current::get_attributes(), since there is only one valid set of credentials in the client at any instance in time. The results could be different if the credentials are not currently in use.

Arguments None.

Returns A variant containing an array of DISecurity_SecAttribute objects. Table 6-4 describes the attribute values returned by get_attributes.

DISecurityLevel2_Credentials.is_valid

Synopsis Checks status of credentials.

MIDL Mapping HRESULT is_valid(
 [out] IDispatch** expiry_time,
 [in,out,optional] VARIANT* exceptionInfo,
 [out,retval] VARIANT_BOOL* returnValue

Automation Mapping Function is_valid(expiry_time As Object,
 [exceptionInfo]) As Boolean

Description This method returns TRUE if the credentials used are active at the time; that is, you did not call `DITobj_PrincipalAuthenticator.logoff` or `destroy_current`. If this method is called after `DITobj_PrincipalAuthenticator.logoff()`, FALSE is returned. If this method is called after `destroy_current()`, the CORBA::BAD_INV_ORDER exception is raised.

Returns The output `expiry_time` as a `DITimeBase_UtcT` object set to max.

Programming Examples

This section provides programming examples that use the Security Service.

Note: In Listing 6-7 and Listing 6-8, notice that the `resolve_initial_references("SecurityCurrent")` method is used to get a reference to the `SecurityCurrent` object. The reference is then narrowed, assigned to `cur`, and used to get `PrincipalAuthenticator`. Refer to “`Tobj_Bootstrap::register_callback_port`” on page 4-14 for a description of this method.

C++ Example: Using WLE Extensions to Log on

Listing 6-7 shows how to program a Netscape Communicator client using the WLE extensions to CORBA security. The WLE extensions enable you to use BEA TUXEDO style authentication. The code in **boldface** shows the OMG method for logging on, which is an alternative to the BEA TUXEDO method. You may prefer the OMG method for log on. Note that the `build_auth_data` method is a BEA-specific method used to prepare data for the OMG method.

Listing 6-7 C++ Client Application Using WLE Extensions to CORBA Security to Log on (BEA TUXEDO Style Authentication)

```
/*      Copyright (c) 1998 BEA Systems, Inc.
All rights reserved
THIS IS PROPRIETARY
SOURCE CODE OF BEA Systems, Inc.
The copyright notice above does not
evidence any actual or intended
publication of such source code.
*/

//*****
// File: SECURITY_EXAMPLE.CPP
// Description: C++ Client Application Program.
//*****
```

```
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <CORBA.h>
#include <Tobj_Bootstrap.h>
#include <tobj_c.h>

#define TMMAXPASSLEN 30 // Currently Supported password length

void
main(int argc, char* argv[])
{
    const char *what = "Start of main";

    try {
        inti = 0;
        const char *orbid = 0;
        // look for the first arg that doesn't begin with
        // a '-', that should be either an empty string
        // indicating we should run as native client, or
        // a URL (/host:port) for a gateway and to run
        // as non-native
        for( i = 1 ; i < argc ; i++ )
            if( argv[i][0] != '-' ) {
                if( argv[i][0] == '\0' )
                    orbid = OBB::BEA_TOBJ_ID;
                else
                    orbid = OBB::BEA_IIOP_ID;
                break;
            }
        // Initialize ORB
        what = "Calling ORB_init";
        CORBA::ORB_var orb = CORBA::ORB_init(argc, argv, orbid);

        // Create bootstrap object
        what = "Instantiating Tobj_Bootstrap object";
        Tobj_Bootstrap Bs(orb, argv[1]);

        // Get the SecurityCurrent Object
        what = "Calling resolve_initial_references(SecurityCurrent)";
        CORBA::Object_var sec_current_obj =
            Bs.resolve_initial_references("SecurityCurrent");

        SecurityLevel2::Current_var sec_current =
            SecurityLevel2::Current::_narrow(sec_current_obj);

        // Get a Principal Authenticator
        what = "getting principal_authenticator";
```

```
SecurityLevel2::PrincipalAuthenticator_var Pa =
    sec_current->principal_authenticator();
// Narrow to a BEA Principal Authenticator
Tobj::PrincipalAuthenticator_var BeaPa =
    Tobj::PrincipalAuthenticator::_narrow(Pa);

// Get Authentication type
what = "calling get_auth_type";
Tobj::AuthType AuthType = BeaPa->get_auth_type();

// Set args according to security level
const char * userName = "guest";
const char * clientName = "simpclt";
char systemPassword[TMMAXPASSLEN + 1] = {'\0'};
char userPassword[TMMAXPASSLEN + 1] = {'\0'};
switch (AuthType)
{
    case Tobj::TOBJ_NOAUTH:
        printf(" No Passwords required\n");
        break;
    case Tobj::TOBJ_SYSAUTH:
        printf("System Password required\n");
        strcpy (systemPassword, "security");
        break;
    case Tobj::TOBJ_APPAUTH:
        printf("Both System and User Password required\n");
        strcpy (systemPassword, "security");
        strcpy (userPassword, "hello");
        break;
}
#ifdef USE_LOGON
    what = "calling logon";
    Security::AuthenticationStatus Status =
        BeaPa->logon(userName, clientName, systemPassword,
            userPassword, NULL);

    if ( Status != Security::SecAuthSuccess){
        printf("logon failed\n");
        exit(1);
    }
    else {
        printf("login succeeded\n");
    }
    // Security precautions
    if ( systemPassword[0] != '\0')
        memset (systemPassword, 0, sizeof(systemPassword));
    if ( userPassword[0] != '\0')
        memset (userPassword, 0, sizeof(userPassword));
#else

```



```

// Prepare args CORBA Seciop style for authentication
Tobj::UserAuthData userData;
Security::Opaque_var authData;
Security::AttributeList_var privileges;
SecurityLevel2::Credentials_var creds;
Security::Opaque_var continueData;
CORBA::ULong method = Tobj::TuxedoSecurity;
Security::Opaque_var authSpecificData;

// Use helper to build the authentication data
what = "calling the build_auth_data helper method";
BeaPa->build_auth_data(userName, clientName, systemPassword,
                      userPassword, userData, authData, privileges);
// Security precautions
if ( systemPassword[0] != '\0')
    memset (systemPassword, 0, sizeof(systemPassword));
if ( userPassword[0] != '\0')
    memset (userPassword, 0, sizeof(userPassword));
// Perform Corba Seciop authentication
what = "calling the authenticate method";
Security::AuthenticationStatus Status =
BeaPa->authenticate(method, userName, authData,
                  privileges, creds, continueData, authSpecificData);

if (Status != Security::SecAuthSuccess) {
    puts("Authenticate failed");
    exit(1);
}
else {
    puts("Authenticate succeeded");
}

#endif

}
catch (CORBA::UserException& e){
    (void)printf("What: %s\n", what);
    (void)printf("User exception: %s\n", e.get_id());
}
catch (CORBA::SystemException& e){
    (void)printf("What: %s\n", what);
    (void)printf("System exception: %s\n", e.get_id());
    (void)printf("System exception description is:\n%s\n",
e.OBB_errortext()?e.OBB_errortext():"\tNo description is available");
}
catch (...){
    (void)printf("What: %s\n", what);
    (void)printf("Unknown exception\n");
}
}

```

C++ Example: Using CORBA Security to Log on

The code shown in **boldface** type in the Listing 6-7 shows how to program a C++ client to log on to a WLE domain using OMG CORBA authentication.

Java Example: Using WLE Extensions to Log on

Listing 6-8 shows how to program a Netscape Communicator client using the WLE extensions to CORBA security. The WLE extensions enable you to use BEA TUXEDO style authentication.

Listing 6-8 Java Client Application Using WLE Extensions to CORBA Security to Log on

```
/*      Copyright (c) 1998 BEA Systems, Inc.
All rights reserved
THIS IS PROPRIETARY
SOURCE CODE OF BEA Systems, Inc.
The copyright notice above does not
evidence any actual or intended
publication of such source code.
*/

//*****
//File: SECURITY_CLIENT_EXAMPLE.JAVA
//Description: JAVA Client program.
//*****

import org.omg.CORBA.*;
import com.beasys.Tobj.*;
import com.beasys.*;
import com.beasys.TobjInternal.*;
import java.io.*;

public class security_client_example {
    public static void main(String args[])
    {
```

```
try {
    String HostPort = args[0];
    // Initialize ORB
    ORB orb = ORB.init();

    // Create Bootstrap Object
    Tobj_Bootstrap bs = new Tobj_Bootstrap(orb, null);

    // Get the SecurityCurrent Object
    org.omg.CORBA.Object secCurObj = bs.resolve_initial_references(
        "SecurityCurrent");

    org.omg.SecurityLevel2.Current secCur =
        org.omg.SecurityLevel2.CurrentHelper.narrow(secCurObj);

    // Get a principalauthenticator
    org.omg.SecurityLevel2.PrincipalAuthenticator authlev2=
        secCur.principal_authenticator();

    com.beasys.Tobj.PrincipalAuthenticator auth =
        (com.beasys.Tobj.PrincipalAuthenticator)

org.omg.SecurityLevel2.PrincipalAuthenticatorHelper.narrow(authlev2);

    // Get the auth type
    com.beasys.Tobj.AuthType authType = auth.get_auth_type();
    System.out.println( "authType =" + authType);
    byte[] userData = new byte[0];
    String userName = "guest";
    String clientName = "simpclt";
    String systemPassword = null;
    String userPassword = null;

    // Set args according to security level
    switch (authType.value())
    {
        case com.beasys.Tobj.AuthType._TOBJ_NOAUTH:
            System.out.println(" No Password Required ");
            break;
        case com.beasys.Tobj.AuthType._TOBJ_SYSAUTH:
            System.out.println("System Password Required ");
            systemPassword = "security";
            break;
        case com.beasys.Tobj.AuthType._TOBJ_APPAUTH:
            System.out.println("System Password Required & ");
            System.out.println("User Password Required ");
            systemPassword = "security";
```

```
        userPassword = "hello";
        break;
    }
    // Perform Tuxedo style logon
    org.omg.Security.AuthenticationStatus status =
        auth.logon(userName, clientName, systemPassword,
            userPassword, userData);
    System.out.println( "logon status =" + status);
    if (status != AuthenticationStatus.SecAuthSuccess)
        System.exit(1);
}
catch (UserException e){
    System.err.println("User exception: " + e);
    e.printStackTrace();
    System.exit(1);
}

catch (SystemException e){
    System.err.println("System exception: " + e);
    e.printStackTrace();
    System.exit(1);
}
}
}
```

Java Example: Getting Information from Privileges

Listing 6-9 shows how to use the Security Service to get information from privileges on a Java client.

Listing 6-9 Getting Information from Privileges

```
try {
    // Build empty attribute list to return all privileges
    org.omg.Security.AttributeType[] type_list =
        new org.omg.Security.AttributeType[0];
    // Get attributes from current
    org.omg.Security.SecAttribute[] privs =
        secCur.get_attributes(type_list);
    // Print attributes contents
    for (int i = 0 ; i < privs.length ; i++){
        switch( privs[i].attribute_type.attribute_type){
```

```
        case org.omg.Security.Public.value:
            // No security was specified.
            // Nothing to print.
            continue;
        case org.omg.Security.AccessId.value:
            // User name
            String user = new String(privs[i].value);
            System.out.println("User = " + user);
            continue;
        case org.omg.Security.PrimaryGroupId.value:
            // Client name
            String client = new String(privs[i].value);
            System.out.println("Client = " + client);
            continue;
    }
}
}
catch (SystemException e){
    System.out.println("Exception while checking attributes");
    System.exit(1);
}
```

Java Example: Checking the Validity of the Credentials Expiration Time

Listing 6-10 shows how to use the Security Service to check the validity of the Credentials expiration time on a Netscape Communicator client.

Listing 6-10 Checking Validity of Credentials Expiration Time on a Java Client

```
try {
    // Get Credentials from current
    org.omg.SecurityLevel2.Credentials cred = secCur.get_credentials(
        org.omg.Security.CredentialType.SecInvocationCredentials);
    // Verify credentials
    org.omg.TimeBase.UtcTHolder expiry_time =
        new org.omg.TimeBase.UtcTHolder();
    if (!cred.is_valid(expiry_time)){
        System.out.println(
            "Credentials are not valid any more");
    }
}
```

```
        System.exit(1);
    }
    // expiry_time contains credentials expiration in
    // 100 nanoseconds since 15 October 1582 00:00
}
catch (SystemException e){
    System.out.println("Exception while checking credentials");
    System.exit(1);
}
```

Java Example: Authentication Using SecurityLevel2::PrincipalAuthenticator

The following code fragment illustrates the use of the CORBA-compliant interfaces to perform authentication.

```
import org.omg.CORBA.*;
import com.beasys.Tobj.*;
import com.beasys.*;
import com.beasys.TobjInternal.*;
import java.io.*;

public class security_client
{
    public static void main(String[] args)
    {
        Tobj.PrincipalAuthenticator auth = null;

        try
        {
            String  HostPort = args[0];

            // Initialize ORB
            ORB orb = ORB.init();

            // Create bootstrap object
            Tobj_Bootstrap bs =
                new Tobj_Bootstrap(orb, "/" + HostPort);

            // Get security current
            org.omg.CORBA.Object secCurObj =
                bs.resolve_initial_references( "SecurityCurrent" );
            org.omg.SecurityLevel2.Current secCur2Obj =
                org.omg.SecurityLevel2.CurrentHelper.narrow(secCurObj);
```

```
// Get Principal Authenticator
org.omg.Security.PrincipalAuthenticator princAuth =
    secCur2Obj.principal_authenticator();
com.beasys.Tobj.PrincipalAuthenticator auth =
    Tobj.PrincipalAuthenticatorHelper.narrow(princAuth);

// Get Authentication type
com.beasys.Tobj.AuthType authType = auth.get_auth_type();

// Initialize arguments
String userName = "John";
String clientName = "Teller";
String systemPassword = null;
String userPassword = null;
byte[] userData = new byte[0];

// Prepare arguments according to security level requested
switch(authType.value())
{
    case com.beasys.Tobj.AuthType._TPNOAUTH:
        break;

    case com.beasys.Tobj.AuthType._TPSYSAUTH:
        systemPassword = "sys_pw";
        break;

    case com.beasys.Tobj.AuthType._TPAPPAUTH:
        systemPassword = "sys_pw";
        userPassword = "john_pw";
        break;
}

// Build security data
org.omg.Security.OpaqueHolder auth_data =
    new org.omg.Security.OpaqueHolder();
org.omg.Security.AttributeListHolder privs =
    new Security.AttributeListHolder();
auth.build_auth_data(userName, clientName, systemPassword,
                    userPassword, userData, authData,
                    privs);

// Authenticate user
org.omg.SecurityLevel2.CredentialsHolder creds =
    new org.omg.SecurityLevel2.CredentialHolder();
org.omg.Security.OpaqueHolder cont_data =
    new org.omg.Security.OpaqueHolder();
org.omg.Security.OpaqueHolder auth_spec_data =
    new org.omg.Security.OpaqueHolder();
```

```
        org.omg.Security.AuthenticationStatus status =
            auth.authenticate(com.beasys.Tobj.TuxedoSecurity.value,
                             0, userName, auth_data.value(),
                             privs.value(), creds, cont_data,
                             auth_spec_data);
        if (status != AuthenticatooinStatus.SecAuthSuccess)
            System.exit(1);
    }

    catch(UserException e )
    {
        System.err.println( "User exception: " + e );
        e.printStackTrace();
        System.exit(1);
    }

    catch(SystemException e )
    {
        System.err.println( "User exception: " + e );
        e.printStackTrace();
        System.exit(1);
    }
}
}
```

Java Example: Authentication Using Tobj::PrincipalAuthenticator

The following code fragment illustrates the use of the BEA extensions to perform authentication.

```
import org.omg.CORBA.*;
import com.beasys.Tobj.*;
import com.beasys.*;
import com.beasys.TobjInternal.*;
import java.io.*;

public class security_client
{
    public static void main(String[] args)
    {
        Tobj.PrincipalAuthenticator auth = null;
```



```
try
{
    String  HostPort = args[0];

    // Initialize ORB
    ORB orb = ORB.init();

    // Create bootstrap object
    Tobj_Bootstrap bs =
        new Tobj_Bootstrap(orb, "/" + HostPort);

    // Get security current
    org.omg.CORBA.Object secCurObj =
        bs.resolve_initial_references( "SecurityCurrent" );
    org.omg.SecurityLevel2.Current secCur2Obj =
        org.omg.SecurityLevel2.CurrentHelper.narrow(secCurObj);

    // Get Principal Authenticator
    org.omg.Security.PrincipalAuthenticator princAuth =
        secCur2Obj.principal_authenticator();
    com.beasys.Tobj.PrincipalAuthenticator auth =
        Tobj.PrincipalAuthenticatorHelper.narrow(princAuth);

    // Get Authentication type
    com.beasys.Tobj.AuthType authType = auth.get_auth_type();

    // Initialize arguments
    String userName = "John";
    String clientName = "Teller";
    String systemPassword = null;
    String userPassword = null;
    byte[] userData = new byte[0];

    // Prepare arguments according to security level requested
    switch(authType.value())
    {
        case com.beasys.Tobj.AuthType._TPNOAUTH:
            break;

        case com.beasys.Tobj.AuthType._TPSYSAUTH:
            systemPassword = "sys_pw";
            break;

        case com.beasys.Tobj.AuthType._TPAPPAUTH:
            systemPassword = "sys_pw";
            userPassword = "john_pw";
            break;
    }
}
```

```
// TUXEDO-style Authentification
org.omg.Security.AuthenticationStatus status =
    auth.login(userName, clientName, systemPassword,
               userPassword, userData);

// Check authentication result
if (status!= Security.AuthenticationStatus._SecAuthSuccess)
    System.exit(1);
}

catch(UserException e )
{
    System.err.println( "User exception: " + e );
    e.printStackTrace();
    System.exit(1);
}

catch(SystemException e )
{
    System.err.println( "User exception: " + e );
    e.printStackTrace();
    System.exit(1);
}

// Can now proceed with application
}
```

Java Example: Logging Off Using Tobj::PrincipalAuthenticator

The following code fragment illustrates the use of the BEA extensions to log off of a domain.

```
// Log off
try
{
    auth.logoff();
}
catch (SystemException e)
{
}
}
```

Java Example: Checking the Validity of Credentials

The following code fragment illustrates the use of the CORBA-compliant interfaces to check the validity of a principal's credentials.

```
try
{
    org.omg.Security.UtcTHolder expiry_time =
        new org.omg.Security.UtcTHolder();

    // Verify credentials
    if (!cred.is_valid(expiry_time))
    {
        System.out.println( "Credentials are not valid any more" );
        System.exit(1);
    }

    // expiry_time contains credentials expiration in
    // 100 nanoseconds since 15 October 1582 00:00
}

catch (SystemException e)
{
    System.out.println( "Exception while checking credentials" );
    System.exit(1);
}
```

Java Example: Getting Principal's Privileges

The following code fragment illustrates the use of the CORBA-compliant interfaces to retrieve the privileges and other attributes from a principal's credentials.

```
try
{
    // Build empty attribute type list to return all privileges
    org.omg.Security.AttributeType[] type_list =
        new org.omg.Security.AttributeType[0];

    // Get attributes from current
    org.omg.Security.SecAttribute[] privs =
        creds.get_attributes(type_list);
}
```

```
// Print attributes contents
for (int i = 0 ; i < privs.length ; i++)
{
    switch(privs[i].attribute_type)
    {
        case org.omg.Security.Public.value:
            // No security was specified – Nothing to print.
            continue;
        case org.omg.Security.AccessId.value:
            // User name
            String user = new String(privs[i].value);
            System.out.println( "User = " + user);
            continue;
        case org.omg.Security.PrimaryGroupId.value:
            // Client name
            String client = new String(privs[i].value);
            System.out.println( "Client = " + client);
            continue;
    }
}
}
catch (SystemException e)
{
    System.out.println( "Exception while getting privileges" );
    System.exit(1);
}
```

Java Example: Copying a Credentials Object

The following code fragment illustrates the use of the CORBA-compliant interfaces to copy a Credentials object. Copying a Credentials object results in a “deep copy,” possibly creating another security association based on the security technology used by the security provided. Copying a Credentials object that is on the SecurityCurrent object’s “own” list does not place the newly create copy on the “own” list. As a result, the newly created copy of the Credentials object can only be used as the default for one or more threads of the application, and will never be used as a default Credentials object for the capsule (process).

```
try
{
    org.omg.SecurityLevel2.Credentials    creds_copy =
        secCur2.copy();
}
```

```
catch
{
    System.out.println( "Exception while copying credential" );
    System.exit(1);
};
```

Java Example: Destroying a Credentials Object

The following code fragment illustrates the use of the CORBA-compliant interfaces to destroy a Credentials object. Typically, a Credentials object exists on the “own” list of the SecurityCurrent object. As a result, it should be removed from the “own” list prior to being destroyed. Destroying a Credentials object always results in the destruction of the security association between the client application and the target object, unless the security association is shared with another Credentials object.

```
try
{
    secCur2.remove_own_credentials( creds );
    secCur2.destroy();
}
catch
{
    System.out.println( "Exception while destroying credential" );
    System.exit(1);
};
```

Java Example: Getting the Principal Authenticator Object

The following code fragment illustrates the use of the CORBA-compliant interfaces to retrieve the Principal Authenticator object.

```
try
{
    org.omg.SecurityLevel2.PrincipalAuthenticator princAuth =
        secCurLev2.principal_authenticator();
}
catch (SystemException e )
{
    }
```

```
System.err.println( "Error getting principal authenticator" );
System.exit(1);
}
```

Java Example: Getting Credentials

The following code fragment illustrates the use of the CORBA-compliant interfaces to retrieve the privileges and other attributes from a principal's credentials.

```
try
{
    org.omg.SecurityLevel2.Credentials creds =
        secCur.get_credentials(
            org.omg.Security.CredentialType.SecInvocatonCredentials);
}
catch (SystemException e)
{
    System.out.println( "Exception while getting credentials" );
    System.exit(1);
}
```

Java Example: Setting Default Credentials

The following code fragment illustrates the use of the CORBA-compliant interfaces to set the privileges and other attributes for a principal's credentials as the credentials to be used for invocations in the current thread.

```
try
{
    secCur.set_credentials(
        org.omg.Security.CredentialType.SecInvocationCredentials,
        creds );
}
catch (SystemException e )
{
    System.out.println( "Exception while setting credentials" );
    System.exit(1);
}
```

Java Example: Getting a Principal's Privileges

The following code fragment illustrates the use of the CORBA-compliant interfaces to retrieve the privileges and other attributes from a principal's credentials.

```
try
{
    // Build empty attribute type list to return all privileges
    org.omg.Security.AttributeType[] type_list =
        new org.omg.Security.AttributeType[0];

    // Get attributes from current
    org.omg.Security.SecAttribute[] privs =
        secCur.get_attributes(type_list);

    // Print attributes contents
    for (int i = 0 ; i < privs.length ; i++)
    {
        switch(privs[i].attribute_type)
        {
            case org.omg.Security.Public.value:
                // No security was specified – Nothing to print.
                continue;
            case org.omg.Security.AccessId.value:
                // User name
                String user = new String(privs[i].value);
                System.out.println( "User = " + user);
                continue;
            case org.omg.Security.PrimaryGroupId.value:
                // Client name
                String client = new String(privs[i].value);
                System.out.println( "Client = " + client);
                continue;
        }
    }
}
catch (SystemException e)
{
    System.out.println( "Exception while getting privileges" );
    System.exit(1);
}
```

Java Example: Removing a Credentials Object from the “Own” List

The following code fragment illustrates the use of the CORBA-compliant interfaces to remove a Credentials object from the list of default Credentials objects for the current capsule (process). Removing a Credentials object from this list eliminates the ability for the removed Credentials object to be used as the capsule default. It does not destroy the Credentials object, or the security association that it represents.

```
try
{
    secCur2.remove_own_credentials( creds );
}
catch
{
    System.out.println( "Exception while removing credential" );
    System.exit(1);
};
```

Java Example: Getting Credentials of the Requesting Principal

The following code fragment illustrates the use of the CORBA-compliant interfaces to retrieve the credentials for the requesting principal.

```
try
{
    org.omg.SecurityLevel2.ReceivedCredentials recCreds =
        secCurLev2.received_credentials();
    org.omg.SecurityLevel2.Credentials creds =
        recCreds.accepting_credentials();
}
catch (SystemException e )
{
    System.err.println( "Exception getting received credentials" );
    System.exit(1);
}
```


Java Example: Getting the Principal's Privileges from Credentials

The following code fragment illustrates the use of the CORBA-compliant interfaces to retrieve the privileges and other attributes from the requesting principal's credentials.

```
try
{
    // Build empty attribute list to return all privileges
    org.omg.Security.AttributeType[] type_list =
        new org.omg.Security.AttributeType[0];

    // Get attributes from Credentials
    org.omg.Security.SecAttribute[] privs =
        creds.get_attributes(type_list);

    // Print attributes contents
    for (int i = 0 ; i < privs.length ; i++)
    {
        switch(privs[i].attribute_type)
        {
            case org.omg.Security.Public.value:
                // No security was specified – Nothing to print.
                continue;
            case org.omg.Security.AccessId.value:
                // User name
                String user = new String(privs[i].value);
                System.out.println( "User = " + user);
                continue;
            case org.omg.Security.PrimaryGroupId.value:
                // Client name
                String client = new String(privs[i].value);
                System.out.println( "Client = " + client);
                continue;
        }
    }
}
catch (SystemException e)
{
    System.out.println( "Exception while getting privileges" );
    System.exit(1);
}
```

Java Example: Getting the Principal's Privileges from the SecurityCurrent Object

The following code fragment illustrates the use of the CORBA-compliant interfaces to retrieve the privileges and other attributes for the requesting principal from the SecurityCurrent object.

```
try
{
    // Build empty attribute list to return all privileges
    org.omg.Security.AttributeType[] type_list =
        new org.omg.Security.AttributeType[0];

    // Get attributes from current
    org.omg.Security.SecAttribute[] privs =
        secCurLev2.get_attributes(type_list);

    // Print attributes contents
    for (int i = 0 ; i < privs.length ; i++)
    {
        switch(privs[i].attribute_type)
        {
            case org.omg.Security.Public.value:
                // No security was specified – Nothing to print.
                continue;
            case org.omg.Security.AccessId.value:
                // User name
                String user = new String(privs[i].value);
                System.out.println( "User = " + user);
                continue;
            case org.omg.Security.PrimaryGroupId.value:
                // Client name
                String client = new String(privs[i].value);
                System.out.println( "Client = " + client);
                continue;
        }
    }
}
catch (SystemException e)
{
    System.out.println( "Exception while getting privileges" );
    System.exit(1);
}
```

Java Example: Obtaining the SecurityCurrent Object

The following code fragment illustrates how a server application can obtain a reference to the securityCurrent object.

```
// Obtain a reference to the bootstrap object
Tobj_Bootstrap bs = TP::bootstrap();
// Get the Security Current
org.omg.CORBA.Object secCurObj =
    bs.resolve_initial_references( "SecurityCurrent" );
org.omg.SecurityLevel2.Current secCurLev2
    org.omg.SecurityLevel2.CurrentHelper.narrow(secCurObj);
```

Java Example: Getting Association Options

The following code fragment illustrates the use of the CORBA-compliant interfaces to get the association options in effect for the secure association with the remote principal.

```
try
{
    short options = recCreds.association_options_used();
}
catch (SystemException e)
{
    System.out.println( "Exception getting association options" );
    System.exit(1);
}
```

Java Example: Getting Delegation State

The following code fragment illustrates the use of the CORBA-compliant interfaces to get the delegation state of the remote principal for these credentials.

```
try
{
    org.omg.Security.DelegationState delState =
        recCreds.delegation_state();
    switch( delState )
    {
```

```
        case org.omg.Security.SecInitiator:
            System.out.println( "Acting on own behalf" );
            break;
        case org.omg.Security.SecDelegate:
            System.out.println( "acting on behalf of another" );
            break;
    }
}
catch (SystemException e)
{
    System.out.println( "Exception getting delegation state" );
    System.exit(1);
}
```

Java Example: Getting Delegation Mode

The following code fragment illustrates the use of the CORBA-compliant interfaces to get the delegation mode of the credentials.

```
try
{
    org.omg.Security.DelegationMode delMode =
        recCreds.delegation_mode();
    switch( delMode )
    {
        case org.omg.Security.SecDelModeNoDelegation:
            System.out.println( "Unusable for invocation" );
            break;
        case org.omg.Security.SecDelModeSimpleDelegation:
            System.out.println( "Usable for simple delegation" );
            break;
        case org.omg.Security.SecDelModeCompositeDelegation:
            System.out.println( "Usable for composite delegation" );
            break;
    }
}
catch (SystemException e)
{
    System.out.println( "Exception getting delegation mode" );
    System.exit(1);
}
```

7 Transaction Service

The WLE software provides an implementation of the CORBAservices Object Transaction Service that is described in Chapter 10 of the *CORBAservices: Common Object Services Specification*. This specification defines the interfaces for an object service that provides transactional functions.

This chapter describes how the WLE software implements that portion of the CORBAservices Object Transaction Service that is described as implementation specific.

This chapter provides the information that programmers need to write transactional applications for the WLE system. It describes the application programming interface (API) that you use to begin or terminate transactions, suspend or resume transactions, and get information about transactions.

Capabilities and Limitations

The following sections describe the capabilities and limitations of the Transaction Service.

Lightweight Clients with Delegated Commit

A lightweight client runs on a single-user, unmanaged desktop system that has irregular availability; that is, the owners may turn their desktop systems off when they are not in use. These single-user, unmanaged desktop systems should not be required to perform network functions like transaction coordination. In particular, unmanaged

systems should not be responsible for ensuring atomicity, consistency, isolation, and durability (ACID) properties across failures for transactions involving server resources. WLE remote clients are lightweight clients.

The Transaction Service allows lightweight clients to do delegated commit. Delegated commit means that the Transaction Service allows lightweight clients to begin and terminate transactions while the responsibility for transaction coordination is delegated to a transaction manager running on a server machine. The lightweight client does not need a local CORBAservices Object Transaction Service transaction manager.

Transaction Propagation

The CORBAservices Object Transaction Service specification states that a client can choose to propagate transaction context either implicitly or explicitly. This implementation of the CORBAservices Object Transaction Service *provides* implicit propagation. Explicit propagation *is strongly discouraged*.

Objects that are related to transaction context that are passed around using explicit transaction propagation *should not* be mixed with implicit transaction propagation APIs. It should be noted, however, that explicit propagation does not place any constraints on when transactional methods can be processed; there is no guarantee that all transactional methods will be completed before the transaction is committed.

Transaction Integrity

Checked transaction behavior provides transaction integrity by guaranteeing that a `commit` will not succeed unless all transactional objects involved in the transaction have completed the processing of their transactional requests. If implicit transaction propagation is used, the Transaction Service *provides* checked transaction behavior that is equivalent to that provided by the request/response interprocess communication models defined by The Open Group. The Transaction Service performs `reply` checks, `commit` checks, and `resume` checks, as described in the *CORBAservices Object Transaction Service Specification*.

Unchecked transaction behavior relies completely on the application to provide transaction integrity. If explicit propagation is used, the Transaction Service *does not* provide checked transaction behavior and transaction integrity *is not* guaranteed.

Transaction Termination

This implementation of the CORBAservices Object Transaction Service allows transactions to be terminated *only* by the client that created the transaction.

Note: The client may be a server object that requests the services of another object.

Flat Transactions

This implementation of the CORBAservices Object Transaction Service implements the flat transaction model.

Interoperability Between Remote Clients and the WLE Domain

This implementation of the CORBAservices Object Transaction Service *does not* support remote clients invoking methods on server objects in *different* WLE domains in the *same* transaction.

Remote clients with multiple connections to the same WLE domain may not make invocations to server objects on these separate connections within the same transaction. A `NO_PERMISSION` standard system exception is returned to the client.

Intradomain Interoperability

The WLE implementation of the CORBAservices Object Transaction Service supports native clients invoking methods on server objects in the WLE domain. In addition, server objects invoking methods on other objects in the same or in different processes in the same WLE domain is supported.

Network Interoperability

This implementation of the CORBAservices Object Transaction Service does not support the export or import of transactions to or from remote WLE domains.

Relationship of the Transaction Service to Transaction Processing

This section describes the relationship of the Transaction Service to various transaction processing servers, interfaces, protocols, and standards, as follows:

- ◆ Support of BEA TUXEDO ATMI servers

Servers using the WLE Transaction Service can make invocations on other BEA TUXEDO Application-to-Transaction Monitor Interface (ATMI) server processes in the same domain. This implementation of the CORBAservices Object Transaction Service *does not* support the following:

- ◆ Remote clients or native clients invoking ATMI services in the WLE domain
- ◆ ATMI services invoking objects

- ◆ Support of The Open Group XA interface

The Open Group Resource Managers are resource managers that can be involved in a distributed transaction by allowing their two-phase commit protocol to be controlled via The Open Group XA interface. This implementation of the CORBAservices Object Transaction Service supports interaction with The Open Group Resource Managers.

- ◆ Support of the OSI TP protocol

Open Systems Interconnect Transaction Processing (OSI TP) is the transactional protocol defined by the International Organization for Standardization (ISO). The WLE implementation of the CORBAservices Object Transaction Service *does not* support interactions with OSI TP transactions.

- ◆ Support of the LU 6.2 protocol

Systems Network Architecture (SNA) LU 6.2 is a transactional protocol defined by IBM. The WLE implementation of the CORBAServices Object Transaction Service *does not* support interactions with LU 6.2 transactions.

◆ Support of the ODMG standard

ODMG-93 is a standard defined by the Object Database Management Group (ODMG) that describes a portable interface to access Object Database Management Systems. The WLE implementation of the CORBAServices Object Transaction Service *does not* support interactions with ODMG transactions.

Process Failure

The Transaction Service monitors the participants in a transaction for failures and inactivity. One of the features that distinguishes the BEA TUXEDO system from other distributed application environments is the management tools for keeping the application running when failures occur. Because the WLE implementation of the CORBAServices Object Transaction Service is built upon the existing BEA TUXEDO transaction management system, it inherits the capabilities of the BEA TUXEDO system for keeping applications running.

Multithreaded Support

The WLE implementation of the CORBAServices Object Transaction Service supports single-threaded implementations *only*. Specifically, a client with an active transaction *cannot* make requests for the same transaction on multiple threads. However, it is possible to have multiple transactions serially active at the same time in a single thread.

OMG Interface Definition Language (IDL)

The CORBAServices Object Transaction Service OMG IDL is described in detail in Chapter 10 of the *CORBAServices: Common Object Services Specification*. The WLE implementation of the CORBAServices Object Transaction Service supports a *functionally complete* subset of the CORBAServices Object Transaction Service OMG IDL interfaces. For details, see the section “Transaction Service API” on page 7-7.

General Constraints

The following constraints apply:

- ◆ The WLE implementation of the CORBAservices Object Transaction Service imposes a limitation on programmers in that a server application object using transactions from the WLE Transaction Service library *needs* the WLE TP Framework functionality. A restriction imposed by the WLE TP Framework is that a client or a server object *cannot* invoke methods on an object that is infected with another transaction. The method invocation issued by the client or the server will return an exception. For further details on the TP Framework, see Chapter 3, “TP Framework”.
- ◆ A return from the rollback method on the Current object is asynchronous. A consequence of this is that the objects that were infected by the rolled back transaction get their states cleared by the WLE TP Framework *a little later*. This implies that *no* other client can infect these objects with a different transaction until the WLE TP Framework clears the states of these objects. This race condition exists for a very short amount of time and is typically not noticeable in a full-fledged application. A simple workaround for this race condition is to try the appropriate operation after a short (typically a 1-second) delay.
- ◆ In the WLE implementation of the CORBAservices Object Transaction Service, clients using other CORBAservices Object Transaction Service implementations *are not* supported.
- ◆ In the WLE implementation, clients may not make oneway method invocations within the context of a transaction to server objects having the NEVER, OPTIONAL or ALWAYS transaction policies. No error or exception will be returned to the client because it is a oneway method invocation; however, the method on the server object will not be executed. Also, an appropriate error message will be written to the log. Clients may make oneway method invocations within the context of a transaction to server objects having the IGNORE transaction policy. In this case, the method on the server object will be executed, but not in the context of a transaction. For further details on the transaction policies, see Chapter 2, “Implementation Configuration File (ICF)”.

Getting Initial References to the TransactionCurrent Object

To access the Transaction Service API and the extension to the Transaction Service API as described later in this chapter, an application needs to issue the following commands.

1. Create a Bootstrap object.
For details on creating a Bootstrap object, see Chapter 4, “Bootstrap Object”.
2. Invoke the `resolve_initial_reference("TransactionCurrent")` method on the Bootstrap object. The invocation returns a standard CORBA object pointer. For a description of this Bootstrap object method, see the section “Tobj_Bootstrap::resolve_initial_references” on page 4-15.
3. If an application is interested in only the Transaction Service APIs, a `CosTransactions::Current::_narrow()` should be issued on the object pointer returned from step 2 above. If an application is interested in the Transaction Service APIs with the extensions, a `Tobj::TransactionCurrent::_narrow()` should be issued on the object pointer returned from step 2 above.

Transaction Service API

The following sections describe the portions of the CosTransactions modules that are based on CORBA that are implemented in the WLE software to support the Transaction Service. For further details, refer to Chapter 10 of the *CORBA services: Common Object Services Specification*.

The definitions and interfaces supported by the Transaction Service in the WLE software are as follows:

- ◆ Data types
- ◆ Exceptions

- ◆ CORBA::Current interface
- ◆ Control interface
- ◆ CosTransactions::TransactionalObject interface

Data Types

Listing 7-1 shows the supported data types.

Listing 7-1 Data Types Supported by the Transaction Service

```
enum Status {  
    StatusActive,  
    StatusMarkedRollback,  
    StatusPrepared,  
    StatusCommitted,  
    StatusRolledBack,  
    StatusUnknown,  
    StatusNoTransaction  
    StatusPreparing,  
    StatusCommitting,  
    StatusRollingBack,  
};  
  
// This information is taken from CORBAServices: Common Object  
// Services Specification, p. 10-15. Revised Edition:  
// March 31, 1995. Updated: March 1997. Used with permission by OMG.
```

Exceptions

Listing 7-2 shows the supported exceptions.

Listing 7-2 Exceptions Supported by the Transaction Service

```
// Heuristic exceptions
exception HeuristicMixed {};
exception HeuristicHazard {};

// Other transaction-specific exceptions
exception SubtransactionsUnavailable {};
exception NoTransaction {};
exception InvalidControl {};
exception Unavailable {};
```

Table 7-1 describes the exceptions.

Note: This information is taken from *CORBA services: Common Object Services Specification*, pp. 10-16, 19, 20. Revised Edition: March 31, 1995. Updated: March 1997. Used with permission by OMG.

Table 7-1 Exceptions Supported by the Transaction Service

Exception	Description
HeuristicMixed	A request raises this exception to report that a heuristic decision was made and that some relevant updates have been committed and others have been rolled back.
HeuristicHazard	A request raises this exception to report that a heuristic decision was made, that the disposition of all relevant updates is not known, and that for those updates whose disposition is known, either all have been committed or all have been rolled back. Therefore, the <code>HeuristicMixed</code> exception takes priority over the <code>HeuristicHazard</code> exception.
SubtransactionsUnavailable	This exception is raised for the <code>Current</code> interface <code>begin</code> method if the client already has an associated transaction.
NoTransaction	This exception is raised for the <code>Current</code> interface <code>rollback</code> and <code>rollback_only</code> methods if there is no transaction associated with the client thread.
InvalidControl	This exception is raised for the <code>Current</code> interface <code>resume</code> method if the parameter is not valid in the current execution environment.
Unavailable	This exception is raised for the <code>Control</code> interface <code>get_terminator</code> and <code>get_coordinator</code> methods if the <code>Control</code> interface cannot provide the requested object.

Current Interface

The `Current` interface defines methods that allow a client of the Transaction Service to explicitly manage the association between threads and transactions. The `Current` interface also defines methods that simplify the use of the Transaction Service for most applications. These methods can be used to begin and end transactions, to suspend and resume transactions, and to obtain information about the current transaction.

The `CosTransactions` module defines the `Current` interface (shown in Listing 7-3).

Listing 7-3 Current Interface

```
// Current transaction
interface Current : CORBA::Current {
    void begin()
        raises(SubtransactionsUnavailable);
    void commit(in boolean report_heuristics)
        raises(
            NoTransaction,
            HeuristicMixed,
            HeuristicHazard
        );
    void rollback()
        raises(NoTransaction);
    void rollback_only()
        raises(NoTransaction);
    Status get_status();
    string get_transaction_name();
    void set_timeout(in unsigned long seconds);
    Control get_control();
    Control suspend();
    void resume(in Control which)
        raises(InvalidControl);
};

// This information is taken from CORBAServices: Common Object
// Services Specification, p. 10-18. Revised Edition:
// March 31, 1995. Updated: November 1997. Used with permission by
// OMG
```

Table 7-2 provides a description of the Current transaction methods.

Note: This information is taken from *CORBAServices: Common Object Services Specification*, pp. 10-18, 19, 20. Revised Edition: March 31, 1995. Updated: November 1997. Used with permission by OMG.

Table 7-2 Current Transaction Methods

Method	Description
<code>begin</code>	<p>Creates a new transaction. The transaction context of the client thread is modified so that the thread is associated with the new transaction. If the client thread is currently associated with a transaction, the <code>SubtransactionsUnavailable</code> exception is raised. If the client thread cannot be placed in transaction mode due to an error while starting the transaction, the <code>INVALID_TRANSACTION</code> standard system exception is raised. If the call was made in an improper context, the <code>BAD_INV_ORDER</code> standard system exception is raised.</p>
<code>commit</code>	<p>If there is no transaction associated with the client thread, the <code>NoTransaction</code> exception is raised.</p> <p>If the call was made in an improper context, the standard system exception <code>BAD_INV_ORDER</code> is raised.</p> <p>If the system decides to roll back the transaction, the standard exception <code>TRANSACTION_ROLLEDBACK</code> is raised and the thread's transaction context is set to null.</p> <p>A <code>HeuristicMixed</code> exception is raised to report that a heuristic decision was made and that some relevant updates have been committed and others have been rolled back. A <code>HeuristicHazard</code> exception is raised to report that a heuristic decision was made, and that the disposition of all relevant updates is not known; for those updates whose disposition is known, either all have been committed or all have been rolled back. The <code>HeuristicMixed</code> exception takes priority over the <code>HeuristicHazard</code> exception. If a heuristic exception is raised or the operation completes normally, the thread's transaction exception context is set to null.</p> <p>If the operation completes normally, the thread's transaction context is set to null.</p>

Table 7-2 Current Transaction Methods (Continued)

Method	Description
<code>rollback</code>	<p>If there is no transaction associated with the client thread, the <code>NoTransaction</code> exception is raised.</p> <p>If the call was made in an improper context, the standard system exception <code>BAD_INV_ORDER</code> is raised.</p> <p>If the operation completes normally, the thread's transaction context is set to null.</p>
<code>rollback_only</code>	<p>If there is no transaction associated with the client thread, the <code>NoTransaction</code> exception is raised. Otherwise, the transaction associated with the client thread is modified so that the only possible outcome is to roll back the transaction.</p>
<code>get_status</code>	<p>If there is no transaction associated with the client thread, the <code>StatusNoTransaction</code> value is returned. Otherwise, this method returns the status of the transaction associated with the client thread.</p>
<code>get_transaction_name</code>	<p>If there is no transaction associated with the client thread, an empty string is returned. Otherwise, this method returns a printable string describing the transaction (specifically, the XID as specified by The Open Group). The returned string is intended to support debugging.</p>
<code>set_timeout</code>	<p>This method modifies a state variable associated with the target object that affects the time-out period associated with transactions created by subsequent invocations of the <code>begin</code> method. If the parameter has a nonzero value <code>n</code>, transactions created by subsequent invocations of <code>begin</code> are subject to being rolled back if they do not complete before <code>n</code> seconds after their creation. If the parameter is zero, no time-out specified by the application is established.</p> <p>Note: The initial transaction timeout value is 300 seconds. If a transaction is started via <code>AUTOTRAN</code> instead of the <code>begin</code> method, then the timeout value is determined by the <code>TRANTIME</code> value in the WLE configuration file. For more information, refer to the <i>Administration Guide</i>.</p>

Table 7-2 Current Transaction Methods (Continued)

Method	Description
<code>get_control</code>	<p>If the client is not associated with a transaction, a null object reference is returned. Otherwise, a Control object is returned that represents the transaction context currently associated with the client thread. This object may be given to the <code>resume</code> method to reestablish this context.</p>
<code>suspend</code>	<p>If the client thread is not associated with a transaction, a null object reference is returned.</p> <p>If the associated transaction is in a state such that the only possible outcome of the transaction is to be rolled back, a <code>TRANSACTION_ROLLEDBACK</code> standard system exception is raised and the client thread becomes associated with no transaction.</p> <p>If the call was made in an improper context, the standard system exception <code>BAD_INV_ORDER</code> is raised. The caller's state with respect to the transaction is not changed.</p> <p>Otherwise, an object is returned that represents the transaction context currently associated with the client thread. The same client can subsequently give this object to the <code>resume</code> method to reestablish this context. In addition, the client thread becomes associated with no transaction.</p> <p>Note: As defined in The Common Object Request Broker: Architecture and Specification, Revision 2.2, February 1998, the <code>TRANSACTION_ROLLEDBACK</code> standard system exception indicates that the transaction associated with the request has already been rolled back or has been marked to roll back. Thus, the requested method either could not be performed or was not performed because further computation on behalf of the transaction would be fruitless.</p>

Table 7-2 Current Transaction Methods (Continued)

Method	Description
<code>resume</code>	<p>If the client thread is already associated with a transaction which is in a state such that the only possible outcome of the transaction is to be rolled back, the <code>TRANSACTION_ROLLEDBACK</code> standard system exception is raised and the client thread becomes associated with no transaction.</p> <p>If the call was made in an improper context, the standard system exception <code>BAD_INV_ORDER</code> is raised.</p> <p>If the system is unable to resume the global transaction because the caller is currently participating in work outside any global transaction with one or more resource managers, the <code>INVALID_TRANSACTION</code> standard system exception is raised.</p> <p>If the parameter is a null object reference, the client thread becomes associated with no transaction. If the parameter is valid in the current execution environment, the client thread becomes associated with that transaction (in place of any previous transaction). Otherwise, the <code>InvalidControl</code> exception is raised.</p> <p>Note: See <code>suspend</code> for a definition of the <code>TRANSACTION_ROLLEDBACK</code> standard system exception.</p>

Control Interface

The Control interface allows a program to explicitly manage or propagate a transaction context. An object that supports the Control interface is implicitly associated with one specific transaction.

The CosTransactions module defines the Control interface (shown in Listing 7-4).

Listing 7-4 Control Interface

```
interface Control {
    Terminator get_terminator()
        raises(Unavailable);
    Coordinator get_coordinator()
        raises(Unavailable);
};

// This information is taken from CORBAservices: Common Object
// Services Specification, p. 10-21. Revised Edition:
// March 31, 1995. Updated: November 1997. Used with permission by
// OMG.
```

Table 7-3 provides descriptions of the Control interface methods.

Note: This information is taken from *CORBAservices: Common Object Services Specification*, p. 10-22. Revised Edition: March 31, 1995. Updated: March 1997. Used with permission by OMG.

Table 7-3 Control Interface Methods

Method	Description
get_terminator	Returns a Terminator object, which supports methods to end the transaction. The object can be used to roll back or commit the transaction associated with the Control interface. The current WLE implementation of the CORBAservices Object Transaction Service always raises the Unavailable exception.

Table 7-3 Control Interface Methods (Continued)

Method	Description
<code>get_coordinator</code>	Returns a Coordinator object, which supports methods needed by resources to participate in the transaction. The object can be used to register resources for the transaction associated with the Control interface. The current WLE implementation of the CORBAServices Object Transaction Service always raises the Unavailable exception.

TransactionalObject Interface

The `CosTransactions::TransactionalObject` interface is used by an object to indicate that it is transactional. By supporting this interface, an object indicates that it wants the transaction context associated with the client thread to be propagated on requests to the object. However, this interface is no longer needed. For details on transaction policies that need to be set to infect objects with transactions, see the sections “ICF Syntax” on page 2-2 and “`CosTransactions::TransactionalObject` Interface Not Enforced” on page 3-68.

The `CosTransactions` module defines the `TransactionalObject` interface (shown in Listing 7-5). The `CosTransactions::TransactionalObject` interface defines no methods. It is simply a marker.

Listing 7-5 TransactionalObject Interface

```
interface TransactionalObject {  
};  
  
// This information is taken from CORBAServices: Common Object  
// Services Specification, p. 10-30. Revised Edition:  
// March 31, 1995. Updated: November 1997. Used with permission by  
// OMG.
```

Other CORBAServices Object Transaction Service Interfaces

All other CORBAServices Object Transaction Service interfaces are not supported. Note that the Current interface described earlier is supported only if it has been obtained from the Bootstrap object. The Control interface described earlier is supported only if it has been obtained using the `get_control()` and the `suspend()` methods on the Current object.

Transaction Service API Extensions

This section describes specific extensions to the CORBAServices Transaction Service API described earlier. The APIs in this section enable an application to open or close an Open Group Resource Manager.

The following APIs help facilitate participation of resource managers in a distributed transaction by allowing their two-phase commit protocol to be controlled via The Open Group XA interface.

The following definitions and interfaces are defined in the Tobj module.

Exception

The following exception is supported:

```
exception RMfailed {};
```

A request raises this exception to report that an attempt to open or close a resource manager failed.

TransactionCurrent Interface

This interface supports all the methods of the Current interface in the CosTransactions module as described in the previous section. Additionally, this interface supports APIs to open and close the resource manager.

The Tobj module defines the TransactionCurrent interface, as shown in Listing 7-6.

Listing 7-6 TransactionCurrent Interface

```
Interface TransactionCurrent: CosTransactions::Current {  
    void open_xa_rm()  
        raises(RMfailed);  
    void close_xa_rm()  
        raises(Rmfailed);  
}
```

Table 7-4 describes APIs that are specific to the resource manager. For more information about these APIs, see sections “TP::close_xa_rm()” on page 3-48 and “TP::open_xa_rm()” on page 3-60.

Table 7-4 Resource Manager APIs for the Current Interface

Method	Description
open_xa_rm	<p>This method opens The Open Group Resource Manager to which this process is linked. A <code>RMfailed</code> exception is raised if there is a failure while opening the Resource Manager.</p> <p>Any attempts to invoke this method by remote clients or the native clients raises a <code>NO_IMPLEMENT</code> standard system exception.</p>

Table 7-4 Resource Manager APIs for the Current Interface (Continued)

Method	Description
<code>close_xa_rm</code>	<p>This method closes The Open Group Resource Manager to which this process is linked. An <code>RMfailed</code> exception is raised if there is a failure while closing the Resource Manager. A <code>BAD_INV_ORDER</code> standard system exception is raised if the function was called in an improper context (for example, the caller is in transaction mode).</p> <p>Any attempts by the remote clients or the native clients to invoke this method raises a <code>NO_IMPLEMENT</code> standard system exception.</p>

8 Interface Repository Interface

This chapter describes the Interface Repository interfaces.

Note: Most of the information in this chapter is taken from Chapter 8 of the *Common Object Request Broker: Architecture and Specification*, Revision 2.2, February 1998. The OMG information has been modified as required to describe the WLE implementation of the Interface Repository interfaces. Used with permission by OMG.

The WLE Interface Repository contains the interface descriptions of the CORBA objects that are implemented within the WLE domain.

The WLE Interface Repository is based on the CORBA definition of an Interface Repository. It offers a proper subset of the interfaces defined by CORBA; that is, the APIs that are exposed to programmers are implemented as defined by the *Common Object Request Broker: Architecture and Specification* Revision 2.2. However, not all interfaces are supported. In general, the interfaces required to read from the Interface Repository are supported, but the interfaces required to write to the Interface Repository are not. Additionally, not all TypeCode interfaces are supported.

Administration of the Interface Repository is done using tools specific to the WLE software. These tools allow the system administrator to create an Interface Repository, populate it with definitions specified in Object Management Group Interface Definition Language (OMG IDL), and then delete interfaces. Additionally, an administrator may need to configure the system to include an Interface Repository server. For a description of the Interface Repository administration commands, see *Administration Guide*.

Several abstract interfaces are used as base interfaces for other objects in the Interface Repository. A common set of operations is used to locate objects within the Interface Repository. These operations are defined in the abstract interfaces `IObject`, `Container`, and `Contained` described in this chapter. All Interface Repository objects inherit from the `IObject` interface, which provides an operation for identifying the actual type of the object. Objects that are containers inherit navigation operations from the `Container` interface. Objects that are contained by other objects inherit navigation operations from the `Contained` interface. The `IDLType` interface is inherited by all Interface Repository objects that represent OMG IDL types, including interfaces, typedefs, and anonymous types. The `TypedefDef` interface is inherited by all named noninterface types.

The `IObject`, `Contained`, `Container`, `IDLType`, and `TypedefDef` interfaces are not instantiable.

All string data in the Interface Repository are encoded as defined by the ISO 8859-1 character set.

Note: The Write interface is not documented in this chapter because the WLE software supports only read access to the Interface Repository. Any attempt to use the Write interface to the Interface Repository will raise the exception `CORBA::NO_IMPLEMENT`.

Structure and Usage

The Interface Repository consists of two distinct components: the database and the server. The server performs operations on the database.

The Interface Repository database is created and populated using the `idl2ir` administrative command. For a description of this command, see the *Administration Guide*. From the programmer's point of view, there is no write access to the Interface Repository. None of the write operations defined by CORBA are supported, nor are set operations on non-readonly attributes.

Read access to the Interface Repository database is always through the Interface Repository server; that is, a client reads from the database by invoking methods that are performed by the server. The read operations as defined by the *CORBA Common Object Request Broker: Architecture and Specification*, Revision 2.2, are described in this chapter.

Programming Information

The interface to a server is defined in the OMG IDL file. How the OMG IDL file is accessed depends on the type of client being built. Three types of clients are considered: stub based, Dynamic Invocation Interface (DII), and ActiveX.

Client applications that use stub-style invocations need the OMG IDL file at build time. The programmer can use the OMG IDL file to generate stubs, and so forth. (For more information, see *Creating Client Applications*.) No other access to the Interface Repository is required.

Client applications that use the Dynamic Invocation Interface (DII) need to access the Interface Repository programmatically. The interface to the Interface Repository is defined in this chapter and is discussed in “Building Client Applications” on page 8-5. The exact steps taken to access the Interface Repository depend on whether the client is seeking information about a specific object, or browsing the Interface Repository to find an interface. To obtain information about a specific object, clients use the `CORBA::Object::_get_interface` method to obtain an `InterfaceDef` object. (Refer to “`CORBA::Object::_get_interface`” on page 1-53 for a description of this method.) Using the `InterfaceDef` object, the client can get complete information about the interface.

Before a DII client can browse the Interface Repository, it needs to obtain the object reference of the Interface Repository to start the search. DII clients use the Bootstrap object to obtain the object reference. (For a description of this method, see the section “`Tobj_Bootstrap::register_callback_port`” on page 4-14.) Once the client has the object reference, it can navigate the Interface Repository, starting at the root.

Note: To use the DII, the OMG IDL file must be stored in the Interface Repository.

Client applications that use ActiveX are not aware that they are using the Interface Repository. From the Interface Repository perspective, an ActiveX client is no different than a DII client. ActiveX clients include the Bootstrap object in the Visual Basic code. Like DII clients, ActiveX clients use the Bootstrap object to obtain the

Interface Repository object reference. (Refer to “Tobj_Bootstrap::register_callback_port” on page 4-14 for a description of this method.) Once the client has the object reference, it can navigate the Interface Repository, starting at the root.

Note: To use an ActiveX client, the OMG IDL file must be stored in the Interface Repository.

Performance Implications

All run-time access to the Interface Repository is via the Interface Repository server. Because there is considerable overhead in making requests of a remote server application, designers need to be aware of this. For example, consider the interaction required to use an object reference to obtain the necessary information to make a DII invocation on the object reference. The steps are as follows:

1. The client application invokes the `_get_interface` operation on the `CORBA::Object` to get the `InterfaceDef` object associated with the object in question. This causes a message to be sent to the ORB that created the object reference.
2. The ORB returns the `InterfaceDef` object to the client.
3. The client invokes one or more `_is_a` operations on the object to determine what type of interface is supported by the object.
4. After the client has identified the interface, it invokes the `describe_interface` operation on the `Interface` object to get a full description of the interface (for example, version number, operations, attributes, and parameters). This causes a message to be sent to the Interface Repository, and a reply is returned.
5. The client is now ready to construct a DII request.

Building Client Applications

Clients that use the Interface Repository need to link in Interface Repository stubs. How this happens is specific to the vendor. If the client application is using the WLE ORB, the WLE software provides the stubs in the form of a library. Therefore, programmers do not need to use the Interface Repository OMG IDL file to build the stubs. The Interface Repository definitions are contained within the `CORBA.h` file, but they are not included by default.

Note: To use the Interface Repository definitions, you must define the `ORB_INCLUDE_REPOSITORY` macro before including `CORBA.h` in your client application code (for example: `#Define ORB_INCLUDE_REPOSITORY`).

If the client application is using a third-party ORB (for example, Orbix) the programmer must use the mechanisms that are provided by that vendor. This might include generating stubs from the OMG IDL file using the IDL compiler supplied by the vendor, simply linking against the stubs provided by the vendor, or some other mechanism.

Some third-party ORBs provide a local Interface Repository capability. In this case, the local Interface Repository is provided by the vendor and is populated with the interface definitions that are needed by that client.

Getting Initial References to the InterfaceRepository Object

You use the Bootstrap object to get an initial reference to the InterfaceRepository object. For a description of the Bootstrap object method, see the command “`Tobj_Bootstrap::register_callback_port`” on page 4-14.

Interface Repository Interfaces

Client applications use the interfaces defined by CORBA to access the Interface Repository. This section contains descriptions of each interface that is implemented in the WLE software.

Supporting Type Definitions

Several types are used throughout the Interface Repository interface definitions.

```
module CORBA {
    typedef string          Identifier;
    typedef string          ScopedName;
    typedef string          RepositoryId;

    enum DefinitionKind {
        dk_none, dk_all,
        dk_Attribute, dk_Constant, dk_Exception, dk_Interface,
        dk_Module, dk_Operation, dk_Typedef,
        dk_Alias, dk_Struct, dk_Union, dk_Enum,
        dk_Primitive, dk_String, dk_Sequence, dk_Array,
        dk_Repository,
    };
};
```

Identifiers are the simple names that identify modules, interfaces, constants, typedefs, exceptions, attributes, and operations. They correspond exactly to OMG IDL identifiers. An **Identifier** is not necessarily unique within an entire Interface Repository; it is unique only within a particular Repository, ModuleDef, InterfaceDef, or OperationDef.

A **ScopedName** is a name made up of one or more identifiers separated by the characters “::”. They correspond to OMG IDL scoped names. An absolute **ScopedName** is one that begins with “::” and unambiguously identifies a definition in a Repository. An absolute **ScopedName** in a Repository corresponds to a global name in an OMG IDL file. A relative **ScopedName** does not begin with “::” and must be resolved relative to some context.

A `RepositoryId` is an identifier used to uniquely and globally identify a module, interface, constant, typedef, exception, attribute, or operation. Because `RepositoryIds` are defined as strings, they can be manipulated (for example, copied and compared) using a language binding's string manipulation routines.

A `DefinitionKind` identifies the type of an Interface Repository object.

IObject Interface

The `IObject` interface (shown below) represents the most generic interface from which all other Interface Repository interfaces are derived, even the `Repository` itself.

```
module CORBA {
    interface IObject {
        readonly attribute DefinitionKind def_kind;
    };
};
```

The `def_kind` attribute identifies the type of the definition.

Contained Interface

The `Contained` interface (shown below) is inherited by all Interface Repository interfaces that are contained by other Interface Repository objects. All objects within the Interface Repository, except the root object (`Repository`) and definitions of anonymous (`ArrayDef`, `StringDef`, and `SequenceDef`), and primitive types are contained by other objects.

```
module CORBA {
    typedef string VersionSpec;

    interface Contained : IObject {
        readonly attribute RepositoryId    id;
        readonly attribute Identifier      name;
        readonly attribute VersionSpec     version;
        readonly attribute Container       defined_in;
        readonly attribute ScopedName     absolute_name;
        readonly attribute Repository     containing_repository;
        struct Description {

```

```

        DefinitionKind          kind;
        any                     value;
    };

    Description describe ();
};
};

```

An object that is contained by another object has an `id` attribute that identifies it globally, and a `name` attribute that identifies it uniquely within the enclosing Container object. It also has a `version` attribute that distinguishes it from other versioned objects with the same name. The WLE Interface Repository does not support simultaneous containment or multiple versions of the same named object.

Contained objects also have a `defined_in` attribute that identifies the Container within which they are defined. Objects can be contained either because they are defined within the containing object (for example, an interface is defined within a module) or because they are inherited by the containing object (for example, an operation may be contained by an interface because the interface inherits the operation from another interface). If an object is contained through inheritance, the `defined_in` attribute identifies the `InterfaceDef` from which the object is inherited.

The `absolute_name` attribute is an absolute `ScopedName` that identifies a Contained object uniquely within its enclosing Repository. If this object's `defined_in` attribute references a Repository, the `absolute_name` is formed by concatenating the string `":"` and this object's `name` attribute. Otherwise, the `absolute_name` is formed by concatenating the `absolute_name` attribute of the object referenced by this object's `defined_in` attribute, the string `":"`, and this object's `name` attribute.

The `containing_repository` attribute identifies the Repository that is eventually reached by recursively following the object's `defined_in` attribute.

The `describe` operation returns a structure containing information about the interface. The description structure associated with each interface is provided below with the interface's definition. The kind of definition described by the structure returned is provided with the returned structure. For example, if the `describe` operation is invoked on an attribute object, the `kind` field contains `dk_Attribute` and the `value` field contains an `any`, which contains the `AttributeDescription` structure.

Container Interface

The Container interface is used to form a containment hierarchy in the Interface Repository. A Container can contain any number of objects derived from the Contained interface. All Containers, except for Repository, are also derived from Contained.

```
module CORBA {
    typedef sequence <Contained> ContainedSeq;

    interface Container : IRObject {
        Contained lookup (in ScopedName search_name);

        ContainedSeq contents (
            in DefinitionKind          limit_type,
            in boolean                  exclude_inherited
        );

        ContainedSeq lookup_name (
            in Identifier               search_name,
            in long                     levels_to_search,
            in DefinitionKind           limit_type,
            in boolean                  exclude_inherited
        );

        struct Description {
            Contained                contained_object;
            DefinitionKind            kind;
            any                       value;
        };

        typedef sequence<Description> DescriptionSeq;

        DescriptionSeq describe_contents (
            in DefinitionKind          limit_type,
            in boolean                  exclude_inherited,
            in long                     max_returned_objs
        );
    };
};
```

The lookup operation locates a definition relative to this container, given a scoped name using the OMG IDL rules for name scoping. An absolute scoped name (beginning with ":" : ") locates the definition relative to the enclosing Repository. If no object is found, a nil object reference is returned.

The `contents` operation returns the list of objects directly contained by or inherited into the object. The operation is used to navigate through the hierarchy of objects. Starting with the Repository object, a client uses this operation to list all of the objects contained by the Repository, all of the objects contained by the modules within the Repository, all of the interfaces within a specific module, and so on.

`limit_type`

If `limit_type` is set to `dk_all`, objects of all types are returned. For example, if this is an `InterfaceDef`, the attribute, operation, and exception objects are all returned. If `limit_type` is set to a specific interface, only objects of that type are returned. For example, only attribute objects are returned if `limit_type` is set to `dk_Attribute`.

`exclude_inherited`

If set to `TRUE`, inherited objects (if there are any) are not returned. If set to `FALSE`, all contained objects (whether contained due to inheritance or because they were defined within the object) are returned.

The `lookup_name` operation is used to locate an object by name within a particular object or within the objects contained by that object. The `describe_contents` operation combines the `contents` operation and the `describe` operation. For each object returned by the `contents` operation, the description of the object is returned (that is, the object's `describe` operation is invoked and the results are returned).

`search_name`

Specifies which name is to be searched for.

`levels_to_search`

Controls whether the lookup is constrained to the object the operation is invoked on, or whether the lookup should search through objects contained by the object as well. Setting `levels_to_search` to `-1` searches the current object and all contained objects. Setting `levels_to_search` to `1` searches only the current object.

`max_returned_objs`

Limits the number of objects that can be returned in an invocation of the call to the number provided. Setting the parameter to `-1` indicates return all contained objects.

IDLType Interface

The IDLType interface (shown below) is an abstract interface inherited by all Interface Repository objects that represent OMG IDL types. It provides access to the TypeCode describing the type, and is used in defining other interfaces wherever definitions of IDL types must be referenced.

```
module CORBA {  
    interface IDLType : IObject {  
        readonly attribute TypeCode          type;  
    };  
};
```

The `type` attribute describes the type defined by an object derived from IDLType.

Repository Interface

Repository (shown below) is an interface that provides global access to the Interface Repository. The Repository object can contain constants, typedefs, exceptions, interfaces, and modules. As it inherits from Container, it can be used to look up any definition (whether globally defined or defined within a module or an interface) either by name or by id.

```
module CORBA {  
    interface Repository : Container {  
        Contained lookup_id (in RepositoryId search_id);  
        PrimitiveDef get_primitive (in PrimitiveKind kind);  
    };  
};
```

The `lookup_id` operation is used to look up an object in a Repository, given its `RepositoryId`. If the Repository does not contain a definition for `search_id`, a nil object reference is returned.

The `get_primitive` operation returns a reference to a `PrimitiveDef` with the specified `kind` attribute. All `PrimitiveDefs` are immutable and are owned by the Repository.

ModuleDef Interface

A ModuleDef (shown below) can contain constants, typedefs, exceptions, interfaces, and other module objects.

```
module CORBA {  
    interface ModuleDef : Container, Contained {  
    };  
  
    struct ModuleDescription {  
        Identifier      name;  
        RepositoryId    id;  
        RepositoryId    defined_in;  
        VersionSpec     version;  
    };  
};
```

The inherited `describe` operation for a ModuleDef object returns a ModuleDescription.

ConstantDef Interface

A ConstantDef object (shown below) defines a named constant.

```
module CORBA {  
    interface ConstantDef : Contained {  
        readonly attribute TypeCode      type;  
        readonly attribute IDLType       type_def;  
        readonly attribute any           value;  
    };  
  
    struct ConstantDescription {  
        Identifier      name;  
        RepositoryId    id;  
        RepositoryId    defined_in;  
        VersionSpec     version;  
        TypeCode        type;  
        any              value;  
    };  
};
```

<code>type</code>	Specifies the <code>TypeCode</code> describing the type of the constant. The type of a constant must be one of the simple types (long, short, float, char, string, octet, and so on).
<code>type_def</code>	Identifies the definition of the type of the constant.
<code>value</code>	Contains the value of the constant, not the computation of the value (for example, the fact that it was defined as “1+2”).

The `describe` operation for a `ConstantDef` object returns a `ConstantDescription`.

TypedefDef Interface

A `TypedefDef` (shown below) is an abstract interface used as a base interface for all named nonobject types (structures, unions, enumerations, and aliases). The `TypedefDef` interface is not inherited by the definition objects for primitive or anonymous types.

```
module CORBA {  
    interface TypedefDef : Contained, IDLType {  
    };  
  
    struct TypeDescription {  
        Identifier           name;  
        RepositoryId        id;  
        RepositoryId        defined_in;  
        VersionSpec         version;  
        TypeCode            type;  
    };  
};
```

The inherited `describe` operation for interfaces derived from `TypedefDef` returns a `TypeDescription`.

StructDef

A StructDef (shown below) represents an OMG IDL structure definition. It contains the members of the struct.

```
module CORBA {
    struct StructMember {
        Identifier    name;
        TypeCode      type;
        IDLType       type_def;
    };
    typedef sequence <StructMember> StructMemberSeq;

    interface StructDef : TypedefDef, Container{
        readonly attribute StructMemberSeq    members;
    };
};
```

The `members` attribute contains a description of each structure member.

The inherited `type` attribute is a `tk_struct` `TypeCode` describing the structure.

UnionDef

A UnionDef (shown below) represents an OMG IDL union definition. It contains the members of the union.

```
module CORBA {
    struct UnionMember {
        Identifier    name;
        any           label;
        TypeCode      type;
        IDLType       type_def;
    };
    typedef sequence <UnionMember> UnionMemberSeq;

    interface UnionDef : TypedefDef, Container {
        readonly attribute TypeCode    discriminator_type;
        readonly attribute IDLType     discriminator_type_def;
        readonly attribute UnionMemberSeq    members;
    };
};
```

`discriminator_type` and `discriminator_type_def`

Describe and identify the union's discriminator type.

`members`

Contains a description of each union member. The label of each `UnionMemberDescription` is a distinct value of the `discriminator_type`. Adjacent members can have the same name. Members with the same name must also have the same type. A label with type `octet` and value 0 (zero) indicates the default union member.

The inherited `type` attribute is a `tk_union` `TypeCode` describing the union.

EnumDef

An `EnumDef` (shown below) represents an OMG IDL enumeration definition.

```
module CORBA {
    typedef sequence <Identifier> EnumMemberSeq;

    interface EnumDef : TypedefDef {
        readonly attribute EnumMemberSeq          members;
    };
};
```

`members`

Contains a distinct name for each possible value of the enumeration.

The inherited `type` attribute is a `tk_enum` `TypeCode` describing the enumeration.

AliasDef

An `AliasDef` (shown below) represents an OMG IDL typedef that aliases another definition.

```
module CORBA {
    interface AliasDef : TypedefDef {
        readonly attribute IDLType original_type_def;
    };
};
```

`original_type_def`

Identifies the type being aliased.

The inherited `type` attribute is a `tk_alias` `TypeCode` describing the alias.

PrimitiveDef

A `PrimitiveDef` (shown below) represents one of the OMG IDL primitive types. Because primitive types are unnamed, this interface is not derived from `TypedefDef` or `Contained`.

```
module CORBA {
    enum PrimitiveKind {
        pk_null, pk_void, pk_short, pk_long, pk_ushort, pk_ulong,
        pk_float, pk_double, pk_boolean, pk_char, pk_octet,
        pk_any, pk_TypeCode, pk_Principal, pk_string, pk_objref,
        pk_longlong, pk_ulonglong, pk_longdouble, pk_wchar, pk_wstring
    };

    interface PrimitiveDef: IDLType {
        readonly attribute PrimitiveKind      kind;
    };
};
```

`kind`

Indicates which primitive type the `PrimitiveDef` represents. There are no `PrimitiveDefs` with kind `pk_null`. A `PrimitiveDef` with kind `pk_string` represents an unbounded string. A `PrimitiveDef` with kind `pk_objref` represents the OMG IDL type Object.

The inherited `type` attribute describes the primitive type.

All `PrimitiveDefs` are owned by the Repository. References to them are obtained using `Repository::get_primitive`.

ExceptionDef

An `ExceptionDef` (shown below) represents an exception definition. It can contain structs, unions, and enums.


```

module CORBA {
    interface ExceptionDef : Contained, Container {
        readonly attribute TypeCode          type;
        readonly attribute StructMemberSeq    members;
    };

    struct ExceptionDescription {
        Identifier          name;
        RepositoryId        id;
        RepositoryId        defined_in;
        VersionSpec         version;
        TypeCode            type;
    };
};

```

type
tk_except TypeCode that describes the exception.

members
Describes any exception members.

The describe operation for a ExceptionDef object returns an ExceptionDescription.

AttributeDef

An AttributeDef (shown below) represents the information that defines an attribute of an interface.

```

module CORBA {
    enum AttributeMode {ATTR_NORMAL, ATTR_READONLY};

    interface AttributeDef : Contained {
        readonly attribute TypeCode          type;
        attribute IDLType                    type_def;
        attribute AttributeMode              mode;
    };

    struct AttributeDescription {
        Identifier          name;
        RepositoryId        id;
        RepositoryId        defined_in;
        VersionSpec         version;
        TypeCode            type;
    };
};

```

```

        AttributeMode      mode;
    };

    type
        Provides the TypeCode describing the type of this attribute.

    type_def
        Identifies the object that defines the type of this attribute.

    mode
        Specifies read only or read/write access for this attribute.

```

OperationDef

An OperationDef (shown below) represents the information needed to define an operation of an interface.

```

module CORBA {
    enum OperationMode {OP_NORMAL, OP_ONEWAY};

    enum ParameterMode {PARAM_IN, PARAM_OUT, PARAM_INOUT};
    struct ParameterDescription {
        Identifier      name;
        TypeCode        type;
        IDLType         type_def;
        ParameterMode   mode;
    };
    typedef sequence <ParameterDescription> ParDescriptionSeq;

    typedef Identifier ContextIdentifier;
    typedef sequence <ContextIdentifier> ContextIdSeq;

    typedef sequence <ExceptionDef> ExceptionDefSeq;
    typedef sequence <ExceptionDescription> ExcDescriptionSeq;

    interface OperationDef : Contained {
        readonly attribute TypeCode      result;
        readonly attribute IDLType        result_def;
        readonly attribute ParDescriptionSeq params;
        readonly attribute OperationMode  mode;
        readonly attribute ContextIdSeq    contexts;
        readonly attribute ExceptionDefSeq exceptions;
    };
}

```

```

};

struct OperationDescription {
    Identifier          name;
    RepositoryId        id;
    RepositoryId        defined_in;
    VersionSpec         version;
    TypeCode            result;
    OperationMode       mode;
    ContextIdSeq        contexts;
    ParDescriptionSeq   parameters;
    ExcDescriptionSeq   exceptions;
};
};

```

`result`

A `TypeCode` that describes the type of the value returned by the operation.

`result_def`

Identifies the definition of the returned type.

`params`

Describes the parameters of the operation. It is a sequence of `ParameterDescription` structures. The order of the `ParameterDescriptions` in the sequence is significant. The `name` member of each structure provides the parameter name. The `type` member is a `TypeCode` describing the type of the parameter. The `type_def` member identifies the definition of the type of the parameter. The `mode` member indicates whether the parameter is an in, out, or inout parameter.

`mode`

The operation's `mode` is either oneway (that is, no output is returned) or normal.

`contexts`

Specifies the list of context identifiers that apply to the operation.

`exceptions`

Specifies the list of exception types that can be raised by the operation.

The inherited `describe` operation for an `OperationDef` object returns an `OperationDescription`.

The inherited `describe_contents` operation provides a complete description of this operation, including a description of each parameter defined for this operation.

InterfaceDef

An `InterfaceDef` object (shown below) represents an interface definition. It can contain constants, typedefs, exceptions, operations, and attributes.

```
module CORBA {
    interface InterfaceDef;
        typedef sequence <InterfaceDef> InterfaceDefSeq;
        typedef sequence <RepositoryId> RepositoryIdSeq;
        typedef sequence <OperationDescription> OpDescriptionSeq;
        typedef sequence <AttributeDescription> AttrDescriptionSeq;

        interface InterfaceDef : Container, Contained, IDLType {

            readonly attribute InterfaceDefSeq    base_interfaces;

            boolean is_a (in RepositoryId interface_id);

            struct FullInterfaceDescription {
                Identifier          name;
                RepositoryId        id;
                RepositoryId        defined_in;
                VersionSpec          version;
                OpDescriptionSeq     operations;
                AttrDescriptionSeq   attributes;
                RepositoryIdSeq      base_interfaces;
                TypeCode             type;
            };

            FullInterfaceDescription describe_interface();

        };

        struct InterfaceDescription {
            Identifier          name;
            RepositoryId        id;
            RepositoryId        defined_in;
            VersionSpec          version;
            RepositoryIdSeq      base_interfaces;
        };
};
```

`base_interfaces`

Lists all the interfaces from which this interface inherits. The `is_a` operation returns `TRUE` if the interface on which it is invoked either is identical to or inherits, directly or indirectly, from the interface identified by its `interface_id` parameter. Otherwise, it returns `FALSE`.

The `describe_interface` operation returns a `FullInterfaceDescription` describing the interface, including its operations and attributes.

The inherited `describe` operation for an `InterfaceDef` returns an `InterfaceDescription`.

The inherited `contents` operation returns the list of constants, typedefs, and exceptions defined in this `InterfaceDef` and the list of attributes and operations either defined or inherited in this `InterfaceDef`. If the `exclude_inherited` parameter is set to `TRUE`, only attributes and operations defined within this interface are returned. If the `exclude_inherited` parameter is set to `FALSE`, all attributes and operations are returned.

9 Joint Client/Servers

This chapter describes programming requirements for joint client/servers and the BEA Wrapper Callbacks API.

For either a WLE client or joint client/server (that is, a client that can receive and process object invocations), the programmer writes the client `main()`. The `main()` uses WLE environmental objects to establish connections, set up security, and start transactions.

WLE clients invoke operations on objects. In the case of DII, client code creates the DII Request object and then invokes one of two operations on the DII Request. In the case of static invocation, client code performs the invocation by performing what looks like an ordinary C++ invocation (which ends up calling code in the generated client stub). Additionally, the client programmer uses ORB interfaces defined by OMG, and WLE environmental objects that are supplied with the WLE software, to perform functions unique to WLE.

For WLE joint client/servers, the client code must be structured so that it can act as a server for callback WLE objects. Such clients do not use the TP Framework and are not subject to WLE system administration. Besides the programming implications, this means that joint client/servers do not have the same scalability and reliability as WLE servers, nor do they have the state management and transaction behavior available in the TP Framework. If a user wants to have those characteristics, the application must be structured in such a way that the object implementations are in a WLE server, rather than in a client.

The following sections describe the mechanisms you use to add callback support to a WLE client. In some cases, the mechanisms are contrasted with the WLE server mechanisms that use the TP Framework.

Main Program and Server Initialization

In a WLE server, you use the `buildobjserver` command to create the main program for the server. That main program takes care of all WLE- and CORBA-related initialization of the server functions. The server main program allows the user to take part in server initialization and shutdown by making invocations on a user-written C++ object, the `Server` class.

In contrast, for a WLE joint client/server (as for a WLE client), you create the main program and are responsible for all initialization. You do not need to provide a `Server` object because you have complete control over the main program and you can provide initialization and shutdown code in any way that is convenient.

The specific initialization needed for a joint client/server is discussed below.

Servants

Servants (method code) for WLE joint client/servers are very similar to servants for WLE servers. All business logic is written the same way. The differences result from not using the TP Framework, which includes the `Server`, `TP`, and `Tobj_ServantBase` interfaces. Therefore, the main difference is that you use CORBA functions directly instead of indirectly through the TP Framework.

The `Server` interface is used in WLE servers to allow the TP Framework to ask the user to create a servant for an object when the ORB receives a request for that object. In WLE joint client/servers, the user program is responsible for creating a servant before any requests arrive; thus, the `Server` interface is not needed. Typically, the program creates a servant and then activates the object (using the servant and an `ObjectId`; the `ObjectId` is possibly system generated) before handing a reference to the object. Such an object might be used to handle callbacks. Thus, the servant already exists and the object is activated before a request for the object arrives.

Instead of invoking the TP interface to perform certain operations, client servants directly invoke the ORB and POA (which is what the TP interface does internally). Alternately, since much of the interaction with the ORB and POA is the same for all applications, for ease of use, the WLE client library provides a convenience wrapper

object that does the same things, using a single operation. For a discussion of how to use the convenience wrapper object, see “Callback Object Models Supported” on page 9-4 and “Preparing Callback Objects Using BEAWrapper Callbacks” on page 9-7.

Servant Inheritance from Skeletons

In a WLE client that supports callbacks, as well as in a WLE server, you write a C++ implementation class that inherits from the same skeleton class name generated by the IDL compiler (the `idl` command). For example, given the IDL:

```
interface Hospital{ ... };
```

The skeleton generated by the `idl` command contains a “skeleton” class, `POA_Hospital`, that the user-written class inherits from, as in:

```
class Hospital_i : public POA_Hospital { ... };
```

In a WLE server, the skeleton class inherits from the TP Framework class `Tobj_ServantBase`, which in turn inherits from the predefined `PortableServer::ServantBase`.

The inheritance tree for a callback object implementation in a joint client/server is different than that in a WLE server. The skeleton class does not inherit from the TP Framework class `Tobj_ServantBase`, but instead inherits directly from `PortableServer::ServantBase`. This behavior is achieved by specifying the `-P` option in the `idl` command.

Not having the `Tobj_ServantBase` class in the inheritance tree for a servant means that the servant does not have `activate_object` and `deactivate_object` methods. In a WLE server, these methods are called by the TP Framework to dynamically initialize and save a servant’s state before invoking a method on the servant. For a WLE client that supports callbacks, you must write code that explicitly creates a servant and initializes a servant’s state.

Callback Object Models Supported

WLE software supports four kinds of callback objects and provides wrappers for the three that are most common. These objects correspond to three combinations of POA policies. The POA policies control both the types of objects and the types of object references that are possible.

The POA policies that are applicable are:

- ◆ The `LifeSpanPolicy`, which controls how long an object reference is valid
- ◆ The `IdAssignmentPolicy`, which controls who assigns the `ObjectId`—the user or the system

These objects are explained primarily in terms of their behavioral characteristics rather than in details about how the ORB and the POA handle them. Those details are discussed in the next sections, using either direct ORB and POA calls (which requires a little extra knowledge of CORBA servers) or using the `BEAWrapperCallbacks` interface, which hides the ORB and POA calls (for users who do not care about the details).

- ◆ **Transient/SystemId**—Object references are valid only for the life of the client process. The `ObjectId` is not assigned by the client application, but is a unique value assigned by the system. This type of object is useful for invocations that a client wants to receive only until the client terminates. (The corresponding POA `LifeSpanPolicy` value is `TRANSIENT` and the `IdAssignmentPolicy` is `SYSTEM_ID`.)
- ◆ **Persistent/SystemId**—Object references are valid across multiple activations. The `ObjectId` is not assigned by the client application, but is a unique value assigned by the system. This type of object and object reference is useful when the client goes up and down over a period of time. When the client is up, it can receive callback objects on that particular object reference.

Typically, the client will create the object reference once, save it in its own permanent storage area, and reactivate the servant for that object every time it comes up. (The corresponding POA policy values are `PERSISTENT` and `SYSTEM_ID`.)

- ◆ **Persistent/UserId**—This is the same as `Persistent/SystemId` with the exception that the `ObjectId` has to be assigned by the client application. Such an

`objectId` might be, for example, a database key meaningful only to the client. (The corresponding POA policy values are `PERSISTENT` and `USER_ID`.)

Note: The Transient/UserId policy combination is not considered particularly important. It is possible for users to provide for themselves by using the POA in a manner analogous to either of the persistent cases, but the WLE wrappers do not provide special help to do so.

Note: For WLE native joint client/servers, neither of the Persistent policies is supported, only the Transient policy.

Preparing Callback Objects Using CORBA

To set up WLE callback objects using CORBA, the client must do the following:

1. Establish a connection with a POA with the appropriate policies for the callback object model. (This can be the root POA, available by default, or it may require creating a new POA.)
2. Create a servant (that is, an instance of the C++ implementation class for the interface).
3. Inform the POA that the servant is ready to accept requests on the callback WLE object. Technically, this means the client *activates* the object in the POA (that is, puts the servant and the `objectId` into the POA's Active Object Map).
4. Tell the POA to start accepting requests from the network (that is, activate the POA itself).
5. Create an object reference for the callback WLE object.
6. Give out the object reference. This usually happens by making an invocation on another object with the callback object reference as a parameter (that is, the parameter is a callback object). That other object will then invoke the callback object (perform a callback invocation) at some later time.

Assuming that the client already has obtained a reference to the ORB, performing this task takes four interactions with the ORB and the POA. It might look like the following for the Transient/SystemId model. In this model, only the Root POA is needed.

```
// Create a servant for the callback Object
Catcher_i* my_catcher_i = new Catcher_i();

// Get root POA reference and activate the POA
1  CORBA::Object_var oref =
    orb->resolve_initial_references("RootPOA");
2  PortableServer::POA_var root_poa =
    PortableServer::POA::_narrow(oref);
3  root_poa -> the_POAManager() -> activate();
4  PortableServer::objectId_var temp_Oid =
    root_poa ->activate_object ( my_catcher_i );
5  oref = root_poa->create_reference_with_id(
    temp_Oid, _tc_Catcher->id() );
6  Catcher_var my_catcher_ref = Catcher::_narrow( oref );
```

To use the **Persistent/UserId** model, there are some additional steps required when creating a POA. Further, the `objectId` is specified by the client, and this requires more steps. It might look like the following.

```
Catcher_i* my_catcher_i = new Catcher_i();
const char* oid_str = "783";
1  PortableServer::objectId_var oid =
    PortableServer::string_to_objectId(oid_str);

// Find root POA
2  CORBA::Object_var oref =
    orb->resolve_initial_references("RootPOA");
3  PortableServer::POA_var root_poa =
    PortableServer::POA::_narrow(oref);

// Create and activate a Persistent/UserId POA
4  CORBA::PolicyList policies(2);
5  policies.length(2);
6  policies[0] = root_poa->create_lifespan_policy(
    PortableServer::PERSISTENT);
7  policies[1] = root_poa->create_id_assignment_policy(
    PortableServer::USER_ID );
8  PortableServer::POA_var my_poa_ref =
    root_poa->create_POA(
    "my_poa_ref", root_poa->the_POAManager(), policies);
9  root_poa->the_POAManager()->activate();

// Create object reference for callback Object
10 oref = my_poa_ref -> create_reference_with_id(
    oid, _tc_Catcher->id());
11 Catcher_var my_catcher_ref = Catcher::_narrow( oref );

// activate object
12 my_poa_ref -> activate_object_with_id( oid, my_catcher_i );
```

```
// Make the call passing the callback ref
foo -> register_callback ( my_catcher_ref );
```

All the interfaces and operations described here are standard CORBA interfaces and operations.

Preparing Callback Objects Using BEAWrapper Callbacks

Since the code required for callback objects is nearly identical for every client that supports callbacks, you may find it convenient to use the BEAWrappers provided in the library provided for joint client/servers.

The BEAwrappers are described in IDL, as follows.

Note: These same wrappers are designed to be used for the WebLogic Enterprise V4.2 (Java) software, where a POA is not yet available, although aspects related to POAs do exist (notably, `PortableServer.Servant`). For a discussion of these for the Java software, see *Java Programming Reference*.

```
// File: BEAWrapper
#ifndef _BEA_WRAPPER_IDL_
#define _BEA_WRAPPER_IDL_
#include <orb.idl>
#include <PortableServer.idl>

#pragma prefix "beasys.com"

module BEAWrapper {
    interface Callbacks
    {
        exception ServantAlreadyActive{ };
        exception ObjectAlreadyActive { };
        exception NotInRequest{ };

        // set up transient callback Object
        // -- prepare POA, activate object, return objref
        Object start_transient(
            in PortableServer::Servant      Servant,
            in CORBA::RepositoryId         rep_id)
            raises (ServantAlreadyActive);
    };
};
```

```

        // set up persistent/systemid callback Object
        Object start_persistent_systemid(
            in PortableServer::Servant      servant,
            in CORBA::Repository            rep_id,
            out string                       stroid)
            raises (ServantAlreadyActive);

        // reinstate set up for persistent/systemid
        // callback object
        Object restart_persistent_systemid(
            in PortableServer::Servant      servant,
            in CORBA::RepositoryId          rep_id,
            in string                       stroid)
            raises (ServantAlreadyActive, ObjectAlreadyActive);

        // set up persistent/userid callback Object
        Object start_persistent_userid(
            in PortableServer::Servant      servant,
            in CORBA::RepositoryId          rep_id,
            in string                       stroid)
            raises (ServantAlreadyActive, ObjectAlreadyActive);

        // stop servicing a particular callback Object
        // with the given servant
        void stop_object( in PortableServer::Servant servant);

        //shutdown Stop all callback Object processing
        void stop_all_objects();

        // get oid string for the current request
        string get_string_oid() raises (NotInRequest);
    };
}
#endif /* _BEA_WRAPPER_IDL_ */

```

The BEAwrappers are described in C++ as follows:

C++ Declarations (in beawrapper.h)

```

#ifndef _BEAWRAPPER_H_
#define _BEAWRAPPER_H_

#include <PortableServer.h>
class BEAWrapper{
class Callbacks{
public:
    Callbacks (CORBA::ORB_ptr init_orb);

```

```
CORBA::Object_ptr start_transient (
    PortableServer::Servant servant,
    const char *      rep_id);

CORBA::Object_ptr start_persistent_systemid (
    PortableServer::Servant servant,
    const char *      rep_id,
    char * & stroid);

CORBA::Object_ptr restart_persistent_systemid (
    PortableServer::Servant servant,
    const char *      rep_id,
    const char *      stroid);

CORBA::Object_ptr start_persistent_userid (
    PortableServer::Servant servant,
    const char *      rep_id,
    const char *      stroid);

void stop_object(PortableServer::Servant servant);
char* get_string_oid ();
void stop_all_objects();
~Callbacks();

private:

    static CORBA::ORB_var orb_ptr;

    static PortableServer::POA_var root_poa;
    static PortableServer::POA_var trasys_poa;
    static PortableServer::POA_var persys_poa;
    static PortableServer::POA_var peruser_poa;
};
#endif // _BEAWRAPPER_H_
```

The description of each operation in the `BEAWrapper::Callbacks` interface follows, in the order declared above.

Callbacks

Synopsis	Returns a reference to the Callbacks interface.
C++ Binding	<code>BEAWrapper::Callbacks(CORBA::ORB_ptr init_orb);</code>
Java Binding	<code>public Callbacks(org.omg.CORBA.Object init_orb);</code>
Argument	<code>init_orb</code> The ORB to be used for all further operations.
Return Value	A reference to the Callbacks object.
Description	The constructor returns a reference to the Callbacks interface. Only one such object should be created for the process, even if multiple threads are used. Using more than one such object will result in undefined behavior.
Exception	<code>CORBA::IMP_LIMIT</code> The <code>BEAWrapper::Callbacks</code> class has already be instantiated with an ORB pointer. Only one instance of this class can be used in a process. Users who need additional flexibility should use the POA directly.

start_transient

Synopsis	Activates an object, sets the ORB and the POA to the proper state, and returns an object reference to the activated object.		
IDL	Object	<pre>start_transient(in PortableServer::Servant servant, in CORBA::RepositoryId rep_id) raises (ServantAlreadyActive);</pre>	
C++ Binding	CORBA::Object_ptr	<pre>start_transient(PortableServer::Servant servant, const char* rep_id);</pre>	
Java Binding	org.omg.CORBA.Object	<pre>start_transient(org.omg.PortableServer.Servant servant, java.lang.String rep_id);</pre>	
Arguments	servant	An instance of the C++ implementation class for the interface.	
	rep_id	The repository id of the interface.	
Return Value	CORBA::Object_ptr	A reference to the object that was created with the ObjectId generated by the system and the rep_id provided by the user. The object reference will need to be converted to a specific object type by invoking the _narrow() operation defined for the specific object. The caller is responsible for releasing the object when the conversion is done.	
Description	This operation performs the following actions:		
	<ul style="list-style-type: none">◆ Activates an object using the Servant supplied to service objects of the type rep_id, using an ObjectId generated by the system.◆ Sets the ORB and the POA into the state in which they will accept requests on that object.◆ Returns an object reference to the activated object. The returned object reference will be valid only until the termination of the client or until the callback servant is halted by the user via the stop_object operation; after that, invocations on that object reference are no longer valid and can never be made valid.		

Exceptions `ServantAlreadyActive`

The servant is already being used for a callback. A servant can be used only for a callback with a single `ObjectId`. To receive callbacks on objects containing different `ObjectIds`, you must create different servants and activate them separately. The same servant can be re-used only if a `stop_object` operation tells the system to stop using the servant for its original `ObjectId`.

`CORBA::BAD_PARAM`

The repository ID was a null string or the servant was a null pointer.

start_persistent_systemid

Synopsis Activates an object, sets the ORB and the POA to the proper state, sets the output parameter `stroid`, and returns an object reference to the activated object.

IDL

```
Object start_persistent_systemid(
                                in PortableServer::Servant    servant,
                                in CORBA::RepositoryId        rep_id,
                                out string                     stroid)
    raises ( ServantAlreadyActive );
```

C++ Binding

```
CORBA::Object_ptr start_persistent_systemid(
                                PortableServer::Servant    servant,
                                const char*                 rep_id,
                                char*&                      stroid);
```

Java Binding

```
org.omg.CORBA.Object start_persistent_systemid(
                                org.omg.PortableServer.Servant    servant,
                                java.lang.String                  rep_id,
                                java.lang.String                  stroid);
```

Arguments

`servant`
An instance of the C++ implementation class for the interface.

`rep_id`
The repository id of the interface.

`stroid`
This argument is set by the system and is opaque to the user. The client will use it when it reactivates the object at a later time (using `restart_persistent_systemid`), most likely after the client process has terminated and restarted.

Return Value `CORBA::Object_ptr`
An object reference created with the `ObjectId` generated by the system and the `rep_id` provided by the user. The object reference will need to be converted to a specific object type by invoking the `_narrow()` operation defined for the specific object. The caller is responsible for releasing the object when the conversion is done.

Description	<p>This operation performs the following actions:</p> <ul style="list-style-type: none"> ◆ Activates an object using the <code>Servant</code> supplied to service objects of the type <code>rep_id</code>, using an <code>ObjectId</code> generated by the system. ◆ Sets the ORB and the POA into the state in which they will accept requests on that object. ◆ Sets the output parameter <code>stroid</code> to the stringified version of an <code>ObjectId</code> assigned by the system. ◆ Returns an object reference to the activated object. The returned object reference will be valid even after termination of the client. That is, if the client terminates, restarts again, and then activates a servant with the same <code>rep_id</code> and for the same <code>ObjectId</code>, the servant will accept requests made on that same object reference. Since the <code>ObjectId</code> was generated by the system, the application has to save that <code>ObjectId</code>.
Exceptions	<p><code>ServantAlreadyActive</code></p> <p>The servant is already being used for a callback. A servant can be used only for a callback with a single <code>ObjectId</code>. To receive callbacks on objects containing different <code>ObjectIds</code>, you must create different servants and activate them separately. The same servant can be re-used only if a <code>stop</code> operation tells the system to stop using the servant for its original <code>ObjectId</code>.</p> <p><code>CORBA::BAD_PARAMETER</code></p> <p>The repository ID was a null string or the servant was a null pointer.</p> <p><code>CORBA::IMP_LIMIT</code></p> <p>In addition to other system reasons for this exception, a reason unique to this situation is that the joint client/server was not initialized with a port number; therefore, a persistent object reference cannot be created.</p>

restart_persistent_systemid

Synopsis Activates an object, sets the ORB and the POA to the proper state, and returns an object reference to the activated object.

IDL

```
Object restart_persistent_systemid(
    in PortableServer::Servant    servant,
    in CORBA::RepositoryId        rep_id,
    in string                      stroid)
    raises (ServantAlreadyActive, ObjectAlreadyActive);
```

C++ Binding

```
CORBA::Object_ptr restart_persistent_systemid(
    PortableServer::Servant    servant,
    const char*                rep_id,
    const char*                stroid);
```

Java Binding

```
org.omg.CORBA.Object restart_persistent_systemid(
    org.omg.PortableServer.Servant    servant,
    java.lang.String                  rep_id,
    java.lang.String                  stroid);
```

Arguments

`servant`
An instance of the C++ implementation class for the interface.

`rep_id`
The repository id of the interface.

`stroid`
The stringified version of the `ObjectId` provided by the user to be set in the object reference being created. It must have been returned from a previous call on `start_persistent_systemid`.

Return Value `CORBA::Object_ptr`
An object reference created with the stringified `ObjectId` `stroid` and the `rep_id` provided by the user. The object reference will need to be converted to a specific object type by invoking the `_narrow()` operation defined for the specific object. The caller is responsible for releasing the object when done.

Description This operation performs the following actions:

- ◆ Activates an object using the `Servant` supplied to service objects of the type `rep_id`, using the supplied `stroid` (a stringified `ObjectId`), which must have been obtained by a previous call on `start_persistent_systemid`.

- ◆ Sets the ORB and the POA into the state in which they will accept requests on that object.
- ◆ Returns an object reference to the object activated.
- ◆ The re-activation would be done using the "restart_persistent_systemid" method.

Exceptions `ServantAlreadyActive`

The servant is already being used for a callback. A servant can be used only for a callback with a single `ObjectId`. To receive callbacks on objects containing different `ObjectIds`, you must create different servants and activate them separately. The same servant can be re-used only if a `stop_object` operation tells the system to stop using the servant for its original `ObjectId`.

`ObjectAlreadyActive`

The stringified `ObjectId` is already being used for a callback. A given `ObjectId` can have only one servant associated with it. If you wish to change to a different servant, you must first invoke `stop_object` with the servant currently in use.

`CORBA::BAD_PARAM`

The repository id was a null string or the servant was a null pointer or the `ObjectId` supplied was not previously assigned by the system.

`CORBA::IMP_LIMIT`

In addition to other system reasons for this exception, a reason unique to this situation is that the joint client/server was not initialized with a port number; therefore, a persistent object reference cannot be created.

start_persistent_userid

Synopsis Activates an object, sets the ORB and the POA to the proper state, and returns an object reference to the activated object.

IDL

```
Object start_persistent_userid(
    portableServer::Servant      a_servant,
    in CORBA::RepositoryId      rep_id,
    in string                    stroid)
    raises ( ServantAlreadyActive, ObjectAlreadyActive );
```

C++ Binding

```
CORBA::Object_ptr start_persistent_userid (
    PortableServer::Servant      servant,
    const char*                  rep_id,
    const char*                  stroid);
```

Java Binding

```
org.omg.CORBA.Object start_persistent_userid(
    org.omg.PortableServer.Servant      servant,
    java.lang.String                    rep_id,
    java.lang.String                    stroid);
```

Arguments

`servant` An instance of the C++ implementation class for the interface.

`rep_id` The repository id of the interface.

`stroid` The stringified version of an `ObjectId` provided by the user to be set in the object reference being created. The `stroid` holds application-specific data and is opaque to the ORB.

Return Value `CORBA::Object_ptr`

An object reference created with the stringified `ObjectId` `stroid` and the `rep_id` provided by the user. The object reference will need to be converted to a specific object type by invoking the `_narrow()` operation defined for the specific object. The caller is responsible for releasing the object when the conversion is done.

Description	<p>This operation performs the following actions:</p> <ul style="list-style-type: none"> ◆ Activates an object using the <code>Servant</code> supplied to service objects of the type <code>rep_id</code>, using the object id <code>stroid</code>. ◆ Sets the ORB and the POA into the state in which they will accept requests on that object. ◆ Returns an object reference to the activated object. The returned object reference will be valid even after termination of the client. That is, if the client terminates, and restarts again, and then activates a servant with the same <code>rep_id</code> and for the same <code>ObjectId</code>, the servant will accept requests made on that same object reference.
Exceptions	<p><code>ServantAlreadyActive</code></p> <p>The servant is already being used for a callback. A servant can be used only for a callback with a single <code>ObjectId</code>. To receive callbacks on objects containing different <code>ObjectIds</code>, you must create different servants and activate them separately. The same servant can be re-used only if a <code>stop_object</code> operation tells the system to stop using the servant for its original <code>ObjectId</code>.</p> <p><code>ObjectAlreadyActive</code></p> <p>The stringified <code>ObjectId</code> is already being used for a callback. A given <code>ObjectId</code> can have only one servant associated with it. If you wish to change to a different servant, you must first invoke <code>stop_object</code> with the servant currently in use.</p> <p><code>CORBA::BAD_PARAM</code></p> <p>The repository ID was a null string or the servant was a null pointer.</p> <p><code>CORBA::IMP_LIMIT</code></p> <p>In addition to other system reasons for this exception, a reason unique to this situation is that the joint client/server was not initialized with a port number; therefore, a persistent object reference cannot be created.</p>

stop_object

Synopsis Tells the ORB to stop accepting requests on the object that is using the given servant.

IDL `void stop_object(in PortableServer::Servant servant);`

C++ Binding `void stop_object(PortableServer::Servant servant);`

Java Binding `void stop_object(org.omg.PortableServer.Servant servant);`

Argument `servant`
An instance of the C++ implementation class for the interface. The association between this servant and its `ObjectId` will be removed from the Active Object Map.

Description This operation tells the ORB to stop accepting requests on the given servant. It does not matter what state the servant is in, activated or deactivated; no error is reported.

Return Value None.

Exceptions None.

stop_all_objects

Synopsis	Tells the ORB to stop accepting requests on all servants.
IDL	<code>void stop_all_objects ();</code>
C++ Binding	<code>void stop_all_objects ();</code>
Java Binding	<code>void stop_all_objects ();</code>
Return Value	None.
Description	This operation tells the ORB to stop accepting requests on all servants that have been set up in this process.
Usage Note	If a client calls the <code>ORB::shutdown</code> method, then it must not subsequently call <code>stop_all_objects</code> .
Exceptions	None.

get_string_oid

Synopsis Requests the string version of the `ObjectId` of the current request.

IDL `string get_string_oid() raises (NotInRequest);`

C++ Binding `char* get_string_oid();`

Java Binding `java.lang.String get_string_oid();`

Return Value `char*`
The string version of the `ObjectId` of the current request. This is the string that was supplied when the object reference was created. The string is meaningful to users only in the case when the object reference was created by the `start_persistent_userid` function. (That is, the `ObjectId` created by `start_transient` and `start_persistent_systemid` were created by the ORB and has no relationship to the user application.)

Description This operation returns the string version of the `ObjectId` of the current request.

Exceptions `NotInRequest`
The function was called when the ORB was not in the context of a request (that is, not while the ORB was servicing a request in method code). Do not call this function from client code. It is legal only during the execution of a method of the callback object (that is, the servant).

~Callbacks

Synopsis	Destroys the callback object.
C++ Binding	<code>BEAWrapper::~~Callbacks() ;</code>
Java Binding	<code>public ~Callbacks() ;</code>
Arguments	None.
Return Value	None.
Description	This destructor destroys the callback object.
Usage Note	If a client wants to get rid of the wrapper, but not shut down the ORB, then the client must call the <code>stop_all_objects</code> method.
Exceptions	None.

10 Development Commands

This chapter is an alphabetical reference that describes each WLE development command and Interface Repository administration command for the Windows NT and UNIX environments. A list of valid parameters and options is shown for each command.

The following commands are described:

- ◆ `buildobjclient`
- ◆ `buildobjserver`
- ◆ `genicf`
- ◆ `idl`

Note: For descriptions of the `idl2ir`, `irdel`, and `ir2idl` commands, see the *Administration Guide*.

Before executing a WLE command, you must ensure that the `bin` directory is in your defined path:

On Windows NT

```
Set Path=%TUXDIR%\Bin;%Path%
```

On UNIX

For c shell (csh): `set path = ($TUXDIR/bin $path)`

For Bourne (sh) or Korn (ksh): `PATH=$TUXDIR/bin:$PATH`
`export PATH`

To set environment variables:

On Windows NT

`set var=value`

On UNIX:

For c shell: `setenv var value`

For Bourne and Korn (sh/ksh): `var=value`
`export var`

buildobjclient

Synopsis	Constructs a WLE client application.
Syntax	<code>buildobjclient [-v][-o <i>name</i>] [-f <i>firstfile-syntax</i>] [-l <i>lastfile-syntax</i>] -P</code>
Description	<p>Use the <code>buildobjclient</code> command to construct a WLE client application. The command combines the files specified in the <code>-f</code> and <code>-l</code> options with the standard WLE libraries to form a client application. The client application is built using the default C++ language compile command defined for the operating system in use.</p> <p>All specified <code>.c</code> and <code>.cpp</code> files are compiled in one invocation of the compilation system for the operating system in use. Users may specify the compiler to invoke by setting the <code>CC</code> environment variable to the name of the compiler. If the <code>CC</code> environment variable is not defined when <code>buildobjclient</code> is invoked, the default C++ language compile command for the operating system in use will be invoked to compile all <code>.c</code> and <code>.cpp</code> files.</p> <p>Users may specify options to be passed to the compiler by setting the <code>CFLAGS</code> or the <code>CPPFLAGS</code> environment variables. If <code>CFLAGS</code> is not defined when <code>buildobjclient</code> is invoked, the <code>buildobjclient</code> command uses the value of <code>CPPFLAGS</code> if that variable is defined.</p>
Options	<p><code>-v</code></p> <p>Specifies that the <code>buildobjclient</code> command should work in verbose mode. In particular, it writes the compile command to its standard output.</p> <p><code>-o <i>name</i></code></p> <p>Specifies the name of the client application generated by this command. If the name is not supplied, the application file is named <code>client<.type></code>, where <code>type</code> is an extension that is dependent on the operating system for an application (for example, on a UNIX system, there would not be a <code>type</code>; on a Windows NT system, the <code>type</code> would be <code>.EXE</code>).</p> <p><code>-f <i>firstfile-syntax</i></code></p> <p>Specifies a file to be included first in the compile and link phases of the <code>buildobjclient</code> command. The specified file is included before the WLE libraries are included. There are three ways of specifying a file or files, as shown in Table 10-1.</p>

Table 10-1 Specifying the First Filename(s)

Filename Specification	Definition
<code>-f firstfile</code>	One file is specified.
<code>-f "file1.cpp file2.cpp file3.cpp ..."</code>	Multiple files may be specified if they are enclosed in quotation marks and are separated by white space.

Note: Filenames that include spaces are not supported

Note: The `-f` option may be specified multiple times.

`-l lastfile-syntax`
Specifies a file to be included last in the compile and link phases of the `buildobjclient` command. The specified file is included after the WLE libraries are included. There are three ways of specifying a file, as shown in Table 10-2.

Table 10-2 Specifying the Last Filename(s)

Filename Specification	Definition
<code>-l lastfile</code>	One file is specified.
<code>-l "file1.cpp file2.cpp file3.cpp ..."</code>	Multiple files may be specified if they are enclosed in quotation marks and are separated by white space.

Note: The `-l` option may be specified multiple times.

`-P`
Specifies that the appropriate POA libraries should be linked into the image (that is, libraries that allow a client to also function as a server). The resulting image can act as a server and can use the `Callbacks` wrapper class for creating objects. The resulting joint client/server cannot take advantage of the object state management and transaction management provided by the WLE TP Framework. The `-P` switch should have been passed to the IDL compiler

when generating the client. Use `buildobjserver` to build a server with all the support provided by the TP Framework. The default is to not link in the server libraries; that is, the default is to create a client only, not a joint client/server.

`-h` or `-?`

Provides help that explains the usage of the `buildobjclient` command. No other action results.

Environment
Variables

`TUXDIR`

Finds the WLE libraries and include files to use when compiling the client applications.

`CC`

Indicates the compiler to use to compile all files with `.c` or `.cpp` file extensions. If not defined, the default C++ language compile command for the operating system in use will be invoked to compile all `.c` and `.cpp` files.

`CFLAGS`

Indicates any arguments that are passed as part of the compiler command line for any files with a `.c` or `.cpp` file extensions. If `CFLAGS` does not exist in the `buildobjclient` command environment, the `buildobjclient` command checks for the `CPPFLAGS` environment variable.

`CPPFLAGS`

Note: Arguments passed by the `CFLAGS` environment variable take priority over the `CPPFLAGS` variable.

Contains a set of arguments that are passed as part of the compiler command line for any files with a `.c` or `.cpp` file extensions.

This is in addition to the command line option `-I$(TUXDIR)/include` for UNIX systems or the command line option `/I%TUXDIR%\include` for Windows NT systems, which is passed automatically by the `buildobjclient` command. If `CPPFLAGS` does not exist in the `buildobjclient` command environment, no compiler commands are added.

`LD_LIBRARY_PATH` (UNIX systems)

Indicates which directories contain shared objects to be used by the compiler, in addition to the objects shared by the WLE software. A colon (`:`) is used to separate the list of directories.

`LIB` (Windows NT systems)

Indicates a list of directories within which to find libraries. A semicolon (`;`) is used to separate the list of directories.

Portability The `buildobjclient` command is not supported on client-only platforms.

Examples The following example builds a WLE client application on an NT system:

```
set CPPFLAGS=-I%APPDIR%\include
buildobjclient -o empclient.exe -f emp_c.cpp -l userlib1.lib
```

The following example builds a WLE client application on a UNIX system using the c shell:

```
setenv CPPFLAGS=$APPDIR/include
buildobjclient -o empclient -f emp_c.cpp -l userlib1.a
```

buildobjserver

Synopsis	Constructs a WLE server application.
Syntax	<code>buildobjserver [-v] [-o name] [-f firstfile-syntax] [-l lastfile-syntax] [-r rmname]</code>
Description	<p>Use the <code>buildobjserver</code> command to construct a WLE server application. The command combines the files specified in the <code>-f</code> and <code>-l</code> options with the main routine and the standard WLE libraries to form a server application. The server application is built using the default C++ compiler provided for the platform.</p> <p>All specified <code>.c</code> and <code>.cpp</code> files are compiled in one invocation of the compilation system for the operating system in use. Users may specify the compiler to invoke by setting the <code>CC</code> environment variable to the name of the compiler. If the <code>CC</code> environment variable is not defined when <code>buildobjserver</code> is invoked, the default C++ language compile command for the operating system in use will be invoked to compile all <code>.c</code> and <code>.cpp</code> files.</p> <p>Users may specify options to be passed to the compiler by setting the <code>CFLAGS</code> or the <code>CPPFLAGS</code> environment variables. If <code>CFLAGS</code> is not defined when <code>buildobjserver</code> is invoked, the <code>buildobjserver</code> command uses the value of <code>CPPFLAGS</code> if that variable is defined.</p>
Options	<p><code>-v</code></p> <p>Specifies that the <code>buildobjserver</code> command should work in verbose mode. In particular, it writes the compile command to its standard output.</p> <p><code>-o name</code></p> <p>Specifies the name of the server application generated by this command. If the name is not supplied, the application file is named <code>server<.type></code>, where <code>type</code> is the extension that is dependent on the operating system for an application (for example, on UNIX systems, there would not be a <code>type</code>; on Windows NT systems, the <code>type</code> would be <code>.EXE</code>).</p> <p><code>-f firstfile-syntax</code></p> <p>Specifies a file to be included first in the compile and link phases of the <code>buildobjserver</code> command. The specified file is included before the WLE libraries are included. For a description of the three ways to specify a file or files, see Table 10-1, “Specifying the First Filename(s),” on page 10-4.</p> <p><code>-l lastfile-syntax</code></p> <p>Specifies a file to be included last in the compile and link phases of the <code>buildobjserver</code> command. The specified file is included after the WLE</p>

libraries are included. For a description of the three ways to specify a file or files, see Table 10-2, “Specifying the Last Filename(s),” on page 10-4.

`-r rmname`

Specifies the resource manager associated with this server. The value `rmname` must appear in the resource manager table located in `$TUXDIR/udataobj/RM` on UNIX systems or `%TUXDIR%\udataobj\RM` on Windows NT systems. Each entry in this file is of the form

`rmname:rmstructure_name:library_names.`

Using the `rmname` value, the entry in `$TUXDIR/udataobj/RM` or `%TUXDIR%\udataobj\RM` automatically includes the associated libraries for the resource manager and properly sets up the interface between the transaction manager and the resource manager. The value `TUXEDO/SQL` includes the libraries for the BEA TUXEDO System/SQL resource manager. Other values can be specified as they are added to the resource manager table. If the `-r` option is not specified, the default is to use the null resource manager.

`-h or -?`

Provides help that explains the usage of the `buildobjserver` command. No other action results.

Environment Variables

`TUXDIR`

Finds the WLE libraries and include files to use when compiling the server application.

`CC`

Indicates the compiler to use to compile all files with `.c` or `.cpp` file extensions that are passed in through the `-l` or `-f` options.

`CFLAGS`

Specifies any arguments that are passed as part of the compiler command line for any files with `.c` or `.cpp` file extensions. If `CFLAGS` does not exist in the `buildobjserver` command environment, the `buildobjserver` command checks for the `CPPFLAGS` environment variable.

`CPPFLAGS`

Note: Arguments passed by the `CFLAGS` environment variable take priority over the `CPPFLAGS` environment variable.

Contains a set of arguments that are passed as part of the compiler command line for any files with a `.c` or `.cpp` file extensions. This is in addition to the command line option `-I$(TUXDIR)/include` for UNIX systems or the command line option `/I%TUXDIR%\include` for Windows NT systems,

which is passed automatically by the `buildobjserver` command. If `CPPFLAGS` does not exist in the `buildobjserver` command environment, no compiler commands are added.

LD_LIBRARY_PATH (UNIX systems)

Indicates which directories contain shared objects to be used by the compiler, in addition to the WLE shared objects. A colon (:) is used to separate the list of directories.

LIB (Windows NT systems)

Indicates a list of directories within which to find libraries. A semicolon (;) is used to separate the list of directories.

Portability The `buildobjserver` command is not supported on client-only WLE systems.

Examples The following example builds a WLE server application on a UNIX system using the `emp_s.cpp` and `emp_i.cpp` files:

```
buildobjserver -r TUXEDO/SQL -o unobserved
               -f "emp_s.cpp emp_i.cpp"
```

The following example shows how to use the `CC` and `CFLAGS` environment variables with the `buildobjserver` command. The example also shows how to link in the `math` library on UNIX systems using the Bourne or Korn shells using the `-f` and `-lm` options:

```
CFLAGS=-g CC=/bin/cc \
buildobjserver -r TUXEDO/SQL -o TLR -f TLR.o -f util.o -l -lm
```

The following example shows how to use the `buildobjserver` command on UNIX systems with no resource manager specified:

```
buildobjserver -o PRINTER -f PRINTER.o
```

Sample RM Files The following are sample RM files for all the supported operating system platforms:

Windows NT

```
Oracle_XA:xaosw;C:\Orant\rdbms73\xa\xa73.lib
C:\Orant\pro22\lib\msvc\sqllib18.lib
```

Solaris

```
Oracle_XA:xaosw:-L$ORACLE_HOME/rdbms/lib
-L$ORACLE_HOME/precomp/lib -lc
-L/home4/m01/app/oracle/product/7.3.2/lib -lsql -lclntsh
-lsqlnet -lnchr -lcommon -lgeneric -lepc -lnlsrtl3 -lc3v6
```

```
-lcore3 -lsocket -lnsl -lm -ldl -lthread
```

Tru64 UNIX

```
Oracle_XA:xaosw:-L${ORACLE_HOME}/lib -lxa  
${ORACLE_HOME}/lib/libsql.a -lsqlnet -lnchr -lsqlnet  
${ORACLE_HOME}/lib/libclient.a -lcommon -lgeneric -lsqlnet  
-lnchr -lsqlnet ${ORACLE_HOME}/lib/libclient.a -lcommon  
-lgeneric -lepc -lepcpt -lnlsrtl3 -lc3v6 -lcore3  
-lnlsrtl3 -lcore3 -lnlsrtl3 -lm
```

AIX

```
Oracle_XA:xaosw:-L${ORACLE_HOME}/lib -lxa -lsql -lsqlnet  
-lnchr -lclient -lcommon -lgeneric -lepc -lnlsrtl3 -lc3v6  
-lcore3 -lm -lld
```

HP-UX : Oracle 8.04

```
Oracle_XA:xaosw:-L${ORACLE_HOME}/lib -lclntsh
```

genicf

Synopsis Generates an ICF file.

Syntax `genicf [options] idl-filename...`

Description Given the `idl-filename(s)`, generates an ICF file that provides the code generation process with additional information about policies on implementations and the relationship between implementations and the interface they implement. If an ICF file is provided as input to the `idl` command, the `idl` command generates server code for only the implementation/interface pairs specified in the ICF file.

The generated ICF file has the same filename as the first `idl-filename` specified on the command line, but with a `.icf` extension.

If incorrect OMG IDL syntax is specified in the `idl-filename(s)` file, appropriate errors are returned.

Options `-D identifier=[definition]`

Performs the same function as the `#define` C++ preprocessor directive; that is, the `-D` option defines a token string or a macro to be substituted for every occurrence of a given identifier in the definition file. If a definition is not specified, the identifier is defined as 1. Multiple `-D` options can be specified. White space between the `-D` option and the identifier is optional.

`-I pathname`

Specifies directories within which to search for include files, in addition to any directories specified with the `#include` OMG IDL preprocessor directive. Multiple directories can be specified by using multiple `-I` options.

There are two types of `#include` OMG IDL preprocessor directives: system (for example, `<a.idl>`) and user (for example, `"a.idl"`). On UNIX systems, the path for system `#include` directories is `/usr/include` and any directories specified with the `-I` option; the path for user `#include` directives is the location of the file containing the `#include` directive, followed by the path specified for the system `#include` directive. On Windows NT systems, no distinction is made between the system `#include` directories and the user `#include` directives.

`-h` and `-?`

Provides help that explains the usage of the `genicf` command. No other action results.

Example This command creates the `emp.icf` file: `genicf emp.idl`.

idl

Synopsis	Compiles the Object Management Group (OMG) Interface Definition Language (IDL) file and generates the files required for the interface.
Syntax	<code>idl [-i] [-Did[=value]] [-I pathname][-h] [-P] [-T] idl-filename... [icf-filename...]</code>
Description	<p>Given the provided <code>idl-filename()</code> file(s) and optional <code>icf-filename()</code> file(s), the <code>idl</code> command generates the following files:</p> <p><code>idl-filename_c.cpp</code> Client stub (includes embedded user-defined data type functions).</p> <p><code>idl-filename_c.h</code> Class definitions for interfaces.</p> <p><code>idl-filename_s.cpp</code> Server skeleton containing an implementation of the <code>POA_skeleton</code> classes.</p> <p><code>idl-filename_s.h</code> <code>POA_skeleton</code> class definitions.</p> <p><code>idl-filename_i.cpp</code> Example version of the implementation. This file is generated only when the <code>-i</code> option is given.</p> <p><code>idl-filename_i.h</code> Class definition of an example implementation that inherits from the <code>POA_skeleton</code> class. This file is generated only when the <code>-i</code> option is given.</p> <p>Note: If any ICF files are specified, the information within the ICF files is used to provide the code generator with information about the interface/implementations that override the defaults. Typically, an activation policy and a transaction policy for an implementation may be specified in the ICF file. If no ICF files are specified, default policies are in effect for all of the interfaces specified in the OMG IDL file, and skeleton code for all of the interfaces is generated. If an <code>icf-filename</code> is provided as input to the <code>idl</code> command, only the implementation/interfaces specified in the <code>icf-filename</code> are generated as part of the server. For information about ICF syntax, see “ICF Syntax” on page 2-2.</p>

The IDL compiler places the generated client stub information in the `filename_c.cpp` and `filename_c.h` files. The generated server skeleton information is placed in the `filename_s.cpp` and `filename_s.h` files.

The IDL compiler overwrites the generated client stub files (`filename_c.cpp` and `filename_c.h`), and the generated server skeleton files (`filename_s.cpp` and `filename_s.h`). Any previous versions are destroyed.

When using the `-i` option, the IDL compiler overwrites the sample implementation class definition file (`filename_i.h`). Previous versions are destroyed. The sample implementation file (`filename_i.cpp`) is overwritten, however, any code contained within the code preservation blocks is preserved and restored in the newly generated file. To avoid the loss of data, it is recommended that you copy the sample implementation files (`filename_i.h` and `filename_i.cpp`) to a safe location before regenerating these files.

If an unknown option is passed to this command, the offending option and a usage message is displayed to the user, and the compile is not performed.

For more information about OMG IDL syntax, see Chapter 1, “OMG IDL Syntax”.

Parameter `idl filename`

The name of one or more files that contain OMG IDL statements.

Options `-D identifier[=definition]`

Performs the same function as the `#define` C++ preprocessor directive; that is, the `-D` option defines a token string or a macro to be substituted for every occurrence of a given identifier in the definition file. If a definition is not specified, the identifier is defined as 1. Multiple `-D` options can be specified. White space between the `-D` option and the name is optional.

`-I pathname`

Specifies directories within which to search for include files, in addition to any directories specified with the `#include` OMG IDL preprocessor directive. Multiple directories can be specified by using multiple `-I` options.

There are two types of `#include` OMG IDL preprocessor directives: `system` (for example, `<a.idl>`) and `user` (for example, `"a.idl"`). The path for `system` `#include` directories is the system include directory and any directories specified with the `-I` option. The path for `user` `#include` directives is the location of the file containing the `#include` directive, followed by the path specified for the `system` `#include` directive.

By default, the text in files included with an `#include` directive is not included in the client and server code that is generated.

`-i`

Results in `idl-filename_i.cpp` files being generated. These files contain example templates for the implementations that implement the interfaces specified in the OMG IDL file.

Note: When using the `idl` command `-i` option to update your implementation files, proceed as follows to update your implementation files:

1. Back up your implementation files.
2. If you are migrating from BEA ObjectBroker to WLE, edit your generated implementation files to change the code preservation block delimiters from `"OBB_PRESERVE_BEGIN"` and `"OBB_PRESERVE_END"` to `"M3_PRESERVE_BEGIN"` and `"M3_PRESERVE_END"`.
3. If you added include files to your method implementation file (`*_i.cpp`), edit the file and move the includes inside the `INCLUDES` preservation block.
4. Regenerate your edited implementation files (using the `idl` command with the `-i` option).
5. If you previously made modifications to the implementation definition file (`*_i.h`), edit the newly generated definition file and add your modifications back in. Be sure to put your modifications inside the code preservation blocks so subsequent updates will automatically retain them. Pay particular attention to the implementation constructor and destructor functions; the function signatures have been changed in this release.
6. If you previously made modifications outside the preservation blocks of the method implementation file (`*_i.cpp`) or to the implementation constructor and destructor functions, edit the newly generated file and add those modifications. Be sure to put the modifications inside a preservation block so subsequent updates will automatically retain them.

`-P`

Generates server code that uses the POA instead of the TP Framework. With this option specified, the skeleton class does not inherit from the TP Framework `Tobj_ServantBase` class, but instead inherits directly from the `PortableServer::ServantBase` class. By default, the skeleton class uses the TP Framework. So you must use this switch when you are developing joint client/servers as these servers do not use the TP framework.

Not having the `Tobj_ServantBase` class in the inheritance tree for a servant means that the servant does not have `activate_object` and

deactivate_object methods. In WLE servers these methods are called by the TP Framework to dynamically initialize and save a servant's state before invoking a method on the servant. For WLE joint client/servers, user-written code must explicitly create a servant and initialize a servant's state; therefore, the Tobj_ServantBase operations are not needed. When using the -P option, ICF files are not used because the TP Framework is not available.

-T

Generates tie-based servant code that allows the use of delegation to tie an instance of a C++ implementation class to the servant. This option allows classes that are not related to skeletons by inheritance to implement CORBA object operations. This option is set to off by default.

Note: When using tie classes with the Digital C++ V6.0 compiler for Tru64 UNIX, include the `-noimplicit_include` option in the definition of the `CFLAGS` and `CPPFLAGS` environment variables used by the `buildobjserver` command. This option prevents the Digital C++ compiler from automatically including the server skeleton definition file (`_s.cpp`) everywhere the server skeleton header file (`_s.h`) is included. This is necessary to avoid multiply-defined symbol errors. See *Using DIGITAL C++ for Tru64 Unix Systems* for additional information on using class templates (such as the tie classes) with Digital C++.

-h or -?

Provides help that explains the usage of the `idl` command. No other action results.

Examples `idl emp.idl`
 `idl emp.idl emp.icf`

11 Mapping of OMG IDL Statements to C++

This chapter discusses the mappings from OMG IDL statements to C++.

Note: Some of the information in this chapter is taken from the *Common Object Request Broker: Architecture and Specification*. Revision 2.2, February 1998, published by the Object Management Group (OMG). Used with permission by OMG.

Mappings

OMG IDL-to-C++ mappings are described for the following:

- ◆ Data types
- ◆ Strings
- ◆ Constants
- ◆ C PIDL
- ◆ Enums
- ◆ Portableserver functions
- ◆ Pseudo-objects
- ◆ Serverless objects

- ◆ Structs
- ◆ Unions
- ◆ Usage
- ◆ Sequences
- ◆ Arrays
- ◆ Exceptions
- ◆ Typedefs
- ◆ Operations (implementing)
- ◆ Operations (interfaces)
- ◆ Attributes
- ◆ Any types

This chapter also describes the generated var classes for user-defined data types.

Data Types

Each OMG IDL data type is mapped to a C++ data type or class.

Basic Data Types

The basic data types in OMG IDL statements are mapped to C++ typedefs in the CORBA module, as shown in Table 11-1.

Table 11-1 Basic OMG IDL and C++ Data Types

OMG IDL	C++	C++ Out Type
short	CORBA::Short	CORBA::Short_out
long	CORBA::Long	CORBA::Long_out
unsigned short	CORBA::UShort	CORBA::UShort_out

Table 11-1 Basic OMG IDL and C++ Data Types

OMG IDL	C++	C++ Out Type
unsigned long	CORBA::ULong	CORBA::ULong_out
float	CORBA::Float	CORBA::Float_out
double	CORBA::Double	CORBA::Double_out
char	CORBA::Char	CORBA::Char_out
boolean	CORBA::Boolean	CORBA::Boolean_out
octet	CORBA::Octet	CORBA::Octet_out

Note: On a 64-bit machine where a long integer is 64 bits, the definition of CORBA::Long would still refer to a 32-bit integer.

Complex Data Types

Object, pseudo-object, and user-defined types are mapped as shown in Table 11-2.

Table 11-2 Object, Pseudo-Object, and User-Defined OMG IDL and C++ Types

OMG IDL	C++
Object	CORBA::Object_ptr
struct	C++ struct
union	C++ class
enum	C++ enum
string	char *
sequence	C++ class
array	C++ array

The mapping for strings and UDTs is described in more detail in the following sections.

Strings

A string in OMG IDL is mapped to `char *` in C++. Both bounded and unbounded strings are mapped to `char *`. CORBA strings in C++ are NULL-terminated and can be used wherever a `char *` type is used.

If a string is contained within another user-defined type, such as a `struct`, it is mapped to a `CORBA::String_var` type. This ensures that each member in the struct manages its own memory.

Strings must be allocated and deallocated using the following member functions in the CORBA class:

- ◆ `string_alloc`
- ◆ `string_dup`
- ◆ `string_free`

Note: The `string_alloc` function allocates `len+1` characters so that the resulting string has enough space to hold a trailing NULL character.

For more information about string member functions, see the section “Strings” on page 1-119.

Constants

A constant in OMG IDL is mapped to a C++ `const` definition. For example, consider the following OMG IDL definition:

```
// OMG IDL

const string CompanyName = "BEA Systems Incorporated";

module INVENT
{
    const string Name = "Inventory Modules";

    interface Order
    {
        const long MAX_ORDER_NUM = 10000;
    }
}
```



```
    };
};
```

This definition maps to C++ as follows:

```
// C++

const char *const
    CompanyName = "BEA Systems Incorporated";
. . .
class INVENT
{
    static const char *const Name;
    . . .

    class Order : public virtual CORBA::Object
    {
        static const CORBA::Long MAX_ORDER_NUM;
        . . .
    };
};
```

Top-level constants are initialized in the generated .h include file, but module and interface constants are initialized in the generated client stub modules.

The following is an example of a valid reference to the `MAX_ORDER_NUM` constant, as defined in the previous example:

```
CORBA::Long acct_id = INVENT::Order::MAX_ORDER_NUM;
```

Enums

An enum in OMG IDL is mapped to an enum in C++. For example, consider the following OMG IDL definition:

```
// OMG IDL

module INVENT
{
    enum Reply {ACCEPT, REFUSE};
}
```

This definition maps to C++ as follows:

```
// C++  
  
class INVENT  
{  
    . . .  
  
    enum Reply {ACCEPT, REFUSE};  
};
```

The following is an example of a valid reference to the enum defined in the previous example. You refer to enum as follows:

```
INVENT::Reply accept_reply;  
accept_reply = INVENT::ACCEPT;
```

Structs

A struct in OMG IDL is mapped to a C++ struct.

The generated code for a struct depends upon whether it is fixed-length or variable-length. For more information about fixed-length versus variable-length types, see the section “Fixed-Length Versus Variable-Length User-Defined Types” on page 11-49.

Fixed-Length Versus Variable-Length Structs

A variable-length struct contains an additional assignment operator member function to handle assignments between two variable-length structs.

For example, consider the following OMG IDL definition:

```
// OMG IDL  
  
module INVENT  
{  
    // Fixed-length  
    struct Date  
    {  
        long year;
```

```
        long month;
        long day;
    };

    // Variable-length
    struct Address
    {
        string aptNum;
        string streetName;
        string city;
        string state;
        string zipCode;
    };
};
```

This definition maps to C++ as follows:

```
// C++

class INVENT
{
    struct Date
    {
        CORBA::Long year;
        CORBA::Long month;
        CORBA::Long day;
    };

    struct Address
    {
        CORBA::String_var aptNum;
        CORBA::String_var streetName;
        CORBA::String_var city;
        CORBA::String_var state;
        CORBA::String_var zipCode;
        Address &operator=(const Address &_obj);
    };

};
```

Member Mapping

Members of a struct are mapped to the appropriate C++ data type. For basic data types (long, short, and so on), see Table 11-1 on page 11-2. For object references, pseudo-object references, and strings, the member is mapped to the appropriate var class:

◆ CORBA::String_var

◆ CORBA::Object_var

All other data types are mapped as shown in Table 11-2, “Object, Pseudo-Object, and User-Defined OMG IDL and C++ Types,” on page 11-3.

No constructor for a generated struct exists, so none of the members are initialized. Fixed-length structs can be initialized using aggregate initialization. For example:

```
INVENT::Date a_date = { 1995, 10, 12 };
```

Variable-length members map to self-managing types; these types have constructors that initialize the member.

Var

A var class is generated for structs. For more information, see the section “Using var Classes” on page 11-50.

Out

An out class is generated for structs. For more information, see the section “Using out Classes” on page 11-56.

Unions

A union in OMG IDL is mapped to a C++ class. The C++ class contains the following:

- ◆ Constructors
- ◆ Destructors
- ◆ Assignment operators

- ◆ Modifiers for the union value
- ◆ Accessors for the union value
- ◆ Modifiers and accessors for the union discriminator

For example, consider the following OMG IDL definition:

```
// OMG IDL

union OrderItem switch (long)
{
    case 1: itemStruct itemInfo;
    case 2: orderStruct orderInfo;
    default: ID idInfo;
};
```

This definition maps to C++ as follows:

```
// C++

class OrderItem
{
public:
    OrderItem();
    OrderItem(const OrderItem &);
    ~OrderItem();

    OrderItem &operator=(const OrderItem&);

    void _d (CORBA::Long);
    CORBA::Long _d () const;

    void itemInfo (const itemStruct &);
    const itemStruct & itemInfo () const;
    itemStruct & itemInfo ();

    void orderInfo (const orderStruct &);
    const orderStruct & orderInfo () const;
    orderStruct & orderInfo ();

    void idInfo (ID);
    ID idInfo () const;
```

```
        . . .
    };
```

The default union constructor does not set a default discriminator value for the union; therefore, you cannot call any union accessor member function until you have set the value of the union. The discriminator is an attribute that is mapped through the `_d` member function.

Union Member Accessor and Modifier Member Function Mapping

For each member in the union, accessor and modifier member functions are generated.

In the following code, taken from the previous example, two member functions are generated for the `ID` member function:

```
void idInfo (ID);
ID idInfo () const;
```

In this example, the first function (the modifier) sets the discriminator to the default value and sets the value of the union to the specified `ID` value. The second function, the accessor, returns the value of the union.

Depending upon the data `type` of the union member, additional modifier functions are generated. The member functions generated for each data `type` are as follows:

- ◆ Basic data types—short, long, unsigned short, unsigned long, float, double, char, boolean, and octet

The following example generates two member functions for a basic data `type` with member name `basictype`:

```
void basictype (TYPE);           // modifier
TYPE basictype () const;        // accessor
```

For the mapping from an OMG IDL data `type` to the C++ data `type` `TYPE`, see Table 11-1, “Basic OMG IDL and C++ Data Types,” on page 11-2.

- ◆ Object and pseudo-object

For object and `Typecode` types with member name `objtype`, member functions are generated as follows:

```
void objtype (TYPE);           // modifier
```

```
TYPE objtype () const;      // accessor
```

For the mapping from an OMG IDL data type to the C++ data type `TYPE`, see Table 11-1, “Basic OMG IDL and C++ Data Types,” on page 11-2.

The modifier member function does not assume ownership of the specified object reference argument. Instead, the modifier duplicates the object reference or pseudo-object reference. You are responsible for releasing the reference when it is no longer required.

◆ Enum

For an enum `TYPE` with member name `enumtype`, member functions are generated as follows:

```
void enumtype (TYPE);      // modifier
TYPE enumtype () const;    // accessor
```

◆ String

For strings, one accessor and three modifier functions are generated, as follows:

```
void stringInfo (char *);           // modifier 1
void stringInfo (const char *);     // modifier 2
void stringInfo (const CORBA::String_var &); // modifier 3
const char * stringInfo () const;    // accessor
```

The first modifier assumes ownership of the `char *` parameter passed to it and the union is responsible for calling the `CORBA::string_free` member function on this string when the union value changes or when the union is destroyed.

The second and third modifiers make a copy of the specified string passed in the parameter or contained in the string var.

The accessor function returns a pointer to internal memory of the union; do not attempt to free this memory, and do not access this memory after the union value has been changed or the union has been destroyed.

◆ Struct, union, sequence, and any

For these data types, modifier and accessor functions are generated with

references to the `type`, as follows:

```
void reftype (TYPE &);           // modifier
const TYPE & reftype () const;  // accessor
TYPE & reftype ();              // accessor
```

The modifier function does not assume ownership of the input `type` parameter; instead, the function makes a copy of the data `type`.

◆ Array

For an array, the modifier member function accepts an array pointer while the accessor returns a pointer to an array slice, as follows:

```
void arraytype (TYPE);           // modifier
TYPE_slice * arraytype () const; // accessor
```

The modifier function does not assume ownership of the input `type` parameter; instead, the function makes a copy of the array.

Var

A var class is generated for a union. For more information, see the section “Using var Classes” on page 11-50 .

Out

An out class is generated for a union. For more information, see the section “Using out Classes” on page 11-56.

Member Functions

In addition to the accessor and modifiers, the following member functions are generated for an OMG IDL union of type `TYPE` with switch (long) discriminator:

```
TYPE( ) ;
```

This is the default constructor for a union. No default discriminator is set by this function, so you cannot access the union until you set the value of the union.

```
TYPE( const TYPE & From);
```


This copy constructor deep copies the specified union. Any data in the union parameter is copied. The `From` argument specifies the union to be copied.

```
~TYPE();
```

This destructor frees the data associated with the union.

```
TYPE &operator=(const TYPE & From);
```

This assignment operator copies the specified union. Any existing value in the current union is freed. The `From` argument specifies the union to be copied.

```
void _d (CORBA::Long Descrip);
```

This member function sets the value of the discriminant and frees the current value. The `Descrip` argument specifies the new discriminant. The data type of the argument is determined by the OMG IDL data type specified in the switch statement of the union. For each OMG IDL data type, see Table 11-1, “Basic OMG IDL and C++ Data Types,” on page 11-2 for the C++ data type.

```
CORBA::Long _d () const;
```

This function returns the current discriminant value. The data type of the return value is determined by the OMG IDL data type specified in the switch statement of the union. For each OMG IDL data type, see Table 11-1, “Basic OMG IDL and C++ Data Types,” on page 11-2 for the C++ data type.

Sequences

A sequence in OMG IDL is mapped to a C++ class. The C++ class contains the following:

- ◆ Constructors

Each sequence has the following:

- ◆ A default constructor
- ◆ A constructor that initializes each element
- ◆ A copy constructor

- ◆ Destructors

- ◆ Modifiers for current length (and for maximum, if the sequence is unbounded)
- ◆ Accessors for current length

- ◆ Operator[] functions to access or modify sequence elements
- ◆ Allocation and deallocation member functions

You *must* set the length before accessing any elements.

For example, consider the following OMG IDL definition:

```
// OMG IDL

module INVENT
{
    . . .
    typedef sequence<LogItem>      LogList;
}
```

This definition maps to C++ as follows:

```
// C++

class LogList
{
public:
    // Default constructor
    LogList();

    // Maximum constructor
    LogList(CORBA::ULong _max);

    // TYPE * data constructor
    LogList
    (
        CORBA::ULong _max,
        CORBA::ULong _length,
        LogItem *_value,
        CORBA::Boolean _relse = CORBA_FALSE
    );

    // Copy constructor
    LogList(const LogList&);

    // Destructor
    ~LogList();

    LogList &operator=(const LogList&);
```

```

CORBA::ULong maximum() const;

void length(CORBA::ULong);
CORBA::ULong length() const;

LogItem &operator[](CORBA::ULong _index);
const LogItem &operator[](CORBA::ULong _index) const;

static LogItem *allocbuf(CORBA::ULong _nelems);
static void freebuf(LogItem *);
};

};

```

Sequence Element Mapping

The `operator[]` functions are used to access or modify the sequence element. These operators return a reference to the sequence element. The OMG IDL sequence base type is mapped to the appropriate C++ data type.

For basic data types, see Table 11-1, “Basic OMG IDL and C++ Data Types,” on page 11-2. For object references, TypeCode references, and strings, the base type is mapped to a generated `_ForSeq_var` class. The `_ForSeq_var` class provides the capability to update a string or an object that is stored within the sequence. This generated class has the same member functions and signatures as the corresponding var class. However, this `_ForSeq_var` class honors the setting of the release parameter in the sequence constructor. The distinction is that the `_ForSeq_var` class lets users specify the Release flag, thereby allowing users to control the freeing of memory.

All other data types are mapped as shown in Table 11-2, “Object, Pseudo-Object, and User-Defined OMG IDL and C++ Types,” on page 11-3.

Vars

A var class is generated for a sequence. For more information, see the section “Using var Classes” on page 11-50.

Out

An out class is generated for a sequence. For more information, see the section “Using out Classes” on page 11-56.

Member Functions

For a given OMG IDL sequence *SEQ* with base type *TYPE*, the member functions for the generated sequence class are described as follows:

`SEQ ();`

This is the default constructor for a sequence. The length is set to 0 (zero). If the sequence is unbounded, the maximum is also set to 0 (zero). If the sequence is bounded, the maximum is specified by the OMG IDL type and cannot be changed.

`SEQ (CORBA::ULong Max);`

This constructor is present only if the sequence is unbounded. This function sets the length of the sequence to 0 (zero) and sets the maximum of the buffer to the specified value. The `Max` argument specifies the maximum length of the sequence.

`SEQ (CORBA::ULong Max, CORBA::ULong Length, TYPE * Value,
CORBA::Boolean Release);`

This constructor sets the maximum, length, and elements of the sequence. The `Release` flag determines whether elements are released when the sequence is destroyed. Explanations of the arguments are as follows:

`Max`

The maximum value of the sequence. This argument is not present in bounded sequences.

`Length`

The current length of the sequence. For bounded sequences, this value must be less than the maximum specified in the OMG IDL type.

`Value`

A pointer to the buffer containing the elements of the sequence.

`Release`

Determines whether elements are released. If this flag has a value of `CORBA_TRUE`, the sequence assumes ownership of the buffer pointed to by the `Value` argument. If the `Release` flag is `CORBA_TRUE`, this buffer must be allocated using the `allocbuf` member function, because it will be freed using the `freebuf` member function when the sequence is destroyed.

`SEQ(const S& From);`

This copy constructor deep copies the sequence from the specified argument. The `From` argument specifies the sequence to be copied.

```
~SEQ();
```

This destructor frees the sequence and, depending upon the `Release` flag, may free the sequence elements.

```
SEQ& operator=(const SEQ& From);
```

This assignment operator deep copies the sequence from the specified sequence argument. Any existing elements in the current sequence are released if the `Release` flag in the current sequence is set to `CORBA_TRUE`. The `From` argument specifies the sequence to be copied.

```
CORBA::ULong maximum() const;
```

This function returns the maximum of the sequence. For a bounded sequence, this is the value set in the OMG IDL type. For an unbounded sequence, this is the current maximum of the sequence.

```
void length(CORBA::ULong Length);
```

This function sets the current length of the sequence. The `Length` argument specifies the new length of the sequence. If the sequence is unbounded and the new length is greater than the current maximum, the buffer is reallocated and the elements are copied to the new buffer. If the new length is greater than the maximum, the maximum is set to the new length.

For a bounded sequence, the length cannot be set to a value greater than the maximum.

```
CORBA::ULong length() const;
```

This function returns the current length of the sequence.

```
TYPE & operator[](CORBA::ULong Index);
```

```
const TYPE & operator[](CORBA::ULong Index) const;
```

These accessor functions return a reference to the sequence element at the specified index. The `Index` argument specifies the index of the element to return. This index cannot be greater than the current sequence length. The length must have been set either using the `TYPE *` constructor or the `length(CORBA::ULong)` modifier. If `TYPE` is an object reference, `TypeCode` reference, or string, the return type will be a `ForSeq_var` class.

```
static TYPE * allocbuf(CORBA::ULong NumElems);
```

This static function allocates a buffer to be used with the `TYPE *` constructor. The `NumElems` argument specifies the number of elements in the buffer to allocate. If the buffer cannot be allocated, `NULL` is returned.

If this buffer is not passed to the `TYPE *` constructor with `release` set to `CORBA_TRUE`, it should be freed using the `freebuf` member function.

```
static void freebuf(TYPE * Value);
```

This static function frees a `TYPE *` sequence buffer allocated by the `allocbuf` function. The `Value` argument specifies the `TYPE *` buffer allocated by the `allocbuf` function. A 0 (zero) pointer is ignored.

Arrays

An array in OMG IDL is mapped to a C++ array definition. For example, consider the following OMG IDL definition:

```
// OMG IDL

module INVENT
{
    . . .
    typedef LogItem    LogArray[10];
};
```

This definition maps to C++ as follows:

```
// C++

module INVENT
{
    . . .
    typedef LogItem LogArray[10];
    typedef LogItem LogArray_slice;
    static LogArray_slice * LogArray_alloc(void);
    static void LogArray_free(LogArray_slice *data);

};
```

Array Slice

A slice of an array is an array with all the dimensions of the original array except the first demension. The member functions for the array-generated classes use a pointer to a slice to return pointers to an array. A typedef for each slice is generated.

For example, consider the following OMG IDL definition:

```
// OMG IDL
typedef LogItem          LogMultiArray[5][10];
```

This definition maps to C++ as follows:

```
// C++
typedef LogItem          LogMultiArray[5][10];
typedef LogItem          LogMultiArray_slice[10];
```

If you have a one-dimensional array, an array slice is just a type. For example, if you had a one-dimensional array of `long`, an array slice would result in a `CORBA::Long` data type.

Array Element Mapping

The type of the OMG IDL array is mapped to the C++ array element type in the same manner as structs. For more information, see the section “Member Mapping” on page 11-8.

Vars

A var class is generated for an array. For more information, see the section “Using var Classes” on page 11-50.

Out

An out class is generated for an array. For more information, see the section “Using out Classes” on page 11-56.

Allocation Member Functions

For each array, there are two static functions for array allocation and deallocation. For a given OMG IDL type `TYPE`, the allocation and deallocation routines are as follows:

```
static TYPE_slice * TYPE_alloc(void);
```

This function allocates a `TYPE` array, returning a pointer to the allocated `TYPE` array. If the array cannot be dynamically allocated, 0 (zero) is returned.

```
static void TYPE_free(TYPE_slice * Value);
```

This function frees a dynamically allocated `TYPE` array. The `Value` argument is a pointer to the dynamically allocated `TYPE` array to be freed.

Exceptions

An exception in OMG IDL is mapped to a C++ class. The C++ class contains the following:

- ◆ Constructors
- ◆ Destructors
- ◆ A static `_narrow` function, to determine the type of exception

The generated class is similar to a variable-length structure, but with an additional constructor to simplify initialization, and with the static `_narrow` member function to determine the type of `UserException`.

For example, consider the following OMG IDL definition:

```
// OMG IDL

module INVENT
{
    exception NonExist
    {
        ID BadId;
    };
};
```

This definition maps to C++ as follows:

```
// C++

class INVENT
{
    . . .

    class NonExist : public CORBA::UserException
    {
    public:
        static NonExist * _narrow(CORBA::Exception_ptr);
```



```

        NonExist (ID _BadId);
        NonExist ();
        NonExist (const NonExist &);
        ~NonExist ();
        NonExist & operator=(const NonExist &);
        void _raise ();
        ID BadId;
    };
};

```

Attributes (data members) of the Exception class are public, so you may access them directly.

Member Mapping

Members of an exception are mapped in the same manner as structs. For more information, see “Member Mapping” on page 11-8.

All exception members are public data in the C++ class, and are accessed directly.

Var

A var class is generated for an exception. For more information, see the section “Using var Classes” on page 11-50.

Out

An out class is generated for an exception. For more information, see the section “Using out Classes” on page 11-56.

Member Functions

For a given OMG IDL exception `TYPE`, the generated member functions are as follows:

```
static TYPE * _narrow(CORBA::Exception_ptr Except);
```

This function returns a pointer to a `TYPE` exception class if the exception can be narrowed to a `TYPE` exception. If the exception cannot be narrowed, 0 (zero) is returned. The `TYPE` pointer is not a pointer to a new class. Instead, it is a typed pointer to the original exception pointer and is valid only as long as the `Except` parameter is valid.

`TYPE ();`

This is the default constructor for the exception. No initialization of members is performed for fixed-length members. Variable-length members map to self-managing types; these types have constructors that initialize the member.

`TYPE(member-parameters);`

This constructor has an argument for each of the members in the exception. The constructor copies each argument and does not assume ownership of the memory for any argument. Building on the previous example, the signature of the constructor is:

`NonExist (ID _BadId);`

There is one argument for each member of the exception. The type and parameter-passing mechanism are identical to the Any insertion operator. For information about the Any insertion operator, see the section to “Insertion into Any” on page 11-39.

`TYPE (const TYPE & From);`

This copy constructor copies the data from the specified `TYPE` exception argument. The `From` argument specifies the exception to be copied.

`~TYPE ();`

This destructor frees the data associated with the exception.

`TYPE & operator=(const TYPE & From);`

This assignment operator copies the data from the specified `TYPE` exception argument. The `From` argument specifies the exception to be copied.

`void _raise ();`

This function causes the exception instance to throw itself. A catch clause can catch it by a more derived type.

Mapping of Pseudo-objects to C++

CORBA pseudo-objects may be implemented either as normal CORBA objects or as serverless objects. In the CORBA specification, the fundamental differences between these strategies are:

- ◆ Serverless object types do not inherit from `CORBA::Object`.
- ◆ Individual serverless objects are not registered with any ORB.

- ◆ Serverless objects do not necessarily follow the same memory management rules as for regular IDL types.

References to serverless objects are not necessarily valid across computational contexts; for example, address spaces. Instead, references to serverless objects that are passed as parameters may result in the construction of independent, functionally identical copies of objects used by receivers of these references. To support this, the otherwise hidden representational properties (such as data layout) of serverless objects are made known to the ORB. Specifications for achieving this are not contained in this chapter; making serverless objects known to the ORB is an implementation detail.

This chapter provides a standard mapping algorithm for all pseudo-object types. This avoids the need for piecemeal mappings for each of the nine CORBA pseudo-object types, and accommodates any pseudo-object types that may be proposed in future revisions of *CORBA*. It also avoids representation dependence in the C mapping, while still allowing implementations that rely on C-compatible representations.

Usage

Rather than C-PIDL, this mapping uses an augmented form of full OMG IDL to describe serverless object types. Interfaces for pseudo-object types follow the same rules as normal OMG IDL interfaces, with the following exceptions:

- ◆ They are prefaced by the keyword `pseudo`.
- ◆ Their declarations may refer to other¹ serverless object types that are not otherwise necessarily allowed in OMG IDL.

The `pseudo` prefix means that the interface may be implemented in either a normal or serverless fashion. That is, apply either the rules described in the following sections, or the normal mapping rules described in this chapter.

Mapping Rules

Serverless objects are mapped in the same way as normal interfaces, except for the differences outlined in this section.

1. In particular, `exception` used as a data type and a function name.

Classes representing serverless object types are *not* subclasses of `CORBA::Object`, and are not necessarily subclasses of any other C++ class. Thus, they do not necessarily support, for example, the `Object::create_request` operation.

For each class representing a serverless object type `T`, overloaded versions of the following functions are provided in the CORBA namespace:

```
// C++
void release(T_ptr);
Boolean is_nil(T_ptr p);
```

The mapped C++ classes are not guaranteed to be usefully subclassable by users, although subclasses can be provided by implementations. Implementations are allowed to make assumptions about internal representations and transport formats that may not apply to subclasses.

The member functions of classes representing serverless object types do not necessarily obey the normal memory management rules. This is because some serverless objects, such as `CORBA::NVList`, are essentially just containers for several levels of other serverless objects. Requiring callers to explicitly free the values returned from accessor functions for the contained serverless objects would be counter to their intended usage.

All other elements of the mapping are the same. In particular:

- ◆ The types of references to serverless objects, `T_ptr`, may or may not simply be a typedef of `T*`.
- ◆ Each mapped class supports the following static member functions:
- ◆

```
// C++
static T_ptr _duplicate(T_ptr p);
static T_ptr _nil();
```
- ◆ Legal implementations of `_duplicate` include simply returning the argument or constructing references to a new instance. Individual implementations may provide stronger guarantees about behavior.
- ◆ The corresponding C++ classes may or may not be directly instantiable or have other instantiation constraints. For portability, users should invoke the appropriate constructive operations.
- ◆ As with normal interfaces, assignment operators are not supported.
- ◆ Although they can transparently employ “copy-style” rather than “reference-style” mechanics, parameter passing signatures and rules as well as memory management rules are identical to those for normal objects, unless otherwise noted.

Relation to the C PIDL Mapping

All serverless object interfaces and declarations that rely on them have direct analogs in the C mapping. The mapped C++ classes can, but need not, be implemented using representations compatible to those chosen for the C mapping. Differences between the pseudo-object specifications for C-PIDL and C++ PIDL are as follows:

- ◆ C++ PIDL calls for removal of representation dependencies through the use of interfaces rather than structs and typedefs.
- ◆ C++ PIDL calls for placement of operations on pseudo-objects in their interfaces, including a few cases of redesignated functionality as noted.
- ◆ In C++ PIDL, `release` performs the role of the associated `free` and `delete` operations in the C mapping, unless otherwise noted.

Brief descriptions and listings of each pseudo-interface and its C++ mapping are provided in the following sections. Further details, including definitions of types referenced but not defined below, may be found in the relevant sections of this document.

Typedefs

A typedef in OMG IDL is mapped to a typedef in C++. Depending upon the OMG IDL data type, additional typedefs and member functions may be defined. The generated code for each data type is as follows:

- ◆ Basic data types (short, long, unsigned short, unsigned long, float, double, char, boolean, and octet)

Basic data types map to a simple typedef. For example:

```
// OMG IDL
typedef long ID;

// C++
typedef CORBA::Long ID;
```

- ◆ string

A string typedef is mapped to a simple typedef. For example:

```
// OMG IDL
typedef string IDStr;

// C++
typedef char * IDStr;
```

- ◆ object, interfaces, TypeCode

Object, interfaces, and TypeCode types are mapped to four typedefs. For example:

```
// OMG IDL
typedef Item Intf;

// C++
typedef Item Intf;
typedef Item_ptr Intf_ptr;
```

```
typedef Item_var Intf_var;
typedef Item_ptr & Intf _out;
```

◆ enum, struct, union, sequence

UDTs are mapped to three typedefs. For example:

```
// OMG IDL
typedef LogList ListRetType;

// C++
typedef LogList ListRetType;
typedef LogList_var ListRetType_var;
typedef LogList_out & ListRetType_out;
```

◆ array

Arrays are mapped to four typedefs and the static member functions to allocate and free memory. For example:

```
// OMG IDL
typedef LogArray ArrayRetType;

// C++
typedef LogArray ArrayRetType;
typedef LogArray_var ArrayRetType_var;
typedef LogArray_forany ArrayRetType_forany;
typedef LogArray_slice ArrayRetType_slice;
ArrayRetType_slice * ArrayRetType_alloc();
void ArrayRetType_free(ArrayRetType_slice *);
```

Implementing Interfaces

An operation in OMG IDL is mapped to a C++ member function.

The name of the member function is the name of the operation. The operation is defined as a member function in both the interface class and the stub class. The interface class is virtual; the stub class inherits from the virtual class and contains the member function code from the client application stub. When an operation is invoked on the object reference, the code contained in the corresponding stub member function executes.

For example, consider the following OMG IDL definition:

```
// OMG IDL

module INVENT
{
    interface Order
    {
        . . .
        ItemList modifyOrder (in ItemList ModifyList);
    };
};
```

This definition maps to C++ as follows:

```
// C++

class INVENT
{
    . . .

    class Order : public virtual CORBA::Object
    {
        . . .
        virtual ItemList * modifyOrder (
            const ItemList & ModifyList) = 0;
    };
};

class Stub_Order : public Order
{
    . . .
    ItemList * modifyOrder (
        const ItemList & ModifyList);
};
```

The generated client application stub then contains the following generated code for the stub class:

```
// ROUTINE NAME:      INVENT::Stub_Order::modifyOrder
//
// FUNCTIONAL DESCRIPTION:
//
// Client application stub routine for operation
```



```
// modifyOrder.
// (Interface : Order)

INVENT::ItemList * INVENT::Stub_Order::modifyOrder (
    const INVENT::ItemList & ModifyList)
{
    . . .
}
```

Argument Mapping

Each of the arguments in an operation is mapped to the corresponding C++ type as described in Table 11-1, “Basic OMG IDL and C++ Data Types,” on page 11-2 and Table 11-2, “Object, Pseudo-Object, and User-Defined OMG IDL and C++ Types,” on page 11-3.

The parameter passing modes for arguments in an operation are described in Table 11-7, “Basic Argument and Result Passing,” on page 11-65 and Table 11-8, “T_var Argument and Result Passing,” on page 11-66.

Implementing Operations

The signature of an implementation member function is the mapped signature of the OMG IDL operation. Unlike the client side, the server-side mapping requires that the function header include the appropriate exception (`throw`) specification. This requirement allows the compiler to detect when an invalid exception is raised, which is necessary in the case of a local C++-to-C++ library call (otherwise, the call would have to go through a wrapper that checks for a valid exception). For example:

```
// IDL
interface A
{
    exception B {};
    void f() raises(B);
};

// C++
class MyA : public virtual POA_A
{
    public:
```

```
void f() throw(A::B, CORBA::SystemException);
...
};
```

Since all operations and attributes may throw CORBA system exceptions, `CORBA::SystemException` must appear in all exception specifications, even when an operation has no `raises` clause.

Within a member function, the “this” pointer refers to the implementation object’s data as defined by the class. In addition to accessing the data, a member function may implicitly call another member function defined by the same class. For example:

```
// IDL
interface A
{
void f();
void g();
};

// C++
class MyA : public virtual POA_A
{
public:

void f() throw(SystemException);
void g() throw(SystemException);
private:
long x_;
};

void
MyA::f() throw(SystemException)
{
this->x_ = 3;
this->g();
}
```

However, when a servant member function is invoked in this manner, it is being called simply as a C++ member function, not as the implementation of an operation on a CORBA object. In such a context, any information available via the `POA_Current` object refers to the CORBA request invocation that performed the C++ member function invocation, not to the member function invocation itself.

Skeleton Derivation from Object

In several existing ORB implementations, each skeleton class derives from the corresponding interface class. For example, for interface `Mod::A`, the skeleton class `POA_Mod::A` is derived from class `Mod::A`. These systems, therefore, allow an object reference for a servant to be implicitly obtained via normal C++ derived-to-base conversion rules:

```
// C++
MyImplOfA my_a;    // declare impl of A
A_ptr a = &my_a;   // obtain its object reference
                  // by C++ derived-to-base conversion
```

Such code can be supported by a conforming ORB implementation, but it is not required, and is thus not portable. The equivalent portable code invokes `_this()` on the implementation object to implicitly register it if it has not yet been registered, and to get its object reference:

```
// C++
MyImplOfA my_a;    // declare impl of A
A_ptr a = my_a._this(); // obtain its object reference
```

PortableServer Functions

Objects registered with POAs use sequences of octet, specifically the `PortableServer::POA::ObjectId` type, as object identifiers. However, because C++ programmers often want to use strings as object identifiers, the C++ mapping provides several conversion functions that convert strings to `ObjectId` and vice versa:

```
// C++
namespace PortableServer
{
char* ObjectId_to_string(const ObjectId&);

ObjectId* string_to_ObjectId(const char*);
}
```

These functions follow the normal C++ mapping rules for parameter passing and memory management.

If conversion of an `ObjectId` to a string would result in illegal characters in the string (such as a NUL), the first two functions throw the `CORBA::BAD_PARAM` exception.

Modules

A module in OMG IDL is mapped to a C++ class. Objects contained in the module are defined within this C++ class. Because interfaces and types are also mapped to classes, nested C++ classes result.

For example, consider the following OMG IDL definition:

```
// OMG IDL

module INVENT
{
    interface Order
    {
        . . .
    };
};
```

This definition maps to C++ as follows:

```
// C++

class INVENT
{
    . . .
    class Order : public virtual CORBA::Object
    {
        . . .
    }; // class Order
}; // class INVENT
```

Multiple nested modules yield multiple nested classes. Anything inside the module will be in the module class. Anything inside the interface will be in the interface class.

OMG IDL allows modules, interfaces, and types to have the same name. However, when generating files for the C++ language, having the same name is not allowed. This restriction is necessary because the OMG IDL names are generated into nested C++ classes with the same name; this is not supported by C++ compilers.

Note: The WLE OMG IDL compiler outputs an informational message if you generate C++ code from OMG IDL with an interface or type with the same name as the current module name. If you ignore this informational message

and do not use unique names to differentiate the interface or type from the module name, the compiler will signal errors when compiling the generated files.

Interfaces

An interface in OMG IDL is mapped to a C++ class. This class contains the definitions of the operations, attributes, constants, and user-defined types (UDTs) contained in the OMG IDL interface.

For an interface *INTF*, the generated interface code contains the following items:

- ◆ Object reference type (*INTF_ptr*)
- ◆ Object reference variable type (*INTF_var*)
- ◆ `_duplicate` static member function
- ◆ `_narrow` static member function
- ◆ `_nil` static member function
- ◆ UDTs
- ◆ Member functions for attributes and operations

For example, consider the following OMG IDL definition:

```
// OMG IDL

module INVENT
{
    interface Order
    {
        void cancelOrder ();
    };
};
```

This definition maps to C++ as follows:

```
// C++
class INVENT
{
```

```
. . .
class    Order;
typedef Order *          Order_ptr;

class Order : public virtual CORBA::Object
{
    . . .
    static Order_ptr _duplicate(Order_ptr obj);
    static Order_ptr _narrow(CORBA::Object_ptr obj);
    static Order_ptr _nil();
    virtual void cancelOrder () = 0;
    . . .
};
};
```

The object reference types and static member functions are described in the following sections, as are UDTs, operations, and attributes.

Generated Static Member Functions

This section describes in detail the generated static member functions: `_duplicate`, `_narrow`, and `_nil` for an interface *INTF*.

```
static INTF_ptr _duplicate (INTF_ptr Obj)
```

This static member function duplicates an existing *INTF* object reference and returns a new *INTF* object reference. The new *INTF* object reference must be released by calling the `CORBA::release` member function. If an error occurs, a reference to the nil *INTF* object is returned. The argument `Obj` specifies the object reference to be duplicated.

```
static INTF_ptr _narrow (CORBA::Object_ptr Obj)
```

This static member function returns a new *INTF* object reference given an existing `CORBA::Object_ptr` object reference. The `Object_ptr` object reference may have been created by a call to the `CORBA::ORB::string_to_object` member function or may have been returned as a parameter from an operation.

The *INTF_ptr* object reference must correspond to an *INTF* object or to an object that inherits from the *INTF* object. The new *INTF* object reference must be released by calling the `CORBA::release` member function. The argument `Obj` specifies the object reference to be narrowed to an *INTF* object

reference. The `Obj` parameter is not modified by this member function and should be released by the user when it is no longer required. If `Obj` cannot be narrowed to an *INTF* object reference, the *INTF* nil object reference is returned.

```
static INTF_ptr _nil ( )
```

This static member function returns the new nil object reference for the *INTF* interface. The new reference does *not* have to be released by calling the `CORBA::release` member function.

Object Reference Types

An interface class (*INTF*) is a virtual class; the CORBA standard does not allow you to:

- ◆ Create or hold an instance of the interface class
- ◆ Use a pointer or a reference to the interface class

Instead, you use one of the object reference types, *INTF_ptr* or *INTF_var* class. You can obtain an object reference by using the `_narrow` static member function. Operations are invoked on these classes using the arrow operator (`->`).

The *INTF_var* class simplifies memory management by automatically releasing the object reference when the *INTF_var* class goes out of scope or is reassigned. Variable types are generated for many of the UDTs and are described in “Using var Classes” on page 11-50.

Attributes

A read-only attribute in OMG IDL is mapped to a C++ function that returns the attribute value. A read-write attribute maps to two overloaded C++ functions, one to return the attribute value and one to set the attribute value. The name of the overloaded member function is the name of the attribute.

Attributes are generated in the same way that operations are generated. They are defined in both the virtual and the stub classes. For example, consider the following OMG IDL definition:

```
// OMG IDL
```

```
module INVENT
{
    interface Order
    {
        . . .
        attribute itemStruct    itemInfo;
    };
};
```

This definition maps to C++ as follows:

```
// C++

class INVENT
{
    . . .

    class Item : public virtual CORBA::Object
    {
        . . .
        virtual itemStruct * itemInfo ( ) = 0;

        virtual void itemInfo (
            const itemStruct & itemInfo) = 0;
    };
};

class Stub_Item : public Item
{
    . . .
    itemStruct * itemInfo ();

    void itemInfo (
        const itemStruct & itemInfo);
};
```

The generated client application stub then contains the following generated code for the stub class:

```
// ROUTINE NAME:          INVENT::Stub_Item::itemInfo
//
// FUNCTIONAL DESCRIPTION:
//
```



```

//      Client application stub routine for attribute
//      INVENT::Stub_Item::itemInfo. (Interface : Item)

INVENT::itemStruct * INVENT::Stub_Item::itemInfo ( )
{
    . . .
}

//
// ROUTINE NAME:          INVENT::Stub_Item::itemInfo
//
// FUNCTIONAL DESCRIPTION:
//
//      Client application stub routine for attribute
//      INVENT::Stub_Item::itemInfo. (Interface : Item)

void INVENT::Stub_Item::itemInfo (
    const INVENT::itemStruct & itemInfo)
{
}

```

Argument Mapping

An attribute is equivalent to two operations, one to return the attribute and one to set the attribute. For example, the `itemInfo` attribute listed above is equivalent to:

```

void itemInfo (in itemStruct itemInfo);
itemStruct itemInfo ();

```

The argument mapping for the attribute is the same as the mapping for an operation argument. The attribute is mapped to the corresponding C++ type as described in Table 11-1, “Basic OMG IDL and C++ Data Types,” on page 11-2 and Table 11-2, “Object, Pseudo-Object, and User-Defined OMG IDL and C++ Types,” on page 11-3. The parameter passing modes for arguments in an operation are described in Table 11-7, “Basic Argument and Result Passing,” on page 11-65 and Table 11-8, “T_var Argument and Result Passing,” on page 11-66.

Any Type

An `any` in OMG IDL is mapped to the `CORBA::Any` class. The `CORBA::Any` class handles C++ types in a type-safe manner.

Handling Typed Values

To decrease the chances of creating an `any` with a mismatched `TypeCode` and value, the C++ function overloading facility is utilized. Specifically, for each distinct type in an OMG IDL specification, overloaded functions to insert and extract values of that type are provided. Overloaded operators are used for these functions to completely avoid any name space pollution. The nature of these functions, which are described in detail below, is that the appropriate `TypeCode` is implied by the C++ type of the value being inserted into or extracted from the `any`.

Since the type-safe `any` interface described below is based upon C++ function overloading, it requires C++ types generated from OMG IDL specifications to be distinct. However, there are special cases in which this requirement is not met:

- ◆ The `boolean`, `octet`, and `char` OMG IDL types are not required to map to distinct C++ types, which means that a separate means of distinguishing them from each other for the purpose of function overloading is necessary. The means of distinguishing these types from each other is described in “Distinguishing `boolean`, `octet`, `char`, and `Bounded Strings`” on page 11-44.
- ◆ Since all strings are mapped to `char*` regardless of whether they are bounded or unbounded, another means of creating or setting an `any` with a bounded string value is necessary. This is described in “Distinguishing `boolean`, `octet`, `char`, and `Bounded Strings`” on page 11-44.
- ◆ In C++, arrays within a function argument list decay into pointers to their first elements. This means that function overloading cannot be used to distinguish between arrays of different sizes. The means for creating or setting an `any` when dealing with arrays is described below and in “Arrays” on page 11-18.

Insertion into Any

To allow a value to be set in an any in a type-safe fashion, the following overloaded operator function is provided for each separate OMG IDL type T:

```
// C++
void operator<=<=(Any&, T);
```

This function signature suffices for the following types, which are usually passed by value:

- ◆ Short, UShort, Long, ULong, Float, Double
- ◆ enumerations
- ◆ unbounded strings (char* passed by value)
- ◆ object references (T_ptr)

For values of type T that are too large to be passed by value efficiently, two forms of the insertion function are provided:

```
// C++
void operator<=<=(Any&, const T&);           // copying form
void operator<=<=(Any&, T*);                 // non-copying form
```

Note that the copying form is largely equivalent to the first form shown, as far as the caller is concerned.

These “left-shift-assign” operators are used to insert a typed value into an any, as follows:

```
// C++
Long value = 42;
Any a;
a <=<= value;
```

In this case, the version of operator<=<= overloaded for type Long sets both the value and the TypeCode properly for the Any variable.

Setting a value in an any using operator<=<= means the following:

- ◆ For the copying version of operator<=<=, the lifetime of the value in the Any is independent of the lifetime of the value passed to operator<=<=. The implementation of the Any does not store its value as a reference or a pointer to the value passed to operator<=<=.

- ◆ For the noncopying version of `operator<=`, the inserted `T*` is consumed by the `Any`. The caller may not use the `T*` to access the pointed-to data after insertion because the `Any` assumes ownership of `T*`, and the `Any` may immediately copy the pointed-to data and destroy the original.
- ◆ With both the copying and noncopying versions of `operator<=`, any previous value held by the `Any` is properly deallocated. For example, if the `Any(TypeCode_ptr, void*, TRUE)` constructor (described in “Handling Untyped Values” on page 11-47) were called to create the `Any`, the `Any` is responsible for deallocating the memory pointed to by the `void*` before copying the new value.

Copying insertion of a string type causes the following function to be invoked:

```
// C++
void operator<=(Any&, const char*);
```

Since all string types are mapped to `char*`, this insertion function assumes that the value being inserted is an unbounded string. “Distinguishing boolean, octet, char, and Bounded Strings” on page 11-44 describes how bounded strings may be correctly inserted into an `Any`. Noncopying insertion of both bounded and unbounded strings can be achieved using the `Any::from_string` helper type described in “Distinguishing boolean, octet, char, and Bounded Strings” on page 11-44.

Type-safe insertion of arrays uses the `Array_forany` types described in “Arrays” on page 11-18. The ORB provides a version of `operator<=` overloaded for each `Array_forany` type. For example:

```
// IDL
typedef long LongArray[4][5];

// C++
typedef Long LongArray[4][5];
typedef Long LongArray_slice[5];
class LongArray_forany { ... };

void operator<=(Any &, const LongArray_forany &);
```

The `Array_forany` types are always passed to `operator<=` by reference to `const`. The `nocopy` flag in the `Array_forany` constructor is used to control whether the inserted value is copied (`nocopy == FALSE`) or consumed (`nocopy == TRUE`). Because the `nocopy` flag defaults to `FALSE`, copying insertion is the default.

Because of the type ambiguity between an array of `T` and a `T*`, it is highly recommended that portable code explicitly use the appropriate `Array_forany` type when inserting an array into an `Any`. For example:

```
// IDL
struct S { ... };
typedef S SA[5];

// C++
struct S { ... };
typedef S SA[5];
typedef S SA_slice;
class SA_forany { ... };

SA s;
// ...initialize s...
Any a;
a <<= s; // line 1
a <<= SA_forany(s); // line 2
```

Line 1 results in the invocation of the noncopying `operator<<=(Any&, S*)` due to the decay of the `SA` array type into a pointer to its first element, rather than the invocation of the copying `SA_forany` insertion operator. Line 2 explicitly constructs the `SA_forany` type and thus results in the desired insertion operator being invoked.

The noncopying version of `operator<<=` for object references takes the address of the `T_ptr` type, as follows:

```
// IDL
interface T { ... };

// C++
void operator<<=(Any&, T_ptr);           // copying
void operator<<=(Any&, T_ptr*);          // non-copying
```

The noncopying object reference insertion consumes the object reference pointed to by `T_ptr*`; therefore, after insertion the caller may not access the object referred to by `T_ptr` because the `Any` may have duplicated and then immediately released the original object reference. The caller maintains ownership of the storage for the `T_ptr` itself.

The copying version of `operator<<=` is also supported on the `Any_var` type.

Extraction from Any

To allow type-safe retrieval of a value from an `any`, the ORB provides the following operators for each OMG IDL type `T`:

```
// C++
Boolean operator>>=(const Any&, T&);
```

This function signature suffices for primitive types that are usually passed by value. For values of type `T` that are too large to be passed by value efficiently, the ORB provides a different signature, as follows:

```
// C++
Boolean operator>>=(const Any&, T*&);
```

The first form of this function is used only for the following types:

- ◆ `Boolean`, `Char`, `Octet`, `Short`, `UShort`, `Long`, `ULong`, `Float`, `Double`
- ◆ enumerations
- ◆ unbounded strings (`char*` passed by reference, i.e., `char*&`)
- ◆ object references (`T_ptr`)

For all other types, the second form of the function is used.

This “right-shift-assign” operator is used to extract a typed value from an `any`, as follows:

```
// C++
Long value;
Any a;
a <=& Long(42);
if (a >>= value) {
    // ... use the value ...
}
```

In this case, the version of `operator>>=` for type `Long` determines whether the `Any` truly does contain a value of type `Long` and, if so, copies its value into the reference variable provided by the caller and returns `TRUE`. If the `Any` does not contain a value of type `Long`, the value of the caller’s reference variable is not changed, and `operator>>=` returns `FALSE`.

For nonprimitive types, extraction is done by pointer. For example, consider the following OMG IDL struct:

```
// IDL
struct MyStruct {
    long lmem;
    short smem;
};
```

Such a struct could be extracted from an Any as follows:

```
// C++
Any a;
// ... a is somehow given a value of type MyStruct ...
MyStruct *struct_ptr;
if (a >=> struct_ptr) {
    // ... use the value ...
}
```

If the extraction is successful, the caller's pointer points to storage managed by the Any, and `operator>=>` returns `TRUE`. The caller must not try to delete or otherwise release this storage. The caller also should not use the storage after the contents of the Any variable are replaced via assignment, insertion, or the `replace` function, or after the Any variable is destroyed. Care must be taken to avoid using `T_var` types with these extraction operators, since they will try to assume responsibility for deleting the storage owned by the Any.

If the extraction is not successful, the value of the caller's pointer is set equal to the null pointer, and `operator>=>` returns `FALSE`.

Correct extraction of array types relies on the `Array_forany` types described in "Arrays" on page 11-18.

An example of the OMG IDL is as follows:

```
// IDL
typedef long A[20];
typedef A B[30][40][50];

// C++
typedef Long A[20];
typedef Long A_slice;
class A_forany { ... };
typedef A B[30][40][50];
typedef A B_slice[40][50];
class B_forany { ... };
```

```
Boolean operator>=>(const Any&, A_forany&);    // for type A
Boolean operator>=>(const Any&, B_forany&);    // for type B
```

The `Array_forany` types are always passed to `operator>=>` by reference.

For strings and arrays, applications are responsible for checking the `TypeCode` of the `Any` to be sure that they do not overstep the bounds of the array or string object when using the extracted value.

The `operator>=>` is also supported on the `Any_var` type.

Distinguishing boolean, octet, char, and Bounded Strings

Since the boolean, octet, and char OMG IDL types are not required to map to distinct C++ types, another means of distinguishing them from each other is necessary so that they can be used with the type-safe `Any` interface. Similarly, since both bounded and unbounded strings map to `char*`, another means of distinguishing them must be provided. This is done by introducing several new helper types nested in the `Any` class interface. For example, this is accomplished as shown below:

```
// C++
class Any
{
public:
    // special helper types needed for boolean, octet,
    // char, and bounded string insertion
    struct from_boolean {
        from_boolean(Boolean b) : val(b) {}
        Boolean val;
    };
    struct from_octet {
        from_octet(Octet o) : val(o) {}
        Octet val;
    };
    struct from_char {
        from_char(Char c) : val(c) {}
        Char val;
    };
    struct from_string {
        from_string(char* s, ULong b,
                    Boolean nocopy = FALSE) :
            val(s), bound(b) {}
        char *val;
        ULong bound;
    };

    void operator<=>= (from_boolean);
```

```

void operator<=<=(from_char);
void operator<=<=(from_octet);
void operator<=<=(from_string);

// special helper types needed for boolean, octet,
// char, and bounded string extraction
struct to_boolean {
    to_boolean(Boolean &b) : ref(b) {}
    Boolean &ref;
};
struct to_char {
    to_char(Char &c) : ref(c) {}
    Char &ref;
};
struct to_octet {
    to_octet(Octet &o) : ref(o) {}
    Octet &ref;
};
struct to_string {
    to_string(char *&s, ULong b) : val(s), bound(b) {}
    char *&val;
    ULong bound;
};

Boolean operator>=>=(to_boolean) const;
Boolean operator>=>=(to_char) const;
Boolean operator>=>=(to_octet) const;
Boolean operator>=>=(to_string) const;

// other public Any details omitted

private:
    // these functions are private and not implemented
    // hiding these causes compile-time errors for
    // unsigned char
    void operator<=<=(unsigned char);
    Boolean operator>=>=(unsigned char &) const;
};

```

The ORB provides the overloaded `operator<=<=` and `operator>=>=` functions for these special helper types. These helper types are used as shown here:

```
// C++
Boolean b = TRUE;
Any any;
any <= Any::from_boolean(b);
// ...
if (any >= Any::to_boolean(b)) {
    // ...any contained a Boolean...
}

char* p = "bounded";
any <= Any::from_string(p, 8);
// ...
if (any >= Any::to_string(p, 8)) {
    // ...any contained a string<8>...
}
```

A bound value of 0 (zero) indicates an unbounded string.

For noncopying insertion of a bounded or unbounded string into an Any, the `nocopy` flag on the `from_string` constructor should be set to `TRUE`:

```
// C++
char* p = string_alloc(8);
// ...initialize string p...
any <= Any::from_string(p, 8, 1);      // any consumes p
```

Assuming that `boolean`, `char`, and `octet` all map the C++ type `unsigned char`, the private and unimplemented `operator<=` and `operator>=` functions for `unsigned char` cause a compile-time error if straight insertion or extraction of any of the `Boolean`, `Char`, or `Octet` types is attempted:

```
// C++
Octet oct = 040;
Any any;
any <= oct;                                // this line will not compile
any <= Any::from_octet(oct);              // but this one will
```

Widening to Object

Sometimes it is desirable to extract an object reference from an `Any` as the base `Object` type. This can be accomplished using a helper type similar to those required for extracting `Boolean`, `Char`, and `Octet`:

```
// C++
class Any
{
    public:
        ...
        struct to_object {
            to_object(Object_ptr &obj) : ref(obj) {}
            Object_ptr &ref;
        };
        Boolean operator>>=(to_object) const;
        ...
};
```

The `to_object` helper type is used to extract an object reference from an `Any` as the base `Object` type. If the `Any` contains a value of an object reference type as indicated by its `TypeCode`, the extraction function `operator>>=(to_object)` explicitly widens its contained object reference to `Object` and returns `true`; otherwise, it returns `false`. This is the only object reference extraction function that performs widening on the extracted object reference. As with regular object reference extraction, no duplication of the object reference is performed by the `to_object` extraction operator.

Handling Untyped Values

Under some circumstances the type-safe interface to `Any` is not sufficient. An example is a situation in which data types are read from a file in binary form and are used to create values of type `Any`. For these cases, the `Any` class provides a constructor with an explicit `TypeCode` and generic pointer:

```
// C++
Any(TypeCode_ptr tc, void *value, Boolean release = FALSE);
```

The constructor duplicates the given `TypeCode` pseudo-object reference. If the `release` parameter is `TRUE`, the `Any` object assumes ownership of the storage pointed to by the `value` parameter. A caller should make no assumptions about the continued lifetime of the `value` parameter once it has been handed to an `Any` with `release=TRUE`, since the `Any` may copy the `value` parameter and immediately free the original pointer. If the `release` parameter is `FALSE` (the default case), the `Any` object assumes that the caller manages the memory pointed to by `value`. The `value` parameter can be a null pointer.

The `Any` class also defines three unsafe operations:

```
// C++
void replace(
    TypeCode_ptr,
    void *value,
    Boolean release = FALSE
);
TypeCode_ptr type() const;
const void *value() const;
```

The `replace` function is intended to be used with types that cannot be used with the type-safe insertion interface, and so is similar to the constructor described above. The existing `TypeCode` is released and value storage is deallocated, if necessary. The `TypeCode` function parameter is duplicated. If the `release` parameter is `TRUE`, the `Any` object assumes ownership for the storage pointed to by the `value` parameter. The `Any` should make no assumptions about the continued lifetime of the `value` parameter once it has been handed to the `Any::replace` function with `release=TRUE`, since the `Any` may copy the `value` parameter and immediately free the original pointer. If the `release` parameter is `FALSE` (the default case), the `Any` object assumes that the caller manages the memory occupied by the `value`. The `value` parameter of the `replace` function can be a null pointer.

Note that neither the constructor shown above nor the `replace` function is type-safe. In particular, no guarantees are made by the compiler at run time as to the consistency between the `TypeCode` and the actual type of the `void*` argument. The behavior of an ORB implementation when presented with an `Any` that is constructed with a mismatched `TypeCode` and `value` is not defined.

The `type` function returns a `TypeCode_ptr` pseudo-object reference to the `TypeCode` associated with the `Any`. Like all object reference return values, the caller must release the reference when it is no longer needed, or assign it to a `TypeCode_var` variable for automatic management.

The `value` function returns a pointer to the data stored in the `Any`. If the `Any` has no associated value, the `value` function returns a null pointer.

Any Constructors, Destructor, Assignment Operator

The default constructor creates an `Any` with a `TypeCode` of type `tk_null`, and no value. The copy constructor calls `_duplicate` on the `TypeCode_ptr` of its `Any` parameter and deep-copies the parameter's value. The assignment operator releases its own `TypeCode_ptr` and deallocates storage for the current value if necessary, then

duplicates the `TypeCode_ptr` of its `Any` parameter and deep-copies the parameter's value. The destructor calls `release` on the `TypeCode_ptr` and deallocates storage for the value, if necessary.

Other constructors are described in the section “Handling Untyped Values” on page 11-47.

The Any Class

The full definition of the `Any` class can be found in the section “Any Class Member Functions” on page 1-8.

Fixed-Length Versus Variable-Length User-Defined Types

The memory management rules and member function signatures for a user-defined type depend upon whether the type is fixed-length or variable-length. A user-defined type is variable-length if it is one of the following:

- ◆ A bounded or unbounded string
- ◆ A bounded or unbounded sequence
- ◆ A struct or union that contains a variable-length member
- ◆ An array with a variable-length element type
- ◆ A typedef to a variable-length type

If a type is not on this list, the type is fixed-length.

Using var Classes

Automatic variables (vars) are provided to simplify memory management. Vars are provided through a var class that assumes ownership for the memory required for the type and frees the memory when the instance of the var object is destroyed or when a new value is assigned to the var object.

The WLE provides var classes for the following types:

- ◆ string (CORBA::String_var)
- ◆ object references (CORBA::Object_var)
- ◆ user-defined OMG IDL types (struct, union, sequence, array, and interface)

The var classes have common member functions, but may support additional operators depending upon the OMG IDL type. For an OMG IDL type `TYPE`, the `TYPE_var` class contains constructors, destructors, assignment operators, and operators to access the underlying `TYPE` type. An example var class is as follows:

```
class TYPE_var
{
public:
    // constructors
    TYPE_var();
    TYPE_var(TYPE *);
    TYPE_var(const TYPE_var &);
    // destructor
    ~TYPE_var();

    // assignment operators
    TYPE_var &operator=(TYPE *);
    TYPE_var &operator=(const TYPE_var &);

    // accessor operators
    TYPE *operator->();
    TYPE *operator->() const;

    TYPE_var_ptr in() const;
    TYPE_var_ptr& inout();
    TYPE_var_ptr& out();
```

```
TYPE_var_ptr _retn();
operator const TYPE_ptr&() const;
operator TYPE_ptr&();
operator TYPE_ptr;
};
```

The details of the member functions are as follows:

`TYPE_var()`

This is the default constructor for the `TYPE_var` class. The constructor initializes to 0 (zero) the `TYPE *` owned by the var class. You may not invoke the `operator->` on a `TYPE_var` class unless a valid `TYPE *` has been assigned to it.

`TYPE_var(TYPE * Value);`

This constructor assumes ownership of the specified `TYPE *` parameter. When the `TYPE_var` is destroyed, the `TYPE` is released. The `Value` argument is a pointer to the `TYPE` to be owned by this var class. This pointer must not be 0 (zero).

`TYPE_var(const TYPE_var & From);`

This copy constructor allocates a new `TYPE` and makes a deep copy of the data contained in the `TYPE` owned by the `From` parameter. When the `TYPE_var` is destroyed, the copy of the `TYPE` is released or deleted. The `From` parameter specifies the var class that points to the `TYPE` to be copied.

`~TYPE_var();`

This destructor uses the appropriate mechanism to release the `TYPE` owned by the var class. For strings, this is the `CORBA::string_free` routine. For object references, this is the `CORBA::release` routine. For other types, this may be `delete` or a generated static routine used to free allocated memory.

`TYPE_var &operator=(TYPE * NewValue);`

This assignment operator assumes ownership of the `TYPE` pointed to by the `NewValue` parameter. If the `TYPE_var` currently owns a `TYPE`, it is released before assuming ownership of the `NewValue` parameter. The `NewValue` argument is a pointer to the `TYPE` to be owned by this var class. This pointer must not be 0 (zero).

`TYPE_var &operator=(const TYPE_var &From);`

This assignment operator allocates a new `TYPE` and makes a deep copy of the data contained in the `TYPE` owned by the `From` `TYPE_var` parameter. If `TYPE_var` currently owns a `TYPE`, it is released. When the `TYPE_var` is

destroyed, the copy of the `TYPE` is released. The `From` parameter specifies the var class that points to the data to be copied.

```
TYPE *operator->();  
TYPE *operator->() const;
```

These operators return a pointer to the `TYPE` owned by the var class. The var class continues to own the `TYPE` and it is the responsibility of the var class to release `TYPE`. You cannot use the `operator->` until the var owns a valid `TYPE`. Do not try to release this return value or access this return value after the `TYPE_var` has been destroyed.

```
TYPE_var_ptr in() const;  
TYPE_var_ptr& inout();  
TYPE_var_ptr& out();  
TYPE_var_ptr _retn();
```

Because implicit conversions can sometimes cause a problem with some C++ compilers and with code readability, the `TYPE_var` types also support member functions that allow them to be explicitly converted for purposes of parameter passing. To pass a `TYPE_var` and an `in` parameter, call the `in()` member function; for `inout` parameters, the `inout()` member function; for `out` parameters, the `out()` member function. To obtain a return value from the `TYPE_var`, call the `_return()` function. For each `TYPE_var` type, the return types of each of these functions will match the type shown in Table 11-7, “Basic Argument and Result Passing,” on page 11-65 for the `in`, `inout`, `out`, and `return` modes for the underlying type `TYPE`, respectively.

Some differences occur in the operators supported for the user-defined data types. Table 11-3 describes the various operators supported by each OMG IDL data type, in the generated C++ code. Because the assignment operators are supported for all of the data types described in Table 11-3, they are not included in the comparison.

Table 11-3 Comparison of Operators Supported for User-Defined Data Type Var Classes

OMG IDL Data Type	operator ->	operator[]
struct	Yes	No
union	Yes	No
sequence	Yes	Yes, non-const only
array	No	Yes

The signatures are as shown in Table 11-4.

Table 11-4 Operator Signatures for _var Classes

OMG IDL Data Type	Operator Member Functions
struct	TYPE * operator-> () TYPE * operator-> () const
union	TYPE * operator-> () TYPE * operator-> () const
sequence	TYPE * operator-> () TYPE * operator-> () const TYPE & operator[] (CORBA::Long index)
array	TYPE_slice & operator[] (CORBA::Long index) TYPE_slice & operator[] (CORBA::Long index) const

Sequence vars

Sequence vars support the following additional `operator[]` member function:

```
TYPE &operator[] (CORBA::ULong Index);
```

This operator invokes the `operator[]` of sequence owned by the var class. The `operator[]` returns a reference to the appropriate element of the sequence at the specified index. The `Index` argument specifies the index of the element to return. This index cannot be greater than the current sequence length.

Array vars

Array vars do not support `operator->`, but do support the following additional `operator[]` member functions to access the array elements:

```
TYPE_slice& operator[] (CORBA::ULong Index);  
const TYPE_slice & operator[] (CORBA::ULong Index) const;
```

These operators return a reference to the array slice at the specified index. An array slice is an array with all the dimensions of the original array except the first dimension. The member functions for the array-generated classes use a pointer to a slice to return pointers to an array. The `Index` argument specifies the index of the slice to return. This index cannot be greater than the array dimension.

String vars

The `String` vars in the member functions described in this section and in the section “Sequence vars” on page 11-53 have a `TYPE` of `char *`. `String` vars support additional member functions, as follows:

```
String_var(char * str)
```

This constructor makes a `String_var` from a string. The `str` argument specifies the string that will be assumed. The user must not use the `str` pointer to access data.

```
String_var(const char * str)
```

```
String_var(const String_var & var)
```

This constructor makes a `String_var` from a `const` string. The `str` argument specifies the `const` string that will be copied. The `var` argument specifies a reference to the string to be copied.

```
String_var & operator=(char * str)
```

This assignment operator first releases the contained string using `CORBA::string_free`, and then assumes ownership of the input string. The `str` argument specifies the string whose ownership will be assumed by this `String_var` object.

```
String_var & operator=(const char * str)
```

```
String_var & operator=(const String_var & var)
```

This assignment operator first releases the contained string using `CORBA::string_free`, and then copies the input string. The `Data` argument specifies the string whose ownership will be assumed by this `String_var` object.

```
char operator[] (Ulong Index)
char operator[] (Ulong Index) const
```

These array operators are superscripting operators that provide access to characters within the string. The `Index` argument specifies the index of the array to use in accessing a particular character within the array. Zero-based indexing is used. The returned value of the `Char operator[] (Ulong Index)` function can be used as an lvalue. The returned value of the `Char operator[] (Ulong Index) const` function cannot be used as an lvalue.

out Classes

Structured types (struct, union, sequence), arrays, and interfaces have a corresponding generated `_out` class. The `out` class is provided for simplifying the memory management of pointers to variable-length and fixed-length types. For more information about `out` classes and the common member functions, see the section “Using out Classes” on page 11-56.

Some differences occur in the operators supported for the user-defined data types. Table 11-5 describes the various operators supported by each OMG IDL data type, in the generated C++ code. Because the assignment operators are supported for all of the data types described in Table 11-3, they are not included in the comparison.

Table 11-5 Comparison of Operators Supported for User-Defined Data Type Out Classes

OMG IDL Data Type	operator ->	operator[]
struct	Yes	No
union	Yes	No
sequence	Yes	Yes, non-const only
array	No	Yes

The signatures are as shown in Table 11-6.

Table 11-6 Operator Signatures for _out Classes

OMG IDL Data Type	Operator Member Functions
struct	TYPE * operator-> () TYPE * operator-> () const
union	TYPE * operator-> () TYPE * operator-> () const
sequence	TYPE * operator-> () TYPE * operator-> () const TYPE & operator[](CORBA::Long index)
array	TYPE_slice & operator[](CORBA::Long index) TYPE_slice & operator[](CORBA::Long index) const

Using out Classes

When a TYPE_var is passed as an out parameter, any previous value it referred to must be implicitly deleted. To give the ORB enough hooks to meet this requirement, each T_var type has a corresponding TYPE_out type that is used solely as the out parameter type.

Note: The _out classes are not intended to be instantiated directly by the programmer. Specify an _out class only in function signatures.

The general form for TYPE_out types for variable-length types is as follows:

```
// C++
class TYPE_out
{
public:
    TYPE_out(TYPE*& p) : ptr_(p) { ptr_ = 0; }
    TYPE_out(TYPE_var& p) : ptr_(p.ptr_) { delete ptr_; ptr_ = 0;}
    TYPE_out(TYPE_out& p) : ptr_(p.ptr_) {}
```

```

    TYPE_out& operator=(TYPE_out& p) { ptr_ = p.ptr_;
                                   return *this;
    }
    Type_out& operator=(Type* p) { ptr_ = p; return *this; }

    operator Type*&() { return ptr_; }
    Type*& ptr() { return ptr_; }

    Type* operator->() { return ptr_; }

private:
    Type*& ptr_;

    // assignment from TYPE_var not allowed
    void operator=(const TYPE_var&):
};

```

The first constructor binds the reference data member with the `T*&` argument and sets the pointer to the zero (0) pointer value. The second constructor binds the reference data member with the pointer held by the `TYPE_var` argument, and then calls `delete` on the pointer (or `string_free()` in the case of the `String_out` type or `TYPE_free()` in the case of a `TYPE_var` for an array type `TYPE`). The third constructor, the copy constructor, binds the reference data member to the same pointer referenced by the data member of the constructor argument.

Assignment from another `TYPE_out` copies the `TYPE*` referenced by the `TYPE_out` argument to the data member. The overloaded assignment operator for `TYPE*` simply assigns the pointer argument to the data member. Note that assignment does not cause any previously held pointer to be deleted; in this regard, the `TYPE_out` type behaves exactly as a `TYPE*`. The `TYPE*&` conversion operator returns the data member. The `ptr()` member function, which can be used to avoid having to rely on implicit conversion, also returns the data member. The overloaded arrow operator (`operator->()`) allows access to members of the data structure pointed to by the `TYPE*` data member. Compliant applications may not call the overloaded `operator->()` unless the `TYPE_out` has been initialized with a valid non-null `TYPE*`.

Assignment to a `TYPE_out` from instances of the corresponding `TYPE_var` type is disallowed because there is no way to determine whether the application developer wants a copy to be performed, or whether the `TYPE_var` should yield ownership of its managed pointer so it can be assigned to the `TYPE_out`. To perform a copy of a `TYPE_var` to a `TYPE_out`, the application should use `new`, as follows:

```
// C++
TYPE_var t = ...;
my_out = new TYPE(t.in());           // heap-allocate a copy
```

The `in()` function called on `t` typically returns a `const TYPE&`, suitable for invoking the copy constructor of the newly allocated `T` instance.

Alternatively, to make the `TYPE_var` yield ownership of its managed pointer so it can be returned in a `T_out` parameter, the application should use the `TYPE_var::_retn()` function, as follows:

```
// C++
TYPE_var t = ...;
my_out = t._retn();                  // t yields ownership, no copy
```

Note that the `TYPE_out` types are not intended to serve as general-purpose data types to be created and destroyed by applications; they are used only as types within operation signatures to allow necessary memory management side-effects to occur properly.

Object Reference out Parameter

When a `_var` is passed as an `out` parameter, any previous value it refers to must be implicitly released. To give C++ mapping implementations enough hooks to meet this requirement, each object reference type results in the generation of an `_out` type that is used solely as the `out` parameter type. For example, interface `TYPE` results in the object reference type `TYPE_ptr`, the helper type `TYPE_var`, and the `out` parameter type `TYPE_out`. The general form for object reference `_out` types is as follows:

```
// C++
class TYPE_out
{
public:
    TYPE_out(TYPE_ptr& p) : ptr_(p) { ptr_ = TYPE::_nil(); }
    TYPE_out(TYPE_var& p) : ptr_(p.ptr_) {
        release(ptr_); ptr_ = TYPE::_nil();
    }
    TYPE_out(TYPE_out& a) : ptr_(a.ptr_) {}
    TYPE_out& operator=(TYPE_out& a) {
        ptr_ = a.ptr_; return *this;
    }
    TYPE_out& operator=(const TYPE_var& a) {
        ptr_ = TYPE::_duplicate(TYPE_ptr(a)); return *this;
    }
};
```

```

    }
    TYPE_out& operator=(TYPE_ptr p) { ptr_ = p; return *this; }
    operator TYPE_ptr&() { return ptr_; }
    TYPE_ptr& ptr() { return ptr_; }
    TYPE_ptr operator->() { return ptr_; }

private:
    TYPE_ptr& ptr_;
};

```

Sequence outs

Sequence outs support the following additional `operator[]` member function:

```
TYPE &operator[](CORBA::ULong Index);
```

This operator invokes the `operator[]` of the sequence owned by the out class. The `operator[]` returns a reference to the appropriate element of the sequence at the specified index. The `Index` argument specifies the index of the element to return. This index cannot be greater than the current sequence length.

Array outs

Array outs do not support `operator->`, but do support the following additional `operator[]` member functions to access the array elements:

```
TYPE_slice& operator[](CORBA::ULong Index);
const TYPE_slice & operator[](CORBA::ULong Index) const;
```

These operators return a reference to the array slice at the specified index. An array slice is an array with all the dimensions of the original array except the first dimension. The member functions for the array-generated classes use a pointer to a slice to return pointers to an array. The `Index` argument specifies the index of the slice to return. This index cannot be greater than the array dimension.

String outs

When a `String_var` is passed as an out parameter, any previous value it refers to must be implicitly freed. To give C++ mapping implementations enough hooks to meet this requirement, the string type also results in the generation of a `String_out` type in the CORBA namespace that is used solely as the string out parameter type. The general form for the `String_out` type is as follows:

```
// C++
class String_out
{
public:
    String_out(char*& p) : ptr_(p) { ptr_ = 0; }
    String_out(String_var& p) : ptr_(p.ptr_) {
        string_free(ptr_); ptr_ = 0;
    }
    String_out(String_out& s) : ptr_(s.ptr_) {}
    String_out& operator=(String_out& s) {
        ptr_ = s.ptr_; return *this;
    }
    String_out& operator=(char* p) {
        ptr_ = p; return *this;
    }
    String_out& operator=(const char* p) {
        ptr_ = string_dup(p); return *this;
    }
    operator char*&() { return ptr_; }
    char*& ptr() { return ptr_; }

private:
    char*& ptr_;

    // assignment from String_var disallowed
    void operator=(const String_var&);
};
```

The first constructor binds the reference data member with the `char*&` argument. The second constructor binds the reference data member with the `char*` held by the `String_var` argument, and then calls `string_free()` on the string. The third constructor, the copy constructor, binds the reference data member to the same `char*` bound to the data member of its argument.

Assignment from another `String_out` copies the `char*` referenced by the argument `String_out` to the `char*` referenced by the data member. The overloaded assignment operator for `char*` simply assigns the `char*` argument to the data member. The overloaded assignment operator for `const char*` duplicates the argument and assigns the result to the data member. Note that the assignment does not cause any previously held string to be freed; in this regard, the `String_out` type behaves exactly as a `char*`. The `char*&` conversion operator returns the data member. The `ptr()` member function, which can be used to avoid having to rely on implicit conversion, also returns the data member.

Assignment from `String_var` to a `String_out` is disallowed because of the memory management ambiguities involved. Specifically, it is not possible to determine whether the string owned by the `String_var` should be taken over by the `String_out` without copying, or if it should be copied. Disallowing assignment from `String_var` forces the application developer to make the choice explicitly, as follows:

```
// C++
void
A::op(String_out arg)
{
    String_var s = string_dup("some string");
    ...
    out = s;                // disallowed; either
    out = string_dup(s);    // 1: copy, or
    out = s._retn();        // 2: adopt
}
```

On the line marked with the comment “1,” the caller is explicitly copying the string held by the `String_var` and assigning the result to the `out` argument. Alternatively, the caller could use the technique shown on the line marked with the comment “2” to force the `String_var` to give up its ownership of the string it holds so that it may be returned in the `out` argument without incurring memory management errors.

Argument Passing Considerations

The mapping of parameter passing modes attempts to balance the need for both efficiency and simplicity. For primitive types, enumerations, and object references, the modes are straightforward, passing the type `P` for primitives and enumerations and the type `A_ptr` for an interface type `A`.

Aggregate types are complicated by the question of when and how parameter memory is allocated and deallocated. Mapping `in` parameters is straightforward because the parameter storage is caller-allocated and read-only. The mapping for `out` and `inout` parameters is more problematic. For variable-length types, the callee must allocate some if not all of the storage. For fixed-length types, such as a *Point* type represented as a struct containing three floating point members, caller allocation is preferable (to allow stack allocation).

To accommodate both kinds of allocation, avoid the potential confusion of split allocation, and eliminate confusion with respect to when copying occurs, the mapping is `T&` for a fixed-length aggregate `T` and `T*&` for a variable-length `T`. This approach has the unfortunate consequence that usage for structs depends on whether the struct is fixed- or variable-length; however, the mapping is consistently `T_var&` if the caller uses the managed type `T_var`.

The mapping for `out` and `inout` parameters additionally requires support for deallocating any previous variable-length data in the parameter when a `T_var` is passed. Even though their initial values are not sent to the operation, the WLE includes `out` parameters because the parameter could contain the result from a previous call. The provision of the `T_out` types is intended to give implementations the hooks necessary to free the inaccessible storage while converting from the `T_var` types. The following examples demonstrate the compliant behavior:

```
// IDL
struct S { string name; float age; };
void f(out S p);

// C++
S_var s;
f(s);
// use s
f(s);           // first result will be freed

S *sp;         // need not initialize before passing to out
f(sp);
// use sp
delete sp;     // cannot assume next call will free old value
f(sp);
```

Note that implicit deallocation of previous values for `out` and `inout` parameters works only with `T_var` types, not with other types:

```
// IDL
void q(out string s);
```

```
// C++
char *s;
for (int i = 0; i < 10; i++)
    q(s);           // memory leak!
```

Each call to the `q` function in the loop results in a memory leak because the caller is not invoking `string_free` on the `out` result. There are two ways to fix this, as shown below:

```
// C++
char *s;
String_var svar;
for (int i = 0 ; i < 10; i++) {
    q(s);
    string_free(s);    // explicit deallocation
    // OR:
    q(svar);           // implicit deallocation
}
```

Using a plain `char*` for the `out` parameter means that the caller must explicitly deallocate its memory before each reuse of the variable as an `out` parameter, while using a `String_var` means that any deallocation is performed implicitly upon each use of the variable as an `out` parameter.

Variable-length data must be explicitly released before being overwritten. For example, before assigning to an `inout` string parameter, the implementor of an operation may first delete the old character data. Similarly, an `inout` interface parameter should be released before being reassigned. One way to ensure that the parameter storage is released is to assign it to a local `T_var` variable with an automatic release, as in the following example:

```
// IDL
interface A;
void f(inout string s, inout A obj);

// C++
void Aimpl::f(char *&s, A_ptr &obj) {
    String_var s_tmp = s;
    s = /* new data */;
    A_var obj_tmp = obj;
    obj = /* new reference */
}
```

For parameters that are passed or returned as a pointer (T^*) or as a reference to a pointer ($T^*\&$), an application is not allowed to pass or return a null pointer; the result of doing so is undefined. In particular, a caller may not pass a null pointer under any of the following circumstances:

- ◆ in and inout string
- ◆ in and inout array (pointer to first element)

However, a caller may pass a reference to a pointer with a null value for out parameters, because the callee does not examine the value, but overwrites it. A callee may not return a null pointer under any of the following circumstances:

- ◆ out and return variable-length struct
- ◆ out and return variable-length union
- ◆ out and return string
- ◆ out and return sequence
- ◆ out and return variable-length array, return fixed-length array
- ◆ out and return any

Operation Parameters and Signatures

Table 11-7, “Basic Argument and Result Passing,” on page 11-65 displays the mapping for the basic OMG IDL parameter passing modes and return type according to the type being passed or returned. Table 11-8, “ T_var Argument and Result Passing,” on page 11-66 displays the same information for T_var types. Table 11-8 is merely for informational purposes; it is expected that operation signatures for both clients and servers will be written in terms of the parameter-passing modes shown in Table 11-7, with the exception that the T_out types will be used as the actual parameter types for all out parameters.

It is also expected that T_var types will support the necessary conversion operators to allow them to be passed directly. Callers should always pass instances of either T_var types or the base types shown in Table 11-7, and callees should treat their T_out parameters as if they were actually the corresponding underlying types shown in Table 11-7.

In Table 11-7, fixed-length arrays are the only case where the type of an `out` parameter differs from a return value, which is necessary because C++ does not allow a function to return an array. The mapping returns a pointer to a *slice* of the array, where a slice is an array with all the dimensions of the original array specified except the first dimension.

Table 11-7 Basic Argument and Result Passing

Data Type	In	Inout	Out	Return
short	Short	Short&	Short&	Short
long	Long	Long&	Long&	Long
unsigned short	UShort	UShort&	UShort&	UShort
unsigned long	ULong	ULong&	ULong&	ULong
float	Float	Float&	Float&	Float
double	Double	Double&	Double&	Double
boolean	Boolean	Boolean&	Boolean&	Boolean
char	Char	Char&	Char&	Char
octet	Octet	Octet&	Octet&	Octet
enum	enum	enum&	enum&	enum
object reference ptr (See Note below.)	objref_ptr	objref_ptr&	objref_ptr&	objref_ptr
struct, fixed	const struct&	struct&	struct&	struct
struct, variable	const struct&	struct&	struct*&	struct*
union, fixed	const union&	union&	union&	union
union, variable	const union&	union&	union*&	union*
string	const char*	char*&	char*&	char*
sequence	const sequence&	sequence&	sequence*&	sequence*
array, fixed	const array	array	array	array slice* (See Note below.)

Table 11-7 Basic Argument and Result Passing (Continued)

Data Type	In	Inout	Out	Return
array, variable	const array	array	array slice*&	array slice*
any	const any&	any&	any*&	any*

Note: The Object reference ptr data type includes pseudo-object references. The array slice return is an array with all the dimensions of the original array except the first dimension.

A caller is responsible for providing storage for all arguments passed as `in` arguments.

Table 11-8 T_var Argument and Result Passing

Data Type	In	Inout	Out	Return
object reference var (See Note below.)	const objref_var&	objref_var&	objref_var&	objref_var
struct_var	const struct_var&	struct_var&	struct_var&	struct_var
union_var	const union_var&	union_var&	union_var&	union_var
string_var	const string_var&	string_var&	string_var&	string_var
sequence_var	const sequence_var&	sequence_var&	sequence_var&	sequence_var
array_var	const array_var&	array_var&	array_var&	array_var
any_var	const any_var&	any_var&	any_var&	any_var

Note: The object reference var data type includes pseudo-object references

Table 11-9 and Table 11-10 describe the caller's responsibility for storage associated with `inout` and `out` parameters and for return results.

Table 11-9 Caller Argument Storage Responsibilities

Type	Inout Param	Out Param	Return Result
short	1	1	1

Table 11-9 Caller Argument Storage Responsibilities (Continued)

Type	Inout Param	Out Param	Return Result
long	1	1	1
unsigned short	1	1	1
unsigned long	1	1	1
float	1	1	1
double	1	1	1
boolean	1	1	1
char	1	1	1
octet	1	1	1
enum	1	1	1
object reference ptr	2	2	2
struct, fixed	1	1	1
struct, variable	1	3	3
union, fixed	1	1	1
union, variable	1	3	3
string	4	3	3
sequence	5	3	3
array, fixed	1	1	6
array, variable	1	6	6
any	5	3	3

Table 11-10 Argument Passing Cases

Case	
1	<p>Caller allocates all necessary storage, except that which may be encapsulated and managed within the parameter itself. For <code>inout</code> parameters, the caller provides the initial value, and the callee may change that value. For <code>out</code> parameters, the caller allocates the storage but need not initialize it, and the callee sets the value. Function returns are by value.</p>
2	<p>Caller allocates storage for the object reference. For <code>inout</code> parameters, the caller provides an initial value; if the callee wants to reassign the <code>inout</code> parameter, it will first call <code>CORBA::release</code> on the original input value. To continue to use an object reference passed in as an <code>inout</code>, the caller must first duplicate the reference. The caller is responsible for the release of all <code>out</code> and return object references. Release of all object references embedded in other structures is performed automatically by the structures themselves.</p>
3	<p>For <code>out</code> parameters, the caller allocates a pointer and passes it by reference to the callee. The callee sets the pointer to point to a valid instance of the parameter's type. For returns, the callee returns a similar pointer. The callee is not allowed to return a null pointer in either case.</p> <p>In both cases, the caller is responsible for releasing the returned storage. To maintain local/remote transparency, the caller must always release the returned storage, regardless of whether the callee is located in the same address space as the caller or is located in a different address space. Following the completion of a request, the caller is not allowed to modify any values in the returned storage—to do so, the caller must first copy the returned instance into a new instance, and modify the new instance.</p>
4	<p>For <code>inout</code> strings, the caller provides storage for both the input string and the <code>char*</code> pointing to it. Since the callee may deallocate the input string and reassign the <code>char*</code> to point to new storage to hold the output value, the caller should allocate the input string using <code>string_alloc()</code>. The size of the <code>out</code> string is, therefore, not limited by the size of the <code>in</code> string. The caller is responsible for deleting the storage for the <code>out</code> using <code>string_free()</code>. The callee is not allowed to return a null pointer for an <code>inout</code>, <code>out</code>, or return value.</p>
5	<p>For <code>inout</code> sequences and <code>anys</code>, assignment or modification of the sequence or <code>any</code> may cause deallocation of owned storage before any reallocation occurs, depending upon the state of the Boolean <code>release</code> parameter with which the sequence or <code>any</code> was constructed.</p>

Table 11-10 Argument Passing Cases (Continued)

Case	
6	<p>For out parameters, the caller allocates a pointer to an array slice, which has all the same dimensions of the original array except the first, and passes the pointer by reference to the callee. The callee sets the pointer to point to a valid instance of the array.</p> <p>For returns, the callee returns a similar pointer. The callee is not allowed to return a null pointer in either case. In both cases, the caller is responsible for releasing the returned storage.</p> <p>To maintain local/remote transparency, the caller must always release the returned storage, regardless of whether the callee is located in the same address space as the callee or is located in a different address space. Following completion of a request, the caller is not allowed to modify any values in the returned storage—to do so, the caller must first copy the returned array instance into a new array instance, and modify the new instance.</p>

12 CORBA API

This chapter describes the BEA WLE implementation of the CORBA core member functions in C++ and their extensions. It also describes pseudo-objects and their relationship to C++ classes. Pseudo-objects are object references that cannot be transmitted across the network. Pseudo-objects are similar to other objects; however, because the ORB owns them, they cannot be extended.

Note: Some of the information in this chapter is taken from Chapter 20 of the *Common Object Request Broker: Architecture and Specification*. Revision 2.2, February 1998, published by the Object Management Group (OMG). Used with permission by OMG.

Global Classes

The following WLE classes are global in scope:

- ◆ CORBA
- ◆ Tobj

These classes contain the predefined types, classes, and functions used in WLE development.

The CORBA class contains the classes, data types, and member functions essential to using an Object Request Broker (ORB) as defined by CORBA. The WLE extensions to CORBA are contained in the Tobj C++ class. The Tobj class contains data types, nested classes, and member functions that WLE provides as an extension to CORBA.

Using CORBA data types and member functions in the WLE product requires the `CORBA::` prefix. For example, a `Long` is a `CORBA::Long`. Likewise, to use `Tobj` nested classes and member functions in the WLE product, you need the `Tobj::` prefix. For example, `FactoryFinder` is `Tobj::FactoryFinder`.

Pseudo-objects

Pseudo-objects are represented as local classes, which reside in the `CORBA` class. A pseudo-object and its corresponding member functions are named using a nested class structure. For example, an `ORB` object is a `CORBA::ORB` and a `Current` object is a `CORBA::Current`.

Any Class Member Functions

This section describes the member functions of the `Any` class.

The mapping of these member functions to C++ is as follows:

```
class CORBA
{
    class Any
    {
    public:

        Any ();
        Any (const Any&);
        Any (TypeCode_ptr tc, void *value, Boolean release =
                                                    CORBA_ FALSE);

        ~Any ();
        Any & operator=(const Any&);

        void    operator<=(Short);
        void    operator<=(UShort);
        void    operator<=(Long);
        void    operator<=(ULong);
        void    operator<=(Float);
```

```

void      operator<=(Double);
void      operator<=(const Any&);
void      operator<=(const char*);
void      operator<=(Object_ptr);
void      operator<=(from_boolean);
void      operator<=(from_char);
void      operator<=(from_octet);
void      operator<=(from_string);
Boolean   operator>=(Short&) const;
Boolean   operator>=(UShort&) const;
Boolean   operator>=(Long&) const;
Boolean   operator>=(ULong&) const;
Boolean   operator>=(Float&) const;
Boolean   operator>=(Double&) const;
Boolean   operator>=(Any&) const;
Boolean   operator>=(char*)& const;
Boolean   operator>=(Object_ptr&) const;
Boolean   operator>=(to_boolean) const;
Boolean   operator>=(to_char) const;
Boolean   operator>=(to_octet) const;
Boolean   operator>=(to_object) const;
Boolean   operator>=(to_string) const;

TypeCode_ptr type()const;
void      replace(TypeCode_ptr, void *, Boolean);
void      replace(TypeCode_ptr, void *);
const void * value() const;

};
}; //CORBA

```

CORBA::Any::Any()

Synopsis	Constructs the <code>Any</code> object.
C++ Binding	<code>CORBA::Any::Any()</code>
Arguments	None.
Description	This is the default constructor for the <code>CORBA::Any</code> class. It creates an <code>Any</code> object with a <code>TypeCode</code> of type <code>tc_null</code> and a value of 0 (zero).
Return Values	None.

CORBA::Any::Any(const CORBA::Any & InitAny)

Synopsis	Constructs the Any object that is a copy of another Any object.
C++ Binding	<code>CORBA::Any::Any(const CORBA::Any & InitAny)</code>
Argument	<p><code>InitAny</code></p> <p>Refers to the <code>CORBA::Any</code> to copy.</p>
Description	<p>This is the copy constructor for the <code>CORBA::Any</code> class. This constructor duplicates the <code>TypeCode</code> reference of the Any that is passed in.</p> <p>The type of copying to be performed is determined by the <code>release</code> flag of the Any object to be copied. If <code>release</code> evaluates as <code>CORBA_TRUE</code>, the constructor deep-copies the parameter's value; if <code>release</code> evaluates as <code>CORBA_FALSE</code>, the constructor shallow-copies the parameter's value. Using a shallow copy gives you more control to optimize memory allocation, but the caller must ensure the Any does not use memory that has been freed.</p>
Return Values	None.

CORBA::Any::Any(TypeCode_ptr TC, void * Value, Boolean Release)

Synopsis	Creates the Any object using a TypeCode and a value.						
C++ Binding	<code>CORBA::Any::Any(TypeCode_ptr TC, void * Value, Boolean Release)</code>						
Arguments	<table><tr><td>TC</td><td>A pointer to a TypeCode pseudo-object reference, specifying the type to be created.</td></tr><tr><td>Value</td><td>A pointer to the data to be used to create the Any object. The data type of this argument must match the TypeCode specified.</td></tr><tr><td>Release</td><td>Determines whether the Any assumes ownership of the memory specified by the Value argument. If Release is CORBA_TRUE, the Any assumes ownership. If Release is CORBA_FALSE, the Any does not assume ownership; the data pointed to by the Value argument is not released upon assignment or destruction.</td></tr></table>	TC	A pointer to a TypeCode pseudo-object reference, specifying the type to be created.	Value	A pointer to the data to be used to create the Any object. The data type of this argument must match the TypeCode specified.	Release	Determines whether the Any assumes ownership of the memory specified by the Value argument. If Release is CORBA_TRUE, the Any assumes ownership. If Release is CORBA_FALSE, the Any does not assume ownership; the data pointed to by the Value argument is not released upon assignment or destruction.
TC	A pointer to a TypeCode pseudo-object reference, specifying the type to be created.						
Value	A pointer to the data to be used to create the Any object. The data type of this argument must match the TypeCode specified.						
Release	Determines whether the Any assumes ownership of the memory specified by the Value argument. If Release is CORBA_TRUE, the Any assumes ownership. If Release is CORBA_FALSE, the Any does not assume ownership; the data pointed to by the Value argument is not released upon assignment or destruction.						
Description	This constructor is used with the nontype-safe Any interface. It duplicates the specified TypeCode object reference and then inserts the data pointed to by value inside the Any object.						
Return Values	None.						

CORBA::Any::~~Any()

Synopsis	Destructor for the Any.
C++ Binding	<code>CORBA::Any::~~Any()</code>
Arguments	None.
Description	This destructor frees the memory that the <code>CORBA::Any</code> holds (if the <code>Release</code> flag is specified as <code>CORBA_TRUE</code>), and releases the <code>TypeCode</code> pseudo-object reference contained in the <code>Any</code> .
Return Values	None.

CORBA::Any & CORBA::Any::operator=(const CORBA::Any & InitAny)

Synopsis	Any assignment operator.
C++ Binding	<code>CORBA::Any & CORBA::Any::operator=(const CORBA::Any & InitAny)</code>
Arguments	<p><code>InitAny</code></p> <p>A reference to an <code>Any</code> to use in the assignment. The <code>Any</code> to use in the assignment determines whether the <code>Any</code> assumes ownership of the memory in <code>Value</code>. If <code>Release</code> is <code>CORBA_TRUE</code>, the <code>Any</code> assumes ownership and deep-copies the <code>InitAny</code> argument's value; if <code>Release</code> is <code>CORBA_FALSE</code>, the <code>Any</code> shallow-copies the <code>InitAny</code> argument's value.</p>
Description	<p>This is the assignment operator for the <code>Any</code> class. Memory management of this member function is determined by the current value of the <code>Release</code> flag. The current value of the <code>Release</code> flag determines whether the current memory is released before the assignment. If the current <code>Release</code> flag is <code>CORBA_TRUE</code>, the <code>Any</code> releases any value previously held; if the current <code>Release</code> flag is <code>CORBA_FALSE</code>, the <code>Any</code> does not release any value previously held.</p>
Return Values	<p>Returns the <code>Any</code>, which holds the copy of the <code>InitAny</code>.</p>

void CORBA::any::operator<<=()

Synopsis Type safe Any insertion operators.

C++ Binding `void CORBA::Any::operator<<=(CORBA::Short Value)`
 `void CORBA::Any::operator<<=(CORBA::UShort Value)`
 `void CORBA::Any::operator<<=(CORBA::Long Value)`
 `void CORBA::Any::operator<<=(CORBA::Ulong Value)`
 `void CORBA::Any::operator<<=(CORBA::Float Value)`
 `void CORBA::Any::operator<<=(CORBA::Double Value)`
 `void CORBA::Any::operator<<=(const CORBA::Any & Value)`
 `void CORBA::Any::operator<<=(const char * Value)`
 `void CORBA::Any::operator<<=(Object_ptr Value)`

Argument `Value`
 Type specific value to be inserted into the Any.

Description This insertion member function performs type-safe insertions. If the Any had a previous value, and the `release` flag is `CORBA_TRUE`, the memory is deallocated and the previous `TypeCode` object reference is freed. The new value is inserted into the Any by copying the value passed in using the `Value` parameter. The appropriate `TypeCode` reference is duplicated.

Return Values None.

CORBA::Boolean CORBA::Any::operator>=()

Synopsis Type safe `Any` extraction operators.

```
C++ Binding CORBA::Boolean CORBA::Any::operator>>=(
                CORBA::Short & Value) const
CORBA::Boolean CORBA::Any::operator>>=(
                CORBA::UShort & Value) const
CORBA::Boolean CORBA::Any::operator>>=(
                CORBA::Long & Value) const
CORBA::Boolean CORBA::Any::operator>>=(
                CORBA::ULong & Value) const
CORBA::Boolean CORBA::Any::operator>>=(
                CORBA::Float & Value) const
CORBA::Boolean CORBA::Any::operator>>=(
                CORBA::Double & Value) const
CORBA::Boolean CORBA::Any::operator>>=(CORBA::Any & Value) const
CORBA::Boolean CORBA::Any::operator>>=(char * & Value) const
CORBA::Boolean CORBA::Any::operator>>=(Object_ptr & Value) const
```

Argument	The <code>value</code> argument is a reference to the relevant object that receives the output of the value contained in the <code>Any</code> object.
----------	---

Description	<p>This extraction member function performs type-safe extractions. If the <code>Any</code> object contains the specified type, this member function assigns the pointer of the <code>Any</code> to the output reference value, <code>Value</code>, and <code>CORBA_TRUE</code> is returned. If the <code>Any</code> does not contain the appropriate type, <code>CORBA_FALSE</code> is returned. The caller must not attempt to release or delete the storage because it is owned and managed by the <code>Any</code> object. The <code>Value</code> argument is a reference to the relevant object that receives the output of the value contained in the <code>Any</code> object. If the <code>Any</code> object does not contain the appropriate type, the value remains unchanged.</p>
-------------	--

Return Values CORBA_TRUE if the Any contained a value of the specific type. CORBA_FALSE if the Any did not contain a value of the specific type.

CORBA::Any::operator<=<=()

Synopsis Type safe insertion operators for Any.

C++ Binding `void CORBA::Any::operator<=<=(from_boolean Value)`
 `void CORBA::Any::operator<=<=(from_char Value)`
 `void CORBA::Any::operator<=<=(from_octet Value)`
 `void CORBA::Any::operator<=<=(from_string Value)`

Argument Value
 A relevant object that contains the value to insert into the Any.

Description These insertion member functions perform a type-safe insertion of a CORBA::Boolean, a CORBA::Char, or a CORBA::Octet reference into an Any. If the Any had a previous value, and its Release flag is CORBA_TRUE, the memory is deallocated and the previous TypeCode object reference is freed. The new value is inserted into the Any object by copying the value passed in using the Value parameter. The appropriate TypeCode reference is duplicated.

Return Values None.

CORBA::Boolean CORBA::Any::operator>>=()

Synopsis	Type-safe extraction operators for Any.
C++ Binding	<pre> CORBA::Boolean CORBA::Any::operator>>=(to_boolean Value) const CORBA::Boolean CORBA::Any::operator>>=(to_char Value) const CORBA::Boolean CORBA::Any::operator>>=(to_octet Value) const CORBA::Boolean CORBA::Any::operator>>=(to_object Value) const CORBA::Boolean CORBA::Any::operator>>=(to_string Value) const </pre>
Argument	<p>Value</p> <p>A reference to the relevant object that receives the output of the value contained in the Any object. If the Any object does not contain the appropriate type, the value remains unchanged.</p>
Description	<p>These extraction member functions perform a type-safe extraction of a CORBA::Boolean, a CORBA::Char, a CORBA::Octet, a CORBA::Object, or a String reference from an Any. These member functions are helpers nested in the Any class. Their purpose is to distinguish extractions of the OMG IDL types: boolean, char, and octet (C++ does not require these to be distinct types).</p>
Return Values	<p>If the Any contains the specified type, this member function assigns the value in the Any object reference to the output variable, Value, and returns CORBA_TRUE. If the Any object does not contain the appropriate type, CORBA_FALSE is returned.</p>

CORBA::TypeCode_ptr CORBA::Any::type() const

Synopsis	TypeCode accessor for Any.
C++ Binding	<code>CORBA::TypeCode_ptr CORBA::Any::type();</code>
Arguments	None.
Description	This function returns the <code>TypeCode_ptr</code> pseudo-object reference of the <code>TypeCode</code> object associated with the <code>Any</code> . The <code>TypeCode_ptr</code> pseudo-object reference must be released by the <code>CORBA::release</code> member function or must be assigned to a <code>TypeCode_var</code> to be automatically released.
Return Values	<code>TypeCode_ptr</code> contained in the <code>Any</code> .

void CORBA::Any::replace()

Synopsis Non-type safe Any “insertion.”

C++ Binding `void CORBA::Any::replace(TypeCode_ptr TC, void * Value,
 Boolean Release = CORBA_FALSE);`

Arguments TC A TypeCode pseudo-object reference specifying the TypeCode value for the replaced Any object. This argument is duplicated.

Value A void pointer specifying the storage pointed to by the Any object.

Release Determines whether the Any manages the specified Value argument. If Release is CORBA_TRUE, the Any assumes ownership. If Release is CORBA_FALSE, the Any does not assume ownership and the data pointed to by the Value parameter is not released upon assignment or destruction.

Description These member functions replace the data and TypeCode value currently contained in the Any with the value of the TC and Value arguments passed in. The functions perform a nontype-safe replacement, which means that the caller is responsible for consistency between the TypeCode value and the data type of the storage pointed to by the Value argument.

If the value of Release is CORBA_TRUE, this function releases the existing TypeCode pseudo-object in the Any object and frees the storage pointed to by the Any object reference.

Return Values None.

Context Member Functions

A Context supplies optional context information associated with a method invocation.

The mapping of these member functions to C++ is as follows:

```
class CORBA
{
    class Context
    {
    public:
        const char *context_name() const;
        Context_ptr parent() const;

        void create_child(const char *, Context_out);

        void set_one_value(const char *, const Any &);
        void set_values(NVList_ptr);
        void delete_values(const char *);
        void get_values(
            const char *,
            Flags,
            const char *,
            NVList_out
        );
    }; // Context
} // CORBA
```

Memory Management

Context has the following special memory management rule:

- ◆ Ownership of the return values of the `context_name` and `parent` functions is maintained by the Context; these return values must not be freed by the caller.

This section describes Context member functions.

CORBA::Context::context_name

Synopsis	Returns the name of a given Context object.
C++ Binding	<code>Const char * CORBA::Context::context_name () const;</code>
Arguments	None.
Description	This member function returns the name of a given Context object. The Context object reference owns the memory for the returned <code>char *</code> . Users should not modify this memory.
Return Values	<p>If the member function succeeds, it returns the name of the Context object. The value may be empty if the Context object is not a child Context created by a call to <code>CORBA::Context::create_child</code>.</p> <p>If the Context object has no name, this is an empty string.</p>

CORBA::Context::create_child

Synopsis Creates a child of the Context object.

C++ Binding

```
void CORBA::Context::create_child (
    const char *          CtxName,
    CORBA::Context_out    CtxObject);
```

Arguments

CtxName
The name to be associated with the child of the Context reference.

CtxObject
The newly created Context object reference.

Description This member function creates a child of the Context object. That is, searches on the child Context object will look for matching property names in the parent context (and so on, up the context tree), if necessary.

Return Values None.

Exception CORBA::NO_MEMORY

See Also CORBA::ORB::get_default_context
CORBA::release

CORBA::Context::delete_values

Synopsis	Deletes the values for a specified attribute in the Context object.
C++ Binding	<pre>void CORBA::Context::delete_values (const char * AttrName);</pre>
Argument	<p>AttrName</p> <p>The name of the attribute whose values are to be deleted. If this argument has a trailing wildcard character (*), all names that match the string preceding the wildcard character are deleted.</p>
Description	This member function deletes named values for an attribute in the Context object. Note that it does not do recursively do the same to its parents, if any.
Return Values	None.
Exceptions	<p>CORBA::BAD_PARAM if attribute is an empty string.</p> <p>CORBA::BAD_CONTEXT if no matching attributes to be deleted were found.</p>
See Also	<p>CORBA::Context::create_child</p> <p>CORBA::ORB::get_default_context</p>

CORBA::Context::get_values

Synopsis Retrieves the values for a given attribute in the Context object within the specified scope.

C++ Binding

```
void CORBA::Context::get_values (
    const char *          StartScope,
    CORBA::Flags          OpFlags,
    const char *          AttrName,
    CORBA::NVList_out     AttrValues);
```

Arguments **StartScope**
The Context object level at which to initiate the search for specified properties. The level is the name of the context, or `parent`, at which the search is started. If the value is 0 (zero), the search begins with the current Context object.

OpFlags
The only valid operation flag is `CORBA::CTX_RESTRICT_SCOPE`. If you specify this flag, the object implementation restricts the property search to the current scope only (that is, the property search is not executed recursively up the chain of the parent context); otherwise, the search continues to a wider scope until a match has been found or until all wider levels have been searched.

AttrName
The name of the attribute whose values are to be returned. If this argument has a trailing wildcard character (*), all names that match the string preceding the wildcard character are returned.

AttrValues
Receives the values for the specified attributes (returns an `NVList` object) where each item in the list is a `NamedValue`.

Description This member function retrieves the values for a specified attribute in the Context object. These values are returned as an `NVList` object, which must be freed when no longer needed using the `CORBA::release` member function.

Return Values None.

12 *CORBA API*

Exceptions `CORBA::BAD_PARAM` if attribute is an empty string.
 `CORBA::BAD_CONTEXT` if no matching attributes were found.
 `CORBA::NO_MEMORY` if dynamic memory allocation failed.

See Also `CORBA::Context::create_child`
 `CORBA::ORB::get_default_context`

CORBA::Context::parent

Synopsis	Returns the parent context of the Context object.
C++ Binding	<code>CORBA::Context_ptr CORBA::Context::parent () const;</code>
Arguments	None.
Description	This member function returns the parent context of the Context object. The parent of the Context object is an attribute owned by the Context and should not be modified or freed by the caller. This parent is nil unless the Context object was created using the <code>CORBA::Context::create_child</code> member function.
Return Values	<p>If the member function succeeds, the parent context of the Context object is returned. The parent context may be nil. Use the <code>CORBA::is_nil</code> member function to test for a nil object reference.</p> <p>If the member function does not succeed, an exception is thrown. Use the <code>CORBA::is_nil</code> member function to test for a nil object reference.</p>

CORBA::Context::set_one_value

Synopsis Sets the value for a given attribute in the Context object.

C++ Binding

```
void CORBA::Context::set_one_value (
    const char *          AttrName,
    const CORBA::Any &    AttrValue);
```

Arguments **AttrName**
 The name of the attribute to set.

AttrValue
 The value of the attribute. Currently, the WLE system supports only the string type; therefore, this parameter must contain a `CORBA::Any` object with a string inside.

Description This member function sets the value for a given attribute in the Context object. Currently, only string values are supported by the Context object. If the Context object already has an attribute with the given name, it is deleted first.

Return Values None.

Exceptions `CORBA::BAD_PARAM` if `AttrName` is an empty string or `AttrValue` does not contain a string type.
 `CORBA::NO_MEMORY` if dynamic memory allocation failed.

See Also `CORBA::Context::get_values`
 `CORBA::Context::set_values`

CORBA::Context::set_values

Synopsis	Sets the values for given attributes in the Context object.
C++ Binding	<pre>void CORBA::Context::set_values (CORBA::NVList_ptr AttrValue);</pre>
Argument	<p>AttrValues</p> <p>The name and value of the attribute. Currently the WLE system supports only the string type; therefore, all NamedValue objects in the list must have CORBA::Any objects with a string inside.</p>
Description	This member function sets the values for given attributes in the Context object. The CORBA::NVList member function contains the property name and value pairs to be set.
Return Values	None.
Exceptions	CORBA::BAD_PARAM if any of the attribute values has a value that is not a string type. CORBA::NO_MEMORY if dynamic memory allocation failed.
See Also	CORBA::Context::get_values CORBA::Context::set_one_value

ContextList Member Functions

The ContextList allows a client or server application to provide a list of context strings that must be supplied with Request invocation. For a description of the Request member functions, see the section “Request Member Functions” on page 12-112.

The ContextList differs from the Context in that the former supplies only the context strings whose values are to be looked up and sent with the request invocation (if applicable), while the latter is where those values are obtained. For a description of the Context member functions, see the section “Context Member Functions” on page 12-15.

The mapping of these member functions to C++ is as follows:

```
class CORBA
{
    class ContextList
    {
    public:
        ULong count ();
        void add(const char* ctxt);
        void add_consume(char* ctxt);
        const char* item(ULong index);
        Status remove(ULong index);
    }; // ContextList
} // CORBA
```

CORBA::ContextList::count

Synopsis	Retrieves the current number of items in the list.
C++ Binding	<code>Ulong count ();</code>
Arguments	None.
Description	This member function retrieves the current number of items in the list.
Return Values	If the function succeeds, the returned value is the number of items in the list. If the list has just been created, and no ContextList objects have been added, this function returns 0 (zero).
Exception	If the function does not succeed, an exception is thrown.
See Also	<code>CORBA::ContextList::add</code> <code>CORBA::ContextList::add_consume</code> <code>CORBA::ContextList::item</code> <code>CORBA::ContextList::remove</code>

CORBA::ContextList::add

Synopsis	Constructs a ContextList object with an unnamed item, setting only the <code>flags</code> attribute.
C++ Binding	<code>void add(const char* ctxt);</code>
Argument	<code>ctxt</code> Defines the memory location referred to by <code>char*</code> .
Description	<p>This member function constructs a ContextList object with an unnamed item, setting only the <code>flags</code> attribute.</p> <p>The ContextList object grows dynamically; your application does not need to track its size.</p>
Return Values	If the function succeeds, the return value is a pointer to the newly created ContextList object.
Exception	If the member function does not succeed, a <code>CORBA::NO_MEMORY</code> exception is thrown.
See Also	<code>CORBA::ContextList::add_consume</code> <code>CORBA::ContextList::count</code> <code>CORBA::ContextList::item</code> <code>CORBA::ContextList::remove</code>

CORBA::ContextList::add_consume

Synopsis	Constructs a ContextList object.
C++ Binding	<code>void add_consume(const char* ctxt);</code>
Argument	<code>ctxt</code> Defines the memory location referred to by <code>char*</code> .
Description	<p>This member function constructs a ContextList object.</p> <p>The ContextList object grows dynamically; your application does not need to track its size.</p>
Return Values	If the function succeeds, the return value is a pointer to the newly created ContextList object.
Exception	If the member function does not succeed, an exception is raised.
See Also	<code>CORBA::ContextList::add</code> <code>CORBA::ContextList::count</code> <code>CORBA::ContextList::item</code> <code>CORBA::ContextList::remove</code>

CORBA::ContextList::item

Synopsis	Retrieves a pointer to the ContextList object, based on the index passed in.
C++ Binding	<code>const char* item(ULong index);</code>
Argument	<code>index</code> The index into the ContextList object. The indexing is zero-based.
Description	This member function retrieves a pointer to a ContextList object, based on the index passed in. The function uses zero-based indexing.
Return Values	If the function succeeds, the return value is a pointer to the ContextList object.
Exceptions	If this function does not succeed, the <code>BAD_PARAM</code> exception is thrown.
See Also	<code>CORBA::ContextList::add</code> <code>CORBA::ContextList::add_consume</code> <code>CORBA::ContextList::count</code> <code>CORBA::ContextList::remove</code>

CORBA::ContextList::remove

Synopsis	Removes the item at the specified index, frees any associated memory, and reorders the remaining items on the list.
C++ Binding	<code>Status remove(ULong index);</code>
Argument	Index The index into the ContextList object. The indexing is zero-based.
Description	This member function removes the item at the specified index, frees any associated memory, and reorders the remaining items on the list.
Return Values	None.
Exceptions	If this function does not succeed, the <code>BAD_PARAM</code> exception is thrown.
See Also	<code>CORBA::ContextList::add</code> <code>CORBA::ContextList::add_consume</code> <code>CORBA::ContextList::count</code> <code>CORBA::ContextList::item</code>

NamedValue Member Functions

NamedValue is used only as an element of NVList, especially in the DII. NamedValue maintains an (optional) name, an any value, and labelling flags. Legal flag values are CORBA::ARG_IN, CORBA::ARG_OUT, and CORBA::ARG_INOUT.

The value in a NamedValue may be manipulated via standard operations on any.

The mapping of these member functions to C++ is as follows:

```
// C++
class NamedValue
{
public:
    Flags          flags() const;
    const char *   name() const;
    Any *          value() const;
};
```

Memory Management

NamedValue has the following special memory management rule:

- ◆ Ownership of the return values of the name() and value() functions is maintained by the NamedValue; these return values must not be freed by the caller.

The following sections describe NamedValue member functions.

CORBA::NamedValue::flags

Synopsis	Retrieves the flags attribute of the NamedValue object.
C++ Binding	<code>CORBA::Flags CORBA::NamedValue::flags () const;</code>
Arguments	None.
Description	This member function retrieves the flags attribute of the NamedValue object.
Return Values	<p>If the function succeeds, the return value is the flags attribute of the NamedValue object.</p> <p>If the function does not succeed, an exception is thrown.</p>

CORBA::NamedValue::name

Synopsis Retrieves the name attribute of the NamedValue object.

C++ Binding `const char * CORBA::NamedValue::name () const;`

Arguments None.

Description This member function retrieves the name attribute of the NamedValue object. The name returned by this member function is owned by the NamedValue object and should not be modified or released.

Return Values If the function succeeds, the value returned is a constant Identifier object representing the name attribute of the NamedValue object.

 If the function does not succeed, an exception is thrown.

CORBA::NamedValue::value

Synopsis	Retrieves a pointer to the value attribute of the NamedValue object.
C++ Binding	<code>CORBA::Any * CORBA::NamedValue::value () const;</code>
Arguments	None.
Description	This member function retrieves a pointer to the <i>Any</i> object that represents the value attribute of the NamedValue object. This attribute is owned by the NamedValue object, and should not be modified or released.
Return Values	<p>If the function succeeds, the return value is a pointer to the <i>Any</i> object contained in the NamedValue object.</p> <p>If the function does not succeed, an exception is thrown.</p>

NVList Member Functions

NVList is a list of NamedValues. A new NVList is constructed using the `ORB::create_list` operation (see “CORBA::ORB::create_list” on page 12-61). New NamedValues may be constructed as part of an NVList, in any of following ways:

- ◆ `add`—creates an unnamed value, initializing only the flags
- ◆ `add_item`—initializes name and flags
- ◆ `add_value`—initializes name, value, and flags

Each of these operations returns the new item.

Elements may be accessed and deleted via zero-based indexing. The `add`, `add_item`, `add_value`, `add_item_consume`, and `add_value_consume` functions lengthen the NVList to hold the new element each time they are called. The `item` function can be used to access existing elements.

```
// C++
class NVList
{
public:
    ULong count() const;
    NamedValue_ptr add(Flags);
    NamedValue_ptr add_item(const char*, Flags);
    NamedValue_ptr add_value(const char*, const Any&, Flags);
    NamedValue_ptr item(ULong);
    void remove(ULong);
};
```

Memory Management

NVList has the following special memory management rules:

- ◆ Ownership of the return values of the `add`, `add_item`, `add_value`, `add_item_consume`, `add_value_consume`, and `item` functions is maintained by the NVList; these return values must not be freed by the caller.

- ◆ The `char*` parameters to the `add_item_consume` and `add_value_consume` functions and the `Any*` parameter to the `add_value_consume` function are consumed by the `NVList`. The caller may not access these data after they have been passed to these functions because the `NVList` may copy them and destroy the originals immediately. The caller should use the `NamedValue::value()` operation to modify the `value` attribute of the underlying `NamedValue`, if desired.
- ◆ The `remove` function also calls `CORBA::release` on the removed `NamedValue`.

The following sections describe `NVList` member functions.

CORBA::NVList::add

Synopsis Constructs a NamedValue object with an unnamed item, setting only the `flags` attribute.

C++ Binding

```
CORBA::NamedValue_ptr CORBA::NVList::add (  
CORBA::Flags Flags);
```

Argument `Flags`
Flags to determine argument passing. Valid values are:

```
CORBA::ARG_IN  
CORBA::ARG_INOUT  
CORBA::ARG_OUT
```

Description This member function constructs a NamedValue object with an unnamed item, setting only the `flags` attribute. The NamedValue object is added to the NVList object that the call was invoked upon.

The NVList object grows dynamically; your application does not need to track its size.

Return Values If the function succeeds, the return value is a pointer to the newly created NamedValue object. The returned NamedValue object reference is owned by the NVList and should not be released.

If the member function does not succeed, a `CORBA::NO_MEMORY` exception is thrown.

See Also

```
CORBA::NVList::add  
CORBA::NVList::add_item  
CORBA::NVList::add_value  
CORBA::NVList::count  
CORBA::NVList::remove
```

CORBA::NVList::add_item

Synopsis Constructs a NamedValue object, creating an empty value attribute and initializing the name and flags attributes.

C++ Binding

```
CORBA::NamedValue_ptr CORBA::NVList::add_item (
    const char *      Name,
    CORBA::Flags      Flags);
```

Arguments

Name

The name of the list item.

Flags

Flags to determine argument passing. Valid values are:

```
CORBA::ARG_IN
CORBA::ARG_INOUT
CORBA::ARG_OUT
```

Description This member function constructs a NamedValue object, creating an empty value attribute and initializing the name and flags attributes that pass in as parameters. The NamedValue object is added to the NVList object that the call was invoked upon.

The NVList object grows dynamically; your application does not need to track its size.

Return Values If the function succeeds, the return value is a pointer to the newly created NamedValue object. The returned NamedValue object reference is owned by the NVList and should not be released.

If the member function does not succeed, an exception is thrown.

See Also

```
CORBA::NVList::add
CORBA::NVList::add_value
CORBA::NVList::count
CORBA::NVList::item
CORBA::NVList::remove
```

CORBA::NVList::add_value

Synopsis Constructs a NamedValue object, initializing the name, value, and flags attribute.

C++ Binding

```
CORBA::NamedValue_ptr CORBA::NVList::add_value (  
    const char *           Name,  
    const CORBA::Any &     Value,  
    CORBA::Flags           Flags);
```

Arguments

Name The name of the list item.

Value The value of the list item.

Flags Flags to determine argument passing. Valid values are:

```
CORBA::ARG_IN  
CORBA::ARG_INOUT  
CORBA::ARG_OUT
```

Description This member function constructs a NamedValue object, initializing the name, value, and flags attributes. The NamedValue object is added to the NVList object that the call was invoked upon.

The NVList object grows dynamically; your application does not need to track its size.

Return Values If the function succeeds, the return value is a pointer to the newly created NamedValue object. The returned NamedValue object reference is owned by the NVList and should not be released.

If the member function does not succeed, an exception is raised.

See Also

```
CORBA::NVList::add  
CORBA::NVList::add_item  
CORBA::NVList::count  
CORBA::NVList::item  
CORBA::NVList::remove
```


CORBA::NVList::count

Synopsis	Retrieves the current number of items in the list.
C++ Binding	<code>CORBA::ULong CORBA::NVList::count () const;</code>
Arguments	None.
Description	This member function retrieves the current number of items in the list.
Return Values	<p>If the function succeeds, the returned value is the number of items in the list. If the list has just been created, and no NamedValue objects have been added, this function returns 0 (zero).</p> <p>If the function does not succeed, an exception is thrown.</p>
See Also	<p> CORBA::NVList::add CORBA::NVList::add_item CORBA::NVList::add_value CORBA::NVList::item CORBA::NVList::remove </p>

CORBA::NVList::item

Synopsis	Retrieves a pointer to the NamedValue object, based on the index passed in.
C++ Binding	<pre>CORBA::NamedValue_ptr CORBA::NVList::item (CORBA::ULong Index);</pre>
Argument	<p>Index</p> <p>The index into the NVList object. The indexing is zero-based.</p>
Description	This member function retrieves a pointer to a NamedValue object, based on the index passed in. The function uses zero-based indexing.
Return Values	If the function succeeds, the return value is a pointer to the NamedValue object. The returned NamedValue object reference is owned by the NVList and should not be released.
Exception	If this function does not succeed, the <code>BAD_PARAM</code> exception is thrown.
See Also	<pre>CORBA::NVList::add CORBA::NVList::add_item CORBA::NVList::add_value CORBA::NVList::count CORBA::NVList::remove</pre>

CORBA::NVList::remove

Synopsis	Removes the item at the specified index, frees any associated memory, and reorders the remaining items on the list.
C++ Binding	<pre>void CORBA::NVList::remove (CORBA::ULong Index);</pre>
Argument	<p>Index</p> <p>The index into the NVList object. The indexing is zero-based.</p>
Description	This member function removes the item at the specified index, frees any associated memory, and reorders the remaining items on the list.
Return Values	None.
Exception	If this function does not succeed, the BAD_PARAM exception is thrown.
See Also	<p>CORBA::NVList::add</p> <p>CORBA::NVList::add_item</p> <p>CORBA::NVList::add_value</p> <p>CORBA::NVList::count</p> <p>CORBA::NVList::item</p>

Object Member Functions

The rules in this section apply to the OMG IDL interface `Object`, which is the base of the OMG IDL interface hierarchy. Interface `Object` defines a normal CORBA object, not a pseudo-object. However, it is included here because it references other pseudo-objects.

In addition to other rules, all operation names in interface `Object` have leading underscores in the mapped C++ class. Also, the mapping for `create_request` is divided into three forms, corresponding to the usage styles described in the section “Request Member Functions” on page 12-112. The `is_nil` and `release` functions are provided in the CORBA namespace, as described in “Object Member Functions” on page 12-42.

The WLE software uses object reference operations that are defined by CORBA Revision 2.2. These operations depend only on type `Object`, so they can be expressed as regular functions within the CORBA namespace.

Note: Because the WLE software uses the POA and not the BOA, the deprecated `get_implementation()` member function is not visible; you will get a compile error if you attempt to reference it.

The mapping of these member functions to C++ is as follows:

```
class CORBA
{
    class Object
    {
    public:
        CORBA::Boolean _is_a(const char *)
        CORBA::Boolean _is_equivalent();
        CORBA::Boolean _nonexistent(Object_ptr);

        static Object_ptr _duplicate(Object_ptr obj);
        static Object_ptr _nil();
        InterfaceDef_ptr _get_interface();
        CORBA::ULong _hass(CORBA::ULong);
        void _create_request(
            Context_ptr ctx,
            const char *operation,
            NVList_ptr arg_list,
            NamedValue_ptr result,
```

```
        Request_out request,
        Flags req_flags
    );
    Status _create_request(
        Context_ptr      ctx,
        const char *      operation,
        NVList_ptr        arg_list,
        NamedValue_ptr     result,
        ExceptionList_ptr  Except_list,
        ContextList_ptr    Context_list,
        Request_out        request,
        Flags              req_flags
    );
    Request_ptr _request(const char* operation);
}; //Object
}; // CORBA
```

The following sections describe the Object member functions.

CORBA::Object::_create_request

Synopsis Creates a request with user-specified information.

C++ Binding `Void CORBA::Object::_create_request (`
 `CORBA::Context_ptr Ctx,`
 `const char * Operation,`
 `CORBA::NVList_ptr Arg_list,`
 `CORBA::NamedValue_ptr Result,`
 `CORBA::ExceptionList_ptr Except_list,`
 `CORBA::ContextList_ptr Context_list,`
 `CORBA::Request_out Request,`
 `CORBA::Flags Req_flags,);`

Arguments `Ctx`
 The Context to be used for this request.

`Operation`
 The operation name for this request.

`Arg_list`
 The argument list for this request.

`Result`
 The NamedValue reference where the return value of this request is to be stored after a successful invocation.

`Except_list`
 The exception list for this request.

`Context_list`
 The context list for this request.

`Request`
 The newly created request reference.

`Req_flags`
 Reserved for future use; the user must pass a value of zero.

Description This member function creates a request that provides information on context, operation name, and other values (long form). To create a request with just the operation name supplied at the time of the call (short form), use the `CORBA::Object::_request` member function. The remainder of the information provided in the long form eventually needs to be supplied.

Return Values None.

See Also `CORBA::Object::_request`

CORBA::Object::_duplicate

Synopsis Duplicates the Object object reference.

C++ Binding

```
CORBA::Object_ptr CORBA::Object::_duplicate(  
                                Object_ptr Obj);
```

Argument `obj`
The object reference to be duplicated.

Description This member function duplicates the specified Object object reference (`Obj`). If the given object reference is nil, the `_duplicate` function returns a nil object reference. The object returned by this call should be freed using `CORBA::release`, or should be assigned to `CORBA::Object_var` for automatic destruction.

This function can throw CORBA system exceptions.

Return Values Returns the duplicate object reference. If the specified object reference is nil, a nil object reference is returned.

Example

```
CORBA::Object_ptr op = TP::create_object_reference(  
                                "IDL:Teller:1.0", "MyTeller");  
CORBA::Object_ptr dop = CORBA::Object::_duplicate(op);
```


CORBA::Object::_get_interface

Synopsis Returns an interface definition for the Repository object.

C++ Binding `CORBA::InterfaceDef_ptr CORBA::Object::_get_interface ();`

Arguments None.

Description Returns an interface definition for the Repository object.

Note: To use the Repository Interface API, define a macro before `CORBA.h` is included. For information about how to define a macro, see *Creating C++ Server Applications*.

Return Values `InterfaceDef_ptr`

CORBA::Object::_is_a

Synopsis	Determines whether an object is of a certain interface.
C++ Binding	<pre>CORBA::Boolean CORBA::Object::_is_a(const char * interface_id);</pre>
Argument	<pre>interface_id</pre> <p>A string that denotes the interface repository ID.</p>
Description	This member function is used to determine if an object is an instance of the interface that you specify in the <code>interface_id</code> parameter. It facilitates maintaining type-safety for object references over the scope of an ORB.
Return Values	Returns TRUE if the object is an instance of the specified type, or if the object is an ancestor of the “most derived” type of that object.
Example	<pre>CORBA::Object_ptr op = TP::create_object_reference("IDL:Teller:1.0", "MyTeller"); CORBA::Boolean b = op->_is_a("IDL:Teller:1.0");</pre>
Exceptions	Can throw a standard CORBA exception.

CORBA::Object::_is_equivalent

Synopsis	Determines if two object references are equivalent.
C++ Binding	<pre>CORBA::Boolean CORBA::Object::_is_equivalent (CORBA::Object_ptr other_obj);</pre>
Argument	<p><code>other_obj</code></p> <p>The object reference for the other object, which is used for comparison with the target object.</p>
Description	This member function is used to determine if two object references are equivalent, so far as the ORB can easily determine. It returns TRUE if your object reference is equivalent to the object reference you pass as a parameter. If two object references are identical, they are equivalent. Two different object references that refer to the same object are also equivalent.
Return Values	Returns TRUE if the target object reference is known to be equivalent to the other object reference passed as a parameter; otherwise, it returns FALSE.
Example	<pre>CORBA::Object_ptr op = TP::create_object_reference("IDL:Teller:1.0", "MyTeller"); CORBA::Object_ptr dop = CORBA::Object::_duplicate(op); CORBA::Boolean b = op->_is_equivalent(dop);</pre>

CORBA::Object::_nil

Synopsis Returns a reference to a nil object.

C++ Binding `CORBA::Object_ptr CORBA::Object::_nil();`

Arguments None.

Description This member function returns a nil object reference. To test whether a given object is nil, use the appropriate `CORBA::is_nil` member function (see the section “CORBA::release” on page 12-54). Calling the `CORBA::is_nil` routine on any `_nil` member function always yields `CORBA_TRUE`.

Return Values Returns a nil object reference.

Example `CORBA::Object_ptr op = CORBA::Object::_nil();`

CORBA::Object::_non_existent

Synopsis	May be used to determine if an object has been destroyed.
C++ Binding	<code>CORBA::Boolean CORBA::Object::_non_existent();</code>
Arguments	None.
Description	This member function may be used to determine if an object has been destroyed. It does this without invoking any application-level operation on the object, and so will never affect the object itself.
Return Values	Returns <code>CORBA_TRUE</code> (rather than raising <code>CORBA::OBJECT_NOT_EXIST</code>) if the ORB knows authoritatively that the object does not exist; otherwise, it returns <code>CORBA_FALSE</code> .

CORBA::Object::_request

Synopsis	Creates a request specifying the operation name.
C++ Binding	<pre>CORBA::Request_ptr CORBA::Object::_request (const char * Operation);</pre>
Argument	<p>Operation</p> <p>The name of the operation for this request.</p>
Description	This member function creates a request specifying the operation name. All other information, such as arguments and results, must be populated using <code>CORBA::Request</code> member functions.
Return Values	<p>If the member function succeeds, the return value is a pointer to the newly created request.</p> <p>If the member function does not succeed, an exception is thrown.</p>
See Also	<code>CORBA::Object::_create_request</code>

CORBA Member Functions

This section describes the Object and Pseudo-Object Reference member functions.

The mapping of these member functions to C++ is as follows:

```
class CORBA {
    void release(Object_ptr);
    void release(Environment_ptr);
    void release(NamedValue_ptr);
    void release(NVList_ptr);
    void release(Request_ptr);
    void release(Context_ptr);
    void release(TypeCode_ptr);
    void release(POA_ptr);
    void release(ORB_ptr);
    void release(ExceptionList_ptr);
    void release(ContextList_ptr);

    Boolean is_nil(Object_ptr);
    Boolean is_nil(Environment_ptr);
    Boolean is_nil(NamedValue_ptr);
    Boolean is_nil(NVList_ptr);
    Boolean is_nil(Request_ptr);
    Boolean is_nil(Context_ptr);
    Boolean is_nil(TypeCode_ptr);
    Boolean is_nil(POA_ptr);
    Boolean is_nil(ORB_ptr);
    Boolean is_nil(ExceptionList_ptr);
    Boolean is_nil(ContextList_ptr);

    hash(maximum);

    resolve_initial_references(identifier);
    ...
};
```

CORBA::release

Synopsis Allows allocated resources to be released for the specified object type.

C++ Binding `void CORBA::release(spec_object_type obj);`

Argument `obj`
The object reference that the caller will no longer access. The specified object type must be one of the types listed in the section “CORBA Member Functions” on page 12-53.

Description This member function indicates that the caller will no longer access the reference so that associated resources may be deallocated. If the specified object reference is nil, the release operation does nothing. If the ORB instance release is the last reference to the ORB, then the ORB will be shutdown prior to its destruction. This is the same as calling `ORB_shutdown` prior to calling `CORBA::release`. This only applies to the release member function called on the ORB.

This member function may not throw CORBA exceptions.

Return Values None.

Example

```
CORBA::Object_ptr op = TP::create_object_reference(
    "IDL:Teller:1.0", "MyTeller");
CORBA::release(op);
```


CORBA::is_nil

Synopsis	Determines if an object exists for the specified object type.
C++ Binding	<code>CORBA::Boolean CORBA::is_nil(spec_object_type obj);</code>
Argument	<code>obj</code> The object reference. The specified object type must be one of the types listed in the section “CORBA Member Functions” on page 12-53.
Description	<p>This member function is used to determine if a specified object reference is nil. It returns TRUE if the object reference contains the special value for a nil object reference as defined by the ORB.</p> <p>This operation may not throw CORBA exceptions.</p>
Return Values	Returns TRUE if the specified object is nil; otherwise, returns FALSE.
Example	<pre>CORBA::Object_ptr op = TP::create_object_reference("IDL:Teller:1.0", "MyTeller"); CORBA::Boolean b = CORBA::is_nil(op);</pre>

CORBA::hash

Synopsis	Provides indirect access to object references using identifiers internal to the ORB.
C++ Binding	<code>CORBA::hash(CORBA::ULong maximum);</code>
Argument	<code>maximum</code> Specifies an upper bound on the hash value returned by the ORB.
Description	<p>Object references are associated with ORB-internal identifiers that may indirectly be accessed by applications using the <code>hash()</code> operation. The value of this identifier does not change during the lifetime of the object reference, and so neither will any hash function of that identifier.</p> <p>The value of this operation is not guaranteed to be unique; that is, another object reference may return the same hash value. However, if two object references hash differently, applications can determine that the two object references are <i>not</i> identical.</p> <p>The <code>maximum</code> parameter to the <code>hash</code> operation specifies an upper bound on the hash value returned by the ORB. The lower bound of that value is zero. Since a typical use of this feature is to construct and access a collision-chained hash table of object references, the more randomly distributed the values are within that range, and the less expensive those values are to compute, the better.</p>
Return Values	None.

CORBA::resolve_initial_references

Synopsis Returns an initial object reference corresponding to an identifier string.

C++ Binding `CORBA::Object_ptr CORBA::resolve_initial_references(
const CORBA::char *identifier);`

Argument identifier
String identifying the object whose reference is required.

Description Returns an initial object reference corresponding to an identifier string. Valid identifiers are "RootPOA" and "POACurrent".

Note: This function is supported only for a joint client/server.

Return Values Returns a `CORBA::Object_ptr`.

Exception InvalidName

Example

```
CORBA::ORB_ptr orb = CORBA::ORB_init(argc, argv);  
CORBA::Object_ptr pobj =  
    orb->resolve_initial_references("RootPOA");  
PortableServer::POA_ptr rootPOA;  
rootPOA = PortableServer::POA::narrow(pobj);
```

ORB Member Functions

The ORB member functions constitute the programming interface to the Object Request Broker.

The mapping of the ORB member functions to C++ is as follows:

```
class CORBA
{
    class ORB
    {
    public:
        char *object_to_string(Object_ptr);
        Object_ptr string_to_object(const char *);
        void create_list(Long, NVList_out);
        void create_operation_list(operationDef_ptr, NVList_out);
        void create_named_value(NamedValue_out);
        void create_environment(Environment_out);
        void create_policy (in PolicyType type, in any val);
        void destroy ();
        void send_multiple_requests_oneway(const requestSeq&);
        void send_multiple_requests_deferred(const requestSeq&);
        void create_exception_list(ExceptionList_out);
        void create_context_list(ContextList_out);
        void get_default_context(Context_out);
        void get_next_response(Request_out);
        void perform_work();
        void run();
        void shutdown(in boolean wait_for_completion);
        Boolean poll_next_response();
        Boolean work_pending( );
    }; //ORB
}; // CORBA
```

Thread-related Operations:

To support single-threaded ORBs, as well as multithreaded ORBs that run multithread-unaware code, four operations (`perform_work`, `run`, `shutdown`, and `work_pending`) are included in the ORB interface. These operations can be used by single-threaded and multithreaded applications. An application that is a pure ORB client would not need to use these operations. Both the `ORB::run()` and `ORB::shutdown()` are useful in fully multithreaded programs.

The following sections describe the ORB member functions.

CORBA::ORB::create_environment

Synopsis Creates an environment.

```
C++ Binding void CORBA::ORB::create_environment (
               CORBA::Environment_out New_env);
```

Argument	<code>New_env</code>	Receives a reference to the newly created environment.
----------	----------------------	--

Description	This member function creates an environment.
-------------	--

Return Values **None.**

See Also

- `CORBA::NVList::add`
- `CORBA::NVList::add_item`
- `CORBA::NVList::add_value`
- `CORBA::release`

CORBA::ORB::create_list

Synopsis Creates and returns an NVList object reference.

C++ Binding `void CORBA::ORB::create_list (`
 `CORBA::Long NumItem,`
 `CORBA::NVList_out List);`

Arguments `NumItem`
 The number of elements to preallocate in the newly created list.

`List`
 Receives the newly created list.

Description This member function creates a list, preallocating a specified number of items. List items may be sequentially added to the list using the `CORBA::NVList_add_item` member function. When no longer needed, this list must be freed using the `CORBA::release` member function.

Return Values `None.`

See Also `CORBA::NVList::add`
 `CORBA::NVList::add_item`
 `CORBA::NVList::add_value`
 `CORBA::release`

CORBA::ORB::create_named_value

Synopsis Creates a NamedValue object reference.

```
C++ Binding    void CORBA::ORB::create_named_value (
                NameValue_out          NewNamedVal);
```

Argument	NewNamedVal
	A reference to the newly created NamedValue object.

Description	This member function creates a NamedValue object. Its intended use is for the result argument of a request that needs a NamedValue object. The extra steps of creating an NVList object are avoided by calling this member function.
-------------	--

When no longer needed, the NamedValue object must be freed using the `CORBA::release` member function.

Return Values **None.**

See Also

- `CORBA::NVList::add`
- `CORBA::NVList::add_item`
- `CORBA::NVList::add_value`
- `CORBA::release`

CORBA::ORB::create_exception_list

Synopsis Returns a list of exceptions.

C++ Binding `void CORBA::ORB::create_exception_list(
CORBA::ExceptionList_out List);`

Argument `List`
Receives a reference to the newly created exception list.

Description This member function creates and returns a list of exceptions in a form that may be used in the Dynamic Invocation Interface (DII). When no longer needed, this list must be freed using the `CORBA::release` member function.

Return Values None.

CORBA::ORB::create_context_list

Synopsis Creates and returns a list of contexts.

C++ Binding `void CORBA::ORB::create_context_list(
 CORBA::ContextList_out List);`

Argument `List`
 Receives a reference to the newly created context list.

Description This member function creates and returns a list of context strings that must be supplied with the Request operation in a form that may be used in the Dynamic Invocation Interface (DII). When no longer needed, this list must be freed using the `CORBA::release` member function.

Return Values None.

CORBA::ORB::create_policy

Synopsis Creates new instances of policy objects of a specific type with specified initial state.

C++ Binding `void CORBA::ORB::create_policy (`
 `in PolicyType type,`
 `in any val);`

Arguments `type`
 `BiDirPolicy::BIDIRECTIONAL_POLICY_TYPE` is the only `PolicyType` value supported for WLE V4.2 .

`val`
 The only `val` value supported for WLE V4.2 is
 `BiDirPolicy::BidirectionalPolicyValue`.

Description This operation can be invoked to create new instances of policy objects of a specific type with specified initial state. If `create_policy` fails to instantiate a new `Policy` object due to its inability to interpret the requested type and content of the policy, it raises the `Policy Error` exception with the appropriate reason. (See Exceptions below.)

The `BidirectionalPolicy` argument is provided for remote clients using callbacks because remote clients use `IOP`. It is not used for native clients using callbacks or for WLE servers because machines inside a WLE domain communicate differently.

Before `GIOP 1.2`, `bidirectional` policy was not available as a choice in `IOP` (which uses `TCP/IP`). Connections in `GIOP 1.0` and `1.1` were one way (that is, a request flowed from a client to a server); only responses flowed from the server back to the client. If the server wanted to make a request back to the client machine (say for a callback), the server machine had to establish another one-way connection. (Be advised that “connections” in this sense mean operating system resources, not physically different wires or communication paths. A connection uses resources, so minimizing connections is desirable.)

Since this release of the WLE C++ software supports `GIOP 1.2`, it supports re-use of the `TCP/IP` connection for both incoming and outgoing requests. Re-using connections saves resources when a remote client sends callback references to a WLE domain. The joint client/server uses a connection to send a request to a WLE domain; that connection can be re-used for the callback request. If the connection is not re-used, the callback request must establish another connection.

Allowing re-use of a connection is a choice of the ORB/POA that creates callback object references. The server for those object references (usually the creator of the references, especially in the callback case) might choose not to allow re-use for

security considerations (that is, the outgoing connection (a client request from this machine to a remote server) may not need security because the remote server does not require it, but the callback server on this machine might require security). Since security is established partly on a connection basis, the incoming security can be established only if a separate connection is used. If the remote server requires security, and if that security involves a mutual authentication, the local server usually feels safe in allowing re-use of the connection.

Since the choice of connection re-use is at the server end, whenever a process acts as a server—in this case a joint client/server—and creates object references, it must inform the ORB that it is willing to re-use connections. The process does this by setting a policy on the POA that creates the object references. The default policy is to not allow re-use (that is, if you do not supply a policy object for re-use, the POA does not allow re-use).

This default allows for backward compatibility with code written before CORBA version 2.3. Such code did not know that re-use was possible so it did not have to take into consideration the security implications of re-use. Thus, that unchanged code should continue to disallow re-use until the user considers security and explicitly makes a decision to the contrary.

To allow re-use, you use the `create_policy` operation to create a policy object that allows re-use, and use that policy object as part of the list of policies for POA creation.

Return Values None.

Exceptions `PolicyError`
 This exception is raised to indicate problems with the parameter values passed to the `ORB::create_policy` operation. The specific exception and reasons are as follows:

Exception	Reason
<code>BAD_POLICY</code>	The requested Policy is not understood by the ORB.
<code>UNSUPPORTED_POLICY</code>	The requested Policy is understood to be valid by the ORB, but is not currently supported.
<code>BAD_POLICY_TYPE</code>	The type of the value requested for the Policy is not valid for that PolicyType.
<code>BAD_POLICY_VALUE</code>	The value requested for the Policy is of a valid type, but is not within the valid range for that type.

Exception	Reason
UNSUPPORTED_POLICY_VALUE	The value requested for the Policy is of a valid type and within the valid range for that type, but this valid value is not currently supported.

Example

```
#include <BiDirPolicy_c.h>
BiDirPolicy::BidirectionalPolicy_var bd_policy;
CORBA::Any allow_reuse;

allow_reuse <= BiDirPolicy::BOTH;

CORBA::Policy_var generic_policy =
    orb->create_policy( BiDirPolicy::BIDIRECTIONAL_POLICY_TYPE,
                      allow_reuse );
bd_policy = BiDirPolicy::BidirectionalPolicy::_narrow(
    generic_policy );
```

In the above example, the `bd_policy` would then be placed in the `PolicyList` passed to the `create_poa` operation.

CORBA::ORB::create_operation_list

Synopsis Creates and returns a list of the arguments of a specified operation.

C++ Binding `void CORBA::ORB::create_operation_list (`
 `CORBA::OperationDef_ptr Oper,`
 `CORBA::NVList_out List);`

Arguments `Oper`
 The operation definition for which the list is being created.

`List`
 Receives a reference to the newly created arguments list.

Description This member function creates and returns a list of the arguments of a specified operation, in a form that may be used with the Dynamic Invocation Interface (DII). When no longer needed, this list must be freed using the `CORBA::release` member function.

Return Values None.

See Also `CORBA::ORB::create_list`
 `CORBA::NVList::add`
 `CORBA::NVList::add_item`
 `CORBA::NVList::add_value`
 `CORBA::release`

CORBA::ORB::get_default_context

Synopsis Returns a reference to the default context.

C++ Binding

```
void CORBA::ORB::get_default_context (
    CORBA::Context_out ContextObj);
```

Argument ContextObj
The reference to the default context.

Description This member function returns a reference to the default context. When no longer needed, this context reference must be freed using the `CORBA::release` member function.

Return Values None.

See Also `CORBA::Context::get_one_value`
`CORBA::Context::get_values`

CORBA::ORB::get_next_response

Synopsis Determines and reports the next deferred synchronous request that completes.

C++ Binding `void CORBA::ORB::get_next_response (`
 `CORBA::Request_out RequestObj);`

Argument `RequestObj`
 The reference to the next completed request.

Description This member function returns a reference to the next request that completes. If no requests have completed, the function waits for a request to complete. This member function returns the next request on the queue, in contrast to the `CORBA::Request::get_response` member function, which waits for a particular request to complete. When no longer needed, this request must be freed using the `CORBA::release` member function.

Return Values None.

See Also `CORBA::ORB::poll_next_response`
 `CORBA::Request::get_reponse`

CORBA::ORB::perform_work

Synopsis Allows the ORB to perform server-related work.

C++ Binding `void CORBA::ORB::perform_work ();`

Arguments None.

Description If called by the main thread, this operation allows the ORB to perform server-related work. Otherwise, it does nothing.

The `work_pending()` and `perform_work()` operations can be used to write a simple polling loop that multiplexes the main thread among the ORB and other activities. Such a loop would most likely be needed in a single-threaded server. A multithreaded server would need a polling loop only if there were both ORB and other code that required use of the main thread. See the example below for such a polling loop.

Return Values None.

Exceptions Once the ORB has shut down, a call to `work_pending` and `perform_work()` raises the `BAD_INV_ORDER` exception. An application can detect this exception to determine when to terminate a polling loop.

See Also `CORBA::ORB::work_pending`

Example Here is an example of a polling loop:

```
// C++
for (;;) {
    if (orb->work_pending()) {
        orb->perform_work();
    }
    // do other things
    // sleep?
}
```

CORBA::ORB::run

Synopsis	Enables the ORB to perform work using the main thread.
C++ Binding	<code>void CORBA::ORB::run();</code>
Argument	None
Description	<p>This operation provides execution resources to the ORB so that it can perform its internal functions. Since the WLE C++ ORB is single-threaded, this essentially turns the process into a pure server.</p> <p>This operation blocks until the ORB has completed the shutdown process, initiated when a thread calls <code>CORBA::ORB::shutdown()</code>.</p>
Return Values	None.
See Also	<code>CORBA::ORB::perform_work</code>

CORBA::ORB::shutdown

Synopsis Instructs the ORB to shut down.

C++ Binding `void shutdown(in boolean wait_for_completion);`

Argument `wait_for_completion`
 A value of TRUE blocks until all ORB processing has completed.

Description This operation instructs the ORB to shut down (that is, to stop processing in preparation for destruction).

Shutting down the ORB causes all POAs to be destroyed, since they cannot exist in the absence of an ORB. Shut down is complete when all ORB processing (including request processing and object deactivation or other operations associated with the POAs) has completed and all POAs have been destroyed.

If the `wait_for_completion` parameter is TRUE, this operation blocks until the shut down is complete. If an application calls `shutdown(TRUE)` in a thread that is currently servicing an invocation, the `BAD_INV_ORDER` system exception will be raised with the OMG minor code 3, since blocking would result in a deadlock.

If the `wait_for_completion` parameter is FALSE, the shutdown may not have completed upon return. This use of FALSE is not recommended.

While the ORB is in the process of shutting down, the ORB operates as normal, servicing incoming and outgoing requests until all requests have been completed.

Once an ORB has shut down, invoking any operation on that ORB or any object reference obtained from that ORB will raise the `BAD_INV_ORDER` system exception with the OMG minor code 4, except for the reference management operations `duplicate()`, `release()`, and `is_nil()`.

Return Values None.

CORBA::ORB::object_to_string

Synopsis Produces a string representation of an object reference.

C++ Binding

```
char * CORBA::ORB::object_to_string (  
    CORBA::Object_ptr  ObjRef);
```

Argument ObjRef
The object reference to represent as a string.

Description This member function produces a string representation of an object reference. The calling program must use the `CORBA::string_free` member function to free the string memory after it is no longer needed.

Return Values The string representing the specified object reference.

Example

```
CORBA::Object_ptr op = TP::create_object_reference(  
    "IDL:Teller:1.0", "MyTeller");  
char* objstr = TP::orb()->object_to_string(op);
```

See Also `CORBA::ORB::string_to_object`
`CORBA::string_free`

CORBA::ORB::poll_next_response

Synopsis Determines whether a completed request is outstanding.

C++ Binding `CORBA::Boolean CORBA::ORB::poll_next_response ();`

Arguments None.

Description This member function reports on whether there is an outstanding (pending) completed request; it does not remove the request. If a completed request is outstanding, the next call to the `CORBA::ORB::get_next_response` member function is guaranteed to return a request without waiting. If there are no completed requests outstanding, the `CORBA::ORB::poll_next_response` member function returns without waiting (blocking).

Return Values If a completed request is outstanding, the function returns `CORBA_TRUE`.

 If no completed request is outstanding, the function returns `CORBA_FALSE`.

See Also `CORBA::ORB::get_next_response`

CORBA::ORB::work_pending

Synopsis	Returns an indication of whether the ORB needs the main thread to perform server-related work.
C++ Binding	<code>CORBA::boolean CORBA::ORB::work_pending ();</code>
Arguments	None.
Description	This operation returns an indication of whether the ORB needs the main thread to perform server-related work.
Return Values	A result of TRUE indicates that the ORB needs the main thread to perform server-related work, and a result of FALSE indicates that the ORB does not need the main thread.
See Also	<code>CORBA::ORB::perform_work</code>

CORBA::ORB::send_multiple_requests_deferred

Synopsis Sends a sequence of deferred synchronous requests.

C++ Binding `void CORBA::ORB::send_multiple_requests_deferred (
 const CORBA::ORB::RequestSeq & Reqs);`

Argument `Reqs`
 The sequence of requests to be sent. For more information about how to populate the sequence with request references, see `CORBA::ORB::RequestSeq` in the section “Usage” on page 11-23.

Description This member function sends out a sequence of requests and returns control to the caller without waiting for the operation to complete. The caller uses `CORBA::ORB::poll_next_response`, `CORBA::ORB::get_next_response`, or `CORBA::Request::get_response` or all three to determine if the operation has completed and if the output arguments have been updated.

Return Values None.

See Also `CORBA::Request::get_response`
 `CORBA::ORB::get_next_response`
 `CORBA::ORB::send_multiple_requests_oneway`

CORBA::ORB::send_multiple_requests_oneway

Synopsis Sends a sequence of one-way, deferred synchronous requests.

C++ Binding `void CORBA::ORB::send_multiple_requests_oneway (
 const CORBA::RequestSeq & Reqs);`

Argument `Reqs`
 The sequence of requests to be sent. For more information about how to populate the sequence with request references, see `CORBA::ORB::RequestSeq` in the section “Usage” on page 11-23.

Description This member function sends out a sequence of requests and returns control to the caller without waiting for the operation to complete. The caller neither intends to wait for a response nor expects any output arguments to be updated.

Return Values None.

See Also `CORBA::ORB::send_multiple_requests_deferred`

CORBA::ORB::string_to_object

Synopsis Creates an object reference, given a specified string.

C++ Binding `CORBA::Object_ptr CORBA::ORB::string_to_object (`
 `const char * ObjRefString);`

Argument `ObjRefString`
 The string to be transformed into an object reference.

Description This member function creates an object reference, given a specified string. Usually the string has been obtained previously by calling the `CORBA::ORB::object_to_string` member function. After you are done with the object reference, use the `CORBA::release` member function to free the associated memory.

Return Values If the member function succeeds, the object reference represented by the specified string is returned.

 If the member function does not succeed, an exception is thrown.

Example `CORBA::Object_ptr op = TP::create_object_reference(`
 `"IDL:Teller:1.0", "MyTeller");`
 `char* objstr = TP::orb()->object_to_string(op);`
 `CORBA::Object_ptr op2 = TP::orb()->string_to_object(objstr);`

See Also `CORBA::ORB::object_to_string`

ORB Initialization Member Function

The mapping of this member function to C++ is as follows:

```
class CORBA {  
    typedef char* ORBId;  
    static CORBA::ORB_ptr ORB_init(int& argc, char** argv,  
                                   const char* orb_identifier = 0,  
                                   const char* -ORBport nnn);  
};
```

CORBA::ORB_init

Synopsis Initializes operations for an ORB.

C++ Binding `static CORBA::ORB_ptr ORB_init(int& argc, char** argv,
 const char* orb_identifier = 0);`

Arguments `argc`

 The number of strings in `argv`.

`argv`

 This argument is defined as an unbound array of strings (`char **`) and the number of strings in the array is passed in the `argc` parameter.

`orb_identifier`

 If the `orb_identifier` parameter is supplied, "BEA_IIOP" explicitly specifies a remote client and "BEA_TOBJ" explicitly specifies a native client, as defined in the section "Tobj_Bootstrap" on page 4-11.

Description

This member function initializes operations for an ORB and returns a pointer to the ORB. When your program is done with the ORB, use the `CORBA::release` member function to free the resources allocated for the ORB pointer returned from `CORBA::ORB_ptr ORB_init`.

The ORB returned has been initialized with two pieces of information to determine how it will operate: client type (remote or native) and server port number. The client type can be specified in the `orb_identifier` argument, in the `argv` argument, or in the system registry. The server port number can be specified in the `argv` argument.

The arguments `argc` and `argv` are typically the same parameters that were passed to the main program. As specified by C++, these parameters contain string tokens from the command line that started the client. The two ORB options can be specified on the command line, each using a pair of tokens, as shown in examples below.

CLIENT TYPE

The `ORB_init` function determines the client type of the ORB by the following steps.

1. If the `orb_identifier` argument is present, `ORB_init` determines the client type, either native or remote, if the string is "BEA_IIOP" or "BEA_TOBJ", respectively. If an `orb_identifier` string is present, all `-ORBid` parameters in the `argv` are ignored (removed).

2. If `orb_identifier` is not present or is explicitly zero, `ORB_init` looks at the entries in `argc/argv`. If `argv` contains an entry with `"-ORBid"`, the next entry should be either `"BEA_IIOP"` or `"BEA_TOBJ"`, again specifying remote or native. This pair of entries occurs if the command line contains either `"-ORBid BEA_IIOP"` or `"-ORBid BEA_TOBJ"`.
3. If no client type is specified in `argc/argv`, `ORB_init` uses the default client type from the system registry (`BEA_IIOP` or `BEA_TOBJ`). The system registry was initialized at the time WLE was installed.

SERVER PORT

In the case of a WLE remote joint client/server, in order to support IIOP, by definition, the object references created for the server part must contain a host and port. For transient object references, any port is sufficient and can be obtained by the ORB dynamically, but this is not sufficient for persistent

object references. Persistent references must be served on the same port after the ORB restarts, that is, the ORB must be prepared to accept requests on the same port with which it created the object reference. Thus, there must be some way to configure the ORB to use a particular port.

Typically, a system administrator assigns the port number for the client from the “user” range of port numbers rather than the dynamic range. This keeps the joint client/servers from using conflicting ports.

To determine port number, `ORB_init` searches the `argv` parameter for the token `"-ORBport"` and a following numeric token. For example, if the client executable is named `sherry`, the command line might specify that the server port should be 937 as follows:

```
sherry -ORBport 937
```

ARGV PARAMETER CONSIDERATIONS

For C++, the order of consumption of `argv` parameters may be significant to an application. To ensure that applications are not required to handle `argv` parameters they do not recognize, the ORB initialization function must be called before the remainder of the parameters are consumed. Therefore, after the `ORB_init` call, the `argv` and `argc` parameters have been modified to remove the ORB understood arguments. It is important to note that the `ORB_init` function can only reorder or remove references to parameters from the `argv` list. This restriction is made to avoid

potential memory management problems caused by trying to free parts of the `argv` list or extending the `argv` list of parameters. This is why `argv` is passed as a `char**` and not as a `char*&`.

Note: Use the `CORBA::release` member function to free the resources allocated for the pointer returned from `CORBA::ORB_init`.

Return Value A pointer to a `CORBA::ORB`.

Exceptions None.

Policy Member Functions

A policy is an object used to communicate certain choices to an ORB regarding its operation. This information is accessed in a structured manner using interfaces derived from the Policy interface defined in the CORBA module.

Note: These `CORBA::Policy` operations and structures are not usually needed by programmers. The derived interfaces usually contain the information relevant to specifications. A Policy object can be constructed by a specific factory or by using the `CORBA::create_policy` operation.

The mapping of this object to C++ is as follows:

```
class CORBA
{
    class Policy
    {
    public:
        copy();
        void destroy();
    }; //Policy
    typedef sequence<Policy>PolicyList;
}; // CORBA
```

`PolicyList` is used the same as any other C++ sequence mapping. For a discussion of sequence usage, see “Sequences” on page 11-13.

See Also: POA Policy and `CORBA::ORB::create_policy`.

CORBA:Policy::copy

Synopsis Copies the policy object.

C++ Binding `CORBA::Policy::copy()` ;

Argument None.

Description This operation copies the policy object. The copy does not retain any relationships that the policy had with any domain or object.

Note: This function is supported only for a joint client/server.

Return Values None.

CORBA::Policy::destroy

Synopsis Destroys the policy object.

C++ Binding `void CORBA::Policy::destroy();`

Argument None.

Description This operation destroys the policy object. It is the responsibility of the policy object to determine whether it can be destroyed.

Note: This function is supported only for a joint client/server.

Return Values None.

Exceptions If the policy object determines that it cannot be destroyed, the `CORBA::NO_PERMISSION` exception is raised.

PortableServer Member Functions

The mapping of the PortableServer member functions to C++ is as follows:

```
// C++
class PortableServer
{
    public:
        class LifespanPolicy;
        class IdAssignmentPolicy;
        class POA::find_POA
        class reference_to_id
        class POAManager;
        class POA;
        class Current;
        class virtual ObjectId
        class ServantBase
};
```

ObjectId—An `ObjectId` is a value that is used by the POA and by the user-supplied implementation to identify a particular abstract CORBA object. `ObjectId` values may be assigned and managed by the POA, or they may be assigned and managed by the implementation. `ObjectId` values are hidden from clients, encapsulated by references. `ObjectIds` have no standard form; they are managed by the POA as uninterpreted octet sequences.

The following sections describe the remaining classes.

PortableServer::POA::activate_object

Synopsis Explicitly activates an individual object.

C++ Binding `ObjectId * activate_object (
 Servant p_servant);`

Argument `p_servant`
 An instance of the C++ implementation class for the interface.

Description This operation explicitly activates an individual object by generating an `ObjectId` and entering the `ObjectId` and the specified servant in the Active Object Map.

Note: This function is supported only for a joint client/server.

Return Values If the function succeeds, the `ObjectId` is returned.

Exceptions If the specified servant is already in the Active Object Map, the `ServantAlreadyActive` exception is raised.

Note: Other exceptions can occur if the POA uses unsupported policies.

Example In the following example, the first struct creates a servant by a user-defined constructor. The second struct tells the POA that the servant can be used to handle requests on an object. The POA returns the `ObjectId` it has created for the object. The third statement assumes that the POA has the `IMPLICIT_ACTIVATION` policy (the only supported policy in version 4.2 of the WLE software) and returns a reference to the object. That reference can then be handed to a client for invocations. When the client invokes on the reference, the request is returned to the servant just created.

```
MyFooServant* afoo = new MyFooServant(poa,27);
PortableServer::ObjectId_var oid =
    poa->activate_object(afoo);
Foo_var foo = afoo->_this();
```

PortableServer::POA::activate_object_with_id

Synopsis Activates an individual object with a specified `ObjectId`.

C++ Binding

```
void activate_object_with_id (
    const ObjectId & id,
    Servant p_servant);
```

Argument `id`
 `ObjectId` that identifies the object on which that operation was invoked.

 `p_servant`
 An instance of the C++ implementation class for the interface.

Description This operation enters an association between the specified `ObjectId` and the specified servant in the Active Object Map.

Note: This function is supported only for a joint client/server.

Return Values None.

Exceptions The `ObjectAlreadyActive` exception is raised if the CORBA object denoted by the `ObjectId` value is already active in this POA.

 The `ServantAlreadyActive` exception is raised if the servant is already in the Active Object Map.

Note: Other exceptions can occur if the POA uses unsupported policies.

 The `BAD_PARAM` system exception may be raised if the POA has the `SYSTEM_ID` policy and it detects that the `ObjectId` value was not generated by the system or for this POA. An ORB is not required to detect all such invalid `ObjectId` values. However, a portable application must not invoke `activate_object_with_id` on a POA if the POA has the `SYSTEM_ID` policy with an `ObjectId` value that was not previously generated by the system for that POA, or, if the POA also has the `PERSISTENT` policy, for a previous instantiation of the same POA.

Example

```
MyFooServant* afoo = new MyFooServant(poa, 27);
PortableServer::ObjectId_var oid =
    PortableServer::string_to_ObjectId("myLittleFoo");
poa->activate_object_with_id(oid.in(), afoo);
Foo_var foo = afoo->_this();
```

PortableServer::POA::create_id_assignment_policy

Synopsis Obtain an object with the `IdAssignmentPolicy` interface so the user can pass the object to the `POA::create_POA` operation.

C++ Binding `IdAssignmentPolicy_ptr`
`PortableServer::POA::create_id_assignment_policy (`
`PortableServer::IdAssignmentPolicyValue value)`

Argument `value`
 A value of either `PortableServer::USER_ID`, indicating `ObjectIds` are assigned only by the application, or `PortableServer::SYSTEM_ID`, indicating `ObjectIds` are assigned only by the system.

Description The `POA::create_id_assignment_policy` operation obtains objects with the `IdAssignmentPolicy` interface. When passed to the `POA::create_POA` operation, this policy specifies whether `ObjectIds` in the created POA are generated by the application or by the ORB. The following values can be supplied:

- ◆ `PortableServer::USER_ID`—objects created with that POA are assigned `ObjectIds` only by the application.
- ◆ `PortableServer::SYSTEM_ID`—objects created with that POA are assigned `ObjectIds` only by the POA. If the POA also has the `PERSISTENT LifespanPolicy`, assigned `ObjectIds` must be unique across all instantiations of the same POA.

If no `IdAssignmentPolicy` is specified at POA creation, the default is `SYSTEM_ID`.

Note: This function is supported only for a joint client/server.

Return Values Returns an `Id Assignment` policy.

PortableServer::POA::create_lifespan_policy

Synopsis Obtain an object with the `LifespanPolicy` interface so the user can pass the object to the `POA::create_POA` operation.

C++ Binding `LifespanPolicy_ptr`
`PortableServer::POA::create_lifespan_policy (`
`PortableServer::LifespanPolicyPolicyValue value)`

Argument `value`
A value of either `PortableServer::USER_ID`, indicating `ObjectIds` are assigned only by the application, or `PortableServer::SYSTEM_ID`, indicating `ObjectIds` are assigned only by the system.

Description Objects with the `LifespanPolicy` interface are obtained using the `POA::create_lifespan_policy` operation and passed to the `POA::create_POA` operation to specify the lifespan of the objects implemented in the created POA. The following values can be supplied.

- ◆ **TRANSIENT**—The objects implemented in the POA cannot outlive the process in which they are first created. Once the POA is deactivated, use of any object references generated from it will result in an `OBJECT_NOT_EXIST` exception.
- ◆ **PERSISTENT**—The objects implemented in the POA can outlive the process in which they are first created.
 - ◆ Persistent objects have a POA associated with them (the POA which created them). When the ORB receives a request on a persistent object, it first searches for the matching POA, based on the names of the POA and all of its ancestors.
 - ◆ Administrative action beyond the scope of this specification may be necessary to inform the ORB's location service of the creation and eventual termination of existence of this POA, and optionally to arrange for on-demand activation of a process implementing this POA.
 - ◆ POA names must be unique within their enclosing scope (the parent POA). A portable program can assume that POA names used in other processes will not conflict with its own POA names. A conforming CORBA implementation will provide a method for ensuring this property.

If no `LifespanPolicy` object is passed to `POA::create_POA`, the lifespan policy defaults to **TRANSIENT**.

Note: This function is supported only for a joint client/server.

Return Values Returns a LifespanPolicy.

PortableServer::POA::create_POA

Synopsis	Creates a new POA as a child of the target POA.
C++ Binding	<pre>POA_ptr PortableServer::create_POA (const char * adapter_name, POAManager_ptr a_POAManager, const CORBA::PolicyList & policies)</pre>
Arguments	<p><code>adapter_name</code> The name of the POA to be created.</p> <p><code>a_POAManager</code> Either a null value, indicating that a new POAManager is to be created and associated with the new POA, or a pointer to an existing POAManager.</p> <p><code>policies</code> Policy objects to be associated with the new POA.</p>
Description	<p>This operation creates a new POA as a child of the target POA. The specified name, which must be unique, identifies the new POA with respect to other POAs with the same parent POA.</p> <p>If the <code>a_POAManager</code> parameter is null, a new <code>PortableServer::POAManager</code> object is created and associated with the new POA. Otherwise, the specified POAManager object is associated with the new POA. The POAManager object can be obtained using the attribute name <code>the_POAManager</code>.</p> <p>The specified policy objects are associated with the POA and are used to control its behavior. The policy objects are effectively copied before this operation returns, so the application is free to destroy them while the POA is in use. Policies are <i>not</i> inherited from the parent POA.</p> <p>Note: This function is supported only for joint client/servers.</p>
Return Values	Returns a pointer to the POA that was created.
Exceptions	<p><code>AdapterAlreadyExists</code> Raised if the target POA already has a child POA with the specified name.</p> <p><code>InvalidPolicy</code> Raised if any of the policy objects specified are not valid for the ORB implementation, if conflicting policy objects are specified, or if any of the specified policy objects require prior administrative action that has not been</p>

performed. This exception contains the index in the policy parameter value of the first offending policy object.

IMP_LIMIT

Raised if the program tries to create a POA with a LifespanPolicy of PERSISTENT without having set a port, as described in the operation “CORBA::ORB_init” on page 12-81.

Examples **Example 1:**

In this example, the child POA would use the same manager as the parent POA; the child POA would then have the same state as the parent (that is, it would be active if the parent is active).

```
CORBA::PolicyList policies(2);
policies.length (1);
policies[0] = rootPOA->create_lifespan_policy(
    PortableServer::LifespanPolicy::TRANSIENT);
PortableServer::POA_ptr poa =
    rootPOA->create_POA("my_little_poa",
        rootPOA->the_POAManager, policies);
```

Example 2:

In this example, a new POA is created as a child of the root POA.

```
CORBA::PolicyList policies(2);
policies.length (1);
policies[0] = rootPOA->create_lifespan_policy(
    PortableServer::LifespanPolicy::TRANSIENT);
PortableServer::POA_ptr poa =
    rootPOA->create_POA("my_little_poa",
        PortableServer::POAManager::_nil(), policies);
```


PortableServer::POA::create_reference_with_id

Synopsis Creates an object reference that encapsulates the specified `ObjectId` and interface repository ID values.

C++ Binding

```
CORBA::Object_ptr create_reference_with_id (  
                        const ObjectId & oid,  
                        const char * intf)
```

Arguments

<code>oid</code>	<code>ObjectId</code> that identifies the object on which that operation was invoked.
<code>intf</code>	The interface repository ID.

Description The `create_reference` operation creates an object reference that encapsulates the specified `ObjectId` and interface repository ID values. This operation collects the necessary information to constitute the reference from information associated with the POA and from parameters to the operation. This operation only creates a reference; it does not associate the reference with an active servant. The resulting reference may be passed to clients, so that subsequent requests on those references cause the invocation to be returned to the same POA with `ObjectId` specified.

Note: This function is supported only for a joint client/server.

Return Values Returns `Object_ptr`.

Exceptions If the POA has a `LifespanPolicy` with value `SYSTEM_ID` and it detects that the `ObjectId` value was not generated by the system or for this POA, the operation will raise the `BAD_PARAM` system exception.

Example

```
PortableServer::ObjectId_var oid =  
    PortableServer::string_to_ObjectId("myLittleFoo");  
CORBA::Object_var obj = poa->create_reference_with_id(  
    oid.in(), "IDL:Foo:1.0");  
Foo_var foo = Foo::_narrow(obj);
```

PortableServer::POA::deactivate_object

Synopsis Removes the `ObjectId` from the Active Object Map.

C++ Binding `void deactivate_object (`
 `const ObjectId & oid)`

Argument `oid`
 `ObjectId` that identifies the object.

Description This operation causes the association of the `ObjectId` specified by the `oid` parameter and its servant to be removed from the Active Object Map.

Note: This function is supported only for a joint client/server.

Return Values None.

Exceptions If there is no active object associated with the specified `ObjectId`, the operation raises an `ObjectNotActive` exception.

PortableServer::POA::destroy

Synopsis Destroys the POA and all descendant POAs.

C++ Binding

```
void destroy (
    CORBA::Boolean etherealize_objects,
    CORBA::Boolean wait_for_completion)
```

Arguments `etherealize_objects`
This argument should be FALSE for this release of WLE.

`wait_for_completion`
This argument indicates whether or not the operation should return immediately.

Description This operation destroys the POA and all descendant POAs. The POA with its name may be re-created later in the same process. (This differs from the `POAManager::deactivate` operation, which does not allow a re-creation of its associated POA in the same process.)

When a POA is destroyed, any requests that have started execution continue to completion. Any requests that have not started execution are processed as if they were newly arrived and there is no POA; that is, they are rejected and the `OBJECT_NON_EXIST` exception is raised.

If the `wait_for_completion` parameter is `TRUE`, the `destroy` operation returns only after all requests in process have completed and all invocations of `etherealize` have completed. Otherwise, the `destroy` operation returns after destroying the POAs.

Note: This release of WLE does not support multithreading. Hence, `wait_for_completion` should not be `TRUE` if the call is made in the context of an object invocation. That is, the POA cannot start destroying itself if it is currently executing.

Note: This function is supported only for a joint client/server.

Return Values None.

PortableServer::POA::find_POA

Synopsis Returns a reference to a child POA with a given name.

C++ Binding `void find_POA(in string adapter_name, in boolean activate_it);`

Argument `adapter_name`

A reference to the target POA.

`activate_it`

In this version of WLE, this parameter must be false.

Description If the POA has a child POA with the specified name, that child POA is returned. If a child POA with the specified name does not exist and the value of the `activate_it` parameter is `FALSE`, the `AdapterNonExistent` exception is raised.

Return Values None.

Exception `AdapterNonExistent`

This exception is raised if the POA does not exist.

PortableServer::POA::reference_to_id

Synopsis Returns the `ObjectId` value encapsulated by the specified `reference`.

C++ Binding `ObjectId reference_to_id(in Object reference);`

Argument `reference`
Specifies the reference to the object.

Description This operation returns the `ObjectId` value encapsulated by the specified `reference`. This operation is valid only if the reference was created by the POA on which the operation is being performed. The object denoted by the reference does not have to be active for this operation to succeed.

Note: This function is supported only for a joint client/server.

Return Values Returns the `ObjectId` value encapsulated by the specified `reference`.

Exceptions `WrongAdapter`
This exception is raised if the reference was not created by that POA.

PortableServer::POA::the_POAManager

Synopsis Identifies the POA manager associated with the POA.

C++ Binding `POAManager_ptr the_POAManager ();`

Argument None.

Description This read-only attribute identifies the POA manager associated with the POA.

Note: This function is supported only for a joint client/server.

Return Values None.

Example `poa->the_POAManager()->activate();`

This statement will set the state of the POAManager for the given POA to active, which is required if the POA is to accept requests. Note that if the POA has a parent, that is, it is not the root POA, all of its parent's POAManagers must also be in the active state for this statement to have any effect.

PortableServer::ServantBase::_default_POA

Synopsis Returns an object reference to the POA associated with the servant.

C++ Binding

```
class PortableServer
{
    class ServantBase
    {
    public:
        virtual POA_ptr _default_POA();
    }
}
```

Argument None.

Description All C++ Servants inherit from `PortableServer::ServantBase`, so they all inherit the `_default_POA` function. In this version of WLE there is usually no reason to use `_default_POA`.

The default implementation of this function returns an object reference to the root POA of the default ORB in this process—the same as the return value of an invocation of `ORB::resolve_initial_references("RootPOA")`. A C++ Servant can override this definition to return the POA of its choice, if desired.

Note: This function is supported only for joint client/servers.

Return Values The default POA associated with the servant.

POA Current Member Functions

The `PortableServer::Current` interface, derived from `CORBA::Current`, provides method implementations with access to the identity of the object on which the method was invoked.

PortableServer::Current::get_object_id

Synopsis Returns the `ObjectId` identifying the object in whose context it is called.

C++ Binding `ObjectId * get_object_id ();`

Arguments None.

Description This operation returns the `PortableServer::ObjectId` identifying the object in whose context it is called.

Note: This function is supported only for a joint client/server.

Return Values This operation returns the `ObjectId` identifying the object in whose context it is called.

Exception If called outside the context of a POA-dispatched operation, a `PortableServer::NoContext` exception is raised.

PortableServer::Current::get_POA

Synopsis Returns a reference to the POA implementing the object in whose context it is called.

C++ Binding `POA_ptr get_POA ();`

Argument None.

Description This operation returns a reference to the POA implementing the object in whose context it is called.

Note: This function is supported only for a joint client/server.

Return Values This operation returns a reference to the POA implementing the object in whose context it is called.

Exceptions If this operation is called outside the context of a POA-dispatched operation, a `PortableServer::NoContext` exception is raised.

POAManager Member Functions

Each POA object has an associated `POAManager` object. A POA manager may be associated with one or more POA objects. A POA manager encapsulates the processing state of the POAs with which it is associated. Using operations on the POA manager, an application can cause requests for those POAs to be queued or discarded, and can cause the POAs to be deactivated.

POA managers are created and destroyed implicitly. Unless an explicit POA manager object is provided at POA creation time, a POA manager is created when a POA is created and is automatically associated with that POA. A POA manager object is implicitly destroyed when all of its associated POAs have been destroyed.

A POA manager has four possible processing states: *active*, *inactive*, *holding*, and *discarding*. The processing state determines the capabilities of the associated POAs and the disposition of requests received by those POAs.

A POA manager is created in the holding state. In that state, any invocations on its POA are queued until the POA manager enters the active state. This version of WLE supports only the ability to enter active and inactive states. That is, this version does not support the ability to return to holding state or to enter discarding state.

PortableServer::POAManager::activate

Synopsis Changes the state of the POA manager to *active*.

C++ Binding `void activate();`

Argument None.

Description This operation changes the state of the POA manager to *active*. Entering the *active* state enables the associated POAs to process requests.

Note: All parent POAs must also have POAManagers in the active state for this POA to process requests.

Note: This function is supported only for a joint client/server.

Return Values None.

Exceptions If this operation is issued while the POA manager is in the *inactive* state, the `PortableServer::POAManager::AdapterInactive` exception is raised.

PortableServer::POAManager::deactivate

Synopsis Changes the state of the POA manager to *inactive*.

C++ Binding `void deactivate (
 CORBA::Boolean etherealize_objects,
 CORBA::Boolean wait_for_completion);`

Argument `etherealize_objects`
 For WLE V4.2 software, this argument should always be set to FALSE.

`wait_for_completion`
 If this argument is TRUE, the deactivate operation returns only after all requests in process have completed. If this argument is FALSE, the deactivate operation returns after changing the state of the associated POAs.

Description This operation changes the state of the POA manager to *inactive*. Entering the inactive state causes the associated POAs to reject requests that have not begun to be executed, as well as any new requests.

Note: This release of WLE does not support multithreading. Hence, `wait_for_completion` should not be TRUE if the call is made in the context of an object invocation. That is, the POAManager cannot be set to inactive state if it is currently executing.

Note: This function is supported only for a joint client/server.

Return Values None.

Exceptions If issued while the POA manager is in the *inactive* state, the `PortableServer::POAManager::AdapterInactive` exception is raised.

POA Policy Member Objects

Interfaces derived from `CORBA::Policy` are used with the `POA::create_POA` operation to specify policies that apply to a POA. Policy objects are created using factory operations on any pre-existing POA, such as the root POA. Policy objects are specified when a POA is created. Policies may not be changed on an existing POA. Policies are *not* inherited from the parent POA.

PortableServer::LifespanPolicy

Synopsis Specifies the life span of objects to the `create_POA` operation.

Description Objects with the `LifespanPolicy` interface are obtained using the `POA::create_lifespan_policy` operation and are passed to the `POA::create_POA` operation to specify the life span of the objects implemented in the created POA. The following values can be supplied:

- ◆ **TRANSIENT**—The objects implemented in the POA cannot outlive the process in which they are first created.
- ◆ **PERSISTENT**—The objects implemented in the POA can outlive the process in which they are first created.

Persistent objects have a POA associated with them (the POA that created them). When the ORB receives a request on a persistent object, it searches for the matching POA, based on the names of the POA and all of its ancestors.

POA names must be unique within their enclosing scope (the parent POA). A portable program can assume that POA names used in other processes will not conflict with its own POA names.

If no `LifespanPolicy` object is passed to `create_POA`, the lifespan policy defaults to **TRANSIENT**.

Note: This function is supported only for a joint client/server.

Exceptions None.

PortableServer::IdAssignmentPolicy

Synopsis Specifies whether `ObjectIds` in the created POA are generated by the application or by the ORB.

Description Objects with the `IdAssignmentPolicy` interface are obtained using the `POA::create_id_assignment_policy` operation and are passed to the `POA::create_POA` operation to specify whether `ObjectIds` in the created POA are generated by the application or by the ORB. The following values can be supplied:

- ◆ `USER_ID`—objects created with that POA are assigned `ObjectIds` only by the application.
- ◆ `SYSTEM_ID`—objects created with that POA are assigned `ObjectIds` only by the POA. If the POA also has the `PERSISTENT` policy, assigned `ObjectIds` must be unique across all instantiations of the same POA.

If no `IdAssignmentPolicy` is specified at POA creation, the default is `SYSTEM_ID`.

Note: This function is supported only for a joint client/server.

Request Member Functions

The mapping of these member functions to C++ is as follows:

```
// C++
class Request
{
public:
    Object_ptr target() const;
    const char *operation() const;
    NamedValue_ptr result();
    NVList_ptr arguments();
    Environment_ptr env();
    ExceptionList_ptr exceptions();
    ContextList_ptr contexts();
    void ctx(Context_ptr);
    Context_ptr ctx() const

    // argument manipulation helper functions
    Any &add_in_arg();
    Any &add_in_arg(const char* name);
    Any &add_inout_arg():
    Any &add_inout_arg(const char* name);
    Any &add_out_arg():
    Any &add_out_arg(const char* name);
    void set_return_type(TypeCode_ptr tc);
    Any &return_value();

    void invoke();
    void send_oneway();
    void send_deferred();
    void get_response();
    Boolean poll_response();
};
```

Note: The `add_*_arg`, `set_return_type`, and `return_value` member functions are added as shortcuts for using the attribute-based accessors.

The following sections describe these member functions.

CORBA::Request::arguments

Synopsis Retrieves the argument list for the request.

C++ Binding `CORBA::NVList_ptr CORBA::Request::arguments () const;`

Arguments None.

Description This member function retrieves the argument list for the request. The arguments can be input, output, or both.

Return Values If the function succeeds, the value returned is a pointer to the list of arguments to the operation for the request. The returned argument list is owned by the Request object reference and should not be released.

 If the function does not succeed, an exception is thrown.

CORBA::Request::ctx(Context_ptr)

Synopsis Sets the Context object for the operation.

C++ Binding `void CORBA::Request::ctx (`
 `CORBA::Context_ptr CtxObject);`

Argument CtxObject
 The new value to which to set the Context object.

Description This member function sets the Context object for the operation.

Return Values None.

See Also `CORBA::Request::ctx()`

CORBA::Request::get_response

Synopsis Retrieves the response of a specific deferred synchronous request.

C++ Binding `void CORBA::Request::get_response ();`

Arguments None.

Description This member function retrieves the response of a specific request; it is used after a call to the `CORBA::Request::send_deferred` function or the `CORBA::Request::send_multiple_requests` function. If the request has not completed, the `CORBA::Request::get_response` function blocks until it does complete.

Return Values None.

See Also `CORBA::Request::send_deferred`

CORBA::Request::invoke

Synopsis Performs an invoke on the operation specified in the request.

C++ Binding `void CORBA::Request::invoke ();`

Arguments None.

Description This member function calls the Object Request Broker (ORB) to send the request to the appropriate server application.

Return Values None.

CORBA::Request::operation

Synopsis Retrieves the operation intended for the request.

C++ Binding `const char * CORBA::Request::operation () const;`

Arguments None.

Description This member function retrieves the operation intended for the request.

Return Values If the function succeeds, the value returned is a pointer to the operation intended for the object; the value can be 0 (zero). The memory returned is owned by the Request object and should not be freed.

 If the function does not succeed, an exception is thrown.

CORBA::Request::poll_response

Synopsis	Determines whether a deferred synchronous request has completed.
C++ Binding	<code>CORBA::Boolean CORBA::Request::poll_response ();</code>
Arguments	None.
Description	This member function determines whether the request has completed and returns immediately. You can use this call to check the state of the request. This member function can also be used to determine whether a call to <code>CORBA::Request::get_response</code> will block.
Return Values	If the function succeeds, the value returned is <code>CORBA_TRUE</code> if the response has already completed, and <code>CORBA_FALSE</code> if the response has not yet completed. If the function does not succeed, an exception is thrown.
See Also	<code>CORBA::ORB::get_next_response</code> <code>CORBA::ORB::poll_next_response</code> <code>CORBA::ORB::send_multiple_requests</code> <code>CORBA::Request::get_response</code> <code>CORBA::Request::send_deferred</code>

CORBA::Request::result

Synopsis Retrieves the result of the request.

C++ Binding `CORBA::NamedValue_ptr CORBA::Request::result ();`

Arguments None.

Description This member function retrieves the result of the request.

Return Values If the function succeeds, the value returned is a pointer to the result of the operation.
The returned result is owned by the Request object and should not be released.

If the function does not succeed, an exception is thrown.

CORBA::Request::env

Synopsis Retrieves the environment of the request.

C++ Binding `CORBA::Environment_ptr CORBA::Request::env ();`

Arguments None.

Description This member function retrieves the environment of the request.

Return Values If the function succeeds, the value returned is a pointer to the environment of the operation. The returned environment is owned by the Request object and should not be released.

 If the function does not succeed, an exception is thrown.

CORBA::Request::ctx

Synopsis Retrieves the context of the request.

C++ Binding `CORBA::context_ptr CORBA::Request::ctx ();`

Arguments None.

Description This member function retrieves the context of the request.

Return Values If the function succeeds, the value returned is a pointer to the context of the operation.
The returned context is owned by the Request object and should not be released.

If the function does not succeed, an exception is thrown.

CORBA::Request::contexts

Synopsis Retrieves the context lists for the request.

C++ Binding `CORBA::ContextList_ptr CORBA::Request::contexts ();`

Arguments None.

Description This member function retrieves the context lists for the request.

Return Values If the function succeeds, the value returned is a pointer to the context lists for the operation. The returned context list is owned by the Request object and should not be released.

 If the function does not succeed, an exception is thrown.

CORBA::Request::exceptions

Synopsis Retrieves the exception lists for the request.

C++ Binding `CORBA::ExceptionList_ptr CORBA::Request::exceptions ();`

Arguments None.

Description This member function retrieves the exception lists for the request.

Return Values If the function succeeds, the value returned is a pointer to the exception list for the request. The returned exception list is owned by the Request object and should not be released.

 If the function does not succeed, an exception is thrown.

CORBA::Request::target

Synopsis Retrieves the target object reference for the request.

C++ Binding `CORBA::Object_ptr CORBA::Request::target () const;`

Arguments None.

Description This member function retrieves the target object reference for the request.

Return Values If the function succeeds, the value returned is a pointer to the target object of the operation. The returned value is owned by the Request object and should not be released.

 If the function does not succeed, an exception is thrown.

CORBA::Request::send_deferred

Synopsis Initiates a deferred synchronous request.

C++ Binding `void CORBA::Request::send_deferred ();`

Arguments None.

Description This member function initiates a deferred synchronous request. You use this function when a response is expected and in conjunction with the `CORBA::Request::get_response` function.

Return Values None.

See Also `CORBA::ORB::get_next_response`
 `CORBA::ORB::poll_next_response`
 `CORBA::ORB::send_multiple_requests`
 `CORBA::Request::get_response`
 `CORBA::Request::poll_response`
 `CORBA::Request::send_oneway`

CORBA::Request::send_oneway

Synopsis Initiates a one-way request.

C++ Binding `void CORBA::Request::send_oneway ();`

Arguments None.

Description This member function initiates a one-way request; it does not expect a response.

Return Values None.

See Also `CORBA::ORB::send_multiple_requests`
 `CORBA::Request::send_deferred`

Strings

The mapping of these functions to C++ is as follows:

```
// C++

namespace CORBA {
    static char *    string_alloc(ULong len);
    static char *    string_dup (const char *);
    static void      string_free(char *);
    ...
}
```

Note: A static array of `char` in C++ decays to a `char*`. Therefore, care must be taken when assigning a static array to a `String_var`, because the `String_var` assumes that the pointer points to data allocated via `string_alloc`, and thus eventually attempts to free it using `string_free`.

This behavior has changed in ANSI/ISO C++, where string literals are `const char*`, not `char*`. However, since most C++ compilers do not yet implement this change, portable programs must heed the advice given here.

The following sections describe the functions that manage memory allocated to strings.

CORBA::string_alloc

Synopsis Allocates memory for a string.

C++ Binding `char * CORBA::string_alloc(ULong len);`

Argument `len`
 The length of the string for which to allocate memory.

Description This member function dynamically allocates memory for a string, or returns a Nil pointer if it cannot perform the allocation. It allocates `len+1` characters so that the resulting string has enough space to hold a trailing NULL character. Free the memory allocated by this member function by calling the `CORBA::string_free` member function.

 This function does not throw CORBA exceptions.

Return Values If the function succeeds, the return value is a pointer to the newly allocated memory for the string object; if the function fails, the return value is a Nil pointer.

Example `char* s = CORBA::string_alloc(10);`

See Also `CORBA::string_free`
 `CORBA::string_dup`

CORBA::string_dup

Synopsis Makes a copy of a string.

C++ Binding `char * CORBA::string_dup (const char * Str);`

Argument `Str`
 The address of the string to be copied.

Description This function dynamically allocates enough memory to hold a copy of its string argument, including the NULL character, copies the string argument into that memory, and returns a pointer to the new string.

 This function does not throw CORBA exceptions.

Return Values If the function succeeds, the return value is a pointer to the new string; if the function fails, the return value is a Nil pointer.

Example `char* s = CORBA::string_dup("hello world");`

See Also `CORBA::string_free`
 `CORBA::string_alloc`

CORBA::string_free

Synopsis Frees memory allocated to a string.

C++ Binding `void CORBA::string_free(char * Str);`

Argument `Str`
 The address of the memory to be deallocated.

Description This member function deallocates memory that was previously allocated to a string using the `CORBA::string_alloc()` or `CORBA::string_dup()` member functions. Passing a Nil pointer to this function is acceptable and results in no action being performed.

 This function may not throw CORBA exceptions.

Return Values None.

Example `char* s = CORBA::string_dup("hello world");
CORBA::string_free(s);`

See Also `CORBA::string_alloc`
 `CORBA::string_dup`

TypeCode Member Functions

A TypeCode represents OMG IDL type information.

No constructors for TypeCodes are defined. However, in addition to the mapped interface, for each basic and defined OMG IDL type, an implementation provides access to a TypeCode pseudo-object reference (TypeCode_ptr) of the form `_tc_<type>` that may be used to set types in Any, as arguments for `equal`, and so on. In the names of these TypeCode reference constants, `<type>` refers to the local name of the type within its defining scope. Each C++ `_tc_<type>` constant is defined at the same scoping level as its matching type.

Like all other serverless objects, the C++ mapping for TypeCode provides a `_nil()` operation that returns a nil object reference for a TypeCode. This operation can be used to initialize TypeCode references embedded within constructed types. However, a nil TypeCode reference may never be passed as an argument to an operation, since TypeCodes are effectively passed as values, not as object references.

The mapping of these member functions to C++ is as follows:

```
class CORBA
{
    class TypeCode
    {
    public:
        class Bounds { ... };
        class BadKind { ... };

        Boolean equal(TypeCode_ptr) const;
        TCKind kind() const;
        Long param_count() const;
        Any *parameter(Long) const;
        RepositoryId id () const;
    }; // TypeCode
}; // CORBA
```

Memory Management

TypeCode has the following special memory management rule:

- ◆ Ownership of the return values of the `id` function is maintained by the TypeCode; these return values must not be freed by the caller.

The following sections describe these member functions.

CORBA::TypeCode::equal

Synopsis	Determines whether two TypeCode objects are equal.
C++ Binding	<pre>CORBA::Boolean CORBA::TypeCode::equal (CORBA::TypeCode_ptr TypeCodeObj) const;</pre>
Argument	<p>TypeCodeObj</p> <p>A pointer to a TypeCode object with which to make the comparison.</p>
Description	This member function determines whether a TypeCode object is equal to the input parameter, TypeCodeObj.
Return Values	<p>If the TypeCode object is equal to the TypeCodeObj parameter, CORBA_TRUE is returned.</p> <p>If the TypeCode object is not equal to the TypeCodeObj parameter, CORBA_FALSE is returned.</p> <p>If the function does not succeed, an exception is thrown.</p>

CORBA::TypeCode::id

Synopsis Returns the ID for the TypeCode.

C++ Binding `CORBA::RepositoryId CORBA::TypeCode::id () const;`

Arguments None.

Description This member function returns the ID for the TypeCode.

Return Values Repository ID for the TypeCode.

CORBA::TypeCode::kind

Synopsis	Retrieves the kind of data contained in the TypeCode object reference.
C++ Binding	<code>CORBA::TCKind CORBA::TypeCode::kind () const;</code>
Arguments	None.
Description	This member function retrieves the <code>kind</code> attribute of the <code>CORBA::TypeCode</code> class, which specifies the kind of data contained in the TypeCode object reference.
Return Values	If the member function succeeds, it returns the kind of data contained in the TypeCode object reference. For a list of the TypeCode kinds and their parameters, see Table 12-1. If the member function does not succeed, an exception is thrown.

Table 12-1 Legal Typecode Kinds and Parameters

Typecode Kind	Parameters List
<code>CORBA::tk_null</code>	*NONE*
<code>CORBA::tk_void</code>	*NONE*
<code>CORBA::tk_short</code>	*NONE*
<code>CORBA::tk_long</code>	*NONE*
<code>CORBA::tk_long</code>	*NONE*
<code>CORBA::tk_ushort</code>	*NONE*
<code>CORBA::tk_ulong</code>	*NONE*
<code>CORBA::tk_float</code>	*NONE*
<code>CORBA::tk_double</code>	*NONE*
<code>CORBA::tk_boolean</code>	*NONE*
<code>CORBA::tk_char</code>	*NONE*
<code>CORBA::tk_octet</code>	*NONE*
<code>CORBA::tk_Typecode</code>	*NONE*
<code>CORBA::tk_Principal</code>	*NONE*

Table 12-1 Legal Typecode Kinds and Parameters (Continued)

Typecode Kind	Parameters List
CORBA::tk_objref	{interface_id}
CORBA::tk_struct	{struct-name, member-name, TypeCode, ... (repeat pairs)}
CORBA::tk_union	{union-name, switch-TypeCode, label-value, member-name, enum-id, ...}
CORBA::tk_enum	{enum-name, enum-id, ...}
CORBA::tk_string	{maxlen-integer}
CORBA::tk_sequence	{TypeCode, maxlen-integer}
CORBA::tk_array	{TypeCode, length-integer}

CORBA::TypeCode::param_count

Synopsis	Retrieves the number of parameters for the TypeCode object reference.
C++ Binding	<code>CORBA::Long CORBA::TypeCode::param_count () const;</code>
Arguments	None.
Description	This member function retrieves the parameter attribute of the <code>CORBA::TypeCode</code> class, which specifies the number of parameters for the TypeCode object reference. For a list of parameters of each kind, see Table 12-1.
Return Values	<p>If the function succeeds, it returns the number of parameters contained in the TypeCode object reference.</p> <p>If the function does not succeed, an exception is thrown.</p>

CORBA::TypeCode::parameter

Synopsis	Retrieves a parameter specified by the index input argument.
C++ Binding	<pre>CORBA::Any * CORBA::TypeCode::parameter (CORBA::Long Index) const;</pre>
Argument	<p>Index</p> <p>An index to the parameter list, used to determine which parameter to retrieve.</p>
Description	This member function retrieves a parameter specified by the index input argument. For a list of parameters of each kind, see Table 12-1.
Return Values	<p>If the member function succeeds, the return value is a pointer to the parameter specified by the index input argument.</p> <p>If the member function does not succeed, an exception is thrown.</p>

Exception Member Functions

The WLE system supports the throwing and catching of exceptions.

Descriptions of exception member functions follow:

```
CORBA::SystemException::SystemException ()
```

This is the default constructor for the `CORBA::SystemException` class. Minor code is initialized to 0 (zero) and the completion status is set to `COMPLETED_NO`.

```
CORBA::SystemException::SystemException (  
    const CORBA::SystemException & Se)
```

This is the copy constructor for the `CORBA::SystemException` class.

```
CORBA::SystemException::SystemException(  
    CORBA::ULong Minor, CORBA::CompletionStatus Status)
```

This constructor for the `CORBA::SystemException` class sets the minor code and completion status.

Explanations of the arguments are as follows:

Minor

The minor code for the Exception object. The minor field is an implementation-specific value used by the ORB to identify the exception. The WLE minor field definitions can be found in the file `orbminor.h`.

Status

The completion status for the Exception object. The values are as follows:

`CORBA::COMPLETED_YES`

`CORBA::COMPLETED_NO`

`CORBA::COMPLETED_MAYBE`

```
CORBA::SystemException::~~SystemException ()
```

This is the destructor for the `CORBA::SystemException` class. It frees any memory used for the Exception object.

```
CORBA::SystemException CORBA::SystemException::operator =  
    const CORBA::SystemException Se)
```

This assignment operator copies exception information from the source exception. The `Se` argument specifies the `SystemException` object that is to be copied by this operator.

```
CORBA::CompletionStatus CORBA::SystemException::completed()
```

This member function returns the completion status for this exception.

```
CORBA::SystemException::completed(  
    CORBA::CompletionStatus Completed)
```

This member function sets the completion status for this exception. The `Completed` argument specifies the completion status for this exception.

```
CORBA::ULong CORBA::SystemException::minor()
```

This member function returns the minor code for this exception.

```
CORBA::SystemException::minor (CORBA::ULong Minor)
```

This member function sets the minor code for this exception. The `minor` argument specifies the new minor code for this exception. The `minor` field is an implementation-specific value used by the application to identify the exception.

```
CORBA::SystemException * CORBA::SystemException::_narrow (  
    CORBA::Exception_ptr Exc)
```

This member function determines whether a specified exception can be narrowed to a system exception. The `Exc` argument specifies the exception to be narrowed.

If the specified exception is a system exception, this member function returns a pointer to the system exception. If the specified exception is not a system exception, the function returns 0 (zero).

```
CORBA::UserException * CORBA::UserException::_narrow(  
    CORBA::Exception_ptr Exc)
```

This member function determines whether a specified exception can be narrowed to a user exception. The `Exc` argument specifies the exception to be narrowed.

If the specified exception is a user exception, this member function returns a pointer to the user exception. If the specified exception is not a user exception, the function returns 0 (zero).

Standard Exceptions

This section presents the standard exceptions defined for the ORB. These exception identifiers may be returned as a result of any operation invocation, regardless of the interface specification. Standard exceptions are not listed in `raises` expressions.

To bound the complexity in handling the standard exceptions, the set of standard exceptions is kept to a tractable size. This constraint forces the definition of equivalence classes of exceptions, rather than enumerating many similar exceptions.

For example, an operation invocation can fail at many different points due to the inability to allocate dynamic memory. Rather than enumerate several different exceptions that correspond to the different ways that memory allocation failure causes the exception (during marshaling, unmarshaling, in the client, in the object implementation, allocating network packets, and so forth), a single exception corresponding to dynamic memory allocation failure is defined. Each standard exception includes a minor code to designate the subcategory of the exception; the assignment of values to the minor codes is left to each ORB implementation.

Each standard exception also includes a `completion_status` code, which takes one of the following values:

CORBA: :COMPLETED_YES

The object implementation completed processing prior to the exception being raised.

CORBA: :COMPLETED_NO

The object implementation was not initiated prior to the exception being raised.

CORBA: :COMPLETED_MAYBE

The status of implementation completion is unknown.

Exception Definitions

The standard exceptions are defined below. Clients must be prepared to handle system exceptions that are not on this list, both because future versions of this specification may define additional standard exceptions, and because ORB implementations may raise nonstandard system exceptions. For more information about exceptions, see *System Messages*.

Table 12-2 defines the exceptions.

Table 12-2 Exception Definitions

Exception	Description
CORBA::UNKNOWN	The unknown exception.
CORBA::BAD_PARAM	An invalid parameter was passed.
CORBA::NO_MEMORY	Dynamic memory allocation failure.
CORBA::IMP_LIMIT	Violated implementation limit.
CORBA::COMM_FAILURE	Communication failure.
CORBA::INV_OBJREF	Invalid object reference.
CORBA::NO_PERMISSION	No permission for attempted operation.
CORBA::INTERNAL	ORB internal error.
CORBA::MARSHAL	Error marshalling parameter/result.
CORBA::INITIALIZE	ORB initialization failure.
CORBA::NO_IMPLEMENT	Operation implementation unavailable.
CORBA::BAD_TYPECODE	Bad typecode.
CORBA::BAD_OPERATION	Invalid operation.
CORBA::NO_RESOURCES	Insufficient resources for request.
CORBA::NO_RESPONSE	Response to request not yet available.
CORBA::PERSIST_STORE	Persistent storage failure.

Table 12-2 Exception Definitions (Continued)

Exception	Description
<code>CORBA::BAD_INV_ORDER</code>	Routine invocations out of order.
<code>CORBA::TRANSIENT</code>	Transient failure; reissue request.
<code>CORBA::FREE_MEM</code>	Cannot free memory.
<code>CORBA::INV_IDENT</code>	Invalid identifier syntax.
<code>CORBA::INV_FLAG</code>	Invalid flag was specified.
<code>CORBA::INTF_REPOS</code>	Error accessing interface repository.
<code>CORBA::BAD_CONTEXT</code>	Error processing context object.
<code>CORBA::OBJ_ADAPTER</code>	Failure detected by object adapter.
<code>CORBA::DATA_CONVERSION</code>	Data conversion error.
<code>CORBA::OBJECT_NOT_EXIST</code>	Non-existent object; delete reference.
<code>CORBA::TRANSACTION_REQUIRED</code>	Transaction required.
<code>CORBA::TRANSACTION_ROLLEDBACK</code>	Transaction rolled back.
<code>CORBA::INVALID_TRANSACTION</code>	Invalid transaction.

Object Nonexistence

The `CORBA::OBJECT_NOT_EXIST` exception is raised whenever an invocation on a deleted object is performed. It is an authoritative “hard” fault report. Anyone receiving it is allowed (even expected) to delete all copies of this object reference and to perform other appropriate “final recovery” style procedures.

Transaction Exceptions

The `CORBA::TRANSACTION_REQUIRED` exception indicates that the request carried a null transaction context, but an active transaction is required.

The `CORBA::TRANSACTION_ROLLEDBACK` exception indicates that the transaction associated with the request has already been rolled back or marked to roll back. Thus, the requested operation either could not be performed or was not performed because further computation on behalf of the transaction would be fruitless.

The `CORBA::INVALID_TRANSACTION` indicates that the request carried an invalid transaction context. For example, this exception could be raised if an error occurred when trying to register a resource.

ExceptionList Member Functions

The `ExceptionList` member functions allow a client or server application to provide a list of `TypeCodes` for all user-defined exceptions that may result when the `Request` is invoked. For a description of the `Request` member functions, see the section “`Request` Member Functions” on page 12-112.

The mapping of these member functions to C++ is as follows:

```
class CORBA
{
    class ExceptionList
    {
    public:
        Ulong count ();
        void add(TypeCode_ptr tc);
        void add_consume(TypeCode_ptr tc);
        TypeCode_ptr item(Ulong index);
        Status remove(Ulong index);
    }; // ExceptionList
} // CORBA
```

CORBA::ExceptionList::count

Synopsis Retrieves the current number of items in the list.

C++ Binding `Ulong count ();`

Arguments None.

Description This member function retrieves the current number of items in the list.

Return Values If the function succeeds, the returned value is the number of items in the list. If the list has just been created, and no ExceptionList objects have been added, this function returns 0 (zero).

Exception If the function does not succeed, an exception is thrown.

CORBA::ExceptionList::add

Synopsis	Constructs a ExceptionList object with an unnamed item, setting only the <code>flags</code> attribute.
C++ Binding	<pre>void add(TypeCode_ptr tc);</pre>
Arguments	<p><code>tc</code></p> <p>Defines the memory location referred to by <code>TypeCode_ptr</code>.</p>
Description	<p>This member function constructs an ExceptionList object with an unnamed item, setting only the <code>flags</code> attribute.</p> <p>The ExceptionList object grows dynamically; your application does not need to track its size.</p>
Return Values	If the function succeeds, the return value is a pointer to the newly created ExceptionList object.
Exception	If the member function does not succeed, a <code>CORBA::NO_MEMORY</code> exception is thrown.
See Also	<p><code>CORBA::ExceptionList::add_consume</code></p> <p><code>CORBA::ExceptionList::count</code></p> <p><code>CORBA::ExceptionList::item</code></p> <p><code>CORBA::ExceptionList::remove</code></p>

CORBA::ExceptionList::add_consume

Synopsis	Constructs an ExceptionList object.
C++ Binding	<code>void add_consume(TypeCode_ptr tc);</code>
Arguments	<code>tc</code> The memory location to be assumed.
Description	This member function constructs an ExceptionList object. The ExceptionList object grows dynamically; your application does not need to track its size.
Return Values	If the function succeeds, the return value is a pointer to the newly created ExceptionList object.
Exceptions	If the member function does not succeed, an exception is raised.
See Also	CORBA::ExceptionList::add CORBA::ExceptionList::count CORBA::ExceptionList::item CORBA::ExceptionList::remove

CORBA::ExceptionList::item

Synopsis	Retrieves a pointer to the ExceptionList object, based on the index passed in.
C++ Binding	<code>TypeCode_ptr item(ULong index);</code>
Argument	<code>index</code> The index into the ExceptionList object. The indexing is zero-based.
Description	This member function retrieves a pointer to an ExceptionList object, based on the index passed in. The function uses zero-based indexing.
Return Values	If the function succeeds, the return value is a pointer to the ExceptionList object.
Exceptions	If the function does not succeed, the <code>BAD_PARAM</code> exception is thrown.
See Also	<code>CORBA::ExceptionList::add</code> <code>CORBA::ExceptionList::add_consume</code> <code>CORBA::ExceptionList::count</code> <code>CORBA::ExceptionList::remove</code>

CORBA::ExceptionList::remove

Synopsis	Removes the item at the specified index, frees any associated memory, and reorders the remaining items on the list.
C++ Binding	<code>Status remove(ULong index);</code>
Argument	Index The index into the ContextList object. The indexing is zero-based.
Description	This member function removes the item at the specified index, frees any associated memory, and reorders the remaining items on the list.
Return Values	None.
Exceptions	If the function does not succeed, the <code>BAD_PARAM</code> exception is thrown.
See Also	<code>CORBA::ExceptionList::add</code> <code>CORBA::ExceptionList::add_consume</code> <code>CORBA::ExceptionList::count</code> <code>CORBA::ExceptionList::item</code>

13 Server-side Mapping

Server-side mapping refers to the portability constraints for an object implementation written in C++. The term *server* is not meant to restrict implementations to situations in which method invocations cross-address space or machine boundaries. This mapping addresses any implementation of an Object Management Group (OMG) Interface Definition Language (IDL) interface.

Note: The information in this chapter is based on the *Common Object Request Broker: Architecture and Specification*. Revision 2.2, February 1998, published by the Object Management Group (OMG). Used with permission by OMG.

Implementing Interfaces

To define an implementation in C++, you define a C++ class with any valid C++ name. For each operation in the interface, the class defines a nonstatic member function with the mapped name of the operation (the mapped name is the same as the OMG IDL identifier).

The server application mapping specifies two alternative relationships between the implementation class supplied by the application and the generated class or classes for the interface. Specifically, the mapping requires support for both inheritance-based relationships and delegation-based relationships. Conforming applications may use either or both of these alternatives. This release of the WLE software supports both inheritance-based and delegation-based relationships.

Inheritance-based Interface Implementation

In the inheritance-based interface implementation approach, the implementation classes are derived from a generated base class based on the OMG IDL interface definition. The generated base classes are known as skeleton classes, and the derived classes are known as implementation classes. Each operation of the interface has a corresponding virtual member function declared in the skeleton class. The generated skeleton class is partially opaque to the programmer, though it will contain a member function corresponding to each operation in the interface. The signature of the member function is identical to that of the generated client stub class.

To implement this interface using inheritance, a programmer must derive from this skeleton class and implement each of the operations in the OMG IDL interface. To allow portable implementations to multiple inheritances from both skeleton classes and implementation classes for other base interfaces without error or ambiguity, the `Tobj_ServantBase` class must be a virtual base class of the skeleton, and the `PortableServer::ServantBase` class must be a virtual base class of the `Tobj_ServantBase` class. The inheritance among the implementation class, the skeleton class, the `Tobj_ServantBase` class, and the `PortableServer::ServantBase` class must all be public virtual.

The implementation class or servant must only derive directly from a single generated skeleton class. Direct derivation from multiple skeleton classes could result in ambiguous errors due to multiple definitions of the `_this()` operation. This should not be a limitation, however, since CORBA objects have only a single most-derived interface. C++ servants that are intended to support multiple interface types can utilize the delegation-based interface implementation approach. See Listing 13-1 for an example of OMG IDL that uses interface inheritance.

Listing 13-1 OMG IDL that Uses Interface Inheritance

```
// IDL
interface A
{
    short op1() ;
    void op2(in long val) ;
};
```

Listing 13-2 Interface Class A

```
// C++
class A : public virtual CORBA::Object
{
    public:
        virtual CORBA::Short op1 ();
        virtual void op2 (CORBA::Long val);
};
```

On the server side, a skeleton class is generated. This class is partially opaque to the programmer, though it does contain a member function corresponding to each operation in the interface.

For the Portable Object Adapter (POA), the name of the skeleton class is formed by prepending the string “POA_” to the fully scoped name of the corresponding interface, and the class is directly derived from the servant base class `Tobj_ServantBase`. The C++ mapping for `Tobj_ServantBase` is as follows:

```
// C++
class Tobj_ServantBase
{
    public:
        virtual void activate_object(const char* stroid);
        virtual void deactivate_object (
            const char* stroid,
            TobjS::DeactivateReasonValue reason
        );
}
```

The `activate_object()` and `deactivate_object()` member functions are described in detail in the sections “`Tobj_ServantBase::activate_object()`” on page 3-35 and “`Tobj_ServantBase::deactivate_object()`” on page 3-38.

The skeleton class for interface A shown above would appear as shown in Listing 13-3.

Listing 13-3 Skeleton Class for Interface A

```
// C++
class POA_A : public Tobj_ServantBase
{
    public:
        // ... server-side ORB-implementation-specific
        // goes here...

        virtual CORBA::Short op1 () = 0;
        virtual void op2 (CORBA::Long val) = 0;
        //...
};
```

If interface A were defined within a module rather than at global scope (for example, `Mod::A`), the name of its skeleton class would be `POA_Mod::A`. This helps to separate server application skeleton declarations and definitions from C++ code generated for the client.

To implement this interface using inheritance, you must derive from this skeleton class and implement each of the operations in the corresponding OMG IDL interface. An implementation class declaration for interface A would take the form shown in Listing 13-4.

Listing 13-4 Interface A Implementation Class Declaration

```
// C++
class A_impl : public POA_A
{
    public:
        CORBA::Short op1();
        void op2(CORBA::Long val);
        ...
};
```

Delegation-based Interface Implementation

The delegation-based interface implementation approach is an alternative to using inheritance when implementing CORBA objects. This approach is used when the overhead of inheritance is too high or cannot be used. For example, due to the invasive nature of inheritance, implementing objects using existing legacy code might be impossible if inheritance for some global class were required. Instead, delegation can be used to solve these types of problems. Delegation is a more natural fit doing object implementations when the Process-Entity design pattern is used. In this pattern, the Process object would delegate operations onto one or more entity objects.

In the delegation-based approach, the implementation does not inherit from a skeleton class. Instead, the implementation can be coded as required for the application, and a wrapper object will delegate upcalls to that implementation. This “wrapper object,” called a *tie*, is generated by the IDL compiler, along with the same skeleton class used for the inheritance approach. The generated *tie* class is partially opaque to the programmer, though, like the skeleton, it provides a method corresponding to each OMG IDL operation for the associated interface. The name of the generated *tie* class is the same as the generated skeleton class with the addition that the string `_tie` is appended to the end of the class name.

An instance of the `tie` class is the servant, not the C++ object being delegated to by the `tie` object, that is passed as the argument to the operations that require a `Servant` argument. It should also be noted that the tied object has no access to the `_this()` operation, nor should it access data members directly.

A type-safe tie class is implemented using C++ templates. The code shown in Listing 13-5 illustrates a tie class generated from the Derived interface in the previous OMG IDL example.

Listing 13-5 tie Class Generated from the Derived Interface

```
// C++
template <class T>
class POA_A_tie : public POA_A {
public:
    POA_A_tie(T& t)
        : _ptr(&t), _poa(PortableServer::POA::_nil()), _rel(0) {}
    POA_A_tie(T& t, PortableServer::POA_ptr poa)
        : _ptr(&t), _poa(PortableServer::POA::_duplicate(poa)), _rel(0) {}
    POA_A_tie(T* tp, CORBA::Boolean release = 1)
        : _ptr(tp), _poa(PortableServer::POA::_nil()), _rel(release) {}
    POA_A_tie(T* tp, PortableServer::POA_ptr poa, CORBA::Boolean release = 1)
        : _ptr(tp), _poa(PortableServer::POA::_duplicate(poa)), _rel(release) {}
    ~POA_A_tie()
    { CORBA::release(_poa);
      if (_rel) delete _ptr;
    }

    // tie-specific functions
    T* _tied_object () {return _ptr;}
    void _tied_object(T& obj)
    { if (_rel) delete _ptr;
      _ptr = &obj;
      _rel = 0;
    }
    void _tied_object(T* obj, CORBA::Boolean release = 1)
    { if (_rel) delete _ptr;
      _ptr = obj;
      _rel = release;
    }

    CORBA::Boolean _is_owner() { return _rel; }
    void _is_owner (CORBA::Boolean b) { _rel = b; }

    // IDL operations*****
    CORBA::Short opl ()
    {
        return _ptr->opl ();
    }
}
```

```

void op2 (CORBA::Long val)
{
    _ptr->op2 (val);
}
// *****

// override ServantBase operations
PortableServer::POA_ptr _default_POA()
{
    if (!CORBA::is_nil(_poa))
    {
        return _poa;
    }
    else {
#ifdef WIN32
        return ServantBase::_default_POA();
#else
        return PortableServer::ServantBase::_default_POA();
#endif
    }
}

private:
    T* _ptr;
    PortableServer::POA_ptr _poa;
    CORBA::Boolean _rel;

    // copy and assignment not allowed
    POA_A_tie (const POA_A_tie<T> &);
    void operator=(const POA_A_tie<T> &);
};

```

This class definition is a template generated by the IDL compiler. You typically use it by first getting a pointer to the legacy class and then instantiating the tie class with that pointer. For example,

```

Old::Legacy * legacy = new Old::Legacy( oid);
POA_A_tie<Old::Legacy> * A_servant_ptr =
    new POA_A_tie<Old::Legacy>( legacy );

```

As you can see, the tie class contains definitions for the op1 and op2 operations of the interface that assume that the legacy class has operations with the same signatures as those given in the IDL. If this is the case, you can use the tie class file as is, letting it delegate exactly. It is more likely, however, that the legacy class will not have identical

signatures or you may have to do more than a single function call. In that case, it is your job to replace the code for `op1` and `op2` in this generated code. The code for each operation typically makes invocations on the legacy class using the tie class variable `_ptr`, which contains the pointer to the legacy class. For example, you might change the following lines:

```
CORBA::Short op1 () {return _ptr->op1 (); }  
void op2 (CORBA::Long val) {_ptr->op2 (val); }
```

to the following:

```
CORBA::Short op1 ()  
{  
    return _ptr->op37 ();  
}  
  
void op2 (CORBA::Long val)  
{  
    CORBA::Long temp;  
    temp = val + 15;  
    _ptr->lookup(val, temp, 43);  
}
```

An instance of this template class performs the task of delegation. When the template is instantiated with a class type that provides the operation of the `Derived` interface, then the `POA_Derived_tie` class will delegate all operations to an instance of that implementation class. A reference or pointer to the actual implementation object is passed to the appropriate *tie* constructor when an instance of the `POA_Derived_tie` class is created. When a request is invoked on it, the *tie* servant will just delegate the request by calling the corresponding method on the implementation class.

The use of templates for *tie* classes allows the application developer to provide specializations for some or all of the template's operations for a given instantiation of the template. This allows the application to use legacy classes for tied object types, where the operation signatures of the tied object will differ from that of the *tie* class.

Implementing Operations

The signature of an implementation member function is the mapped signature of the OMG IDL operation. Unlike the client-side mapping, the OMG specifies that the function header for the server-side mapping include the appropriate exception specification. An example of this is shown in Listing 13-6.

Listing 13-6 Exception Specification

```
// IDL
interface A
{
    exception B {};
    void f() raises(B);
};

// C++
class MyA : public virtual POA_A
{
public:
    void f();
    ...
};
```

Since all operations and attributes may raise CORBA system exceptions, `CORBA::SystemException` must appear in all exception specifications, even when an operation has no `raises` clause.

Note: Because of the differences in C++ compilers, it is best to leave out the "throw declaration" in the method signature. Some systems cause the application server to crash if an undeclared exception is thrown in a method that has declared the exceptions it will throw.

Within a member function, the "this" pointer refers to the implementation object's data as defined by the class. In addition to accessing the data, a member function may implicitly call another member function defined by the same class. An example of this is shown in Listing 13-7.

Listing 13-7 Calling Another Member Function

```
// IDL
interface A
{
    void f();
    void g();
};

// C++
class MyA : public virtual POA_A
{
public:
    void f();
    void g();
private:
    long x_;
};

void
MyA::f()
{
    x_ = 3;
    g();
}
```

When a servant member function is invoked in this manner, it is being called simply as a C++ member function, not as the implementation of an operation on a CORBA object.