

# Objectivity/C++ Programmer's Reference

Release 6.0

#### **Objectivity/C++ Programmer's Reference**

Part Number: 60-CPPRF-0

Release 6.0, September 22, 2000

The information in this document is subject to change without notice. Objectivity, Inc. assumes no responsibility for any errors that may appear in this document.

Copyright 2000 by Objectivity, Inc. All rights reserved. This document may not be copied, photocopied, reproduced, translated, or converted to any electronic or machine-readable form in whole or in part without prior written approval of Objectivity, Inc.

Objectivity and Objectivity/DB are registered trademarks of Objectivity, Inc. Objectivity/DB Fault Tolerant Option, Objectivity/FTO, Objectivity/DB Data Replication Option, Objectivity/DRO, Objectivity/DB Hot Failover, Objectivity/DB In-Process Lock Server, Objectivity/IPLS, Objectivity/DB Open File System, Objectivity/OFS, Objectivity/DB Secure Framework, Objectivity/Secure, Objectivity/C++, Objectivity/C++ Data Definition Language, Objectivity/DDL, Objectivity/C++ Active Schema, Objectivity/C++ Standard Template Library, Objectivity/C++ STL, Objectivity/C++ Spatial Index Framework, Objectivity/Spatial, Objectivity for Java, Objectivity/Smalltalk, Objectivity/SQL++, Objectivity/SQL++ ODBC Driver, Objectivity/ODBC, and Objectivity Event Notification Services are trademarks of Objectivity, Inc. Standards<ToolKit> is a trademark of ObjectSpace, Inc. Other trademarks and products are the property of their respective owners.

ODMG information in this document is based in whole or in part on material from *The Object Database Standard: ODMG 2.0*, edited by R.G.G. Cattell, and is reprinted with permission of Morgan Kaufmann Publishers. Copyright 1997 by Morgan Kaufmann Publishers.

The software and information contained herein are proprietary to, and comprise valuable trade secrets of, Objectivity, Inc., which intends to preserve as trade secrets such software and information. This software is furnished pursuant to a written license agreement and may be used, copied, transmitted, and stored only in accordance with the terms of such license and with the inclusion of the above copyright notice. This software and information or any other copies thereof may not be provided or otherwise made available to any other person.

U. S. Government Restricted Rights: Use, duplication or disclosure of the software or other information by the U. S. Government or any unit or agency thereof is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and the Government is acquiring only restricted rights in the software and limited rights in any technical data provided (as such terms are defined in such clause of the DFARS). If the software or other information is supplied to any unit or agency of the U. S. other than the Department of Defense, the Government's rights will be as defined in clause 52.227-19(c)(2) of the FAR or, in the case of NASA, in clause 18-52.227-86 (d) of the NASA Supplement to the FAR.

# Contents

About This Book	7
Audience	7
Organization	7
<b>Conventions and Abbreviations</b>	8
Getting Help	9
Objectivity/C++ Programming Interface	11
Global Names	23
appClass Class	81
d_Database Class	93
d_Date Class	99
d_Interval Class	111
d_Iterator< <i>element_type</i> > Class	119
d_Ref_Any Class	125
d_Time Class	131
d_Timestamp Class	145
ooAdmin Class	155
ooAPObj Class	157
ooBTree Class	163
ooCollection Class	173
ooCollectionIterator Class	185
ooCompare Class	197
ooContext Class	203
ooContObj Class	207

ooConvertInObject Class	217
ooConvertInOutObject Class	225
ooDBObj Class	233
ooDefaultContObj Class	239
ooEqualLookupField Class	241
ooFDObj Class	245
ooGCContObj Class	247
ooGeneObj Class	251
ooGreaterThanEqualLookupField Class	261
ooGreaterThanLookupField Class	265
ooHashAdmin Class	269
ooHashMap Class	273
ooHashSet Class	283
ooltr( <i>appClass</i> ) Class	293
ooltr(ooAPObj) Class	299
ooltr(ooContObj) Class	303
ooltr(ooDBObj) Class	309
ooltr(ooObj) Class	313
oojArray Class	319
oojArrayOfBoolean Class	323
oojArrayOfCharacter Class	327
oojArrayOfDouble Class	331
oojArrayOfFloat Class	335
oojArrayOfInt8 Class	339
oojArrayOfInt16 Class	343
oojArrayOfInt32 Class	347
oojArrayOfInt64 Class	351
oojArrayOfObject Class	355
oojDate Class	359
oojString Class	363
oojTime Class	365
oojTimestamp Class	369

ooKeyDesc Class	373
ooKeyField Class	379
ooLessThanEqualLookupField Class	387
ooLessThanLookupField Class	391
ooLookupFieldBase Class	395
ooLookupKey Class	399
ooMap Class	409
ooMapElem Class	423
ooMapItr Class	427
ooObj Class	431
ooOperatorSet Class	465
ooQuery Class	469
ooRefHandle(appClass) Classes	471
ooRefHandle(ooAPObj) Classes	489
ooRefHandle(ooContObj) Classes	509
<i>ooRefHandle</i> (ooDBObj) Classes	537
<i>ooRefHandle</i> (ooFDObj) Classes	571
<i>ooRefHandle</i> (ooObj) Classes	593
ooShortRef( <i>appClass</i> ) Class	631
ooShortRef(ooObj) Class	637
ooString(N) Class	645
ooTrans Class	653
ooTreeAdmin Class	661
ooTreeList Class	667
ooTreeMap Class	677
ooTreeSet Class	687
ooTVArrayT< <i>element_type</i> > Class	695
ooUtf8String Class	703
ooVArrayT< <i>element_type</i> > Class	707
ooVString Class	721

Topic Index	733
Classes Index	763
Functions and Macros Index	773
Types and Constants Index	787

# **About This Book**

This book, *Objectivity/C++ Programmer's Reference*, provides a detailed description of each public construct in the Objectivity/C++ programming interface. Objectivity/C++ enables a C++ application to create, store, and manipulate persistent data in an Objectivity/DB federated database.

You should use this book in conjunction with the Objectivity/C++ programmer's guide, which describes the various tasks an application can perform with the programming interface.

## Audience

This book assumes that you know how to program in C++.

# Organization

- The chapter "Objectivity/C++ Programming Interface" describes the interface's naming conventions and provides an overview of the public classes and global names defined by Objectivity/C++.
- The chapter "Global Names" provides detailed descriptions of the Objectivity/C++ global types, functions, macros, and variables.
- Each subsequent chapter describes an Objectivity/C++ class, providing detailed descriptions of its public member functions. The chapters are organized in alphabetical order by class name.
- The "Topic Index" lists the topics that are presented in this book.
- The "Classes Index" contains an alphabetical list of classes, with member functions under each class.
- The "Functions and Macros Index" contains an alphabetical list of all functions, including member functions.
- The "Types and Constants Index" lists all non-class types and constants.

# **Conventions and Abbreviations**

#### **Navigation**

Table of contents entries, index entries, cross-references, and <u>underlined</u> text are hypertext links.

### **Typographical Conventions**

oobackup	Command, literal parameter, code sample, filename, pathname, output on your screen, or Objectivity-defined identifier
installDir	Variable element (such as a filename or a parameter) for which you must substitute a value
Browse FD	Graphical user-interface label for a menu item or button
lock server	New term, book title, or emphasized word

## Abbreviations

(administration)	Feature intended for database administration tasks
(FTO)	Feature of the Objectivity/DB Fault Tolerant Option product
(DRO)	Feature of the Objectivity/DB Data Replication Option product
(IPLS)	Feature of the Objectivity/DB In-Process Lock Server Option product
(ODMG)	Feature conforming to the Object Database Management Group interface

#### **Command Syntax Symbols**

[]	Optional item. You may either enter or omit the enclosed item.
{}	Item that can be repeated.
	Alternative items. You should enter only one of the items separated by this symbol.
()	Logical group of items. The parentheses themselves are not part of the command syntax; do not type them.

## **Command and Code Conventions**

In code examples or commands, the continuation of a long line is indented. Omitted code is indicated with the ellipsis (...) symbol. "Enter" refers to the standard key (labelled either Enter or Return) for terminating a line of input.

# **Getting Help**

We have done our best to make sure all the information you need to install and operate Objectivity products is provided in the product documentation. However, we also realize problems requiring special attention sometimes occur.

## **Technical Support Web Site**

You can find answers to frequently asked questions, supported platforms, known bugs, and bug fixes on the Objectivity Technical Support web site. Send electronic mail or call Objectivity Customer Support to gain access to the site.

## How to Reach Objectivity Customer Support

You can contact Objectivity Customer Support by:

■ **Telephone:** Call 1.650.254.7100 *or* 1.800.SOS.OBJY (1.800.767.6259) Monday through Friday between 6:00 A.M. and 6:00 P.M. Pacific Time, and ask for Customer Support.

The toll-free 800 number can be dialed *only* within the 48 contiguous states of the United States and Canada.

- **Fax:** Send a fax to Objectivity at 1.650.254.7171.
- **Electronic Mail:** Send electronic mail to *help@objectivity.com*.

#### Before You Call

If you need help from Customer Support, please have the following information ready before you contact Objectivity:

- Your name, company name, address, telephone number, fax number, and email address
- Description of your workstation environment, including the type of workstation, its operating system version, compiler or interpreter, and windowing environment
- Information about the Objectivity product you are using, including the version of the Objectivity/DB libraries
- Detailed description of the problem you have encountered

# **Objectivity/C++ Programming Interface**

The Objectivity/C++ programming interface enables an application to represent Objectivity/DB objects and to manage the interaction between an application and Objectivity/DB. The Objectivity/C++ programming interface also has several extensions that are available when you purchase the corresponding add-on products.

This chapter provides:

- A list of the <u>extensions</u> to the core Objectivity/C++ programming interface.
- Descriptions of naming and usage <u>conventions</u> within the interface.
- An <u>overview</u> of the Objectivity/C++ programming constructs, according to their intended use.

## Interface Extensions

Each of the add-on products listed in the following table defines an extension to the Objectivity/C++ interface, allowing you to access the add-on product's capabilities in addition to Objectivity/DB's core capabilities. The classes, functions, and types belonging to these extensions are described in this book, although they are available only when you purchase the corresponding add-on product. When this book describes a features of a product-specific extension, the description is preceded by the indicator shown in the table:

Add-On Product	Extension Indicator
Objectivity/DB Fault Tolerant Option	(FTO)
Objectivity/DB Data Replication Option	(DRO)
Objectivity/DB In-Process Lock Server Option	(IPLS)

The following products also extend the Objectivity/C++ interface; these extensions are described in separate books, as listed below:

- Objectivity/C++ Active Schema
- *Objectivity/C++ Standard Template Library*

# **Interface Conventions**

This section describes the naming and usage conventions of the various programming constructs within the Objectivity/C++ programming interface.

#### **Naming Conventions**

The names of Objectivity/C++ programming constructs begin with the oo prefix. Multiple words in a name are indicated by mixed case, for example, ooGetOfflineMode. Names of non-ODMG items follow these conventions:

Item	Prefix	Description
Classes and non-class types	00	Names of all classes and most types begin with the oo prefix.The primitive types (for example, uint8, int16, float32) are the exception, but their alternative names do include the prefix (for example, ooUint8, ooInt16, ooFloat32).
Constants	000	Names of all constants begin with the $ooc$ prefix. The "c" indicates "constant".
Variables	00V	Names of all variables begin with the $00v$ prefix. The "v" indicates "variable".
Macros and functions	00	Names of all macros and global functions begin with the $\infty$ prefix.
Member functions of classes derived from oo0bj	00	Most of the member functions of classes representing Objectivity/DB objects have names that begin with the $\infty$ prefix.
Member functions of handle and object-reference classes		Names of member functions of handle and object-reference classes do not begin with a prefix. If the name of a member function contains the underscore character, then the member function gets or sets the contents of the object reference or handle itself.

Classes and types that conform to the Object Database Management Group specification (ODMG-93) have names that begin with the d\_ prefix. Multiple words in a name are separated by underscore characters (\_), for example, d\_Object.

The name *appClass* is used to represent the name of any application-defined class; *appClass* also appears in the name of each parameterized class that is generated for an application-defined class—for example, <code>ooltr(appClass)</code> is the name of the object-iterator class that is generated for the class *appClass*.

#### **Global Names**

The Objectivity/C++ programming interface includes functions, macros, types, constants, and variables that are defined in the global scope; they are described in detail in "Global Names" on page 23.

The global functions and macros are used for the following purposes:

- Initialization and cleanup of threads and processes.
- Establishing settings for the application as a whole.
- Establishing settings for the current Objectivity context.
- Error handling.
- Performance tuning.
- Administrative operations on the federated database.

The functions and macros are listed in "Functions and Macros Index" on page 773. The types, constants, and variables are listed in "Types and Constants Index" on page 787.

#### Classes

The classes in the Objectivity/C++ programming interface can be classified as follows:

- Classes for application objects, which control the application's interaction with a federated database
- Classes for Objectivity/DB objects
- Handle classes
- Object-reference classes
- Classes that can be used as types for attributes of persistent objects
- Iterator classes
- Classes to support context-based filtering
- Classes for ODMG applications

In general, detailed descriptions of the classes appear in alphabetical order, one per chapter, starting on page 93. In the following cases, descriptions for multiple classes are combined:

- A handle class ooHandle(className) and the corresponding object-reference class ooRef(className) are combined under the name ooRefHandle(className). For example, the classes ooHandle(ooObj) and ooRef(ooObj) are described in one chapter called ooRefHandle(ooObj).
- When an ODMG class consists of a renamed Objectivity/C++ class, the ODMG class is indicated in the chapter introduction for the Objectivity/C++ class.

The classes are listed alphabetically in "Classes Index" on page 763, with member functions listed under each class. All functions, including member functions, are listed alphabetically in "Functions and Macros Index" on page 773.

#### Handle and Object-Reference Parameters

Objectivity/C++ programming interface considers handles and object references to be syntactically interchangeable. Because of implicit type conversion, Objectivity/C++ functions can accept a handle wherever an object reference is expected, and vice versa:

- You can specify a variable of type ooRef(className) to a function that accepts a parameter of type const ooHandle(className) &.
- You can specify a variable of type <code>ooHandle(className)</code> to a function that accepts a parameter of type <code>const ooRef(className) &</code>.

In some cases, function overloading is used to achieve the same effect—that is, the function has one declaration that expects a handle, and a second declaration that expects an object reference.

**NOTE** When two declarations of an overloaded function are the same except for a parameter for specifying a handle or an object reference, they are collapsed into a single declaration, and the parameter type is given as the abbreviation <code>ooRefHandle(className)</code>.

# **Reference Summary**

The tables in the following subsections present an overview of the Objectivity/C++ programming constructs, according to their intended use.

### Applications

Processes and Threads	<u>ooContext</u> class <u>ooInit</u> function <u>ooInitThread</u> function <u>ooTermThread</u> function <u>ooExitCleanup</u> function
Transactions	<u>ooTrans</u> class <u>ooMode</u> type <u>ooIndexMode</u> type <u>ooDowngradeMode</u> type <u>ooUpdateIndexes</u> function
Settings for the Application	ooCheckVTablePointer functionooAMSUsage typeooSetAMSUsage functionooOfflineMode type (FTO)ooGetOfflineMode function (FTO)ooSetOfflineMode function (FTO)ooStartInternalLS function (IPLS)ooStopInternalLS function (IPLS)ooCheckLS function
Settings for an Objectivity Context	<u>ooSetRpcTimeout</u> function <u>ooSetLockWait</u> function <u>ooUseIndex</u> function <u>ooSetHotMode</u> function <u>ooSetLargeObjectMemoryLimit</u> function <u>ooNoLock</u> function <u>Error handling</u> functions and variables

Error Handling	<u>ooStatus</u> type
-	<u>ooError</u> type
	<u>oovLastError</u> variable
	<u>ooErrorLevel</u> type
	oovLastErrorLevel variable
	<u>oovNError</u> variable
	<u>ooResetError</u> macro
	<u>ooSignal</u> function
	ooSetErrorFile function
	<u>ooErrorHandlerPtr</u> type
	<u>ooRegErrorHandler</u> macro
	ooGetErrorHandler macro
	ooMsgHandlerPtr <b>type</b>
	ooRegMsgHandler macro
	ooGetMsgHandler macro
	ooCheckVTablePointer function
Performance Tuning	ooRunStatus function
	<u>ooSetHotMode</u> function
	<u>ooNoLock</u> function
	<u>ooSetLargeObjectMemoryLimit</u> function
	<u>ooShortRef(appClass)</u> class
Administration of Federated	ooCleanup function
Administration of Federated Database	<u>ooCleanup</u> function ooGetActiveTrans function
Administration of Federated Database	<u>ooCleanup</u> function <u>ooGetActiveTrans</u> function ooGetResourceOwners function
Administration of Federated Database	<u>ooCleanup</u> function <u>ooGetActiveTrans</u> function <u>ooGetResourceOwners</u> function ooTransId type
Administration of Federated Database	<u>ooCleanup</u> function <u>ooGetActiveTrans</u> function <u>ooGetResourceOwners</u> function <u>ooTransId</u> type ooTransInfo type
Administration of Federated Database	<u>ooCleanup</u> function <u>ooGetActiveTrans</u> function <u>ooGetResourceOwners</u> function <u>ooTransId</u> type <u>ooTransInfo</u> type ooResource type
Administration of Federated Database	<u>ooCleanup</u> function <u>ooGetActiveTrans</u> function <u>ooGetResourceOwners</u> function <u>ooTransId</u> type <u>ooResource</u> type <u>ooPurgeAps</u> function (ETO)
Administration of Federated Database	<u>ooCleanup</u> function <u>ooGetActiveTrans</u> function <u>ooGetResourceOwners</u> function <u>ooTransId</u> type <u>ooTransInfo</u> type <u>ooResource</u> type <u>ooPurgeAps</u> function ( <i>FTO</i> ) <u>ooFileNameFormat</u> type
Administration of Federated Database	<u>ooCleanup</u> function <u>ooGetActiveTrans</u> function <u>ooGetResourceOwners</u> function <u>ooTransId</u> type <u>ooResource</u> type <u>ooPurgeAps</u> function ( <i>FTO</i> ) <u>ooFileNameFormat</u> type <u>ooRefHandle(ooDBObi)</u> class members
Administration of Federated Database	<u>ooCleanup</u> function <u>ooGetActiveTrans</u> function <u>ooGetResourceOwners</u> function <u>ooTransId</u> type <u>ooTransInfo</u> type <u>ooResource</u> type <u>ooPurgeAps</u> function ( <i>FTO</i> ) <u>ooFileNameFormat</u> type <u>ooRefHandle(ooDBObj)</u> class members <u>ooRefHandle(ooFDObj)</u> class members
Administration of Federated Database	ooCleanupfunctionooGetActiveTransfunctionooGetResourceOwnersfunctionooTransIdtypeooTransInfotypeooResourcetypeooPurgeApsfunction (FTO)ooFileNameFormattypeooRefHandle(ooFDObj)class membersooRefHandle(ooFDObj)class members
Administration of Federated Database Application-Defined Functions	ooCleanupfunctionooGetActiveTransfunctionooGetResourceOwnersfunctionooTransIdtypeooTransInfotypeooResourcetypeooPurgeApsfunction (FTO)ooFileNameFormattypeooRefHandle(ooDBObj)class membersooRefHandle(ooFDObj)class membersooConvertFunctiontype
Administration of Federated Database Application-Defined Functions	ooCleanupfunctionooGetActiveTransfunctionooGetResourceOwnersfunctionooTransIdtypeooTransInfotypeooResourcetypeooPurgeApsfunction (FTO)ooFileNameFormattypeooRefHandle(ooDBObj)class membersooRefHandle(ooFDObj)class membersooConvertFunctiontypeooErrorHandlerPtrtype
Administration of Federated Database Application-Defined Functions	ooCleanup function         ooGetActiveTrans function         ooGetResourceOwners function         ooTransId type         ooTransInfo type         ooResource type         ooPurgeAps function (FTO)         ooRefHandle(ooDBObj) class members         ooConvertFunction type         ooConvertFunction type         ooRegErrorHandler macro
Administration of Federated Database Application-Defined Functions	ooCleanup function         ooGetActiveTrans function         ooGetResourceOwners function         ooTransId type         ooTransInfo type         ooResource type         ooPurgeAps function (FTO)         ooRefHandle(ooDBObj) class members         ooConvertFunction type         ooConvertFunction type         ooRegErrorHandler macro         ooGetErrorHandler macro
Administration of Federated Database Application-Defined Functions	ooCleanup functionooGetActiveTrans functionooGetResourceOwners functionooTransId typeooTransInfo typeooResource typeooFileNameFormat typeooRefHandle(ooDBObj) class membersooRefHandle(ooFDObj) class membersooConvertFunction typeooRegErrorHandler macroooGetErrorHandler macroooMsgHandlerPtr type
Administration of Federated Database Application-Defined Functions	ooCleanup function         ooGetActiveTrans function         ooGetResourceOwners function         ooTransId type         ooTransInfo type         ooResource type         ooFileNameFormat type         ooRefHandle(ooBObj) class members         ooConvertFunction type         ooRegErrorHandler macro         ooGetErrorHandler type         ooRegErrorHandler macro         ooMsgHandlerPtr type         ooRegMsgHandler macro
Administration of Federated Database Application-Defined Functions	ooCleanup function         ooGetActiveTrans function         ooGetResourceOwners function         ooTransId type         ooTransInfo type         ooResource type         ooFileNameFormat type         ooRefHandle(ooBObj)         class members         ooConvertFunction type         ooRegErrorHandler macro         ooRegtHandlerPtr         ooRegErrorHandler macro         ooRegMsgHandler macro         ooRegMsgHandler macro
Administration of Federated Database Application-Defined Functions	ooCleanup functionooGetActiveTrans functionooGetResourceOwners functionooTransId typeooTransInfo typeooResource typeooPurgeAps function (FTO)ooFileNameFormat typeooRefHandle(ooDBObj) class membersooRefHandle(ooFDObj) class membersooConvertFunction typeooRegErrorHandler macroooGetErrorHandler macroooRegMandlerPtr typeooRegMsgHandler macroooGetMsgHandler macroooRegMsgHandler macroooRegMsgHandler macroooNameHashFuncPtr type
Administration of Federated Database Application-Defined Functions	ooCleanup functionooGetActiveTrans functionooGetResourceOwners functionooTransId typeooTransInfo typeooResource typeooFileNameFormat typeooRefHandle(ooBObj) class membersooRefHandle(ooFDObj) class membersooConvertFunction typeooRegErrorHandler macroooGetErrorHandler macroooRegMsgHandler macroooRegMsgHandler macroooRegMsgHandler macroooRegMsgHandler macroooRegMsgHandler typeooNameHashFuncPtr typeooQueryOperatorPtr type
Administration of Federated Database	ooCleanup functionooGetActiveTrans functionooGetResourceOwners functionooTransId typeooTransInfo typeooResource typeooPurgeAps function (FTO)ooFileNameFormat typeooRefHandle(ooDBObj) class membersooRefHandle(ooFDObj) class membersooConvertFunction typeooRegErrorHandler macroooGetErrorHandler macroooRegMsgHandler macroooRegMsgHandler macroooRegMsgHandler macroooRegMsgHandler typeooNameHashFuncPtr typeooQueryOperatorPtr typeooTwoMachineHandlerPtr type (DRO)

## **Objectivity/DB Objects**

All	Objectivity/DB Objects	
	Deleting	ooDelete function
	Accessing	<u>ooMode</u> type <u>ooLockMode</u> type
	Getting Class Information	<u>ooTypeN</u> function ooTypeNumber type
Fee	derated Databases	<u>ooFDObj</u> class <u>ooHandle(ooFDObj)</u> class <u>ooRef(ooFDObj)</u> class
Dat	tabases	<u>ooDBObj</u> class <u>ooHandle(ooDBObj)</u> class <u>ooRef(ooDBObj)</u> class <u>ooItr(ooDBObj)</u> class <u>ooReplace</u> function <u>ooContainsFilter</u> type ( <i>FTO</i> )
Со	ntainers	<u>ooContObj</u> class <u>ooDefaultContObj</u> class <u>ooGCContObj</u> class <u>ooHandle(ooContObj)</u> class <u>ooRef(ooContObj)</u> class <u>ooItr(ooContObj)</u> class <u>ooNewConts</u> function
All	Persistent Objects	
	Deleting	<u>ooDeleteNoProp</u> function
	Object Conversion	<u>ooConvertFunction</u> type <u>ooConvertInObject</u> class <u>ooConvertInOutObject</u> class <u>ooRefHandle(ooFDObj)</u> class members <u>ooTrans</u> class members

Basic Objects	
General	<u>ooObj</u> class <u>ooHandle(ooObj)</u> class <u>ooRef(ooObj)</u> class <u>ooShortRef(ooObj)</u> class <u>ooItr(ooObj)</u> class <u>appClass</u> classes <u>ooHandle(appClass)</u> classes <u>ooShortRef(appClass)</u> classes <u>ooItr(appClass)</u> classes <u>ooItr(appClass)</u> classes
Scalable Persistent Collections	ooCollectionclassooBTreeclassooTreeListclassooTreeSetclassooTreeMapclassooHashSetclassooCollectionIteratorclassooCompareclassooCompareclassooTreeAdminclassooTreeAdminclassooHashAdminclass
Nonscalable Persistent Collections	<u>ooMap</u> class <u>ooMapElem</u> class <u>ooMapItr</u> class <u>ooNameHashFuncPtr</u> type
Keyed Objects	<u>ooKey</u> type <u>ooNewKey</u> function <u>ooKeyType</u> type <u>ooGetMemberOffset</u> function <u>ooGetMemberSize</u> function
Versions of Basic Object	ooGeneObj class ooVersMode type
Autonomous Partitions ( <i>FTO</i> )	<u>ooAPObj</u> class <u>ooHandle(ooAPObj)</u> class <u>ooRef(ooAPObj)</u> class <u>ooItr(ooAPObj)</u> class <u>ooOfflineMode</u> type <u>ooGetOfflineMode</u> function <u>ooSetOfflineMode</u> function <u>ooPurgeAps</u> function

### Handles and Object References

Handles	<u>ooHandle(ooFDObj)</u> class <u>ooHandle(ooDBObj)</u> class <u>ooHandle(ooContObj)</u> class <u>ooHandle(ooObj)</u> class <u>ooHandle(appClass)</u> classes
Object References	
Standard	<u>ooRef(ooFDObj)</u> class <u>ooRef(ooDBObj)</u> class <u>ooRef(ooContObj)</u> class <u>ooRef(ooObj)</u> class <u>ooRef(appClass)</u> classes
Short	<u>ooShortRef(ooObj)</u> class ooShortRef(appClass)classes

## **Common Data Types**

Primitive Data Types	<pre>int8 type int16 type int32 type int64 type uint8 type uint16 type uint32 type uint64 type float32 type float64 type char type ooBoolean type</pre>
Strings	<u>ooVString</u> class ooString(N) class
Variable-Size Arrays (VArrays)	<pre>ooVArrayT<element_type> classes ooTVArrayT<element_type> classes</element_type></element_type></pre>
Object References	<u>ooRef(appClass)</u> classes <u>ooRef(ooObj)</u> class <u>ooRef(ooContObj)</u> class <u>ooShortRef(appClass)</u> classes <u>ooShortRef(ooObj)</u> class

Dat	es and Times	<u>d_Date</u> class <u>d_Time</u> class <u>d_Timestamp</u> class <u>d_Interval</u> class
Jav	a compatibility	
	Strings	ooUtf8String class
-	Dates and Times	<u>oojDate</u> class <u>oojTime</u> class <u>oojTimestamp</u> class
	Arrays	<u>oojArrayOfInt8</u> class <u>oojArrayOfInt16</u> class <u>oojArrayOfInt32</u> class <u>oojArrayOfInt64</u> class <u>oojArrayOfFloat</u> class <u>oojArrayOfDouble</u> class <u>oojArrayOfCharacter</u> class <u>oojArrayOfBoolean</u> class <u>oojArrayOfObject</u> class <u>oojString</u> class
Sta (ST	ndard Template Library 'L) Strings and Containers	See Objectivity/C++ Standard Template Library.

### Iterators

Object Iterators	<u>ooItr(appClass)</u> classes <u>ooItr(ooObj)</u> class <u>ooItr(ooContObj)</u> class <u>ooItr(ooDBObj)</u> class
Name-Map Iterators	<u>ooMapItr</u> class
Scalable-Collection Iterators	ooCollectionIterator class
VArray Iterators	<u>d Iterator<element type=""> classes</element></u>

## **Content-Based Filtering**

Creating Indexes	<u>ooKeyDesc</u> class ooKeyField class
Using Indexes	<u>ooUseIndex</u> function <u>ooIndexMode</u> type <u>ooUpdateIndexes</u> function
Query Objects	ooQuery class
Lookup Keys	<u>ooLookupKey</u> Class <u>ooLookupFieldBase</u> Class <u>ooEqualLookupField</u> Class <u>ooGreaterThanLookupField</u> Class <u>ooGreaterThanEqualLookupField</u> Class <u>ooLessThanLookupField</u> Class <u>ooLessThanEqualLookupField</u> Class
Extending the Predicate Query Language	<u>ooOperatorSet</u> class <u>ooUserDefinedOperators</u> variable <u>ooQueryOperatorPtr</u> type <u>ooDataType</u> type

## **ODMG** Applications

	-
Transactions	<u>d Transaction</u> class
Objectivity/DB Objects	<u>d_Database</u> class <u>d_Object</u> class
Object References	<u>d_Ref<d_object></d_object></u> class <u>d_Ref<appclass></appclass></u> classes <u>d_Ref_Any</u> class
Primitive Data Types	<u>d Boolean</u> <u>d Char</u> <u>d Double</u> <u>d Float</u> <u>d Long</u> <u>d Octet</u> <u>d Short</u> <u>d ULong</u> <u>d UShort</u>

Dates and Times	<u>d Date</u> class <u>d Time</u> class <u>d Timestamp</u> class <u>d Interval</u> class
Strings and Arrays	<u>d_String</u> class <u>d_Varray<element_type></element_type></u> classes
VArray Iterator	<u>d_Iterator<element_type> classes</element_type></u>

# **Global Names**

Objectivity/C++ global names include functions, macros, types, constants, and variables defined in the global scope. Although all of these names are syntactically global, some variables and functions manage settings that are global to the entire application, while others manage settings that are specific to the current Objectivity context.

See:

- "Reference Index" on page 26 for an alphabetical list of global names
- "Reference Descriptions" on page 31 for individual descriptions

## **Global Functions and Macros**

#### Error Conditions

Most Objectivity/C++ global functions and macros recognize a number of error conditions. If a function fails due to one of these error conditions, it puts the condition's error number and error message in the variable <code>oovLastError</code>, where the currently registered error handler can obtain them. See Chapter 23, "Error Handling," of the Objectivity/C++ programmer's guide for more information about the Objectivity/DB error handling facility.

#### Macro Expansion

Global macros exist to accommodate multiple platforms. The DDL processor replaces each occurrence of a macro with the code given in its definition in the Objectivity/C++ header files. Because the macro is replaced prior to compilation, the compiler does not generate any error messages using the macro name and does not perform type checking on the macro arguments. The compiler sees only the code resulting from the DDL processor's expansion of the macro.

### **Global and Context Settings**

Some Objectivity/C++ global functions and macros manage settings that pertain globally to an entire application; others manage settings that are specific to the current Objectivity context. The following table lists these functions and macros, and the scope to which they apply:

ooCheckVTablePointerooGetErrorHandlerooExitCleanupooGetMsgHandler	Global Settings	Context Settings
ooGetOfflineMode       ooInit (file descriptors setting)         ooInit (signal handler setting)       ooInit (file descriptors setting)         ooRegTwoMachineHandler       ooNoLock         ooSetAMSUsage       ooRegErrorHandler         ooStartInternalLS       ooResetError         ooSetAmsus       ooResetErrorFile         ooStopInternalLS       ooSetLargeObjectMemoryLimit         ooSetLockWait       ooSetRpcTimeout         ooSignal       ooTermThread         ooUseIndex       ooUseIndex	poCheckVTablePointer poExitCleanup poGetOfflineMode poInit (signal handler setting) poRegTwoMachineHandler poSetAMSUsage poSetOfflineMode poStartInternalLS poStopInternalLS	ooGetErrorHandler         ooGetMsgHandler         ooInit (file descriptors setting)         ooInitThread         ooNoLock         ooRegErrorHandler         ooRegMsgHandler         ooResetError         ooRunStatus         ooSetLargeObjectMemoryLimit         ooSetRpcTimeout         ooSignal         ooTermThread

# **Global Types**

Objectivity/C++ global types include:

- Primitive data types
- Types that define constants used by various Objectivity/C++ functions
- Function-pointer types that define the syntax for application-defined error handling, message handling, query operator, and object conversion functions

## **Primitive Type Names**

An alternate name exists for most of the primitive data types. Many of these types have equivalent names in the ODMG standard. The type name, alternate name, and ODMG name for a given primitive type can be used interchangeably. The following table summarizes the primitive types and their names. For more information about primitive types, see the Objectivity/C++ Data Definition Language book.

Category	Type Name	Alternate Name	ODMG Name	Brief Description
Integer	<u>int8</u>	ooInt8	(None)	8-bit signed integer type
	<u>uint8</u>	ooUInt8	d_Octet	8-bit unsigned integer type
	<u>int16</u>	ooInt16	d_Short	16-bit signed integer type
	<u>uint16</u>	ooUInt16	d_UShort	16-bit unsigned integer type
	<u>int32</u>	ooInt32	d_Long	32-bit signed integer type
	<u>uint32</u>	ooUInt32	d_ULong	32-bit unsigned integer type
	<u>int64</u>	ooInt64	(None)	64-bit signed integer type
	<u>uint64</u>	ooUInt64	(None)	64-bit unsigned integer type
Floating point	<u>float32</u>	ooFloat32	d_Float	32-bit floating-point type
	<u>float64</u>	ooFloat64	d_Double	64-bit floating-point type
Character	<u>char</u>	ooChar	d_Char	8-bit character type
Boolean	<u>ooBoolean</u>	(None)	d_Boolean	8-bit unsigned integer type

# **Global and Context Variables**

Objectivity/C++ defines several variables that you can use throughout your application to refer to data, such as the most recently signaled error condition. The following table lists these variables by their semantic scope—that is, whether they manage settings that are global to the entire application or specifc to the current Objectivity context:

Variables That Store Global State	Variables That Store Context State
ooUserDefinedOperators	<u>oovLastError</u> <u>oovLastErrorLevel</u> <u>oovNError</u>

In a single-threaded application, Objectivity/C++ context variables behave exactly like global variables. In a multithreaded application, they are expressions whose values are specific to the current Objectivity context. Thus, changing the value of a global variable in one thread does not affect its value in other threads.

# **Reference Index**

char	Objectivity/C++ 8-bit character type.
<u>float32</u>	Objectivity/C++ 32-bit floating-point type.
float64	Objectivity/C++ 64-bit floating-point type.
int8	Objectivity/C++ 8-bit signed integer type.
int16	Objectivity/C++ 16-bit signed integer type.
int32	Objectivity/C++ 32-bit signed integer type.
int64	Objectivity/C++ 64-bit signed integer type.
<u>ooAccessMode</u>	Type for specifying which data members an iterator can test in a predicate query.
ooAMSUsage	(administration) Type for specifying how an application uses the Advanced Multithreaded Server (AMS) as data server software on remote hosts.
ooBoolean	Objectivity/C++ Boolean type.
ooChar	See <u>char</u> .

<u>ooCheckLS</u>	Tests whether a lock server is running on the specified host machine.
<u>ooCheckVTablePointer</u>	Suppresses or re-enables the warning messages issued for classes that do not have virtual-function tables.
ooCleanup	(administration) Recovers the specified transaction; primarily used in creating database recovery tools.
<u>ooContainsFilter</u>	(FTO) Type for specifying the objects to iterate over in containers or databases.
<u>ooConvertFunction</u>	Function-pointer type for application-defined conversion functions.
<u>ooDataType</u>	Type for indicating the data types of operands in the definitions of custom relational operators.
ooDelete	Removes the Objectivity/DB object that is referenced by the specified handle, propagating deletion along all associations that have delete propagation enabled.
<u>ooDeleteNoProp</u>	Removes only the persistent object that is referenced by the specified handle, without removing any associated objects.
ooDowngradeMode	Type for specifying how update locks are treated when checkpointing a transaction with a commit-and-hold operation.
<u>ooError</u>	Type for representing error identifiers for application-defined and Objectivity/C++-defined errors.
<u>ooErrorHandlerPtr</u>	Function-pointer type for application-defined error handler functions.
<u>ooErrorLevel</u>	Type for specifying the severity level of Objectivity/C++ errors.
<u>ooExitCleanup</u>	Leaves Objectivity/DB in a safe state for process termination.
<u>ooFileNameFormat</u>	(administration) Type for specifying the output format of filenames.
<u>ooGetActiveTrans</u>	(administration) Returns an array of transaction-information structures describing the active transactions for a particular federated database or autonomous partition; used in database recovery tools.
<u>ooGetErrorHandler</u>	Returns a pointer to the currently registered error handler function.

<u>ooGetMemberOffset</u>	Provides information required for initializing a key structure for a class—specifically, the byte offset from the start of the class to the member field that is to serve as the key.
<u>ooGetMemberSize</u>	Provides information required for initializing a key structure for a class—specifically, the size of the member field in the class.
<u>ooGetMsgHandler</u>	Returns a pointer to the currently registered message handler function.
<u>ooGetOfflineMode</u>	(FTO) Gets a process' current offline mode. The offline mode determines whether an application ignores or enforces the offline status of all autonomous partitions.
<u>ooGetResourceOwners</u>	(administration) Returns information about the resource for which the specified transaction is waiting, along with the transaction(s) currently holding that resource; used in database recovery tools.
<u>ooIndexMode</u>	Type for specifying how a transaction updates indexes after creating or modifying objects of an indexed class.
<u>ooInit</u>	Initializes Objectivity/DB in a single-threaded or multithreaded application.
<u>ooInitThread</u>	Initializes the current thread so that it can execute Objectivity/DB operations.
<u>ooKey</u>	Type for specifying the key structure for a keyed object.
<u>ooKeyType</u>	Type for specifying the data type of a key field within a keyed object. Constants of this type are used in creating key structures.
<u>ooLockMode</u>	Type for specifying how to lock an Objectivity/DB object.
<u>ooMode</u>	Type specifying the intended level of access to an Objectivity/DB object or for specifying the concurrent access policy for a transaction.
<u>ooMsgHandlerPtr</u>	Function-pointer type for application-defined message handler functions.
<u>ooNameHashFuncPtr</u>	Function-pointer type for an application-defined hash function to be used by instances of ooMap.
<u>ooNewConts</u>	Creates a batch of containers in the specified database. This macro provides better performance than calling the $new$ operator repeatedly.

<u>ooNewKey</u>	Creates a keyed object for the specified class.
ooNoLock	Disables the Objectivity/DB locking facilities, removing concurrent access protection.
<u>ooOfflineMode</u>	(FTO) Type for specifying how an application should respond to the offline status of all autonomous partitions.
<u>ooPurgeAps</u>	(FTO) Purges the specified autonomous partitions from the federated database.
<u>ooQueryOperatorPtr</u>	Function-pointer type for application-defined relational-operator functions.
<u>ooRegErrorHandler</u>	Registers the specified error handler function with Objectivity/DB, replacing the previously registered error handler.
<u>ooRegMsgHandler</u>	Registers the specified message handler function with Objectivity/DB.
<u>ooRegTwoMachineHandler</u>	(DRO) Registers the specified two-machine handler function with Objectivity/DB.
<u>ooReplace</u>	Creates a database with the specified name in the specified federated database, replacing (deleting) any existing database with the same name.
ooResetError	Clears Objectivity/DB error flags.
ooResource	Type for representing information about a locked resource (typically a container).
<u>ooRunStatus</u>	Prints a summary of Objectivity/DB internal statistics to stdout; used primarily for performance tuning.
ooSetAMSUsage	(administration) Sets the application's policy for using the Advanced Multithreaded Server (AMS).
<u>ooSetErrorFile</u>	Sets the error message output file.
<u>ooSetHotMode</u>	Enables or disables hot mode, which controls the timing of certain internal overhead operations in the Objectivity/DB cache, for purposes of improving performance.
<u>ooSetLargeObjectMemory</u> <u>Limit</u>	Sets the maximum amount of dynamically allocated memory that is available for caching large persistent objects.
<u>ooSetLockWait</u>	Sets the default lock-waiting option for a series of transactions.

<u>ooSetOfflineMode</u>	(FTO) Sets the offline mode for the current process.
<u>ooSetRpcTimeout</u>	Sets how long an application is to wait for an Objectivity server to respond before signaling a timeout error.
<u>ooSignal</u>	Signals the specified error and reports additional error information.
<u>ooStartInternalLS</u>	(IPLS) Starts an in-process lock server within an application.
ooStatus	General return type for Objectivity/C++ global functions and member functions.
<u>ooStopInternalLS</u>	(IPLS) Shuts down an in-process lock server within an application.
<u>ooTermThread</u>	Terminates the current thread's ability to invoke Objectivity/DB operations.
<u>ooTransId</u>	Unique identifier for a transaction.
<u>ooTransInfo</u>	Type for representing information about a transaction.
<u>ooTwoMachineHandlerPtr</u>	(DRO) Function-pointer type for application-defined message handler functions.
<u>ooTypeN</u>	Gets the type number of the specified class of Objectivity/DB objects.
<u>ooTypeNumber</u>	Type number of a class in the ooObj inheritance hierarchy.
<u>ooUpdateIndexes</u>	Explicitly updates all applicable indexes to reflect a new or modified object.
<u>ooUseIndex</u>	Enables or disables the use of indexes during a predicate query.
ooUserDefinedOperators	Operator set consulted by the predicate query mechanism to resolve application-defined relational operators used in a predicate.
<u>ooVersMode</u>	Type for representing the versioning behavior of a basic object.
<u>oovLastError</u>	Pointer to the error identifier structure for the most recent
	error condition.
<u>oovLastErrorLevel</u>	error condition. Severity level of the most recent error condition.

<u>ooVoidFuncPtr</u>	Function-pointer type for a function that has no parameters and that returns no result.
uint8	Objectivity/C++ 8-bit unsigned integer type.
uint16	Objectivity/C++ 16-bit unsigned integer type.
uint32	Objectivity/C++ 32-bit unsigned integer type.
uint64	Objectivity/C++ 64-bit unsigned integer type.
Wait Options	Integer constants that specify whether to wait for locks when starting a transaction.

# **Reference Descriptions**

char	global type
	Objectivity/C++ 8-bit character type. This type is signed on architectures where the C++ primitive type char is signed, and unsigned on architectures where the C++ primitive type char is unsigned. This type is also called ooChar.
float32	global type
	Objectivity/C++ 32-bit floating-point type. The range and precision of this type vary for each architecture. Objectivity/DB stores numbers of this type in the native format of the architecture on which they are instantiated or modified, and automatically converts the format when these numbers are accessed from a different architecture. This type is also called $ooFloat32$ .
float64	global type
	Objectivity/C++ 64-bit floating-point type. The range and precision of this type vary for each architecture. Objectivity/DB stores numbers of this type in the native format of the architecture on which they are instantiated or modified, and automatically converts the format when these numbers are accessed from a different architecture. This type is also called $ooFloat64$ .
int8	global type
	Objectivity/C++ 8-bit signed integer type. Values of this primitive type range from -128 to +127. This type is portable across all architectures supported by Objectivity/C++. This type is also called $ooInt8$ .

Global Names

int16	global type
	Objectivity/C++ 16-bit signed integer type. Values of this primitive type may range from -32,768 to +32,767. This type is portable across all architectures supported by Objectivity/C++. This type is also called <code>ooInt16</code> .
int32	global type
	Objectivity/C++ 32-bit signed integer type. Values of this primitive type may range from -2,147,483,648 to +2,147,483,647. This type is portable across all architectures supported by Objectivity/C++. This type is also called <code>ooInt32</code> .
int64	global type
	Objectivity/C++ 64-bit signed integer type. Values of this primitive type may range from $-(2^{63})$ to $+2^{63}-1$ . This type is portable across all architectures supported by Objectivity/C++. This type is also called ooInt64.
ooAccessN	fode global type
	Type for specifying which data members an iterator can test in a predicate query.
Constants	oocPublic The predicate query tests only public data members, preserving encapsulation.
	oocAll
	The predicate query tests all data members (public, protected, or private). To preserve encapsulation, you should use this mode only within member functions of the class you are querying.
See also	ooItr(ooObj):: <u>scan</u> ooItr(ooContObj):: <u>scan</u>
ooAMSUsa	global type
	( <i>administration</i> ) Type for specifying how an application uses the Advanced Multithreaded Server (AMS) as data server software on remote hosts.
Constants	oocAMSPreferred The application uses AMS whenever AMS is available.

The application uses only AMS; an error is signaled if AMS is not available.

OOCNOAMS The application does not use AMS even if it is available. See also ooSetAMSUsage ooBoolean global type Objectivity/C++ Boolean type. Constants oocTrue The true value. Its integer value is 1. oocFalse The false value. Its integer value is 0. ooCheckLS global function Tests whether a lock server is running on the specified host machine. ooBoolean ooCheckLS(const char \*host = NULL); Parameters host Name of the host machine to be checked for a running lock server. If you omit the *hostName* parameter or specify NULL, the current host is checked. Returns oocTrue if a lock server is running on the specified host; otherwise, oocFalse. Discussion (*IPLS*) An application can call this function to decide whether to start an in-process lock server. An in-process lock server can be started only if no other lock server process is running on the same host as the application. The application could take various actions based on the result of this function. For example, if another lock server is running on the current host, the application could report it and allow the user to choose whether to quit the application or

See also <u>ooStartInternalLS</u>

#### ooCheckVTablePointer

global function

Suppresses or re-enables the warning messages issued for classes that do not have virtual-function tables.

continue using the other lock server.

#### Parameters checkVptr

Specifies whether to issue a warning message whenever a persistent object is opened whose class does not have a virtual table pointer. If you omit this parameter, your application may issue such warning messages. If you specify <code>oocFalse</code>, the warning messages are suppressed.

Discussion Sometimes when an application opens a persistent object from the database, Objectivity/DB cannot obtain the virtual function table for the object's class. This can happen when an existing application accesses instances of recently added subclasses. For example, assume that an existing application Appl iterates over persistent instances of class A and calls a virtual member function on each instance. A more recent application App2 uses class A as the base class for a new derived class B that implements the virtual member function; App2 stores persistent instances of B in the federated database. The next time App1 iterates over instances of A, it finds instances of both A and B, and attempts to invoke the virtual function on each found object. But because class B was added to the schema after App1 was compiled and linked, App1 has no registered virtual-function table for dispatching virtual member-function calls to their implementations in class B. The absence of a virtual-function table results in a fatal error, usually with an illegal memory access, invalid memory address, or similar diagnostic.

The system default is for an application such as App1 to issue a warning message when it opens an object for which a virtual function table cannot be obtained. Depending on how the object is opened, the application can trap such warnings and handle the problem gracefully. If, however, the object is opened implicitly by a call to a virtual member function, graceful recovery is not possible.

The best solution is to relink an old application (such as App1) when instances of a new subclass are added to the federated database. The old application should link with the object file corresponding to the method implementation file for the new subclass. (A method implementation file is generated by the DDL processor for each persistence-capable class; if a class implements virtual functions, the file contains code for registering a virtual-function table.) Of course, relinking is necessary only for applications that will encounter instances of the new subclass (for example, through iteration over the base class) *and* call virtual member functions implemented by the subclass.

If you can guarantee that an application does not call any virtual member functions on any potential base classes, you can use the ooCheckVTablePointer function to suppress the warning messages that result from opening instances of a class for which no virtual-function table exists. For example, this might be useful in a conversion application that simply opens every existing object of a particular class without invoking any virtual member functions.

#### ooCleanup

global function

(*administration*) Recovers the specified transaction; primarily used in creating database recovery tools.

```
ooStatus ooCleanup(
    char **ppBootFilePath,
    ooTransId tId,
    int ignoreHost,
    int standalone,
    int resetLock,
    ooTransInfo *pLockOwner,
    oo2PCTransState reservedParameter);
```

Parameters ppE

*ppBootFilePath* 

Pointer to the string name (a char\*) of the boot file for the federated database or any autonomous partition that is accessed by the specified transaction. If *\*ppBootFilePath* is 0, the boot file path is obtained from the environment variable OO\_FD\_BOOT and the pointed-to string is set to the obtained name.

tId

Transaction identifier of the transaction to recover. Transaction identifiers are typically obtained through the ooGetActiveTrans function.

#### ignoreHost

Specifies whether to permit the recovery of a transaction that started on another node. If you specify 0, the specified transaction is recovered only if it started on the same node from which ooCleanup was invoked. If you specify a number other than 0, the transaction is recovered regardless of where it was started.

#### standalone

Specifies whether to contact the lock server during the recovery process. If you specify 0, ooCleanup contacts the lock server to release any locks left by the transaction. You must specify a number other than 0 if you want to run ooCleanup when the lock server isn't running.

#### resetLock

Specifies ooCleanup's response when it encounters another ooCleanup process. If you specify 0 and another ooCleanup process owns the recovery lock on the transaction:

- Your ooCleanup process fails.
- The data pointed to by *pLockOwner* is set to information about the competing ooCleanup process.

If you specify a number other than 0, ooCleanup resets any existing recovery lock on the transaction and recovers the transaction.

#### pLockOwner

Pointer to a transaction-information structure that, on return, contains information about any competing ooCleanup process that already has a recovery lock on the transaction. You can specify 0 to prevent *pLockOwner* from being updated with such information.

#### reservedParameter

Reserved for future development. You must specify 0.

Returns oocSuccess if successful; otherwise oocError.

Discussion You use this function only in special-purpose recovery tools. When you use this function, you must not call the <code>ooInit</code> function or perform any other nonrecovery Objectivity/DB operations (such as starting a transaction) in the same application. You may, however, use <code>ooCleanup</code> along with <code>ooGetActiveTrans</code> or <code>ooGetResourceOwners</code>. You should use <code>ooCleanup</code> only in single-threaded applications.

ooCleanup uses both the specified transaction identifier and the lock server named in the specified boot file to determine which transaction to recover. ooCleanup recovers a transaction only if the process owning that transaction is inactive.

When ooCleanup recovers a transaction, it rolls back the transaction's uncommitted changes, restoring the federated database to the logical state it was in before the transaction started. If the transaction left uncommitted changes in multiple autonomous partitions, the changes are rolled back in all of the available partitions. In particular, if the transaction left uncommitted changes in a replicated database, the changes are rolled back in all of the available images.

If the lock server is still running and *standalone* is 0, all locks held by the transaction are released. If the lock server stops—or stops and restarts—the transaction's locks are lost, so ooCleanup just rolls back changes. (If ooCleanup is to run while the lock server is stopped, *standalone* must be nonzero.)

ooCleanup puts a *recovery lock* on the specified transaction before proceeding. The recovery lock helps you determine whether additional ooCleanup processes are running. You typically run ooCleanup with resetLock set to 0 to find out whether another ooCleanup process is recovering the specified transaction. If so, you can find out about this process by inspecting the information pointed to by pLockOwner. Based on this information, you may decide to rerun ooCleanup with resetLock set to a nonzero integer value.

When you use the ooCleanup function in an application, you must include the Objectivity/C++ header file ooRecover.h. UNIX applications additionally need
Reference Descriptions

to be linked with the Objectivity/DB administration library (see *Installation and Platform Notes for UNIX*).

See also "Creating a Recovery Application" on page 528 of the Objectivity/C++ programmer's guide

# ooContainsFilter

global type

(FTO) Type for specifying the objects to iterate over in containers or databases.

Constants oocAllObjs

Includes all objects in a federated database; none are filtered out.

### oocNotTransferred

Includes only objects in the boot autonomous partition (local objects); remote objects are filtered out.

### oocTransferred

Includes only objects *not* in the boot autonomous partition (remote objects); local objects are filtered out.

See also ooRefHandle(ooDBObj)::contains

# ooConvertFunction

global type

Function-pointer type for application-defined conversion functions.

typedef void (*ooConvertFunction)(
<pre>const ooConvertInObject &amp;existObj,</pre>
ooConvertInOutObj & <i>convObj</i> );

Parameters existObj

Existing object before it has been converted.

convObj

Existing object after it has been converted.

Discussion Use *ooRefHandle*(ooFDObj)::setConversion to register this function to be invoked automatically on each converted object of a changed class. You can have no more than one conversion function for each changed persistence-capable class, and the function must not access any other persistent object.

See also *ooRefHandle*(ooFDObj)::<u>setConversion</u>

### ooDataType

global type

Type for indicating the data types of operands in the definitions of custom relational operators.

### Constants oocInt64T

Indicates an operand of type int8, int16, int32, or int64.

oocUint64T

Indicates an operand of type uint16, uint32, or uint64.

oocFloat64T

Indicates an operand of type float32 or float64.

### oocCharPtrT

Indicates an operand of type char \*, char[], ooString(N), ooVString, ooUtf8String, or ooVArrayT<char>.

#### oocBooleanT

Indicates an operand of type ooBoolean.

### oocInvalidTypeT

Indicates a type that is not handled by the predicate query mechanism.

- Discussion When you define your own relational operators for use in predicate queries, you use the constants of this type to provide appropriate return values for the possible operand types.
- See also <u>ooQueryOperatorPtr</u>

### ooDelete

global function

Removes the Objectivity/DB object that is referenced by the specified handle, propagating deletion along all associations that have delete propagation enabled.

ooStatus ooDelete(ooRefHandle(ooObj) &object);

Parameters object

Object reference or handle to the object to be deleted. The object may be a database, an autonomous partition, a container, or a basic object. The object may have other open handles to it.

- Returns oocSuccess if successful; otherwise oocError.
- Discussion Deleting a basic object or container:
  - Calls the object's destructor, if any, before deallocating storage.

Deletes all associations from the object to destination objects. If an
association has <u>delete propagation</u> enabled, the destination objects are
deleted as well (see the Objectivity/C++ Data Definition Language book).

For basic objects and containers, the ooDelete function is equivalent to <u>operator delete</u>, which is available for each persistence-capable class. However, ooDelete is recommended because it is simpler to use.

Deleting a container or database:

- Deletes all of the persistent objects in it.
- Deletes all associations from each contained object to destination objects; if delete propagation is enabled for any of these associations, the destination objects are deleted as well.
- Does not call the destructors of the contained objects for performance reasons. To ensure that destructors are called, you must iterate through the contained objects and explicitly delete them before deleting the container or database.

Each deleted object is automatically removed from any bidirectional associations to maintain referential integrity. The delete operation must therefore be able to obtain an update lock on every object that is bidirectionally associated with a deleted object. Note that if another perisistent object references a deleted object through a unidirectional association or directly in one of its data members, you are responsible for removing that reference.

Deleting a database removes its file from the file system. You may not delete a database that has been made read-only; you change the database back to read-write before you can delete it (see <code>ooRefHandle(ooDBObj)::setReadOnly</code>).

### **WARNING** Deleting a database cannot be undone by aborting the transaction.

(*FTO*) Deleting an autonomous partition:

- Clears all previously transferred containers from the autonomous partition, implicitly invoking ooRefHandle(ooAPObj):::returnAll (see the Objectivity/FTO and Objectivity/DRO book).
- Removes the system database and boot file for the autonomous partition from the file system. This cannot be undone by aborting the transaction.
- Removes the database images controlled by the autonomous partition from the file system. If the autonomous partition controls the last or only image of any database, an error is signalled and the autonomous partition is left unchanged.

An autonomous partition cannot be deleted if any journal files exist in its journal directory. You can correct this by using the oocleanup administration tool to recover the incomplete transactions.

 WARNING
 Deleting an autonomous partition cannot be undone by aborting the transaction.

 See also
 ooDeleteNoProp

# ooDeleteNoProp

global function

Removes only the persistent object that is referenced by the specified handle, without removing any associated objects.

ooStatus ooDeleteNoProp(ooRefHandle(ooObj) &object);

Parameters object

Object reference or handle to the object to be deleted. This object may be a container or a basic object. This object may have other open handles to it.

Returns oocSuccess if successful; otherwise oocError.

Discussion This function is identical to <u>ooDelete</u>, except that it:

- Applies only to containers and basic objects, and not to databases or autonomous partitions.
- Does not propagate deletion, even along associations that have delete propagation enabled. Thus, only the specified object is deleted.

The deleted object is automatically removed from any bidirectional associations to maintain referential integrity. The operation must therefore be able to obtain an update lock on every object that is bidirectionally associated with a deleted object. Know that if another perisistent object references the deleted object through a unidirectional association or directly in one of its data members, you are responsible for removing that reference.

See also <u>ooDelete</u>

# ooDowngradeMode

global type

Type for specifying how update locks are treated when checkpointing a transaction with a commit-and-hold operation.

Constants oocNoDowngrade

All locks obtained during the transaction are preserved as is.

	oocDowngradeAll All update locks obtained during the transaction are downgraded to read locks (MROW read locks if the transaction is an MROW transaction, and normal read locks otherwise), thereby improving concurrency.
See also	ooTrans:: <u>commitAndHold</u>
ooError	global type
	Type for representing error identifiers for application-defined and Objectivity/C++-defined errors.
	<pre>struct ooError {     uint32 errorN;     char *message; };</pre>
Members	<pre>errorN Error number. Error numbers 0 through 999,999 are reserved for Objectivity/DB system error codes. Within your application, make sure that you assign modules unique error numbers greater than 999,999. message Identical to a printf format string, allowing specification of a variable number of arguments.</pre>
Discussion	You declare and initialize all error identifier structures in an error message header file, and you include this file in your application source file.
See also	<u>ooErrorLevel</u> <u>ooSignal</u>

# ooErrorHandlerPtr

global type

Function-pointer type for application-defined error handler functions.

```
typedef ooStatus (*ooErrorHandlerPtr)(
    ooErrorLevel errorLevel,
    ooError &errorID,
    ooHandle(ooObj) *contextObj,
    char *errorMsg);
```

Parameters errorLevel

Error level for the event being handled.

	<i>errorID</i> Symbolic name of the error identifier for the event being handled.
	<pre>contextObj Pointer to the handle of an object that may help the system error handler to pinpoint the context of the error. Set this parameter to 0 if you do not want to use context information.</pre>
	errorMsg String created by running vsprintf over the message part of the error identifier and its arguments.
Returns	oocSuccess if the error level of the event is acceptable, otherwise oocError.
Discussion	A registered error handler is a function that is invoked whenever the ooSignal function signals an error for an event. Objectivity/DB provides a predefined error handler that is automatically registered when you start your application. You can replace the predefined error handler by defining a custom error handler and using the ooRegErrorHandler function to register it.
	Your custom error handler must be able to accept the arguments passed to it by the ooSignal function. Therefore, your error handler must conform to the calling interface defined by this function-pointer type.
See also	Chapter 23, "Error Handling," of the Objectivity/C++ programmer's guide <u>ooRegErrorHandler</u> <u>ooSignal</u>

# ooErrorLevel

global type

Type for specifying the severity level of Objectivity/C++ errors.

#### Constants oocNoError

Indicates an event with no severity level.

### oocWarning

Indicates an abnormal event. No action should be taken beyond notifying the user of the condition.

### oocUserError

Indicates a nonfatal user error detected directly by the programming interface. Such errors can be directly attributable to the application programmer-for example, inconsistent user data passed to the programming interface.

#### oocSystemError

Indicates a nonfatal system error detected by an Objectivity/DB operation, or a nonfatal user error that slipped through the consistency checks in the programming interface. In-depth analysis is usually required to determine whether the error is attributable to Objectivity/DB or to the user.

#### oocFatalError

Indicates the most severe type of error that can occur in the system. Such errors are signalled when Objectivity/DB detects an unrecoverable internal inconsistency that might have already caused data corruption. An application should respond to such errors by aborting the active transaction and shutting down as quickly as possible.

See also <u>ooError</u> ooSignal

### ooExitCleanup

global function

Leaves Objectivity/DB in a safe state for process termination.

```
void ooExitCleanup();
```

Discussion During normal process termination, Objectivity/DB calls various destructors to shut itself down. To prepare Objectivity/DB for shutdown, a multithreaded application running on a Windows platform must call the <code>ooExitCleanup</code> function before the application exits. This function leaves Objectivity/DB in a safe state for process termination. On a Windows platform, <code>ooExitCleanup</code> ensures that the Objectivity/C++ dynamic load libraries (DLLs) terminate properly.

You should call this function before returning from your main function and before any call to exit (for example, in an application-defined signal handler). It is good programming practice to terminate all threads that perform Objectivity/DB operations before you invoke <code>ooExitCleanup</code>.

**WARNING** Executing <code>ooExitCleanup</code> must be the *last* Objectivity/DB operation in an application. In particular, an application must call <code>ooExitCleanup</code> after all threads have finished performing Objectivity/DB operations and after all instances of Objectivity/C++ classes have been destructed (see the discussion below). If <code>ooExitCleanup</code> is not the last Objectivity/DB operation in an application, undefined results (such as an access violation or data corruption) may occur.

To ensure that all instances of Objectivity/C++ classes are destructed before your call to <code>ooExitCleanup</code>:

- You must explicitly delete any dynamically allocated instances of Objectivity/C++ classes before you call ooExitCleanup.
- You must not use global instances of Objectivity/C++ classes.
- You must not declare instances of Objectivity/C++ classes inside the same block that contains a call to ooExitCleanup.

You *must* call <code>ooExitCleanup</code> explicitly in any multithreaded application that will run on a Windows platform and any multithreaded application that must support future portability to Windows. An application that is intended to run *only* on UNIX platforms can omit the call to <code>ooExitCleanup</code>, because it is invoked implicitly during the current UNIX shutdown process.

Although <code>ooExistCleanup</code> is not required in a single-threaded application, you can call it from any application. In a single-threaded application, <code>ooExitCleanup</code> aborts the active transaction, if any.

In a multithreaded application, the <code>ooExitCleanup</code> function:

- Aborts all active transactions in all threads.
- Leaves the calling thread executing, along with any thread that has no Objectivity context, or whose Objectivity context has been set to null.
- Suspends or terminates any other thread (that is, any thread with a nonnull Objectivity context). Whether threads are suspended or terminated depends on the platform (see *Installation and Platform Notes* for your platform). You must not attempt to restart threads suspended by ooExitCleanup.

# ooFileNameFormat

global type

(administration) Type for specifying the output format of filenames.

Constants oocNative

Specify filenames as full pathnames—for example, /net/mach3/usr/mnt/project/myfd.FDDB.

oocHostLocal

Specify filenames in host format hostName::localPath—for example, mach3::/mnt/fred/project/myfd.FDDB.

# ooGetActiveTrans

### global function

(*administration*) Returns an array of transaction-information structures describing the active transactions for a particular federated database or autonomous partition; used in database recovery tools.

```
ooStatus ooGetActiveTrans(
    ooTransInfo **ppTrans,
    char **ppBootFilePath,
    char *pHost,
    unsigned int *pUid);
```

### Parameters

ppTrans

Pointer to the returned array (an ooTransInfo\*) of transaction-information structures describing the active transactions. This array is allocated by ooGetActiveTrans. The array is terminated by an entry whose transaction identifier is oocInValidTransId. The array of transaction-information structures is valid until a subsequent call to ooGetActiveTrans.

### *ppBootFilePath*

Pointer to the string name (a char\*) of a federated-database or autonomous-partition boot file. If \*ppBootFilePath is 0, the boot file path is obtained from the environment variable OO\_FD\_BOOT and the pointed-to string is set to the obtained name.

### pHost

Host for filtering the active transactions. If specified, the returned array describes only transactions started on the specified host, subject to filtering based on other parameters. Specify 0 to include transactions started on any node.

### pUid

User identifier for filtering the active transactions. If specified, the returned array describes only transactions started by the specified user, subject to filtering based on other parameters. Specify 0 to include transactions started by any user.

Returns oocSuccess if successful; otherwise oocError.

Discussion This function uses the journal files of the specified federated database or autonomous partition to identify the active transactions to be included in the array.

You use this function only in special-purpose recovery tools. When you use this function, you must not call the <code>ooInit</code> function or perform any other nonrecovery Objectivity/DB operations (such as starting a transaction) in the same application. You may, however, use <code>ooGetActiveTrans</code> along with

	<u>ooCleanup</u> or <u>ooGetResourceOwners</u> . You should use ooGetActiveTrans only in single-threaded applications.
	When you use the <code>ooGetActiveTrans</code> function in an application, you must include the Objectivity/C++ header file <code>ooRecover.h.</code> UNIX applications additionally need to be linked with the Objectivity/DB administration library (see <i>Installation and Platform Notes for UNIX</i> ).
See also	"Creating a Recovery Application" on page 528 of the Objectivity/C++ programmer's guide
<b>A</b> 1 <b>F</b>	

# ooGetErrorHandler

### global macro

Returns a pointer to the currently registered error handler function.

ooErrorHandlerPtr ooGetErrorHandler();

Returns Pointer to the currently registered error handler function.

Discussion In a multithreaded application, this macro returns a pointer to the error handler that is registered in the current Objectivity context.

See also <u>ooRegErrorHandler</u>

# ooGetMemberOffset

global macro

Provides information required for initializing a key structure for a class—specifically, the byte offset from the start of the class to the member field that is to serve as the key.

uint32 ooGetMemberOffset(className, memberName);

Parameters	className

Name of the persistence-capable class for which a key structure is to be initialized.

memberName

Name of the nonstatic data member that is to serve as the key. This data member must be public.

Returns Offset to the member field in number of bytes from the start of the class.

See also <u>ooKey</u>

<u>ooNewKey</u>

# ooGetMemberSize

global macro

Provides information required for initializing a key structure for a class—specifically, the size of the member field in the class.

```
uint32 ooGetMemberSize(className, memberName);
```

 Parameters
 className

 Name of the persistence-capable class for which a key structure is to be initialized.

 memberName

 Name of the nonstatic data member that is to serve as the key. This data member must be public.

 Returns
 Size of the member field in bytes.

 See also
 OOKey

 OONewKey

# ooGetMsgHandler

global macro

Returns a pointer to the currently registered message handler function.

ooMsgHandlerPtr ooGetMsgHandler();

Returns Pointer to the currently registered message handler function.

Discussion In a multithreaded application, this macro returns a pointer to the message handler that is registered in the current Objectivity context.

See also <u>ooRegMsgHandler</u>

# ooGetOfflineMode

global function

(*FTO*) Gets a process' current offline mode. The offline mode determines whether an application ignores or enforces the offline status of all autonomous partitions.

ooOfflineMode ooGetOfflineMode();

Returns Either oocIgnore or oocEnforce.

### ooGetResourceOwners

### global function

(*administration*) Returns information about the resource for which the specified transaction is waiting, along with the transaction(s) currently holding that resource; used in database recovery tools.

```
ooStatus ooGetResourceOwners(
    ooTransInfo **ppOwners,
    ooResource *pResource,
    char **ppBootFilePath,
    ooTransId tId);
```

#### Parameters ppOwners

Pointer to an array (an <code>ooTransInfo\*</code>) of the transaction-information structures for the transactions that own the resource in question. This array is allocated by <code>ooGetResourceOwners</code>. The array is terminated by an entry whose transaction identifier (*tId* member) is set to <code>oocInValidTransId</code>. The array of transaction-information structures is valid until a subsequent call to <code>ooGetResourceOwners</code>.

#### pResource

Pointer to the resource-information structure describing the resource for which the transaction is waiting.

#### ppBootFilePath

Pointer to the string name (a char\*) of the boot file that names the lock server to be queried. You may specify a federated database or autonomous partition boot file. If *\*ppBootFilePath* is 0, the boot file path is obtained from the environment variable OO\_FD\_BOOT and the pointed-to string is set to the obtained name.

#### tId

Transaction identifier of the waiting transaction for which resource information is to be returned.

Returns oocSuccess if successful; otherwise oocError.

### Discussion This function queries the lock server that is listed in the boot file to identify:

- The resource for which the specified transaction is waiting. A resource is an Objectivity/DB object, typically a container.
- The resource's owner(s). This may be a single transaction with a read or update lock on the resource, or one or more transactions with MROW read locks and possibly a transaction with an update lock.

You use this function only in special-purpose recovery tools. When you use this function, you must not call the <code>oolnit</code> function or perform any other non-recovery Objectivity/DB operations (such as starting a transaction) in the

	<pre>same application. You may, however, use ooGetResourceOwners along with ooCleanup or ooGetActiveTrans. You should use ooGetResourceOwners only in single-threaded applications.</pre>
	When you use the <code>ooGetResourceOwners</code> function in an application, you must include the Objectivity/C++ header file <code>ooRecover.h</code> . UNIX applications additionally need to be linked with the Objectivity/DB administration library (see <i>Installation and Platform Notes for UNIX</i> ).
See also	"Creating a Recovery Application" on page 528 of the Objectivity/C++ programmer's guide
ooIndexMo	de global type
	Type for specifying how a transaction updates indexes after creating or modifying objects of an indexed class.
Constants	oocInsensitive
	All applicable indexes are updated automatically when the transaction commits.
	oocSensitive
	All applicable indexes are updated automatically when the next predicate scan is performed in the transaction or, if no scans are performed, when the transaction commits. This allows you to change indexed objects and then scan them in the same transaction using any applicable index.
	oocExplicitUpdate
	Applicable indexes are updated only by explicit calls to the ooUpdateIndexes function.
See also	<u>ooUpdateIndexes</u> ooTrans:: <u>start</u>
oolnit	global function
	Initializes Objectivity/DB in a single-threaded or multithreaded application.
	<pre>ooStatus ooInit( uint32 nFiles = 12, uint32 nPages = 200, uint32 nMaxPages = 500, ooBoolean installSigHandler = oocTrue);</pre>

#### Parameters *nFiles*

Number of file descriptors to be reserved for Objectivity/DB. The reserved file descriptors are a subset of those allocated for the entire process. *nFiles* limits the number of files that Objectivity/DB can have open concurrently in an Objectivity context. When this limit is reached, Objectivity/DB must close files before opening additional files. (Note that applications have no explicit way to control when files are closed.) In a multithreaded application, the *nFiles* limit applies to *each* Objectivity context created by the application.

#### nPages

Initial number of buffer pages to be allocated for each buffer pool in the Objectivity/DB cache. This parameter affects the cache in the main thread; the cache sizes in additional Objectivity contexts are set by the ooContext constructor.

### nMaxPages

Maximum number of buffer pages that can be allocated for each buffer pool in the Objectivity/DB cache. This number is limited by the amount of available swap space. If you specify 0, the cache may grow as needed up to the amount of available swap space. This parameter affects the cache in the main thread; the cache sizes in additional Objectivity contexts are set by the ooContext constructor.

### installSigHandler

Specifies whether to register the predefined Objectivity/DB signal handler. If you omit this parameter, the predefined signal handler is registered. The signal handler is global to the process; the same signal handler is used in every thread.

Discussion Your application must call <code>ooInit</code> early in <code>main()</code> and before using any Objectivity/DB services. If your application has multiple threads, it must call <code>ooInit</code> in the main thread (the thread that starts implicitly when you start the application) before starting any other threads. You should call <code>ooInit</code> only once in an application. Objectivity/DB ignores subsequent calls to this function. Do not call <code>ooInit</code> before <code>main()</code>—for example, do not call it in a global constructor.

The ooInit function implicitly calls the ooInitThread function for the main thread of the application, initializing that thread and creating an Objectivity context for it. A single-threaded application has just this single Objectivity context; a multithreaded application normally creates additional Objectivity contexts.

Every Objectivity context has an Objectivity/DB cache, which is a portion of the process' virtual memory that is reserved for caching persistent objects. The *nPages* and *nMaxPages* parameters of the ooInit function affect the size of the

Objectivity/DB cache in the Objectivity context of the main thread. The cache sizes in other, subsequently-created Objectivity contexts are set by similar parameters of the ooContext constructor.

An Objectivity/DB cache consists of *buffer pages*, which are the same size as the storage pages in the federated database opened by the application. Within the cache, buffer pages are organized into two *buffer pools*, plus a memory pool:

- The small-object buffer pool of pages containing small objects (objects that are a page in size or smaller).
- The large-object buffer pool of header pages for large objects (objects that span multiple pages).
- The large-object memory pool of dynamically allocated memory blocks for large objects.

The *nPages* and *nMaxPages* parameters determine the minimum and maximum sizes of each of the two buffer pools within the cache. If no large objects are opened, the maximum size of the Objectivity/DB cache is the maximum size of the two buffer pools, given by the formula: 2 \* *nMaxPages* \* *pageSize* bytes, where *pageSize* is the number of bytes per storage page in the federated database. When large objects are opened, the cache increases by the size of the large-object memory pool, whose limit is controlled by the <u>ooSetLargeObjectMemoryLimit</u> function.

The default values for *nPages* and *nMaxPages* are usually sufficient for early versions of an application. However, because cache size affects performance, you may want to try different values during performance tuning.

Do not call the ooInit function in applications that call ooCleanup, ooGetActiveTrans, or ooGetResourceOwners.

Returns oocSuccess if successful; otherwise oocError.

See also <u>ooSetLargeObjectMemoryLimit</u> <u>ooInitThread</u>

# oolnitThread

### global function

Initializes the current thread so that it can execute Objectivity/DB operations.

void ooInitThread (ooContext \*context = new ooContext());

Parameters context

Objectivity context to be set for the thread. If you omit this parameter, a new Objectivity context is created for the thread. If you specify 0, a null Objectivity context is set for the thread; you must set a nonnull context explicitly before invoking any Objectivity/DB operations in that thread. If

	you specify an existing Objectivity context, that context is set unless it is already the context for another thread; in this case, an error is signaled.
Discussion	Every thread that invokes Objectivity/DB operations must call this function before any of those operations are invoked.
	It is <i>not</i> necessary to call this function:
	<ul> <li>In the process' main thread. This thread is initialized automatically by the ooInit function, which implicitly calls the ooInitThread function.</li> <li>In threads that do not invoke Objectivity/DB operations.</li> </ul>
See also	<u>ooTermThread</u> ooContext:: <u>setCurrent</u>
ooKey	global type
	Type for specifying the key structure for a keyed object.
	<pre>struct ooKey {     ooKeyType type;     uint32 offset;     uint32 size;     void *value; };</pre>
Members	<i>type</i> Type of the key field in the keyed object.
	offset Offset of the key field (in bytes) from the beginning of the keyed object's class.
	size Size of the key field (in bytes) in the keyed object's class.
	value Pointer to the value for which you want to search.
Discussion	A key structure provides Objectivity/DB with information about a specific key field, including the value for which you want to search.
	When <i>type</i> is oocString, <i>value</i> must be a character array ending in a null terminator. If you specify a <i>value</i> that is shorter than <i>size</i> , it is padded with null characters. If you specify a <i>value</i> that is longer than <i>size</i> , an error results. When <i>type</i> is oocCharArray, the character array <i>value</i> has length <i>size</i> and may include null characters.

See also <u>ooKeyType</u> <u>ooNewKey</u> <u>ooGetMemberOffset</u> <u>ooGetMemberSize</u>

### ooKeyType

global type

Type for specifying the data type of a key field within a keyed object. Constants of this type are used in creating key structures.

Constants oocUint16

Objectivity/C++ type uint16.

oocUint32

Objectivity/C++ type uint32.

oocInt16

Objectivity/C++ type int16.

oocInt32

Objectivity/C++ type int32.

oocFloat32

Objectivity/C++ type float32.

oocFloat64

Objectivity/C++ type float64.

oocString

A fixed array of characters (type char \*). In a key structure, a value of this type must have a null terminator. If you specify a value containing multiple null characters, only the characters up to and including the first null character are used.

#### oocCharArray

Character array with no particular terminator character. In a key structure, a value of this type may include null characters. The length of the character array is determined by the size of the key field, which is specified in the key structure.

See also <u>ooKey</u>

ooLockMoo	de global type
	Type for specifying how to lock an Objectivity/DB object.
Constants	oocLockRead The object is locked for read. Other processes may also obtain read locks on the object. If all read locks on an object were obtained within multiple reader, one writer (MROW) transactions, another process may obtain an update lock; however, if a read lock on an object was obtained within a standard (non-MROW) transaction, no other process may obtain an update lock.
	oocLockUpdate The object is locked for update. No other process may obtain an update lock on the object. A process running an MROW transaction may obtain a read lock on the object; a process running a standard (non-MROW) transaction may not obtain a read lock.
ooMode	global type
	Type specifying the intended level of access to an Objectivity/DB object or for specifying the <u>concurrent access policy</u> for a transaction.
Constants	oocNoOpen (Used by various open, lookup, and scan operations.) A reference is set to the object without opening or locking it.
	oocRead (Used by various open, lookup, and scan operations.) The object is opened only for read.
	oocUpdate (Used by various open, lookup, and scan operations.) The object is opened for update (read and write).
	oocNoMROW (Used by ooTrans::start.) The <u>standard</u> concurrent access policy is activated for the transaction being started.
	OOCMROW (Used by OOTrans::start.) The multiple reader, one writer ( <u>MROW</u> ) concurrent access policy is activated for the transaction being started.

# ooMsgHandlerPtr

global type

Function-pointer type for application-defined message handler functions.

```
typedef void (*ooMsgHandlerPtr)(
     char *message);
```

Parameters message

String containing the message to be displayed.

Discussion A registered message handler is a function that is called by the Objectivity/DB-defined error handler to write a signaled error's message to an output device. Objectivity/DB provides a predefined message handler that is automatically registered when you start your application. This predefined message handler prints error messages to stderr. You can replace the predefined message handler by defining a custom message handler and using the ooRegMsgHandler function to register it.

> Your custom message handler must be able to accept the argument passed to it by the predefined error handler. Therefore, your message handler must conform to the calling interface defined by this function-pointer type.

See also <u>ooRegMsgHandler</u>

# ooNameHashFuncPtr

global type

Function-pointer type for an application-defined hash function to be used by instances of ooMap.

```
typedef uint32 (*ooNameHashFuncPtr)(
    const char *name,
    uint32 size);
```

Parameters name

String name to be hashed.

size

Number of bins in the hash table.

See also ooMap::<u>set\_nameHashFunction</u>

### global macro

Creates a batch of containers in the specified database. This macro provides better performance than calling the new operator repeatedly.

```
void ooNewConts(
    className,
    uint32 numberOfConts,
    const ooHandle(ooObj) &near,
    uint32 hash,
    uint32 initPages,
    uint32 percentGrowth,
    ooBoolean open,
    ooHandle(ooObj) *pHandle);
```

Parameters className

Class name of the containers to be created. The class must have a default constructor.

### number Of Conts

Number of containers to be created. This number is limited to 32766.

near

Clustering directive that specifies where to locate the new containers:

- If *near* is a handle to a database, the new containers are created in that database.
- If *near* is a handle to a container or basic object, the new containers are created in the database that contains the referenced container or basic object.

An error is signalled if *near* is 0.

hash

Determines whether to create hashed containers. You must create hashed containers if you intend to use them or any objects in them as scopes for naming objects, or if you intend to create keyed objects in the containers.

- Specify 0 to create nonhashed containers.
- Specify 1 or greater to create hashed containers.

The number you specify is the clustering factor for any keyed objects created in the containers. A clustering factor is the number of sequentially keyed objects to be placed onto a page. A clustering factor of 1 maximally distributes keyed objects across pages. A higher number means fewer pages need to be read when finding sequences of keyed objects.

#### init Pages

Initial number of logical pages to be allocated for each container. Specify 0 to use the system default value (4 pages for hashed containers, and 2 pages for nonhashed containers). The maximum value is 65535.

#### percentGrowth

Amount by which each container may grow, expressed as a percentage of its current size. Specify 0 to use the system default value (10%).

open

Specifies whether to open or close the batch containers after creation. Specify <code>oocTrue</code> to leave all the containers open; specify <code>oocFalse</code> to close all the containers. It is better to close all the containers if they will not be used immediately.

### pHandle

Pointer to an array of container handles that, on return, stores the handles to the newly created containers. The number of handles in this array should be greater than or equal to *numberOfConts*.

Discussion You can use the first handle in the handle array for error checking—if the operation is successful, the first handle is not 0; otherwise, the first handle is 0.

This macro does not provide a way to give a system name to each container. You may not use this macro as an expression—for example, in an assignment statement.

See also ooContObj::operator new

### ooNewKey

global macro

Creates a keyed object for the specified class.

```
ooHandle(ooObj) ooNewKey(
    className,
    (initializer),
    const ooHandle(ooObj) &contH,
    ooKey keyStruct);
```

### Parameters className

Name of the class for which a keyed object is to be created. The class must be a basic object class.

### initializer

List of arguments to be passed to the constructor for class *className*.

	<pre>contH Handle to the container in which the keyed object is to be creat container must be a hashed container. If you specify the handle the keyed object is created in the default container for the data</pre>	ted. The e to a database, base.
	<i>keyStruct</i> The key structure with which the keyed object is created.	
Returns	Handle to a newly created keyed object.	
See also	<u>ooKey</u> <u>ooGetMemberOffset</u> <u>ooGetMemberSize</u> <i>ooRefHandle</i> (ooObj):: <u>lookupObj</u>	
ooNoLock		global function

Disables the Objectivity/DB locking facilities, removing concurrent access protection.

ooStatus ooNoLock();

Returns oocSuccess if successful; otherwise oocError.

Discussion If your application is guaranteed exclusive access to a federated database and it requires maximum performance, you may consider disabling locking to remove the runtime overhead associated with managing locks. However, if another process has access to the same data as your application, unpredictable results may occur, including corruption of your data. Therefore, use ooNoLock with caution. For most applications, benefits such as data integrity and concurrent access far outweigh the slight performance gain obtained by disabling locking.

The ooNolock function is called after ooInit, but before any transactions are started.

In a multithreaded application, this function disables locking for transactions in the current Objectivity context.

### ooOfflineMode

global type

(*FTO*) Type for specifying how an application should respond to the offline status of all autonomous partitions.

### Constants oocIgnore

The application ignores the offline status of autonomous partitions.

### oocEnforce

The application enforces the offline status of autonomous partitions.

See also	<u>ooGetOfflineMode</u>
	<u>ooSetOfflineMode</u>

# ooPurgeAps

global function

(FTO) Purges the specified autonomous partitions from the federated database.

	ooStatus ooPurgeAps(char ** <i>apNames</i> , int <i>numOfAps</i> );
Parameters	apNames An array of strings, where each string is the system name of an autonomous partition to be purged. numOfAps The number of numbers.
Poturos	The number of system names in apNames.
Returns	oocsuccess if successful, otherwise oocerror.
Discussion	You use this function only in special-purpose recovery tools. You must call this function from within an update transaction (that is, after starting a transaction and opening the federated database for update).
	This function is called to remove permanently inaccessible partitions from the federated database. Purging a partition is similar to deleting it; the difference is that you can purge partitions that are not accessible, but you cannot delete an inaccessible partition.
	If any currently unavailable partition is <i>not</i> specified for purging, this function returns oocSuccess; however, the transaction cannot commit because it can't propagate the catalog changes to the unavailable partition.
	If a quorum of images for a particular database still exists after removal of all the unavailable partitions, this function removes any catalog reference to that database's images and/or tie breakers in the purged partitions. If a quorum of images for a particular database is no longer available, this function deletes that database from the federated-database catalog. The database files, however, are not deleted.
WARNING	To use this function safely, you must be <i>absolutely certain</i> that no process is active in any of the partitions to be purged. Purging partitions cannot be undone.

# ooQueryOperatorPtr

global type

Function-pointer type for application-defined relational-operator functions.

	<pre>typedef ooBoolean (*ooQueryOperatorPtr)(     const void *lPtr,     const void *rPtr,     ooDataType lAType,     ooDataType rAType);</pre>
Parameters	<i>lPtr</i> Pointer to the left operand of the operator.
	<i>rPtr</i> Pointer to the right operand of the operator.
	<i>LAType</i> Indicates the data type of the left operand.
	<i>rAType</i> Indicates the data type of the right operand.
Returns	occTrue if the relation expressed by the operator is satisfied by the operands; otherwise occFalse.
Discussion	An application can define its own relational operators for use in predicates. To do this, the application defines an operator function that tests whether two operands are of an appropriate type and whether the desired relation holds between them. When registered along with a token, the operator function is called by the Objectivity/DB predicate-query mechanism each time the token is encountered in a predicate.
	The operator function must conform to the calling interface of this function-pointer type so that the predicate-query mechanism can pass it the operands to be compared and a pair of constants indicating the operands' data types. The operator function should return $oocTrue$ if the two operands are of the expected type and the desired relation holds between them; otherwise, the function should return $oocTalse$ . For information about the types that are recognized by the predicate query mechanism, see the Objectivity/C++ programmer's guide.
See also	<u>ooDataType</u> global type <u>ooUserDefinedOperators</u> global variable ooOperatorSet:: <u>registerOperator</u>

# ooRegErrorHandler

global macro

Registers the specified error handler function with Objectivity/DB, replacing the previously registered error handler.

```
Parameters handlerName
```

Name of the error handler function to be registered.

### Returns A function pointer to the previously registered error handler.

Discussion A registered error handler is a function that is invoked whenever the ooSignal function signals an error. Objectivity/DB provides a predefined error handler that is automatically registered when you start your application. You can replace the predefined error handler by defining a custom error handler and using ooRegErrorHandler to register it.

In a multithreaded application, the Objectivity/DB-defined error handler is automatically registered in each Objectivity context that your application creates. You can register a custom error handler in a thread's Objectivity context by invoking <code>ooRegErrorHandler</code> in that thread.

ooRegErrorHandler returns a pointer to the replaced error handler so that you can register it again when desired. When an error is signaled, the most recently registered error handler is invoked.

See also Chapter 23, "Error Handling," of the Objectivity/C++ programmer's guide

# ooRegMsgHandler

global macro

Registers the specified message handler function with Objectivity/DB.

ooMsgHandlerPtr ooRegMsgHandler(ooMsgHandlerPtr msgHandler);

Parameters msgHandler

Pointer to the message handler function to be registered.

- Returns A function pointer to the previously registered message handler.
- Discussion A registered message handler is a function that is called by the Objectivity/DB-defined error handler to write a signaled error's message to an output device. Objectivity/DB provides a predefined message handler that is automatically registered when you start your application. This predefined message handler prints error messages to stderr. You can replace the

predefined message handler by defining a custom message handler and using ooRegMsgHandler to register it.

In a multithreaded application, the Objectivity/DB-defined message handler is automatically registered in each Objectivity context that your application creates. You can register a custom message handler in a thread's Objectivity context by invoking <code>ooRegMsgHandler</code> in that thread.

ooRegMsgHandler returns a pointer to the replaced message handler so that you can register it again when desired. The Objectivity/DB-defined error handler invokes the most recently registered message handler.

See also Chapter 23, "Error Handling," of the Objectivity/C++ programmer's guide

# ooRegTwoMachineHandler

global function

(*DRO*) Registers the specified two-machine handler function with Objectivity/DB.

Parameters	<i>twoMachineHandler</i> Pointer to the two-machine handler function to be registered.
Returns	A function pointer to the previously registered two-machine handler.
Discussion	You can call this function after calling <code>ooInit</code> in a single-threaded application or after you call <code>ooInitThread</code> in each thread of a multithreaded application.
	A two-machine handler is relevant only in a two-machine hot-failover configuration in which a hard link between the two machines enables applications on each machine to check on the status of the other machine.
	An application's two-machine handler function is called when only one of two equally weighted database images is available. The function determines whether the application can write to the available database image (typically a local image). If the function returns oocTrue, the application can proceed as if the local database image by itself constituted a quorum; if it returns oocFalse, the local image does not constitute a quorum and the application cannot access the database. An application with no registered two-machine handler function is equivalent to an application with a handler function that always returns oocFalse.
	A two-machine handler function should check whether the other machine is still running. If not, a machine failure has occurred and applications on the remaining machine should be able to access the database. If the other machine is running,

however, a network failure has occurred; applications on one, but not both, of the two machines should be able to continue.

It is your responsibility to ensure coordinated behavior of the two-machine handler functions registered by applications running on the two machines. Typically the same handler function is registered with all applications on a given machine.

**NOTE** In any given network failure, if at least one application on one machine has a two-machine handler function that returns oocTrue, then no application on the other machine should have a two-machine handler function that returns oocTrue.

# ooReplace

global macro

Creates a database with the specified name in the specified federated database, replacing (deleting) any existing database with the same name.

```
ooHandle(ooDBObj) ooReplace(
        ooDBObj,
        (char *dbSysName,
        uint32 defaultContInitPages = 0,
        uint32 defaultContPercentGrow = 0,
        const char *hostName = 0,
        const char *pathName = 0,
        uint32 weight = 1),
        const ooHandle(ooFDObj) &fdH);
```

### Parameters dbSysName

System name of the database to create. Any existing database with this system name is removed before the new database is created. *dbSysName* follows the same naming rules as files of your operating system, and must be unique within the federated database.

### *defaultContInitPages*

Initial number of logical pages to allocate for the default container in the new database. The maximum value is 65535. If you omit this parameter or specify 0, the system default value (4) is used.

defaultContPercentGrow

Amount by which the default container may grow, expressed as a percentage of its current size. If you omit this parameter or specify 0, the system default value (10%) is used.

	hostName
	Name of the host system on which to create the database file. If you specify a <i>hostName</i> , you must also specify the <i>pathName</i> parameter. If you omit the <i>hostName</i> parameter or specify 0, the database file is created on the same host as the federated database's system-database file.
	pathName
	Full pathname of the directory in which to create the database file. If you specify a <i>pathName</i> , you must also specify the <i>hostName</i> parameter. If you omit the <i>pathName</i> parameter or specify 0, the database file is created in the same directory as the federated database's system-database file. The <i>pathName</i> you specify may, but need not, include the filename for the database file.
	weight
	( <i>DRO</i> ) Weight of the first database image. <i>weight</i> must be an integer greater than zero. If you omit this parameter, the weight is 1.
	fdH
	Handle to the federated database in which to create the new database.
	( <i>FTO</i> ) In a partitioned federated database, the new database is created in the federated database's initial partition.
Returns	If successful, a handle to the newly created database; otherwise, a null handle.
Discussion	This macro is generally used to clean up a federated database for repeated tests.
WARNING	The removed database is <i>not</i> recovered if the transaction is aborted.
	Note that all parameters except the first and last correspond to parameters of the ooDBObj constructor.
Example	The syntax for using ooReplace is somewhat unusual. This example shows how to use ooReplace to replace a database whose system name is testDB. The new database is created in the same location as the federated database's system-database file.
	ooHandle(ooFDObj) fdH; ooHandle(ooDBObj) dbH; // Start transaction, open FD, initialize fdH
	dbH = ooReplace(ooDBObj, ("testDB"), fdH);

# ooResetError

global macro

Clears Objectivity/DB error flags.

void ooResetError();

Discussion This function sets the following global variables:

- oovLastErrorLevel = oocNoError
- oovLastError = 0

In a multithreaded application, <code>ooResetError</code> resets the error flags in the current Objectivity context.

See also Chapter 23, "Error Handling," in the Objectivity/C++ programmer's guide.

### ooResource

global type

Type for representing information about a locked resource (typically a container).

	struct ooResource {
	···· };
Members	 Members contain information about the locked resource, an Objectivity/DB object. These members are subject to change from release to release. See the header file ooRecover.h for the current definition of this structure.
Discussion	If a transaction is waiting for a lock on an Objectivity/DB object (typically a container), a structure of this type, called a <i>resource-information structure</i> , identifies the object and its lock status.
	You should not create structures of this type. Instead, you obtain a resource-information structure by calling $ooGetResourceOwners$ ; you must include the Objectivity/C++ header file $ooRecover.h$ .
ooRunSta	itus global function
	Prints a summary of Objectivity/DB internal statistics to stdout; used primarily for performance tuning.
	<pre>void ooRunStatus();</pre>
Discussion	In a multithreaded application, this function obtains statistics that pertain to the current Objectivity context.
See also	Chapter 24, "Performance," in the Objectivity/C++ programmer's guide

# ooSetAMSUsage

Global Names

global function

(*administration*) Sets the application's policy for using the Advanced Multithreaded Server (AMS).

void ooSetAMSUsage(ooAMSUsage amsUsage = oocAMSPreferred);

Parameters amsUsage

Policy for using AMS. If you omit this parameter, the application uses AMS whenever possible. See <u>ooAMSUsage</u> for other values you can specify.

Discussion The system default is for an application to use AMS whenever possible (that is, the policy is implicitly oocAMSPreferred). If you choose to invoke this function, you should call it after ooInit and before any other Objectivity/DB operation. An error is signaled if you specify oocAMSOnly and AMS is unavailable on the host that contains the requested data. The policy set by this function applies to the entire application, including all of its threads.

# ooSetErrorFile

global function

Sets the error message output file.

ooStatus ooSetErrorFile(char \*errorFileName);

Parameters errorFileName

Name of the error message output file.

Discussion In a multithreaded application, this function sets the error message output file for the current Objectivity context.

# ooSetHotMode

global function

Enables or disables *hot mode*, which controls the timing of certain internal overhead operations in the Objectivity/DB cache, for purposes of improving performance.

```
void ooSetHotMode(ooBoolean hotMode = oocTrue);
```

 Parameters
 hotMode

 Specifies whether to enable or disable hot mode. If you omit this parameter, hot mode is enabled.

 Discussion
 You can call this function anytime after calling ooInit and as often as desired to

Discussion You can call this function anytime after calling <code>ooInit</code> and as often as desired to enable or disable hot mode. In a multithreaded application, this function affects only the Objectivity/DB cache in the current Objectivity context. The system default is for hot mode to be disabled in an Objectivity context unless you explicitly enable it.

Hot mode can improve the performance of an application that repeatedly opens, closes, and reopens persistent objects that were created by applications running on other architectures. For a discussion of performance tradeoffs, see "Use Hot Mode" in Chapter 24 of the Objectivity/C++ programmer's guide.

### ooSetLargeObjectMemoryLimit

global function

Sets the maximum amount of dynamically allocated memory that is available for caching large persistent objects.

```
void ooSetLargeObjectMemoryLimit(uint32 size);
```

Parameters size

Number of bytes in the Objectivity/DB large-object memory pool.

Discussion The Objectivity/DB cache has a both a buffer pool and a memory pool for handling *large objects* (persistent objects that span multiple storage pages). When a large object is created or opened, Objectivity/DB:

- Reads the object's header page into the large-object buffer pool. This page contains overhead information about the object.
- Reads the object's storage pages into the large-object memory pool. These
  pages form a single dynamically allocated block of contiguous buffer pages
  pointed to by the header page.

The ooSetLargeObjectMemoryLimit function sets the suggested limit on the total number of bytes in the large-object memory pool. When this limit is reached, Objectivity/DB attempts to swap out the pages of closed large objects before opening additional large objects. However, if Objectivity/DB cannot find enough closed large objects to swap out, it will ignore the specified limit and allocate additional pages as needed. Thus, the limit you specify is a soft limit that affects the amount of swapping performed on behalf of large objects.

If you do not call the <code>ooSetLargeObjectMemoryLimit</code> function, the default limit on the large-object memory pool the equal to the maximum size of the large-object buffer pool (see <code>ooInit</code>).

You can call the <code>ooSetLargeObjectMemoryLimit</code> function anytime after calling <code>ooInit</code> and as often as needed to increase or decrease the limit on the large-object memory pool. In a multithreaded application, this function sets the limit for the current Objectivity context.

# ooSetLockWait

global function

Sets the default lock-waiting option for a series of transactions.

```
void ooSetLockWait(int32 waitOption = oocNoWait);
```

### Parameters waitOption

Specifies whether transactions are to wait to obtain locks, and if so, for how long:

- Specify oocNoWait or 0 to turn off lock waiting.
- Specify oocWait to cause transactions to wait indefinitely for locks.
- Specify a timeout period of *n* seconds, where *n* is an integer in the range 1 <= *n* <= 14400. If *n* is less than 0 or greater than 14400, it is treated as oocWait.
- Discussion By default, the transactions in an application do not wait for locks. You can override the default behavior on a per-transaction basis by setting the *waitOption* parameter of the ooTrans::<u>start</u> member function. Alternatively, you can change the default behavior by using ooSetLockWait.

Calling <code>ooSetLockWait</code> within a transaction overrides any lock-waiting option you specified through the <code>ooTrans::start</code> member function when you started that transaction. The option you set with <code>ooSetLockWait</code> becomes the new default, which subsequent transactions can choose to use or override.

Lock waiting does not apply to MROW read transactions. Therefore, the lock-waiting option you specify through ooSetLockWait is ignored by such transactions.

In a multithreaded application, this function sets the default lock-waiting option for the current Objectivity context.

# ooSetOfflineMode

global function

(FTO) Sets the offline mode for the current process.

ooStatus ooSetOfflineMode(ooOfflineMode offlineMode);

Parameters offlineMode

Offline mode to be set. Specify occlgnore to ignore the offline status of autonomous partitions. Specify occEnforce to enforce the offline status.

Returns oocSuccess if successful; otherwise oocError.

Discussion Offline mode determines whether a process ignores or enforces the offline status of all autonomous partitions. If you do not call this function, offline status is enforced.

# ooSetRpcTimeout

global function

Sets how long an application is to wait for an Objectivity server to respond before signaling a timeout error.

void ooSetRpcTimeout(long seconds);

Parameters seconds

Number of seconds to wait before signaling a timeout error.

Discussion By default, an application waits 25 seconds for the lock server or AMS to respond to a request. However, an Objectivity server running on a busy host machine may need more time to respond. If your application consistently signals a lock server or AMS timeout error, you can call this function to increase the timeout period; alternatively, you can consider running the Objectivity server on a less congested host.

In a multithreaded application, this function sets the timeout period for the current Objectivity context.

# ooSignal

global function

Signals the specified error and reports additional error information.

1.	<pre>ooStatus ooSignal(   const ooErrorLevel errorLevel,   const ooError &amp;errorID,</pre>
	<pre>const ooHandle(ooObj) *relevantObj,);</pre>
2.	<pre>ooStatus ooSignal( const ooErrorLevel errorLevel, const ooError &amp;errorID, const ooHandle(ooObj) relevantObj, );</pre>

### Parameters errorLevel

Error level for this event.

### errorID

Error identifier for this event. Error identifiers must be declared and initialized in an error message header file, which is included in your

application source file. An error identifier consists of an error number and an error message.

```
relevant0bj
```

Handle or pointer to the handle of an Objectivity/DB object that may help the system error handler to pinpoint the situation in which the error condition occurred. Set this parameter to 0 if there is no relevant Objectivity/DB object.

...
 Variable number of parameters required by the format string stored in the message part of the error identifier. These parameters provide information about the situation in which the error condition occurred. They are substituted for conversion specifications (such as %s or %d) in the format string to construct the error message.
 Returns oocSuccess if successful; otherwise oocError.
 Discussion ooSignal signals the specified error, invokes the currently registered error

handler, and returns the result from the error handler. It also sets Objectivity/C++ global variables as follows:

- oovLastErrorLevel is set to *errorLevel*.
- oovLastError is set with a pointer to errorID.
- oovNError is incremented.

In a multithreaded application, this function signals an error (and invokes the registered error handler) in the current Objectivity context.

# ooStartInternalLS

global function

(IPLS) Starts an in-process lock server within an application.

```
ooStatus ooStartInternalLS(
    const char *fdName = 0
    void (*threadFn)() = 0);
```

Parameters fdName

Pathname of the boot file of the federated database or autonomous partition that the application intends to open. This parameter controls the conditions under which an in-process lock server is started:

- If you specify 0 (the default), an in-process lock server is started whether or not it will be used by the calling application.
- If you specify a boot file, an in-process lock server is started only if the calling application will use it—that is, only if the lock server host shown

in the boot file is the current host. If the lock server host is *not* the current host, an error is signalled and an in-process lock server is not started.

#### threadFn

Pointer to an application-defined function with no arguments that will be called from the in-process lock server's listener thread. A listener thread is created to service lock requests from external applications.

The specified function is called after the listener thread has initialized its Objectivity context and before it begins servicing requests. The function could be used to customize the behavior of the thread, such as adjusting its priority or installing an application-defined error handler to direct how error messages originating in that thread will be handled.

Returns oocSuccess if the in-process lock server was successfully started; otherwise, oocError—for example, if a separate lock server process is already running on the same host machine, or if some operating-system resource could not be obtained.

> If *fdName* is nonzero, *oocSuccess* indicates the in-process lock server is started and will be used by the application; *oocError* indicates that the boot file specified by *fdName* cannot be opened or specifies a lock server host that is different from the current host.

# Discussion This function must be called after ooInit is called and before the federated database is opened for the first time.

An application cannot start an in-process lock server if another lock server process is already running on the same host. If the application chooses to continue running, it will use the lock server that is already running on that host. An application can call the <code>ooCheckLS</code> function to test for a running lock server.

When an application successfully starts an in-process lock server:

- The application becomes the lock server process for the workstation on which it is running. If a federated database names that workstation as its lock server host, all applications accessing the federated database will send their lock requests to the application running the in-process lock server. The in-process lock server uses a separate *listener thread* to service requests from external applications.
- The application uses its own in-process lock server only if the opened federated database names the application's host as the lock server host. That is, an application always uses the lock server specified by the federated database, whether or not the application is running an in-process lock server. You can find out whether an application will use its in-process lock server by specifying the *fdName* parameter to <code>ooStartInternalLS</code>.

An in-process lock server improves performance only if most or all of the lock requests for a given federated database originate from a single multithreaded application process. The in-process lock server can then coordinate locking through simple function calls instead of servicing lock requests over the network.

See also <u>ooCheckLS</u> <u>ooStopInternalLS</u>

### ooStatus

global type

General return type for Objectivity/C++ global functions and member functions.

Constants oocSuccess

Indicates a successful outcome. Its value is a nonzero integer.

oocError

Indicates that an error occurred. Its value is the integer 0.

# ooStopInternalLS

global function

(IPLS) Shuts down an in-process lock server within an application.

```
ooStatus ooStopInternalLS(
    int wait = INT_MAX,
    ooBoolean force = oocFalse);
```

### Parameters wait

Number of seconds to wait for active transactions to terminate. The default value, INT\_MAX, means there is no time limit, so this function will wait indefinitely. If active transactions terminate within the specified wait period, the in-process lock server is shut down. Otherwise, if transactions are still active when the specified wait period expires, this function takes the action specified by *force*.

force

Specifies whether to shut down the in-process lock server if active transactions have not yet terminated by the end of the *wait* period:

- If you omit this parameter or specify oocFalse, the in-process lock server continues running and this function returns oocError.
- If you specify oocTrue, the in-process lock server is shut down, even if transactions are active, and this function returns oocSuccess.

Returns oocSuccess if the in-process lock server is successfully shut down, or if it was not running; otherwise oocError.
#### Global Names

Reference Descriptions

Discussion This function safely shuts down an in-process lock server so that you can terminate the application in which it is running without harming any other applications that may be using its in-process lock server.

This function should be called at the end of the application that is running the in-process lock server, after committing or aborting any transactions, and before calling <code>ooExitCleanup</code> or <code>exit</code>.

This function causes the in-process lock server to refuse any new transactions from external client applications, and to wait for any active transactions to finish during the period specified by *wait*. During the wait period, the in-process lock server may accept new client connections (for example, to allow administration tools to run), but continues to disallow new transactions. When all active transactions have finished, the in-process lock server is shut down. At this point, the application may safely exit.

If active transactions do not finish and the wait period expires, this function either shuts down the in-process lock server or allows it to continue running, according to the *force* parameter. If the in-process lock server continues running, the application can repeat the attempted shutdown by calling this function again later.

See also <u>ooStartInternalLS</u>

### ooTermThread

global function

Terminates the current thread's ability to invoke Objectivity/DB operations.

ooStatus ooTermThread();

Discussion Threads initialized with the <code>ooInitThread</code> or <code>ooInit</code> function must call the <code>ooTermThread</code> function after completing their Objectivity/DB operations and before termination. It is not, however, necessary to call <code>ooTermThread</code> in threads that are terminated due to process termination. So, for example, the main thread need invoke <code>ooTermThread</code> only if it terminates before the process terminates.

You may not reinitialize a thread after calling ooTermThread. You may not invoke ooTermThread more than once per thread.

The ooTermThread function deletes the thread's current Objectivity context, unless this context is the null context. Because Objectivity contexts are usually allocated dynamically, ooTermThread invokes the delete operator to delete the current context. Consequently, if the current Objectivity context is either statically allocated or a dynamically allocated member of a larger object, you must set the current context to 0 before using ooTermThread, and then arrange to delete the Objectivity context as appropriate.

See also	<u>ooInitThread</u>
	ooContext:: <u>setCurrent</u>

## ooTransId

global type

Unique identifier for a transaction.

Constants oocInValidTransId
 An invalid transaction identifier. This value is used as the identifier in a transaction-information structure that signals the end of an array of transaction-information structures.
 Discussion Every transaction has a *transaction identifier* of type ooTransId that uniquely identifies it to the lock server. Recovery functions use parameters of this type to identify a transaction of interest. Administration tools such as oolockmon and oolistwait display transaction identifiers in their output.

You can get a transaction identifier of an active transaction by calling the getID member function on a transaction object. Alternatively, you can obtain the transaction identifier for a particular transaction from a transaction-information structure of type ooTransInfo that describes the transaction; you must include the Objectivity/C++ header file ooRecover.h.

See also ooTrans::getID

## ooTransInfo

global type

Type for representing information about a transaction.

	struct ooTransInfo { ooTransId <i>tid;</i>
	}; 
Members	tid
	Transaction identifier for the described transaction.
	Additional members contain information about the described transaction. These members are subject to change from release to release. See the header file ooRecover.h for the current definition of this structure.
Discussion	A structure of this type, called a <i>transaction-information structure</i> , describes a particular transaction. You should not create structures of this type. Instead, you obtain a transaction-information structure by calling recovery functions; you must include the Objectivity/C++ header file <code>ooRecover.h</code> .

See also	<u>ooCleanup</u>
	<u>ooGetActiveTrans</u>

## ooTwoMachineHandlerPtr

global type

	(DRO) Function-pointer type for application-defined message handler functions.
	<pre>typedef ooBoolean (*ooTwoMachineHandlerPtr)();</pre>
Returns	oocTrue if this application can safely access a database even though only one of two equally weighted images is available; otherwise, oocFalse.
Discussion	A registered two-machine handler is a function that is called when only one of two equally weighted database images is available. The function determines whether the application can write to the available database image (typically a local image). If the function returns <code>oocTrue</code> , the application can proceed as if the local database image by itself constituted a quorum; if it returns <code>oocFalse</code> , the local image does not constitute a quorum and the application cannot access the database. An application with no registered two-machine handler function is equivalent to an application with a handler function that always returns <code>oocFalse</code> .
See also	<u>ooRegTwoMachineHandler</u>
ооТуреN	global macro
	Gets the type number of the specified class of Objectivity/DB objects.
	ooTypeNumber ooTypeN( <i>className</i> );
Parameters	className
	Name of the class whose type number is to be returned. The specified class must be either ooObj or derived from ooObj—that is, a basic-object class, a container class, the database class (ooDBObj), the federated-database class (ooFDObj), or the autonomous-partition class (ooAPObj).
Returns	Type number of the class.
Discussion	Every class in the $000bj$ inheritance hierarchy has a unique type number that identifies the class within the federated-database schema. You can use this number to determine whether an object is an instance of a particular class.
	An error is signalled if the specified class is not in the $000bj$ inheritance hierarchy.

Because ooTypeN is expanded into a variable name, you cannot use it as a label in a switch statement; to do so causes a compiler error. However, you can use ooTypeN as part of a conditional expression in an if-else statement.

Example This example obtains the type number of a referenced object and tests it against the type number of several different classes:

```
ooHandle(Fruit) fruitH;
ooTypeNumber typeNum;
fruitH = ... // Set fruitH to reference some kind of fruit
// Set typeNum to the type number of the referenced fruit
typeNum = fruitH.typeN();
if (typeNum == ooTypeN(Apple)) {
    ... // perform operation for apples
}
else if (typeNum == ooTypeN(Orange)) {
    ... // perform operation for oranges
}
else if (fruitH->ooIsKindOf(ooTypeN(Berry))) {
    ... // perform operation for berries
}
```

See also <u>ooTypeNumber</u> ooObj::<u>ooIsKindOf</u> *ooRefHandle*(classname)::typeN

### ooTypeNumber

global type

Type number of a class in the ooObj inheritance hierarchy.

Discussion The ooObj class and every class that derives from it has a unique type number that identifies the class within the federated-database schema. You can use this number to determine whether an object is an instance of a particular class.

See also <u>ooTypeN</u> ooObj::<u>ooIsKindOf</u> *ooRefHandle(classname)*::<u>typeN</u>

## ooUpdateIndexes

global function

Explicitly updates all applicable indexes to reflect a new or modified object.

ooStatus ooUpdateIndexes(ooHandle(ooObj) &pHandle);

Parameters *pHandle* 

Handle to the new or modified object for which indexes are to be updated.

#### Returns oocSuccess if successful; otherwise oocError.

Discussion Each transaction has an index mode that determines when indexes are to be updated for objects that are created or changed during the transaction. If a transaction's index mode is set to occExplicitUpdate, indexes are updated only when you call the ooUpdateIndexes function. You must call this function once for each indexed object that is created or modified during the transaction—for example, after you have created a new object of an indexed class and initialized all its key fields, or after you have changed any of an existing object's key-field values. Explicit updates are recommended for update-intensive applications that use indexes over database and federated-database scopes.

See also <u>ooIndexMode</u>

#### ooUseIndex

global function

Enables or disables the use of indexes during a predicate query.

void ooUseIndex(ooBoolean useIndex = oocTrue);

Parameters useIndex

Specifies whether the scan member function of an iterator can use indexes during a predicate query. Specify oocTrue to use indexes during scan; specify oocFalse to disable the use of indexes.

#### Discussion Indexes are used in the optimization of certain kinds of predicate queries. Disabling the use of indexes may be desirable when:

- A query is going to scan the entire range of the type and sorting is not necessary. In such a case, indexes do not speed up your query.
- The data to be scanned has been created or modified since the last time indexes were updated.

If you do not call this function, the use of indexes is disabled. In a multithreaded application, this function controls the use of indexes in transactions in the current Objectivity context.

### ooUserDefinedOperators

global variable

Operator set consulted by the predicate query mechanism to resolve application-defined relational operators used in a predicate.

ooOperatorSet \*ooUserDefinedOperators;

Discussion When an application defines its own relational operators, it registers them with an operator set. The operator set must be assigned to this variable to make the

registered relational operators available to the Objectivity/DB predicate query mechanism. By default, this variable points to an initial operator set that is created by Objectivity/DB when the application starts.

In a multithreaded application, the operator set to which this variable refers may be consulted for predicate queries in any Objectivity context. This variable is not thread-safe, however, so your application must ensure that:

- Only one thread updates the global operator set at a time.
- If a thread is updating the global operator set, no other thread can be making a predicate query at the same time.

See also

<u>ooQueryOperatorPtr</u> function-pointer type <u>ooOperatorSet</u> class

## ooVersMode

global type

Type for representing the versioning behavior of a basic object.

#### Constants oocNoVers

The versioning behavior of the object is disabled—that is, no new version can be created from it.

#### oocLinearVers

Linear versioning is enabled, which allows exactly one new version to be created from the object. When a new version is created, the versioning status of the original object is automatically set to oocNoVers so no further versions can be made from it. The new version's versioning status is oocLinearVers, allowing it to be the source of the next new version.

#### oocBranchVers

Branch versioning is enabled, which allows any number of new versions to be created from the object.

See also ooRefHandle(ooObj)::<u>setVersStatus</u>

### oovLastError

context variable

Pointer to the error identifier structure for the most recent error condition.

ooError \*oovLastError;

Discussion In a multithreaded application, this variable refers to the most recent error condition in the current Objectivity context.

oovLastErrorLevel context varia		context variable
	Severity level of the most recent error condition.	
	ooErrorLevel oovLastErrorLevel;	
Discussion	In a multithreaded application, this variable refers to the severity most recent error condition in the current Objectivity context.	level of the
oovNError		context variable
	Count of the total number of errors (not including warnings) that so far.	have occurred
	uint32 oovNError;	
Discussion	In a multithreaded application, this variable refers to the number have occurred in the current Objectivity context.	of errors that
ooVoidFuncPtr global t		global type
	Function-pointer type for a function that has no parameters and t result.	hat returns no
	<pre>typedef void (*ooVoidFuncPtr)(void);</pre>	
uint8		global type
	Objectivity/C++ 8-bit unsigned integer type. Values of this primi range from 0 to +255. This type is portable across all architectures Objectivity/C++. This type is also called <code>ooUInt8</code> .	tive type may supported by
uint16		global type
	Objectivity/C++ 16-bit unsigned integer type. Values of this primrange from 0 to +65,535. This type is portable across all architectuby Objectivity/C++. This type is also called <code>ooUInt16</code> .	itive type may res supported
uint32		global type
	Objectivity/C++ 32-bit unsigned integer type. Values of this prim range from 0 to +4,294,967,295. This type is portable across all arc supported by Objectivity/C++. This type is also called <code>ooUInt32</code>	ittive type may hitectures

## uint64

global type

Objectivity/C++ 64-bit unsigned integer type. Values of this primitive type may range from 0 to  $+2^{64}$ -1. This type is portable across all architectures supported by Objectivity/C++. This type is also called <code>ooUInt64</code>.

## Wait Options

Integer constants that specify whether to wait for locks when starting a transaction.

#### Constants oocNoWait

Turns off lock waiting for a transaction.

#### oocTransNoWait

Causes a transaction to use the default lock-waiting option currently in effect for the Objectivity context.

#### oocWait

Causes a transaction to wait indefinitely for locks.

See also <u>ooSetLockWait</u> ooTrans::<u>start</u>

## appClass Class

Inheritance: ooObj->appClass

Inheritance: ooObj->ooContObj->appClass

Handle Class: <u>ooHandle(appClass)</u>

**Object-Reference Class**: <u>ooRef(appClass)</u>

Every application-defined persistence-capable class *appClass* represents a particular kind of basic object or container. *appClass* and its corresponding handle and object-reference classes together define the behavior for *appClass* instances.

See:

- "Reference Summary" on page 84 for an overview of the member functions generated for *appClass*
- "Reference Index" on page 85 for a list of the generated member functions

For operations performed through a handle or object reference, see:

"Reference Summary" on page 478

## **About Application-Defined Classes**

An application that wants to save objects in a federated database must define a persistence-capable class *appClass* for each kind of object to be saved. An application can define its own:

- Basic-object classes, which are derived from ooObj (but not from ooContObj).
- Container classes, which are derived from ooContObj.

Every *appClass* normally defines data members representing application-specific attributes and associations, as well as member functions that provide application-specific behavior. In addition, an *appClass* can optionally override various member functions inherited from ooObj to customize behavior:

- Basic-object and container classes can override ooValidate if they need a way to perform tests that validate instances.
- Basic-object classes can override ooCopyInit, ooPreMoveInit, ooPostMoveInit, or ooVersInit if they need to customize copy, move, or versioning operations.

Every  $a_{PP}Class$  is defined using the Objectivity/C++ Data Definition Language (DDL), and is added to the federated-database schema by the DDL processor. Besides updating the schema, the DDL processor produces C++ header and implementation files containing:

- Classes for working with appClass instances—specifically, a handle class (ooHandle(appClass)), an object-reference class (ooRef(appClass)), and an iterator class (ooItr(appClass)). If appClass is a basic-object class, a short object-reference class (ooShortRef(appClass)) is generated as well.
- The definition of *appClass* with the following additions:
  - □ Member functions that redefine some of the functions inherited from ooObj or ooContObj (see "Redefinitions of Inherited Member Functions" on page 82).
  - Member functions for creating, deleting, and accessing each association defined by the class (see "Association Member Functions" on page 83).

An application must include the DDL-generated header files before it can use the generated classes and functions; the application must also be compiled with the corresponding implementation file. See the Objectivity/C++ Data Definition Language book for complete information about persistence-capable class definitions, the DDL processor, and the generated files.

### **Redefinitions of Inherited Member Functions**

The DDL processor generates definitions in *appClass* for various member functions that are defined in oo0bj:

- The function ooThis is redefined with *appClass*-specific parameter types.
- The virtual functions ooGetTypeN, ooGetTypeName, and oolsKindOf are overridden for internal purposes; their interfaces and behavior are unchanged, so they can be regarded as inherited from ooObj.
- The functions operator new and operator delete are overloaded for internal purposes; their interfaces and behavior are unchanged, so they can be regarded as inherited from ooObj.

If *appClass* is a container class, the container-specific variants of operator new are also overloaded for internal purposes and can be regarded as inherited from in ooContObj.

*appClass* also inherits member functions that are not redefined by generated definitions:

- If *appClass* is a basic-object or container class, it inherits the ooUpdate member function, and can optionally override the inherited virtual member function ooValidate.
- If *appClass* is a basic-object class, it inherits the member functions for managing advanced versioning associations, and can optionally override the inherited virtual member functions ooCopyInit, ooPreMoveInit, ooPostMoveInit, or ooVersInit.

### **Association Member Functions**

The DDL processor generates a set of member functions in *appClass* for each association defined in the class. These member functions enable you to create, delete, and access the corresponding association. The generated member-function names are constructed from the association's unique name. Various parameters of the generated functions accept object references, handles, or iterators to objects of the association's destination class.

A slightly different set of member functions is generated, depending on whether the association is *to-one* (that is, one-to-one or many-to-one) or *to-many* (that is, one-to-many or many-to-many).

The descriptions in this chapter use the following conventions:

linkName	Name of an association defined in <i>appClass</i> .
className	Name of the destination class specified in the <i>linkName</i> association definition—the class whose instances are to be associated with instances of any <i>Glass</i> .
	instances of appciass.

## Working With appClass Instances

An application works with instances of *appClass* as described in "Working With Basic Objects" on page 432 and "Working With Containers" on page 209. In brief, after the application includes the appropriate DDL-generated header files, it can:

 Create a new persistent instance using the application-defined (or default) appClass constructor and the inherited operator new.

- Reference an instance of *appClass* through a handle of class
   ooHandle(*appClass*) or an object reference of class ooRef(*appClass*).
- Perform Objectivity/DB operations on a referenced instance of *appClass* by calling various member functions on the referencing handle or object reference.
- Perform various operations on a referenced instance of *appClass* by accessing the instance's defined, generated, or inherited members. These members can be accessed directly from within a member function of *appClass* or indirectly through the indirect member-access operator (->) on an *appClass* handle or object reference.

## **Reference Summary**

The following table lists the member functions generated by the DDL processor in the definition of each application-defined persistence-capable class appClass. Member functions indicated as *(inherited)* are redefined in appClass with no change in behavior. They are documented with the base class, along with the other inherited member functions not listed here; see the <u>ooObj</u> class (page 431) or the <u>ooContObj</u> class (page 207).

Creating and Deleting an <i>appClass</i> Object	<u>operator new</u> (inherited) <u>operator new</u> (inherited, containers only) <u>operator delete</u> (inherited)
Managing To-One Associations	<u>linkName</u> <u>exist linkName</u> <u>set linkName</u> <u>del linkName</u>
Managing To-Many Associations	<u>linkName</u> exist <u>linkName</u> add <u>linkName</u> sub <u>linkName</u> del_linkName
Working With This <i>appClass</i> Object	<u>ooGetTypeN</u> (inherited) <u>ooGetTypeName</u> (inherited) <u>ooIsKindOf</u> (inherited) <u>ooThis</u>

## **Reference Index**

add_linkName	Creates a linkName association from this object to the specified destination object.
<u>del_linkName</u>	Deletes all linkName associations that exist for this object.
<u>exist_linkName</u>	Tests whether this object has a linkName association to the specified destination object.
<u>linkName</u>	Finds, and optionally opens, the destination object linked to this object by the to-one association linkName.
<u>linkName</u>	Initializes an iterator to find, and optionally open, all destination objects that are linked to this object by the to-many association linkName, and that satisfy any specified selection criteria.
<u>ooThis</u>	Sets an object reference or handle to reference this ${\tt appClass}$ object.
<u>set_linkName</u>	Creates a linkName association from this object to the specified destination object.
<u>sub_linkName</u>	Deletes one or more linkName associations from this object to the specified destination object.

## **Generated Member Functions**

### add\_linkName

Creates a *linkName* association from this object to the specified destination object.

ooStatus add\_linkName(const ooHandle(className) &objH);

ParametersobjH<br/>Handle to the destination object to be linked to this object. The destination<br/>object must be an instance of the destination class className or any of its<br/>derived classes. If objH is a null handle, no action is performed.ReturnsoocSuccess if successful; otherwise oocError.DiscussionThis member function is generated for the to-many association linkName<br/>defined by this object's class.<br/>The application must be able to obtain an update lock for this object.

If *linkName* is a bidirectional association, this operation also creates the inverse association from the specified destination object to this object. In this case, the application must be able to obtain update locks on both objects.

No error is signaled if a *linkName* association already exists between this object and the specified destination object. That is, you can create duplicate associations between the two objects (even though it could be semantically meaningless to do so).

### del\_linkName

Deletes all *linkName* associations that exist for this object.

ooStatus del\_linkName();

Returns oocSuccess if successful; otherwise oocError.

Discussion This member function is generated for the to-one or to-many association *linkName* defined by this object's class.

The application must be able to obtain an update lock for this object.

If *linkName* is a bidirectional association, this operation also deletes the inverse association to this object from each of the former destination objects. In this case, the application must be able to obtain update locks on all of the affected objects.

You should call the exist\_linkName member function to test whether linkName associations exist before you try to delete them.

### exist\_linkName

Tests whether this object has a *linkName* association to the specified destination object.

ooBoolean	exist_ <i>linkName</i> (		
const	ooHandle( <i>className</i> )	&objH)	const;

Parameters objH

Handle to the destination object to be tested for association. The destination object must be an instance of the destination class *className* or any of its derived classes. You can specify 0 to test whether this object has any *linkName* association(s) at all.

```
Returns oocTrue if an association exists, otherwise oocFalse.
```

Discussion This member function is generated for the to-one or to-many association *linkName* defined by this object's class.

## linkName

Finds, and optionally opens, the destination object linked to this object by the to-one association <code>linkName</code>.

	<pre>1. ooHandle(className) linkName(</pre>
	<pre>2. ooRef(className) &amp;linkName(</pre>
	<pre>3. ooHandle(className) &amp;linkName(</pre>
Parameters	openMode
	Intended level of access to the destination object:
	<ul> <li>Specify oocNoOpen (the default) to set the returned object reference or handle to the destination object without opening it.</li> </ul>
	<ul> <li>Specify oocRead to open the destination object for read.</li> </ul>
	<ul> <li>Specify occupdate to open the destination object for update.</li> </ul>
	object
	Object reference or handle to be set to the destination object, which is an instance of the destination class <i>className</i> or one of its derived classes.
Returns	Object reference or handle to the destination object. A null handle is returned if no <i>linkName</i> association exists.
Discussion	This member function is generated for the to-one association <code>linkName</code> defined by this object's class.
	When called without an <i>object</i> parameter, this member function allocates a new <i>className</i> handle and returns it. Otherwise, this member function returns the object reference or handle that is passed to it.

## linkName

Initializes an iterator to find, and optionally open, all destination objects that are linked to this object by the to-many association *linkName*, and that satisfy any specified selection criteria.

#### Parameters

iterator

Iterator for finding the destination objects, which are instances of the destination class *className* or any of its derived classes.

#### openMode

Intended level of access to each destination object found by the iterator's next member function:

- oocNoOpen (the default in variant 1) causes next to set the iterator to the next destination object without opening it.
- oocRead causes next to open the next destination object for read.
- oocUpdate causes next to open the next destination object for update.

#### predicate

String expression in predicate query language. Specifies the condition to be met by the found destination objects. The iterator is initialized to find only those destination objects that match *predicate*.

#### access

Level of access control of the data members that *predicate* can test:

- Specify oocPublic to permit the predicate to test only public data members, preserving encapsulation.
- Specify oocAll to permit the predicate to test any data member. To
  preserve encapsulation, you should use this mode only within member
  functions of the class you are querying.

or.
c

Discussion This member function is generated for the to-many association *linkName* defined by this object's class.

Variant 2 finds destination objects without opening them; the other variants may open the found objects, as specified by the <code>openMode</code> parameter.

If no *linkName* associations currently exist from this object, the iterator is set to null, and the *linkName* member function returns oocSuccess. (The iterator's next member function will return oocFalse, however.)

An error is signaled if *predicate* tries to test a non-existent data member, or if *predicate* tries to test a protected or private data member when the *access* parameter is oocPublic.

The *linkName* member function does *not* use indexes to optimize a search, even if the *predicate* parameter is specified.

### ooThis

Sets an object reference or handle to reference this *appClass* object.

- 1. ooHandle(*appClass*) ooThis() const;

#### Parameters object

Object reference or handle to be set to this object.

Returns Object reference or handle to this object.

Discussion You normally use ooThis in a member function of *appClass*; when such a member function is called on an *appClass* object, ooThis provides the member function with an *appClass* handle or object reference to the object. The member function can then perform operations on the object that are available only through a handle or object reference.

When called without an *object* parameter, <code>ooThis</code> allocates a new handle and returns it. Otherwise, <code>ooThis</code> returns the object reference or handle that is passed to it.

Calling <code>ooThis</code> in a member function of a persistence-capable class serves the same purpose as using the C++ this keyword in a member function of a non-persistence-capable class—both <code>ooThis</code> and this enable you to access the object on which the member function is called. However, <code>ooThis</code> and this differ in the following ways:

- The this keyword is the pointer to the object. Syntactically, the this keyword is a name, and can be used in expressions such as this->get();
- The ooThis member function returns an object reference or handle to the object. Syntactically, ooThis is a function and can be used in expressions such as ooThis()->get();

Member functions of persistence-capable classes should use ooThis (and not this) because the only safe way to operate on a persistent object is through an object reference or handle.

**WARNING** ooThis is the only safe way to obtain an object reference or handle within a member function; do *not* attempt to initialize an object reference or handle by assigning the this pointer to it.

#### set\_linkName

Creates a *linkName* association from this object to the specified destination object.

Parameters	objH
	Handle to the destination object to be associated. The destination object must be an instance of the destination class <i>className</i> or any of its derived classes. If <i>objH</i> is a null handle, any existing <i>linkName</i> association from this object is deleted.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	This member function is generated for the to-one association <code>linkName</code> defined by this object's class.
	The application must be able to obtain an update lock for this object.
	If <i>linkName</i> is a bidirectional association, this operation also creates the inverse association from the specified destination object to this object. In this case, the application must be able to obtain update locks on both objects.
	Because this member function creates to-one associations, an error is signaled if this object already has a <i>linkName</i> association to some destination object.

### sub\_linkName

Deletes one or more *linkName* associations from this object to the specified destination object.

	<pre>2. ooStatus sub_linkName( const ooHandle(className) &amp;objH, const unit32 number = 1);</pre>
Parameters	objH Handle to the destination object. The destination object must be an instance of the destination class <i>className</i> or any of its derived classes. If <i>objH</i> is a null handle, no action is performed.
	number
	Number of <i>linkName</i> associations to delete between this object and the destination object:
	■ If you specify 0, all such associations are deleted.
	<ul> <li>If you specify 1 (the default), the first or only such association is deleted.</li> <li>If you specify a number greater than 1, this member function deletes the first <i>number</i> associations encountered.</li> </ul>
	Do not use this parameter if <i>linkName</i> is a one-to-many bidirectional association.
Returns	oocSuccess if successful; otherwise oocError. oocSuccess is returned even if no association exists to be deleted or if <i>number</i> exceeds the number of existing associations. oocError results from internal errors or locking errors.
Discussion	Only a single variant of this member function is generated for any given <i>linkName</i> association:
	Variant 1 is generated if <i>linkName</i> is a one-to-many bidirectional association. Only one <i>linkName</i> association can exist from the source object to a given destination object, so there is no need for the second parameter ( <i>number</i> ).
	Variant 2 is generated if <i>linkName</i> is a one-to-many unidirectional association and for a many-to-many bidirectional association. It is possible for multiple <i>linkName</i> associations to exist from the source object to a given destination object. In that situation, you can specify the <i>number</i> parameter to delete some or all of those associations.
	The application must be able to obtain an update lock for this object. If <code>linkName</code> is a bidirectional association, this operation also deletes the inverse association(s) from the specified destination object to this object. In this case, the application must be able to obtain update locks on both objects.
See also	del_linkName

## d\_Database Class

#### Inheritance: d\_Database

The non-persistence-capable ODMG class d\_Database represents an *ODMG database*, which is the ODMG equivalent of an Objectivity/DB federated database.

See:

- "Reference Summary" on page 94 for an overview of member functions
- "Reference Index" on page 94 for a list of member functions

To add this and other ODMG types and definitions to your application, you must run the DDL processor and your C++ compiler with the  $-DOO_ODMG$  flag.

## About ODMG Databases

In the storage hierarchy, an ODMG database is equivalent to an *Objectivity/DB federated database*. Therefore, an ODMG-compliant application uses an instance of d\_Database to refer to and manipulate the top-level storage object, whereas a standard Objectivity/C++ application uses a federated-database handle (an instance of *ooRefHandle*(ooFDObj)).

Although analogous in purpose, d\_Database and *ooRefHandle*(ooFDObj) define different behavior. As specified by ODMG, d\_Database member functions enable you to open, close, and find objects named in the scope of an ODMG database (Objectivity/DB federated database). See the documentation for <u>ooRefHandle(ooFDObj)</u> for a list of operations on federated-database handles.

# **NOTE** Do not confuse an *ODMG database* with an *Objectivity/DB database*; they refer to different storage levels.

The ODMG standard has no storage level equivalent to an Objectivity/DB database. Consequently, a strictly ODMG-compliant application should avoid manipulating Objectivity/DB databases explicitly. To do this:

- 1. Before you start your application, create a default Objectivity/DB database:
  - **a.** Use the oonewdb tool create an Objectivity/DB database.
  - **b.** Set the system name of the new Objectivity/DB database as the value of the environment variable OO\_DB\_NAME.
- 2. In your application, create persistent objects by calling <u>operator new</u> with a d\_Database object as the clustering directive. This directive causes new objects to be placed in the Objectivity/DB database specified by the OO\_DB\_NAME environment variable.
- **NOTE** If the OO\_DB\_NAME environment variable is not set, an Objectivity/DB database named default\_odmg\_db is automatically created and used.

## **Reference Summary**

Accessing an ODMG Database	<u>access_status</u> <u>close</u> <u>open</u>
Working With Scope Names	<u>get_object_name</u> <u>lookup_object</u> <u>rename_object</u> <u>set_object_name</u>
Creating Transient Objects	transient_memory

## **Reference Index**

<u>access status</u>	(ODMG) Type specifying the intended level of access to an ODMG database.
<u>close</u>	(ODMG) Closes an ODMG database.
get_object_name	(ODMG) Gets the name defined in the scope of this ODMG database for the specified object.

<u>lookup object</u>	(ODMG) Finds the persistent object with the specified name in the scope of this ODMG database.
open	(ODMG) Opens an ODMG database.
<u>rename_object</u>	(ODMG) Changes the name of a persistent object within the scope of this ODMG database.
<u>set_object_name</u>	(ODMG) Assigns a scope name to a persistent object within the scope of this ODMG database.
transient memory	(ODMG) Clustering directive for creating a transient object from a persistence-capable class.

## **Types and Constants**

#### access\_status

(ODMG) Type specifying the intended level of access to an ODMG database.

Constants not\_open

Equal to the constant oocNoOpen of the Objectivity/C++ global type ooMode.

read\_write

Equal to the constant oocUpdate of the Objectivity/C++ global type ooMode.

read\_only

Equal to the constant oocRead of the Objectivity/C++ global type ooMode.

exclusive

Reserved for future development.

Discussion This type corresponds to the global type  $\underline{ooMode}$ .

#### transient\_memory

(*ODMG*) Clustering directive for creating a transient object from a persistence-capable class.

static const d\_Database \*const transient\_memory;

Discussion This static data member is set to 0 by Objectivity/DB. You can specify it as a clustering directive to <u>operator new</u> to create a transient object from a persistence-capable class.

## **Member Functions**

## close

(ODMG) Closes an ODMG database.

void close();

Discussion This member function calls the <u>close</u> member function on a federated-database handle.

## get\_object\_name

	( <i>ODMG</i> ) Gets the name defined in the scope of this ODMG database for the specified object.
	<pre>const char *get_object_name(const d_Ref_Any &amp;object) const;</pre>
Parameters	object An ODMG generic reference to the object whose scope name you want to get.
Returns	Pointer to a string containing the scope name. If the referenced object does not have a scope name in the specified scope, the returned pointer is null.
Discussion	The string is statically allocated by the member function and overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.
	The application must be able to get a read lock on the hashed container used by this ODMG database for storing scope names.

### lookup\_object

(*ODMG*) Finds the persistent object with the specified name in the scope of this ODMG database.

d\_Ref\_Any lookup\_object(const char \*name) const;

Parameters name Scope name to look up in this ODMG database. Returns An ODMG generic reference to the found object; a null generic reference is returned if no object with the specified scope name is found in the ODMG database. Discussion The application must be able to obtain a read lock on the hashed container used by the ODMG database for storing scope names.

#### open

(ODMG) Opens an ODMG database.

```
void open(
    const char *databaseName,
    access_status status = read_write);
```

Parameters databaseName

Path to the boot file of the ODMG database to be opened. You can omit this parameter if you set the OO\_FD\_BOOT environment variable to the path. You can specify this path with or without a host name. If you specify it as a host path, use the format host::path.

#### status

Intended level of access to the ODMG database:

- Specify read\_write (the default) to open the ODMG database for update and to designate the transaction as an update transaction.
- Specify read\_only to open the ODMG database for read and to designate the transaction as a read transaction.
- Discussion This member function calls the <u>open</u> member function on a federated-database handle (*without* performing local automatic recovery).

#### rename\_object

(*ODMG*) Changes the name of a persistent object within the scope of this ODMG database.

```
void rename_object(
    const char *oldName,
    const char *newName);
```

Parameters oldName

Scope name of the object to be renamed.

newName

New scope name to assign. This name:

- Must be a null-terminated string that can contain any nonnull character.
- Must be unique within the name scope defined by this ODMG database.
- May contain up to 500 characters.

### set\_object\_name

(*ODMG*) Assigns a scope name to a persistent object within the scope of this ODMG database.

```
void set_object_name(
    const d_Ref_Any &object,
    const char *name);
```

Parameters object

An ODMG generic reference to the object to be named.

name

Scope name to assign. This name:

- Must be a null-terminated string that can contain any nonnull character.
- Must be unique within the name scope defined by this ODMG database.
- May contain up to 500 characters.

## d\_Date Class

Inheritance: d\_Date

The non-persistence-capable ODMG class d\_Date represents a *date*, which consists of three components—a year, a month, and a day.

See:

- "Reference Summary" on page 99 for an overview of member functions
- "Reference Index" on page 100 for a list of member functions

To use this class, your application must include the ooTime.h header file. No extra linking is required.

## **Reference Summary**

Creating a Date	<u>d_Date</u>
Date Components	<u>Month</u> Weekday
Setting a Date	<u>next</u> operator= previous

Getting Information	current day day_of_week day_of_year days_in_month days_in_year is_between is_leap_year is_valid_date month year
Comparing Dates	<pre>::operator== ::operator!= ::operator&lt; ::operator&lt;= ::operator&gt;= ::operator&gt;=</pre>
Arithmetic Operations on Dates	::operator+ operator++ operator+= ::operator- operator-= operator-=
Determining Whether Time Periods Overlap	::overlaps

# **Reference Index**

<u>current</u>	(ODMG) Returns the current date.
day	(ODMG) Returns this date's day of the month.
<u>day_of_week</u>	(ODMG) Returns this date's day of the week.
<u>day of year</u>	(ODMG) Returns this date's day of the year.
<u>days in month</u>	(ODMG) Returns the number of days in this date's month or in the specified month.
<u>days in year</u>	(ODMG) Returns the number of days in this date's year or in the specified year.
<u>d Date</u>	(ODMG) Default constructor that constructs a new date set to the current date.

<u>d Date</u>	(ODMG) Constructs a new date from the specified date or timestamp.
<u>d_Date</u>	(ODMG) Constructs a new date from the specified information.
<u>is between</u>	(ODMG) Returns 1 if this date is within the specified period.
<u>is leap year</u>	(ODMG) Returns 1 if this date's year or the specified year is a leap year.
<u>is_valid_date</u>	(ODMG) Returns 1 if the specified components make a valid date.
Month	(ODMG) Enumerated type for the months of the year.
month	(ODMG) Returns this date's month.
next	(ODMG) Advances this date to the next occurrence of the specified day of type Weekday.
::operator+	(ODMG) Addition operator; allocates a new date that is set to the sum of the specified date and interval of days.
operator++	(ODMG) Increment operator; increments this date by a day.
<u>operator+=</u>	(ODMG) Increment operator; adds the specified interval or number of days to this date.
::operator-	(ODMG) Subtraction operator; allocates a new date that is set to the specified date minus the specified interval of days.
operator	(ODMG) Decrement operator; decrements this date by a day.
<u>operator-=</u>	(ODMG) Decrement operator; subtracts the specified interval or number of days from this date.
<u>operator=</u>	(ODMG) Assignment operator; sets this date to the specified date or timestamp.
::operator==	(ODMG) Equality operator; returns 1 if every component of one date matches the corresponding component of the other.
::operator!=	(ODMG) Inequality operator; returns 1 if any component of one date differs from the corresponding component of the other.
::operator<	(ODMG) Less-than operator; returns 1 if one date is less than another.
::operator<=	(ODMG) Less-than-or-equal-to operator; returns 1 if one date is less than or equal to another.
::operator>	(ODMG) Greater-than operator; returns 1 if one date is greater than another.

::operator>=	(ODMG) Greater-than-or-equal-to operator; returns 1 if one date is greater than or equal to another.
::overlaps	(ODMG) Returns 1 if two time periods overlap, where one or both periods are specified using dates.
previous	(ODMG) Moves this date to the previous occurrence of the specified day of type Weekday.
Weekday	(ODMG) Enumerated type for the days of the week.
year	(ODMG) Returns this date's year.

## Types

### Month

(ODMG) Enumerated type for the months of the year.

```
enum Month {
    January = 1,
    February = 2,
    March = 3,
    April = 4,
    May = 5,
    June = 6,
    July = 7,
    August = 8,
    September = 9,
    October = 10,
    November = 11,
    December = 12};
```

### Weekday

(ODMG) Enumerated type for the days of the week.

```
enum Weekday {
   Sunday = 0,
   Monday = 1,
   Tuesday = 2,
   Wednesday = 3,
   Thursday = 4,
   Friday = 5,
   Saturday = 6};
```

## Constructors

## d\_Date

(*ODMG*) Default constructor that constructs a new date set to the current date.

d\_Date();

## d\_Date

(ODMG) Constructs a new date from the specified date or timestamp.

- 1. d\_Date(const d\_Date &date);
- 2. d\_Date(const d\_Timestamp &timeStamp);

## d\_Date

(ODMG) Constructs a new date from the specified information.

1.	d_Date(		
	unsigned	short	year,
	unsigned	short	dayOfYear);
2.	d Date(		

```
unsigned short year,
unsigned short month = 1,
unsigned short day = 1);
```

#### Parameters

Year component of the new date.

#### *dayOfYear*

Sequence number of a day in a year. The maximum value is 365 (366 in a leap year). This value is used to calculate the month and day of the new date.

#### month

year

Month component of the new date. The maximum value is 12.

#### day

Day component of the new date. The maximum value is the number of days in the specified month for the specified year.

## **Operators**

#### ::operator+

global function

(*ODMG*) Addition operator; allocates a new date that is set to the sum of the specified date and interval of days.

1.	d_Date ::operator+(
	const d_Date & <i>left</i> ,
	<pre>const d_Interval &amp;right);</pre>
2.	<pre>d_Date ::operator+(     const d_Interval &amp;left,     const d_Date &amp;right);</pre>

**WARNING** Addition that causes the year component to exceed 65535 results in an invalid date.

#### operator++

(ODMG) Increment operator; increments this date by a day.

	1.	d_Date &operator++();			
	2.	<pre>d_Date operator++(int n);</pre>			
Parameters	n T fi	his parameter is not used in calling this operator; its presence in the unction declaration specifies a postfix operator.			
Returns	(Variant 1) This date, incremented by a day. (Variant 2) A new date whose value is this date, before this date is incremented.				
Discussion	Variant 1 is the prefix increment operator, which increments this date and then returns it. Variant 2 is the postfix increment operator, which allocates and returns a new date set to this date, and then increments this date.				
Example	d_Dat	te d;			
	++d;	// Prefix			
	d++;	// Postfix			

#### operator+=

(*ODMG*) Increment operator; adds the specified interval or number of days to this date.

```
    d_Date & operator += (const d_Interval &);
```

```
2. d_Date &operator+=(int ndays);
```

*WARNING* Addition that causes the year component to exceed 65535 results in an invalid date.

#### ::operator-

global function

(*ODMG*) Subtraction operator; allocates a new date that is set to the specified date minus the specified interval of days.

```
d_Date ::operator-(
const d_Date &left,
const d_Interval &right);
```

#### operator--

(ODMG) Decrement operator; decrements this date by a day.

1. d_Date	&operator	(	)	į
-----------	-----------	---	---	---

2. d\_Date operator--(int);

# Returns (Variant 1) This date, decremented by a day. (Variant 2) A new date whose value is this date, before this date is decremented.

Discussion Variant 1 is the prefix decrement operator, which decrements this date and then returns it. Variant 2 is the postfix decrement operator, which allocates and returns a new date set to this date, and then decrements this date.

#### operator-=

(*ODMG*) Decrement operator; subtracts the specified interval or number of days from this date.

- d\_Date & operator -= (const d\_Interval &);
- 2. d\_Date &operator-=(int ndays);

#### operator=

(ODMG) Assignment operator; sets this date to the specified date or timestamp.

- 1. d\_Date &operator=(const d\_Date &date);
- 2. d\_Date &operator =(const d\_Timestamp &timeStamp);

#### ::operator==

(*ODMG*) Equality operator; returns 1 if every component of one date matches the corresponding component of the other.

```
int ::operator==(
    const d_Date &left,
    const d_Date &right);
```

### ::operator!=

(*ODMG*) Inequality operator; returns 1 if any component of one date differs from the corresponding component of the other.

```
int ::operator!= (
    const d_Date &left,
    const d_Date &right);
```

### ::operator<

global function

(ODMG) Less-than operator; returns 1 if one date is less than another.

```
int ::operator< (
    const d_Date &left,
    const d_Date &right);</pre>
```

#### ::operator<=

global function

(*ODMG*) Less-than-or-equal-to operator; returns 1 if one date is less than or equal to another.

```
int ::operator<=(
    const d_Date &left,
    const d_Date &right);</pre>
```

global function

global function

### ::operator>

global function

(ODMG) Greater-than operator; returns 1 if one date is greater than another.

```
int ::operator> (
    const d_Date &left,
    const d_Date &right);
```

## ::operator>=

global function

(*ODMG*) Greater-than-or-equal-to operator; returns 1 if one date is greater than or equal to another.

```
int ::operator>=(
    const d_Date &left,
    const d_Date &right);
```

## **Member Functions**

### current

(ODMG) Returns the current date.

static d\_Date current();

### day

(ODMG) Returns this date's day of the month.

unsigned short day() const;

## day\_of\_week

(ODMG) Returns this date's day of the week.

Weekday day\_of\_week() const;

## day\_of\_year

(ODMG) Returns this date's day of the year.

unsigned short day\_of\_year() const;

### days\_in\_month

(*ODMG*) Returns the number of days in this date's month or in the specified month.

- unsigned int days\_in\_month() const;

#### days\_in\_year

(ODMG) Returns the number of days in this date's year or in the specified year.

- unsigned int days\_in\_year() const;
- 2. static unsigned int days\_in\_year(unsigned short year);

#### is\_between

(ODMG) Returns 1 if this date is within the specified period.

```
int is_between(
    const d_Date &date,
    const d_Date &date) const;
```

#### is\_leap\_year

(ODMG) Returns 1 if this date's year or the specified year is a leap year.

- 1. int is\_leap\_year() const;
- 2. static int is\_leap\_year(unsigned short year);

#### is\_valid\_date

(ODMG) Returns 1 if the specified components make a valid date.

```
static int is_valid_date(
    unsigned short year,
    unsigned short month,
    unsigned short day);
```

#### month

(ODMG) Returns this date's month.

```
Month month() const;
```
#### next

(ODMG) Advances this date to the next occurrence of the specified day of type Weekday.

d\_Date &next(Weekday day);

## previous

(ODMG) Moves this date to the previous occurrence of the specified day of type Weekday.

d\_Date &previous(Weekday day);

## year

(ODMG) Returns this date's year.

unsigned short year() const;

# **Related Global Functions**

## ::overlaps

(*ODMG*) Returns 1 if two time periods overlap, where one or both periods are specified using dates.

```
1.
    int ::overlaps(
      const d Date &startLeft,
      const d_Date & endLeft,
      const d_Date &startRight,
      const d_Date &endRight);
2.
    int ::overlaps(
      const d_Timestamp &startLeft,
      const d_Timestamp & endLeft,
      const d_Date &startRight,
      const d_Date &endRight);
3.
    int ::overlaps(
      const d_Date &startLeft,
      const d_Date & endLeft,
      const d_Timestamp &startRight,
      const d_Timestamp & endRight);
```

Discussion Each time period is specified by a start and end time. You can specify the time periods using dates (variant 1), or using dates and timestamps (variants 2 and 3).

# d\_Interval Class

#### Inheritance: d\_Interval

The non-persistence-capable ODMG time class d\_Interval represents an *interval*, or duration of time, conforming to the day-time interval defined in the SQL standard.

See:

- "Reference Summary" on page 112 for an overview of member functions
- "Reference Index" on page 113 for a list of member functions

To use this class, your application must include the ooTime.h header file. No extra linking is required.

# About d\_Interval

A d\_Interval object consists of four time components—days, hours, minutes, and seconds. The size of these components depends on the total duration to be represented:

- For intervals of less than a day, the d\_Interval class accepts nonnormalized input, and then normalizes the time components when they are accessed. For example, you can construct an interval of 68 minutes; calling the hour function returns a value of 1, and calling the minute function returns a value of 8.
- For intervals of greater than a day, the d\_Interval class requires that you
  partially normalize the input time components to specify an integer number
  of days, plus some (nonnormalized) number of hours, minutes, and seconds.

For example, instead of specifying 28 hours, you could specify 1 day, 2 hours, and 120 minutes, which would be normalized on access to 1 day and 4 hours.

The Objectivity/C++ implementation of the d\_Interval class stores each time component as an integer number of milliseconds. Within a d\_Interval object:

- The total number of milliseconds for all time components must not exceed 2,147,483,648.
- The total number of milliseconds for the hour, day, and seconds time components must not exceed 86,400,000 (the number of milliseconds in a day).

For seconds, which are input as floating-point numbers, conversion to milliseconds may result in a loss of precision. For example, if you specify 25.532962 seconds, only 25532 milliseconds are stored, and the exact value returned by the seconds function (such as 25.532) is architecture-dependent.

# **Reference Summary**

Creating an Interval	<u>d_Interval</u>
Setting an Interval	operator=
Getting Information	day hour is_zero minute second
Comparing Intervals	<pre>::operator== ::operator!= ::operator&lt; ::operator&lt;= ::operator&gt;=</pre>
Arithmetic Operations on Intervals	<pre>::operator+ operator+= operator- ::operator- operator-= ::operator* operator*= ::operator/ operator/=</pre>

# **Reference Index**

day	(ODMG) Returns the normalized day component of this interval.
<u>d_Interval</u>	(ODMG) Constructs a new interval of time from the specified components.
<u>d_Interval</u>	(ODMG) Copy constructor that constructs a new interval from the specified interval.
hour	(ODMG) Returns the normalized hour component of this interval.
<u>is zero</u>	(ODMG) Returns 1 if the duration of this interval is zero.
<u>minute</u>	(ODMG) Returns the normalized minute component of this interval.
second	(ODMG) Returns the normalized seconds component of this interval.
::operator+	(ODMG) Addition operator; allocates a new interval that is set to the sum of the specified intervals.
<u>operator+=</u>	(ODMG) Increment operator; adds the specified interval to this interval.
<u>operator-</u>	(ODMG) Unary minus operator; returns the negative of this interval.
::operator-	(ODMG) Subtraction operator; allocates a new interval that is set to the difference between the specified intervals.
<u>operator-=</u>	(ODMG) Decrement operator; subtracts the specified interval from this interval.
::operator*	(ODMG) Multiplication operator; multiplies the specified interval and an integer value, and then returns the result in a new interval.
<u>operator*=</u>	(ODMG) Multiplication operator; multiplies this interval by the specified integer value.
::operator/	(ODMG) Division operator; divides the specified interval by an integer value.
<u>operator/=</u>	(ODMG) Division operator; divides this interval by the specified integer value.
<u>operator=</u>	(ODMG) Assignment operator; sets this interval to the specified interval.

::operator==	(ODMG) Equality operator; returns 1 if the specified intervals are of the same duration.
::operator!=	(ODMG) Inequality operator; returns 1 if the specified intervals are of different durations.
::operator<	(ODMG) Less-than operator; returns 1 if one interval is less than another.
::operator<=	(ODMG) Less-than-or-equal-to operator; returns 1 if one interval is less than or equal to another.
::operator>	(ODMG) Greater-than operator; returns 1 if one interval is greater than another.
::operator>=	(ODMG) Greater-than-or-equal-to operator; returns 1 if one interval is greater than or equal to another.

# Constructors

## d\_Interval

(ODMG) Constructs a new interval of time from the specified components.

```
d_Interval(
    int day=0,
    int hour=0,
    int minute=0,
    float second=0.0);
```

#### Parameters day

Number of whole days in the interval.

hour

Number of hours in the interval. This number must be less than 24; if the interval is longer than a day, you must represent the excess hours as a number of whole days.

#### minute

Number of minutes in the interval.

second

Number of seconds in the interval.

Discussion When combined, *hour*, *minute*, and *second* may not exceed 86,400,000 (the number of milliseconds in a day). For more information about permitted component values, see "About d\_Interval" on page 111.

# d\_Interval

(*ODMG*) Copy constructor that constructs a new interval from the specified interval.

```
d_Interval(const d_Interval &interval);
```

# **Operators**

## ::operator+

global function

(*ODMG*) Addition operator; allocates a new interval that is set to the sum of the specified intervals.

```
d_Interval ::operator+(
    const d_Interval &left,
    const d_Interval &right);
```

#### operator+=

(ODMG) Increment operator; adds the specified interval to this interval.

d\_Interval &operator+=(const d\_Interval &interval);

#### operator-

(ODMG) Unary minus operator; returns the negative of this interval.

d\_Interval operator-() const;

#### ::operator-

global function

(*ODMG*) Subtraction operator; allocates a new interval that is set to the difference between the specified intervals.

```
d_Interval ::operator-(
    const d_Interval &left,
    const d_Interval &right);
```

#### operator-=

(ODMG) Decrement operator; subtracts the specified interval from this interval.

```
d_Interval &operator-=(const d_Interval &interval);
```

## ::operator\*

global function

(*ODMG*) Multiplication operator; multiplies the specified interval and an integer value, and then returns the result in a new interval.

# operator\*=

(*ODMG*) Multiplication operator; multiplies this interval by the specified integer value.

```
d_Interval &operator*=(int);
```

# ::operator/

global function

(ODMG) Division operator; divides the specified interval by an integer value.

# operator/=

(*ODMG*) Division operator; divides this interval by the specified integer value.

```
d_Interval &operator/=(int);
```

# operator=

(ODMG) Assignment operator; sets this interval to the specified interval.

d\_Interval &operator=(const d\_Interval &interval);

#### ::operator==

Operators

global function

(*ODMG*) Equality operator; returns 1 if the specified intervals are of the same duration.

```
int ::operator==(
    const d_Interval &left,
    const d_Interval &right);
```

## ::operator!=

global function

(*ODMG*) Inequality operator; returns 1 if the specified intervals are of different durations.

```
int ::operator!= (
    const d_Interval &left,
    const d_Interval &right);
```

## ::operator<

global function

(ODMG) Less-than operator; returns 1 if one interval is less than another.

```
int ::operator< (
    const d_Interval &left,
    const d_Interval &right);</pre>
```

#### ::operator<=

global function

(*ODMG*) Less-than-or-equal-to operator; returns 1 if one interval is less than or equal to another.

int ::operator<=(
 const d\_Interval &left,
 const d\_Interval &right);</pre>

## ::operator>

global function

(ODMG) Greater-than operator; returns 1 if one interval is greater than another.

```
int ::operator> (
    const d_Interval &left,
    const d_Interval &right);
```

## ::operator>=

global function

(*ODMG*) Greater-than-or-equal-to operator; returns 1 if one interval is greater than or equal to another.

```
int ::operator>=(
    const d_Interval &left,
    const d_Interval &right);
```

# **Member Functions**

# day (ODMG) Returns the normalized day component of this interval. int day() const; hour (ODMG) Returns the normalized hour component of this interval. int hour() const; is\_zero (ODMG) Returns 1 if the duration of this interval is zero. int is\_zero() const; minute (ODMG) Returns the normalized minute component of this interval. int minute() const; second (ODMG) Returns the normalized seconds component of this interval. float second() const;

# d\_lterator<element\_type> Class

Inheritance: d\_Iterator<element\_type>

The non-persistence-capable ODMG class d\_Iterator<element\_type> represents a VArray iterator, which supports iteration over the elements of a VArray of the same element type.

See:

- "Reference Summary" on page 120 for an overview of member functions
- "Reference Index" on page 120 for a list of member functions

See also:

- <u>d\_Varray<element\_type></u>
- ooVArrayT<element\_type>

# About VArray Iterators

You create an iterator to find elements of a particular VArray by calling the <u>create\_iterator</u> member function on that VArray. The new iterator is initialized to point to the first element of that VArray. The initialized iterator pins the VArray in memory.

An initialized VArray iterator supports two alternative iteration styles for finding the VArray's elements:

- The <u>next</u> member function tests for a next element, sets a parameter to the current element, and then advances the iterator, all as a single operation.
   You typically use next to control a while loop that executes the same statements once for each element in the VArray.
- The <u>not\_done</u>, <u>get\_element</u>, and <u>advance</u> member functions perform the iteration actions as separate operations.

You typically use not\_done and advance as the expressions that control a for loop; within the loop, you can use get\_element to get the current element.

# **Reference Summary**

Creating a VArray Iterator	<u>d_Iterator</u>
Assigning a VArray Iterator	operator=
Positioning a VArray Iterator	advance next operator++ operator reset
Finding a VArray Element	get_element next
Testing for the Next Element	next not_done

# **Reference Index**

<u>advance</u>	(ODMG) Advances this VArray iterator to the next element in the VArray.
<u>d_Iterator</u>	(ODMG) Default constructor that constructs a new uninitialized VArray iterator.
<u>d_Iterator</u>	(ODMG) Copy constructor that constructs a new VArray iterator initialized with the specified VArray iterator.
<u>get_element</u>	(ODMG) Returns the element at this VArray iterator's current position in the VArray.
<u>next</u>	(ODMG) Tests whether to continue iterating and, if so, assigns the current element to the output parameter and advances this VArray iterator to the next element.
<u>not_done</u>	(ODMG) Tests whether to continue iterating.

<u>reset</u>	(ODMG) Reinitializes this VArray iterator to the start of iteration for the same VArray.
<u>operator++</u>	(ODMG) Increment operator; advances this VArray iterator forward to the next element in the VArray.
<u>operator</u>	(ODMG) Decrement operator; moves this VArray iterator backward to the previous element in the VArray.
<u>operator=</u>	(ODMG) Assignment operator; assigns the specified VArray iterator to this VArray iterator.

# **Constructors and Destructors**

# d\_lterator

(*ODMG*) Default constructor that constructs a new uninitialized VArray iterator.

d\_Iterator();

# Discussion You initialize the new VArray iterator by assigning another (initialized) VArray iterator to it. (You normally obtain an initialized VArray iterator by calling the <u>create\_iterator</u> member function on a VArray of the same *element\_type*.)

# d\_lterator

(*ODMG*) Copy constructor that constructs a new VArray iterator initialized with the specified VArray iterator.

d\_Iterator(const d\_Iterator<element\_type> &iterator);

#### Parameters iterator

VArray iterator of the same *element\_type* as the VArray iterator you are constructing.

# Operators

# operator++

(*ODMG*) Increment operator; advances this VArray iterator forward to the next element in the VArray.

	<pre>1. d_Iterator<element_type> &amp;operator++();</element_type></pre>
	<pre>2. d_Iterator<element_type> operator++(int n);</element_type></pre>
Parameters	n This parameter is not used in calling this operator; its presence in the function declaration specifies a postfix operator.
Returns	(Variant 1) This VArray iterator, advanced to the next element. (Variant 2) A new VArray iterator set to this iterator, before this iterator is advanced.
Discussion	Variant 1 is the prefix increment operator, which advances this VArray iterator and then returns it.
	Variant 2 is the postfix increment operator, which allocates and returns a new VArray iterator set to this iterator, and then advances this iterator.
	An error is signaled if you attempt to advance the VArray iterator once it has already reached the end of the VArray.
	The increment operator is equivalent to the <u>advance</u> member function.
operator	

(*ODMG*) Decrement operator; moves this VArray iterator backward to the previous element in the VArray.

	<pre>1. d_Iterator<element_type> &amp;operator();</element_type></pre>
	<pre>2. d_Iterator<element_type> operator(int n);</element_type></pre>
Parameters	n This parameter is not used in calling this operator; its presence in the function declaration specifies a postfix operator.
Returns	(Variant 1) This VArray iterator, moved to the previous element. (Variant 2) A new VArray iterator set to this iterator, before this iterator is decremented.

Discussion Variant 1 is the prefix decrement operator, which moves this VArray iterator and then returns it.

Variant 2 is the postfix decrement operator, which allocates and returns a new VArray iterator set to this iterator, and then moves this iterator.

An error is signaled if you attempt to move the VArray iterator backward once it has already reached the first element of the VArray.

## operator=

(*ODMG*) Assignment operator; assigns the specified VArray iterator to this VArray iterator.

Parameters iterator

VArray iterator of the same *element\_type* as this VArray iterator.

# **Member Functions**

## advance

(ODMG) Advances this VArray iterator to the next element in the VArray.

```
void advance();
```

# get\_element

(*ODMG*) Returns the element at this VArray iterator's current position in the VArray.

element\_type get\_element() const;

Returns The current element in the VArray.

Discussion An error is signaled if there is no current element—for example, if the iteration is completed or if the VArray has no elements.

## next

(*ODMG*) Tests whether to continue iterating and, if so, assigns the current element to the output parameter and advances this VArray iterator to the next element.

int next(element\_type &found);

Parameters	found
	Output parameter to which to assign the current element in the VArray.
Returns	1, if there is a next element in the VArray; 0, if iteration is complete.
Discussion	This member function is implemented using the not_done, get_element, and advance member functions.

## not\_done

(ODMG) Tests whether to continue iterating.		
<pre>int not_done() const;</pre>		
1, if there is a next element in the VArray; 0, if iteration is complete		

#### reset

Returns

(*ODMG*) Reinitializes this VArray iterator to the start of iteration for the same VArray.

void reset();

This member function positions the iterator just before the first element of the VArray. Consequently, after a reset, the <u>advance</u> member function positions the iterator at the first element. However, you do *not* need to advance the iterator before calling either <u>get\_element</u> or <u>next</u>, because both of these member functions treat a reset iterator as if it started at the first element.

# d\_Ref\_Any Class

#### Inheritance: d\_Ref\_Any

The non-persistence-capable ODMG class d\_Ref\_Any represents an *ODMG* generic reference to a persistent object of any type.

See:

- "Reference Summary" on page 125 for an overview of member functions
- "Reference Index" on page 126 for a list of member functions

# About ODMG Generic References

An ODMG generic reference allows conversions among object references in the type hierarchy. An ODMG generic reference can be used as an intermediary between any two types  $d_{Ref}<x>$  and  $d_{Ref}<y>$ , where x and y are different types. An Objectivity/C++ object reference ooRef(x) can always be converted to or constructed from a  $d_{Ref}$ -Any; each object-reference class has an appropriate constructor and a conversion operator.

# **Reference Summary**

Creating a Generic Reference	<u>d_Ref_Any</u>
Setting a Generic Reference	<u>clear</u> <u>operator=</u>
Getting Information	<u>is null</u>

Comparing Generic References	::operator== ::operator!=
Deleting the Referenced Object	<u>delete_object</u>

# **Reference Index**

<u>clear</u>	(ODMG) Sets this ODMG generic reference to null.
<u>delete_object</u>	(ODMG) Deletes the persistent object referenced by this ODMG generic reference.
<u>d Ref Any</u>	(ODMG) Default constructor that constructs a new null ODMG generic reference.
<u>d Ref Any</u>	(ODMG) Constructs a new ODMG generic reference that refers to the same persistent object as the specified ODMG generic reference or pointer.
<u>is_null</u>	(ODMG) Tests whether this ODMG generic reference is null.
<u>operator=</u>	(ODMG) Assignment operator; sets this ODMG generic reference to refer to the same persistent object as the specified ODMG generic reference or pointer.
::operator==	(ODMG) Equality operator; tests whether the same persistent object is referenced by both of the specified items.
::operator!=	(ODMG) Inequality operator; tests whether different persistent objects are referenced by the two specified items.

# **Constructors and Destructors**

# d\_Ref\_Any

(ODMG) Default constructor that constructs a new null ODMG generic reference.

d\_Ref\_Any();

# d\_Ref\_Any

(*ODMG*) Constructs a new ODMG generic reference that refers to the same persistent object as the specified ODMG generic reference or pointer.

- 1. d\_Ref\_Any(const d\_Ref\_Any &);
- 2. d\_Ref\_Any(d\_Persistent\_Object \*);

# Operators

## operator=

(*ODMG*) Assignment operator; sets this ODMG generic reference to refer to the same persistent object as the specified ODMG generic reference or pointer.

- 1. d\_Ref\_Any & operator=(const d\_Ref\_Any &);
- 2. d\_Ref\_Any &operator=(d\_Persistent\_Object \*);

#### ::operator==

global function

(*ODMG*) Equality operator; tests whether the same persistent object is referenced by both of the specified items.

1.	<pre>int ::operator==(    const d_Ref_Any &amp;left,    const d_Ref_Any &amp;right);</pre>
2.	<pre>int ::operator==(     const d_Ref_Any &amp;left,     const d_Persistent_Object *right);</pre>
3.	<pre>int ::operator==(     const d_Persistent_Object *left,     const d_Ref_Any &amp;right);</pre>
4.	<pre>int ::operator==(     const ooRefHandle(ooObj) &amp;left,     const d_Ref_Any &amp;right);</pre>
5.	<pre>int ::operator==(     const d_Ref_Any &amp;left,     const ooRefHandle(ooObj) &amp;right);</pre>

Returns 1, if *left* and *right* refer to the same persistent object; otherwise 0.

Discussion You can compare two ODMG generic references (variant 1), an ODMG generic reference and a pointer to a persistent object (variants 2 and 3), or an ODMG generic reference and an object reference or handle to a persistent object (variants 4 and 5).

## ::operator!=

global function

(*ODMG*) Inequality operator; tests whether different persistent objects are referenced by the two specified items.

1.	<pre>int ::operator!=(    const d_Ref_Any &amp;left,    const d_Ref_Any &amp;right);</pre>
2.	<pre>int ::operator!=(    const d_Ref_Any &amp;left,    const d_Persistent_Object *right);</pre>
3.	<pre>int ::operator!=(    const d_Persistent_Object *left,    const d_Ref_Any &amp;right);</pre>
4.	<pre>int ::operator!=(    const ooRefHandle(ooObj) &amp;left,    const d_Ref_Any &amp;right);</pre>
5.	<pre>int ::operator!=(    const d_Ref_Any &amp;left,    const ooRefHandle(ooObj) &amp;right);</pre>
1, if 1	eft and right do not refer to the same persistent object; otherwise 0.

Discussion You can compare two ODMG generic references (variant 1), an ODMG generic reference and a pointer to a persistent object (variants 2 and 3), or an ODMG generic reference and an object reference or handle to a persistent object (variants 4 and 5).

# **Member Functions**

# clear

Returns

(ODMG) Sets this ODMG generic reference to null.

void clear();

# delete\_object

(ODMG) Deletes the persistent object referenced by this ODMG generic reference.

void delete\_object();

# is\_null

(ODMG) Tests whether this ODMG generic reference is null. int is\_null() const;

Returns 1, if this generic reference is null; 0, if it is not null.

Member Functions

# d\_Time Class

Inheritance: d\_Time

The non-persistence-capable ODMG class d\_Time represents a specific *time value*, which is internally stored in Greenwich Mean Time (GMT).

See:

- "Reference Summary" on page 135 for an overview of member functions
- "Reference Index" on page 136 for a list of member functions

To use this class, your application must include the ooTime.h header file. No extra linking is required.

# **About Time Values**

A time value represents an instant in time expressed as a number of hours, minutes, and seconds since midnight. All arithmetic on time values is done on a modulo 24-hour basis.

An application can create a time value from the current time or from a specified number of hours, minutes, and seconds. In either case, the number of hours is expressed locally to the time zone specified by the application. The hours, minutes, and seconds components are then normalized so that they are expressed as the number of milliseconds after midnight in Greenwich Mean Time (GMT), and the time value stores the difference in hours between the specified time zone and GMT.

An application can specify a time zone explicitly for each time value it creates. More typically, however, an application sets a default time zone and then creates its time values without explicitly specifying their time zones. Initially, the default time zone is set to GMT; the application uses a static member function to change its default time zone (for example, to the computer's local time zone). Because the default time zone is static data, it is compiled into the application, but not stored in the database.

An application can specify a time zone as a constant of type Time\_Zone. Each time zone is numbered according to the number of hours that must be added or subtracted from local time to get GMT. For example, GMT6 indicates a time of 6 hours ahead of GMT, so 6 must be subtracted from it to get GMT. Conversely, GMT\_6 means that the time is 6 hours earlier than GMT (read the underscore as a minus), so 6 hours must be added to get GMT.

It is the application's responsibility to adjust for daylight savings time, if necessary; the d\_Time class does not provide any mechanisms for doing so.

# **Examples**

The following program is a prototype for simulating a simple roaming application and a simple billing application for cellular telephone calls. The first two transactions of the program each simulate the roaming application, which records the time and duration of a cellular telephone call. The third transaction simulates the centralized billing application that calculates the cost of each call.

```
// DDL file
class eventObject : public ooObj {
public:
   int32 event_num;
                                  // unique event number
  d_Time event_time;
                                   // beginning of a call
                                  // duration of a call
   int32 duration in seconds;
};
// Application code
int main() {
   ooHandle(ooFBObj) fdH;
   ooHandle(ooDBObi) dbH;
   ooHandle(ooContObj) contH;
   ooHandle(eventObject) eventH;
   ooTrans trans;
   ooInit();
// The first transaction simulates the roaming application at
// a local receiving station, which records a call made by
// customer1
   trans.start();
   fdH.open("Example1", oocUpdate);
```

```
if(!dbH.exist(fdH, "ROAMING_CALLS") ) {
      dbH = new(fdH) ooDBObj("ROAMING CALLS");
   }
   dbH.open(fdH, "ROAMING_CALLS", oocUpdate);
   if(!contH.exist(dbH,"Customer1")) {
      contH = new ("Customer1",0,0,0,dbH) cont ("Customer1");
   }
   contH.open(dbH,"Customer1",oocUpdate);
   // Assume customer1 makes a call from zone GMT_5 (-5 hours),
   // starting at 11.30 p.m. and lasting 50 minutes.
   // Set the application's default time zone to GMT_5
   d_Time::set_default_Time_Zone(d_Time::Time_Zone::GMT_5);
   // Create a time value at the time of call activation
   // The new time value is local to GMT_5
   d Time call N 1;
   // Store the time value in an event that represents the call
   eventH = new(contH) eventObject();
   eventH->event_num = 1;
   eventH->event_time = call_N_1;
   eventH->duration_in_seconds = 300;
   trans.commit();
// The second transaction simulates the roaming application at
// a second local receiving station, which records the second
// call made by customer
   trans.start();
   // open the federated database
   // open the ROAMING CALLS database
   // open Customer1 container
   ...
   // Assume customer1 travels for a day and makes a call
   // from zone GMT10 (+10 hours), starting at 11.50 a.m. and
   // lasting 30 minutes.
   // Set the application's default time zone to GMT10
   d_Time::set_default_Time_Zone(d_Time::Time_Zone::GMT10);
```

```
// Create a time value at the time of call activation
   // The new time value is local to GMT10
   d_Time call_N_2;
   // Store the time value in an event that represents the call
   eventH = new(contH) eventObject();
   eventH->event_num = 2;
   eventH->event_time = call_N_2;
   eventH->duration_in_seconds = 180;
   trans.commit();
// The third transaction simulates a central billing application,
// which calculates a statement for customer1
   trans.start();
   // open the federated database
   // open the ROAMING CALLS database
   // open Customer1 container
   . . .
   ooItr(eventObject) eventI;
   unsigned short local_hour_of_call, GMT_hour_of_call;
   short local to GMT difference;
   // Find each call event for customer1 and calculate its cost
   eventI.scan(contH);
   while(eventI.next()) {
      // Get the call's time value expressed in GMT
      GMT_hour_of_call = eventI->event_time.hour();
      // Get the number of hours between the call's time zone
      // and GMT
      local to GMT difference =
            eventI->event_time.tz_hour();
      // Reconstruct the time of the call in its time zone
      local_hour_of_call =
            GMT_hour_of_call + local_to_GMT_difference;
      // Use local_hour_of_call, the call's duration, and the
      // call's event time->minute() value to determine whether
      // to use a peak-time rate, an off-peak rate, or a
      // combination of rates. Get the cost of the call by
      // calculating the number of minutes at each rate.
   } // End while
```

```
// Print a statement showing each call and its cost.
...
trans.commit();
return 1;
} // End main
```

# **Reference Summary**

Creating a Time Value	<u>d_Time</u>
Time Zones	Time_Zone
Setting a Time Value	operator=
Setting a Time Zone	<u>set default Time Zone</u> <u>set default Time Zone to local</u>
Getting Information	<u>current</u> <u>hour</u> <u>minute</u> <u>second</u> <u>tz_hour</u> <u>tz_minute</u>
Comparing Time Values	<pre>::operator== ::operator!= ::operator&lt; ::operator&lt;= ::operator&gt;=</pre>
Arithmetic Operations on Time Values	<u>::operator+</u> <u>operator+=</u> <u>::operator-</u> <u>operator-=</u>
Determining Whether Time Periods Overlap	::overlaps

# **Reference Index**

<u>current</u>	(ODMG) Returns the current time in the default time zone.
<u>d Time</u>	(ODMG) Default constructor that constructs a new time value set to the current time in the default time zone.
<u>d Time</u>	(ODMG) Constructs a new time value from the specified time value.
<u>d Time</u>	(ODMG) Constructs a new time value from the specified components.
hour	(ODMG) Returns the hours component of this time value, expressed in GMT.
minute	(ODMG) Returns the minutes component of this time value.
::operator+	(ODMG) Addition operator; adds the specified time value and interval and returns the sum in a newly allocated time value.
<u>operator+=</u>	(ODMG) Increment operator; adds the specified interval to this time value.
::operator-	(ODMG) Subtraction operator; subtracts a time value or interval from a time value and returns the difference in a newly allocated time value.
<u>operator-=</u>	(ODMG) Decrement operator; subtracts the specified interval from this time value.
<u>operator=</u>	(ODMG) Assignment operator; sets this time value to the specified time value.
::operator==	(ODMG) Equality operator; returns 1 if every component of one time value matches the corresponding component of the other.
::operator!=	(ODMG) Inequality operator; returns 1 if any component of one time value differs from the corresponding component of the other.
::operator<	(ODMG) Less-than operator; returns 1 if one time value is less than another.

::operator<=	(ODMG) Less-than-or-equal-to operator; returns 1 if one time value is less than or equal to another.
::operator>	(ODMG) Greater-than operator; returns 1 if one time value is greater than another.
::operator>=	(ODMG) Greater-than-or-equal-to operator; returns 1 if one time value is greater than or equal to another.
::overlaps	(ODMG) Returns 1 if two specified time periods overlap, where one or both periods are specified using time values.
second	(ODMG) Returns the seconds component of this time value.
<u>set_default_Time_Zone</u>	(ODMG) Sets the default time zone.
set_default_Time_Zone_to_local	(ODMG) Resets the default time zone to your computer's local time zone.
<u>Time Zone</u>	(ODMG) Enumerated type used to denote a specific time zone.

value and GMT.

time value and GMT.

(ODMG) Returns the number of hours between the time zone stored in this time

(ODMG) Returns the number of minutes between the time-zone minutes stored in this

# **Types**

# Time\_Zone

(ODMG) Enumerated type used to denote a specific time zone.

```
enum Time_Zone {
     GMT = 0,
     GMT12 = 12,
     GMT_{12} = -12,
     GMT1 = 1,
     GMT_1 = -1,
     GMT2 = 2,
     GMT_2 = -2,
```

tz hour

tz\_minute

GMT3 = 3,  $GMT_3 = -3$ , GMT4 = 4,  $GMT_4 = -4,$ GMT5 = 5,  $GMT_5 = -5$ , GMT6 = 6, GMT = -6, GMT7 = 7,  $GMT_7 = -7,$ GMT8 = 8,  $GMT_8 = -8,$ GMT9 = 9,  $GMT_9 = -9,$ GMT10 = 10,  $GMT_{10} = -10$ , GMT11 = 11,  $GMT_{11} = -11$ , USeastern = -5, UScentral = -6, USmountain = -7, USpacific = -8;

Discussion Time zones are numbered according to the number of hours that must be added or subtracted from the local time to get the time in Greenwich, England (GMT). For example, GMT6 indicates a time of 6 hours greater than GMT, so 6 must be subtracted from it to get GMT. Conversely, GMT\_6 means that the time is 6 hours earlier than GMT (read the underscore as a minus).

# Constructors

# d\_Time

(*ODMG*) Default constructor that constructs a new time value set to the current time in the default time zone.

d\_Time();

# d\_Time

(ODMG) Constructs a new time value from the specified time value.

- 1. d\_Time(const d\_Time &);
- 2. d\_Time(const d\_Timestamp &);

# d\_Time

(ODMG) Constructs a new time value from the specified components.

1.	d_Time(
	unsigned short hour,
	unsigned short minute,
	<pre>float second);</pre>
2.	d_Time(
	unsigned short hour,
	unsigned short minute,
	float <i>second</i> ,
	short <i>tzhour</i> ,
	<pre>short tzminute);</pre>

#### Parameters

Hours component of the new time value; the number of hours past midnight.

#### minute

hour

Minutes component of the new time value; the number of minutes in any fractional hour remaining after whole hours are subtracted out.

#### second

Seconds component of the new time value; the number of seconds in any fractional minute remaining after whole minutes are subtracted out.

#### tzhour

The time zone for the new time value. You can specify the time zone using a constant of type  $\underline{\texttt{Time Zone}}$ . Alternatively, you can specify the integer number of hours to subtract from the hours component of the new time value to get the hours component of GMT:

- Specify a positive number if the new time value is later than GMT.
- Specify a negative number if the new time value is earlier than GMT.

#### tzminute

Number of minutes to subtract from the minutes component of the new time value to get the minutes component of GMT; normally 0.

Discussion Variant 1 constructs a new time value that is local to the default time zone, which implicitly sets the *tzhour* and *tzminute* components.

Variant 2 allows you to construct a time value that is local to a nondefault time zone.

## ::operator+

global function

(*ODMG*) Addition operator; adds the specified time value and interval and returns the sum in a newly allocated time value.

 d\_Time ::operator+( const d\_Time &left, const d\_Interval &right);
 d\_Time ::operator+( const d\_Interval &left, const d\_Time &right);

## operator+=

(ODMG) Increment operator; adds the specified interval to this time value.

d\_Time &d\_Time::operator+=(const d\_Interval &interval);

## ::operator-

global function

(*ODMG*) Subtraction operator; subtracts a time value or interval from a time value and returns the difference in a newly allocated time value.

- 2. d\_Time ::operator -(
   const d\_Time &left,
   const d\_Interval &right);

# operator-=

(*ODMG*) Decrement operator; subtracts the specified interval from this time value.

d\_Time &operator = (const d\_Interval &interval);

# operator=

(ODMG) Assignment operator; sets this time value to the specified time value.

- 1. d\_Time & operator = (const d\_Time & time);
- 2. d\_Time operator=(const d\_Timestamp &timeStamp);

## ::operator==

(*ODMG*) Equality operator; returns 1 if every component of one time value matches the corresponding component of the other.

```
int ::operator==(
    const d_Time &left,
    const d_Time &right);
```

# ::operator!=

(*ODMG*) Inequality operator; returns 1 if any component of one time value differs from the corresponding component of the other.

```
int ::operator!= (
    const d_Time &left,
    const d_Time &right);
```

# ::operator<

(ODMG) Less-than operator; returns 1 if one time value is less than another.

int ::operator< (
 const d\_Time &left,
 const d\_Time &right);</pre>

## ::operator<=

(*ODMG*) Less-than-or-equal-to operator; returns 1 if one time value is less than or equal to another.

int ::operator<=(
 const d\_Time &left,
 const d\_Time &right);</pre>

global function

global function

# global function

global function

## ::operator>

global function

(ODMG) Greater-than operator; returns 1 if one time value is greater than another.

```
int ::operator> (
    const d_Time &left,
    const d_Time &right);
```

#### ::operator>=

global function

(*ODMG*) Greater-than-or-equal-to operator; returns 1 if one time value is greater than or equal to another.

```
int ::operator>=(
    const d_Time &left,
    const d_Time &right);
```

# **Member Functions**

## current

(*ODMG*) Returns the current time in the default time zone.

static d\_Time current();

# hour

(ODMG) Returns the hours component of this time value, expressed in GMT.

unsigned short hour() const;

Discussion This member function expresses the hours component of this time value as a number of hours past midnight in GMT. To get the hours component expressed locally to the time value's time zone, you must add the value returned by <u>tz\_hour</u>.

# minute

(ODMG) Returns the minutes component of this time value.

unsigned short minute() const;

## second

(ODMG) Returns the seconds component of this time value.

```
float second() const;
```

# set\_default\_Time\_Zone

(ODMG) Sets the default time zone.

static void set\_default\_Time\_Zone(Time\_Zone zone);

Discussion The default time zone value is initially set to GMT. (This behavior differs from the ODMG standard, in which the default time zone is initially set to your computer's local time.) Changing the default time zone affects subsequently created time values, but does not affect any existing time values.

# set\_default\_Time\_Zone\_to\_local

(ODMG) Resets the default time zone to your computer's local time zone.

static void set\_default\_Time\_Zone\_to\_local();

Discussion Changing the default time zone affects subsequently created time values, but does not affect any existing time values.

## tz\_hour

(*ODMG*) Returns the number of hours between the time zone stored in this time value and GMT.

short tz\_hour() const;

Returns The number of hours add to the value returned by <u>hour</u> to get the hours component in this time value's time zone. A positive number indicates this time value is later than GMT; a negative number indicates this time value is earlier than GMT.

# tz\_minute

(*ODMG*) Returns the number of minutes between the time-zone minutes stored in this time value and GMT.

short tz\_minute() const;

Returns The number of minutes to add to the value returned by <u>minute</u> to get the minutes component in this time value's time zone; normally 0.

# **Related Global Functions**

# ::overlaps

(*ODMG*) Returns 1 if two specified time periods overlap, where one or both periods are specified using time values.

1.	<pre>int ::overlaps( const d_Time &amp;startLeft, const d_Time &amp;endLeft, const d_Time &amp;startRight, const d_Time &amp;endRight);</pre>
2.	<pre>int ::overlaps(     const d_Timestamp &amp;startLeft,     const d_Timestamp &amp;endLeft,     const d_Time &amp;startRight,     const d_Time &amp;endRight);</pre>
3.	<pre>int ::overlaps(     const d_Time &amp;startLeft,     const d_Time &amp;endLeft,     const d_Timestamp &amp;startRight,     const d_Timestamp &amp;endRight);</pre>

Discussion Each time period is specified by a start and end time. Variant 1 allows you to specify the time periods using time values. Variants 2 and 3 allow you to specify time periods using time values and timestamps.
# d\_Timestamp Class

#### Inheritance: d\_Timestamp

The non-persistence-capable ODMG class d\_Timestamp represents a timestamp, which contains both a date (see <u>d\_Date</u>) and a time value (see <u>d\_Time</u>).

See:

- "Reference Summary" on page 145 for an overview of member functions
- "Reference Index" on page 146 for a list of member functions

To use this class, your application must include the ooTime.h header file. No extra linking is required.

## **Reference Summary**

Creating a Timestamp	<u>d_Timestamp</u>
Setting a Timestamp	<u>operator=</u>
Getting Information	<u>current</u> <u>date</u> <u>day</u> <u>hour</u> <u>minute</u> <u>month</u> <u>second</u> <u>time</u> <u>tz hour</u> <u>tz minute</u> <u>year</u>

Comparing Timestamps	<pre>::operator== ::operator!= ::operator&lt; ::operator&lt;=</pre>
	::operator> ::operator>=
Arithmetic Operations on Timestamps	::operator+ operator+= ::operator- operator-=
Determining Whether Time Periods Overlap	::overlaps

# **Reference Index**

<u>current</u>	(ODMG) Returns a timestamp containing the current date and time in the default time zone.
date	(ODMG) Returns the date component of this timestamp.
day	(ODMG) Returns the day component of this timestamp.
<u>d Timestamp</u>	(ODMG) Default constructor that constructs a new timestamp set to the current date and time.
<u>d_Timestamp</u>	(ODMG) Constructs a new timestamp from the specified object(s).
<u>d Timestamp</u>	(ODMG) Constructs a new timestamp with the specified components.
hour	(ODMG) Returns the hours component of this timestamp, expressed in GMT.
<u>minute</u>	(ODMG) Returns the minutes component of this timestamp.
month	(ODMG) Returns the month component of this timestamp.
::operator+	(ODMG) Addition operator; allocates a new timestamp that is set to the sum of the specified timestamp and interval.
<u>operator+=</u>	(ODMG) Increment operator; adds the specified interval to this timestamp.
::operator-	(ODMG) Subtraction operator; allocates a new timestamp that is set to the specified timestamp minus the specified interval.

<u>operator-=</u>	(ODMG) Decrement operator; subtracts the specified interval from this timestamp.
<u>operator=</u>	(ODMG) Assignment operator; assigns the specified timestamp or date to this timestamp.
::operator==	(ODMG) Equality operator; returns 1 if every component of one timestamp matches the corresponding component of the other.
::operator!=	(ODMG) Inequality operator; returns 1 if any component of one timestamp differs from the corresponding component of the other.
::operator<	(ODMG) Less-than operator; returns 1 if one timestamp is less than another.
::operator<=	(ODMG) Less-than-or-equal-to operator; returns 1 if one timestamp is less than or equal to another.
::operator>	(ODMG) Greater-than operator; returns 1 if one timestamp is greater than another.
::operator>=	(ODMG) Greater-than-or-equal-to operator; returns 1 if one timestamp is greater than or equal to another.
::overlaps	(ODMG) Returns 1 if two time periods overlap, where one or both periods are specified using timestamps.
second	(ODMG) Returns the seconds component of this timestamp.
time	(ODMG) Returns the time component of this timestamp.
<u>tz hour</u>	(ODMG) Returns the number of hours between the time zone stored in this timestamp and GMT.
<u>tz_minute</u>	(ODMG) Returns the number of minutes between the time-zone minutes stored in this timestamp and GMT.
year	(ODMG) Returns the year component of this timestamp.

# Constructors

## d\_Timestamp

(*ODMG*) Default constructor that constructs a new timestamp set to the current date and time.

d\_Timestamp();

## d\_Timestamp

(ODMG) Constructs a new timestamp from the specified object(s).

- 1. d\_Timestamp(const d\_Date &);
- 2. d\_Timestamp(const d\_Date &, const d\_Time &);
- 3. d\_Timestamp(const d\_Timestamp &);

Discussion You can construct a timestamp from:

- A date plus the current time (variant 1)
- A date and a time value (variant 2)
- Another timestamp (variant 3)

## d\_Timestamp

(ODMG) Constructs a new timestamp with the specified components.

```
d_Timestamp(
    unsigned short year,
    unsigned short month= 1,
    unsigned short day = 1,
    unsigned short hour = 0,
    unsigned short minute = 0,
    float second = 0.0);
```

Parameters year

Corresponds to the year component of a d\_Date. Maximum value is 65535.

#### month

Corresponds to the month component of a d\_Date.

day

Corresponds to the day component of a d\_Date.

hour

Corresponds to the hours component of a d\_Time.

#### minute

Corresponds to the minutes component of a d\_Time.

second

Corresponds to the seconds component of a d\_Time.

Discussion The time value in the new timestamp is local to the application's default time zone, which is set by d\_Time::<u>set\_default\_Time\_Zone</u>. The timestamp stores

the time zone as the number of hours between the local time and Greenwich Mean Time (GMT).

## **Operators**

### ::operator+

global function

(*ODMG*) Addition operator; allocates a new timestamp that is set to the sum of the specified timestamp and interval.

1.	d_Timestamp ::operator+(
	const d_Timestamp & <i>left</i> ,
	<pre>const d_Interval &amp;right);</pre>
2.	d_Timestamp ::operator+(
	<pre>const d_Interval &amp;left,</pre>
	<pre>const d_Timestamp &amp;right);</pre>

*WARNING* Addition that causes the year component to exceed 65535 results in an invalid timestamp.

### operator+=

(ODMG) Increment operator; adds the specified interval to this timestamp.

d\_Timestamp & operator+=(const d\_Interval &);

*WARNING* Addition that causes the year component to exceed 65535 results in an invalid timestamp.

### ::operator-

global function

(*ODMG*) Subtraction operator; allocates a new timestamp that is set to the specified timestamp minus the specified interval.

```
d_Timestamp ::operator-(
    const d_Timestamp &left,
    const d_Interval &right);
```

#### operator-=

(*ODMG*) Decrement operator; subtracts the specified interval from this timestamp.

d\_Timestamp & operator -= (const d\_Interval &);

#### operator=

(*ODMG*) Assignment operator; assigns the specified timestamp or date to this timestamp.

```
1. d_Timestamp & operator=(const d_Timestamp &);
```

```
2. d_Timestamp & operator=(const d_Date &);
```

#### ::operator==

global function

(*ODMG*) Equality operator; returns 1 if every component of one timestamp matches the corresponding component of the other.

```
int ::operator==(
    const d_Timestamp &left,
    const d_Timestamp &right);
```

### ::operator!=

global function

(*ODMG*) Inequality operator; returns 1 if any component of one timestamp differs from the corresponding component of the other.

```
int ::operator!= (
    const d_Timestamp &left,
    const d_Timestamp &right);
```

### ::operator<

global function

(ODMG) Less-than operator; returns 1 if one timestamp is less than another.

```
int ::operator< (
    const d_Timestamp &left,
    const d_Timestamp &right);</pre>
```

### ::operator<=

Member Functions

global function

(*ODMG*) Less-than-or-equal-to operator; returns 1 if one timestamp is less than or equal to another.

```
int ::operator<=(
    const d_Timestamp &left,
    const d_Timestamp &right);</pre>
```

### ::operator>

global function

(*ODMG*) Greater-than operator; returns 1 if one timestamp is greater than another.

```
int ::operator> (
    const d_Timestamp &left,
    const d_Timestamp &right);
```

### ::operator>=

global function

(*ODMG*) Greater-than-or-equal-to operator; returns 1 if one timestamp is greater than or equal to another.

```
int ::operator>=(
    const d_Timestamp &left,
    const d_Timestamp &right);
```

## **Member Functions**

### current

(*ODMG*) Returns a timestamp containing the current date and time in the default time zone.

static d\_Timestamp current();

### date

(ODMG) Returns the date component of this timestamp.

const d\_Date &date() const;

day	
	(ODMG) Returns the day component of this timestamp.
	unsigned short day() const;
bour	
nour	(ODMG) Returns the hours component of this timestamp, expressed in GMT
	(opine) notation in the second component of the intestant, expressed in anti-
Discussion	This member function expresses the hours component of this timestamp as a number of hours past midnight in GMT. To get the hours component expressed locally to the time value's time zone, you must add the value returned by $\underline{tz}$ hour.
minute	
	(ODMG) Returns the minutes component of this timestamp.
	unsigned short minute() const;
w o w th	
month	(ODMC) Deturns the month component of this timestown
	(ODMG) Returns the month component of this timestamp.
	unsigned short month() const;
second	
	(ODMG) Returns the seconds component of this timestamp.
	<pre>float second() const;</pre>
timo	
	(ODMC) Returns the time component of this timestamp
	(ODMO) Returns the time component of this timestamp.
	Const d_Time &time() Const;
tz_hour	
	( <i>ODMG</i> ) Returns the number of hours between the time zone stored in this timestamp and GMT.
	<pre>short tz_hour() const;</pre>

Returns The number of hours to add to the value returned by <u>hour</u> to get the hours component in this timestamp's time zone. A positive number indicates this time value is later than GMT; a negative number indicates this time value is earlier than GMT.

### tz\_minute

(*ODMG*) Returns the number of minutes between the time-zone minutes stored in this timestamp and GMT.

short tz\_minute() const;

Returns The number of minutes to add to the value returned by <u>minute</u> to get the minutes component in this timestamp's time zone; normally 0.

#### year

(ODMG) Returns the year component of this timestamp.

unsigned short year() const;

# **Related Global Functions**

### ::overlaps

global function

(*ODMG*) Returns 1 if two time periods overlap, where one or both periods are specified using timestamps.

```
1.
    int ::overlaps(
      const d_Timestamp &startLeft,
      const d_Timestamp & endLeft,
      const d_Timestamp &startRight,
      const d_Timestamp & endRight);
2.
    int ::overlaps(
      const d_Timestamp &startLeft,
      const d_Timestamp & endLeft,
      const d_Date &startRight,
      const d Date &endRight);
3.
    int ::overlaps(
      const d Date &startLeft,
      const d_Date & endLeft,
      const d_Timestamp &startRight,
      const d_Timestamp &endRight);
```

4. int ::overlaps(
const d_Timestamp & <i>startLeft</i> ,
const d_Timestamp & <i>endLeft</i> ,
const d_Time & <i>startRight</i> ,
<pre>const d_Time &amp;endRight);</pre>
5. int :: overlaps (
const d_Time & <i>startLeft</i> ,
const d_Time & <i>endLeft</i> ,
const d_Timestamp & <i>startRight</i> ,
<pre>const d_Timestamp &amp;endRight);</pre>

Discussion Each time period is specified by a start and end time. You can specify the time periods using timestamps (variant 1), using timestamps and dates (variants 2 and 3), or using timestamps and time values (variants 4 and 5).

# ooAdmin Class

Inheritance: ooObj->ooAdmin

Handle Class: ooHandle(ooAdmin)

**Object-Reference Class:** ooRef(ooAdmin)

The persistence-capable class ooAdmin is the abstract base class for all *administrator* classes.

To use this class, your application must include the ooCollections.h header file. For UNIX linking information, see *Installation and Platform Notes for UNIX*.

## **About Administrators**

An administrator is an object of a concrete class derived from ooAdmin. Each collection derived from ooCollection has an administrator that manages the containers for the collection's internal objects. A collection's administrator is created when the collection itself is created.

Because the ooAdmin class is abstract, you never create instances of it; instead, you work with instances of its concrete derived classes:

- Ordered collections use tree administrators of the <u>ooTreeAdmin</u> class.
- Unordered collections use hash administrators of the <u>ooHashAdmin</u> class.

You should not create your own subclasses of this class.

Like other persistent objects, administrators are normally manipulated through handles or object references. You call a collection's admin member function to obtain an object reference to the collection's administrator.

# ooAPObj Class

Inheritance: ooObj->ooAPObj

Handle Class: <u>ooHandle(ooAPObj)</u>

**Object-Reference Class:** <u>ooRef(ooAPObj)</u>

The persistence-capable class ooAPObj represents an *autonomous partition*. You use ooAPObj only for autonomous-partition creation; you work with an existing autonomous partition through a handle or object reference.

You can create and work with autonomous partitions only if you have bought and installed Objectivity/DB Fault Tolerant Option (Objectivity/FTO).

See:

■ "Reference Index" on page 159 for a list of ooAPObj member functions

For operations performed through an autonomous-partition handle or object reference, see:

■ "Reference Summary" on page 493

## **About Autonomous Partitions**

An autonomous partition is an independent piece of a federated database. Each autonomous partition is self-sufficient in case a network or system failure occurs in another partition. Although data physically resides in database files, each autonomous partition *controls* access to particular databases and containers. This capability is available through the Objectivity/FTO product.

Every autonomous partition has all the system resources necessary to run an Objectivity/DB application, including a boot file, a lock server, and a system-database file. The system-database file contains schema information and

a global catalog of all autonomous partitions, their locations, and the databases they contain.

Each autonomous partition has a system name, which is its logical name within the federated database. The system name of each autonomous partition must be unique among all the system names of databases and autonomous partitions in the federated database.

An autonomous partition also has a unique integer identifier that is assigned by Objectivity/DB when the partition is created.

See Chapter 27, "Autonomous Partitions," in the the Objectivity/C++ programmer's guide for additional information about autonomous partitions.

## **Working With Autonomous Partitions**

Autonomous partitions can be created and deleted with administration tools or from within an application. An application:

- Creates an autonomous partition using the ooAPObj class constructor and operator new. The new autonomous partition is represented in memory as an instance of ooAPObj, which serves as a proxy for the actual system-database file on disk.
- Deletes an autonomous using the global function ooDelete; operator delete is not available on class ooAPObj.

To work with a new autonomous partition, an application must assign the result of operator new to an autonomous-partition handle or object reference (an instance of ooHandle(ooAPObj) or ooRef(ooAPObj)). Similarly, to work with an existing autonomous partition, the application must open it through an autonomous-partition handle or object reference; multiple handles and object references can be set to reference the same autonomous partition. The application then operates on the referenced autonomous partition by:

- Calling a member function on an appropriate handle or object reference.
- Passing an appropriate handle or object reference to a global function.

Member functions for operating on existing autonomous partitions are defined in the *ooRefHandle*(*ooAPObj*) classes, not in the *ooAPObj* class itself. Although *ooAPObj* inherits member functions from *ooObj*, these generally do not apply to autonomous partitions; in fact, the inherited member functions cannot be called, because there is no indirect member-access operator (->) on the *ooRefHandle*(*ooAPObj*) classes.

You may not derive classes from ooAPObj.

## **Reference Index**

<u>ooAPObj</u>	Constructs a new autonomous partition with the specified system name, lock server host, and file locations.
<u>operator new</u>	Creates a new autonomous partition in the currently open federated database.

## Constructors

## ooAPObj

Constructs a new autonomous partition with the specified system name, lock server host, and file locations.

#### ooAPObj(

const	char	*apSysName,	
const	char	*lockServer,	
const	char	*apFileHost,	
const	char	*apFilePath,	
const	char	*bootFileHost	= 0,
const	char	*bootFilePath	= 0,
const	char	*jnlDirHost =	Ο,
const	char	*jnlDirPath =	0);

#### Parameters apSysName

System name of the autonomous partition to create. The specified name:

- Must follow the same naming rules as files of your operating system.
- Must be unique among all the system names of databases and autonomous partitions in the federated database.

The system name implicitly specifies the name of the autonomous-partition boot file.

#### lockServer

Name of the host that runs the lock server for the new autonomous partition.

#### apFileHost

Name of the data server host on which to create the autonomous partition's system-database file.

#### apFilePath

Fully qualified pathname (including the filename) of the autonomous partition's system-database file on *apFileHost*. By convention, you construct the filename from the system name with the extension . AP.

#### bootFileHost

Name of the data server host on which to create the autonomous partition's boot file. If you specify 0 (the default) and *bootFilePath* has a nondefault value, the boot file is created on the current host. If you specify 0 for both *bootFileHost* and *bootFilePath*, the boot file is created on *apFileHost*.

#### bootFilePath

Fully qualified pathname for the directory on *bootFileHost* in which to create the autonomous partition's boot file. The filename for the boot file is *apSysName.boot*. If you specify 0 (the default), the boot file is created in the same directory as the autonomous partition's system-database file on *apFileHost*.

#### jnlDirHost

Name of the data server host on which to create the journal directory for the autonomous partition. If you specify 0 (the default) and *jnlDirPath* has a nondefault value, the journal directory is created on the current host. If you specify 0 for both *jnlDirHost* and *jnlDirPath*, the journal directory is created on *apFileHost*.

#### jnlDirPath

Fully qualified pathname for the autonomous partition's journal directory on *jnlDirHost*. If you specify 0 (the default), the journal directory is created in the directory that contains the autonomous partition's system-database file on *apFileHost*.

## **Operators**

### operator new

Creates a new autonomous partition in the currently open federated database.

```
void *operator new(
    size_t,
    const ooRefHandle(ooFDObj) &containingFd);
```

Parameters size\_t

Do not specify; this parameter is automatically initialized by the compiler with the size of the class type in bytes.

#### containingFd

Federated database in which to create the new autonomous partition. Because an autonomous partition can be created only in the currently open federated database, you can omit this parameter, or, for explicitness, you can specify an object reference or handle to the open federated database.

- Returns Memory pointer to the new autonomous partition. This pointer is null if an error occurs during the creation of the autonomous partition or if an autonomous partition with the specified system name already exists.
- Discussion This operator must be used in an update transaction. When you commit or checkpoint the transaction, the new autonomous partition is made permanent on disk. If the transaction is aborted, the autonomous partition is not created.

Expressions containing this operator are of the following form:

new(containingFD) ooAPObj(initializers)

The *initializers* you specify are a list of values to be passed as parameters to the ooAPObj <u>constructor</u>. At a minimum, you must specify the system name, lock server host, and file location for the new autonomous partition.

You normally assign the result of operator new directly to an autonomous partition handle. You can verify the creation of the autonomous partition by checking whether the handle is null.

Operators

# ooBTree Class

Inheritance: ooObj->ooCollection->ooBTree

Handle Class: ooHandle(ooBTree)

**Object-Reference Class:** ooRef(ooBTree)

The persistence-capable class ooBTree is the abstract base class for classes that represent *ordered collections* containing persistent objects.

See:

- "Reference Summary" on page 164 for an overview of member functions
- "Reference Index" on page 164 for a list of member functions

To use this class, your application must include the ooCollections.h header file. For UNIX linking information, see *Installation and Platform Notes for UNIX*.

## **About Ordered Collections**

Ordered collections are *scalable*, that is, they can increase in size with minimal performance degradation. They are implemented using a B-tree data structure so that elements can be added, deleted, and retrieved efficiently. Each node in the B-tree has a corresponding array. The array for a non-leaf node contains references to the nodes at the leaf level of the B-tree; the array for a leaf node contains references to elements of the collection whose indexes are within a particular range. A newly created B-tree consists of a single node (the root of the B-tree) and its array. As the collection grows, additional nodes are created as necessary. When each node is created, its corresponding array is also created.

Concrete classes derived from ooBTree represent more specific kinds of ordered collections:

ooTreeSet represents sorted sets of persistent objects.

- ooTreeList represents lists of persistent objects.
- ooTreeMap represents sorted object maps.

Because the ooBTree class is abstract, you never create instances of it; instead, you work with instances of its concrete derived classes. You should not create your own subclasses of this class.

Like other persistent objects, ordered collections are normally manipulated through handles and object references.

For additional information about the the various persistence-capable collection classes, see Chapter 11, "Persistent Collections," in the Objectivity/C++ programmer's guide.

## **Reference Summary**

Adding and Removing Elements	remove
Getting Elements	<u>first</u> get <u>iterator</u> last
Getting Indexes	<u>indexOf</u> <u>lastIndexOf</u>
Getting Information	<u>depth</u> <u>size</u>
Testing	<u>contains</u> isEmpty
Maintaining the B-Tree	compact
Viewing in an MROW Transaction	refresh

## **Reference Index**

<u>compact</u>	Minimizes the number of nodes in this ordered collection's B-tree.
<u>contains</u>	Tests whether this ordered collection contains the specified element (or key).

<u>depth</u>	Gets the depth of this ordered collection's B-tree.
<u>first</u>	Finds the first element (or key) in this ordered collection.
get	Finds the specified element of this ordered collection.
<u>index0f</u>	Searches this ordered collection for the first element (or key) that is equal to the specified object.
isEmpty	Tests whether this ordered collection is empty.
iterator	Initializes a scalable-collection iterator to find the elements of this ordered collection.
last	Finds the last element (or key) in this ordered collection.
<u>lastIndexOf</u>	Searches this ordered collection for the last element (or key) that is equal to the specified object.
refresh	Refreshes each container used internally by this ordered collection,
	except for the container in which the collection itself is stored.
remove	except for the container in which the collection itself is stored. Removes the first occurrence of the specified object from this ordered collection.
<u>remove</u> <u>size</u>	except for the container in which the collection itself is stored. Removes the first occurrence of the specified object from this ordered collection. Gets the size of this ordered collection.

## **Member Functions**

### compact

Minimizes the number of nodes in this ordered collection's B-tree.

virtual void compact();

Discussion After you have added all elements that you expect this collection to have, you can call this member function to minimize the number of nodes in its B-tree. Doing so saves space and improves read performance. Indexes of elements within the collection remain unchanged.

If you call this member function before all elements have been added, insert (add) performance will not necessarily improve. After the B-tree has been compacted, adding an element will very likely cause one or more nodes to be added to the B-tree.

## contains

	Tests whether this ordered collection contains the specified element (or key).
	<pre>1. virtual ooBoolean contains(</pre>
	<pre>2. ooBoolean contains(</pre>
Parameters	objH Handle to the object to be tested for containment in this ordered collection.
	<i>lookupVal</i> Pointer to data that identifies the object to be tested for containment in this ordered collection.
Returns	ocTrue if this ordered collection contains an element (or key) equal to the specified object; otherwise, occFalse.
Discussion	Variant 2 tests whether any element (or key) of this ordered collection is "equal" to the specified lookup data, as determined by the comparator for this collection. It is useful if this ordered collection has an application-defined comparator that can identify an element (or key) based on class-specific data. Because a list has no comparator, variant 2 is not relevant for lists.
depth	
	Gets the depth of this ordered collection's B-tree.
	<pre>int depth() const;</pre>
Returns	The number of nodes between the root and a leaf node in this ordered collection's B-tree.
first	
	Finds the first element (or key) in this ordered collection.
	<pre>virtual ooRef(ooObj) first() const;</pre>
Returns	If the elements of this ordered collection are persistent objects, an object reference to the first element; if the elements are key-value pairs, an object reference to the key of the first element.
See also	last

get	
	Finds the specified element of this ordered collection.
	<pre>1. ooRef(ooObj) get(const ooInt32 index) const;</pre>
	<pre>2. virtual ooRef(ooObj) get(const void *lookupVal) const = 0;</pre>
Parameters	<i>index</i> The zero-based index of the desired element.
	lookupVal
	Pointer to data that identifies the desired element (or key).
Returns	(Variant 1) Finds the element at the specified index.
	<ul> <li>If the elements of this ordered collection are persistent objects, variant 1 returns an object reference to the element whose index is <i>index</i>.</li> </ul>
	■ If the elements are key-value pairs, variant 1 returns an object reference to the key of the element whose index is <i>index</i> .
	(Variant 2) Finds the element (or key) that is equal to the specified lookup data, as determined by the comparator for this ordered collection.
	If the elements of this ordered collection are persistent objects, variant 2 returns an object reference to the element that is equal to <i>lookupVal</i> , or a null object reference if this ordered collection does not contain such an element.
	■ If the elements are key-value pairs, variant 2 returns an object reference to the value of the element whose key is equal to <i>lookupVal</i> , or a null object reference if this ordered collection does not contain such an element.
Discussion	Variant 2 is useful if this ordered collection has an application-defined comparator that can identify an element (or key) based on class-specific data. Because a list has no comparator, variant 2 is not relevant for lists.
See also	<u>first</u> <u>last</u>
indexOf	
	Searches this ordered collection for the first element (or key) that is equal to the specified object.

```
virtual int indexOf(
    const ooHandle(ooObj) &objH,
    const int index = 0) const;
```

Parameters	objH
	Handle to the object whose index is to be found.
	index
	The zero-based index at which search should start.
Returns	The index of the first element in this ordered collection, at or after the starting position, that is equal to (or whose key is equal to) the specified object. If no such element is found, this member function returns -1.
Discussion	This member function searches forward in the collection, starting the search at the element whose index is <i>index</i> .
	If elements of this ordered collection are persistent objects, this member function compares each element with the specified object; if elements are key-value pairs, this member function compares the key of each element with the specified object.
	Search stops when a matching element is found; if more than one element matches, this member function finds the one closest to the beginning of the collection (but at or after the starting index).
See also	<u>lastIndexOf</u>
isEmpty	
	Tests whether this ordered collection is empty.
	Tests whether this ordered collection is empty. virtual ooBoolean isEmpty() const;
Returns	Tests whether this ordered collection is empty. virtual ooBoolean isEmpty() const; oocTrue if this ordered collection has no elements; otherwise, oocFalse.
Returns iterator	Tests whether this ordered collection is empty. virtual ooBoolean isEmpty() const; oocTrue if this ordered collection has no elements; otherwise, oocFalse.
Returns <b>iterator</b>	Tests whether this ordered collection is empty. <pre>virtual ooBoolean isEmpty() const; oocTrue if this ordered collection has no elements; otherwise, oocFalse.</pre> Initializes a scalable-collection iterator to find the elements of this ordered collection.
Returns <b>iterator</b>	Tests whether this ordered collection is empty. <pre>virtual ooBoolean isEmpty() const; oocTrue if this ordered collection has no elements; otherwise, oocFalse. Initializes a scalable-collection iterator to find the elements of this ordered collection. virtual ooCollectionIterator *iterator() const;</pre>
Returns <b>iterator</b> Returns	Tests whether this ordered collection is empty. virtual ooBoolean isEmpty() const; oocTrue if this ordered collection has no elements; otherwise, oocFalse. Initializes a scalable-collection iterator to find the elements of this ordered collection. virtual ooCollectionIterator *iterator() const; A pointer to a scalable-collection iterator for finding the elements of this ordered collection; the caller is responsible for deleting the iterator when it is no longer needed.
Returns <b>iterator</b> Returns Discussion	<pre>Tests whether this ordered collection is empty. virtual ooBoolean isEmpty() const; oocTrue if this ordered collection has no elements; otherwise, oocFalse.</pre> Initializes a scalable-collection iterator to find the elements of this ordered collection. virtual ooCollectionIterator *iterator() const; A pointer to a scalable-collection iterator for finding the elements of this ordered collection; the caller is responsible for deleting the iterator when it is no longer needed. The returned iterator finds the elements as ordered in the collection.

last	
	Finds the last element (or key) in this ordered collection.
	<pre>ooRef(ooObj) last() const;</pre>
Returns	If the elements of this ordered collection are persistent objects, an object reference to the last element; if the elements are key-value pairs, an object reference to the key of the last element.
See also	<u>first</u> get

## lastIndexOf

Searches this ordered collection for the last element (or key) that is equal to the specified object.

	<pre>1. int lastIndexOf(     const ooHandle(ooObj) &amp;objH) const;</pre>
	<pre>2. int lastIndexOf(     const ooHandle(ooObj) &amp;objH,     const int index) const;</pre>
Parameters	objH
	Handle to the object whose index is to be found.
	index
	The zero-based index at which search should start.
Returns	The index of the last element in this ordered collection, at or before the starting position, that is equal to (or whose key is equal to) the specified object. If no such element is found, this member function returns -1.
Discussion	This member function searches backward in the collection. Variant 1 starts searching at the last element; variant 2 starts searching at the element whose index is <i>index</i> .
	If elements of this ordered collection are persistent objects, this member function compares each element with the specified object; if elements are key-value pairs, this member function compares the key of each element with the specified object.
	Search stops when a matching element is found; if more than one element matches, this member function finds the one closest to the end of the collection (but at or before the starting index).
See also	indexOf

### refresh

Refreshes each container used internally by this ordered collection, except for the container in which the collection itself is stored.

virtual ooStatus refresh(ooMode & openMode) const;

Parameters openMode

Intended level of access to each refreshed container:

- Specify occRead to open the container for read. This implicitly requests a read lock on the container.
- Specify occUpdate to open the container for update (read and write). This implicitly requests an update lock on the container.

#### Returns oocSuccess if every container can be refreshed; otherwise, oocError.

Discussion You typically call this member function when you need to refresh your view of an ordered collection that you are reading in an MROW transaction. This member function calls <u>refreshOpen</u> on each container that is used internally by the ordered collection—that is, on the administrator, node, and array containers maintained by the ordered collection. This member function does not refresh the container in which the ordered collection itself is stored, nor does it necessarily refresh the containers that store the collection's elements.

#### remove

Removes the first occurrence of the specified object from this ordered collection.

Parameters	оbјн Handle to the object to be removed.
Returns	oocTrue if an element was removed; otherwise, oocFalse.
Discussion	If the elements of this ordered collection are persistent objects, this member function removes the first element (if any) that is equal to the specified object. If the elements are key-value pairs, this member function removes the element (if any) whose key is the specified object.

### size

Gets the size of this ordered collection.

virtual int size() const;

Returns The number of elements in this ordered collection

Member Functions

# ooCollection Class

Inheritance: ooObj->ooCollection

Handle Class: ooHandle(ooCollection)

```
Object-Reference Class: ooRef(ooCollection)
```

The persistence-capable class ooCollection is the abstract base class for classes that represent *scalable collections* containing persistent objects.

See:

- "Reference Summary" on page 174 for an overview of member functions
- "Reference Index" on page 175 for a list of member functions

To use this class, your application must include the ooCollections.h header file. For UNIX linking information, see *Installation and Platform Notes for UNIX*.

## **About Scalable Collections**

A *collection* is an aggregate that can contain a variable number of elements. A scalable collection can increase in size with minimal performance degradation.

Concrete classes derived from this class represent more specific kinds of collections of persistent objects and collections of key-value pairs. A collection of key-value pairs in which both the key and the value are persistent objects is called an *object map*.

- ooTreeList represents lists of persistent objects.
- ooHashSet represents unordered sets of persistent objects.
- ooTreeSet represents sorted sets of persistent objects.
- ooHashMap represents unordered object maps.
- ooTreeMap represents sorted object maps.

Because the ooCollection class is abstract, you never create instances of it; instead, you work with instances of its concrete derived classes. You should not create your own subclasses of this class.

Like other persistent objects, collections are normally manipulated through handles and object references.

## **Related Classes**

Every collection derived from ooCollection uses a persistent object, called an *administrator*, an instance of a concrete class derived from ooAdmin. A collection's administrator manages the collections used by the collection's interal objects.

An instance of any derived class except <code>ooTreeList</code> uses a persistent object, called a *comparator*; an instance of a concrete class derived from <code>ooCompare</code>.

Objectivity/C++ includes one additional persistence-capable collection class that is not derived from ooCollection. ooMap represents name maps; a *name map* is a collection of key-value pairs in which the key is a string and the value is a persistent object.

For additional information about the difference among the various persistence-capable collection classes, see Chapter 11, "Persistent Collections," in the Objectivity/C++ programmer's guide.

Adding and Removing Elements	add addAll clear remove removeAll removeAllDeleted
Getting Elements	<u>get</u> <u>iterator</u> <u>keyIterator</u> valueIterator
Getting Information	size

## **Reference Summary**

Finding Auxiliary Objects	admin comparator
Testing	<u>contains</u> <u>containsAll</u> <u>isEmpty</u>
Viewing in an MROW Transaction	refresh

# **Reference Index**

add	Adds the specified object to this collection.
<u>addAll</u>	Adds all elements (or keys) in the specified collection to this collection.
<u>admin</u>	Finds the administrator for this collection.
<u>clear</u>	Sets this collection to an empty collection, removing any elements it contains.
comparator	Finds the comparator for this collection.
<u>contains</u>	Tests whether this collection contains the specified object.
<u>containsAll</u>	Tests whether this collection contains all elements (or keys) in the specified collection.
get	Finds the element of this collection that is equal to the specified lookup data, as determined by the comparator for this collection.
isEmpty	Tests whether this collection is empty.
<u>iterator</u>	Initializes a scalable-collection iterator to find the elements of this collection.
<u>keyIterator</u>	Initializes a scalable-collection iterator to find the keys or elements of this collection.
<u>refresh</u>	Refreshes each container used internally by this collection, except for the container in which the collection itself is stored.
remove	Removes the first occurrence of the specified object from this collection.
<u>removeAll</u>	Removes all elements (or keys) of the specified collection from this collection.

removeAllDeleted	Removes from this collection all persistent objects that have been deleted from the federated database.
<u>retainAll</u>	Retains all elements of this collection that are also in the specified collection, removing all other elements.
<u>size</u>	Gets the size of this collection.
valueIterator	Initializes a scalable-collection iterator to find the values or elements of this collection.

# **Member Functions**

## add

	Adds the specified object to this collection.
	<pre>virtual ooBoolean add(     const ooHandle(ooObj) &amp;objH) = 0;</pre>
Parameters	оbjн Handle to the object to be added.
Returns	oocTrue if an element was added; otherwise, oocFalse.
See also	addAll remove
addAll	
	Adds all elements (or keys) in the specified collection to this collection.
	<pre>virtual ooBoolean addAll(     const ooHandle(ooCollection) &amp;collectionH);</pre>
Parameters	<pre>collectionH Handle to the collection whose elements are to be added to this collection. If collectionH is the handle to a collection of key-value pairs and this is a collection of objects, only the keys of the specified collection are added to this collection.</pre>
Returns	oocTrue if any elements were added; otherwise, oocFalse.
See also	add removeAll

## admin

	Finds the administrator for this collection.
	<pre>virtual ooRef(ooAdmin) admin() const = 0;</pre>
Returns	Object reference to the administrator for this collection.
clear	
	Sets this collection to an empty collection, removing any elements it contains.
	<pre>virtual void clear();</pre>

See also <u>remove</u> <u>removeAll</u> <u>retainAll</u>

### comparator

Finds the comparator for this collection.

virtual ooRef(ooCompare) comparator() const = 0;

Returns Object reference to the comparator for this collection, or null if this collection has a default comparator.

### contains

	Tests whether this collection contains the specified object.
	<pre>1. virtual ooBoolean contains(</pre>
	<pre>2. virtual ooBoolean contains(</pre>
Parameters	objH

Handle to the element to be tested for containment in this collection.

#### lookupVal

Pointer to data that identifies the element to be tested for containment in this collection.

Returns oocTrue if this collection contains an element equal to the specified object; otherwise, oocFalse.

Discussion Variant 2 tests whether any element (or key) of this collection is "equal" to the specified lookup data, as determined by the comparator for this collection. It is useful if this collection has an application-defined comparator that can identify an element (or key) based on class-specific data.

See also <u>containsAll</u>

### containsAll

Tests whether this collection contains all elements (or keys) in the specified collection.

Parameters collectionH

Handle to the collection whose elements (or keys) are to be tested for containment in this collection. Two elements (or keys) are equal if they are references for the same persistent object.

- Returns occTrue if this collection contains an element (or key) equal to each element (or key) of the specified collection; otherwise, occFalse.
- Discussion The meaning of this member function depends on whether the elements of the collections being compared are persistent objects or key-value pairs.

Elements of This Collection	Elements of Other Collection	Meaning
Persistent objects	Persistent objects	Tests whether this collection contains all elements of the specified collection
Persistent objects	Key-value pairs	Tests whether this collection contains all keys of the specified collection
Key-value pairs	Persistent objects	Tests whether all elements of the specified collection are keys of this collection
Key-value pairs	Key-value pairs	Tests whether all keys of the specified collection are also keys of this collection

See also <u>contains</u>

get	
	Finds the element of this collection that is equal to the specified lookup data, as determined by the comparator for this collection.
	<pre>virtual ooRef(ooObj) get(const void *lookupVal) const = 0;</pre>
Parameters	<i>lookupVal</i> Pointer to data that identifies the desired element.
Returns	If the elements of this collection are persistent objects, this member function returns an object reference to the element that is equal to <i>lookupVal</i> , or a null object reference if this collection does not contain such an element.
	If the elements are key-value pairs, this member function returns an object reference to the value of the element whose key is equal to <i>lookupVal</i> , or a null object reference if this collection does not contain such an element.
Discussion	This member function is useful if this collection has an application-defined comparator that can identify an element (or key) based on class-specific data.
isEmpty	
	Tests whether this collection is empty.
	<pre>virtual ooBoolean isEmpty() const = 0;</pre>
Returns	oocTrue if this collection has no elements; otherwise, oocFalse.
iterator	
	Initializes a scalable-collection iterator to find the elements of this collection.
	<pre>virtual ooCollectionIterator *iterator() const = 0;</pre>
Returns	A pointer to a scalable-collection iterator for finding the elements of this collection; the caller is responsible for deleting the iterator when it is no longer needed.
Discussion	If this collection is ordered, the returned iterator finds the elements as ordered in the collection; if this collection is unordered, the iterator finds the elements in an undefined order.
	You must delete the iterator when you have finished using it.
See also	<u>keyIterator</u> <u>valueIterator</u>

## keylterator

	Initializes a scalable-collection iterator to find the keys or elements of this collection.
	virtual ooCollectionIterator *keyIterator() const;
Returns	A pointer to a scalable-collection iterator for finding the keys or elements of this collection; the caller is responsible for deleting the iterator when it is no longer needed.
Discussion	If the elements of this collection are key-value pairs, the returned iterator finds the keys of this collection; if the elements are persistent objects, the iterator finds this collection's elements.
	If this collection is ordered, the returned iterator finds the keys or elements as ordered in the collection; if the collection is unordered, the iterator finds the keys or elements in an undefined order.
	You must delete the iterator when you have finished using it.
See also	<u>iterator</u> <u>valueIterator</u>
refresh	
	Refreshes each container used internally by this collection, except for the container in which the collection itself is stored.
	<pre>virtual ooStatus refresh(ooMode &amp;openMode) const = 0;</pre>
Parameters	<ul> <li>openMode</li> <li>Intended level of access to each refreshed container:</li> <li>Specify occRead to open the container for read. This implicitly requests a read lock on the container.</li> <li>Specify occUpdate to open the container for update (read and write). This implicitly requests an update lock on the container.</li> </ul>
Returns	oocSuccess if every container can be refreshed; otherwise, oocError.
Discussion	You typically call this member function when you need to refresh your view of a collection that you are reading in an MROW transaction. This member function calls <u>refreshOpen</u> on each container that is used internally by the collection—that is, on the administrator, node, and array containers maintained by an ordered collection, or the administrator and hash-bucket containers maintained by an unordered collection. This member function does not refresh
the container in which the collection itself is stored, nor does it necessarily refresh the containers that store the collection's elements.

#### remove

Removes the first occurrence of the specified object from this collection.

Parameters objH

Handle to the object to be removed.

- Returns oocTrue if an element was removed; otherwise, oocFalse.
- Discussion If the elements of this collection are persistent objects, this member function removes the first element (if any) that is equal to the specified object. If the elements are key-value pairs, this member function removes the element (if any) whose key is the specified object.
- See also <u>add</u> <u>clear</u> <u>removeAll</u> retainAll

### removeAll

Removes all elements (or keys) of the specified collection from this collection.

Parameters collectionH Handle to the collection whose elements are to be removed from this collection.

#### Returns oocTrue if any elements were removed; otherwise, oocFalse.

Discussion Which elements are removed depends on whether elements of the two collections are persistent objects or key-value pairs.

Elements of This Collection	Elements of Other Collection	Removes From This Collection
Persistent objects	Persistent objects	All elements that are also elements of the specified collection
Persistent objects	Key-value pairs	All elements that are keys of the specified collection
Key-value pairs	Persistent objects	All elements whose keys are elements of the specified collection
Key-value pairs	Key-value pairs	All elements whose keys are also keys of the specified collection

See also

<u>addAll</u> <u>clear</u> <u>remove</u> <u>removeAllDeleted</u> <u>retainAll</u>

#### removeAllDeleted

Removes from this collection all persistent objects that have been deleted from the federated database.

virtual void removeAllDeleted();

Discussion You can call this member function to restore this collection's referential integrity.

See also

<u>remove</u> <u>removeAll</u> <u>retainAll</u>

clear

### retainAll

Retains all elements of this collection that are also in the specified collection, removing all other elements.

# Parameters collectionH Handle to the collection whose elements are to be retained in this collection. Returns oocTrue if any elements were removed; otherwise, oocFalse.

Discussion Which elements are removed depends on whether elements of the t

Which elements are removed depends on whether elements of the two collections are persistent objects or key-value pairs.

Elements of This Collection	Elements of Other Collection	Removes From This Collection
Persistent objects	Persistent objects	All elements that are not also elements of the specified collection
Persistent objects	Key-value pairs	All elements that are not keys of the specified collection
Key-value pairs	Persistent objects	All elements whose keys are not elements of the specified collection
Key-value pairs	Key-value pairs	All elements whose keys are not also keys of the specified collection

See also

<u>clear</u> <u>remove</u> removeAll

### size

Gets the size of this collection.

virtual int size() const = 0;

Returns The number of elements in this collection.

## valuelterator

Initializes a scalable-collection iterator to find the values or elements of this collection.

virtual ooCollectionIterator \*valueIterator() const;

Returns A pointer to a scalable-collection iterator for finding the values or elements of this collection; the caller is responsible for deleting the iterator when it is no longer needed.

Discussion If the elements of this collection are key-value pairs, the returned iterator finds the values of the collection; if the elements are persistent objects, the iterator finds the collection's elements.

If this collection is ordered, the iterator finds the values or elements as ordered in the collection; if the collection is unordered, the iterator finds the values or elements in an undefined order.

You must delete the iterator when you have finished using it.

See also <u>iterator</u> keyIterator

# ooCollectionIterator Class

Inheritance: ooCollectionIterator

The non-persistence-capable class ooCollectionIterator is the abstract base class for *scalable-collection iterators*, which step through the objects in scalable persistent collections.

See:

- "Reference Summary" on page 187 for an overview of member functions
- "Reference Index" on page 188 for a list of member functions

To use this class, your application must include the ooCollections.h header file. For UNIX linking information, see *Installation and Platform Notes for UNIX*.

# About Scalable-Collection Iterators

A scalable-collection iterator finds elements, keys, or values in a particular scalable collection. This group of objects is called the iterator's *iteration set*.

A scalable-collection iterator has a *current index*, which gives its zero-based position within the iteration set; the *current element* is the element of the iteration set at the scalable-collection iterator's current index. When the iterator is created, it is positioned before the first element. That is, its index is -1 and it has no current element.

# Working With a Scalable-Collection Iterator

Because the ooCollectionIterator class is abstract, you never create instances of it; instead, you work with instances of its concrete derived classes. You should not create your own subclasses of this class.

## **Obtaining a Scalable-Collection Iterator**

During a transaction, you obtain a pointer to a scalable-collection iterator by calling a member function of a collection.

- Call the iterator member function of a collection of objects (list, unordered set, or sorted set) to obtain a scalable-collection iterator that finds the elements of that collection. If the collection is ordered, the iterator finds elements as ordered by the collection; if the collection is unordered, the iterator finds the elements in an undefined order.
- Call the keyIterator member function of a collection of key-value pairs (unordered object map or sorted object map) to obtain a scalable-collection iterator that finds the keys of that collection. If the collection is a sorted object map, the iterator finds the keys in their sorted order; if it is an unordered object map, the iterator finds the keys in an undefined order.
- Call the valueIterator member function of a collection of key-value pairs to obtain a scalable-collection iterator that finds the values of that collection. If the collection is a sorted object map, the iterator finds the value in the order in which their keys are sorted; if it is an unordered object map, the iterator finds the keys in an undefined order. In either case, the iterator returned by valueIterator finds the elements of the collection in the same order as does the iterator returned by keyIterator.

### Using a Scalable-Collection Iterator

After obtaining an initialized scalable-collection iterator, you can use it in a loop that processes each element of the iteration set in turn.

To step through the iteration set from beginning to end, call the <u>hasNext</u> member function for loop control to test whether additional elements remain. Within the loop, you make successive calls to the iterator's <u>next</u> member function to find each element. As you step through the iteration set from beginning to end, the current index increases until the scalable-collection iterator is positioned after the last element. At that point, the current index is the size of the iteration set and the scalable-collection iterator has no current element.

You can reposition the iterator within the iteration set. To position the iterator at a particular index, you call gotoIndex; to position the iterator at a particular object, you call goto. (The latter is useful when iterating through a sorted collection.)

After iterating forward (or repositioning the iterator), you can reverse the direction of iteration. To iterate backward from the current index, you call the <u>hasPrevious</u> member function for loop control to test whether additional elements remain. Within the loop, you make successive calls to the iterator's <u>previous</u> member function to find each element. As you step backward through the iteration set, the current index decreases until the scalable-collection iterator

is positioned before the first element. At that point, the current index is back to -1 and the iterator has no current element.

Various member functions, such as next and previous, find persistent objects in the iterator's corresponding collection. These functions return type-independent object references or handles, which you can then cast to the appropriate type.

### **Modifying the Collection**

Methods of the scalable-collection iterator allow you to modify the corresponding collection. If the iterator is currently positioned at an element of the iteration set (that is, not before the first element or after the last element), the <u>remove</u> member function removes the current element of the iteration set from the corresponding collection; the <u>set</u> member function replaces the current element with a specified object.

**NOTE** You should not continue to use a scalable-collection iterator after you add elements to the collection from which you obtained the iterator. Instead, you should delete the old iterator and get a new scalable-collection iterator after you add elements to the collection.

Testing	<u>hasNext</u> <u>hasPrevious</u>
Finding Elements	<u>next</u> <u>current</u> <u>previous</u> <u>goToIndex</u>
Getting Indexes	<u>currentIndex</u> <u>nextIndex</u> previousIndex
Repositioning the Iterator	goTo
Modifying the Collection	<u>remove</u> <u>set</u>

# **Reference Summary**

Moving an Element	moveCurrentTo
Getting Information	collection currentValue

# **Reference Index**

<u>collection</u>	Finds this scalable-collection iterator's corresponding collection.
<u>current</u>	Finds the current object in this scalable-collection iterator's iteration set.
<u>currentValue</u>	Finds the value object paired with the current (key) object in this scalable-collection iterator's iteration set.
<u>currentIndex</u>	Gets the current index of this scalable-collection iterator.
<u>goTo</u>	Positions this scalable-collection iterator at the first occurrence of the specified object.
goToIndex	Positions this scalable-collection iterator at the specified index and finds the object at that position in the iteration set.
<u>hasNext</u>	Tests whether this scalable-collection iterator has any elements after its current location.
<u>hasPrevious</u>	Tests whether this scalable-collection iterator has any elements before its current location.
<u>moveCurrentTo</u>	Moves the current element of this scalable-collection iterator, clustering it with the specified object.
<u>next</u>	Finds the next object in this scalable-collection iterator's iteration set.
<u>nextIndex</u>	Gets the index of this scalable-collection iterator's next element.
previous	Finds the previous object in this scalable-collection iterator's iteration set.
<u>previousIndex</u>	Gets the index of this scalable-collection iterator's previous element.
remove	Removes the current element from this scalable-collection iterator's corresponding collection.
<u>set</u>	Modifies this scalable-collection iterator's corresponding collection, replacing the current element with the specified object.

# **Member Functions**

## collection

	Finds this scalable-collection iterator's corresponding collection.
	<pre>virtual ooHandle(ooCollection) collection() = 0;</pre>
Returns	Handle to the collection whose elements, keys, or values this scalable-collection iterator finds.
current	
	Finds the current object in this scalable-collection iterator's iteration set.
	ooRef(ooObj) current() const;
Returns	Object reference to this scalable-collection iterator's current element, or a null object reference if there is no current element (because the iterator is positioned before the first element or after the last element).

### currentValue

Finds the value object paired with the current (key) object in this scalable-collection iterator's iteration set.

virtual ooRef(ooObj) currentValue() const;

- Returns Object reference to the value of this scalable-collection iterator's current element, or a null object reference if there is no current element (because the iterator is positioned before the first element or after the last element).
- Discussion You typically use this member function only if you are iterating over the keys of an unordered object map or a sorted object map, and you want to find the value that is paired with the current key. This function is equivalent to the current member function if you are iterating over a set or list, whose elements are not key-value pairs.

## currentIndex

Gets the current index of this scalable-collection iterator.

virtual int currentIndex() = 0;

Returns The index in the iteration set at which this iterator is currently positioned.

# goTo

Positions this scalable-collection iterator at the first occurrence of the specified object.

virtual ooBoolean goTo(const ooHandle(ooObj) &objH) = 0; 1. 2. virtual ooBoolean goTo(const void \*lookupVal) = 0; Parameters objH Handle to the object to find in the iteration set. lookupVal Pointer to data that identifies the object to find in the iteration set. Returns oocTrue if the object was found, otherwise, oocFalse. Discussion You typically use this member function only if you are iterating over the elements of a sorted set or the keys of a sorted object map. In those cases, you are assured that the iteration set contains at most one occurrence of the specified object; in addition, the position of that object is meaningful because the elements of the iteration set are sorted. For example, if a sorted set has a comparator that sorts objects by a name data member, you might want to iterate starting with the objects whose names begin with the letter C. To do so, you could create an object with the name "C", use that object as the parameter to  $g_{OTO}$ , then iterate from the new position. Variant 2 finds the element that is equal to the specified lookup data, as determined by the comparator for the collection. It is useful if the collection has an application-defined comparator that can identify an element (or key) based on class-specific data. Because a list has no comparator, variant 2 is not relevant when iterating through the elements of a list. If the specified object is found, this scalable-collection iterator is positioned at that object (as if the object had just been returned by next). To iterate starting with the specified object, you need to call current followed by successive calls to next. If the specified object is not found, this scalable-collection iterator is positioned at the index just past where that object belongs in the sorted collection. In the preceding example, if the set contained two successive elements with names "Byrnes" and "Cabbot", the current index would be set to the index of the element named "Cabbot.".

# goToIndex

Positions this scalable-collection iterator at the specified index and finds the object at that position in the iteration set.

ooRef(ooObj) goToIndex(int index);

Parameters index Zero-based index at which to position this scalable-collection iterator. Returns Object reference to the object at the specified index in the iteration set, or a null object reference if *index* is out of bounds. Discussion You typically use this member function only if this scalable-collection iterator is initialized to find objects in an ordered collection; indexes are not meaningful when iterating through an unordered collection. If the specified index is valid, it becomes this iterator's new current index. If the specified index is less than -1, this iterator is positioned before the first element. If the specified index is too large, this iterator is positioned after the last element. This member function can be inefficient because it requires traversing the collection to find the indicated index. It locks all the collection's B-tree nodes up to and including the one containing the specified index. Whenever possible, you should use goto instead of this member function.

## hasNext

	Tests whether this scalable-collection iterator has any elements after its current location.
	<pre>virtual ooBoolean hasNext() const = 0;</pre>
Returns	oocTrue if the iteration set has at least one element after the current location; otherwise, oocFalse.
See also	<u>next</u> hasPrevious

# hasPrevious

Tests whether this scalable-collection iterator has any elements before its current location.

virtual ooBoolean hasPrevious();

Returns occTrue if the iteration set has at least one element before the current location; otherwise, occFalse.

See also	<u>previous</u>	
	<u>hasNext</u>	

#### moveCurrentTo

Moves the current element of this scalable-collection iterator, clustering it with the specified object.

virtual ooStatus moveCurrentTo(const ooHandle(ooObj) & objH) = 0; Parameters objH A handle to the object with which to cluster the current element; must be a database, a persistent container, or a persistent basic object. If the clustering object is a database, the current element is moved to the default container of that database. If the clustering object is a persistent container, the current element is moved to that container. If the clustering object is a persistent basic object, the current element is moved to the container in which that object is stored. If possible, it will be stored on the same page as the clustering object or on a page close to the clustering object. oocSuccess if successful; otherwise, oocError. Returns The current element must be a basic object, not a container. This member Discussion function changes the object identifier of the current element. If the collection has a default comparator, this member function first removes the current element from the collection; then it moves that object and adds it back to the collection. You should not call this member function in either of the following conditions: The current element is an element, key, or value of another persistent collection. The collection has an application-specific comparator that uses the object identifiers of the elements or keys.

#### next

Finds the next object in this scalable-collection iterator's iteration set.

ooRef(ooObj) next();

Returns Object reference to the next element in the iteration set, or a null object reference if this scalable-collection iterator is positioned after the last element in the iteration set.

- Discussion If a next element exists, this member function returns a type-independent object reference to that element, which you must then cast to the appropriate type. It also increments this scalable-collection iterator's index.
- See also <u>hasNext</u> <u>nextIndex</u> <u>previous</u>

## nextIndex

	Gets the index of this scalable-collection iterator's next element.
	virtual int nextIndex();
Returns	The index of the element that would be returned by the subsequent call to ${\tt next}.$
Discussion	The returned index is not necessarily a valid index for the iteration set. If this scalable-collection iterator is currently positioned at the last element of the iteration set, this member function returns the number of elements in the iteration set.
See also	<u>next</u> <u>hasNext</u> previousIndex
previous	
	Finds the previous object in this scalable-collection iterator's iteration set.
	ooRef(ooObj) previous();
Returns	Object reference to the previous element in the iteration set, or a null object reference if this scalable-collection iterator is positioned before the first element in the iteration set.
Discussion	If a previous element exists, this member function returns a type-independent object reference to that element, which you must then cast to the appropriate type. It also decrements this scalable-collection iterator's index.
See also	hasPrevious

next

<u>previousIndex</u>

#### previousIndex

Gets the index of this scalable-collection iterator's previous element.

virtual int previousIndex();

- Returns The index of the element that would be returned by the subsequent call to previous.
- Discussion The returned index is not necessarily a valid index for the iteration set. If this scalable-collection iterator is currently positioned at the first element of the iteration set, this member function returns -1.

See also <u>previous</u> <u>hasPrevious</u> nextIndex

#### remove

Removes the current element from this scalable-collection iterator's corresponding collection.

virtual ooStatus remove() = 0;

#### Returns oocSuccess if successful; otherwise, oocError.

Discussion If this scalable-collection iterator is positioned before the first element or after the last element of the iteration set, this member function does nothing.

If this scalable-collection iterator is positioned at an element of the iteration set, the operation of this member function is as follows:

- If iterating through the elements of a collection of objects, remove the current element from the collection and position the iterator at the previous element. (If you removed the first element, the iterator is then positioned before the new first element.)
- If iterating through the keys of an object map, remove the element of the object map whose key is the current element of the iteration set. Then position the iterator at the previous element of the iteration set.
- If iterating through the values of an object map, set the value to null in the key-value pair whose value is the current element of the iteration set. Leave this iterator positioned at this same element.

set	
	Modifies this scalable-collection iterator's corresponding collection, replacing the current element with the specified object.
	<pre>virtual ooStatus set(const ooHandle(ooObj) &amp;objH);</pre>
Parameters	оbјн Handle to the object that is to replace the current element.
Returns	oocSuccess if successful; otherwise, oocError.
Discussion	If this scalable-collection iterator is positioned before the first element or after the last element of the iteration set, this member function does nothing.
	If this scalable-collection iterator is positioned at an element of the iteration set, the operation of this member function is as follows.
	■ If iterating through the elements of a list, replace the current element of the collection with the specified object.
	<ul> <li>If iterating through the values of an object map, set the value to the specified object in the key-value pair whose value is the current element of the iteration set.</li> </ul>
	<ul> <li>If iterating through the elements of a set, the set operation is not supported; any call to set returns oocError.</li> </ul>
	<ul> <li>If iterating through the keys of an object map, the set operation is not supported; any call to set returns occError.</li> </ul>

# ooCompare Class

Inheritance: ooObj->ooCompare

Handle Class: ooHandle(ooCompare)

**Object-Reference Class:** ooRef(ooCompare)

The persistence-capable class ooCompare is the abstract base class for all *comparator* classes.

See:

"Reference Index" on page 199 for list of member functions

To use this class, your application must include the ooCollections.h header file. For UNIX linking information, see *Installation and Platform Notes for UNIX*.

# **About Comparators**

A comparator is an object of a concrete derived class of ooCompare. It provides a comparison function for ordering elements of sorted collections and a hashing function for computing the hash codes for elements of unordered collections.

Every sorted or unordered collection uses the default comparator unless you assign an application-defined comparator to it; see "Working With a Comparator" on page 198. All application-defined comparators must be classes derived from ooCompare.

#### **Comparators for Sorted Collections**

Every sorted set and every sorted object map has a comparator. The comparator of a sorted collection defines a total ordering used by the underlying B-tree. Its <u>compare</u> member function is used to compare two persistent objects and indicate

their relative position in the total ordering. A sorted collection with a default comparator sorts persistent objects by their object identifiers (OIDs).

If elements of the sorted collection are persistent objects, the elements themselves are compared; if elements are key-value pairs, the elements' keys are compared.

For additional information about defining a comparate class for sorted collections, see "Comparator Class for Sorted Collections" on page 255 in the Objectivity/C++ programmer's guide.

## **Comparators for Unordered Collections**

Every unordered set and every unordered object map has a comparator. The comparator of a scalable unordered collection supplies the hash function used by the underlying extendible hash table. Its <u>hash</u> member function computes the hash value for a persistent object; its <u>compare</u> member function tests two persistent objects for equality. An unordered collection with a default comparator computes hash values for persistent objects from their OIDs and compares objects for equality by comparing their OIDs.

If elements of the unordered collection are persistent objects, hash values are computed from the elements themselves and elements are compared for equality. If elements are key-value pairs, hash values are computed from the elements' keys and the keys are compared for equality.

For additional information about defining a comparate class for sorted collections, see "Comparator Class for Unordered Collections" on page 260 in the Objectivity/C++ programmer's guide.

### **Unique Identification of Collection Elements**

An application-defined comparator class can optionally provide the ability to identify an element (or key) based on class-specific data. For additional information, see "Supporting Content-Based Lookup in a Sorted Collection" on page 258 and "Supporting Content-Based Lookup in an Unordered Collection" on page 264 in the Objectivity/C++ programmer's guide.

# Working With a Comparator

Because the ooCompare class is abstract, you never create instances of it; instead, you work with instances of its concrete derived classes.

You may define your own comparator class with custom comparator and hashing functions. All application-defined comparator classes must derive from the ooCompare class.

If your application uses an application-defined comparator class, you create an instance of that class with a call to the new operator. As is the case for any basic object, you specify whether a comparator is to be transient or persistent when you create it; comparators *must be persistent*. The clustering directive in new operator specifies where in the federated database to store the new comparator.

A comparator is locked whenever you access its corresponding collection. To avoid locking conflicts, you typically cluster the comparator in a separate container. If the comparator is stored in the same container as the collection, applications may fail to get the necessary read lock on the comparator when another process is updating the collection.

After creating a comparator, you assign it to a collection by passing a handle to the comparator as a parameter to the constructor that creates the collection. Typically, an application creates a comparator for each collection; you should not assign the same comparator to more than one collection.

**NOTE** The persistent data for a persistent collection includes an object reference to its comparator. Your application should not explicitly save object references to any comparator. For example, you should not add a comparator to a collection or save an object reference to a comparator in a persistent data member of any persistent object.

Like other persistent objects, comparators are normally manipulated through handles or object references.

# **Reference Index**

This class overloads <u>operator new</u> and <u>operator delete</u>, which behave as described for the <u>ooObj</u> class (page 431).

<u>compare</u>	Compares two persistent objects.
hash	Computes the hash value for a persistent object in a persistent collection that uses this comparator.

# **Member Functions**

#### compare

#### Compares two persistent objects.

	<pre>1. virtual int compare( const ooHandle(ooObj) &amp;obj1H, const ooHandle(ooObj) &amp;obj2H) const;</pre>
	<pre>2. virtual int compare( const ooHandle(ooObj) &amp;obj1H, const void *&amp;lookupVal) const;</pre>
Parameters	<i>obj1H</i> Handle to the first of the two persistent objects to be compared. This object is an element (or key) of a persistent collection that uses this comparator.
	obj2H
	Handle to the second of the two persistent objects to be compared.
	<pre>lookupVal Pointer to data that identifies an element (or key) of the persistent collection containing obj1H; the identified element (or key) is the second of the two persistent objects to be compared.</pre>
Returns	A negative integer if the first object is less than (sorts before) the second; zero if the two objects are equal; a positive integer if the first object is greater than (sorts after) the second.
Discussion	Variant 1 compares the two specified persistent objects; the default implementation of variant 1 compares their OIDs.
	Variant 2 is called by member functions that accept arbitrary data to identify an element. The default implementation of variant 2 casts its void *& parameter to a handle of type <code>ooHandle(ooObj)</code> and compares the OIDs of the two objects.
hash	
	Computes the hash value for a persistent object in a persistent collection that uses this comparator.
	<pre>1. virtual int hash(</pre>
	2. virtual int hash(

const void \*&lookupVal) const;

objH
Handle to the persistent object from which to calculate the hash value. This object is an element (or key) of a persistent collection that uses this comparator.
<i>lookupVal</i> Pointer to data that identifies a persistent object whose hash value is to be computed.
The hash value for the specified object.
Variant 1 computes the hash value for the specified persistent object. The default implementation of this variant computes the hash value from the object's OID.
Variant 2 is called by member functions that accept arbitrary data to identify an element. The specified data identifies the element whose hash value is to be computed. The default implementation of variant 1 casts its void *& parameter to a handle of type ooHandle(ooObj) and computes the hash value from the object's OID.

# ooContext Class

#### Inheritance: ooContext

The non-persistence-capable class ooContext represents an *Objectivity context*—the operating context that enables a thread to execute a series of transactions with an Objectivity/DB federated database.

See:

- "Reference Summary" on page 204 for an overview of member functions
- "Reference Index" on page 204 for a list of member functions

# About Objectivity Contexts

An Objectivity context is a complete set of Objectivity/DB-managed data and memory resources, including an Objectivity/DB cache, the current values of Objectivity/C++ global variables, current error and message handlers, and so on. Every Objectivity/DB application has at least one Objectivity context, which belongs to the main thread (the thread that starts implicitly when you start the application). A multithreaded application typically provides an Objectivity context for each additional thread that executes Objectivity/DB transactions, although it is possible for the same Objectivity context to be reused by several threads in turn. Objectivity contexts are the only kind of transient Objectivity/DB object that you can pass between threads.

The ooContext class provides a public constructor and destructor for creating and destroying Objectivity contexts. Instances of the class are also created and destroyed implicitly by the <u>ooInitThread</u> and <u>ooTermThread</u> global functions, respectively.

Every thread that is initialized for Objectivity/DB has a *current Objectivity context*, which is the context that Objectivity/DB uses when the thread executes database operations. The <code>ooInitThread</code> global function sets a thread's current Objectivity context for the first time; the <code>ooContext</code> class provides an interface

for changing it. An uninitialized thread has no current Objectivity context associated with it.

Objectivity contexts are typically referenced through pointers. A null ooContext pointer is called a *null Objectivity context*. A thread's current Objectivity context may be null or nonnull; however, a thread with a null context cannot execute Objectivity/DB operations other than context operations and destructors.

For more information about Objectivity contexts, see Chapter 3, "Objectivity/DB Initialization," and Chapter 5, "Multithreaded Objectivity/C++ Applications," in the Objectivity/C++ programmer's guide.

# **Reference Summary**

Creating and Destroying an Objectivity Context	<u>ooContext</u> <u>~ooContext</u>
Accessing an Objectivity Context	<u>current</u> <u>setCurrent</u> <u>setCurrentShared</u>

# **Reference Index**

ooContext	Constructs a new Objectivity context whose Objectivity/DB cache has the specified initial and maximum numbers of buffer pages.
~ooContext	Destructor for the class ooContext.
current	Returns a pointer to the current Objectivity context for the executing thread; returns 0 if this thread has a null Objectivity context.
setCurrent	Sets the current Objectivity context for the executing thread.
setCurrentShared	For experienced users only. Sets the current Objectivity context for the executing thread, allowing multiple threads to share the same Objectivity context.

# **Constructors and Destructors**

## ooContext

Constructs a new Objectivity context whose Objectivity/DB cache has the specified initial and maximum numbers of buffer pages.

```
ooContext(uint32 nPages = 200, uint32 nMaxPages = 500);
```

Parameters npages

Initial number of buffer pages to be allocated for each buffer pool in the Objectivity/DB cache.

#### nMaxPages

Maximum number of buffer pages that can be allocated for each buffer pool in the Objectivity/DB cache. This number is limited by the amount of available swap space. Specifying 0 allows the cache to grow as needed, up to the amount of available swap space.

Discussion You use the <code>ooInit</code> global function to specify nondefault cache sizes for the Objectivity context in the main thread; you use the <code>ooContext</code> constructor to specify nondefault cache sizes for any additional Objectivity context. See <code>ooInit</code> for more information about initial and maximum Objectivity/DB cache sizes.

### ~ooContext

Destructor for the class ooContext.

~ooContext();

Discussion An error is signaled if you call this destructor for an Objectivity context that is the current context for any thread.

# **Member Functions**

#### current

Returns a pointer to the current Objectivity context for the executing thread; returns 0 if this thread has a null Objectivity context.

```
static ooContext *current();
```

### setCurrent

Sets the current Objectivity context for the executing thread.

static void setCurrent(ooContext \*context);

 Parameters
 context

 Pointer to the Objectivity context to be set. You can specify 0 to set the thread's current Objectivity context to the null context.

 Discussion
 An error is signaled if the specified Objectivity context is the current context for another thread.

A thread may not invoke Objectivity/DB operations (other than context operations and destructors) until a nonnull context is set.

## setCurrentShared

For experienced users only. Sets the current Objectivity context for the executing<br/>thread, allowing multiple threads to share the same Objectivity context.<br/>static void setCurrentShared(ooContext \*context);Parameterscontext<br/>Pointer to the Objectivity context to be set.DiscussionThis member function is the same as setCurrent, except that the specified<br/>Objectivity context may be the current context for another thread.WARNINGAllowing multiple threads to share the same Objectivity context is a dangerous<br/>operation. To prevent data corruption, you must guarantee that no two threads<br/>execute Objectivity context at the same time. For example, you can use mutexes<br/>to achieve this.

# ooContObj Class

Inheritance: ooObj->ooContObj

Handle Class: <u>ooHandle(ooContObj)</u>

Object-Reference Class: <a href="mailto:ooRef(ooContObj">ooRef(ooContObj</a>)

The persistence-capable class <code>ooContObj</code> represents a standard Objectivity/DB *container*, and serves as the base class for all container classes. Together, the <code>ooContObj</code> class and its corresponding handle and object-reference classes define the behavior of containers.

See:

- "Reference Summary" on page 211 for an overview of ooContObj member functions
- "Reference Index" on page 211 for a list of ooContObj member functions

For operations performed through a container handle or object reference, see:

■ "Reference Summary" on page 515

# **About Containers**

A container is the third highest level in the Objectivity/DB storage hierarchy; every database contains one or more containers, and every container can contain one or more basic objects.

Containers serve a number of purposes. They are used:

 To group basic objects. Basic objects within a container are physically clustered together in memory and on disk, so access to basic objects in a single container is very efficient.

- As the unit of locking. When a basic object is locked, its container and all other objects in the container are also locked. This enables the lock server to manage relatively few container-level locks rather than potentially millions or billions of object-level locks.
- Optionally, to maintain application-specific data.

Containers function both as *storage objects* (because basic objects can be clustered in them) and as *persistent objects* (for example, because they can be referenced in attributes, elements of a persistent collection, scope-named, and so on).

A container is physically maintained within a database file, and may optionally have a system name in addition to any number of scope names; a container's system name cannot be changed.

(*FTO*) Every container is controlled by an autonomous partition—usually the partition in which the container's database resides. If an application moves the control of a container to a different partition, the container is physically moved to the system database file in that partition.

#### **Kinds of Container**

An application normally creates *standard containers* to control the clustering and locking of basic objects. Standard containers are created as instances of class ooContObj.

An application can also derive its own persistence-capable container classes from ooContObj. Instances of such classes, called *application-defined containers*, can have associations and attributes for persistent data). Like any persistence-capable class, an application-defined container class must be defined in a DDL file and its definition must be processed by the DDL processor.

An application that interoperates with Objectivity for Java or Objectivity/Smalltalk applications may choose to create *garbage-collectible containers* instead of standard containers; garbage-collectible containers are created as instances of class <u>ooGCContObj</u>, which is a predefined class derived from ooContObj.

Exactly one *default container* is created automatically by Objectivity/DB for each database; default containers are created as instances of class <u>ooDefaultContObj</u>.

Every container is created as either *hashed* or *nonhashed*. If a container is hashed, the container (and any object in it) can be used as a scope for naming objects. A non-hashed container occupies less storage than a hashed container, but does not support scope naming.

## **Container Structure**

A container consists of a *container object* that manages a particular set of *storage pages* allocated within a particular database file. The storage pages contain the basic objects that are clustered in the container:

- Every small basic object resides entirely within a single storage page, possibly along with other small basic objects.
- Every large basic object spans multiple storage pages; the first of these pages (called a header page) contains certain housekeeping information and the remaining pages contain the object's persistent data.

A storage page is considered a *logical page* if it contains one or more small objects *or* the header information for a large object. Within each container, the logical pages are numbered; the object identifier of a basic object includes the number of the logical page on which the object resides.

The container object has various responsibilities, which include:

- Maintaining a *page map* that records the physical location of each logical page. The page map enables the container object to use object identifiers to quickly locate basic objects in the container.
- Managing the allocation of pages according to the initial number and growth characteristics you specify when the container is created.
- Maintaining a hash table, if the container is hashed for scope naming.
- Storing persistent data, if the container is an instance of an application-defined container class that defines attributes or associations.

# **Working With Containers**

Standard, application-defined, and garbage-collectible containers are created and deleted from within an application. An application:

Creates a container using the constructor and operator new defined by the appropriate container class—namely, ooContObj, the application-defined class appClass, or ooGCContObj.

The resulting container-class instance represents the container object, which, in turn, manages the container's pages in the database.

Deletes a container by calling the ooDelete global function.

As is the case for any persistent object, you specify whether a container is to be transient or persistent when you create it. Containers should normally be persistent; a clustering directive on operator new specifies where to locate a new persistent container in the federated database. A transient container is

simply the container object without any allocated storage pages; it can have application-specific persistent data, but basic objects cannot be clustered in it.

To work with a new container, an application *must* assign the result of operator new to a handle. For example, a new standard container is normally assigned to an instance of ooHandle(ooContObj); a new application-defined container of class *appClass* is normally assigned to an instance of the corresponding handle class ooHandle(*appClass*).

Similarly, to identify and work with an existing container, the application must open it through an appropriate container handle or object reference; multiple handles and object references can be set to reference the same container. Note that opening a container reads just the container object into memory; pages managed by the container are fetched only when the basic objects on those pages are opened.

The application operates on an open container by:

- Calling various member functions on any of the referencing handles or object references.
- Passing any of the referencing handles or object references to various global functions or member functions of other classes.

The general handle and object-reference classes for containers (ooRefHandle(ooContObj)) provide the primary interface for operating on existing containers, including member functions for opening, locking, getting information, and finding objects from a referenced container. The type-specific handle and object-reference classes ooRefHandle(appClass) inherit these member functions from ooRefHandle(ooContObj), redefining them wherever type-specific behavior or parameters are required.

In addition, every application-defined container class *appClass* has:

- DDL-generated member functions for creating, deleting, and accessing any associations defined by the class.
- An <u>operator new</u> defined by ooContObj.
- The member functions defined by ooObj that apply to containers (see "Reference Summary" on page 211). These include member functions for identifying a container's type, for obtaining a handle to the container, or for opening the container for update.
- Data members and member functions specific to *appClass*.

You can call these members directly from within a member function of *appClass* or indirectly through the indirect member-access operator (->) on an *appClass* handle or object reference.

**NOTE** 000bj also defines member functions for moving, copying, and versioning basic objects. These member functions are *not* available on containers, which cannot be moved, copied, or versioned.

# **Reference Summary**

In the following table, the member functions indicated as *(inherited)* are defined by ooObj and are available for containers; they are documented with the ooObjclass (page 431). They are the *only* ooObj member functions that apply to containers.

Creating and Deleting a Container	<u>ooContObj</u> <u>operator new</u> <u>operator delete</u> (inherited)
Working With the Container	<u>ooGetTypeN</u> (inherited) <u>ooGetTypeName</u> (inherited) <u>ooIsKindOf</u> (inherited) <u>ooThis</u> <u>ooUpdate</u> (inherited) <u>ooValidate</u> (inherited)
ODMG Interface	operator new

# **Reference Index**

ooThis

<u>ooContObj</u>	Default constructor	that constructs a new	standard container.
------------------	---------------------	-----------------------	---------------------

Sets an object reference or handle to reference this container.

<u>operator new</u> Creates a new standard container.

# Constructors

# ooContObj

Default constructor that constructs a new standard container.

ooContObj();

# Operators

#### operator new

Creates a new standard container.

1.	<pre>void *operator new(     size_t);</pre>
2.	<pre>void *operator new(    size_t,    const ooRefHandle(ooObj) &amp;near);</pre>
3.	<pre>void *operator new(    size_t,    const ooObj *near);</pre>
(ODMG)4.	<pre>void *operator new(    size_t,    ooRefHandle(ooObj) near,    const char *type);</pre>
(ODMG)5.	<pre>void *operator new(    size_t,    d_Database *near,    const char *type = 0);</pre>
6.	<pre>void *operator new(    size_t,    const char *contSysName,    const uint32 hash,    const uint32 initPages,    const uint32 percentGrowth);</pre>

```
7.
    void *operator new(
      size t,
      const char * contSysName,
      const uint32 hash,
      const uint32 initPages,
      const uint32 percentGrowth,
      const ooRefHandle(ooObj) &near);
8.
    void *operator new(
      size t,
      const char * contSysName,
      const uint32 hash,
      const uint32 initPages,
      const uint32 percentGrowth,
      const ooObj *near);
```

Parameters

size\_t

Do not specify; this parameter is automatically initialized by the compiler with the size of the class type in bytes.

near

Specifies whether the new container is to be persistent or transient. If the new container is persistent, *near* is the clustering directive that specifies the container's location in the federated database.

To create a transient container, specify 0 or a pointer to a transient object.

To create a persistent container, choose one of the following alternatives as the clustering directive:

- Omit *near* (variants 1 and 6)
- Specify an object reference, handle, or pointer to a database, container, or persistent basic object (variants 2, 3, 4, 7, and 8). *near* may not reference a federated database or an autonomous partition.
- (*ODMG*) Specify a pointer to a d\_Database object that references the federated database (variant 5).

When you create a persistent container:

- If you omit *near*, the new container is created in the most recently opened or created database. An error is signalled if no such database exists.
- If *near* references a database, the new container is created in that database.
- If near references a container or basic object, the new container is created in the database that contains the referenced container or basic object.

 (ODMG) If near specifies a valid d\_Database object, the new container is created in the most recently opened or created database in the federation. If no such database exists, a database called default\_odmg\_db is used; if necessary, this database is created.

type

(*ODMG*) This parameter is ignored. By convention it is the name of the class you are instantiating.

#### contSysName

System name of the new container. Specify 0 to omit the system name. This name must be unique within the database.

hash

Determines whether to create a hashed container. You must create a hashed container if you intend to use it or any object in it as a scope for naming objects, or if you intend to create keyed objects in the container.

- Specify 0 to create a nonhashed container.
- Specify 1 or greater to create a hashed container.

The number you specify is the clustering factor for any keyed objects created in the container. A clustering factor is the number of sequentially keyed objects to be placed onto a page. A clustering factor of 1 maximally distributes keyed objects across pages. A higher number means fewer pages need to be read when finding sequences of keyed objects.

#### initPages

Initial number of logical pages to allocate for the new container. Specify 0 to use the system default value (4 pages for a hashed container, and 2 pages for a nonhashed container). The maximum value is 65535.

percentGrowth
---------------

Amount by which the new container may grow when it needs to accommodate more basic objects, expressed as a percentage of its current size. Specify 0 to use the system default value (10%).

# Returns Memory pointer to the new container. This pointer is null if an error occurs during the creation of the container.

Discussion Variants 1 through 5 each create an unnamed, nonhashed container with an initial size of 2 pages and a growth factor of 10%—as if parameters *contSysName*, *hash*, *initPages*, and *percentGrowth* were each set to 0.

When you create a new persistent container, you must:

 Use operator new in an update transaction. The new container is made permanent on disk when the transaction commits or is checkpointed. If the transaction is aborted, the container is not created. Assign the result of operator new directly to a handle. You can verify the creation of the container by checking whether the handle is null. (If an object reference is desired, you can then assign the handle to an object reference.)
 WARNING Although direct assignment to a pointer or object reference does not raise compile-time or runtime errors, such assignments can eventually cause the Objectivity/DB cache to run out of memory.
 See also ooNewConts global macro

# **Member Functions**

## ooThis

	Sets an object reference or handle to reference this container.	
	1. ooHandle(ooContObj) ooThis() const;	
	<pre>2. ooRef(ooContObj) &amp;ooThis(</pre>	
	<pre>3. ooHandle(ooContObj) &amp;ooThis(</pre>	
Parameters	<i>object</i> Object reference or handle to be set to this object.	
Returns	Object reference or handle to this object.	
Discussion	You normally use ooThis in a member function of an application-defined container class; when such a member function is called on a container of the class, ooThis provides the member function with an object reference or handle to the container. The member function can then perform operations on the container that are available only through an object reference or handle.	
	When called without a <i>container</i> parameter, ooThis allocates a new handle and returns it. Otherwise, ooThis returns the object reference or handle that is passed to it.	
See also	appClass::ooThis	
# ooConvertInObject Class

Inheritance: ooConvertInObject

The non-persistence-capable class <code>ooConvertInObject</code> represents an *unconverted object*—an object that is affected by schema evolution and that is about to be converted to its new shape.

See:

- "Reference Summary" on page 218 for an overview of member functions
- "Reference Index" on page 218 for a list of member functions

## About ooConvertInObject Instances

The ooConvertInObject class enables you to get the values of primitive data members in unconverted objects. You use instances of this class in any *conversion functions* you create to augment the object conversion process. When a conversion function is called during object conversion, two objects are passed to it:

- A representation of an existing, unconverted object (an instance of ooConvertInObject)
- A representation of the object after it has been converted (an instance of ooConvertInOutObject)

You use member functions of <code>ooConvertInObject</code> to get values from the unconverted object's primitive data members. From these pre-conversion values, you can compute data-member values to be set after the object has been converted. You use member functions of <code>ooConvertInOutObject</code> to set data-member values in the converted object.

The ooConvertInObject class provides access to primitive data members, but not to variable-length data members (VArrays), associations, or object references. You can access primitive data members that are defined:

- Directly in the class whose objects are being converted (see the get *Type* member functions).
- In a base class of the class whose objects are being converted (see the getOldBaseClass member function).
- In an embedded class of the class whose objects are being converted (see the getOldDataMember member function).

See Chapter 19, "Object Conversion," of the Objectivity/C++ programmer's guide for an example that uses this class.

Accessing Primitive Data Members	<pre>getFloat32 getFloat64 getInt8 getInt16 getInt32 getUInt64 getUInt16 getUInt32 getUInt32 getUInt64</pre>
Accessing a Base or Embedded Class	getOldBaseClass getOldDataMember

## **Reference Summary**

## **Reference Index**

getFloat32	Gets the value of the specified ${\tt float32}$ data member in an unconverted object.
getFloat64	Gets the value of the specified ${\tt float64}$ data member in an unconverted object.
getInt8	Gets the value of the specified int8 data member in an unconverted object.

getInt16	Gets the value of the specified int16 data member in an unconverted object.
getInt32	Gets the value of the specified int 32 data member in an unconverted object.
getInt64	Gets the value of the specified int64 data member in an unconverted object.
<u>getOldBaseClass</u>	Gets the part of an unconverted object that is inherited from the specified base class.
<u>getOldDataMember</u>	Gets the part of an unconverted object that belongs to the specified embedded object.
getUInt8	Gets the value of the specified uint8 data member in an unconverted object.
getUInt16	Gets the value of the specified uint16 data member in an unconverted object.
getUInt32	Gets the value of the specified uint 32 data member in an unconverted object.
getUInt64	Gets the value of the specified uint 64 data member in an unconverted object.

## **Member Functions**

### getFloat32

Gets the value of the specified float32 data member in an unconverted object.

ooStatus getFloat32( const char \*memberName, float32 &value);

Parameters memberName

Name of the data member whose value is to be obtained.

value

Name of the float 32 variable in which the value is to be returned.

Returns oocSuccess if successful, or oocError if the data-member name is not valid.

### getFloat64

Gets the value of the specified float64 data member in an unconverted object.

```
ooStatus getFloat64(
        const char *memberName,
        float64 &value);
```

Parameters	memberName Name of the data member whose value is to be obtained.
	value Name of the float64 variable in which the value is to be returned.
Returns	oocSuccess if successful, or oocError if the data-member name is not valid.

### getInt8

Gets the value of the specified int8 data member in an unconverted object.

```
ooStatus getInt8(
    const char *memberName,
    int8 &value);
```

#### Parameters memberName

Name of the data member whose value is to be obtained.

value

Name of the int8 variable in which the value is to be returned.

Returns oocSuccess if successful, or oocError if the data-member name is not valid.

### getInt16

Gets the value of the specified int16 data member in an unconverted object.

Parameters	memberName Name of the data member whose value is to be obtained.
	value Name of the int16 variable in which the value is to be returned.
Returns	oocSuccess if successful, or oocError if the data-member name is not valid.

### getInt32

Gets the value of the specified int32 data member in an unconverted object.

Parameters	memberName
	Name of the data member whose value is to be obtained.
	value
	Name of the int32 variable in which the value is to be returned.
Returns	oocSuccess if successful, or oocError if the data-member name is not valid.

### getInt64

Gets the value of the specified int64 data member in an unconverted object.

```
ooStatus getInt64(
    const char *memberName,
    int64 &value);
```

```
Parameters memberName
```

Name of the data member whose value is to be obtained.

value

Name of the int64 variable in which the value is to be returned.

Returns oocSuccess if successful, or oocError if the data-member name is not valid.

## getOldBaseClass

Gets the part of an unconverted object that is inherited from the specified base class.

```
ooStatus getOldBaseClass(
    const char *baseClassName,
    ooConvertInObject &basePartOfExistObj);
```

#### Parameters baseClassName

Name of a base class of the class whose objects are being converted.

basePartOfExistObj

Name of the <code>ooConvertInObject</code> variable in which the inherited part is to be returned.

Returns oocSuccess if successful, or oocError if the path is not valid.

Discussion You use this member function when you need to get the value of an inherited primitive data member of an unconverted object. First, you use this member function to get the part of the object that is inherited from the specified base class. Then, you use the appropriate getType member function on the inherited part (accessed through the basePartOfExistObj variable) to get the value of the desired data member.

### getOldDataMember

Gets the part of an unconverted object that belongs to the specified embedded object.

ooStatus getOldDataMember(
 const char \*memberName,
 ooConvertInObject &embeddedPartOfExistObj);

Parameters memberName

Name of the data member that specifies the desired embedded object.

embeddedPartOfExistObj

Name of the ooConvertInObject variable in which the embedded part is to be returned.

- Returns oocSuccess if successful, or oocError if the memberName is not valid.
- Discussion You use this member function when you need to get the value of a primitive data member of an object that is embedded as a data member of an unconverted object. First, you use this member function to get the desired embedded part of the unconverted object. Then, you use the appropriate getType member function on the embedded part (accessed through the *embeddedPartOfExistObj* variable) to get the value of the desired primitive data member.

### getUInt8

Gets the value of the specified uint8 data member in an unconverted object.

ooStatus getUInt8( const char \*memberName, uint8 &value);

Parameters memberName

Name of the data member whose value is to be obtained.

	value Name of the uint8 variable in which the value is to be returned.
Returns	oocSuccess if successful, or oocError if the data-member name is not valid.
getUInt16	Gets the value of the specified uint 16 data member in an unconverted object

	1	3
	ooStatus getUInt16( const char * <i>memberName</i> , uint16 & <i>value</i> );	
Parameters	memberName Name of the data member whose value is to be obtained.	
	value	
	Name of the uint16 variable in which the value is to be returned.	
Returns	oocSuccess if successful, or oocError if the data-member name is not	valid

### getUInt32

Gets the value of the specified uint 32 data member in an unconverted object.

ooStatus getUInt32( const char \*memberName, uint32 &value);

Parameters	memberName Name of the data member whose value is to be obtained.
	value Name of the uint32 variable in which the value is to be returned.

Returns oocSuccess if successful, or oocError if the data-member name is not valid.

## getUInt64

Gets the value of the specified  ${\tt uint64}$  data member in an unconverted object.

```
ooStatus getUInt64(
    const char *memberName,
    uint64 &value);
```

Parameters memberName

Name of the data member whose value is to be obtained.

value

Name of the uint64 variable in which the value is to be returned.

Returns oocSuccess if successful, or oocError if the data-member name is not valid.

# ooConvertInOutObject Class

#### Inheritance: ooConvertInObject->ooConvertInOutObject

The non-persistence-capable class <code>ooConvertInOutObject</code> represents a *converted object*—an object that is affected by schema evolution and that has been converted to its new shape.

See:

- "Reference Summary" on page 226 for an overview of member functions
- "Reference Index" on page 226 for a list of member functions

## About ooConvertInOutObject Instances

The ooConvertInOutObject class enables you to set the values of data members in objects that have been converted as a result of schema evolution. You use instances of this class in any *conversion functions* you create to augment the object conversion process. When a conversion function is called during object conversion, two objects are passed to it:

- A representation of an existing, unconverted object (an instance of ooConvertInObject)
- A representation of the object after it has been converted (an instance of ooConvertInOutObject)

You use member functions of <code>ooConvertInObject</code> to get values from the unconverted object's primitive data members. From these pre-conversion values, you can compute data-member values to be set after the object has been converted. You use member functions of <code>ooConvertInOutObject</code> to set data-member values in the converted object.

The ooConvertInOutObject class allows you to set primitive data members. You can access data members that are defined:

- Directly in the class whose objects are being converted (see the set *Type* member functions).
- In a base class of the class whose objects are being converted (see the getNewBaseClass member function).
- In an embedded class of the class whose objects are being converted (see the getNewDataMember member function).

See Chapter 19, "Object Conversion," in the Objectivity/C++ programmer's guide for an example that uses this class.

Setting Primitive Data Members	setFloat32
	<u>setFloat64</u>
	<u>setInt8</u>
	<u>setInt16</u>
	<u>setInt32</u>
	<u>setInt64</u>
	<u>setUInt8</u>
	<u>setUInt16</u>
	setUInt32
	<u>setUInt64</u>
Accessing a Base or Embedded Class	<u>getNewBaseClass</u>
	getNewDataMember

## **Reference Summary**

# **Reference Index**

<u>getNewBaseClass</u>	Gets the part of a converted object that is inherited from the specified base class.
<u>getNewDataMember</u>	Gets the part of a converted object that belongs to the specified embedded object.
<u>setFloat32</u>	Sets the value of the specified float32 data member in a converted object.
<u>setFloat64</u>	Sets the value of the specified float64 data member in a converted object.

<u>setInt8</u>	Sets the value of the specified int8 data member in a converted object.
<u>setInt16</u>	Sets the value of the specified int16 data member in a converted object.
<u>setInt32</u>	Sets the value of the specified int32 data member in a converted object.
<u>setInt64</u>	Sets the value of the specified int64 data member in a converted object.
<u>setUInt8</u>	Sets the value of the specified uint8 data member in a converted object.
<u>setUInt16</u>	Sets the value of the specified uint16 data member in a converted object.
<u>setUInt32</u>	Sets the value of the specified uint32 data member in a converted object.
<u>setUInt64</u>	Sets the value of the specified uint64 data member in a converted object.

# **Member Functions**

## getNewBaseClass

	Gets the part of a converted object that is inherited from the specified base class.		
	ooStatus getNewBaseClass( const char * <i>baseClassName</i> , ooConvertInOutObject & <i>basePartOfConvObj</i> );		
Parameters	<i>baseClassName</i> Name of a base class of the class whose objects are being converted.		
	<pre>basePartOfConvObj Name of the ooConvertInOutObject variable in which the inherited part is to be returned.</pre>		
Returns	oocSuccess if successful, or oocError if the base class path is not valid.		
Discussion	You use this member function when you need to set the value of an inherited primitive data member of a converted object. First, you use this member function to get the part of the object that is inherited from the specified base class. Then, you use the appropriate set $Type$ member function on the inherited part		

(accessed through the  ${\tt basePartOfConvObj}$  variable) to set the value of the desired data member.

See also ooConvertInObject::getOldBaseClass

### getNewDataMember

Gets the part of a converted object that belongs to the specified embedded object.

	ooStatus getNewDataMember(
	const char *memberName,
	ooConvertInOutObject &embeddedPartOfConvObj);
Parameters	memberName
	Name of the data member that specifies the desired embedded object.
	embeddedPartOfConvObj
	Name of the ooConvertInOutObject variable in which the embedded part is to be returned.
Returns	oocSuccess if successful, or oocError if the data-member name is not valid.
Discussion	You use this member function when you need to set the value of a primitive data member of an object that is embedded as a data member of a converted object. First, you use this member function to get the desired embedded part of the converted object. Then, you use the appropriate $setType$ member function on the embedded part (accessed through the <i>embeddedPartOfConvObj</i> variable) to set the value of the desired primitive data member.

### setFloat32

Sets the value of the specified float32 data member in a converted object.

ooStatus setFloat32( const char \*memberName, const float32 value);

Parameters	memberName Name of the data member whose value is to be set.
	value Variable containing the desired value.
Returns	oocSuccess if successful, or oocError if the data-member name

is not valid.

### setFloat64

Sets the value of the specified float64 data member in a converted object.

```
ooStatus setFloat64(
    const char *memberName,
    const float64 value);
```

Parameters	memberName		
	Name of the data member whose value is to be set.		
	value		
	Variable containing the desired value.		
Returns	oocSuccess if successful, or oocError if the data-member name is not valid.		

### setInt8

Sets the value of the specified int8 data member in a converted object.

```
ooStatus setInt8(
    const char *memberName,
    const int8 value);
```

```
Parameters memberName
```

Name of the data member whose value is to be set.

value

Variable containing the desired value.

Returns oocSuccess if successful, or oocError if the data-member name is not valid.

### setInt16

Sets the value of the specified int16 data member in a converted object.

ooStatus setInt16( const char \*memberName, const int16 value);

Parameters	memberName
	Name of the data member whose value is to be set.

value

Variable containing the desired value.

Returns oocSuccess if successful, or oocError if the data-member name is not valid.

### setInt32

Sets the value of the specified int 32 data member in a converted object.

Parameters memberName Name of the data member whose value is to be set. value Variable containing the desired value.

Returns oocSuccess if successful, or oocError if the data-member name is not valid.

### setInt64

Sets the value of the specified int64 data member in a converted object.

```
ooStatus setInt64(
    const char *memberName,
    const int64 value);
```

#### Parameters memberName

Name of the data member whose value is to be set.

value

Variable containing the desired value.

Returns oocSuccess if successful, or oocError if the data-member name is not valid.

### setUInt8

Sets the value of the specified uint8 data member in a converted object.

ooStatus setUInt8( const char \*memberName, const uint8 value);

Parameters	memberName Name of the data member whose value is to be set.
	value Variable containing the desired value.
Returns	oocSuccess if successful, or oocError if the data-member name is not valid.

### setUInt16

Sets the value of the specified uint16 data member in a converted object.

```
ooStatus setUInt16(
    const char *memberName,
    const uint16 value);
```

Parameters memberName

Name of the data member whose value is to be set.

```
value
```

Variable containing the desired value.

Returns oocSuccess if successful, or oocError if the data-member name is not valid.

### setUInt32

Sets the value of the specified uint 32 data member in a converted object.

ooStatus	s	etUIn	t3:	2 (	
cons	t	char	* n	nemberName	,
cons	t	uint3	32	value);	

#### Parameters memberName

Name of the data member whose value is to be set.

value

Variable containing the desired value.

Returns oocSuccess if successful, or oocError if the data-member name is not valid.

### setUInt64

Sets the value of the specified uint64 data member in a converted object.

Parameters	memberName Name of the data member whose value is to be set.
	value Variable containing the desired value.
Returns	oocSuccess if successful, or oocError if the data-member name is not valid.

# ooDBObj Class

Inheritance: ooObj->ooDBObj

Handle Class: <u>ooHandle(ooDBObj)</u>

**Object-Reference Class:** <u>ooRef(ooDBObj)</u>

The persistence-capable class ooDBObj represents an Objectivity/DB *database*. You use ooDBObj only for database creation; you work with an existing database through a handle or object reference.

(*DRO*) The oodbobj class also represents a *database image*, which you can create if you have bought and installed both Objectivity/DB Data Replication Option (Objectivity/DRO) and Objectivity/DB Fault Tolerant Option (Objectivity/FTO). You use oodbobj to create an initial database image; additional images are created through a handle or object reference to an existing database.

See:

■ "Reference Index" on page 235 for a list of ooDBObj member functions

For operations performed through a database handle or object reference, see:

• "Reference Summary" on page 541

## About Databases and Database Images

A database is the second highest level in the Objectivity/DB storage hierarchy; a federated database contains one or more databases and every database contains one or more containers.

Physically, a database is maintained in a *database file*, which contains the database and all the persistent data stored in it. Each database is attached to exactly one federated database and is listed in the federated database's catalog. Database

files may reside on different machines than the system-database file for the federated database to which they are attached.

In addition to having a physical filename, each database has a system name, which is its logical name within the federated database. The system name of each database must be unique among all the system names of databases and autonomous partitions in the federated database. A database also has a unique integer identifier within its federated database. When a database is created, you *must* specify a system name for it; you may additionally specify an identifier or allow it to be assigned by Objectivity/DB.

Every database contains a *default container* that is created automatically; default containers are created as instances of class <u>ooDefaultContObj</u>.

(*DRO*) In a partitioned federated database, you can use Objectivity/DRO to create and manage multiple *images* of a database to replicate its data across multiple autonomous partitions. Each image of a database is a separate database file that contains all the data in the database. Location and order of creation do not distinguish an image in any way. Each image is controlled by a single autonomous partition; each partition can control at most one image of any given database. If one image of a particular database becomes unavailable due to a network or machine failure, work may continue with a different available image.

All images of a database share the same system name and database identifier; each image is distinguished by the autonomous partition that controls it. Each image has a *weight*, which is used to determine whether a *quorum* of replicated images exists. In general, tasks affecting database images require that a quorum of the database images be available (an image is available if the containing partition is available).

See:

- Chapter 8, "Storage Objects," of the Objectivity/C++ programmer's guide for additional information about databases.
- Chapter 28, "Database Images," of the Objectivity/C++ programmer's guide for additional information about database images.

## Working With Databases and Database Images

Databases are normally created and deleted from within an application, although these tasks can also be performed with administration tools. An application:

Creates a database using the ooDBObj class constructor and operator new. The new database is represented in memory as an instance of ooDBObj, which serves as a proxy for the actual database file on disk.

- Deletes a database using the global function ooDelete; operator delete is not available on class ooDBObj.
- **NOTE** (*DRO*) You use ooDBObj only to create an initial database; you create and work with additional images through a database handle or object reference.

To work with a new database, an application must assign the result of operator new to a database handle or object reference (an instance of ooHandle(ooDBObj) or ooRef(ooDBObj)). Similarly, to identify and work with an existing database, the application must open it through a database handle or object reference; multiple handles and object references can be set to reference the same database. The application then operates on the referenced database by:

- Calling various member functions on any of the referencing handles or object references.
- Passing one of the referencing handles or object references to various global functions or member functions of other classes.

Member functions for operating on existing databases are defined in the ooRefHandle(ooDBObj) classes, not in the ooDBObj class itself. Although ooDBObj inherits member functions from ooObj, these generally do not apply to databases; in fact, the inherited member functions cannot be called, because there is no indirect member-access operator (->) on the ooRefHandle(ooDBObj) classes.

You may not derive classes from ooDBObj.

## **Reference Index**

<u>ooDBObj</u>	Constructs a new database with the specified system name, database file, image weight, and default container parameters.
<u>operator new</u>	Creates a new database in the currently open federated database.

## Constructors

### ooDBObj

Constructs a new database with the specified system name, database file, image weight, and default container parameters.

#### ooDBObj(

```
const char *dbSysName,
const uint32 defContInitPages = 0,
const uint32 defContGrowth = 0,
const char *hostName = 0,
const char *pathName = 0,
uint32 weight = 1,
uint16 userDBID = 0);
```

Parameters

#### dbSysName

System name of the database to create. The specified name:

- Must follow the same naming rules as files of your operating system.
- Must be unique among all the system names of databases and autonomous partitions in the federated database.

The specified name is also used as the database filename if one is not explicitly specified as part of *pathname*.

#### *defContInitPages*

Initial number of logical pages to allocate for the default container in the new database. The maximum value is 65535. If you omit this parameter or specify 0, the system default value (4) is used.

#### defContGrowth

Amount by which the default container may grow, expressed as a percentage of its current size. If you omit this parameter or specify 0, the system default value (10%) is used.

#### hostName

Name of the data server host on which to create the database file. If you specify a *hostName*, you must also specify the *pathName* parameter. If you omit the *hostName* parameter or specify 0, the database file is created on the same host as the federated database's system database file.

#### pathName

Fully qualified pathname of the database file on *hostName*. If you specify a *pathName*, you must also specify the *hostName* parameter. If you omit the *pathName* parameter or specify 0, the database file is created in the same directory as the federated database's system database file.

The *pathName* may optionally include a filename for the database file. If you omit a filename, the system name is used.

weight

(*DRO*) Weight of the first database image. *weight* must be an integer greater than zero. If you omit this parameter, the weight is 1.

#### userDBID

Application-assigned database identifier, expressed as a single integer. If you omit this parameter or specify 0, the identifier is assigned by Objectivity/DB.

See also <u>ooReplace</u> global macro

## **Operators**

### operator new

Creates a new database in the currently open federated database.

1.	<pre>void *operator new(    size t,</pre>
	const <i>ooRefHandle</i> (ooFDObj) & <i>containingFD</i> );
2.	void *operator new(
	size_t,
	const <i>ooRefHandle</i> (ooFDObj) & <i>containingFD</i> ,
	<pre>const ooRefHandle(ooAPObj) &amp;containingAP);</pre>

Parameters size\_t

Do not specify; this parameter is automatically initialized by the compiler with the size of the class type in bytes.

#### containingFD

Federated database in which to create the new database. A database can be created only in the currently open federated database. As a consequence, you can omit this parameter in variant 1. In either variant, the specified parameter must be an object reference or handle to the open federated database.

#### containingAP

(*FTO*) Autonomous partition in which to create the new database. You can specify an object reference or handle to the desired autonomous partition. If you omit this parameter, the database is created in the boot autonomous partition.

- Returns Memory pointer to the new database. This pointer is null if an error occurs during the creation of the database or if a database with the specified system name already exists.
- Discussion This operator must be used in an update transaction. When you commit or checkpoint the transaction, the new database is made permanent on disk. If the transaction is aborted, the database is not created.

Expressions containing this operator are of the following form:

new(locatingDirectives) ooDBObj(initializers)

The *locatingDirectives* are the *containingFD* and optional *containingAP* parameters.

The *initializers* are a list of values to be passed as parameters to the ooDBObj <u>constructor</u>. At a minimum, you must specify the system name of the new database. If you omit the host and pathname, the database file is located in the directory containing the system database file of the specified federated database or autonomous partition.

You normally assign the result of operator new directly to a database handle. You can verify the creation of the database by checking whether the handle is null.

# ooDefaultContObj Class

Inheritance: ooObj->ooContObj->ooDefaultContObj

Handle Class: ooHandle(ooDefaultContObj)

**Object-Reference Class:** ooRef(ooDefaultContObj)

The persistence-capable class ooDefaultContObj represents the *default container* in a database.

See:

"Reference Summary" on page 240 for an overview of member functions

## About Default Containers

A default container is created automatically by Objectivity/DB for each database. The default container holds:

- Basic objects that are clustered with the database but not explicitly assigned to an application-defined container.
- Scope names of objects that are named in the scope of the database.
   Consequently every default container is hashed, with a hash value of 1.

Every default container is an instance of ooDefaultContObj and has the system name \_ooDefaultContObj.

## **Working With Default Containers**

A default container is created automatically for each database when the database is created. You should not create additional default containers in a database, nor shoud you delete a database's default container. You specify a default container's initial size and growth factor when you create the database. An application should not create derived classes from <code>ooDefaultContObj</code>.

Like other storage objects, default containers are manipulated through handles or object references. The handle and object-reference classes <code>ooHandle(ooDefaultContObj)</code> and <code>ooRef(ooDefaultContObj)</code> are derived from <code>ooHandle(ooContObj)</code> and <code>ooRef(ooContObj)</code>, respectively. You can find a default container in a database using the default container's system name.

Default containers are found by operations that initialize an object iterator of class ooltr(ooContObj) or ooltr(ooObj). When advancing such an iterator, you should test for the default container before performing any operation that deletes or moves the found container. For example, when the iterator references the next found container, you could call the container's inherited oolsKindOf member function.

# **Reference Summary**

In the following table, the member functions indicated as *(inherited)* are defined by ooObj and are available for default containers; they are documented with the <u>ooObj</u> class (page 431) and are the *only* ooObj member functions that apply to default containers.

Working With the Container	<u>ooGetTypeN</u> (inherited) <u>ooGetTypeName</u> (inherited) <u>ooIsKindOf</u> (inherited) <u>ooUpdate</u> (inherited) <u>ooValidate</u> (inherited)
----------------------------	--

# ooEqualLookupField Class

Inheritance: ooLookupFieldBase->ooEqualLookupField

The non-persistence-capable indexing class <code>ooEqualLookupField</code> represents a *lookup field* that tests whether the value of an indexed object's key field is equal to (=) a specified comparison value.

See:

"Reference Index" on page 242 for a list of member functions

Applications that use indexes must include the ooIndex.h header file.

## About Equal-To Lookup Fields

A lookup field is part of a *lookup key* (an instance of the <u>ooLookupKey</u> class that is used for looking up persistent objects in an index). A lookup key consists of one or more lookup fields, each representing a condition that a found object must satisfy. Lookup fields are instances of the concrete classes derived from the abstract base class <u>ooLookupFieldBase</u>.

The concrete class <code>ooEqualLookupField</code> represents a condition that tests whether the values of a particular key field are equal to a particular comparison value. You use the class constructor to specify the key field and the value.

For complete information about using lookup fields in lookup keys, see "About Lookup Keys" on page 399

# **Reference Index**

<u>ooEqualLookupField</u>

Constructs a new lookup field for testing whether values of the specified key field are equal to the specified comparison value.

## Constructors

### ooEqualLookupField

Constructs a new lookup field for testing whether values of the specified key field are equal to the specified comparison value.

- 2. ooEqualLookupField( const ooTypeNumber typeN, const char \*memberName, const void \*valuePtr);

#### Parameters

field

Key-field object representing a key field of the indexed class. The values of the specified key field will be tested by this lookup field.

#### typeN

Type number of the indexed class.

#### memberName

Name of a data member that serves as a key field of the indexed class. The values of the specified key field will be tested by this lookup field.

The specified data member must be defined or inherited by the class specified by *typeN*. You can qualify the name of an inherited data member using the following notation (where *baseClassName* is the name of the base class that defines the inherited data member):

baseClassName::dataMemberName

You *must* qualify the name of an inherited data member if the member name is ambiguous (for example, the same name is defined in both the base class and the class to be indexed) or if the member name is not visible due to access control.

#### valuePtr

Data value to which key-field values are to be compared. The type of this value must match the key-field type (the type of the data member you specified to the key-field object). Using data of any other type may have unpredictable results.

Discussion This lookup field causes the values of the specified key field to be tested when an iterator scans the index; an indexed object is found if its tested value is equal to the value specified by *valuePtr*.

This lookup field will be ignored if *field* does not specify a key field of the index being searched.

# ooFDObj Class

Inheritance: ooObj->ooFDObj

 Handle Class:
 ooHandle(ooFDObj)

 Object-Reference Class:
 ooRef(ooFDObj)

 d\_Database

The persistence-capable ooFDObj class represents an Objectivity/DB *federated database*. The ooFDObj class is used internally by Objectivity/DB; you work with a federated database through a handle or object reference.

See:

 "Reference Summary" on page 575 for operations performed through a federated-database handle or object reference

(*ODMG*) An Objectivity/DB federated database corresponds to an ODMG database (see class <u>d\_Database</u>).

## **About Federated Databases**

A federated database is the highest level in the Objectivity/DB storage hierarchy; each federated database logically contains one or more databases.

Physically, an Objectivity/DB federated database is maintained in a *system-database file*, which stores the schema for the federated database, a catalog of all the databases, and the scope names of any objects named in the scope of the federated database. Configuration information for the federated database is maintained in a second file (the *boot file*), along with various other attributes. The simple name of the boot file serves as the federated database's system name.

# **Working With Federated Databases**

A federated database can be created and deleted only with administration tools. An application obtains a single instance of ooFDObj by opening the federated database through a federated-database handle or object reference (an instance of ooHandle(ooFDObj) or ooRef(ooFDObj)). The obtained ooFDObj instance serves as a proxy for the actual system-database file on disk.

The application operates on the federated database by:

- Calling various member functions on the federated-database handle or object reference.
- Passing the federated-database handle or object reference to various global functions or member functions of other classes.

Member functions for operating on federated databases are defined in the ooRefHandle(ooFDObj) classes, not in the ooFDObj class itself. Although ooFDObj inherits member functions from ooObj, these generally do not apply to federated databases; in fact, the inherited member functions cannot be called, because there is no indirect member-access operator (->) on the ooRefHandle(ooFDObj) classes.

You may not derive classes from ooFDObj.

See the Objectivity/DB administration book for information about federated database creation, files, and attributes.

# ooGCContObj Class

Inheritance: ooObj->ooContObj->ooGCContObj

Handle Class: ooHandle(ooGCContObj)

**Object-Reference Class:** ooRef(ooGCContObj)

The persistence-capable class <code>ooGCContObj</code> represents a garbage-collectible container, and serves as the base class for application-defined garbage-collectible container classes. The <code>ooGCContObj</code> class and its corresponding handle and object-reference classes together define the behavior of garbage-collectible containers.

See:

- "Reference Summary" on page 248 for an overview of member functions
- "Reference Index" on page 249 for a list of member functions

## **About Garbage-Collectible Containers**

As the name implies, garbage-collectible containers adhere to a garbage-collection paradigm analogous to the memory of languages such as Java and Smalltalk. These languages consider an object to be "garbage" if it is not referenced by any other object, and periodically clean up memory by deleting garbage objects. In an Objectivity/DB garbage-collectible container, a persistent object becomes "garbage" if it is no longer linked to a named root either directly or indirectly through associations, reference attributes, or membership in persistent collections. However, "garbage" persistent objects are not deleted automatically; you must use the oogc administration tool when you want to delete invalid objects from the garbage-collectible containers of a federated database.

Because Objectivity/C++ does not directly support named roots, Objectivity/C++ applications generally use standard containers (instances of ooContObj) instead of garbage-collectible containers. In a standard container, each unwanted object must be deleted explicitly (for example, using the ooDelete global function). However an Objectivity/C++ application can choose to create garbage-collectible containers if it interoperates with Objectivity for Java or Objectivity/Smalltalk applications.

**NOTE** The ooGCContObj class has a predefined derived class called ooGCRootsCont. Do not use this derived class in any way—for example, do not instanatiate it or or derive any classes from it.

# **Working With Garbage-Collectible Containers**

You can work with instances of <code>oogCContObj</code> or you can work with instances of application-defined classes derived from <code>oogCContObj</code>.

As is the case for any container, you specify whether a garbage-collectible container is to be transient or persistent when you create it; garbage-collectible containers should be persistent. You create a garbage-collectible container with a call to operator new; the clustering directive in that call specifies where in the federated database to store the new container.

Like other persistent objects, garbage-collectible containers are normally manipulated through handles or object references. You can store and find a garbage-collectible container in the database just as you would any other persistent object.

# **Reference Summary**

In the following table, the member functions indicated as *(inherited)* are defined by oobj and are available for containers; they are documented with the <u>oobj</u> class (page 431) and are the *only* oo0bj member functions that apply to containers. operator new is documented with <u>ooContObj</u> (page 207).

Creating and Deleting a Garbage-Collectible Container	<u>ooGCContObj</u> <u>operator new</u> (inherited from ooContObj) <u>operator delete</u> (inherited)
Working With a Garbage-Collectible Container	<u>ooGetTypeN</u> (inherited) <u>ooGetTypeName</u> (inherited) <u>ooIsKindOf</u> (inherited) <u>ooThis</u> <u>ooUpdate</u> (inherited) <u>ooValidate</u> (inherited)
ODMG Interface	operator new

## **Reference Index**

<u>ooGCContObj</u>	Default constructor that constructs a new garbage-collectible container.
<u>ooThis</u>	Sets an object reference or handle to reference this garbage-collectible container.

# Constructors

## ooGCContObj

Default constructor that constructs a new garbage-collectible container.

```
ooGCContObj();
```

# **Member Functions**

### ooThis

Sets an object reference or handle to reference this garbage-collectible container.

- 1. ooHandle(ooGCContObj) ooThis() const;

	<pre>3. ooHandle(ooGCContObj) &amp;ooThis(</pre>
Parameters	<i>container</i> Object reference or handle to be set to this garbage-collectible container.
Returns	Object reference or handle to this garbage-collectible container.
Discussion	You normally use $ooThis$ in a member function of an application-defined container class; when such a member function is called on a container of the class, $ooThis$ provides the member function with an object reference or handle to the container. The member function can then perform operations on the container that are available only through an object reference or handle.
	When called without a <i>container</i> parameter, this member function allocates a new handle and returns it. Otherwise, this member function returns the object reference or handle that is passed to it.

# ooGeneObj Class

Inheritance: ooObj->ooGeneObj

Handle Class: ooHandle(ooGeneObj)

**Object-Reference Class:** ooRef(ooGeneObj)

The persistence-capable class <code>ooGeneObj</code> represents a *genealogy*. The <code>ooGeneObj</code> class and its corresponding handle and object-reference classes together define the behavior of genealogies.

See:

- "Reference Summary" on page 253 for an overview of member functions
- "Reference Index" on page 253 for a list of member functions

## **About Genealogies**

A genealogy manages a set of next and previous versions of a particular basic object, allowing you to appoint one of these versions as the default version. The semantics of the default version are application-specific—some applications may successively appoint each new version as the default; other applications may use an older default version until a later version has been developed, tested, approved, and finally appointed as the new default.

A genealogy maintains:

- A bidirectional to-many association to each of the versions it manages.
- A bidirectional one-to-one association to its default version.

```
class ooGeneObj : public ooObj {
```

```
// Links a genealogy to the versions in it.
ooRef(ooObj) allVers[ ] <-> geneObj;
```

}

```
// Links a genealogy to the default version.
ooRef(ooObj) defaultVers <-> defaultToGeneObj;
...
```

Because of these associations, you can find the default version from any other version in the genealogy. Furthermore, you can easily find all of the versions in the genealogy by iterating over the destination objects of the allVers association. This technique of finding versions is more convenient than the alternative, which is to recursively request the next version(s) from each previous one.

A genealogy may, but need not, track an entire set of next and previous versions. After the first default version is appointed, all subsequent versions created from it are automatically linked to the genealogy. If you appoint an object to be the default version before you create any versions from it, then all of the object's versions will be linked to the genealogy. If, however, you create a few versions and then appoint a default, the pre-existing versions are not automatically added to the genealogy; you can add them explicitly by setting them as destination objects of the genealogy's allVers association.

## **Working With Genealogies**

Genealogies can be created either implicitly or explicitly. An instance of <code>ooGeneObj</code> is created implicitly the first time you appoint a default version by calling the <code>setDefaultVers</code> member function on a handle to a basic object. Alternatively, you can create a genealogy explicitly as you would any other basic object and then appoint its default version by calling the genealogy's <code>set\_defaultVers</code> member function.

As is the case for any basic object, you specify whether a genealogy is to be transient or persistent when you create it; genealogies *must be persistent*. You create a genealogy with a call to the new operator; the clustering directive in that call specifies where in the federated database to store the new genealogy.

You may create and work with instances of the <code>ooGeneObj</code> class. To support more complex versioning semantics, an application can define its own custom genealogy class derived from <code>ooGeneObj</code>, and then work with instances of the derived class.

Like other persistent objects, genealogies are normally manipulated through handles or object references. You can store and find a genealogy in the database just as you would any other persistent object.
# **Reference Summary**

In the following table, operators indicated as *(inherited)* are overloaded in this class with no change in behavior; they are documented with the 000bj class (page 431), along with the other inherited member functions not listed here.

Creating and Deleting a Genealogy	<u>ooGeneObj</u> <u>operator new</u> (inherited) <u>operator delete</u> (inherited)
Referencing This Genealogy	<u>ooThis</u>
Managing the Default Version	<u>defaultVers</u> <u>del_defaultVers</u> <u>exist_defaultVers</u> <u>set_defaultVers</u>
Managing All Versions	add_allVers allVers exist_allVers sub_allVers

# **Reference Index**

allVers	Initializes an object iterator to find, and optionally open, all versions in this genealogy that satisfy any specified selection criteria.
<u>add_allVers</u>	Links this genealogy with the specified version, adding the version to this genealogy.
<u>defaultVers</u>	Finds, and optionally opens, the default version in this genealogy.
<u>del_allVers</u>	Removes all versions from this genealogy by deleting all of its allVers associations.
<u>del_defaultVers</u>	Deletes any defaultVers association that exists for this genealogy, leaving this genealogy without a default version.
<u>exist_allVers</u>	Tests whether this genealogy has any versions; if a parameter is given, tests whether the specified object is a version in this genealogy.

<u>exist defaultVers</u>	Tests whether this genealogy has a default version; if a parameter is given, tests whether the specified object is the default version in this genealogy.
<u>ooGeneObj</u>	Default constructor that constructs a new genealogy with no associated versions.
<u>ooThis</u>	Gets an object reference or handle to this genealogy.
<u>set_defaultVers</u>	Links this genealogy to the specified object, setting the specified object as the default version in the genealogy.
sub allVers	Removes the specified version from this genealogy by deleting the allVers association link from this genealogy to the specified version.

## **Constructors and Destructors**

## ooGeneObj

Default constructor that constructs a new genealogy with no associated versions.

ooGeneObj();

## **Member Functions**

## allVers

Initializes an object iterator to find, and optionally open, all versions in this genealogy that satisfy any specified selection criteria.

```
    ooStatus allVers(
        ooItr(ooObj) &iterator,
        const ooMode openMode = oocNoOpen) const;
    ooStatus allVers(
        ooItr(ooObj) &iterator,
        const char *predicate) const;
    ooStatus allVers(
        ooItr(ooObj) &iterator,
        const ooMode openMode,
        const ooAccessMode access,
        const char *predicate) const;
```

#### Parameters iterator

Object iterator for finding all versions associated with this genealogy.

#### openMode

Intended level of access to each version found by the iterator's next member function:

- oocNoOpen (the default in variant 1) causes next to set the iterator to the next found version without opening it.
- oocRead causes next to open the next found version for read.
- oocUpdate causes next to open the next found version for update.

Warning: If versioning is enabled for one or more found objects, specifying oocUpdate means that next will create a new version of each such object.

#### predicate

String expression in predicate query language; specifies the condition to be met by the found versions. The iterator is initialized only with versions that match *predicate*.

#### access

Level of access control of the data members that *predicate* can test:

- Specify oocPublic to permit the predicate to test only public data members, preserving encapsulation.
- Specify oocAll to permit the predicate to test any data member. To
  preserve encapsulation, you should use this mode only within member
  functions of the class you are querying.

Returns oocSuccess if successful; otherwise oocError.

See also <u>add allVers</u> <u>sub allVers</u> <u>del allVers</u>

#### add\_allVers

Links this genealogy with the specified version, adding the version to this genealogy.

ooStatus add\_allVers(const ooHandle(ooObj) &object);

Parameters object

Handle to a basic object.

Returns oocSuccess if successful; otherwise oocError.

Discussion The operation links this genealogy to the specified version by:

- Creating an allVers association link in this genealogy.
- Creating the inverse geneObj association link in the specified version.

The application must be able to obtain update locks on both objects.

See also <u>allVers</u> <u>sub\_allVers</u> <u>del\_allVers</u>

## defaultVers

Finds, and optionally opens, the default version in this genealogy.

	<pre>1. ooHandle(ooObj) defaultVers(</pre>		
	<pre>2. ooRef(ooObj) &amp;defaultVers(</pre>		
	<pre>3. ooHandle(ooObj) &amp;defaultVers(</pre>		
Parameters	defaultVersion		
	Object reference or handle to set to the default version.		
	openmode		
	Intended level of access to the default version:		
	<ul> <li>Specify occNoOpen (the default) to set the returned object reference or handle to the object without opening it.</li> </ul>		
	Specify oocRead to open the object for read.		
	<ul> <li>Specify oocUpdate to open the object for update.</li> </ul>		
Returns	Object reference or handle to the default version. A null object reference or handle is returned if this genealogy has no default version.		
Discussion	When called without an object reference or handle argument, this member function allocates a new handle and returns it. Otherwise, this member function returns the result in the object reference or handle that is passed to it.		
See also	<u>set_defaultVers</u> <u>del_defaultVers</u>		

## del\_allVers

Removes all versions from this genealogy by deleting all of its allVers associations.

ooStatus del\_allVers();

Returns oocSuccess if successful; otherwise oocError.

Discussion This operation also deletes the inverse geneObj association link from each of the former versions. The application must be able to obtain update locks on all of the affected objects.

You should call the <u>exist\_allVers</u> member function to test whether any associations exist before you try to delete them.

See also <u>allVers</u> <u>add\_allVers</u> <u>sub\_allVers</u>

## del\_defaultVers

Deletes any defaultVers association that exists for this genealogy, leaving this genealogy without a default version.

ooStatus del\_defaultVers();

Returns oocSuccess if successful; otherwise oocError.

Discussion This operation also deletes the inverse defaultToGeneObj association link from the former default version. The application must be able to obtain update locks on both objects.

You typically call this function to remove a genealogy's previous default version before setting a new one. You shouldn't leave a genealogy without a default version.

You should call the <u>exist\_defaultVers</u> member function to test whether an association exists before you try to delete it.

See also <u>defaultVers</u> set\_defaultVers

## exist\_allVers

Tests whether this genealogy has any versions; if a parameter is given, tests whether the specified object is a version in this genealogy.

	<pre>1. ooBoolean exist_allVers() const;</pre>
	<pre>2. ooBoolean exist_allVers(</pre>
Parameters	object Handle to the object to be tested as a version.
Returns	(Variant 1) OOCTrue if this genealogy has any versions, otherwise OOCFalse. (Variant 2) OOCTrue if the specified object is a version in this genealogy,

## exist\_defaultVers

otherwise oocFalse.

Tests whether this genealogy has a default version; if a parameter is given, tests whether the specified object is the default version in this genealogy.

1.	ooBoolean exist_defaultVers() const;
2.	ooBoolean exist_defaultVers(
	<pre>const ooHandle(ooObj) &amp;object) const;</pre>

Parameters object

Handle to the object to be tested as the default version.

Returns (Variant 1) OOCTrue if this genealogy has a default version, otherwise OOCFalse. (Variant 2) OOCTrue if the specified object is the default version in this genealogy, otherwise OOCFalse.

## ooThis

Gets an object reference or handle to this genealogy.

- 1. ooHandle(ooGeneObj) ooThis() const;

Parameters genealogy

#### Object reference or handle to be set to this genealogy.

#### Returns Object reference or handle to this genealogy.

Discussion You normally use ooThis in a member function of a custom genealogy class. When such a member function is called on a genealogy of the class, ooThis provides the member function with an object reference or handle to the genealogy. The member function can then perform operations on the genealogy that are available only through an object reference or handle. (**Note:** You actually use the type-specific version of ooThis that is generated for your custom genealogy class by the DDL processor.)

> When called without a *genealogy* parameter, *ooThis* allocates a new handle and returns it. Otherwise, *ooThis* returns the object reference or handle that is passed to it.

#### set\_defaultVers

Links this genealogy to the specified object, setting the specified object as the default version in the genealogy.

ooStatus set\_defaultVers(const ooHandle(ooObj) &object);

Parameters object Handle to a basic object.

Returns oocSuccess if successful; otherwise oocError.

Discussion All subsequently created versions of the specified object are automatically added to this genealogy.

The operation links this genealogy with the specified object by:

- Creating a defaultVers association link in this genealogy.
- Creating the inverse defaultToGeneObj association link in the specified object.

The application must be able to obtain update locks on both objects.

Because defaultVers is a to-one association, an error is signaled if this object already has a link for the association. That is, you must remove any existing default version from the genealogy before you set a new default version.

See also <u>defaultVers</u> <u>del\_defaultVers</u>

## sub\_allVers

Removes the specified version from this genealogy by deleting the allVers association link from this genealogy to the specified version.

ooStatus sub\_allVers(const ooHandle(ooObj) &object);

Parameters object

Handle to the version to be removed.

Returns oocSuccess if successful; otherwise oocError.

See also <u>add allVers</u> <u>allVers</u> <u>del allVers</u>

# ooGreaterThanEqualLookupField Class

#### Inheritance: ooLookupFieldBase->ooGreaterThanEqualLookupField

The non-persistence-capable indexing class <code>ooGreaterThanEqualLookupField</code> represents a *lookup field* that tests whether the value of an indexed object's key field is greater than or equal to (>=) a specified comparison value.

See:

• "Reference Index" on page 262 for a list of member functions

Applications that use indexes must include the ooIndex.h header file.

## About Greater-Than-Or-Equal-To Lookup Fields

A lookup field is part of a *lookup key* (an instance of the <u>ooLookupKey</u> class that is used for looking up persistent objects in an index). A lookup key consists of one or more lookup fields, each representing a condition that a found object must satisfy. Lookup fields are instances of the concrete classes derived from the abstract base class <u>ooLookupFieldBase</u>.

The concrete class <code>ooGreaterThanEqualLookupField</code> represents a condition that tests whether the values of a particular key field are greater than or equal to a particular comparison value. You use the class constructor to specify the key field and the value.

For complete information about using lookup fields in lookup keys, see "About Lookup Keys" on page 399.

# **Reference Index**

ooGreaterThanEqualLookupField

Constructs a new lookup field for testing whether values of the specified key field are greater than or equal to the specified comparison value.

# Constructors

## ooGreaterThanEqualLookupField

Constructs a new lookup field for testing whether values of the specified key field are greater than or equal to the specified comparison value.

1.	ooGreaterThanEqualLookupField			LookupField(
	const	ооКеу	/Field	&field,
	const	void	*value	ePtr);

2. ooGreaterThanEqualLookupField( const ooTypeNumber typeN, const char \*memberName, const void \*valuePtr);

#### Parameters field

Key-field object representing a key field of the indexed class. The values of the specified key field will be tested by this lookup field.

#### typeN

Type number of the indexed class.

#### memberName

Name of a data member that serves as a key field of the indexed class. The values of the specified key field will be tested by this lookup field.

The specified data member must be defined or inherited by the class specified by *typeN*. You can qualify the name of an inherited data member using the following notation (where *baseClassName* is the name of the base class that defines the inherited data member):

baseClassName::dataMemberName

You *must* qualify the name of an inherited data member if the member name is ambiguous (for example, the same name is defined in both the base class and the class to be indexed) or if the member name is not visible due to access control.

#### valuePtr

Data value to which key-field values are to be compared. The type of this value must match the key-field type (the type of the data member you specified to the key-field object). Using data of any other type may have unpredictable results.

Discussion This lookup field causes the values of the specified key field to be tested when an iterator scans the index; an indexed object is found if its tested value is greater than or equal to the value specified by *valuePtr*.

This lookup field will be ignored if *field* does not specify a key field of the index being searched.

# ooGreaterThanLookupField Class

#### Inheritance: ooLookupFieldBase->ooGreaterThanLookupField

The non-persistence-capable indexing class <code>ooGreaterThanLookupField</code> represents a *lookup field* that tests whether the value of an indexed object's key field is greater than (>) the specified comparison value.

See:

• "Reference Index" on page 266 for a list of member functions

Applications that use indexes must include the ooIndex.h header file.

## **About Greater-Than Lookup Fields**

A lookup field is part of a *lookup key* (an instance of the <u>ooLookupKey</u> class that is used for looking up persistent objects in an index). A lookup key consists of one or more lookup fields, each representing a condition that a found object must satisfy. Lookup fields are instances of the concrete classes derived from the abstract base class <u>ooLookupFieldBase</u>.

The concrete class <code>ooGreaterThanLookupField</code> represents a condition that tests whether the values of a particular key field are greater than a particular comparison value. You use the class constructor to specify the key field and the value.

For complete information about using lookup fields in lookup keys, see "About Lookup Keys" on page 399.

# **Reference Index**

<u>ooGreaterThanLookupField</u>

Constructs a new lookup field for testing whether values of the specified key field are greater than the specified comparison value.

# Constructors

## ooGreaterThanLookupField

Constructs a new lookup field for testing whether values of the specified key field are greater than the specified comparison value.

- 1. ooGreaterThanLookupField(
   const ooKeyField &field,
   const void \*valuePtr);
- 2. ooGreaterThanLookupField(
   const ooTypeNumber typeN,
   const char \*memberName,
   const void \*valuePtr);

#### Parameters

field

Key-field object representing a key field of the indexed class. The values of the specified key field will be tested by this lookup field.

#### typeN

Type number of the indexed class.

#### memberName

Name of a data member that serves as a key field of the indexed class. The values of the specified key field will be tested by this lookup field.

The specified data member must be defined or inherited by the class specified by *typeN*. You can qualify the name of an inherited data member using the following notation (where *baseClassName* is the name of the base class that defines the inherited data member):

baseClassName::dataMemberName

You *must* qualify the name of an inherited data member if the member name is ambiguous (for example, the same name is defined in both the base class and the class to be indexed) or if the member name is not visible due to access control.

#### valuePtr

Data value to which key-field values are to be compared. The type of this value must match the key-field type (the type of the data member you specified to the key-field object). Using data of any other type may have unpredictable results.

Discussion This lookup field causes the values of the specified key field to be tested when an iterator scans the index; an indexed object is found if its tested value is greater than the value specified by *valuePtr*.

This lookup field will be ignored if *field* does not specify a key field of the index being searched.

# ooHashAdmin Class

Inheritance: ooObj->ooAdmin->ooHashAdmin

Handle Class: ooHandle(ooHashAdmin)

**Object-Reference Class:** ooRef(ooHashAdmin)

The persistence-capable class ooHashAdmin represents hash administrators.

See:

- "Reference Summary" on page 270 for an overview of member functions
- "Reference Index" on page 270 for a list of member functions

To use this class, your application must include the ooCollections.h header file. For UNIX linking information, see *Installation and Platform Notes for UNIX*.

## **About Hash Administrators**

Each unordered collection derived from <code>ooHashSet</code> has a hash administrator that manages the containers used by the collection's internal objects, namely the hash buckets of the collection's extendible hash table. An unordered collection's hash administrator is created when the collection itself is created.

A hash administrator has a property that you can set to control when the current hash-bucket container is considered "full." The *maximum buckets per container* property specifies how many hash buckets can be clustered together in the same container. It is typical for a hash bucket to be updated frequently. The default value for this property is 1, which minimizes lock conflicts. If you know that a particular collection will be used by a single user, locking is not an issue. In that case, a larger value may be appropriate for the collection's hash administrator.

For additional information, see "Hash Administrator" on page 253 in the Objectivity/C++ programmer's guide.

# Working With a Hash Administrator

Like other persistent objects, hash administrators are normally manipulated through handles or object references.

You call an unordered collection's admin member function to obtain an object reference to the collection's hash administrator; you must then cast the returned object reference to type <code>ooRef(ooHashAdmin)</code> before you access the hash administrator's data member.

# **Reference Summary**

Getting Information	<u>bucketContainer</u> maxBucketsPerContainer
Setting Information	setMaxBucketsPerContainer

# **Reference Index**

<u>bucketContainer</u>	Gets this hash administrator's current hash-bucket container.
<u>maxBucketsPerContainer</u>	Gets the maximum number of hash buckets per container for this hash administrator.
<u>setMaxBucketsPerContainer</u>	Sets the maximum number of hash buckets per container for this hash administrator.

# **Member Functions**

## bucketContainer

Gets this hash administrator's current hash-bucket container.

ooRef(ooContObj) bucketContainer();

Returns Object reference to this hash administrator's current hash-bucket container.

## maxBucketsPerContainer

Gets the maximum number of hash buckets per container for this hash administrator.

ooInt32 maxBucketsPerContainer();

Returns The maximum number of hash buckets that can be stored in a single container.

## setMaxBucketsPerContainer

Sets the maximum number of hash buckets per container for this hash administrator.
 void setMaxBucketsPerContainer(ooInt32 max);
 Parameters max

 The maximum number of hash buckets that can be stored in a single container.

 Discussion Changing the maximum buckets per container affects only the collection's current hash-bucket container and any hash-bucket containers created in the future. If you increase the number of hash buckets per containing only one hash bucket each.

# ooHashMap Class

Inheritance: ooObj->ooCollection->ooHashSet->ooHashMap

Handle Class: ooHandle(ooHashMap)

**Object-Reference Class:** ooRef(ooHashMap)

The persistence-capable class ooHashMap represents unordered object maps.

See:

- "Reference Summary" on page 274 for an overview of member functions
- "Reference Index" on page 275 for a list of member functions

To use this class, your application must include the ooCollections.h header file. For UNIX linking information, see *Installation and Platform Notes for UNIX*.

## **About Unordered Object Maps**

An object map is a collection of key-value pairs; each key and each value is a persistent object. No two elements of the object map may have the same key. As the name implies, each element of an object map is a mapping from its key object to its value object.

Unordered object maps are *scalable* collections, that is, they can increase in size with minimal performance degradation. They are implemented using an extendible hashing mechanism so that elements can be added, deleted, and retrieved efficiently. The hash value for each element is computed from its key.

The hash values for keys of an unordered object map are computed by the object map's corresponding comparator. If an unordered object map has a default comparator, the hash values are computed from the object identifiers (OIDs) of the keys.

For additional information, see Chapter 11, "Persistent Collections," in the Objectivity/C++ programmer's guide.

# Working With an Unordered Object Map

As is the case for any basic object, you specify whether an unordered object map is to be transient or persistent when you create it; unordered object maps *must be persistent*. You create an unordered object map with a call to the new operator; the clustering directive in that call specifies where in the federated database to store the new unordered object map.

Like other persistent objects, unordered object maps are normally manipulated through handles or object references. You can store and find an unordered object map in the database just as you would any other persistent object.

# **Related Classes**

Two additional classes represent persistent collections of key-value pairs:

- OOTreeMap represents a *sorted* object map. It is implemented using a B-tree data structure.
- OOMap represents an unordered name map, that is, a collection of key-value pairs in which *the key is a string* and the value is an object reference to a persistent object. It uses a traditional (non-extendible) hashing mechanism.

# **Reference Summary**

In the following table:

- Operators indicated as *(inherited)* are overloaded in this class with no change in behavior; they are documented with the <u>ooObj</u> class (page 431), along with the other inherited member functions not listed here.
- Member function indicated as (inherited) are inherited from the <u>ooHashSet</u> class (page 283) or the <u>ooCollection</u> class (page 173) and are documented with the defining class.

Creating and Deleting	<u>ooHashMap</u> <u>operator new</u> (inherited) <u>operator delete</u> (inherited)
Adding, Removing, and Changing Elements	add addAll clear (inherited) put remove removeAll (inherited) removeAllDeleted retainAll (inherited)
Getting Elements	get <u>keyIterator</u> valueIterator
Getting Information	<u>hashOf</u> (inherited) <u>size</u> (inherited)
Finding Auxiliary Objects	<u>admin</u> (inherited) comparator (inherited)
Testing	<u>containsKey</u> <u>containsValue</u> <u>containsAll</u> (inherited) <u>isEmpty</u> (inherited)
Viewing in an MROW Transaction	refresh (inherited)

# **Reference Index**

add	Adds the specified object to this unordered object map.
<u>addAll</u>	Adds all elements in the specified collection to this unordered object map.
<u>containsKey</u>	Tests whether this unordered object map contains an element with the specified key.
<u>containsValue</u>	Tests whether this unordered object map contains an element with the specified value.
get	Finds the value paired with the specified key in this unordered object map.

<u>keyIterator</u>	Initializes a scalable-collection iterator to find all the keys of this unordered object map.
<u>ooHashMap</u>	Constructs a new empty unordered object map.
put	Maps the specified key to the specified value in this unordered object map.
remove	Removes the element, if any, with the specified key from this unordered object map.
removeAllDeleted	Removes from this unordered object map all elements in which either the key or the value has been deleted from the federated database.
valueIterator	Initializes a scalable-collection iterator to find all values in this unordered object map.

# Constructors

## ooHashMap

Constructs a new empty unordered object map.

```
1. ooHashMap(
    int bucketSize = 30011,
    int initialBuckets = 1,
    ooHandle(ooContObj) contAdminH = 0,
    ooHandle(ooContObj) contBucketH = 0);
2. ooHashMap(
    ooHandle(ooCompare) &compH,
    int bucketSize = 30011,
    int initialBuckets = 1,
    ooHandle(ooContObj) contAdminH = 0,
    ooHandle(ooContObj) contBucketH = 0);
```

Parameters bucketSize

The size of a hash bucket in the new unordered object map's hash table. The size of a hash bucket is the number of elements that can be hashed into each bucket. For optimal performance, the bucket size should be a prime number. If you specify a bucket size that is not a prime number, the next higher prime number is computed and used as the actual bucket size.

#### initialBuckets

The minimum number of initial hash buckets to create for the new unordered object map. The actual number of hash buckets created is the smallest power of 2 that is greater than *initialBuckets*. For example, if *initialBuckets* is 10, the number of hash buckets is 2\*\*4, or 16.

Preallocating multiple hash buckets increases the speed of adding and finding map elements. If each hash bucket is stored in a separate container (the default behavior), preallocating hash buckets also reduces the chance of lock conflicts. However, an unordered object map with a large number of initial hash buckets requires more disk space, more memory for the directory, and more time to create.

By default, each initial hash bucket is stored in its own newly created container. If you specify a container in the *contBucketH* parameter, however, all initial hash buckets are instead created in that container. In that case, you reduce the storage overhead for the object map, but you also reduce concurrent access.

#### contAdminH

Handle to the container in which to store the hash administrator for the new unordered object map. The default value (0) creates a container for the hash administrator in the same database as the unordered object map itself.

#### contBucketH

Handle to the container in which to store the initial hash bucket(s) for the new unordered object map. The default value (0) creates a separate container for each initial hash bucket, all in the same database as the unordered object map itself.

#### сотрН

Handle to the comparator for the new unordered object map; must be an instance of an application-specific derived class of ooCompare.

# Discussion Variant 1 creates an empty unordered object map with a default comparator and the specified bucket size.

Variant 2 creates an empty unordered object map with the specified comparator and the specified bucket size.

# **Member Functions**

## add

	Adds the specified object to this unordered object map.
	<pre>virtual ooBoolean add(     const ooHandle(ooObj) &amp;objH);</pre>
Parameters	оbjн Handle to the object to be added.
Returns	oocTrue if an element was added; otherwise, oocFalse.
Discussion	This member function adds a new key-value pair to this unordered object map with the specified object as its key and a null value. If this unordered object map currently has an element whose key is the specified object, the value of the existing element is replaced by null.
See also	addAll remove
addAll	
	Adds all elements in the specified collection to this unordered object map.
	<pre>virtual ooBoolean addAll(     const ooHandle(ooCollection) &amp;collectionH);</pre>
Parameters	<i>collectionH</i> Handle to the unordered or sorted object map whose elements are to be added to this unordered object map.
Returns	oocTrue if any elements were added; otherwise, oocFalse.
Discussion	If the specified collection is an object map, its elements are added to this unordered object map. If this unordered object map currently has an element with the same key as an element of the specified collection, the existing element is replaced by the element of the specified collection.
	If the specified collection is a collection of persistent objects, each of its elements is added as a key to this unordered object map; a null value is paired with each key. If this unordered object map currently has an element whose key is an

element of the specified collection, the value of the existing element is replaced by null.

See also <u>add</u>

#### containsKey

Tests whether this unordered object map contains an element with the specified key.

1.	<pre>ooBoolean containsKey(     const ooHandle(ooObj) &amp;keyH) const;</pre>
2.	ooBoolean containsKey(
	const void * <i>lookupVal</i> ) const;

Parameters keyH

Handle to the key to be tested for containment in this unordered object map.

lookupVal

Pointer to data that identifies the key to be tested for containment in this unordered object map.

- Returns occTrue if this unordered object map contains an element with the specified key; otherwise, occFalse.
- Discussion You can call this member function to check whether this unordered object map maps the specified key to some value.

Variant 2 tests whether any key is "equal" to the specified lookup data, as determined by the comparator for this unordered object map. It is useful if this unordered object map has an application-defined comparator that can identify a key based on class-specific data.

See also <u>containsValue</u>

## containsValue

Tests whether this unordered object map contains an element with the specified value.

Parameters valueH

Handle to the value to be tested for containment in this unordered object map.

Returns	oocTrue if this unordered object map contains an element whose value is the specified object; otherwise, oocFalse.
Discussion	You can call this member function to check whether this unordered object map maps at least one key to the specified value.
See also	containsKey
get	
	Finds the value paired with the specified key in this unordered object map.
	<pre>1. ooRef(ooObj) get(const ooHandle(ooObj) &amp;keyH) const;</pre>
	<pre>2. ooRef(ooObj) get(const void *lookupVal) const;</pre>
Parameters	keyH
	Handle to the key to be looked up.
	lookupVal
	Pointer to data that identifies the desired key.
Returns	Object reference to the value in the element with the specified key, or a null object reference if this unordered object map contains no mapping for that key.
Discussion	A return value of null does not necessarily indicate that no element has the specified key. It is possible that this unordered object map explicitly maps the key to null. You can use the <u>containsKey</u> member function to distinguish these two cases.
	Variant 2 finds the element whose key is "equal" to the specified lookup data, as determined by the comparator for this unordered object map. It is useful if this unordered object map has an application-defined comparator that can identify a key based on class-specific data.
See also	put addAll

## keylterator

Initializes a scalable-collection iterator to find all the keys of this unordered object map.

```
virtual ooCollectionIterator *keyIterator() const;
```

Returns	A pointer to a scalable-collection iterator for finding the keys of this unordered object map; the caller is responsible for deleting the iterator when it is no longer needed.
Discussion	The returned iterator finds the keys in an undefined order; however, it iterates through the key-value pairs of this unordered object map in the same order as does an iterator returned by <u>valueIterator</u> .
	You must delete the iterator when you have finished using it.
put	
	Maps the specified key to the specified value in this unordered object map.
	ooStatus put( const ooHandle(ooObj) & <i>keyH</i> , const ooHandle(ooObj) & <i>valueH</i> );
Parameters	<i>keyH</i> Handle to the key.
	<i>valueH</i> Handle to the value.
Returns	oocSuccess if successful; otherwise, oocError.
Discussion	If this unordered object map already contains an element with the specified key, this member function replaces the value in that element. Otherwise, this member function adds a new element with the specified key and value.
See also	<u>get</u> <u>addAll</u>
remove	
	Removes the element, if any, with the specified key from this unordered object map.
	<pre>virtual ooBoolean remove(</pre>
Parameters	keyH Handle to the key of the element to be removed.
Returns	oocTrue if an element was removed; otherwise, oocFalse.
See also	add

## removeAllDeleted

Removes from this unordered object map all elements in which either the key or the value has been deleted from the federated database.

virtual void removeAllDeleted();

Discussion You can call this member function to restore this unordered object map's referential integrity.

#### valuelterator

Initializes a scalable-collection iterator to find all values in this unordered object map.

ooCollectionIterator \*valueIterator() const;

- Returns A pointer to a scalable-collection iterator for finding all the persistent objects used as values in elements of this unordered object map.
- Discussion The returned iterator finds the keys in an undefined order; however, it iterates through the key-value pairs of this unordered object map in the same order as does an iterator returned by <u>keyIterator</u>.

You must delete the iterator when you have finished using it.

# ooHashSet Class

Inheritance: ooObj->ooCollection->ooHashSet

Handle Class: ooHandle(ooHashSet)

```
Object-Reference Class: ooRef(ooHashSet)
```

The persistence-capable class <code>ooHashSet</code> represents *unordered* sets of persistent objects with no duplicate elements.

See:

- "Reference Summary" on page 284 for an overview of member functions
- "Reference Index" on page 285 for a list of member functions

To use this class, your application must include the ooCollections.h header file. For UNIX linking information, see *Installation and Platform Notes for UNIX*.

## **About Unordered Sets**

An unordered set is an unordered collection of persistent objects with no duplicate elements. Unordered sets are *scalable* collections, that is, they can increase in size with minimal performance degradation. They are implemented using an extendible hashing mechanism so that elements can be added, deleted, and retrieved efficiently. The hash value for each element is computed from the element itself.

The hash values for elements of an unordered set are computed by the set's corresponding comparator. If an unordered set has a default comparator, the hash values are computed from the OIDs of the elements.

For additional information, see Chapter 11, "Persistent Collections," in the Objectivity/C++ programmer's guide.

# Working With an Unordered Set

As is the case for any basic object, you specify whether an unordered set is to be transient or persistent when you create it; unordered sets *must be persistent*. You create an unordered set with a call to the new operator; the clustering directive in that call specifies where in the federated database to store the unordered set.

Like other persistent objects, unordered sets are normally manipulated through handles or object references. You can store and find an unordered set in the database just as you would any other persistent object.

# **Related Classes**

Two additional classes represent persistent collections of persistent objects:

- ooTreeSet represents a *sorted* collection of persistent objects with no duplicate elements.
- ooTreeList represents a collection of persistent objects that are maintained in the order specified when they are added to the collection. A list can contain duplicate elements.

Both sorted sets and lists are implemented using a B-tree data structure.

# **Reference Summary**

In the following table:

- Operators indicated as *(inherited)* are overloaded in this class with no change in behavior; they are documented with the <u>ooObj</u> class (page 431), along with the other inherited member functions not listed here.
- Member function indicated as *(inherited)* are inherited from the <u>ooCollection</u> class and are documented with that class (page 173).

Creating and Deleting	<u>ooHashSet</u> <u>operator new</u> (inherited) <u>operator delete</u> (inherited)
Adding and Removing Elements	<u>add</u> <u>addAll (inherited)</u> <u>clear (inherited)</u> <u>removeAll (inherited)</u> <u>removeAllDeleted (inherited)</u> <u>retainAll (inherited)</u>
Getting Elements	<u>get</u> <u>iterator</u>
Getting Information	hashOf size
Finding Auxiliary Objects	<u>admin</u> <u>comparator</u>
Testing	<u>contains</u> <u>containsAll</u> (inherited) <u>isEmpty</u>
Viewing in an MROW Transaction	refresh

# **Reference Index**

add	Adds the specified object to this unordered set.
<u>admin</u>	Gets the hash administrator for this unordered collection.
<u>comparator</u>	Finds the comparator for this unordered collection.
<u>contains</u>	Tests whether this unordered set contains the specified object.
get	Finds the specified element of this unordered set.
<u>hash0f</u>	Gets the hash code for the specified object in this unordered collection.
isEmpty	Tests whether this unordered collection is empty.
<u>iterator</u>	Initializes a scalable-collection iterator to find the elements of this unordered set.

<u>ooHashSet</u>	Constructs a new empty unordered set.
refresh	Refreshes each container used internally by this unordered collection, except for the container in which the unordered collection itself is stored.
remove	Removes the specified object from this unordered set.
<u>size</u>	Gets the size of this unordered collection.

# Constructors

## ooHashSet

Constructs a new empty unordered set.

	<pre>1. ooHashSet( int bucketSize = 30011, int initialBuckets = 1, ooHandle(ooContObj) contAdminH = 0, ooHandle(ooContObj) contBucketH = 0);</pre>
	<pre>2. ooHashSet( ooHandle(ooCompare) &amp; compH, int bucketSize = 30011, int initialBuckets = 1, ooHandle(ooContObj) contAdminH = 0, ooHandle(ooContObj) contBucketH = 0);</pre>
Parameters	bucketSize The size of a hash bucket in the new unordered set's hash table. The size of a hash bucket is the number of elements that can be hashed into each bucket. For optimal performance, the bucket size should be a prime number. If you specify a bucket size that is not a prime number, the next higher prime number is computed and used as the actual bucket size.

```
initialBuckets
```

The minimum number of initial hash buckets to create for the new unordered set. The actual number of hash buckets created is the smallest power of 2 that is greater than *initialBuckets*. For example, if *initialBuckets* is 10, the number of hash buckets is 2\*\*4, or 16.

Preallocating multiple hash buckets increases the speed of adding and finding elements. If each hash bucket is stored in a separate container (the default behavior), preallocating hash buckets also reduces the chance of lock conflicts. However, an unordered set with a large number of initial hash buckets requires more disk space, more memory for the directory, and more time to create.

By default, each initial hash bucket is stored in its own newly created container. If you specify a container in the *contBucketH* parameter, however, all initial hash buckets are instead created in that container. In that case, you reduce the storage overhead for the set, but you also reduce concurrent access.

All initial hash buckets are created in the container indicated by *contBucketH*.

contAdminH

Handle to the container in which to store the hash administrator for the new unordered set. The default value (0) creates a container for the hash administrator in the same database as the unordered set itself.

#### contBucketH

Handle to the container in which to store the initial hash bucket(s) for the new unordered set. The default value (0) creates a separate container for each initial hash bucket, all in the same database as the unordered set itself.

#### сотрН

Handle to the comparator for the new unordered set; must be an instance of an application-specific derived class of ooCompare.

# Discussion Variant 1 creates an empty unordered set with a default comparator and the specified bucket size.

Variant 2 creates an empty unordered set with the specified comparator and the specified bucket size.

## **Member Functions**

#### add

#### Adds the specified object to this unordered set.

	<pre>virtual ooBoolean add(     const ooHandle(ooObj) &amp;objH);</pre>
Parameters	оbјн Handle to the object to be added.
Returns	oocTrue if an element was added; otherwise, oocFalse.

Discussion This member function returns false if the specified object is already an element of this unordered set.

See also <u>remove</u>

## admin

	Gets the hash administrator for this unordered collection.
	<pre>virtual ooRef(ooAdmin) admin() const;</pre>
Returns	Object reference to the hash administrator for this unordered collection.
Discussion	You typically call this member function when you want to change the way that this unordered collection's internal objects (hash buckets) are assigned to containers. Before you do so, you must cast the returned object reference to <code>ooRef(ooHashAdmin)</code> .

#### comparator

Finds the comparator for this unordered collection.

virtual ooRef(ooCompare) comparator() const;

Returns Object reference to the comparator for this unordered collection, or null if this unordered collection has a default comparator.

## contains

Tests whether this unordered set contains the specified object.

# Parameters objH Handle to the element to be tested for containment in this unordered set. lookupVal Pointer to data that identifies the object to be tested for containment in this unordered set.

# Returns oocTrue if this unordered set contains an element equal to the specified object; otherwise, oocFalse.
Discussion	Variant 2 tests whether any element is "equal" to the specified lookup data, as determined by the comparator for this unordered set. It is useful if this unordered set has an application-defined comparator that can identify an element based on class-specific data.	
get		
	Finds the specified element of this unordered set.	
	<pre>virtual ooRef(ooObj) get(const void *lookupVal) const;</pre>	
Parameters	<i>lookupVal</i> Pointer to data that identifies the desired element.	
Returns	Object reference to the element that is "equal" to the specified lookup data, as determined by the comparator for this unordered set, or a null object reference if this unordered set does not contain such an element.	
Discussion	This member function is useful if this unordered set has an application-defined comparator that can identify an element based on class-specific data.	
hashOf		
	Gets the hash code for the specified object in this unordered collection.	
	<pre>int hashOf(const ooHandle(ooObj) &amp;objH) const;</pre>	
Parameters	objH	
	Handle to the object whose hash value is to be computed.	
Returns	The hash code for the specified object.	
Discussion	If this unordered collection has a comparator of an application-defined class, that comparator computes the hash code; otherwise, the hash code is computed from the object's OID.	
isEmpty		
	Tests whether this unordered collection is empty.	
	virtual ooBoolean isEmpty() const;	
Returns	oocTrue if this unordered collection has no elements; otherwise, oocFalse.	

### iterator

	Initializes a scalable-collection iterator to find the elements of this unordered set.
	<pre>virtual ooCollectionIterator *iterator() const;</pre>
Returns	A pointer to a scalable-collection iterator for finding the elements of this unordered set; the caller is responsible for deleting the iterator when it is no longer needed.
Discussion	The returned iterator finds the elements in an undefined order.
	You must delete the iterator when you have finished using it.
refresh	
	Refreshes each container used internally by this unordered collection, except for the container in which the unordered collection itself is stored.
	<pre>virtual ooStatus refresh(ooMode &amp;openMode) const;</pre>
Parameters	<ul> <li>openMode</li> <li>Intended level of access to each refreshed container:</li> <li>Specify oocRead to open the container for read. This implicitly requests a read lock on the container.</li> <li>Specify oocUpdate to open the container for update (read and write).</li> </ul>
	This implicitly requests an update lock on the container.
Returns	oocSuccess if every container can be refreshed; otherwise, oocError.
Discussion	You typically call this member function when you need to refresh your view of an unordered collection that you are reading in an MROW transaction. This member function calls <u>refreshOpen</u> on each container that is used internally by the unordered collection—that is, on the administrator and hash-bucket containers maintained by the unordered collection. This member function does not refresh the container in which the unordered collection itself is stored, nor does it necessarily refresh the containers that store the collection's elements.

### remove

Removes the specified object from this unordered set.

Parameters	оbјн Handle to the object to be removed.
Returns	oocTrue if an element was removed; otherwise, oocFalse.
See also	add
size	
	Gets the size of this unordered collection.
	<pre>virtual int size() const;</pre>
Returns	The number of elements in this unordered collection.

# ooltr(appClass) Class

Inheritance: ooHandle(appClass)->ooItr(appClass)

The non-persistence-capable class <code>ooltr(appClass)</code> represents an *object iterator* for finding instances of the application-defined persistence-capable class <code>appClass.appClass</code> can be either a basic-object class (derived from <code>ooObj</code>) or a container class (derived from <code>ooContObj</code>).

See:

- "Reference Summary" on page 294 for an overview of member functions
- "Reference Index" on page 295 for a list of member functions

To use the ooltr(*appClass*) class, you must include and compile with files generated by the DDL processor, as described in "Obtaining Generated Class Definitions" on page 294.

## **About Class-Specific Object Iterators**

When an application defines a persistence-capable class *appClass* and adds it to the federated-database schema, the DDL processor generates the corresponding iterator class <code>ooltr(appClass)</code>.

An object iterator of class <code>ooltr(appClass)</code> finds objects of class <code>appClass</code> and its derived classes.

- You can initialize the object iterator to find objects below a given storage object in the storage hierarchy. To do so, call the object iterator's scan member function.
- If *appClass* is the destination class for the to-many association *linkName*, you can initialize the object iterator to find the destination objects linked to a given source object by this association. To do so, cast the object iterator to type ooltr(ooObj) and pass it as a parameter when you call the *linkName* member function on the source object.

You advance a class-specific object iterator, use it to reference objects in the iteration set, and termination the iteration as you would with any object iterator. For more information, see the ooltr(ooObj) class.

### **Obtaining Generated Class Definitions**

To use the ooltr(*appClass*) class, you must include either the primary header file or the references header file generated by the DDL processor for *appClass*. Thus, if *appClass* is defined in the DDL file *classDefFile*.ddl, you must include one of the following files:

- The primary header file classDefFile.h
- The references header file classDefFile\_ref.h

Furthermore, you must compile the method implementation file *classDefFile\_ddl.cxx* with your application code files.

For more information about DDL-generated files and how to use them, see the Objectivity/C++ Data Definition Language book.

### When appClass is a Template Class

When *appClass* is a persistence-capable template class with multiple parameters, the name of the generated object-iterator class contains the symbol OO\_COMMA to separate the template parameters. For example, for a persistence-capable template class Example<Float, Node>, the generated object-iterator class is ooltr(Example<Float OO\_COMMA Node>). This is because the macro syntax of the generated class name interprets embedded commas as separators between the as macro parameters instead of as separators between the template parameters.

# **Reference Summary**

Creating an Object Iterator	<u>ooItr(appClass)</u>
Finding Objects	<u>end</u> <u>next</u> <u>scan</u>

## **Reference Index**

end	Explicitly terminates iteration by this object iterator.
<u>next</u>	Advances this object iterator to the next object in the iteration set.
<u>ooItr(appClass)</u>	Default constructor that constructs a new uninitialized object iterator for finding instances of $appClass$ and its derived classes.
<u>scan</u>	Initializes this object iterator to find all instances of appClass and its derived classes in the specified storage object, satisfying the specified condition, if any.

## Constructors

### ooltr(appClass)

Default constructor that constructs a new uninitialized object iterator for finding instances of *appClass* and its derived classes.

```
ooItr(appClass)();
```

## **Member Functions**

### end

	Explicitly terminates iteration by this object iterator.
	ooStatus end();
Discussion	You can call this member function to signal that you are finished using this object iterator even though it has not yet found all the objects in the iteration set.
next	
	Advances this object iterator to the next object in the iteration set.

ooBoolean next();

Returns occTrue if another object is found; occFalse if all objects in the iteration set have been found or if an error occurred.

#### scan

Initializes this object iterator to find all instances of *appClass* and its derived classes in the specified storage object, satisfying the specified condition, if any.

1.	<pre>ooStatus scan(     const ooRefHandle(ooObj) &amp;storageObject,     const ooMode openMode = oocNoOpen);</pre>
2.	<pre>ooStatus scan(     const ooRefHandle(ooObj) &amp;storageObject,     const ooLookupKey &amp;lookupKey,     const ooMode openMode = oocNoOpen);</pre>
3.	<pre>ooStatus scan(     const ooRefHandle(ooObj) &amp;storageObject,     const char *predicate);</pre>
4.	<pre>ooStatus scan(    const ooRefHandle(ooObj) &amp;storageObject,    const ooMode openMode,    const ooAccessMode access,    const char *predicate);</pre>

#### Parameters storageObject

Object reference or handle to the container, database, or federated database to be scanned. If you specify a basic object, its container is scanned.

If *appClass* is a container class, and *storageObject* specifies a container or basic object, this object iterator is set to null (initialized with an empty iteration set).

(FTO) You may not scan an autonomous partition.

#### lookupKey

Lookup key representing the condition that each found object must satisfy.

#### openMode

Intended level of access to each object found by the next member function:

- oocNoOpen (the default in variants 1 and 2) causes next to obtain a reference to the next object without opening it.
- oocRead causes next to open the next object for read.
- oocUpdate causes next to open the next object for update.

access

Limits the data members that can be specified in *predicate*:

- oocPublic permits the predicate to test only public data members, thus preserving encapsulation.
- oocAll permits the predicate to test any data member, thus decreasing encapsulation. To preserve maximum encapsulation, you should specify oocAll only within member functions of the class you are querying.

#### predicate

Condition that objects must satisfy to be found by this object iterator. This string must be a valid expression in the predicate query language. This object iterator finds only those objects that satisfy the *predicate*.

Returns oocSuccess if successful; otherwise oocError.

Discussion This scan operation initializes this object iterator to find objects of class *appClass* or any class derived from *appClass*. If no such objects exist in the specified storage object, this object iterator is set to null and scan returns oocSuccess. (This object iterator's next member function will return oocFalse, however.)

Variant 3 finds persistent objects without opening them; the other variants may open the found objects, as specified by the <code>openMode</code> parameter.

If you specify the *predicate* parameter and the specified storage object has an index over the appropriate objects, scan can use the index to optimize its search. (You may need to enable index usage; see the <u>ooUseIndex</u> global function.) If no index is available, scan finds objects by inspection.

An error is signaled if *predicate* tries to test a non-existent data member, or if *predicate* tries to test a protected or private data member when the *access* parameter is oocPublic.

If you specify the *lookupKey* parameter, scan finds objects *only* if a compatible index exists in specified storage object; otherwise, this object iterator is set to null. For performance reasons, you should specify the *lookupKey* parameter only if you know the specified storage object contains such an index. You can test the storage object for compatible indexes by calling the <code>oolookupKey::anyIndex</code> member function. (Indexes need not be enabled to scan with a lookup key.)

See also <u>next</u>

# ooltr(ooAPObj) Class

Inheritance: **ooItr(ooAPObj)** 

The non-persistence-capable class <code>ooltr(ooAPObj)</code> represents an *object iterator* for finding the autonomous partitions in a federated database.

See:

- "Reference Summary" on page 300 for an overview of member functions
- "Reference Index" on page 300 for a list of member functions

## **About Autonomous-Partition Iterators**

An iterator of class <code>ooltr(ooAPObj)</code> finds the autonomous partitions of the open federated database. You can initialize an autonomous-partition iterator in either of two ways:

- Call its <u>scan</u> member function, passing a handle to the federated database as a parameter.
- Call the contains member function on a handle to the federated database, passing the autonomous-partition iterator as a parameter.

You advance an autonomous-partition iterator, use it to reference autonomous partitions in the iteration set, and termination the iteration as you would with any object iterator. For more information, see the ooltr(oolbj) class.

# **Reference Summary**

Creating an Object Iterator	<u>ooItr(ooAPObj)</u>
Finding Autonomous Partitions	end next scan

## **Reference Index**

end	Explicitly terminates iteration by this object iterator.
next	Advances this object iterator to the next autonomous partition in the iteration set.
<u>ooItr(ooAPObj)</u>	Default constructor that constructs a new uninitialized object iterator for finding autonomous partitions.
<u>scan</u>	Initializes this object iterator to find all autonomous partitions in the specified federated database.

## **Constructors and Destructors**

## ooltr(ooAPObj)

Default constructor that constructs a new uninitialized object iterator for finding autonomous partitions.

```
ooItr(ooAPObj)();
```

## **Member Functions**

end

Explicitly terminates iteration by this object iterator.

```
ooStatus end();
```

Discussion	You can call this member function to signal that you are finished using this autonomous-partition iterator even though it has not yet found all the autonomous partitions in the iteration set.	
next		
	Advances this object iterator to the next autonomous partition in the iteration set.	
	ooBoolean next();	
Returns	oocTrue if another autonomous partition is found; oocFalse if all autonomous partitions in the iteration set have been found or if an error occurred.	
scan		
	Initializes this object iterator to find all autonomous partitions in the specified federated database.	
	<pre>ooStatus scan( const ooHandle(ooFDObj) &amp;federation, const ooMode openMode = oocNoOpen);</pre>	
Parameters	federation Handle to the federated database to be scanned for autonomous partitions.	
	openMode	
	Intended level of access to each autonomous partition found by the <code>next</code> member function:	
	<ul> <li>oocNoOpen (the default) causes next to obtain a reference to the next partition without opening it.</li> </ul>	
	<ul> <li>oocRead causes next to open the next partition for read.</li> </ul>	
	<ul> <li>oocUpdate causes next to open the next partition for update.</li> </ul>	
Returns	oocSuccess if successful; otherwise oocError.	
Discussion	If no autonomous partitions exist within the federated database, this object iterator is set to null and scan returns oocSuccess. (This object iterator's next member function will return oocFalse, however.)	
	Calling this member function on an autonomous-partition iterator is equivalent to passing the iterator to the <u>contains</u> member function on an object reference or handle to the federated database.	

# ooltr(ooContObj) Class

Inheritance: ooHandle(ooContObj)->ooItr(ooContObj)

The non-persistence-capable class <code>ooltr(ooContObj)</code> represents a *container iterator*—that is, an object iterator for finding groups of containers in a database or federated database.

See:

- "Reference Summary" on page 304 for an overview of member functions
- "Reference Index" on page 304 for a list of member functions

## **About Container Iterators**

An iterator of class <code>ooltr(ooContObj)</code> finds the containers in the open federated database or a particular database. You can initialize a container iterator in either of two ways:

- Call its <u>scan</u> member function, passing a handle to the federated database or a database as a parameter.
- Call the contains member function on a handle to a database, passing the container iterator as a parameter.

You advance a container iterator, use it to reference containers in the iteration set, and termination the iteration as you would with any object iterator. For more information, see the ooltr(ooObj) class.

**NOTE** Operations that initialize a container iterator find default containers (instances of <u>ooDefaultContObj</u>). When advancing such an iterator, you should test for the default container before performing any operation that deletes or moves the found container. For example, when the iterator references the next found container, you could call the container's inherited oolsKindOf member function.

# **Reference Summary**

Creating an Object Iterator	<u>ooItr(ooContObj)</u>
Finding Containers	<u>end</u> <u>next</u> scan

# **Reference Index**

end	Explicitly terminates iteration by this container iterator.
next	Advances this container iterator to the next container in the iteration set.
<u>ooItr(ooContObj)</u>	Default constructor that constructs a new uninitialized container iterator.
scan	Initializes this container iterator to find all containers in the specified storage object, satisfying the specified condition, if any.

## **Constructors and Destructors**

## ooltr(ooContObj)

Default constructor that constructs a new uninitialized container iterator.

```
ooltr(ooContObj)();
```

## **Member Functions**

end

Explicitly terminates iteration by this container iterator.

```
ooStatus end();
```

Discussion You can call this member function to signal that you are finished using this container iterator even though it has not yet found all the containers in the iteration set.

#### next

Advances this container iterator to the next container in the iteration set.

ooBoolean next();

Returns oocTrue if another container is found; oocFalse if all containers in the iteration set have been found or if an error occurred.

#### scan

Initializes this container iterator to find all containers in the specified storage object, satisfying the specified condition, if any.

1.	<pre>ooStatus scan(     const ooRefHandle(ooObj) &amp;storageObject,     const ooMode openmode = oocNoOpen);</pre>
2.	<pre>ooStatus scan(     const ooRefHandle(ooObj) &amp;storageObject,     const ooLookupKey &amp;lookupKey,     const ooMode openmode = oocNoOpen);</pre>
3.	<pre>ooStatus scan(     const ooRefHandle(ooObj) &amp;storageObject,     const char *predicate);</pre>
4.	<pre>ooStatus scan( const ooRefHandle(ooObj) &amp;storageObject, const ooMode openMode, const ooAccessMode access, const char *predicate);</pre>

Parameters storageObject

Object reference or handle to the database or federated database to be scanned for containers. If you specify a container or basic object, this container iterator is set to null (initialized with an empty iteration set). (*FTO*) You may not scan an autonomous partition.

#### lookupKey

Condition that each found container must satisfy.

#### openMode

Intended level of access to each container found by the next member function:

- oocNoOpen (the default in variants 1 and 2) causes next to obtain a reference to the next container without opening it.
- oocRead causes next to open the next container for read.
- oocUpdate causes next to open the next container for update.

#### access

Limits the data members that can be specified in *predicate*:

- oocPublic permits the predicate to test only public data members, thus preserving encapsulation.
- oocAll permits the predicate to test any data member, thus decreasing encapsulation. To preserve maximum encapsulation, you should specify oocAll only within member functions of the class you are querying.

#### predicate

Condition that each retrieved container must satisfy. This string must be a valid expression in the predicate query language. This container iterator finds only those containers that satisfy the *predicate*.

Returns oocSuccess if successful; otherwise oocError.

Discussion This scan operation initializes this container iterator to find objects of class ocContObj or any class derived from ocContObj. If no such objects exist in the specified storage object, this container iterator is set to null and scan returns ocCuccess. (This container iterator's next member function will return occFalse, however.)

Variant 3 finds containers without opening them; the other variants may open the found objects, as specified by the <code>openMode</code> parameter.

If you specify the *predicate* parameter and the specified storage object has an index over the appropriate objects, scan can use the index to optimize its search. (You may need to enable index usage; see the <u>ooUseIndex</u> global function.) If no index is available, scan finds objects by inspection.

An error is signaled if *predicate* tries to test a non-existent data member, or if *predicate* tries to test a protected or private data member when the *access* parameter is oocPublic.

If you specify the *lookupKey* parameter, the scan member function finds objects only if a compatible index exists in the specified scope; otherwise this container iterator is set to null. For performance reasons, you should specify the *lookupKey* parameter only if you know the specified scope contains such an index. You can test the scope for compatible indexes by calling the oolookupKey::anyIndex member function. (Indexes need not be enabled to scan with a lookup key.)

# ooltr(ooDBObj) Class

Inheritance: **ooHandle(ooObj)-**>**ooItr(ooDBObj)** 

The non-persistence-capable class <code>ooltr(ooDBObj)</code> represents a *database iterator*—that is, an object iterator for finding the databases in a federated database.

See:

- "Reference Summary" on page 310 for an overview of member functions
- "Reference Index" on page 310 for a list of member functions

## **About Database Iterators**

An iterator of class ooltr(ooDBObj) finds the databases in the open federated database. You can initialize a database iterator in either of two ways:

- Call its <u>scan</u> member function, passing a handle to the federated database as a parameter.
- Call the contains member function on a handle to the federated database, passing the database iterator as a parameter.

You advance a database iterator, use it to reference databases in the iteration set, and termination the iteration as you would with any object iterator. For more information, see the ooltr(oobj) class.

## **Reference Summary**

Creating an Object Iterator	<u>ooItr(ooDBObj)</u>
Finding Databases	end next scan

# **Reference Index**

end	Explicitly terminates iteration by this database iterator.
next	Advances this database iterator to the next database in the iteration set.
<u>ooItr(ooDBObj)</u>	Default constructor that constructs a new uninitialized database iterator.
scan	Initializes this database iterator to find all databases in the specified federated database.

# **Constructors and Destructors**

## ooltr(ooDBObj)

Default constructor that constructs a new uninitialized database iterator.

```
ooItr(ooDBObj)();
```

# **Member Functions**

end

Explicitly terminates iteration by this database iterator.

ooStatus end();

Discussion	You can call this member function to signal that you are finished using this database iterator even though it has not yet found all the databases in the iteration set.
next	
	Advances this database iterator to the next database in the iteration set.
	ooBoolean next();
Returns	occTrue if another database is found; occFalse if all databases in the iteration set have been found or if an error occurred.
scan	
	Initializes this database iterator to find all databases in the specified federated database.
	<pre>ooStatus scan( const ooHandle(ooFDObj) &amp;federation, const ooMode openmode = oocNoOpen);</pre>
Parameters	federation
	Handle to the federated database to be scanned for databases.
	openMode
	Intended level of access to each database found by the $mext$ member function:
	<ul> <li>oocNoOpen (the default) causes next to obtain a reference to the next database without opening it.</li> </ul>
	<ul> <li>oocRead causes next to open the next database for read.</li> </ul>
	<ul> <li>oocUpdate causes next to open the next database for update.</li> </ul>
Returns	oocSuccess if successful; otherwise oocError.
Discussion	If no databases exist within the federated database, this database iterator is set to null and scan returns oocSuccess. (The database iterator's next member function will return oocFalse, however.)
	Calling this member function on a database iterator is equivalent to passing the database iterator to the <u>contains</u> member function on an object reference or handle to the federated database.

# ooltr(ooObj) Class

Inheritance: **ooHandle(ooObj)->ooItr(ooObj)** 

The non-persistence-capable class ooltr(ooObj) represents an *object iterator* for finding Objectivity/DB objects—that is, instances of ooObj or classes derived from ooObj. The ooltr(ooObj) class is also the base class for all of object-iterator classes.

See:

- "Reference Summary" on page 315 for an overview of member functions
- "Reference Index" on page 315 for a list of member functions

## **About Object Iterators**

An object iterator steps through a group of objects found in the federated database. During a transaction, you can create and initialize an object iterator to find a specified group of Objectivity/DB objects. For example, you can use an object iterator to find all the basic objects in a particular container. The group of objects to be found is called an *iteration set*.

An object iterator consists of:

- A description of an iteration set and the object iterator's current location in this set.
- Other information, such as the intended level of access to each found object.
- A handle for referencing the object at the current location in the iteration set.

When you create a null object iterator, its state is undefined and its handle is null. Initializing the object iterator identifies an iteration set and locates the object iterator just before the first object in the set. If the set is nonnull, the first call to the <u>next</u> member function advances the object iterator to the first object in the set and initializes the handle to reference it. Successive calls to next advance through the set, so that the handle references each object in turn. An object

iterator makes a single pass through the iteration set, returning the objects in the set in an undefined order. When the end of the iteration set is reached, the next member function returns oocFalse.

An application can initialize an object iterator of type ooltr(ooObj) in a number of ways. Most Objectivity/C++ operations that initialize iterators of this type do so to find persistent objects (containers and basic objects), although occasionally, such iterators are initialized to find Objectivity/DB objects of any class. The most common ways of initializing an object iterator include:

- Calling the object iterator's <u>scan</u> member function to find the persistent objects that reside in a particular storage object.
- Passing the object iterator to the *linkName* member function on a persistent object to find the destination objects of a to-many association.
- Passing the object iterator to the contains member function on a container to find the basic objects in the container.

Note that initializing an object iterator does not result in an intermediate collection in memory. Objects can be added, moved, or deleted from the database while the object iterator is active, and such changes can affect the set of objects returned by the object iterator.

Iteration is terminated automatically after the object iterator has found all objects in an iteration set. An object iterator is valid only during the transaction in which it was initialized. Committing, aborting, or checkpointing the transaction terminates the iteration automatically, even if the iteration set has not yet been exhausted.

If you finish using a particular object iterator without advancing through the entire iteration set, you can terminate the iteration explicitly by calling its <u>end</u> member function. Doing so signals that you will not use the object iterator again and that its data structures can be deleted.

Terminating the iteration makes the object iterator a null iterator, which has no iteration set. After iteration has terminated, you should not attempt to use the object iterator without reinitializing it. If you do so, an error occurs.

## **Related Classes**

In addition to object iterators of class ooltr(className), Objectivity/C++ provides three other kinds of iterators:

■ A name-map iterator is an instance of OoMapItr; it steps through the key-value pairs in a name map.

- A scalable-collection iterator is an instance of a class derived from ooCollectionIterator; it steps through the objects in a scalable persistent collection.
- A VArray iterator is an instance of a class created from the class template d\_Iterator<element\_type>; it steps through the elements of a VArray. The element\_type parameter specifies the type of elements in the VArray.

## **Reference Summary**

Creating an Object Iterator	<u>ooItr(ooObj)</u>
Finding Objects	end next scan

## **Reference Index**

end	Explicitly terminates iteration by this object iterator.
next	Advances this object iterator to the next Objectivity/DB object in the iteration set.
<u>ooItr(ooObj)</u>	Default constructor that constructs a new uninitialized object iterator for finding Objectivity/DB objects.
scan	Initializes this object iterator to find all persistent objects in the specified storage object, satisfying the specified condition, if any.

## Constructors

### ooltr(ooObj)

Default constructor that constructs a new uninitialized object iterator for finding Objectivity/DB objects.

```
ooItr(ooObj)();
```

## **Member Functions**

### end Explicitly terminates iteration by this object iterator. ooStatus end(); Discussion You can call this member function to signal that you are finished using this object iterator even though it has not yet found all the objects in the iteration set. next Advances this object iterator to the next Objectivity/DB object in the iteration set. ooBoolean next(); Returns oocTrue if another object is found; oocFalse if the iteration set is null, if all the objects in the iteration set have been found, or if an error occurred. scan Initializes this object iterator to find all persistent objects in the specified storage object, satisfying the specified condition, if any. 1. ooStatus scan( const ooRefHandle(ooObj) &storageObject, const ooMode openmode = oocNoOpen); 2. ooStatus scan( const *ooRefHandle*(ooObj) & *storageObject*, const ooLookupKey &lookupKey, const ooMode openmode = oocNoOpen); 3. ooStatus scan( const ooRefHandle(ooObj) &storageObject, const char \*predicate); ooStatus scan( 4. const *ooRefHandle*(ooObj) & *storageObject*, const ooMode openMode, const ooAccessMode access, const char \*predicate); Parameters storageObject Object reference or handle to the container, database, or federated database to be scanned for persistent objects:

- If you specify a container, the iteration set includes all basic objects in the container.
- If you specify a database or federated database, the iteration set includes all persistent objects (that is, basic objects and containers) in the database or federated database.

If you specify a basic object, its container is scanned for basic objects.

(FTO) You may not scan an autonomous partition.

#### lookupKey

Lookup key that represents the condition that each found object must satisfy.

#### openMode

Intended level of access to each object found by the next member function:

- oocNoOpen (the default in variants 1 and 2) causes next to obtain a reference to the next object without opening it.
- oocRead causes next to open the next object for read.
- oocUpdate causes next to open the next object for update.

#### access

Limits the data members that can be specified in *predicate*:

- oocPublic permits the predicate to test only public data members, thus preserving encapsulation.
- oocAll permits the predicate to test any data member, thus decreasing encapsulation. To preserve maximum encapsulation, you should specify oocAll only within member functions of the class you are querying.

#### predicate

Condition that each found object must satisfy. This string must be a valid expression in the predicate query language. This object iterator finds only those objects that satisfy the *predicate*.

Returns oocSuccess if successful; otherwise oocError.

Discussion This scan operation initializes this object iterator to find objects of class ooObj, class ooContObj, any basic-object class derived from ooObj, or any container class derived from ooContObj. If no such objects exist in the specified storage object, this object iterator is set to null and scan returns oocSuccess. (The object iterator's next member function will return oocFalse, however.)

Variant 3 finds persistent objects without opening them; the other variants may open the found objects, as specified by the <code>openMode</code> parameter.

If you specify the *predicate* parameter and the specified storage object has an index over the appropriate objects, scan can use the index to optimize its search.

(You may need to enable index usage; see the oouseIndex global function.) If no index is available, scan finds objects by inspection.

An error is signaled if *predicate* tries to test a non-existent data member, or if *predicate* tries to test a protected or private data member when the *access* parameter is oocPublic.

If you specify the *lookupKey* parameter, scan finds objects *only* if a compatible index exists in specified storage object; otherwise, this object iterator is set to null. For performance reasons, you should specify the *lookupKey* parameter only if you know the specified storage object contains such an index. You can test the storage object for compatible indexes by calling the <code>oolookupKey::anyIndex</code> member function. (Indexes need not be enabled to scan with a lookup key.)

See also <u>next</u>

# oojArray Class

Inheritance: ooObj->oojArray

Handle Class: ooHandle(oojArray)

**Object-Reference Class:** ooRef(oojArray)

The persistence-capable class oojArray is the abstract base class for all Java-compatibility classes that represent variable-size arrays. Each concrete derived class represents variable-size arrays with elements of a particular type.

See:

• "Reference Index" on page 321 for a list of member functions

To use the Java-compatibility classes, your application source must include the <code>javaBuiltins.h</code> header file.

## **About Java Arrays**

Unlike a C++ object, a Java object cannot contain another embedded object. As a consequence, a Java object cannot have an embedded variable-size array (VArray) the way a C++ object can. A variable-sized array in a Java persistent object is stored in an Objectivity/DB federated database as an object reference to a persistent array object of some Java-compatibility class. The persistent array object is a wrapper for a VArray with the same number of elements of the appropriate type. For example, a Java array of 16-bit integers (of the Java type <code>short[]</code>) is stored as an object reference of the type <code>ooRef(oojArrayOfInt16)</code>.

**NOTE** Java does not support fixed-size arrays.

Because the oojArray class is abstract, you never create instances of it; instead, you work with instances of its concrete derived classes.

You should not create your own subclasses of this class.

## **Multidimensional Arrays**

Unlike a VArray, a Java array can be multidimensional.

### **Element Order**

When the a multidimensional Java array is "flattened" into a one-dimensional VArray, elements are written to the VArray in row-major order, that is, with the last array index varying first.

**EXAMPLE** Consider the following two-dimensional 2-by-3 array of 16-bit integers:

111213212223

This array could be represented in Java by a short[] array (called ary for illustration) whose individual elements have the values:

ary[0][0]	=	11	ary[0][1]	=	12	ary[0][2]	=	13
ary[1][0]	=	21	ary[1][1]	=	22	ary[1][2]	=	23

If that array is stored in a database and accessed by a C++ application as an oojArrayOfInt16, the array object's corresponding VArray (called vary for illustration) would have six elements with the values:

vary[0] = 11 vary[1] = 12 vary[2] = 13 vary[3] = 21 vary[4] = 22 vary[5] = 23

### Array Dimensions

The persistent array's <u>getDimensionsArray</u> member function returns a VArray containing the dimensions of the Java array.

- The number of elements in the VArray is the number of dimensions in the Java array.
- The value of each element is the size of the corresponding array.

For example, consider a 3-dimensional Java array with dimensions 10 by 10 by 3. The corresponding C++ persistent array object wraps a VArray with 300 elements. The array object's dimensions array is a VArray of 3 elements: 10, 10, and 3.

**WARNING** You must not change the size of the VArray for a multidimensional array object. If you do so, the stored dimensions for the array object will be incorrect and applications will not be able to index the elements correctly.

## **Reference Index**

<u>getDimensionsArray</u> Gets the dimensions of the Java array corresponding to this array object.

## **Member Functions**

### getDimensionsArray

Gets the dimensions of the Java array corresponding to this array object.

ooVArrayT<int32> &getDimensionsArray();

Returns A VArray containing the dimensions of the Java array. If the Java array is unidimensional, the returned VArray will contain a single element, which is the size (number of elements) in the Java array. Otherwise, the returned VArray will contain one element from each dimension of the Java array. Member Functions

# oojArrayOfBoolean Class

Inheritance: ooObj->oojArray->oojArrayOfBoolean

Handle Class: ooHandle(oojArrayOfBoolean)

**Object-Reference Class:** ooRef(oojArrayOfBoolean)

The persistence-capable class <code>oojArrayOfBoolean</code> is a Java-compatibility class that represents a variable-size array of Boolean elements.

See:

• "Reference Index" on page 324 for a list of member functions

To use the Java-compatibility classes, your application source must include the <code>javaBuiltins.h</code> header file.

## **About Boolean Arrays**

A Java array of Boolean values (of the Java type boolean[]) is stored in an Objectivity/DB federated database as an object reference to a persistent array object of the class <code>oojArrayOfBoolean</code>. If your application interoperates with a Java application to access objects with a field that contains an array of Boolean values, you can define the corresponding data member of your C++ class to be of type <code>ooRef(oojArrayOfBoolean)</code>.

An instance of <code>oojArrayOfBoolean</code> is a wrapper for a variable-size array (VArray) with elements of the type <code>uint8</code>. You can obtain the VArray by calling the array object's <code>getBooleanArray</code> member function.

As is the case for all persistence-capable classes, you specify whether an instance of <code>oojArrayOfBoolean</code> is to be transient or persistent when you create it. You create the Boolean array object with a call to the <code>new</code> operator; the clustering

directive in that call specifies whether to make the new array object persistent and, if so, where to locate it.

Like other persistent objects, persistent Boolean arrays are normally manipulated through handles or object references.

## **Reference Index**

This class overloads  $\underline{operator new}$  and  $\underline{operator delete}$ , which behave as described for the  $\underline{ooObj}$  class (page 431).

<u>oojArrayOfBoolean</u>	Constructs a new Boolean array.
getBooleanArray	Gets this Boolean array's VArray.

## Constructors

## oojArrayOfBoolean

Constructs a new Boolean array.

	<ol> <li>oojArrayOfBoolean();</li> <li>oojArrayOfBoolean(int <i>initialSize</i>);</li> </ol>			
Parameters	<i>initialSize</i> Initial number of elements for which space should be allocated.			
Discussion	Variant 1 is the default constructor. It creates a Boolean array whose VArray has no element vector.			
	Variant 2 creates a Boolean array whose VArray's element vector contains the specified number of elements. If <i>initialSize</i> is 0, no element vector is			
#### getBooleanArray

Gets this Boolean array's VArray.

ooVArrayT<uint8> &getBooleanArray();

Returns This Boolean array's VArray.

## oojArrayOfCharacter Class

#### Inheritance: ooObj->oojArray->oojArrayOfCharacter

Handle Class: ooHandle(oojArrayOfCharacter)

**Object-Reference Class:** ooRef(oojArrayOfCharacter)

The persistence-capable class <code>oojArrayOfCharacter</code> is a Java-compatibility class that represents a variable-size array of characters.

See:

• "Reference Index" on page 328 for a list of member functions

To use the Java-compatibility classes, your application source must include the <code>javaBuiltins.h</code> header file.

### About Character Arrays

A Java array of characters (of the Java type char[]) is stored in an Objectivity/DB federated database as an object reference to a persistent array object of the class <code>oojArrayOfCharacter</code>. If your application interoperates with a Java application to access objects with a field that contains an array of characters, you can define the corresponding data member of your C++ class to be of type <code>ooRef(oojArrayOfCharacter)</code>.

An instance of <code>oojArrayOfCharacter</code> is a wrapper for a variable-size array (VArray) with elements of the type <code>uint16</code>. You can obtain the VArray by calling the array object's <code>getCharacterArray</code> member function.

As is the case for all persistence-capable classes, you specify whether an instance of <code>oojArrayOfCharacter</code> is to be transient or persistent when you create it. You create the character array object with a call to the <code>new</code> operator; the clustering directive in that call specifies whether to make the new array object persistent and, if so, where to locate it.

Like other persistent objects, persistent character arrays are normally manipulated through handles or object references.

## **Reference Index**

This class overloads  $\underline{operator new}$  and  $\underline{operator delete}$ , which behave as described for the  $\underline{ooObj}$  class (page 431).

getCharacterArray	Gets this character array's VArray.
<u>oojArrayOfCharacter</u>	Constructs a new character array.

## Constructors

### oojArrayOfCharacter

Constructs a new character array.

	<ol> <li>oojArrayOfCharacter();</li> <li>oojArrayOfCharacter(int <i>initialSize</i>);</li> </ol>	
Parameters	<i>initialSize</i> Initial number of elements for which space should be allocated.	
Discussion	Variant 1 is the default constructor. It creates a character array whose VArray has no element vector.	
	Variant 2 creates a character array whose VArray's element vector contains the specified number of elements. If <i>initialSize</i> is 0, no element vector is allocated for the VArray.	

#### getCharacterArray

Gets this character array's VArray.

ooVArrayT<uint16> &getCharacterArray();

Returns This character array's VArray.

# oojArrayOfDouble Class

Inheritance: ooObj->oojArray->oojArrayOfDouble

Handle Class: ooHandle(oojArrayOfDouble)

```
Object-Reference Class: ooRef(oojArrayOfDouble)
```

The persistence-capable class <code>oojArrayOfDouble</code> is a Java-compatibility class that represents a variable-size array of double-precision floating-point numbers.

See:

• "Reference Index" on page 332 for a list of member functions

To use the Java-compatibility classes, your application source must include the  ${\tt javaBuiltins.h}$  header file.

## **About Arrays of Double**

A Java array of double-precision floating-point numbers (of the Java type double[]) is stored in an Objectivity/DB federated database as an object reference to a persistent array object of the class <code>oojArrayOfDouble</code>. If your application interoperates with a Java application to access objects with a field that contains an array of double-precision floating-point numbers, you can define the corresponding data member of your C++ class to be of type <code>ooRef(oojArrayOfDouble)</code>.

An instance of <code>oojArrayOfDouble</code> is a wrapper for a variable-size array (VArray) with elements of the type <code>float64</code>. You can obtain the VArray by calling the array object's <code>getDoubleArray</code> member function.

As is the case for all persistence-capable classes, you specify whether an instance of <code>oojArrayOfDouble</code> is to be transient or persistent when you create it. You create the array of double with a call to the <code>new</code> operator; the clustering directive

in that call specifies whether to make the new array object persistent and, if so, where to locate it.

Like other persistent objects, persistent arrays of double are normally manipulated through handles or object references.

## **Reference Index**

This class overloads  $\underline{operator new}$  and  $\underline{operator delete}$ , which behave as described for the  $\underline{ooObj}$  class (page 431).

getDoubleArray	Gets this array of double's VArray.
<u>oojArrayOfDouble</u>	Constructs a new array of double.

## Constructors

### oojArrayOfDouble

Constructs a new array of double.

	<pre>1. oojArrayOfDouble();</pre>
	<pre>2. oojArrayOfDouble(int initialSize);</pre>
Parameters	<i>initialSize</i> Initial number of elements for which space should be allocated.
Discussion	Variant 1 is the default constructor. It creates an array of double whose VArray has no element vector.
	Variant 2 creates an array of double whose VArray's element vector contains the specified number of elements. If <i>initialSize</i> is 0, no element vector is allocated for the VArray.

#### getDoubleArray

Gets this array of double's VArray.

ooVArrayT<float64> &getDoubleArray();

Returns This array of double's VArray.

## oojArrayOfFloat Class

Inheritance: ooObj->oojArray->oojArrayOfFloat

Handle Class: ooHandle(oojArrayOfFloat)

```
Object-Reference Class: ooRef(oojArrayOfFloat)
```

The persistence-capable class <code>oojArrayOfFloat</code> is a Java-compatibility class that represents a variable-size array of single-precision floating-point numbers.

See:

• "Reference Index" on page 336 for a list of member functions

To use the Java-compatibility classes, your application source must include the  ${\tt javaBuiltins.h}$  header file.

## **About Arrays of Float**

A Java array of floating-point numbers (of the Java type float[]) is stored in an Objectivity/DB federated database as an object reference to a persistent array object of the class oojArrayOfFloat. If your application interoperates with a Java application to access objects with a field that contains an array of floating-point numbers, you can define the corresponding data member of your C++ class to be of type ooRef(oojArrayOfFloat).

An instance of <code>oojArrayOfFloat</code> is a wrapper for a variable-size array (VArray) with elements of the type <code>float32</code>. You can obtain the VArray by calling the array object's <code>getFloatArray</code> member function.

As is the case for all persistence-capable classes, you specify whether an instance of <code>oojArrayOfFloat</code> is to be transient or persistent when you create it. You create the array of float with a call to the <code>new operator</code>; the clustering directive in

that call specifies whether to make the new array object persistent and, if so, where to locate it.

Like other persistent objects, persistent arrays of float are normally manipulated through handles or object references.

## **Reference Index**

This class overloads  $\underline{operator new}$  and  $\underline{operator delete}$ , which behave as described for the  $\underline{ooObj}$  class (page 431).

getFloatArray	Gets this array of float's VArray.
<u>oojArrayOfFloat</u>	Constructs a new array of float.

## Constructors

### oojArrayOfFloat

Constructs a new array of float.

	<ol> <li>oojArrayOfFloat();</li> <li>oojArrayOfFloat(int <i>initialSize</i>);</li> </ol>
Parameters	<i>initialSize</i> Initial number of elements for which space should be allocated.
Discussion	Variant 1 is the default constructor. It creates an array of float whose VA rray has
DISCUSSION	no element vector.

#### getFloatArray

Gets this array of float's VArray.

ooVArrayT<float32> &getFloatArray();

Returns This array of float's VArray.

# oojArrayOfInt8 Class

#### Inheritance: ooObj->oojArray->oojArrayOfInt8

Handle Class: ooHandle(oojArrayOfInt8)

```
Object-Reference Class: ooRef(oojArrayOfInt8)
```

The persistence-capable class <code>oojArrayOfInt8</code> is a Java-compatibility class that represents a variable-size array of 8-bit integers.

See:

• "Reference Index" on page 340 for a list of member functions

To use the Java-compatibility classes, your application source must include the  ${\tt javaBuiltins.h}$  header file.

## **About 8-Bit Integer Arrays**

A Java array of 8-bit integers (of the Java type byte[]) is stored in an Objectivity/DB federated database as an object reference to a persistent array object of the class <code>oojArrayOfInt8</code>. If your application interoperates with a Java application to access objects with a field that contains an array of 8-bit integers, you can define the corresponding data member of your C++ class to be of type <code>ooRef(oojArrayOfInt8)</code>.

An instance of <code>oojArrayOfInt8</code> is a wrapper for a variable-size array (VArray) with elements of the type <code>int8</code>. You can obtain the VArray by calling the array object's <code>getInt8Array</code> member function.

As is the case for all persistence-capable classes, you specify whether an instance of <code>oojArrayOfInt8</code> is to be transient or persistent when you create it. You create the 8-bit integer array with a call to the <code>new</code> operator; the clustering directive in

that call specifies whether to make the new array object persistent and, if so, where to locate it.

Like other persistent objects, persistent 8-bit integer arrays are normally manipulated through handles or object references.

## **Reference Index**

This class overloads  $\underline{operator new}$  and  $\underline{operator delete}$ , which behave as described for the  $\underline{ooObj}$  class (page 431).

getInt8Array	Gets this 8-bit integer array's VArray.
oojArrayOfInt8	Constructs a new 8-bit integer array.

## Constructors

#### oojArrayOfInt8

Constructs a new 8-bit integer array.

	<pre>1. oojArrayOfInt8();</pre>
	<pre>2. oojArrayOfInt8(int initialSize);</pre>
Parameters	<i>initialSize</i> Initial number of elements for which space should be allocated.
Discussion	Variant 1 is the default constructor. It creates an 8-bit integer array whose VArray has no element vector.

#### getInt8Array

Gets this 8-bit integer array's VArray.

ooVArrayT<int8> &getInt8Array();

Returns This 8-bit integer array's VArray.

# oojArrayOfInt16 Class

#### Inheritance: ooObj->oojArray->oojArrayOfInt16

Handle Class: ooHandle(oojArrayOfInt16)

```
Object-Reference Class: ooRef(oojArrayOfInt16)
```

The persistence-capable class <code>oojArrayOfInt16</code> is a Java-compatibility class that represents a variable-size array of 16-bit integers.

See:

• "Reference Index" on page 344 for a list of member functions

To use the Java-compatibility classes, your application source must include the  ${\tt javaBuiltins.h}$  header file.

## About 16-Bit Integer Arrays

A Java array of 16-bit integers (of the Java type byte[]) is stored in an Objectivity/DB federated database as an object reference to a persistent array object of the class <code>oojArrayOfInt16</code>. If your application interoperates with a Java application to access objects with a field that contains an array of 16-bit integers, you can define the corresponding data member of your C++ class to be of type <code>ooRef(oojArrayOfInt16)</code>.

An instance of <code>oojArrayOfInt16</code> is a wrapper for a variable-size array (VArray) with elements of the type <code>int16</code>. You can obtain the VArray by calling the array object's <code>getInt16Array</code> member function.

As is the case for all persistence-capable classes, you specify whether an instance of <code>oojArrayOfInt16</code> is to be transient or persistent when you create it. You create the 16-bit integer array with a call to the <code>new</code> operator; the clustering

directive in that call specifies whether to make the new array object persistent and, if so, where to locate it.

Like other persistent objects, persistent 16-bit integer arrays are normally manipulated through handles or object references.

## **Reference Index**

This class overloads  $\underline{operator new}$  and  $\underline{operator delete}$ , which behave as described for the  $\underline{ooObj}$  class (page 431).

<u>getInt16Array</u>	Gets this 16-bit integer array's VArray.
<u>oojArrayOfInt16</u>	Constructs a new 16-bit integer array.

## Constructors

### oojArrayOfInt16

Constructs a new 16-bit integer array.

	<ol> <li>oojArrayOfInt16();</li> <li>oojArrayOfInt16(int <i>initialSize</i>);</li> </ol>	
Parameters	<i>initialSize</i> Initial number of elements for which space should be allocated.	
Discussion	Variant 1 is the default constructor. It creates a 16-bit integer array whose VArray has no element vector.	
	Variant 2 creates a 16-bit integer array whose VArray's element vector contains the specified number of elements. If <i>initialSize</i> is 0, no element vector is allocated for the VArray.	

#### getInt16Array

Gets this 16-bit integer array's VArray.

ooVArrayT<int16> &getInt16Array();

Returns This 16-bit integer array's VArray.

# oojArrayOfInt32 Class

Inheritance: ooObj->oojArray->oojArrayOfInt32

Handle Class: ooHandle(oojArrayOfInt32)

```
Object-Reference Class: ooRef(oojArrayOfInt32)
```

The persistence-capable class <code>oojArrayOfInt32</code> is a Java-compatibility class that represents a variable-size array of 32-bit integers.

See:

• "Reference Index" on page 348 for a list of member functions

To use the Java-compatibility classes, your application source must include the  ${\tt javaBuiltins.h}$  header file.

## About 32-Bit Integer Arrays

A Java array of 32-bit integers (of the Java type byte[]) is stored in an Objectivity/DB federated database as an object reference to a persistent array object of the class <code>oojArrayOfInt32</code>. If your application interoperates with a Java application to access objects with a field that contains an array of 32-bit integers, you can define the corresponding data member of your C++ class to be of type <code>ooRef(oojArrayOfInt32)</code>.

An instance of <code>oojArrayOfInt32</code> is a wrapper for a variable-size array (VArray) with elements of the type <code>int32</code>. You can obtain the VArray by calling the array object's <code>getInt32Array</code> member function.

As is the case for all persistence-capable classes, you specify whether an instance of <code>oojArrayOfInt32</code> is to be transient or persistent when you create it. You create the 32-bit integer array with a call to the <code>new</code> operator; the clustering

directive in that call specifies whether to make the new array object persistent and, if so, where to locate it.

Like other persistent objects, persistent 32-bit integer arrays are normally manipulated through handles or object references.

## **Reference Index**

This class overloads  $\underline{operator new}$  and  $\underline{operator delete}$ , which behave as described for the  $\underline{ooObj}$  class (page 431).

getInt32Array	Gets this 32-bit integer array's VArray.
<u>oojArrayOfInt32</u>	Constructs a new 32-bit integer array.

## Constructors

### oojArrayOfInt32

Constructs a new 32-bit integer array.

	<ol> <li>oojArrayOfInt32();</li> <li>oojArrayOfInt32(int <i>initialSize</i>);</li> </ol>	
Parameters	<i>initialSize</i> Initial number of elements for which space should be allocated.	
Discussion	Variant 1 is the default constructor. It creates a 32-bit integer array whose VArray has no element vector.	
	Variant 2 creates a 32-bit integer array whose VArray's element vector contains the specified number of elements. If <i>initialSize</i> is 0, no element vector is allocated for the VArray.	

#### getInt32Array

Gets this 32-bit integer array's VArray.

ooVArrayT<int32> &getInt32Array();

Returns This 32-bit integer array's VArray.

# oojArrayOfInt64 Class

#### Inheritance: ooObj->oojArray->oojArrayOfInt64

Handle Class: ooHandle(oojArrayOfInt64)

```
Object-Reference Class: ooRef(oojArrayOfInt64)
```

The persistence-capable class <code>oojArrayOfInt64</code> is a Java-compatibility class that represents a variable-size array of 64-bit integers.

See:

• "Reference Index" on page 352 for a list of member functions

To use the Java-compatibility classes, your application source must include the  ${\tt javaBuiltins.h}$  header file.

## About 64-Bit Integer Arrays

A Java array of 64-bit integers (of the Java type byte[]) is stored in an Objectivity/DB federated database as an object reference to a persistent array object of the class <code>oojArrayOfInt64</code>. If your application interoperates with a Java application to access objects with a field that contains an array of 64-bit integers, you can define the corresponding data member of your C++ class to be of type <code>ooRef(oojArrayOfInt64)</code>.

An instance of <code>oojArrayOfInt64</code> is a wrapper for a variable-size array (VArray) with elements of the type <code>int64</code>. You can obtain the VArray by calling the array object's <code>getInt64Array</code> member function.

As is the case for all persistence-capable classes, you specify whether an instance of <code>oojArrayOfInt64</code> is to be transient or persistent when you create it. You create the 64-bit integer array with a call to the <code>new</code> operator; the clustering

directive in that call specifies whether to make the new array object persistent and, if so, where to locate it.

Like other persistent objects, persistent 64-bit integer arrays are normally manipulated through handles or object references.

## **Reference Index**

This class overloads  $\underline{operator new}$  and  $\underline{operator delete}$ , which behave as described for the  $\underline{ooObj}$  class (page 431).

getInt64Array	Gets this 64-bit integer array's VArray.
<u>oojArrayOfInt64</u>	Constructs a new 64-bit integer array.

## Constructors

### oojArrayOfInt64

Constructs a new 64-bit integer array.

	1.	oojArrayOfInt64();
	2.	<pre>oojArrayOfInt64(int initialSize);</pre>
Parameters	init Iı	<i>ialSize</i> nitial number of elements for which space should be allocated.
	Variant 1 is the default constructor. It creates a 64-bit integer array whose VArray has no element vector.	
Discussion	Varia has n	nt 1 is the default constructor. It creates a 64-bit integer array whose VArray o element vector.

#### getInt64Array

Gets this 64-bit integer array's VArray.

ooVArrayT<int64> &getInt64Array();

Returns This 64-bit integer array's VArray.

# oojArrayOfObject Class

Inheritance: ooObj->oojArray->oojArrayOfObject

Handle Class: ooHandle(oojArrayOfObject)

```
Object-Reference Class: ooRef(oojArrayOfObject)
```

The persistence-capable class <code>oojArrayOfObject</code> is a Java-compatibility class that represents a persistent variable-size array of object references.

See:

• "Reference Index" on page 356 for a list of member functions

To use the Java-compatibility classes, your application source must include the  ${\tt javaBuiltins.h}$  header file.

## About Object-Reference Arrays

A Java array of objects is stored in an Objectivity/DB federated database as an object reference to a persistent array object of the class <code>oojArrayOfObject</code>. The Java array may be of types shown in the following table.

Java Array Type	Schema Class of Objects Referenced by Array Elements
java.lang.String[]	oojString
java.util.Date[]	oojDate
java.sql.Date[]	oojDate
java.sql.Time[]	oojTime
java.sql.Timestamp[]	oojTimestamp

Java Array Type	Schema Class of Objects Referenced by Array Elements
<i>AppClass</i> [], where <i>AppClass</i> is an application-defined persistence-capable class whose schema class name is <i>PCclass</i> .	PCclass
APIclass[], where APIclass is a persistence-capable class in the public Objectivity for Java programmer interface (for example, ooContObj or ooMap).	APIclass
<i>PCinterface</i> [], where <i>PCinterface</i> is an interface (implemented by one or more persistence-capable classes).	The schema class corresponding to a Java class that implements <i>PCinterface</i> . (Different elements of the array may reference objects of different classes.)

If your application interoperates with a Java application to access objects with a field that contains an array of object references, you can define the corresponding data member of your C++ class to be of type ooRef(oojArrayOfObject).

An instance of <code>oojArrayOfObject</code> is a wrapper for a variable-size array (VArray) with elements of the type <code>ooRef(ooObj)</code>. You can obtain the VArray by calling the array object's <code>getObjectArray</code> member function.

As is the case for all persistence-capable classes, you specify whether an instance of <code>oojArrayOfObject</code> is to be transient or persistent when you create it. You create the object-reference array object with a call to the <code>new</code> operator; the clustering directive in that call specifies whether to make the new array object persistent and, if so, where to locate it.

Like other persistent objects, persistent object-reference arrays are normally manipulated through handles or object references.

## **Reference Index**

This class overloads  $\underline{operator new}$  and  $\underline{operator delete}$ , which behave as described for the  $\underline{ooObj}$  class (page 431).

<u>getObjectArray</u>	Gets this object-reference array's VArray.
<u>oojArrayOfObject</u>	Constructs a new object-reference array.

## Constructors

### oojArrayOfObject

Constructs a new object-reference array.

oojArrayOfObject(int initialSize);

Parameters initialSize
 Initial number of elements for which space should be allocated.

 Discussion Variant 1 is the default constructor. It creates an object-reference array whose VArray has no element vector.
 Variant 2 creates an object-reference array whose VArray's element vector contains the specified number of elements. If initialSize is 0, no element vector is allocated for the VArray.

## **Member Functions**

### getObjectArray

Gets this object-reference array's VArray.

ooVArrayT<ooRef(ooObj)> &getObjectArray();

Returns This object-reference array's VArray.

## oojDate Class

Inheritance: ooObj->oojDate

Handle Class: ooHandle(oojDate)

**Object-Reference Class:** ooRef(oojDate)

The persistence-capable class oojDate is a Java-compatibility class that represents an instant in time with millisecond precision.

See:

• "Reference Index" on page 360 for a list of member functions

To use the Java-compatibility classes, your application source must include the  ${\tt javaBuiltins.h}$  header file.

## **About Dates**

A Java date (of the Java class java.util.Date or java.sql.Date) is stored in an Objectivity/DB federated database as an object reference to a persistent object of the class oojDate. If your application interoperates with a Java application to access objects with a field that contains a date, you can define the corresponding data member of your C++ class to be of type ooRef(oojDate).

The class oojDate represents an instant in time as the number of milliseconds since January 1, 1970 00:00:00.000 GMT.

As is the case for all persistence-capable classes, you specify whether an instance of oojDate is to be transient or persistent when you create it. You create the date object with a call to the new operator; the clustering directive in that call specifies whether to make the new date object persistent and, if so, where to locate it.

Like other persistent objects, persistent dates are normally manipulated through handles or object references.

## **Reference Index**

This class overloads <u>operator new</u> and <u>operator delete</u>, which behave as described for the <u>ooObj</u> class (page 431).

<u>getMillis</u>	Gets the millisecond representation of this date.
<u>oojDate</u>	Constructs a new date object.
<u>setMillis</u>	Sets the millisecond representation of this date.

## Constructors

#### oojDate

Constructs a new date object.

	1. oojDate();
	<pre>2. oojDate(int64 millisecs);</pre>
Parameters	millisecs The millisecond representation of the new date.
Discussion	Variant 1 is the default constructor, which creates a date representing January 1, 1970 00:00:00.000 GMT.
	Variant 2 creates a date with the specified millisecond representation.

## **Member Functions**

#### getMillis

Gets the millisecond representation of this date.

int64 getMillis();

Returns The millisecond representation of this date.
### setMillis

Sets the millisecond representation of this date.

void setMillis(int64 millisecs);

#### Parameters *millisecs*

The new millisecond representation of this date.

Member Functions

# oojString Class

Inheritance: ooObj->oojString

Handle Class: ooHandle(oojString)

**Object-Reference Class:** ooRef(oojString)

The persistence-capable class <code>oojString</code> is a Java-compatibility class that represents a string element of a persistent object-reference array.

See:

• "Reference Index" on page 364 for a list of member functions

To use the Java-compatibility classes, your application source must include the <code>javaBuiltins.h</code> header file.

## **About String Elements**

A Java string array (of the Java type <code>String[]</code>) is stored in an Objectivity/DB federated database as an object reference to a persistent array object of the class <code>oojArrayOfObject</code>. Elements of the array are object references to instances of <code>oojString</code>.

This class is a wrapper for a Unicode string of the <u>ooUtf8String</u> class; you can obtain the string from a string element by calling the <u>getStringValue</u> member function.

As is the case for all persistence-capable classes, you specify whether an instance of oojString is to be transient or persistent when you create it. You create the string element with a call to the new operator; the clustering directive in that call specifies whether to make the new string element persistent and, if so, where to locate it.

Like other persistent objects, persistent string elements are normally manipulated through handles or object references.

## **Reference Index**

This class overloads <u>operator new</u> and <u>operator delete</u>, which behave as described for the <u>ooObj</u> class (page 431).

<u>getStringValue</u>	Gets the Unicode string in this string element.
oojString	Default constructor that creates a new string element with an empty string.

# Constructors

## oojString

Default constructor that creates a new string element with an empty string.

oojString();

# **Member Functions**

### getStringValue

Gets the Unicode string in this string element.

ooUtf8String &getStringValue();

Returns The Unicode string in this string element.

Discussion Modifying the return string modifies this string element.

# oojTime Class

Inheritance: ooObj->oojTime

Handle Class: ooHandle(oojTime)

**Object-Reference Class:** ooRef(oojTime)

The class oojTime is a Java-compatibility class that represents a time with millisecond precision.

See:

• "Reference Index" on page 366 for a list of member functions

To use the Java-compatibility classes, your application source must include the <code>javaBuiltins.h</code> header file.

## **About Times**

A Java time object (of the Java class java.sql.Time) is stored in an Objectivity/DB federated database as an object reference to a persistent object of the class oojTime. If your application interoperates with a Java application to access objects with a field that contains a time, you can define the corresponding data member of your C++ class to be of type <code>ooRef(oojTime)</code>.

The class oojTime represents an instant in time as the number of milliseconds since midnight. Thus 1000 milliseconds corresponds to the time one second after midnight, namely, 12:00:01 AM.

As is the case for all persistence-capable classes, you specify whether an instance of oojTime is to be transient or persistent when you create it. You create the time object with a call to the new operator; the clustering directive in that call specifies whether to make the new time object persistent and, if so, where to locate it.

Like other persistent objects, persistent times are normally manipulated through handles or object references.

## **Reference Index**

This class overloads <u>operator new</u> and <u>operator delete</u>, which behave as described for the <u>ooObj</u> class (page 431).

getMillis	Gets the millisecond representation of this time.
<u>oojTime</u>	Constructs a new time object.
<u>setMillis</u>	Sets the millisecond representation of this time.

# Constructors

## oojTime

Constructs a new time object.

<pre>1. oojTime();</pre>	
--------------------------	--

2. oojTime(int64 millisecs);

Parameters millisecs

The millisecond representation of the new time.

DiscussionVariant 1 is the default constructor, which creates a time representing midnight.Variant 2 creates a time with the specified millisecond representation.

# **Member Functions**

### getMillis

Gets the millisecond representation of this time.

int64 getMillis();

Returns The millisecond representation of this time.

### setMillis

Sets the millisecond representation of this time.

void setMillis(int64 millisecs);

#### Parameters *millisecs*

The new millisecond representation of this time.

Member Functions

# oojTimestamp Class

Inheritance: ooObj->oojTimestamp

Handle Class: ooHandle(oojTimestamp)

```
Object-Reference Class: ooRef(oojTimestamp)
```

The class oojTimestamp is a Java-compatibility class that represents an instant in time with nanosecond precision.

See:

- "Reference Summary" on page 370 for an overview of member functions
- "Reference Index" on page 370 for a list of member functions

To use the Java-compatibility classes, your application source must include the <code>javaBuiltins.h</code> header file.

## About Timestamps

A Java timestamp (of the Java class java.sql.Timestamp) is stored in an Objectivity/DB federated database as an object reference to a persistent object of the class oojTimestamp. If your application interoperates with a Java application to access objects with a field that contains a timestamp, you can define the corresponding data member of your C++ class to be of type <code>ooRef(oojTimestamp)</code>.

The class oojTimestamp represents an instant in time as the number of nanoseconds since January 1, 1970 00:00:00.000000000 GMT, given as an integral part and a fractional part:

The integral part specifies the number of seconds since January 1, 1970 00:00:00 GMT; however, it is expressed in *milliseconds* rather than seconds. • The fractional part specifies the number of nanoseconds.

For example, a timestamp representing January 1, 1970 00:00:01.111222333 GMT consists of the integral part 1000 and the fractional part 111222333.

As is the case for all persistence-capable classes, you specify whether an instance of oojTimestamp is to be transient or persistent when you create it. You create the timestamp object with a call to the new operator; the clustering directive in that call specifies whether to make the new timestamp object persistent and, if so, where to locate it.

Like other persistent objects, persistent timestamps are normally manipulated through handles or object references.

## **Reference Summary**

In the following table, operators indicated as *(inherited)* are overloaded in this class with no change in behavior; they are documented with the <u>ooObj</u> class (page 431).

Creating and Deleting a Timestamp	<u>oojTimestamp</u> <u>operator_new</u> (inherited) <u>operator_delete</u> (inherited)
Getting Information About the Timestamp	<u>getMillis</u> <u>getNanos</u>
Setting Information About the Timestamp	<u>setMillis</u> <u>setNanos</u>

# **Reference Index**

getMillis	Gets the integral part of this timestamp.
getNanos	Gets the fractional part of this timestamp.
<u>oojTimestamp</u>	Constructs a new timestamp object.
<u>setMillis</u>	Sets the integral part of this timestamp.
<u>setNanos</u>	Sets the fractional part of this timestamp.

# Constructors

## oojTimestamp

Constructs a new timestamp object.

	<ol> <li>oojTimestamp();</li> <li>oojTimestamp(int64 millisecs, int nanosecs);</li> </ol>	
Parameters	millisecs The integral part of the new timestamp.	
	nanosecs The fractional part of the new timestamp.	
Discussion	Variant 1 is the default constructor, which creates a timestamp representing January 1, 1970 00:00:00.000000000 GMT.	
	Variant 2 creates a timestamp with the specified integral and fractional parts.	

# **Member Functions**

## getMillis

Cets the integral	nart	of this	timestam	n
Gets the integral	part	or uns	unnestann	μ

int64 getMillis();

Returns The integral part of this timestamp.

### getNanos

Gets the fractional part of this timestamp.

int32 getNanos();

Returns The fractional part of this timestamp

## setMillis

	Sets the integral part of this timestamp.
	<pre>void setMillis(int64 millisecs);</pre>
Parameters	millisecs The new integral part of this timestamp.
setNanos	

Sets the fractional part of this timestamp.

void setNanos(int32 nanosecs);

#### Parameters nanosecs

The new fractional part of this timestamp.

# ooKeyDesc Class

Inheritance: ooObj->...->ooKeyDesc

Handle Class: ooHandle(ooKeyDesc)

**Object-Reference Class:** ooRef(ooKeyDesc)

The persistence-capable class ookeyDesc represents a *key description* from which one or more indexes can be created.

See:

- "Reference Summary" on page 374 for an overview of member functions
- "Reference Index" on page 374 for a list of member functions

Applications that use indexes must include the ooIndex.h header file.

## **About Key Descriptions**

A key description is a persistent object from which an index is created. Each key description identifies the class of objects to be indexed, the key fields on which to sort, and whether the index is to be unique. The key description is then used to create an index for a particular storage object, sometimes called the *scope* of the index. The storage object you choose for an index limits the objects referenced by that index—for example, an index created for a container references the objects of the indexed class that reside in that container. When an index is no longer needed, you use its key description to remove or *drop* it from the storage object.

You can use a particular key description to create multiple indexes, provided that each index is in a different storage object. Multiple indexes can exist in the same storage object only if each was created from a different key description.

# **Working With Key Descriptions**

As is the case for any basic object, you specify whether a key description is to be transient or persistent when you create it; key descriptions *must be persistent*. You create a key description with a call to the new operator; the clustering directive in that call specifies where in the federated database to store the new key description.

Like other persistent objects, key descriptions are normally manipulated through handles or object references. You can store and find a key description in the database just as you would any other persistent object.

# **Reference Summary**

In the following table, operators indicated as *(inherited)* are overloaded in this class with no change in behavior; they are documented with the  $\underline{ooObj}$  class (page 431), along with the other inherited member functions not listed here.

Creating and Deleting a Key Description	<u>ooKeyDesc</u> <u>operator_new</u> (inherited) <u>operator_delete</u> (inherited)
Working With Key-Field Objects	addField nField
Creating and Dropping Indexes	<u>createIndex</u> <u>dropIndex</u> <u>removeIndexes</u>
Getting Information	getTypeN getTypeName isConsistent isUnique

## **Reference Index**

<u>addField</u>	Adds the referenced key-field object to this key description.
<u>createIndex</u>	Creates an index for the specified storage object from this key description.

Drops the index associated with this key description from the referenced storage object.
Gets the type number of the class that is indexed by this key description.
Gets the name of the class indexed by this key description.
Checks whether the key-field objects in this key description are still consistent with the corresponding data members of the indexed class.
Determines whether a particular index was created to be unique. When an index is unique, each of its indexed objects must have unique key field values.
Gets the number of fields in the index.
Constructs a new key description with the specified characteristics.
Removes from the federated database all indexes that were created from this key description.

## Constructors

#### ooKeyDesc

Constructs a new key description with the specified characteristics.

```
ooKeyDesc(
    const ooTypeNumber typeN,
    const ooBoolean unique = oocFalse);
```

#### Parameters *typeN*

Type number of the persistence-capable class whose objects are to be indexed. This class, and all of its subclasses, will be included in the index.

#### unique

Specifies whether the indexes created from this key description are to be unique. If you omit this parameter, this key description will create nonunique indexes, in which multiple indexed objects may have the same combination of key field values.

# **Member Functions**

### addField

Adds the referenced key-field object to this key description.

ooStatus addField(const ooHandle(ooKeyField) &fieldH);

Parameters	fieldH Handle to the key-field object to be added.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	You may add a key-field object to a key description only if both are defined on the same class (the class to be indexed). The key-field object being added must refer to a data member of the indexed class or one of its base classes.
	You can add multiple key-field objects to the same description. The order in which key-field objects are added to a key description determines the sorting order of the indexed objects.
	The same key-field object can be added to multiple key descriptions.

### createIndex

Creates an index for the specified storage object from this key description.

	<pre>1. ooStatus createIndex(</pre>	
	<pre>2. ooStatus createIndex(</pre>	
	<pre>3. ooStatus createIndex(</pre>	
Parameters	storageObject	
	Handle to the container, database, or federated database for which the index is to be created. The specified storage object limits the objects that are referenced by the new index—for example, an index created for a container references only the objects of the indexed class that reside in that container.	
	storageObject may not specify a container if this key description specifies a container class.	
Returns	oocSuccess if successful; otherwise oocError.	

Discussion You can create indexes for different storage objects from the same key description.

If an index cannot be created, a message is displayed and the transaction is aborted.

### dropIndex

Drops the index associated with this key description from the referenced storage object.

Returns	oocSuccess if successful; otherwise oocError.
Parameters	storageObject Handle to the container, database, or federated database from which the index is to be dropped.
	<pre>3. ooStatus dropIndex(     const ooHandle(ooFDObj) &amp;storageObject);</pre>
	<pre>2. ooStatus dropIndex(</pre>
	<pre>1. ooStatus dropIndex(</pre>

## getTypeN

Gets the type number of the class that is indexed by this key description.

ooTypeNumber getTypeN() const;

Returns Type number of the indexed class.

### getTypeName

Gets the name of the class indexed by this key description.

const char \*getTypeName() const;

Returns String name of the indexed class.

### isConsistent

Checks whether the key-field objects in this key description are still consistent with the corresponding data members of the indexed class.

```
ooBoolean isConsistent();
```

Returns	Returns oocTrue if all the key-field objects in this key description are consistent; returns oocFalse if any key-field object in this key description is inconsistent.
Discussion	If this member function returns oocFalse, you can call the <u>isConsistent</u> member function on each key-field object to determine which one is inconsistent.
	You should test for consistency after performing a schema-evolution operation that affects an indexed class or a data member of an indexed class.
isUnique	
	Determines whether a particular index was created to be unique. When an index is unique, each of its indexed objects must have unique key field values.
	ooBoolean isUnique();

Returns oocTrue if the target index is unique and oocFalse if it is not.

### nField

Gets the number of fields in the index.

unit32 nField() const;

Returns Integer giving the number of fields in the index.

#### removeIndexes

Removes from the federated database all indexes that were created from this key description.

ooStatus removeIndexes();

Returns oocSuccess if successful; otherwise oocError.

Discussion This member function deletes indexes while preserving the key description itself. Alternatively, if you no longer need the key description, you can delete it with the ooDelete function, which automatically deletes all its indexes.

# ooKeyField Class

Inheritance: ooObj->ooKeyField

Handle Class: ooHandle(ooKeyField)

```
Object-Reference Class: ooRef(ooKeyField)
```

The persistence-capable class <code>ooKeyField</code> represents a *key field* of an index.

See:

- "Reference Summary" on page 382 for an overview of member functions
- "Reference Index" on page 382 for a list of member functions

Applications that use indexes must include the ooIndex.h header file.

## **About Key-Field Objects**

*Key-field objects* are added to a key description (an instance of <u>ooKeyDesc</u>) to define an index's key fields. Each key-field object in a key description corresponds to a particular data member of the indexed class, and designates that data member as a key field of the index. The order in which key field objects are added to a key description determines the order in which the resulting index considers the corresponding key fields when sorting the indexed objects.

For example, assume you want to create an index over Person objects, sorted by key fields name and age. You achieve this by creating key-field objects corresponding to the name and age data members of class Person, and then adding them, in that order, to an appropriate key description. In the resulting index, objects are first sorted by name; if two or more objects have the same name, the one with the lowest age comes first in the indexed order. You can create different sorts over the same objects by defining different key descriptions on the same indexed class and adding the same key-field objects in different orders.

Each key-field object identifies one of the following:

- A particular attribute data member of the indexed class.
- A particular data member of the embedded class of an embedded-class attribute of the indexed class.

The data member may be C++ private, protected, or public. Its data type must be a numeric, character, Boolean, or string type. In particular:

■ A primitive type such as uint16, int16, char, or ooBoolean.

See the Objectivity/C++ Data Definition Language book for a complete list of primitive types, or see "Primitive Type Names" on page 25 of the Objectivity/C++ programmer's reference.

- A string class: ooString(N), ooVString, or ooUtf8String
- A fixed-size character array of type char[] (treated as a null-terminated string)
- A VArray of characters (treated as a null-terminated string)

Indexes can optimize tests that compare a key field with a literal numeric or string value. Indexes can optimize tests that compare a key field with a literal numeric or string value. If you define a key field of type <code>ooBoolean</code>, an optimized predicate must use an *integer* literal (1 for <code>oocTrue</code>, 0 for <code>oocFalse</code>).

# Working With a Key-Field Object

As is the case for any basic object, you specify whether a key-field object is to be transient or persistent when you create it; key-field objects *must be persistent*. You create a key-field object with a call to the new operator; the clustering directive in that call specifies where in the federated database to store the new key-field object.

Like other persistent objects, key-field objects are normally manipulated through handles or object references. You can store and find a key-field object in the database just as you would any other persistent object.

# **Optimizing String Storage and Lookup**

When you create a key-field object for a string-typed data member (for example, char[], ooString(N), ooVString, or ooUtf8String), you can optimize the space required to store each string key in the index, as well as the processing time

required to access it during lookup. You adjust these characteristics by setting the *ooKeyField* constructor's *fixed* and *maxstrlen* parameters.

Choosing the default values for these parameters (*fixed*=oocFalse and *maxstrlen*=24) causes the index to allocate 24 bytes per indexed object to store a string key for that object (a copy of the string value of the relevant key field). Shorter string keys are padded with null characters. If a string key is longer, only the first 24 bytes are stored in the index. During lookup, these bytes are compared to the first 24 bytes of the comparison value; if they match, the index obtains the entire string value from the indexed object to complete the comparison.

You can use the *maxstrlen* parameter to reduce the size of the index or to improve performance:

- If most of the string keys are smaller than 24 bytes, you can set maxstrlen to
  a smaller value to reduce the amount of extra space occupied by null
  characters.
- If most of the string keys are larger than 24 bytes, you can set maxstrlen to a larger value to reduce number of times the index must open an indexed object to get a complete string value.

If the string values of the indexed objects are guaranteed to be of a fixed (or limited) size, you can improve performance by setting the *fixed* parameter to oocTrue. This prevents an object from being indexed if its string value is greater than *maxstrlen*, so every string key in the index is complete; during lookup, comparisons are performed without accessing the actual string value in the indexed object.

**WARNING** Do not set *fixed* to oocTrue if any string value might be longer than *maxstrlen*. If the string value of an object's key field exceeds *maxstrlen*, that object is omitted from the index.

For optimal results, use the following guidelines:

- If the key field is a fixed-size string of length *N*, you should set *fixed* to oocTrue and *maxstrlen* to *N*.
- If the key field is a variable-size string such as an ooVString and you know that most of the string values (for example, 90%) are of length *N* or less, you should set *fixed* to oocFalse and *maxstrlen* to *N*.
- **EXAMPLE** This example creates a key-field object on an optimized string of length 8 (ooString(8)), a type that was chosen because the length of most strings in the class is 7 bytes or less. You can save space in the index by setting the maxstrlen parameter of the ooKeyField constructor to 8. However, because it is possible

for some string values to be longer than 8 bytes, the *fixed* parameter is oocFalse.

```
// DDL file
class Person : public ooObj {
    ooString(8) name; // Most strings are < 8 bytes long
}
// Application code
ooHandle(ooKeyField) keyFieldH;
ooHandle(ooKeyDesc) keyDescH;
ooHandle(ooContObj) contH;
keyDescH = new(contH) ooKeyDesc(ooTypeN(Person),oocTrue);
keyFieldH = new(keyDescH) ooKeyField (ooTypeN(Person),"name",
    oocFalse, 8); // Allows a variable length</pre>
```

## **Reference Summary**

In the following table, operators indicated as *(inherited)* are overloaded in this class with no change in behavior; they are documented with the 000bj class (page 431), along with the other inherited member functions not listed here.

Creating and Deleting	<u>ooKeyField</u> <u>operator new</u> (inherited) <u>operator delete</u> (inherited)
Getting Information	getName getTypeN isConsistent isNamed

# **Reference Index**

<u>getName</u>	Gets the name of the data member to which this key-field object corresponds.
getTypeN	Gets the type number of the indexed class for which the key-field object was created.

<u>isConsistent</u>	Checks whether this key-field object is still consistent with the corresponding data member of the indexed class.
<u>isNamed</u>	Tests whether the data member corresponding to this key-field object has the specified name.
ooKeyField	Constructs a new key-field object with the specified characteristics.

## Constructors

### ooKeyField

Constructs a new key-field object with the specified characteristics.

```
ooKeyField(
```

```
const ooTypeNumber typeN,
const char *memberName,
ooBoolean fixed = oocFalse,
uint32 maxstrlen = 24);
```

Parameters

Type number of the indexed class—that is, persistence-capable class whose objects are to be indexed.

#### memberName

typeN

Name of the data member to be represented by this key-field object. The values of this data member will be used to sort the indexed objects.

• To indicate the data member *attributeName* of the indexed class, use a string of the form:

"attributeName"

If the name of an inherited data member is ambiguous (for example, the same name is defined in both the base class and the indexed class) or if the member name is not visible to the indexed class due to access control, use a string of the following form to indicate the data member *inheritedAttribute*, which the indexed class inherits from the base class baseClassName:

"baseClassName::inheritedAttribute"

To indicate the data member fieldName of the embedded object (or struct) in the embedded-class attribute attributeName of the indexed class, use a string of the form:

"attributeName.fieldName"

#### fixed

Specifies whether the values of this key field are to be treated as fixed or variable in length; applies only if *memberName* specifies a data member whose type is char[], ooString(N), or ooVString.

If you specify oocFalse (the default), the index allows string keys of any length:

- If a string key is *maxstrlen* bytes or shorter, it is stored in the index (padded with null characters if necessary).
- If a string key is longer than maxstrlen bytes, the first maxstrlen bytes are stored in the index for preliminary comparisons during lookup; when the index needs the complete string value, it accesses the indexed object.

If you specify oocTrue, the index allows only string keys of length *maxstrlen* or shorter. If the string key for an object is longer than *maxstrlen* bytes, that object is omitted from the index and an error is signalled. If you know the maximum length of the string keys for all objects, specifying oocTrue can optimize the index's performance during lookup, because it never has to consult an indexed object to obtain a complete string value.

maxstrlen

Specifies the number of bytes to be allocated by the index for storing a copy of each value of this key field; applies only if *memberName* specifies a data member whose type is char[], ooString(N), or ooVString.

The index allocates *maxstrlen* bytes for each string key. Shorter string keys are padded with null characters. Longer string keys are handled only if *fixed* is set to oocFalse (see *fixed* above).

# **Member Functions**

### getName

Gets the name of the data member to which this key-field object corresponds.

virtual const char \*getName() const;

Returns String name of a data member of the indexed class.

### getTypeN

Gets the type number of the indexed class for which the key-field object was created.

ooTypeNumber getTypeN() const;

### isConsistent

Checks whether this key-field object is still consistent with the corresponding data member of the indexed class.

ooBoolean isConsistent();

- Returns Returns occTrue if this key-field object is still consistent with the data member with which it was initialized. Returns occFalse if any of the following is true:
  - The type of the data member has been changed.
  - The data member has been deleted
  - The indexed class has been deleted.
- Discussion You should test for consistency after performing a schema-evolution operation that affects an indexed class or a data member of an indexed class.

### isNamed

Tests whether the data member corresponding to this key-field object has the specified name.

virtual ooBoolean isNamed(const char \*name) const;

Parameters name

String name to be matched against the data-member name.

Returns oocSuccess if successful; otherwise oocError.

# ooLessThanEqualLookupField Class

#### Inheritance: ooLookupFieldBase->ooLessThanEqualLookupField

The non-persistence-capable indexing class <code>ooLessThanEqualLookupField</code> represents a *lookup field* that tests whether the value of an indexed object's key field is less than or equal to (<=) the specified comparison value.

See:

• "Reference Index" on page 388 for a list of member functions

Applications that use indexes must include the ooIndex.h header file.

## About Less-Than-Or-Equal-To Lookup Fields

A lookup field is part of a *lookup key* (an instance of the <u>ooLookupKey</u> class that is used for looking up persistent objects in an index). A lookup key consists of one or more lookup fields, each representing a condition that a found object must satisfy. Lookup fields are instances of the concrete classes derived from the abstract base class <u>ooLookupFieldBase</u>.

The concrete class <code>ooLessThanEqualLookupField</code> represents a condition that tests whether the values of a particular key field are less than or equal to a particular comparison value. You use the class constructor to specify the key field and the value.

For complete information about using lookup fields in lookup keys, see "About Lookup Keys" on page 399

# **Reference Index**

<u>ooLessThanEqualLookupField</u>

Constructs a new lookup field for testing whether values of the specified key field are less than or equal to the specified comparison value.

# Constructors

## ooLessThanEqualLookupField

Constructs a new lookup field for testing whether values of the specified key field are less than or equal to the specified comparison value.

- 1. ooLessThanEqualLookupField(
   const ooKeyField &field,
   const void \*valuePtr);
- 2. ooLessThanEqualLookupField( const ooTypeNumber typeN, const char \*memberName, const void \*valuePtr);

#### Parameters

field

Key-field object representing a key field of the indexed class. The values of the specified key field will be tested by this lookup field.

#### typeN

Type number of the indexed class.

#### memberName

Name of a data member that serves as a key field of the indexed class. The values of the specified key field will be tested by this lookup field.

The specified data member must be defined or inherited by the class specified by *typeN*. You can qualify the name of an inherited data member using the following notation (where *baseClassName* is the name of the base class that defines the inherited data member):

baseClassName::dataMemberName

You *must* qualify the name of an inherited data member if the member name is ambiguous (for example, the same name is defined in both the base class and the class to be indexed) or if the member name is not visible due to access control.

#### valuePtr

Data value to which key-field values are to be compared. The type of this value must match the key-field type (the type of the data member you specified to the key-field object). Using data of any other type may have unpredictable results.

Discussion This lookup field causes the values of the specified key field to be tested when an iterator scans the index; an indexed object is found if its tested value is less than or equal to the value specified by *valuePtr*.

This lookup field will be ignored if *field* does not specify a key field of the index being searched.

# ooLessThanLookupField Class

#### Inheritance: ooLookupFieldBase->ooLessThanLookupField

The non-persistence-capable indexing class <code>ooLessThanLookupField</code> represents a *lookup field* that tests whether the value of an indexed object's key field is less than (<) the specified comparison value.

See:

• "Reference Index" on page 392 for a list of member functions

Applications that use indexes must include the ooIndex.h header file.

## **About Less-Than Lookup Fields**

A lookup field is part of a *lookup key* (an instance of the <u>ooLookupKey</u> class that is used for looking up persistent objects in an index). A lookup key consists of one or more lookup fields, each representing a condition that a found object must satisfy. Lookup fields are instances of the concrete classes derived from the abstract base class <u>ooLookupFieldBase</u>.

The concrete class <code>ooLessThanLookupField</code> represents a condition that tests whether the values of a particular key field are less than a particular comparison value. You use the class constructor specify the key field and the value.

For complete information about using lookup fields in lookup keys, see "About Lookup Keys" on page 399

## **Reference Index**

<u>ooLessThanLookupField</u>

Constructs a new lookup field for testing whether values of the specified key field are less than the specified comparison value.

# Constructors

### ooLessThanLookupField

Constructs a new lookup field for testing whether values of the specified key field are less than the specified comparison value.

- 1. ooLessThanLookupField(
   const ooKeyField &field,
   const void \*valuePtr);
- 2. ooLessThanLookupField( const ooTypeNumber typeN, const char \*memberName, const void \*valuePtr);

#### Parameters

field

Key-field object representing a key field of the indexed class. The values of the specified key field will be tested by this lookup field.

#### typeN

Type number of the indexed class.

#### memberName

Name of a data member that serves as a key field of the indexed class. The values of the specified key field will be tested by this lookup field.

The specified data member must be defined or inherited by the class specified by *typeN*. You can qualify the name of an inherited data member using the following notation (where *baseClassName* is the name of the base class that defines the inherited data member):

baseClassName::dataMemberName

You *must* qualify the name of an inherited data member if the member name is ambiguous (for example, the same name is defined in both the base class and the class to be indexed) or if the member name is not visible due to access control.

#### valuePtr

Data value to which key-field values are to be compared. The type of this value must match the key-field type (the type of the data member you specified to the key-field object). Using data of any other type may have unpredictable results.

Discussion This lookup field causes the values of the specified key field to be tested when an iterator scans the index; an indexed object is found if its tested value is less than the value specified by *valuePtr*.

This lookup field will be ignored if *field* does not specify a key field of the index being searched.

# ooLookupFieldBase Class

Inheritance: ooLookupFieldBase

The non-persistence-capable class ooLookupFieldBase is the abstract base class for classes that represent *lookup fields*.

See:

"Reference Index" on page 396 for a list of member functions

Applications that use indexes must include the ooIndex.h header file.

## **About Lookup Fields**

A lookup field is part of a *lookup key* (an instance of the <u>ooLookupKey</u> class that is used for looking up persistent objects in an index). A lookup key consists of one or more lookup fields, each representing a condition that a found object must satisfy. When a lookup key contains multiple lookup fields, every found object must satisfy all of the conditions.

The concrete classes derived from the ooLookupFieldBase class each represent a specific condition to be satisfied. These are: ooEqualLookupField, ooGreaterThanEqualLookupField, ooGreaterThanLookupField, ooLessThanEqualLookupField, and ooLessThanLookupField.

Because the oolookupFieldBase class is abstract, you never instantiate it; instead, you work with instances of its concrete derived classes. You should not create your own subclasses of this class.

## **Reference Index**

isNamed	Compares the specified name to the name of the key field whose values will be tested by this lookup field.
<u>ooLookupFieldBase</u>	Constructs a new lookup field.

# Constructors

### ooLookupFieldBase

Constructs a new lookup field.

```
ooLookupFieldBase(
    const ooKeyField &field,
    const void *valuePtr,
    ooRelatOp relatOp = oocEQ);
```

Parameters field

Key-field object representing a key field of the indexed class. The values of the specified key field will be tested by this lookup field.

#### valuePtr

Data value to which key-field values are to be compared. The type of this value must match the key-field type (the type of the data member you specified to the key-field object). Using data of any other type may have unpredictable results.

```
relat0p
```

Relational operator used by this lookup field. This parameter value is supplied by the predefined concrete derived classes.

## **Member Functions**

#### isNamed

Compares the specified name to the name of the key field whose values will be tested by this lookup field.

virtual ooBoolean isNamed(const char \*name) const;
Returns oocTrue if the specified name matches the key-field name; otherwise oocFalse.

Parameters name

Name to be matched by the key-field name.

# ooLookupKey Class

#### Inheritance: ooLookupKey

The non-persistence-capable class oolookupKey represents a *lookup key* used for looking up persistent objects in an index.

See:

- "Reference Summary" on page 405 for an overview of member functions
- "Reference Index" on page 406 for a list of member functions

Applications that use indexes must include the ooIndex.h header file.

## About Lookup Keys

A lookup key is a transient object for looking up persistent objects in an index. A lookup key consists of one or more *lookup fields*, each representing a condition that the found objects must satisfy. When a lookup key contains multiple lookup fields, every found object must satisfy all of the conditions. Lookup fields are instances of the concrete classes derived from <u>ooLookupFieldBase</u>.

You create a lookup key and add lookup fields to it in a transaction. You then use the lookup key to initialize an object iterator to scan a particular storage object. Because lookup keys are transient, they become invalid after you commit the transaction in which they were created.

As an alternative to performing a predicate scan, an application can scan a storage object with a lookup key that represents the condition to be met. Such a scan, called an *index scan*, initializes an object iterator to find objects by searching a compatible index on the storage object. If no such index exists, no objects are found. In contrast, an object iterator initialized by a predicate scan can find objects in the scanned storage object even if no compatible index exists. An index scan can search an index even if indexes are disabled; disabling indexes affect only predicate scans, not index scans.

#### **Lookup Fields**

Every lookup field in a lookup key specifies a relational operator (such as equal or greater than) and a comparison value for testing a particular key field of an index. For example, to search an index of Person objects whose key fields are age and weight, you could create lookup fields that represent the conditions age = 40 and weight > 150.

A lookup key can contain two conditions that apply to the same key field—for example, age > 20 and age < 40. Because the found objects must satisfy all of the conditions (a conjunction), you may not specify constraints that could only be satisfied by a disjunction—for example, age < 20 and age > 40.

Lookup fields are instances of classes derived from <code>ooLookupFieldBase</code>. Each such derived class corresponds to a permitted relational operator, as shown in the following table:

Relational Operator	Lookup-Field Class
=	ooEqualLookupField
>	ooGreaterThanLookupField
<	<u>ooLessThanLookupField</u>
>=	ooGreaterThanEqualLookupField
<=	ooLessThanEqualLookupField

As the table illustrates, a lookup field cannot perform string matching or test an application-defined relational operator.

To create a lookup field, instantiate the lookup-field class that corresponds to the desired relational operator. For example, instantiate <code>ooEqualLookupField</code> to represent the relational operator =.

When you create a lookup field, parameters to the constructor specify the indexed class, the attribute to be tested, and the value to be compared with the attribute value. You can specify both the indexed class and attribute with a key-field object; alternatively, you can specify the class by its type number and the attribute by its name.

#### **Compatible Indexes**

A lookup key can search only *compatible indexes*—that is, indexes created from a compatible key description. A lookup key and a key description are compatible only if both of the following are true:

- The lookup key was created for the same indexed class as the key description—that is, the type number specified to the lookup key constructor must match the type number specified to the key description constructor.
- The lookup key contains at least one lookup field that corresponds to the primary key field of the index—that is, at least one lookup field specifies the first key-field object that was added to the key description.

An object iterator initialized with a lookup key finds objects only if they are referenced by a compatible index within the storage object being scanned. If the storage object has no compatible index, the object iterator is initialized to null, so it will find no objects. This is true even if the container, database, or federated database being scanned actually contains objects that would satisfy the lookup key's conditions. You can use the <u>anyIndex</u> member function on a lookup key to test whether a given storage object contains a compatible index.

#### Lookup Fields Used in an Index Scan

When a compatible index is searched by a lookup key with multiple lookup fields, the lookup key uses at least one lookup field during the search (the lookup field corresponding to the index's primary key field). However, the other lookup fields may or may not be used, depending on how they correspond to the index's key fields. In general, a lookup key uses only the lookup fields that correspond to the first *n* consecutive key fields in the index's key description, where *n* is an integer from 1 to the number of key fields in the index. Any other lookup fields are ignored.

For example, assume an index of Person objects has key fields age, weight, and height, and you create five lookup keys whose lookup fields correspond to the combinations of key fields listed in the following table. If you then use each lookup key to search the index, certain lookup fields would be used or ignored, as indicated in the table:

Key Fields Referenced by the Lookup Key	Key Fields Tested by the Lookup Key
age	age; the sole lookup field is used because it corresponds to the primary key field.
age,weight	age, weight; both lookup fields are used because they correspond to the first 2 key fields of the index.

Key Fields Referenced by the Lookup Key	Key Fields Tested by the Lookup Key
age, weight, height	age, weight, height; all lookup fields are used because they correspond to the first 3 key fields of the index.
age, height	age; the lookup field corresponding to height is ignored, because height is the index's third key field, and the second key field (weight) is not represented in the lookup key.
weight, height	None; both lookup fields are ignored because the primary key field (age) is not represented in the lookup key.

A lookup key ignores a lookup field that does not correspond to any key field of the index being searched. For example, assume you want to search an index of Person objects with key fields age, weight, and height. If you search the index using a lookup key that has fields corresponding to age, weight, and eyecolor, the lookup field for testing eyecolor is ignored.

# **WARNING** If lookup fields are ignored by a particular scan operation, the found objects are not guaranteed to match the entire condition represented by the lookup key. Instead, they match a conjunction of the conditions represented by the lookup fields that were tested.

Whenever a lookup field is ignored, the application is responsible for explicitly testing the appropriate data member of each found object. To avoid confusion, you should use a lookup key only if you know that a compatible index exists in which *all* lookup fields can be tested.

#### Finding Indexed Objects With a Lookup Key

To use a lookup key to find an object in a particular container, database, or federated database:

- 1. Create a lookup key by calling the ooLookupKey class constructor within a transaction.
- **2.** Create a lookup field for each condition to be satisfied by the found objects. You must create at least one lookup field that tests the primary key field of the indexed objects. To create a lookup field:
  - **a.** Choose the lookup-field class that corresponds to the desired relational operator. For example, choose ooLessThanLookupField to represent the relational operator <=.

**b.** Invoke the constructor of the chosen class, specifying the desired comparison value and the key field to be tested.

**Note:** You specify a range of values for testing a single key field by creating two lookup fields. For example, to find indexed objects whose age key field is between 15 and 30, you create lookup fields that express the conditions age  $\geq 15$  and age  $\leq 30$ .

- **3.** Add each lookup field to the lookup key by calling the addField member function on the lookup key. **Note:** The order in which lookup fields are added determines the order in which key fields are tested. For optimal performance:
  - Keep the number of lookup fields small.
  - Add lookup fields to the lookup key in the same order in which the corresponding key fields were added to the key description, starting with the most major key field.
  - Specify as many equality (=) comparisons as possible. If you are adding multiple lookup fields, only the last one should represent a relational operator other than =.
- 4. Test whether the lookup key is compatible with any indexes in the storage object you plan to search. To do this, call the anyIndex member function on the lookup key.
- 5. Use the lookup key to initialize an object iterator of type <code>ooltr(className)</code>, where <code>className</code> is the name of the indexed class. To initialize the object iterator, call the <code>scan</code> member function on it, specifying the lookup key and the container, database, or federated database to be searched.

If no compatible index exists in the specified storage object, the object iterator is initialized with a null iteration set. If a compatible index exists, the object iterator is initialized to find objects satisfying the applicable lookup fields, in ascending order.

The lookup key and its lookup fields must exist (remain in scope) while you are initializing and advancing the object iterator; if they have been destructed, the iteration will fail.

**EXAMPLE** Assume an index is created from a key description whose key-field objects were added in the following order: plantNumber, deptNumber, groupNumber, projLeaderId.

A lookup key whose lookup fields represent the following conditions would produce a fast search, because they match the key-field order and because only the last condition is not an equality comparison.

```
plantNumber = 1
deptNumber = 6
groupNumber > 15
```

The following search key looks for groups with group numbers ranging from 15 to 30:

```
plantNumber = 1
deptNumber = 6
groupNumber >= 15
groupNumber <= 30</pre>
```

```
EXAMPLE This example defines the persistence-capable class Point in a DDL file. The application code creates an index of Point objects with xCoord and yCoord as the primary and secondary key fields.
```

```
// DDL file
class Point : public ooObj {
public:
   char name[32];
   int32 xCoord;
   int32 yCoord;
};
// Application code
ooHandle(ooKeyField) keyFieldH;
ooHandle(ooKeyDesc) keyDescH;
ooHandle(ooContObj) contH;
                   // Set contH to the desired container.
...
// Create the key description in the container.
keyDescH = new(contH) ooKeyDesc(ooTypeN(Point), oocTrue);
// Define the primary key field; add it to the key description.
keyFieldH = new(keyDescH) ooKeyField(ooTypeN(Point), "xCoord");
keyDescH->addField(keyFieldH);
// Define the secondary key field; add it to the key description.
keyFieldH = new(keyDescH) ooKeyField(ooTypeN(Point), "yCoord");
keyDescH->addField(keyFieldH);
// Create an index in the container.
ooStatus status = keyDescH->createIndex(contH);
```

Another transaction (or a different program) creates a lookup key to search for indexed <code>Point</code> objects that have an x-coordinate equal to 100 and a y-coordinate

greater than 50. The lookup key is used to initialize an object iterator to find such points.

```
// Application code
ooHandle(ooContObj) contH;
ooItr(Point) pointI;
const int xPoint = 100;
                           // x == 100
const int yMinPoint = 50;
                            //v > 50
            // Set contH to the container containing the index.
...
// Create a lookup key.
ooLookupKey lookupKey(ooTypeN(Point), 2);
// Create the lookup fields and add them to the lookup key.
ooEqualLookupField xLookupField(
   ooTypeN(Point), "xCoord", &xPoint);
lookupKey.addField(xLookupField);
ooGreaterThanLookupField yLookupField(
   ooTypeN(Point), "yCoord", &yMinPoint);
lookupKey.addField(yLookupField);
// Initialize the object iterator by scanning the container
// with the lookup key.
if (pointI.scan(contH, lookupKey)) {
   // Advance the object iterator
   while (pointI.next()) {
                   "Found point at (" << pointI->xCoord
      cout <<
                   ", " << pointI->yCoord << ")." << endl;
             <<
   }
}
```

## **Reference Summary**

Creating	ooLookupKey
Modifying	<u>addField</u>
Getting Information	anyIndex nField

## **Reference Index**

addField	Adds the specified lookup field to this lookup key.
anyIndex	Tests whether this lookup key is compatible with one or more indexes in the specified storage object.
nField	Gets the number of lookup fields that have been added to this lookup key .
<u>ooLookupKey</u>	Constructs a new lookup key with the specified characteristics.

## Constructors

## ooLookupKey

Constructs a new lookup key with the specified characteristics.

```
ooLookupKey(
    const ooTypeNumber typeN,
    const uint32 number);
```

#### Parameters typeN

Type number of the persistence-capable class of indexed objects. For the lookup key to be valid and usable,  $t_{YP}eN$  must match exactly the  $t_{YP}eN$  that was used in creating the key description.

number

Number of lookup fields you intend to add to the lookup key. You cannot add more fields to the lookup key than specified by *number*, but you can add fewer fields. This number cannot be changed after the lookup key is created. Every lookup field counts separately, even if multiple lookup fields test the values of the same key field.

## **Member Functions**

#### addField

Adds the specified lookup field to this lookup key.

```
ooStatus addField(const ooLookupFieldBase &field);
```

Parameters	field
	Lookup field to be added.
Returns	oocSuccess, if the lookup field is added successfully; otherwise, oocError if the lookup key is full.
Discussion	You may add a lookup field to a lookup key only if both are defined on the same class.
	You can add multiple lookup fields to the same lookup key.
	The order in which lookup fields are added to a lookup key affects the performance of the lookup. For optimal performance, you should add lookup fields to the lookup key in the same order in which the corresponding key fields were added to the key description, starting with the primary key field.
	The same lookup field can be added to multiple lookup keys.

## anyIndex

Tests whether this lookup key is compatible with one or more indexes in the specified storage object.

	<pre>ooBoolean anyIndex(     const ooHandle(ooContObj) &amp;storageObject) const;</pre>	
	<pre>ooBoolean anyIndex(     const ooHandle(ooDBObj) &amp;storageObject) const;</pre>	
	<pre>ooBoolean anyIndex(     const ooHandle(ooFDObj) &amp;storageObject) const;</pre>	
Parameters	torageObject Handle to the container, database, or federated database you want to te a compatible index. An error is reported if any other type of handle is specified.	st for
Returns	ocTrue if there is a compatible index, oocFalse if there is no compatible	index.
Discussion	If this member function returns <code>oocTrue</code> , then you can use <code>scan</code> on an object iterator to look up objects using an index.	

Key Fields	Lookup Fields	anyIndex Return Value
x	x	oocTrue
x	У	oocFalse
х, у	x	oocTrue
х, у	х, у	oocTrue
х, у	У	oocFalse
x, y, z	x	oocTrue
x, y, z	х, у	oocTrue
x, y, z	х, у, z	oocTrue
х, у, z	X, Z	oocTrue (but only the lookup field for x will be used; the application must test z values explicitly)

The following table shows the return values for various combinations of indexed key fields (x, y, and z) and corresponding lookup fields.

## nField

Gets the number of lookup fields that have been added to this lookup key.

uint32 nField() const;

Returns An integer number of lookup fields.

# ooMap Class

Inheritance: ooObj->ooMap

Handle Class: ooHandle(ooMap)

```
Object-Reference Class: ooRef(ooMap)
```

The persistence-capable class ooMap represents unordered name maps.

See:

- "Reference Summary" on page 412 for an overview of member functions
- "Reference Index" on page 413 for a list of member functions

To use this class, your application must include the ooMap.h header file. No extra linking is required.

## **About Name Maps**

A name map is a collection whose elements (instances of <u>ooMapElem</u>) are key-value pairs, where each key is a string and each value is an object reference to a persistent object. No two elements of the name map may have the same key.

Each element of a name map defines a mapping from a name to an object with that name. The word *name* is used as a synonym for *key* because an object's key is often a name; however, a key can be any valid C++ string that identifies the object. Keys can be strings of any length.

A name map provides an efficient way to assign identifying keys to objects and to look up objects by their keys. Objectivity/DB uses name maps to implement its dictionaries of scope-named objects. In addition, you can instantiate this class to create your own application-specific dictionary for objects.

#### **Growth Characteristics**

A name map is implemented as a hash table using a traditional hashing mechanism. The hash table can grow dynamically; however, increasing its size requires rehashing the entire hash table.

When you create a name map, you specify:

- The *initial number of bins* (hash buckets). For optimal performance, the number of bins should always be a prime number.
- The maximum average density, that is, the average number of elements per bin allowed before the hash table must be resized. The hash table is resized whenever:

totalElements >= numberBins \*maximumAverageDensity

■ The *growth factor*. This number gives the percentage by which the hash table grows when it is resized. Each time the hash table is resized, the number of bins is increased by the growth factor, then rounded up to the nearest prime number.

#### **Performance Considerations**

Name maps are good for looking up one object at a time. They give better performance than indexes on data involving frequent updates. Indexes perform better on read-only data.

To avoid rehashing the name map, you need a good estimate of the initial number of hash bins. Also, you should put the name map in a container by itself (its <code>ooMapElem</code> elements will automatically be clustered with it). Clustering other objects with a name map and its elements will result in poor performance of the name map.

#### **Runtime Statistics**

Name maps accumulate statistical information about runtime usage. You can print a summary of this information or clear the statistical parameters, resetting them to zero. The statistical parameters record:

- The number of elements added.
- The number of elements deleted.
- The number of times the hash table was rehashed.

#### **Referential Integrity**

By default, when you add an object to a name map, Objectivity/C++ automatically creates a unidirectional association from the added object to the name-map element that references it. Because this association has delete propagation enabled, deleting the object automatically causes the name-map element to be deleted as well, preserving the referential integrity of the name map.

After you create a name map and before you add any elements, you can call its <u>set\_refEnable</u> member function to disable the automatic maintenance of its referential integrity. When you do so, you reduce the overhead in adding and deleting elements; however, you become responsible for ensuring that the name map does not contain any dangling references to deleted objects.

#### **Hash Function**

For efficient lookup, elements of name maps are stored in a hash table; hash values are computed from the key of each element. You can define your own hashing function if you desire.

A hash function takes two parameters: the string from which to compute the hash value, and the number of bins in the hash table. It returns the hash value for the specified string, which must be between 0 and one less than the number of bins.

All name maps that the application accesses use the same hash function. To install an application-defined hash function for name maps, call the static member function <code>ooMap::set\_nameHashFunction</code>.

```
EXAMPLE This example shows how to install the function myNameHash as the hash function for all name maps in the current application.
```

```
uint32 myNameHash(const char *name, const uint32 modulus)
{
    // This function should return a value between 0 and modulus-1
    ...
}
// Set the name hash function used in name map
// to be myNameHash
ooMap::set_nameHashFunction(myNameHash);
```

## Working With a Name Map

As is the case for any basic object, you specify whether a name map is to be transient or persistent when you create it; name maps *must be persistent*. You create a name map with a call to the new operator; the clustering directive in that call specifies where in the federated database to store the new name map.

Like other persistent objects, name maps are normally manipulated through handles or object references. You can store and find an unordered name map in the database just as you would any other persistent object.

## **Related Classes**

The class <code>ooMapElem</code> represents an element in a name map. Each element is a key-value pair that associates a name with a persistent object. The key is a C++ string and the value is an object reference to a persistent object.

The class ooMapItr defines an iterator for name maps; you can use an iterator to obtain each element of the name map.

Two additional classes represent persistent collections of key-value pairs.

- ooHashMap represents an unordered object map, that is, a collection of key-value pairs in which both the key and the value are persistent objects. It uses an extendible hashing mechanism that allows the hash table to grow without being completely rehashed.
- ooTreeMap represents a sorted object map.

## **Reference Summary**

In the following table, operators indicated as *(inherited)* are overloaded in this class with no change in behavior; they are documented with the 000bi class (page 431), along with the other inherited member functions not listed here.

Creating and Deleting	<u>ooMap</u> <u>operator new</u> (inherited) <u>operator delete</u> (inherited)
Adding and Removing Elements	<u>add</u> <u>forceAdd</u> <u>remove</u> <u>replace</u>
Finding Objects	lookup

Getting Information	<pre>maxAvgDensity nameHashFunction nBin nElement percentGrow printStat</pre>
Testing	<u>isMember</u> <u>refEnable</u>
Working With Runtime Statistics	<u>clearParam</u> printStat
Maintaining the Hash Table	<u>refEnable</u> <u>rehash</u> <u>set_nameHashFunction</u> <u>set_refEnable</u>
Static Utilities	set_nameHashFunction nameHashFunction

# **Reference Index**

add	Adds a new element to this name map.
clearParam	Resets this name map's statistical parameters to zero.
forceAdd	Forces the addition of an element to the name map without checking for the prior existence of an element with the same name.
isMember	Tests whether any element in the name map contains the key name.
lookup	Looks up the specified name in this name map and returns either the corresponding object or the name-map element.
maxAvgDensity	Returns the allowable maximum average density of this name map's hash table.
nameHashFunction	Gets the hash function used by all name maps in the federated database.
nBin	Returns the number of bins in this name map's hash table.

<u>nElement</u>	Returns the total number of elements in this name map's hash table.
ooMap	Constructs a new name map with the specified growth characteristics.
percentGrow	Returns the percentage by which this name map's hash table grows when it is resized.
<u>printStat</u>	Prints runtime statistical information about this name map to the specified file.
<u>refEnable</u>	Tests whether Objectivity/DB automatically maintains the referential integrity of this name map.
<u>rehash</u>	Resizes this name map's hash table.
remove	Removes the specified element from this name map.
<u>replace</u>	Replaces the object associated with the specified name in this name map.
set_nameHashFunction	Sets the hash function for all instances of $ooMap$ in the federated database.
<u>set_refEnable</u>	Enables or disables maintenance of referential integrity for this name map.

## Constructors

#### ооМар

Constructs a new name map with the specified growth characteristics.

```
ooMap(
```

```
const uint32 nBin = oocMapInitHashBinSize,
const uint32 maxAvgDensity = oocMapMaxAvgDensity,
const uint32 percentGrowth = oocMapPercentGrow);
```

#### Parameters nBin

Initial number of bins in the new name map's hash table. The default value is 11.

#### maxAvgDensity

Average number of elements per bin allowed before hash table is resized. For example, if the total number of elements in the table is greater than or equal

to nBins \*maxAvgDensity, the table is resized. The default value of this parameter is 5.

percentGrowth

Parameter (expressed as percent) used to resize the hash table. The default value is 100. This growth factor is used to multiply the original bin size to get the new bin size. The system rounds the number of bins after rehashing to a prime number.

## **Member Functions**

#### add

Adds a new element to this name map.

	<pre>ooStatus add(     const char *name,     const ooRef(ooObj) &amp;objR,     ooRefHandle(ooMapElem) &amp;elem = oocDefaultMapElemHandle);</pre>
Parameters	name
	String name; specifies the key in the added element.
	objR
	Object reference to a persistent object; specifies the value in the added element.
	elem
	Object reference or handle to set to the new element. By default, <i>elem</i> is a null map-element object reference or handle.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	If you specify <i>elem</i> , you can use the initialized object reference or handle to perform operations such as setting a unidirectional association to the new element.
See also	forceAdd remove

#### clearParam

Resets this name map's statistical parameters to zero.

ooStatus clearParam();

Returns oocSuccess if successful; otherwise oocError.

#### forceAdd

Forces the addition of an element to the name map without checking for the prior existence of an element with the same name.

	ooStatus forceAdd(
	const char * <i>name</i> ,
	<pre>const ooRef(ooObj) &amp;objR,</pre>
	<pre>ooRefHandle(ooMapElem) &amp;elem = oocDefaultMapElemHandle);</pre>
Parameters	name
	String name; specifies the key in the added element.
	objR
	Object reference to a persistent object; specifies the value in the added element.
	elem
	Object reference or handle to set to the new element. The default is a null reference to a map-element handle.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	If there is already an element with this name in the table, two elements will contain the same name after a successful forceAdd operation. The element that is found by lookup is undefined. You should only use this member function if you know there is no element with the same name.
	If you specify <i>elem</i> , you can use the initialized object reference or handle to perform operations such as setting a unidirectional association to the new element.
See also	add

## isMember

Tests whether any element in the name map contains the key name.

ooBoolean isMember(const char \*name);

Parameters	name
	Name to be checked.
Returns	oocTrue if any element contains the name, otherwise oocFalse.
lookup	
	Looks up the specified name in this name map and returns either the corresponding object or the name-map element.
	<pre>1. ooRef(ooObj) lookup( const char *name);</pre>
	<pre>2. ooStatus lookup( const char *name, ooRefHandle(ooObj) &amp;object, const ooMode openMode = oocRead);</pre>
	<pre>3. ooStatus lookup( const char *name, ooRefHandle(ooMapElem) &amp;elem);</pre>
Parameters	name
	Name (key) of the object to be looked up.
	object
	Object reference or handle to set to the found object. <i>object</i> is set to null if the specified name is not found or if an error occurs.
	openMode
	Mode in which to open the found object. Specify $\verb+oocNoOpen+$ to prevent the object from being opened.
	elem
	Object reference or handle to set to the found name-map element.
Returns	(Variant 1) A type-independent object reference to the found object; you must cast the returned object reference to the appropriate type.
	(Variants 2 and 3) oocSuccess if successful; otherwise oocError.

## maxAvgDensity

Returns the allowable maximum average density of this name map's hash table.

uint32 maxAvgDensity() const;

#### nameHashFunction

 Gets the hash function used by all name maps in the federated database.

 static ooNameHashFuncPtr nameHashFunction();

 Returns
 A pointer to the hash function currently in use.

 See also
 set nameHashFunction

 nBin
 Returns the number of bins in this name map's hash table.

 uint32 nBin() const;

#### nElement

Returns the total number of elements in this name map's hash table.

```
uint32 nElement() const;
```

#### percentGrow

Returns the percentage by which this name map's hash table grows when it is resized.

uint32 percentGrow() const;

Discussion Each time the hash table is resized, the number of bins is increased by the growth factor, then rounded up to the nearest prime number.

#### printStat

	Prints runtime statistical information about th	is name map to the specified file.
	<pre>void printStat(FILE *outFile = stdou</pre>	ut) const;
Parameters	outfile Name of the output file.	
Discussion	The output from this member function is similar Run statistics of ooMap #2-3-3-3 (Fr	<b>lar to:</b> ri Aug 25 15:22:45 PDT 1992)
	** Number of elements added ** Number of elements removed ** Number of rehashes	=> 0 => 0 => 0

* *	Current state:	
	Number of bins	=> 23
	Number of elements	=> 100
	Maximum average density	=> 5
	Percent Growth	=> 100
	Average length per bin	=> 4.34783
	Maximum length	=> 7
	Standard deviation of length	=> 1.08783
	Maintain referential integrity	=> Yes

#### refEnable

	Tests whether Objectivity/DB automatically maintains the referential integrity of this name map.		
	ooBoolean refEnable() const;		
Returns	oocTrue <b>if referential integrity is being maintained automatically; otherwise</b> oocFalse.		
Discussion	By default, Objectivity/DB maintains the referential integrity of name maps. You can disable the automatic maintenance of referential integrity for a particular name map by calling its <u>set_refEnable</u> member function.		
See also	<u>set refEnable</u>		
rehash			
	Resizes this name map's hash table.		
	<pre>ooStatus rehash(     const uint32 binSize = oocInitMapHashBinSize);</pre>		
Parameters	<i>binSize</i> New number of bins in the hash table. If this parameter is 0, the default		

New number of bins in the hash table. If this parameter is 0, the default number of bins (11) is used. For optimal performance, the number of bins should be a prime number.

- Returns oocSuccess if successful; otherwise oocError.
- Discussion Objectivity/DB automatically increases the size of the hash table as necessary. The only time you should need to call this member function is when the number of bins is very large and the total number of elements is relatively small. This situation can occur when many elements are deleted after the hash table has grown to its peak size.

remove	
	Removes the specified element from this name map.
	1. ooStatus remove(const char *name);
	<pre>2. ooStatus remove(const ooHandle(ooMapElem) &amp;objH);</pre>
Parameters	name Name of the element to be removed.
	objH Handle to the element to be removed.
Returns	oocSuccess if successful; otherwise oocError. This function returns oocSuccess if the specified element does not exist.
Discussion	If forceAdd was used to add more than one element with the same name, remove only deletes the first such element it finds and gives no indication that there are other elements with the same name.
See also	add
replace	
	Replaces the object associated with the specified name in this name map.
	<pre>ooStatus replace(     const char *name,     const ooRef(ooObj) &amp;objR,     ooRefHandle(ooMapElem) &amp;elem = oocDefaultMapElemHandle);</pre>
Parameters	name
	Name of the element to be replaced.
	objR Object reference of the referenced element
	Object reference or handle to set to the replaced element. The default is a null reference to a map-element handle.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	If this name map contains an element with the key <i>name</i> , this member function replaces that element's value with the specified object reference. Otherwise, this

member function adds a new element whose key is *name* and whose value is the specified object reference.

If *elem* is given, it is initialized to reference the element whose key is *name*. You can use this initialized object reference or handle to perform operations such as setting a unidirectional association to the new element.

#### set\_nameHashFunction

Sets the hash function for all instances of OoMap in the federated database.

static	void	set_	_nameHa	shFunction(
00	NameH	lashF	<sup>r</sup> uncPtr	<pre>hashFunction);</pre>

Parameters	hashFunction
	Function pointer to the new hash function.

Discussion After a new hash function is set, existing name maps must be rehashed before they can be accessed.

The hash function is set as an attribute of the <code>ooMap</code> class and so is not stored as part of any persistent instance. Consequently, the same hash function must be set by every application that is to access a name map in the federated database.

See also <u>nameHashFunction</u> <u>ooNameHashFuncPtr</u> global type

#### set\_refEnable

	Enables or disables maintenance of referential integrity for this name map.		
	<pre>ooStatus set_refEnable(ooBoolean refEnable = oocTrue);</pre>		
Parameters	<i>refEnable</i> Specifies whether to enable (oocTrue) or disable (oocFalse) referential integrity.		
Returns	oocSuccess if successful; otherwise oocError.		
Discussion	By default, Objectivity/DB maintains the referential integrity of name maps. You may disable referential integrity to reduce the overhead of adding and deleting elements. However, if you disable referential integrity, you are responsible for deleting elements from the name map when the corresponding objects are deleted from the database.		

If you call this member function, you must call it before any elements are added to this name map.

See also <u>refEnable</u>

# ooMapElem Class

Inheritance: ooObj->ooMapElem

Handle Class: ooHandle(ooMapElem)

```
Object-Reference Class: ooRef(ooMapElem)
```

The persistence-capable class <code>ooMapElem</code> represents individual elements in a name map (an instance of <u>ooMap</u>).

See:

- "Reference Summary" on page 424 for an overview of member functions
- "Reference Index" on page 424 for a list of member functions

To use this class, your application must include the ooMap.h header file. No extra linking is required.

## **About Name-Map Elements**

Each element of a name map is a key-value pair that associates a name with a persistent object. The key is a C++ string and the value is an object reference to a persistent object. The ooMapElem class has member functions to get the key and value from a name-map element and to replace the value in a name-map element.

## Working With a Name-Map Element

The ooMapElem class is transparent to applications in most cases; the class ooMap provides the interface to managing a name map's elements. Thus, instances of the ooMapElem class are not created or deleted by an application directly, but

through calls to a name map's <u>add</u> and <u>remove</u> member functions. You can obtain an existing name-map element from the name map's <u>lookup</u> member function.

Like other persistent objects, name-map elements are manipulated through handles and object references. The iterator class <code>ooMapItr</code> iterates through the name-map elements in a name map.

# **Reference Summary**

Getting a Key or Value	name oid
Setting a Value	set_oid

## **Reference Index**

name	Gets the key from this name-map element.
oid	Gets the value from this name-map element.
<u>set_oid</u>	Sets the value in this name-map element.

# **Member Functions**

name	
	Gets the key from this name-map element.
	<pre>const char *name() const;</pre>
Returns	The string that is the key in this name-map element.
oid	
	Gets the value from this name-map element.
	ooRef(ooObj) oid() const;
Returns	The object reference that is the value in this name-map element.

Discussion This member function returns a type-independent object reference, which you must cast to the appropriate type.

## set\_oid

	Sets the value in this name-map element.
	<pre>ooStatus set_oid(const ooRef(ooObj) &amp;objR);</pre>
Parameters	<i>objR</i> The object reference that is to be the new value in this name-map element.
Returns	oocSuccess if successful; otherwise oocError.

# ooMapltr Class

Inheritance: ooHandle(ooMapElem)->ooItr(ooMapElem)->ooMapItr

The non-persistence-capable class <code>ooMapItr</code> defines a *name-map iterator*—that is, an iterator for finding the elements of a name map.

See:

- "Reference Summary" on page 427 for an overview of member functions
- "Reference Index" on page 428 for a list of member functions

To use this class, your application must include the ooMap.h header file. No extra linking is required.

## **About Name-Map Iterators**

You initialize a name-map iterator to find the elements of a particular name map. The elements of the name map constitute the name-map iterator's iteration set. You can initialize a name-map iterator in either of two ways:

- Construct an initialized name-map iterators with the ooMapItr constructor, passing a handle to the name map as the parameter.
- Initialize a name-map iterator with the assignment operator (=), specifying a handle to the name map as right-hand operand.

Initializing	<u>ooMapItr</u> <u>operator=</u>
Advancing	next

## **Reference Summary**

## **Reference Index**

<u>ooMapItr</u>	Default constructor that constructs a new uninitialized iterator for finding elements in a name map.
<u>ooMapItr</u>	Constructs a new name-map iterator and initializes it with the specified name map.
<u>operator=</u>	Assignment operator; initializes this name-map iterator with the referenced map.
<u>next</u>	Advances this name-map iterator to the next name-map element.

## Constructors

#### ooMapltr

Default constructor that constructs a new uninitialized iterator for finding elements in a name map.

```
ooMapItr();
```

#### ooMapltr

	Constructs a new name-map iterator and initializes it with the specified name map.
	<pre>ooMapItr(const ooRefHandle(ooMap) ↦);</pre>
Parameters	map Object reference or handle to the map with which to initialize the name-map iterator.
Discussion	The name-map iterator is initialized so that the $next$ member function will open each name-map element for read-only access.

# **Operators**

#### operator=

Assignment operator; initializes this name-map iterator with the referenced map.

ooMapItr &operator=(ooRefHandle(ooMap) &map);

 Parameters
 map

 Object reference or handle to the map with which to initialize the name-map iterator.

 Discussion
 The name-map iterator is initialized so that the next member function will open each name-map element for read-only access.

## **Member Functions**

#### next

Advances this name-map iterator to the next name-map element.
ooBoolean next();
oocTrue if another element is found in the map; oocFalse if all of the elements in the map have been found or if an error occurred.
This member function finds each successive name-map element in a map. To access the object to which a name-map element refers, you must use the oid member function on the name-map element.
This member function opens each found name-map element for read only (oocRead), so you cannot modify a found name-map element during an iteration. You can, however, modify an object obtained from a name-map element if you first explicitly open the object for update.
You must not modify the map during an iteration—for example, by adding or deleting name-map elements. This means that, if referential integrity is enabled, you must not delete any object that is referenced by a name-map element, because the deletion will propagate to the name-map element.

Member Functions

# ooObj Class

Inheritance: ooObj

Handle Class: <u>ooHandle(ooObj)</u>

Object-Reference Class: <a href="mailto:ooRef(ooObj">ooRef(ooObj)</a>

The persistence-capable class ooObj is the base class for all classes of Objectivity/DB *objects*. The ooObj class and its corresponding handle and object-reference classes together define persistence behavior for various kinds of Objectivity/DB objects.

See:

- "Reference Summary" on page 435 for an overview of oo0bj member functions
- "Reference Index" on page 436 for a list of ooObj member functions

For operations performed through a handle or object reference, see:

"Reference Summary" on page 601

(*ODMG*) The ooObj class is equivalent to the ODMG standard class d\_Object. You can use the name d\_Object interchangeably with ooObj.

## About Objectivity/DB Objects

All Objectivity/DB objects are instances of classes that are derived from ooObj. Objectivity/DB objects include:

■ *Basic objects*, which are the fundamental units stored by Objectivity/DB. An application defines persistence-capable classes for basic objects by deriving them from oo0bj.

- Persistent objects, which include basic objects and containers. An application defines its own persistence-capable classes for containers by deriving them from ooContObj.
- Storage objects, which include standard containers, databases, and federated databases. These are instances of ooContObj, ooDBObj, and ooFDObj, respectively.
- Autonomous partitions, which partition storage objects into independent units that continue to function even if separated by network or system failures. Autonomous partitions are instances of ooAPObj.

Because ooObj is the abstract base class for all Objectivity/DB objects, you can use handles and object references of type *ooRefHandle*(ooObj) to reference any kind of Objectivity/DB object.

Because of the member functions it defines, however, ooObj should also be thought of in two more specific roles:

- As the base class for basic objects. Most of the member functions defined by oo0bj apply *only* to basic objects—for example, member functions that support copying, moving, and versioning.
- As the base class for persistent objects. Member functions that apply to both basic objects and containers include operator delete and member functions that enable you to operate on "this" object within a member function you are defining on a derived persistence-capable class.

None of the member functions defined by 000bj apply to databases, federated databases, or autonomous partitions.

# **Working With Basic Objects**

**NOTE** This section focuses on the role of ooObj as the base class for basic objects. See the <u>ooContObj</u>, <u>ooDBObj</u>, <u>ooFDObj</u>, and <u>ooAPObj</u> classes for information about working with containers, databases, federated databases, and autonomous partitions, respectively.

A basic object is an instance of any application-defined class that is derived from ooObj (but not through ooContObj). Like any persistence-capable class, an application-defined basic-object class must be defined in a DDL file and its definition must be processed by the DDL processor.

An application can create instances of 000bj, although such instances are of limited use because they can contain no application-specific data. However,
instances of ooObj can be useful for populating test databases—for example, to estimate file size, establish performance limits, and so on.

Basic objects are created and deleted from within an application. An application:

- Creates a basic object using the constructor and operator new of the appropriate basic-object class—namely, ooObj or an application-defined class appClass derived from ooObj.
- Deletes a basic object using either the ooDelete global function or the operator delete defined by ooObj.

As is the case for any persistent object, you specify whether a basic object is to be transient or persistent when you create it. A clustering directive on <code>operator</code> new specifies where to locate a new persistent basic object in the federated database.

To work with a new basic object, an application must assign the result of operator new to a handle—for example, an instance of <code>ooHandle(ooObj)</code>. A new basic object of an application-defined class <code>appClass</code> is normally assigned to an instance of the type-specific handle class <code>ooHandle(appClass)</code>.

Similarly, to identify and work with an existing basic object, the application must open it through an appropriate handle or object reference; multiple handles and object references can be set to reference the same basic object. The application then operates on an open basic object by:

- Calling various member functions on any of the referencing handles or object references.
- Passing any of the referencing handles or object references to various global functions or member functions of other classes.

The general handle and object-reference classes *ooRefHandle*(ooObj) provide the primary interface for operating on existing basic objects, with member functions for opening, locking, getting information, moving, copying, versioning, and finding referenced objects. The type-specific handle and object-reference classes *ooRefHandle*(*appClass*) inherit these member functions from *ooRefHandle*(ooObj), redefining them wherever type-specific behavior or parameters are required.

In addition, every application-defined basic-object class appClass has:

- DDL-generated member functions for creating, deleting, and accessing any associations defined by the class.
- Member functions defined by ooObj (see "Reference Summary" on page 435), which are inherited or redefined by *appClass*.
- Data members and member functions specific to *appClass*.

You can call these members directly from within a member function of *appClass* or indirectly through the indirect member-access operator (->) on a handle or object reference of type *ooRefHandle(appClass)*.

# **Support for Versioning Basic Objects**

The ooObj class defines the following bidirectional associations that support the versioning of basic objects:

```
class ooObj {
    ...
    // Links a previous version with one or more next versions.
    ooRef(ooObj) nextVers[] <-> prevVers ;
    ooRef(ooObj) prevVers <-> nextVers[] ;
    // Links the default version with its genealogy.
    ooRef(ooGeneObj) defaultToGeneObj <-> defaultVers ;
    // Links a nondefault version with its genealogy.
    ooRef(ooGeneObj) geneObj <-> allVers[] : version(copy);
    // Links a derived version with the ancestor from which it
    // is derived; represents merged version branches.
    ooRef(ooObj) derivatives[] <-> derivedFrom[];
    ooRef(ooObj) derivedFrom[] <-> derivatives[];
};
```

You enable versioning on an object by calling <u>setVersStatus</u> on a handle to the object. When versioning is enabled, opening the object for update creates a new *version* of it. A version is a bit-wise copy of the object that has its own object identifier; Objectivity/DB automatically links each new version with its previous version by setting their prevVers and nextVers associations. When linear versioning is enabled, a previous version can have at most one next version; when branch versioning is enabled, a previous version can have multiple next versions (consequently, the nextVers association is to-many).

A basic interface for versioning is provided by the object-reference and handle classes *ooRefHandle*(*ooObj*). For example, if you have a handle to an object for which versioning is enabled, you can use the handle to find other versions of the referenced object; you can set up a genealogy of versions with the referenced object as the default version; or you can find the default version (if one exists) for a particular referenced object. This basic interface automatically manages the prevVers, nextVers, defaultToGeneObj, and geneObj associations.

The 000bj class supports application-defined versioning semantics by making public the member functions that create, delete, and access the versioning

associations. These functions correspond to the standard set of functions generated for each association *linkName* defined in a persistence-capable class <u>appClass</u> (see page 81).

While Objectivity/DB continues to manage the prevVers and nextVers associations whenever new versions are created, you can use the provided functions to:

- Maintain a sequence of next and previous versions in a linear genealogy—for example, you can reset the nextVers and prevVers association to connect adjacent versions after a version is deleted.
- Manage the branches of a branched genealogy—for example, you can set the derivatives and derivedFrom associations to connect the most recent versions of two different branches, merging those branches.
- Set up a custom genealogy—for example, you can derive a genealogy class from the standard <u>ooGeneObj</u> class, and then use the defaultToGeneObj and geneObj associations to maintain a default version for the custom genealogy.

Creating and Deleting Persistent Objects	<u>ooObj</u> <u>operator new</u> <u>operator delete</u>
Working With This Persistent Object	<u>ooGetTypeN</u> <u>ooGetTypeName</u> <u>ooIsKindOf</u> <u>ooThis</u> <u>ooUpdate</u> <u>ooValidate</u>
Customizing Behavior for Basic Objects	<u>ooCopyInit</u> <u>ooPreMoveInit</u> <u>ooPostMoveInit</u> <u>ooNewVersInit</u>

# **Reference Summary**

Customized Versioning of Basic Objects— Previous and Next Versions	nextVers exist_nextVers add_nextVers sub_nextVers del_nextVers prevVers exist_prevVers set_prevVers del_prevVers
Customized Versioning of Basic Objects— Derived Versions	derivedFrom <u>exist_derivedFrom</u> <u>add_derivedFrom</u> <u>sub_derivedFrom</u> <u>del_derivedFrom</u> <u>derivatives</u> <u>exist_derivatives</u> <u>add_derivatives</u> <u>sub_derivatives</u> <u>del_derivatives</u>
Customized Versioning of Basic Objects— Default Versions	geneObj exist_geneObj set_geneObj del_geneObj defaultToGeneObj exist_defaultToGeneObj set_defaultToGeneObj del_defaultToGeneObj
ODMG Interface	<u>operator new</u> <u>mark_modified</u>

# **Reference Index**

<u>add derivatives</u>	Links this object with the specified object, making the specified object a derivative (successor version) of this object.
add_derivedFrom	Links this object with the specified object, making the specified object an ancestor version of this object.

<u>add_nextVers</u>	Links this object with the specified object, making the specified object a next version of this object.
<u>defaultToGeneObj</u>	Finds, and optionally opens, the genealogy associated with this object, if this object is the default version in a genealogy.
<u>del_defaultToGeneObj</u>	Deletes any defaultToGeneObj association from this object, indicating this object is no longer the default version.
<u>del_derivatives</u>	Deletes all derivatives associations from this object, indicating this object has no derivatives (successor versions).
<u>del_derivedFrom</u>	Deletes all derivedFrom associations from this object, indicating this object is not derived from any ancestor versions.
<u>del_geneObj</u>	Deletes any geneObj association that exists for this object, indicating this object is no longer in a genealogy.
<u>del_nextVers</u>	Deletes all nextVers associations that exist for this object, indicating that this object has no next versions.
<u>del_prevVers</u>	Deletes any prevVers association that exists for this object, indicating this object has no previous version.
<u>derivatives</u>	Initializes an object iterator to find, and optionally open, all versions that are derived from this object and that satisfy any specified selection criteria.
<u>derivedFrom</u>	Initializes an object iterator to find, and optionally open, all ancestor versions from which this object is derived, and that satisfy any specified selection criteria.
<u>exist_defaultToGeneObj</u>	Tests whether this object is the default version in any genealogy or, if a parameter is given, in the specified genealogy.
<u>exist derivatives</u>	Tests whether this object has any derivatives (successor versions); if a parameter is given, tests whether the specified object is a derivative of this object.
<u>exist_derivedFrom</u>	Tests whether this object is derived from any other objects (ancestor versions), or, if a parameter is given, from the specified object.

<u>exist geneObj</u>	Tests whether this object is a version within any genealogy, or, if a parameter is given, within the specified genealogy.
<u>exist nextVers</u>	Tests whether this object has a next version; if a parameter is given, tests whether the specified object is the next version.
<u>exist_prevVers</u>	Tests whether this object has a previous version; if a parameter is given, tests whether the specified object is the previous version.
<u>geneObj</u>	Finds, and optionally opens, the genealogy associated with this object.
<pre>mark_modified</pre>	(ODMG) Opens this persistent object for update.
<u>nextVers</u>	Initializes an object iterator to find, and optionally open, all next versions that are created from this object and that satisfy any specified selection criteria.
<u>ooCopyInit</u>	Default implementation for performing custom postprocessing when this object is copied.
<u>ooGetTypeN</u>	Gets the type number of this persistent object's class.
ooGetTypeName	Gets the name of this persistent object's class.
<u>oolsKindOf</u>	Tests whether this persistent object belongs to the class with the specified type number or to a class derived from that class.
<u>ooNewVersInit</u>	Default implementation for performing custom postprocessing when a new version of this object is created.
<u>oo0bj</u>	Default constructor that constructs a new basic object.
<u>ooPreMoveInit</u>	Default implementation for performing custom preprocessing when this object is moved.
<u>ooPostMoveInit</u>	Default implementation for performing custom postprocessing when this object is moved.
<u>ooThis</u>	Sets an object reference or handle to reference this persistent object.
<u>ooUpdate</u>	Opens this persistent object for update.
<u>ooValidate</u>	Default implementation for testing whether this object is valid.

<u>operator delete</u>	Deletes this persistent object.
operator new	Creates a new basic object using the specified clustering directive, if any.
prevVers	Finds, and optionally opens, the previous version of this object.
<u>set_defaultToGeneObj</u>	Links this object with the specified genealogy, making this object the default version in the genealogy.
<u>set_geneObj</u>	Links this object with the specified genealogy, making this object a version in the genealogy.
<u>set_prevVers</u>	Links this object to the specified object, making the specified object the previous version of this object.
<u>sub derivatives</u>	Deletes one or more derivatives associations from this object to the specified destination object, indicating that the specified object is no longer a derivative of this object.
<u>sub_derivedFrom</u>	Deletes one or more derivedFrom associations that link this object to the specified object, indicating that the specified object is no longer an ancestor version of this object.
sub_nextVers	Deletes the nextVers association that links this object to the specified object, so the specified object is no longer a next version of this object.

# Constructors

## ooObj

Default constructor that constructs a new basic object.

ooObj();

# **Operators**

#### operator delete

Deletes this persistent object.

```
void operator delete(void *objP);
```

Parameters	objP
	Pointer to the persistent object to be deleted. This pointer must be extracted from a handle to the object—for example, through <u>operator*</u> on that handle.
Discussion	This operator is an alternative to the <u>ooDelete</u> global function. Note that ooDelete is the recommended choice, however, because it doesn't require you to extract a pointer to the object to be deleted (you specify a handle or object reference instead).
	Deleting a basic object or container:
	<ul> <li>Calls the object's destructor, if any, before deallocating storage.</li> <li>Deletes all associations from the object to destination objects. If an association has <u>delete propagation</u> enabled, the destination objects are deleted as well.</li> </ul>
	Deleting a container:
	<ul> <li>Deletes all of the basic objects in it.</li> </ul>
	<ul> <li>Deletes all associations from each contained object to destination objects; if delete propagation is enabled for any of these associations, the destination objects are deleted as well.</li> </ul>
	Does not call the destructors of the contained objects for performance reasons. To ensure that destructors are called, you must iterate through the contained objects and explicitly delete them before deleting the container.
	Each deleted object is automatically removed from any bidirectional associations to maintain referential integrity. The delete operation must therefore be able to obtain an update lock on every object that is bidirectionally associated with a deleted object. Note that if another perisistent object references the deleted object through a unidirectional association or directly in one of its data members, you are responsible for removing that reference.
Example	Assume dbh is a valid handle to a database, and that Net is a persistence-capable class, which therefore inherits operator delete from ooObj.
	<pre>ooHandle(Net) netH = new(dbh) Net("vdd");// Create net object</pre>
	 Net *pVdd = netH; // Extract pointer to net object from handle delete pVdd; // Delete net object

#### operator new

Creates a new basic object using the specified clustering directive, if any.

;

1.	<pre>void *operator new(     size_t);</pre>
2.	<pre>void *operator new(     size_t,     const ooRefHandle(ooObj) &amp;near)</pre>
3.	<pre>void *operator new(    size_t,    const oo0bj *near);</pre>
(ODMG)4.	<pre>void *operator new(    size_t,    ooRefHandle(ooObj) near,    const char *type);</pre>
(ODMG) 5.	<pre>void *operator new(    size_t,    d_Database *near,    const char *type = 0);</pre>

#### Parameters size\_t

Do not specify; this parameter is automatically initialized by the compiler with the size of the class type in bytes.

#### near

Specifies whether the new basic object is to be persistent or transient. If the new basic object is persistent, *near* is the clustering directive that specifies the object's location in the federated database.

To create a transient basic object, specify 0 or a pointer to a transient object.

To create a persistent basic object, choose one of the following alternatives as the clustering directive:

- Omit *near* (variant 1).
- Specify an object reference, handle, or pointer to a database, container, or persistent basic object (variants 2 through 4). *near* may not reference a federated database or an autonomous partition.
- (ODMG) Specify a pointer to a d\_Database object that references the federated database (variant 5).

When you create a persistent basic object:

 If you omit *near*, the new basic object is created in the default container of the most recently opened or created database. An error is signalled if no such database exists.

	<ul> <li>If <i>near</i> references a database, the new basic object is created in the default container of that database.</li> </ul>
	<ul> <li>If <i>near</i> references a container, the new basic object is created in that container.</li> </ul>
	<ul> <li>If <i>near</i> references a basic object, the new basic object is put into the same container as the specified basic object. If possible, the new object will be put on the same page as the specified object or on a nearby page.</li> </ul>
	<ul> <li>(ODMG) If near specifies a valid d_Database object, the new basic object is created in the default container of the most recently opened or created database in the federation. If no such database exists, the object is created in the default container of a database called default_odmg_db; if necessary, this database is created.</li> </ul>
	type
	( <i>ODMG</i> ) This parameter is ignored. By convention it is the name of the class you are instantiating.
Returns	Memory pointer to the new basic object. This pointer is null if an error occurs during the creation of the object.
Discussion	When you create a new persistent basic object, you must:
	<ul> <li>Use operator new in an update transaction. The new basic object is made permanent on disk when the transaction commits or is checkpointed. If the transaction is aborted, the object is not created.</li> </ul>
	Assign the result of operator new directly to a handle. You can verify the creation of the basic object by checking whether the handle is null. (If an object reference is desired, you can then assign the handle to an object reference.)
Warning	Although direct assignment to a pointer or object reference does not raise compile-time or runtime errors, such assignments can eventually cause the Objectivity/DB cache to run out of memory.

# **Member Functions**

### add\_derivatives

Links this object with the specified object, making the specified object a derivative (successor version) of this object.

ooStatus add\_derivatives(const ooHandle(ooObj) &object);

Parameters	object Handle to a basic object.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	This operation is normally used when merging two version branches into a single branch.
	The operation links the specified object to this object by:
	<ul> <li>Creating a derivatives association link in this object.</li> <li>Creating the inverse derivedFrom association link in the specified object.</li> </ul>
	The application must be able to obtain update locks on both objects.
See also	derivatives sub_derivatives del_derivatives

## add\_derivedFrom

	Links this object with the specified object, making the specified object an ancestor version of this object.
	<pre>ooStatus add_derivedFrom(const ooHandle(ooObj) &amp;object);</pre>
Parameters	object Handle to a basic object.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	This operation is normally used when merging two version branches into a single branch.
	The operation links the specified object to this object by:
	<ul> <li>Creating a derivedFrom association link in this object.</li> </ul>
	<ul> <li>Creating an inverse derivatives association link in the specified object.</li> </ul>
	The application must be able to obtain update locks on both objects.
See also	derivedFrom sub_derivedFrom del_derivedFrom

#### add\_nextVers

Links this object with the specified object, making the specified object a next version of this object.

ooStatus add\_nextVers(const ooHandle(ooObj) &object);

 

 Parameters
 object Handle to a basic object.

 Returns
 oocSuccess if successful; otherwise oocError.

 Discussion
 The operation links the specified object to this object by:

 Creating a nextVers association link in this object.
 Creating an inverse prevVers association link in the specified object.
 The application must be able to obtain update locks on both objects.

 See also
 nextVers sub\_nextVers

 NextVers
 Subject

 Discussion
 NextVers

 See also
 NextVers

### defaultToGeneObj

del nextVers

Finds, and optionally opens, the genealogy associated with this object, if this object is the default version in a genealogy.

	<pre>1. ooHandle(ooGeneObj) defaultToGeneObj     (const ooMode openMode = oocNoOpen) const;</pre>
	<pre>2. ooRef(ooGeneObj) &amp;defaultToGeneObj(</pre>
	<pre>3. ooHandle(ooGeneObj) &amp;defaultToGeneObj(</pre>
Parameters	genealogy Object reference or handle to set to the found genealogy.
	openMode
	Intended level of access to the found genealogy:
	<ul> <li>Specify oocNoOpen (the default) to set the returned object reference or handle to the object without opening it.</li> </ul>
	<ul> <li>Specify oocRead to open the object for read.</li> </ul>
	Specify oocUpdate to open the object for update.

Returns	Object reference or handle to the found genealogy. A null object reference or handle is returned if no genealogy is found.
Discussion	When called without a <i>genealogy</i> parameter, this member function allocates a new genealogy handle and returns it. Otherwise, this member function returns the result in the object reference or handle that is passed to it.
	You can call this member function on an object to test whether it is the default version in a genealogy. The object is a default version if an associated genealogy is found; otherwise, if null is returned, the object being tested is not a default version.
See also	<u>set_defaultToGeneObj</u> <u>del_defaultToGeneObj</u>

#### del\_defaultToGeneObj

Deletes any defaultToGeneObj association from this object, indicating this object is no longer the default version.

ooStatus del\_defaultToGeneObj();

Returns oocSuccess if successful; otherwise oocError.

Discussion This operation also deletes the inverse defaultVers association link from the formerly associated genealogy. The application must be able to obtain update locks on both objects.

You typically call this function to remove a genealogy's previous default version before setting a new one. You shouldn't leave a genealogy without a default version.

You should call the <u>exist\_defaultToGeneObj</u> member function to test whether an association exists before you try to delete it.

See also <u>defaultToGeneObj</u> <u>set\_defaultToGeneObj</u>

#### del\_derivatives

Deletes all derivatives associations from this object, indicating this object has no derivatives (successor versions).

ooStatus del\_derivatives();

Returns oocSuccess if successful; otherwise oocError.

 Discussion
 This operation also deletes the inverse derivedFrom association link from each of the former derivatives. The application must be able to obtain update locks on all of the affected objects.

 You should call the <u>exist derivatives</u> member function to test whether any associations exist before you try to delete them.

 See also
 <u>derivatives</u> add derivatives sub derivatives

## del\_derivedFrom

	Deletes all derivedFrom associations from this object, indicating this object is not derived from any ancestor versions.			
	<pre>ooStatus del_derivedFrom();</pre>			
Returns	oocSuccess if successful; otherwise oocError.			
Discussion	This operation also deletes the inverse derivatives association link from each of the former ancestor versions. The application must be able to obtain update locks on all of the affected objects.			
	You should call the <u>exist_derivedFrom</u> member function to test whether any associations exist before you try to delete them.			
See also	<u>derivedFrom</u> <u>add_derivedFrom</u> <u>sub_derivedFrom</u>			

## del\_geneObj

Deletes any geneObj association that exists for this object, indicating this object is no longer in a genealogy.

ooStatus del\_geneObj();

Returns oocSuccess if successful; otherwise oocError.

Discussion This operation also deletes the inverse allVers association link from the formerly associated genealogy. The application must be able to obtain update locks on both objects.

You should call the  $\underline{exist geneObj}$  member function to test whether an association exists before you try to delete it.

See also	<u>geneObj</u>		
	set geneObj		

#### del\_nextVers

Deletes all nextVers associations that exist for this object, indicating that this object has no next versions.

ooStatus del\_nextVers();

Returns oocSuccess if successful; otherwise oocError.

Discussion This operation also deletes the inverse prevVers association link from each of the former next versions. The application must be able to obtain update locks on all of the affected objects.

You should call the <u>exist nextVers</u> member function to test whether any associations exist before you try to delete them.

See also <u>nextVers</u> <u>add\_nextVers</u> <u>sub\_nextVers</u>

#### del\_prevVers

Deletes any prevVers association that exists for this object, indicating this object has no previous version.

ooStatus del\_prevVers();

Returns oocSuccess if successful; otherwise oocError.

Discussion This operation also deletes the inverse nextVers association from the former previous version. The application must be able to obtain update locks on both objects.

You should call the <u>exist prevVers</u> member function to test whether an association exists before you try to delete it.

See also <u>prevVers</u> <u>set\_prevVers</u>

#### derivatives

Initializes an object iterator to find, and optionally open, all versions that are derived from this object and that satisfy any specified selection criteria.

1.	<pre>ooStatus derivatives( ooItr(ooObj) &amp;iterator, const ooMode openMode = oocNoOpen) const;</pre>
2.	<pre>oStatus derivatives( ooItr(ooObj) &amp;iterator,</pre>
	const char * <i>predicate</i> ) const;

3. ooStatus derivatives( ooItr(ooObj) &iterator, const ooMode openMode, const ooAccessMode access, const char \*predicate) const;

Parameters iterator

Object iterator for finding all derivative versions.

#### openMode

Intended level of access to each derivative object found by the iterator's next member function:

- oocNoOpen (the default in variant 1) causes next to set the iterator to the next derivative object without opening it.
- OOCRead causes next to open the next derivative object for read.
- oocUpdate causes next to open the next derivative object for update.

**Warning:** If versioning is enabled for one or more found objects, specifying oocUpdate means that next will create a new version of each such object.

#### predicate

String expression in predicate query language; specifies the condition to be met by the found objects. The iterator is initialized only with derivative objects that match *predicate*.

access

Level of access control of the data members that *predicate* can test:

- Specify oocPublic to permit the predicate to test only public data members, preserving encapsulation.
- Specify ocall to permit the predicate to test any data member. To preserve encapsulation, you should use this mode only within member functions of the class you are querying.

Returns oocSuccess if successful; otherwise oocError.

See also del\_derivatives sub\_derivatives add\_derivatives

#### derivedFrom

Initializes an object iterator to find, and optionally open, all ancestor versions from which this object is derived, and that satisfy any specified selection criteria.

```
    ooStatus derivedFrom(
ooItr(ooObj) &iterator,
const ooMode openMode = oocNoOpen) const;
    ooStatus derivedFrom(
ooItr(ooObj) &iterator,
const char *predicate) const;
    ooStatus derivedFrom(
ooItr(ooObj) &iterator,
const ooMode openMode,
const ooAccessMode access,
const char *predicate) const;
```

Parameters iterator

Object iterator for finding all ancestor versions.

#### openMode

Intended level of access to each object found by the iterator's next member function:

- oocNoOpen (the default in variant 1) causes next to set the iterator to the next found object without opening it.
- oocRead causes next to open the next found object for read.
- oocUpdate causes next to open the next found object for update.

**Warning:** If versioning is enabled for one or more found objects, specifying oocUpdate means that next will create a new version of each such object.

#### predicate

String expression in predicate query language; specifies the condition to be met by the found objects. The iterator is initialized only with found objects that match *predicate*.

access			
Level of access control of the data members that predicate can test:			
<ul> <li>Specify occPublic to permit the predicate to test only public data members, preserving encapsulation.</li> </ul>			
<ul> <li>Specify oocAll to permit the predicate to test any data member. To preserve encapsulation, you should use this mode only within member functions of the class you are querying.</li> </ul>			
oocSuccess if successful; otherwise oocError.			
add_derivedFrom sub_derivedFrom del_derivedFrom			

## exist\_defaultToGeneObj

Tests whether this object is the default version in any genealogy or, if a parameter is given, in the specified genealogy.

	<ol> <li>ooBoolean exist_defaultToGeneObj() const;</li> </ol>		
	<pre>2. ooBoolean exist_defaultToGeneObj(</pre>		
Parameters	genealogy Handle to the genealogy to be tested.		
Returns	(Variant 1) OCTTUE if this object is the default version in any genealogy, otherwise OCFAlse. (Variant 2) OCTTUE if this object is the default version in the specified genealogy, otherwise OCFAlse.		
Discussion	This member function tests whether this object has a defaultToGeneObj association link to a genealogy. If so, this object is by convention the default version in the corresponding genealogy.		

#### exist\_derivatives

Tests whether this object has any derivatives (successor versions); if a parameter is given, tests whether the specified object is a derivative of this object.

- ooBoolean exist\_derivatives() const;

Parameters	object
	Handle to the object to be tested as a derivative.
Returns	(Variant 1) occTrue if this object has any derivatives, otherwise occFalse. (Variant 2) occTrue if the specified object is a derivative of this object, otherwise
	oocFalse.

### exist\_derivedFrom

Tests whether this object is derived from any other objects (ancestor versions), or, if a parameter is given, from the specified object.

Parameters	object
	Handle to the object to be tested.
Returns	(Variant 1) oocTrue if this object is derived from any other object, otherwise oocFalse.
	oocFalse.

## exist\_geneObj

Tests whether this object is a version within any genealogy, or, if a parameter is given, within the specified genealogy.

	1. ooBoolean exist_geneObj() const;			
	<pre>2. ooBoolean exist_geneObj(</pre>			
Parameters	genealogy Handle to the genealogy to be tested.			
Returns	(Variant 1) oocTrue if this object belongs to any genealogy, otherwise oocFalse (Variant 2) oocTrue if this object belongs to the specified genealogy, otherwise oocFalse.			
Discussion	This member function tests whether this object has a geneObj association to a genealogy. If so, this object is by convention a version within the corresponding genealogy.			

#### exist\_nextVers

Tests whether this object has a next version; if a parameter is given, tests whether the specified object is the next version.

	<pre>1. ooBoolean exist_nextVers() const;</pre>				
	<pre>2. ooBoolean exist_nextVers(</pre>				
Parameters object Handle to the object to be tested as a next version.					
Returns	(Variant 1) oocTrue if this object has any next version, otherwise oocFalse				

# (Variant 2) oocTrue if the specified object is a next version of this object, otherwise oocFalse.

#### exist\_prevVers

Tests whether this object has a previous version; if a parameter is given, tests whether the specified object is the previous version.

1.	ooBoolean exist_prevVers() const;				
2.	ooBoolean exist_prevVers(				
	const ooHandle(ooObj) & <i>object</i> ) const;				

Parameters object

Handle to the object to be tested as a previous version.

Returns (Variant 1) oocTrue if this object has any previous version, otherwise oocFalse. (Variant 2) oocTrue if the specified object is a previous version of this object, otherwise oocFalse.

## geneObj

Finds, and optionally opens, the genealogy associated with this object.

1.	ooHandl	e(ooGen	eObj) gen	e0	bj(	
	const	ooMode	openMode	=	oocNoOpen)	const;

Parameters	genealogy	
	Object reference or handle to set to the found genealogy.	
	openMode	
	Intended level of access to the found genealogy:	
	<ul> <li>Specify occNoOpen (the default) to set the returned object reference or handle to the object without opening it.</li> </ul>	
	<ul> <li>Specify oocRead to open the object for read.</li> </ul>	
	<ul> <li>Specify oocUpdate to open the object for update.</li> </ul>	
Returns	Object reference or handle to the found genealogy. A null object reference or handle is returned if no genealogy is found.	
Discussion	When called without a <i>genealogy</i> parameter, this member function allocates a new genealogy-object handle and returns it. Otherwise, this member function returns the result in the object reference or handle that is passed to it.	
See also	<u>set_geneObj</u> del geneObj	

#### mark\_modified

(ODMG) Opens this persistent object for update.

```
void mark_modified();
```

Discussion This member function is an ODMG-equivalent name for the member function ooUpdate.

#### nextVers

Initializes an object iterator to find, and optionally open, all next versions that are created from this object and that satisfy any specified selection criteria.

 ooStatus nextVers( ooItr(ooObj) &*iterator*, const ooMode *openMode* = oocNoOpen) const;
 ooStatus nextVers( ooItr(ooObj) &*iterator*, const char \**predicate*) const; 

#### Parameters iterator

Object iterator for finding the next version(s) created from this object:

- Under linear versioning, this iterator finds at most one version.
- Under branch versioning, this iterator may find multiple versions.

#### openMode

Intended level of access to the versions found by the iterator's  ${\tt next}$  member function:

- oocNoOpen (the default in variant 1) causes next to set the iterator to the next version without opening it.
- oocRead causes next to open the next version for read.
- oocUpdate causes next to open the next version for update.

**Warning:** If versioning is enabled for one or more found objects, specifying oocUpdate means that next will create a new version of each such object.

#### predicate

String expression in predicate query language; specifies the condition to be met by the found objects. The iterator is initialized only with versions that match *predicate*.

access

Level of access control of the data members that *predicate* can test:

- Specify oocPublic to permit the predicate to test only public data members, preserving encapsulation.
- Specify occAll to permit the predicate to test any data member. To
  preserve encapsulation, you should use this mode only within member
  functions of the class you are querying.
- Returns oocSuccess if successful; otherwise oocError.

See also <u>add\_nextVers</u> <u>sub\_nextVers</u> <u>del\_nextVers</u>

#### ooCopyInit

Default implementation for performing custom postprocessing when this object is copied.

virtual ooStatus ooCopyInit();

Returns oocSuccess if postprocessing is successful; otherwise oocError.

Discussion This member function is invoked automatically after a basic object is copied—for example, by the copy member function on a handle to the object. The default implementation simply returns oocSuccess.

An application can override ooCopyInit in a basic-object class to perform custom postprocessing after an instance of the class is copied—typically:

- Performing any necessary operations on attribute data members for which bit-wise copying is inadequate.
- Propagating the copy operation to associated or referenced objects (creating a deep copy).

### ooGetTypeN

Gets the type number of this persistent object's class.

	virtual ooTypeNumber ooGetTypeN() const;	
Returns	Type number of this object's class.	
Discussion	Every persistence-capable class has a schema-defined type number that uniquely identifies the class to the federated database. You can use this number to determine whether an object is an instance of a particular class.	
	You normally use $ooGetTypeN$ in a member function of an application-defined persistence-capable class; when such a member function is called on an object of the class, $ooGetTypeN$ provides the member function with the type number of the object's class.	
	If you already have a handle or object reference to a persistent object, you can equivalently call the $\underline{typeN}$ member function on the handle or object reference.	
See also	<u>ooGetTypeName</u> <u>ooIsKindOf</u> <u>ooTypeN</u> global macro <i>ooRefHandle</i> (ooObj)::: <u>typeN</u> member function	

#### ooGetTypeName

	Gets the name of this persistent object's class.
	<pre>virtual char *ooGetTypeName() const;</pre>
Returns	String containing the name of this object's class.
Discussion	You normally use <code>ooGetTypeName</code> in a member function of an application-defined persistence-capable class; when such a member function is called on an object of the class, <code>ooGetTypeName</code> provides the member function with the name of the object's class.
	If you already have a handle or object reference to a persistent object, you can equivalently call the $\underline{typeName}$ member function on the handle or object reference.
WARNING	Do not modify the returned string in any manner. Doing so may result in unexpected program errors. This string is used internally by Objectivity/DB.
See also	<u>ooGetTypeN</u> <u>ooIsKindOf</u> <u>ooTypeN</u> global macro <i>ooRefHandle</i> (ooObj):: <u>typeName</u> member function
oolsKindO	f
	Tests whether this persistent object belongs to the class with the specified type number or to a class derived from that class.
	<pre>virtual ooBoolean ooIsKindOf(     const ooTypeNumber typeN) const;</pre>

Parameters $t_{YP}eN$ Type number of a persistence-capable class. If you know the name of the<br/>desired class, you can call the ootypeN global macro to obtain its type<br/>number. If you have a handle to an object of the desired class, you can call<br/>the typeN member function on the handle to obtain the class's type number.Returnsooctrue if this object is an instance of the class with type number  $t_{YP}eN$  or to a<br/>class derived from that class; otherwise oocFalse.DiscussionYou normally call oolsKindOf to perform runtime type identification—for<br/>example, to determine whether it is safe to downcast a handle that references this<br/>object.

#### See also <u>ooGetTypeN</u> <u>ooGetTypeName</u> <u>ooTypeN</u> global macro *ooRefHandle*(ooObj)::typeN member function *ooRefHandle*(ooObj)::typeName member function

#### ooNewVersInit

Default implementation for performing custom postprocessing when a new version of this object is created.

virtual ooStatus ooNewVersInit();

- Returns oocSuccess if postprocessing is successful; otherwise oocError.
- Discussion This member function is invoked automatically when a new version of a basic object is created. A new version is created when you open an object for update after versioning has been enabled for it. The default implementation simply returns oocSuccess.

A new version is a bit-wise copy of the existing object. An application can override <code>ooNewVersInit</code> in a basic-object class to perform postprocessing after an instance of the class is versioned (copied)—typically:

- Performing any necessary operations on attribute data members for which bit-wise copying is inadequate.
- Propagating the versioning operation to associated or referenced objects (creating a deep copy).

#### ooPostMoveInit

Default implementation for performing custom postprocessing when this object is moved.

virtual ooStatus ooPostMoveInit();

Returns oocSuccess if postprocessing is successful; otherwise oocError.

Discussion This member function is invoked automatically after a basic object is moved—for example, by the move member function on a handle to the object. The default implementation simply returns oocSuccess.

An application can override <code>ooPostMoveInit</code> in a basic-object class to perform postprocessing after an instance of the class is moved. Such postprocessing typically reestablishes references to an object after it acquires a new object identifier—for example, by updating any referencing attributes, unidirectional associations, persistent collections, or scope names. (Objectivity/DB automatically maintains referential integrity for bidirectional associations.)

#### ooPreMoveInit

Default implementation for performing custom preprocessing when this object is moved.

virtual ooStatus ooPreMoveInit();

Returns oocSuccess if preprocessing is successful; otherwise oocError.

Discussion This member function is invoked automatically before a basic object is moved—for example, by the move member function on a handle to the object. The default implementation simply returns oocSuccess.

> An application can override <code>ooPreMoveInit</code> in a basic-object class to perform preprocessing before an instance of the class is moved. Such preprocessing typically removes existing references to an object before its object identifier becomes invalid—for example, to maintain referential integrity for any referencing attributes, unidirectional associations, persistent collections, or scope names. (Objectivity/DB automatically maintains referential integrity for bidirectional associations.)

#### ooThis

Sets an object reference or handle to reference this persistent object.

- 1. ooHandle(ooObj) ooThis() const;

Parameters	object Object reference or handle to be set to this object.
Returns	Object reference or handle to this object.
Discussion	You normally use ooThis in a member function of an application-defined persistence-capable class. When such a member function is called on an object of the class, ooThis provides the member function with an object reference or handle to the object. The member function can then perform operations on the object that are available only through an object reference or handle.

When called without an *object* parameter, ooThis allocates a new handle and returns it. Otherwise, ooThis returns the object reference or handle that is passed to it.

Calling ooThis in a member function of a persistence-capable class serves the same purpose as using the C++ this keyword in a member function of a non-persistence-capable class—both ooThis and this enable you to access the object on which the member function is called. However, ooThis and this differ in the following ways:

- The this keyword is the pointer to the object. Syntactically, the this keyword is a name, and can be used in expressions such as this->get();
- The ooThis member function returns an object reference or handle to the object. Syntactically, ooThis is a function and can be used in expressions such as ooThis()->get();

Member functions of persistence-capable classes should use ooThis (and not this) because the only safe way to operate on a persistent object is through an object reference or handle.

WARNING	ooThis is the only safe way to obtain an object reference or handle within a
	member function; do not attempt to initialize an object reference or handle by
	assigning the this pointer to it.

See also appClass::ooThis

#### ooUpdate

Opens this persistent object for update.ooStatus ooUpdate();ReturnsoocSuccess if successful; otherwise oocError.DiscussionYou normally use ooUpdate in a member function of an application-defined<br/>persistence-capable class; doing so informs Objectivity/DB that the member<br/>function intends to modify one or more attributes of the object on which it is<br/>called. For example, each accessor member function that sets a data-member<br/>value should call ooUpdate, to ensure that Objectivity/DB will save the<br/>modification persistently.If you already have a handle or object reference to the persistent object to be<br/>modified, you can equivalently call the update member function on the handle<br/>or object reference.

### ooValidate

Default implementation for testing whether this object is valid.

virtual ooBoolean ooValidate();

Returns oocTrue if this object is valid; otherwise, oocFalse.

Discussion The ooValidate member function provides a framework for application-specific validation of persistent objects. The default implementation simply returns oocTrue.

An application can override this function in a persistence-capable class to perform whatever checks are necessary to test whether an object of the class is valid.

#### prevVers

Finds, and optionally opens, the previous version of this object.

	<pre>1. ooHandle(ooObj) prevVers(     const ooMode openmode = oocNoOpen) const;</pre>		
	<pre>2. ooHandle(ooObj) &amp;prevVers(</pre>		
	<pre>3. ooRef(ooObj) &amp;prevVers(</pre>		
Parameters	object		
	Object reference or handle to be set to the previous version.		
	openMode		
	Intended level of access to the previous version:		
	<ul> <li>Specify occNoOpen (the default) to set the returned object reference or handle to the object without opening it.</li> </ul>		
	<ul> <li>Specify oocRead to open the object for read.</li> </ul>		
	<ul> <li>Specify oocUpdate to open the object for update.</li> </ul>		
Returns	Object reference or handle to the previous version. A null object reference or handle is returned if no previous version exists.		
Discussion	When called without an <i>object</i> parameter, this member function allocates a new handle and returns it. Otherwise, this member function returns the result in the object reference or handle that is passed to it.		

See also	<u>set</u>	_prevVers
	<u>del</u>	prevVers

## set\_defaultToGeneObj

Links this object with the specified genealogy, making this object the default version in the genealogy.

	<pre>ooStatus set_defaultToGeneObj(     const ooHandle(ooGeneObj) &amp;genealogy);</pre>
Parameters	genealogy Handle to a genealogy.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	<ul> <li>The operation links the specified genealogy to this object by:</li> <li>Creating a defaultToGeneObj association link in this object.</li> <li>Creating the inverse defaultVers association link in the specified genealogy.</li> <li>The application must be able to obtain update locks on both objects.</li> <li>Because defaultToGeneObj is a to-one association, an error is signaled if this object already has a link for the association. That is, you must remove any existing default version from the genealogy before you set a new default version.</li> </ul>
See also	<u>defaultToGeneObj</u> <u>del defaultToGeneObj</u>
set_geneOl	bj
	Links this object with the specified genealogy, making this object a version in the genealogy.
	<pre>ooStatus set_geneObj(const ooHandle(ooGeneObj) &amp;genealogy);</pre>
Parameters	genealogy

Handle to a genealogy.

Returns oocSuccess if successful; otherwise oocError.

Discussion The operation links the specified genealogy to this object by:

- Creating a geneObj association link in this object.
- Creating the inverse allVers association link in the specified genealogy.

The application must be able to obtain update locks on both objects.

Because geneObj is a to-one association, an error is signaled if this object already has a link for the association. That is, you must remove this object from its current genealogy before you make it a version of the specified genealogy.

See also <u>geneObj</u> <u>del geneObj</u>

#### set\_prevVers

	Links this object to the specified object, making the specified object the previous version of this object.
	<pre>ooStatus set_prevVers(const ooHandle(ooObj) &amp;object);</pre>
Parameters	object Handle to a basic object.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	<ul> <li>The operation links the specified object to this object by:</li> <li>Creating a prevVers association link in this object.</li> <li>Creating the inverse nextVers association link in the specified object.</li> <li>The application must be able to obtain update locks on both objects.</li> <li>Because prevVers is a to-one association, an error is signaled if this object already has a link for the association. That is, you must remove the current previous version from this object before you set a new previous version.</li> </ul>
See also	<u>prevVers</u> <u>del prevVers</u>

#### sub\_derivatives

Deletes one or more derivatives associations from this object to the specified destination object, indicating that the specified object is no longer a derivative of this object.

```
ooStatus sub_derivatives(
    const ooHandle(ooObj) &object)
    const uint32 number = 1);
```

Parameters object

Handle to the unwanted derivative version.

	number
	Number of derivatives association links to delete between this object and the specified object:
	<ul> <li>If you specify 0, all such links are deleted.</li> </ul>
	<ul> <li>If you specify 1 (the default), the first or only such link is deleted.</li> </ul>
	<ul> <li>If you specify a number greater than 1, this member function deletes the first number links encountered.</li> </ul>
Returns	oocSuccess if successful; otherwise oocError. oocSuccess is returned even if no link exists to be deleted or if <i>number</i> exceeds the number of existing links. oocError results from internal errors or locking errors.
Discussion	Because derivatives is a many-to-many bidirectional association, it is possible for multiple links to exist between this object and the specified destination object. The <i>number</i> parameter allows you specify how many such links to delete.
See also	derivatives
	del_derivatives
	add_derivatives

#### sub\_derivedFrom

Deletes one or more derivedFrom associations that link this object to the specified object, indicating that the specified object is no longer an ancestor version of this object.

```
ooStatus sub_derivedFrom(
    const ooHandle(ooObj) &object)
    const uint32 number = 1);
```

Parameters

Handle to the unwanted ancestor version.

number

object

Number of derivedFrom association links to delete between this object and the specified object.

- If you specify 0, all such links are deleted.
- If you specify 1 (the default), the first or only such link is deleted.
- If you specify a number greater than 1, this member function deletes the first *number* links encountered.

Returns oocSuccess if successful; otherwise oocError.oocSuccess is returned even if no link exists to be deleted or if *number* exceeds the number of existing links. oocError results from internal errors or locking errors.

Discussion Because derivedFrom is a many-to-many bidirectional association, it is possible for multiple links to exist between this object and the specified ancestor version. The *number* parameter allows you specify how many such links to delete.

See also derivedFrom del\_derivedFrom add\_derivedFrom

#### sub\_nextVers

Deletes the nextVers association that links this object to the specified object, so the specified object is no longer a next version of this object.

```
ooStatus sub_nextVers(
    const ooHandle(ooObj) &object);
```

 Parameters
 object

 Handle to the unwanted next version.

 Returns
 oocSuccess if successful; otherwise oocError.

See also <u>add\_nextVers</u> <u>nextVers</u> <u>del\_nextVers</u>

# ooOperatorSet Class

#### Inheritance: ooOperatorSet

The non-persistence-capable class ooOperatorSet represents an *operator set*—that is, a set of application-defined relational operators.

See:

- "Reference Summary" on page 466 for an overview of member functions
- "Reference Index" on page 466 for a list of member functions

## **About Operator Sets**

An operator set is a collection of application-defined relational operators, where each relational operator consists of both of the following:

- An application-defined *operator function* that conforms to the calling interface defined by the <u>ooQueryOperatorPtr</u> function pointer type.
- A *token* that represents the operator function in predicates.

When an operator set is assigned to the <u>ooUserDefinedOperators</u> global variable, the predicate query mechanism consults it to interpret any nonstandard tokens in a predicate.

An application normally uses the default operator set that is created and assigned to <code>ooUserDefinedOperators</code> when the application starts. An application can also create additional operator sets (for example, to define alternative sets of relational operators). However, only one operator set can be in effect at a time; if you create an operator set, you must assign it to the <code>ooUserDefinedOperators</code> global variable to make its operators available to Objectivity/DB.

An operator set is empty at creation, and you add relational operators to it by calling the <u>registerOperator</u> member function on the operator set. At any

point in the application, you can call the <u>clear</u> member function to remove all the currently registered operators from the set.

The ooUserDefinedOperators global variable is not thread-safe, so your application must ensure that:

- Only one thread updates the operator set at a time.
- If a thread is updating the operator set, no other thread can be making a predicate query at the same time.

# **Reference Summary**

Creating	<u>ooOperatorSet</u>
Adding and Removing Operator Functions	<u>registerOperator</u> <u>clear</u>

# **Reference Index**

<u>clear</u>	Clears all application-defined relational operators from this operator set.
<u>ooOperatorSet</u>	Default constructor that constructs a new empty operator set.
<u>registerOperator</u>	Registers the specified operator function and its token with this operator set.

# **Constructors and Destructors**

#### ooOperatorSet

Default constructor that constructs a new empty operator set.

```
ooOperatorSet();
```

# **Member Functions**

## registerOperator

Registers the specified operator function and its token with this operator set.

	<pre>ooStatus registerOperator( const char *name, ooQueryOperatorPtr funcPtr);</pre>
Parameters	name
	Token that is to represent the operator function in a predicate. This token may not begin or end with the following symbols or symbol combinations: ) ( $\&\&$    ! , . If the token is the same as an existing operator, the new operator will override the standard behavior.
	funcPtr
	Pointer to the operator function to be registered.
Returns	oocSuccess if successful; otherwise oocError.
See also	<u>ooQueryOperatorPtr</u>
clear	

Clears all application-defined relational operators from this operator set.

clear();
## ooQuery Class

Inheritance: ooQuery

The non-persistence-capable class <code>ooQuery</code> represents a *query object*, which evaluates whether a persistent object matches a given predicate string.

See:

• "Reference Index" on page 469 for a list of member functions

## **About Query Objects**

You use a query object to filter arbitrary groups of objects that cannot be scanned directly with a predicate string. For example, you can use a query object to perform a predicate query over the group of objects referenced by a VArray of object references.

### **Reference Index**

<u>evaluate</u>	Evaluates whether the referenced object matches this query object's predicate string.
<u>setup</u>	Sets up a query object for evaluating the specified type of objects against the specified predicate string.

## **Member Functions**

### evaluate

	Evaluates whether the referenced object matches this query object's predicate string.
	<pre>ooBoolean evaluate(ooHandle(ooObj) &amp;objH);</pre>
Parameters	оbјн Handle to the object to be evaluated.
Returns	oocTrue if the object is a match for the predicate string specified by <code>setup</code> ; otherwise <code>oocFalse</code> .
setup	
	Sets up a query object for evaluating the specified type of objects against the specified predicate string.
	<pre>ooStatus setup(     char *predicate,     ooTypeNumber typeN,     ooOperatorSet *operatorSet = 0);</pre>
Parameters	predicate
	Predicate string for the query object.
	typeN
	Type number of the class whose instances are to be evaluated.
	operatorSet
	The operator set to be consulted if <i>predicate</i> uses tokens for any application-defined relational operators. You must specify the operator set explicitly, even if it is the default operator set (that is, the current value of the <u>ooUserDefinedOperators</u> variable).
	Specifying 0 causes all tokens in <i>predicate</i> to be interpreted as Objectivity/DB standard relational operators.
Returns	oocSuccess if successful; otherwise oocError.

# ooRefHandle(appClass) Classes

Inheritance:	ooRef(ooObj)->ooRef( <i>appClass</i> )
	ooRef(ooObj)->ooRef(ooContObj)->ooRef( <i>appClass</i> )
Inheritance:	ooHandle(ooObj)->ooHandle( <i>appClass</i> )
	ooHandle(ooObj)->ooHandle(ooContObj)->ooHandle( <i>appClass</i> )
	The abbreviation <i>ooRefHandle(appClass)</i> refers to two non-persistence-capable classes:
	• ooRef( <i>appClass</i> ), which represents an <i>object reference</i> to a persistent object of the application-defined class <i>appClass</i> .
	• ooHandle( <i>appClass</i> ), which represents a <i>handle</i> to a persistent object of the application-defined class <i>appClass</i> .
	These two classes are documented together because they define almost identical sets of member functions (exceptions are listed in the "Reference Summary"). These classes provide the primary interface for operating on persistent objects of the application-defined persistence-capable class $appClass$ and its derived classes. Note that $appClass$ can be either a basic-object class (derived from $ooObj$ ) or a container class (derived from $ooContObj$ ).
	See:
	■ "Reference Summary" on page 478 for an overview of member functions
	<ul> <li>"Reference Index" on page 479 for a list of member functions</li> </ul>
	To use the <i>ooRefHandle(appClass)</i> classes, you must include and compile with files generated by the DDL processor, as described in "Obtaining Generated Class Definitions" on page 474.
	( <i>ODMG</i> ) You can use the ODMG standard class name d_Ref< <i>appClass</i> > interchangeably with ooRef( <i>appClass</i> ).

## About appClass Handles and References

When an application defines a persistence-capable class *appClass* and adds it to the federated-database schema, the DDL processor generates the corresponding handle and object-reference classes *ooRefHandle(appClass)*. The application works with each persistent object of class *appClass* indirectly through one or more instances of *ooRefHandle(appClass)* that are set to reference the object. A handle or object reference to a persistent object serves as a type-safe smart pointer that:

- Identifies the persistent object to the application or to another object.
- Provides an interface for operating on the persistent object.
- Manages the memory pointer to the persistent object.
- Provides an indirect member-access operator (->) for accessing the persistent object's public member functions.

It is sometimes more appropriate to use a handle rather than an object reference, and vice versa; the choice is described in "Structure and Behavior" on page 473. A simple guideline is to use handles in function definitions and object references as data member types in persistence-capable class definitions.

#### Interface

The *ooRefHandle(appClass)* classes provide the primary interface for operating on a referenced instance of *appClass*.

If *appClass* is a basic-object class, the entire *ooRefHandle(appClass)* interface consists of member functions defined by the *ooRefHandle(ooObj)* base classes. These member functions are either inherited by the *ooRefHandle(appClass)* classes or redefined wherever type-specific parameters or behavior are required.

If *appClass* is a container class, the *ooRefHandle(appClass)* interface provides the same interface as *ooRefHandle(ooContObj)* classes, with a few redefined member functions for type-specific behavior or parameters. This interface includes:

- Public member functions defined by the ooRefHandle(ooContObj) classes for specialized container operations, such as finding the basic objects in a container.
- Public member functions defined by the *ooRefHandle(ooObj)* classes for general Objectivity/DB operations, such as opening, locking, printing object identifiers, and so on.

Note: The member functions defined by the *ooRefHandle*(ooObj) classes for moving, copying, and versioning basic objects are disallowed for containers because the *ooRefHandle*(ooContObj) classes redefine them as private.

If other classes are derived from *appClass*, the *ooRefHandle(appClass)* classes themselves serve as base classes for the corresponding handle and object-reference classes. All such derived handle and object-reference classes provide the same interface as *ooRefHandle(appClass)*, with the usual redefined member functions for type-specific behavior or parameters.

#### Structure and Behavior

Although both handles and object references provide a way to reference and operate on a persistent object, they are optimized for different purposes:

- Handles are optimized for accessing persistent objects in memory—that is, for performing multiple operations on a referenced object or repeatedly accessing the object's members.
- Object references are optimized for linking persistent objects—that is, for storing object identifiers persistently in reference attributes, in associations, or as elements of a collection.

### Handles

Handles are optimized for efficient in-memory access because they can automatically obtain and manipulate pointers to referenced persistent objects. Thus, when a handle is set to reference a particular persistent object, the handle stores the object identifier for that object. The first time the persistent object is accessed through the handle, the handle is automatically *opened*—that is, the handle obtains a pointer to the object's representation in memory. This memory pointer enables the handle to access the referenced object quickly during subsequent operations performed through the handle. When the handle is *closed*, it invalidates the pointer but keeps the object identifier, so the application can reuse the handle (without resetting it) to access the same object.

Besides maintaining a pointer to the referenced persistent object, an open handle also *pins* the object's memory representation in the Objectivity/DB cache. Pinning guarantees that the persistent object is readily available in memory for as long as it is needed. Closing the handle removes its particular "pin"; when the last open handle to that persistent object is closed, the last pin is removed and the object itself is closed. Closing the last (or only) object on a buffer page permits Objectivity/DB to swap the page out of the cache as needed to make room for other open objects.

A handle to a persistent object, like any handle, has cache-related state that is associated with the Objectivity context in which it was created. Therefore, a handle:

 Cannot be stored persistently—for example, as an attribute value. In fact, the DDL processor does not accept ooHandle(appClass) as a data member type in a persistence-capable class definition. • Cannot be passed between Objectivity contexts.

### **Object References**

Object references are optimized for implementing persistent links because they are essentially wrappers for object identifiers. Thus, setting an object reference to a particular persistent object causes the object reference to store the object's identifier. The object reference never acquires a pointer to the persistent object in memory; instead, whenever the persistent object is accessed through the object reference, the operation is delegated to a temporary handle that provides the necessary pointer.

Because it has no bulky cache-specific state, an object reference to a persistent object:

- Can be stored persistently—for example, as an attribute value. The DDL processor accepts ooRef(appClass) as a data member type in a persistence-capable class definition.
- Can be passed between Objectivity contexts.

For convenience, an application can use an object reference (instead of a handle) to perform an operation on a referenced persistent object or to access one of the object's members. However, poor performance results when an object reference is used for multiple such operations on the same persistent object, because *each* operation causes a temporary handle to be created, used, and discarded. Performance may also be affected by swapping, because the object reference does not pin the persistent object's memory representation in the Objectivity/DB cache.

An object reference of type <code>ooRef(appClass)</code> is sometimes called a *standard object reference* because it contains the complete object identifier for the referenced object. An alternative for referencing a basic object under certain circumstances is the *short object reference* of type <code>ooShortRef(appClass)</code>, which saves space by storing object identifiers in a truncated format.

#### **Obtaining Generated Class Definitions**

To use the *ooRefHandle(appClass)* classes, you must include either the primary header file or the references header file generated by the DDL processor for *appClass*. Thus, if *appClass* is defined in the DDL file *classDefFile.ddl*, you must include one of the following files:

- The primary header file classDefFile.h
- The references header file classDefFile\_ref.h

Furthermore, you must compile the method implementation file *classDefFile\_ddl.cxx* with your application code files.

For more information about DDL-generated files and how to use them, see the Objectivity/C++ Data Definition Language book.

#### When appClass is a Template Class

When *appClass* is a persistence-capable template class with multiple parameters, the names of the generated handle and object-reference classes contain the symbol OO\_COMMA to separate the template parameters. For example, for a persistence-capable template class <code>Example<Float</code>, <code>Node></code>, the generated object-reference class is <code>ooRef(Example<Float OO\_COMMA Node>)</code>. This is because the macro syntax of the generated class name interprets embedded commas as separators between the as macro parameters instead of as separators between the template parameters.

## Working With appClass Handles

**NOTE** For simplicity, this section describes how to work with handles. Except where noted, the same information applies to object references.

An application normally creates a handle as a local variable on the stack, rather than allocating it on the heap. A handle should not be declared as const, because its internal state is changed by any operation that accesses a persistent object through it. (Object references may be declared as const.) An application should not explicitly define subclasses of the *ooRefHandle(appClass)* classes; any necessary subclasses are generated automatically by the DDL processor if the application defines any subclasses of *appClass*.

A new handle is normally *null*—that is, it contains the value 0 instead of an object identifier for a persistent object. The application can then set the handle to reference a particular persistent object in any of the following ways:

- By creating a new persistent object with <u>operator new</u> of the *appClass* class and assigning the result to the handle.
- By finding an existing persistent object with the handle's <u>lookupObj</u> member function. (If *appClass* is a container class, you can use the handle's <u>exist</u> or <u>open</u> member function.)
- By passing the handle to a member function that sets it, such as the <u>linkName</u> member function of a persistent object that has a <u>linkName</u> association. The <u>linkName</u> member function finds the associated destination object and sets the specified handle to the found object.
- By assignment or initialization from another handle or object reference.

An object reference may be set in any of these ways, with the following exception—the result of operator new may not be assigned to an object reference.

A handle continues to reference the same persistent object until it is set to another persistent object or to null. Furthermore, multiple handles and object references can be set to the same persistent object. A handle of class <code>ooHandle(appClass)</code> can be set to an instance of <code>appClass</code> or any class derived from <code>appClass</code>.

An application operates on a persistent object by calling:

 Member functions of a handle that references the object. As indicated in "Reference Summary" on page 478, such member functions allow you to copy the referenced object, open it for update, and so on.

To call a member function of a handle, you use the direct member-access operator (.). For example, appH.lookupObj calls the <u>lookupObj</u> member function of the handle appH.

 Member functions of the referenced object itself. As indicated in "Reference Summary" on page 84, these include generated member functions and member functions inherited from ooObj or ooContObj, as well as member functions defined by *appClass*.

To call a member function of a referenced object, you use the handle's overloaded indirect member-access operator (<u>operator-></u>). For example, appH->oolsKindOf calls the <u>oolsKindOf</u> member function on the persistent object that is referenced by the handle appH.

Although most of a handle's member functions operate on the referenced persistent object, some functions operate on the handle itself. For example, you use:

- The assignment operator <u>operator=</u> to set a handle from another handle or from an object reference.
- The inherited comparison operators <u>operator==</u> and <u>operator!=</u> to test whether a handle references the same persistent object as another handle or object reference.
- The inherited member function <u>isNull</u> to test whether a handle is null. (Alternatively, you can use the overloaded <u>operator==</u> to compare a handle to 0.)

A handle preserves its reference to a persistent object across transaction boundaries, provided that the handle does not go out of scope and is not set to null as the result of an abort operation. Before reusing the handle in a new transaction, however, the application should call <u>isValid</u> to test whether the handle is still *valid*—that is, whether it still references an existing persistent object. A handle becomes invalid if it is set to null *or* if the referenced persistent object has been deleted by another process between transactions.

### Specifying an appClass Handle to a Function

Objectivity/C++ functions that require a persistent object as input typically obtain the object through a parameter of type const ooHandle(ooObj) &. You can pass a handle of type ooHandle(appClass) through such a parameter because ooHandle(appClass) is derived from ooHandle(ooObj).

Similarly, Objectivity/C++ functions that require containers typically obtain them through parameters of type of const <code>ooHandle(ooContObj) &</code>. If <code>appClass</code> is a container class, you can pass an <code>appClass</code> handle through such a parameter.

### Opening and Closing an appClass Handle

**NOTE** This subsection applies only to handles, not to object references, which are in effect always closed.

A handle is automatically opened when a persistent object is opened through it. The open persistent object is both locked and represented in memory; the open handle manages a pointer to the persistent object, pinning the object in memory until the handle is closed. A closed handle, which has an object identifier instead of a pointer, can reference either an open or a closed persistent object.

The most common way to open a persistent object through a handle is to do so implicitly by using the handle's indirect member-access operator (<u>operator-></u>) to access a member of the referenced object. Alternatively, a referenced persistent object can be opened by explicit request—for example, by calling the handle's <u>open or update</u> member function. Another way to explicitly open a persistent object is by finding it with a function whose *openMode* parameter is either <code>oocRead</code> or <code>oocUpdate</code>. (Most functions that set a handle to a found persistent object provide an *openMode* parameter for specifying the desired level of access through that handle.) In all cases, if the found or referenced persistent object is already open (for example, because another operation opened it earlier in the transaction), the accessing handle gets a pointer to the existing memory representation and adds a pin.

You obtain a closed handle to a persistent object by finding the object with a function whose *openMode* parameter is set to oocNoOpen. Such operations simply provide the handle with a persistent object's object identifier without adding a pin, even if the object is already open through another handle.

Objectivity/DB automatically closes an open handle when the handle is destroyed (for example, by going out of scope), when it is set to reference another persistent object, or when the transaction that opened it commits or aborts. An application can close a handle explicitly by calling the handle's <u>close</u> member

function. Closing the last open handle to a particular persistent object unpins and closes the object.

## **Reference Summary**

The following table summarizes just the member functions that are redefined by *ooRefHandle(appClass)* to provide type-specific parameters. For descriptions of inherited member functions:

- See the <u>ooRefHandle(ooObj)</u> classes (page 593) if *appClass* is a basic object class.
- See the <u>ooRefHandle(ooContObj)</u> classes (page 509) if *appClass* is a container class.

The summarized member functions are defined on both the object-reference class and the handle class. Two operators are defined on only the handle class, namely, <u>operator\*</u> and <u>operator appClass\*</u>.

Creating a Handle or Object Reference	<u>ooHandle(appClass)</u> ooRef(appClass)
Setting a Handle or Object Reference	<u>operator=</u> lookupObj
Accessing the Referenced Persistent Object	<u>operator-&gt;</u> <u>operator*</u> <u>operator appClass*</u> <u>ptr(ODMG</u> )
Copying a Basic Object	copy
Testing a Handle	operator appClass*
Finding a Persistent Object	<u>lookupObj</u>
ODMG Interface	<u>operator d_Ref_Any</u> ptr

## **Reference Index**

сору	Creates a copy of the referenced basic object, clustering the new copy near the specified object.
<u>lookupObj</u>	Finds the persistent object with the specified scope name (or the basic object matching the specified key structure) within the specified scope, and sets this object reference or handle to reference the found object.
<u>ooHandle(appClass)</u>	Default constructor that constructs a null handle.
<u>ooHandle(appClass)</u>	Constructs a handle that references the same persistent object as the specified object reference, handle, pointer, or ODMG generic reference.
<u>ooRef(appClass)</u>	Default constructor that constructs a null object reference.
<u>ooRef(appClass)</u>	Constructs an object reference that references the same persistent object as the specified object reference, handle, pointer, or ODMG generic reference.
<u>operator-&gt;</u>	Indirect member-access operator; accesses a member of the referenced persistent object.
<u>operator*</u>	Handle class only. Dereference operator; returns the persistent object referenced by this handle.
<u>operator=</u>	Assignment operator; sets this object reference or handle to reference the same persistent object as the specified object reference, handle, or pointer.
<u>operator d Ref Any</u>	(ODMG) Conversion operator that returns an ODMG generic reference to the referenced persistent object.
<u>operator appClass*</u>	Handle class only. Conversion operator that returns an appClass pointer to the referenced persistent object.
ptr	(ODMG) Returns a C++ pointer to the referenced persistent object.

## **Constructors and Destructors**

## ooHandle(appClass)

Default constructor that constructs a null handle.

```
ooHandle(appClass)();
```

## ooHandle(appClass)

Constructs a handle that references the same persistent object as the specified object reference, handle, pointer, or ODMG generic reference.

	<pre>1. ooHandle(appClass)(     const ooRefHandle(appClass) &amp;objectRH);</pre>
	<pre>2. ooHandle(appClass)(const appClass *objectP);</pre>
(ODMG)	<pre>3. ooHandle(appClass)(const d_Ref_Any &amp; from);</pre>
Parameters	objectRH
	Object reference or handle to an instance of <i>appClass</i> .
	objectP
	Pointer to an instance of <i>appClass</i> . A pointer <i>must</i> be the result of operator new on <i>appClass</i> or a derived class.
	from
	(ODMG) An ODMG generic reference to an instance of <i>appClass</i> . An error is signalled if <i>from</i> references an object that is not an instance of <i>appClass</i> or a class derived from <i>appClass</i> .
Discussion	Variants 1 and 3 allow a new handle to be constructed from an existing object reference, handle, or ODMG generic reference. If the new handle is constructed from an existing open handle, the new handle is open; in all other cases, the new handle is closed.
	Variant 2, which constructs a handle from the specified pointer, has a narrower purpose—to obtain an open handle to a newly created persistent object so you can perform persistence operations on it, and so the object can eventually be unpinned when it is no longer needed in memory.

## ooRef(appClass)

Default constructor that constructs a null object reference.

```
ooRef(appClass)();
```

### ooRef(appClass)

Constructs an object reference that references the same persistent object as the specified object reference, handle, pointer, or ODMG generic reference.

1.	ooRef( <i>appClass</i> )(	
	<pre>const ooRefHandle(appClass) &amp;objectRH);</pre>	
2.	ooRef( <i>appClass</i> )(const <i>appClass</i> * <i>objectP</i> );	
(ODMG) 3.	ooRef( <i>appClass</i> )(const d_Ref_Any & <i>from</i> );	

Parameters objectRH

Object reference or handle to an instance of appClass.

#### objectP

Pointer to an instance of *appClass*. The pointer may *not* be the result of operator new. Instead, the pointer must be the result of using either <u>operator appClass\*</u> on a handle *or* <u>ptr</u> on an object reference or handle earlier in the same transaction. If the specified pointer was originally extracted from a handle, that handle must still exist and reference the same persistent object.

from

(ODMG) An ODMG generic reference to an instance of *appClass*. An error is signalled if *from* references an object that is not an instance of *appClass* or a class derived from *appClass*.

Discussion Variants 1 and 3 allow a new object reference to be constructed from an existing object reference, handle, or ODMG generic reference.

Variant 2 has a narrower purpose, which is to allow you to resume persistence operations on an existing persistent object after manipulating it through a pointer. The use of this variant should be rare, however, because pointers are not normally used to manipulate persistent objects.

## **Operators**

#### operator->

Indirect member-access operator; accesses a member of the referenced persistent object.

appClass \*operator->();

Returns Pointer to the referenced instance of *appClass*.

Discussion The accessed persistent object is opened for read, if it is not already open.

If the referenced object is an instance of *appClass*, operator-> can access any of the object's public members. If the referenced object is an instance of a derived class, operator-> accesses only the members that are defined in *appClass*.

You use operator-> in an expression handle->member, where handle is an instance of ooHandle(appClass) and member is the name of a public member defined on class appClass. As for any overloading of the C++ member-access operator (->), the expression handle->member is interpreted as (handle.operator->())->member. That is, the overloaded operator-> returns a pointer to the referenced object, and then the ordinary C++ operator-> selects the specified member of that object, returning the value of that member.

Although you can obtain the returned pointer (for example, from the this keyword in an accessed member function), you should view the pointer as an intermediate by-product of the access operation. The same is true for pointers to accessed data members obtained through expressions such as &(handle->dataMember). You should not use such pointers in any further operations (for example, do not assign the returned pointer to a handle, object reference, or pointer variable).

**Warning:** The pointer returned by operator-> is guaranteed valid for only a limited time:

- When you access an object through an object reference, the returned pointer is guaranteed valid only for the duration of the -> operation.
- When you access an object through a handle, the returned pointer is guaranteed valid only as long as the handle exists and references the same object.

#### operator\*

*Handle class only.* Dereference operator; returns the persistent object referenced by this handle.

appClass &operator\*();

Returns C++ reference to the persistent object referenced by this handle.

Discussion This operator enables you to pass a handle to a function that accepts a persistent object by reference. This operator is analogous to the C++ operator\* for dereferencing a pointer.

The persistent object referenced by this handle is opened for read, if it is not already open.

WARNING	The returned reference is guaranteed valid only as long as the handle exists, remains open, and references the same object.	
Example	Assume that Book is a persistence-capable class. You use <code>operator*</code> to pass a book handle to <code>helperFunction</code> , which accepts a C++ reference to a Book.	
	<pre>void helperFunction(Book &amp;aBook);</pre>	
	<pre>void processBook(ooHandle(Book) &amp;bookH) { helperFunction(*bookH);</pre>	

#### operator=

Assignment operator; sets this object reference or handle to reference the same persistent object as the specified object reference, handle, or pointer.

1.	ooRefHa	ndle(appClass)	&operat	or=(
	const	ooRefHandle(ap	pClass)	&objectRH);

- 2. ooRefHandle(appClass) &operator=(
   const ooShortRef(appClass) &shortObjR);

#### Parameters objectRH

Object reference or handle to an instance of appClass.

#### shortObjR

Short object reference to an instance of *appClass* that resides in the same container as the object that is referenced by this object reference or handle. (If this object reference or handle is null, you can set it to a container with the <u>set\_container</u> member function.)

A short object reference specifies just the lower half of an object identifier (corresponding to the object's logical page and slot numbers). The upper half of the object identifier (corresponding to the database and container) is taken from this object reference or handle.

#### objectP

0, or a nonnull pointer to an instance of *appClass*.

■ If you are assigning to a handle, the specified pointer *must* be the result of operator new on *appClass* or a derived class.

- If you are assigning to an object reference, the specified pointer may not be the result of operator new. Instead, the pointer must be the result of using either <u>operator appClass\*</u> on a handle or <u>ptr</u> on an object reference or handle earlier in the same transaction. If the specified pointer was originally extracted from a handle, that handle must still exist and reference the same persistent object.
- Returns This object reference or handle, after it has been set to reference the specified object.
- Discussion Variants 1 and 2 allow you to use the specified object reference or handle to produce another object reference or handle to the same persistent object. If you are assigning to a handle from an open handle, the returned handle is open; in all other cases, the returned handle is closed.

Variant 3 allows you to set this object reference or handle to null. Otherwise, assignment-from-pointer has two specific purposes, depending on whether you are assigning to a handle or to an object reference:

- Pointer-to-handle assignment enables you to obtain an open handle to a newly created persistent object, so you can perform persistence operations on it, and so the object can eventually be unpinned when it is no longer needed in memory.
- Pointer-to-object-reference assignment enables you to resume persistence operations on a persistent object after manipulating it through a pointer. This usage of variant 3 is rare, because pointers are not normally used to manipulate persistent objects.

### operator d\_Ref\_Any

(*ODMG*) Conversion operator that returns an ODMG generic reference to the referenced persistent object.

```
operator d_Ref_Any() const;
```

### operator appClass\*

*Handle class only*. Conversion operator that returns an *appClass* pointer to the referenced persistent object.

operator appClass\*();

Returns Pointer to the persistent object referenced by this handle. Returns a null pointer if the handle is a null handle.

Discussion This conversion operator enables you to:

- Pass a handle to a function that accepts a pointer to a persistent object.
- Assign a handle to an *appClass*\* variable (for example, to pass to the overloaded <u>operator delete</u>).
- Use a handle as the conditional expression in an if or while statement to test whether the handle is null.

The persistent object referenced by this handle is opened for read, if it is not already open.

**WARNING** The returned pointer is guaranteed valid only as long as this handle exists, remains open, and references the same object.

An application generally relies on handles to provide memory management for persistent objects, and avoids the explicit use of pointers to such objects. On occasion, explicit use of pointers is required for performance reasons or for compatibility with functions that are indifferent to persistence (for example, in legacy code or a third-party library). However, you should be careful when manipulating a pointer from a handle because the validity of the pointer depends on the state of the handle. You should not use the returned pointer in other persistence operations (for example, do not pass it to any Objectivity/C++ member function other than the overloaded operator delete).

Example The handle objectH is used as a conditional expression which evaluates to 0 if the handle is null.

ooHandle(myClass) objectH; ... // Set objectH to some object if (objectH) { ... // Do something interesting if initialization was successful }

```
See also <u>ptr</u>
ooRef(ooObj)::<u>operator int</u>
```

## **Member Functions**

#### сору

Creates a copy of the referenced basic object, clustering the new copy near the specified object.

	<pre>1. ooHandle(appClass) copy(</pre>	
	<pre>2. ooRef(appClass) ©(</pre>	
	<pre>3. ooHandle(appClass) ©( const ooHandle(ooObj) &amp;near, ooHandle(appClass) &amp;newCopy) const;</pre>	
Parameters	<i>near</i> Handle to the object with which to cluster the new copy. <i>near</i> may be a handle to a database, a container, or a basic object:	
	<ul> <li>If <i>near</i> is a database handle, the new copy is stored in the default container of that database.</li> </ul>	
	If near is a container handle the new copy is stored in that container	
	<ul> <li>If <i>near</i> is a container handle, the new copy is stored in that container.</li> <li>If <i>near</i> is a basic object handle, the new copy is stored in the same container as the referenced basic object. If possible, the copy will be put on the same page as the referenced basic object or on a nearby page.</li> </ul>	
	newCopy	
	Object reference or handle to set to the new copy.	
Returns	oocSuccess if successful; otherwise oocError.	
Discussion	Copying applies only to basic objects. Therefore, you should call this member function only if <i>appClass</i> is derived from ooObj but not ooContObj. This member function signals an error if you attempt to copy a container.	
	When called without a <i>newcopy</i> parameter, copy allocates a new handle and returns it. Otherwise, copy returns the object reference or handle that is passed to it.	
	The application must be able to lock the container of the original object for read and the container of the new copy for update.	

### lookupObj

Finds the persistent object with the specified scope name (or the basic object matching the specified key structure) within the specified scope, and sets this object reference or handle to reference the found object.

1.	ooStatus lookupObj (
	const ooHandle(ooObj) & <i>scope</i> ,
	const char * <i>scopeName</i> ,
	<pre>const ooMode openMode = oocRead);</pre>
2.	ooStatus lookupObj (
2.	<pre>ooStatus lookupObj (     const ooHandle(ooObj) &amp;scope,</pre>
2.	ooStatus lookupObj ( const ooHandle(ooObj) & <i>scope,</i> const ooKey & <i>keyStruct</i> ,

Parameters scope

Handle to an object that defines the scope of the lookup:

- (Variant 1) When the lookup is by scope name, *scope* specifies the scope object that defines the name scope to search. *scope* can reference the federated database, a database, a persistent container, a persistent basic object, or an autonomous partition.
- (Variant 1) When the lookup is by key structure, *scope* specifies the container to search. *scope* can reference the container itself or the database whose default container is to be searched.

#### scopeName

Scope name to look up in the scope specified by scope.

#### openMode

Intended level of access to the found object:

- Specify oocRead (the default) to open the object for read.
- Specify oocUpdate to open the object for update.
- (Variant 2 only) Specify oocNoOpen to set this object reference or handle to the object without opening it. oocNoOpen is valid only for scope-name lookup, because scope-named objects can be found without being opened, whereas keyed objects must be opened during the search.

#### keyStruct

Key structure specifying the key field and key field value to match.

- Returns oocSuccess if an object of the appropriate class is found; otherwise oocError.
- Discussion Scope-name lookup applies either to basic objects or containers. Therefore, you can call variant 1 regardless of whether appClass is derived from ooObj or ooContObj.

Keyed-object lookup applies only to basic objects. Therefore, you should call variant 2 only if *appClass* is derived from ooObj (but not ooContObj).

In any case, an object is found only if it is an instance of *appClass* or a class derived from *appClass*.

The application must be able to obtain a read lock on the hashed container used by the scope object.

#### ptr

(*ODMG*) Returns a C++ pointer to the referenced persistent object.

```
appClass *ptr();
```

Returns Pointer to the referenced basic object or container.

Discussion You use this member function to obtain a pointer to an application-defined basic object or container—for example, to pass to a function that accepts a pointer instead of a handle or object reference.

If ptr is called on an object reference, the referenced persistent object is opened for update. If ptr is called on a handle, the referenced persistent object is opened for read.

Warning: The returned pointer is guaranteed valid for only a limited time:

- If ptr is called on an object reference, the returned pointer is valid and the persistent object is pinned in memory until the end of the transaction.
- If ptr is called on a handle, the returned pointer is valid only as long as the handle exists, remains open, and references the same persistent object (equivalent to <u>operator appClass\*</u>).

An application generally relies on handles to provide memory management for persistent objects, and avoids the explicit use of pointers to such objects. On occasion, explicit use of pointers is required for performance reasons or for compatibility with functions that are indifferent to persistence (for example, in legacy code or a third-party library). However, you should be careful when manipulating persistent objects through pointers:

- Pointers extracted from handles become invalid if the handles change or go out of scope.
- Pointers extracted from object references can cause the Objectivity/DB cache to run out of memory if too many objects are pinned until the end of the transaction.

You should not use the returned pointer in other persistence operations (for example, do not pass it to any Objectivity/C++ member function other than the overloaded operator delete).

## ooRefHandle(ooAPObj) Classes

Inheritance: ooRef(ooObj)->ooRef(ooAPObj)

Inheritance: ooHandle(ooObj)->ooHandle(ooAPObj)

The abbreviation *ooRefHandle*(ooAPObj) refers to two non-persistence-capable classes:

- ooRef(ooAPObj), which represents an object reference to an autonomous partition.
- ooHandle(ooAPObj), which represents a handle to an autonomous partition.

The two classes ooRef(ooAPObj) and ooHandle(ooAPObj) are documented together because they define the same set of member functions. These member functions provide the primary interface for operating on autonomous partitions (instances of ooAPObj).

You can create and work with autonomous partitions only if you have bought and installed Objectivity/DB Fault Tolerant Option (Objectivity/FTO).

See:

- "Reference Summary" on page 493 for an overview of member functions
- "Reference Index" on page 494 for a list of member functions

## **About Autonomous-Partition Handles and References**

An application works with an autonomous partition indirectly through one or more handles or object references—that is, through instances of *ooRefHandle*(*ooAPObj*) that are set to reference the desired partition. A handle or object reference to an autonomous partition both identifies the partition and provides the complete public interface for operating on it. Handles and object references to autonomous partitions do not support indirect member access; the *ooRefHandle*(ooAPObj) classes provide no indirect member-access operator (->), and the ooAPObj class defines no public member functions other than a constructor and operator new.

You can work with an autonomous partition through either a handle or an object reference—the choice is arbitrary, except as described in "Structure and Behavior" on page 490. Most applications use handles rather than object references.

#### Interface

The *ooRefHandle*(*ooAPObj*) classes provide the primary interface for operating on a referenced autonomous partition. Part of this interface consists of member functions defined by *ooRefHandle*(*ooAPObj*) for specialized operations such as finding the containers that a partition controls. The other part of this interface consists of member functions defined by the *ooRefHandle*(*ooObj*) base classes for more general Objectivity/DB operations, such as opening, printing object identifiers, and so on. These member functions are either inherited by the *ooRefHandle*(*ooAPObj*) classes or redefined wherever type-specific parameters or behavior are required.

Some of the member functions defined by the *ooRefHandle*(ooObj) base classes are *not* available to instances of the *ooRefHandle*(ooAPObj) classes. These include member functions for moving, copying, versioning, scope-naming, and member-access operations, which apply only to basic objects or persistent objects. The disallowed member functions and operators are redefined as private members of the *ooRefHandle*(ooAPObj) classes.

#### **Structure and Behavior**

Handles and object references to autonomous partitions are essentially wrappers for autonomous-partition identifiers. For example, when you set a handle to reference a particular partition, the handle stores the object identifier of that partition. If the partition is then opened through the handle, Objectivity/DB uses the identifier to locate the partition's system-database file on disk. Subsequent member-function calls on the handle are applied to the instance of <code>ooAPObj</code> that represents the identified partition in memory.

In general, object references are optimized for implementing links among related persistent objects, while handles are optimized for memory management and member-access. When an autonomous partition is referenced, however, these optimizations are largely irrelevant, because partitions (unlike persistent objects):

- Cannot be linked (for example, through associations).
- Are not subject to memory management (they have no attributes for persistent data and are therefore not manipulated through pointers).

Have no accessible members.

One significant exception is that a handle to an autonomous partition, like any handle, contains cache-related state that is associated with the Objectivity context in which it was created. Therefore, only object references (but not handles) can be passed between Objectivity contexts. Otherwise, a handle to a partition and an object reference to a partition are functionally equivalent.

## **Working With Autonomous-Partition Handles**

**NOTE** For simplicity, this section describes how to work with handles. Except where noted, the same information applies to object references.

An application normally creates a handle as a local variable on the stack, rather than allocating it on the heap. A handle should not be declared as const, because its internal state may be changed when an autonomous partition is accessed through it. (Object references may be declared as const). Applications should not create subclasses of the <code>ooRefHandle(ooAPObj)</code> classes.

A new handle is normally *null*—that is, it contains the value 0 instead of the object identifier for an autonomous partition. The application can then set the handle to reference a particular partition in any of the following ways:

- By creating a new partition with <u>operator new</u> of the ooAPObj class and assigning the result to the handle.
- By finding an existing partition with the handle's <u>exist</u> or <u>open</u> member function.
- By passing the partition to a member function that sets it, such as the <u>controlledBy</u> member function of a container handle, which finds the partition that controls the container.
- By assignment or initialization from another handle or object reference.

A handle continues to reference the same autonomous partition until it is set to another partition or to null. Furthermore, multiple handles and object references can be set to the same partition.

An application operates on an autonomous partition by calling member functions on a handle that references it. To call a member function of a handle, you use the direct member-access operator (.). For example, apH.name calls the <u>change</u> member function of the handle apH.

As indicated in "Reference Summary" on page 493, an application can get and change a referenced partition's attributes, find its containers and database

images, and so on. For more information about operating on autonomous partitions, see Chapter 27, "Autonomous Partitions," in the Objectivity/C++ programmer's guide.

Although most of a partition handle's member functions operate on the referenced partition, some functions operate on the handle itself. For example, you use:

- The assignment operator <u>operator=</u> to set a handle from another handle or from an object reference.
- The inherited comparison operators <u>operator==</u> and <u>operator!=</u> to test whether a handle references the same partition as another handle or object reference.
- The inherited member function <u>isNull</u> to test whether a handle is null. (Alternatively, you can use the overloaded <u>operator==</u> to compare a handle to 0.)

A handle preserves its reference to an autonomous partition across transaction boundaries, provided that the handle does not go out of scope and is not set to null as the result of an abort operation. Before reusing the handle in a new transaction, however, the application should call <u>isValid</u> to test whether the handle is still *valid*—that is, whether it still references an existing autonomous partition. A handle becomes invalid if it is set to null *or* if the referenced partition has been deleted by another process between transactions.

Objectivity/C++ functions that require an autonomous partition as input normally obtain the partition through a parameter of type const ooHandle(ooAPObj) &. If a function manipulates other types of Objectivity/DB objects in addition to autonomous partitions, the parameter type may be specified as const ooHandle(ooObj) &. You can pass a partition handle to a parameter of this type, because ooHandle(ooAPObj) is derived from ooHandle(ooObj). In practice, however, relatively few functions accept a partition handle where a more general handle is requested; these include ooDelete and functions that require a storage object for a name scope. In most cases, a function that requests a general-purpose handle operates only on basic objects or persistent objects, and signals an error if you specify a partition handle.

Any operation that affects an autonomous partition implicitly opens it if it is not already open; an application does not need to open an autonomous partition explicitly unless it is used as the entry point into the data or unless access must be guaranteed in advance. A handle is automatically closed when it is destroyed (for example, when it goes out of scope). However, closing one or more handles to a particular partition has no effect on that partition, which remains open until the transaction commits or aborts.

## **Reference Summary**

The following table summarizes all the member functions that are available to instances of ooRefHandle(ooAPObj). Member functions indicated as *(inherited)* documented with the <u>ooRefHandle(ooObj)</u> classes (page 593).

Creating a Handle or Object Reference	<u>ooHandle(ooAPObj)</u> ooRef(ooAPObj)
Setting the Handle or Object Reference	<u>operator=</u> <u>open</u> <u>exist</u>
Comparing Handles and Object References	<pre>operator== (inherited) operator!= (inherited)</pre>
Opening and Closing the Autonomous Partition	<u>open</u> <u>update</u> <u>openMode</u> <u>close</u>
Modifying the Autonomous Partition	<u>update</u> <u>change</u> <u>markOffline</u> <u>markOnline</u> <u>returnAll</u>
Getting Information About the Autonomous Partition	<pre>name bootFileHost bootFilePath jnlDirHost jnlDirPath lockServerHost sysDBFileHost sysDBFilePath typeN typeName print (inherited) sprint (inherited)</pre>
Testing the Autonomous Partition	<u>exist</u> <u>isAvailable</u> <u>isOffline</u>

Testing the Handle or Object Reference	<u>is_null (inherited)</u> <u>isNull (inherited)</u> <u>isValid</u> <u>operator_int (inherited)</u> <u>operator_ooObj*</u> (inherited)
Finding Objects	<u>exist</u> <u>containedIn</u> <u>containersControlledBy</u> <u>imagesContainedIn</u> ( <i>DRO</i> )

## **Reference Index**

<u>bootFileHost</u>	Gets the network name of the data server host containing the boot file for the referenced autonomous partition.
bootFilePath	Gets the fully qualified pathname of the boot file for the referenced autonomous partition.
change	Changes catalog information for the referenced autonomous partition.
close	Internal use only. Objectivity/DB closes an autonomous partition automatically when the transaction that opened it commits or aborts.
<u>containedIn</u>	Finds the federated database that contains the referenced autonomous partition.
<u>containersControlledBy</u>	Initializes an object iterator to find the containers that are controlled by the referenced autonomous partition.
<u>exist</u>	Tests whether the specified autonomous partition exists in the federated database; if successful, sets this object reference or handle to reference the partition.
imagesContainedIn	(DRO) Initializes an object iterator to find all the database images that are contained in the referenced autonomous partition.
<u>isAvailable</u>	Tests whether the current process can access the referenced autonomous partition.

<u>isOffline</u>	Tests whether the referenced autonomous partition is offline (inaccessible to tools and applications).
<u>isValid</u>	Tests whether this object reference or handle is valid—that is, whether it references an existing autonomous partition.
jnlDirHost	Gets the network name of the host that contains the journal directory for the referenced autonomous partition.
jnlDirPath	Gets the fully qualified pathname of the journal directory for the referenced autonomous partition.
lockServerHost	Gets the network name of the host running the lock server for the referenced autonomous partition.
markOffline	Makes the referenced autonomous partition inaccessible to tools and applications.
<u>markOnline</u>	Makes the referenced autonomous partition accessible to tools and applications.
name	Gets the system name of the referenced autonomous partition.
<u>ooHandle(ooAPObj)</u>	Default constructor that constructs a null handle.
<u>ooHandle(ooAPObj)</u>	Constructs a handle that references the same autonomous partition as the specified object reference or handle.
<u>ooRef(ooAPObj)</u>	Default constructor that constructs a null object reference.
<u>ooRef(ooAPObj)</u>	Constructs an object reference that references the same autonomous partition as the specified object reference or handle.
open	Explicitly opens the referenced or specified autonomous partition, preparing the partition for the specified level of access.
openMode	Gets the current level of access to the referenced autonomous partition.
<u>operator=</u>	Assignment operator; sets this object reference or handle to reference the specified autonomous partition.

<u>returnAll</u>	Clears the referenced autonomous partition, returning all controlled containers to their home autonomous partitions.
<u>sysDBFileHost</u>	Gets the network name of the host containing the system-database file of the referenced autonomous partition.
<u>sysDBFilePath</u>	Gets the fully qualified pathname of the system-database file of the referenced autonomous partition.
typeN	Gets the type number of the autonomous-partition class ooAPObj.
typeName	Gets the name of the autonomous-partition class
update	Opens the referenced autonomous partition for update access.

## Constructors

### ooHandle(ooAPObj)

Default constructor that constructs a null handle.

```
ooHandle(ooAPObj)();
```

### ooHandle(ooAPObj)

Constructs a handle that references the same autonomous partition as the specified object reference or handle.

```
ooHandle(ooAPObj)(
     const ooRefHandle(ooAPObj) &existing);
```

Parameters existing

Object reference or handle to an existing autonomous partition.

### ooRef(ooAPObj)

Default constructor that constructs a null object reference.

```
ooRef(ooAPObj)();
```

### ooRef(ooAPObj)

Constructs an object reference that references the same autonomous partition as the specified object reference or handle.

Parameters existing

Object reference or handle to an existing autonomous partition.

## Operators

#### operator=

Assignment operator; sets this object reference or handle to reference the specified autonomous partition.

```
ooRefHandle(ooAPObj) &operator=(
    const ooRefHandle(ooAPObj) &existing);
```

Parameters	existing
------------	----------

Object reference or handle to an existing autonomous partition.

Returns This object reference or handle.

## **Member Functions**

### bootFileHost

Gets the network name of the data server host containing the boot file for the referenced autonomous partition.

char \*bootFileHost() const;

- Returns Pointer to a string containing the network name of the data server host. If an error occurs, a null pointer is returned and the error is signalled.
- Discussion The returned string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.

### bootFilePath

Gets the fully qualified pathname of the boot file for the referenced autonomous partition.

char \*bootFilePath() const;

- Returns Pointer to a string containing the pathname of the boot file. If an error occurs, a null pointer is returned and the error is signalled.
- Discussion The returned string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.

### change

Changes catalog information for the referenced autonomous partition.

```
ooStatus change(
    const char *lockServer = 0,
    const char *apFileHost = 0,
    const char *apFilePath = 0,
    const char *bootFileHost = 0,
    const char *bootFilePath = 0,
    const char *jnlDirHost = 0,
    const char *jnlDirPath = 0) const;
```

#### Parameters lockServer

Name of the new lock server host (the host that runs the lock server for this autonomous partition). Specify 0 (the default) to leave the lock server host unchanged.

#### apFileHost

Name of the data server host on which the autonomous partition's system-database file is to reside. Specify 0 (the default) to leave the autonomous-partition host unchanged. If you specify this parameter, you must also specify a nondefault value for *apFilePath*.

```
apFilePath
```

New pathname for the autonomous partition's system-database file on *apFileHost*, including the filename. Specify 0 (the default) to leave the autonomous-partition pathname unchanged.

```
bootFileHost
```

Name of the new data server host on which the autonomous-partition boot file is to reside. Specify 0 (the default) to leave the boot file host unchanged.

If you specify this parameter, you must also specify a nondefault value for *bootFilePath*.

#### bootFilePath

New pathname for the autonomous-partition boot file on *bootFilePath*, including the filename. Specify 0 (the default) to leave the boot file pathname unchanged.

#### jnlDirHost

New host machine for the autonomous-partition journal directory. Specify 0 (the default) to leave the journal-directory host unchanged. If you specify this parameter, you must also specify a nondefault value for *jnlDirPath*.

#### jnlDirPath

New path for the new autonomous-partition journal directory. Specify 0 (the default) to leave the journal-directory pathname unchanged.

Returns oocSuccess if successful; otherwise oocError.

Discussion You cannot change the system name of an autonomous partition. If you specify a new boot file location, the updated boot file is written to the new location, but the old boot file remains; you must delete this file using appropriate operating system commands. If you specify a new location for the autonomous partition's system-database file, the catalog is updated; however, you must use appropriate operating system commands to actually move the file.

You use this member function in a special-purpose application that consists of a single update transaction. The application must exit immediately after the transaction commits. This is because the new state of the autonomous partition is inconsistent with information cached by the executing application.

#### close

Internal use only. Objectivity/DB closes an autonomous partition automatically when the transaction that opened it commits or aborts.

```
ooStatus close() const;
```

#### containedIn

Finds the federated database that contains the referenced autonomous partition.

- 1. ooHandle(ooFDObj) containedIn() const;

	<pre>3. ooHandle(ooFDObj) &amp;containedIn(</pre>
Parameters	returnedFD Object reference or handle to be set to the federated database.
Returns	Object reference or handle to the federated database.
Discussion	When called without a parameter, <code>containedIn</code> allocates a new federated-database handle and returns it. Otherwise, <code>containedIn</code> returns the object reference or handle that was passed to it.
containersControlledBy	
	Initializes an object iterator to find the containers that are controlled by the referenced autonomous partition.

ooStatus containersControlledBy(
 ooItr(ooContObj) &controlledConts,
 const ooMode openMode = oocNoOpen) const;

#### Parameters controlledConts

Object iterator for finding the controlled containers.

#### openMode

Intended level of access to the containers found by the iterator's next member function:

- oocNoOpen (the default) causes next to set the iterator to the next container without opening it.
- oocRead causes next to open the next container for read.
- oocUpdate causes next to open the next container for update.

#### Returns oocSuccess if successful; otherwise oocError.

# Discussion The iterator finds only those containers that were explicitly transferred to the control of the referenced autonomous partition.

#### exist

Tests whether the specified autonomous partition exists in the federated database; if successful, sets this object reference or handle to reference the partition.

```
ooBoolean exist(
    const ooHandle(ooFDObj) &containingFD,
    const char *apSysName,
    const ooMode openMode = oocNoOpen);
```

#### Parameters containingFD

Handle to the federated database.

#### apSysName

System name of the autonomous partition to be found.

#### openMode

Intended level of access to the autonomous partition, if it exists:

- Specify oocNoOpen (the default), to set this object reference or handle to the autonomous partition without opening it.
- Specify oocRead to open the autonomous partition for read.
- Specify occUpdate to open the autonomous partition for update.
- Returns occTrue if the specified autonomous partition exists, or occFalse if the database does not exist or if it is not accessible.
- Discussion If the specified autonomous partition exists, this object reference or handle is set to reference it; otherwise, this object reference or handle is set to null.

If you specifically want to test for existence, you use the *openMode* parameter's default value (*oocNoOpen*). Otherwise, a return value of *oocFalse* could mean either that the partition doesn't exist, or that it does exist, but cannot be opened.

### imagesContainedIn

(*DRO*) Initializes an object iterator to find all the database images that are contained in the referenced autonomous partition.

Parameters containedDBs

Object iterator for finding the database images in the autonomous partition.

	openMode
	Intended level of access to the database images found by the iterator's $\mathtt{next}$ member function:
	<ul> <li>oocNoOpen (the default) causes next to set the iterator to the next database image without opening it.</li> </ul>
	<ul> <li>oocRead causes next to open the next database image for read.</li> </ul>
	<ul> <li>oocUpdate causes next to open the next database image for update.</li> </ul>
Returns	oocSuccess if successful; otherwise oocError.
isAvailable	
	Tests whether the current process can access the referenced autonomous partition.
	ooBoolean isAvailable() const;
Returns	oocTrue if the partition is available to the current process; otherwise oocFalse. This information can change at any time during the process.
isOffline	
	Tests whether the referenced autonomous partition is offline (inaccessible to tools and applications).
	<pre>ooBoolean isOffline() const;</pre>
Returns	oocTrue if the referenced autonomous partition is offline.
See also	markOffline
	markOnline
	<u>ooGetOfflineMode</u> global function
	ooSetOfflineMode global function
isValid	
	Tests whether this object reference or handle is valid—that is, whether it references an existing autonomous partition.
	ooBoolean isValid() const;
Returns	oocTrue if the object reference or handle references an existing autonomous partition; <code>oocFalse</code> if this object reference or handle is null or has a stale identifier.

- Discussion You can use isValid to determine whether it is safe to use an object reference or handle that was set in a previous transaction. Such an object reference or handle still retains its reference to a partition; however, between transactions, that reference may have become invalid (for example, because another process has deleted the partition).
  - **NOTE** isValid checks only for the existence of a partition with a particular identifier, but has no way of knowing whether it is the same partition. It is possible, although very unlikely, for another process to have deleted the original partition and created a new one with the same identifier.

If your purpose is simply to test whether an object reference or handle has been initialized, it is more efficient to use <u>isNull</u>, which performs its test entirely in memory without having to access files on disk.

### jnlDirHost

Gets the network name of the host that contains the journal directory for the referenced autonomous partition.

char \*jnlDirHost() const;

- Returns Pointer to a string containing the journal-directory host name. If an error occurs, a null pointer is returned and the error is signalled.
- Discussion The string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.

### **jnlDirPath**

Gets the fully qualified pathname of the journal directory for the referenced autonomous partition.

char \*jnlDirPath() const;

- Returns Pointer to a string containing the journal-directory pathname.
- Discussion The string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.

### **lockServerHost**

Gets the network name of the host running the lock server for the referenced autonomous partition.

char \*lockServerHost() const;

- Returns Pointer to a string containing the lock server host name. If an error occurs, a null pointer is returned.
- Discussion The string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.

### markOffline

Makes the referenced autonomous partition inaccessible to tools and applications.

ooStatus markOffline() const;

Returns oocSuccess if successful; otherwise oocError.

- Discussion This member function must be invoked in a transaction that includes no other operations.
- See also <u>isOffline</u> <u>ooOfflineMode</u> global type <u>ooGetOfflineMode</u> global function <u>ooSetOfflineMode</u> global function

### markOnline

Makes the referenced autonomous partition accessible to tools and applications.

ooStatus markOnline() const;

Returns oocSuccess if successful; otherwise oocError.

See also <u>isOffline</u> <u>ooOfflineMode</u> global type <u>ooGetOfflineMode</u> global function <u>ooSetOfflineMode</u> global function
name	
	Gets the system name of the referenced autonomous partition.
	char *name() const;
Returns	Pointer to a string containing the system name.
Discussion	The string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.
open	
	Explicitly opens the referenced or specified autonomous partition, preparing the partition for the specified level of access.
	<pre>1. ooStatus open(     const ooMode openMode = oocRead);</pre>
	<pre>2. ooStatus open( const ooHandle(ooFDObj) &amp;containingFD, const char *apSysName, const ooMode openMode = oocRead);</pre>
Parameters	openMode
	Intended level of access to the opened autonomous partition:
	<ul> <li>Specify occRead (the default) to open and implicitly lock the partition for read.</li> </ul>
	<ul> <li>Specify occupdate to open and implicitly lock the partition for update (read and write).</li> </ul>
	containingFD
	Handle to the currently open federated database.
	apSysName
	System name of the autonomous partition to open.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	Variant 1 opens the autonomous partition referenced by this object reference or handle.
	Variant 2 finds and opens the autonomous partition with the specified system name, and sets this object reference or handle to reference it. An error is signalled if no partition exists with the specified system name or if the partition's system-database file cannot be found or accessed. This variant is especially useful

when you want to use the partition as an entry point into your data. For example, you might find and open a partition so you can iterate over the containers it controls.

Opening an autonomous partition makes it available to an application by locating and opening the partition's system-database file, provided that appropriate access permissions are set on it.

It is normally not necessary to open partitions explicitly because they are usually opened automatically by operations that access them or their contents. For example, once you have a reference to a partition, creating a database in it automatically opens it for update. In general, you open a referenced partition explicitly only when you want to guarantee access to the partition in advance—for example, before starting a complex operation.

Any number of transactions can concurrently open the same partition in any mode.

You must be in an update transaction to open a partition for update. If necessary, you can promote a read transaction to an update transaction by promoting the open mode of the federated database.

## openMode

Gets the current level of access to the referenced autonomous partition.

ooMode openMode() const;

Returns One of the following constants:

- oocNoOpen—the partition is not open in this transaction or an error has been signalled.
- oocRead—the partition is open for read in this transaction.
- oocUpdate—the partition is open for update in this transaction.

### returnAll

Clears the referenced autonomous partition, returning all controlled containers to their home autonomous partitions.

```
ooStatus returnAll() const;
```

- Returns oocSuccess if successful; otherwise oocError.
- Discussion Clearing an autonomous partition clears it of the containers it controls. Specifically:
  - Each container is returned to the control of the autonomous partition that contains the container's database.

• Each container is physically moved from the autonomous partition's system-database file to the file of the container's database.

## sysDBFileHost

Gets the network name of the host containing the system-database file of the referenced autonomous partition.

char \*sysDBFileHost() const;

Returns Pointer to a string containing the network name of the data server host. If an error occurs, a null pointer is returned and the error is signalled.

Discussion The returned string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.

## sysDBFilePath

Gets the fully qualified pathname of the system-database file of the referenced autonomous partition.

char \*sysDBFilePath() const;

- Returns Pointer to a string containing the pathname of the autonomous partition's system-database file. If an error occurs, a null pointer is returned and the error is signalled.
- Discussion The returned string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.

## typeN

Gets the type number of the autonomous-partition class ooAPObj.

ooTypeNumber typeN() const;

Returns Type number of the autonomous-partition class ooAPObj.

## typeName

Gets the name of the autonomous-partition class ooAPObj.

char \*typeName() const;

Returns The string "ooAPObj".

## Discussion The returned string must be treated as read only.

## update

	Opens the referenced autonomous partition for update access.	
	ooStatus update();	
Returns	oocSuccess if successful; otherwise oocError.	
Discussion	This member function is equivalent to calling <u>open(oocUpdate)</u> .	

# ooRefHandle(ooContObj) Classes

Inheritance:	ooRef(ooObj)->ooRef(ooContObj)
Inheritance:	ooHandle(ooObj)->ooHandle(ooContObj)
	The abbreviation <i>ooRefHandle</i> (ooContObj) refers to two non-persistence-capable classes:
	■ ooRef(ooContObj), which represents an <i>object reference</i> to a container.
	■ ooHandle(ooContObj), which represents a <i>handle</i> to a container.
	These two classes are documented together because they define almost identical sets of member functions (exceptions are listed in the "Reference Summary"). These classes, along with the <u>ooContObj</u> class, define the behavior of all containers (instances of ooContObj and its derived classes).
	See:
	<ul> <li>"Reference Summary" on page 515 for an overview of member functions</li> </ul>

• "Reference Index" on page 517 for a list of member functions

# **About Container Handles and References**

An application works with a container indirectly through one or more handles or object references—that is, through instances of <code>ooRefHandle(ooContObj)</code> that are set to reference the desired container. A handle or object reference to a container serves as a type-safe smart pointer that:

- Identifies the container to the application or to another object.
- Provides an interface for operating on the container.
- Manages the memory pointer to the container.
- Provides an indirect member-access operator (->) for accessing the container's public member functions.

It is sometimes more appropriate to use a handle rather than an object reference, and vice versa; the choice is described in "Structure and Behavior" on page 510. A simple guideline is to use handles in function definitions and object references as data member types in persistence-capable class definitions.

## Interface

The ooRefHandle(ooContObj) classes provide the primary interface for operating on a referenced container. Part of this interface consists of member functions defined by ooRefHandle(ooContObj) for specialized operations such as finding the basic objects in a container. The other part of this interface consists of member functions defined by the ooRefHandle(ooCbj) base classes for more general Objectivity/DB operations, such as opening, locking, printing object identifiers, and so on. These member functions are either inherited by the ooRefHandle(ooContObj) classes or redefined wherever type-specific parameters or behavior are required.

Some of the member functions defined by the *ooRefHandle*(ooObj) base classes are *not* available to instances of the *ooRefHandle*(ooContObj) classes. These include member functions for moving, copying, and versioning operations, which apply only to basic objects. The disallowed member functions and operators are redefined as private members of the *ooRefHandle*(ooContObj) classes.

The ooRefHandle(ooContObj) classes themselves serve as base classes for other handle and object-reference classes. Some of these derived classes, such as ooRefHandle(ooGCContObj) and ooRefHandle(ooDefaultContObj), are predefined by Objectivity/C++. The other derived handle and object-reference classes—namely, the ooRefHandle(appClass) classes—are generated by the DDL processor for every application-defined container class appClass. All of the derived handle and object-reference classes provide the same interface as ooRefHandle(ooContObj), with a few redefined member functions for type-specific behavior or parameters; see <u>ooRefHandle(appClass)</u> (page 509).

## **Structure and Behavior**

Although both handles and object references provide a way to reference and operate on a container, they are optimized for different purposes:

- Handles are optimized for accessing persistent objects in memory—that is, for performing multiple operations on a referenced object or repeatedly accessing the object's members.
- Object references are optimized for linking persistent objects—that is, for storing object identifiers persistently in reference attributes, in associations, or as elements of a collection.

## Handles

Handles are optimized for efficient in-memory access because they can automatically obtain and manipulate pointers to referenced objects. Thus, when a handle is set to reference a particular container, the handle stores the object identifier for that container. The first time the container is accessed through the handle, the handle is automatically *opened*—that is, the handle obtains a pointer to the container's representation in memory. This memory pointer enables the handle to access the referenced container quickly during subsequent operations performed through the handle. When the handle is *closed*, it invalidates the pointer but keeps the container's object identifier, so the application can reuse the handle (without resetting it) to access the same container.

Besides maintaining a pointer to the referenced container, an open handle also *pins* the container's memory representation in the Objectivity/DB cache; that is, the handle pins the persistent *container object* itself. (Among other things, the container object is where an application-defined container stores its persistent data.) Pinning guarantees that the container object is readily available in memory for as long as it is needed. Closing the handle removes its particular "pin"; when the last open handle to the container is closed, the last pin is removed and the container object is closed. Objectivity/DB is permitted to swap the page containing the closed container object out of the cache to make room for other open persistent objects. (This affects only the container object itself, and has no effect any open basic objects in the container.)

A handle to a container, like any handle, has cache-related state that is associated with the Objectivity context in which it was created. Therefore, a handle:

- Cannot be stored persistently—for example, as an attribute value. In fact, the DDL processor does not accept ooHandle(ooContObj) as a data member type in a persistence-capable class definition.
- Cannot be passed between Objectivity contexts.

## **Object References**

Object references are optimized for implementing persistent links because they are essentially wrappers for object identifiers. Thus, when an object reference is set to reference a particular container, the object reference stores the object identifier of that container. The object reference never acquires a pointer to the container in memory; instead, whenever the container is accessed through the object reference, the operation is delegated to a temporary handle that provides the necessary pointer.

Because it has no bulky cache-specific state, an object reference to a container:

Can be stored persistently—for example, as an attribute value. The DDL processor accepts ooRef(ooContObj) as a data member type in a persistence-capable class definition.

• Can be passed between Objectivity contexts.

For convenience, an application can use an object reference (instead of a handle) to perform an operation on a referenced container or to access one of the container's members. However, poor performance results when an object reference is used for multiple such operations on the same container, because *each* operation causes a temporary handle to be created, used, and discarded. Performance may also be affected by swapping, because the object reference does not pin the container's memory representation in the Objectivity/DB cache.

# **Working With Container Handles**

**NOTE** For simplicity, this section describes how to work with handles. Except where noted, the same information applies to object references.

An application normally creates a handle as a local variable on the stack, rather than allocating it on the heap. A handle should not be declared as const, because its internal state is changed by any operation that accesses a container through it. (Object references may be declared as const.) Applications should not explicitly define subclasses of the *ooRefHandle*(ooContObj) classes; any necessary subclasses are generated automatically by the DDL processor when an application defines its own subclasses of ooContObj.

A new handle is normally *null*—that is, it contains the value 0 instead of an object identifier for a container. The application can then set the handle to reference a particular container in any of the following ways:

- By creating a new container with <u>operator new</u> of the ooContObj class and assigning the result to the handle.
- By finding an existing container with the handle's <u>exist</u>, <u>lookupObj</u>, or <u>open</u> member function.
- By passing the handle to a member function that sets it, such as the <u>containedIn</u> member function of a basic-object handle, which finds the container where the basic object is located.
- By assignment or initialization from another handle or object reference.

An object reference may be set in any of these ways, with the following exception—the result of operator new may not be assigned to an object reference.

A handle continues to reference the same container until it is set to another container or to null. Furthermore, multiple handles and object references can be set to the same container. A handle of class <code>ooHandle(ooContObj)</code> can be set to

reference any kind of container—that is, an instance of ooContObj or any predefined or application-defined class derived from ooContObj.

An application operates on a container by calling:

- Member functions of a handle that references the container. As indicated in "Reference Summary" on page 515, such member functions allow you to get the attributes of the referenced container, find its basic objects, and so on. To call a member function of a handle, you use the direct member-access operator (.). For example, contH.name calls the <u>name</u> member function of the container handle contH.
- Member functions of the referenced container itself (including member functions inherited from oo0bj).

To call a member function of a referenced container, you use the handle's overloaded indirect member-access operator (<u>operator-></u>). For example, contH->oolsKindOf calls the <u>oolsKindOf</u> member function on the container that is referenced by the handle contH.

Although most of a container handle's member functions operate on the referenced container, some functions operate on the handle itself. For example, you use:

- The assignment operator <u>operator=</u> to set a handle from another handle or from an object reference.
- The inherited comparison operators <u>operator==</u> and <u>operator!=</u> to test whether a handle references the same container as another handle or object reference.
- The inherited member function <u>isNull</u> to test whether a handle is null. (Alternatively, you can use the overloaded <u>operator==</u> to compare a handle to 0.)

A handle preserves its reference to a container across transaction boundaries, provided that the handle does not go out of scope and is not set to null as the result of an abort operation. Before reusing the handle in a new transaction, however, the application should call <u>isValid</u> to test whether the handle is still *valid*—that is, whether it still references an existing container. A handle becomes invalid if it is set to null *or* if the referenced container has been deleted by another process between transactions.

## Specifying a Container Handle to a Function

Objectivity/C++ functions that require a container as input typically obtain the container through a parameter of type const ooHandle(ooContObj) &. You can pass any type of container handle through such a parameter, because ooHandle(ooContObj) is the base class for all container-handle classes. For example, you could pass a handle of type ooHandle(ooGContObj) or

ooHandle(appClass), where appClass is an application-defined container class.

If a function can accept other types of Objectivity/DB objects in addition to containers, the function may do so through a parameter of type const ooHandle(ooObj) &. You can pass a container handle to a parameter of this type because ooHandle(ooContObj) is derived from ooHandle(ooObj). Many functions accept a container handle where a more general handle is requested, including:

- Functions that operate on any persistent object, such as the persistent-collection functions for managing elements.
- Functions that operate on (nearly) any Objectivity/DB object, such as ooDelete and various functions that scan, cluster, or define a scope name within the object.

Some functions do *not* accept a container handle where a more general handle is requested—for example, functions that manage the versioning of basic objects.

## **Opening and Closing a Container Handle**

**NOTE** This subsection applies only to handles, not to object references, which are in effect always closed.

A handle is automatically opened when a container is opened through it. The open container is both locked and represented in memory; the open handle manages a pointer to the container, pinning the container in memory until the handle is closed. A closed handle, which has an object identifier instead of a pointer, can reference either an open or a closed container.

The most common way to open a container through a handle is to do so implicitly by using the handle's indirect member-access operator (<u>operator-></u>) to access a member of the referenced container. Alternatively, a referenced container can be opened by explicit request—for example, by calling the handle's <u>open or update</u> member function. Another way to explicitly open a container is by finding it with a function whose *openMode* parameter is either oocRead or oocUpdate. (Most functions that set a handle to a found container provide an *openMode* parameter for specifying the desired level of access through that handle.) In all cases, if the found or referenced container is already open (for example, because a basic object in it was opened earlier in the transaction), the accessing handle gets a pointer to the container's existing memory representation and adds a pin.

You obtain a closed handle to a container by finding the container with a function whose *openMode* parameter is set to oocNoOpen. Such operations

simply provide the handle with a container's object identifier without adding a pin, even if the container is already open through another handle.

Objectivity/DB automatically closes an open container handle when the handle is destroyed (for example, by going out of scope), when it is set to reference another container, or when the transaction that opened it commits or aborts. An application can close a handle explicitly by calling the handle's <u>close</u> member function.

# **Reference Summary**

The following table summarizes all the member functions that are available to instances of <code>ooRefHandle(ooContObj)</code>. Member functions indicated as (inherited) are documented with the <code>ooRefHandle(ooObj)</code> classes (page 593).

The summarized member functions are defined on both the object-reference class and the handle class. Two operators are defined on only the handle class, namely, <u>operator\*</u> and <u>operator ooContObj\*</u>.

Creating a Handle or Object Reference	<u>ooHandle(ooContObj)</u> ooRef(ooContObj)
Setting the Handle or Object Reference	<u>operator=</u> <u>open</u> <u>exist</u> <u>lookupObj</u>
Comparing Handles and Object References	<u>operator==</u> (inherited) operator!= (inherited)
Accessing the Container	<u>operator-&gt;</u> <u>operator*</u> <u>operator ooContObj*</u> <u>ptr(ODMG</u> )
Opening, Closing, and Locking the Container	open openMode <u>lock</u> (inherited) <u>lockNoProp</u> <u>refreshOpen</u> close
Modifying the Container	<u>update</u> (inherited) <u>delete_object</u> (inherited)

Getting Information About the Container	<u>name</u> <u>hash</u> <u>nPage</u> <u>numLogicalPages</u> <u>percentGrow</u> <u>typeN (inherited)</u> <u>typeName (inherited)</u> <u>print (inherited)</u> <u>sprint (inherited)</u>
Testing the Container	<u>exist</u> isUpdated
Testing the Handle or Object Reference	<u>is null (inherited)</u> <u>isNull (inherited)</u> <u>isValid (inherited)</u> <u>operator int (inherited)</u> <u>operator ooContObj*</u> <u>operator ooObj* (inherited)</u>
Working With Scope Names	<u>nameObj</u> (inherited) <u>getObjName</u> (inherited) <u>unnameObj</u> (inherited) <u>getNameObj</u> (inherited) <u>getNameScope</u> (inherited)
Finding Objects	<u>exist</u> <u>open</u> <u>lookupObj</u> <u>contains</u> <u>containedIn</u> <u>controlledBy</u> ( <i>FTO</i> ) <u>getNameObj</u> ( <i>inherited</i> ) <u>getNameScope</u> ( <i>inherited</i> )
Converting Objects	<u>convertObjects</u>
Working With Autonomous Partitions ( <i>FTO</i> )	transferControl controlledBy returnControl
ODMG Interface	operator d Ref Any ptr

# **Reference Index**

close	Explicitly closes this handle.
<u>containedIn</u>	Finds the database that contains the referenced container.
<u>contains</u>	Initializes an object iterator to find all basic objects stored in the referenced container.
<u>controlledBy</u>	(FTO) Finds the autonomous partition, if any, that controls the referenced container.
<u>convertObjects</u>	Performs on-demand object conversion on any affected objects in the referenced container.
exist	Tests whether the specified container exists in the specified database; if successful, sets this object reference or handle to reference the container.
hash	Gets the hash value for the referenced container.
isUpdated	Tests whether the referenced container has already been updated and committed by another transaction.
lockNoProp	Explicitly locks the referenced container, without propagating locks to associated destination objects.
<u>lookupObj</u>	Finds the container with the specified scope name in the specified scope; if successful, sets this object reference or handle to reference the found container.
name	Gets the system name of the referenced container.
<u>nPage</u>	Gets the current number of storage pages in the referenced container.
<u>numLogicalPages</u>	Gets the current number of logical pages in the referenced container.
<u>ooHandle(ooContObj)</u>	Default constructor that constructs a null handle.
<u>ooHandle(ooContObj)</u>	Constructs a handle that references the same container as the specified object reference, handle, pointer, or ODMG generic reference.
<u>ooRef(ooContObj)</u>	Default constructor that constructs a null object reference.

<u>ooRef(ooContObj)</u>	Constructs an object reference that references the same container as the specified object reference, handle, pointer, or ODMG generic reference.
open	Explicitly opens the referenced or specified container, preparing the container for the specified level of access.
openMode	Gets the current level of access to the referenced container.
operator->	Indirect member-access operator; accesses a member of the referenced container.
operator*	Handle class only. Dereference operator; returns the container referenced by this handle.
<u>operator=</u>	Assignment operator; sets this object reference or handle to reference the same container as the specified object reference, handle, or pointer.
<u>operator d Ref Any</u>	(ODMG) Conversion operator that returns an ODMG generic reference to the referenced container.
<u>operator ooContObj*</u>	Handle class only. Conversion operator that returns a pointer to the referenced container.
percentGrow	Gets the growth factor for the referenced container.
ptr	(ODMG) Returns a C++ pointer to the referenced container.
<u>refresh0pen</u>	Reopens the referenced container, refreshing the view of the container in the MROW transaction reading it.
<u>returnControl</u>	(FTO) Releases the referenced container from the currently controlling autonomous partition and returns control to the autonomous partition of the container's database.
<u>transferControl</u>	(FTO) Transfers the referenced container into the control of the specified autonomous partition.

# Constructors

## ooHandle(ooContObj)

Default constructor that constructs a null handle.

```
ooHandle(ooContObj)();
```

;

## ooHandle(ooContObj)

Constructs a handle that references the same container as the specified object reference, handle, pointer, or ODMG generic reference.

1.	ooHandle(ooContObj)(
	<pre>const ooRefHandle(ooContObj) &amp;containerRH);</pre>
2.	<pre>ooHandle(ooContObj)(const ooContObj *containerP)</pre>
(ODMG) 3.	ooHandle(ooContObj)(const d_Ref_Any &from);

#### Parameters containerRH

Object reference or handle to a container.

#### containerP

Pointer to a container. The pointer *must* be the result of <u>operator new</u> on a container class.

#### from

(ODMG) An ODMG generic reference to a container. An error is signalled if *from* references an object that is not an instance of ooContObj or a class derived from ooContObj.

Discussion Variants 1 and 3 allow a new handle to be constructed from an existing object reference, handle, or ODMG generic reference. If the new handle is constructed from an existing open handle, the new handle is open; in all other cases, the new handle is closed.

Variant 2, which constructs a handle from the specified pointer, has a narrower purpose—to obtain an open handle to a newly created container so you can perform persistence operations on it, and so the container can eventually be unpinned when it is no longer needed in memory.

## ooRef(ooContObj)

Default constructor that constructs a null object reference.

```
ooRef(ooContObj)();
```

## ooRef(ooContObj)

Constructs an object reference that references the same container as the specified object reference, handle, pointer, or ODMG generic reference.

```
1. ooRef(ooContObj)(
```

const ooRefHandle(ooContObj) &containerRH);

2. ooRef(ooContObj)(const ooContObj \*containerP);

(ODMG) 3. ooRef(ooContObj)(const d\_Ref\_Any & from);

#### Parameters containerRH

Object reference or handle to a container.

#### containerP

Pointer to a container. The pointer may *not* be the result of operator new. Instead, the pointer must be the result of using either <u>operator\_ooContObj\*</u> on a handle *or* <u>ptr</u> on an object reference or handle earlier in the same transaction. If the specified pointer was originally extracted from a handle, that handle must still exist and reference the same container.

#### from

(ODMG) An ODMG generic reference to a container. An error is signalled if *from* references an object that is not an instance of ooContObj or a class derived from ooContObj.

Discussion Variants 1 and 3 allow a new object reference to be constructed from an existing object reference, handle, or ODMG generic reference.

Variant 2 has a narrower purpose, which is to allow you to resume persistence operations on a container after manipulating it through a pointer. The use of this variant should be rare, however, because pointers are not normally used to manipulate containers.

# Operators

### operator->

Indirect member-access operator; accesses a member of the referenced container.

```
const ooContObj *operator->();
```

Returns Pointer to the referenced container.

Discussion The accessed container is opened for read, if it is not already open.

You use operator-> in an expression handle->member, where handle is an instance of ooHandle(ooContObj) and member is the name of a public member defined on class ooContObj. As for any overloading of the C++ member-access operator (->), the expression handle->member is interpreted as (handle.operator->())->member. That is, the overloaded operator-> returns a pointer to the referenced container, and then the ordinary C++ operator-> selects the specified member of that container, returning the value of that member.

If the referenced object is an instance of <code>ooContObj</code>, <code>operator-></code> can access any of the object's public members. If the referenced object is an instance of a derived class, <code>operator-></code> accesses only the members that are defined in <code>ooContObj</code>.

### operator\*

	<i>Handle class only.</i> Dereference operator; returns the container referenced by this handle.
	ooContObj &operator*();
Returns	C++ reference to the container referenced by this handle.
Discussion	This operator enables you to pass a handle to a function that accepts a container by reference. This operator is analogous to the $C++$ operator* for dereferencing a pointer.
	The container referenced by this handle is opened for read, if it is not already open.
	<b>Warning:</b> The returned C++ reference is guaranteed valid only as long as the handle exists, remains open, and references the same container.
Example	This example uses operator* to pass a container handle to helperFunction, which accepts a C++ reference to a container.
	<pre>void helperFunction(ooContObj &amp;aContainer);</pre>
	<pre>void processContainer(ooHandle(ooContObj) &amp;contH) {</pre>
	helperFunction(*contH);
	}

## operator=

Assignment operator; sets this object reference or handle to reference the same container as the specified object reference, handle, or pointer.

#### Parameters containerRH

Object reference or handle to a container.

#### containerP

0, or a nonnull pointer to a container:

- If you are assigning to a handle, the specified pointer *must* be the result of <u>operator new</u> on ooContObj or an application-defined derived class.
- If you are assigning to an object reference, the specified pointer may not be the result of operator new. Instead, the pointer must be the result of using either <u>operator ooContObj\*</u> on a handle or <u>ptr</u> on an object reference or handle earlier in the same transaction. If the specified pointer was originally extracted from a handle, that handle must still exist and reference the same object.

# Returns This object reference or handle, after it has been set to reference the specified container.

Discussion Variant 1 allows you to use the specified object reference or handle to produce another object reference or handle to the same container. If you are assigning to a handle from an open handle, the returned handle is open; in all other cases, the returned handle is closed.

Variant 2 allows you to set this object reference or handle to null. Otherwise, assignment-from-pointer has two specific purposes, depending on whether you are assigning to a handle or to an object reference:

- Pointer-to-handle assignment enables you to obtain an open handle to a newly created container, so you can perform persistence operations on it, and so the container can eventually be unpinned when it is no longer needed in memory.
- Pointer-to-object-reference assignment enables you to resume persistence operations on a container after manipulating it through a pointer. This usage of variant 2 is rare, because pointers are not normally used to manipulate containers.

## operator d\_Ref\_Any

(*ODMG*) Conversion operator that returns an ODMG generic reference to the referenced container.

```
operator d_Ref_Any() const;
```

## operator ooContObj\*

*Handle class only.* Conversion operator that returns a pointer to the referenced container.

```
operator ooContObj*();
```

Returns Pointer to the container referenced by this handle. Returns a null pointer if this handle is a null handle.

#### Discussion This conversion operator enables you to:

- Pass a handle to a function that accepts a pointer to a container.
- Assign a handle to an ooContObj\* variable (for example, to pass to the overloaded operator delete).
- Use a handle as the conditional expression in an if or while statement to test whether the handle is null.

The container referenced by this handle is opened for read, if it is not already open.

**WARNING** The returned pointer is guaranteed valid only as long as the handle exists, remains open, and references the same container.

An application generally relies on handles to provide memory management for persistent objects, and avoids the explicit use of pointers to such objects. On occasion, explicit use of pointers is required for performance reasons or for compatibility with functions that are indifferent to persistence (for example, in legacy code or a third-party library). However, you should be careful when manipulating a pointer from a handle because the validity of the pointer depends on the state of the handle. You should not use the returned pointer in other persistence operations (for example, do not pass it to any Objectivity/C++ member function other than the overloaded operator delete).

# Example The container handle contH is used as a conditional expression which evaluates to 0 if the handle is null.

ooHandle(ooContObj) contH; ... // Set contH to some container if (contH) { ... //Do something interesting if initialization was successful }

See also <u>ptr</u>

# **Member Functions**

## close

Explicitly closes this handle. ooStatus close() const; Returns oocSuccess if successful: otherwise oocError. Discussion This member function is redundant for object references, which are, in effect, always closed. Therefore, you should use this member function only on handles. Objectivity/DB automatically closes container handles when they go out of scope, when they are set to reference other containers, or when the transaction that opened them commits or aborts. You can use the close member function to close a container handle explicitly. This informs Objectivity/DB that the application no longer requires access to the referenced container through this handle. Closing does not, however, affect any open objects in the container, nor does it release the lock on the container; locks are released only by committing or aborting a transaction. When closed explicitly, a container handle retains the object identifier of the container to which it refers, so you can reopen it without reinitializing. Note, however, that a retained object identifier can become invalid between transactions (for example, because a concurrent process has deleted the corresponding container), and opening a handle with an invalid object identifier signals an error. Closing a container handle invalidates its pointer to the container's representation in the Objectivity/DB cache. Closing the last open handle to a particular container unpins and closes the container object; closing the last open object on a buffer page permits Objectivity/DB to swap the page out of the cache as needed. Note that this affects only the persistent data defined for application-specific containers, and does not affect any open basic objects in the container.

## containedIn

Finds the database that contains the referenced container.

- ooHandle(ooDBObj) containedIn() const;

3.	ooHandle(ooDBObj)&c	ontainedIn(	
	ooHandle(ooDBObj)	&database)	const;

Parameters	<i>database</i> Object reference or handle to be set to the found database.
Returns	Object reference or handle to the found database.
Discussion	When called without a <i>database</i> parameter, containedIn allocates a new database handle and returns it. Otherwise, containedIn returns the object reference or handle that is passed to it.

## contains

Initializes an object iterator to find all basic objects stored in the referenced container.

ooStatus contains( ooItr(ooObj) &objI, const ooMode openMode = oocNoOpen) const;

#### Parameters obj1

Object iterator for finding the contained basic objects.

#### openMode

Intended level of access to the basic objects found by the iterator's next member function:

- oocNoOpen (the default) causes next to set the iterator to the next basic object without opening it.
- oocRead causes next to open the next basic object for read.
- oocUpdate causes next to open the next basic object for update.

Returns oocSuccess if successful; otherwise oocError.

## controlledBy

(*FTO*) Finds the autonomous partition, if any, that controls the referenced container.

- 1. ooHandle(ooAPObj) controlledBy();

Parameters	<i>partition</i> Object reference or handle to set to the controlling autonomous partition.
Returns	Object reference or handle to the controlling autonomous partition. If the container has not been transferred to the control of a partition, a null object reference or handle is returned.
Discussion	When called without a parameter, <code>controlledBy</code> allocates a new autonomous-partition handle and returns it. Otherwise, <code>controlledBy</code> returns the object reference or handle that is passed to it.
See also	<u>returnControl</u> <u>transferControl</u>

## convertObjects

Performs on-demand object conversion on any affected objects in the referenced container.

ooStatus convertObjects();

- Returns oocSuccess if successful, or oocError if the federated database is opened only for read.
- Discussion Object conversion is the process of making existing persistent objects consistent with class definition changes introduced by schema evolution. Certain schema evolution operations affect how instances of a class should be laid out in storage. After you perform such operations, existing objects of the changed classes are rendered out-of-date until they are converted to their new representations.

In general, you can allow each affected object to be converted automatically the first time it is accessed after schema evolution, potentially distributing the performance impact of conversion across many transactions. Alternatively, you can concentrate the performance impact of conversion into fewer transactions by converting all the affected objects in a container, a database, or a federated database *on demand*. You use this member function in an update transaction to convert the affected objects in a container on demand. This member function has no effect if the affected objects in the container have already been converted.

**Note:** On-demand object conversion cannot be used for schema operations that require an upgrade application; see ooTrans::<u>upgrade</u>.

The convertObjects member function automatically drops any index that is invalidated by a schema evolution change. Specifically, if you changed the type or deleted a data member that is a key field in a key description, the corresponding indexes are dropped.

See also	Chapter 19, "Object Conversion," in the Objectivity/C++ programmer's guide		
exist			
	Tests whether the specified container exists in the specified database; if successful, sets this object reference or handle to reference the container.		
	<pre>ooBoolean exist( const ooHandle(ooDBObj) &amp;database, const char *contSysName, const ooMode openMode = oocNoOpen);</pre>		
Parameters	database Handle to the database to search.		
	contSysName System name of the desired container.		
	<ul> <li>openMode</li> <li>Intended level of access to the container, if it exists:</li> <li>Specify oocNoOpen (the default), to set this object reference or handle to the container without opening it.</li> <li>Specify oocRead to open the container for read.</li> <li>Specify oocUpdate to open the container for update.</li> </ul>		
Returns	oocTrue if the specified container exists, or oocFalse if the container does not exist or if it is not accessible.		
Discussion	If the specified container exists, this object reference or handle is set to reference it; otherwise, this object reference or handle is set to null.		
	If you specifically want to test for existence, you use the <i>openMode</i> parameter's default value ( <i>oocNoOpen</i> ). Otherwise, a return value of <i>oocFalse</i> could mean either that the container doesn't exist, or that it does exist, but cannot be opened.		
Νοτε	This member function identifies a container using its system name. To find a container by its scope name, use <u>lookupObj</u> .		
hash			

Gets the hash value for the referenced container.

int32 hash() const;

Returns	-1 if this object reference or handle is null or if the referenced container cannot be opened; otherwise, returns the referenced container's hash value:		
	<ul> <li>0 indicates a nonhashed container.</li> <li>1 or greater indicates a hashed container, and is also the clustering factor for keyed objects.</li> </ul>		
Discussion	A container's hash value is set when the container is created. A hashed container and the objects in it can be used as scopes for naming objects; furthermore, a hashed container can contain keyed objects. A clustering factor is the number of sequentially keyed objects to be placed onto a page when keyed objects are created within the container.		
See also	ooContObj:: <u>operator_new</u> .		

## isUpdated

	Tests whether the referenced container has already been updated and committed by another transaction.
	ooBoolean isUpdated() const;
Returns	ocTrue if the referenced container has been updated and committed by another transaction since being locked for read by the current <u>MROW</u> transaction; otherwise, oocFalse.
Discussion	You can use this member function within an MROW transaction to determine whether to call <u>refreshOpen</u> .

## lockNoProp

Explicitly locks the referenced container, without propagating locks to associated destination objects. ooStatus lockNoProp(const ooLockMode lockMode) const; Parameters lockMode Type of lock to request: oocLockRead requests a read lock. oocLockUpdate requests an update lock. Returns oocSuccess if the requested lock is obtained; otherwise oocError. Discussion Objectivity/DB operations request and obtain locks implicitly as they are needed. You use this member function to obtain a lock explicitly when you want to reserve access to an object in advance, but you do not want to lock any

associated destination objects, even along associations that have lock propagation enabled.

Locking a container locks all the basic objects in it.

Whenever a lock is requested on a container, Objectivity/DB applies the transaction's <u>concurrent access policy</u> to determine whether the requested lock is compatible with other existing locks. An error is signaled if a requested lock cannot be obtained.

See also

*ooRefHandle*(ooObj)::<u>lock</u> <u>ooLockMode</u> **global type** 

## lookupObj

Finds the container with the specified scope name in the specified scope; if successful, sets this object reference or handle to reference the found container.

ooStatus lookupObj(	
const ooHandle(ooObj	) &scope,
const char * <i>name</i> ,	
const ooMode <i>openMode</i>	e = oocRead) const;

#### Parameters scope

Handle to the scope object that defines the name scope to search. *scope* can reference the federated database, a database, a persistent container, a persistent basic object, or an autonomous partition.

name

Scope name to look up in the scope specified by *scope*.

#### openMode

Intended level of access to the found container:

- Specify oocRead (the default) to open the object for read.
- Specify oocUpdate to open the object for update.
- Specify oocNoOpen to set this object reference or handle to the object without opening it.

Returns oocSuccess if a container is found; otherwise oocError.

# Discussion Scope-named objects are found only if they are instances of ooContObj or application-defined classes derived from ooContObj.

name	
	Gets the system name of the referenced container.
	char *name() const;
Returns	Pointer to a string containing the system name. If the container does not have a system name, the returned pointer is null.
Discussion	The string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.
nPage	
	Gets the current number of storage pages in the referenced container.
	uint32 nPage();
Returns	The current number of storage pages in the container; returns 0 if this object reference or handle is null or if the container cannot be opened.
Discussion	You can use this function to help you calculate the size (in bytes) of a container within a database.
	A storage page is the minimum unit of transfer to and from disk and across networks. The size of a storage page is configurable for each federated database and is set with the federated database is created. A federated database's storage pages are usually sized so that one or more typical persistent objects (called <i>small</i> <i>objects</i> ) will fit within a single storage page; by definition, a <i>large object</i> spans multiple storage pages. You can get the page size of a federated database by calling the <u>pageSize</u> member function on a federated-database handle.
	A subset of a container's storage pages are also logical pages; logical pages are storage pages that contain either one or more small objects, or the header information for a large object. If no large objects reside in a container, the number of storage pages will be close to the number of logical pages. A container with very large objects will have many more storage pages than logical pages. You can get the number of logical pages by calling the <u>numLogicalPages</u> member function.

## numLogicalPages

Gets the current number of logical pages in the referenced container.

```
uint32 numLogicalPages();
```

- Returns The current number of logical pages in the container; returns 0 if this object reference or handle is null or if the container cannot be opened.
- Discussion You can use this function to monitor when the container will reach its maximum logical page limit (65535).

A logical page is a storage page that contains either one or more small objects, or the header information for a large object. Logical pages are numbered within each container; a logical page number appears as one of the fields within the object identifier of a persistent object.

Logical pages are a subset of a container's storage pages. If no large objects reside in a container, the number of logical pages will be close to the number of storage pages. A container with very large objects will have many fewer logical pages than storage pages. You can get the number of storage pages by calling the <u>nPage</u> member function.

#### open

Explicitly opens the referenced or specified container, preparing the container for the specified level of access.

	<pre>1. ooStatus open(const ooMode openMode = oocRead);</pre>		
	<pre>2. ooStatus open(     const ooHandle(ooDBObj) &amp;database,     const char *contSysName,     const ooMode openMode = oocRead);</pre>		
Parameters	openMode		
	Intended level of access to the opened container:		
	<ul> <li>Specify occRead (the default) to open the container for read. This implicitly locks the container (and the basic objects in it) for read.</li> </ul>		
	<ul> <li>Specify occupdate to open the container for update (read and write).</li> <li>This implicitly locks the container (and the basic objects in it) for update.</li> </ul>		
	database		
	Handle to the database in which to find the specified container.		
	contSysName		
	System name of the container to open.		
Returns	oocSuccess if successful; otherwise oocError.		
Discussion	Variant 1 assumes that this object reference or handle already references a container, and opens the referenced container.		

Variant 2 finds and explicitly opens the container with the specified system name, and sets this object reference or handle to reference it. An error is signalled if no container exists with the specified system name in the specified database. This variant is especially useful when you want to use the container as an entry point into your data. For example, you might find and open a container so you can iterate over the basic objects in it.

Opening a container makes it available to an application by:

- Implicitly locking the container for read or update, as specified by openMode.
- Obtaining a representation of the container in memory, either by fetching storage pages from the database or reusing an existing memory representation that is guaranteed current. This memory representation includes pages describing the container and any data defined for it by the application (if the container is an instance of an application-defined container class).

Opening a container does not open any of the basic objects in it.

Opening a container for update additionally marks it as modified, causing any application-specific data to be written to the database when the transaction commits, whether or not that data was actually modified. You must be in an update transaction to open a container for update. If necessary, you can promote a read transaction to an update transaction by promoting the open mode of the federated database.

The open operation fails if the container cannot be locked—for example, due to a lock conflict. Objectivity/DB applies the transaction's <u>concurrent access policy</u> to determine whether the requested lock is compatible with other existing locks. Once a lock is obtained, it is kept until the transaction either commits or aborts.

A container is opened automatically when you open one of the basic objects in it. Furthermore, using <u>operator-></u> to access a member of a container implicitly opens the container for read. You open a container explicitly when:

- You require update access so you can modify the container's application-specific data.
- You want to reserve either read or update access to the container in advance—for example, before starting a complex operation.

## openMode

Gets the current level of access to the referenced container.

ooMode openMode() const;

Returns One of the following constants:

- oocNoOpen—the container is not open in this transaction.
- oocRead—the container is open for read in this transaction.
- oocUpdate—the container is open for update in this transaction.

### percentGrow

Gets the growth factor for the referenced container.

uint32 percentGrow() const;

- Returns The growth factor for the container; returns 0 if this object reference or handle is null or if the container cannot be opened.
- Discussion A container's growth factor is set when the container is created. The growth factor is the amount by which the container may grow when it needs to accommodate more basic objects, expressed as a percentage of its current size.
- See also ooContObj::<u>operator new</u>.

#### ptr

(ODMG) Returns	A C++ pointer to the referenced container.	
( = = = = = = = = = = = = = = = = = = =	· · · · · · · · · · · · · · · · · · ·	

ooContObj \*ptr();

Discussion You use this member function to obtain a pointer to a container—for example, to pass to a function that accepts a pointer instead of a handle or object reference.

If ptr is called on an object reference, the referenced container is opened for update. If ptr is called on a handle, the referenced container is opened for read.

Warning: The returned pointer is guaranteed valid for only a limited time:

- If ptr is called on an object reference, the returned pointer is valid and the container is pinned in memory until the end of the transaction.
- If ptr is called on a handle, the returned pointer is valid only as long as the handle exists, remains open, and references the same container (equivalent to <u>operator ooContObj\*</u>).

An application generally relies on handles to provide memory management for persistent objects, and avoids the explicit use of pointers to such objects. On occasion, explicit use of pointers is required for performance reasons or for compatibility with functions that are indifferent to persistence (for example, in legacy code or a third-party library). However, you should be careful when manipulating persistent objects through pointers:

- Pointers extracted from handles become invalid if the handles change or go out of scope.
- Pointers extracted from object references can cause the Objectivity/DB cache to run out of memory if too many objects are pinned until the end of the transaction.

You should not use the returned pointer in other persistence operations (for example, do not pass it to any Objectivity/C++ member function other than the overloaded <code>operator delete</code>).

## refreshOpen

Reopens the referenced container, refreshing the view of the container in the MROW transaction reading it.

```
ooStatus refreshOpen(
    const ooMode openMode,
    ooBoolean *pIsUpdated,
    ooBoolean closeHandles = oocFalse);
```

Parameters openMode

Intended level of access to the reopened container:

- Specify occRead to open the container for read. This implicitly requests a read lock on the container.
- Specify occupdate to open the container for update (read and write). This implicitly requests an update lock on the container.
- pIsUpdated

Pointer to a value that, on return, indicates whether the referenced container has been updated and committed by another transaction since being locked for read by the current MROW transaction. The returned value corresponds to the result of the <u>isUpdated</u> member function.

#### closeHandles

Action to take if the transaction has any open handles to objects in the container:

- Specify oocTrue to close the all open handles.
- Specify oocFalse (the default) to signal an error and take no other action.

Returns oocSuccess if successful; otherwise oocError.

Discussion When you use an <u>MROW</u> transaction to read a container, one other transaction is allowed to concurrently update that container. If the updating transaction commits, your view is rendered out-of-date. You use the refreshOpen member

function to open the most recently committed version of an updated container within an MROW transaction. This means that each open object in the container is implicitly closed and then reopened the next time you access it.

## returnControl

(*FTO*) Releases the referenced container from the currently controlling autonomous partition and returns control to the autonomous partition of the container's database.

ooStatus returnControl() const;

Returns oocSuccess if successful; otherwise oocError.

- Discussion This member function physically moves the container back to the file of its database.
- See also <u>controlledBy</u> <u>transferControl</u>

## transferControl

(*FTO*) Transfers the referenced container into the control of the specified autonomous partition.

ooStatus	t	ransferControl(		
cons	t	ooHandle(ooAPObj)	&newControllingAP)	const;

 

 Parameters
 newControllingAP Handle to the autonomous partition that is to control the container.

 Returns
 oocSuccess if successful; otherwise oocError.

 Discussion
 This member function physically moves the referenced container into the system database file of the specified autonomous partition.

See also <u>controlledBy</u> <u>returnControl</u>

# ooRefHandle(ooDBObj) Classes

Inheritance:	ooRef(ooObj)->ooRef(ooDBObj)			
Inheritance:	ooHandle(ooObj)->ooHandle(ooDBObj)			
	The abbreviation <i>ooRefHandle</i> (ooDBObj) refers to two non-persistence-capable classes:			
	<ul> <li>ooRef(ooDBObj), which represents an <i>object reference</i> to a database.</li> <li>ooHandle(ooDBObj), which represents a <i>handle</i> to a database.</li> </ul>			
	The two classes <code>ooRef(ooDBObj)</code> and <code>ooHandle(ooDBObj)</code> are documented together because they define the same set of member functions. These member functions provide the primary interface for operating on Objectivity/DB databases (instances of <code>ooDBObj</code> ).			
	( <i>DRO</i> ) The <i>ooRefHandle</i> (ooDBObj) classes also provide the primary interface for managing database images (also instances of <u>ooDBObj</u> ), which you can create if you have bought and installed both Objectivity/DB Data Replication Option (Objectivity/DRO) and Objectivity/DB Fault Tolerant Option (Objectivity/FTO).			

See:

- "Reference Summary" on page 541 for an overview of member functions
- "Reference Index" on page 543 for a list of member functions

## **About Database Handles and References**

An application works with a database indirectly through one or more handles or object references—that is, through instances of ooRefHandle(ooDBObj) that are set to reference the desired database. A handle or object reference to a database both identifies the database and provides the complete public interface for operating on it.

Handles and object references to databases do not support indirect member access: the <code>ooRefHandle(ooDBObj)</code> classes provide no indirect member-access operator (->), and the <code>ooDBObj</code> class defines no public member functions other than a constructor and <code>operator</code> new.

You can work with a database through either a handle or an object reference—the choice is arbitrary, except as described in "Structure and Behavior" on page 538. Most applications use handles rather than object references.

## Interface

The *ooRefHandle*(*ooDBObj*) classes provide the primary interface for operating on a referenced database. Part of this interface consists of member functions defined by *ooRefHandle*(*ooDBObj*) for specialized operations such as getting the number of containers in a database. The other part of this interface consists of member functions defined by the *ooRefHandle*(*ooObj*) base classes for more general Objectivity/DB operations, such as opening, locking, printing object identifiers, and so on. These member functions are either inherited by the *ooRefHandle*(*ooDBObj*) classes or redefined wherever type-specific parameters or behavior are required.

Some of the member functions defined by the *ooRefHandle*(ooObj) base classes are *not* available to instances of the *ooRefHandle*(ooDBObj) classes. These include member functions for moving, copying, versioning, scope-naming, and member-access operations, which apply only to basic objects or persistent objects. The disallowed member functions and operators are redefined as private members of the *ooRefHandle*(ooDBObj) classes.

## **Structure and Behavior**

Handles and object references to databases are essentially wrappers for database identifiers. For example, when you set a handle to reference a particular database, the handle stores the object identifier of that database. If the database is then opened through the handle, Objectivity/DB uses the identifier to locate the database file on disk. Subsequent member-function calls on the handle operate on the instance of ooDBObj that represents the identified database in memory.

In general, object references are optimized for implementing links among related persistent objects, while handles are optimized for memory management and member-access. When a database is referenced, however, these optimizations are largely irrelevant, because databases (unlike persistent objects):

- Cannot be linked (for example, through associations).
- Are not subject to memory management (they have no attributes for persistent data and are therefore not manipulated through pointers).
- Have no accessible members.

One significant exception is that a handle to a database, like any handle, contains cache-related state that is associated with the Objectivity context in which it was created. Therefore, only object references (but not handles) can be passed between Objectivity contexts. Otherwise, a handle to a database and an object reference to a database are functionally equivalent.

# **Working With Database Handles**

**NOTE** For simplicity, this section describes how to work with handles. Except where noted, the same information applies to object references.

An application normally creates a handle as a local variable on the stack, rather than allocating it on the heap. A handle should not be declared as const, because its internal state may be changed when a database is accessed through it. (Object references may be declared as const). Applications should not create subclasses of the *ooRefHandle*(ooDBObj) classes.

A new handle is normally *null*—that is, it contains the value 0 instead of the object identifier for a database. The application can then set the handle to reference a particular database in any of the following ways:

- By creating a new database with <u>operator new</u> of the ooDBObj class and assigning the result to the handle.
- By finding an existing database with the handle's <u>exist</u> or <u>open</u> member function.
- By passing the handle to a member function that sets it, such as the <u>containedIn</u> member function of a container handle, which finds the database where a container is located.
- By assignment or initialization from another handle or object reference.

A handle continues to reference the same database until it is set to another database or to null. Furthermore, multiple handles and object references can be set to the same database.

An application operates on a database by calling member functions on a handle that references it. To call a member function of a handle, you use the direct member-access operator (.). For example, dbH.name calls the <u>name</u> member function of the handle dbH.

As indicated in "Reference Summary" on page 541, an application can use the handle's member functions to get and change a referenced database's attributes, find its containers, tidy its disk space, and so on. If Objectivity/FTO and Objectivity/DRO are installed, the application can also create and manage

multiple images of a database. For more information about operating on databases and database images, see Chapter 8, "Storage Objects," and Chapter 28, "Database Images," respectively, in the Objectivity/C++ programmer's guide.

Although most of a database handle's member functions operate on the referenced database, some functions operate on the handle itself. For example, you use:

- The assignment operator <u>operator=</u> to set a handle from another handle or from an object reference.
- The inherited comparison operators <u>operator==</u> and <u>operator!=</u> to test whether a handle references the same database as another handle or object reference.
- The inherited member function <u>isNull</u> to test whether a handle is null. (Alternatively, you can use the overloaded <u>operator==</u> to compare a handle to 0.)

A handle preserves its reference to a database across transaction boundaries, provided that the handle does not go out of scope and is not set to null as the result of an abort operation. Before reusing the handle in a new transaction, however, the application should call <u>isValid</u> to test whether the handle is still *valid*—that is, whether it still references an existing database. A handle becomes invalid if it is set to null *or* if the referenced database has been deleted by another process between transactions.

Objectivity/C++ functions that require a database as input normally obtain the database through a parameter of type const ooHandle(ooDBObj) &. If a function manipulates other types of Objectivity/DB objects as well, the parameter type may be specified as const ooHandle(ooObj) &. You can pass a database handle to a parameter of this type, because ooHandle(ooDBObj) is derived from ooHandle(ooObj). In practice, however, relatively few functions accept a database handle where a more general handle is requested; these include ooDelete and various functions that require a storage object for scanning, clustering, or as a name scope. In most cases, a function that requests a general-purpose handle operates only on basic objects or persistent objects, and signals an error if you specify a database handle.

Any operation that affects a database opens it implicitly if it is not already open; an application does not need to open a database explicitly unless it is used as the entry point into the data or unless access must be guaranteed in advance. A handle is automatically closed when it is destroyed (for example, when it goes out of scope). However, closing one or more handles to a particular database has no effect on that database, which remains open until the transaction commits or aborts.
# **Reference Summary**

The following table summarizes all the member functions that are available to instances of ooRefHandle(ooDBObj). Member functions indicated as *(inherited)* documented with the <u>ooRefHandle(ooObj)</u> classes (page 593).

Creating a Handle or Object Reference	<u>ooHandle(ooDBObj)</u> ooRef(ooDBObj)
Setting the Handle or Object Reference	<u>operator=</u> <u>open</u> <u>exist</u>
Comparing Handles and Object References	<u>operator==</u> (inherited) <u>operator!=</u> (inherited)
Opening, Closing, and Locking the Database	open update openMode isReadOnly setReadOnly lock close
Modifying the Database	update setReadOnly change changePartition (FTO) tidy
Getting Information About the Database	name fileName pathName hostName numContObjs getImageFileName (DRO) getImageHostName (DRO) getImagePathName (DRO) getImageWeight (DRO) numImages (DRO) typeN typeName print (inherited) sprint (inherited)

Testing the Database	<u>exist</u> <u>isReadOnly</u> <u>hasImageIn</u> ( <i>DRO</i> ) <u>isAvailable</u> ( <i>FTO</i> ) <u>isImageAvailable</u> ( <i>DRO</i> ) <u>isReplicated</u> ( <i>DRO</i> )
Testing the Handle or Object Reference	<u>is_null (inherited)</u> <u>isNull (inherited)</u> <u>isValid</u> <u>operator_int (inherited)</u> <u>operator_ooObj* (inherited)</u>
Finding Objects	<u>exist</u> <u>open</u> <u>contains</u> <u>getDefaultContObj</u> <u>containedIn</u> <u>containingPartition</u> ( <i>FTO</i> ) <u>getTieBreaker</u> ( <i>DRO</i> )
Converting Objects	<u>convertObjects</u>
Working With Containers	<u>contains</u> <u>getDefaultContObj</u> <u>numContObjs</u>
Working With Autonomous Partitions	<u>changePartition</u> ( <i>FTO</i> ) <u>containingPartition</u> ( <i>FTO</i> ) <u>isAvailable</u> ( <i>FTO</i> )
Working With Database Images	<pre>replicate (DRO) isReplicated (DRO) numImages (DRO) setImageWeight (DRO) setTieBreaker (DRO) negotiateQuorum (DRO) isNonQuorumRead (DRO) getAllowNonQuorumRead (DRO) setAllowNonQuorumRead (DRO) partitionsContainingImage (DRO) hasImageIn (DRO) getImageFileName (DRO) getImageFileName (DRO) getImageHostName (DRO) getImagePathName (DRO) getImagePathName (DRO) getImageAthName (DRO) getImageAthName (DRO) getImageAthName (DRO) getImageAthName (DRO) getImageAthName (DRO) getImageAthName (DRO) getImageAthName (DRO)</pre>

# **Reference Index**

<u>change</u>	(administration) Changes the location attributes of the referenced database.
<u>changePartition</u>	(FTO) Changes the autonomous partition that contains the referenced database.
close	Internal use only. Objectivity/DB closes a database automatically when the transaction that opened it commits or aborts.
<u>containedIn</u>	Finds the federated database that contains the referenced database.
<u>containingPartition</u>	(FTO) Finds the autonomous partition that contains the referenced database.
<u>contains</u>	(FTO) Initializes an object iterator to find all containers in the referenced database.
<u>convertObjects</u>	Performs on-demand object conversion on any affected objects in the referenced database.
<u>deleteImage</u>	(DRO) Deletes the specified autonomous partition's image of the referenced database.
<u>exist</u>	Tests whether the specified database exists in the federated database; if successful, sets this object reference or handle to reference the database.
<u>fileName</u>	Gets the fully qualified filename of the referenced database.
getAllowNonQuorumRead	(DRO) Tests whether this application is allowed to read the referenced database when a quorum of images is not available.
getDefaultContObj	Finds the default container of the referenced database.
getImageFileName	(DRO) Gets the fully qualified filename of the specified image of the referenced database.
<u>getImageHostName</u>	(DRO) Gets the name of the data server host that contains the specified image of the referenced database.

<u>getImagePathName</u>	(DRO) Gets the pathname of the directory that contains the specified image of the referenced database.
getImageWeight	(DRO) Gets the weight of the specified image of the referenced database.
getTieBreaker	(DRO) Finds the tie-breaker partition for the referenced database.
<u>hasImageIn</u>	(DRO) Tests whether the specified autonomous partition contains an image of the referenced database.
hostName	Gets the network name of the data server host that contains the referenced database.
<u>isAvailable</u>	(FTO) Tests whether the current process can access the referenced database.
<u>isImageAvailable</u>	(DRO) Tests whether the current process can access the specified image of the referenced database.
<u>isNonQuorumRead</u>	(DRO) Tests whether the application is currently reading the referenced database without having a quorum of images.
<u>isReadOnly</u>	Tests whether the referenced database is a read-only database.
isReplicated	(DRO) Tests whether the referenced database has more than one image.
<u>isValid</u>	Tests whether this object reference or handle is valid—that is, whether it references an existing database.
lock	Explicitly locks the referenced database.
name	Gets the system name of the referenced database.
negotiateQuorum	(DRO) Forces recalculation of the quorum for the referenced database.
<u>numContObjs</u>	Gets the number of containers in the referenced database.
numImages	(DRO) Gets the number of images of the referenced database.
<u>ooHandle(ooDBObj)</u>	Default constructor that constructs a null handle.

<u>ooHandle(ooDBObj)</u>	Constructs a handle that references the same database as the specified object reference or handle.
<u>ooRef(ooDBObj)</u>	Default constructor that constructs a null object reference.
<u>ooRef(ooDBObj)</u>	Constructs an object reference that references the same database as the specified object reference or handle.
open	Explicitly opens the referenced or specified database, preparing the database for the specified level of access.
openMode	Gets the current level of access to the referenced database.
operator=	Assignment operator; sets this object reference or handle to reference the same database as the specified object reference or handle.
partitionsContainingImage	(DRO) Initializes an object iterator to find all the autonomous partitions that contain an image of the referenced database.
pathName	Gets the pathname of the directory that contains the referenced database.
replicate	(DRO) Creates an image of the referenced database.
<u>setAllowNonQuorumRead</u>	(DRO) Specifies whether the application can read the referenced database when a quorum of images is not available.
<u>setImageWeight</u>	(DRO) Sets the weight of the specified image of the referenced database.
setReadOnly	Sets the access status of the referenced database so that it is either read-only or read-write.
<u>setTieBreaker</u>	(DRO) Sets the tie-breaker autonomous partition for the referenced database.
tidy	(administration) Consolidates fragmented storage space in the referenced database; used only in custom administration tools.
typeN	Gets the type number of the database class ooDBObj.

<u>typeName</u>

Gets the name of the database class ooDBObj.

<u>update</u>

Opens the referenced database for update access.

# Constructors

## ooHandle(ooDBObj)

Default constructor that constructs a null handle.

```
ooHandle(ooDBObj)();
```

## ooHandle(ooDBObj)

Constructs a handle that references the same database as the specified object reference or handle.

```
ooHandle(ooDBObj)(
     const ooRefHandle(ooDBObj) &existing);
```

Parameters existing

Object reference or handle to an existing database.

# ooRef(ooDBObj)

Default constructor that constructs a null object reference.

```
ooRef(ooDBObj)();
```

# ooRef(ooDBObj)

Constructs an object reference that references the same database as the specified object reference or handle.

Parameters existing

Object reference or handle to an existing database.

#### operator=

Assignment operator; sets this object reference or handle to reference the same database as the specified object reference or handle.

```
ooRefHandle(ooDBObj) &operator=(
    const ooRefHandle(ooDBObj) &existing);
```

Parameters existing

Object reference or handle to an existing database.

Returns This object reference or handle.

# **Member Functions**

## change

(administration) Changes the location attributes of the referenced database.

ooStatus cl	hange	(	
const	char	*dbSysName = 0,	
const	char	*dbHostName,	
const	char	*dbPathName,	
FILE*	outpi	<i>utFile</i> = stdout)	const;

#### Parameters dbSysName

New system name of the database. *This feature is currently not implemented.* Always pass the value of zero for this parameter.

#### dbHostName

Name of the data server host on which the database is to reside. You must specify this parameter, even when relocating the database file on the same host.

#### dbPathName

New pathname for the database file on *dbHostName*. You must include the database's filename as the last component of the pathname.

#### outputFile

Pointer to a transcript file in which to report the original and changed database attributes. The default is standard output.

Returns	oocSuccess if successful; otherwise oocError.
Discussion	Together, the parameters <i>dbHostName</i> and <i>dbPathName</i> determine the network address of the database file.
	This member function updates the database's host and pathname in the federated database catalog, but does not actually move or rename the database file on your file system. You must use appropriate operating system commands to actually move the database file. The change member function is similar to the oochangedb tool, except that the tool also moves or renames the database file in the file system (see the Objectivity/DB administration book).
	This member function requests an exclusive update lock on the referenced database. If it cannot obtain the lock, the member function both returns occerror and signals an error.
	If the referenced database is read-only, you must use the <u>setReadOnly</u> member function to change it back to read-write before you can change its attributes.
	( <i>DRO</i> ) An error is signalled if the database is replicated (has more than one image).
Example	This code sets the database's host name to myHost and pathname to /mnt/john/design/adder.ecad.DB.
	ooHandle(ooDBObj) dbH; // Set dbH to reference a database dbH.change(0, "myHost", "/mnt/john/design/adder.ecad.DB");

# changePartition

(FTO) Changes the autonomous partition that contains the referenced database.

	<pre>ooStatus changePartition(     const ooRefHandle(ooAPObj) &amp;newPartition) const;</pre>
Parameters	newPartition Object reference or handle to the autonomous partition that is to contain the database.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	This member function assigns the referenced database to the specified autonomous partition, removing the database from the current containing partition.
	This member function updates the containment information in the catalogs of the autonomous partitions. If you also want to change the physical location of a database file, you must do so using the oochangedb tool.

You may not combine this operation with other updates on the database in the same transaction. You should consider executing changePartition in a separate transaction.

If the referenced database is read-only, you must use the <u>setReadOnly</u> member function to change it back to read-write before you can change its partition.

(*DRO*) If multiple images of the database exist, an error is signalled, and the current containing partition is left unchanged.

#### close

Internal use only. Objectivity/DB closes a database automatically when the transaction that opened it commits or aborts.

```
ooStatus close() const;
```

#### containedIn

Finds the federated database that contains the referenced database.

- 1. ooHandle(ooFDObj) containedIn() const;

Parameters returnedFD

Object reference or handle to be set to the federated database.

Returns Object reference or handle to the federated database.

Discussion When called without a *returnedFD* parameter, containedIn allocates a new federated-database handle and returns it. Otherwise, containedIn returns the object reference or handle that was passed to it.

#### containingPartition

(FTO) Finds the autonomous partition that contains the referenced database.

- 1. ooHandle(ooAPObj) containingPartition() const;
- 3. ooHandle(ooAPObj) &containingPartition( ooHandle(ooAPObj) &returnedAP) const;

Parameters	<i>returnedAP</i> Object reference or handle to be set to the containing autonomous partition.
Returns	An object reference or handle to the autonomous partition that contains the referenced database.
	( <i>DRO</i> ) If multiple images of the database exist, a null handle is returned and an error is signalled.
Discussion	When called without a parameter, containingPartition allocates a new autonomous-partition handle and returns it. Otherwise, containingPartition returns the object reference or handle that was passed to it.
contains	
	(FTO) Initializes an object iterator to find all containers in the referenced database.
	<pre>ooStatus contains( ooItr(ooContObj) &amp;returnedConts, const ooMode openMode = oocNoOpen, const ooContainsFilter whichConts = oocAllObjs);</pre>
Parameters	returnedConts Object iterator for finding the database's containers
	-
	openMode Intended level of access to the containers found by the iterator's next member function:
	<ul> <li>oocNoOpen (the default) causes next to set the iterator to the next container without opening it.</li> </ul>
	<ul> <li>oocRead causes next to open the next container for read.</li> </ul>
	oocUpdate causes next to open the next container for update.
	whichConts
	Filters the containers to be found by the iterator:
	<ul> <li>Specify oocAllObjs (the default) to initialize the iterator with all containers contained in the database.</li> </ul>
	<ul> <li>(FTO) Specify oocNotTransferred to initialize the iterator with containers that have the same controlling autonomous partition as the database.</li> </ul>
	<ul> <li>(FTO) Specify occTransferred to initialize the iterator with those containers that have a different controlling autonomous partition than the database.</li> </ul>
Returns	oocSuccess if successful; otherwise oocError.

#### convertObjects

Performs on-demand object conversion on any affected objects in the referenced database.

ooStatus convertObjects();

- Returns oocSuccess if successful, or oocError if the federated database is opened only for read.
- Discussion Object conversion is the process of making existing persistent objects consistent with class definition changes introduced by schema evolution. Certain schema evolution operations affect how instances of a class should be laid out in storage. After you perform such operations, existing objects of the changed classes are rendered out-of-date until they are converted to their new representations.

In general, you can allow each affected object to be converted automatically the first time it is accessed after schema evolution, potentially distributing the performance impact of conversion across many transactions. Alternatively, you can concentrate the performance impact of conversion into fewer transactions by converting all the affected objects in a container, a database, or a federated database *on demand*. You use this member function in an update transaction to convert the affected objects in a database on demand. This member function has no effect if the affected objects in the database have already been converted.

**Note:** On-demand object conversion cannot be used for schema operations that require an upgrade application; see ooTrans::upgrade.

The convertObjects member function automatically drops any index that is invalidated by a schema evolution change. Specifically, if you changed the type or deleted a data member that is a key field in a key description, the corresponding indexes are dropped.

If the referenced database is read-only, you must use the <u>setReadOnly</u> member function to change it back to read-write before you can convert objects in it.

See also Chapter 19, "Object Conversion," in the Objectivity/C++ programmer's guide

#### deletelmage

(*DRO*) Deletes the specified autonomous partition's image of the referenced database.

```
ooStatus deleteImage(
    const ooRefHandle(ooAPObj) &partition,
    ooBoolean deleteDBifLast = oocFalse) const;
```

Parameters	partition
	Object reference or handle to the autonomous partition containing the database image.
	deleteDBifLast
	Specifies whether to delete the database image even if no other images remain:
	<ul> <li>Specify occTrue to delete the last image and the database itself.</li> </ul>
	<ul> <li>Specify occFalse (the default) to preserve the last image and signal an error.</li> </ul>
Returns	oocSuccess if successful; otherwise oocError.
Discussion	An error is signalled if either of the following are true:
	• The referenced database has no image in the specified autonomous partition.
	■ If the referenced database is read-only. You must use the <u>setReadOnly</u> member function to change the database back to read-write before you can delete an individual image.
exist	
	Tests whether the specified database exists in the federated database; if successful, sets this object reference or handle to reference the database.
	<pre>ooBoolean exist( const ooHandle(ooFDObj) &amp;fdH, const char *dbSysName, const ooMode openMode = oocNoOpen);</pre>
Parameters	fdH
	Handle to the federated database.
	dbSysName
	System name of the database to be found.
	openMode
	Intended level of access to the database, if it exists:
	<ul> <li>Specify oocNoOpen (the default), to set this object reference or handle to the database without opening it.</li> </ul>
	<ul> <li>Specify oocRead to open the database for read.</li> </ul>
	<ul> <li>Specify oocUpdate to open the database for update.</li> </ul>
Returns	oocTrue if the specified database exists, or oocFalse if the database does not exist or if it is not accessible.

# Discussion If the specified database exists, this object reference or handle is set to reference it; otherwise, this object reference or handle is set to null.

If you specifically want to test for existence, you use the *openMode* parameter's default value (*oocNoOpen*). Otherwise, a return value of *oocFalse* could mean either that the database doesn't exist, or that it does exist, but cannot be opened.

### fileName

Gets the fully qualified filename of the referenced database.

char \*fileName() const;

- Returns Pointer to a string containing the filename.
- Discussion The returned filename includes the directory pathname as well as the simple name of the database file. Use <u>pathName</u> to obtain just the directory pathname.

The returned string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.

(DRO) An error is signalled if multiple images of the database exist. Use getImageFileName instead.

Example This example prints the filename of the database referenced by dbH. On UNIX, the output of this example might be: /mnt/john/design/testDb.testFd.DB.

ooHandle(ooFDObj) fdH = fdH.open("testFd", oocRead); ooHandle(ooDBObj) dbH = dbH.open(fdH, "testDb", oocRead);

printf("filename: %s\n", dbH.fileName());

### getAllowNonQuorumRead

(*DRO*) Tests whether this application is allowed to read the referenced database when a quorum of images is not available.

ooBoolean getAllowNonQuorumRead() const;

Returns oocTrue if nonquorum reads are allowed; otherwise oocFalse.

See also <u>isNonQuorumRead</u> <u>setAllowNonQuorumRead</u>

## getDefaultContObj

#### Finds the default container of the referenced database.

	<pre>1. ooHandle(ooContObj) getDefaultContObj(</pre>
	<pre>2. ooRef(ooContObj) &amp;getDefaultContObj(</pre>
	<pre>3. ooHandle(ooContObj) &amp;getDefaultContObj(</pre>
Parameters	<i>returnedCont</i> Object reference or handle to be set to the default container.
	openMode
	Intended level of access to the default container:
	<ul> <li>Specify occRead (the default) to open the default container for read.</li> </ul>
	<ul> <li>Specify occUpdate to open the default container for update.</li> </ul>
	<ul> <li>Specify oocNoOpen to set the object reference or handle reference to the default container without opening it.</li> </ul>
Returns	Object reference or handle to the default container of the database.
Discussion	When called without a <i>returnedCont</i> parameter, this member function allocates a new container handle and returns it. Otherwise, this member function returns the object reference or handle that was passed to it.

### getImageFileName

(*DRO*) Gets the fully qualified filename of the specified image of the referenced database.

Parameters partition

Object reference or handle to the autonomous partition that contains the desired database image.

Returns Pointer to a string containing the database image's filename.

Discussion The returned filename includes the directory pathname as well as the simple name of the database image file. Use <u>getImagePathName</u> to obtain just the directory pathname.

The returned string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.

An error is signalled if the specified autonomous partition does not contain an image of the referenced database.

#### getImageHostName

(*DRO*) Gets the name of the data server host that contains the specified image of the referenced database.

Parameters partition

Object reference or handle to the autonomous partition that contains the desired database image.

- Returns Pointer to a string containing the network name of the data server host where the database image file is located.
- Discussion The string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.

An error is signalled if the specified autonomous partition does not contain an image of the referenced database.

### getImagePathName

(*DRO*) Gets the pathname of the directory that contains the specified image of the referenced database.

Parameters partition

Object reference or handle to the autonomous partition that contains the desired database image.

Returns Pointer to a string containing the directory pathname.

Discussion The returned pathname does not include the simple name of the database file; use getImageFileName to obtain a path name that includes the filename.

The string is statically allocated by the member function and is overwritten with each invocation. You should copy the string if you wish to use it later in your application.

An error is signalled if the specified autonomous partition does not contain an image of the referenced database.

## getImageWeight

(DRO) Gets the weight of the specified image of the referenced database.

Parameters partition

Object reference or handle to the autonomous partition that contains the desired database image.

Returns The weight of the database image; returns 0 if the referenced database is not replicated in the specified autonomous partition.

#### getTieBreaker

(DRO) Finds the tie-breaker partition for the referenced database.

ooHandle(ooAPObj) getTieBreaker() const;

- Returns Handle to the tie-breaker partition if one exists; otherwise, a null handle.
- Discussion This member function allocates an autonomous-partition handle and returns it.

#### hasImageIn

(*DRO*) Tests whether the specified autonomous partition contains an image of the referenced database.

- Parameters *partition* Object reference or handle to the autonomous partition to be searched.
- Returns oocTrue if the partition contains an image of the referenced database; otherwise oocFalse.

### hostName

Gets the network name of the data server host that contains the referenced database. char \*hostName() const; Pointer to a string containing the network name of the data server host. Returns Discussion The returned string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application. (DRO) An error is signalled if multiple images of the database exist. Use getImageHostName instead. Example This example prints the host name for the database referenced by dbH. The output of this example might be: myMachine. ooHandle(ooFDObi) fdH; ooHandle(ooDBObj) dbH; fdH.open("testFd", oocRead); dbH.open(fdH, "testDb", oocRead); printf("hostname: %s\n", dbH.hostName());

#### isAvailable

(FTO) Tests whether the current process can access the referenced database.

ooBoolean isAvailable() const;

Returns oocTrue if the current process can access the database; otherwise oocFalse.

(*FTO*) If the database is not replicated, it is accessible if the process can access its containing partition.

(*DRO*) If the database is replicated, it is accessible if the process can access a quorum of its images.

#### isImageAvailable

(*DRO*) Tests whether the current process can access the specified image of the referenced database.

Parameters	partition	
	Object reference or handle to the autonomous partition containing the desired image.	
Returns	oocTrue if the image in the specified autonomous partition is available; otherwise oocFalse.	

## isNonQuorumRead

	( <i>DRO</i> ) Tests whether the application is currently reading the referenced database without having a quorum of images.
	ooBoolean isNonQuorumRead() const;
Returns	oocTrue if the application is reading the referenced database without having a quorum of images; otherwise oocFalse.
Discussion	Nonquorum reading of a database is possible only after an application has called the setAllowNonQuorumRead member function for the database during the current transaction.
WARNING	If this member function returns oocTrue, your application may be reading stale data from the database.
See also	getAllowNonQuorumRead setAllowNonQuorumRead
isReadOnl	У
	Tests whether the referenced database is a read-only database.
	ooBoolean isReadOnly() const;
Returns	oocTrue if this object reference or handle references a read-only database; otherwise oocFalse.
Discussion	A read-only database can be opened only for read; any attempt to implicitly or explicitly open the database for update will fail as if there were a lock conflict.
See also	setReadOnly

#### isReplicated

ooBoolean isReplicated() const;

Returns oocTrue if the referenced database is replicated; otherwise oocFalse.

## isValid

Tests whether this object reference or handle is valid—that is, whether it references an existing database.

ooBoolean isValid() const;

- Returns occTrue if this object reference or handle references an existing database; occFalse if this object reference or handle is null or has a stale identifier, or if the application cannot obtain a read lock on the database to be checked.
- Discussion You can use isValid to determine whether it is safe to use an object reference or handle that was set in a previous transaction. Such an object reference or handle still retains its reference to a database; however, between transactions, that reference may have become invalid (for example, because another process has deleted the database).
  - **NOTE** isValid checks only for the existence of a database with a particular identifier, but has no way of knowing whether it is the same database. It is possible, although very unlikely, for another process to have deleted the original database and created a new one with the same identifier.

If your purpose is simply to test whether an object reference or handle has been initialized, it is more efficient to use <u>isNull</u>, which performs its test entirely in memory without having to access files on disk.

### lock

Explicitly locks the referenced database.

ooStatus lock(const ooLockMode lockMode) const;

#### Parameters lockMode

Type of lock to request:

- Specify oocLockRead to request a read lock.
- Specify oocLockUpdate to request an update lock.

- Returns oocSuccess if the requested lock is obtained; otherwise oocError.
- Discussion Objectivity/DB operations request and obtain locks implicitly as they are needed. You use this member function to obtain a lock explicitly when you want to reserve access to a database in advance.

Explicitly locking a database essentially limits the level of concurrent access to its containers, allowing them to be read, but guaranteeing they will not change while the lock is held. That is, explicitly locking a database for either read or update:

- Prevents any other transaction from concurrently opening the database or a container in it for update.
- Allows any other transaction to concurrently open the database or a container in it for read.

A database cannot be locked if it, or any container in it, is already opened for update.

**NOTE** Holding an update lock on a database does *not* guarantee update access to the database's individual containers. For example, if transaction T1 locks a database for update, and then a standard (non-MROW) transaction T2 opens a container in the database for read, T1 cannot concurrently get an update lock on the container being read.

When a database is locked for read, other transactions can concurrently lock it for read. When a database is locked for update, MROW transactions can concurrently lock it for read, but standard transactions cannot. Two transactions cannot lock the same database for update.

#### name

Gets the system name of the referenced database.

char \*name() const;

- Returns Pointer to a string containing the system name.
- Discussion The string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.

#### negotiateQuorum

(DRO) Forces recalculation of the quorum for the referenced database.

ooStatus negotiateQuorum(ooMode openMode);

Parameters openMode

Intended level of access to this database:

- Specify occRead to open and implicitly lock the database for read.
- Specify occupdate to open and implicitly lock the database for update (read and write).

Returns oocSuccess if successful; otherwise oocError.

Discussion If the referenced database has any images in previously unavailable partitions, this member function adds those images back into the quorum for this database. The reacquired images are resynchronized—that is, they are updated to be consistent with the images in the quorum.

If your application includes a reinitialization procedure that is executed when an autonomous partition that was down is brought back up, that procedure should call this member function for every database with an image in the restored autonomous partition.

This member function returns OOCError and performs no action if any of the following is true:

- If the application does not have access to a quorum of images for the referenced database.
- If any partition containing an image for this database is still unavailable—for example, if the partition's lock server or AMS has not been restarted yet.

If the referenced database is not replicated, this member function returns oocSuccess and performs no other action.

You may not recalculate a quorum if nonquorum reading has been enabled for the referenced database (see <u>setAllowNonQuorumRead</u>). To recalculate a quorum in this case, you must end the current transaction (to disable nonquorum reading), and then call negotiateQuorum from within a new transaction.

#### numContObjs

Gets the number of containers in the referenced database.

unsigned long numContObjs() const;

Returns Number of containers.

#### numImages

(DRO) Gets the number of images of the referenced database.

uint32 numImages() const;

Returns Number of images.

Discussion For any existing database, there is always at least one image.

#### open

Explicitly opens the referenced or specified database, preparing the database for the specified level of access.

	<pre>1. ooStatus open(     const ooMode openMode = oocRead);</pre>	
	<pre>2. ooStatus open( const ooHandle(ooFDObj) &amp;fdH, const char *dbSysName, const ooMode openMode = oocRead);</pre>	
Parameters	openMode	
	Intended level of access to the opened database:	
	<ul> <li>Specify occreat (the default) to open and implicitly lock the database for read.</li> </ul>	
	<ul> <li>Specify occupdate to open and implicitly lock the database for update (read and write).</li> </ul>	
	fdH	
	Handle to the currently open federated database.	
	dbSysName	
	System name of the database to open.	
Returns	oocSuccess if successful; otherwise oocError.	
Discussion	Variant 1 assumes that this object reference or handle already references a database, and opens the referenced database.	
	Variant 2 finds and explicitly opens the database with the specified system name and sets this object reference or handle to reference it. An error is signalled if no database exists with the specified system name or if the database file cannot be found or accessed. This variant is especially useful when you want to use the database as an entry point into your data. For example, you might find and open a database so you can iterate over the containers in it.	

Opening a database makes it available to an application by locating and opening the database file, provided that appropriate access permissions are set on it.

It is normally not necessary to open databases explicitly because they are usually opened automatically by operations that access them or their contents. For example, once you have a reference to a database, creating a container in the database automatically opens it for update. In general, you open a referenced database explicitly only when you want to guarantee access to the database in advance—for example, before starting a complex operation.

Any number of transactions can concurrently open the same database in any mode. However, a database cannot be opened for update if another transaction already has a read or update lock on it.

You must be in an update transaction to open a database for update. If necessary, you can promote a read transaction to an update transaction by promoting the open mode of the federated database.

See also <u>update</u>

#### openMode

Gets the current level of access to the referenced database.

ooMode openMode() const;

Returns

One of the following constants:

- oocNoOpen—the database is not open in this transaction.
- oocRead—the database is open for read in this transaction.
- oocUpdate—the database is open for update in this transaction.

#### partitionsContainingImage

(*DRO*) Initializes an object iterator to find all the autonomous partitions that contain an image of the referenced database.

#### Parameters apI

Object iterator for finding the containing autonomous partitions.

Returns oocSuccess if successful; otherwise oocError.

Discussion If the referenced database is not replicated in any other autonomous partition, the iterator finds the partition in which the database resides. The iterator's next member function finds each autonomous partition without opening or locking it.

#### pathName

Gets the pathname of the directory that contains the referenced database.

char \*pathName() const;

Returns Pointer to a string containing the pathname of the directory.

Discussion The returned pathname does not include the simple name of the database file; use  $\underline{fileName}$  to obtain a path name that includes the filename.

The string is statically allocated by the member function and is overwritten with each invocation. You should copy the string if you wish to use it later in your application.

(DRO) An error is signalled if multiple images of the database exist. Use getImagePathName instead.

Example This example prints the directory pathname of the database referenced by dbH. On UNIX, the output of this example might be: /mnt/john/design.

> ooHandle(ooFDObj) fdH; ooHandle(ooDBObj) dbH;

fdH.open("testFd", oocRead); dbH.open(fdH, "testDb", oocRead);

printf("pathname: %s\n", dbH.pathName());

### replicate

(DRO) Creates an image of the referenced database.

```
ooStatus replicate(
    const ooRefHandle(ooAPObj) &partition,
    const char *hostName = 0,
    const char *pathName = 0,
    uint32 weight = 1) const;
```

Parameters partition

Object reference or handle to the autonomous partition in which to create the new database image. This partition may not already contain an image of this database.

#### hostName

Name of the data server host on which to create the new database image. If this parameter is omitted, the host will be the same as that hosting the partition's system database file. The host machine must be running AMS.

#### pathName

Fully qualified pathname (including the filename) of the new database image. The format of the pathname must follow the naming conventions of the specified host. If this parameter is omitted, the image's file is created with a name based on this database's system name and is placed in the same directory as the partition's system database file.

#### weight

Weight of the new database image. weight must be 1 or greater.

Returns oocSuccess if successful; otherwise oocError.

Discussion If neither *hostName* or *pathName* is specified, the database image is created in the same directory as the autonomous partition's system database file, and the image's file name is generated from the database's system name.

An error is signalled if any of the following are true:

- If the referenced database already has an image in the specified autonomous partition.
- If AMS is not running on the data server hosts where the original and the new database images reside. For information about AMS, see the Objectivity/DB administration book.
- If the referenced database has been made read-only with the <u>setReadOnly</u> member function. You must change the database back to read-write before you can add a new image.

#### setAllowNonQuorumRead

(*DRO*) Specifies whether the application can read the referenced database when a quorum of images is not available.

Parameters value

Specify occTrue (the default) to allow this application to read the referenced database even if a quorum is not available; specify occFalse to prevent this application from reading the database when a quorum is not available.

Returns oocSuccess if successful; otherwise oocError.

Discussion	By default, nonquorum reading is disabled for all databases. You must call this member function explicitly to permit nonquorum reading for a particular database during a particular transaction. Nonquorum reading is automatically disabled at the end of the transaction in which this member function is called.
Warning	If you enable nonquorum reading, your application may read stale data from the database.
See also	<u>getAllowNonQuorumRead</u> isNonQuorumRead

### setImageWeight

(DRO) Sets the weight of the specified image of the referenced database.

	<pre>ooStatus setImageWeight( const ooRefHandle(ooAPObj) &amp;partition, uint32 weight) const;</pre>
Parameters	<pre>partition Object reference or handle to the autonomous partition containing the desired image. weight Weight to be set for the database image. weight must be 1 or greater.</pre>
Returns	oocSuccess if successful; otherwise oocError.
Discussion	<ul> <li>An error is signalled if either of the following are true:</li> <li>The referenced database is not replicated in the specified autonomous partition.</li> <li>If the referenced database is read-only. You must use the <u>setReadOnly</u> member function to change the database back to read-write before you can set an image's weight.</li> </ul>
setReadOn	ly

Sets the access status of the referenced database so that it is either read-only or read-write.

ooStatus setReadOnly(ooBoolean value) const;

Returns oocSuccess if successful; otherwise oocError.

Parameters	value
	Specify oocTrue to mark the database as read-only; specify oocFalse to change the database back to read-write.
Discussion	A read-only database can be opened only for read. Any attempt to implicitly or explicitly open the database for update will fail as if there were a lock conflict. Using a read-only database can improve performance by allowing the application to grant read locks and deny update locks without consulting the lock server.
	You can change a read-only database back to read-write only if no application or tool is currently reading either that database or any other read-only database in the same federated database. You must change the database back to read-write before you can perform any other operation on it.
	( <i>DRO</i> ) If one image of a database is made read-only, all images are automatically made read-only. Similarly, if one image is changed back to read-write, all images are changed back to read-write.
See also	isReadOnly

#### setTieBreaker

(DRO) Sets the tie-breaker autonomous partition for the referenced database.

Parameters	partition
	Object reference or handle to the tie-breaker autonomous partition. This partition may not already contain an image of this database. If <i>partition</i> is 0, any existing tie-breaker is eliminated.
Returns	oocSuccess if successful; otherwise oocError.

### tidy

(*administration*) Consolidates fragmented storage space in the referenced database; used only in custom administration tools.

```
ooStatus tidy(
    FILE* outputFile = stdout,
    const char *hostName = 0,
    const char *pathName = 0) const;
```

Parameters	outputFile Pointer to the transcript file in which to write the generated report. The
	default is standard output.
	hostName
	Name of the data server nost on which to create the temporary file.
	PathName Pathname of the directory in which to create the temporary file
	Faumanie of the directory in which to create the temporary me.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	This member function creates a temporary database file during execution and therefore requires free disk space equal to the size of the database file being tidied. The temporary file is created in the directory that contains the database file unless you specify <i>hostName</i> and <i>pathName</i> .
	If either <i>hostName</i> or <i>pathName</i> is zero or empty (""), the temporary file is created in the default directory.
	You should call tidy in a single-purpose update transaction. That is, you must not manipulate any database, container, or basic object before calling tidy in the same transaction, and you must commit the transaction immediately after tidy completes. This is because compacting and relocating physical storage renders the database inconsistent with any system data that was cached during the transaction, and committing the transaction discards the obsolete cached data.
	You must not abort the transaction after calling this member function.
	This member function performs the same function as the $ootidy$ tool (see the Objectivity/DB administration book).
typeN	
	Gets the type number of the database class ooDBObj.
	ooTypeNumber typeN() const;
Returns	Type number of the database class ooDBObj.
typeName	
	Gets the name of the database class ooDBObj.
	char *typeName() const;
Returns	The string "ooDBObj".
Discussion	The returned string must be treated as read-only.

## update

Opens the referenced database for update access.

ooStatus update();

- Returns oocSuccess if successful; otherwise oocError.
- Discussion This member function is equivalent to calling <u>open(oocUpdate)</u>.
- See also <u>open</u>

# ooRefHandle(ooFDObj) Classes

• ooHandle(ooFDObj), which represents a *handle* to the federated database.

The two classes ooRef(ooFDObj) and ooHandle(ooFDObj) are documented together because they define the same set of member functions. These member functions provide the primary interface for operating on federated databases (instances of ooFDObj).

See:

- "Reference Summary" on page 575 for an overview of member functions
- "Reference Index" on page 576 for a list of member functions

# **About Federated-Database Handles and References**

An application that is connected to a particular federated database can work with that federated database indirectly through one or more handles or object references—that is, through instances of *ooRefHandle*(ooFDObj) that are set to reference the federated database being accessed. Such handles and object references both identify the federated database and provide the complete public interface for operating on it.

Handles and object references to federated databases do not support indirect member access; the *ooRefHandle*(ooFDObj) classes provide no indirect

member-access operator (->), and the ooFDObj class defines no public member functions other than a constructor and operator new.

You can work with a federated database through either a handle or an object reference—the choice is arbitrary, except as described in "Structure and Behavior" on page 572. Most applications use handles rather than object references.

#### Interface

The *ooRefHandle*(*ooFDObj*) classes provide the primary interface for operating on a federated database. Part of this interface consists of member functions defined by *ooRefHandle*(*ooFDObj*) for specialized operations such as tidying. The other part of this interface consists of member functions defined by the *ooRefHandle*(*ooObj*) base classes for more general Objectivity/DB operations, such as opening, locking, printing object identifiers, and so on. These member functions are either inherited by the *ooRefHandle*(*ooFDObj*) classes or redefined wherever type-specific parameters or behavior are required.

Some of the member functions defined by the *ooRefHandle*(ooObj) base classes are *not* available to instances of the *ooRefHandle*(ooFDObj) classes. These include member functions for moving, copying, versioning, scope-naming, and member-access operations, which apply only to basic objects or persistent objects. The disallowed member functions and operators are redefined as private members of the *ooRefHandle*(ooFDObj) classes.

#### **Structure and Behavior**

Handles and object references to a federated database are essentially wrappers for the federated database's identifier. Thus, when an application opens a federated database through a handle, the handle stores that federated database's identifier. The handle uses the identifier to find the instance of ooFDObj that represents the federated database in memory.

In general, object references are optimized for implementing links among related persistent objects, while handles are optimized for memory management and member-access. When a federated database is referenced, however, these optimizations are largely irrelevant, because federated databases (unlike persistent objects):

- Cannot be linked (for example, through associations).
- Are not subject to memory management (they have no attributes for persistent data and are therefore not manipulated through pointers).
- Have no accessible members.

One significant exception is that a handle to a federated database, like any handle, contains cache-related state that is associated with the Objectivity context

in which it was created. Therefore, only object references (but not handles) can be passed between Objectivity contexts. Otherwise, a handle to a federated database and an object reference to a federated database are functionally equivalent.

# **Working With Federated-Database Handles**

**NOTE** For simplicity, this section describes how to work with handles. Except where noted, the same information applies to object references.

An application normally creates a handle as a local variable on the stack, rather than allocating it on the heap. A handle should not be declared as const, because its internal state may be changed when a federated database is accessed through it. (Object references may be declared as const). Applications should not create subclasses of the *ooRefHandle*(ooFDObj) classes.

A new handle is normally *null*—that is, it contains the value 0 instead of a federated database identifier. The application can set the handle to reference a federated database in any of the following ways:

- By connecting to the federated database with the handle's <u>exist</u> or <u>open</u> member function. (Within a particular application, the *same* federated database must be specified in all such operations.)
- By passing the handle to a member function that sets it, such as the <u>containedIn</u> member function of a database handle, which finds the federated database from a database in it.
- By assignment or initialization from another handle or object reference.

A handle continues to reference the federated database unless it is set to null; it cannot be set to any other federated database during the application. Multiple handles and object references can be set to the same federated database.

An application operates on the federated database by calling member functions on a handle that references it. To call a member function of a handle, you use the direct member-access operator (.). For example, fdH.name calls the <u>dumpCatalog</u> member function of the handle fdH.

As indicated in "Reference Summary" on page 575, an application can get and change the federated database's attributes, find its databases, tidy its disk space, and so on.

Although most of the handle's member functions operate on the referenced federated database, some functions operate on the handle itself. For example, you use:

- The assignment operator <u>operator=</u> to set a handle from another handle or from an object reference.
- The inherited comparison operators <u>operator==</u> and <u>operator!=</u> to test whether a handle is equal to another handle or object reference.
- The inherited member function <u>isNull</u> to test whether a handle is null. (Alternatively, you can use the overloaded <u>operator==</u> to compare a handle to 0.)

A handle preserves its reference to a federated database across transaction boundaries, provided that the handle does not go out of scope and is not set to null as the result of an abort operation.

Objectivity/C++ functions that require a federated database as input normally obtain the federated database through a parameter of type const ooHandle(ooFDObj) &. If a function manipulates other types of Objectivity/DB objects as well, the parameter type may be specified as const ooHandle(ooObj) &. You can pass a federated-database handle to a parameter of this type, because ooHandle(ooFDObj) is derived from ooHandle(ooObj). In practice, however, relatively few functions accept a federated-database handle where a more general handle is requested; these include the functions that require a storage object for scanning or as a name scope. In most cases, a function that requests a general-purpose handle operates only on basic objects or persistent objects, and signals an error if you specify a federated-database handle.

An application must call the <u>open</u> member function on a federated-database handle to explicitly open the federated database at the beginning of every transaction; doing so verifies the connection and designates the transaction as either a read or update transaction. The handle is automatically closed when it is destroyed (for example, by going out of scope). However, closing one or more handles to the federated database has no particular effect on the federated database, which remains open until the transaction commits or aborts.

# **Reference Summary**

The following table summarizes all the member functions that are available to instances of ooRefHandle(ooFDObj). Member functions indicated as *(inherited)* documented with the <u>ooRefHandle(ooObj)</u> classes (page 593).

Creating a Handle or Object Reference	<u>ooHandle(ooFDObj)</u> ooRef(ooFDObj)
Setting the Handle or Object Reference	<u>operator=</u> <u>open</u> <u>exist</u>
Comparing Handles and Object References	<u>operator==</u> (inherited) <u>operator!=</u> (inherited)
Opening, Closing, and Locking the Federated Database	<u>open</u> <u>update</u> <u>openMode</u> <u>lock</u> <u>close</u>
Modifying the Federated Database	update change
Getting Information About the Federated Database	name number pageSize lockServerName dumpCatalog typeN typeName print (inherited) sprint (inherited)
Testing the Federated Database	<u>exist</u>
Testing the Handle or Object Reference	<u>is null (inherited)</u> <u>isNull (inherited)</u> <u>isValid</u> <u>operator int (inherited)</u> <u>operator ooObj* (inherited)</u>
Finding Objects	<u>exist</u> <u>open</u> bootAP
	contains

Converting Objects	<u>convertObjects</u> <u>setConversion</u> upgradeObjects
Working With Databases	<u>contains</u> tidy
Working With Autonomous Partitions ( <i>FTO</i> )	<u>bootAP</u> <u>contains</u>

# **Reference Index**

bootAP	(FTO) Finds the boot autonomous partition for the current application.	
<u>change</u>	(administration) Changes the attributes of the referenced federated database; used only in custom administration tools.	
close	Closes any open Objectivity/DB objects in the referenced federated database.	
<u>contains</u>	Initializes an object iterator to find all the databases or all the autonomous partitions in the referenced federated database.	
<u>convertObjects</u>	Performs on-demand object conversion on any affected objects in the referenced federated database.	
<u>dumpCatalog</u>	(administration) Prints out a list of the files associated with the referenced federated database; primarily used in administration tools.	
<u>exist</u>	Tests for the existence of the specified federated database; if successful, sets this object reference or handle to reference the federated database and optionally opens it.	
<u>isValid</u>	Checks whether this object reference or handle is valid—that is, whether it references the federated database.	
lock	Explicitly locks the referenced federated database.	
<u>lockServerName</u>	Gets the network name of the host running the lock server for the referenced federated database.	
name	Gets the system name of the referenced federated database.	
--------------------------	---	--
number	(administration) Gets the identifier of the referenced federated database.	
<u>ooHandle(ooFDObj)</u>	Default constructor that constructs a null handle.	
<u>ooHandle(ooFDObj)</u>	Constructs a handle that references the same federated database as the specified object reference or handle.	
<u>ooRef(ooFDObj)</u>	Default constructor that constructs a null object reference.	
<u>ooRef(ooFDObj)</u>	Constructs an object reference that references the same federated database as the specified object reference or handle.	
open	Opens the specified federated database for the specified level of access and sets this object reference or handle to reference the opened federated database.	
<u>operator=</u>	Assignment operator; sets this object reference or handle to reference the specified federated database.	
<u>openMode</u>	Gets the current level of access to the referenced federated database.	
pageSize	(administration) Gets the storage page size of the referenced federated database.	
setConversion	Registers a conversion function for objects of the specified class in the referenced federated database.	
tidy	(administration) Consolidates fragmented storage space in all the databases in the referenced federated database; used only in custom administration tools.	
typeN	Gets the type number of the federated-database class ooFDObj.	
typeName	Gets the name of the federated-database class $ooFDObj$ .	
update	Opens the referenced federated database for update access.	
<u>upgradeObjects</u>	Performs object conversion throughout the referenced federated database; used only in a special-purpose upgrade application.	

# **Constructors and Destructors**

# ooHandle(ooFDObj)

Default constructor that constructs a null handle.

```
ooHandle(ooFDObj)();
```

# ooHandle(ooFDObj)

Constructs a handle that references the same federated database as the specified object reference or handle.

```
ooHandle(ooFDObj)(
     const ooRefHandle(ooFDObj) &federation);
```

Parameters federation

Object reference or handle to the federated database.

# ooRef(ooFDObj)

Default constructor that constructs a null object reference.

```
ooRef(ooFDObj)();
```

# ooRef(ooFDObj)

Constructs an object reference that references the same federated database as the specified object reference or handle.

Parameters federation

Object reference or handle to the federated database.

# Operators

### operator=

Assignment operator; sets this object reference or handle to reference the specified federated database.

```
ooRefHandle(ooFDObj) &operator=(
    const ooRefHandle(ooFDObj) federation);
```

Parameters federation

Object reference or handle to the federated database.

Returns This object reference or handle.

# **Member Functions**

## bootAP

(FTO) Finds the boot autonomous partition for the current application.

	1. ooHandle(ooAPObj) bootAP() const;	
	<pre>2. ooRef(ooAPObj) &amp;bootAP(</pre>	
	<pre>3. ooHandle(ooAPObj) &amp;bootAP(</pre>	
Parameters	partition Object reference or handle to be set to the boot autonomous partition.	
Returns	Dbject reference or handle to the boot autonomous partition, if successful; otherwise, returns a null handle.	
Discussion	When called without a parameter, bootAP allocates a new autonomous-partition handle and returns it. Otherwise, bootAP returns the object reference or handle that was passed to it.	

### change

(*administration*) Changes the attributes of the referenced federated database; used only in custom administration tools.

obboardap offattige (	
const char * <i>bootFilePath</i> = 0,	,
<pre>const char *lockServer = 0,</pre>	
const uint32 <i>fdNumber</i> = 0,	
<pre>FILE* outputFile = stdout) co</pre>	onst;

#### Parameters bootFilePath

New pathname for the boot file of the federated database. You must include the filename as the last component of the pathname. Specify 0 (the default) to leave the boot file path unchanged.

#### lockServer

Name of the new lock server host (the host that runs the lock server for this federated database). Specify 0 (the default) to leave the lock server host unchanged.

fdNumber

New federated database identifier. Specify 0 (the default) to leave the identifier unchanged.

#### outputFile

Pointer to the transcript file in which to report the original and changed federated-database attributes. The default value is standard output.

- Returns oocSuccess if successful; otherwise oocError.
- Discussion You cannot change the system name of the federated database or its storage page size. If you specify a new boot file location, the updated boot file is written to the new location, but the old boot file remains; you must delete this file using appropriate operating system commands.

You use this member function in a special-purpose application that consists of a single update transaction. The application must exit immediately after the transaction commits. This is because the new state of the federated database is inconsistent with information cached by the executing application.

You should call  $\underline{ooNolock}$  so that the application will run in single-user mode. You then stop the lock server before running the application and restart it after the application completes.

WARNING	You must guarantee that no other transaction has access to the federated
	database or data corruption could result.

Example This example changes the lock server host for the federated database to moon.

```
ooTrans trans;
ooHandle(ooFDObj) fdH;
...
ooInit();
ooNoLock();
trans.start();
fdH.open("Documentation", oocUpdate);
fdH.change(0, "moon"); // change the lock server
trans.commit();
exit(0);
```

### close

Closes any open Objectivity/DB objects in the referenced federated database.

ooStatus close() const;

Returns oocSuccess if successful; otherwise oocError.

Discussion This member function closes all Objectivity/DB objects that have been opened in the current transaction. Closing persistent objects may be useful for purposes of memory management, because Objectivity/DB can swap the pages of closed objects out of the Objectivity/DB cache. Note, however, that locks on the closed objects are retained until the transaction commits or aborts.

> You do not need to explicitly reopen the federated database before performing further operations within the same transaction. However, you cannot open a different federated database within a process even if you have closed the first one.

### contains

Initializes an object iterator to find all the databases *or* all the autonomous partitions in the referenced federated database.

```
1. ooStatus contains(
        ooItr(ooDBObj) &containedDBs,
        const ooMode openMode = oocNoOpen) const;
```

#### Parameters containedDBs

Object iterator for finding the databases in the federated database.

(*FTO*) If the federated database has more than one autonomous partition, the iterator finds the databases contained in all the autonomous partitions.

containedAPs

(*FTO*) Object iterator for finding the autonomous partitions of the federated database.

#### openMode

Intended level of access to the databases or partitions found by the iterator's next member function:

- oocNoOpen (the default) causes next to set the iterator to the next found object without opening it.
- OOCRead causes next to open the next found object for read.
- oocUpdate causes next to open the next found object for update.

Returns oocSuccess if successful; otherwise oocError.

### convertObjects

Performs on-demand object conversion on any affected objects in the referenced federated database.

ooStatus convertObjects(ooBoolean purge\_schema = oocFalse);

```
Parameters purge_schema
```

Specifies whether to remove schema-evolution history after all objects are converted. By default, this history is not purged. **Warning:** Purging the schema may delete information that is required for distributing schema changes to deployed federated databases (see Discussion below). You should take a backup of the federated database before purging its schema.

- Returns oocSuccess if successful, or oocError if the federated database is opened for read-only access.
- Discussion Object conversion is the process of making existing persistent objects consistent with class definition changes introduced by schema evolution. Certain schema evolution operations affect how instances of a class should be laid out in storage. After you perform such operations, existing objects of the changed classes are rendered out-of-date until they are converted to their new representations.

In general, you can allow each affected object to be converted automatically the first time it is accessed after schema evolution, potentially distributing the performance impact of conversion across many transactions. Alternatively, you can concentrate the performance impact of conversion into fewer transactions by converting all the affected objects in a container, a database, or a federated database *on demand*.

You use this member function in an update transaction to convert the affected objects in a federated database on demand. This member function has no effect if the affected objects in the federated database have already been converted.

**NOTE** On-demand object conversion cannot be used for schema operations that require an upgrade application; see ooTrans::upgrade.

The convertObjects member function automatically drops any index that is invalidated by a schema evolution change. Specifically, if you changed the type or deleted a data member that is a key field in a key description, the corresponding indexes are dropped.

If you make schema changes in a development federated database and then distribute these to deployed federated databases (using the oschemadump and oschemaupgrade tools), you can safely purge schema-evolution history from the development federated database *only after* you have distributed *all* schema changes and then converted *all* affected objects in *all* of the deployed federated databases.

**WARNING** If you use the ooschemadump tool to write a purged schema to an output file, the resulting file will not be accepted by any deployed federated database that still contains unconverted objects from earlier schema-evolution operations (that is, objects whose shapes have been purged from the schema you are distributing). To update the schema of such a deployed federated database, your only recourse is to restore the development federated database to an earlier state, perform the schema-evolution operation again, and distribute the evolved schema without purging its history.

See also Chapter 19, "Object Conversion," in the Objectivity/C++ programmer's guide

## dumpCatalog

(*administration*) Prints out a list of the files associated with the referenced federated database; primarily used in administration tools.

```
ooStatus dumpCatalog(
    FILE *outputFile = stdout,
    const ooFileNameFormat format = oocHostLocal,
    const ooBoolean printLabels = oocTrue) const;
```

Parameters *outputFile* 

Pointer to a transcript file in which to print out the list of files. The default is standard output.

format

Format in which to print each filename:

- Specify oocHostLocal (the default) to use host:localPath—for example, object:/mnt/ed/design/up.FDB).
- Specify oocNative to use full network pathnames—for example, /net/object/usr/mnt/ed/dsgn/up.FDB.

#### printLabels

Specifies whether to identify each file listed in the output. By default, each filename is labeled; specify oocFalse to suppress labels.

- Returns oocSuccess if successful; otherwise oocError.
- Discussion This member function lists all of the files associated with the federated database, including the system database file, the boot file, the journal directory, files used by the lock server, and all database files. (*FTO*) In a partitioned federated database, the output includes system database files and boot files of all autonomous partitions. (*DRO*) If replication is used, the output includes all database image files.

This member function performs the same function as the oodumpcatalog tool (see the Objectivity/DB administration book).

### exist

Tests for the existence of the specified federated database; if successful, sets this object reference or handle to reference the federated database and optionally opens it.

```
ooBoolean exist(
    const char *bootFilePath,
    const ooMode openMode = oocNoOpen);
```

Parameters	<ul> <li>bootFilePath         <ul> <li>Path to the boot file of the federated database or an autonomous partition.</li> <li>(FTO) If the specified file is the boot file of an autonomous partition, that partition is the boot autonomous partition for the application.</li> </ul> </li> <li>openMode         <ul> <li>Mode in which to open the federated database if it exists:</li> <li>Specify oocNoOpen (the default), to set this object reference or handle to the federated database without opening it.</li> <li>Specify oocRead to open the federated database for read.</li> <li>Specify oocUpdate to open the federated database for update.</li> </ul> </li> </ul>
Returns	occTrue if the specified federated database exists, or occFalse if the federated database does not exist or if it is not accessible.
Discussion	If the specified federated database exists, this object reference or handle is set to the federated database; otherwise, this object reference or handle is set to null.
	If you specifically want to test for existence, you use the <code>openMode</code> parameter's default value ( <code>oocNoOpen</code> ). Otherwise, a return value of <code>oocFalse</code> could mean either that the federation doesn't exist, or that it does exist, but cannot be opened.
See also	open
isValid	
	Checks whether this object reference or handle is valid—that is, whether it references the federated database.
	ooBoolean isValid() const;
Returns	oocTrue if the object reference or handle references the federated database; oocFalse if this object reference or handle is null.
lock	
	Explicitly locks the referenced federated database.
	<pre>ooStatus lock(const ooLockMode lockMode) const;</pre>
Parameters	<pre>lockMode Type of lock to request:     Specify oocLockRead to request a read lock.     Specify oocLockUpdate to request an update lock.</pre>

Returns oocSuccess if the requested lock is obtained; otherwise oocError.

Discussion Objectivity/DB operations request and obtain locks implicitly as they are needed. You use this member function to obtain a lock explicitly when you want to reserve access to a federated database in advance.

Explicitly locking a federated database for read essentially allows concurrent read transactions and prevents concurrent update transactions against the same federated database. More specifically, locking a federated database for read in a read transaction:

- Prevents any other transaction from concurrently opening the federated database for update.
- Allows any other transaction to concurrently open the federated database (and its databases and containers) for read.

Explicitly locking a federated database for update in an update transaction prevents all concurrent access; an exclusive lock is placed on the federated database, which prevents any other transaction from opening it for read or update.

A federated database cannot be locked:

- For read, if another transaction has already opened it for update.
- For update, if another transaction has already locked it for read or opened it for update.

### lockServerName

Gets the network name of the host running the lock server for the referenced federated database.

char \*lockServerName() const;

Returns Pointer to a string containing the lock server host name.

Discussion The string is statically allocated by the member function and is overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.

#### name

Gets the system name of the referenced federated database.

char \*name() const;

Returns Pointer to a string containing the system name.

Discussion	The string is statically allocated by the member function and is overwritten with
	each invocation. You should make a local copy of the returned string if you
	intend to use it later in the application.

### number

(*administration*) Gets the identifier of the referenced federated database.

uint32 number() const;

Returns Integer representing the federated-database identifier.

#### open

Opens the specified federated database for the specified level of access and sets this object reference or handle to reference the opened federated database.

```
ooStatus open(
    const char *bootFilePath,
    const ooMode openMode = oocRead,
    ooBoolean recover = oocFalse);
```

#### Parameters *bootFilePath*

Path to the boot file of the federated database or autonomous partition to be opened. You can omit this parameter if you set the  $OO_FD_BOOT$  environment variable to the path. You can specify this path with or without a host name. If you specify it as a host path, use the format *host*::*path*.

(*FTO*) If the specified file is the boot file of an autonomous partition, that partition is the boot autonomous partition for the application.

#### openMode

Intended level of access to the federated database:

- Specify occRead (the default) to open the federated database for read and to designate the transaction as a read transaction.
- Specify occUpdate to open the federated database for update (read and write) and to designate the transaction as an update transaction.

#### recover

Specifies whether to perform local automatic recovery when opening the federated database. If you specify occTrue, Objectivity/DB rolls back any incomplete *local* transactions against the federated database (transactions that were started by other applications running on the same host).

Returns oocSuccess if successful; otherwise oocError.

Discussion You must call open at the beginning of each transaction in an application. In the first transaction, opening the federated database helps to initialize Objectivity/DB with schema information and storage page size. In subsequent transactions, this operation has minimal performance impact because it simply verifies that the same federated database is being accessed. *Only one* federated database can be open in a process.

Any number of transactions may concurrently open the same federated database for read or update access, subject to existing locks (see the <u>lock</u> member function):

- A federated database can be opened for read but not update if another transaction has already locked it for read.
- A federated database cannot be opened in any mode if another transaction has already locked it for update.

If you intend to modify any object in a federated database, you must open the federated database for update, or the changes you make will be lost. Within a transaction, you can promote the open mode from read to update by calling open again with *openMode* set to oocUpdate. You do not need to close the federated database first. However, you may not demote the open mode from update to read; if you try, oocSuccess is returned, but the open mode remains unchanged.

You use the recover parameter to enable automatic recovery for the application. For performance reasons, you should arrange for this parameter to be set to oocTrue only once in an application (during the first transaction). For more information about automatic recovery, see the Objectivity/DB administration book.

See also <u>exist</u>

### openMode

Gets the current level of access to the referenced federated database.

ooMode openMode() const;

Returns One of the following constants:

- oocNoOpen—the federated database is closed in this transaction.
- oocRead—the federated database is open for read in this transaction.
- oocUpdate—the federated database is open for update in this transaction.

## pageSize

(administration) Gets the storage page size of the referenced federated database.

```
uint32 pageSize() const;
```

Returns Integer representing the federated database's storage page size. This size was set when the federated database was created.

### setConversion

Registers a conversion function for objects of the specified class in the referenced federated database.

ooStatus setConversion(	
const char * <i>className</i> ,	
ooConvertFunction convFunction)	const;

Parameters className

Class of objects to be converted by the conversion function.

#### convFunction

Pointer to an application-defined conversion function.

Discussion A conversion function augments any kind of object conversion performed during or after schema evolution. You can register a conversion function in any application that is to trigger object conversion. When a conversion function is registered for a particular evolved class, the application invokes the function each time an object of that class is converted. A conversion function typically gets data member values from an existing unconverted object and then sets new data member values in the object's converted representation.

> In a given application, you can register no more than one conversion function for each evolved persistence-capable class. Registering a second conversion function for a class replaces the previously registered function.

- Returns oocSuccess if successful; otherwise oocError.
- See also Chapter 19, "Object Conversion," in the Objectivity/C++ programmer's guide

### tidy

(*administration*) Consolidates fragmented storage space in all the databases in the referenced federated database; used only in custom administration tools.

```
ooStatus tidy(
   FILE *outputFile = stdout,
   const char *hostName = 0,
   const char *pathName = 0) const;
```

Parameters	outputFile		
	Pointer to the transcript file in which to write the generated report. The default is standard output.		
	hostName		
	Name of the data server host on which to create the temporary file. Specify 0 to create the file on the current host.		
	pathName		
	Pathname of the directory on <i>hostName</i> in which to create the temporary file.		
Returns	Returns a non-zero value for ooStatus if successful.		
Discussion	This member function creates a temporary file for intermediate data during execution and therefore requires free disk space equal to the size of the largest database file in the federation. The temporary file is created in the directory that contains the federated database's system database file unless you specify <i>hostName</i> and <i>pathName</i> .		
	You should call tidy in a single-purpose update transaction. That is, you must not manipulate any database, container, or basic object before calling tidy in the same transaction, and you must commit the transaction immediately after tidy completes. This is because compacting and relocating physical storage renders the databases inconsistent with any system data that was cached during the transaction, and commiting the transaction discards the obsolete cached data.		
	You must not abort the transaction after calling the $\mathtt{tidy}$ member function.		
Warning	To prevent database corruption, make sure no other transactions are concurrently accessing a federated database being tidied.		
	This member function performs the same function as the <code>ootidy</code> tool (see the Objectivity/DB administration book).		
typeN			
	Gets the type number of the federated-database class ooFDObj.		
	ooTypeNumber typeN() const;		
Returns	Type number of the federated-database class ooFDObj.		

## typeName

	Gets the name of the federated-database class ooFDObj.
	char *typeName() const;
Returns	The string "ooFDObj".
Discussion	The returned string must be treated as read-only.
update	
	Opens the referenced federated database for update access.
	ooStatus update();
Returns	oocSuccess if successful; otherwise oocError.
Discussion	This member function is equivalent to calling $\underline{open}(oocUpdate)$ .

## upgradeObjects

Performs object conversion throughout the referenced federated database; used only in a special-purpose upgrade application.

ooStatus upgradeObjects(ooBoolean purge\_schema = oocFalse);

Parameters	purge_schema
	Specifies whether to remove schema-evolution history after objects are converted. By default, this history is not purged. <b>Warning:</b> Purging the schema may delete information that is required for distributing schema changes to deployed federated databases (see Discussion below). You should take a backup of the federated database before purging its schema.
Returns	oocSuccess if successful, or oocError if no classes have been changed by schema evolution.
Discussion	Certain schema-evolution operations require that you create a special-purpose upgrade application to perform object conversion. An upgrade application takes the place of on-demand or deferred conversion. You can, however, register one or more conversion functions in an upgrade application.
	An upgrade application is identified by a invoking <code>ooTrans::upgrade</code> . You then call the <code>upgradeObjects</code> member function in an update transaction.
	If you make schema changes in a development federated database and then distribute these to deployed federated databases (using the ooschemadump and ooschemaupgrade tools), you can safely purge schema-evolution history from

the development federated database *only after* you have distributed *all* schema changes and then converted *all* affected objects in *all* of the deployed federated databases.

**WARNING** If you use the ooschemadump tool to write a purged schema to an output file, the resulting file will not be accepted by any deployed federated database that still contains unconverted objects from earlier schema-evolution operations (that is, objects whose shapes have been purged from the schema you are distributing). To update the schema of such a deployed federated database, your only recourse is to restore the development federated database to an earlier state, perform the schema-evolution operation again, and distribute the evolved schema without purging its history.

See also Chapter 19, "Object Conversion," in the Objectivity/C++ programmer's guide

# ooRefHandle(ooObj) Classes

Inheritance: **ooRef(ooObj)** 

Inheritance: **ooHandle(ooObj)** 

The abbreviation *ooRefHandle*(ooObj) refers to two non-persistence-capable classes:

- ooRef(ooObj), which represents an object reference to an Objectivity/DB object.
- ooHandle(ooObj), which represents a *handle* to an Objectivity/DB object.

These two classes are documented together because they define almost identical sets of member functions (exceptions are listed in the "Reference Summary"). These classes, along with the 000bj class, define persistence behavior for various kinds of *Objectivity/DB objects* (instances of 000bj and its derived classes).

See:

- "Reference Summary" on page 601 for an overview of member functions
- "Reference Index" on page 603 for a list of member functions

(ODMG) You can use either of the ODMG standard class names d\_Ref<ooObj> or d\_Ref<d\_Object> interchangeably with ooRef(ooObj).

# **About Handles and Object References**

An application works with an Objectivity/DB object indirectly through one or more handles or object references that are set to reference the object. An object of a particular type is usually referenced by a handle or object reference of the corresponding type—for example, a standard container is referenced by instances of *ooRefHandle*(ooContObj), a database is referenced by instances of *ooRefHandle*(ooDBObj), a basic object of an application-defined class appClass is referenced by instances of *ooRefHandle(appClass)*, and so on. All of the various type-specific handle and object-reference classes are derived from the *ooRefHandle(ooObj)* classes.

In general, every handle or object reference serves to:

- Identify the referenced object to the application or to another object.
- Provide an interface for operating on the referenced object.

In addition, a handle or object reference to a persistent object serves as a type-safe smart pointer that:

- Manages the memory pointer to the object.
- Provides an indirect member-access operator (->) for accessing the object's public member functions.

It is sometimes more appropriate to use a handle rather than an object reference, and vice versa; the choice is described in "Structure and Behavior" on page 597. A simple guideline is to use handles in function definitions and object references as data member types in persistence-capable class definitions.

### **Inheritance Hierarchy**

The ooRefHandle(ooObj) classes are the base classes for all type-specific handle and object-reference classes, forming a pair of inheritance hierarchies that parallel the class derivation from <u>ooObj</u>. That is:

- Just as class ooObj is the base class for all Objectivity/DB objects, ooHandle(ooObj) is the base class for all handle classes, and ooRef(ooObj) is the base class for all object-reference classes.
- For every class className that derives from ooObj, a corresponding handle class ooHandle(className) derives from ooHandle(ooObj), and a corresponding object-reference class ooRef(className) derives from ooRef(ooObj).

The *ooRefHandle*(ooObj) classes are therefore the base classes for:

- Handles and object references to storage objects and autonomous partitions (instances of the <u>ooRefHandle(ooContObj)</u>, <u>ooRefHandle(ooDBObj)</u>, <u>ooRefHandle(ooFDObj)</u>, and <u>ooRefHandle(ooAPObj)</u> classes).
- Handles and object references to persistent objects, including:
  - □ Instances of the <u>ooRefHandle(appClass)</u> classes, which are generated by the DDL processor for every application-defined persistence-capable class appClass derived from ooObj or ooContObj.
  - Instances of predefined classes such as ooRefHandle(ooMap), ooRefHandle(ooGeneObj), and ooRefHandle(ooGCContObj), which exist for every persistence-capable class defined by Objectivity/C++.

**NOTE** The predefined handle and reference classes for most persistence-capable classes are not documented separately; you can think of them as a kind of <a href="https://ooRefHandle(appClass">ooRefHandle(appClass)</a>.

### Interface

The *ooRefHandle*(ooObj) base classes define three overlapping interfaces:

- The core interface for operating on Objectivity/DB objects of any type. This interface includes member functions for general Objectivity/DB operations, such as opening, locking, printing object identifiers, and so on.
- An interface for operating on persistent objects (basic object or container only). This interface extends the core interface to include an indirect member-access operator (->) and member functions for managing scope names.
- The complete interface for operating on basic objects. This interface extends the persistent-object interface to include member functions for moving, copying, and versioning operations.

Every derived handle or object-reference class acquires the member functions of the applicable interface by either inheriting them or redefining them wherever type-specific parameters or behavior are required. (Because these member functions are not virtual, they are hidden but not overridden in the derived classes that redefine them.)

Furthermore, the derived handle and object-reference classes disallow any inapplicable member functions by redefining them as private members. For example, the *ooRefHandle*(ooContObj) classes redefine the inherited copy member function as private because the copy operation is intended only for basic objects.

The majority of member functions defined by <code>ooRefHandle(ooObj)</code> are inherited by the handle and object-reference classes for persistent objects; comparatively few member functions are reimplemented or disallowed. In contrast, the <code>ooRefHandle(ooDBObj)</code>, <code>ooRefHandle(ooFDObj)</code>, and <code>ooRefHandle(ooAPObj)</code> classes disallow or reimplement a majority of the <code>ooRefHandle(ooObj)</code> member functions, and so diverge considerably from the base interface.

### Usage

Handles and object references of type ooRefHandle(ooObj) can reference instances of ooObj or of any class derived from ooObj. Such handles and object references provide an alternative to using type-specific handles and object

references. For example, a basic object of an application-defined class
RentalFleet can be referenced by instances of ooRefHandle(RentalFleet)
and by instances of ooRefHandle(ooObj).

Because of the interfaces of the various handle and object-reference classes, instances of the *ooRefHandle*(ooObj) classes are used in two distinct ways:

- As *persistent-object handles and object references*, which can be set to reference either basic objects or containers. This is the more common usage.
- As *general-purpose handles and object references*, which can be set to reference any type of Objectivity/DB object. This usage is fairly rare.

## Persistent-Object Handles and Object References

Instances of the *ooRefHandle*(ooObj) base classes are normally used for referencing just persistent objects (basic objects and containers) rather than all types of Objectivity/DB objects. This is because:

- Most member functions defined by ooRefHandle(ooObj) are implemented for operating on persistent objects. Thus, a persistent-object handle of type ooHandle(ooObj) can perform the same persistence operations as can a handle of the type-specific class ooHandle(appClass), when referencing an appClass object.
- Most Objectivity/C++ functions that use a ooRefHandle(ooObj) & parameter perform operations that are available only to persistent objects or only to basic objects, but not to the federated database, databases, or autonomous partitions.

Your application can use variables of type <code>ooRefHandle(ooObj)</code> as persistent-object handles and object references when the class of the referenced persistent object can't be known until run time.

## **General-Purpose Handles and Object References**

Instances of the ooRefHandle(ooObj) classes are used as general-purpose handles and object references by certain Objectivity/C++ functions that accept or return several kinds of referenced Objectivity/DB object. These functions use such handles and object references as parameters or return values. For example, because any kind of Objectivity/DB object can be a scope object for defining scope names, the functions that set, get, and remove scope names have a parameter of type ooHandle(ooObj) & to specify the scope object.

**Note** Although a parameter of type <code>ooRefHandle(ooObj) & syntactically accepts a handle or object reference of any type, the function itself need not operate on all types of referenced object. For example, <code>ooDelete</code> has a parameter of type <code>ooHandle(ooObj) & for specifying the object to be deleted. This function does</code></code>

not delete federated databases, however, so it signals an error if it receives a handle to a federated database.

It is convenient and appropriate to use a general-purpose handle or object reference when both of the following conditions are met:

- The type of the referenced object can't be known until runtime.
- The specific type isn't relevant because only core persistence operations are to be performed on the referenced object.

In all other cases, your application should avoid using variables of type *ooRefHandle*(*ooDbj*) as general-purpose handles and object references. This is because a general-purpose handle or object reference to a database, federated database, or autonomous partition cannot perform the same persistence operations on the referenced object as can a handle or object reference of the appropriate type-specific class.

Instead, your application should:

- Use variables of type *ooRefHandle*(ooObj) only as persistent-object handles and object references.
- Reference the federated database, databases, and autonomous partitions with type-specific handles and object references.

### Structure and Behavior

Regardless of type, all handles and object references provide a way to reference and operate on Objectivity/DB objects. The key difference between handles and object references is:

- Handles contain both an identifier for the referenced object and cache-related state that is associated with the Objectivity context in which it was created.
- Object references contain an identifier for the referenced object, but no bulky cache-specific state.

Therefore, only object references (but not handles) can be passed between Objectivity contexts and saved persistently.

When databases, federated databases, or autonomous partitions are referenced, handles are nearly equivalent to object references—both are essentially wrappers for a referenced object's identifier, as described in the "Structure and Behavior" sections of <u>ooRefHandle(ooDBObj)</u>, <u>ooRefHandle(ooFDObj)</u>, and <u>ooRefHandle(ooAPObj)</u>.

When persistent objects are referenced, handles and object references are optimized for different purposes:

- Handles are optimized for accessing persistent objects in memory—that is, for performing multiple operations on a referenced object or repeatedly accessing the object's members.
- Object references are optimized for linking persistent objects—that is, for storing object identifiers persistently in reference attributes, in associations, or as elements of a collection.

### Handles to Persistent Objects

Handles are optimized for efficient in-memory access because they can automatically obtain and manipulate pointers to referenced persistent objects. Thus, when a handle is set to reference a particular persistent object, the handle stores the object identifier for that object. The first time the persistent object is accessed through the handle, the handle is automatically *opened*—that is, the handle obtains a pointer to the object's representation in memory. This memory pointer enables the handle to access the referenced object quickly during subsequent operations performed through the handle. When the handle is *closed*, it invalidates the pointer but keeps the object identifier, so the application can reuse the handle (without resetting it) to access the same object.

Besides maintaining a pointer to the referenced persistent object, an open handle also *pins* the object's memory representation in the Objectivity/DB cache. Pinning guarantees that the persistent object is readily available in memory for as long as it is needed. Closing the handle removes its particular "pin"; when the last open handle to that persistent object is closed, the last pin is removed and the object itself is closed. Closing the last (or only) object on a buffer page permits Objectivity/DB to swap the page out of the cache as needed.

## **Object References to Persistent Objects**

Object references are optimized for implementing persistent links because they are essentially wrappers for object identifiers. Thus, setting an object reference to a particular persistent object causes the object reference to store the object's identifier. The object reference never acquires a pointer to the persistent object in memory; instead, whenever the persistent object is accessed through the object reference, the operation is delegated to a temporary handle that provides the necessary pointer.

For convenience, an application can use an object reference (instead of a handle) to perform an operation on a referenced persistent object or to access one of the object's members. However, poor performance results when an object reference is used for multiple such operations on the same persistent object, because *each* operation causes a temporary handle to be created, used, and discarded. Performance may also be affected by swapping, because the object reference does

not pin the persistent object's memory representation in the Objectivity/DB cache.

An object reference of type <code>ooRef(ooObj)</code> is sometimes called a *standard object reference* because it contains the complete object identifier for the referenced object. An alternative for referencing a basic object under certain circumstances is the *short object reference* of type <code>ooShortRef(ooObj)</code>, which saves space by storing object identifiers in a truncated format.

# **Working With Persistent-Object Handles**

**NOTE** For simplicity, this section describes how to work with handles to persistent objects. Except where noted, the same information applies to object references.

An application normally creates a persistent-object handle as a local variable on the stack, rather than allocating it on the heap. A handle should not be declared as const, because its internal state is changed by any operation that accesses a persistent object through it. (Object references may be declared as const.) An application should not explicitly define subclasses of the ooRefHandle(ooObj)classes; any necessary subclasses are generated automatically by the DDL processor if the application defines any subclasses of ooObj.

A new persistent-object handle is normally *null*—that is, it contains the value 0 instead of an object identifier. The application can then set the handle to reference a particular persistent object in any of the following ways:

- By creating a new persistent object with operator new of the desired class and assigning the result to the handle.
- By finding an existing persistent object with the handle's <u>lookupObj</u> member function.
- By passing the handle to a member function that sets it, such as the <u>linkName</u> member function of a persistent object that has a <u>linkName</u> association. The <u>linkName</u> member function finds the associated destination object and sets the specified handle to the found object.
- By assignment or initialization from another handle or object reference.

An object reference may be set in any of these ways, with the following exception—the result of operator new may not be assigned to an object reference.

A handle continues to reference the same persistent object until it is set to another persistent object or to null. Furthermore, multiple handles and object references can be set to the same persistent object. An application operates on a persistent object by calling:

- Member functions of a handle that references the object. As indicated in "Reference Summary" on page 601, such member functions allow you to assign a scope name to the referenced object, open it for update, and so on. To call a member function of a handle, you use the direct member-access operator (.). For example, objH.lookupObj calls the <u>lookupObj</u> member function of the handle objH.
- Member functions of the referenced object itself.

To call a member function of a referenced object, you use the handle's overloaded indirect member-access operator (<u>operator-></u>). For example, objH->oolsKindOf calls the <u>oolsKindOf</u> member function on the persistent object that is referenced by the handle objH. Only members inherited from ooObj can be accessed from a persistent-object handle.

Although most of a handle's member functions operate on the referenced persistent object, some functions operate on the handle itself. For example, you use:

- The assignment operator <u>operator=</u> to set a handle from another handle or from an object reference.
- The comparison operators <u>operator==</u> and <u>operator!=</u> to test whether a handle references the same persistent object as another handle or object reference.
- The member function <u>isNull</u> to test whether a handle is null. (Alternatively, you can use the overloaded <u>operator==</u> to compare a handle to 0.)

A handle preserves its reference to a persistent object across transaction boundaries, provided that the handle does not go out of scope and is not set to null as the result of an abort operation. Before reusing the handle in a new transaction, however, the application should call <u>isValid</u> to test whether the handle is still *valid*—that is, whether it still references an existing persistent object. A handle becomes invalid if it is set to null *or* if the referenced persistent object has been deleted by another process between transactions.

## **Opening and Closing a Persistent-Object Handle**

**NOTE** This subsection applies only to handles, not to object references, which are in effect always closed.

A handle is automatically opened when a persistent object is opened through it. The open persistent object is both locked and represented in memory; the open handle manages a pointer to the persistent object, pinning the object in memory until the handle is closed. A closed handle, which has an object identifier instead of a pointer, can reference either an open or a closed persistent object.

The most common way to open a persistent object through a handle is to do so implicitly by using the handle's indirect member-access operator (<u>operator-></u>) to access a member of the referenced object. Alternatively, a referenced persistent object can be opened by explicit request—for example, by calling the handle's <u>open or update</u> member function. Another way to explicitly open a persistent object is by finding it with a function whose *openMode* parameter is either oocRead or oocUpdate. (Most functions that set a handle to a found persistent object provide an *openMode* parameter for specifying the desired level of access through that handle.) In all cases, if the found or referenced persistent object is already open (for example, because another operation opened it earlier in the transaction), the accessing handle gets a pointer to the existing memory representation and adds a pin.

You obtain a closed handle to a persistent object by finding the object with a function whose *openMode* parameter is set to oocNoOpen. Such operations simply provide the handle with a persistent object's object identifier without adding a pin, even if the object is already open through another handle.

Objectivity/DB automatically closes an open handle when the handle is destroyed (for example, by going out of scope), when it is set to reference another persistent object, or when the transaction that opened it commits or aborts. An application can close a handle explicitly by calling the handle's <u>close</u> member function. Closing the last open handle to a particular persistent object unpins and closes the object.

# **Reference Summary**

The summarized member functions are defined on both the object-reference class and the handle class. Two operators are defined on only the handle class, namely, <u>operator\*</u> and <u>operator ooObj\*</u>. One operator is defined on only the object-reference class, namely, <u>operator int</u>.

Creating a Handle or Object Reference	<u>ooRef(ooObj)</u> ooHandle(ooObj)
Setting the Handle or Object Reference	<u>operator= lookup0bj set_container</u>
Comparing Handles and Object References	<u>operator==</u> operator!=

Accessing the Persistent Object	<u>operator-&gt;</u> <u>operator*</u> <u>operator ooObj*</u> <u>ptr</u> ( <i>ODMG</i> )
Opening, Closing, and Locking the Persistent Object	open openMode update close lock lockNoProp
Modifying the Persistent Object	<u>delete object</u> ( <i>ODMG</i> ) <u>update</u>
Copying or Moving the Basic Object	<u>copy</u> move
Getting Information About the Objectivity/DB Object	<u>print</u> <u>sprint</u> <u>typeN</u> <u>typeName</u>
Testing the Handle or Object Reference	<u>openMode</u> <u>is null</u> (ODMG) <u>isNull</u> <u>isValid</u> <u>operator int</u> <u>operator ooObj*</u>
Working With Scope Names	nameObj lookupObj getObjName unnameObj getNameObj getNameScope
Finding Objects	<u>lookupObj</u> <u>containedIn</u> <u>getNameObj</u> getNameScope

Versioning the Basic Object	<u>getDefaultVers</u> <u>setDefaultVers</u> <u>getNextVers</u> <u>getPrevVers</u> <u>getVersStatus</u> <u>setVersStatus</u>
ODMG Interface	<u>operator d Ref_Any</u> <u>delete_object</u> <u>is_null</u> ptr

# **Reference Index**

<u>close</u>	Explicitly closes this handle.
<u>containedIn</u>	Finds the container that contains the referenced basic object.
copy	Creates a copy of the referenced basic object, clustering the new copy near the specified object.
<u>delete_object</u>	(ODMG) Deletes the referenced persistent object.
<u>getDefaultVers</u>	Finds the default version of the referenced basic object.
<u>getNameObj</u>	Initializes an object iterator to find all objects named in the scope of the referenced persistent object.
getNameScope	Initializes an object iterator to find all scope objects in the federated database that define a scope name for the referenced persistent object.
<u>getNextVers</u>	Initializes an object iterator to find the next version(s) of the referenced basic object.
<u>getObjName</u>	Gets the name defined in the specified scope for the referenced persistent object.
getPrevVers	Finds the previous version of the referenced basic object.
<u>getVersStatus</u>	Gets the current versioning mode of the referenced basic object.
<u>is null</u>	(ODMG) Tests whether this object reference or handle is null.
<u>isNull</u>	Tests whether this object reference or handle is null.

<u>isValid</u>	Tests whether this object reference or handle is valid—that is, whether it references an existing Objectivity/DB object.
lock	Explicitly locks the referenced persistent object; propagates locks along associations that have lock propagation enabled.
lockNoProp	Explicitly locks the referenced persistent object, without propagating locks to associated destination objects.
<u>lookupObj</u>	Finds the persistent object with the specified scope name (or the basic object matching the specified key structure) within the specified scope, and sets this object reference or handle to reference the found object.
move	Moves the referenced basic object to a different container.
<u>nameObj</u>	Names the referenced persistent object in the specified scope.
<u>ooHandle(ooObj)</u>	Default constructor that constructs a null handle.
<u>ooHandle(ooObj)</u>	Constructs a handle that references the same Objectivity/DB object as the specified object reference, handle, pointer, or ODMG generic reference.
<u>ooRef(ooObj)</u>	Default constructor that constructs a null object reference.
<u>ooRef(ooObj)</u>	Constructs an object reference that references the same Objectivity/DB object as the specified object reference, handle, pointer, or ODMG generic reference.
open	Explicitly opens the referenced persistent object, preparing the object for the specified level of access.
<u>openMode</u>	Tests whether this handle is open, and, if so, gets the current level of access to the referenced basic object.
<u>operator-&gt;</u>	Indirect member-access operator; accesses a member of the referenced persistent object.
operator*	Handle class only. Dereference operator; returns the persistent object referenced by this handle.
<u>operator=</u>	Assignment operator; sets this object reference or handle to reference the same Objectivity/DB object as the specified object reference, handle, or pointer.
<u>operator==</u>	Equality operator; tests whether this object reference or handle has the same value as the specified item.

<u>operator!=</u>	Inequality operator; tests whether this object reference or handle has a different value from the specified item.
<u>operator d Ref Any</u>	(ODMG) Conversion operator that returns an ODMG generic reference to the referenced object.
<u>operator int</u>	Object-reference class only. Conversion operator that tests whether this object reference is null.
<u>operator ooObj*</u>	Handle class only. Conversion operator that returns a C++ pointer to the Objectivity/DB object that is referenced by this handle.
<u>print</u>	Prints the object identifier of the referenced Objectivity/DB object.
ptr	(ODMG) Returns a C++ pointer to the referenced persistent object.
<u>set_container</u>	Provides this object reference or handle with container information so you can assign a short object reference to it.
<u>setDefaultVers</u>	Sets the referenced basic object as the default version of its genealogy.
<u>setVersStatus</u>	Enables or disables versioning for the referenced basic object by setting the object's versioning mode.
<u>sprint</u>	Returns a string containing the object identifier of the referenced Objectivity/DB object.
typeN	Gets the type number of the class of the referenced Objectivity/DB object.
typeName	Gets the name of the class of the referenced Objectivity/DB object.
<u>unnameObj</u>	Deletes the name in the specified scope for the referenced persistent object.
update	Opens the referenced persistent object for update access.

# Constructors

# ooHandle(ooObj)

Default constructor that constructs a null handle.

ooHandle(ooObj)();

## ooHandle(ooObj)

Constructs a handle that references the same Objectivity/DB object as the specified object reference, handle, pointer, or ODMG generic reference.

1.	ooHandle(ooObj)(const	<i>ooRefHandle</i> (ooObj)	&objectRH);
2.	ooHandle(ooObj)(const	ooObj * <i>objectP</i> );	
(ODMG) 3.	ooHandle(ooObj)(const	<pre>d_Ref_Any &amp;from);</pre>	

Parameters

objectRH

Object reference or handle to any Objectivity/DB object (basic object, container, database, federated database, or autonomous partition).

objectP

Pointer to a basic object, container, database, or autonomous partition. The pointer *must* be the result of <u>operator new</u> on ooObj or a derived class.

from

(ODMG) An ODMG generic reference to a basic object or container.

Discussion Variants 1 and 3 allow a new handle to be constructed from an existing object reference, handle, or ODMG generic reference. If the new handle is constructed from an existing open handle, the new handle is open; in all other cases, the new handle is closed.

Variant 2, which constructs a handle from the specified pointer, has a narrower purpose—to obtain an open handle to a newly created Objectivity/DB object so you can perform persistence operations on the object and so the object can eventually be unpinned when it is no longer needed in memory.

## ooRef(ooObj)

Default constructor that constructs a null object reference.

```
ooRef(ooObj)();
```

## ooRef(ooObj)

Constructs an object reference that references the same Objectivity/DB object as the specified object reference, handle, pointer, or ODMG generic reference.

```
1. ooRef(ooObj)(const ooRefHandle(ooObj) &objectRH);
```

```
2. ooRef(ooObj)(const ooObj *objectP);
```

```
(ODMG)3. ooRef(ooObj)(const d_Ref_Any & from);
```

Parameters objectRH

Object reference or handle to any Objectivity/DB object (basic object, container, database, federated database, or autonomous partition).

objectP

Pointer to a persistent object (basic object or container). The pointer may *not* be the result of <code>operator new</code>. Instead, the pointer must be the result of using either <code>operator new</code>. Instead, the pointer must be the result of using either <code>operator new</code>. Instead, the pointer must be the result of using either <code>operator new</code>. Instead, the pointer must be the result of using either <code>operator new</code>. Instead, the pointer must be the result of using either <code>operator new</code>. Instead, the pointer must be the result of using either <code>operator new</code>. Instead, the pointer must be the result of using either <code>operator new</code>. Instead, the pointer must be the result of using either <code>operator new</code>. Instead, the pointer must be the result of using either <code>operator new</code>. Instead, the pointer must be the result of using either <code>operator new</code>. Instead, the pointer must be the result of using either <code>operator new</code>. Instead, the pointer must be the result of using either <code>operator new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must be the result of using either <code>new</code>. Instead, the pointer must

```
from
```

(*ODMG*) An ODMG generic reference to a persistent object (basic object or container).

Discussion Variants 1 and 3 allow a new object reference to be constructed from an existing object reference, handle, or ODMG generic reference.

Variant 2 has a narrower purpose, which is to allow you to resume persistence operations on a persistent object after manipulating it through a pointer. The use of this variant should be rare, however, because pointers are not normally used to manipulate persistent objects.

# Operators

### operator->

Indirect member-access operator; accesses a member of the referenced persistent object.

ooObj \*operator->();

Returns Pointer to the referenced basic object or container.

Discussion The accessed persistent object is opened for read, if it is not already open.

You use operator-> in an expression handle->member, where handle is an instance of ooHandle(ooObj) and member is the name of a public member defined on class ooObj. As for any overloading of the C++ member-access operator (->), the expression handle->member is interpreted as (handle.operator->())->member. That is, the overloaded operator-> returns a pointer to the referenced object, and then the ordinary C++ operator-> selects the specified member of that object, returning the value of that member. If the referenced object is an instance of ooObj, operator-> can access any of the object's public members. If the referenced object is an instance of a derived class, operator-> accesses only the members that are defined in ooObj.

### operator\*

*Handle class only.* Dereference operator; returns the persistent object referenced by this handle.

ooObj &operator\*();

Returns C++ reference to the persistent object referenced by this handle.

Discussion This operator enables you to pass a handle to a function that accepts a basic object or container by reference. This operator is analogous to the C++ operator\* for dereferencing a pointer.

The persistent object referenced by this handle is opened for read, if it is not already open.

- **WARNING** The returned reference is guaranteed valid only as long as the handle exists, remains open, and references the same object.
- Example This examples uses operator\* to pass a handle to helperFunction, which accepts a C++ reference to a persistent object.

```
void helperFunction(ooObj &anObject);
```

```
void processObject(ooHandle(ooObj) &objH) {
```

```
helperFunction(*objH);
```

### operator=

}

Assignment operator; sets this object reference or handle to reference the same Objectivity/DB object as the specified object reference, handle, or pointer.

1.	ooRefHandle(ooObj) & operator=(
	Const Obkerhandre(OOOD) &ODJectkh),
2.	oo <i>RefHandle</i> (ooObj) &operator=( const ooShortRef(ooObj) & <i>shortObjR</i> );
3.	<pre>ooRefHandle(ooObj) &amp;operator=(const ooObj *objectP);</pre>

#### Parameters

objectRH

Object reference or handle to an Objectivity/DB object (basic object, container, database, federated database, or autonomous partition).

#### short0bjR

Short object reference to a basic object that resides in the same container as the object that is referenced by this object reference or handle. (If this object reference or handle is null, you can set it to a container with the <u>set\_container</u> member function.)

A short object reference specifies just the lower half of an object identifier (corresponding to the object's logical page and slot numbers). The upper half of the object identifier (corresponding to the database and container) must be supplied by this object reference or handle.

#### objectP

0, or a nonnull pointer to one of several types of Objectivity/DB object:

- If you are assigning to a handle, the pointer may reference a basic object, container, database, or autonomous partition. The specified pointer *must* be the result of <u>operator new</u> on oo0bj or a derived class.
- If you are assigning to an object reference, the pointer may reference a persistent object (basic object or container). The specified pointer may not be the result of operator new. Instead, the pointer must be the result of using either <u>operator ooObj\*</u> on a handle or <u>ptr</u> on an object reference or handle earlier in the same transaction. If the specified pointer was originally extracted from a handle, that handle must still exist and reference the same persistent object.
- Returns This object reference or handle, after it has been set to reference the specified object.
- Discussion Variants 1 and 2 allow you to use the specified object reference or handle to produce another object reference or handle to the same object. If you are assigning to a handle from an open handle, the returned handle is open; in all other cases, the returned handle is closed.

Variant 3 allows you to set this object reference or handle to null. Otherwise, assignment-from-pointer has two specific purposes, depending on whether you are assigning to a handle or to an object reference:

- Pointer-to-handle assignment enables you to obtain an open handle to a newly created Objectivity/DB object, so you can perform persistence operations on it, and so the object can be unpinned when it is no longer needed in memory.
- Pointer-to-object-reference assignment enables you to resume persistence operations on a persistent object after manipulating it through a pointer. This

usage of variant 3 is rare, because pointers are not normally used to manipulate persistent objects.

### operator==

Equality operator; tests whether this object reference or handle has the same value as the specified item.

	<pre>1. ooBoolean operator==(     const ooRefHandle(ooObj) &amp;compare) const;</pre>
	<pre>2. ooBoolean operator==(     const ooObj *compare) const;</pre>
	<pre>3. ooBoolean operator==(     const ooShortRef(ooObj) &amp;compare) const;</pre>
	4. ooBoolean operator==(int zero) const;
Parameters	compare
	Any of the following:
	<ul> <li>Object reference or handle to an Objectivity/DB object (basic object, container, database, federated database, or autonomous partition).</li> </ul>
	<ul> <li>Pointer to a basic object or container.</li> </ul>
	<ul> <li>Short object reference to a basic object. A short object reference must refer to a basic object that resides in the same container as the object referenced by this object reference or handle.</li> </ul>
	zero
	Literal 0. This value allows you to use <code>operator==</code> as an alternative for <u>isNull</u> .
Returns	Variants 1, 2, and 3 return oocTrue if this object reference or handle contains the same identifier as <i>compare</i> ; otherwise oocFalse.
	Variant 4 returns oocTrue if this object reference or handle is null; otherwise oocFalse.
Discussion	Variants 1, 2, and 3 test whether this object reference or handle references the same object as <i>compare</i> . Variant 4 tests whether this object reference or handle is null.
See also	<u>isNull</u> (as an alternative to variant 4)

Inequality operator; tests whether this object reference or handle has a different value from the specified item.

	<pre>1. ooBoolean operator!=(     const ooRefHandle(ooObj) &amp;compare) const;</pre>	
	<pre>2. ooBoolean operator!=(     const ooObj *compare) const;</pre>	
	<pre>3. ooBoolean operator!=(     const ooShortRef(ooObj) &amp;compare) const;</pre>	
	4. ooBoolean operator!=(int zero) const;	
Parameters	<i>compare</i> Any of the following:	
	<ul> <li>Object reference or handle to an Objectivity/DB object (basic object, container, database, federated database, or autonomous partition).</li> </ul>	
	<ul> <li>Pointer to a basic object or container.</li> </ul>	
	<ul> <li>Short object reference to a basic object. A short object reference must refer to a basic object that resides in the same container as the object referenced by this object reference or handle.</li> </ul>	
	zero Literal 0.	
Returns	Variants 1, 2, and 3 return oocTrue if this object reference or handle contains a different identifier than <i>compare</i> ; otherwise oocFalse. Variant 4 returns oocTrue if this object reference or handle is nonnull; otherwise oocFalse.	
Discussion	Variants 1, 2, and 3 test whether this object reference or handle references a different object than <i>compare</i> . Variant 4 tests whether this object reference or handle is nonnull.	

## operator d\_Ref\_Any

(*ODMG*) Conversion operator that returns an ODMG generic reference to the referenced object.

```
operator d_Ref_Any() const;
```

### operator int

	<i>Object-reference class only.</i> Conversion operator that tests whether this object reference is null.
	operator int() const;
Returns	0 if this object reference is null; otherwise, returns a nonzero integer.
Discussion	This conversion operator enables you to use an object reference as the conditional expression in an if or while statement to test whether the object reference is null.
Example	ooRef(ooObj) objectR; // Set objectR to some object if (objectR) { //Do something interesting if initialization was successful }
See also	operator ooObi*

## operator ooObj\*

*Handle class only.* Conversion operator that returns a C++ pointer to the Objectivity/DB object that is referenced by this handle.

operator ooObj\*();

Returns The return value depends on what this handle references—specifically:

- A null pointer, if this handle is null.
- A nonnull pointer, if this handle references a database, federated database, or autonomous partition. This pointer should be used only in expressions that test for null.
- Otherwise, a pointer to the persistent object referenced by this handle.

Discussion

This conversion operator enables you to:

- Pass a handle to a function that accepts a pointer to a persistent object.
- Assign a handle to an ooObj\* variable (for example, to pass to the overloaded <u>operator delete</u>).
- Use a handle as the conditional expression in an if or while statement to test whether the handle is null.

If this handle references a persistent object, the object is opened for read (unless it is already open).
WARNING	The returned pointer is guaranteed valid only as long as this handle exists, remains open, and references the same object.	
	An application generally relies on handles to provide memory management for persistent objects, and avoids the explicit use of pointers to such objects. On occasion, explicit use of pointers is required for performance reasons or for compatibility with functions that are indifferent to persistence (for example, in legacy code or a third-party library). However, you should be careful when manipulating a pointer from a handle because the validity of the pointer depends on the state of the handle. You should not use the returned pointer in other persistence operations (for example, do not pass it to any Objectivity/C++ member function other than the overloaded operator delete).	
Example	The handle ${\tt objectH}$ is used as a conditional expression which evaluates to 0 if the handle is null.	
	<pre>ooHandle(ooObj) objectH;  // Set objectH to some object if (objectH) {  // Do something interesting if initialization was successful }</pre>	
See also	<u>operator int</u> <u>ptr</u>	

# **Member Functions**

### close

	Explicitly closes this handle.
	ooStatus close() const;
Returns	oocSuccess if successful; otherwise oocError.
Discussion	This member function is redundant for object references, which are, in effect, always closed. Therefore, you should use this member function only on handles.
	Objectivity/DB automatically closes handles when they go out of scope, when they are set to reference other objects, or when the transaction that opened them commits or aborts.
	You can use the close member function to close a handle explicitly. This informs Objectivity/DB that the application no longer requires access to the referenced

object through this handle. Closing does not, however, release any locks; locks are released only by committing or aborting a transaction.

A closed handle retains the object identifier of the Objectivity/DB object to which it refers, so you can reopen it without reinitializing. Note, however, that a retained object identifier can become invalid between transactions (for example, because a concurrent process has deleted the corresponding object), and opening a handle with an invalid object identifier signals an error.

Closing a handle to a persistent object invalidates the pointer to the object's representation in the Objectivity/DB cache. Closing the last open handle to a particular persistent object unpins and closes that object; closing the last open object on a buffer page permits Objectivity/DB to swap the page out of the cache as needed.

#### containedIn

Finds the container that contains the referenced basic object.

1.	ooHandle(ooContObj)	containedIn()	const;
	· · · · · · · · · · · · · · · · · · ·		

#### Parameters container

Object reference or handle to set to the found container.

#### Returns Object reference or handle to the found container.

Discussion When called without a *container* parameter, *containedIn* allocates a new container handle and returns it. Otherwise, *containedIn* returns the object reference or handle that is passed to it.

#### сору

Creates a copy of the referenced basic object, clustering the new copy near the specified object.

 ooHandle(ooObj) copy( const ooHandle(ooObj) &near) const;
 ooRef(ooObj) &copy( const ooHandle(ooObj) &near, ooRef(ooObj) &newCopy) const;

	<pre>3. ooHandle(ooObj) ©(</pre>
	ooHandle(ooObj) & <i>newCopy</i> ) const;
Parameters	near
	Handle to the object with which to cluster the new copy. <i>near</i> may be a handle to a database, a container, or a basic object:
	<ul> <li>If <i>near</i> is a database handle, the new copy is stored in the default container of that database.</li> </ul>
	■ If <i>near</i> is a container handle, the new copy is stored in that container.
	<ul> <li>If <i>near</i> is a basic object handle, the new copy is stored in the same container as the referenced basic object. If possible, the copy will be put on the same page as the referenced basic object or on a nearby page.</li> </ul>
	newCopy
	Object reference or handle to set to the new copy.
Returns	Object reference or handle to the new copy.
Discussion	When called without a <i>newCopy</i> parameter, <i>copy</i> allocates a new handle and returns it. Otherwise, <i>copy</i> returns the object reference or handle that is passed to it.
	You can copy basic objects only, not containers. The application must be able to lock the container of the original object for read and the container of the new copy for update.

#### delete\_object

(ODMG) Deletes the referenced persistent object.

void delete\_object();

Discussion This member function is equivalent to ooObj::operator delete.

### getDefaultVers

Finds the default version of the referenced basic object.

Parameters default

Object reference or handle to set to the default version.

Returns oocSuccess if successful; otherwise oocError.

Discussion This member function identifies the genealogy of the referenced basic object, and finds the default version in the genealogy, provided that the referenced object was created *after* the default version was set. (Versions created before a default version is set cannot know about the default version.)

If a default version for the genealogy is found, getDefaultVers sets the specified object reference or handle to reference it. If no default version exists, or if the default version was specified after the referenced version was created, the specified object reference or handle is set to null.

See also <u>setDefaultVers</u>

## getNameObj

Initializes an object iterator to find all objects named in the scope of the referenced persistent object.

Parameters	objI Object iterator for finding the named objects
Returns	oocSuccess if successful; otherwise oocError.
Discussion	The application must be able to obtain a read lock on the hashed container used by the scope object.

### getNameScope

Initializes an object iterator to find all scope objects in the federated database that define a scope name for the referenced persistent object.

 Parameters
 obj1

 Object iterator for finding the scope objects.

 Returns
 oocSuccess if successful; otherwise oocError.

 Discussion
 The application must be able to obtain read locks on all the hashed containers used by the scope objects.

Objectivity/C++ Programmer's Reference

### getNextVers

Initializes an object iterator to find the next version(s) of the referenced basic object.

	<pre>ooStatus getNextVers( ooItr(ooObj) &amp;nextVers, const ooMode openMode = oocNoOpen) const;</pre>
Parameters	<i>nextVers</i> Object iterator for finding the next version(s).
	<ul> <li>openMode</li> <li>Intended level of access to the versions found by the iterator's next member function:         <ul> <li>oocNoOpen (the default) causes next to set the iterator to the next version without opening it.</li> <li>oocRead causes next to open the next version for read.</li> <li>oocUpdate causes next to open the next version for update.</li> <li>Warning: If versioning is enabled for one or more found objects, specifying oocUpdate means that next will create a new version of each such object.</li> </ul> </li> </ul>
Returns	oocSuccess if successful; otherwise oocError.
Discussion	This member function initializes the specified iterator to find the next version (if linear versioning is used) or versions (if branch versioning is used) of the referenced object.
	If no next version(s) exist, the iterator is set to null, so that invoking the iterator's next member function will return oocFalse.
See also	getPrevVers

### getObjName

Gets the name defined in the specified scope for the referenced persistent object.

```
char *getObjName(
    const ooHandle(ooObj) &scope) const;
```

#### Parameters scope

Handle to the scope object that defines the name scope to be searched. The scope object can be the federated database, a database, a persistent container, a persistent basic object, or an autonomous partition.

Returns	Pointer to a string containing the scope name. If the referenced object does not have a scope name in the specified scope, the returned pointer is null.
Discussion	The string is statically allocated by the member function and overwritten with each invocation. You should make a local copy of the returned string if you intend to use it later in the application.
	The application must be able to get a read lock on the hashed container used by the scope object.
See also	lookupObj

### getPrevVers

Finds the previous version of the referenced basic object.

	ooStatus getPrevVers( ooRefHandle(ooObj) &previous) const;
Parameters	<i>previous</i> Object reference or handle to set to the previous version.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	If a previous version is found, getPrevVers sets the specified object reference or handle to reference it. Otherwise, the specified object reference or handle is set to null.
See also	getNextVers

### getVersStatus

Gets the current versioning mode of the referenced basic object.

ooVersMode getVersStatus() const;

Returns One of the following constants:

- oocNoVers—the versioning behavior of the object is disabled.
- oocLinearVers—linear versioning is enabled for the object.
- oocBranchVers—branch versioning is enabled for the object.

See also <u>setVersStatus</u> <u>ooVersMode</u> global type

is_null	
	(ODMG) Tests whether this object reference or handle is null.
	<pre>int is_null();</pre>
Returns	Non-zero integer (true) if this object reference or handle is null; otherwise 0 (false).
isNull	
	Tests whether this object reference or handle is null.
	ooBoolean isNull() const;
Returns	oocTrue if this object reference or handle is null; otherwise oocFalse.
Discussion	You can use this member function as an alternative to comparing this handle or object reference to 0 with operator==.
See also	operator==
isValid	
	Tests whether this object reference or handle is valid—that is, whether it references an existing Objectivity/DB object.
	ooBoolean isValid() const;
Returns	occTrue if this object reference or handle references an existing Objectivity/DB object; occFalse if this object reference or handle is null or has a stale object identifier, or if the application cannot obtain a read lock on the container and database to be checked.
Discussion	You can use isValid to determine whether it is safe to use an object reference or handle that was set in a previous transaction. Such an object reference or handle still retains its reference to an Objectivity/DB object; however, between transactions, that reference may have become invalid (for example, because another process has deleted the referenced object).
Νοτε	isValid checks only for the existence of a referenced object, but does <i>not</i> check whether the class of the referenced object corresponds to the class of the object reference or handle. An object reference or handle can have a valid but incorrect object identifier after another process has deleted the originally referenced object and created another object in the same location, where it is addressible by the same object identifier. Therefore, if isValid returns oocTrue, you can verify the

	referenced object's class (for example, using the <u>typeN</u> member function), but you must write application-specific code to verify the identity of the referenced object.
	If your purpose is simply to test whether an object reference or handle has been initialized, it is more efficient to use $\underline{isNull}$ , which performs its test entirely in memory without having to access files on disk.
lock	
	Explicitly locks the referenced persistent object; propagates locks along associations that have lock propagation enabled.
	<pre>ooStatus lock(const ooLockMode lockMode) const;</pre>
Parameters	<ul> <li><i>lockMode</i></li> <li>Type of lock to request:</li> <li>Specify oocLockRead to request a read lock.</li> <li>Specify oocLockUpdate to request an update lock.</li> </ul>
Returns	oocSuccess if all requested locks are obtained; otherwise oocError.
Discussion	<ul> <li>Objectivity/DB operations request and obtain locks implicitly as they are needed. You use this member function to obtain a lock explicitly when:</li> <li>You want to reserve access to an object in advance—for example, before</li> </ul>
	<ul> <li>You want to lock an entire composite object (a group of associated objects whose associations have lock propagation enabled). When you explicitly lock an object with such associations, the associated destination objects are locked as well. Lock propagation is enabled by <u>behavior specifiers</u> in an association's definition.</li> </ul>
	Because containers are the fundamental unit of locking within Objectivity/DB, locking a basic object causes its container to be locked. This effectively locks all the basic objects in the same container.
	Whenever a lock is requested on a container, Objectivity/DB applies the transaction's <u>concurrent access policy</u> to determine whether the requested lock is compatible with other existing locks. An error is signaled if a requested lock cannot be obtained.
See also	<u>lockNoProp</u> <u>ooLockMode</u> global type

### lockNoProp

Explicitly locks the referenced persistent object, without propagating locks to associated destination objects.

ooStatus lockNoProp(const ooLockMode lockMode) const;

Parameters lockMode

Type of lock to request:

- oocLockRead requests a read lock.
- oocLockUpdate requests an update lock.

Returns oocSuccess if the requested lock is obtained; otherwise oocError.

Discussion Objectivity/DB operations request and obtain locks implicitly as they are needed. You use this member function to obtain a lock explicitly when you want to reserve access to an object in advance, but you do *not* want to lock any associated destination objects, even along associations that have lock propagation enabled.

Because containers are the fundamental unit of locking within Objectivity/DB, locking a basic object causes its container to be locked. This effectively locks all the basic objects in the same container.

Whenever a lock is requested on a container, Objectivity/DB applies the transaction's <u>concurrent access policy</u> to determine whether the requested lock is compatible with other existing locks. An error is signaled if a requested lock cannot be obtained.

See also <u>lock</u> <u>ooLockMode</u> global type

### lookupObj

Finds the persistent object with the specified scope name (or the basic object matching the specified key structure) within the specified scope, and sets this object reference or handle to reference the found object.

 ooStatus lookupObj ( const ooHandle(ooObj) &scope, const char \*scopeName, const ooMode openMode = oocRead);
 ooStatus lookupObj ( const ooHandle(ooObj) &scope, const ooKey &keyStruct, const ooMode openMode = oocRead);

#### Parameters scope Handle to an object that defines the scope of the lookup: (Variant 1) When the lookup is by scope name, *scope* specifies the scope object that defines the name scope to search. *scope* can reference the federated database, a database, a persistent container, a persistent basic object, or an autonomous partition. • (Variant 2) When the lookup is by key structure, *scope* specifies the container to search. *scope* can reference the container itself or the database whose default container is to be searched. scopeName Scope name to look up in the scope specified by *scope*. openMode Intended level of access to the found object: Specify oocRead (the default) to open the object for read. Specify oocUpdate to open the object for update. (Variant 1 only.) Specify oocNoOpen to set this object reference or handle to the object without opening it. oocNoOpen is valid only for scope-name lookup, because scope-named objects can be found without being opened, whereas keyed objects must be opened during the search. keyStruct Key structure specifying the key field and key field value to match. Returns oocSuccess if an object is found; otherwise oocError. Discussion The application must be able to obtain a read lock on the hashed container used by the scope object. See also getObiName move Moves the referenced basic object to a different container. ooStatus move(const ooHandle(ooObj) &target);

#### Hand

target

Handle of the object with which to cluster the basic object. *target* may be a handle to a database, a container, or a persistent basic object:

■ If *target* is a database handle, the basic object is moved to the default container of that database.

Parameters

	<ul> <li>If <i>target</i> is a container handle, the basic object is moved to that container.</li> </ul>
	<ul> <li>If <i>target</i> is a basic object handle, the moved basic object is put into the same container as the specified target object. If possible, the moved object will be put on the same page as the target object or on a nearby page.</li> </ul>
Returns	oocSuccess if successful; otherwise oocError.
Discussion	You can move basic objects only, not containers. The application must be able to obtain an update lock on the object's container and the container to which the object will be moved.
Warning	This member function changes the object identifier of the moved object, and all references containing the old identifier become invalid. When you move an object you should, within the same transaction, update references to the object within all relevant attributes, persistent collections, name scopes, indexes, and unidirectional associations. (Objectivity/DB automatically maintains referential integrity for bidirectional associations). See "Moving a Basic Object" in the Objectivity/C++ programmer's guide for further information.

### nameObj

Names the referenced persistent object in the specified scope.

```
ooStatus nameObj(
    const ooHandle(ooObj) &scope,
    const char *name) const;
```

#### Parameters scope

Handle to the scope object that defines the name scope for *name*. The scope object can be the federated database, a database, a persistent container, a persistent basic object, or an autonomous partition. If *scope* references a persistent container, that container must be hashed. If *scope* references a persistent basic object, that object must be in a hashed container.

name

Scope name to assign. This name:

- Must be a null-terminated string that can contain any non-null character.
- Must be unique within the name scope defined by *scope*.
- May contain up to 487 characters.

Returns oocSuccess if successful; otherwise oocError.

Discussion	You may give an object only a single name within a given scope. If the object already has a scope name in the specified scope, you must remove that name with <u>unnameObj</u> before you can assign a new name.
	The application must be able to obtain an update lock on the hashed container used by the scope object.
See also	<u>lookupObj</u>
open	
	Explicitly opens the referenced persistent object, preparing the object for the specified level of access.
	<pre>ooStatus open(const ooMode openMode = oocRead);</pre>
Parameters	<ul> <li>openMode         Intended level of access to the opened object:         Specify oocRead (the default) to open the object for read. This implicitly requests a read lock on the relevant container.         Specify oocUpdate to open the object for update (read and write). This implicitly requests an update lock on the relevant container.     </li> </ul>
Returns	oocSuccess if successful: otherwise oocError.
Discussion	<ul> <li>Opening a basic object makes it available to your application by:</li> <li>Implicitly locking the basic object's container for read or update, as specified by <i>openMode</i>.</li> <li>Obtaining a representation of the basic object in memory, either by fetching the page(s) containing the object from the database or by reusing an existing memory representation that is guaranteed current.</li> <li>(If open is called on a handle) Providing this handle with a memory pointer to the opened object.</li> <li>Opening a container makes it available to an application by:</li> <li>Implicitly locking the container for read or update, as specified by <i>openMode</i>.</li> <li>Obtaining a representation of the container in memory, as for basic objects. The representation for a container includes its page map and any persistent data (if the container is an instance of an application-defined container class). Opening a container does not open any of the basic objects in it.</li> <li>Opening an object for update additionally marks it as modified, causing it to be</li> </ul>
	written to the database when the transaction commits, whether or not the object

was actually modified. You must be in an update transaction to open an object

for update. If necessary, you can promote a read transaction to an update transaction by promoting the open mode of the federated database.

The open operation fails if the relevant container cannot be locked—for example, due to a lock conflict. Objectivity/DB applies the transaction's <u>concurrent access</u> <u>policy</u> to determine whether the requested lock is compatible with other existing locks on the container. Once a lock is obtained, it is kept until the transaction either commits or aborts.

You rarely need to explicitly open an object for read because accessing the object through an object reference or handle accomplishes this implicitly (see <a href="mailto:operator->">operator-></a>). You open an object explicitly when:

- You require update access so you can modify the object.
- You want to reserve either read or update access to the object in advance—for example, before starting a complex operation.

If the referenced object is a basic object for which versioning is enabled (see <u>setVersStatus</u>), opening it for update causes a new version to be created, and sets this object reference or handle to reference the new version.

See also <u>update</u>

### openMode

Tests whether this handle is open, and, if so, gets the current level of access to the referenced basic object.

ooMode openMode() const;

Returns

One of the following constants:

- oocNoOpen—this handle is not open in this transaction (although the referenced basic object may be open through another handle).
- oocRead—this handle is open, and the referenced basic object is open for read in this transaction.
- oocUpdate—this handle is open, and the referenced basic object is open for update in this transaction.

### print

Prints the object identifier of the referenced Objectivity/DB object.

void print(FILE \*outputFile = stdout) const;

Parameters	outputFile		
	Pointer to the file in which to print the object identifier. The default is standard output.		
Discussion	The object identifier for a basic object or container is printed in $\#D$ - <i>C</i> - <i>P</i> - <i>S</i> format, which identifies the database ( <i>D</i> ), container ( <i>C</i> ), logical page number ( <i>P</i> ), and logical slot number ( <i>S</i> ) of the object—for example, $\#2$ -3-3-12. (For a container, the page and slot numbers refer to the location of the container object itself.) The object identifier for a database, autonomous partition, or federated database is printed as <i>I</i> -0-0-0, where <i>I</i> is the integer identifier of the object—for example, $\#2$ -0-0-0.		
ptr			
	( <i>ODMG</i> ) Returns a C++ pointer to the referenced persistent object.		
	ooObj *ptr();		
Returns	Pointer to the referenced basic object or container.		
Discussion	You use this member function to obtain a pointer to a basic object or container—for example, to pass to a function that accepts a pointer instead of a handle or object reference.		
	If ptr is called on an object reference, the referenced persistent object is opened for update. If ptr is called on a handle, the referenced persistent object is opened for read.		
	Warning: The returned pointer is guaranteed valid for only a limited time:		
	<ul> <li>If ptr is called on an object reference, the returned pointer is valid and the persistent object is pinned in memory until the end of the transaction.</li> </ul>		
	<ul> <li>If ptr is called on a handle, the returned pointer is valid only as long as the handle exists, remains open, and references the same persistent object (equivalent to <u>operator ooObj*</u>).</li> </ul>		
	An application generally relies on handles to provide memory management for persistent objects, and avoids the explicit use of pointers to such objects. On occasion, explicit use of pointers is required for performance reasons or for compatibility with functions that are indifferent to persistence (for example, in legacy code or a third-party library). However, you should be careful when manipulating persistent objects through pointers:		
	<ul> <li>Pointers extracted from handles become invalid if the handles change or go</li> </ul>		

out of scope.

Pointers extracted from object references can cause the Objectivity/DB cache to run out of memory if too many objects are pinned until the end of the transaction.

You should not use the returned pointer in other persistence operations (for example, do not pass it to any Objectivity/C++ member function other than the overloaded <code>operator delete</code>).

#### set\_container

Provides this object reference or handle with container information so you can assign a short object reference to it.

Parameters object

Object reference or handle to a basic object or container.

Returns oocSuccess if successful; otherwise oocError.

- Discussion You use this member function to set a null object reference or handle before assigning a short object reference to it. A short object reference specifies just the lower half of an object identifier (corresponding to the object's logical page and slot numbers); set\_container supplies the upper half of the object identifier (corresponding to the database and container), and sets the page and slot components to null. That is:
  - If you specify a container, this object reference or handle is set to the upper half of the object identifier for the container.
  - If you specify a basic object, this object reference or handle is set to the upper half of the object identifier for the object.
  - **WARNING** After set\_container is invoked, this object reference or handle is not valid for referencing either a basic object or a container, because it contains only partial identifier information.

### setDefaultVers

Sets the referenced basic object as the default version of its genealogy.

ooStatus setDefaultVers() const;

Returns oocSuccess if successful; otherwise oocError.

DiscussionIf the referenced basic object does not already belong to a genealogy, this member<br/>function implicitly creates one (an instance of ooGeneObj) and adds the<br/>referenced object to it. All subsequently created versions of the same object are<br/>automatically added to this genealogy.If you want all versions of an object to belong to a genealogy, you should set the<br/>object to be the initial default version as soon as you enable versioning for it. You<br/>can use this member function to reset the default version of a genealogy at any<br/>time.See alsogetDefaultVers

### setVersStatus

Enables or disables versioning for the referenced basic object by setting the object's versioning mode.

ooStatus setVersStatus(const ooVersMode versMode) const;

#### Parameters versMode

Versioning behavior to set for the referenced object:

- Specify oocNoVers to disable versioning for the object.
- Specify oocLinearVers to enable linear versioning for the object. This allows exactly one new version to be created from the object.
- Specify oocBranchVers to enable branch versioning for the object. This allows any number of new versions to be created from the object.
- Returns oocSuccess if successful; otherwise oocError.
- Discussion When versioning is enabled for a basic object, a new version of the object is created each time the object is opened for update. By default, the new version is a bit-wise copy of the opened object. An application can customize the copy operation for versioning instances of a basic-object class by overriding the ooNewVersInit member function in that class.

```
See also <u>getVersStatus</u>
ooObj::<u>ooNewVersInit</u>
<u>ooVersMode</u> global type
```

### sprint

Returns a string containing the object identifier of the referenced Objectivity/DB object.

```
char *sprint(char *buffer = 0) const;
```

Parameters	buffer	
	String in which to return the object identifier. If you omit this parameter, sprint statically allocates a new string.	
Returns	String representing the object identifier.	
Discussion	The object identifier for a basic object or container is printed in $\#D$ - <i>C</i> - <i>P</i> - <i>S</i> format which identifies the database ( <i>D</i> ), container ( <i>C</i> ), logical page number ( <i>P</i> ), and logical slot number ( <i>S</i> ) of the object—for example, $\#2$ -3-3-12. (For a container, t page and slot numbers refer to the location of the container object itself.) The object identifier for a database, autonomous partition, or federated database is printed as <i>I</i> -0-0-0, where <i>I</i> is the integer identifier of the object—for example, #2-0-0-0.	
	Each successive invocation of this member function without the parameter overwrites the statically allocated string. You should make a local copy of the returned string if you intend to use it later in the application.	
typeN		
	Gets the type number of the class of the referenced Objectivity/DB object.	
	ooTypeNumber typeN() const;	
Returns	Type number for the referenced basic object, container, database, federated database, or autonomous partition.	
Discussion	Every class in the $000$ bj inheritance hierarchy has a uniqe type number that identifies the class within a federated-database schema. This member function gets the type number of the referenced object's class.	
See also	<u>ooTypeN global macro</u> <u>ooTypeNumber</u> global type ooObj:: <u>ooGetTypeN</u>	
typeName		

Gets the name of the class of the referenced Objectivity/DB object.

char \*typeName() const;

Returns String containing the class name of the referenced basic object, container, database, federated database, or autonomous partition.

WARNING	Do not modify the returned string in any manner. Doing so may result in unexpected program errors. This string is used internally by Objectivity/DB.		
See also	<u>ooTypeN</u> global macro ooObj:: <u>ooGetTypeName</u>		
unnameOb	oj		
	Deletes the name in the specified scope for the referenced persistent object.		
	ooStatus unnameObj( const ooHandle(ooObj) & <i>scope,</i> const char * <i>name</i> = 0) const;		
Parameters	<i>scope</i> Handle to the scope object that defines the scope name to be deleted. The scope object can be the federated database, a database, a persistent container, a persistent basic object, or an autonomous partition.		
	name This parameter is ignored.		
Returns	oocSuccess if successful; otherwise oocError.		
Discussion	The application must be able to obtain an update lock on the persistent object, the scope object, and the hashed container used by the scope object.		
See also	<u>lookupObj</u> , <u>nameObj</u>		
update			
	Opens the referenced persistent object for update access.		
	ooStatus update();		
Returns	oocSuccess if successful; otherwise oocError.		
Discussion	This member function is equivalent to calling <u>open(oocUpdate)</u> .		

# ooShortRef(appClass) Class

Inheritance: **ooShortRef(ooObj)->ooShortRef(***appClass***)** 

The non-persistence-capable class <code>ooShortRef(appClass)</code> represents a *short object reference* to an instance of the application-defined basic-object class <code>appClass</code>.

See:

• "Reference Index" on page 633 for a list of member functions

To use the ooShortRef(*appClass*) class, you must include and compile with files generated by the DDL processor, as described in "Obtaining Generated Class Definitions" on page 632.

## About Short Object References

When an application defines a basic-object class appClass and adds it to the federated-database schema, the DDL processor generates the corresponding short-object-reference class ooShortRef(appClass). Short object references enable you to save space when storing persistent references to basic objects in attributes and associations. Short object references are alternatives to standard object references (instances of ooRef(appClass)); both kinds of object reference store the object identifier of a referenced basic object, with the following difference:

- A standard object reference stores all components of the object identifier, including the database, container, logical page, and logical slot numbers.
- A short object reference truncates the object identifier, storing just its lower half (the logical page and slot numbers).

Short object references occupy about half the space of standard object references, so they can be useful when disk usage is a concern for a federated database that must maintain a large number of references to basic objects. However, a

truncated object identifier provides only enough information to locate an object within its container. Consequently, a short object reference must used as an attribute or in an association of a persistent object, which determines the necessary container and database information.

You can use a short object reference of type <code>ooShortRef(ooObj)</code> to reference a basic object of any type. Normally, however, a basic object of a particular type is referenced by a short object reference of the corresponding type—that is, a basic object of an application-defined class <code>appClass</code> is referenced by instances of <code>ooShortRef(appClass)</code>.

### **Obtaining Generated Class Definitions**

To use the <code>ooShortRef(appClass)</code> class, you must include either the primary header file or the references header file generated by the DDL processor for <code>appClass</code>. Thus, if <code>appClass</code> is defined in the DDL file <code>classDefFile.ddl</code>, you must include one of the following files:

- The primary header file *classDefFile*.h
- The references header file classDefFile\_ref.h

Furthermore, you must compile the method implementation file *classDefFile\_ddl.cxx* with your application code files.

For more information about DDL-generated files and how to use them, see the Objectivity/C++ Data Definition Language book.

#### When appClass is a Template Class

When *appClass* is a persistence-capable template class with multiple parameters, the name of the generated short-object-reference class contains the symbol OO\_COMMA to separate the template parameters. For example, for a persistence-capable template class <code>Example<Float</code>, <code>Node></code>, the generated class is <code>ooShortRef(Example<Float OO\_COMMA Node>)</code>. This is because the macro syntax of the generated class name interprets embedded commas as separators between the as macro parameters instead of as separators between the template parameters.

# **Working With Short Object References**

An application normally creates a short object reference implicitly as a data member or a short inline association of a persistent object. An application should not explicitly define subclasses of the <code>ooShortRef(appClass)</code> classes; any necessary subclasses are generated automatically by the DDL processor if the application defines any subclasses of <code>appClass</code>.

A new short object reference is normally *null*—that is, it contains the value 0 instead of an object identifier. The application can then set the short object reference a particular basic object in any of the following ways:

- By assignment or initialization from a standard object reference or handle. Only the lower half of the object identifier is assigned to the short object reference; the upper half is ignored.
- By assignment or initialization from another short object reference.

Unlike standard object references, short object references cannot be used as smart pointers. That is, you cannot use a short object reference to operate on the referenced basic object or access its members. Rather, the member functions inherited from the <code>ooShortRef(ooObj)</code> base class are limited to operations that apply to short object references themselves, such as comparing, testing, and printing.

You can assign a short object reference to a standard object reference or handle. The upper half of the object identifier is typically supplied by calling <u>set\_container</u> on the standard object reference or handle.

## **Reference Index**

The following table summarizes just the member functions that are redefined by *ooRefHandle(appClass)* to provide type-specific parameters. For descriptions of inherited member functions, see the <u>ooShortRef(ooObj)</u> classes (page 637).

<u>ooShortRef(appClass)</u>	Default constructor that constructs a null short object reference.
<u>ooShortRef(appClass)</u>	Constructs a new short object reference that references the same basic object as the specified short object reference, standard object reference, handle, or pointer.
<u>operator=</u>	Assignment operator; sets this short object reference to the same basic object as the specified short object reference, standard object reference, handle, or pointer.

## **Constructors and Destructors**

## ooShortRef(appClass)

Default constructor that constructs a null short object reference.

```
ooShortRef(appClass)();
```

## ooShortRef(appClass)

Constructs a new short object reference that references the same basic object as the specified short object reference, standard object reference, handle, or pointer.

	<pre>1. ooShortRef(appClass)(         const ooShortRef(appClass) &amp;object);</pre>		
	<pre>2. ooShortRef(appClass)( const ooRefHandle(appClass) &amp;object);</pre>		
	<pre>3. ooShortRef(appClass)(</pre>		
Parameters	object Short object reference, standard object reference, or handle to an instance of appClass.		
	<ul> <li>objectP</li> <li>Pointer to an instance of appClass. The pointer may not be the result of operator new on appClass. Instead, the pointer must be the result of using operator appClass* on a handle or ptr on an object reference or handle earlier in the same transaction. If the specified pointer was originally extracted from a handle, that handle must still exist and reference the same object.</li> </ul>		
Discussion	If you use Variant 2 to construct a short object reference from a standard object reference or handle, the database and container portion of the object identifier are discarded.		
	Variant 3 allows you to resume persistence operations on a basic object after manipulating it through a pointer. The use of this variant should be rare, however, because pointers are not normally used to manipulate basic objects.		

# **Operators**

### operator=

Assignment operator; sets this short object reference to the same basic object as the specified short object reference, standard object reference, handle, or pointer.

1.	ooShort	Ref( <i>appClass</i> )	&operat	or=(
	const	ooShortRef(ap	<i>pClass</i> )	&object);

 

#### Parameters object

Short object reference, standard object reference, or handle to an instance of *appClass*.

#### objectP

0, or a nonnull pointer to an instance of *appClass*. The pointer may *not* be the result of operator new on *appClass*. Instead, the pointer must be the result of using <u>operator appClass</u>\* on a handle *or* <u>ptr</u> on an object reference or handle earlier in the same transaction. If the specified pointer was originally extracted from a handle, that handle must still exist and reference the same object.

- Returns This short object reference.
- Discussion If you use Variant 2 to assign a standard object reference or handle to a short object reference, the database and container portion of the object identifier are discarded.

Variant 3 allows you to set this short object reference to null. Otherwise, assignment-from-pointer enables you to resume persistence operations on a basic object after manipulating it through a pointer. The use of this variant should be rare, however, because pointers are not normally used to manipulate basic objects. Operators

# ooShortRef(ooObj) Class

Inheritance: **ooShortRef(ooObj)** 

The non-persistence-capable class <code>ooShortRef(ooObj)</code> is the base class for all classes of *short object references*.

See:

- "Reference Summary" on page 638 for an overview of member functions
- "Reference Index" on page 639 for a list of member functions

## About Short Object References

Short object references enable you to save space when storing persistent references to basic objects in attributes and associations. Short object references are alternatives to standard object references (instances of ooRef(ooObj) and its derived classes); both kinds of object reference store the object identifier of a referenced basic object, with the following difference:

- A standard object reference stores all components of the object identifier, including the database, container, logical page, and logical slot numbers.
- A short object reference truncates the object identifier, storing just its lower half (the logical page and slot numbers).

Short object references occupy about half the space of standard object references, so they can be useful when disk usage is a concern for a federated database that must maintain a large number of references to basic objects. However, a truncated object identifier provides only enough information to locate a basic object within its container. Consequently, a short object reference must be used as an attribute or in an association of a persistent object, which determines the necessary container and database information:

■ If the referencing persistent object is a basic object, the referenced object is assumed to be in the same container.

 If the referencing persistent object is a container, the referenced object is assumed to be in that container.

You can use a short object reference of type <code>ooShortRef(ooObj)</code> to reference a basic object of any type. Normally, however, a basic object of a particular type is referenced by a short object reference of the corresponding type—that is, a basic object of an application-defined class <code>appClass</code> is referenced by instances of <code>ooShortRef(appClass)</code>.

# **Working With Short Object References**

An application normally creates a short object reference implicitly as a data member or a short inline association of a persistent object. An application should not explicitly define subclasses of the <code>ooShortRef(ooObj)</code> classes; any necessary subclasses are generated automatically by the DDL processor if the application defines any subclasses of <code>ooObj</code>.

A new short object reference is normally *null*—that is, it contains the value 0 instead of an object identifier. The application can then set the short object reference to reference a particular basic object in any of the following ways:

- By assignment or initialization from a standard object reference or handle. Only the lower half of the object identifier is assigned to the short object reference; the upper half is ignored.
- By assignment or initialization from another short object reference.

Unlike standard object references, short object references cannot be used as smart pointers. That is, you cannot use a short object reference to operate on the referenced basic object or to access its members. Rather, the member functions defined by the short object reference classes are limited to operations that apply to short object references themselves, such as comparing, testing, and printing.

You can assign a short object reference to a standard object reference or handle. The upper half of the object identifier is typically supplied by calling <u>set\_container</u> on the standard object reference or handle.

# **Reference Summary**

Creating	<u>ooShortRef(ooObj)</u>
Setting	<u>operator=</u>

Comparing	<u>operator==</u> operator!=
Getting Information	<u>print</u> <u>sprint</u>
Testing	<u>isNull</u> operator int

# **Reference Index**

<u>isNull</u>	Tests whether this short object reference is null.
<u>ooShortRef(ooObj)</u>	Default constructor that constructs a null short object reference.
<u>ooShortRef(ooObj)</u>	Constructs a new short object reference that references the same basic object as the specified short object reference, standard object reference, handle, or pointer.
<u>operator=</u>	Assignment operator; sets this short object reference to the same basic object as the specified short object reference, standard object reference, handle, or pointer.
<u>operator==</u>	Equality operator; tests whether this short object reference has the same value as the specified item.
<u>operator!=</u>	Inequality operator; tests whether this short object reference has a different value from the specified item.
<u>operator int</u>	Conversion operator that tests whether this short object reference is null.
print	Prints the object identifier of the referenced basic object.
sprint	Returns a string containing the object identifier of the referenced basic object.

# Constructors

## ooShortRef(ooObj)

Default constructor that constructs a null short object reference.

```
ooShortRef(ooObj)();
```

## ooShortRef(ooObj)

Constructs a new short object reference that references the same basic object as the specified short object reference, standard object reference, handle, or pointer.

	<pre>1. ooShortRef(ooObj)(     const ooShortRef(ooObj) &amp;object);</pre>		
	<pre>2. ooShortRef(ooObj)(     const ooRefHandle(ooObj) &amp;object);</pre>		
	<pre>3. ooShortRef(ooObj)(</pre>		
Parameters	object Short object reference, standard object reference, or handle to an instance of		
	<pre>objectP Pointer to an instance of appClass. The pointer may not be the result of operator new on appClass. Instead, the pointer must be the result of using operator appClass* on a handle or ptr on an object reference or handle earlier in the same transaction. If the specified pointer was originally extracted from a handle, that handle must still exist and reference the same object.</pre>		
Discussion	If you use Variant 2 to construct a short object reference from a standard object reference or handle, the database and container portion of the object identifier are discarded.		
	Variant 3 allows you to resume persistence operations on a basic object after manipulating it through a pointer. The use of this variant should be rare, however, because pointers are not normally used to manipulate basic objects.		

# Operators

### operator=

Assignment operator; sets this short object reference to the same basic object as the specified short object reference, standard object reference, handle, or pointer.

1.	ooShort	Ref(ooObj)	&operat	or=(
	const	ooShortRef	(oo0bj)	&object);

 

#### Parameters object

Short object reference, standard object reference, or handle to an instance of *appClass*.

#### objectP

0, or a nonnull pointer to an instance of *appClass*. The pointer may *not* be the result of operator new on *appClass*. Instead, the pointer must be the result of using <u>operator appClass</u>\* on a handle *or* <u>ptr</u> on an object reference or handle earlier in the same transaction. If the specified pointer was originally extracted from a handle, that handle must still exist and reference the same object.

- Returns This short object reference.
- Discussion If you use Variant 2 to assign a standard object reference or handle to a short object reference, the database and container portion of the object identifier are discarded.

Variant 3 allows you to set this short object reference to null. Otherwise, assignment-from-pointer enables you to resume persistence operations on a basic object after manipulating it through a pointer. The use of this variant should be rare, however, because pointers are not normally used to manipulate basic objects.

#### operator==

Equality operator; tests whether this short object reference has the same value as the specified item.

1.	<pre>ooBoolean operator==(    const ooShortRef(ooObj) &amp;compare) const;</pre>
2.	<pre>ooBoolean operator==(    const ooRefHandle(ooObj) &amp;compare) const;</pre>
3.	ooBoolean operator==( const ooObj * <i>compare</i> ) const;
4.	ooBoolean operator==(int zero) const;

Parameters compare

Short object reference, object reference, handle, or pointer to a basic object.

	<i>zero</i> Literal 0. This value allows you to use operator== as an alternative for <u>isNull</u> .
Returns	Variants 1, 2, and 3 return oocTrue if this short object reference references the same object as <i>compare</i> ; otherwise oocFalse. Variant 4 returns oocTrue if this short object reference is null; otherwise oocFalse.
Discussion	Variants 1, 2, and 3 test whether this object reference or handle references the same basic object as <i>compare</i> . If you specify a standard object reference or handle, this operator considers only the lower half of the object identifier, ignoring the database and container information. The comparison is most meaningful when the object referenced by <i>compare</i> resides in the same container as the object referenced by this short object reference.
	Variant 4 tests whether this object reference or handle is null.
See also	<u>isNull</u> (as an alternative to variant 4)

### operator!=

Inequality operator; tests whether this short object reference has a different value from the specified item.

<pre>1. ooBoolean operator!=(     const ooShortRef(ooObj) &amp;compare) const;</pre>	
<pre>2. ooBoolean operator!=(     const ooRefHandle(ooObj) &amp;compare) const;</pre>	
<pre>3. ooBoolean operator!=(     const ooObj* compare) const;</pre>	
4. ooBoolean operator!=(int zero) const;	
<i>compare</i> Short object reference, object reference, handle, or pointer to a basic object.	
zero Literal 0.	
Variants 1, 2, and 3 return oocTrue if this short object reference does not reference the same object as <i>compare</i> ; otherwise oocFalse. Variant 4 returns oocTrue if this short object reference is not null; otherwise	

Discussion Variants 1, 2, and 3 test whether this object reference or handle references a different object than *compare*. If you specify a standard object reference or handle, this operator considers only the lower half of the object identifier, ignoring the database and container information. The comparison is most meaningful when the object referenced by *compare* resides in the same container as the object referenced by this short object reference.

Variant 4 tests whether this object reference or handle is nonnull.

### operator int

	Conversion operator that tests whether this short object reference is null.	
	operator int() const;	
Returns	0 if this short object reference is null; otherwise, returns a nonzero integer.	
Discussion	This conversion operator enables you to use a short object reference as the conditional expression in an if or while statement to test whether the short object reference is null.	
Example	<pre>ooShortRef(ooObj) objectR;  // Set objectR to some object if (objectR) {  // Do something interesting if initialization was successful }</pre>	

# **Member Functions**

### isNull

	Tests whether this short object reference is null.
	ooBoolean isNull() const;
Returns	oocTrue if this short object reference is null; otherwise oocFalse.
Discussion	You can use this member function as an alternative to comparing this short object reference to 0 with $operator==$ .
See also	operator==

print	
	Prints the object identifier of the referenced basic object.
	<pre>void print(FILE *outputFile = stdout) const;</pre>
Parameters	outputFile Pointer to the file in which to print the object identifier. The default is standard output.
Discussion	The object identifier is printed in $\#P-S$ format, which identifies the logical page number ( <i>P</i> ) and logical slot number ( <i>S</i> ) of the basic object—for example, $\#3-12$ .
sprint	
	Returns a string containing the object identifier of the referenced basic object.
	<pre>char *sprint(char *buffer = 0) const;</pre>
Parameters	buffer
	String in which to return the object identifier. If you omit this parameter, sprint statically allocates a new string.
Returns	String representing the object identifier.
Discussion	The object identifier is printed in $\#P-S$ format, which identifies the logical page number ( <i>P</i> ) and logical slot number ( <i>S</i> ) of the basic object—for example, $\#3-12$ .
	Each successive invocation of this member function without the parameter overwrites the statically allocated string. You should make a local copy of the returned string if you intend to use it later in the application.

# ooString(N) Class

#### Inheritance: **ooString(N)**

The non-persistence-capable class ooString(N) represents an *optimized string* based on a VArray of characters and a fixed character array of length *N*.

See:

- "Reference Summary" on page 647 for an overview of member functions
- "Reference Index" on page 648 for a list of member functions
- **NOTE** The name ooString(N) is a macro that expands to a template class whose parameter is N.

## **About Optimized Strings**

An optimized string is a character string of any length that can be stored in a persistent object. Although an optimized string is able to store any number of characters, it is optimized for strings that are less than a particular length. You can use optimized strings anywhere in an application; however, their primary purpose is to serve as string attributes of persistence-capable classes (in place of C++ char \* strings, which cannot be stored persistently). You can convert transparently between an optimized string and a const char \* string, enabling optimized strings to be passed to functions as parameters of type const char \* and vice versa.

#### **Choosing Optimized Strings**

Objectivity/C++ provides two kinds of strings that can be stored persistently—variable-size strings and optimized strings. If you know that the strings to be stored are generally less than N characters long, you should choose

optimized strings of class oostring(N). Otherwise, you should use variable-size strings if you cannot predict the lengths of the strings to be stored or if you know these lengths will vary widely.

### **Structure of Optimized Strings**

An optimized string of class oostring(N) contains a VArray of characters and a fixed character array whose length is the integer *N*, where *N* > 0. If an optimized string contains fewer than *N* characters, these characters are stored in the fixed array, and the vector portion of the VArray is not allocated. Otherwise, if the number of characters is greater than or equal to *N*, the vector is allocated and all of the characters are stored in it.

An optimized string always contains space for the fixed array (whether or not it is used) and for the VArray's reference to its vector (whether or not the vector is actually allocated). When the number of characters is N or greater, space for the vector is added.

The fixed portion of the optimized string is embedded in the containing persistent object; the vector, if any, is external to the object and may be relocated by certain operations. Elements in the vector are guaranteed contiguous within virtual memory.

### **Efficient Use of Optimized Strings**

An optimized string allows you to avoid the overhead of VArrays when operating on strings whose size you can predict, and still have the flexibility to use VArrays if an occasional large string occurs. For example, if you are defining a class that contains mostly strings of less than 8 bytes, you might want to use the ooString(8) class. This class provides maximum efficiency for most of your strings (avoiding VArray overhead when the VArray is not needed) and uses VArray for the occasional occurrence of strings of length greater than 7. Furthermore, performance is better for the shorter strings whose characters are directly embedded in the containing persistent object; when a VArray is used, a dereference operation is performed to find the vector containing the characters.

You should choose a value for N (the length of the fixed character array in the class) so that a high percentage (for example, 90%) of the strings in the class have a length less than N. It is preferable that N be an even number. N must take into account the terminating null needed by C++ strings.

An optimized string allocates the fixed character array whether or not it is used. If the number N is not properly chosen, then the fixed part of the optimized string could be too big to be fully utilized or be too small to store the string in most cases. In either case, significant storage space may be wasted. You should perform an analysis of usage patterns before selecting N.

# **Working With Optimized Strings**

Like instances of any other non-persistence-capable class, optimized strings are not independently persistent. However, when an optimized string is an attribute of a persistent object, it is saved in the federated database when the persistent object is saved.

You use public constructors to create optimized strings with 0 or more characters. You can initialize an empty optimized string by assigning a C++ string to it, and you can use the <u>resize</u> member function to grow or truncate the optimized string dynamically. Resizing an optimized string to 0 removes all its characters.

You get each character by its position in the optimized string. Characters are numbered starting with 0; the position number is the character's *index* or *subscript*. Operations on an optimized string verify that any specified indexes are valid based on the string's current size. The <u>length</u> of an optimized string is the number of characters up to the first null character.

Because of the way an optimized string is represented, you cannot get the first character by dereferencing the string; that is, the expression  $*m_YVString$  does not get the first element of  $m_YVString$ . Instead, you can specify the index 0 to the subscript operator (<code>operator[]</code>) to get the first character; alternatively, you can call the <u>head</u> member function to get a pointer to the first character.

Creating	<u>ooString(N)</u>
Assigning	operator=
Type Conversion	operator const char * operator ooVString
Modifying	<u>operator=</u> <u>operator+=</u> <u>resize</u>
Finding Characters	<u>head</u> operator[_]
Getting Information	length
Testing	<u>operator==</u> <u>operator!=</u>

# **Reference Summary**

## **Reference Index**

<u>ooString(N)</u>	Default constructor that constructs a new optimized string whose size is 0.
<u>ooString(N)</u>	Constructs a new optimized string containing a copy of the characters in the specified string.
<u>operator[]</u>	Subscript operator; gets the specified character of this optimized string.
<u>operator+=</u>	Append-to operator; concatenates this optimized string with the specified C++ string.
<u>operator=</u>	Assignment operator; assigns a copy of the specified C++ string to this optimized string.
<u>operator==</u>	Equality operator; tests whether this optimized string matches the specified string.
<u>operator!=</u>	Inequality operator; tests whether this optimized string is different from the specified string.
operator const char *	Conversion operator that accesses this optimized string as an object of type const char *.
operator ooVString	Conversion operator that accesses this optimized string as an object of type ooVString.
head	Gets a pointer to the first character of this optimized string.
<u>length</u>	Gets the number of characters in this optimized string.
resize	Extends or truncates this optimized string to the specified number of characters.

## **Constructors and Destructors**

### ooString(N)

Default constructor that constructs a new optimized string whose size is 0.

ooString(N)();
## ooString(N)

Constructs a new optimized string containing a copy of the characters in the specified string.

	1.	<pre>ooString(N)(const char *p);</pre>
	2.	<pre>ooString(N)(const ooString(N) &amp;s);</pre>
	3.	<pre>ooString(N)(const ooVString &amp;s);</pre>
Parameters	p E	Existing C++ string from which to construct the new optimized string.
	s E r	Existing optimized string or variable-size string from which to construct the new optimized string.
Discussion	If p is	s null or if the length of ${\scriptscriptstyle {\cal S}}$ is 0, an uninitialized string of size 0 is created.

# Operators

## operator[]

Subscript operator; gets the specified character of this optimized string.

char &operator[](const uint32 index) const;

Parameters	index
	Index of the character to get. Specify 0 to get the first character.
Returns	index'th character of this optimized string.
Discussion	An error is reported if the index is not within the allocated size of the optimized string (including the terminating null character).

### operator+=

Append-to operator; concatenates this optimized string with the specified C++ string.

ooString(N) &operator+=(const char \*p);

#### Parameters

C++ string whose characters are to be concatenated.

р

Returns	This optimized string.
Discussion	The concatenation operator adds the characters pointed to by ${\ensuremath{\mathcal{P}}}$ to the end of this optimized string.
	You can use this operator to concatenate another optimized string to this one, because <u>operator const char</u> * automatically converts the string being concatenated to const char *.

### operator=

Assignment operator; assigns a copy of the specified C++ string to this optimized string.

ooString(N) &operator=(const char \*p);

### Parameters

C++ string whose characters are to be assigned.

#### Returns This optimized string.

р

Discussion The assignment operation resizes this optimized string to be the same size as  $_{P}$ , and then copies the characters of  $_{P}$  into this optimized string. Any characters already in this optimized string are overwritten. If  $_{P}$  is null, this optimized string is in effect deleted and replaced with an empty string.

You can use this operator to assign another optimized string to this one, because <u>operator const char \*</u> automatically converts the string being assigned to const char \*.

### operator==

Equality operator; tests whether this optimized string matches the specified string.

- ooBoolean operator==(const char \*p) const;
- 2. ooBoolean operator==(const ooString(N) &s) const;
- 3. ooBoolean operator==(const ooVString &s) const;
- Returns occTrue if every character of this string matches the corresponding character of the other string; otherwise, occFalse.
- Discussion You can compare this optimized string to a C++ string (variant 1), to another optimized string (variant 2) or to a variable-size string (variant 3).

### operator!=

Inequality operator; tests whether this optimized string is different from the specified string.

	<ol> <li>ooBoolean operator!=(const char *p) const;</li> </ol>	
	<ol> <li>ooBoolean operator!=(const ooString(N) &amp;s) const;</li> </ol>	
	3. ooBoolean operator!=(const ooVString &s) const;	
Returns	oocTrue if any character of this string differs from the corresponding charact of the other string; otherwise, oocFalse.	
Discussion	You can compare this optimized string to a C++ string (variant 1), to another optimized string (variant 2), or to a variable-size string (variant 3).	

### operator const char \*

Conversion operator that accesses this optimized string as an object of type const char \*.

operator const char \*() const;

## operator ooVString

Conversion operator that accesses this optimized string as an object of type ooVString.

operator const ooVString() const;

Discussion This operator copies the characters into the resulting variable-size string instead of sharing the same storage.

# **Member Functions**

### head

Gets a pointer to the first character of this optimized string.

char \*head() const;

Returns Pointer to the first character of this optimized string.

length	
	Gets the number of characters in this optimized string.
	<pre>uint32 length() const;</pre>
Returns	Integer number of characters in this string preceding the first null terminating character as computed by strlen. If the string contains no characters, returns 0.
Discussion	The actual number of bytes allocated is at least $length() + 1$ because an extra byte is reserved for the null terminating character. (If the string contains an embedded null character, the bytes beyond that null character are not included in the returned length.)
resize	
	Extends or truncates this optimized string to the specified number of characters.
	<pre>ooStatus resize(const uint32 newLength);</pre>
Parameters	newLength Number of characters this optimized string is to have. Specify 0 to remove all the characters. If newLength is less than N, only the fixed-array portion of the optimized string is used, freeing any storage allocated to the optimized string's vector.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	The actual number of bytes allocated is $newLength + 1$ because an extra byte is reserved for the terminating null character, 0, which is automatically inserted.

# ooTrans Class

#### Inheritance: ooTrans

The non-persistence-capable class ooTrans represents a *transaction object*, which you can use to start and terminate a series of transactions against an Objectivity/DB federated database.

See:

- "Reference Summary" on page 654 for an overview of member functions
- "Reference Index" on page 654 for a list of member functions

(ODMG) The ooTrans class is equivalent to the ODMG standard class d\_Transaction. You can use the name d\_Transaction interchangeably with ooTrans.

## **About Transaction Objects**

An application uses one or more transaction objects to start and stop its transactions. A single-threaded application normally creates a single transaction object; a multithreaded application normally creates one transaction object in each Objectivity context that is to execute transactions. A particular transaction object can be used to start and stop any number of transactions.

You may create additional transaction objects for programming convenience—for example, in each of several local scopes. However, in a given Objectivity context, only one transaction object may be *active* (used to start a transaction) at a time. If you have defined several transaction objects in the same Objectivity context, and you have started a transaction from one of them, you must commit or abort that transaction before starting another transaction, whether from the same or a different transaction object.

Every transaction has an integer identifier that uniquely identifies it to the lock server. A transaction's identifier is assigned when the federated database is opened in that transaction. Administration tools such as oolockmon and oolistwait refer to a transaction using its identifier. Recovery functions use parameters of this type to identify a transaction of interest.

The way you start a transaction determines its <u>concurrent access policy</u>—that is, whether it is a *standard transaction* or a *multiple readers, one writer* (MROW) transaction. The concurrent access policy determines whether a requested read lock is considered compatible with an existing update lock on a container. However, a separate operation (namely, opening the federated database) determines whether the transaction is a *read transaction* or an *update transaction;* see <code>ooRefHandle(ooFDObj)::open</code>.

# **Reference Summary**

Creating a Transaction Object	ooTrans
Controlling Transactions	<u>abort</u> <u>commit</u> <u>commitAndHold</u> <u>start</u>
Getting Information	getID
Testing the Transaction Object	isActive
Object Conversion	upgrade
ODMG Interface	<u>begin</u> <u>checkpoint</u>

# **Reference Index**

abort	Terminates the currently active transaction on this transaction object, and aborts (does not apply) changes to the federated database.
<u>begin</u>	(ODMG) Starts a new transaction on this transaction object.
<u>checkpoint</u>	(ODMG) Checkpoints the currently active transaction on this transaction object.
<u>commit</u>	Terminates the currently active transaction on this transaction object, and commits all changes to the database.

object.
Gets the identifier of the currently active transaction on this transaction object.
Tests whether a transaction has been started on this transaction object.
Default constructor that constructs a new transaction object.
Starts a new transaction on this transaction object.
Identifies the application containing this transaction object as a special-purpose upgrade application for converting objects after schema evolution.

## Constructors

### ooTrans

Default constructor that constructs a new transaction object.

ooTrans();

# **Member Functions**

### abort

Terminates the currently active transaction on this transaction object, and aborts (does not apply) changes to the federated database.

ooStatus abort(ooHandleMode mode = oocHandleToNull);

Parameters mode

Determines what happens to the transaction's open handles after the transaction is aborted:

- Omit this parameter (or specify oocHandleToNull) to set all open handles to null.
- Specify oocHandleToOID to close all open handles, which invalidates their pointers, but preserves the object identifiers they contain. You should test each closed handle for validity before reusing it in another transaction, however, to ensure that the referenced object has not been

	deleted or moved by another process between your process's transactions.
Returns	oocSuccess if successful; otherwise oocError.
begin	
	(ODMG) Starts a new transaction on this transaction object.
	<pre>void begin();</pre>
Discussion	This member function calls $\underline{start}$ with default parameter values.
checkpoint	
	(ODMG) Checkpoints the currently active transaction on this transaction object.
	<pre>void checkpoint();</pre>
Discussion	Checkpointing a transaction terminates it, commits all changes to the database, and implicitly starts a new transaction.
	This member function calls commitAndHold with the default parameter value.
commit	
	Terminates the currently active transaction on this transaction object, and commits all changes to the database.
	ooStatus commit();
Returns	oocSuccess if successful; otherwise oocError.
Discussion	If the federated database is open, this member function closes it.

## commitAndHold

Checkpoints the currently active transaction on this transaction object.

```
ooStatus commitAndHold(
ooDowngradeMode mode = oocNoDowngrade);
```

Parameters mode

Mode in which update locks are to be treated:

• Omit this parameter (or specify oocNoDowngrade) to preserve all locks held by the transaction as is.

	<ul> <li>Specify oocDowngradeAll to change all locks to read locks (MROW read if the transaction is an MROW transaction, and normal read locks otherwise).</li> </ul>
Returns	oocSuccess if successful; otherwise oocError.
Discussion	Checkpointing a transaction terminates it, commits all changes to the database, and implicitly starts a new transaction. By default, all locks acquired during the transaction are preserved as is.
getID	
	Gets the identifier of the currently active transaction on this transaction object.
	ooTransID getID() const;
Returns	Identifier of the current transaction, if the identifier has been assigned; otherwise 0.
Discussion	Every transaction has an integer identifier that uniquely identifies it to the lock server. A transaction's identifier is assigned when the federated database is opened in that transaction. Consequently, you should call getID after you call <code>open</code> on a handle to the federated database. If getID is called before a the transaction's identifier is assigned or after the transaction ends, the value 0 is returned.
	Administration tools such as oolockmon and oolistwait refer to a transaction using its identifier. Recovery functions use parameters of this type to identify a transaction of interest.
isActive	
	Tests whether a transaction has been started on this transaction object.
	ooBoolean isActive();
Returns	oocTrue if this transaction object is active; otherwise oocFalse.
start	
	Starts a new transaction on this transaction object.
	<pre>ooStatus start( ooMode openMode = oocNoMROW, const int32 waitOption = oocTransNoWait, ooIndexMode indexMode = oocInsensitive);</pre>

#### Parameters openMode

<u>Concurrent access policy</u> for the newly started transaction. Objectivity/DB uses this policy to determine whether the locks requested or held by this transaction are compatible with those of other transactions:

- Specify oocNoMROW (the default) to enable the <u>standard</u> concurrent access policy. This policy allows multiple transactions to lock the same container for read, but prevents concurrent read and update locks on a container.
- Specify OOCMROW to enable the <u>MROW</u> concurrent access policy. This policy allows multiple transactions to lock the same container for read while one transaction locks it for update.

#### waitoption

Lock-waiting behavior for the newly started transaction:

- Specify oocTransNoWait (the default) to use the default lock-waiting option currently in effect for the Objectivity context (see the <u>ooSetLockWait</u> global function).
- Specify oocNoWait or 0 to turn off lock waiting for the transaction.
- Specify oocWait to cause the transaction to wait indefinitely for locks.
- Specify a timeout period of *n* seconds for the transaction, where *n* is an integer in the range 1 <= *n* <= 14400. If *n* is less than 0 or greater than 14400, it is treated as oocWait.

Lock waiting does not apply to MROW read transactions. Therefore, any *waitOption* value you specify is ignored when you set the *openMode* parameter to oocMROW.

#### indexMode

Policy for updating indexes when objects of an indexed class are created or key-field values are modified:

- Specify oocInsensitive (the default) to update all applicable indexes automatically when the transaction commits.
- Specify oocSensitive to update all applicable indexes automatically when the next predicate scan is performed in the transaction or, if no scans are performed, when the transaction commits. This allows you to change indexed objects and then scan them in the same transaction using any applicable index. (Note, however, that the transaction must commit before the updates are available to other transactions.)
- Specify oocExplicitUpdate to suppress automatic index updates; the application must update indexes explicitly by calling the <u>ooUpdateIndexes</u> global function after every relevant change.

Returns oocSuccess if successful; otherwise oocError.

Discussion	After you call this member function, the first Objectivity/DB operation of the transaction must be to open the federated database by calling the <u>open</u> member function on a federated-database handle.
upgrade	
	Identifies the application containing this transaction object as a special-purpose <i>upgrade application</i> for converting objects after schema evolution.
	ooStatus upgrade();
Discussion	You must call this member function before starting the first (and only) transaction in an upgrade application. This transaction must be an update transaction (that is, the transaction must open the federated database for update) and it must also call the <i>ooRefHandle</i> (ooFDObj)::: <u>upgradeObjects</u> member function to initiate the upgrade process.
	Object conversion is the process of making existing persistent objects consistent with class definition changes introduced by schema evolution. Certain schema evolution operations affect how instances of a class should be laid out in storage. After you perform such operations, existing objects of the changed classes are rendered out-of-date until they are converted to their new representations.
	In general, affected objects are converted automatically when they are accessed after schema evolution. However, some schema changes require that you convert objects explicitly using an upgrade application to ensure referential integrity.
See also	Chapter 19, "Object Conversion," in the Objectivity/C++ programmer's guide

Member Functions

# ooTreeAdmin Class

Inheritance: ooObj->ooAdmin->ooTreeAdmin

Handle Class: ooHandle(ooTreeAdmin)

```
Object-Reference Class: ooRef(ooTreeAdmin)
```

The persistence-capable class ooTreeAdmin represents tree administrators.

See:

- "Reference Summary" on page 662 for an overview of member functions
- "Reference Index" on page 663 for a list of member functions

To use this class, your application must include the ooCollections.h header file. For UNIX linking information, see *Installation and Platform Notes for UNIX*.

## **About Tree Administrators**

Each ordered collection has a tree administrator that manages the containers used by the collection's internal objects, namely the nodes of the collection's B-tree and the array of references for each node. An ordered collection's tree administrator is created when the collection itself is created.

A tree administrator has two properties that you can set to control when the ordered collection's current node container and the current array container are considered "full."

The maximum nodes per container property specifies how many B-tree nodes can be clustered together in the same container. The default value for this property depends on the federated database's storage page size; it is calculated as:

```
pageSize / 47
```

Because B-tree nodes are small objects, many of them can fit on a single storage page. Because nodes are not updated frequently, many can be clustered in the same container without causing locking problems.

• The *maximum arrays per container* property specifies how many arrays can be clustered together in the same container.

One array fills up an entire storage page in the federated database. It is typical for a node's array to be updated frequently; the default value of 1 for this property minimizes lock conflicts. If you know that a particular collection will be used by a single user, locking is not an issue. In that case, a larger value, such as 5000, may be appropriate for the collection's tree administrator.

For additional information, see "Tree Administrator" on page 250 in the Objectivity/C++ programmer's guide.

# Working With a Tree Administrator

Like other persistent objects, tree administrators are normally manipulated through handles or object references.

You call an ordered collection's admin member function to obtain an object reference to the collection's tree administrator; you must then cast the returned object reference to type <code>ooRef(ooTreeAdmin)</code> before you access the tree administrator's data members.

# **Reference Summary**

Getting Information	<u>maxNodesPerContainer</u> <u>maxVArraysPerContainer</u> <u>nodeContainer</u> <u>vArrayContainer</u>
Setting Information	<u>setMaxNodesPerContainer</u> <u>setMaxVArraysPerContainer</u>

# **Reference Index**

maxNodesPerContainer	Gets the maximum number of B-tree nodes per container for this tree administrator.
<u>maxVArraysPerContainer</u>	Gets the maximum number of arrays per container for this tree administrator.
nodeContainer	Gets this tree administrator's current node container.
<u>setMaxNodesPerContainer</u>	Sets the maximum number of B-tree nodes per container for this tree administrator.
<u>setMaxVArraysPerContainer</u>	Sets the maximum number of arrays per container for this tree administrator.
vArrayContainer	Gets this tree administrator's current array container.

# **Member Functions**

### maxNodesPerContainer

Gets the maximum number of B-tree nodes per container for this tree administrator.

ooInt32 maxNodesPerContainer();

Returns The maximum number of B-tree nodes that can be stored in a single container.

See also <u>setMaxNodesPerContainer</u>

## maxVArraysPerContainer

Gets the maximum number of arrays per container for this tree administrator.

ooInt32 maxVArraysPerContainer();

Returns The maximum number of arrays that can be stored in a single container.

See also <u>setMaxVArraysPerContainer</u>

### nodeContainer

Gets this tree administrator's current node container.

ooRef(ooContObj) nodeContainer();

Returns Object reference to this tree administrator's current node container.

### setMaxNodesPerContainer

Sets the maximum number of B-tree nodes per container for this tree administrator.

```
void setMaxNodesPerContainer(ooInt32 max);
```

## Parameters max The maximum number of B-tree nodes that can be stored in a single container.

Discussion Changing the maximum nodes per container affects only the collection's current node container and any node containers created in the future. If you reduce the number of nodes per container, existing node containers are left with more nodes than the new maximum; if you increase the number, existing node containers are left with fewer nodes than the new maximum.

See also <u>maxNodesPerContainer</u>

### setMaxVArraysPerContainer

	Sets the maximum number of arrays per container for this tree administrator.
	<pre>void setMaxVArraysPerContainer(ooInt32 max);</pre>
Parameters	max The maximum number of arrays that can be stored in a single container.
Discussion	One array fills up an entire storage page in the federated database. It is typical for a node's array to be updated frequently; the default value of 1 for this property minimizes lock conflicts. If you know that a particular collection will be used by a single user, locking is not an issue. In that case, a larger value, such as 5000, may be appropriate for the collection's tree administrator.
	Changing the maximum arrays per container affects only the collection's current array container and any array containers created in the future.
See also	setMaxVArraysPerContainer

## vArrayContainer

Gets this tree administrator's current array container.

ooRef(ooContObj) vArrayContainer();

Returns Object reference to this tree administrator's current array container.

# ooTreeList Class

Inheritance: ooObj->ooCollection->ooBTree->ooTreeList

Handle Class: ooHandle(ooTreeList)

```
Object-Reference Class: ooRef(ooTreeList)
```

The persistence-capable class ooTreeList represents *lists* of persistent objects.

See:

- "Reference Summary" on page 668 for an overview of member functions
- "Reference Index" on page 669 for a list of member functions

To use this class, your application must include the ooCollections.h header file. For UNIX linking information, see *Installation and Platform Notes for UNIX*.

## **About Lists**

A list is an ordered collection; unlike a set, a list can contain duplicate elements and null elements. An element of a list can be located by position, given as a zero-based index. Like all collections implemented with B-trees, lists are *scalable* collections, that is, they can increase in size with minimal performance degradation.

The ooTreeList class overrides the inherited add and addAll member functions to add elements to the end of the list. Member functions defined in this class add elements to the list, or insert elements into it, at an indicated position. The elements of a list are kept in the order in which they were added or inserted. Note, however, that an element's position may change as elements are inserted in front of it in the list. For example, the element that was at index 2 will be at index 4 after two elements are added to the front of the list. For additional information, see Chapter 11, "Persistent Collections," in the Objectivity/C++ programmer's guide.

# Working With a List

As is the case for any basic object, you specify whether a list is to be transient or persistent when you create it; lists *must be persistent*. You create a list with a call to the new operator; the clustering directive in that call specifies where in the federated database to store the new list.

Like other persistent objects, lists are normally manipulated through handles or object references. You can store and find a list in the database just as you would any other persistent object.

## **Related Classes**

Two additional classes represent persistent collections of persistent objects:

- OOHashSet represents an *unordered* collection of persistent objects *with no duplicate elements*. It uses an extendible hashing mechanism.
- ooTreeSet represents a sorted collection of persistent objects with no duplicate elements; elements are sorted by the corresponding comparator (or by increasing OID if the sorted set uses the default comparator). Like this class, ooTreeSet is implemented using a B-tree data structure.

# **Reference Summary**

In the following table:

- Operators indicated as *(inherited)* are overloaded in this class with no change in behavior; they are documented with the <u>ooObj</u> class (page 431), along with the other inherited member functions not listed here.
- Member function indicated as *(inherited)* are inherited from the <u>ooBTree</u> class (page 283) or the <u>ooCollection</u> class (page 173) and are documented with the defining class.

Creating and Deleting	<u>ooTreeList</u> <u>operator new (inherited)</u> <u>operator delete (inherited)</u>
Adding, Removing, and Changing Elements	add addAll addFirst addLast clear (inherited) remove removeAll (inherited) removeRange retainAll (inherited) set
Getting Elements	<u>first</u> <u>get</u> <u>iterator</u> <u>last</u> (inherited)
Getting Indexes	<u>indexOf</u> (inherited) <u>lastIndexOf</u> (inherited)
Getting Information	<u>depth</u> (inherited) <u>size</u> (inherited)
Finding Auxiliary Objects	<u>admin</u> <u>comparator</u>
Testing	<u>contains</u> <u>containsAll</u> (inherited) <u>isEmpty</u> (inherited)
Maintaining the B-Tree	<u>compact</u> (inherited)
Viewing in an MROW Transaction	refresh (inherited)

# **Reference Index**

add	Adds the specified object to this list.
addAll	Adds all elements (or keys) in the specified collection to this list.

<u>addFirst</u>	Adds the specified object to the beginning of this list.
addLast	Adds the specified object to the end of this list.
<u>admin</u>	Finds the tree administrator for this list.
<u>comparator</u>	Overrides the inherited member function; disallows finding the comparator for lists.
<u>contains</u>	Tests whether this list contains the specified object.
first	Finds the first element in this list.
get	Finds the specified element of this list.
iterator	Initializes a scalable-collection iterator to find the elements of this list.
<u>ooTreeList</u>	Constructs a new list.
remove	Removes the first occurrence of the specified object from this list.
<u>removeAllDeleted</u>	Removes from this list all persistent objects that have been deleted from the federated database.
<u>removeRange</u>	Removes from this list all elements with indexes in the specified range.
set	Replaces the element at the specified index with the specified object.

## Constructors

### ooTreeList

Constructs a new list.

```
ooTreeList(
    int maxNodeSize = (oomGetPageSize() - 92) / 8,
    ooHandle(ooContObj) contAdminH = 0,
    ooHandle(ooContObj) contVarrayH = 0);
```

Parameters maxNodeSize

The node size for the new list's B-tree.

#### contAdminH

Handle to the container in which to store the tree administrator for the new lists.

#### contVarrayH

Handle to the initial array container for the new lists.

Discussion The constructor creates an empty list. The optional parameters allow the caller to:

- Override the default node size.
- Prevent creation of a container for the list's tree administrator.
- Prevent creation of the list's initial array container.

## **Member Functions**

### add

Adds the specified object to this list.

	<pre>1. virtual ooBoolean add(</pre>	
	<pre>2. ooStatus add( const ooInt32 index, const ooHandle(ooObj) &amp;objH);</pre>	
Parameters	о <i>bjн</i> Handle to the object to be added.	
	<i>index</i> The index where the new element is to be inserted.	
Returns	(Variant 1) occTrue if an element was added; otherwise, occFalse.	
	(Variant 2) oocSuccess if successful; otherwise, oocError.	
Discussion	Variant 1 overrides the inherited member function to add the specified object at the end of this list.	
	Variant 2 inserts the specified object into this list at the specified index, increasing the size of the list by one and effectively incrementing the index of all subsequent elements. For example, if <i>index</i> is 2, this member function inserts the new element before the third existing element. The new element now has index 2 and what used to be the third element is now the fourth element (and has index 3).	

See also	<u>addAll</u>
	<u>addFirst</u>
	<u>addLast</u>
	<u>remove</u>
	set

### addAll

Adds all elements (or keys) in the specified collection to this list.

	<pre>1. ooStatus addAll( int index, const ooHandle(ooCollection) &amp;collectionH);</pre>
	<pre>2. ooBoolean addAll( const ooHandle(ooCollection) &amp;collectionH);</pre>
Parameters	index
	The index where the first of the new elements is to be inserted.
	collectionH
	Handle to the collection whose elements are to be added to this list. If the elements of the collection are key-value pairs, its keys are added to this list.
	Note: <i>collectionH</i> may done be a handle to this list. That is, this method cannot be used to add another copy of this list's elements to this list.
Returns	(Variant 1) oocSuccess if successful; otherwise, oocError.
	(Variant 2) <code>oocTrue</code> if any elements were added; otherwise, <code>oocFalse</code> .
Discussion	Variant 1 inserts elements (or keys) of the specified collection into this list at the indicated index. It fails if <i>index</i> is not a valid index for this list.
	Variant 2 overrides the inherited method to add the new elements to the end of this list.
See also	<u>add</u> <u>addFirst</u> <u>addLast</u>
addFirst	

Adds the specified object to the beginning of this list.

```
ooStatus addFirst(const ooHandle(ooObj) &objH);
```

Parameters	objH
	Handle to the object to be added.
Returns	oocSuccess if successful; otherwise, oocError.
Discussion	The new element becomes the first element of the list, effectively incrementing the index of all existing elements.
See also	add addAll addLast
addLast	
	Adds the specified object to the end of this list.
	<pre>ooStatus addLast(const ooHandle(ooObj) &amp;objH);</pre>
Parameters	оbjн Handle to the object to be added.
Returns	oocSuccess if successful; otherwise, oocError.
See also	add addAll addFirst
admin	
	Finds the tree administrator for this list.
	<pre>virtual ooRef(ooAdmin) admin() const;</pre>
Returns	Object reference to the tree administrator for this list.
Discussion	You typically call this member function when you want to change the way that this list's internal objects (B-tree nodes and the arrays they reference) are assigned to containers. Before you do so, you must cast the returned object reference to $ooRef(ooTreeAdmin)$ .

### comparator

Overrides the inherited member function; disallows finding the comparator for lists.

virtual ooRef(ooCompare) comparator() const;

Returns	A null object reference.
Discussion	Lists do not use comparators, so this member function always returns null.
contains	
	Tests whether this list contains the specified object.
	<pre>virtual ooBoolean contains(     const ooHandle(ooObj) &amp;objH) const;</pre>
Parameters	оbjн Handle to the element to be tested for containment in this list.
Returns	oocTrue if this list contains an element equal to the specified object; otherwise, oocFalse.
first	
	Finds the first element in this list.
	<pre>virtual ooRef(ooObj) first() const;</pre>
Returns	Object reference to the first element of this list.
See also	get
get	
	Finds the specified element of this list.
	<pre>1. ooRef(ooObj) get(const ooInt32 index) const;</pre>
	<pre>2. virtual ooRef(ooObj) get(const void *lookupVal) const;</pre>
Parameters	<i>index</i> The zero-based index of the desired element.
	<i>lookupVal</i> <b>Pointer to data that identifies the desired element.</b>
Returns	(Variant 1) Finds the element at the specified index; it returns an object reference to the element whose index is <i>index</i> .
	(Variant 2) Overrides the inherited member function to disallow looking up a list element by data that identifies the element. It always returns a null object reference.

Discussion Variant 2 requires the collection to use an application-defined comparator that can identify an element based on class-specific data. Because a list has no comparator, variant 2 is not relevant for lists.

```
See also <u>first</u>
```

### iterator

	Initializes a scalable-collection iterator to find the elements of this list.
	<pre>virtual ooCollectionIterator *iterator() const;</pre>
Returns	A pointer to a scalable-collection iterator for finding the elements of this list; the caller is responsible for deleting the iterator when it is no longer needed.
Discussion	The returned iterator finds the elements as ordered in this list.
	You must delete the iterator when you have finished using it.

### remove

Removes the first occurrence of the specified object from this list.

Parameters	objH
	Handle to the object to be removed.
Returns	oocTrue if an element was removed; otherwise, oocFalse.
See also	<u>add</u> <u>removeAllDeleted</u> removeRange

### removeAllDeleted

Removes from this list all persistent objects that have been deleted from the federated database.

virtual void removeAllDeleted();

Discussion You can calling this member function to restore this list's referential integrity.

### removeRange

Removes from this list all elements with indexes in the specified range.

	<pre>ooStatus removeRange(     int fromIndex,     int toIndex);</pre>
Parameters	fromIndex The index of the first element to be removed.
	toIndex The index of the last element to be removed.
Returns	oocSuccess if successful; otherwise, oocError.
See also	remove
set	
	Replaces the element at the specified index with the specified object.
	<pre>Replaces the element at the specified index with the specified object. ooStatus set(     const int index,     const ooHandle(ooObj) &amp;objH);</pre>
Parameters	<pre>Replaces the element at the specified index with the specified object. ooStatus set(     const int index,     const ooHandle(ooObj) &amp;objH); index The index of the element to be replaced.</pre>
Parameters	<pre>Replaces the element at the specified index with the specified object. ooStatus set(     const int index,     const ooHandle(ooObj) &amp;objH); index The index of the element to be replaced. objH Handle to the object that is to replace the existing element at the specified index.</pre>
Parameters	<pre>Replaces the element at the specified index with the specified object. ooStatus set(     const int index,     const ooHandle(ooObj) &amp;objH); index The index of the element to be replaced. objH Handle to the object that is to replace the existing element at the specified index. oocSuccess if successful; otherwise, oocError.</pre>

# ooTreeMap Class

Inheritance: ooObj->ooCollection->ooBTree->ooTreeSet->ooTreeMap

Handle Class: ooHandle(ooTreeMap)

**Object-Reference Class:** ooRef(ooTreeMap)

The persistence-capable class ooTreeMap represents sorted object maps.

See:

- "Reference Summary" on page 678 for an overview of member functions
- "Reference Index" on page 679 for a list of member functions

To use this class, your application must include the ooCollections.h header file. For UNIX linking information, see *Installation and Platform Notes for UNIX*.

## About Sorted Object Maps

An object map is a collection of key-value pairs; each key and each value is a persistent object. No two elements of the object map may have the same key. As the name implies, each element of an object map is a mapping from its key object to its value object.

Like all collections implemented with B-trees, sorted object maps are *scalable* collections, that is, they can increase in size with minimal performance degradation.

The elements of a sorted object map are sorted by their keys according to the ordering implemented by the object map's corresponding comparator. If a sorted object map has a default comparator, its elements are sorted by the object identifiers (OIDs) of their keys.

For additional information, see Chapter 11, "Persistent Collections," in the Objectivity/C++ programmer's guide.

# Working With a Sorted Object Map

As is the case for any basic object, you specify whether a sorted object map is to be transient or persistent when you create it; sorted object maps *must be persistent*. You create a sorted object map with a call to the new operator; the clustering directive in that call specifies where in the federated database to store the new sorted object map.

Like other persistent objects, sorted object maps are normally manipulated through handles or object references. You can store and find a sorted object map in the database just as you would any other persistent object.

# **Related Classes**

Two additional classes represent persistent collections of key-value pairs:

- OOHashMap represents an *unordered* object map. It uses an extendible hashing mechanism.
- OOMap represents an *unordered* name map, that is, a collection of key-value pairs in which *the key is a string* and the value is an object reference to a persistent object. It uses a traditional (non-extendible) hashing mechanism.

# **Reference Summary**

In the following table:

- Operators indicated as *(inherited)* are overloaded in this class with no change in behavior; they are documented with the <u>ooObj</u> class (page 431), along with the other inherited member functions not listed here.
- Member function indicated as *(inherited)* are inherited from the <u>ooTreeSet</u> class (page 283), the <u>ooBTree</u> class (page 283), or the <u>ooCollection</u> class (page 173) and are documented with the defining class.

Creating and Deleting	<u>ooTreeMap</u> <u>operator new</u> (inherited) <u>operator delete</u> (inherited)
Adding, Removing, and Changing Elements	add addAll clear (inherited) put remove removeAll (inherited) removeAllDeleted retainAll (inherited)
Getting Elements	<u>first</u> (inherited) <u>get</u> <u>keyIterator</u> (inherited) <u>valueIterator</u> <u>last (inherited)</u>
Getting Indexes	<u>indexOf</u> (inherited) <u>lastIndexOf</u> (inherited)
Getting Information	<u>depth</u> (inherited) <u>size</u> (inherited)
Finding Auxiliary Objects	<u>admin</u> (inherited) <u>comparator</u> (inherited)
Testing	<u>containsKey</u> <u>containsValue</u> <u>containsAll</u> (inherited) <u>isEmpty</u> (inherited)
Maintaining the B-Tree	compact
Viewing in an MROW Transaction	refresh (inherited)

# **Reference Index**

add	Adds the specified object to this sorted object map.
<u>addAll</u>	Adds all elements in the specified object map to this sorted object map.
compact	Minimizes the number of nodes in this sorted object map's B-tree.

<u>containsKey</u>	Tests whether this sorted object map contains an element with the specified key.
<u>containsValue</u>	Tests whether this sorted object map contains an element with the specified value.
<u>get</u>	Finds the value paired with the specified key in this sorted object map.
<u>ooTreeMap</u>	Constructs a new sorted object map.
put	Maps the specified key to the specified value in this sorted object map.
remove	Removes the element, if any, with the specified key from this sorted object map.
removeAllDeleted	Removes from this sorted object map all elements in which either the key or the value has been deleted from the federated database.
valueIterator	Initializes a scalable collection iterator to find all values in this sorted object map.

## Constructors

### ooTreeMap

### Constructs a new sorted object map.

1.	<pre>ooTreeMap(     int maxNodeSize = (oomGetPageSize() - 92) / 8,     ooHandle(ooContObj) contAdminH = 0,     ooHandle(ooContObj) contVarrayH = 0);</pre>
2.	<pre>ooTreeMap(    ooHandle(ooCompare) &amp; compH,    int maxNodeSize = (oomGetPageSize() - 92) / 8,    ooHandle(ooContObj) contAdminH = 0,    ooHandle(ooContObj) contVarrayH = 0);</pre>

#### Parameters maxNodeSize

The node size for the new sorted object map's B-tree.

#### contAdminH

Handle to the container in which to store the tree administrator for the new sorted object map.

#### contVarrayH

Handle to the initial array container for the new sorted object map.

сотрН

Handle to the comparator for the new sorted object map; must be an instance of an application-specific derived class of ooCompare.

### Discussion Variant 1 creates an empty sorted object map with a default comparator; Variant 2 creates an empty sorted object map with the specified comparator. The optional parameters to both constructors allow the caller to:

- Override the default node size.
- Prevent creation of a container for the sorted object map's tree administrator.
- Prevent creation of the sorted object map's initial array container.

## **Member Functions**

### add

	Adds the specified object to this sorted object map.
	<pre>virtual ooBoolean add( const ooHandle(ooObj) &amp;objH);</pre>
Parameters	оbjн Handle to the object to be added.
Returns	oocTrue if an element was added; otherwise, oocFalse.
Discussion	This member function adds a new key-value pair to this sorted object map with the specified object as its key and a null value. If this sorted object map currently has an element whose key is the specified object, the value of the existing element is replaced by null.
See also	addAll remove
addAll	
	Adds all elements in the specified object map to this sorted object map.
	<pre>virtual ooBoolean addAll(     const ooHandle(ooCollection) &amp;collectionH);</pre>

Parameters	collectionH
	Handle to the sorted or unordered object map whose elements are to be added to this sorted object map.
Returns	oocTrue if any elements were added; otherwise, oocFalse.
Discussion	If the specified collection is an object map, its elements are added to this sorted object map. If this sorted object map currently has an element with the same key as an element of the specified collection, the existing element is replaced by the element of the specified collection.
	If the specified collection is a collection of persistent objects, each of its elements is added as a key to this sorted object map; a null value is paired with each key. If this sorted object map currently has an element whose key is an element of the specified collection, the value of the existing element is replaced by null.
See also	<u>add</u> get
compact	
	Minimizes the number of nodes in this sorted object map's B-tree.
	<pre>virtual void compact();</pre>

Discussion After you have added all elements that you expect this sorted object map to have, you can call this member function to minimize the number of nodes in its B-tree. Doing so saves space and improves read performance. Indexes of elements within the sorted object map remain unchanged.

> If you call this member function before all elements have been added, insert (add) performance will not necessarily improve. After the B-tree has been compacted, adding an element will very likely cause one or more nodes to be added to the B-tree.

### containsKey

Tests whether this sorted object map contains an element with the specified key.

1.	ooBoolean containsKey( const ooHandle(ooObj) & <i>keyH</i> );
2.	ooBoolean containsKey( const void * <i>lookupVal</i> );
keyh	I

Parameters

Handle to the key to be tested for containment in this sorted object map.

	lookupVal
	Pointer to data that identifies the key to be tested for containment in this sorted object map.
Returns	oocTrue if this sorted object map contains an element with the specified key; otherwise, oocFalse.
Discussion	You can call this member function to check whether this sorted object map maps the specified key to some value.
	Variant 2 tests whether any key is "equal" to the specified lookup data, as determined by the comparator for this sorted object map. It is useful if this sorted object map has an application-defined comparator that can identify a key based on class-specific data.
See also	<u>containsValue</u>

### containsValue

Tests whether this sorted object map contains an element with the specified value.

ooBoolean	containsValue(		
const	ooHandle(ooObj)	&valueH)	const;

Parameters valueH

Handle to the value to be tested for containment in this sorted object map.

- Returns occTrue if this sorted object map contains an element whose value is the specified object; otherwise, occFalse.
- Discussion You can call this member function to check whether this sorted object map maps at least one key to the specified value.

See also <u>containsKey</u>

### get

Finds the value paired with the specified key in this sorted object map.

- ooRef(ooObj) get(const ooHandle(ooObj) &keyH) const;
- ooRef(ooObj) get(const void \*lookupVal) const;

#### Parameters keyH

Handle to the key to be looked up.

	lookupVal Pointer to data that identifies the desired key.
Returns	Object reference to the value in the element with the specified key, or a null object reference if this sorted object map contains no mapping for that key.
Discussion	A return value of null does not necessarily indicate that no element has the specified key. It is possible that this sorted object map explicitly maps the key to null. You can use the <u>containsKey</u> member function to distinguish these two cases.
	Variant 2 finds the element whose key is "equal" to the specified lookup data, as determined by the comparator for this sorted object map. It is useful if this sorted object map has an application-defined comparator that can identify a key based on class-specific data.
See also	<u>put</u> addAll
put	
	Maps the specified key to the specified value in this sorted object map.
	<pre>ooStatus put( const ooHandle(ooObj) &amp;keyH, const ooHandle(ooObj) &amp;valueH);</pre>
Parameters	keyH Handle to the key.
	<i>valueH</i> Handle to the value.
Returns	oocSuccess if successful; otherwise, oocError.
Discussion	If this sorted object map already contains an element with the specified key, this member function replaces the value in that element. Otherwise, this member function adds a new element with the specified key and value.
See also	<u>get</u> addall
#### remove

Removes the element, if any, with the specified key from this sorted object map.

	<pre>1. virtual ooBoolean remove(</pre>
	<pre>2. ooBoolean remove(     const void *lookupVal);</pre>
Parameters	keyH Handle to the key of the element to be removed.
	<i>lookupVal</i> Pointer to data that identifies the key of the element to be removed.
Returns	oocTrue if an element was removed; otherwise, oocFalse.
	Variant 2 removes the element whose key is "equal" to the specified lookup data, as determined by the comparator for this sorted object map. It is useful if this sorted object map has an application-defined comparator that can identify a key based on class-specific data.
See also	add

#### removeAllDeleted

Removes from this sorted object map all elements in which either the key or the value has been deleted from the federated database.

virtual void removeAllDeleted();

Discussion You can call this member function to restore this sorted object map's referential integrity.

See also remove

#### valuelterator

Initializes a scalable collection iterator to find all values in this sorted object map.

ooCollectionIterator \*valueIterator() const;

A pointer to a scalable collection iterator for finding all the persistent objects Returns used as values in elements of this sorted object map. The iterator finds the values in an unspecified order.

Discussion The returned iterator finds the values in the order in which the elements are sorted; that is, it finds the values in the sorted order of their keys.

You must delete the iterator when you have finished using it.

# ooTreeSet Class

Inheritance: ooObj->ooCollection->ooBTree->ooTreeSet

Handle Class: ooHandle(ooTreeSet)

```
Object-Reference Class: ooRef(ooTreeSet)
```

The persistence-capable class <code>ooTreeSet</code> represents *sorted sets* of persistent objects with no duplicate elements.

See:

- "Reference Summary" on page 688 for an overview of member functions
- "Reference Index" on page 689 for a list of member functions

To use this class, your application must include the ooCollections.h header file. For UNIX linking information, see *Installation and Platform Notes for UNIX*.

## **About Sorted Sets**

The elements of a sorted set are sorted according to the ordering implemented by the set's corresponding comparator. If a sorted set has a default comparator, its elements are sorted by their object identifiers (OIDs).

Like all collections implemented with B-trees, sorted sets are *scalable* collections, that is, they can increase in size with minimal performance degradation.

For additional information, see Chapter 11, "Persistent Collections," in the Objectivity/C++ programmer's guide.

## Working With a Sorted Set

As is the case for any basic object, you specify whether a sorted set is to be transient or persistent when you create it; sorted sets *must be persistent*. You create a sorted set with a call to the new operator; the clustering directive in that call specifies where in the federated database to store the new sorted set.

Like other persistent objects, sorted sets are normally manipulated through handles or object references. You can store and find a sorted set in the database just as you would any other persistent object.

## **Related Classes**

Two additional classes represent persistent collections of persistent objects:

- ooHashSet represents an *unordered* collection of persistent objects with no duplicate elements. It uses an extendible hashing mechanism.
- ooTreeList represents a collection of persistent objects that are maintained in the order specified when they are added to the collection. A list can contain duplicate elements. Like this class, ooTreeList is implemented using a B-tree data structure.

## **Reference Summary**

In the following table:

- Operators indicated as *(inherited)* are overloaded in this class with no change in behavior; they are documented with the <u>ooObj</u> class (page 431), along with the other inherited member functions not listed here.
- Member function indicated as *(inherited)* are inherited from the <u>ooBTree</u> class (page 283) or the <u>ooCollection</u> class (page 173) and are documented with the defining class.

Creating and Deleting	<u>ooTreeSet</u> <u>operator new</u> (inherited) <u>operator delete</u> (inherited)
Adding and Removing Elements	<u>add</u> <u>addAll</u> (inherited) <u>clear</u> (inherited) <u>remove</u> <u>removeAll</u> (inherited) <u>removeAllDeleted</u> (inherited) <u>retainAll</u> (inherited)
Getting Elements	<u>first</u> (inherited) get <u>iterator</u> (inherited) <u>last</u> (inherited)
Getting Indexes	<u>indexOf</u> (inherited) <u>lastIndexOf</u> (inherited)
Getting Information	<u>depth</u> (inherited) <u>size</u> (inherited)
Finding Auxiliary Objects	admin comparator
Testing	<u>contains</u> <u>containsAll</u> (inherited) <u>isEmpty</u> (inherited)
Maintaining the B-Tree	<u>compact</u> (inherited)
Viewing in an MROW Transaction	<u>refresh</u> (inherited)

# **Reference Index**

<u>add</u>	Adds the specified object to this sorted set.
<u>admin</u>	Gets the tree administrator for this sorted collection.
<u>comparator</u>	Finds the comparator for this sorted collection.
<u>contains</u>	Tests whether this sorted set contains the specified object.
get	Gets the specified element of this sorted set.

<u>ooTreeSet</u>	Constructs a new sorted set.
remove	Removes the specified object from this sorted set.

## Constructors

## ooTreeSet

#### Constructs a new sorted set.

	<pre>1. ooTreeSet(     int maxNodeSize = (oomGetPageSize() - 92) / 8,     ooHandle(ooContObj) contAdminH = 0,     ooHandle(ooContObj) contVarrayH = 0);</pre>		
	<pre>2. ooTreeSet(     ooHandle(ooCompare) &amp; compH,     int maxNodeSize = (oomGetPageSize() - 92) / 8,     ooHandle(ooContObj) contAdminH = 0,     ooHandle(ooContObj) contVarrayH = 0);</pre>		
Parameters	maxNodeSize The node size for the new sorted set's B-tree.		
	contAdminH Handle to the container in which to store the tree administrator for the new sorted set.		
	<i>contVarrayH</i> Handle to the initial array container for the new sorted set.		
	<i>compH</i> Handle to the comparator for the new sorted set; must be an instance of an application-specific derived class of ooCompare.		
Discussion	Variant 1 creates an empty sorted set with a default comparator; Variant 2 creates an empty sorted set with the specified comparator. The optional parameters to both constructors allow the caller to:		

- Override the default node size.
- Prevent creation of a container for the sorted set's tree administrator.
- Prevent creation of the sorted set's initial array container.

# **Member Functions**

#### add

	Adds the specified object to this sorted set.
	<pre>virtual ooBoolean add(     const ooHandle(ooObj) &amp;objH);</pre>
Parameters	оbјн Handle to the object to be added.
Returns	oocTrue if an element was added; otherwise, oocFalse.
Discussion	This member function returns false if the specified object is already an element of this sorted set.
See also	remove
admin	
	Gets the tree administrator for this sorted collection.
	<pre>virtual ooRef(ooAdmin) admin() const;</pre>
Returns	Object reference to the tree administrator for this sorted collection.
Discussion	You typically call this member function when you want to change the way that this sorted collection's internal objects (B-tree nodes and the arrays they reference) are assigned to containers. Before you do so, you must cast the returned object reference to <code>ooRef(ooTreeAdmin)</code> .
comparato	r

Finds the comparator for this sorted collection.

virtual ooRef(ooCompare) comparator() const;

Returns Object reference to the comparator for this sorted collection, or null if this sorted collection has a default comparator.

## contains

	Tests whether this sorted set contains the specified object.	
	<pre>1. virtual ooBoolean contains(</pre>	
	<pre>2. ooBoolean contains(</pre>	
Parameters	objH Handle to the element to be tested for containment in this sorted set.	
	lookupVal	
	Pointer to data that identifies the object to be tested for containment in this sorted set.	
Returns	oocTrue if this sorted set contains an element equal to the specified object; otherwise, oocFalse.	
Discussion	Variant 2 tests whether any element is "equal" to the specified lookup data, as determined by the comparator for this sorted set. It is useful if this sorted set has an application-defined comparator that can identify an element based on class-specific data.	
get		
	Gets the specified element of this sorted set.	
	<pre>1. ooRef(ooObj) get(const ooInt32 index) const;</pre>	
	<ol> <li>virtual ooRef(ooObj) get(const void *lookupVal) const;</li> </ol>	
Parameters	<i>index</i> The zero-based index of the desired element.	
	<i>lookupVal</i> Pointer to data that identifies the desired element.	
Returns	Variant 1 finds the element at the specified index; it returns an object reference to that element, or a null object reference if <i>index</i> is out of bounds.	
	Variant 2 finds the element that is "equal" to the specified lookup data, as determined by the comparator for this sorted set; it returns an object reference to that element, or a null object reference if this sorted set does not contain such an element.	

Discussion	Variant 2 is useful if this sorted set has an application-defined comparator that can identify an element based on class-specific data.
remove	
	Removes the specified object from this sorted set.
	<pre>1. virtual ooBoolean remove(</pre>
	<pre>2. virtual ooBoolean remove(</pre>
Parameters	objH
	Handle to the object to be removed.
	lookupVal
	Pointer to data that identifies the element to be removed.
Returns	oocTrue if an element was removed; otherwise, oocFalse.
Discussion	Variant 2 is useful if this sorted set has an application-defined comparator that can identify an element based on class-specific data.
See also	add

# ooTVArrayT<element\_type> Class

Inheritance: ooTVArrayT<element\_type>

The non-persistence-capable template class <code>ooTVArrayT<element\_type></code> represents a *temporary variable-size array* (VArray) whose elements are of type <code>element\_type</code>.

See:

- "Reference Summary" on page 697 for an overview of member functions
- "Reference Index" on page 697 for a list of member functions

For backward compatibility, you can use the macro-style name ooTVArray(element\_type) instead of the name ooTVArrayT<element\_type>.

## **About Temporary VArrays**

A temporary VArray is a variable-size array that can only be transient. Unlike *standard VArrays* (instances of <u>ooVArrayT<element\_type></u>), temporary VArrays cannot be incorporated in a persistent object, either through inheritance or embedding as a data member; only standard VArrays can be made persistent.

#### **Elements of a Temporary VArray**

Elements of temporary VArrays can be of any type, including non-persistence-capable class types. For example, <code>ooTVArrayT<Point></code> is a template class representing temporary VArrays whose elements are instances of the non-persistence-capable class <code>Point</code>.

*Note* Temporary VArrays cannot contain other VArrays, either directly or indirectly.

Unlike standard VArrays, elements of temporary VArrays can be handles, iterators, or, more generally, elements that contain memory pointers to other elements.

Like the elements of fixed C++ arrays, the *element\_type* of a temporary VArray must have a default constructor (a constructor that can take no arguments).

#### **Structure and Behavior**

A temporary VArray is a compound object consisting of a reference to a vector of elements. The reference portion of the temporary VArray occupies a fixed amount of space; the vector portion occupies a variable amount of space and may be relocated by certain operations. Elements in the vector are guaranteed contiguous within virtual memory.

You use public constructors to create temporary VArrays with 0 or more elements. An empty temporary VArray has no vector allocated for it until you add elements using the  $\underline{resize}$  or  $\underline{extend}$  member function. You use these member functions to grow or truncate the vector dynamically. Resizing a temporary VArray to 0 elements deallocates the vector.

You access each element by its position in the array. Elements are numbered starting with 0; the position number is the element's *index* or *subscript*. VArray operations verify that any specified indexes are valid based on the VArray's current size; the <u>elem</u> access function bypasses this subscript range checking.

Because of the way temporary VArrays are represented, you cannot access the first element of an array by dereferencing it; that is, the expression *\*myArray* does not access the first element of *myArray*. You must use member functions such as <u>operator[]</u>, <u>elem</u>, or <u>head</u> to access the first element of a temporary VArray.

#### Effect of Resizing

A resizing operation may relocate the vector portion of a VArray to keep the elements contiguous in virtual memory. This relocation is performed differently by standard and temporary VArrays, which affects the permitted element types:

- When a standard VArray is relocated, its elements are *bit-wise* copied, which preserves the element data exactly, but invalidates any element data that consists of memory pointers to other (now relocated) elements.
- When a temporary VArray is relocated, its elements are copied element-by-element, which invokes the element\_type default constructor to create new, empty elements and then uses element\_type assignment to assign each original element value to a corresponding new element. This preserves the validity of any elements that point to other elements, provided

that the default constructor and destructor for *element\_type* manage the pointer linkage as appropriate.

# **Reference Summary**

Creating	<u>ooTVArrayT<element_type></element_type></u>
Assigning	<u>operator=</u>
Modifying	<u>operator=</u> <u>extend</u> <u>resize</u> <u>set</u> <u>update</u>
Finding Elements	operator[] elem head
Getting Information	size

# **Reference Index**

<u>elem</u>	Gets the specified element of this temporary VArray, without performing subscript boundary checking.
<u>extend</u>	Allocates a new element at the end of this temporary VArray, and sets it to the specified value.
head	Gets the first element of this temporary VArray.
<u>ooTVArrayT<element_type></element_type></u>	Default constructor that constructs a new temporary VArray.
<u>ooTVArrayT<element_type></element_type></u>	Constructs a new temporary VArray of the specified size.
<u>ooTVArrayT<element_type></element_type></u>	Copy constructor that constructs a new copy of the specified temporary VArray.

<u>operator=</u>	Assignment operator; assigns the specified temporary VArray to this temporary VArray, automatically adjusting the size of this temporary VArray.
<u>operator[]</u>	Subscript operator; accesses the specified element of this temporary VArray.
resize	Extends or truncates this temporary VArray to the specified number of elements.
set	Sets the specified element of this VArray to be the indicated value.
size	Gets the current number of elements in this temporary VArray.
update	Returns oocSuccess because a temporary VArray can always be updated.

## **Constructors and Destructors**

#### ooTVArrayT<element\_type>

Default constructor that constructs a new temporary VArray.

ooTVArrayT<element\_type> ();

Discussion Constructs an uninitialized temporary VArray whose size is 0. No vector is allocated until you add elements using the resize or extend member function.

#### ooTVArrayT<element\_type>

Constructs a new temporary VArray of the specified size.

ooTVArrayT<element\_type> (uint32 initSize);

Parameters initSize Initial number of elements to allocate. If you specify 0, no vector is allocated until you add elements using the resize or extend member function.

## ooTVArrayT<element\_type>

Copy constructor that constructs a new copy of the specified temporary VArray.

ooTVArrayT<element\_type> (ooTVArrayT<element\_type> &array);

# Parameters array Temporary VArray of the same type as this temporary VArray. Discussion Constructs a temporary VArray whose size is equal to the size of array, then performs an element-by-element copy from array into the newly created

VArray. The constructor operation populates the new temporary VArray by using the *element\_type* default constructor to create new, empty elements, and then using *element\_type* assignment to assign each element of *array* to a corresponding new element.

## Operators

#### operator=

Assignment operator; assigns the specified temporary VArray to this temporary VArray, automatically adjusting the size of this temporary VArray.

Parameters array

Temporary VArray of the same type as this temporary VArray.

Returns This temporary VArray.

Discussion The assignment operation resizes this temporary VArray to be the same size as *array*, and then performs an element-by-element copy from *array* into this temporary VArray. The assignment operation populates this temporary VArray by using the *element\_type* default constructor to create new, empty elements, and then using *element\_type* assignment to assign each element of *array* to a corresponding new element.

## operator[]

Subscript operator; accesses the specified element of this temporary VArray.

element\_type &operator[] (uint32 index);

Parameters	index Index of the element to access. Specify 0 to access the first element.
Returns	index'th element of this temporary VArray.

Discussion You can use the subscript operator to either get or change the specified element of this temporary VArray (that is, you can use the subscript operator on either the right or left side of an assignment operation).

Subscript boundaries are checked to ensure integrity. You can use the elem member function to bypass subscript range checking.

# **Member Functions**

#### elem

	Gets the specified element of this temporary VArray, without performing subscript boundary checking.
	<pre>element_type &amp;elem(uint32 index);</pre>
Parameters	<i>index</i> Index of the element to access. Specify 0 to access the first element.
Returns	Reference to the <i>index</i> 'th element of the temporary VArray.
Discussion	You can use elem to either get or set the specified element of this temporary VArray (that is, you can use elem on either the right or left side of an assignment operation).
	The elem member function bypasses subscript range checking. You should use <pre>operator[_]</pre> if you want subscript boundaries to be checked.
extend	
	Allocates a new element at the end of this temporary VArray, and sets it to the specified value.
	<pre>ooStatus extend(element_type &amp;newValue);</pre>
Parameters	newValue Value to be assigned to the new element, passed by reference.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	Extending a temporary VArray implicitly resizes it, which is a potentially expensive operation. You should therefore use extend as a convenient way to add only a single element to a VArray. If you need to add multiple elements in a single transaction, you should consider using <u>resize</u> to allocate all the elements in one operation.

#### head

	Gets the first element of this temporary VArray.
	<pre>element_type *head() const;</pre>
Returns	Pointer to the first element of this temporary VArray. If this VArray contains no elements, a null pointer is returned.
resize	
	Extends or truncates this temporary VArray to the specified number of elements.
	<pre>ooStatus resize(uint32 newSize);</pre>
Parameters	newSize Total number of elements that this temporary VArray is to have. Specify 0 to remove all the elements, freeing the storage allocated to the element vector.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	If the new size is larger than the current size, resize allocates storage for the additional elements and invokes the $element_type$ default constructor to create new, empty elements.
	If the new size is smaller than the current size, resize invokes the $element\_type$ destructor for the elements to be truncated (the elements from index $newSize + 1$ to the end) and then truncates the temporary VArray to the new size.
	To keep the temporary VArray's elements contiguous, resize may relocate the vector portion of this VArray in virtual memory. If this happens, the elements of this VArray are copied <i>element-by-element</i> to the new location. Element-by-element copying invokes the $element\_type$ default constructor to create new, empty elements and then uses $element\_type$ assignment to assign each original element value to a corresponding new element. This preserves the validity of any elements, provided that the $element\_type$ default constructor and the $element\_type$ default constructor and the $element\_type$ destructor manage the pointer linkage as appropriate. This kind of copying causes resizing a temporary VArray to be slower than resizing a standard VArray.

#### set

Sets the specified element of this VArray to be the indicated value.

```
ooStatus set(uint32 index, element_type &newValue);
```

Parameters	index	
	Index of the element to access. Specify 0 to access the first element.	
	newValue	
	Value to be assigned to the specified element, passed by reference.	
Returns	oocSuccess if successful; otherwise oocError.	
size		
	Gets the current number of elements in this temporary VArray.	
	uint32 size();	
Returns	Number of elements in this temporary VArray.	
update		
	$Returns \ \texttt{oocSuccess} \ because a \ temporary \ VArray \ can \ always \ be \ updated.$	
	ooStatus update();	
Returns	oocSuccess	
Discussion	This member function is defined for consistency with other classes.	

# ooUtf8String Class

Inheritance: ooVArrayT<ooChar> -> ooUtf8String

The class ooutf8String is a Java compatibility class that represents a *Unicode string*—a sequence of Unicode characters in UTF-8 encoding.

See:

- "Reference Summary" on page 704 for an overview of member functions
- "Reference Index" on page 704 for a list of member functions

To use the Java-compatibility classes, your application source must include the <code>javaBuiltins.h</code> header file.

## **About Unicode Strings**

A Java string (of the Java class java.lang.String) is stored in an Objectivity/DB federated database as an embedded object of the class <code>ooUtf8String</code>. If your application interoperates with a Java application to access objects with a field that contains a string, you can define the corresponding data member of your C++ class to be of type <code>ooUtf8String</code>.

The class <code>ooUtf8String</code> represents a Unicode string as a variable-size array (VArray) whose elements are the component bytes of the string. This class simply allows a C++ application to store and retrieve the binary representation of a Java string; it is the application's responsibility to parse the sequence of bytes into Unicode characters. If your application renders a Unicode string, it is responsible for selecting the appropriate glyph for any non-ASCII character in the string.

# **Working With Unicode Strings**

Because this class is derived from <code>ooVArrayT<ooChar></code>, you work with a Unicode string just as you would work with a character VArray. Alternatively, you can work with an object of class <code>ooUtf8String</code> using member functions defined for the Objectivity/C++ variable-size string class <code>ooVString</code>. To do this, you cast the <code>ooUtf8String</code> object to a <code>const char \*</code> and then pass the <code>const char \*</code> to the <code>ooVString</code> constructor.

## **Related Classes**

The persistence-capable class <u>oojString</u> represents a string element of a persistent array. It is a wrapper for a string of the ooUtf8String class.

A Java string array (of the Java type <code>String[]</code>) is stored in an Objectivity/DB federated database as an object reference to a persistent array object of the class <code>oojArrayOfObject</code>. Elements of the array are object references to instances of <code>oojString</code>.

# **Reference Summary**

Creating	<u>ooUtf8String</u>
Assigning	<u>operator=</u>
Type Conversion	operator const char *

## **Reference Index**

<u>ooUtf8String</u>	Constructs a new Unicode string.
<u>operator=</u>	Assignment operator; assigns the specified C++ string to this Unicode string.
operator const char *	Conversion operator that accesses this Unicode string as an object of type const char *.

## Constructors

## ooUtf8String

Constructs a new Unicode string.

	<ol> <li>ooUtf8String();</li> <li>ooUtf8String(const char *p);</li> </ol>
Parameters	P Existing C++ string from which to construct the new Unicode string.
Discussion	Variant 1 is the default constructor. It constructs an empty Unicode string. Variant 2 constructs a new Unicode string containing a copy of the characters in the specified C++ string.

# **Operators**

#### operator=

Assignment operator; assigns the specified C++ string to this Unicode string.

ooUtf8String &operator=(const char \*p);

Parameters	<i>p</i> C++ string whose characters are to be assigned.
Returns	This Unicode string.
Discussion	The assignment operation resizes this Unicode string to be the same size as $p$ , and then copies the characters of $p$ into this Unicode string. Any characters already in this Unicode string are overwritten. If $p$ is null, the effect is to delete the string.

#### operator const char \*

Conversion operator that accesses this Unicode string as an object of type const char \*.

```
operator const char *() const;
```

Operators

# ooVArrayT<element\_type> Class

Inheritance: ooVArrayT<element\_type>

The non-persistence-capable template class <code>ooVArrayT<element\_type></code> represents a *standard variable-size array (VArray)* whose elements are of type <code>element\_type</code>.

See:

- "Reference Summary" on page 710 for an overview of member functions
- "Reference Index" on page 711 for a list of member functions

(ODMG) The ooVArrayT<element\_type> class is equivalent to the ODMG standard class d\_Varray<element\_type>, but without the base class d\_Collection<element\_type>, which is not implemented by Objectivity/C++.

For backward compatibility, you can use the macro-style name ooVArray(element\_type) instead of the name ooVArrayT<element\_type>.

## **About Standard VArrays**

A standard VArray is a variable-size array that can be used transiently or saved persistently—for example, as an embedded data member of a persistent object. Because standard VArrays can be persistent, their element types are subject to restrictions. In contrast, *temporary VArrays* are only transient and can therefore be of any transient element type (see <u>ootVArrayT<element\_type></u>).

#### **Elements of a Standard VArray**

Elements of standard VArrays can be of most data types permitted in a persistence-capable class definition. For example, given a non-persistence-capable class Point, you can use the template class ooVArrayT<Point> to define VArrays whose elements are instances of class

Point. However, if class Point is persistence-capable, you must use ooVArrayT<ooRef(Point)> to produce VArrays whose elements are object references to points.

**NOTE** Standard VArrays cannot contain other VArrays, either directly or indirectly.

Specifically, the *element\_type* of a standard VArray can be any of the following:

- A primitive type
- An object-reference type for a persistence-capable class or structure
- An embedded-class type—specifically, a non-persistence-capable class or structure for which all of the following are true:
  - □ None of its base classes are virtual.
  - □ Each of its data members is of a primitive type, an object-reference type, or a valid embedded-class type.
  - □ None of its data members is of a prohibited element type.

The following element types are prohibited:

- VArrays
- Unions
- Bit fields
- Member pointers
- Persistence-capable structures and classes (however, object references to these are permitted)
- Handles and iterators for persistence-capable structures or classes

These restrictions apply recursively to standard VArray elements.

If you need to create VArrays of handles, iterators, or elements that contain memory pointers to other elements, you must use instances of the temporary VArray class <code>ooTVArrayT<element\_type></code>.

Like the elements of fixed C++ arrays, the *element\_type* of a VArray must have a default constructor (a constructor that can take no arguments).

#### **Structure and Behavior**

A standard VArray is a compound object consisting of a reference to a vector of elements. The reference portion of the VArray occupies a fixed amount of space; the vector portion occupies a variable amount of space and may be relocated by certain operations. Elements in the vector are guaranteed contiguous within virtual memory.

You use public constructors to create VArrays with 0 or more elements. An empty VArray has no vector allocated for it until you assign another VArray to it or add elements using the <u>resize</u> or <u>extend</u> member function. You use these member functions to grow or truncate the vector dynamically. Resizing a VArray to 0 elements deallocates the vector.

You access each element by its position in the VArray. Elements are numbered starting with 0; the position number is the element's *index* or *subscript*. VArray operations verify that any specified indexes are valid based on the VArray's current size; the <u>elem</u> access function bypasses this subscript range checking.

Because of the way VArrays are represented, you cannot access the first element of an array by dereferencing it; that is, the expression *\*myArray* does not access the first element of *myArray*. You must use member functions such as <u>operator[]</u>, elem, or <u>head</u> to access the first element of a VArray.

#### **Effect of Resizing**

A resizing operation may relocate the vector portion of a VArray to keep the elements contiguous in virtual memory. This relocation is performed differently by standard and temporary VArrays, which affects the permitted element types:

- When a standard VArray is relocated, its elements are *bit-wise* copied, which preserves the element data exactly, but invalidates any element data that consists of memory pointers to other (now relocated) elements.
- When a temporary VArray is relocated, its elements are copied element-by-element, which invokes the element\_type default constructor to create new, empty elements and then uses element\_type assignment to assign each original element value to a corresponding new element. This preserves the validity of any elements that point to other elements, provided that the default constructor and destructor for element\_type manage the pointer linkage as appropriate.

## **Working With Persistent VArrays**

A VArray is transient unless you incorporate it in a persistent object, typically by embedding the VArray as a data member of the persistent object or one of its base classes. A VArray can be stored persistently through multiple levels of inheritance and embedding (a member of a member of a base class, and so on). **EXAMPLE** This example uses a VArray class as a data member type in a persistence-capable class (Polygon), resulting in a persistent VArray embedded in a persistent object.

```
//DDL file
struct Point{
    int32 xCoord;
    int32 yCoord;
};
class Polygon : public ooObj {
    ...
    ooVArrayT<Point> vertices; // Define VArray data member
    ...
};
//Application code file
...
ooHandle(Polygon) polygonH;
polygonH = new (contH) Polygon(); // VArray embedded in Polygon
polygonH->vertices->resize(10); // Set the Varray size
...
```

When you incorporate a VArray in a persistent object, storage for the reference portion of the VArray is embedded in the object, and storage for the vector portion is allocated outside the object but within the same container.

Opening a persistent object allows you to get the reference portion of any member VArray, but does not automatically open the vector of elements. The vector of a persistent VArray is opened when you call a member function on the VArray.

Opening a persistent object implicitly locks that object and any member VArrays, because they are stored in the same container. If a persistent object is open (and locked) for read, operations that modify a member VArray will implicitly attempt to promote the read lock on the container to update.

## **Reference Summary**

Creating	ooVArrayT <element_type></element_type>
Assigning	<u>operator=</u>

Modifying	<u>operator=</u> <u>extend</u> <u>resize</u> <u>set</u> <u>update</u>
Finding Elements	operator[ ] elem head
Getting Information	size
ODMG Interface	cardinality create iterator insert element is empty remove all replace element at retrieve element at upper bound

# **Reference Index**

<u>cardinality</u>	(ODMG) Gets the current number of elements in this VArray; equivalent to the size member function.
<u>create_iterator</u>	(ODMG) Creates an iterator for finding the elements of this VArray.
elem	Accesses the specified element of this VArray, without performing subscript boundary checking.
extend	Allocates a new element at the end of this VArray, and sets it to the specified value.
head	Gets the first element of this VArray.
<u>insert_element</u>	(ODMG) Allocates a new element at the end of this VArray, and sets it to the specified value; equivalent to the extend member function.
is_empty	(ODMG) Checks whether the size of this VArray is 0.
<u>ooVArrayT<element_type></element_type></u>	Default constructor that constructs a new VArray.
<u>ooVArrayT<element_type></element_type></u>	Constructs a new VArray of the specified size.

<u>ooVArrayT<element type=""></element></u>	Copy constructor that constructs a new copy of the specified VArray.
<u>operator=</u>	Assignment operator; assigns the specified VArray to this VArray, automatically adjusting the size of this VArray.
<u>operator[]</u>	Subscript operator; accesses the specified element of this VArray.
remove_all	(ODMG) Removes all the elements from this VArray, changing its size to 0.
<u>replace_element_at</u>	(ODMG) Replaces the specified element of this VArray with the indicated value.
resize	Extends or truncates this VArray to the specified number of elements.
<u>retrieve_element_at</u>	(ODMG) Gets the specified element of this VArray.
set	Sets the specified element of this VArray to be the indicated value.
<u>size</u>	Gets the current number of elements in this VArray.
update	Explicitly opens this VArray for update.
upper_bound	(ODMG) Gets the current number of elements in this VArray; equivalent to the size member function.

## **Constructors and Destructors**

## ooVArrayT<element\_type>

Default constructor that constructs a new VArray.

```
ooVArrayT<element_type>();
```

Discussion Constructs an uninitialized VArray whose size is 0. No vector is allocated until you assign another VArray or add elements using the resize or extend member function.

## ooVArrayT<element\_type>

Constructs a new VArray of the specified size.

```
ooVArrayT<element_type>(uint32 initSize);
```

Parameters initSize

Initial number of elements to allocate. If you specify 0, no vector is allocated until you assign another VArray or add elements using the resize or extend member function.

#### ooVArrayT<element\_type>

Copy constructor that constructs a new copy of the specified VArray.

ooVArrayT<element\_type>(ooVArrayT<element\_type> &array);

- Parameters array VArray of the same type as this VArray.
- Discussion Constructs a VArray whose size is equal to the size of *array*, then performs an element-by-element copy from *array* into the newly created VArray. The constructor operation populates the new VArray by using the *element\_type* default constructor to create new, empty elements, and then using *element\_type* assignment to assign each element of *array* to a corresponding new element.

## Operators

#### operator=

	Assignment operator; assigns the specified VArray to this VArray, automatically adjusting the size of this VArray.
	ooVArrayT <element_type> &amp;operator=( ooVArrayT<element_type> &amp;array);</element_type></element_type>
Parameters	array VArray of the same type as this VArray. array may be either a persistent or transient instance of ooVArrayT <element_type>.</element_type>
Returns	This VArray.
Discussion	The assignment operation resizes this VArray to be the same size as <i>array</i> , and then performs an element-by-element copy from <i>array</i> into this VArray. The assignment operation populates this VArray by using the <i>element_type</i> default constructor to create new, empty elements, and then using <i>element_type</i> assignment to assign each element of <i>array</i> to a corresponding new element.

If this VArray is persistent, it is implicitly opened for update, and the lock on the container is upgraded, if necessary. If *array* is persistent, it is implicitly opened for read.

## operator[]

Subscript operator; accesses the specified element of this VArray.

```
element_type &operator[](uint32 index);
```

Parameters	<i>index</i> Index of the element to access. Specify 0 to access the first element.
Returns	index'th element of this VArray.
Discussion	You can use the subscript operator to either get or change the specified element of this VArray (that is, you can use the subscript operator on either the right or left side of an assignment operation).
	If this VArray is persistent, the subscript operator implicitly opens the VArray for read. However, you must call the $update$ member function on the VArray before using the subscript operator to modify an element. (Alternatively, you can use the <u>set</u> member function to open the VArray for update and modify an element in a single operation.)
	Subscript boundaries are checked to ensure integrity. You can use the $elem$ member function to bypass subscript range checking.
See also	elem

# **Member Functions**

## cardinality

(ODMG) Gets the current number of elements in this VArray; equivalent to the size member function.

uint32 cardinality() const;

Returns Number of elements in this VArray.

See also <u>size</u>

#### create\_iterator

(ODMG) Creates an iterator for finding the elements of this VArray.
<pre>d_Iterator<element_type> create_iterator() const;</element_type></pre>
Iterator of the same type as the VArray elements.
The created iterator is initialized to point to the first element of this VArray.
Accesses the specified element of this VArray, without performing subscript boundary checking.
<pre>element_type &amp;elem(uint32 index);</pre>
<i>index</i> Index of the element to access. Specify 0 to access the first element.
index'th element of this VArray.
You can use elem to either get or set the specified element of this VArray (that is, you can use elem on either the right or left side of an assignment operation).
If this VArray is persistent, it is implicitly opened for read. You must call the $\underline{update}$ member function on the VArray before using elem to modify an element. (Alternatively, you can use the <u>set</u> member function to open the VArray for update and modify an element in a single operation.)
The elem member function bypasses subscript range checking. You should use <pre>operator[]</pre> if you want subscript boundaries to be checked.
operator[]
Allocates a new element at the end of this VArray, and sets it to the specified value.
<pre>ooStatus extend(element_type &amp;newValue);</pre>
newValue Value to be assigned to the new element, passed by reference.
oocSuccess if successful; otherwise oocError.

Discussion	If this VArray is persistent, it is implicitly opened for update, and the lock on the container is upgraded, if necessary.
	Extending a VArray implicitly resizes it, which is a potentially expensive operation. You should therefore use extend as a convenient way to add only a single element to a VArray. If you need to add multiple elements in a single transaction, you should consider using <u>resize</u> to allocate all the elements in one operation.
See also	resize
head	
	Gets the first element of this VArray.
	<pre>element_type *head();</pre>
Returns	Pointer to the first element of this VArray. If the VArray contains no elements, a null pointer is returned.
Discussion	If this VArray is persistent, it is implicitly opened for read.

## insert\_element

	( <i>ODMG</i> ) Allocates a new element at the end of this VArray, and sets it to the specified value; equivalent to the extend member function.
	<pre>void insert_element(const element_type &amp;element);</pre>
Parameters	element Value to be assigned to the new element, passed by reference.
Discussion	If this VArray is persistent, it is implicitly opened for update, and the lock on the container is upgraded, if necessary. Extending a persistent VArray causes the entire VArray to be written to disk when the transaction commits.
	Extending a VArray implicitly resizes it, which is an expensive operation. You should therefore use this operation as a convenient way to add a single element to a VArray. If you need to add a large number of elements in a single transaction, you should consider using <u>resize</u> to allocate all the elements in one operation.
See also	resize

Objectivity/C++ Programmer's Reference

#### is\_empty

(ODMG) Checks whether the size of this VArray is 0.

int is\_empty() const;

Returns True (a nonzero integer) if this VArray has no elements; otherwise, false (the integer 0).

#### remove\_all

(ODMG) Removes all the elements from this VArray, changing its size to 0.

void remove\_all();

Discussion If this VArray is persistent, it is implicitly opened for update, and the lock on the container is upgraded, if necessary. This member function is equivalent to calling <u>resize</u> with a value of 0.

See also resize

#### replace\_element\_at

(ODMG) Replaces the specified element of this VArray with the indicated value.

	<pre>1. void replace_element_at(     const element_type &amp;newValue,     uint32 index);</pre>
	<pre>2. void replace_element_at( const element_type &amp;newValue, const Iterator<element_type> &amp;iterator);</element_type></pre>
Parameters	newValue Value to be assigned to the specified element, passed by reference.
	index Index of the VArray element to replace. Specify 0 to access the first element.
	Iterator indicating the element to replace.
Discussion	If this VArray is persistent, it is implicitly opened for update, and the lock on the container is upgraded, if necessary.

#### resize

Extends or truncates this VArray to the specified number of elements.

ooStatus resize(uint32 newSize);

Parameters newSize

Total number of elements that this VArray is to have. Specify 0 to remove all the elements, freeing the storage allocated to the element vector.

Returns oocSuccess if successful; otherwise oocError.

Discussion If this VArray is persistent, it is implicitly opened for update, and the lock on the container is upgraded, if necessary. Resizing a persistent VArray causes the entire VArray to be written to disk when the transaction commits.

If the new size is larger than the current size, resize allocates storage for the additional elements and invokes the *element\_type* default constructor to create new, empty elements.

If the new size is smaller than the current size, resize invokes the *element\_type* destructor for the elements to be truncated (the elements from index *newSize* + 1 to the end) and then truncates the VArray to the new size.

To keep the VArray elements contiguous, resize may relocate the vector portion of this VArray in virtual memory. If this happens, the elements of this VArray are *bit-wise* copied to the new location. Bit-wise copying preserves the element data exactly, which is efficient, but invalidates any element data that consists of memory pointers to other (now relocated) elements. You should use instances of the transient-only class  $ooTVArrayT<element_type>$  for temporary VArrays of elements that contain memory pointers to each other.

See also <u>extend</u> <u>insert\_element</u> remove all

#### retrieve\_element\_at

	(ODMG) Gets the specified element of this VArray.
	<pre>const element_type &amp;retrieve_element_at(uint32 index) const;</pre>
Parameters	<i>index</i> Index of the VArray element to access. Specify 0 to access the first element.
Returns	index'th element of this VArray.

Discussion	If this VArray is persistent, it is implicitly opened for read.
set	
	Sets the specified element of this VArray to be the indicated value.
	<pre>ooStatus set(uint32 index, element_type &amp;newValue);</pre>
Parameters	<pre>index Index of the VArray element to access. Specify 0 to access the first element. newValue Value to be assigned to the specified element, passed by reference</pre>
Deturne	Value to be assigned to the specified element, passed by reference.
Returns	oocSuccess II Successiui; otherwise oocError.
Discussion	If this VArray is persistent, it is implicitly opened for update, and the lock on the container is upgraded, if necessary.
size	
	Gets the current number of elements in this VArray.
	uint32 size();
Returns	Number of elements in this VArray.
See also	<u>cardinality</u> <u>upper bound</u>
update	
	Explicitly opens this VArray for update.
	ooStatus update();
Returns	oocSuccess if successful; otherwise oocError.
Discussion	This operation applies only to persistent VArrays; it has no effect on transient VArrays.
	You must explicitly open a persistent VArray for update before you can modify any element using <u>elem</u> or <u>operator[]</u> on the left side of an assignment operation. (Alternatively, you can use the <u>set</u> member function to both open the VArray for update and modify an element in a single operation.)
	Explicitly opening a VArray for update locks the enclosing container for update and causes the entire VArray to be written to disk when the transaction commits.

You should use update primarily for changing a large number of elements in a single transaction.

#### upper\_bound

(ODMG) Gets the current number of elements in this VArray; equivalent to the size member function.

uint32 upper\_bound() const;

Returns Number of elements in this VArray.

See also <u>size</u>
# ooVString Class

Inheritance: ooVArrayT<ooChar> -> ooVString

The non-persistence-capable class ooVString represents a variable-size string based on a VArray of 8-bit characters.

See:

- "Reference Summary" on page 723 for an overview of member functions
- "Reference Index" on page 724 for a list of member functions

(*ODMG*) The ooVString class is equivalent to the ODMG standard class d\_String. You can use the name d\_String interchangeably with ooVString.

# About Variable-Size Strings

A variable-size string is a character string of any length that can be stored in a persistent object. Although you can use variable-size strings anywhere in an application, their primary purpose is to serve as string attributes of persistence-capable classes (in place of C++ char \* strings, which cannot be stored persistently). You can convert transparently between a variable-size string and a const char \* string, enabling variable-size strings to be passed to functions as parameters of type const char \* and vice versa.

**EXAMPLE** This example shows a C++ class Loan that was adapted to create the persistence-capable class Loan in the DDL file Loan.ddl. In the DDL file, the dueDate data member has been changed to an ooVString from a pointer to an array of characters.

```
// Original C++ header file Loan.h
class Loan {
    ...
    char *dueDate; // Due date of loan
```

```
...
};
Loan::Loan (char *date) {
   dueDate = malloc(strlen(date) + 1);
   strcpy(dueDate, date); // Allocate memory
   . . .
}
// DDL file Loan.ddl
class Loan : public ooObj {
   ...
   ooVString dueDate;
                                    // Due date of loan
   ...
};
Loan::Loan (char *date) {
   . . .
   dueDate = date;
   . . .
}
```

### **Choosing Variable-Size Strings**

Objectivity/C++ provides two kinds of strings that can be stored persistently—variable-size strings and optimized strings. You should choose variable-size strings if you cannot predict the lengths of the strings to be stored or if you know these lengths will vary widely. In contrast, if you know that the strings to be stored are generally less than *N* characters long, you should use an optimized string of class ooString(N) for greater efficiency; see "ooString(N) Class" on page 645.

### Structure of Variable-Size Strings

A variable-size string is a VArray of character elements. Consequently, a variable-size string is a compound object consisting of a reference to a vector of elements. The reference portion of the variable-size string occupies a fixed amount of space; the vector portion occupies a variable amount of space and may be relocated by certain operations. Elements in the vector are guaranteed contiguous within virtual memory.

# Working With Variable-Size Strings

Like instances of any other non-persistence-capable class, variable-size strings are not independently persistent. However, when a variable-size string is an attribute of a persistent object, it is saved in the federated database when the persistent object is saved.

You use public constructors to create variable-size strings with 0 or more elements. An empty variable-size string has no vector allocated for it until you assign another variable-size string to it or grow it using the  $\underline{resize}$  member function. You can use the  $\underline{resize}$  member function to grow or truncate the vector dynamically. Resizing a variable-size string to 0 elements deallocates the vector.

You get each character by its position in the variable-size string. Characters are numbered starting with 0; the position number is the character's *index* or *subscript*. Operations on a variable-size string verify that any specified indexes are valid based on the string's current size. The <u>length</u> of a variable-size string is the number of characters in the VArray, not including the null terminating character that is automatically added.

Because of the way a variable-size string is represented, you cannot get the first character by dereferencing the string; that is, the expression *\*myVString* does not get the first element of *myVString*. Instead, you can specify the index 0 to the subscript operator (<u>operator[]</u>) to get the first character; alternatively, you can call the <u>head</u> member function to get a pointer to the first character.

Creating	ooVString
Assigning	operator=
Type Conversion	operator const char *
Modifying	<u>operator=</u> <u>operator+=</u> <u>resize</u>
Getting Characters	<u>operator[]</u> <u>head</u>
Getting Information	length

# **Reference Summary**

Testing	<pre>::operator== ::operator!= ::operator&lt; ::operator&lt;= ::operator&gt;=</pre>
ODMG Interface	<pre>::operator== ::operator!= ::operator&lt; ::operator&lt;= ::operator&gt;=</pre>

# **Reference Index**

head	Gets a pointer to the first character of this variable-size string.
length	Gets the number of characters in this variable-size string.
resize	Extends or truncates this variable-size string to the specified number of characters.
ooVString	Default constructor that constructs a new variable-size string whose size is 0.
ooVString	Constructs a new variable-size string containing a copy of the characters in the specified string.
<u>operator[]</u>	Subscript operator; gets the specified character of this variable-size string.
<u>operator+=</u>	Append-to operator; concatenates this variable-size string with the specified C++ string.
<u>operator=</u>	Assignment operator; assigns a copy of the specified C++ string to this variable-size string.
::operator==	(ODMG) Equality operator; tests whether the specified strings match.
::operator!=	(ODMG) Inequality operator; tests whether the specified strings are different.
::operator<	(ODMG) Less-than operator; tests whether one string is less than another.

::operator<=	(ODMG) Less-than-or-equal-to operator; tests whether one string is less than or equal to another.
::operator>	(ODMG) Greater-than operator; tests whether one string is greater than another.
::operator>=	(ODMG) Greater-than-or-equal-to operator; tests whether one string is greater than or equal to another.
operator const char *	Conversion operator that accesses this variable-size string as an object of type const char *.

# **Constructors and Destructors**

### ooVString

Default constructor that constructs a new variable-size string whose size is 0.

ooVString();

Discussion No vector is allocated until you assign a value or add elements using the resize member function.

# ooVString

Constructs a new variable-size string containing a copy of the characters in the specified string.

- 1. ooVString(const char \*p);
- 2. ooVString(const ooVString &s);

#### Parameters

Existing C++ string from which to construct the new variable-size string.

 ${}^{s}$ 

р

Existing variable-size string from which to construct the new variable-size string.

Discussion If *p* is null or if the length of *s* is 0, an uninitialized string of size 0 is created. No vector is allocated until you assign a new value or add elements using the resize member function.

# Operators

# operator[]

Subscript operator; gets the specified character of this variable-size string.

char &operator[](const uint32 index) const;

Parameters	<i>index</i> Index of the character to get. Specify 0 to get the first character.
Returns	index'th character of this variable-size string.
Discussion	An error is reported if the index is not within the allocated size of the vector (including the terminating null character).

### operator+=

	Append-to operator; concatenates this variable-size string with the specified C++ string.
	ooVString &operator+=(const char *p);
Parameters	P C++ string whose characters are to be concatenated.
Returns	This variable-size string.
Discussion	The concatenation operator adds the characters pointed to by ${}_{\mathcal{P}}$ to the end of this variable-size string.
	You can use this operator to concatenate another variable-size string to this one, because <u>operator const char *</u> automatically converts the string being concatenated to const char *.
operator=	
	Assignment operator; assigns a copy of the specified C++ string to this variable-size string.
	ooVString &operator=(const char *p);
Parameters	<i>P</i> C++ string whose characters are to be assigned.
Returns	This variable-size string.

Discussion The assignment operation resizes this variable-size string to be the same size as p, and then copies the characters of p into this variable-size string. Any characters already in this variable-size string are overwritten. If p is null, the vector is deallocated—in effect, deleting the string.

You can use this operator to assign another variable-size string to this one, because <u>operator const char \*</u> automatically converts the string being assigned to const char \*.

#### ::operator==

#### global function

(ODMG) Equality operator; tests whether the specified strings match.

	1. int ::operator==(
	const ooVString &left,
	<pre>const ooVString &amp;right);</pre>
	<pre>2. int ::operator==(     const ooVString &amp;left,     const char *right);</pre>
	<pre>3. int ::operator==(     const char *left,     const ooVString &amp;right);</pre>
eturns	1, if every character of one string matches the con-

- Returns 1, if every character of one string matches the corresponding character of the other; 0, if the two strings are lexicographically unequal.
- Discussion You can compare two variable-size strings (variant 1) or a C++ string and a variable-size string (variants 2 and 3).

### ::operator!=

global function

(ODMG) Inequality operator; tests whether the specified strings are different.

1.	<pre>int ::operator!=(     const ooVString &amp;left,     const ooVString &amp;right);</pre>
2.	<pre>int ::operator!=(</pre>
	const ooVString & <i>left</i> ,
	const char *right);
3.	<pre>int ::operator!=(</pre>
	const char *left,
	const ooVString &right);

Returns 1, if any character of one string differs from the corresponding character of the other; 0, if strings are lexicographically equal to each other.

Discussion You can compare two variable-size strings (variant 1) or a C++ string and a variable-size string (variants 2 and 3).

#### ::operator<

global function

(ODMG) Less-than operator; tests whether one string is less than another.

1.	int ::operator<(		
	const	ooVString	&left,
	const	ooVString	&right);

- 2. int ::operator<(
   const ooVString &left,
   const char \*right);</pre>
- 3. int ::operator<(
   const char \*left,
   const ooVString &right);</pre>

# Returns 1, if *left* is lexicographically less than *right*, or if *left* contains no characters or is a null pointer; otherwise, returns 0.

Discussion You can compare two variable-size strings (variant 1) or a C++ string and a variable-size string (variants 2 and 3).

### ::operator<=

global function

(*ODMG*) Less-than-or-equal-to operator; tests whether one string is less than or equal to another.

1.	int ::operator<=(
	const ooVString & <i>left</i> ,
	<pre>const ooVString &amp;right);</pre>
2.	int ::operator<=(
	const ooVString &left,
	const char *right);
3.	<pre>int ::operator&lt;=(</pre>
	const char * <i>left</i> ,
	const ooVString &right);
1, if 1	eft is lexicographically less than or equal

# Returns 1, if *left* is lexicographically less than or equal to *right*, or if *left* contains no characters or is a null pointer; otherwise, returns 0.

Discussion You can compare two variable-size strings (variant 1) or a C++ string and a variable-size string (variants 2 and 3).

#### ::operator>

global function

(ODMG) Greater-than operator; tests whether one string is greater than another.

- 1. int ::operator>(
   const ooVString &left,
   const ooVString &right);
- 2. int ::operator>(
   const ooVString &left,
   const char \*right);
- 3. int ::operator>(
   const char \*left,
   const ooVString &right);
- Returns 1, if *left* is lexicographically greater than *right*, or if *right* contains no characters or is a null pointer; otherwise, returns 0.
- Discussion You can compare two variable-size strings (variant 1) or a C++ string and a variable-size string (variants 2 and 3).

#### ::operator>=

global function

(*ODMG*) Greater-than-or-equal-to operator; tests whether one string is greater than or equal to another.

	2	<pre>const ooVString &amp;left, const ooVString &amp;right); int ::operator&gt;=(</pre>
	2.	<pre>const ooVString &amp;left, const char *right);</pre>
	3.	<pre>int ::operator&gt;=(     const char *left,     const ooVString &amp;right);</pre>
Returns	1, if <i>left</i> is lexicographically greater than or equal to <i>right</i> , or if <i>right</i> contains no characters or is a null pointer; otherwise, returns 0.	

Discussion You can compare two variable-size strings (variant 1) or a C++ string and a variable-size string (variants 2 and 3).

### operator const char \*

	Conversion operator that accesses this variable-size string as an object of type const char *.
	operator const char *() const;
Discussion	This operator results in a null pointer if the vector is not allocated.

# **Member Functions**

# head Gets a pointer to the first character of this variable-size string. char \*head() const; Returns Pointer to the first character of this variable-size string. If the string contains no characters, returns a null pointer. length Gets the number of characters in this variable-size string. uint32 length() const; Returns Integer number of characters in this string preceding the first null terminating character as computed by strlen. If the string contains no characters, returns 0. Discussion The actual number of bytes allocated is at least length() + 1 because an extra byte is reserved for the null terminating character. (If the string contains an embedded null character, the bytes beyond that null character are not included in the returned length.) You can use the inherited <u>size</u> member function to determine the actual allocated size. resize Extends or truncates this variable-size string to the specified number of characters.

ooStatus resize(const uint32 newLength);

Parameters	newLength
	Number of characters this variable-size string is to have. Specify 0 to remove all the characters, freeing the storage allocated to the variable-size string's vector.
Returns	oocSuccess if successful; otherwise oocError.
Discussion	The actual number of bytes allocated is $newLength + 1$ because an extra byte is automatically reserved for the null terminating character. However, this member function does not actually store the required null character; the caller is responsible for updating the contents of the string.
Example	The following code fragment resizes a variable-size string and provide it with the null terminating character:
	<pre>ooVString s; s.resize(n); s[n] = '\0';</pre>

# **Topic Index**

This index lists topics that are discussed in this book. For a list of classes, see "Classes Index" on page 763. For a list of functions, including member functions, see "Functions and Macros Index" on page 773. For a list of non-class types and constants, see "Types and Constants Index" on page 787.

### Symbols

- \_ooDefaultContObj system name 239
- [] (see subscript operator)
- + (see addition operator)
- ++ (see increment operator)
- += (see append-to operator)
- += (see increment operator)
- (see subtraction operator)
- (see unary minus operator)
- -- (see decrement operator)
- -= (see decrement operator)
- -> (see indirect member-access operator)
- \* (see dereference operator)
- \* (see multiplication operator)
- \*= (see multiplication operator)
- / (see division operator)
- /= (see division operator)
- = (see assignment operator)
- == (see equality operator)
- != (see inequality operator)
- < (see less-than operator)
- <= (see less-than-or-equal-to operator)
- > (see greater-than operator)
- >= (see greater-than-or-equal-to operator)

## Numerics

8-bit integer array (see Java 8-bit integer array) 8-bit integer type signed 31 unsigned 79 16-bit integer array (see Java 16-bit integer array) **16-bit integer type** signed 32 unsigned 79 32-bit integer array (see Java 32-bit integer array) 32-bit integer type signed 32 unsigned 79 64-bit integer array (see Java 64-bit integer array) 64-bit integer type signed 32 unsigned 80

## A

aborting before Objectivity/DB shutdown 44 transaction 655 access mode (see also open mode) in predicate query 32 adding date 104. 105 element to list 671, 672, 673 to name map 415, 416 to scalable collection 176 to sorted object map 681, 684 to sorted set 691 to unordered object map 278, 281 to unordered set 287 to VArray 715, 716 interval 115. 140. 149 key field to key description 376 lookup field to lookup key 403, 406 time value 140 timestamp 149 addition operator (+) date 104 interval 115 time 140 timestamp 149 administrator 155 (see also ooAdmin in the Classes Index) hash administrator 269 tree administrator 661 Advanced Multithreaded Server (see AMS) AMS setting usage policy 66 timeout errors 69 usage policy type 32 appClass class 81 (see also application-defined class) append-to operator (+=) optimized string 649 variable-size string 726 application (see also multithreaded application) initializing Objectivity/DB in 49 standalone 58 terminating 43

application-defined class 81 (see also appClass in the Classes Index) (see also basic object) (see also container) adding association link to-many 85 to-one 90 deleting association link to-many 86, 90 to-one 86 finding destination objects to-many 87 to-one 87 generated member functions 82 handle for 471 iterator for 293 object reference for 471 referencing this instance of 89 short object reference for 631 testing association 86 application-defined functions conversion function 37 error handler 41 hash function for ooMap 55 message handler 55 relational operator 60 two-machine handler 75 application-defined operators (see operator set) array, variable-size (see VArray) assignment operator (=) date 106 handle for application-defined class 483 for autonomous partition 497 for container 521 for database 547 for federated database 579 for Objectivity/DB object 608 interval 116 name-map iterator 429 object reference for application-defined class 483 for autonomous partition 497

for container 521 for database 547 for federated database 579 for Objectivity/DB object 608 ODMG generic reference 127 optimized string 650 short object reference 634, 640 time 141 timestamp 150 Unicode string 705 variable-size string 726 VArray 713 iterator 123 temporary 699 association adding link to-many 85 to-one 90 deleting link to-many 86, 90 to-one 86 finding destination objects to-many 87 to-one 87 for versioning 434 generated member functions for add linkname 85 del linkname 86 exist linkname 86 linkname 87 set linkname 90 sub linkname 90 testing for existence of 86 autonomous partition 157 (see also ooAPObj in the Classes Index) (see also Objectivity/DB object) boot file 157 changing attributes 498 creating 159, 160 deleting 38 enforcing offline status of 68 finding all containers controlled by 500 all database images in 501 all in federated database 581

all that contain an image of database 563 boot autonomous partition 579 by system name 501, 505 federated database from 499 from controlled container 525 from database image 549 tie-breaker for database 556 getting boot-file host name 497 boot-file pathname 498 class name 507 iournal-file host name 503 journal-file pathname 503 lock-server host name 504 open mode 506 system name 505 system-database file host name 507 system-database file pathname 507 type name 507 type number 507 handle for 489 identifier 158 iterator for 299 marking 504 object reference for 489 offline status enforcing 68 marking 504 testing whether enforced 47 opening 501, 505 purging from federated database 59 returning container control 506 setting boot-file location 160 iournal-file location 160 lock server host 159 system name 159 system-database file location 159 tie-breaker partition for database 567 system name 158, 159 system-database file 157, 160 testing for database image in 556 for existence 501

whether available 502 whether offline 502 whether offline status is enforced 47 updating 508

## В

# B-tree (see scalable ordered collection) basic object

(see also appClass in the Classes Index) (see also ooObj in the Classes Index) (see also application-defined class) (see also persistent object) accessing a member of 481, 607 clustering 441 copying 486, 614 creating 439, 441 transient 441 deleting 38, 40, 439 without propagation 40 enabling versioning 628 finding all in container 525 container from 614 default version of 615 derivative versions 448, 449 genealogy for 452 next version(s) of 453, 617 previous version of 460, 618 getting open mode 625 handle for 593 moving 192, 622 object reference for 593 postprocessing for copying 455 for moving 457 for versioning 457 preprocessing, for moving 458 referencing this basic object 458 removing derivative versions 446, 463 from genealogy 446 next version 447, 464 previous version 447

testing for derivatives 451 for existence of derivative versions 450 for next version 452 for previous version 452 Boolean array (see Java Boolean array) **Boolean type** 33 boot autonomous partition 579 boot file of autonomous partition 157 changing host and path 498 getting host 497 getting path 498 setting host and path 160 of federated database 245 changing path 580 **buffer pages** 51 buffer pool (see Objectivity/DB cache) by opening database 552, 562 by opening federated database 587

# С

character array (see Java character array) character type 31 checkpointing lock mode 40 transaction 656 class getting type name of autonomous partition 507 of database 568 of federated database 591 of Objectivity/DB object 629 of persistent object 456 getting type number for persistent object 455 of autonomous partition 507 of database 568 of federated database 590 of Objectivity/DB object 629

classes (see also the Classes Index) (see also generated classes) (see also persistence-capable classes) clearing application-defined operators 467 error flags 65 closing database (ODMG) 96 federated database 581 handle to persistent object 524, 613 persistent object 598, 614 clustering basic object 239, 441 container 213 collection (see persistent collection) committing transaction 656 compacting B-tree 165, 682 comparator 197 (see also ooCompare in the Classes Index) comparison function 200 hashing function 200 of sorted collection 197 of unordered collection 198 comparing date 106 handle 610, 611 interval 117, 118 key field objects for consistency 377, 385 object reference 610, 611 ODMG generic reference 127, 128 optimized string 650, 651 persistent objects 200 predicate string 470 short object reference 641, 642 time 141, 142, 144 timestamp 150, 151 variable-size string 727, 728, 729 concurrent access policy **MROW 658** standard 658 constants 12 (see also the Types and Constants Index) constraints, checking 460

container 207 (see also appClass in the Classes Index) (see also ooContObj in the Classes Index) (see also application-defined class) (see also persistent object) accessing a member of 520 application-defined 208 closing 524 clustering 213 container object in 209 controlling partition of 208 converting objects after schema evolution 526 creating 212 multiple 56 transient 213 default in database 208, 234, 239 setting characteristics 236 deleting 38, 40, 439 without propagation 40 finding all basic objects in 525 all controlled by an autonomous partition 500 all in a database 550 autonomous partition that controls 525 by scope name 529 by system name 527, 531 database from 524 default in database 554 from basic object 614 garbage-collectible 208, 247 getting growth factor 533 hash value 527 number of logical pages 530 number of storage pages 530 open mode 532 pointer to 533 system name 530 handle for 509 hashed 208 creating 214 iterator for 303 kinds of 208

limit on 531 locking 528, 620 looking up 529 nonhashed 208 object reference for 509 opening 527, 531 to refresh view 534 pages in getting number of logical 530 getting number of storage 530 initial number 214 logical 209 page map 209 storage 209 referencing this container 215 refreshing view of 534 returning control of 535 setting growth factor 214 hashed 214 initial number of pages 214 location in federated database 213 system name 214 standard 208 system name 208, 214 testing for existence 527 whether updated 528 transferring control of 535 transient 209 creating 213 updating 534 context variables 26 conversion function 217, 225 getting member values 217 registering 589 setting member values 225 syntax for 37 conversion operator handle to d\_Ref\_Any type 484, 522, 611 to pointer type 484, 523, 612

object reference to d\_Ref\_Any type 484, 522, 611 to integer type 612 optimized string to C++ string type 651 to ooVString type 651 short object reference to integer type 643 Unicode string to C++ string type 705 variable-size string to C++ string type 730 conversion, object (see object conversion) converted object 225 (see also ooConvertInOutObject in the Classes Index) (see also object conversion) getting embedded part 228 inherited part 227 setting data-member values 228 copying basic object 486, 614 creating association (see association, adding) autonomous partition 159, 160 basic object 439, 441 container 212 garbage-collectible 249 hashed 214 multiple 56 database 63, 236, 237 database image 564 handle for autonomous partition 496 for basic object 605, 606 for container 518, 519 for database 546 for federated database 578 for Objectivity/DB object 605, 606 for persistent object 480, 605, 606 index 376 keyed object 57 name-map iterator 428 object iterator for autonomous partitions 300 for containers 304

for databases 310 for persistent objects 295, 315 object reference for autonomous partition 496 for basic object 606 for container 519 for database 546 for federated database 578 for Objectivity/DB object 606 for persistent object 480, 606 **Objectivity context 205** ODMG generic reference 126 short object reference 633, 639 string optimized 648 Unicode 705 variable-size 725 transaction object 655 transient object 441 transient object (ODMG) 95 VArray 712 iterator 121 temporary 698 current Objectivity context 203 getting 205 setting 206 customer support 9

# D

d\_Boolean 25 d\_Char 25 d\_Database class 93 (see also database (ODMG)) d\_Date class 99 (see also date) d\_Double 25 d\_Float 25 d\_Float 25 d\_Interval class 111 (see also interval) d\_Iterator<element\_type> class 119 (see also VArray iterator) d\_Long 25

d\_Object class 431 (see also Objectivity/DB object) d\_Octet 25 d\_Ref\_Any class 125 (see also ODMG generic reference) d\_Ref<appClass> class 471 (see also object reference) d\_Ref<d\_Object> class 593 (see also object reference) d Short 25 d String class 721 (see also string) d Time class 131 (see also time) d Timestamp class 145 (see also timestamp) d Transaction class 653 (see also transaction) d\_ULong 25 d UShort 25 d\_Varray<element\_type> class 707 Data Definition Language (DDL) 82 data member getting value during object conversion 217 setting value during object conversion 225 data replication option (see Objectivity/DRO) database 233 (see also ooDBObj in the Classes Index) (see also database image) (see also Objectivity/DB object) changing containing partition 548 file location 547 converting objects after schema evolution 551 creating 63, 236, 237 an image of 564 default container 234 setting characteristics of 236 deleting 38 enabling nonquorum reads 565

file 64. 233 setting host and path 236 finding all containers in 550 all in federated database 581 by system name 552, 562 default container in 554 federated database from 549 from container 524 getting class name 568 filename 553 host name 557 number of containers in 561 number of images 562 open mode 563 pathname 564 system name 560 type name 568 type number 568 handle for 537. 541 identifier 234. 538 image (see database image) iterator for 309 locking 559 object reference for 537, 541 opening 562 read-only 566 recalculating quorum 561 replacing 63 replicating 564 setting characteristics of default container 236 file location 236 identifier 237 read-only 566 system name 236 weight 237 system name 234, 236 testing for existence 552 for multiple images 559 whether available 557 whether nonquorum reads are allowed 553

whether read-only 558 whether reading without a quorum 558 tidying 567 updating 569 database image 234 counting 562 creating 564 deleting 551 finding all containing partitions 563 all in an autonomous partition 501 autonomous partition that contains 549 tie-breaker partition for 556 getting filename 554 host name 555 pathname 555 weight 556 identifier 234 quorum 234 allowing reads without 553, 565 recalculating 561 testing whether reading without 558 setting weight 566 of first 237 system name 234 testing an autonomous partition for 556 whether available 557 tie-breaker partition finding 556 setting 567 weight 234 database (ODMG) 93 (see also d\_Database in the Classes Index) closing 96 creating transient object 95 finding persistent object 96 Objectivity/DB database and 94 open mode 95 opening 97 scope name in changing 97 getting 96

looking up 96 setting 98 date 99 (see also d Date in the Classes Index) (see also Java date) adding 104, 105 assigning 106 creating 103 decrementing 105, 109 getting current 107 day of the month 107 day of the week 107 day of the year 107 month 108 number of days in a month 108 number of days in a year 108 year 109 incrementing 104, 109 months, type for 102 subtracting 105 testing equality 106 for leap year 108 for overlap 108, 109 for validity 108 valid 108 weekdays, type for 102 date and time classes d date 99 d Interval 111 d\_Time 131 d Timestamp 145 oojDate 359 oojTime 365 oojTimestamp 369 decrement operator (--, -=) date 105 interval 116 time 140 timestamp 150 VArray iterator 122

(see also ooDefaultContObj in the Classes Index) finding from database 554 default\_odmg\_db system name 94 deinitializing Objectivity/DB DLL 43 delete operator 439 deleting application-defined operators from operator set 467 association link 86.90 autonomous partition 38 database 38 database image 551 element from list 675. 676 from name map 420 from scalable collection 177, 181, 182 from scalable ordered collection 170 from sorted object map 685 from sorted set 693 from unordered object map 281, 282 from unordered set 290 from VArray 717 error flags 65 index 378 persistent object 38, 40, 439 persistent object (ODMG) 129, 615 dereference operator (\*) handle for application-defined class 482 for container 521 for persistent object 608 derivative versions adding 442, 443 association for 434 finding all 448, 449 removing 445, 446, 462, 463 testing for existence 450, 451 destroying Objectivity context 205 dividing interval 116 division operator (/=, /) interval 116 double array (see Java array of double)

default container 239

**DRO abbreviation** 8, 11 **dropping index** 377, 378

### Ε

environment variables OO DB NAME 94 OO FD BOOT 587 equal lookup field 241 (see also ooEqualLookupField in the Classes Index) creating 242 equality operator (==) date 106 handle 610 interval 117. 118 object reference 610 ODMG generic reference 127 optimized string 650 short object reference 641 time 141 variable-size string 727 error flags 78, 79 clearing 65 handler application-defined 41 getting pointer to 46 registering 61 identifier 41 level, indicating 42 message output file 66 signal, raising 69 exiting process 43 extending temporary VArray 701

# F

fault tolerant option (see Objectivity/FTO) federated database 245 (see also ooFDObj in the Classes Index) (see also Objectivity/DB object) (see also database (ODMG))

boot file 245 changing boot file path 580 identifier 580 lock server host 580 closing 581 converting objects after schema evolution 582, 591 creating 246 deleting 246 finding all autonomous partitions in 581 all databases in 581 from autonomous partition 499 from database 549 getting catalog information 584 class name 591 identifier 587 lock server host 586 open mode 588 page size 588 system name 586 type name 591 type number 590 handle for 571 listing files 584 locking 585 object reference for 571 opening 584, 587, 659 printing information 584 system name 245 system-database file 245 testing for existence 584 tidying 589 updating 591 upgrading after schema evolution 591 file boot (see also boot file) of autonomous partition 157 of federated database 245 database 233. 553

system-database file of autonomous partition 157 of federated database 245 file descriptors 50 filename, output format type 44 finding autonomous partitions in a federated database 581 database images in an autonomous partition 501 databases in a federated database 581 destination objects to-many 87 to-one 87 genealogy 452 from a default version 444 keyed object 621 persistent object by scope name 487, 621 by scope name (ODMG) 96 scope objects 616 versions from genealogy 254, 256 float array (see Java array of float) floating-point types 31 FTO abbreviation 8.11 function-pointer types ooConvertFunction 37 ooErrorHandlerPtr 41 ooMsgHandlerPtr 55 ooNameHashFuncPtr 55 ooQueryOperatorPtr 60 ooTwoMachineHandlerPtr 75 ooVoidFuncPtr 79 functions 12 (see the Functions and Macros Index)

# G

garbage-collectible container 247 (see also ooGCContObj in the Classes Index) creating 249 genealogy 251 (see also ooGeneObj in the Classes Index)

adding version to 255 association to ooObj 434 creating 254 finding all versions in 254 default version in 256 from any version 452 from default version 444 referencing this genealogy 258 removing all versions from 257 default version from 257 specified version from 260, 446 setting default version 259, 461 testing for existence of any versions 258 of default version 258, 450 testing for membership in 451 generated classes name, for template appClass 294, 475, 632 ooItr(appClass) 293 ooRefHandle(appClass) 471 ooShortRef(appClass) 631 getting current Objectivity context 205 element temporary VArray 701 VArray 715 handle for container 215 for garbage-collectible container 249 for instance of application-defined class 89 for persistent object 458 object reference for container 215 for garbage-collectible container 249 for instance of application-defined class 89 for persistent object 458 global functions 23 (see the Functions and Macros Index) global macros 23 (see the Functions and Macros Index)

global types 25 (see also the Types and Constants Index) global variables 26 (see also the Types and Constants Index) greater-than lookup field 265 (see also ooGreaterThanLookupField in the Classes Index) creating 266 greater-than operator (>) timestamp 151 variable-size string 729 greater-than-equal lookup field 261 (see also ooGreaterThanEqualLookupField in the Classes Index) creating 262 greater-than-or-equal-to operator (>=) timestamp 151 variable-size string 729

# Η

handle 471, 489, 509, 537, 571, 593 (see also the Classes Index for: ooRefHandle(appClass) 471 ooRefHandle(ooAPObj) 489 ooRefHandle(ooContObj) 509 ooRefHandle(ooDBObj) 537 ooRefHandle(ooFDObj) 571 ooRefHandle(ooObj)) 593 accessing a member of referenced container 520 of referenced persistent object 481, 607 assigning to 497, 521, 547, 579, 608 class name of, for template appClass 475 closing 524, 613 creating for autonomous partition 496 for basic object 605, 606 for container 518. 519 for database 546 for federated database 578 for Objectivity/DB object 605, 606 for persistent object 480, 605, 606

dereferencing 521, 608 extracting a pointer from 533, 626 inheritance hierarchy of classes 594 open and closed states 511, 598 setting by assignment 497, 521, 547, 579 by looking up a container 529 by looking up a persistent object 621 by opening a container 527, 531 by opening an autonomous partition 501.505 by opening the federated database 584, 587 container information 627 to null 522. 609 structure and behavior 597 testing for equality 610 for inequality 611 for null 612 for validity 502, 559, 585, 619 whether null 619 whether null (ODMG) 619 type conversion to appClass pointer type 484 to d\_Ref\_Any type 484, 522, 611 to ooContObj pointer type 523 to ooObj pointer type 612 using in a conditional expression 612 hash administrator 269 (see also ooHashAdmin in the Classes Index) getting current hash-bucket container 270 maximum buckets per container 271 hash function pointer type 55 hash table extendible hash buckets 269 containers for 269, 270, 271 hash function 198, 200

of name map growth characteristics 410 growth factor 410 getting 418 setting 415 hash buckets getting number of 418 initial number of 410 setting 414 hash function 411 getting 418 setting 421 maximum average density 410 getting 417 setting 414 hashed container 208 header files javaBuiltins.h 319 ooCollections.h 173, 185, 197 ooMap.h 409, 423, 427 ooRecover.h 36, 46, 49, 65, 74

ooTime.h 99, 111, 131, 145

hot mode, setting 66

# 

identifier (see also object identifier) of autonomous partition 158 of database 234, 538 of transaction 74, 653, 657 image (see database image) include files (see header files) increment operator (++, +=)date 104. 105 interval 115 time 140 timestamp 149 VArray iterator 122 index adding lookup field 403 creating 373, 376 lookup field 402 disabling use 77

dropping 377, 378 enabling use 77 finding objects with a lookup key 402 initializing iterator for 403 key description 373 key field, strings 380 looking up 407 lookup key 399 number of fields 378 scope 373 uniqueness 378 update mode 49 updating explicitly 76 indirect member-access operator (->) handle for application-defined class 481 for container 520 for persistent object 607 object reference for application-defined class 481 for container 520 for persistent object 607 inequality operator (!=) date 106 handle 611 interval 117 object reference 611 ODMG generic reference 128 optimized string 651 short object reference 642 time 141, 142, 144 timestamp 150 variable-size string 727 initializing key structure 46, 47 object iterator for autonomous partitions 301 for containers 305 for databases 311 for persistent objects 296, 316 **Objectivity/DB 49** thread 51

Topic Index

in-process lock server starting 70 stopping 72 testing for running lock servers 33 integer types signed 31 unsigned 79 interoperating with Java or Smalltalk 248 interval 111 (see also d\_Interval in the Classes Index) adding 115 assigning 116 creating 114 dividing 116 getting day component 118 hour component 118 minute component 118 seconds component 118 multiplying 116 returning negative 115 subtracting 115, 116 testing equality 117, 118 for zero duration 118 **IPLS** abbreviation 8 iteration set name-map iterator 427 object iterator 313 type for filtering by partition 37 iterator (see also object iterator) (see name-map iterator) (see object iterator) (see scalable-collection iterator) (see VArray iterator)

### J

Java 8-bit integer array 339 (see also oojArrayOfInt8 in the Classes Index) (see also Java persistent array) creating 340 getting VArray 341 Java 16-bit integer array 343 (see also oojArrayOfInt16 in the Classes Index) (see also Java persistent array) creating 344 getting VArray 345 Java 32-bit integer array 347 (see also oojArrayOfInt32 in the Classes Index) (see also Java persistent array) creating 348 getting VArray 349 Java 64-bit integer array 351 (see also oojArrayOfInt64 in the Classes Index) (see also Java persistent array) creating 352 getting VArray 353 Java array of double 331 (see also oojArrayOfDouble in the Classes Index) (see also Java persistent array) creating 332 getting VArray 333 Java array of float 335 (see also oojArrayOfFloat in the Classes Index) (see also Java persistent array) creating 336 getting VArray 337 Java Boolean array 323 (see also oojArrayOfBoolean in the Classes Index) (see also Java persistent array) creating 324 getting VArray 325

Java character array 327 (see also oojArrayOfCharacter in the Classes Index) (see also Java persistent array) creating 328 getting VArray 329 Java compatibility classes oojArray 319 ooiArravOfBoolean 323 oojArrayOfCharacter 327 oojArrayOfDouble 331 oojArrayOfFloat 335 oojArrayOfInt8 339 oojArrayOfInt16 343 oojArrayOfInt32 347 oojArrayOfInt64 351 oojArrayOfObject 355 oojDate 359 oojString 363 oojTime 365 oojTimestamp 369 ooUtf8String 703 Java date 359 (see also oojDate in the Classes Index) creating 360 millisecond representation 359 getting 360 setting 361 Java interoperability 248 Java object-reference array 355 (see also oojArrayOfObject in the Classes Index) (see also Java persistent array) creating 357 getting VArray 357 Java persistent array 319 (see also oojArray in the Classes Index) getting dimensions 321 of Boolean elements 323 of characters 327 of floating-point numbers double-precision 331 single-precision 335

of integers 8-bit 339 16-bit 343 32-bit 347 64-bit 351 of object references 355 Java string element 363 (see also oojString in the Classes Index) creating 364 getting Unicode string 364 Java time 365 (see also oojTime in the Classes Index) creating 366 millisecond representation 365 getting 366 setting 367 Java timestamp 369 (see also oojTimestamp in the Classes Index) creating 371 fractional part 370 getting 371 setting 372 integral part 369 getting 371 setting 372 javaBuiltins.h header file 319

# K

key description 373 (see also ooKeyDesc in the Classes Index) adding key field 376 creating 375 creating index 376 deleting index 378 dropping index 377 getting name of indexed class 377 number of fields 378 type number of indexed class 377 testing for consistency 377 index for uniqueness 378 key field 379 (see also ooKeyField in the Classes Index) adding to key description 376 creating 383 getting data-member name 384 type number 385 testing data-member name 385 for consistency 385 keyed object creating 57 finding 621 hash clustering factor 56, 214, 527 key field 53 key structure 46, 47, 52 size of member field 47

# L

large objects dynamically-allocated memory for 51, 67 limiting memory for 67 less-than lookup field 391 (see also ooLessThanLookupField in the Classes Index) creating 392 less-than operator (<) timestamp 150 variable-size string 728 less-than-equal lookup field 387 (see also ooLessThanEqualLookupField in the Classes Index) creating 388 less-than-or-equal-to operator (<=) timestamp 151 variable-size string 728 list 667 (see also ooTreeList in the Classes Index) (see also persistent collection) adding elements 671, 672, 673 creating 670

finding object first element 674 looking up data 674 looking up index 674 tree administrator 673 getting iterator for elements 675 removing elements 675, 676 replacing element 676 testing for contained elements 674 lock server changing host 580 disabling use of 58 in-process (see in-process lock server) testing whether running 33 timeout errors 69 locking container 528 database 559 federated database 585 lock mode 54 while checkpointing 40 lock wait, setting 68 persistent object 620, 621 lookup field 395 base class 395 creating 242, 262, 266, 388, 392, 396 equal-to 241 greater-than 265 greater-than-or-equal-to 261 less-than 391 less-than-or-equal-to 387 lookup key 399 (see also ooLookupKey in the Classes Index) adding lookup field 406 creating 402, 406 getting number of added fields 408 testing for compatible index 407 lookup-field base class 395 (see also ooLookupFieldBase in the Classes Index) comparing name 396 creating lookup field 396

## Μ

macros 12 (see the Functions and Macros Index) preprocessing 23 member field offset for keyed object 46 size for keyed object 47 member functions (see the Functions and Macros Index) message handler application-defined 55, 75 getting pointer to 47 registering 61 mode open (see open mode) versioning (see versioning mode) moving basic object 622 MROW disabling 658 enabling 658 multiplication operator (\*=, \*) interval 116 multithreaded application (see also application) (see also Objectivity context) initializing threads 51 **Objectivity contexts and 203** Objectivity/C+ variables and 26 terminating thread 73

### Ν

name map 409 (see also ooMap in the Classes Index) (see also persistent collection) adding new element 415 checking for name 416 creating 414 elements of 423 finding object by looking up name 417 forcing addition of element 416 getting number of elements 418

hash table getting hash function 418 number of hash buckets 418 growth characteristics 410 growth factor 410 getting 418 setting 415 initial number of hash buckets 410 setting 414 maximum average density 410 getting 417 setting 414 resizing 419 iterator for 427 referential integrity 410 checking status 419 setting status 421 removing element 420 replacingvalue 420 runtime statistics 410 getting 418 resetting 416 name-map element 423 (see also ooMapElem in the Classes Index) getting the key from 424 getting the value from 424 setting the value in 425 name-map iterator 427 (see also ooMapItr in the Classes Index) advancing to next name-map element 429 creating 428 initializing 427, 429 iteration set 427 named root 248 naming conventions for system name of autonomous partition 159 of database 236 Objectivity/C++ 12 system-database file name for autonomous partition 160

new operator autonomous partition 160 basic object 441 container 212 database 237 next version adding 444 association for 434 finding 453 removing 447, 464 testing for existence 452 nonhashed container 208 nonquorum reads enabling 565 testing for 553, 558 null Objectivity context defined 204 setting 206

# 0

object 431 (see also ooObj in the Classes Index) (see also persistent object) (see also autonomous partition) (see also basic object) (see also container) (see also database) (see also federated database) (see also federated database) (see also large objects) (see also Objectivity/DB object) (see also persistent object) (see also transient object)

#### object conversion

conversion function 37, 217, 225 getting member values 217 in container 526 in database 551 in federated database 582, 589 by upgrade application 591 registering conversion functions 589 setting member values 225 object identifier (OID) getting as string 628 printing 625, 644 object iterator 293, 299, 303, 309, 313 (see also the Classes Index for: ooItr(appClass) ooItr(ooAPObj) ooItr(ooContObj) ooItr(ooDBObj) ooItr(ooObj)) advancing for autonomous partitions 301 for containers 305 for databases 311 for persistent objects 295, 316 creating for autonomous partitions 299, 300 for containers 304 for databases 310 for persistent objects 295, 315 generated, for template appClass 294 initializing for autonomous partitions 301 for containers 305 for databases 311 for persistent objects 87, 296, 316 iteration set 313 terminating iteration 314 for autonomous partitions 300 for containers 304 for databases 310 for persistent objects 295, 316 object map 273 sorted object map 677 unordered object map 273 object reference 471, 489, 509, 537, 571, 593 (see also the Classes Index for: ooRefHandle(appClass) 471 ooRefHandle(ooAPObj) 489 ooRefHandle(ooContObj) 509 ooRefHandle(ooDBObj) 537 ooRefHandle(ooFDObj) 571 ooRefHandle(ooObj)) 593

accessing a member of referenced container 520 of referenced persistent object 481, 607 assigning to 497, 521, 547, 579, 608 class name of, for template appClass 475 creating for autonomous partition 496 for basic object 606 for container 519 for database 546 for federated database 578 for Objectivity/DB object 606 for persistent object 480, 606 extracting a pointer from 533, 626 inheritance hierarchy of classes 594 setting by assignment 497, 521, 547, 579 by looking up a container 529 by looking up a persistent object 621 by opening a container 527, 531 by opening an autonomous partition 501.505 by opening database 552, 562 by opening the federated database 584, 587 container information 627 to null 522, 609 structure and behavior 597 testing for equality 610 for inequality 611 for null 612 for validity 502, 559, 585, 619 whether null 619 whether null (ODMG) 619 type conversion to d\_Ref\_Any type 484, 522, 611 to integer type 612 using in a conditional expression 612 object-reference array (see Java object-reference array) **Objectivity context** 203 (see also ooContext in the Classes Index) creating 205 current 203

destroying 205 getting current 205 managing settings global functions for 24 variables for 26 null 204, 206 setting current 206 to null 206 **Objectivity/C++** classes (see also the Classes Index) naming conventions 12 **Objectivity/DB** initializing 49 internal statistics 65 preparing for shutdown 43 final Objectivity/DB operation before 43 multithreaded application 44 platform-specific considerations 43, 44 single-threaded application 44 terminating DLL 43 **Objectivity/DB cache 50** buffer pages 51 in main thread 51 large objects 67 large-object buffer pool 51, 67 large-object memory pool 51, 67 size 51 setting 50, 205 small-object buffer pool 51 **Objectivity/DB Data Replication Option** (see Objectivity/DRO) **Objectivity/DB Fault Tolerant Option** (see Objectivity/FTO) **Objectivity/DB object** 431 getting object identifier, as string 628 object identifier, printed to file 625 type name 629 type number 629 handle for 593 object reference for 593

**Objectivity/DRO** 233, 537 (see also database image) abbreviation 11 **Objectivity/FTO** 157, 489 (see also autonomous partition) abbreviation 11 **ODMG** abbreviation 8 **ODMG classes** d Database 93 d Date d\_Interval 111 d Iterator<element type>119 d Object 431 d Ref Any 125 d\_Ref<appClass> 471 d Ref<d Object> 593 d String 721 d Time 131 d\_Timestamp 145 d\_Transaction 653 d\_Varray<element\_type>707 **ODMG generic reference** 125 (see also d\_Ref\_Any in the Classes Index) assigning 127 creating 126 deleting persistent object 129 setting to null 128 testing for equality 127, 128 for null 129 **ODMG** primitive data types (see type, primitive) offline mode getting 47 response type 58 setting 68 offline status marking autonomous partition 504 testing autonomous partition 502 offset. member field 46 online status, marking autonomous partition 504 **OO\_COMMA symbol** 294, 475, 632 **OO DB NAME environment variable** 94

**OO\_FD\_BOOT environment variable 587** ooAdmin class 155 (see also administrator) ooAPObj class 157 (see also autonomous partition) ooBTree class 163 (see also scalable ordered collection) ooCollection class 173 (see also scalable collection) ooCollectionIterator class 185 (see also scalable-collection iterator) ooCollections.h header file 173, 185, 197 ooCompare class 197 (see also comparator) ooContext class 203 (see also Objectivity context) ooContObi class 207 (see also container) ooConvertInObject class 217 (see also unconverted object) ooConvertInOutObject class 225 (see also converted object) ooDBObj class 233 (see also database) ooDefaultContObj class 239 (see also default container) ooEqualLookupField class 241 (see also equal lookup field) ooFDObj class 245 (see also federated database) ooGCContObj class 247 (see also garbage-collectible container) ooGCRootsCont class 248 ooGeneObj class 251 (see also genealogy) ooGreaterThanEqualLookupField class 261 (see also greater-than-equal lookup field) ooGreaterThanLookupField class 265 (see also greater-than lookup field) ooHandle(appClass) class 471 (see also handle) ooHandle(ooAPObj) class 489 (see also handle)

0

ooHandle(ooContObj) class 509 (see also handle) ooHandle(ooDBObj) class 537 (see also handle) ooHandle(ooFDObj) class 571 (see also handle) ooHandle(ooObj) class 593 (see also handle) ooHashAdmin class 269 (see also hash administrator) ooHashMap class 273 (see also unordered object map) ooHashSet class 283 (see also unordered set) ooItr(appClass) class 293 (see also object iterator) ooItr(ooAPObj) class 299 (see also object iterator) ooItr(ooContObj) class 303 (see also object iterator) ooItr(ooDBObj) class 309 (see also object iterator) ooItr(ooObj) class 313 (see also object iterator) ooiArrav class 319 (see also Java persistent array) oojArrayOfBoolean class 323 (see also Java Boolean array) oojArrayOfCharacter class 327 (see also Java character array) oojArrayOfDouble class 331 (see also Java array of double) ooiArravOfFloat class 335 (see also Java array of float) oojArrayOfInt8 class 339 (see also Java 8-bit integer array) ooiArravOfInt16 class 343 (see also Java 16-bit integer array) oojArrayOfInt32 class 347 (see also Java 32-bit integer array) oojArrayOfInt64 class 351 (see also Java 64-bit integer array) oojArrayOfObject class 355

(see also Java object-reference array) oojDate class 359 (see also Java date) oojString class 363 (see also Java string element) ooiTime class 365 (see also Java time) oojTimestamp class 369 (see also Java timestamp) ooKeyDesc class 373 (see also key description) ooKevField class 379 (see also key field) ooLessThanEqualLookupField class 387 (see also less-than-equal lookup field) ooLessThanLookupField class 391 (see also less-than lookup field) ooLookupFieldBase class 395 (see also lookup-field base class) ooLookupKev class 399 (see also lookup key) ooMap class 409 (see also name map) ooMapElem class 423 (see also name-map element) ooMapItr class 427 (see also name-map iterator) **ooMap.h header file** 409, 423, 427 ooObj class 431 (see also object) ooOperatorSet class 465 (see also operator set) ooQuery class 469 (see also query object) ooRecover.h header file 36, 46, 49, 65, 74 ooRefHandle abbreviation 14 ooRef(appClass) class 471 (see also object reference) ooRef(ooAPObj) class 489 (see also object reference) ooRef(ooContObi) class 509 (see also object reference)

ooRef(ooDBObj) class 537 (see also object reference) ooRef(ooFDObj) class 571 (see also object reference) ooRef(ooObj) class 593 (see also object reference) ooschemadump tool 583, 592 ooShortRef(appClass) class 631 (see also short object reference) ooShortRef(ooObj) class 637 (see also short object reference) ooString(N) class 645, 722 (see also optimized string) **ooTime.h header file** 99, 111, 131, 145 ooTrans class 653 (see also transaction) ooTreeAdmin class 661 (see also tree administrator) ooTreeList class 667 (see also list) ooTreeMap class 677 (see also sorted object map) ooTreeSet class 687 (see also sorted set) ooTVArrayT<element\_type> class 695 (see also temporary VArray) ooTVArray(element type) class 695 ooUtf8String class 703 (see also Unicode string) ooVArrayT<element\_type> class 707 (see also VArray) ooVArray(element\_type) class 707 ooVString class 721 (see also variable-size string) open mode data type 54 for ODMG 95 getting for autonomous partition 506 for basic object 625 for container 532 for database 563 for federated database 588

opening autonomous partition 501, 505 container 527. 531 database 552, 562 database (ODMG) 97 federated database 584, 587, 659 persistent object 624 operator delete (see delete operator) operator new (see new operator) operator set 465 (see also ooOperatorSet in the Classes Index) clearing application-defined operators 467 creating 466 default 465 defining operators for 60 registering application-defined operator 467 variable containing 77 optimized string 645 (see also ooString(N) in the Classes Index) adding characters 649 appending to 649 assigning to 650 creating 648 getting character 649 first character 651 length 652 replacing characters 650 resizing 652 testing for equality 650 for inequality 651 type conversion to C++ string type 651 to ooVString type 651 ordered collection array containers getting current container 665 maximum arrays per container 662, 663, 664

node containers getting current container 664 maximum nodes per container 661, 663, 664

## Ρ

page map 209 pages buffer 51 logical 209 storage 209 persistence-capable classes appClass 81 ooAdmin 155 ooBTree 163 ooCollection 173 ooCompare 197 ooContObj 207 ooDefaultContObj 239 ooGCContObj 247 ooGeneObj 251 ooHashAdmin 269 ooHashMap 273 ooHashSet 283 oojArray 319 oojArrayOfBoolean 323 oojArrayOfCharacter 327 oojArrayOfDouble 331 oojArrayOfFloat 335 oojArrayOfInt8 339 oojArrayOfInt16 343 oojArrayOfInt32 347 oojArrayOfInt64 351 oojArrayOfObject 355 oojDate 359 oojString 363 oojTime 365 oojTimestamp 369 ooKeyDesc 373 ooKeyField 379 ooMap 409 ooMapElem 423 ooObj 431 ooTreeAdmin 661

ooTreeList 667 ooTreeMap 677 ooTreeSet 687 persistent collection nonscalable unordered name map 409 scalable 173 list 667 ordered 163 sorted object map 677 sorted set 687 unordered object map 273 unordered set 283 persistent object (see also ooObj in the Classes Index) (see also basic object) (see also container) (see also Objectivity/DB object) accessing a member of 481, 607 closing 598, 614 converted 225 creating 439, 441 deleting 38, 439 without propagation 40 deleting (ODMG) 129, 615 finding all in scope 616 by scope name 487, 621 getting class name 456 pointer to 626 reference to 608 scope name 617 type number 455 handle for 593 identifying type of 456 iterator for 293, 313 locking 620, 621 looking up 487, 621 naming 623 naming (ODMG) 98 object reference for 593 opening 624

testing for validity 460 type of 456 unconverted 217 unnaming 630 updating 453, 459, 630 virtual-function table 33 persistent-collection classes ooBTree 163 ooCollection 173 ooHashMap 273 ooHashSet 283 ooMap 409 ooMapElem 423 ooTreeList 667 ooTreeMap 677 ooTreeSet 687 pointer to persistent object extracting from handle or object reference 626 pool, buffer (see Objectivity/DB cache) predicate query access mode 32 enabling use of index 77 previous version adding 462 association for 434 finding 460 removing 447 testing for existence 452 process termination 43 programmer-defined functions (see application-defined functions) purging autonomous partitions 59

# Q

query object 469
 (see also ooQuery in the Classes Index)
 comparing predicate string 470
 setting up 470
quorum of database images 234

### R

read-only database 566 recovering transaction 35, 45, 48 referential integrity of list, restoring 675 of name map 410 checking status 419 setting status 421 of scalable collection, restoring 182 of sorted object map, restoring 685 of unordered object map, restoring 282 registering application-defined operator 467 conversion function 589 error handler 61 message handler 61 predefined signal handler 50 two-machine handler 62 registration code virtual-function table 34 relational operator functions application-defined 60 replacing database 63 element list 676 VArray 717 value name map 420 sorted object map 684 unordered object map 281 resizing hash table of name map 419 optimized string 652 temporary VArray 700, 701 variable-size string 730 VArray 715, 716, 718 return type, general 72 **RPC**, setting timeout period 69 runtime statistics, name map 416 runtime type identification (RTTI) 456
### S

scalable collection 173 (see also ooCollection in the Classes Index) adding elements 176 administrator 174 finding 177 comparator 174 finding 177 finding object administrator 177 comparator 177 looking up data 179 getting iterator for elements 179 iterator for keys 180 iterator for values 183 number of elements 183 making empty 177 refreshing internal containers 180 removing elements 181, 182 testing for contained elements 177, 178 for empty collection 179 scalable ordered collection 163 (see also ooBTree in the Classes Index) (see also scalable collection) B-tree 163 compacting 165 getting depth 166 finding object first element 166 last element 169 looking up data 167 searching backward 169 searching forward 167 getting iterator for elements 168 number of elements 171 refreshing internal containers 170 removing element 170 testing for contained element 166 for empty collection 168

scalable-collection iterator 185 (see also ooCollectionIterator in the Classes Index) current element 185 finding 189 finding value object for 189 moving 192 current index 185 getting 189 setting to position of object 190 to specified index 191 finding objects corresponding collection 189 current element 189 element at specified index 191 next element 192 previous element 193 value object for current key 189 getting index of next element 193 index of previous element 194 initializing for elements 168, 179, 290, 675 for keys 180, 280 for values 183, 282, 685 iteration set 185 moving current element 192 removing element from corresponding collection 194 replacing element in corresponding collection 195 testing for next element 191 for previous element 191 scope name default container and 239 getting 617 getting (ODMG) 96 hashed container and 208 looking up a persistent object by 621 removing 630 setting 623 setting (ODMG) 98

valid 623 valid (ODMG) 97, 98 scope objects, finding 616 set (see also scalable collection) sorted set 687 unordered set 283 setting AMS usage policy 66 current Objectivity context 206 error message output file 66 hot mode 66 lock wait 68 offline mode 68 **RPC timeout period 69** scope name 623 scope name (ODMG) 98 space for large objects 67 tie-breaker partition 567 short object reference 631, 637 (see also ooShortRef(appClass) in the Classes Index) (see also ooShortRef(ooObj) in the Classes Index) assigning to 634, 640 class name of, for template class 632 conversion to integer type 643 creating 639 for application-defined class 633 getting object identifier, as string 644 object identifier, printed to file 644 testing for equality 641 for inequality 642 for null 643 signal handler, predefined registering 50 suppressing 50 signed integer types 31 Smalltalk interoperability 248 sorted object map 677 (see also ooTreeMap in the Classes Index) (see also persistent collection)

adding elements 681, 684 compacting the B-tree 682 creating 680 finding object looking up data 683 looking up key 683 getting iterator for values 685 removing elements 685 replacing value 684 testing for contained keys 682 for contained values 683 sorted set 687 (see also ooTreeSet in the Classes Index) (see also persistent collection) adding elements 691 creating 690 finding object comparator 691 looking up data 692 looking up index 692 tree administrator 691 removing elements 693 testing for contained elements 692 standalone application 58 standard access, enabling 658 standard container 208 standard VArray (see VArray) starting transaction 656, 657 statistics. internal 65 status, return type 72 string (see Java string element) (see optimized string) (see Unicode string) (see variable-size string) string element (see Java string element) subscript operator ([]) optimized string 649 temporary VArray 699 variable-size string 726 VArray 714

subtraction operator (-) date 105 interval 115 time 140 timestamp 149 suppressing predefined signal handler 50 system name default\_odmg\_db 94 of autonomous partition 158 getting 505 naming conventions for 159 of container 208, 530 of database 234 getting 560 naming conventions for 236 of default container 239 of federated database 245 getting 586 system-database file of autonomous partition 157, 507 of federated database 245 system-defined class (see the Classes Index)

### Т

template classes name of generated class handle 475 object iterator 294 object reference 475 short object reference 632 temporary VArray 695 (see also ooTVArrayT<element\_type> in the Classes Index) (see also VArray) assigning to 699 creating 698 default constructor 696 element type, valid 695 extending 700 getting current number of elements 702 first element 701 specified element 699, 700

resizing 700, 701 effect of 696 setting element value 701 structure and behavior 696 updating 702 terminating application 43 Objectivity/DB DLL 43 thread 73 transaction 655, 656 thread initializing 51 setting Objectivity context 206 terminating 73 tie-breaker partition finding 556 setting 567 time 131 (see also d\_Time in the Classes Index) (see also Java time) adding 140 assigning 141 creating 138 decrementing 140 getting current 142 hours 142 minutes 142 seconds 143 time zone hour 143 time zone minute 143 incrementing 140 setting default time zone 143 subtracting 140 testing equality 141, 142, 144 for overlap 144 time zone type 137 time and date classes (see date and time classes) timeout period, setting RPC 69 timestamp 145 (see also d\_Timestamp in the Classes Index)

(see also Java timestamp) adding 149 assigning 150 creating 147 decrementing 150 getting current 151 date 151 day 152 hour 152 minute 152 month 152 seconds 152 time value 152 time zone hour 152 time zone minute 153 vear 153 incrementing 149 subtracting 149 testing equality 150 for overlap 153 greater-than-or-equal-to 151 less-than-or-equal-to 150, 151 tools ooschemadump 583, 592 transaction 653 (see also ooTrans in the Classes Index) aborting 655 active 653 checkpointing 656 committing 656 committing and holding 656 creating transaction object 655 disabling locks 58 getting identifier of 657 holding resources 48, 65 identifier 74, 653, 657

information about 74

listing active 45

lock wait policy 68

mode for MROW 54 recovering 35, 45, 48

specifying upgrade application 659

terminating 655, 656 testing whether active 657 transaction object 653 creating 653 transaction-information structure 74 transaction, starting 656, 657 transferring control of container 535 transient object, creating 441 tree administrator 661 (see also ooTreeAdmin in the Classes Index) getting current array container 665 current node container 664 maximum arrays per container 663, 664 maximum nodes per container 663, 664 two-machine handler application-defined 75 registering 62 type permitted for index 380 primitive alternative names for 25 summary of 25 type conversion (see conversion operator) type name getting autonomous partition 507 database 568 federated database 591 Objectivity/DB object 629 persistent object 456 type number 76 getting 75 autonomous partition 507 database 568 federated database 590 Objectivity/DB object 629 persistent object 455 types (see the Types and Constants Index)

### U

unary minus operator (-) interval 115 unconverted object 217 (see also ooConvertInObject in the Classes Index) (see also object conversion) getting embedded part 222 inherited part 221 member values 219 **Unicode string** 703 (see also ooUtf8String in the Classes Index) assigning to 705 creating 705 type conversion to C++ string type 705 unordered collection comparator 198 hash-bucket containers current container getting 270 maximum buckets in 269, 271 unordered object map 273 (see also ooHashMap in the Classes Index) (see also scalable collection) adding elements 278, 281 creating 276 finding object looking up data 280 looking up key 280 getting iterator for keys 280 iterator for values 282 removing elements 281, 282 replacing a value 281 testing for contained keys 279 for contained values 279 unordered set 283 (see also ooHashSet in the Classes Index) (see also scalable collection) adding elements 287 creating 286 finding object

comparator 288 hash administrator 288 looking up data 289 getting iterator for elements 290 number of elements 291 object hash value 289 refreshing internal containers 290 removing elements 290 testing for contained elements 288 for empty set 289 unsigned integer types 79 update mode for indexes 49 updating indexes explicitly 76 persistent object 453, 459 VArray 719 temporary 702 upgrade application, identifying as 659

#### ۷

variable-size array (see VArray) variable-size string 721 (see also ooVString in the Classes Index) appending to 726 assigning to 726 creating 725 getting character 726 first character 730 length 730 resizing 730 testing for equality 727 for inequality 727 whether greater-than-or-equal-to 729 whether less-than-or-equal-to 728 type conversion to C++ string type 730 variables 12 ooUserDefinedOperators 77 oovLastError 78

oovLastErrorLevel 79 oovNError 79 VArray 707 (see also ooVArrayT<element\_type> in the Classes Index) (see also temporary VArray) adding element 715, 716 assigning to 713 creating 712 element type, valid 708 embedded in persistent object 709 getting current number of elements 714, 719, 720 first element 716 iterator for finding elements 715 specified element 714, 715, 718 removing element 717 replacing element 717 resizing 715, 716, 718 effect of 709 setting element value 719 standard 707 structure and behavior 708 testing for empty 717 updating 719 VArray iterator 119 (see also d Iterator<element type> in the Classes Index) advancing 122, 123, 124 assigning 123 creating 121 finding element 123 moving backward 122 moving forward 122 reinitializing 124 testing for completion 124 versioning adding any version to a genealogy 461

#### any version to a genealogy 461 default version to a genealogy 259, 461, 627 derivative version 442, 443 next version 444

previous version 462 versions to a genealogy 255 associations for 434 behavior 78 defining copy semantics for 457 enabling for an object 628 finding all versions in a genealogy 254 default version in a genealogy 256, 615 derivative versions 448, 449 genealogy from any version 452 genealogy from default version 444 next version(s) 453, 617 previous version 460, 618 mode 78 getting 618 setting 628 removing any version from a genealogy 446 default version from a genealogy 257, 445 derivative versions 445, 446, 462, 463 next version 447, 464 previous version 447 setting default version 259, 461, 627 testing for existence any version in a genealogy 451 default version in a genealogy 450 derivative versions 450, 451 next version 452 previous version 452 testing whether enabled 618 virtual memory allocated for large objects 51, 67 virtual-function table controlling warning messages about 33 registration code for 34 Visual C++ exit function 43

#### W

weight of database image 234 setting 566

W

# **Classes Index**

This index contains an alphabetical list of classes, with member functions listed under each class. For a list of topics that are discussed in this book, see "Topic Index" on page 733. For an alphabetical list of all functions, including member functions, see "Functions and Macros Index" on page 773. For a list of non-class types and constants, see "Types and Constants Index" on page 787.

> is\_between 108 is\_leap\_vear 108

#### Α

pClass (application-defined class) 81	is_valid_date 108
add_linkName 85	month 108
del_linkName 86	next 109
exist_linkName 86	operator++ 104
linkName 87	operator + = 105
set_linkName 90	operator - = 105
sub_linkName 90	operator - = 105
D	previous 109
d_Database class 93	related global operators 104, 105, 106, 107
<pre>close 96 get_object_name 96 lookup_object 96 open 97 rename_object 97 set_object_name 98 d_Date class 99 constructor 103 current 107 day_of_week 107 day_of_wear 107 days_in_month 108 days_in_year 108</pre>	year 109 d_Interval class 111 constructor 114 day 118 hour 118 is_zero 118 minute 118 operator+= 115 operator-= 116 operator*= 116 operator/= 116 operator= 116

0

related global operators 115, 116, 117, 118 second 118 d\_Iterator<element\_type> class 119 constructor 121 advance 123 get element 123 next 124 not done 124 operator++ 122 operator-- 122 operator = 123reset 124 d\_Object class 431 d\_Ref\_Any class 125 constructor 126 clear 128 delete\_object 129 is null 129 operator= 127 related global operators 127, 128 d\_Ref<appClass> class 471 d\_Ref<d\_Object> class 593 d\_String class 721 d Time class 131 constructor 138 current 142 hour 142 minute 142 operator += 140operator = 140operator= 141 related global functions 142 related global operators 140, 141, 142 second 143 set default Time Zone 143 set default Time Zone to local 143 tz hour 143 tz minute 143 d\_Timestamp class 145 constructor 147 current 151 date 151 day 152 hour 152

minute 152 month 152 operator+= 149 operator= 150 related global operators 149, 150, 151 second 152 time 152 tz\_hour 152 tz\_minute 153 year 153 **d\_Transaction class** 653 **d\_Varray<element\_type> class** 707

# 0

ooAdmin class 155 ooAPObj class 157 constructor 159 operator new 160 ooBTree class 163 compact 165 contains 166 depth 166 first 166 get 167 indexOf 167 isEmpty 168 iterator 168 last 169 lastIndexOf 169 refresh 170 remove 170 size 171 ooCollection class 173 add 176 addAll 176 admin 177 clear 177 comparator 177 contains 177 containsAll 178 get 179 isEmpty 179

iterator 179 kevIterator 180 refresh 180 remove 181 removeAll 181 removeAllDeleted 182 retainAll 182 size 183 valueIterator 183 ooCollectionIterator class 185 collection 189 current 189 currentIndex 189 currentValue 189 goTo 190 goToIndex 191 hasNext 191 hasPrevious 191 moveCurrentTo 192 next 192 nextIndex 193 previous 193 previousIndex 194 remove 194 set 195 ooCompare class 197 compare 200 hash 200 ooContext class 203 constructor 205 destructor 205 current 205 setCurrent 206 setCurrentShared 206 ooContObj class 207 constructor 212 ooThis 215 operator new 212 ooConvertInObject class 217 getFloat32 219 getFloat64 220 getInt8 220 getInt16 220 getInt32 221

getInt64 221 getOldBaseClass 221 getOldDataMember 222 getUInt8 222 getUInt16 223 getUInt32 223 getUInt64 223 ooConvertInOutObject class 225 getNewBaseClass 227 getNewDataMember 228 setFloat32 228 setFloat64 229 setInt8 229 setInt16 229 setInt32 230 setInt64 230 setUInt8 230 setUInt16 231 setUInt32 231 setUInt64 231 ooDBObj class 233 constructor 236 operator new 237 ooDefaultContObj class 239 ooEqualLookupField class 241 constructor 242 ooFDObj class 245 ooGCContObi class 247 constructor 249 ooThis 250 ooGCRootsCont class 248 ooGeneObj class 251 constructor 254 add allVers 255 allVers 254 defaultVers 256 del allVers 257 del defaultVers 257 exist allVers 258 exist defaultVers 258 ooThis 258 set defaultVers 259 sub allVers 260

ooGreaterThanEqualLookupField class 261 constructor 262 ooGreaterThanLookupField class 265 constructor 266 ooHandle(className) classes (see ooRefHandle(appClass) classes) (see ooRefHandle(ooAPObj) classes) (see ooRefHandle(ooContObj) classes) (see ooRefHandle(ooDBObj) classes) (see ooRefHandle(ooFDObj) classes) (see ooRefHandle(ooObj) classes) ooHashAdmin class 269 bucketContainer 270 maxBucketsPerContainer 271 setMaxBucketsPerContainer 271 ooHashMap class 273 constructor 276 add 278 addAll 278 containsKev 279 containsValue 279 get 280 keyIterator 280 put 281 remove 281 removeAllDeleted 282 valueIterator 282 ooHashSet class 283 constructor 286 add 287 admin 288 comparator 288 contains 288 get 289 hashOf 289 isEmpty 289 iterator 290 refresh 290 remove 290 size 291 ooItr(appClass) class 293 constructor 295 end 295

next 295 scan 296 ooItr(ooAPObj) class 299 constructor 300 end 300 next 301 scan 301 ooItr(ooContObj) class 303 constructor 304 end 304 next 305 scan 305 ooItr(ooDBObj) class 309 constructor 310 end 310 next 311 scan 311 ooItr(ooObj) class 313 constructor 315 end 316 next 316 scan 316 oojArray class 319 getDimensionsArray 321 oojArrayOfBoolean class 323 constructor 324 getBooleanArray 325 ooiArravOfCharacter class 327 constructor 328 getCharacterArray 329 oojArrayOfDouble class 331 constructor 332 getDoubleArray 333 oojArrayOfFloat class 335 constructor 336 getFloatArray 337 oojArrayOfInt8 class 339 constructor 340 getInt8Array 341 oojArrayOfInt16 class 343 constructor 344 getInt16Array 345

oojArrayOfInt32 class 347 constructor 348 getInt32Array 349 oojArrayOfInt64 class 351 constructor 352 getInt64Array 353 oojArrayOfObject class 355 constructor 357 getObjectArray 357 ooiDate class 359 constructor 360 getMillis 360 setMillis 361 oojString class 363 constructor 364 getStringValue 364 oojTime class 365 constructor 366 getMillis 366 setMillis 367 oojTimestamp class 369 constructor 371 getMillis 371 getNanos 371 setMillis 372 setNanos 372 ooKevDesc class 373 constructor 375 addField 376 createIndex 376 dropIndex 377 getTypeN 377 getTypeName 377 isConsistent 377 isUnique 378 nField 378 removeIndexes 378 ooKeyField class 379 constructor 383 getName 384 getTypeN 385 isConsistent 385

ooLessThanEqualLookupField class 387 constructor 388 ooLessThanLookupField class 391 constructor 392 ooLookupFieldBase class 395 constructor 396 isNamed 396 ooLookupKey class 399 constructor 406 addField 406 anvIndex 407 nField 408 ooMap class 409 constructor 414 add 415 clearParam 416 forceAdd 416 isMember 416 lookup 417 maxAvgDensitv 417 nameHashFunction 418 nBin 418 nElement 418 percentGrow 418 printStat 418 refEnable 419 rehash 419 remove 420 replace 420 set nameHashFunction 421 set refEnable 421 ooMapElem class 423 name 424 oid 424 set oid 425 ooMapItr class 427 constructor 428 next 429 operator= 429 ooObi class 431 constructor 439 add derivatives 442 add derivedFrom 443 add nextVers 444

isNamed 385

defaultToGeneObj 444 del defaultToGeneObj 445 del derivatives 445 del derivedFrom 446 del\_geneObj 446 del nextVers 447 del prevVers 447 derivatives 448 derivedFrom 449 exist defaultToGeneObj 450 exist derivatives 450 exist derivedFrom 451 exist\_geneObj 451 exist nextVers 452 exist\_prevVers 452 geneObj 452 mark modified 453 nextVers 453 ooCopyInit 455 ooGetTypeN 455 ooGetTypeName 456 oolsKindOf 456 ooNewVersInit 457 ooPostMoveInit 457 ooPreMoveInit 458 ooThis 89. 458 ooUpdate 459 ooValidate 460 operator delete 439 operator new 441 prevVers 460 set\_defaultToGeneObj 461 set\_geneObj 461 set prevVers 462 sub derivatives 462 sub derivedFrom 463 sub\_nextVers 464 ooOperatorSet class 465 constructor 466 clear 467 registerOperator 467 **ooQuery class** 469 evaluate 470 setup 470

ooRefHandle(appClass) classes 471 constructor. handle 480 constructor, object reference 480, 481 copy 486 lookupObj 487 operator appClass\* (handle only) 484 operator d Ref Any 484 operator-> 481 operator\* (handle only) 482 operator= 483 ptr 488 ooRefHandle(ooAPObj) classes 489 constructor, handle 496 constructor, object reference 496, 497 bootFileHost 497 bootFilePath 498 change 498 close 499 containedIn 499 containersControlledBy 500 exist 501 imagesContainedIn 501 isAvailable 502 isOffline 502 isValid 502 jnlDirHost 503 inlDirPath 503 lockServerHost 504 markOffline 504 markOnline 504 name 505 open 505 openMode 506 operator = 497returnAll 506 svsDBFileHost 507 sysDBFilePath 507 typeN 507 typeName 507 update 508 ooRefHandle(ooContObj) classes 509 constructor, handle 518, 519 constructor, object reference 519 close 524 containedIn 524

contains 525 controlledBy 525 convertObjects 526 exist 527 hash 527 isUpdated 528 lockNoProp 528 lookupObj 529 name 530 nPage 530 numLogicalPages 530 open 531 openMode 532 operator d Ref Any 522 operator ooContObj\* (handle only) 523 operator-> 520 operator\* (handle only) 521 operator= 521 percentGrow 533 ptr 533 refreshOpen 534 returnControl 535 transferControl 535 ooRefHandle(ooDBObj) classes 537 constructor. handle 546 constructor, object reference 546 change 547 changePartition 548 close 549 containedIn 549 containingPartition 549 contains 550 convertObjects 551 deleteImage 551 exist 552 fileName 553 getAllowNonQuorumRead 553 getDefaultContObj 554 getImageFileName 554 getImageHostName 555 getImagePathName 555 getImageWeight 556 getTieBreaker 556 hasImageIn 556 hostName 557

isAvailable 557 isImageAvailable 557 isNonQuorumRead 558 isReadOnly 558 isReplicated 559 isValid 559 lock 559 name 560 negotiateQuorum 561 numContObjs 561 numImages 562 open 562 openMode 563 operator= 547 partitionsContainingImage 563 pathName 564 replicate 564 setAllowNonQuorumRead 565 setImageWeight 566 setReadOnly 566 setTieBreaker 567 tidy 567 typeN 568 typeName 568 update 569 ooRefHandle(ooFDObj) classes 571 constructor, handle 578 constructor, object reference 578 bootAP 579 change 580 close 581 contains 581 convertObjects 582 dumpCatalog 584 exist 584 isValid 585 lock 585 lockServerName 586 name 586 number 587 open 587 openMode 588 operator= 579 pageSize 588 setConversion 589

tidy 589 typeN 590 typeName 591 update 591 upgradeObjects 591, 659 ooRefHandle(ooObj) classes 593 constructor, handle 605, 606 constructor, object reference 606 close 613 containedIn 614 copy 614 delete\_object 615 getDefaultVers 615 getNameObj 616 getNameScope 616 getNextVers 617 getObjName 617 getPrevVers 618 getVersStatus 618 is null 619 isNull 619 isValid 619 lock 620 lockNoProp 621 lookupObj 621 move 622 nameObi 623 open 624 openMode 625 operator d\_Ref\_Any 611 operator int (object reference only) 612 operator ooObj\* (handle only) 612 operator-> 607 operator\* (handle only) 608 operator = 608operator = 610operator! = 611print 625 ptr 626 set\_container 627 setDefaultVers 627 setVersStatus 628 sprint 628 typeN 629 typeName 629

unnameObj 630 update 630 ooRef(className) classes (see ooRefHandle(appClass) classes) (see ooRefHandle(ooAPObj) classes) (see ooRefHandle(ooContObj) classes) (see ooRefHandle(ooDBObj) classes) (see ooRefHandle(ooFDObj) classes) (see ooRefHandle(ooObj) classes) ooShortRef(appClass) class 631 constructor 633 operator= 634 ooShortRef(ooObj) class 637 constructor 639 isNull 643 operator int 643 operator = 640operator = 641operator! = 642print 644 sprint 644 ooString(N) class 645 constructor 648 head 651 length 652 operator const char \* 651 operator ooVString 651 operator[] 649 operator = 649operator = 650operator = 650operator! = 651resize 652 ooTrans class 653 constructor 655 abort 655 begin 656 checkpoint 656 commit 656 commitAndHold 656 getID 657 isActive 657 start 657 upgrade 659

ooTreeAdmin class 661 maxNodesPerContainer 663 maxVArravsPerContainer 663. 664 nodeContainer 664 setMaxNodesPerContainer 664 vArrayContainer 665 ooTreeList class 667 constructor 670 add 671 addAll 672 addFirst 672 addLast 673 admin 673 comparator 673 contains 674 first 674 get 674 iterator 675 remove 675 removeAllDeleted 675 removeRange 676 set 676 ooTreeMap class 677 constructor 680 add 681 addAll 681 compact 682 containsKey 682 containsValue 683 get 683 put 684 remove 685 removeAllDeleted 685 valueIterator 685 ooTreeSet class 687

constructor 690 add 691 admin 691 comparator 691 contains 692 get 692 remove 693 ooTVArrayT<element\_type> class 695 constructor 698 elem 700 extend 700 head 701 operator[] 699 operator = 699resize 701 set 701 size 702 update 702 ooTVArray(element\_type) class 695 ooUtf8String class 703 operator const char \* 705 operator= 705 constructor 705 ooVArrayT<element\_type> class 707 constructor 712 cardinality 714 create iterator 715 elem 715 extend 715 head 716 insert element 716 is empty 717 operator[] 714 operator = 713remove\_all 717 replace element at 717 resize 718 retrieve element at 718 set 719 size 719 update 719 upper\_bound 720 ooVArray(element\_type) class 707 ooVString class 721 constructor 725 head 730 length 730 operator const char \* 730 operator[] 726

operator+= 726 operator= 726 related global operators 727, 728, 729 resize 730

# **Functions and Macros Index**

This index contains an alphabetical list of all global functions, member functions, and global macros. For a list of topics that are discussed in this book, see "Topic Index" on page 733. For an alphabetical list of classes, with member functions listed under each class, see "Classes Index" on page 763. For a list of non-class types and constants, see "Types and Constants Index" on page 787

#### **Symbols**

[] (see operator[]) + (see operator+) ++ (see operator++) += (see operator+=) - (see operator-) -- (see operator--) -= (see operator-=) -> (see operator->) \* (see operator\*) \*= (see operator\*=) / (see operator/) /= (see operator/=) = (see operator=) == (see operator==) != (see operator!=) < (see operator<) <= (see operator<=) > (see operator>) >= (see operator>=)

### Α

abort member function of ooTrans class 655 add member function of ooCollection class 176 of ooHashMap class 278 of ooHashSet class 287 of ooMap class 415 of ooTreeList class 671 of ooTreeMap class 681 of ooTreeSet class 691 add allVers member function of ooGeneObj class 255 add derivatives member function of ooObj class 442 add\_derivedFrom member function of ooObj class 443 add linkName member function of application-defined class 85 add nextVers member function of ooObj class 444 addAll member function of ooCollection class 176 of ooHashMap class 278

of ooTreeList class 672 of ooTreeMap class 681 addField member function of ooKeyDesc class 376 of ooLookupKey class 406 addFirst member function of ooTreeList class 672 addLast member function of ooTreeList class 673 admin member function of ooCollection class 177 of ooHashSet class 288 of ooTreeList class 673 of ooTreeSet class 691 advance member function of d\_Iterator<element\_type> class 123 allVers member function of ooGeneObj class 254 anyIndex member function of ooLookupKey class 407

### В

begin member function of ooTrans class 656
bootAP member function of ooRefHandle(ooFDObj) classes 579
bootFileHost member function of ooRefHandle(ooAPObj) classes 497
bootFilePath member function of ooRefHandle(ooAPObj) classes 498
bucketContainer member function of ooHashAdmin class 270

# С

cardinality member function of ooVArrayT<element\_type> class 714 change member function of ooRefHandle(ooAPObj) classes 498 of ooRefHandle(ooDBObj) classes 547 of ooRefHandle(ooFDObj) classes 580 changePartition member function of ooRefHandle(ooDBObj) classes 548 checkpoint member function of ooTrans class 656 clear member function of d Ref Any class 128 of ooCollection class 177 of ooOperatorSet class 467 clearParam member function of ooMap class 416 close member function of d Database class 96 of ooRefHandle(ooAPObj) classes 499 of ooRefHandle(ooContObj) classes 524 of ooRefHandle(ooDBObj) classes 549 of ooRefHandle(ooFDObj) classes 581 of ooRefHandle(ooObj) classes 613 collection member function of ooCollectionIterator class 189 commit member function of ooTrans class 656 commitAndHold member function of ooTrans class 656 compact member function of ooBTree class 165 of ooTreeMap class 682 comparator member function of ooCollection class 177 of ooHashSet class 288 of ooTreeList class 673 of ooTreeSet class 691 compare member function of ooCompare class 200 containedIn member function of ooRefHandle(ooAPObi) classes 499 of ooRefHandle(ooContObj) classes 524 of ooRefHandle(ooDBObj) classes 549 of ooRefHandle(ooObj) classes 614 containersControlledBy member function of ooRefHandle(ooAPObj) classes 500 containingPartition member function of ooRefHandle(ooDBObj) classes 549 contains member function of ooBTree class 166 of ooCollection class 177 of ooHashSet class 288

of ooRefHandle(ooContObj) classes 525 of ooRefHandle(ooDBObj) classes 550 of ooRefHandle(ooFDObj) classes 581 of ooTreeList class 674 of ooTreeSet class 692 containsAll member function of ooCollection class 178 containsKey member function of ooHashMap class 279 of ooTreeMap class 682 containsValue member function of ooHashMap class 279 of ooTreeMap class 683 controlledBy member function of ooRefHandle(ooContObj) classes 525 convertObjects member function of ooRefHandle(ooContObj) classes 526 of ooRefHandle(ooDBObj) classes 551 of ooRefHandle(ooFDObj) classes 582 copy member function of ooRefHandle(appClass) classes 486 of ooRefHandle(ooObj) classes 614 create iterator member function of ooVArrayT<element\_type> class 715 createIndex member function of ooKeyDesc class 376 current member function of d Date class 107 of d\_Time class 142 of d Timestamp class 151 of ooCollectionIterator class 189 of ooContext class 205 currentIndex member function of ooCollectionIterator class 189 currentValue member function of ooCollectionIterator class 189

#### D

d\_Date constructor 103 d\_Interval constructor 114 d\_Iterator<element\_type> constructor 121 d\_Ref\_Any constructor 126 d\_Time constructor 138 d\_Timestamp constructor 147 date member function of d\_Timestamp class 151 day member function of d Date class 107 of d Interval class 118 of d Timestamp class 152 day\_of\_week member function of d Date class 107 day\_of\_year member function of d Date class 107 defaultToGeneObj member function of ooObj class 444 defaultVers member function of ooGeneObi class 256 del\_allVers member function of ooGeneObj class 257 del\_defaultToGeneObj member function of ooObj class 445 del\_defaultVers member function of ooGeneObj class 257 del\_derivatives member function of ooObj class 445 del derivedFrom member function of ooObj class 446 del\_geneObj member function of ooObj class 446 del linkName member function of application-defined class 86 del nextVers member function of ooObj class 447 del\_prevVers member function of ooObj class 447 delete operator (see operator delete) delete object member function of d\_Ref\_Any class 129 of ooRefHandle(ooObj) classes 615 deleteImage member function of ooRefHandle(ooDBObj) classes 551 depth member function of ooBTree class 166 derivatives member function of ooObj class 448

derivedFrom member function of ooObj class 449 dropIndex member function of ooKeyDesc class 377 dumpCatalog member function of ooRefHandle(ooFDObj) classes 584

### Ε

elem member function of ooTVArrayT<element\_type> class 700 of ooVArrayT<element\_type> class 715 end member function of ooItr(appClass) class 295 of ooItr(ooAPObj) class 300 of ooItr(ooContObj) class 304 of ooItr(ooDBObj) class 310 of ooItr(ooObj) class 316 evaluate member function of ooQuery class 470 exist member function of ooRefHandle(ooAPObi) classes 501 of ooRefHandle(ooContObj) classes 527 of ooRefHandle(ooDBObj) classes 552 of ooRefHandle(ooFDObj) classes 584 exist allVers member function of ooGeneObj class 258 exist\_defaultToGeneObj member function of ooObj class 450 exist defaultVers member function of ooGeneObj class 258 exist derivatives member function of ooObj class 450 exist derivedFrom member function of ooObj class 451 exist\_geneObj member function of ooObj class 451 exist linkName member function of application-defined class 86 exist nextVers member function of ooObj class 452 exist\_prevVers member function of ooObj class 452

extend member function of ooTVArrayT<element\_type> class 700 of ooVArrayT<element\_type> class 715

#### F

fileName member function of ooRefHandle(ooDBObj) classes 553 first member function of ooBTree class 166 of ooTreeList class 674 forceAdd member function of ooMap class 416

### G

geneObj member function of ooObj class 452 get member function of ooBTree class 167 of ooCollection class 179 of ooHashMap class 280 of ooHashSet class 289 of ooTreeList class 674 of ooTreeMap class 683 of ooTreeSet class 692 get element member function of d Iterator<element type> class 123 get\_object\_name member function of d Database class 96 getAllowNonQuorumRead member function of ooRefHandle(ooDBObj) classes 553 getBooleanArray member function of oojArrayOfBoolean class 325 getCharacterArray member function of oojArrayOfCharacter class 329 getDefaultContObj member function of ooRefHandle(ooDBObj) classes 554 getDefaultVers member function of ooRefHandle(ooObj) classes 615 getDimensionsArray member function of oojArray class 321 getDoubleArray member function of oojArrayOfDouble class 333

getFloat32 member function of ooConvertInObject class 219 getFloat64 member function of ooConvertInObject class 220 getFloatArray member function of oojArrayOfFloat class 337 getID member function of ooTrans class 657 getImageFileName member function of ooRefHandle(ooDBObj) classes 554 getImageHostName member function of ooRefHandle(ooDBObj) classes 555 getImagePathName member function of ooRefHandle(ooDBObj) classes 555 getImageWeight member function of ooRefHandle(ooDBObj) classes 556 getInt8 member function of ooConvertInObject class 220 getInt8Array member function of oojArrayOfInt8 class 341 getInt16 member function of ooConvertInObject class 220 getInt16Array member function of oojArrayOfInt16 class 345 getInt32 member function of ooConvertInObject class 221 getInt32Array member function of oojArrayOfInt32 class 349 getInt64 member function of ooConvertInObject class 221 getInt64Array member function of oojArrayOfInt64 class 353 getMillis member function of oojDate class 360 of oojTime class 366 of oojTimestamp class 371 getName member function of ooKeyField class 384 getNameObj member function of ooRefHandle(ooObj) classes 616 getNameScope member function of ooRefHandle(ooObj) classes 616 getNanos member function

of oojTimestamp class 371 getNewBaseClass member function of ooConvertInOutObject class 227 getNewDataMember member function of ooConvertInOutObject class 228 getNextVers member function of ooRefHandle(ooObj) classes 617 getObjectArray member function of oojArrayOfObject class 357 getObjName member function of ooRefHandle(ooObj) classes 617 getOldBaseClass member function of ooConvertInObject class 221 getOldDataMember member function of ooConvertInObject class 222 getPrevVers member function of ooRefHandle(ooObj) classes 618 getStringValue member function of oojString class 364 getTieBreaker member function of ooRefHandle(ooDBObj) classes 556 getTypeN member function of ooKeyDesc class 377 of ooKeyField class 385 getTypeName member function of ooKeyDesc class 377 getUInt8 member function of ooConvertInObject class 222 getUInt16 member function of ooConvertInObject class 223 getUInt32 member function of ooConvertInObject class 223 getUInt64 member function of ooConvertInObject class 223 getVersStatus member function of ooRefHandle(ooObj) classes 618 goTo member function of ooCollectionIterator class 190 goToIndex member function of ooCollectionIterator class 191

### Н

hash member function of ooCompare class 200 of ooRefHandle(ooContObj) classes 527 hashOf member function of ooHashSet class 289 hasImageIn member function of ooRefHandle(ooDBObj) classes 556 hasNext member function of ooCollectionIterator class 191 hasPrevious member function of ooCollectionIterator class 191 head member function of ooString(N) class 651 of ooTVArrayT<element\_type> class 701 of ooVArrayT<element\_type> class 716 of ooVString class 730 hostName member function of ooRefHandle(ooDBObj) classes 557 hour member function of d Interval class 118 of d Time class 142 of d\_Timestamp class 152

### I

imagesContainedIn member function of ooRefHandle(ooAPObj) classes 501 indexOf member function of ooBTree class 167 insert element member function of ooVArrayT<element\_type> class 716 is empty member function of ooVArrayT<element\_type> class 717 is\_leap\_year member function of d Date class 108 is null member function of d Ref Any class 129 of ooRefHandle(ooObj) classes 619 is valid date member function of d Date class 108 is zero member function of d\_Interval class 118

isActive member function of ooTrans class 657 isAvailable member function of ooRefHandle(ooAPObj) classes 502 of ooRefHandle(ooDBObj) classes 557 isConsistent member function of ooKeyDesc class 377 of ooKeyField class 385 isEmpty member function of ooBTree class 168 of ooCollection class 179 of ooHashSet class 289 isImageAvailable member function of ooRefHandle(ooDBObj) classes 557 isMember member function of ooMap class 416 isNamed member function of ooKeyField class 385 of ooLookupFieldBase class 396 isNonQuorumRead member function of ooRefHandle(ooDBObj) classes 558 isNull member function of ooRefHandle(ooObj) classes 619 of ooShortRef(ooObj) class 643 isOffline member function of ooRefHandle(ooAPObj) classes 502 isReadOnly member function of ooRefHandle(ooDBObj) classes 558 isReplicated member function of ooRefHandle(ooDBObj) classes 559 isUnique member function of KeyDesc class 378 isUpdated member function of ooRefHandle(ooContObj) classes 528 isValid member function of ooRefHandle(ooAPObj) classes 502 of ooRefHandle(ooDBObj) classes 559 of ooRefHandle(ooFDObj) classes 585 of ooRefHandle(ooObj) classes 619 iterator member function of ooBTree class 168 of ooCollection class 179

of ooHashSet class 290 of ooTreeList class 675

#### J

jnlDirHost member function of ooRefHandle(ooAPObj) classes 503 jnlDirPath member function of ooRefHandle(ooAPObj) classes 503

#### Κ

keyIterator member function of ooCollection class 180 of ooHashMap class 280

#### L

last member function of ooBTree class 169 lastIndexOf member function of ooBTree class 169 length member function of ooString(N) class 652 of ooVString class 730 linkName member function of application-defined class 87 lock member function of ooRefHandle(ooDBObj) classes 559 of ooRefHandle(ooFDObj) classes 585 of ooRefHandle(ooObj) classes 620 lockNoProp member function of ooRefHandle(ooContObj) classes 528 of ooRefHandle(ooObi) classes 621 lockServerHost member function of ooRefHandle(ooAPObi) classes 504 lockServerName member function of ooRefHandle(ooFDObj) classes 586 lookup member function of ooMap class 417 lookup\_object member function of d Database class 96 lookupObj member function of ooRefHandle(appClass) classes 487

of ooRefHandle(ooContObj) classes 529 of ooRefHandle(ooObj) classes 621

#### Μ

mark\_modified member function of ooObi class 453 markOffline member function of ooRefHandle(ooAPObj) classes 504 markOnline member function of ooRefHandle(ooAPObi) classes 504 maxAvgDensity member function of ooMap class 417 maxBucketsPerContainer member function of ooHashAdmin class 271 maxNodesPerContainer member function of ooTreeAdmin class 663 maxVArraysPerContainer member function of ooTreeAdmin class 663 minute member function of d Interval class 118 of d Time class 142 of d\_Timestamp class 152 month member function of d\_Date class 108 of d\_Timestamp class 152 move member function of ooRefHandle(ooObi) classes 622 moveCurrentTo member function of ooCollectionIterator class 192

#### Ν

name member function of ooMapElem class 424 of ooRefHandle(ooAPObj) classes 505 of ooRefHandle(ooContObj) classes 530 of ooRefHandle(ooDBObj) classes 560 of ooRefHandle(ooFDObj) classes 586 nameHashFunction member function of ooMap class 418 nameObj member function of ooRefHandle(ooObj) classes 623 nBin member function of ooMap class 418 negotiateQuorum member function of ooRefHandle(ooDBObj) classes 561 nElement member function of ooMap class 418 new operator (see operator new) next member function of d Date class 109 of d Iterator<element type> class 124 of ooCollectionIterator class 192 of ooItr(appClass) class 295 of ooItr(ooAPObj) class 301 of ooItr(ooContObj) class 305 of ooItr(ooDBObj) class 311 of ooItr(ooObi) class 316 of ooMapItr class 429 nextIndex member function of ooCollectionIterator class 193 nextVers member function of ooObi class 453 nField member function of ooKevDesc class 378 of ooLookupKey class 408 nodeContainer member function of ooTreeAdmin class 664 not done member function of d\_Iterator<element\_type> class 124 nPage member function of ooRefHandle(ooContObj) classes 530 number member function of ooRefHandle(ooFDObj) classes 587 numContObis member function of ooRefHandle(ooDBObj) classes 561 numImages member function of ooRefHandle(ooDBObj) classes 562 numLogicalPages member function of ooRefHandle(ooContObj) classes 530

#### 0

oid member function of ooMapElem class 424 ooAPObj constructor 159 ooCheckLS function 33 ooCheckVTablePointer function 33 ooCleanup function 35 ooContext constructor 205 ooContext destructor 205 ooContObj constructor 212 ooCopyInit member function of ooObj class 455 ooDBObj constructor 236 **ooDelete function** 38, 158, 209, 235, 433 ooDeleteNoProp function 40 ooEqualLookupField constructor 242 ooExitCleanup function 43 ooGCContObj constructor 249 ooGeneObj constructor 254 ooGetActiveTrans function 45 ooGetErrorHandler macro 46 ooGetMemberOffset macro 46 ooGetMemberSize macro 47 ooGetMsgHandler macro 47 ooGetOfflineMode function 47 ooGetResourceOwners function 48 ooGetTypeN member function of ooObj class 455 ooGetTypeName member function of ooObj class 456 ooGreaterThanEqualLookupField constructor 262 ooGreaterThanLookupField constructor 266 ooHandle(appClass) constructor 480 ooHandle(ooAPObj) constructor 496 ooHandle(ooContObj) constructor 518, 519 ooHandle(ooDBObj) constructor 546 ooHandle(ooFDObj) constructor 578 ooHandle(ooObj) constructor 605, 606 ooHashMap constructor 276 ooHashSet constructor 286 ooInit function 49 ooInitThread function 51 oolsKindOf member function of ooObj class 456 ooItr(appClass) constructor 295

ooItr(ooAPObj) constructor 300 ooItr(ooContObj) constructor 304 ooItr(ooDBObj) constructor 310 ooItr(ooObj) constructor 315 oojArrayOfBoolean constructor 324 oojArrayOfCharacter constructor 328 oojArrayOfDouble constructor 332 oojArrayOfFloat constructor 336 oojArrayOfInt8 constructor 340 oojArrayOfInt16 constructor 344 oojArrayOfInt32 constructor 348 oojArrayOfInt64 constructor 352 oojArrayOfObject constructor 357 oojDate constructor 360 oojString constructor 364 oojTime constructor 366 oojTimestamp constructor 371 ooKeyDesc constructor 375 ooKeyField constructor 383 ooLessThanEqualLookupField constructor 388 ooLessThanLookupField constructor 392 ooLookupFieldBase constructor 396 ooLookupKey constructor 406 ooMap constructor 414 ooMapItr constructor 428 ooNewConts macro 56 ooNewKey macro 57 ooNewVersInit member function of ooObj class 457 ooNoLock function 58 ooObj constructor 439 ooOperatorSet constructor 466 ooPostMoveInit member function of ooObj class 457 ooPreMoveInit member function of ooObj class 458 ooPurgeAps function 59 ooRef(appClass) constructor 480, 481 ooRef(ooAPObj) constructor 496, 497 ooRef(ooContObj) constructor 519 ooRef(ooDBObj) constructor 546

ooRef(ooFDObj) constructor 578 ooRef(ooObj) constructor 606 ooRegErrorHandler macro 61 ooRegMsgHandler macro 61 ooRegTwoMachineHandler function 62 ooReplace macro 63 ooResetError function 65 ooRunStatus function 65 ooSetAMSUsage function 66 ooSetErrorFile function 66 ooSetHotMode function 66 ooSetLargeObjectMemoryLimit function 67 ooSetLockMode function 658 ooSetLockWait function 68 ooSetOfflineMode function 68 ooSetRpcTimeout function 69 ooShortRef(appClass) constructor 633 ooShortRef(ooObj) constructor 639 ooSignal function 69 ooStartInternalLS function 70 ooStopInternalLS function 72 ooString constructor 648 ooTermThread function 73 ooThis member function of ooContObj class 215 of ooGCContObj class 250 of ooGeneObj class 258 of ooObj class 89, 458 ooTrans constructor 655 ooTreeList constructor 670 ooTreeMap constructor 680 ooTreeSet constructor 690 ooTVArrayT<element\_type> constructor 698 ooTypeN macro 75 ooUpdate member function of ooObj class 459 ooUpdateIndexes function 76 ooUseIndex function 77 ooUtf8String constructor 705 ooValidate member function of ooObj class 460 ooVArrayT<element\_type> constructor 712

ooVString constructor 725 open member function of d\_Database class 97 of ooRefHandle(ooAPObj) classes 505 of ooRefHandle(ooContObj) classes 531 of ooRefHandle(ooDBObj) classes 562 of ooRefHandle(ooFDObj) classes 587 of ooRefHandle(ooObj) classes 624 openMode member function of ooRefHandle(ooAPObj) classes 506 of ooRefHandle(ooContObj) classes 532 of ooRefHandle(ooDBObj) classes 563 of ooRefHandle(ooFDObj) classes 588 of ooRefHandle(ooObj) classes 625 operator appClass\* of ooHandle(appClass) classes 484 operator const char \* of ooString(N) class 651 of ooUtf8String class 705 of ooVString class 730 operator d Ref Any of ooRefHandle(appClass) classes 484 of ooRefHandle(ooContObj) classes 522 of ooRefHandle(ooObj) classes 611 operator days\_in\_month of d Date class 108 operator days\_in\_year of d\_Date class 108 operator delete of ooObj class 439 operator int of ooRef(ooObj) class 612 of ooShortRef(ooObj) class 643 operator is\_between of d\_Date class 108 operator new of ooAPObj class 160 of ooContObj class 212 of ooDBObj class 237 of ooObj class 441 operator ooContObj\* of ooHandle(ooContObj) class 523 operator ooObj\*

of ooHandle(ooObj) class 612

operator ooVString of ooString(N) class 651 operator[] of ooString(N) class 649 of ooTVArrayT<element\_type> class 699 of ooVArrayT<element\_type> class 714 of ooVString class 726 operator+ global, for d\_Date objects 104 global, for d Interval objects 115 global, for d\_Time objects 140 global, for d\_Timestamp objects 149 operator++ of d Date class 104 of d\_Iterator<element\_type> class 122 operator+= of d Date class 105 of d Interval class 115 of d\_Time class 140 of d Timestamp class 149 of ooString(N) class 649 of ooVString class 726 operatorglobal, for d\_Date objects 105 global, for d\_Interval objects 115 global, for d\_Time objects 140 global, for d\_Timestamp objects 149 of d Interval class 115 operator-of d Date class 105 of d Iterator<element type> class 122 operator-= of d Date class 105 of d Interval class 116 of d\_Time class 140 of d Timestamp class 150 operator-> of ooRefHandle(appClass) classes 481 of ooRefHandle(ooContObj) classes 520 of ooRefHandle(ooObj) classes 607 operator\* global, for d Interval objects 116 of ooHandle(appClass) class 482

of ooHandle(ooContObj) class 521 of ooHandle(ooObj) class 608 operator\*= of d\_Interval class 116 operator/ global, for d\_Interval objects 116 operator/= of d Interval class 116 operator= of d\_Date class 106 of d Interval class 116 of d\_Iterator<element\_type> class 123 of d\_Ref\_Any class 127 of d\_Time class 141 of d Timestamp class 150 of ooMapItr class 429 of ooRefHandle(appClass) classes 483 of ooRefHandle(ooAPObj) classes 497 of ooRefHandle(ooContObj) classes 521 of ooRefHandle(ooDBObj) classes 547 of ooRefHandle(ooFDObj) classes 579 of ooRefHandle(ooObj) classes 608 of ooShortRef(appClass) class 634 of ooShortRef(ooObj) class 640 of ooString(N) class 650 of ooTVArrayT<element\_type> class 699 of ooUtf8String class 705 of ooVArrayT<element\_type> class 713 of ooVString class 726

operator==

global, for d\_Date objects 106 global, for d\_Interval objects 117 global, for d\_Ref\_Any objects 127 global, for d\_Time objects 141 global, for d\_Timestamp objects 150 global, for ooVString objects 727 of ooRefHandle(ooObj) classes 610 of ooShortRef(ooObj) class 641 of ooString(N) class 650

operator!=

global, for d\_Date objects 106 global, for d\_Interval objects 117 global, for d\_Ref\_Any objects 128 global, for d\_Time objects 141

global, for d\_Timestamp objects 150 global, for ooVString objects 727 of ooRefHandle(ooObj) classes 611 of ooShortRef(ooObj) class 642 of ooString(N) class 651 operator< global, for d\_Date objects 106 global, for d Interval objects 117 global, for d\_Time objects 141 global, for d\_Timestamp objects 150 global, for ooVString objects 728 operator<= global, for d\_Date objects 106 global, for d\_Interval objects 117 global, for d Time objects 141 global, for d\_Timestamp objects 151 global, for ooVString objects 728 operator> global, for d\_Date objects 107 global, for d Interval objects 117 global, for d\_Time objects 142 global, for d\_Timestamp objects 151 global, for ooVString objects 729 operator>= global, for d\_Date objects 107 global, for d Interval objects 118 global, for d\_Time objects 142 global, for d\_Timestamp objects 151 global, for ooVString objects 729 overlaps function global, for d Date objects 109 global, for d\_Time objects 144 global, for d\_Timestamp objects 153

#### Ρ

pageSize member function of ooRefHandle(ooFDObj) classes 588 partitionsContainingImage member function of ooRefHandle(ooDBObj) classes 563 pathName member function of ooRefHandle(ooDBObj) classes 564 percentGrow member function of ooMap class 418

of ooRefHandle(ooContObj) classes 533 previous member function of d Date class 109 of ooCollectionIterator class 193 previousIndex member function of ooCollectionIterator class 194 prevVers member function of ooObj class 460 print member function of ooRefHandle(ooObj) classes 625 of ooShortRef(ooObi) class 644 printStat member function of ooMap class 418 ptr member function of ooRefHandle(appClass) classes 488 of ooRefHandle(ooContObj) classes 533 of ooRefHandle(ooObj) classes 626 put member function of ooHashMap class 281 of ooTreeMap class 684

### R

refEnable member function of ooMap class 419 refresh member function of ooBTree class 170 of ooCollection class 180 of ooHashSet class 290 refreshOpen member function of ooRefHandle(ooContObj) classes 534 registerOperator member function of ooOperatorSet class 467 rehash member function of ooMap class 419 remove member function of ooBTree class 170 of ooCollection class 181 of ooCollectionIterator class 194 of ooHashMap class 281 of ooHashSet class 290 of ooMap class 420 of ooTreeList class 675

of ooTreeMap class 685 of ooTreeSet class 693 remove all member function of ooVArrayT<element\_type> class 717 removeAll member function of ooCollection class 181 removeAllDeleted member function of ooCollection class 182 of ooHashMap class 282 of ooTreeList class 675 of ooTreeMap class 685 removeIndexes member function of ooKeyDesc class 378 removeRange member function of ooTreeList class 676 rename\_object member function of d Database class 97 replace member function of ooMap class 420 replace element at member function of ooVArrayT<element\_type> class 717 replicate member function of ooRefHandle(ooDBObj) classes 564 reset member function of d Iterator<element type> class 124 resize member function of ooString(N) class 652 of ooTVArrayT<element\_type> class 701 of ooVArrayT<element\_type> class 718 of ooVString class 730 retainAll member function of ooCollection class 182 retrieve\_element\_at member function of ooVArrayT<element\_type> class 718 returnAll member function of ooRefHandle(ooAPObj) classes 506 returnControl member function of ooRefHandle(ooContObj) classes 535

### S

scan member function of ooAPObj class 301 of ooItr(appClass) class 296

of ooItr(ooContObj) class 305 of ooItr(ooDBObi) class 311 of ooItr(ooObi) class 316 second member function of d Interval class 118 of d Time class 143 of d\_Timestamp class 152 set member function of ooCollectionIterator class 195 of ooTreeList class 676 of ooTVArrayT<element type> class 701 of ooVArrayT<element\_type> class 719 set container member function of ooRefHandle(ooObj) classes 627 set default\_Time\_Zone member function of d Time class 143 set default Time Zone to local member function of d Time class 143 set\_defaultToGeneObj member function of ooObj class 461 set defaultVers member function of ooGeneObj class 259 set geneObj member function of ooObj class 461 set linkName member function of application-defined class 90 set\_nameHashFunction member function of ooMap class 421 set object name member function of d Database class 98 set oid member function of ooMapElem class 425 set\_prevVers member function of ooObj class 462 set\_refEnable member function of ooMap class 421 setAllowNonQuorumRead member function of ooRefHandle(ooDBObi) classes 565 setConversion member function of ooRefHandle(ooFDObj) classes 589 setCurrent member function of ooContext class 206

setCurrentShared member function of ooContext class 206 setDefaultVers member function of ooRefHandle(ooObj) classes 627 setFloat32 member function of ooConvertInOutObject class 228 setFloat64 member function of ooConvertInOutObject class 229 setImageWeight member function of ooRefHandle(ooDBObj) classes 566 setInt8 member function of ooConvertInOutObject class 229 setInt16 member function of ooConvertInOutObject class 229 setInt32 member function of ooConvertInOutObject class 230 setInt64 member function of ooConvertInOutObject class 230 setMaxBucketsPerContainer member function of ooHashAdmin class 271 setMaxNodesPerContainer member function of ooTreeAdmin class 664 setMaxVArraysPerContainer member function of ooTreeAdmin class 664 setMillis member function of ooiDate class 361 of ooiTime class 367 of oojTimestamp class 372 setNanos member function of oojTimestamp class 372 setReadOnly member function of ooRefHandle(ooDBObj) classes 566 setTieBreaker member function of ooRefHandle(ooDBObi) classes 567 setUInt8 member function of ooConvertInOutObject class 230 setUInt16 member function of ooConvertInOutObject class 231 setUInt32 member function of ooConvertInOutObject class 231 setUInt64 member function

of ooConvertInOutObject class 231 setup member function of ooQuery class 470 setVersStatus member function of ooRefHandle(ooObj) classes 628 size member function of ooBTree class 171 of ooCollection class 183 of ooHashSet class 291 of ooTVArrayT<element\_type> class 702 of ooVArrayT<element\_type> class 719 sprint member function of ooRefHandle(ooObj) classes 628 of ooShortRef(ooObj) class 644 start member function of ooTrans class 657 sub allVers member function of ooGeneObj class 260 sub derivatives member function of ooObj class 462 sub\_derivedFrom member function of ooObj class 463 sub\_linkName member function of application-defined class 90 sub nextVers member function of ooObj class 464 sysDBFileHost member function of ooRefHandle(ooAPObj) classes 507 svsDBFilePath member function of ooRefHandle(ooAPObj) classes 507

# Т

tidy member function of ooRefHandle(ooDBObj) classes 567 of ooRefHandle(ooFDObj) classes 589 time member function of d\_Timestamp class 152 transferControl member function of ooRefHandle(ooContObj) classes 535 typeN member function of ooRefHandle(ooAPObj) classes 507 of ooRefHandle(ooFDObj) classes 568 of ooRefHandle(ooFDObj) classes 590

786

of ooRefHandle(ooObj) classes 629 **typeName member function** of ooRefHandle(ooAPObj) classes 507 of ooRefHandle(ooDBObj) classes 568 of ooRefHandle(ooFDObj) classes 591 of ooRefHandle(ooObj) classes 629 **tz\_hour member function** 

of d\_Time class 143 of d\_Timestamp class 152

tz\_minute member function of d\_Time class 143 of d\_Timestamp class 153

#### U

unnameObj member function of ooRefHandle(ooObj) classes 630 update member function of ooRefHandle(ooAPObj) classes 508 of ooRefHandle(ooDBObj) classes 569 of ooRefHandle(ooFDObj) classes 591 of ooRefHandle(ooObj) classes 630 of ooTVArrayT<element\_type> class 702 of ooVArrayT<element\_type> class 719 upgrade member function of ooTrans class 659 upgradeObjects member function of ooRefHandle(ooFDObj) classes 591, 659 upper\_bound member function

of ooVArrayT<element\_type> class 720

#### V

valueIterator member function of ooCollection class 183 of ooHashMap class 282 of ooTreeMap class 685 vArrayContainer member function of ooTreeAdmin class 665

#### Υ

year member function of d\_Date class 109 of d\_Timestamp class 153

Т

# **Types and Constants Index**

This index lists non-class types and constants. For a list of topics that are discussed in this book, see "Topic Index" on page 733. For an alphabetical list of classes, see "Classes Index" on page 763. For a list of functions, including member functions, see "Functions and Macros Index" on page 773.

#### Α

access\_status type (ODMG) 95

#### С

char type 31

#### Е

exclusive constant 95

#### F

float32 type 31 float64 type 31

#### I

int8 type 31 int16 type 32 int32 type 32 int64 type 32

#### Μ

Month type of d\_Date class 102

#### Ν

not\_open constant 95

#### 0

ooAccessMode type 32 ooAMSUsage type 32 ooBoolean type 33 oocAll constant 32 oocAllObjs constant 37 oocAMSOnly constant 32 oocAMSPreferred constant 32 oocBooleanT constant 38 oocBranchVers constant 78 oocCharArray constant 53 oocCharPtrT constant 38 oocDowngradeAll constant 41 oocEnforce constant 59 oocError constant 72 oocExplicitUpdate constant 49, 658 oocFalse constant 33 oocFatalError constant 43 oocFloat32 constant 53 oocFloat64 constant 53 oocFloat64T constant 38

ooChar type 25 oocHostLocal constant 44 oocIgnore constant 58 oocInsensitive constant 49, 657, 658 oocInt16 constant 53 oocInt32 constant 53 oocInt32T constant 38 oocInValidTransId constant 74 oocInvalidTypeT constant 38 oocLinearVers constant 78 oocLockRead constant 54 oocLockUpdate constant 54 oocMROW constant 54 oocNative constant 44 oocNoAMS constant 33 oocNoDowngrade constant 40 oocNoError constant 42 oocNoMROW constant 54 oocNoOpen constant 54 oocNotTransferred constant 37 oocNoVers constant 78 oocNoWait constant 68, 80, 658 ooContainsFilter type 37 **ooConvertFunction function-pointer type** 37 oocPublic constant 32 oocRead constant 54 oocSensitive constant 49, 658 oocString constant 53 oocSuccess constant 72 oocSystemError constant 43 oocTransferred constant 37 oocTransNoWait constant 80, 658 oocTrue constant 33 oocUint16 constant 53 oocUint32 constant 53 oocUint32T constant 38 oocUpdate constant 54 oocUserError constant 42 oocWait constant 68, 80, 658 oocWarning constant 42 ooDataType type 38

ooDowngradeMode type 40 **ooError type** 41 ooErrorHandlerPtr function-pointer type 41 **ooErrorLevel type 42 ooFileNameFormat type 44** ooFloat32 type 25 ooFloat64 type 25 ooIndexMode type 49 ooInt8 type 25 ooInt16 type 25 ooInt32 type 25 ooInt64 type 25 ooKey type 52 ooKeyType type 53 ooLockMode type 54 **ooMode type** 54 ooMsgHandlerPtr function-pointer type 55 ooNameHashFuncPtr function-pointer type 55 **ooOfflineMode type** 58 ooQueryOperatorPtr function-pointer type 60 **ooResource type** 65 ooStatus type 72 ooTransId type 74 ooTransInfo type 74 ooTwoMachineHandlerPtr function-pointer type 75 ooTypeNumber type 76 ooUInt8 type 25 ooUInt16 type 25 ooUInt32 type 25 ooUInt64 type 25 ooUserDefinedOperators variable 77 ooVersMode type 78 oovLastError variable 78 oovLastErrorLevel variable 79 **oovNError variable** 79 ooVoidFuncPtr function-pointer type 79

### R

read\_only constant 95 read\_write constant 95 relational operator functions data types of operands in 38

### Т

Time\_Zone type 137 transient\_memory type (ODMG) 95

#### U

uint8 type 79 uint16 type 79 uint32 type 79 uint64 type 80

#### W

Weekday type of d\_Date class 102